Efficient Photonic Coding: a Considered Revision

Yury Audzevich *
Computer Laboratory,
University of Cambridge

Philip Watts
Department of Electronic and Electrical
Engineering, UCL

Sean Kilmurray
Department of Electronic and Electrical
Engineering, UCL

Andrew W. Moore
Computer Laboratory,
University of Cambridge

ABSTRACT

In this paper we reconsider the energy consumption of traditional DC-balanced physical line coding schemes applied to optical communication. We demonstrate that not only does an implementation of the popular 8B10B coding scheme have higher power consumption than the optical power requirement, but actually has higher power consumption when transmitting idle sequences than for real data packets. Furthermore, we show that simple codes retain the DC balance performance of 8B10B and hence do not increase the optical power requirement. We propose the use of a coding scheme that permits a default-off transmission system through the addition of a preamble. By analysis of trace data taken from a network covering a 24 hour period, we show that the power saving is up to 93%. The proposed approach not only enables energy proportional links but is fully compatible with future low power optical switched networks.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network communications; C.2.2 [**Network Protocols**]: Protocol architecture

General Terms

Measurement, Performance

Keywords

Energy-efficiency, physical line coding

1. INTRODUCTION

The energy performance of networked systems has become a *1st class* property, of interest to industry and researchers. It has been shown that computer systems must be made to 'do nothing well' for large energy savings to be made by achieving energy proportionality [2]. Multiple research groups have indicated that the network-equipment community lags behind on this goal of energy-proportionality. An example of such a profligate approach is optical Ethernet at 1Gb/s [10]; this approach maintains clock-synchronization between transmitter and receiver by sending a continuous stream of idle frames with a known transition density when there is no data to be transmitted.

Increasing use is being made of optical links within computer facilities due to high signal integrity at high bit rates and power consumption that is (on the scale of a machine room or data centre) independent of link distance. In the longer term, switched optical networks will provide reduced hop count transactions between servers, further reducing energy requirements [8, 18]. Quite apart from considerations of the energy which can be saved by switching off optical transceivers when not transmitting, the switched optical network scenario requires so-called burst-mode receivers[1] capable of fast locking to incoming packets which have different frequency, phase and amplitude. Conventional optical receivers use an AC coupled input stage which significantly simplifies the design of the amplification and data recovery circuits at high bit rates. However, for burst mode receivers, AC coupling results in large baseline wander (BLW). Previous burst mode receivers at low bit rates (1 Gb/s) have solved the BLW issue by using a DC coupled receiver [16] or by employing 8B10B coding [21]. 10 Gb/s burst mode operation without DC balance coding has been demonstrated [22], but as the technique requires a 20 GSample/s analog-to-digital converter (ADC) running continuously along with considerable digital signal processing (DSP), the energy characteristics are not favourable. The use of a coding scheme specifically designed for burst mode optical with a lower complexity receiver will lead to energy proportionality for the optical network and the ability to fully exploit low energy properties of future switched optical networks.

In this paper we examine the energy characteristics of one aspect of future protocols for optical networking: DC balance coding. Firstly, we quantify the optical power requirements of 10 Gb/s point-to-point links using a variety of coding schemes using numerical simulations. Next we estimate the power consumption of the encoder and decoder for

*Correspondence author: yury.audzevich@cl.cam.ac.uk

GreenNet'11, August 19, 2011, Toronto, Ontario, Canada.

[1]Optical burst mode receivers should not be confused with optical burst switching networks.

the popular 8B10B coding scheme using an 90 nm CMOS ASIC synthesis flow for comparison with the optical power requirements. As the 8B10B code is logically simple, we use it as a demonstration even though it is not part of current 10 Gb/s standards. We show that the power consumption of the encoder/decoder is considerably higher than the reduction in optical power that can be achieved using 8B10B coding. Furthermore, we show that simple coding schemes can achieve the same DC balance performance as more sophisticated codes, demonstrating that there is scope for reducing power consumption in this area.

In Section 2, we describe the optical simulation, 8B10B Verilog design and ASIC flow in detail. Section 3 shows the results. In Section 4, we discuss the implications of these results for future coding schemes and outline plans for a physical and data link layer protocol optimised for low power operation of switched optical networks. Related work is discussed in Section 5. Finally, the conclusions are summarised in Section 6.

2. METHOD

2.1 Line Codes

A number of mBnB line codes[2] are in widespread use today. For example, 4B5B was popularised by fibre distributed data interface (FDDI), and was later adopted by 100BASE-TX [9]. The 8B10B code is widely used in optical 1 Gb/s Ethernet [10] as well as other standards such as Fiber Channel and PCI-Express. An alternative to 8B10B was proposed to the IEEE during the 10 Gigabit Ethernet standardisation process that maintained the features of the existing 8B10B code but used half the bandwidth [3] but this alternative code was never adopted. Widmer has also proposed 7B/8B, 9B/10B partitioned DC balanced codes that can be used with 3B4B and 5B6B to generate 12B/14B, 17B/20B, and 16B/18B transmission codes [23] although there are alternative novel codes with specific properties for optical communications in [20] [6].

An alternative to an mBnB line code is to use a scrambler to generate a pseudo-random sequence. A hybrid block-codec and scrambler approach has been adopted for 10 Gigabit Ethernet, where 64 data bits are scrambled using the polynomial $x^{58} + x^{39} + 1$. Two control bits are added to the 64 data bits to indicate whether or not that particular 64 bit block contains data or control codes. The 2 bits every 64 also guarantee that a transition will occur every 66 bits [11].

2.2 Optical Simulations

A 10 Gb/s optical transmission link was simulated using a semi-analytical technique. The simulation setup is shown in Figure 1. A pseudo-random sequence (PRBS) of 2^{19} bits was used as the input to an encoder whose output was used to generate an electrical signal waveform with 64 samples per binary bit. This gave a spectral resolution of approximately 10–20 KHz depending on the extra redundant bits that were added in the line coding process.

Table 1 lists the line codes investigated for this section. Each line code was coded using the look-up tables found in the the relevant standard. Each encoder had its disparity and maximum run length verified by counting the individual

[2]An mBnB line code is one that maps blocks of length m bits into representations of length n.

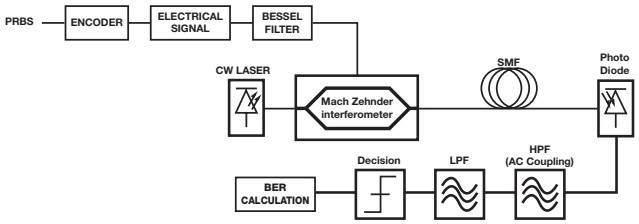

Figure 1: Main components of the simulated optical link

bits. The same PRBS sequence was used as the input to each line coder and the transmission baud rate was adjusted to be in proportion to the redundancy added by the encoder, thereby maintaining the original bit rate. This waveform was then filtered using a 5th order Bessel filter with a cut-off frequency of 80% of the bit rate to simulate the driver circuits electrical bandwidth. The electrical signal was used to drive a Mach Zehnder interferometer in push-pull configuration to intensity modulate a continuous wave Laser at 1550nm. The optical signal was transmitted through 100 metres of standard single mode fibre (SMF) with fibre parameters chosen to meet the 10 Gb/s Ethernet standard over SMF as set out in section three of the 802.3, 2008 IEEE Ethernet standards document.

The optical receiver used an ideal direct detector where thermal noise is assumed to dominate. The AC coupling is achieved using a high pass filter (HPF), whose cut-off frequency is varied in the simulations. The electrical bandwidth of the receiver is simulated using Bessel low-pass filter with the same characteristics as the transmitter. The optimal sampling point was set at half the bit period and the bit error rate (BER) is calculated for a range of received optical power using the complementary error function:

$$BER = \frac{1}{4}\left[erfc\left(\frac{I_1 - I_D}{\sigma_T\sqrt{2}}\right) + erfc\left(\frac{I_0 - I_D}{\sigma_T\sqrt{2}}\right)\right] \quad (1)$$

where I_1 and I_0 are the sampled photo detector currents for detected 1 and 0, I_D is the decision threshold and σ_T is the RMS thermal noise current. The decision threshold was set at the mean level of the output of the HPF used for the AC coupler.

2.3 8B10B Encoder/Decoder

8B10B code [24] belongs to the class of parity-disparity DC-balanced codes. 8-bit symbols coming from the data buffer over a parallel bus are encoded in 10-bit words, while consolidating 5B6B and 3B4B schemes.

To avoid ambiguity at the receiver side and keep the required density of zeros and ones for clock recovery, the encoder must send specific control words over the link even in the absence of Media Access Control (MAC) data. The ordered sets, *IDLE1* and *IDLE2*, present the combinations of two code groups and are usually transmitted to control the running disparity of the stream and keep the optical transceivers of the receiving circuitry alive. In order to obtain power characteristics of the 8B10B encoder and decoder, the behavioural-level design was implemented in Verilog HDL. All the combinatorial and sequential logic used for 8-bit sequence classification, disparity control, encoding and decoding was extracted from the original paper [24].

Table 1: Line codes used in the simulations

	1B2B	3B4B	4B5B	7B8B	8B10B	9B10B	NRZ-Scr.	64B66B	2^{19} PRBS
Overhead (%)	50	33	25	14.29	25	11.1	0	3.2	0
Line Rate(Gb/s)	20	13.33	12.5	11.43	12.5	11.1	10	10.3125	10
Max. run length	2	4	8	7	5	7		16	
Disparity / code	0	0	0 ± 2	0, ± 2, ± 4	0, ± 2	0, ± 1,± 2	NA	NA	NA

ASIC Synthesis. Synthesis of the behavioural-level design into a gate-level Verilog netlist was performed by means of Synopsys Design Vision (SDV) tool and 90nm CMOS UMC90 library, that provides definitions, logical descriptions and timing information of various logical gates (standard cells). Synopsys Design Compiler (SDC) translates behavioural-level design into a gate-level semantic while imposing specific design rules (defined in UMC90 library), optimization constrains (timing requirements for all design paths and area requirements) and environmental description (variations in voltage, temperature and modelling wire loads). The final result produced by SDC is optimized on the gate-level and mapped to the library cells of 90nm UMC technology process (saved as Verilog HDL netlist).

Verification. In order to verify the correctness of produced hardware during the synthesis process and its correspondence to the initial HDL description of 8B10B encoder/decoder we created a test bench that can instantiate our modules with custom input patterns. The tests comparing the behaviour of synthesised modules against results obtained from corresponding simulations were performed by means of Mentor Graphics ModelSim. Additionally, we used ModelSim to produce Value Charge Dump (VCD) files during simulation process to log switching activity needed for power estimation.

Power estimations. Integrated gate-level power estimation was performed by means of Synopsys PrimeTime PX suite that requires the netlist data, switching activity recorded by the VCD file and 90 nm UMC library specifications as input parameters. The total power consumed by a design (8B10B encoder/decoder) is classified into static and dynamic power groups, and is evaluated on the basis of the declared number of combinatorial and sequential elements composing the circuit. While the static or leakage power is considered to be constant over the entire simulation period (it is only dependent on voltage, temperature and state of transistors), the dynamic power tends to depend on the frequency of the logic state transitions. Further, the dynamic power analysis is performed on the cell basis. The internal power model of the CMOS cells is defined by UMC90 library, while the dynamic switching power is estimated as a result of frequent transitions on cells outputs. Latter power group classes are evaluated over the entire simulation period and can be represented as power waveforms.

3. RESULTS

3.1 Optical Simulation Results

Figure 2 shows the receiver sensitivity as a function of HPF cut-off frequency. The receiver sensitivity sets the minimum optical power requirement of the link. Figure 3 shows the bit error rate (BER) versus received power for a AC coupling HPF frequency of 100 MHz. It is striking, how little difference there is in receiver sensitivity between codes, with the exception of the high overhead 1B2B code. The consistent 2 dB lower receiver sensitivity of the the 4B5B line code can be explained by the fact that when only data symbols and no control symbols are sent, which was the case in these simulations, the code is unbalanced over time and reduces the power used to transmit a binary one in comparison to the average transmitted power.

If we take 100 MHz as a typical HPF cut-off frequency and assume that the optical power will be dominated by the laser power requirement, then for a 20 dB link budget, we require -0.46 dBm (0.9 mW) input power for the uncoded signal compared with -2.2 dBm (0.6 mW) for 8B10B coding. Thus, we save 0.3 mW of optical power using coding. The assumption of the dominance of laser power is justified by the fact that new silicon photonic components are offering very low capacitance modulators [19] and detectors [15] which have drive requirements on the order of fJ/bit.

3.2 Power Estimates of 8B10B Coding Circuit

It is proposed that the power consumption of implemented 8B10B encoder and decoder are highly workload dependent. In order to verify this statement we created a test bench in Verilog HDL that used a set of realistic traffic patterns for power consumption estimation over the entire operation period. Ethernet trace files captured at 10 Gb/s were obtained from monitoring locations within the JANET network. To present these files into a bit stream understandable for the encoder, we used a custom trace processing program to insert packets delimiters as part of the conversion from bytes to a stream of symbols. According to the Gigabit Ethernet standard (IEEE 802.3-2008) at least five specific symbols should be inserted to separate the consecutive packets. These symbols represent the start of packet (**S**), the end of packet (**T**), the carrier extend (**R**) and at least one idle ordered set that is given as a combination of two code symbols. As an additional experiment we performed power estimations for the case when there is no MAC data present in the system. Due to the properties of 8B10B encoding scheme it is expected that during the idle periods of operation the specific coded ordered sets are generated and sent over the link for receiver synchronization. Specifically, we created an idle operation pattern that was using only IDLE1 (K28.5 & D5.6 code-words) and IDLE2 (K28.5 & D16.2 code-words) ordered sets as an encoder inputs. To test 8B10B decoder power consumption, the module was placed right after the encoder. Each of these simulations was performed for a 30 microseconds period using an internal encoding clock frequency of 1.25GHz.

Table 2 presents the power estimates for encoder and decoder modules. The power consumption of both encoder and decoder is significantly higher for the idle state operation. Considerable difference in power consumption can be explained by the fact that IDLE ordered sets cause the disparity control check every octet involving extra logic for eval-

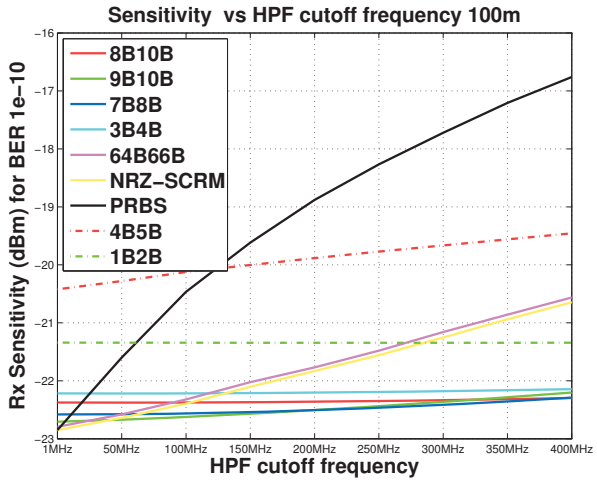

Figure 2: Receiver sensitivity for a BER of 10^{-10} at 10 Gb/s versus AC coupling cut-off frequency

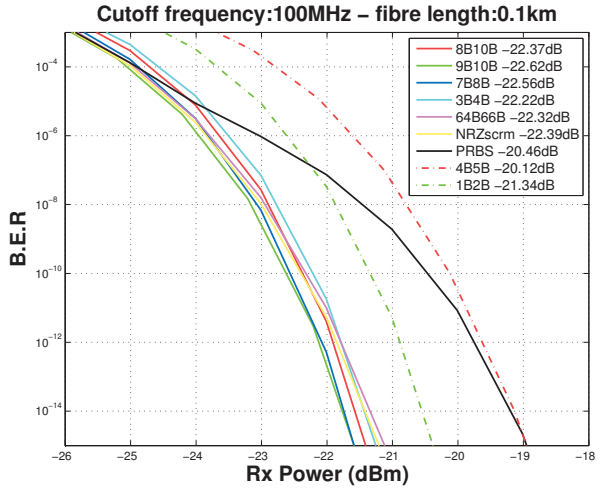

Figure 3: Bit error rate versus received power for an AC coupling cut-off frequency of 100 MHz

uations, whereas the data sequences may use more balanced words than unbalanced, totalling with lower power consumption. The second observation is related to the asymmetry in power consumption of the encoder and decoder. The energy advantage of encoder in this case is due to the logics' complexity, since the encoder contains extra functionality for disparity control of sub-blocks.

Despite the fact that the estimates of optical power obtained in the previous section are very dependent on the AC coupling frequency and link budget, it is clear that the power consumption of the coding circuits dominates over the optical power. We believe that this result will still hold for the case when coding circuits are implemented in a more recent, low leakage CMOS process.

4. IMPLICATIONS

In this section we show how even a modest change in the packet format and MAC protocol will make a significant impact on the power utilization. Drawing from the power consumption figures described in Section 3.2 it is clear how energy efficient shutting down the coding block when unused may be. We propose the combination of reintroducing a preamble for each Ethernet frame, and the complete shutdown of the coding/decoding and transmission subsystems as an effective mechanism for saving energy. The remainder of this section discusses the implications of such changes to a MAC and illustrates the significant saving that this proposal would make.

Prior to the deprecation of the 10BASE5 and 10BASE2 standards, a preamble of 64 bit clocks was provided to permit receiving systems to synchronize timing circuits and to reliably identify the beginning of the received frame. We would promote the use of any appropriate preamble that would permit the speedy recovery of clock information and the robust identification of the start of each frame.

In the case of a scrambled code this may simply be a regular sequence of 1 and 0, in the block codec we have used in the previous section a preamble might consist of a sequence of idle symbols: coded as they are with good DC balance and optimal 1/0 transition rates. Drawing on work from the domain of asynchronous logic design, a preamble no larger than the original Ethernet preamble (64 bit times) would provide ample material to permit the re-establishment of both bit-clock and symbol-clock. Table 3 summarizes possible preamble overheads as a function of Ethernet frame sizes. However, the burst mode receiver described in [22] uses only 32 bit periods whereas novel techniques such as injection locking can reduce clock synchronisation to a single bit period [13]. These techniques can reduce the transmission energy and latency while increasing throughput and enabling optical switching.

To compute the advantage in power consumption of the novel MAC algorithm, we analysed a trace file that was obtained from the monitoring location of JANET network and contained the data captured during a 24 hour period. Simple evaluations showed that our data-set link utilization average is 8.79%, which is within 5-10% and consistent with other experiences of data-center network utilization [5]. Based on this result and power estimate obtained for the 8B10B encoder and decoder in section 3.2, we evaluated energy spent by these modules within a 24 hour period for two variations of the MAC protocol.

Table 3: Ethernet frame overhead with synchronization preamble

Frame size	Preamble size	Overhead %
1518bytes	64bits	5.27×10^{-3}
1150bytes	64bits	6.9×10^{-3}
64bytes	64bits	0.125
1518bytes	128bits	1.05×10^{-2}
1150bytes	128bits	1.4×10^{-2}
64bytes	128bits	0.25

Table 4 lists the advantage in power consumption for both MAC protocols, while assuming 1150 bytes Ethernet frames are used and synchronization preamble is set to 64 bits. It is important to note that significant power savings of the second approach are achieved due to a complete shutdown of the encoding/decoding circuitry during periods of inactivity.

In Table 4 the column entitled *Total* refers to a accumu-

Table 2: Power estimates for 8B10B encoder/decoder

Patterns	Group	Encoder power (W)	Decoder power (W)
IDLE	Average:	2.62×10^{-3} (100.00%)	1.192×10^{-3} (100.00%)
	Peak:	0.0529	0.0298
	Static (Leakage):	5.21×10^{-6} (0.20%)	3.60×10^{-6} (0.19%)
	Dynamic (Net Switching):	7.0×10^{-4} (26.75%)	5.89×10^{-4} (30.69%)
	(Cell Internal):	1.91×10^{-3} (73.05%)	1.33×10^{-3} (69.13%)
20 PACKETS	Average:	1.93×10^{-3} (100.00%)	1.46×10^{-3} (100.00%)
	Peak:	0.0607	0.0305
	Static (Leakage):	5.21×10^{-6} (0.27%)	3.56×10^{-6} (0.24%)
	Dynamic (Net Switching):	4.64×10^{-4} (23.98%)	4.14×10^{-4} (28.37%)
	(Cell Internal):	1.47×10^{-3} (75.75%)	1.04×10^{-3} (71.38%)

lated total of consumption for the coder and decoder in each case: one based upon 8B10B and one based upon our revised Media Access Control protocol. While an idealised value, at 90.55 mW in the original scheme versus 7.21 mW after revision, it is clear the potential savings are noteworthy and significant.

Codes such as 8B10B are designed to perform other functions such as error detection and correction. However, it has been shown that these functions are not used in practice [1]. As a consequence, over-specification of the standard with such features simply wastes energy resources. As an alternative, we propose using a simpler, more energy-efficient error coding instead.

5. RELATED WORK

Bolla *et al.* have produced an extensive survey on green networking in which they classify the base concepts of the approaches taken to date into three main categories: re-engineering, dynamic adaption, sleeping and standby [17]. Falling into the third category is the initial ground-breaking research into the energy efficiency of the Internet and networked equipment carried out by Gupta *et al.* [7] in 2003 and by Christensen *et al.* [14] in 2004. These works (as this paper does) focused on the fact that networked devices use almost as much electricity when they are idle as they do when they are in use. Thus substantial energy saving can be made by changing the architecture and protocols of networked devices to allow the implementation of a low power or sleep mode applicable in idle periods.

The growing importance of network energy efficiency is outlined by the fact that there is an IEEE working group specifically looking at dynamic adaption mechanisms to save energy (802.3az, Energy Efficient Ethernet Task Force [12]). Gunaratne *et al.*show that show that 1Gbps Ethernet consumes about 4W more than 100Mbps while 10Gbs consumes tens of Watts [4]. They propose an Adaptive Link Rate (ALR) protocol as a means of reducing the energy consumption of a typical Ethernet link by adaptively varying the link data rate in response to traffic demands (as opposed to auto-negotiation when the connection is established which has been part of the standard for some time). Chen *et al.* [25] undertook the first studies into redesigning line codes for the energy efficiency of Ethernet. By increasing the number of symbols in a zero energy state it was shown that an 18% saving could be made in the energy consumption of an Ethernet NIC.

6. CONCLUSION

This work has presented some important and exciting conclusions when energy-efficiency becomes a *1st order* consideration in design and engineering.

We initially consider what the range of power - utilizations might be possible if we varied the scheme for physical line coding of data onto a 10 Gb/s service. Our conclusion is that while the range of power-utilization is not trivial, a range of up to 2.5 dBm, such power savings in the transmitter itself are not sufficient in most cases to reconsider the coding on this basis alone. However, Section 3.2 allow a much stronger recommendation. A consideration of the coding circuitry shows that it consumes significant power and mechanisms to reduce unnecessary use can provide significant savings in power utilization. Furthermore, we showed that the idle stream commonly used to maintain transmitter-receiver synchronization can consume more power than a non-idle data stream.

Two conclusions immediately present themselves: firstly, operational practice of leaving unused (idle) links permanently *up* uses more power than a link in-use. Secondly, we are motivated to consider a modified Media Access Control protocol.

Our modified MAC, discussed in Section 4, uses the power utilization values to dramatically demonstrate the power saving by using a coding schema that permits a default-off transmission system through the addition of a preamble. For a link utilization of 8.8% and a modest preamble overhead amounting to 7×10^{-3}%; we can have the transmission system (including codecs) inactive for nearly 22 hours in 24. In real terms this made a comparative reduction in power of 93%. While we admit this is an idealised value it is a clear indicator of the scope of our approach. The proposed approach is fully consistent with future low-power switched optical networks.

Future Work

It is clear this paper presents one of a first few steps along the road of reconsidering how best to send data in an energy efficient manner. The scrambler-hybrid schema utilized in the current 10 Gb/s Ethernet standard is an obvious next-step for further investigation. We do not consider how our ideas might be incorporated into existing standards for Energy Efficient Ethernet but there is a clear future for such consideration.

We plan implementation-based comparisons of our modi-

Table 4: Effective work of encoder/decoder 24h - idealised

Trace period	Util.%	Time busy(h)	Time idle(h)	Work(W·24h) Encoder	Work(W·24h) Decoder	Total(mW)
24h as is	8.796	2.11	21.89	Idle: 5.73×10^{-2} Busy: 4.072×10^{-3}	Idle: 2.61×10^{-2} Busy: 3.08×10^{-3}	90.55
24h, 64bit preamble	8.796	2.11	1.46×10^{-2}	Idle: 3.85×10^{-5} Busy: 4.072×10^{-3}	Idle: 1.75×10^{-5} Busy: 3.08×10^{-3}	7.21

fied physical coding layers using a fully programmable hardware subsystem, such as that provided by the NetFPGA [3].

Acknowledgements

Our thanks to Stelios Timotheou, Adrian Wonfor and Charalampos Rotsos for their thoughtful comments and technical input throughout this work. This work was supported in part by the EPSRC INTelligent Energy awaRe NETworks (INTERNET) Project and an EPSRC Career Acceleration Research Fellowship award to Philip Watts.

7. REFERENCES

[1] A. W. Moore et al. Optical network packet error rate due to physical layer coding. *IEEE/OSA Journal of Lightwave Technology*, 23(10):3056–3065, 2005.

[2] L. A. Barroso and U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33–37, 2007.

[3] C. G. Lee et al. A new line code for 10-Gigabit Ethernet: MB810. In *Proceedings of the IEEE ICC*, volume 3, pages 1774–1777, New Orleans, USA, 2000.

[4] C. Gunaratne et al. Reducing the energy consumption of Ethernet with Adaptive Link Rate (ALR). *IEEE Transactions on Computers*, 57:448–461, 2008.

[5] D. Abts et al. Energy proportional datacenter networks. In *Proceedings of ACM ISCA*, pages 338–347, Saint-Malo, France, 2010.

[6] F. Effenberger et al. A 9b10b line code for 2.5Gb/s upstream PONs. In *Proceedings of the OFC*, pages 1–3, San Diego, USA, 2009.

[7] M. Gupta and S. Singh. Greening of the Internet. In *Proceedings of the ACM SIGCOMM*, pages 19–26, Karlsruhe, Germany, 2003.

[8] I. H. White et al. Scalable optical switches for computing applications. *OSA Journal of Optical Networking*, 8(2):215–224, 2009.

[9] IEEE. IEEE 802.3u — Fast Ethernet at 100Mbit/s, 1995. Standard.

[10] IEEE. IEEE 802.3z — Gigabit Ethernet, 1998. Standard.

[11] IEEE. IEEE 802.3ae — 10 Gb/s Ethernet, 2002. Standard.

[12] IEEE. 802.3az Energy Efficient Ethernet, meeting materials. Technical report, January 2008.

[13] L. Jri and L. Mingchung. A 20Gb/s burst-mode CDR circuit using injection-locking technique. In *IEEE ISSCC*, pages 46–586, San Francisco, USA, 2007.

[14] Ken Christensen et al. Power management in networked devices. *IEEE Computer*, 37:91–93, 2004.

[15] M. R. Reshotko et al. Waveguide Coupled Ge-on-Oxide Photodetectors for Integrated Optical Links. In *Proceedings of the IEEE Group IV Photonics*, pages 182–184, Sorrento, Italy, 2008.

[16] Y. Ota and R. Swartz. Burst-mode compatible optical receiver with a large dynamic range. *IEEE/OSA Lightwave Technology*, 8(12):1897–1903, 1990.

[17] R. Bolla et al. Energy efficiency in the future Internet: A survey of existing approaches and trends in energy-aware fixed network infrastructures. *IEEE Communications Surveys Tutorials*, 13(2):223 –244, 2010.

[18] R. Luijten et al. Viable opto-electronic HPC interconnect fabrics. In *Proceedings of the ACM/IEEE SC*, pages 18–27, Washington, USA, 2005.

[19] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson. Silicon optical modulators. *Nature Photonics*, 4(8):518–526, 2010.

[20] M. Reinecke and K. Hansen. Advanced 3B4B channel coding for low error-rate optical links at 2.488 Gb/s. In *Proceedings of the IEEE ITCC*, pages 524–528, Las Vegas , USA, 2001.

[21] C. Su, L. Chen, and K. Cheung. BER performance of digital optical burst-mode receiver in TDMA all optical multiaccess network. *IEEE Photonics Technology Letters*, 7(1):132–134, 1995.

[22] B. C. Thomsen, B. J. Puttnam, and P. Bayvel. Optically equalized 10 Gb/s NRZ digital burst-mode receiver for dynamic optical networks. *OSA Optics Express*, 15(15):9520–9526, 2007.

[23] A. X. Widmer. DC balanced 7B/8B, 9B/10B, and partitioned DC balanced 12B/14B, 17B/20B, and 16B/18B transmission codes. *U.S. Patent 6614369 B1*, 2003.

[24] A. X. Widmer and P. A. Franaszek. A DC-balanced, partitioned-block, 8B/10B transmission code. *IBM Journal of Research and Development*, 27:440–451, 1983.

[25] Yanpei Chen et al. Energy efficient Ethernet encodings. In *Proceedings of the IEEE LCN*, pages 122–129, Montreal, Canada, 2008.

[3] www.netfpga.org