

**Dieses Dokument ist eine Zweitveröffentlichung (Verlagsversion) /
This is a self-archiving document (published version):**

Ronny Henker, Jan Pliva, Mahdi Khafaji, Frank Ellinger, Thomas Toifl, Bert Offrein, Alessandro Cevrero, Ilter Oezkaya, Marc Seifried, Nikolay Ledentsov, Joerg-R. Kropp, Vitaly Shchukin, Martin Zoldak, Leos Halmo, Jaroslaw Turkiewicz, Wyn Meredith, Iain Eddie, Michael Georgiades, Savvas Charalambides, Jeroen Duis, Pieter van Leeuwen

Adaptive optical interconnects: the ADDAPT project

Erstveröffentlichung in / First published in:

XXXVI Symposium on Photonics Applications in Astronomy, Communications, Industry, and High-Energy Physics Experiments. Wilga, 2015. Bellingham: SPIE, Vol. 9662 [Zugriff am: 23.05.2019].

DOI: <https://doi.org/10.1117/12.2202564>

Diese Version ist verfügbar / This version is available on:

<https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-347931>

„Dieser Beitrag ist mit Zustimmung des Rechteinhabers aufgrund einer (DFGgeförderten) Allianz- bzw. Nationallizenz frei zugänglich.“

This publication is openly accessible with the permission of the copyright owner. The permission is granted within a nationwide license, supported by the German Research Foundation (abbr. in German DFG).

www.nationallizenzen.de/

PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

Adaptive optical interconnects: the ADDAPT project

Ronny Henker, Jan Pliva, Mahdi Khafaji, Frank Ellinger, Thomas Toifl, et al.

Ronny Henker, Jan Pliva, Mahdi Khafaji, Frank Ellinger, Thomas Toifl, Bert Offrein, Alessandro Cevrero, Ilter Oezkaya, Marc Seifried, Nikolay Ledentsov, Joerg-R. Kropp, Vitaly Shchukin, Martin Zoldak, Leos Halmo, Jaroslaw Turkiewicz, Wyn Meredith, Iain Eddie, Michael Georgiades, Savvas Charalambides, Jeroen Duis, Pieter van Leeuwen, "Adaptive optical interconnects: the ADDAPT project," Proc. SPIE 9662, Photonics Applications in Astronomy, Communications, Industry, and High-Energy Physics Experiments 2015, 966207 (11 September 2015); doi: 10.1117/12.2202564

Event: XXXVI Symposium on Photonics Applications in Astronomy, Communications, Industry, and High-Energy Physics Experiments (Wilga 2015), 2015, Wilga, Poland

SPIE.

Adaptive Optical Interconnects – The ADDAPT Project

Ronny Henker^a, Jan Pliva^a, Mahdi Khafaji^a, Frank Ellinger^a, Thomas Toifl^b, Bert Offrein^b, Alessandro Cevrero^b, Ilter Oezkaya^b, Marc Seifried^b, Nikolay Ledentsov^c, Joerg-R. Kropp^c, Vitaly Shchukin^c, Martin Zoldak^d, Leos Halmo^d, Jaroslaw Turkiewicz^e, Wyn Meredith^f, Iain Eddie^f, Michael Georgiades^g, Savvas Charalambides^g, Jeroen Duis^h, Pieter van Leeuwen^h

^aTechnische Universität Dresden, Chair for Circuit Design and Network Theory, 01062 Dresden, Germany;

^bIBM Research GmbH, Zurich Research Laboratory, Säeumerstrasse 4, 8803 Rüschlikon, Switzerland;

^cVI-Systems GmbH, Hardenbergstrasse 7, Berlin 10623, Germany;

^dArgotech a.s., Holubova 978, Nachod 547 01, Czech Republic;

^eInstitute of Telecommunications, Warsaw University of Technology, ul. Nowowiejska 15/19, 00-661 Warsaw, Poland;

^fCompound Semiconductor Technologies, Maryhill Road G20 OSP Glasgow, United Kingdom;

^gPrimeTel PLC, Omonias 141, 3045 Limassol, Cyprus;

^hTE Connectivity, Rietveldweg 32, 5222AR s'Hertogenbosch, The Netherlands

ABSTRACT

Existing optical networks are driven by dynamic user and application demands but operate statically at their maximum performance. Thus, optical links do not offer much adaptability and are not very energy-efficient. In this paper a novel approach of implementing performance and power adaptivity from system down to optical device, electrical circuit and transistor level is proposed. Depending on the actual data load, the number of activated link paths and individual device parameters like bandwidth, clock rate, modulation format and gain are adapted to enable lowering the components supply power. This enables flexible energy-efficient optical transmission links which pave the way for massive reductions of CO₂ emission and operating costs in data center and high performance computing applications. Within the FP7 research project Adaptive Data and Power Aware Transceivers for Optical Communications (ADDAPT) dynamic high-speed energy-efficient transceiver subsystems are developed for short-range optical interconnects taking up new adaptive technologies and methods. The research of eight partners from industry, research and education spanning seven European countries includes the investigation of several adaptive control types and algorithms, the development of a full transceiver system, the design and fabrication of optical components and integrated circuits as well as the development of high-speed, low-loss packaging solutions. This paper describes and discusses the idea of ADDAPT and provides an overview about the latest research results in this field.

Keywords: Optical interconnects, adaptive optical links, energy efficiency, high-speed data transmission

1. INTRODUCTION

1.1 Optical Interconnects

Future communication networks, which include data centers, high performance computing (HPC), and core networks, demand a significant improvement of its performance and flexibility while the power consumption for their operation needs to be drastically reduced. As more and more data is generated and processed in the processor chips, the challenge is to transmit this huge amount of data outside the chips. Therefore, short-transmission links and server in-/outputs (I/Os) with a high accumulated data rate and low energy consumption

Further information: (Send correspondence to R. Henker)

E-mail: ronny.henker@tu-dresden.de

Web-site: www.addapt-fp7.eu

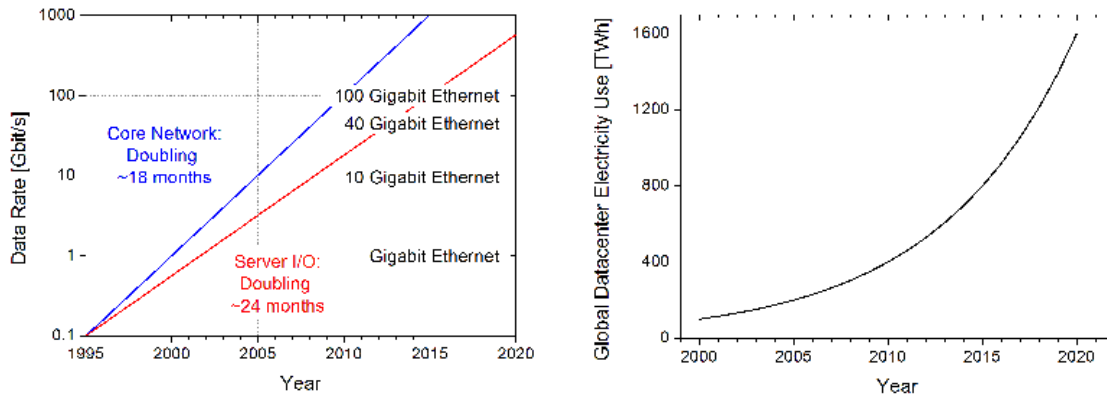


Figure 1. IEEE roadmap for computing network bandwidth growth⁵ and data center electricity use.⁶

are required. This can be provided by optical interconnects. Especially at high data rates and large link distances, optical links have lower losses and consume less power than copper links for distances above 10 cm at 28 Gbit/s, and for 1 m at 4 Gbit/s.¹

According to Moore's law silicon scaling leads to a doubling of the transistor density on a chip each two years pushing forward the computational power of processor integrated circuits (ICs). As a consequence, the speed in the interconnect network has to double each two years as well, as shown on the left hand side of Fig. 1. For example, the *Roadrunner* supercomputer with a 1 PF operational capacity uses approximately 50,000 interconnects each at 5 Gbit/s. To realize exascale computing capacity in 2020, billions of communication links at data rates of ≥ 25 Gbit/s are required.² However, with increasing compute power and interconnect throughput, the power consumption of data processing and storage systems becomes enormous. For example, in the United States (US) a power consumption of 100 TWh with associated energy costs of approximately 5.5 billion Euro were estimated for 2011.³ As can be seen on the right hand side of Fig. 1, data centers worldwide used electricity of more than 400 TWh in this year. It is assumed that the power consumption increases further exponentially to the fourfold in the next five years. In 2010, around 200 times more energy was required to transport than to process the data.⁴ The energy needed for a floating point operation amounted to roughly 0.05 - 0.1 pJ/bit, whereas the energy needed for (electronic) datatransport on a circuit card was in the range of 2 - 10 pJ/bit. Therefore, in a big multistage network with a 50 m diameter and multiple transceiver hops, the transport consumes around 1000 times more energy compared to the processing of the same data. Such a high power consumption comes with the several drawbacks: very high energy costs, environmental pollution, strong self-heating, and additional energy consumption for cooling. Hence, a massive reduction of the power consumption for data transport is mandatory in the high-speed networks and interconnects. Most links in today's optical communication systems are statically driven with its maximum performance and provide the maximum data speed required for data traffic peaks. Therefore, almost the same high power is consumed independently of the data load. This makes the links and networks inflexible which is in conflict with time variable user demands and data loads. Thus, introducing adaptivity to the optical interconnect where performance and power consumption are scaled with the actual work load enables a huge potential to decrease the power consumption of optical networks.

Significant research regarding performance and power adaptivity has been performed primarily in the microprocessor domain. It was shown that power can be saved in low speed conditions by a scaling of the supply voltage.⁷⁻⁹ Driven by supply limitations of batteries, adaptivity is also an important issue for mobile wireless communication systems.^{10,11} For example, power savings by a factor of 2 have been demonstrated for UMTS power amplifiers.¹² Several works on adaptivity in electrically wired links were completed, e.g. by the IEEE 802.3az Energy-Efficient Ethernet study group.¹³ Here, power is saved by a transmission protocol allowing the link to be temporarily put into an idle state with reduced supply power. Further electrical links connecting processors and DRAMs with power-on and off path capabilities were proposed as well.^{14,15}

Also for optical interconnects, first adaptivity studies were performed. For example, power saving strategies

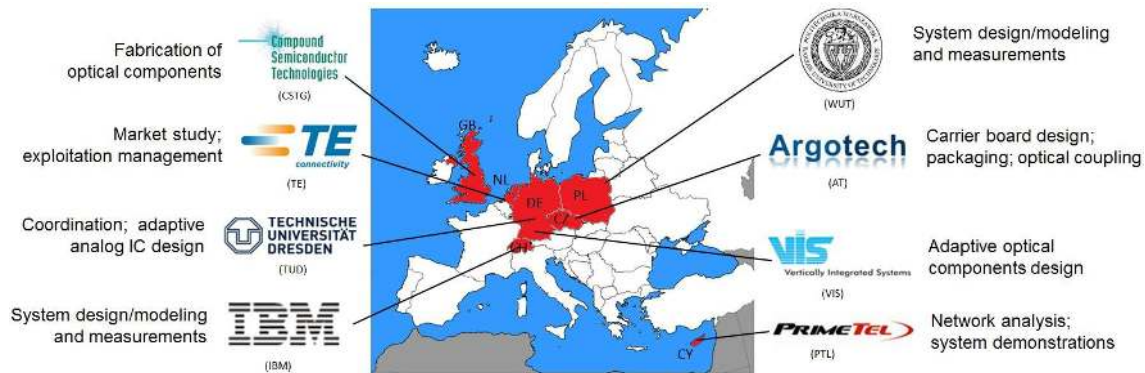


Figure 2. ADDAPT partners and tasks.

in the digital part of an OFDM based optical transceiver were achieved by bandwidth tuning of the OFDM signal and the associated sampling rate.¹⁶ However, just the switching of the digital modules is considered but the component parameters are fixed and are not adjusted. For large-range Telecom backbone networks, power is saved by switching off links with low data load and rerouting their data to other paths.¹⁷ The design space of power-aware opto-electronic networked systems is discussed just theoretically and by simulation on circuit and system level only with focus on voltage tuning.¹⁸

In this paper, the European Unions Seventh Framework Programme project *Adaptive Data and Power Aware Transceivers for Optical Communications (ADDAPT)* is described. The main focus of ADDAPT is to introduce a novel and comprehensive approach of power and performance adaptivity into optical interconnects for short-range data communication. Adaptivity from system down to optical device, electrical circuit and transistor level is proposed. Depending on the actual data load, the number of activated link paths and individual device parameters like bandwidth, clock rate, modulation format and gain are adapted at the same time to enable lowering the components' supply power.

1.2 The ADDAPT Project

The ADDAPT project aims to adjust the performance and in turn the power consumption of the multiple optical links from system down to optical device, electrical circuit and transistor level according to the actual required data load and link conditions. To achieve this, a high-speed electro-optical transceiver module is developed whose parameters like bandwidth, modulation format, clock rate, amplitudes can be adapted. This leads to a reconfiguration of the system according to the actual transmission requirement which in turn reduces the system power consumption. To realize this, a smart adaptivity control is implemented that decides how and when the system parameters need to or may change. The transceiver design includes novel high-speed directly modulated lasers and photodetectors equipped with low-loss optical coupling, novel adaptive integrated circuits like laserdiode driver (LDD), transimpedance/limiting amplifier (TIA/LA), clock data recoveries (CDR) in advanced 14 nm CMOS technology and high-speed low-loss packaging solutions using glass or ceramic substrates. A transceiver system with 4 link paths each with adaptive data rates from 7 Gbit/s up to 56 Gbit/s and maximum 10 m link distance is targeted. Further goals are low power consumption and high energy efficiency of the transceiver and its components as well as low latency data transmission. The development of such an adaptive optical interconnect paves the way to build flexible energy-efficient optical transmission links and networks coping with varying bitrate demands which enables massive reductions of CO2 emission and costs.

To achieve these project goals, the ADDAPT consortium involves a full supply chain from semiconductor technologies, component and system design over packaging, assembling and characterization to user requirements, interconnect applications and commercial markets. As shown in Fig. 2, complementary competences and tasks of eight consortium partners (3 large companies, 3 SMEs and 2 universities from seven European countries) including research and education institutions, device developers and manufacturers, suppliers of communication equipment as well as network operators are combined. Key applications of ADDAPT are seen in optical interconnects for



Figure 3. Traffic load (top), supply voltage (middle) and Tx bias current (bottom) of 10 Gbit/s link.

short-range data communication as used in data-centers or high performance computing (HPC) for rack-to-rack, server-to-server and board-to-board connections. One possibility would be to replace standard active optical cables (AOC) with fixed performance and power consumption by ADDAPTive transceivers.

2. ANALYSIS OF EXISTING DATA CENTER NETWORKS

Recently, various network analysis of data centers and their traffic characteristics were performed in order to optimize the data centers' operation with regard to network flow.¹⁹⁻²¹ Different network configurations, measurement times, data levels and measurement opportunities were used. Most of the studies focused on traffic patterns inside data centers, academic networks or specific nature of traffic like peer-to-peer traffic inside an ISPs network. In contrast to those measurements, the ADDAPT consortium partner PrimeTel, a network operator located in Cyprus, performed an analysis within its core network, to also investigate the potential and opportunities to introduce performance and power adaptivity to optical interconnects.²² Therefore, high-speed optical links in the optical backbone and core network, which carry real time traffic with data rates of 10 Gbit/s, were monitored over a period of one year. To identify different traffic patterns and to collect a variety of network statistics, the links were polled every 5 minutes. The main volume of data was collected via the Simple Network Management Protocol (SNMP). A variety of network statistics such as traffic load, transceivers supply voltage and temperature, Tx bias current and optical power as well as the Rx optical power were collected by monitoring the network devices. This allowed to gain an overview of the traffic load that traverses a core subset of PrimeTel's network and monitor the operation of the optical transceivers that transmit data over those links. In addition,

flow data were collected using the Cisco NetFlow protocol for collecting data time frame of seven hours each day over a three day period. The statistics of this analysis help to constructively exploit the ADDAPTive approaches.

As an example of the analysis the data traffic load, supply voltage and Tx bias current of a 10 Gbit/s link is shown in Fig. 3. In conclusion, the traffic load observed at Primetel's optical network, exhibited variations both between hours of a day as well as days of a week. Namely, the traffic loads diurnal pattern exhibited a maximum during the afternoon to early night hours and a minimum in the early morning hours. Moreover, the traffic load pattern that was observed during weekdays was different than that observed during the weekend. In addition, public (bank) holidays exhibited traffic patterns that were similar to that of the weekend rather than that of the weekdays. Nevertheless, although the traffic load exhibited fluctuations in different time granularities, the optical transceivers operating parameters such as the Tx bias current as well as the Rx/Tx power did not exhibit similar fluctuations. Based on these findings it can be assumed that the optical transceiver modules power consumption also does not present large fluctuations. Thus, this can be exploited by an adaptive approach such as the one developed within the frames of ADDAPT in order to reduce power consumption. In addition, average and maximum link utilization was found to be lower than 50 % for the vast majority of the monitored links. This gives additional exploitation opportunities for the approach envisioned in ADDAPT with the smart tuning of link parameters and lane on/off switching by adapting accordingly to the utilization requirements with the lowest power consumption possible. Finally, the NetFlow data analysis provides a characterization of the traffic that traverses modern optical networks which will be utilized in the testing, validation and development of the AOC.

3. ADAPTIVE OPTICAL INTERCONNECTS

3.1 ADDAPT concept

The performance requirements of future networks, especially optical links and interconnects, are not static and changes over time as the individual needs of the users, applications and boundary conditions lead to a strong varying network traffic behavior. To be energy-efficient, the networks and their system components have to be able to flexibly reset to those changes. Therefore, one of the main innovations treated by ADDAPT is to adjust the performance and in turn the power consumption of the multiple optical links to the actual required data load and link conditions. The complete ADDAPT concept is shown in Fig. 4.

The focus of ADDAPT is the development of a four lane transceiver module and its components. These components include vertical-cavity surface-emitting lasers (VCSEL), photodiodes (PD) and electrical circuitry such as laserdiode driver (LDD), transimpedance/limiting amplifier (TIA/LA) as well as clock-data-recovery (CDR). All those components are designed for high-speed data transmission up to 56 Gbit/s and lowest power consumption. However, the primary goal is to implement performance and energy adaptivity on component level as well as on system/module level. For this, two basic methods are investigated: First, smart switching of the link lanes and second, smart tuning of individual link paths and their components with regard to bandwidth, modulation format, clock rate and amplitude. This can be adjusted by the adaptivity supply variables, namely the bias currents and voltages of the components. Smart switching allows step-wise (coarse) adaption, whereas smart tuning allows continuous (fine) tuning of parameters.

While the adaption to different bitrates, power consumptions and performances will be realized by the several transceiver devices, a higher-level control system has to be implemented for the decision when and how the adaption will be performed. This control can be conducted by microcontrollers, DSPs or FPGAs and interface controllers assisted by software and protocols. The driving parameters for the different operating points of each transceiver component can be stored in look-up tables for instance. Using a data protocol enables to maintain data integrity with variable data rates and system parameters. Furthermore, it defines the data frame for improved RAS (reliability, availability, serviceability) aspects and forward error correction (FEC) if required. From an overall system point of view the following adaptivity categories are possible:

1. *Cognitive adaption based on beforehand known data load profiles*

Beforehand known load profiles are exploited to optimize the performance versus power trade-off. For this the evaluation of the network analysis (see Section 2) can be used to divide into load scenarios with high

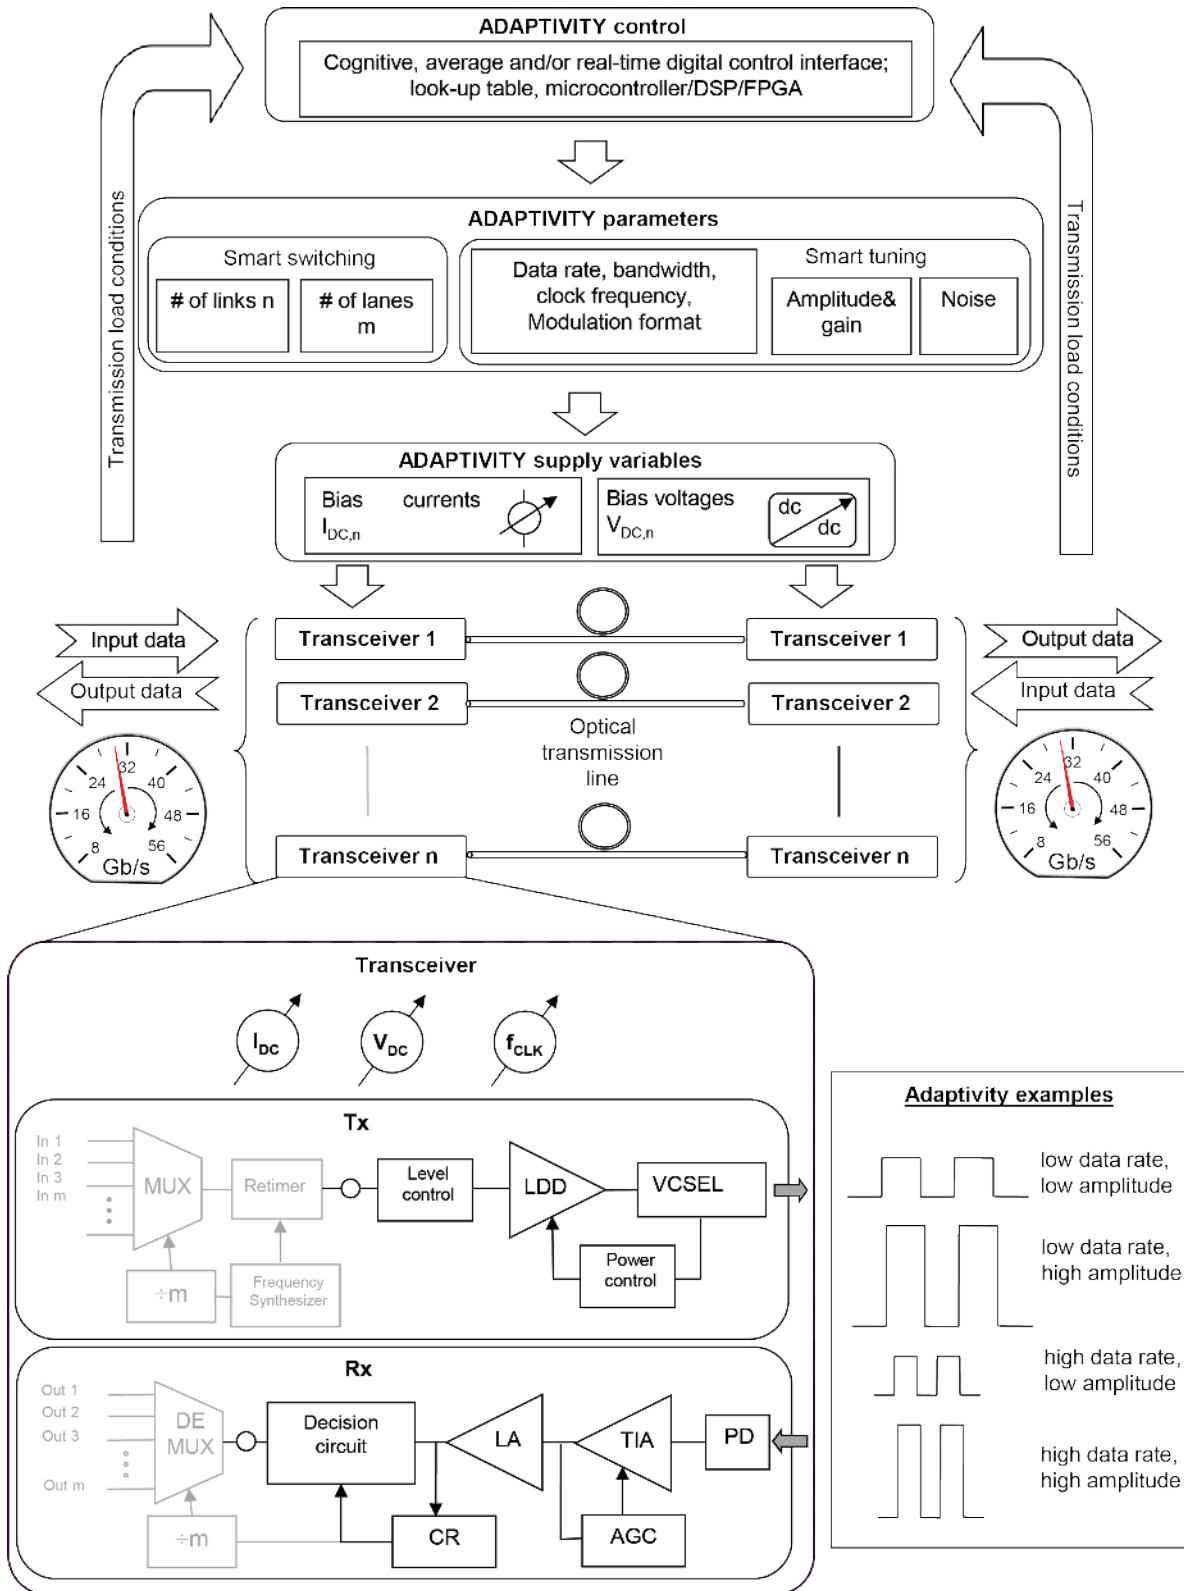


Figure 4. ADDAPT concept.

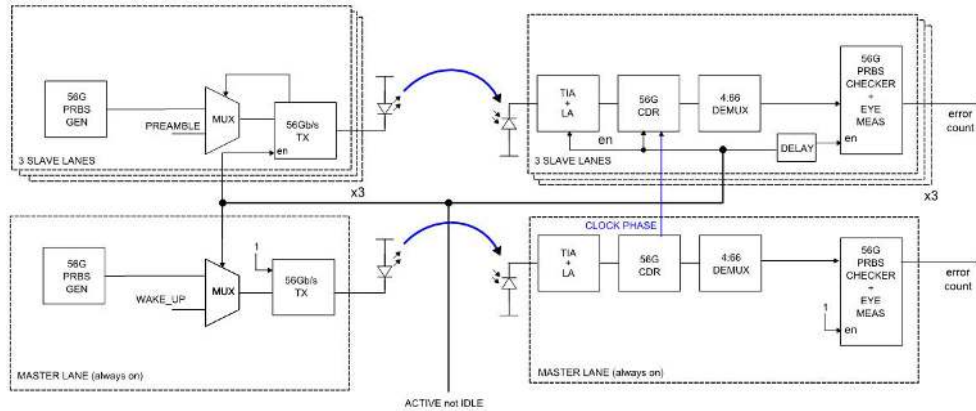


Figure 5. ADDAPT demonstrator.

(e.g. at daytime), medium and low (e.g. at night) performance requirements. Thus, within predefined time periods, the system is supposed to run in different operating profiles with specific performances and consequently adaptable supply powers. This approach is relatively simple to implement but will have a drastic impact on the power consumption of future links. This approach will mainly be implemented in ADDAPT.

2. *Real-time adaption considering the data load in a dynamic fashion*

Here, the actual data load is monitored and the system is adapted immediately. From a theoretical point of view this approach is very attractive since it may always enable ideal adaption. However, this approach is very challenging in practice since there may be data transmissions during adaption, where e.g. switching of components causes nonlinearities and signal degradations. First evaluations reveal very fast control latencies of <20 ns to make this approach attractive. Hence, speed and latency requirements for circuits and algorithms are very ambitious.

3. *Averaged adaption based on the data load averaged over time intervals*

Also this approach is based on recently monitored data load. However, to mitigate the problems associated with real-time demands, the adaption is performed after predefined time intervals to achieve a practical trade-off between data load adaptivity and other conflicting boundaries. This means that the system reacts in accordance to the load averaged over certain time intervals that are to be determined.

Although the decision and setting on how and when the system parameters need to or may change will be executed by a smart higher-level adaptivity control system including software, the tuning of the performance and power consumption has to be realized inside the devices and components. An envisioned demonstrator block diagram of the four lane ADDAPT transceiver is shown in Fig. 5. To the best of the authors' knowledge, such a comprehensive approach has not yet been implemented into optical interconnects.

3.2 Smart Adaptive Switching

Referring to Fig. 4, a communication system consists of several parallel transceiver link paths. All n links have to be activated at maximum required data throughput. If lower speed is sufficient, a certain number of paths can be switched off, e.g. by low loss DC switches, thereby massively saving power. As first order estimation, it can be assumed that the power saving due to path switching is given by $(\text{total number of paths}/\text{number of activated paths}) \times 0.8$. The latter factor takes into account the additional power needed for the system control and the losses of the required switches. Assuming the demonstrator with four paths yielding a maximum data rate of 224 Gbit/s, maximum 3 out of 4 paths can be switched off if e.g. a total data rate of 56 Gbit/s is sufficient. This results in a massive power saving by a factor of approximately 3 based on smart switching. Future systems with performance of tens of Tbit/s will have higher numbers of paths. Hence, for future ultra-high speed systems, the

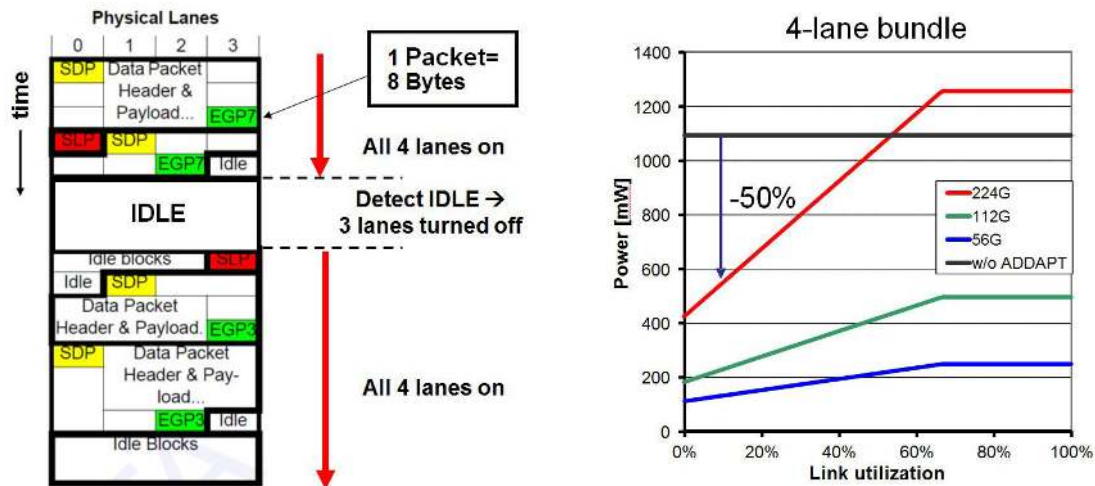


Figure 6. Data and Idle packets in a 4 x Infiniband protocol (left); power consumption in ADDAPT link as a function of link utilization for a link consisting of 4 lanes, assuming 50 % overhead from wake-up transmission (right).

potential for power saving will be tremendous. Optionally, for some system architectures, each link path may have m lanes generating further optimization potential by optimum lane switching.

The biggest challenge for the switching approach is to achieve fast switching times and low latencies especially when waking up a link path. During the wake-up time the incoming data stream can be stored in a FIFO in order to not lose any data. However, a big FIFO includes additional power consumption and latency to the system. Therefore, a wake-up time of <20 ns is targeted. Recently, a rapid on/off transceiver with 20 ns power-on time but with a data rate of only 7 Gbit/s has been presented.²³ Combining the low latency switching at high-speed data rates in ADDAPT would result in a significant breakthrough for HPC applications. The main challenge to achieve a rapid on/off switching is the fast locking of the CDR when a link path is reactivated. Therefore, the ADDAPT concept proposes the implementation of one master and three slave lanes, as shown in Fig. 5. If the link is in an idle mode, all the lanes except the master lane can be switched off to save power. The master lane always maintains phase locking of the link bundle and have an always ready connection to transmit wake-up command to the slave lanes. If active data is detected, the link is woken up by rapidly providing the maintained phase locking from the master to the slave lanes. The on/off switching itself can be realized by a detection of IDLE patterns within the protocol, e.g. a 4 x Infiniband protocol as shown on the left hand side of Fig. 6. The power that can be saved in the proposed ADDAPT switching concept depends on several factors, including link utilization, average package size, wake-up time, and the overhead power for the FIFO and the IDLE detection logic. On the right hand side of Fig. 6 a first conservative estimation of the power consumption for a link consisting of four lanes is shown. If link utilization is close to 0 %, the minimum consumed power is the power of a single optical lane plus the overhead for internal clock generation and distribution. For an average package size of 1 kByte (as indicated by the network analysis by PrimeTel), and a wake-up time of 20 ns at 56 Gbit/s link speed, the overhead value results to approximately 50 %. Typical link utilization values in data centers are way below 10 %.¹⁹ As can be seen in Fig.6, for a 10 % link utilization, this results in a power consumption saving of 50 % compared to a link without ADDAPT.

3.3 Smart Adaptive Tuning

The second type of optimization is based on smart tuning of links and the associated components. By using dynamic speed adaptation, the data rates are scaled down continuously or in case of ADDAPT by a constant factor (e.g. two or four). On the one hand, this approach is obviously most important for systems generally having only a single link path. But on the other hand, it is also relevant for complex systems with parallel link paths, where only one link path is sufficient in the moment and the other link paths are switched off, since the recently required data rate is very low (e.g. if non-time-critical software updates or just synchronisation and

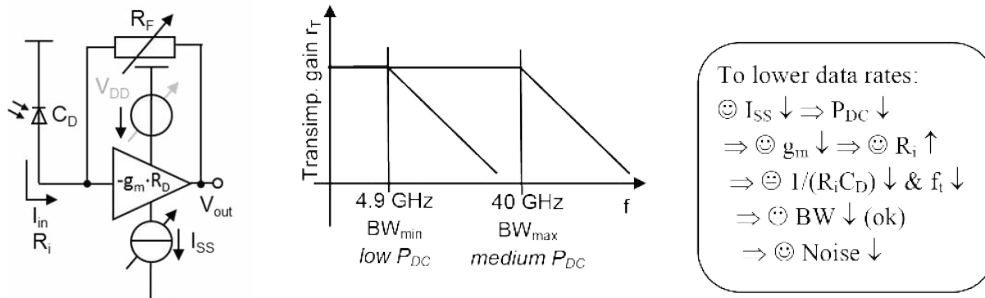


Figure 7. Adaptive tuning of a TIA; concept (left), frequency response (middle), functional chain (right).

simple control signals are transmitted). The key idea is to provide only the level of path performance that is needed by fine tuning of the individual link paths. The key adaptivity performance parameters which can be controlled to reach this goal are:

- bandwidth and clock frequency (and modulation format if appropriate),
- signal amplitudes and consequently also signal gains, and
- noise (this applies mainly to TIA and receiver circuits),

whereas the key adaptivity supply variables are:

- bias currents, and
- bias voltages.

Due to the V_{DD}^2 dependency of the circuits power is mostly saved by reducing the supply voltage. However, this is generated by an external voltage converter or regulation module which are rather slow (in microsecond to millisecond range) and may lead to further losses as there exists a trade-off between the converter speed and efficiency. Therefore, the voltage adaptation method is several orders of magnitude slower than the rapid on/off switching method. In contrast, tuning the bias currents is faster but the bandwidth and power scaling is linear and not as efficient as the voltage tuning. Therefore, the plan is to use both opportunities to realize a dynamic speed adaptation at a slow timescale responding to slowly varying load patterns of the network traffic (e.g. reducing network speeds during night hours and over the week-end).

The concept of the adaptive tuning of the circuit's data rates in steps of 56-28-14-7 Gbit/s is illustrated on the example of a TIA circuit in Fig. 7. If lower data rates are sufficient, the supply current I_{SS} can be reduced which obviously decreases the power consumption of the circuit. At the same time, the bandwidth (BW), which is mainly given by $1/(\text{amplifier input impedance } R_i \times \text{photodiode capacitance } C_D)$ decreases since $R_i \approx R_F/(1 + g_m R_D)$, g_m and f_t falls with I_{SS} . As a side benefit this decreases the noise which is proportional to BW and g_m . If there is any decrease of the effective transimpedance gain $r_T \approx R_F$, this may be compensated by increasing R_F leading to a further acceptable BW reduction. Multi-stage architectures are applied to optimizing bandwidth, transimpedance gain, noise and supply power simultaneously. By implementing several bias operating modes the circuits can be tuned (switched) between the different performance and power consumption modes. Similar approaches can be used to tune further transceiver components such as voltage and limiting amplifier, laserdiode driver, digital gates or oscillators for instance. The goal is 50 % power saving when tuning one path from 56 Gbit/s (maximum data rate per path) to 7 Gbit/s (minimum data rate per path).

Preliminary simulations on a basic TIA (in a BiCMOS technology) have been performed to evaluate the potential of the current tuning for the scaling of the bandwidth and power consumption. The simulation results are shown in Fig. 8. From the diagram it can be revealed that the power consumption reduces significantly when the bandwidth of the circuits is decreased by scaling down the supply current while the gain remains constant. By reducing the bandwidth to approximately 30 % of the peak bandwidth, a reduction of the power consumption by at least 50 % can be estimated.




Figure 8. Simulation of adaptive tuning for a single TIA stage.

4. CONCLUSIONS AND OUTLOOK

In this paper, the FP7 ADDAPT project is presented. It introduces novel methods for realizing high-speed energy-efficient adaptive optical interconnects. By applying adaptivity on system, components and transistor level the performance and power consumption of the optical link can be significantly reduced. A transceiver module with four link paths is developed and novel high-speed adaptive optical components as well as ICs are designed and high-speed low-loss packaging techniques are investigated. Techniques for using IDLE packet detection and rapid, low-latency on/off switching in combination with a master/slave configuration of an optical link with multiple lanes are introduced and enable a power saving factor of approximately 3 for a 4-lane link. Further energy reductions of at least 50 % can be achieved by tuning the performance, i.e. the bandwidth and signal amplitudes, of the individual link components by bias scaling while the components' output levels are held constant. Thus, for an optical link with four lanes an overall power saving factor of approximately 7 is assumed when switching from full speed of 4×56 Gbit/s to 1×7 Gbit/s. Such an energy reduction paves the way for massively decreasing both, operational costs and the carbon footprint of data centers. Promising application for such flexible high-speed low-power and low-latency optical interconnects can be found in data center communication and high performance computing. Within ADDAPT the first designs are finished and components are in fabrication which needs to be tested in order to verify the adaptivity concepts and methods.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Unions Seventh Framework Programme (FP7/2007-2013) under grant agreement № 619197.

REFERENCES

1. Saraswat, K., Cho, H., Kapur, P., and Koo, K.-H., "Performance comparison between copper, carbon nanotube, and optical interconnects," in [*IEEE International Symposium on Circuits and Systems (ISCAS)*], 2781–2784 (18-21 May 2008).
2. Benner, A., "Optical interconnects for HPC," in [*Short-Distance High-Density Optical Interconnects - An OIDA Roadmapping Workshop*], (12-13 April 2011).
3. EPA Energy Star, "Report to congress on server and data center energy efficiency," tech. rep., U.S. Environmental Protection Agency (2007).
4. Dorren, H., Duan, P., Raz, O., and Luitjen, R., "Fundamental bounds for photonic interconnects," in [*16th Opto-Electronics and Communications Conference (OECC)*], 470–472 (4-8 July 2011).
5. Ledentsov, N. N., Lott, J. A., Shchukin, V. A., Bimberg, D., Mutig, A., Germann, T. D., Kropp, J. R., Karachinsky, L. Y., Blokhin, S. A., and Nadochiy, A. M., "Optical components for very short reach applications at 40 Gb/s and beyond," in [*Proc. SPIE, Physics and Simulation of Optoelectronic Devices XVIII*], **7597**, 75971F–75971F–10 (2010).

6. Belady, C. L., "Projecting annual new datacenter construction market size," tech. rep., Microsoft Corporation, Global Foundation Services (2011).
7. Wu, Q., Juang, P., Martonosi, M., Peh, L.-S., and Clark, D., "Formal control techniques for power-performance management," *IEEE Micro* **25**, 52–62 (Sept 2005).
8. Burd, T. and Brodersen, R., "Design issues for dynamic voltage scaling," in [*Proceedings of the 2000 International Symposium on Low Power Electronics and Design (ISLPED)*], 9–14 (2000).
9. Qu, G., "What is the limit of energy saving by dynamic voltage scaling?," in [*IEEE/ACM International Conference on Computer Aided Design (ICCAD)*], 560–563 (Nov 2001).
10. Hassler, F., Ellinger, F., and Carls, J., "Analysis of buck-converters for efficiency enhancements in power amplifiers for wireless communication," in [*SBMO/IEEE MTT-S International Microwave and Optoelectronics Conference (IMOC)*], 616–620 (Oct 2007).
11. Wolf, R., Ellinger, F., and Eickhoff, R., "On the maximum efficiency of power amplifiers in ofdm broadcast systems with envelope following," in [*The 2nd International ICST Conference on Mobile Lightweight Wireless Systems (MobiLight)*], (2010).
12. Hassler, F., Ellinger, F., Jrges, U., Wolf, R., and Lindner, B., "A high-speed buck converter for efficiency enhancement of W-CDMA power amplifiers," *International Journal of Microwave and Wireless Technologies* **4**, 505–514 (10 2012).
13. Bennet, M., "Energy efficient ethernet overview," in [*ITU-T SG 15 Workshop*], (15 February 2008).
14. Leibowitz, B., Palmer, R., Poulton, J., Frans, Y., Li, S., Wilson, J., Bucher, M., Fuller, A., Eyles, J., Aleksic, M., Greer, T., and Nguyen, N., "A 4.3 GB/s mobile memory interface with power-efficient bandwidth scaling," *IEEE Journal of Solid-State Circuits* **45**, 889–898 (April 2010).
15. Zerbe, J., Daly, B., Dettloff, W., Stone, T., Stonecypher, W., Venkatesan, P., Prabhu, K., Su, B., Ren, J., Tsang, B., Leibowitz, B., Dunwell, D., Carusone, A., and Eble, J., "A 5.6Gb/s 2.4mW/Gb/s bidirectional link with 8ns power-on," in [*Symposium on VLSI Circuits (VLSIC)*], 82–83 (June 2011).
16. Zhang, J., Hu, J., Qian, D., and Wang, T., "Energy efficient OFDM transceiver design based on traffic tracking and adaptive bandwidth adjustment," *Opt. Express* **19**, B983–B988 (Dec 2011).
17. Idzikowski, F., Orłowski, S., Raack, C., Woesner, H., and Wolisz, A., "Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios," in [*14th Conference on Optical Network Design and Modeling (ONDM)*], 1–6 (Feb 2010).
18. Chen, X., Peh, L.-S., Wei, G.-Y., Huang, Y.-K., and Prucnal, P., "Exploring the design space of power-aware opto-electronic networked systems," in [*11th International Symposium on High-Performance Computer Architecture (HPCA)*], 120–131 (Feb 2005).
19. Benson, T., Akella, A., and Maltz, D. A., "Network traffic characteristics of data centers in the wild," in [*Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*], IMC '10, 267–280, ACM, New York, NY, USA (2010).
20. Kandula, S., Sengupta, S., Greenberg, A., Patel, P., and Chaiken, R., "The nature of data center traffic: Measurements & analysis," in [*Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*], IMC '09, 202–208, ACM, New York, NY, USA (2009).
21. Murray, D. and Koziniec, T., "The state of enterprise network traffic in 2012," in [*18th Asia-Pacific Conference on Communications (APCC)*], 179–184 (Oct 2012).
22. Charalambides, S., Georgiades, M., Dimosthenous, G., and Christofi, D., "Investigating an ISPs core network in 2014," in [*submitted to 10th International Conference on Broadband and Wireless Computing, Communication and Applications (BWCCA)*], (4-6 November 2015).
23. Anand, T., Talegaonkar, M., Elkholy, A., Saxena, S., Elshazly, A., and Hanumolu, P., "A 7Gb/s rapid on/off embedded-clock serial-link transceiver with 20ns power-on time, 740 μ m off-state power for energy-proportional links in 65nm CMOS," in [*Solid-State Circuits Conference - (ISSCC), 2015 IEEE International*], 1–3 (Feb 2015).