

# Pathways to Servers of the Future

## Highly Adaptive Energy Efficient Computing (HAEC)

Gerhard Fettweis<sup>1</sup>, Wolfgang Nagel<sup>2</sup>, Wolfgang Lehner<sup>3</sup>

<sup>1</sup>Vodafone Chair Mobile Communications Systems, <sup>2</sup>Chair for Computer Architecture, <sup>3</sup>Database Technology Group  
Technische Universität Dresden  
01062 Dresden, Germany  
{gerhard.fettweis, wolfgang.nagel, wolfgang.lehner}@tu-dresden.de

**Abstract**—The Special Session on “Pathways to Servers of the Future” outlines a new research program set up at Technische Universität Dresden addressing the increasing energy demand of global internet usage and the resulting ecological impact of it. The program pursues a novel holistic approach that considers hardware as well as software adaptivity to significantly increase energy efficiency, while suitably addressing application demands. The session presents the research challenges and industry perspective.

**Keywords**—energy efficiency; interconnects; software architecture; data center; server; computing; optical; wireless; adaptivity

### I. MOTIVATION AND GOALS

*Dear reader - when does your bad conscience turn worse with respect to energy consumption and carbon dioxide pollution - if you are traveling by aircraft or if you browse through the Internet?*

It may be surprising, however, fact is that since 2007, the energy consumption and corresponding carbon dioxide pollution of the data servers and networks required for global internet is higher than that of the worldwide air traffic [1]. Considering the strong increase of the required data rates and number of users, this trend will become even more significant in the future [2, 3]. Today, global internet and web applications such as upload and download of video content (video streaming), exchange of large files, and Web 2.0 applications require enormous computing and networking resources. As illustrated in Figure 1, the total amount of internet traffic is estimated to grow about 34 % per year until 2014 to more than a total of 63 exabytes (EB) per month. Among these, different types of video on demand (internet-based high definition television, internet video, and peer-to-peer video) will account for approximately 50% of total IP traffic in 2011.

This massive growth of internet traffic heavily correlates with a growth in total energy consumption both in the network as well as at the edge of the network. More network links and intermediate content distribution servers are needed to keep up with the internet traffic growth caused by video on demand and other internet and web applications. While these applications offer new means of individual entertainment, collaboration, and interactivity, they increase total energy consumption in the

same way. Server farms will grow, consuming energy for operation and cooling. This means that energy efficiency is not only an important goal but an absolute necessity at all levels, particularly within network and content servers. Many other high-end computing scenarios yield similar characteristics and requirements.




Figure 1. Cisco’s Global IP Traffic Forecast 2010-2015, [4]

While energy efficiency has already been considered at certain levels of the above mentioned scenario, such as energy-aware optimization of routing algorithms, what is needed is a comprehensive approach that takes the overall applications’ demand and users’ context as well as the particular aspects of computing resources (software and hardware) into account. At the moment, optimizations are considered on a component level, mostly in isolation, and do not take the demands of applications, user communities, and contexts into account in an integral holistic manner.

Exactly this holistic approach is taken by the newly started Collaborative Research Center 912 HAEC – “Highly Adaptive Energy-Efficient Computing” of Technische Universität Dresden. It plans to address energy efficiency at all system levels ranging from lowest levels of interconnection technologies to highest levels of software engineering

This work is supported by the German Research Foundation (DFG) in the Collaborative Research Center 912 “Highly Adaptive Energy-Efficient Computing”.

approaches. The guiding principles are simple and uniform:

- 1) Ensure flexibility at all levels that allows to provide alternatives tailored to applications and situations.
- 2) Establish – mostly off-line – the added value (utility) of spending energy in one or the other alternative.
- 3) Monitor current energy levels and other activity.
- 4) Take action at run time, based on the monitoring results and the available alternatives.

The development of concepts for a computer system based on a holistic approach of energy adaptivity can only be achieved by considering hard- and software aspects together with their inter-relations. HAEC will perform basic research enabling highly adaptive energy-efficient computing, incorporating an energy-adaptive software architecture, based on a high-performance energy-adaptive computing architecture. The goal is to provide solutions for scalable, energy-adaptive computing and content distribution for the next decades.

The research is structured in two interconnected parts (see Figure 2): part A: being responsible for the hardware technological activities, part B: for the software activities.




Figure 2. Overview of HAEC

## II. PART A: ENERGY-ADAPTIVE HIGH-SPEED COMPUTING PLATFORM (HAEC BOX)

At the hardware level, high-speed energy-adaptive computing demands a radically new hardware connectivity architecture for computers and their components. This architecture will complement current wired connections for high-speed applications. Today’s copper based connections have reached their fundamental physical limitations. It is expected that they will not be able any more to transmit the increasing data rates demanded between integrated circuits and to connect hundreds of nodes in typical high performance cluster servers [5].

Up to now the design of interfaces at board as well as backplane level has focused on delivering an ever increasing data rate. This, however, has come along with an energy demand for operating the interfaces such that it is becoming the bottleneck for implementation [6]. Hence, new energy efficient interface designs as well as new interconnect architectures must be designed to find solutions which meet the increasing demand in communications bandwidth at reasonable energy levels.

In order to enable the targeted breakthrough in the area of computer connectivity, an aggregated bandwidth in Tb/s regime must be achieved by improving the data rate per link

and simultaneously the channel density for parallel links, while overall the specific power consumption per transferred data rate needs to be minimized. To meet this goal we propose to employ wireless beam-steered links for board-to-board, and on-board optical links for chip-to-chip communication. This hybrid link technology will be employed for enabling a novel server architecture, the fat tree with hyper connections which we refer to as hyper-fat tree, Figure 3. It has the potential to provide an excellent hardware infrastructure for energy-adaptive software and networks.

The spectral efficiency of the on-board and board-to-board communications is increased by clever routing and network coding design which achieves the maximum multicast network flow. Considering security aspects of the network coding schemes is of essential importance for ensuring both confidentiality and integrity of the data to be transmitted even in the case of intended attacks. Introducing security mechanisms often increases the costs in terms of additional operations to be performed or additional data to be sent and, hence, influences the energy efficiency. One goal within the project is, therefore, to analyze the security of existing network coding schemes and investigate possible improvements of their security considering energy constraints. We refer to this novel computing architecture approach as the HAEC box. The energy efficiency is increased by providing direct links reducing the amount of switches between nodes, as well as by realizing the capability to adapt dynamically to the required data rates: hardware will be turned off when not needed or operated at lower power when lower data rates are sufficient.




Figure 3. Hyper-fat tree architecture

With respect to the platform, one important innovation of the HAEC program consists in the multidisciplinary approach to the scientific and technical challenges of the two critical interconnect bottlenecks of the hyper-fat tree architecture: the board-to-board and the chip-to-chip on-board communications. We propose to rely on sub-THz wireless links for board-to-board and optical links for on-board communications, as shown in Figure 4.

### A. Wireless Communication Between Computer Boards

Already early on the vision of chips not having to be wired but to be freely connected with untethered links has existed [7]. Recently, first results have been shown for building high-speed mm-wave transceivers, integrating the transmit/receive antenna in silicon [8]. With new results in active antenna design allowing for very small-scale antenna sizes [9] it will be possible to design integrated antenna arrays allowing for setting-up directed controlled beam-steered communications channels.

Based on these first promising results, as a significant novelty, we employ beam steered wireless communication techniques for data connection between the computer boards. Since no wires are required, wireless communication allows significantly lower routing complexity. Hence, novel and significantly more complex network architectures are possible. At the speed required for board nodes, conventional network architectures can be easily extended, e.g. with connections within the horizontal plane. These ring-extended trees, which can be further extended towards hyper-trees, significantly reduce the latency and increase the throughput on the boards, since the data does not have to pass the previous parent nodes typically acting as bottlenecks. Conventional wireless systems operating in the lower GHz range are not sufficient to deliver the data rates required in HAEC systems: we will therefore exploit the frequency range between 100 and 300 GHz. Within this frequency range, bandwidths and data rates above 25 GHz and 50 Gbit/s respectively are expected to be feasible under line of sight conditions up to 30 cm. Our plan is to research polarization multiplexing to enable data rates up to 100 Gbit/s per link, or data rates not realized to date. Wavelengths in the millimeter and sub-millimeter range allow the realization of very compact high-gain antenna structures on several board nodes. Using modern beamforming and -steering techniques, different network topologies including adaptive channel allocation are applicable. This wideband wireless communication concept provides extremely flexible, robust and weight optimized data links between computer boards.




Figure 4. Simplified schematics of novel network topology for server farms combining the ad-vantages of short range wireless and E/O wave guide data communication enabling hyper and ring connections at relevant nodes.

However, the challenges regarding digital modulation schemes, algorithms, circuits, antennas, interconnections,

packaging, semiconductor technologies, modelling and electromagnetic interactions/shielding are significant. Transistors with transit frequencies of around 0.5 THz are required. Studies will be performed to evaluate the suitability of advanced semiconductor technologies, device models, and implementation of the electronic interfaces to improve the system-efficient integration and to reduce the link dimensions. Achieving speeds up to sub-Tbit/s per series link is very challenging since it requires the electronic circuitry to operate at frequencies close to the cut-off of most advanced silicon technologies. In this context, we plan to consider novel circuit topologies, suitable to extend the maximal frequency of operation, and circuit models accurate also in the breakdown regime, as well as on comparing different process technologies with respect to their performance and feasibility. This will include the exploration of ultimate technology limitations via computer models and experimental characterization.

The architecture of the wireless system will impact critically the energy efficiency. Thus system-level decisions will be made across projects with the primary goal of achieving the required performance at the minimum power consumption. For example, it is planned to implement the beam-steering of each phased array in the analog domain, by controlling amplitude and phase at each antenna element with compact millimeter-wave components. This approach will enable energy efficiency improvements through a massive reduction in the number of integrated active blocks as only one ADC per phased array are needed. The digital processing blocks will still have full control of the beam direction, while only a few dB of receiver sensitivity and transmit power will be sacrificed. This drawback is - however - acceptable, given the intrinsic short-range nature of board-to-board communications.

### B. Optical Inter-Chip Onboard Communication

Since a decade it is already widely accepted that onboard optical communication will displace electrical interconnects for very high data rates - only the when and how remained open [10]. The HAEC approach now combines the flexibility of the wireless for board-to-board communication with the surpassing performance and bandwidth density of optical chip-to-board communication in order to simultaneously overcome both major intra-rack system bottlenecks. Moreover, in this interplay the optical onboard communication will work interference-free and allows for conversion-free system-spanning rack-to-rack interconnects, where optical yet has replaced electrical interconnects. In consequence, the traditional backplane could furthermore only exist for mechanical electrical power means and comparable low data rate control signaling.

Compared to competing approaches for optical onboard communication a major innovation within HAEC is the energy adaptivity. The newly designed transceivers can adjust their link parameters (e.g. trading data rate off for power consumption) with respect to the actual load condition and energy requirements. Existing integrated circuits (ICs) are either statically trimmed for highest performance or lowest energy consumption and do not provide such energy adaptivity. In HAEC, the dynamic link requirements will be determined by the energy control loops of the computing management as the focus of HAEC-B. As a consequence, a part of the parallel

links can be switched off to save power, and the remaining active links can be optimized for lowest possible power. Due to the lower link bandwidth requirements at lower data rates, the circuit bandwidths, the clock frequencies, and hence the corresponding supply current (partly also the supply voltage) can be reduced for a link path. Besides the dynamic speed adaptivity also the advanced modulation schemes will be researched and assessed for power optimization to enable a certain data rate at different modulation schemes and symbol rates. The challenge is to find the optimal trade-off between number of active links, bandwidth per link and complexity of the data format and to quickly switch between these optima as to the requirements. In later program phases, in addition to BiCMOS and CMOS technologies, the suitability of emerging technologies, such as nanotubes and nanowires will be investigated.

Since up to now there is no light-source capable for silicon integration, novel packaging and assembly technologies are needed to build up hybrid silicon /SiGe/ III-V transceivers with high-speed 3D stacking. Even though these lasers (VCSELs) can improve the bandwidth density by a higher per-link data rate over the electrical counter parts still massive parallelization is needed. Hence, these optical I/Os will be arranged in arrays. As a consequence, arrayed driver electronics have to be designed and then stacked for better high-speed performance than the already demonstrated lateral or even peripheral integration concepts. Only the close collaboration in the form of a joint project between packaging and chip design can guarantee the optimum system performance because assembling techniques can significantly affect the circuits' behavior. Prospectively SiGe-compatible on-chip light-sources from the joint lab with the Leibniz institute IHP (Innovations for High Performance Microelectronics) can reduce this packaging complexity. Nevertheless, the challenge of the robust optical E/O device-to-waveguide coupling remains and is further complicated by the arrayed arrangement of the optical I/Os underneath the chip. The development of this complete electro-optical subassembly is also addressed. It will comprise of a coupling micro-optic that can redirect the light by 90 degree from the E/O socket on the board surface into the embedded waveguides. To allow further bandwidth density improvements the optical I/O pitches will be downscaled. Therefore, the micro-optics to be newly developed must provide an optical pitch conversion capability (fan-out) in order to passively couple the light into electro-optical circuit boards with relaxed accuracy requirements and low loss.

Due to the arrayed optical I/Os the board-level-waveguides will have to be a multi-layer structure to reroute the dense parallel optical bus. Packaging technologies will be developed for the patterning, stacking and integration of multi-layer optical PCBs for high reliability despite their compliance with PCB typical misalignments. Only this way repairing in the field will be enabled by using a low-cost substrate, which can achieve competitive solutions compared to current electrical solutions. The new technologies will drive the implementation of silicon photonics with their associated waveguide requirements for the later phases of the HAEC program.

### C. HAEC Box Analysis

A central part of this program is the analysis of the new means of wireless and optical chip-to-chip communication from the view of the computing system using them. With this holistic approach, it is adequate to define energy efficiency as the amount of energy the HAEC box needs to finish a specific task set.

An event-based simulator will be developed to emulate the interactions between the HAEC hard- and software models. The event trace-based simulator will be designed to emulate runs of energy-aware software on the HAEC operation system using the HAEC communication system providing sensor data for the HAEC monitoring system. For this, the HAEC box simulator has to combine models of different abstraction levels (for energy consumption of components, communication, operation system, energy-aware software and monitoring). This combination is of central significance since, per design, HAEC hard- and software will be tightly coupled by energy/utility functions and energy control loops, see below.

The connected models include detailed information on energy consumption for all their relevant processes and system states. With this, the simulator becomes an environment to verify the HAEC design ideas and to analyze their impact on the energy efficiency at large. The systematic analysis of simulation runs will allow the specification of design requirements and optimizations for the hard- and software development.

### III. (B) ENERGY-ADAPTIVE COMPUTING MANAGEMENT

Program part B will contribute the software to make use of the adaptive hardware architecture of the HAEC box for highly energy-efficient systems. To achieve this goal, few fundamental principles are used in an integrated holistic manner. At the core are Energy Control Loops which we call HAEC Loops in the following and Energy/Utility Functions.

Analogous to traditional control technology, the system state of controlled objects in the HAEC Loops are continuously monitored and controlled (cf. Figure 5). Controlled objects in the HAEC Loops are the physical and logical resources provided by HAEC hardware and software. HAEC control actions select, activate and deactivate such physical and logical resources. Information about these resources is collected by HAEC monitoring techniques. HAEC algorithms implement the control law and optimize the usage of energy based on an analysis of monitoring results. For example, DVFS (dynamic voltage and frequency scaling of CPUs) could be called a HAEC Loop.

Optimizing the usage of energy comes in several strategies. One is to maximize the utility to be achieved if a given fixed amount of energy is available. Utility can be defined as e.g., the number of transactions per second in a data base system or a web service, the execution time of an application in high performance computing or the quality of service in a multimedia application. Another strategy is to minimize the energy required to obtain a given fixed utility. Many more strategies result from a combination and variation of these two. For example, one other strategy is to minimize the average

energy requirement that suffices to enable a system to achieve a predefined, seldom required peak performance within very short time.

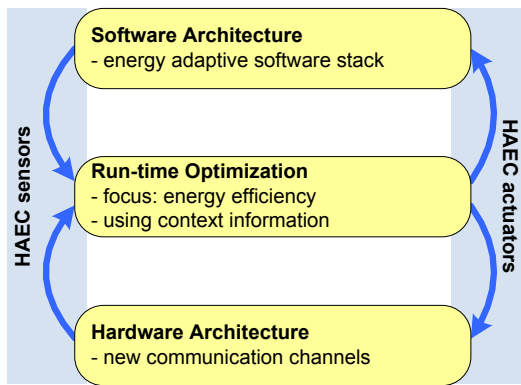


Figure 5. Energy-optimization means energy and context awareness, as well as energy adaptivity under the global strategy of energy optimization. This spans up a dual control loop comprising hard- and software architectures.

For control actions to be useful for energy savings, HAEC hardware and software must provide highly flexible choices between alternatives of differing energy characteristics. While part A takes care of such flexibility in hardware especially in the communication structure, part B must provide such flexibility in software. A common approach to software flexibility is to provide several implementations of one specification. In general, many such techniques have been developed in the disciplines of systems software, compiler construction and software engineering. But only a few of them address alternatives with differing energy characteristics. Examples of such general techniques are hotspot detection and tuning, usage of fundamentally different algorithms, and so called specialization. Once these techniques are combined with HAEC Loops, they allow a characterization by means of currently popular terminology such as self-adaptive software, auto-tuning and so forth. In analogy to hardware resources, we sometimes refer to functionality provided by software as logical resources. For example, sorting is a logical resource with several alternative implementations of different characteristics. One of the major challenges will be to invent and evaluate mutually alternative logical resources and their energy characteristics.

The major challenge for monitoring of resources to be useful is the consideration of trade-offs between size, freshness, accuracy, and overhead induced. Since at the current state of technology energy measurements are not available directly, techniques must be used to deduce the energy usage from other state variables. Monitoring is also needed for system characteristics other than energy, since such characteristics can become constraints for energy optimization. For example, monitoring of CPU load is needed since a fully loaded CPU may become a constraint for the selection of a logical resource requiring many CPU cycles.

While the terminology of control technology served well so far for explaining the fundamental approach of HAEC-B, its usage for more technically detailed descriptions becomes cumbersome and non-typical, especially if several components

of HAEC Loops are intertwined tightly. In such cases we use common terminology from informatics. For example, in resource monitoring the pre-processing of raw measurement data and supervision are often tightly integrated. Resource managers (for example CPU schedulers or disk drivers) are special, highly integrated control loops (though in general not with energy in mind).

Control algorithms to optimize the usage of energy require knowledge of the energy characteristics of physical and logical resources. In contrast to classical scheduling, the well-studied discipline of planning and managing resources, energy-aware scheduling adds an additional indirection: while in classical scheduling all resources are assumed to be just there, resources in energy-aware scheduling become available only if energy is spent to activate the resources. To cope with this challenge, HAEC-B will develop and use Energy/Utility Functions as the second fundamental principle. Energy/Utility Functions characterize the utility of energy for some application, if used for a particular physical or logical resource in a specific context. Application and utility are used here in a fairly general sense.

As there are many classical schedulers in computer systems, there will be many HAEC Loops in the HAEC box. Examples range from low-level CPU schedulers up to the rearrangement of server farms. The HAEC Loops will be able to achieve local optima, but there are cases where cross-layer control is required. For example, if a part of a server farm is shut down, all lower-level HAEC Loops must follow suit accordingly. For these cases we envision a data plane that can be accessed (read from and written to) by HAEC Loops at all layers.

#### IV. HAEC DEMONSTRATORS AND CASE STUDY SCENARIOS

The research program is planned to be performed in three phases. Within the first phase (year 1-4) single technology demonstrators will be built to illustrate and verify the achieved research results. The specification of each of these demonstrators will be based on boundary conditions given by the other demonstrators to which an immediate interfacing is planned. Hence, interactions between demonstrators will be already kept in mind in this early stage to ensure their interoperability later on. At the end of this period, the communication parameters will be available for an expert in computer design in order to start building a specification of the HAEC box. In the second phase (year 5-8) integrated technology demonstrators by collaboration between several projects are intended. For the last phase (year 9 to 12) a comprehensive evaluation is planned, providing an overall picture of the energy reduction that can be gained by the integration of all contributions made in the program. The goal is to show a significant power reduction clearly above 50%.

The technologies and methods developed within the program will be tested against real-life case studies. One currently planned case study scenario is energy-adaptive multi-site video conferencing, as it exposes high power demands yielding great potential to show the benefits of the HAEC design, architecture and the technologies.

The problems that need to be addressed within this case study are real-time requirements (e.g., maximum latency), high scalability (e.g., with regard to bandwidth requirements, computing power, and storage access), synchronization (multi-party audio and video) and transcoding (e.g., to support different stakeholders with different mobile devices). The goal of the case study is to show that these problems, which are shared by most energy-costly distributed applications, are addressed by the program.

Another currently planned case study is the web analysis system Fedseeko [11], which has been developed with the aim to create an all-embracing view of particular products of desire using product information from different online information sources (online malls, producers, consumer rating sites, etc.). The system consists of several web services for gathering, filtering, aggregating, and integrating such information. These web services use a number of algorithms from different domains, such as information extraction, natural language processing, ontology matching, and ontology learning. At present, Fedseeko is neither adaptive nor energy-aware. In the context of the program, the idea is to re-design some of the web services such that they can be reconfigured, migrated or dynamically bound-to to demonstrate some the usefulness of the adaptation and energy-aware strategies that will be developed.

## V. INDUSTRY FEEDBACK

As the research results are expected to affect future server system solutions and to achieve significant industry impact, an early involvement of industry partners is being intended and organized. In the first phase this collaboration will focus on feedback with respect to research directions and to scientific as well as industrial (economical and manufacturing) challenges.

As a world leading IT corporation, IBM can supply the program with valuable feedback regarding boundary conditions of large server systems. So, IBM will share its vision of hardware platforms of the future.

With respect to operating systems of future computing, the AMD OSRC will support HAEC with valuable feedback regarding the energy-adaptive software platform. Also feedback regarding the hardware platform is expected as AMD is a world leading corporation in delivering microprocessor products, in particular for high performance and low-power server hardware.

Additionally, the new on-chip or in-package communication links must be implemented in the future technologies available. Hence, the partner Globalfoundries is instrumental in adding its view on semiconductor technology and manufacturing aspects of the envisioned hardware components.

## VI. THE HAEC PROGRAM, & ACKNOWLEDGEMENT

The described research program brings together 16 chairs of the Technische Universität Dresden from the departments of Electrical Engineering and Computer Science. It will be funded by the German Research Foundation for 4 years (7/2011 – 6/2015) with approximately 8 Mio. Euro.

Special thanks are acknowledged for all partners involved in setting this research up and who contributed with their ideas in putting this program together. In particular to Uwe Assmann, Franz Baader, Christel Baier, Waltenegeus Dargie, Frank Ellinger, Andreas Fischer, Elke Franz, Hermann Härtig, Eduard Jorswieck, Mathias Müller, Andreas Pfitzmann, Dirk Plettemeier, Michael Schroeter, Alexander Schill, Klaus Wolter.

## REFERENCES

- [1] T. Heuzeroth, „Das Internet ist der wahre Klimakiller“, [online] <http://www.welt.de/webwelt/article1203605/>, Sep. 2007.
- [2] G. Fettweis and E. Zimmermann, “ICT Energy Consumption - Trends and Challenges,” Proceedings of the 11th International Symposium on Wireless Personal Multimedia Communications (WPMC'08), Lapland, Finland, 8.9. - 11.9.2008
- [3] A. Fehske, G. Fettweis, J. Malmudin, G. Biczok, "The global footprint of mobile communications: The ecological and economic perspective," Communications Magazine, IEEE , vol.49, no.8, pp.55-62, August 2011
- [4] White paper, “Entering the Zettabyte Era, Visual Networking Index: Forecast and Methodology, 2010-2015”, Cisco Systems, Inc, San Jose, CA, 2011, available online at <http://www.cisco.com/>
- [5] I. A. Young, E. Mohammed, J. T. S. Liao, A. M. Kern, S. Palermo, B. A. Block, M. R. Reshotko, and P. L. D. Chang, “Optical I/O Technology for Tera-Scale Computing,” Solid-State Circuits, IEEE Journal of, 45(1):235–248, 2010.
- [6] Evening Session 3: Energy-efficient High-speed Interface, Organizers: Naresh Shanbhag and Koichi Yamaguchi, Chair: Robert Payne, 2010 IEEE International Solid-State Circuits Conference Dig. Tech. Papers, pp. 524-525, IEEE, 2010
- [7] J.W. Goodman, F.J. Leonberger, S.-Y. Kung, R. Athale,. “Optical Interconnections for VLSI Systems,” Proc. IEEE, vol. 72, pp. 850-866, July 1984.
- [8] Chi-Hsueh Wang et. al., “A 60GHz transmitter with integrated antenna in 0.18 $\mu$ m SiGe BiCMOS technology,” ISSCC Dig. Tech. Papers, pp.659–660, Feb. 2006
- [9] A. Arbabian, B. Afshar, J.-C. Chien, S. Kang, S. Callender, E. Adabi, S. Dal Toso, R. Pilard, D. Gloria, A. Niknejad. "A 90GHz-Carrier 30GHz-Bandwidth Hybrid Switching Transmitter with Integrated Antenna". 2010 IEEE International Solid-State Circuits Conference Dig. Tech. Papers, paper no. 23.4, IEEE, 2010
- [10] D.A.B. Miller. Optical interconnects to silicon. Selected Topics in Quantum Electronics, IEEE Journal of, vol. 6(6): pp. 1312-1317, Nov/Dec 2000
- [11] M. Walther, L. Hähne, D. Schuster, and A. Schill, „Locating and Extracting Product Specifications from Producer Websites,” Int. Conf. on Enterprise Information Systems, 2010