

Entering The Path Towards Terabit/s Wireless Links

Gerhard Fettweis, Falko Guderian, Stefan Krone

Vodafone Chair Mobile Communications Systems

Technische Universität Dresden, 01062 Dresden, Germany

E-mail: {fettweis, falko.guderian, stefan.krone}@ifn.et.tu-dresden.de

Abstract—Wireless communications has been a hot area of technology advancement for the past two decades. As long as memory sizes increase, the demand for higher data rates of communications increases on the same scale. This means that one must understand today's high-end 10 Gbit/s wireless technology to get prepared for 100 Gbit/s and 1 Tbit/s data rates of tomorrow. This paper presents key boundary conditions learned by understanding leading edge wireless links of today to prepare for the Tbit/s technology of the year 2020.

I. INTRODUCTION

During the last years it has become evident that wireless data rates are increasing over time. This is mainly driven by the need for having to move files of ever increasing size, as well as for downloading streaming / podcast services. Reviewing the data rate scaling over time (for non-voice), the main driver of wireless technology development can be identified. The ITRS roadmap for flash memories has shown a history as well as projection of a factor of 10 in memory increase every 5 years, i.e. doubling every 18 months. Flash memories are the important storage technology for most "wireless gadgets" which require broadband connectivity, as e.g. cellular (smart) phones, game consoles, cameras, camcorders, sub-notebooks, e-books, as well as modern laptops. Hence, the storage size increase of flash memories directly drives the size / amount of data stored. And as stored data needs to be moved from one device to another, this size increase also drives the need for communications data rate.

Wireless gadgets need connectivity at different levels. Firstly, an easy way of moving data quickly between two devices that are locally closely positioned at approximately 1 m distance to each other. Up to now cabled USB has proven to be the transport media of choice. Starting with USB 1.0 at 2 Mbit/s, data rates have now reached 4.8 Gbit/s with the introduction of first USB 3.0 interfaces during Q1 of 2009. Wireless USB was considered many times as an alternative, in particular for connecting devices which have an own power supply and do not need the 2 W powering capability via the USB cable (e.g. downloading pictures from a digital camera to a laptop, both having batteries for operations). So far, however, wireless USB has not been economically successful. The main reason for that is easily found by the fact that IEEE 802.11 WLAN standards have also been providing increasing data rates, as can be seen in Fig. 1. So far, the difference in data rates between wireless USB and IEEE 802.11 is not large enough to make a noticeable




Fig. 1: Wireless roadmap.

difference in user experience for the customer. In addition, the cost difference for a chip-set is negligible, and the IEEE 802.11 protocol stack is well integrated into PC and cellular phone operating systems, whereas the wireless USB protocol stack would require a whole new design effort with unknown engineering and debugging problems. Therefore, even though WLAN systems are specified for range of operations of 10 m and are intended for client-server applications in offices, hot-spots, or homes, they have been also able to destroy the business case for wireless USB. The negligible speed improvement of wireless USB over WLAN is too small to have a chance in the market [1].

Cellular, on the other hand, is a very different business. The very most important feature of cellular is to provide coverage, i.e. connectivity no matter where the customer is positioned. Once connectivity is available and a connection has been established, the speed of data communications becomes the next prominent feature to be provided. This has driven a constant demand for higher data rates also for cellular, as can be seen in Fig. 1. Currently, the next generation cellular, LTE, is expected to provide data rates on the order of 36 Mbit/s at its time of large scale market introduction in 2012 [2]. This clearly fits very much in line with the projected 10x data rate increase every 5 years mentioned above.

Remarkably, the data rates provided by USB, WLAN, as well as cellular, all increase by a factor of 10 every 5 years, exactly at the same rate as the increase in size of flash memories according to the ITRS roadmap. The difference in data rates between cellular and WLAN of 100x has, so far, created a large enough differentiator for both technologies to be able to succeed in the market place. In case a 1 Tbit/s wireless link was available before the year 2020, a new market for wireless USB could be created. Fig. 1 shows this clearly.

One open question remains: Is there a need for data rates well beyond 100Mbit/s in future? Today already more than 50% of the data volume measured in cellular networks is generated by users consuming streaming applications. As high definition 3D video streaming with user enabled vision angle control requires on the order of 100Mbit/s, and users want quick downloads of typically above 100 x real-time of multiple streams, we will see 10Gbit/s - 100Gbit/s wireless connectivity as a requirement coming up in future, growing up to 1 Tbit/s! Obviously, this does not lead to a need for a continuously sustainable ultra-high data rate for one user over long periods of time. Instead, 100Gbit/s to 1 Tbit/s data rates will be shared via the wireless medium.

Every line within the wireless roadmap presented above has its unique market need, but interestingly, also its own technical requirements. The short link boundary conditions differ very much from WLAN which again differ much from cellular. In case of cellular the basic technical challenge is to provide coverage. Only once coverage has been achieved, does it matter to provide the maximum data rate possible. The wireless channel is predominately non-line-of-sight (NLOS) and must cope with vehicular velocities. Not only does a multiple access technique need to be designed, but also a hand-off scheme between different cells. In case of WLAN, the users can be mobile at pedestrian speed, and the channel clearly is often also NLOS. And since the market requires multiple terminals to be connected to one access point, the wireless channel must be shared and requires a multiple access technique. In case of short links, the wireless channel can be assumed to be point-to-point with line-of-sight (LOS), connecting two devices with each other. The mobility of the devices is zero. The wireless channel will only have variations over time due to reflectors moving. This is the basis for the following analysis given once for a short-range scenario of USB-like data transfer, i.e. a wireless data kiosk, followed by an outreaching approach of board-to-board communications via a wireless backplane within electronic computing devices.

II. WIRELESS DATA KIOSK: 10 GBIT/S AT 60 GHz

Novel short-range applications like the wireless data kiosk are pushing the data rate requirements of wireless technology towards tens of Gbits/s. Such high data rates require a re-thinking of traditionally digital-oriented system designs and serialized signal processing. New physical layer concepts have to account for hardware-related design constraints such as the coarse resolution of high-speed data converters and a limited clock rate of the digital baseband processing. By shifting parts of the digital baseband processing to the analog domain, and with a proper parallelization of the physical layer, monolithic mm-wave transceivers that allow for 10Gbits/s are feasible with today's CMOS technologies. This section reviews the respective physical layer design and digital baseband implementation as proposed in [3] and [4].

A. Scenario, Requirements and Design Constraints

A wireless data kiosk is a vendor machine that allows for ultra-fast download of large amounts of bulk data (e.g. movies




Fig. 2: Wireless data kiosk: 10 Gbit/s content download.

in HDTV quality) to a user terminal. The user is situated in front of the machine at a distance of less than 1 m. Transmitter and receiver have a LOS connection, as depicted in Fig. 2. The transferred data will be stored on the user terminal and is further processed (e.g. by a movie player) once the download has been completed. The user terminal is a battery powered mobile device with corresponding constraints on the form factor and energy consumption. Both constraints are less severe for the data kiosk which has access to mains supply.

A download should only take an acceptable amount of time, typically less than 10 s. With a data rate of 10 Gbit/s, the content of one DVD (4.7 Gbytes) can be downloaded in less than 4 s, whereas several minutes or even hours would have to be spent when relying on conventional systems (see Table I).

TABLE I: Download times with different data rates.

	1 Gbyte	5 Gbytes
3 Mbit/s (Bluetooth® V2.0)	44 min 27 s	3 h 42 min 13 s
54 Mbit/s (IEEE 802.11a)	2 min 28 s	12 min 21 s
10 Gbit/s (60 GHz)	0.8 s	4 s

The anticipated data rate requires a system bandwidth in the GHz-range. The 60 GHz ISM band provides up to 10 GHz of unlicensed bandwidth. Considering a bandwidth of several GHz, the unavailability of energy-efficient high-resolution data converters (see [5]) calls for an analog-oriented system design that relies on low-order single-carrier modulation. Parts of the baseband processing have to be carried out in the analog domain, which includes the transmit and receive filtering, the carrier / phase synchronization and the symbol clock recovery at the receiver. Such an analog-oriented design is less flexible but very much suited for static scenarios with stringent requirements on the receiver complexity and energy consumption.

The 60 GHz LOS channel found in a wireless data kiosk setup is mostly static and frequency-flat, since the amount of scatterers in the direct vicinity is strongly limited. This avoids channel adaptive equalization, but comes with the drawback that spatial diversity is limited to polarization diversity. There is no interference from multiple active links in the same frequency band, which further simplifies the system design. Low receiver complexity and energy consumption can be achieved with 1-bit data converters and QPSK modulation, which are the basis for the proposed physical layer (see Sec. II-B). Very long data frames of several kbytes reduce the physical layer overhead to a minimum. Latency requirements are relaxed, as large amounts of data can be transmitted at a stretch, without requiring an immediate response from the receiver. Time division duplexing is best suited to separate uplink and downlink, due to the small transmission range and highly asymmetric traffic.

The main challenge for the digital baseband implementation is the processing speed. A throughput of 10 Gbit/s can only be

achieved with massive parallelization. Low-power digital signal processors have a clock rate of at most 300 MHz, which requires at least 30- to 40-fold parallelization for 10 Gbit/s. Another issue is the data storage. The storage capacity at both the transmitter and receiver has to be on the order of hundreds of Gbytes. The (non-volatile) memory must support a read/write speed in excess of 1.25 Gbyte/s for a data transfer at 10 Gbit/s. Today's solid state disks achieve write access at 250 Mbyte/s, and the underlying SATA interface is specified for up to 3 Gbit/s. The performance bottleneck of non-volatile storage devices can be tackled by using multiple devices in parallel.

B. Physical Layer Design

Concerning the hardware-related need for parallelization, it is convenient to parallelize the physical layer itself, such that the data sequence to be transmitted is split into a number of lower rate sequences. With QPSK and polarization multiplexing, a 10 Gbit/s data sequence can be split into four 2.5 Gbit/s sequences, which are mapped to the two complex and two spatial dimensions. To further increase the parallelization, the payload of the data frames can be split into 32 parallel (convolutional) codewords. 128 parallel en-/decoders are then only processing data sequences at 78.125 Mbit/s each. Three types of parallelism are thus jointly used (see Fig. 3): (a) parallel codewords, (b) inphase and quadrature phase, (c) two linear polarizations.

The exact physical layer parameters are tuned for low-complexity transceivers with 1-bit data conversion and symbol-rate sampling at 3.456 GHz. The differential QPSK modulation enables frequency flat phase equalization at the receiver using the analog carrier recovery proposed in [6]. A residual phase-ambiguity of multiples of 90° is resolved in the digital domain. The bandwidth follows the channelization plan of the ECMA-387 standard [7] with bonding of two channels. The physical layer comprises four configurations with different data rates deriving from two different convolutional codes and the usage of polarization multiplexing as shown in Table II.

TABLE II: Physical layer configurations for the wireless data kiosk.

Configuration		SISO-A	SISO-B	MIMO-A	MIMO-B
Carrier frequency	GHz		61.56		
Bandwidth	GHz		4.320		
Symbol rate	GHz		3.456		
Modulation	-	SC - DQPSK			
Polarizations	-	1			2
Parallel frames	-	2			4
Code rate	-	1/2	3/4	1/2	3/4
Payload length	symbols	163840			
Preamble length	symbols	4094			
Data volume per frame	kbyte	20.16	30.40	20.16	30.40
Total frame duration	μ s	48.593			
Data rate	Gbit/s	3.32	5.00	6.64	10.01

The convolutional coding is based on a mother code of rate $1/3$ with termination and puncturing. For both code rates, the codewords length is 1024 bits. The preamble in front of each data frame consists of an analog training part with alternating 1 and 0 for symbol clock recovery, a digital training part that is an M-sequence of length 1023+1 for packet detection and for resolving phase ambiguities of multiples of 90° , and a signaling field with 32 MAC header bits that are encoded with a spreading sequence of length 32.




Fig. 3: Physical layer parallelization for 10 Gbit/s.

Link budget calculations have shown that this physical layer has an operating margin of around 10 dB for a bit error rate of 10^{-6} at 1 m distance, when considering 9 dBm transmit power per polarization, antenna gains of 9 dBi and state-of-the-art SiGe-BiCMOS technology for the RF frontend integration. The maximum distance to support 10 Gbit/s is 3.1 m [4].

C. Digital Baseband Implementation

The digital baseband processing for the proposed physical layer is divided into a transmitter and a receiver data path (each with a dedicated controller). Both data paths are based on a linear data flow structure and clocked at 108 MHz.

Fig. 4 shows the block diagram of the transmitter data path for 10 Gbit/s (MIMO-B). The MIMO input module splits the 128-bit data words from the MAC layer into four 32-bit words that feed the four parallel processing chains. The convolutional encoding with rate $1/3$ and the termination are carried out by 128 parallel encoders. The puncturing and the differential modulation can be realized as a single module as follows: For each polarization, the two 96-bit output words of the encoders are de-multiplexed according to the puncturing pattern and fed into a FIFO. The FIFO is read by a DQPSK modulation unit that joins the inphase and quadrature phase bits for DQPSK. The unit consists of four parallel modulators for the four possible DQPSK states and a subsequent logic that selects the 32-bit output words of the inphase and quadrature phase according to the DQPSK state. The framing module generates the preamble in front of the payload and encodes the signaling field. At each 108 MHz clock cycle, 32 output bits are fed into one of the four 1-bit parallel-to-serial converters (1-bit DACs) that interface the analog frontend (AFE) hardware.

The receiver data path realizes complementary functionality, as shown in Fig. 5. The incoming binary inphase and quadrature phase symbols of the two polarizations are de-serialized, which yields four 32-bit input words at each 108 MHz clock cycle. The clock synchronization between transmitter and receiver is implemented in the analog domain. The four 32-bit words are fed into an acquisition module that realizes frame detection, phase and timing synchronization. For each of the two polarizations, the acquisition module consists of four parallel cross-correlation units that correlate the received bits with phase-shifted versions of the digital training sequences to resolve 90° phase ambiguities, a threshold detector and a multiplexer. Each cross-correlation unit comprises 32 parallel cross-correlators (shifted by 0 to 31 bits) to detect the correct




Fig. 4: 10 Gbit/s transmitter data path.




Fig. 5: 10 Gbit/s receiver data path.

timing instant within the 32-bit input words. If one cross-correlator exceeds a threshold, a frame is detected. The received bits are then multiplexed and delayed to match the parallel structure of the transmitter. The acquisition module also extracts the MAC header bits of the signaling field bits. The differential demodulation and depuncturing are realized as a single module, in the same manner as the puncturing and modulation at the transmitter. 128 parallel Viterbi decoders are decoding the received bits with code rate 1/3. The MIMO output module combines the 32-bit output words of the four receive processing chains to 128-bit data words which are delivered to the MAC layer.

The digital baseband has been implemented and tested on a prototype hardware platform using Altera Stratix III/II-GX FPGAs. The Stratix II-GX Multi-Gigabit Transceivers act as 1-bit DACs / ADCs with parallel-to-serial and serial-to-parallel conversion, respectively. The integrated clock recovery unit of the Multi-Gigabit Transceivers has been used for the symbol clock synchronization between transmitter and receiver. Table III shows the resource usage of the transmitter and receiver data paths (including the respective controllers) for the 10 Gbit/s implementation on Altera Stratix III devices. The first column refers to the logic usage, while the other columns reflect the memory usage. An ALUT is an adaptive lookup table with 4 to 6 inputs. The high logic and memory usage of the receiver data path (requiring more than one Altera Stratix III) is due to the acquisition module (256 parallel cross-correlators) and the decoding module (128 parallel Viterbi decoders). Considering the latency when transmitting a single frame, the transmitter data path introduces a delay of 107 clock cycles ($0.99 \mu s$) without the Framing module. This is less than the preamble duration and therefore not critical, as the preamble is generated and transmitted when the encoding and modulation of the payload bits starts. The delay of the transmitter data path is hence given by the preamble duration ($1.19 \mu s$). The receiver data path introduces a delay of 279 clock cycles ($2.59 \mu s$), where the main part originates from the Viterbi decoding.

TABLE III: Altera Stratix III resource usage for the 10 Gbit/s digital baseband.

	Combinational ALUTs	Memory ALUTs	Dedicated logic registers	Block memory bits
Transmitter data path	10000 (4 %)	4000 (2 %)	14000 (6 %)	132000 (1 %)
Receiver data path	414000 (154 %)	0	290000 (108 %)	2878000 (18 %)

III. WIRELESS BACKPLANE: 100 GBIT/S AT 100-300 GHz

Ultra-high rate wireless links do not only allow for new consumer applications like the wireless data kiosk. They can also overcome technological limits of wired transmission in future computer racks. Today's copper-based connections from one computer board to another will soon become a bottleneck given the constantly increasing data rates and resulting complexity of the connections within a high performance cluster server. 100 Gbit/s wireless links in the frequency range of 100 GHz to 300 GHz are a promising solution. The high frequencies allow for very compact antenna structures that can be integrated on the computer chips. This enables several parallel links between different computer chips that can be allocated in a flexible way and switched by means of beamsteering.

Wireless technology for high-speed board-to-board interconnects is a new territory that has rarely been investigated so far. Most of the few related literature deals with antenna and analog frontend concepts but does not consider the overall system design, i.e. the required physical layer and baseband configuration. A concept for wireless data transmission between computer chips on a single board with beamforming antennas and a reflector in the 10 GHz to 20 GHz frequency range has been introduced in [8]. Sony Corp. has recently announced an integrated solution for wireless board-to-board interconnects relying on inductive coupling [9]. A data rate of 11 Gbit/s can be achieved, but the coupling limits the distance to 14 mm.

This section discusses constraints, requirements and challenges for the physical layer and digital baseband design of a 100 Gbit/s wireless backplane, i.e. a network of 100 Gbit/s wireless board-to-board links based on beamsteering with coverage on the order of 10 cm.

A. Scenario and General Constraints

A wireless board-to-board link is in principle very similar to the 60 GHz link of a wireless data kiosk, except for the higher carrier frequency and a 10 times larger data rate at a 10 times smaller distance. A board-to-board link should be designed for ultra-fast bidirectional data transmission rather than asynchronous data transfer. Transmissions that require an immediate response from the receiver are more likely than the transfer of large amounts of data at a stretch. Latency is therefore a critical issue for the physical layer and digital baseband design of wireless board-to-board links. An example that clearly illustrates the high demands on very short latency is the wireless connection between a CPU and RAM. The requirements on the transceiver complexity and energy consumption are less severe as compared to battery-powered mobile devices, but the trend in ICT energy consumption clearly asks for energy-efficiency.

With highly directional beamsteering antennas, the frequency selectivity of the wireless channel at 100 GHz to 300 GHz will mainly result from the antennas themselves and from the analog frontend hardware. Like the 60 GHz channel of the wireless data kiosk, the radio channel of a single board-to-board link can be assumed to be static and largely frequency flat, i.e., adaptive channel equalization is not required. Multiple active links in a wireless board-to-board network will, however,




Fig. 6: Wireless backplane: 100 Gbit/s board-to-board links with beamsteering.

cause interference. The steered antenna beams consist of a main lobe and undesired side lobes as depicted in Fig. 6. The side lobes and the penetration of the main lobe through boards lead to interference to/from different links. A computer chip can have more than one antenna to serve multiple links in parallel. Each antenna consists of multiple elements. The (analog) beamsteering is based on adjustable amplifiers and phase shifters that are connected to the antenna elements and controlled by the digital baseband to adjust the antenna weights. Digital beamsteering would require a complete RF transmit/receive path per antenna element and would hence only allow for a small number of elements with limited directivity but rather high energy consumption. It would also require high-resolution data converters in each transmit and receive path, which is again not feasible.

100 Gbit/s wireless links clearly involve higher hardware-related challenges for the digital baseband design than 10 Gbit/s links and call for sophisticated concepts that trade off analog versus digital signal processing. The major challenge is to find an optimal placing of the data converters (with limited resolution and sampling rate) in the transmitter and receiver paths. An analog reference clock signal for synchronous sampling and up-/down-conversion can be distributed to all boards over the remaining wired backplane. This avoids sampling and carrier frequency offset compensation. Phase offsets, phase noise, clock jitter and further impairments can, however, still be present and require either analog compensation or robust modulation and error correction in the digital domain.

B. Physical Layer Aspects

As the radio channel of a wireless board-to-board link at 100 GHz to 300 GHz can be assumed to be frequency flat, single-carrier modulation will be the natural choice. There are, however, still degrees of freedom in choosing an optimal modulation alphabet depending on the quantization resolution of the data converters. It has been shown in [10] that even with 1-bit ADCs at the receiver, it is still reasonable to use higher-order alphabets that maximize the data rate if moderate oversampling is applied. Besides the lowest energy consumption of 1-bit ADCs, an automatic gain control can be omitted, which further reduces the receiver complexity and energy consumption. High-order alphabets will, however, increase the peak-to-average-power ratio of the transmitted signals and thus reduce the energy efficiency due to higher losses of the transmit power amplifiers. An optimal physical layer design will have to maximize the overall energy efficiency of both the transmitter and receiver.

A data rate increase of a factor of 10 as compared to the physical layer proposed in Sec. II can be realized in different ways. A straight-forward solution would be to scale up the symbol rate by a factor of 10, which leads to a bandwidth on the order of 50 GHz. Such a large bandwidth requires a free frequency band of corresponding size, e.g. in the 100 GHz to 300 GHz range. The bandwidth can be significantly reduced with higher-order modulation, while still relying on low-resolution (e.g. 1-bit) ADCs with moderate oversampling as indicated above. Another approach to further reduce the bandwidth is the usage of parallel pencil beams between two communicating computer chips with multiple antenna arrays. The best system bandwidth may result from any combination of these approaches and certainly depends on the frequency regulations and on an optimal system design to maximize energy efficiency. The physical layer will require significant parallelization (on the order of 1000) to enable an implementation with state-of-the-art digital processing devices.

When scaling up the symbol rate by a factor of 10, an approximate link budget can be calculated for a single 100 Gbit/s link at e.g. 100 GHz carrier frequency. Considering pure LOS, similar analog frontend characteristics as for the 10 Gbit/s case and a target bit-error-rate of 10^{-6} , the operating margin at 10 cm distance is on the order of 10 dB. The maximum coverage derives as 31 cm. When using carrier frequencies up to 300 GHz, larger antenna gains or a higher transmit power are required to compensate for the higher path loss.

The physical layer frame size has to be kept at minimum to satisfy the latency requirements. The payload and the preamble length should be minimized on the same scale to ensure small physical layer overhead. An analog training part as required for the preamble of the wireless data kiosk scenario can be omitted in a wireless board-to-board setup, since sampling clock synchronization is not necessary when a reference signal is available to all transmitters and receivers.

The most important issue in terms of latency and link performance is the selection of a suitable channel coding scheme. From [11, 12] it is known that convolutional codes are most attractive for low latency en-/decoding, whereas LDPC codes have a better bit-error-rate performance when higher latency can be accepted. LDPC convolutional codes can combine both advantages [13], which makes them suitable for latency-constrained high-performance error correction. With termination and tail-biting, LDPC convolutional codes can be perfectly tailored to codewords of small length [14], as required for low-latency wireless board-to-board links.

C. Digital Baseband Aspects

The key challenge for the digital baseband implementation of 100 Gbit/s wireless links is to find a trade-off between processing delay and energy consumption. A digital baseband clock above 1 GHz would be required when scaling up the 10 Gbit/s design of Sec. II by a factor of 10 without further parallelization. The clock frequency of current energy-efficient ASIC/ASIP designs for digital processing is typically not larger than 500 MHz [15]. The clock distribution network of integrated circuits consumes a significant portion of the energy, area,

and metal resources, all of which scale with the clock frequency [16]. A clock frequency below 500 MHz is therefore required, which can only be achieved by means of parallelization, i.e. at the expense of higher complexity and larger chip area.

A natural approach would be to scale the parallelization of the 10 Gbit/s design by another factor of 10. Then, as a result, the implementation complexity and meeting the delay constraints are the dominant issues. To illustrate the resulting logic and memory requirements, consider the 10-fold up-scaling of the receiver data path proposed in Sec. II-C. At least 13 Altera Stratix V devices would be required to accommodate the receiver data path of a single 100 Gbit/s link. Note that the Multi-Gigabit Transceivers of upcoming Altera and Xilinx FPGAs support 1-bit data conversion at 28 Gbit/s [17, 18], which does already allow for a prototype implementation of 80 Gbit/s wireless links based on an up-scaled version of the design proposed in Sec. II.

Concerning practical realizations, a throughput of 100 Gbit/s calls for new concepts that jointly optimize the implementation complexity, processing delay and energy consumption. A promising approach is to consider pipelined processing rather than pure parallelization. To optimize the energy efficiency it is important to consider both the signal processing as well as the transmit power. Both items depend on each other and can have a significant impact on the overall energy consumption for a target link performance. For LDPC decoding, it is known that the energy consumption scales linearly with the number of iterations. Simplified coding can reduce this energy consumption but may require a higher signal-to-noise ratio and thus a higher transmit power to achieve the same link performance.

For a network of 100 Gbit/s links in a wireless backplane setup, one has to cope with the interference of multiple active links. The simplest way to ensure low complexity and small latency is interference avoidance, i.e. the antenna weights of all potential links are pre-computed to always minimize the interference to / from non-intended receivers / transmitters, and stored in a look-up-table on the computer chip. Techniques that always compute the best antenna weights from the present interference may allow for a better link performance, but only at the expense of increased transceiver complexity and with a delay that will most probably not satisfy the latency constraints of a wireless link between e.g. a CPU and RAM.

IV. CONCLUSIONS

During the last two decades, wireless technology has followed the ITRS roadmap for flash memories with a 10 x data rate increase every 5 years. Tbit/s wireless links will soon be required to keep track with the technology development of flash memories and to satisfy the ever increasing demand for "wireless gadgets" that improve everyday life.

Reviewing latest and upcoming developments, this paper has provided insight into the constraints and challenges for the design and implementation of close-to-Tbit/s wireless links with state-of-the-art technologies. A wireless data kiosk that achieves up to 10 Gbit/s at 1 m distance has been considered from the perspective of a prototype implementation, and a

wireless backplane with multiple 100 Gbit/s links at 10 cm distance has been analyzed as an extended case with focus on the feasibility and further requirements.

The lessons learned are as follows: Hardware-constraints as the limited resolution of energy-efficient data converters call for system concepts that trade-off digital versus analog signal processing. Modulation and processing schemes than can cope with 1-bit data converters are becoming of significant importance to exploit large bandwidths while still satisfying energy-efficiency constraints. The digital baseband processing will typically be clocked at a rate that is much less than the symbol rate. This requires parallelization or pipelined processing which has to be mapped to the physical layer design. These perceptions do not only apply to the considered LOS scenarios but also to NLOS applications, where analog or digital channel equalization is required under the mentioned constraints.

REFERENCES

- [1] C. Christensen, *The Innovator's Dilemma*. Harvard Business Press, 1997.
- [2] Website: www.fiercebroadbandwireless.com/story/vodafone-no-lte-europe-2010/2009-07-09.
- [3] B. Schulte, M. Peter, R. Felbecker, W. Keusgen, R. Steffen, H. Schumacher, M. Hellfeld, A. Barghouthi, S. Krone, F. Guderian, G. Fettweis, and V. Ziegler, "60 GHz WLAN Applications and Implementation Aspects," submitted to *EuMA International Journal of Microwave and Wireless Technologies. Special Issue on 60-GHz-Communication-Systems*, 2011.
- [4] S. Krone, F. Guderian, G. Fettweis, M. Petri, M. Piz, M. Marinkovic, M. Peter, R. Felbecker, and W. Kesugen, "Physical Layer Design, Link Budget Analysis and Digital Baseband Implementation for 60 GHz Short-Range Applications," submitted to *EuMA International Journal of Microwave and Wireless Technologies. Special Issue on Special Issue on 60-GHz-Communication-Systems*, 2011.
- [5] B. Le, T. W. Rondeau, J. H. Reed, and C. W. Bostian, "Analog-to-Digital Converters," *IEEE Signal Processing Magazine*, vol. 22, no. 6, 2005.
- [6] A. C. Ulusoy, G. Liu, M. Peter, R. Felbecker, H. Y. Abdine, and H. Schumacher, "A BPSK/QPSK Receiver Architecture Suitable for Low-Cost Ultra-High Rate 60 GHz Wireless Communications," *Proceedings of the European Microwave Conference (EuMC10)*, September 2010.
- [7] Standard ECMA-387: High Rate 60 GHz PHY, MAC and HDMI PAL, 2008.
- [8] M. K. Iyer, R. D. Seager, and J. C. Vardaxoglou, "Wireless Chip to Chip Interconnections for Multichip Modules Using Leaky Wave Antennas," *IEEE Electronic Letters*, vol. 29, no. 23, 1993.
- [9] K. Kawasaki, Y. Akiyama, K. Komori, M. Uno, H. Takeuchi, T. Itagaki, Y. Hino, Y. Kawasaki, K. Ito, and A. Hajimiri, "A Millimeter-Wave Intra-Connect Solution," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'10)*, February 2010.
- [10] S. Krone and G. Fettweis, "Achievable Rate with 1-Bit Quantization and Oversampling at the Receiver," *IEEE Communication Theory Workshop (CTW'10)*, May 2010.
- [11] T. Hehn and J. B. Huber, "LDPC Codes and Convolutional Codes with Equal Structural Delay: A Comparison," *IEEE Transactions on Communications*, vol. 57, no. 6, 2009.
- [12] S. V. Maiya and D. J. Costello, "Coding with a Latency Constraint: The Benefits of Sequential Decoding," *Proceedings of the Annual Allerton Conference (Allerton'10)*, September / October 2010.
- [13] A. J. Felström and K. S. Zigangirov, "Periodic Time-Varying Convolutional Codes with Low-Density Parity-Check Matrices," *IEEE Transactions on Information Theory*, vol. 45, no. 5, 1999.
- [14] M. B. S. Tavares, K. S. Zigangirov, and G. Fettweis, "Tail-Biting LDPC Convolutional Codes," *IEEE International Symposium on Information Theory*, June 2007.
- [15] F. Kienle, N. Wehn, and H. Meyr, "On Complexity, Energy- and Implementation-Efficiency of Channel Decoders," submitted to *IEEE Transactions on Communications*, arXiv:1003.3792v1 [cs.IT].
- [16] V. Kursun and E. G. Friedman, *Multi-Voltage CMOS Circuit Design*. John Wiley & Sons, Inc., 2006.
- [17] Website: www.altera.com.
- [18] Website: www.xilinx.com.