G.fast: Evolving the Copper Access Network

Michael Timmers, Mamoun Guenach, Carl Nuzman, and Jochen Maes, Alcatel-Lucent Bell Labs

ABSTRACT

Continuous improvements in communication technologies over copper pair access networks have enabled hybrid fiber/copper access networks that provide ever increasing broadband data rates. Most recently, field trials have demonstrated that crosstalk cancellation, also referred to as vectoring, and bonding two pairs to push data rates of very high speed DSL above 100 Mb/s on a single copper pair and above 200 Mb/s on two copper pairs. Looking ahead, the G.fast project group of the ITU-T is currently defining a new copper access technology to provide aggregate (upstream and downstream) data rates of up to 1 Gb/s over distances up to 250 m. This article provides an overview of some key challenges that need to be addressed for G.fast.

INTRODUCTION

Since their introduction in the early 1990s, wireline broadband networks based on twisted copper pair, coaxial cable, and optical fiber media have continuously evolved to keep pace with consumer demand for reliable high-data-rate connectivity. Numerous deployment models have been introduced to marry the tremendous bandwidth and reach of fiber systems to the compelling economics obtainable by exploiting existing copper infrastructure. Key models being used today and in the foreseeable future include fiber to the node (FTTN), fiber to the distribution point (FTTdp), fiber to the building (FTTB), and fiber to the home (FTTH), where typical distances from the access point to the end user are on the order of 500 m for FTTN, 100 m for FTTdp and FTTB, and 0 m for FTTH (Fig. 1). The term FTTdp can be thought of as including, for example, fiber to the building, curb, front yard, and front door. As the fiber portion of the network gets longer and the copper portion gets shorter, the information carrying capacity of the copper portion increases, motivating development of physical layer technologies capable of exploiting this capacity.

To optimize their access networks, operators want a toolkit of FTTx solutions that allows them to select the fastest and most cost-effective way to connect to a subscriber. For communication over copper pairs, the most recent addition to this toolkit is very high rate digital subscriber line (VDSL2) vectoring, a technology that uses multi-user processing at the access node to suppress crosstalk interference between pairs. The application of vectoring technology significantly boosts the data rates obtainable when using VDSL2 in FTTN or FTTdp networks to well above 100 Mb/s [1]. The success of this new technology is further driving the interest in a copper interface standard tailored specifically for FTTdp and FTTB, G.fast (fast access to subscriber terminals), which aims to deliver up to 1 Gb/s and work over loops up to 250 m in length.

The concept of developing a high bandwidth copper technology specifically for FTTdp was introduced in *IEEE Communications Magazine* in January 2009 [2], sparking substantial research interest from both academia and industry. Initial feasibility studies and market analysis led to the launch of several related standardization activities that are currently investigating the design of the physical layer and the end-to-end system. The G.fast project was initiated at the International Telecommunication Union (ITU) in early 2011 to standardize the physical layer protocol, and around the same time the FTTdp architecture was included in a European Telecommunications Standards Institute (ETSI) technical report on reverse powering of access equipment. These were followed in 2012 by the Broadband Forum's project on FTTdp architectures.

In this article, we give an overview of the research challenges in the physical layer for the copper portion of the next-generation access network. We highlight the need for accurate characterization of the twisted pair and noise environment at high frequencies as critical input for research on modulation and coding design. Next, we discuss the medium access control of G.fast and the impact on vectoring. A section is dedicated to the coding design. We further elaborate the need for a system that takes the practical analog design limitations into account. Finally, we outline challenges of reverse powering in FTTdp networks and present our conclusions.

OBJECTIVES

In this section, we present our views on the main objectives for G.fast, related to high data rates, facilitation of technology deployment, and powering of remote equipment.[1] To be future-proof,

---

[1] *As G.fast is a standardization project in progress at the time of writing, the objectives expressed may not be final; nor may they be supported by all parties contributing to the project.*

0163-6804/13/$25.00 © 2013 IEEE

G.fast targets a rate as high as 1 Gb/s aggregated over upstream (US) and downstream (DS). Although the residential market demand is still far below this limit, it can be expected that data consumption will continue its steady growth, and peak rates of 1 Gb/s will become desired. The rate that G.fast should be able to sustain at all times can be lower than the peak rate, but should exceed 150 Mb/s for all G.fast users. Apart from the residential market, G.fast should aim to serve enterprises and to backhaul mobile data. We have shown through simulations that these incredible speeds are feasible while still being spectrally compatible with VDSL2 (i.e., when using non-overlapping frequency bands as indicated in Fig. 2). Crosstalk cancellation, also called vectoring, remains vital when multiple customers share the same binder [3]. These results were validated using a proof-of-concept platform that has successfully been used to cancel crosstalk up to 212 MHz.

An example of the expected aggregate net data rate for different bandwidth utilizations is shown in Fig. 2 for loops up to 250 m and with a gauge of 0.5 mm (the reader is referred to [3] for details on the assumptions). It shows the feasibility of obtaining up close to 1 Gb/s aggregate over a single short loop on the legacy copper network when a 2.2–106 MHz bandwidth is used. When G.fast is deployed in coexistence with VDSL2 17 MHz lines, the net data rate reduces significantly, but rates well above 500 Mb/s remain obtainable. The curves indicated by “w/o XT” correspond to scenarios where no interference from neighboring lines (crosstalk) is present, or when it is entirely removed through vectoring, assuming perfect cancellation [1, 3]. When a single 99 percent worst case crosstalker is present, the net data rate reduces significantly in the loop length region of interest (0 to 200 m). This indicates that in practice, G.fast vectoring is a required technology to achieve the targeted data rates.

The ITU Telecommunication Standardization Sector (ITU-T) aims to consent to a 106 MHz bandwidth profile in December 2013, with a carrier spacing of 51.75 kHz. This corresponds to a DMT symbol rate of 48 kSymbols/s, a 12-times increase compared to a VDSL2 17 MHz profile. In a later stage, a 212 MHz profile will be defined.

Due to the abundance of distribution point units (DPUs) in FTTdp and FTTB deployments, to which G.fast is tailored, operators want to limit overhead and management of these DPUs (Fig. 1). Therefore, G.fast should enable customer self-install of the customer premises equipment (CPE) and provide an install-and-forget architecture of the DPU, such that a single visit from a technician is required to install the DPU at the distribution point. Hence, sufficient intelligence must be present to enable self-diagnosis and remote management. While mains powering may be readily available in some deployment scenarios such as FTTB, for other scenarios powering the DPU from neighboring network elements is more attractive. This also implies that energy-efficient operation is critical to the success of the G.fast technology, due to the limited power budget when reverse powering.

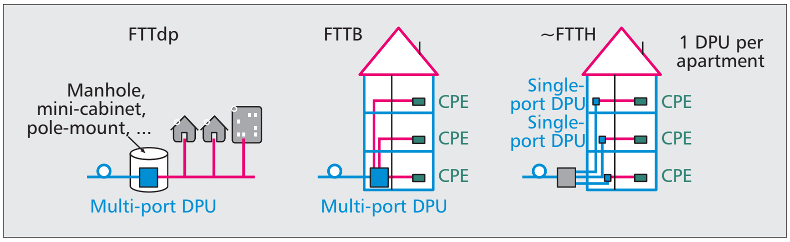

Figure 1. *G.fast is tailored to FTTdP and FTTB deployments where rates up to 1 Gb/s are achievable over copper loops shorter than 250 m.*

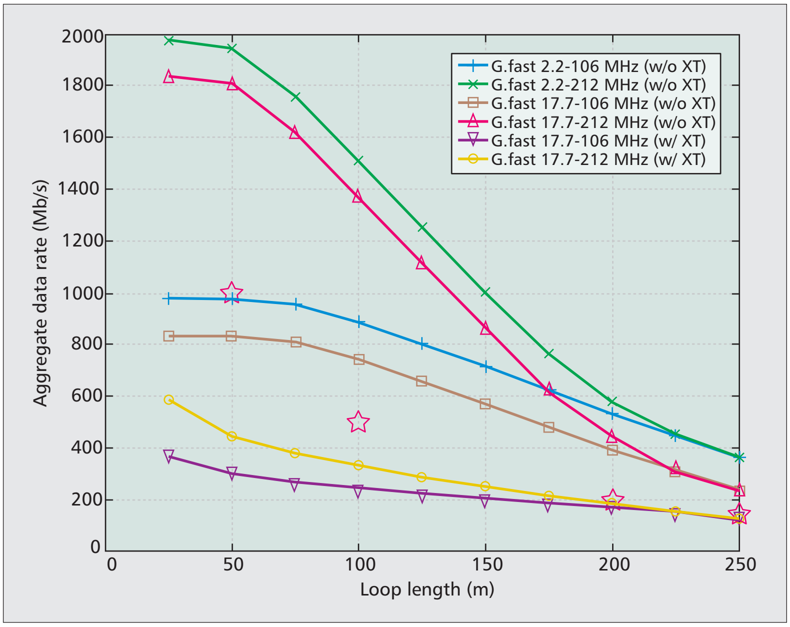

Figure 2. *Gigabit speeds are in reach when vectoring is used to cancel crosstalk interference (agreed G.fast data rate targets are shown as stars).*

CHANNEL CHARACTERIZATION

Before designing a physical layer, the target channels must be well understood and characterized. In order to achieve the desired data rates, it is essential that G.fast must expand the utilized bandwidth by one order of magnitude compared to VDSL2 (Fig. 2). Apart from early exemplary measurements [4], the behavior of the telephony twisted pair channel in the region 30–212 MHz remains uncharted territory. Notable research contributions originate from some operators and research institutes. First, noise ingress and impulsive noise characterization reveal statistics on impulsive noise duration and interarrival times [4–6]. The impulse noise spectral density remains to be studied. Despite these efforts, substantial additional research is needed to establish a reference impulse noise model from an augmented set of measurements. Second, direct and crosstalk channel modeling has led to notable improvements to the direct channel model [7, 8]. Recent measurements have shown deviations from far-end crosstalk (FEXT) models that are extrapolations of existing models for frequencies below 30 MHz. Research is still ongoing on whether these observed deviations are restricted to intra-quad crosstalk, for which the behavior is well under-

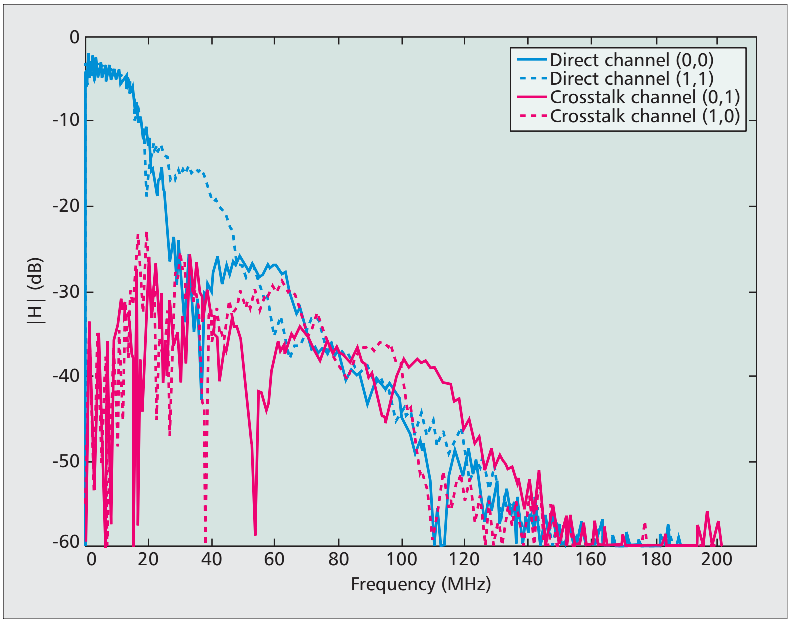

Figure 3. *This measurement shows that from 40 MHz on, crosstalk channels can become dominant with respect to the direct channel (worst case cable).*

stood [7]. The availability of additional measurements will allow creating a wideband DSL channel model. Such a model should be causal to allow time-domain simulations and support high frequencies with few parameters that are easy to fit against measurements. The Broadband Forum has started a project on cable modeling above 30 MHz, in collaboration with ITU-T, ETSI, and the Alliance for Telecommunications Industry Solutions (ATIS).

The channel characterization, modeling, and creation of reference scenarios directly impacts research on physical layer design in general, and the modulation and coding design in particular. By shortening the copper distance and increasing the utilized bandwidth, the DSL crosstalk channels will strengthen relatively to the direct channel, which is shown in the measurement presented in Fig. 3. This strong crosstalk will impact the design of G.fast vectoring, as we discuss later.

Also, in terms of frequency and time selectivity, characterizing the channel is essential in the definition of a suitable modulation to simplify the transceiver operations. Noise characterization is crucial to determine the optimal coding design for reaching a target error probability and better line stability. Furthermore, the exact characterization of the dynamism of the total noise incurred by the lines will determine the speed of the adaptivity in the transceivers to provide a robust service.

MEDIUM ACCESS

The G.fast framing structure is shown in Fig. 4. The first frame of a superframe, which groups eight frames, carries one synchronization symbol for downstream and one for upstream. G.fast uses time-division duplexing (TDD), rather than frequency-division duplexing (FDD) as in VDSL2. This allows the analog front-end to be simplified and accommodates a more flexible allocation of the ratio between the US and DS rates. To avoid near-end crosstalk (NEXT), the US/DS ratio should be the same for all lines served by the DPU, but can be variable in time (e.g., variable per superframe), taking into account the current service requests of all users. Compared to FDD, TDD has a longer round-trip time. However, with a careful choice of the TDD frame length (less than 1 ms), the TDD round-trip time of G.fast should still be shorter than the FDD round-trip time of VDSL2. The number of symbols in a TDD frame should also be variable to allow applications that require small round-trip times, such as mobile backhaul.

Environmental, economic, and political forces are all moving in the direction of accelerating improvements in energy efficiency per bit of communication systems. A key design principle for efficient communication systems is that the energy consumption should be proportional to the information carried; that is, the system should consume much less energy when there is little information being sent. In DSL, a limited form of proportionality has been introduced by introducing low-power modes in ADSL2, which are now also being considered as amendments to the VDSL2 Recommendation. With G.fast, an opportunity exists to incorporate the hooks required for energy proportionality from the start.

The most natural way to introduce energy proportionality to G.fast is for nodes to transmit symbols only when they contain user data, rather than continuously sending symbols containing idle data. In this way, many transceiver elements, including analog-to-digital interfaces and physical layer processing elements, can be put to sleep when symbols are not being sent, making the energy consumption of those elements nearly rate proportional. In G.fast, this mode of operation is referred to as discontinuous operation. In addition, a dynamic resource allocation (DRA) entity will control the transmit opportunities (TxOP — the maximum number of symbols that can be transmitted in one TDD frame) for all lines in both downstream and upstream. By controlling TxOP, the maximum power consumption can be controlled.

To maintain synchronization and exchange control information, it is expected that lines exploiting discontinuous operation will transmit some symbols on a regular basis (e.g., at least one active symbol in each TDD frame in each direction). Two different types of discontinuous operation have been discussed. In static discontinuous operation the number of active symbols transmitted per frame is configured well in advance based on a desired service rate and kept constant over many frames. In dynamic discontinuous operation the number of active symbols may vary from frame to frame based on the amount of downstream data queued at the DPU and upstream data queued at the CPE, further improving the energy efficiency.

To achieve energy proportional operation with discontinuous mode, several challenges need to be addressed. The hardware elements must have low energy sleep states, with sufficiently short transition times. Synchronization and control mechanisms need to be designed to work with a variable number of active symbols per frame. If the protocol calls for the number of active symbols in a frame to be declared in advance, methods for reporting or predicting load levels need to be incorporated. As we discuss later, discontinuous operation will create challenges to cancel crosstalk efficiently.

CROSSTALK CANCELLATION

DEALING WITH STRONG CROSSTALK

Crosstalk cancellation, also called vectoring, actively removes the interference from neighboring lines in DS by adding compensation signals to the desired signal, and in US by combining signals from multiple lines. In VDSL2 systems, which use frequencies below 30 MHz, the crosstalk channel strength is well below the direct channel. Therefore, the power of these compensation signals is small compared to the power of the transmit signal and can be neglected. However, as shown earlier, the crosstalk channels in G.fast are much stronger. Therefore, the power of these compensation signals can no longer be neglected in G.fast.

To stay within the allowed transmit power, a linear precoder needs to reduce the transmit power of the direct signal to compensate for the transmit power of the precompensation signal. This power reduction forms a power penalty that is frequency- and time-dependent, and different for each user line [9]. Multiple ways exist to decrease the transmit power. A fair implementation would give each communication channel an equal transmit budget, regardless of whether this transmit budget is dominated by direct channel components or precoding components [10].

The power penalty can be reduced by nonlinear precoding [9, 11]. A nonlinear precoder limits the transmit power by performing a modulo operation at the transmit and receive sides. This modulo operation limits the peak amplitude per tone of the sum direct signal and the precompensation signal, thereby reducing the average transmit transmitted power per tone. In [11] the nonlinear precoder is constructed using a QR decomposition of the channel matrix, where Q is a unitary matrix and R an upper-triangular matrix. As R is an upper-triangular matrix, the modulo operation can be performed and taken into account by precoding the lines sequentially. After the modulo operation, the signals are rotated using the Q matrix. However, because Q is a unitary matrix, this operation does not increase the transmit power of the signals.

DISCONTINUOUS OPERATION

Vectoring presents particular challenges for discontinuous operation. In discontinuous operation, some lines are active (subset A), while the remaining lines are deactivated (subset D) during a given symbol interval. If vectoring is not used, the transceivers in subset D can simply be put to sleep during this interval. However, in a vectored G.fast system with N lines, the vector of signals to be transmitted on the N lines is multiplied by a precoder matrix, P, normally the inverse of the $N \times N$ channel matrix H. Even if the precoder inputs are zero for lines in subset D, after multiplication by the precoder, the precoder output signal of the deactivated lines will typically be non-zero. As a result, at least the transceiver elements following the precoder (including the inverse fast Fourier transform [IFFT] unit and the analog front-end) need to remain active. This significantly limits the opportunity for energy savings. Alternatively, one could simply mute the precoder output for lines without data to send. This is equivalent to precoding the active lines with a submatrix of the precoder (P_{AA}) and transmitting the resulting signals over a submatrix (H_{AA}) of the original channel. For the high crosstalk channels typical of G.fast subcarrier frequencies, the precoder submatrix is typically a poor approximation to the inverse of the channel submatrix. Therefore, this approach results in significant residual crosstalk and loss of data rate [12–14].

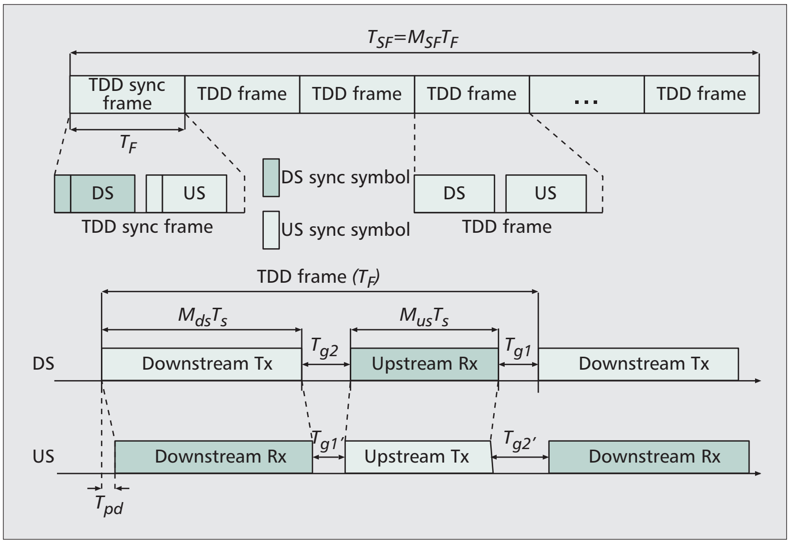

Figure 4. *The superframe and frame structure currently proposed for G.fast.*

In order to fully suppress crosstalk interference while still allowing the deactivated lines to sleep, the precoder used when set A is active must match the inverse of the submatrix HAA. Several techniques for obtaining this inverse efficiently are under study. These include using padding to limit the number of different subsets used [14] and applying correction terms derived from the full $N \times N$ precoder [12, 13].

IMPACT OF ANALOG LIMITATIONS ON MODULATION

While the increase in digital complexity of consecutive DSL standards have stayed within bounds derived from Moore's law, the bandwidth demand in G.fast requires an analog design that overcomes the slower pace at which analog technology advances. Traditionally, DSL uses FDD of US and DS traffic, a high spectral efficiency of 15 b/s/Hz, high aggregate transmit power of up to 20.5 dBm, and the use of frequencies down to voice bands (25 kHz). However, these characteristics are not cost- and energy-effective for G.fast with bandwidths up to 212 MHz. Therefore, the G.fast analog front-end and system architecture should be designed to overcome the analog limitations with minimal sacrifices.

First, G.fast adopts TDD of US and DS traffic, which, among other effects, reduces the need

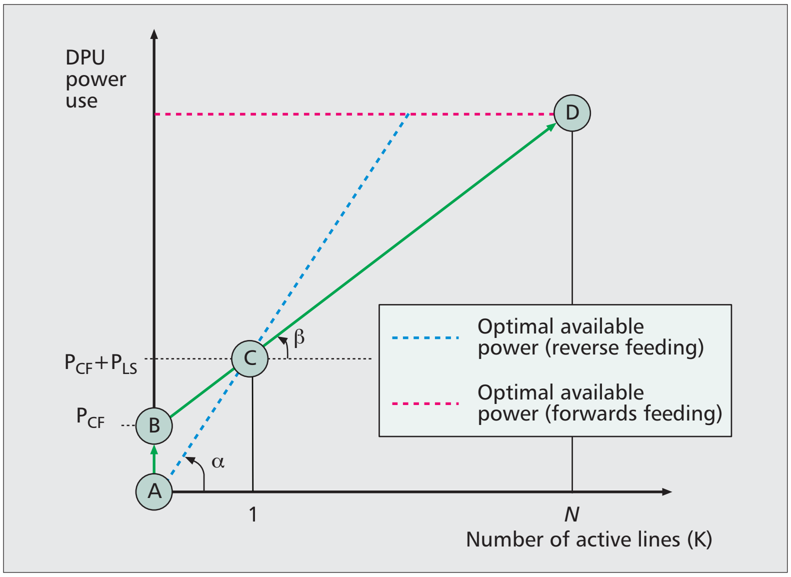

Figure 5. *When reverse feeding a G.fast DPU, the constraint is the power use when one port is active: (A) when not receiving external power; (B) when powered, but all ports disabled; (C) with one active port; (D) when all ports are active.*

for a hybrid between the line interface and the analog-to-digital converter (ADC).

Second, it has been shown that a spectral efficiency of 12 b/s/Hz, under typical G.fast loop conditions (e.g., a maximum power spectral density of –76 dBm/Hz), achieves similar data rates compared to a maximal efficiency of 15 b/s/Hz. This lower spectral efficiency allows reducing the required dynamic range of the ADC, and therefore its complexity and power consumption.

Third, not using frequencies below 2.2 MHz reduces the linearity constraint of the analog components by nearly two orders of magnitude. As G.fast aims to use bands up to 212 MHz, dropping these asymmetric DSL (ADSL) bands will not significantly impact the achievable data rates.

Hence, we can conclude that the analog front-end will become the bottleneck for G.fast and that the system architecture should take this limitation into account. Further research should target energy-effective analog hardware and optimize the trade-offs between analog complexity and system architecture.

CODING

A good design of coding starts from a profound understanding of the channel and noise environment. The DSL direct channel transfer function is relatively stable and does not suffer from fading as in wireless. This reduces the need for error control schemes with low code rate that have proven to be robust in the wireless context. Among the dominant DSL noise impairments are transient noises from alien disturbers, in-house generated impulsive noises, and radio frequency interference (RFI). In a first generation of VDSL2, lines were protected for a scenario close to the worst case. This is achieved by masking changes in the noise level through the configuration of a large signal-to-noise ratio (SNR) margin, and configuration of conservative impulse noise protection (INP) settings, leading to large interleaving depth and Reed Solomon forward error correcting (FEC) overhead. This static approach helps to ensure stable reliable communication sessions, but at the expense of significant data rate penalties. One important research challenge is to avoid an overly protective DSL configuration by designing the physical layer with flexible and dynamic coding and modulation that can provide the protection needed for the encountered scenario at minimal overhead. To offer better INP at minimal overhead, automated retransmission request (ARQ) was added to the ADSL2 and VDSL2 technologies through ITU-T Recommendation G.998.4. Vectored VDSL2 uses retransmission by default.

Two FEC codes were evaluated in ITU G.fast with the goal to select a single coding scheme for inclusion in G.fast. Both schemes use an outer and an inner code, and rely on ARQ for protection against impulsive noise as part of the outer error control. The incoming information bits are grouped into data transfer units (DTUs), which includes means for error detection with a cyclic redundancy check (CRC). Each DTU, after being transmitted, is then stored in a buffer for possible retransmission until the receiver acknowledges its correct reception. The outer error control at the receiver makes use of the CRC to detect whether a DTU has been received in error, and to initiate a retransmission to recover the erroneous DTUs. In combination with FEC, a DTU comprises an integer number of FEC codewords.

The first proposal was based on the VDSL2 coding scheme in which an FEC codeword consists of a block-interleaved Reed-Solomon (RS) code. The inner code consists of trellis coded modulation (TCM), and is a combination of a binary convolutional code and a proper mapping onto an quadrature amplitude modulated constellation of size *M* (*M*-QAM).

The second proposed coding scheme consisted of a quasi-cyclic low density parity check block code (QC-LDPC-BC) that has proven its merits in the in-home network technology ITU G.9960 (a.k.a. G.hn).

The two proposals were evaluated in realistic G.fast environments, and the RS+TCM coding scheme was selected due to lower complexity at an acceptable performance loss in white noise conditions compared to QC-LDPC-BC.

REVERSE POWERING

Both the FTTdp and FTTB scenarios to which G.fast is tailored lead to a multitude of small nodes, which must all be powered. Relying on availability of local powering in all those in-the-field locations can lead to lengthy and complex negotiations with utility companies and building owners. On the other hand, remote powering of the DPUs from the central office over unused twisted pair is a well established method, but still requires the presence of twisted pairs, and non-negligible dissipation exists due to the large distance to be bridged (up to several kilometers). An alternative powering approach is to provide the power from the end-user premises by means of a dedicated power injector at each home network, which may or may not be integrated with the G.fast CPE. The DPU will then be reverse-powered over the active lines.

Although the concept of reverse powering is conceptually simple, in practice several technical challenges need to be addressed. A first challenge is the powering itself. On one hand, the power that the DPU can receive from the lines is limited by the fact that the power injector must comply to in-home safety rules in terms of maximal voltage and current (Safety Extra-Low Voltage), and further reduced by some dissipation in the line. On the other hand, the DPU must remain operational even in the event of a single active line. It becomes clear that the power consumption in the DPU must be as low as possible and scale as linearly as possible with the number of active lines, which we illustrate in Fig. 5 [15]. When reverse powering the DPU, the bottleneck point shifts from (D) to (C), where a CPE needs to power its line-specific functionality, as well as all common functionality. We have demonstrated that a 16-port VDSL2 unit can be reverse powered over a distance of 250 m, even with a single line active, while complying with safety requirements that can be expected in a home environment.

A second challenge is posed in case the G.fast service needs to coexist with legacy plain old telephone service (POTS) on the same access line. To make the reverse powering compatible with POTS, the POTS DC powering and low-frequency signaling have to be decoupled from the line. An additional requirement for POTS lifeline service would add extra complexity. An alternative to POTS is overlay using VoIP up to the in-house gateway and POTS phones directly connected to the gateway (or DECT). ETSI is currently working on voltage and current limits and related aspects for reverse powering of FTTdp equipment.

CONCLUSIONS

In this article we have shown that the currently ongoing standardization of G.fast access technology needs to go hand in hand with research and innovation to overcome emerging challenges. Research opportunities have been identified in the modeling of cable characteristics and noise environment, and the impact of high crosstalk and discontinuous operation on crosstalk cancellation. We indicate how the analog domain has become the predominant bottleneck for G.fast and how the design of the digital modulation can relax constraints on the analog front-end. We have shown the adopted research methodology used in the standard to select the optimal coding scheme. The need for reverse powering in the FTTdp and FTTB deployment scenario is discussed. Digital subscriber line technology remains an important access technology and should leverage innovative ideas to maintain this position. We predict that this trend will continue over the coming years.

Digital subscriber line technology remains an important access technology and should leverage innovative ideas to maintain this position. We predict that this trend will still continue over the coming years.

ACKNOWLEDGMENTS

The authors would like to thank Francois Fredricx, Bart Hillaert, Paul Spruyt, Danny Van Bruyssel, and Frank Van der Putten for their contributions and careful review.

REFERENCES

[1] V. Oksman *et al.*, “The ITU-T’s New G.vector Standard Proliferates 100 Mb/s DSL,” *IEEE Commun. Mag.*, Oct. 2010.

[2] P. Ödling *et al.*, “The Fourth Generation Broadband Concept,” *IEEE Commun. Mag.*, Jan. 2009, pp. 63–69.

[3] J. Maes *et al.*, “Pushing the Limits of Copper — Paving the Road to FTTH,” ICC’2012.

[4] T. Magesacher *et al.*, “Modeling and Measurement of Short Copper Cables for Ultra-Wideband Communications,” *Proc. SPIE OpticsEast Broadband Access Commun. Tech.*, Boston, MA, 2006.

[5] H. Mariotte, “G.fast: Background Colored Noise Measurements,” ITU-T SG15 Q4 Contrib. 11RV-031, Nov. 2011.

[6] T. Morsman, “G.fast Impulsive Noise Characteristics,” ITU-T SG15 Q4a Contrib. 2012-07-4A-024, July 2012.

[7] R. van den Brink and B. van den Heuvel, “G.fast: Wideband Modelling of Twisted Pair Cables as Two-Ports,” ITU-T SG15 Q4 Contrib. 11GS3-028, Sept. 2011.

[8] E. Berndt and P. Muggenthaler, “G.fast: EL-FEXT Analysis,” ITU-T SG15 Q4a Contrib. 2012-06-4A-041, June 2012.

[9] Ikanos Communications, “G.fast Comparison of Linear and Non-Linear Pre-coding for G.fast on 100 m BT Cable,” ITU-T Contrib. 2013-01-Q4-031, Geneva, Switzerland, Jan. 2013.

[10] Alcatel-Lucent, “G.fast: Comparison of Linear and Non-Linear Pre-Coding,” Contrib. ITU-T 2013-01-Q4-046, Geneva, Switzerland, Jan. 2013

[11] G. Ginis and J. M. Cioffi, “A Multi-User Precoding Scheme Achieving Crosstalk Cancellation with Application to DSL Systems,” 34th Asilomar Conf. Signals, Sys., and Computers, Oct. 29–Nov. 1, 2000.

[12] Alcatel-Lucent, “G.fast: Solutions for Precoding in Discontinuous Operation,” ITU-T SG15 Q4 Contrib. 2013-03-Q4-052, Red Bank, NJ, Mar. 2013.

[13] Lantiq, “G.fast: Precoder Update in Support of Discontinuous Operation,” ITU-T SG15 Q4a Contrib. 2013-01-Q4-068, Geneva, Switzerland, Jan. 2013.

[14] Sckipio, “Power Saving Implications on Vectoring — Static Allocation Case.” Contrib. ITU-T SG15/Q4 2012-11-4A-043, Chengdu, China, Nov. 2012.

[15] M. Timmers *et al.*, “System Design of Reverse-Powered G.fast,” IEEE ICC, Wksp. Telecommun.: From Research to Standards, 2012.

BIOGRAPHIES

MICHAEL TIMMERS (Michael.Timmers@alcatel-lucent.com) is a senior research engineer in Bell Labs’ Access Research Domain in Antwerp, Belgium. He received his M.S. in engineering from Katholieke Universiteit, Leuven, in 2005, and obtained his Ph.D. in 2009 at the Interuniversitary Micro-Electronics Center (imec) in Belgium. His doctoral thesis was on dynamic and cognitive spectrum access. For the past three years, his research has targeted next-generation fixed access networks with a focus on DSL.

MAMOUN GUENACH (Mamoun.Guenach@alcatel-lucent.com) is a research scientist with the Alcatel-Lucent Bell Labs Access Node Technology & DSL team in Antwerp, Belgium. He received his degree of engineer in electronics and communications from the Ecole Mohamadia d’Ingénieurs in Morocco. Following that, he moved to the Faculty of Applied Sciences at the Universite Catholique de Louvain (UCL), Louvain-La-Neuve, Belgium, where he received an M.S. degree in electricity and a Ph.D. degree in applied sciences. He served as a post-doctoral researcher in the Department of Telecommunications and Information Processing at the Universiteit Gent, Belgium, prior to joining Alcatel-Lucent.

CARL NUZMAN (Carl.Nuzman@alcatel-lucent.com) received Bachelor’s degrees in electrical engineering and mathematics at the University of Maryland College Park in 1996, and a Ph. D. in electrical engineering from Princeton University in 2000. He is currently a Distinguished Member of Technical Staff in the Enabling Computing Technologies Department at Bell Labs, Alcatel-Lucent, located in Murray Hill, New Jersey. Since joining Bell Labs in 2000, he has performed research in fields of probability, optimization, and dynamical systems, with applications to optical networking, micro electro-mechanical systems, power control, traffic modeling, and digital subscriber line systems.

JOCHEN MAES (Jochen.Maes@alcatel-lucent.com) received his Master’s degree in physics and Ph.D. degree in science from the Katholieke Universiteit Leuven in 2000 and 2004, respectively. He is a member of technical staff in the Access Research Domain of Bell Labs. He manages the research and standardization of copper access technology.