

The ITU-T's New G.fast Standard Brings DSL into the Gigabit Era

Vladimir Oksman, Rainer Strobel, Xiang Wang, Dong Wei, Rami Verbin, Richard Goodson, and Massimo Sorbara

The standardized G.fast transmission method and advanced crosstalk cancellation techniques are presented by the authors, with specific performance projections and measurement results achieved during the first demonstrations and trials, showing bit rates of 500 Mb/s over 250 m and available reach up to 400 m.

ABSTRACT

This article explores the recently issued ITU-T Recommendations specifying “G.fast” (G.9701 [1] and G.9700 [2]) that bring user bit rates up to 1 Gb/s over twisted pairs from the distribution point to customer premises. The overview and some key research challenges of G.fast are discussed in [3]. The standardized G.fast transmission method and advanced crosstalk cancellation techniques are presented here with specific performance projections and measurement results achieved during the first demonstrations and trials, showing bit rates of 500 Mb/s over 250 m and available reach up to 400 m. A description of standardized tools for dynamic performance maintenance, resource allocation, and power saving enhancing G.fast applications concludes this article.

INTRODUCTION

Modern life depends on the Internet, and thus the demand for high-speed Internet access is rapidly growing. Digital subscriber line (DSL) technology, accordingly, keeps up with both customer demand and the progress in competing access technologies, such as DOCSIS, WiMAX/Long Term Evolution (LTE), and gigabit passive optical networking (G-PON). In 2010, the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) developed Recommendation G.993.5 [4], which set a 100 Mb/s benchmark in DSL services [5]. The new G.fast Recommendations [1, 2] specify 1 Gb/s access over copper.

To reach such high bit rates, G.fast uses only the last leg of the existing copper access network and in-premises wiring. These wires are usually unshielded, non-conditioned twisted pairs, flat pairs, or quads (four twisted wires) and known for very strong crosstalk, especially inside quads [3]. Reaching high bit rates over such low-quality copper is a difficult task that requires a substantially new approach. The main challenges encountered by engineers and the potential technical choices are discussed in [3], while this article describes the adopted technical solutions.

The G.fast-based access network uses the fiber-to-the-distribution-point (FTTdp) architecture [6], which comprises a distribution point

unit (DPU) connected to the central office (CO) by fiber (PON or point-to-point fiber). DPUs are installed close to the customer premises, typically in mini-cabinets mounted in basements of multi-dwelling units, on electrical poles, in curb boxes, or in manholes [3, 6], and connected to customer premises equipment (CPE) via copper pairs. A DPU typically serves 4–20 lines, but bigger DPUs are expected in the future. A single-line DPU may serve as a fiber-to-the-home (FTTH) copper extension. DPUs can be powered locally, remotely, or by subscribers from the customer premises using reverse power feeding (RPF) [7]; the latter is very convenient for small DPUs. The achievable bit rate over a particular line depends on its length and wire type. The maximum reach is 400 m, but the majority of installations are expected to be within 100 m.

An example of a typical G.fast installation using RPF is shown in Fig. 1. The G.fast transceivers in the DPU (FTU-O) and in the CPE (FTU-R) are connected via a copper pair; the FTU-R resides in the network termination unit (NTU). The broadband services delivered to the DPU via a PON feeder are conveyed to the NTU and further distributed to various broadband applications via a high-speed in-premises network (e.g., WiFi). The RPF power sourcing equipment (PSE) in the NTU generates sufficient power to supply the associated FTU-O and common functions of the DPU via the copper pair. The DPU power supply unit (PSU) gathers the power sourced by PSEs of all active lines through corresponding power extractors (PEs). In other installations, the PSE may be separate from the NTU and also feed the NTU. The PSE can work during power outages using a backup battery. More RPF details can be found in [3, 7].

Analog phones are connected through adapters because RPF uses DC: the Foreign Exchange Office (FXO) adapter receives plain old telephone service (POTS) signaling derived from voice over IP (VoIP) service by the analog telephone adapter (ATA) and generates alternative signaling capable of running over in-premises wiring with RPF; the Foreign Exchange Subscriber (FXS) adapter further recovers the original POTS signaling. Both the FXO and FXS

Vladimir Oksman and Rainer Strobel are with Lantiq, an Intel Company; Xiang Wang and Dong Wei are with Huawei Technologies; Rami Verbin is with Skopio Technologies; Richard Goodson is with AADTRAN Inc.; Massimo Sorbara is with Qualcomm.

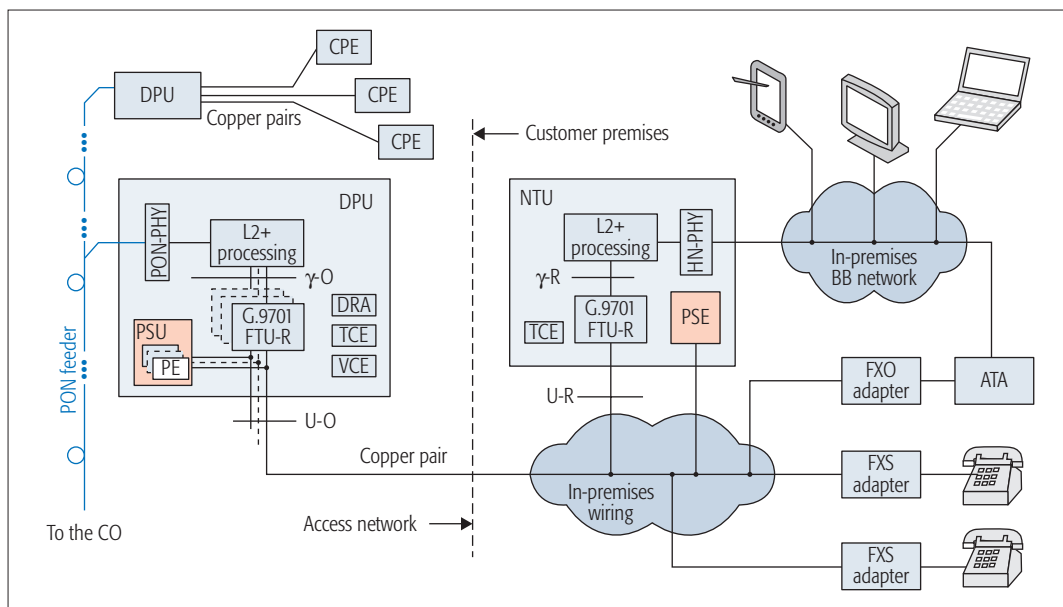


Figure 1. Example of G.fast deployment using RPF and derived POTS.

G.fast is a state-of-the-art copper access technology that provides fiber-grade transmission speed over existing copper, minimizes energy consumption, reduces maintenance cost, and provides great robustness and flexibility for the customers.

are fed by the PSE. To avoid G.fast performance loss, no other devices should be connected to the in-premises wiring.

In other installations, the ATA may reside in the NTU, while derived voice service can be distributed throughout the premises via cordless phone technology or by using smartphone connection to in-premises WiFi; no particular option for voice distribution is implied by G.fast.

G.9701 uses the frequency spectrum from 2.2 to 106 MHz with full crosstalk cancellation between the lines sourced by a DPU: near-end crosstalk (NEXT) is avoided by using synchronized time-division duplexing (STDD), and far-end crosstalk (FEXT) is cancelled using vectoring. No alien crosstalk cancellation is defined. Other G.fast innovations include dynamic allocation of resources between DPU sourced lines, efficient energy-saving techniques, and dynamic performance maintenance. Pair bonding is defined to allow multiplication of customers' bit rate.

G.fast customer installations vary by signal attenuation and noise, and may be influenced by other technologies. For reliable self-installation provisioning, which reduces the operator's cost, G.fast defines a flexible and robust transmission protocol, online reconfiguration, and dynamic adaptation of bit rate, all maintained via robust management channels. Zero-touch management further reduces the operator's cost by avoiding truck rolls for future equipment upgrades and adding new subscribers.

G.fast is a state-of-the-art copper access technology that offers fiber-grade transmission speed over existing copper, minimizes energy consumption, reduces maintenance cost, and provides great robustness and flexibility for customers.

FUNDAMENTALS OF G.9701 TECHNOLOGY

TRANSMISSION METHOD

G.9701 specifies the functionality of G.fast transceivers (FTU-O and FTU-R) that establish a high-speed transmission path between γ -O and γ -R reference points (Fig. 1). The user's data

packets from upper layers (L2+) are mapped into data transmission units (DTUs) that are conveyed transparently over the line. Reed-Solomon forward error correction [8] improves noise immunity: each DTU is assembled from multiple Reed-Solomon codewords, and DTU bytes are interleaved. The number of codewords and their size are configured to fit the throughput of the line. Noise is further mitigated by retransmission of DTUs received in error; the number of retransmissions for a DTU is limited by the latency bound. For retransmission and latency control, each DTU contains a sequence number and associated timestamp.

Discrete multi-tone (DMT) modulation [8] is used for passing DTUs and management data over the line. The advantages of DMT are well known, especially its capability to operate on lines with multiple bridged taps such as in-premises wiring. The specified tone spacing of 51.75 kHz is 12 times that of very-high rate DSL 2 (VDSL2) [8] because G.fast loops are much shorter. Thus, 2048-tone DMT is sufficient to cover the current G.fast frequency spectrum, simplifying the design. Each DMT symbol is cyclically extended using both prefix and suffix. The prefix mitigates inter-symbol interference and is configurable to address a wide range of loop lengths. The suffix is applied for transmit spectrum shaping and overlaps with the following symbol to improve efficiency; suffix size always fits the size of the windowing [8]. The default cyclic extension yields a symbol duration of 20.83 μ s. Up to 14 bits can be loaded per tone. To further increase bit rates, future versions of G.fast may extend the frequency spectrum to 211.968 MHz.

The advantages of the G.fast duplexing scheme (STDD) compared to FDD are described in [3]. With STDD, the upstream and downstream sets of DMT tones can be selected independently, making STDD flexible in the frequency domain. A particular selection depends on the deployment scenario, and involves channel characteristics and spectrum compatibility issues (see below).

To minimize the performance loss, downstream transmit PSDs are optimized across all the lines with precoder updates, including transmit power reduction of tones causing high crosstalk.

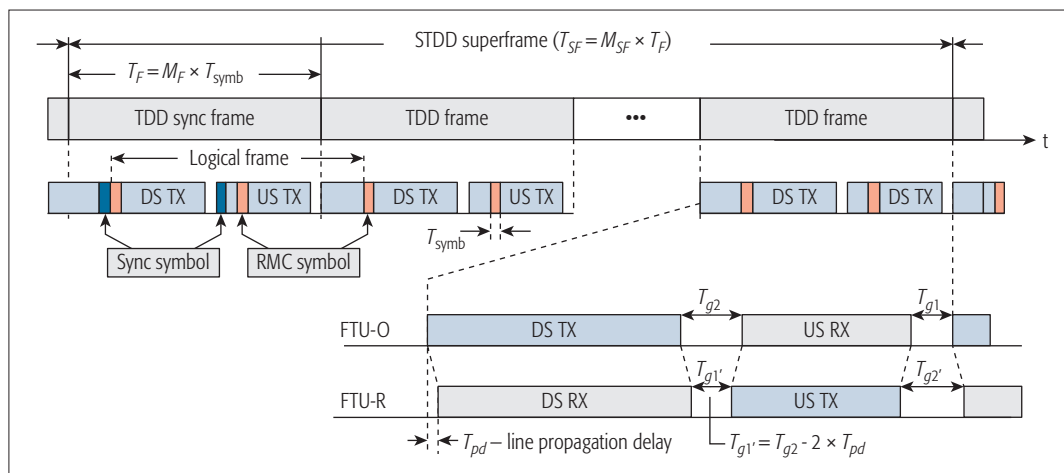


Figure 2. G.fast transmission format.

The G.fast transmission format comprises superframes, each composed of M_{SF} TDD frames (Fig. 2). Each TDD frame contains M_F symbol periods (T_{symb}). One set of contiguous symbol periods is assigned for downstream transmission and another one for upstream transmission. The sum of guard times between upstream and downstream transmissions ($T_{g1} + T_{g2}$) is one symbol period.

Superframes follow each other with no gaps; their boundaries are identified by downstream sync symbols. Both downstream and upstream sync symbols reside in a TDD sync frame, and carry probe sequences used for channel estimation and other purposes (described below).

The maximum duration of a TDD frame is bounded by the propagation delay limit to $M_F = 36$ symbol periods. A setting of $M_F = 23$ reduces round-trip delay. A superframe contains 8 and 12 TDD frames, respectively, so its duration is always about 6 ms, which allows the superframe period to be used as a time base for initialization and management procedures.

The transmission path is maintained by the embedded operations channel (eoc) and robust management channel (RMC). The eoc is multiplexed into DTUs; it has a flexible bit rate that can support high-volume management data, but its robustness is about the same as user data. The RMC, in contrast, is defined to carry short messages and is much more robust due to high-redundancy Reed-Solomon coding, tone selection, and conservative bit loading. Multiple repetitions are applied for critical RMC commands.

One RMC symbol per direction is sent in each TDD frame (Fig. 2). It carries the RMC on dedicated tones and DTU bytes on other tones. The positions of RMC symbols in a TDD frame and the sets of RMC tones are configured at initialization. Symbol positions from one to the next RMC symbol of the same direction represent a logical frame. The RMC carries acknowledgments of received DTUs, supporting DTU retransmission, and conveys management commands facilitating logical frame configuration, online reconfiguration (OLR), and transitions into and out of low-power states.

With STDD, the time positions of superframes, TDD frames, sync symbols, and RMC symbols are aligned across all lines sourced by

a DPU (vectors group). The alignment is by symbol boundaries, and only a small deviation is tolerable to avoid NEXT and facilitate FEXT cancellation, discontinuous operation, and fast reconfiguration.

FEXT CANCELLATION

FEXT cancellation is imperative for reaching high bit rates. Similar to G.993.5, G.fast performs FEXT cancellation at the DPU: the downstream transmit signals are precoded by adding FEXT pre-compensation signals, and a post-processor subtracts FEXT components from the received upstream signal [3, 5]. The vectoring control entity (VCE) at the DPU performs channel estimation, and computes precoder and post-processor matrices for all connected lines. The particular methods of channel estimation, matrix computation, and FEXT cancellation are vendor discretionary. For downstream channel estimation, the VCE may assign the same or different precoder matrices for sync symbols and data symbols (including the use of non-precoded sync symbols).

Like G.993.5, G.fast uses linear precoding. However, the FEXT behavior in G.fast is fundamentally different, especially for quad-twisted cables, due to G.fast's much wider frequency spectrum. Figure 3 shows a typical G.fast FEXT channel: the in-quad crosstalk can be stronger than the direct channel at high frequencies. This complicates precoding because the added pre-compensation signals can substantially increase the transmit PSD. Thus, with a given PSD limit, pre-compensation signals associated with a line generating high crosstalk can suppress the power of the direct signal in other lines of the vectored group, causing substantial performance loss [3].

To minimize this performance loss, downstream transmit PSDs are optimized across all the lines with precoder updates, including transmit power reduction of tones causing high crosstalk. A transmitter-initiated gain adjustment (TIGA) is used to accommodate the change of precoder gain in the peer FTU-R receiver. Figure 3 shows that PSD optimization substantially improves the achievable signal-to-noise ratio (SNR). It may even exceed the single-line SNR due to




Figure 3. Left: direct and crosstalk channels of PE 0.5 mm quad cable (Germany); right: achievable downstream SNR.

additional direct signal propagation via FEXT channels. The latter also has a negative effect: turning off a line in a vectored group requires a precoder update and likely changes performance of the other lines.

CHANNEL ESTIMATION

Channel estimation is necessary to compute the FEXT channel matrix, and to derive precoder and post-processor matrices. Similar to G.993.5, probe sequences are carried by upstream and downstream sync symbols (Fig. 2). The values of the received probe signals are delivered to the VCE, which computes the channel matrix. The VCE assigns appropriate probe sequences (orthogonal, like Walsh-Hadamard sequences, or pseudo-orthogonal) to the lines of the vectored group using its own channel estimation strategy. Probe sequences are repeated periodically, allowing efficient averaging for noise mitigation. Unlike G.993.5, probe sequences may include 0-elements (no transmission of a sync symbol). By using 0-elements, the vectored group can be virtually divided into sub-groups; this reduces the aggregate FEXT inside each sub-group and speeds up channel estimation.

The assigned probe sequences are communicated to the FTU-R at initialization and can be updated via eoc during showtime (the state when user data is communicated). The VCE may request of the FTU-R to report DFT-output samples that represent the received probe signal in the frequency domain, or error samples that represent a normalized error vector between this received probe signal and the associated reference constellation point. The FTU-R uses the communicated downstream probe sequence to identify this constellation point, since the error may be comparable or even exceed the received signal due to strong FEXT.

The FTU-R report (vectoring feedback) is sent to the VCE via the eoc during showtime and via the special operations channel (SOC) during initialization. If available upstream capacity is insufficient, tone interpolation and decimation

of the feedback in time and frequency are used, spreading the transmission over several probe sequence cycles.

OPERATION OF A VECTORED GROUP

Vectored group operation comprises three phases: tracking, joining, and leaving. During tracking, all lines of the group are in showtime: lines neither join nor leave the group, and each line tracks channel variations caused mainly by temperature changes and discontinuous operation. The latter may require frequent precoder updates facilitated by OLR procedures (TIGA, SRA, see below).

In the joining phase, new lines are added to the vectored group. During initialization of new lines, channel matrices, precoders, and transmit PSDs of new and showtime lines are jointly optimized, improving overall performance. This also requires multiple OLR procedures and high-volume transfers of vectoring feedback in showtime lines; both are supported by the eoc.

Lines can leave the group in an orderly or disorderly manner. Orderly leaving first terminates transmission in both directions, then updates the channel matrices of the remaining lines, allowing the FTU-R to then safely disconnect. A disorderly-disconnected FTU-R (e.g., unplugging) usually disturbs other lines substantially due to associated changes in their direct channel and increase of residual FEXT, which is obviously undesirable. Fast rate adaptation (FRA) and retransmission help mitigate error bursts until channel matrices are updated; the performance is restored by associated OLR procedures.

Details of the joining procedure are shown in Fig. 4. After a G.994.1 handshake [9], during which the two sides exchange capabilities, agree on a common operational mode, and set necessary parameters to facilitate STDD, the FTU-O starts training by transmitting superframes containing only sync symbols modulated by a probe sequence (during O-VECTOR-1). This allows the VCE to learn and cancel downstream

During the channel analysis and exchange phase, FTUs establish their desired showtime settings, such as bit loading, DTU size, and RMC tone sets. After CA&E, lines transition into showtime. The expected joining time of a single line is significantly less than in VDSL2 due to the shorter probe sequence cycle.




Figure 4. Joining timeline and examples of precoder and post-processor conversion in a 16-line DPU for tone #1160 (60 MHz).

crosstalk from joining lines into showtime lines without disturbing showtime lines. After the precoders of showtime lines are updated and the downstream crosstalk from joining lines is cancelled, the FTU-O turns on the downstream SOC and sends to the FTU-R the necessary upstream initialization data. Since crosstalk between joining lines is not cancelled, data transmitted over the SOC is scrambled using a unique scrambling seed in each joining line to avoid reception from a non-peer FTU-O. Furthermore, to improve robustness, the data is transmitted using repetitions and is modulated by an ID-sequence, which is orthogonal relative to the ID-sequences of other joining lines.

During R-VECTOR 1, the FTU-R transmits sync symbols modulated by a probe sequence, and the VCE estimates the upstream channel. After upstream crosstalk between all joining and showtime lines is mutually cancelled, a high-speed upstream SOC is established to convey vectoring feedback for precoder training (during O-VECTOR-2). After O-VECTOR-2, the downstream crosstalk between all joining and showtime lines is also cancelled.

During PRM-UPDATE, both FTU-O and FTU-R optimize their transmit PSDs in conditions when crosstalk is cancelled. One goal of optimization is to reduce the transmit power on tones with extra SNR margin. Another is to suppress tones generating very high crosstalk (to avoid performance loss in showtime lines). Other criteria, such as total power reduction, may also be applied [10].

During the channel analysis and exchange (CA&E) phase, FTUs establish their desired showtime settings, such as bit loading, DTU size, and RMC tone sets. After CA&E, lines transition into showtime. The expected joining time of a single line is significantly less than in VDSL2 due to the shorter probe sequence cycle.

POWER SAVING

Low power consumption is vital for G.fast, driven by limited heat dissipation, remote/reverse powering of the DPU, and battery-fed opera-

tion of the CPE. Power-saving mechanisms are discontinuous operation (DO) and low-power states.

Discontinuous Operation: DO scales a transceiver's power consumption with actual data throughput by transmitting only when data is available. During the remaining (quiet) symbol slots, essential analog and digital processing may be turned off, bringing substantial power savings.

In a vectored group with high crosstalk, turning off one line may change the direct channel of other lines, causing performance degradation. Therefore, strict coordination of turning slots quiet is applied across all vectored lines. Specifically, each logical frame is divided into a normal operation interval (NOI, the first TTR slots) and a discontinuous operation interval (DOI, the remaining slots). During NOI, no quiet slots are allowed, while in DOI the first TA slots are quiet, and the remaining slots may be quiet if no user data is available or active otherwise. In some NOI and DOI slots with no data available, an FTU may transmit only pre-compensation signals (idle slots). The DO parameters TTR, TA, and total number of active slots (TBUDGET, Fig. 5) are determined by the DRA and VCE based on DPU dynamic resource allocation (see below). They are configured per logical frame and coordinated across the vectored group via the RMC. Bit loadings and transmit PSDs during NOI and DOI are updated independently using OLR procedures. Figure 5 shows an example of downstream DO in a four-line DPU. The NOI includes five slots (TTR = 5), and values of TA are set so that during DOI only one line is transmitting at a time, facilitating crosstalk avoidance; no transmission in other lines and reduced vector processing saves power.

Power Saving States: During low-power states, FTUs save power by transmitting only sync symbols and RMC symbols in a few assigned TDD frames; all other slots are quiet. Furthermore, the transmit PSD and number of active tones in

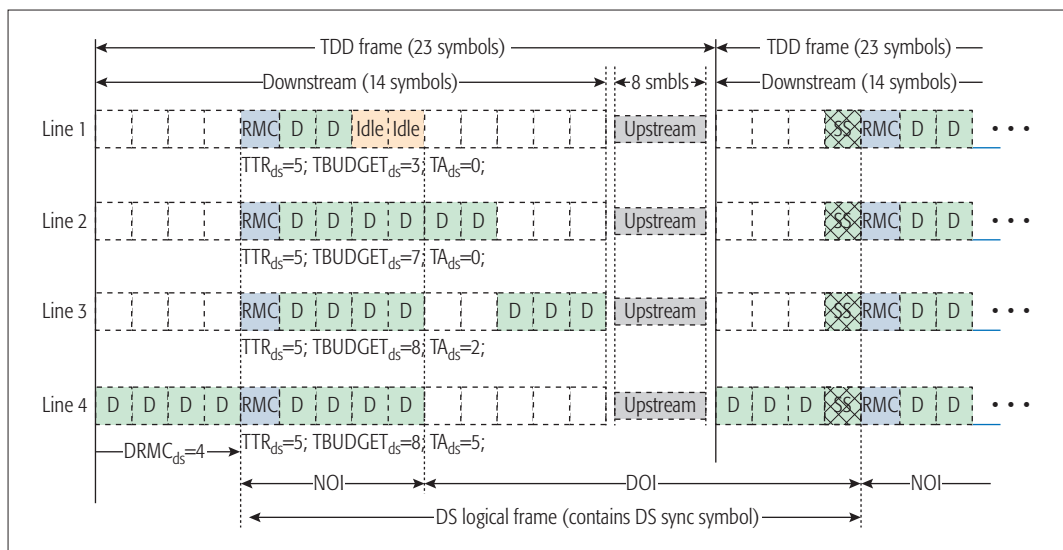


Figure 5. Example of downstream DO in a 4-line group (23-symbol TDD frame).

RMC symbols may be reduced. Two low-power states called L2.1 and L2.2 are defined. L2.1 efficiently saves power when broadband services are off, while continuing VoIP services. When broadband service restarts, the line transitions back to normal in less than 1 s.

In case of a power outage at a customer premises equipped with battery backup, L2.2 maintains only keep-alive traffic (no broadband or VoIP services). Keep-alive is very slow, but provides fast recovery of services. An NTU in L2.2 also detects incoming phone calls and temporarily transitions into L2.1 to support them. When the power is restored, the line moves into L2.1 (to support phone calls) and is ready to restart broadband service.

Transitions between states are triggered from the upper layers of the DPU, and by the DRA that monitors broadband traffic and CPE battery status. They are facilitated by eoc/RMC commands, similar to OLR transitions.

DYNAMIC PERFORMANCE MAINTENANCE

G.fast maintains services under varying channel conditions, overcoming unpredictable changes in the channel response and noise without dropping the link. Re-acquisition of the channel per transmission frame, as in IEEE 802.11n, is infeasible for G.fast due to its relatively long channel estimation time. Instead, G.fast uses a combination of two OLR types: seamless rate adaptation (SRA) and fast rate adaptation (FRA). SRA is accurate but slow and facilitates steady-state performance optimization, while FRA is coarse and fast, and keeps the link stable under sudden deep drops in SNR. In addition, bit swapping is used to permanently adjust the SNR margin.

For SRA, similar to VDSL2, the FTU receiver computes the optimum bit loading per tone and communicates it to the peer transmitter via the eoc. The bit loading update at the peer FTUs is synchronized to the start of a particular superframe by an RMC command. For robustness, this command is repeated multiple times and contains a count-down to the targeted superframe, which allows the receiving FTU to identify it even in very harsh noise conditions.

For FRA, the G.fast spectrum is divided into up to eight contiguous sub-bands. An FRA command determines a coarse bit loading trim per sub-band. A trim-down avoids line drop upon unexpected substantial SNR loss. The trim request is generated by the receiver and communicated to the peer transmitter via the RMC; the transmitter quickly activates a bit loading update by sending an RMC synchronization command.

Since the RMC is much more robust than the data channel, it remains functional to recover the line from a temporary loss of data connectivity. As a receiver senses critical degradation of the channel, it requests via the RMC to trim down the bit loading in the affected sub-bands. After both the transmitter and receiver synchronously lower their bit loading according to the new channel conditions, data and eoc connectivity is recovered, and vectoring feedback is restored. Once vectoring coefficients are updated, the receiver optimizes bit loading by a trim-up FRA and/or SRA. This way, a stable link is maintained under harsh temporary conditions without sacrificing the steady-state performance by using extra SNR margin. The FRA timing is configured to avoid reacting to impulse noise, which is handled by retransmission.

The TIGA procedure is introduced into G.fast to facilitate updates of the precoder upon joining, leaving, tracking, and DO events, which in a high-crosstalk environment usually requires a change in the downstream transmit PSD and bit loading (see above). A TIGA eoc command is sent by the FTU-O to convey the necessary changes (bit loading and complex gain) to the receiver. The following RMC command provides a synchronous update of the precoder and parameters of the FTU-O transmitter and FTU-R receiver across the vectored group.

SPECTRAL COMPATIBILITY AND COEXISTENCE WITH CURRENT DEPLOYMENTS

G.fast spectral compatibility is determined by ITU-T G.9700 [2], which defines the PSD mask, transmit power limit, and a variety of spectrum

Keep-alive is very slow, but provides fast recovery of services. An NTU in L2.2 also detects incoming phone calls and temporarily transitions into L2.1 to support them. When the power is restored, the line moves into L2.1 (to support phone calls) and is ready to restart broadband service.




Figure 6. Simulated rate-reach curves (left) and actual measurement results (right) for 0.5 mm cables, 4 dBm TX power.

management tools established to reduce RFI egress into other DSL and radio services. The G.fast transmit power limit is 4 dBm. The PSD mask is -65 dBm/Hz below 30 MHz, drops to -73 dBm/Hz at 30 MHz, and further slopes down to -79 dBm/Hz at 212 MHz. G.fast suffers from RFI ingress, especially from FM radio, and crosstalk generated by VDSL2 and in-premises networks using power line technologies (PLT) [11].

COMPATIBILITY WITH DSL DEPLOYMENTS

Taking into account the expected migration from DSL to G.fast services and unbundling, it is highly desirable that G.fast be spectrally compatible with asymmetrical DSL (ADSL)/ADSL2plus and with VDSL2 deployed from an exchange or a cabinet.

The most practical and reliable way to maintain spectral compatibility between DSL and G.fast is by spectral separation. G.fast is by design compatible with ADSL/ADSL2plus since the lowest frequency of G.fast is 2.2 MHz. For compatibility with VDSL2, the start frequency of G.fast is set above the VDSL2 spectrum; the latter depends on the used VDSL2 profile. Use of spectral overlap between VDSL2 and G.fast is studied in [12].

COMPATIBILITY WITH BROADCAST, AMATEUR RADIO, AND PLT

For compatibility with broadcast radio, G.fast transmits substantially reduced PSD on frequencies above 30 MHz. If no transmission is allowed inside international amateur-radio bands or the FM-radio band (87.5 MHz — 108 MHz), these frequencies are notched out from the G.fast spectrum.

In-premises PLT networks may impact G.fast performance due to crosstalk between in-premises electrical wiring and phone wiring. This crosstalk is difficult to predict, although recently studied statistical models [11, 13] show that in 90 percent of the cases crosstalk attenuation is 60 dB or even less on frequen-

cies above 30 MHz. ITU-T G.9977 defines a mechanism to reduce this crosstalk by adjusting transmission parameters of G.fast and PLT network nodes via an arbitration device controlled by the operator.

PERFORMANCE OF G.FAST

THEORETICAL EVALUATION AND MEASUREMENT RESULTS

Capacity evaluations for G.fast 106 MHz and 212 MHz profiles over a sample of PE-0.5 mm cable are shown in Fig. 6 (left). The simulation shows bit rates averaged over five 24-pair cable binders, each with line lengths uniformly distributed between 10 m and 400 m, using a start frequency of 2.2 MHz (dots show performance on different pairs). The "optimized" option reflects PSD optimization described above, which is also used for 212 MHz performance projection with linear precoding.

The wideband transmit power and PSD limits meet G.9700. Flat spectrum background noise with a PSD of -140 dBm/Hz is applied to model the effect of QLN and receiver noise factor. For the 106 MHz profile, the aggregate (upstream plus downstream) bit rate of 500 Mb/s is achieved for lines up to 340 m. A similar distance is expected for 1 Gb/s service with 2-line bonding. For the 212 MHz profile, a 1 Gb/s aggregate bit rate can be reached for loops up to 220 m.

The measurement results in Fig. 6 (right) for a 16-pair group of another 0.5 mm 200 m cable with G.fast using the frequency range 23–106 MHz also shows the clear advantage of vectoring; a 500 Mb/s bit rate is supported on all tested lines.

MANAGEMENT

MANAGEMENT INTERFACES

Operators manage G.fast through the management information base (MIB) established at the DPU and in some cases at the NTU. At installation, the operator sets system configuration parameters, and during showtime the operator reads out performance and test

parameters, reported events, and collected statistics using management objects defined in [14]. Operators can access the DPU MIB remotely via the network management system (NMS). If the DPU is unpowered, a persistent management agent (PMA) acts as a proxy for the NMS. Relevant FTU-R management data are retrieved via the eoc. The NMS can access the NTU MIB, if established, using TR-069 [15].

DIAGNOSTIC AND PERFORMANCE PREDICTION

Means for line diagnostics include collection of line attenuation (HLOG), quiet-line noise (QLN), and signal attenuation (SATN) for both upstream and downstream. Unlike VDSL2, these parameters can be monitored during showtime and, together with reports on crosstalk coupling and SNR margin, provide a detailed picture of the current line status.

A method to measure the downstream HLOG and QLN is based on reporting of the discrete Fourier transform (DFT) samples of the received signal for specific elements of a probe sequence. For instance, setting a certain element to 0 in all lines allows measurement of the QLN. Furthermore, setting a particular element to 1 in one line and to 0 in all other lines allows measurement of the HLOG. The vectoring feedback (from the FTU) can be configured to report DFT samples for selected probe sequence elements, while error samples are reported on other elements. Thus, by adding a few elements to a probe sequence, HLOG and QLN can be monitored without interrupting the service and channel estimation.

DYNAMIC RESOURCE ALLOCATION

Dynamic resource allocation (DRA) controls the number of transmission slots in each logical frame for all lines of the DPU as a function of traffic loading, environmental conditions, power state, and battery status. By using DRA, service level agreements (SLAs) may be met using minimized DPU power consumption.

The main inputs to the DRA function are:

- Downstream and upstream traffic load indicators for every line in the DPU. Currently, these are simply per-traffic-class buffer occupancies, much like in G-PON. Using these indicators, the DRA adjusts the number of allowed transmission slots in a logical frame as a function of traffic load, thereby scaling the power consumption with traffic. The DRA can also switch a particular line to a low-power state if no broadband traffic is required.
- Environmental conditions, such as DPU housing temperature. This keeps DPU power dissipation within acceptable limits.
- Battery status indicator. The DRA can switch a particular line to a low-power state if the line is battery-fed as a result of power outage.

Based on these inputs, the DRA determines DO parameters (2.5.1) for each line. The parameter TBUDGET is coordinated across all DPU lines to control power dissipation and to address vectoring processing constraints.

CONCLUSIONS

G.fast brings DSL technology to a new level, comparable to the FTTH grade of service. It allows operators to offer their customers multiple broadband services, of both constant bit rate and variable bit rate, with total aggregated bit rate up to 1 Gb/s and low propagation delay. Such services include multi-channel HD video, high-quality audio and voice, and modern multi-user interactive gaming. G.fast assumes operation over low-grade copper drop cables and in-premises wiring. Reverse power feeding resolves the issue of cabinet powering, and readiness for customer self-install substantially improves cost effectiveness, simplifies system management, and brings convenience to the customer.

ACKNOWLEDGMENTS

The authors wish to acknowledge the great role in project development of Frank Van der Putten (Alcatel-Lucent), Chair, Les Brown (Huawei), Editor, and colleagues Les Humphrey (British Telecom), who led us with his vision of the project, Ian Horsley (British Telecom), Angus Carrick (Swisscom), Hubert Mariotte (France Telecom), Marko Loeffelholz (Deutsche Telekom), Miguel Peeters (Broadcom), Danny Van Bruyssel (Alcatel-Lucent), Ken Kerpez (ASSIA), and Chris Chang (Metanoia). We also thank Tom Starr (AT&T), who initiated development of G.fast in ITU-T, for helpful comments and guidance in this work.

REFERENCES

- [1] ITU-T Rec. G.9701-2014, "Fast Access to Subscriber Terminals (G.fast) Physical Layer Specification."
- [2] ITU-T Rec. G.9700-2014, "Fast Access to Subscriber Terminals (G.fast) Power Spectral Density Specification."
- [3] M. Timmers, C. Nuzman, and J. Maes, "G.fast: Evolving the Copper Access Network," *IEEE Commun. Mag.*, Aug. 2013, pp. 74–79.
- [4] ITU-T Rec. G.993.5-2010, "Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers."
- [5] V. Oksman *et al.*, "The ITU-Ts New G.vector Standard Proliferates 100 Mb/s DSL," *IEEE Commun. Mag.*, Oct. 2010, pp. 140–48.
- [6] Broadband Forum, "Architecture and Requirements for Fiber to the Distribution Point," tech. rep. TR-301.
- [7] ETSI, "European Requirements for Reverse Powering of Remote Access Equipment," ETSI tech. spec. TS 101 548, V1.2.1, Nov. 2014.
- [8] T. Starr *et al.*, *DSL Advances*, Prentice Hall, 2003.
- [9] ITU-T Rec. G.994.1-2012, "Handshake Procedures for Digital Subscriber Line Transceivers."
- [10] R. Strobel, A. Barthelme, and W. Utschick, "Zero-Forcing and MMSE Precoding for G.fast," *Proc. IEEE GLOBECOM*, 2015.
- [11] S. Galli *et al.*, "The Impact of PLC-to-DSL Interference on VDSL2, Vectored VDSL2, and G.fast," *Proc. IEEE ISPLC 2015*.
- [12] R. Strobel and W. Utschick, "Coexistence of G.fast and VDSL in FTTP and FTTC Deployments," *Proc. EUSIPCO*, 2015.
- [13] S. Galli *et al.*, "Statistical Modeling of PLC-to-DSL Interference," *Proc. IEEE ISPLC 2015*.
- [14] ITU-T Rec. G.997.2-2015, "Physical Layer Management for G.fast Transceivers."
- [15] Broadband Forum, "CPE WAN Management Protocol," tech. rep. TR-69.

BIOGRAPHIES

VLADIMIR OKSMAN (vladimir.oksman@intel.com) received his M.S. and Ph.D. degrees from the Leningrad Radio and Telecommunications College (USSR) in 1976 and 1987, respectively; he currently directs DSL standardization technical marketing at Intel Inc. He actively participated in development and standardization of emerging DSL technologies starting in 1991 as a researcher, project manager, and technical advisor. Since 2011 his main focus is concept engineering and standardization of G.fast.

RAINER STROBEL (rainer.strobel@intel.com) works for Concept Engineering at Lantix, now Intel, since 2011. He received his Dipl.-Ing.(FH) degree from the University of Applied Sciences Augsburg in 2009 and his M.Sc. degree in electrical engineering from Technische Universität München (TUM) in 2011.

G.fast brings DSL technology to a new level, comparable to FTTH grade of service. It allows operators to offer their customers multiple broadband services, of both constant bit rate and variable bit rate, with total aggregated bit rate up to 1 Gb/s and low propagation delay.

He is currently working toward a Ph.D. degree at Intel in cooperation with the Associate Institute for Signal Processing, TUM. His research interests include optimization of wired MIMO communication systems, copper-fiber hybrid networks, and wireline channel modeling.

XIANG WANG (Wangxiang@huawei.com) got his Ph.D degree in computational mathematics from Nanjing University in 2009. He is interested in numerical optimization, linear algebra application, and communications. His work at Huawei has been focused on DSL technology research since 2010. He also represented Huawei at the ITU-T Q4/SG15 meeting on DSL and G.fast.

DONG WEI (weidong@huawei.com) received his Ph.D. degree in electrical engineering from the University of Texas at Austin in 1998. He worked as an assistant professor at Drexel University, Philadelphia, Pennsylvania, and as a principal MTS at SBC Labs, with research interests in signal processing for communications and advanced access technologies. Since 2006, he has been a senior expert at the U.S. R&D Center of Huawei Technologies, working on advanced DSL technologies, actively contributing to various international SDOs.

RAMI VERBIN (rami@skipio.com) is a CTO and co-founder of Skipio Technologies, developing G.fast solutions. Prior to Skipio he was CTO of CopperGate and Sigma Designs. Since 1994 he has been developing wire-line access solutions and leading research in various communication fields. He holds an M.Sc. in electrical engineering and an M.B.A..

RICHARD GOODSON (msorbara@adtran.com) has over 30 years of advanced communication technology experience, from spread spectrum radios to xDSL and PON. He holds a B.S.E.E. from the University of Alabama and an M.S.E.E. from the University of Florida. He participates in the Broadband Forum, ITU-T, FSAN, and ATIS standards organizations. He has been at ADTRAN since 1995 where he is currently director of industry standards and technology analysis in the CTO Office.

MASSIMO SORBARA (msorbara@qca.qualcomm.com) is senior director of technical standards at Ikanos Communications and now at Qualcomm in Red Bank, New Jersey. He has worked on the standardization of DSL systems since 1990, holding various leadership positions in ITU-T, Broadband Forum, and ATIS. His research interests are in signal processing for advanced network access technologies. He received a B.S.E.E. from Manhattan College (1978) and an M.S.E.E. from the University of Santa Clara (1982).