The ITU-T's New G.vector Standard Proliferates 100 Mb/s DSL

Vladimir Oksman, Heinrich Schenk, and Axel Clausen, Lantiq Inc.

John M. Cioffi, Mehdi Mohseni, and George Ginis, ASSIA Inc.

Carl Nuzman and Jochen Maes Alcatel-Lucent

Miguel Peeters, Broadcom Corporation

Kevin Fisher, Ikanos Communications

Per-Erik Eriksson, Ericsson AB

ABSTRACT

This article explores the recently issued ITU-T G.vector (G.993.5) [1] that allows expanded use of 100 Mb/s DSL. A tutorial description on G.vector's crosstalk noise reduction methods leads to specific projections and measurements of expanded DSL 100 Mb/s reach. A discussion on dynamic maintenance to enhance G.vector's practical application then concludes this article.

INTRODUCTION

Modern life increasingly depends on faster Internet access for applications such as email, voice, information search, and a variety of social interconnections, while IPTV is increasing Internet video use. Growing wireless connectivity will substantially increase access bandwidth demand over the next decade.

Figure 1 shows digital subscriber line (DSL) technology as the current undisputed leader in broadband access. With more than 300 million worldwide subscribers, DSL use is consistently higher than cable modems and passive optical networks. A substantial part of Fig. 1's fiber to the X (FTTX) subscribers are actually very high rate DSL (VDSL2) connections where a fiber runs to an intermediate network point, usually within a neighborhood or in the basement of a multidwelling unit, and the DSL comprises the remaining connection to the customer. G.vector raises VDSL2 connection speeds up to 100 Mb/s at distances beyond 500 m from the fiber termination point with no transmit power increase and no Shannon-Law violation — G.vector simply removes most of DSL's crosstalk noise, thus providing a very high throughput. A similar principle is used in well-known gigabit Ethernet connections.

A copper twisted pair's throughput is unshared, and each individual DSL customer's speed often exceeds those of many other broadband connections. Furthermore, DSL's *fiber to the cabinet* architecture significantly reduces fiber deployment cost by sharing it between hundreds of customers connected to the cabinet via existing copper, enabling a more profitable DSL broadband business case. However, current VDSL2 provisions 100 Mb/s bit rates only over very short distances and, accordingly, requires too many cabinets. G.vector provides the business case's missing ingredient; it significantly increases the 100 Mb/s range, and thus enables reasonable broadband connection cost. We present some test results and illustrate how G.vector supports such a low-capital-cost VDSL2 deployment of 100 Mb/s.

The existing copper connections are ready for use immediately, but can exhibit enormous variability in signal attenuation and noises. Additionally, the line's state depends on the particular installation, and customers' touching of lines or moving of equipment. Today, the operational costs associated with DSL trouble call response, dispatch of technicians (*truck rolls*), and customer service drop/change (*churn*) dominate DSL operational costs, and are a concern for higher-speed DSL network enhancements. As G.vector removes most of the far-end crosstalk (FEXT) noise of other VDSL2 lines sharing the cable, the nonstationary noise and other remaining uncancelled line noises become increasingly important. Thus, link management using line monitoring and optimization techniques, known as dynamic spectrum management (DSM), becomes crucial to retaining G.vector's fundamental gains.

The shorter-line high-speed DSL concept using fiber-fed street cabinets or curb boxes first appeared in standards in 1994 [2], culminating in International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Recommendations G.993.1 (VDSL1, 2004) and G.993.2 (VDSL2, 2006). VDSL enlarges asymmetric DSL (ADSL) architecture by increasing the number of subcarriers to cover wider bandwidth [3]. Vectoring evolves the multiple-input multiple-output (MIMO) signal processing first suggested by Paulraj [4] for multiple-antenna wireless. Adaptation of this vectoring technology to more stationary channels

first appears in [5] and was continually refined for DSL in 1999 and 2000 [6]. A reduced complexity near-optimal implementation of linear vectoring was proposed in [7]. Complexity of vectoring is still challenging, especially for large numbers of vectored subscribers.

Vectored DSL proposals first appeared in the American National Standard Institute (ANSI) in the 2001 DSM project [8] as the highest of three levels of crosstalk noise control methods for copper loop management. Vectored DSLs have physically separated customer premises (CP) locations that cannot be coordinated directly by a common customer-side controller. Recommendation G.993.5 describes the necessary interoperable line coordination functions at the DSL access multiplexer (DSLAM) and the individual lines' training protocols within the coordinated vector group.

Figure 1. *Broadband access-connection for 2006–2009 (source: Point Topic).*

G.993.5 STANDARD VECTORED TECHNOLOGY

FEXT CANCELATION PRINCIPLES

The dominant VDSL2 noises are near-end crosstalk (NEXT, between transmitters and receivers connected to different pairs of a multipair cable at the same end), far-end crosstalk (FEXT, between transmitters and receivers connected to different pairs of a multipair cable at opposite ends; Fig. 2), and background Gaussian noise [3]. NEXT coupling is usually so strong above 1–2 MHz that VDSL2 systems use non-overlapping downstream/upstream frequency bands multiplexing up to 30 MHz. Thus, with NEXT largely eliminated, FEXT may dominate the remaining noise. G.vector technology cancels VDSL2's mutual FEXT, thus effecting a performance improvement.

In a typical VDSL deployment, multiple VDSL2 lines connect the DSLAM to the VDSL transceiver unit — remote terminals (VTU-Rs) at physically separated individual residences' CPs. VDSL2 uses digital multitone (DMT) modulation [9] with up to 4096 subcarriers located on frequencies f_i, spaced by $\Delta f = 4.3125$ kHz or $\Delta f = 8.625$ kHz ($f_i = i \times \Delta f$, $i = 0, 1, \ldots, 4095$). Each subcarrier carries a certain number of bits that depends on this subcarrier's signal-to-noise ratio (SNR). FEXT cancellation reduces noise and thus can increase these SNRs. This allows carriage of more bits and therefore increases the line's data rate.

FEXT generated by a particular twisted pair m into a victim twisted pair l can be cancelled at the subcarrier frequency f_i by subtracting from the received signal the value of the signal $Um(f_i)$, transmitted over the pair m, multiplied by the FEXT transfer function from the pair m into the pair l, $H_{l-m}(f_i)$. If the cable binder includes N pairs, and FEXT has to be cancelled on M subcarriers, the subtracted crosstalk signal comprises $(N - 1) \times M$ components, where each component corresponds to a particular binder pair's FEXT into the victim line on the particular subcarrier frequency. Each subcarrier's (f_i) crosstalk signal is obtained by multiplication of the signal vector **u** and the FEXT coupling vector **H**, both of size N; hence the name *vectoring*.
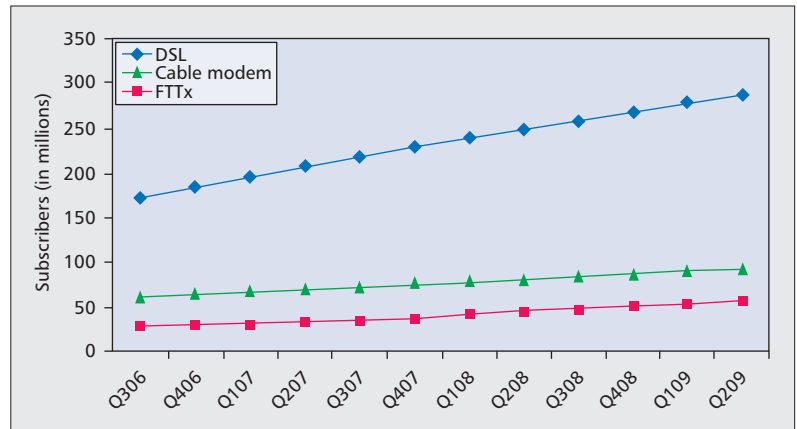
The DSLAM side must perform vectoring for both downstream and upstream, because DSLAM alone contains all the cable's DSL signals (Fig. 2a). A downstream FEXT precoder precedes the modulation, and an upstream canceller follows the demodulation. The vectoring control entity (VCE) supplies updated channel matrices and controls the FEXT cancellation process, such as indicating from which lines to cancel FEXT. Such control allows efficient use of the available processing power to cancel dominant crosstalk.

G.vector defines interoperability only for downstream vectoring, since upstream vector processing needs no interoperability specification. However, G.vector also specifies certain VTU-R control signals and timing that facilitate upstream vectoring: Figure 2a illustrates a potential vectoring implementation.

Figure 2b defines each downstream line's *precoding* for FEXT cancellation. The precoder, for each subcarrier frequency f_i, multiplies the precanceller matrix $\mathbf{C}_p^{(f_i)}$ corresponding row by the signal vector $\mathbf{u}^{(f_i)}$ of N lines. The VTU-R's received vector, $\mathbf{y}^{(f_i)}$, will be: $\mathbf{y} = \mathbf{H} \cdot \mathbf{C}_p \cdot \mathbf{u} + \mathbf{r}$, where **H** is the channel transfer matrix, **r** is the received non-FEXT noise vector, and the precanceller matrix, $\mathbf{C}_p$, is equal to the inverted normalized channel matrix, $\mathbf{H}_0 = [\mathbf{diag}(\mathbf{H})]^{-1} \cdot \mathbf{H}$, and $\mathbf{C}_p = \mathbf{H}_0^{-1}$. Similarly, for upstream FEXT cancellation, the received signal vector can be described as $\mathbf{y} = \mathbf{C}_c \cdot \mathbf{H} \cdot \mathbf{u} + \mathbf{r}$, and differs from the downstream case since the **H** is different.

Vectoring also involves timing requirements. The DSLAM strictly aligns all lines' upstream and downstream DMT symbols to typically within 1 μs. The DSLAM adjusts each vector group VTU-R's upstream timing advance value to meet this small tolerance. Also, all vectored lines transmit downstream sync symbols at the same time, while VTU-Rs transmit all upstream sync symbols also at the same time. This alignment eliminates sync-symbols' FEXT into data symbols, as in Fig. 3. To align all upstream vectored lines' sync symbols, the DSLAM sends the VTU-R a special time marker during the line's initialization, which indicates the time offset between the upstream and downstream sync symbols, and thus allows the new line's VTU-R to align its transmitted sync-symbol position with that of other lines.
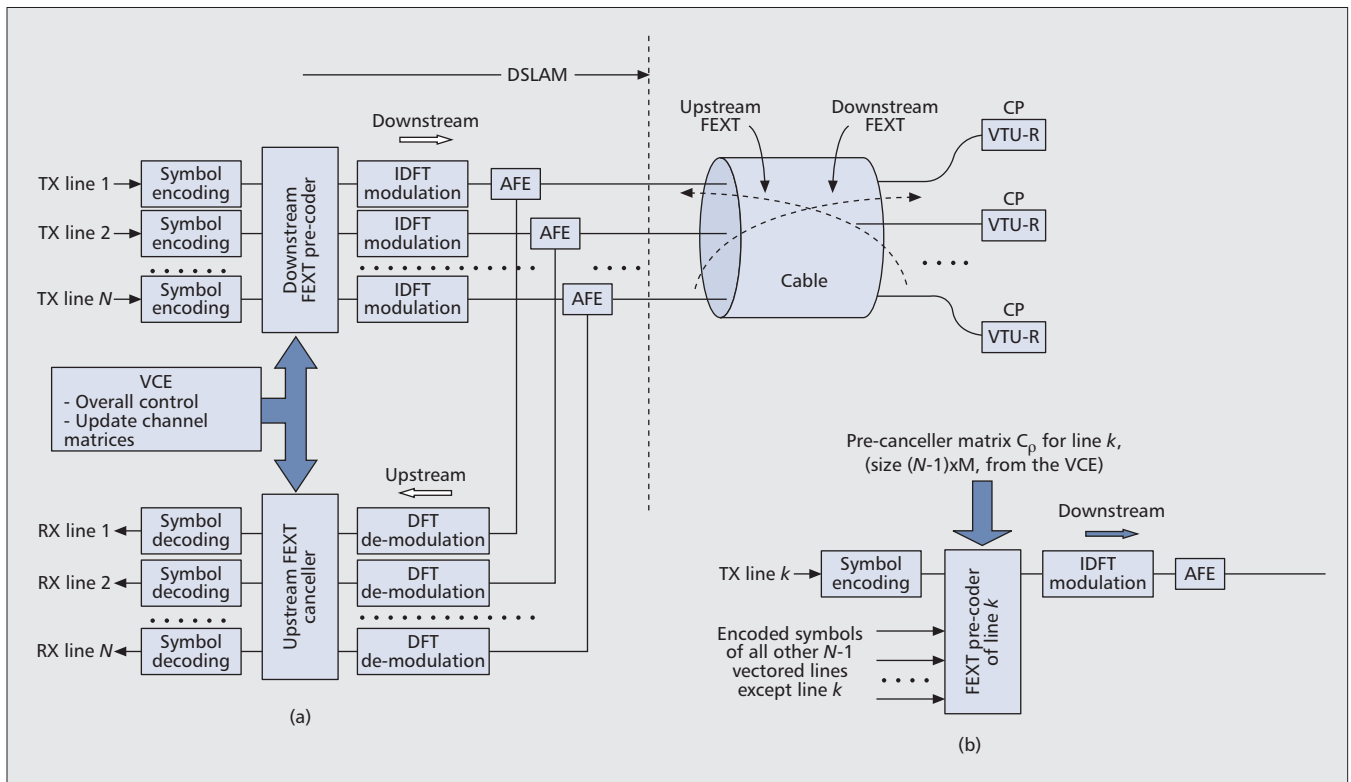
Figure 2. *Functional description of upstream and downstream vectoring.*

FEXT ESTIMATION

FEXT cancellation requires the FEXT coupling coefficients between all pairs of the vectored group at each subcarrier. FEXT coupling estimation between two pairs uses repeating pilot signals to determine the other pair's FEXT component. Since FEXT that is not yet cancelled during joining may cause unacceptably high noise, the pilot signal is transmitted only during sync-symbols. All vectored lines' sync-symbols thus are aligned in time and carry no user data; therefore, even full-power pilot signals sent during sync-symbols do not disturb vectored lines' data transmission (Fig. 3).

VDSL2 transmits a sync symbol every 256 data symbols. For FEXT estimation, a special binary pilot sequence modulates each line's sync symbol on pre-assigned *probe-tone* subcarriers with indices equal to $10n$, $10n + 2$, $10n + 3$, $10n + 4$, $10n + 5$, $10n + 6$, $10n + 8$, and $10n + 9$ ($n = 0, 1, 2, …$). *Flag-tone* subcarriers with indices equal to $10n + 1$ or $10n + 7$ allow communication of standard VDSL2 online reconfiguration (OLR) signals. The reduced number of OLR-carrying subcarriers does not impact OLR robustness because the number of flag tones is still large enough (vectored lines are relatively short and use wide spectrum). The value of FEXT on flag tones is interpolated from adjacent probe tones.

The DSLAM assigns different vector-group lines' binary pilot sequences, which are usually selected as mutually orthogonal. The orthogonality speeds up and simplifies FEXT estimation that correlates victim lines' measured receiver error values with the disturbing lines' orthogonal sequences. One popular class of orthogonal sequences are Walsh-Hadamard sequences, for which the length can be any power of 2.

Apart from crosstalk estimation's use of orthogonal pilot sequences, direct FEXT estimation methods like a least mean square (LMS) algorithm can be used successfully . Selection of a particular pilot sequences is vendor discretionary; each DSLAM vendor can fit the sequence to their preferred FEXT estimation algorithm. G.vector also describes an alternative FEXT estimation method that uses each vector group's reported SNR values, as measured at the VTU-R [1].

REPORTING OF ERROR SAMPLES

The VCE needs feedback from the VTU-R on errors that are signal distortions caused by uncompensated downstream FEXT in the pilot sequences of received downstream sync-symbols on subcarriers assigned for downstream FEXT precoding. Samples of errors help the adaptive determination of the FEXT coupling coefficients used in the precanceller matrix. Thus, each VTU-R's measured receiver error samples are reported via the backchannel to the VCE for FEXT estimation and precanceller matrix ($\mathbf{C}_p$) computation. G.vector assures CP-to-DSLAM interoperability by standardization of backchannel parameters and formats of the reported error samples. Backchannel capacity and reported error sample accuracy are critical for the FEXT cancellation algorithm's convergence speed.

Error Samples — A given subcarrier's error sample is the difference $E = Z – D$ between the received complex signal Z and its intended constellation point D, normalized to the scale of a unit-magnitude constellation point. The VCE
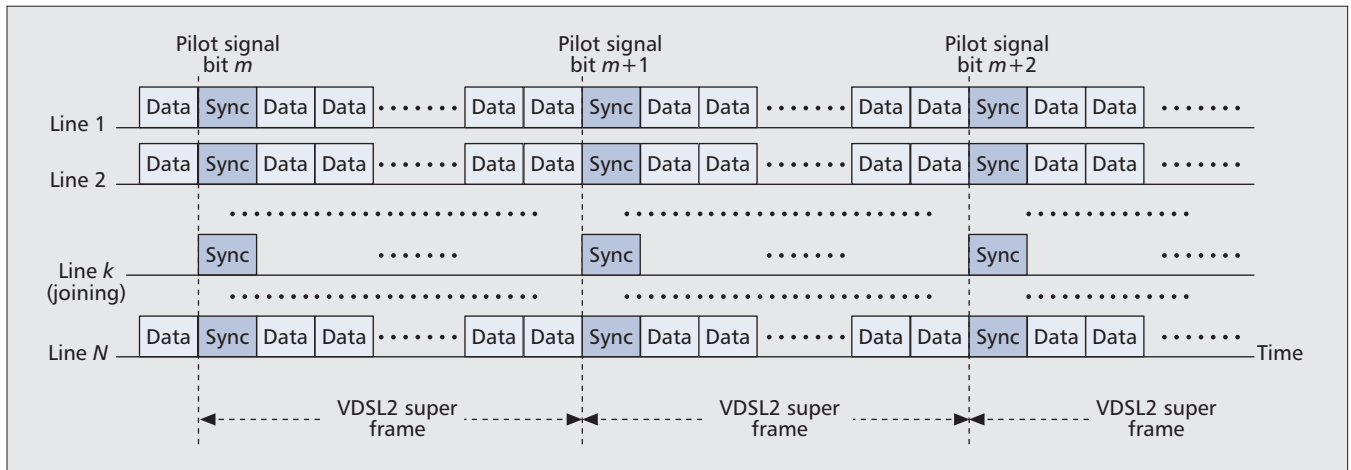
Figure 3. *Symbol alignment, sync-symbol alignment, and pilot bits in a vectored group. (Same in upstream and downstream; line* k *is at the beginning of the joining procedure).*

divides the downstream frequency spectrum into up to eight non-overlapping bands for flexible reporting. A particular band's sync-symbol error samples are reported in a block floating point format that supports different backchannel capacities, FEXT estimation algorithms, and deployment scenarios.

Grouping error samples into blocks saves backchannel bandwidth, since a common 4-bit exponent applies to both real and imaginary parts of all the block's samples; this is effective because neighboring subcarriers' error samples typically have similar magnitudes. A block may contain one error sample, 32 error samples, or all error samples of the band. Further bandwidth may be saved by frequency and time subsampling of the reported error samples (e.g., reporting only even subcarriers' error samples or from only every third sync-symbol). Clipping of error sample components to a specified maximum magnitude to avoid a particular block's exponent is unduly influenced by powerful impairments on certain subcarriers, such as radio frequency ingress (RFI), causing a loss of precision for the remaining subcarriers. The VCE configures the size of the mantissa between 0 (sign only) and 8 bits.

The VCE configures the bands and format of error sample reporting, and can adjust the error feedback accuracy by selecting an appropriate mantissa length and matching the backchannel bandwidth through grouping and subsampling. The VTU-R may also flag a particular report as potentially corrupted (e.g., by impulse noise) and report the error's mean value over the vectored band, thereby assisting the DSLAM assessment of FEXT estimation completeness. With all these means, the DSLAM can effectively address different corner situations, such as excessive SNR variations over the vectored band (due to strong RFI ingress or bridged taps) or very limited bandwidth of the backchannel, and efficiently use the available DSLAM processing power.

Backchannel — The backchannel conveys error samples from the VTU-R to the DSLAM. Three transport mechanisms are specified for flexible backchannel operation: over the special operations channel (SOC) during line initialization,

and over the embedded operations channel (EOC) or Ethernet channel during showtime. All three mechanisms use the same error sample formats presented earlier.

Later we describe the O-P-VECTOR 2 stage initialization's use of the SOC-based transport mechanism. Each error report is transmitted as a message, encapsulated in a high-level data link control (HDLC) frame. This transport mechanism is a high-speed VDSL2 SOC that conveys the initialization messages. The DSLAM configures the backchannel data rate and the VTU-R's used error report format via SOC messages.

A DMT symbol's repetition of the same information on several subcarriers achieves SOC protocol robustness. The repetition rate determines the capacity of the SOC, which can range from 16 to 192 b/DMT symbol, providing bit rates from 64 kb/s to 768 kb/s, respectively, for 4000 symbols/s transmission. Transmission with a reduced number of repetitions (relative to VDSL2) is still reliable for vectored VDSL2 due to shorter distances.

VDSL2 uses EOC to convey OLR and operations, administration, and maintenance (OAM) messages. The EOC bit rate is fixed and at a set level below 256 kb/s. For transportation over EOC, error reports are mapped into a standard EOC message and transported with high priority. The format of the error samples is configured through an EOC command sent from the DSLAM.

The Ethernet backchannel has a flexible data rate: the error reports are encapsulated in Ethernet frames and multiplexed with the upstream user data. At the DSLAM, the received Ethernet packets are identified by the CP's medium access control (MAC) address and delivered to the VCE's assigned MAC address. When multiple lines connect the CP to the DSLAM (bonded connection), a CP-assigned ID marks each line. The error samples' reported format is configured as with the EOC-based backchannel. The VCE and CP MAC addresses and the line ID are set during initialization.

By specifying both Ethernet and EOC mechanisms, G.993.5 offers a choice of cancellation algorithms and system architecture. Some cancellation algorithms may require a higher

backchannel peak bit rate than available through the EOC, while some system architectures may exclude reported-error-sample multiplexing among the user data.

OPERATION OF A VECTORED GROUP

Vector group operation comprises three phases: tracking, joining, and leaving. In tracking, no new lines join or leave the group; each line tracks routine FEXT-coupling variations, mainly caused by temperature changes. Tracking's FEXT matrix update is usually very slow, because low-accuracy infrequent error sample reports suffice as derived from a low-speed backchannel.

A vectored group transitions to the joining phase when one or more lines initialize to join the group. First, the joining line transmits only sync-symbols carrying pilot sequences, as in Fig. 3, line *k*. Existing lines then estimate each joining line's FEXT and update their FEXT cancellation matrices without disruption. Furthermore, joining lines accommodate existing lines' FEXT. A joining event requires much higher backchannel throughput than tracking, so the backchannel throughput flexibility accommodates the required change. For that, Ethernet backchannel packets are assigned high priority or the EOC is pre-configures to provide sufficient throughput. Some auxiliary EOC transactions, such as performance monitoring, may be deferred to provide faster joining. Lower backchannel capacity slows the joining event.

Leaving events may be orderly or disorderly. Orderly leaving first terminates transmission in both directions, thus allowing the CP modem to safely disconnect. Disorderly leaving corresponds to sudden CP modem disconnect, power off, or line disconnect, which can potentially increase other vectored DSLs' residual FEXT during several seconds, until the DSLAM and CP modem both detect the disconnect and stop transmitting. Disorderly leaving is currently under study.

INITIALIZATION

G.vector line initialization adds new stages to VDSL2 initialization (darker blue boxes in Fig. 4), allowing new lines to join seamlessly. During initialization, prior to steady-state transmission (showtime), VDSL2 modems estimate the channel transfer function (*channel discovery phase*), adapt receiver parameters to the channel (*training phase*), and further compute and exchange bit loading tables between the VTU-O and VTU-R (*channel analysis and exchange phase*). At showtime, small adaptations compensate for natural channel and noise changes.

The handshake phase starts initialization, when the two sides exchange capabilities and agree on a common operational mode. During the VECTOR-1 stage, crosstalk from the joining lines into vectored lines is estimated and compensated; after this stage the initializing line can use standard full-power VDSL2 training signals without disrupting vectored lines. The VECTOR-1 signals consist only of modulated sync symbols transmitted simultaneously with other vectored lines' sync symbols (Fig. 3, line *k*). The VECTOR-1-1 stage is similar to VECTOR-1 and readjusts the crosstalk cancellation from the joining lines into vectored lines after potential changes in FEXT coupling during the channel discovery phase (usually due to deviations in impedances of the modems).

During VECTOR-2 stage joining, lines estimate and compensate crosstalk from vectored lines. After VECTOR-2, all joining lines are ready to compute SNR and FEXT-compensated bit loading. VECTOR-2 signals include sync-symbols modulated by pilot sequences and initialization data symbols at all other symbol positions. These sync symbols allow FEXT estimation from vectored lines into each of the joining lines. Data symbols carry the joining CP-line's reported error samples via the SOC backchannel, extended to higher capacity. The DSLAM sets the error sample format based on the actual SOC throughput and required error precision.

To join several lines (or for a full startup), the VCE synchronizes all joining lines' VECTOR initialization stages, providing a sufficiently long time period of simultaneously transmitted sync-symbols with pilot sequences, and then terminates them at the same time. The VCE thereby acquires the same FEXT estimation data from all vectored lines during VECTOR-1 and VECTOR-1-1, and from all joining lines during VECTOR-2.

Figure 4's two plots illustrate SNR convergence to its steady-state value (FEXT cancelled) for the case of one line joining a group of 31 vectored lines (downstream sync-symbols' sub-carrier #684). The upper plot shows the vectored lines' worst SNR (on line 26), and the lower plot shows the joining line's SNR. In both cases, convergence times indicate the required duration of VECTOR-1 and VECTOR-2 transmission, respectively. This duration grows with the number of joining lines and the overall size of the vectored group. The convergence of the upstream SNR looks very similar.

PERFORMANCE OF VECTORED SYSTEM

THEORETICAL CAPACITY OF VECTORED VDSL2

Figure 5 shows downstream bit rates vs. VDSL2 loop length for a North American residential deployment (Profile 17a) operating over 17 MHz with and without vectoring. As each disturber's FEXT magnitude into each subscriber varies based on the particular binder-pair location, cable type, manufacturing tolerances, and deployment practices, the expected vectored VDSL2 gains appear from a statistical perspective using four families of curves. Besides FEXT, simulations also include –140 dBm/Hz additive white Gaussian noise.

Upper and lower performance bounds appear with upward and downward facing triangles. The upward facing triangles show each loop length's FEXT-free bit rate. Downward facing triangles show the bit rate using a standard (conservative) 99 percent worst-case FEXT model. These upper and lower bounds reflect VDSL2's potential capacity range.

The blue crosses show a large non-vectored VDSL2 group's simulated bit rates, assuming a statistically representative loop-length distribution and the NIPP-NAI statistical crosstalk model [10], for 100-pair (4 × 25-pair binders) North
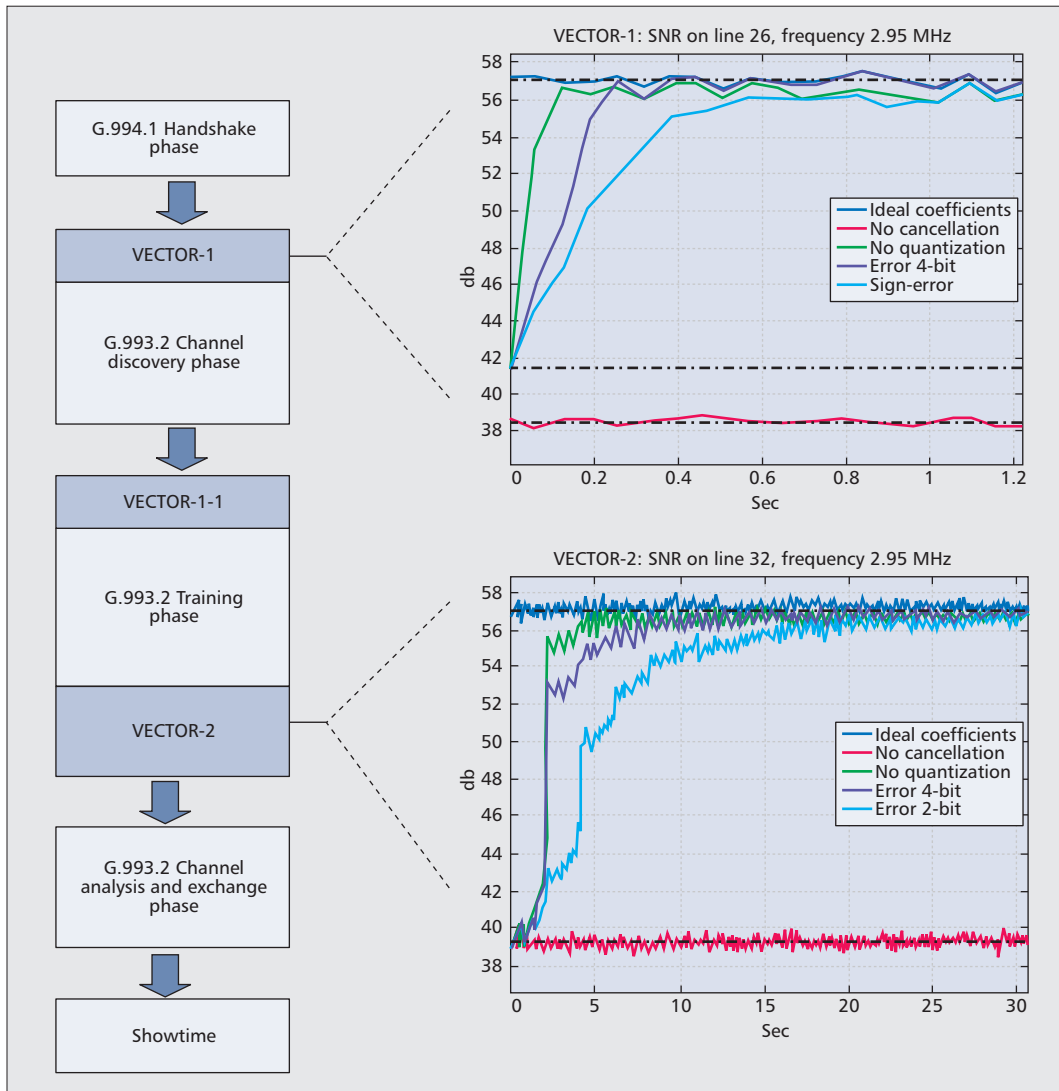
Figure 4. *G.995.3 initialization timeline and simulation results for SNR in vectored and joining line (1 lines joins a group of 31 lines AWG-24 500 m).*

American wire gauge (AWG) 26 cable. Each blue cross represents an individual cable line's data point, where 96 of the 100 pairs carry VDSL2 signals and 4 pairs are unused. Most lines achieve a bit rate higher than the worst-case bit rate; however, there is no easy advance prediction of a particular line's achievable bit rate. Most service providers consequently provision their service based on the worst-case assumption.

The red circles in Fig. 5 show bit rates for the same vectored VDSL2 lines. Since the entire cable employs vectoring, FEXT-free rates are closely achieved. Thus, vectoring can significantly increase FEXT-limited downstream bit rates on loops shorter than 800 m. Similar upstream bit rate improvements are also possible.

Cancellation of all cable crosstalk is usually not necessary, but only the dominant disturbers. Partial cancellation suggests balancing the implementation complexity with the desired deployed bit rate. The presented example cancelled only each line's 63 most significant FEXT sources from 95 (approximately 2/3 of disturbers). Each individual line has its own unique set of signifi-

cant disturbers based on the line's relative position in the cable, as well as on the cable's construction. Thus, partial cancellation necessarily pre-observes crosstalk across the entire cable to identify each line's unique set of top disturbers.

Practical vectored performance gains depend heavily on implementation specifics and service provider deployment practices. Preservation of vectoring performance gains critically involves all active cable pairs at least to ascertain which pairs actively contribute to any specific lines' crosstalk.

LABORATORY TEST RESULTS

These vectored VDSL2 17-MHz profile (17a) prototype measurements were taken on 26 AWG and European Telecommunications Standards Institute (ETSI) 0.5 mm cables to validate the following characteristics:
• The bit rate as a function of loop length, with and without crosstalk
• The number of a binder's crosstalking lines that have to be cancelled to achieve significant performance increase
• The impact of adjacent-binder FEXT
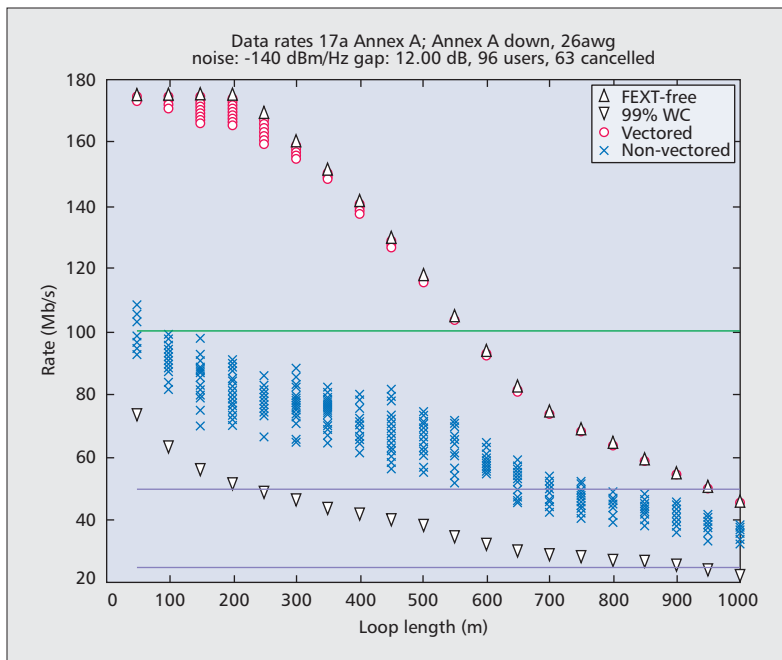
Figure 5. *Downstream bit rates for vectored and regular VDSL2 (Profile 17a, bandplan EU32).*

A 24-port vectored-VDSL2 DSLAM prototype was connected to 50-pair 26-AWG cables of different lengths, comprising two binders with 25 pairs each. Only 24 pairs of 50 were simultaneously accessed. The cable was divided into 500-ft sections, allowing sectional cable extension, while keeping pair consistency at the connection points. Measurements were taken for three test setups:

1. Successive single-line activation to measure each line's FEXT-free performance
2. All lines activated with no vectoring
3. All lines activated with vectoring to record each line's performance with FEXT from other 23 lines cancelled

Figure 6 shows measured vectored results that achieve almost FEXT-free performance. The difference in performance between the 0.5 mm ETSI and 26 AWG cables is caused by different impedance (around 120 Ω for ETSI cable, whereas the 26 AWG impedance is 100 Ω). The rate cap around 100 Mb/s is due to the particular prototype's implementation. Tests for VDSL2 30-MHz profiles are expected in the near future.

Properties for a Vectored VDSL2 System — One vectored VDSL2 property is that its steady-state performance becomes more predictable than legacy VDSL2. One example is the reduced bit rate variation caused by FEXT variation. Another example of the increased predictability is negligible performance variation when lines join or leave. In legacy VDSL2 systems, new joining lines cause unpredictable performance reduction in other lines due to FEXT, and even loss of synchronization. Vectored systems no longer exhibit this instability.

Numbers of Dominant Disturbers and Impact from Adjacent Binders — Measurements also showed how many lines need to be

FEXT-cancelled for a significant performance increase. Best vectored performance often requires cancellation of most disturbers in the same binder; cancelling five or six dominant disturbers is not sufficient. The average FEXT from adjacent binders can be expected to be about 10 dB lower than intra-binder crosstalk, although adjacent binders can still contain a number of lines that must be treated as strong crosstalkers. The overall conclusion is that the strongest crosstalk occurs within the same binder with less crosstalk from the adjacent binder.

MANAGEMENT

The management techniques and interfaces for vectored DSL relate to DSM level 3.

MANAGEMENT CHALLENGES FOR VECTORED DSL

First, it is important to manage performance trade-offs between customers' lines. One reason is that the VCE's computational resources can be insufficient to cancel all disturbers. Service providers' management systems can then set higher priority for lines with bandwidth-sensitive services. Also, algorithm settings can favor certain lines at the expense of others (e.g., the upstream vectoring decoding order).

Second, it is important to manage noise sources that become dominant after FEXT among the vectored lines is cancelled. Other noise sources, such as impulse noise, may become dominant and reduce the performance benefits. Sudden external noise changes may lead lines to re-initialize. Management systems can avert such disruptions by appropriate impulse noise protection configuration, or by other means of mitigating abrupt noise changes.

Finally, the management system can improve overall performance in mixed-binder cases where vectored and non-vectored VDSL2 lines coexist in the same cable or binder. If left unmanaged, such non-vectored lines' FEXT may eliminate significant vectoring benefit. The management system can adjust the non-vectored lines' transmitted power based on their service requirements to limit their impact on vectored lines, [11]. ADSL lines have much less impact due to the narrow bandwidth they use.

MANAGEMENT INTERFACE OF G.VECTOR

The G.vector management interface has been standardized in ITU-T Recommendation G.997.1, which includes enhancements required for effective G.vector management and defines new management parameters.

A new XLOG test parameter reports crosstalk transfer function on downstream subcarrier basis. XLOG indicates the FEXT coupling strength between any two vector-group lines (for a detailed XLOG definition, see [1]). The XLOG enables some very useful management functions.

Crosstalk diagnosis: XLOG allows the identification of lines creating excessive FEXT. Typically, such lines are characterized by faults (e.g., poor balance) that lead to poor performance and require maintenance action. Additionally,

Figure 6. *Measured bit rate for 17a profile: left: downstream, 0.5 mm ETSI; right: downstream and upstream, 26 AWG; the curves for single-line DS and vector DS overlap.*

even though it may be possible to mitigate the FEXT induced by such pairs into vectored lines, high FEXT is likely to also be generated into non-vectored lines, which cannot be cancelled. By identifying such extreme crosstalk *polluters*, the copper network is improved over time.

Performance prediction: XLOG allows expected FEXT cancellation gain prediction for specific pairs. G.vector supports configuration parameters for allocation of vectoring computational resources to prioritized lines. Knowing the expected gains provides essential guidance for priority-level choice or even vector enablement choice.

The G.vector configuration parameters allow:
• Enabling or disabling a particular line's FEXT cancellation
• Selecting FEXT cancellation frequency bands
• Assigning each line's FEXT cancellation priorities
• Assigning each line's target data rate

A particular line's disabled-vectored capability allows service providers allocation of vectoring to maximize benefit (i.e., assisting high-end services to perform maximally). These higher-priority lines can benefit the most when there are limited computational resources.

Control of vectoring frequencies is useful when certain frequencies may suffer from *uncancellable* noises, like radio frequency interference (RFI) or other types of DSL (ADSL or legacy VDSL). The vectored system can then be instructed to disable vectoring on the corresponding subcarriers.

Finally, setting FEXT cancellation priorities is essential for allocating vector computational resources within a vectored group. Since the VCE's computational resources may be limited, the VCE may not be able to cancel all lines' FEXT into every vectored-group line. External guidance on the resource allocation may be based on each line's service requirements. As for legacy VDSL2, a vectored line must always main-

Figure 7. *Setting of line priorities and data rate configuration parameters.*

tain its bit rate between the configured minimum and maximum data rate. In addition, a G.vector system should attempt to allocate enough resources to at least achieve the target data rate. If resources are still available, the VCE should allocate those remaining resources to the lines with line priority set to HIGH. Lines with priority set to LOW might maintain a lower bit rate than the target data rate, as described in Fig. 7. Lines with different priority levels may also use different bit loading algorithms to help the VCE mitigate FEXT.

CONCLUSIONS

G.vector can extend the reach of standard high-speed VDSL2 systems significantly. This reach extension enables service providers to offer higher-speed services to more customers at a lower cost than was previously anticipated in DSL access networks. G.vector or G.993.5 will enable a large interoperable market for equipment and its associated dynamic management across an

array of vendors, ensuring cost-effective availability of high-speed DSL access networks worldwide. G.vector's large DSL benefit is expected to accelerate video, voice, wireless (through backhaul of increasingly smaller cells offering more bandwidth to mobile users), and other highly revenue-generating telecommunications services at a time when such services are of particular interest.

ACKNOWLEDGMENTS

REFERENCES

[1] ITU-T Rec. G.993.5-2010, "Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers."
[2] J. M. Cioffi and J. A. C. Bingham, "A Proposal for Consideration of a VADSL Standard Project," ANSI contrib. T1E1.4/94-183, Dec. 1994
[3] J. Cioffi et al., "Very High-Speed Digital Subscriber Lines," IEEE Commun. Mag., Aug. 2004.
[4] A. Paulraj and T. Kailath, "Increasing Capacity in Wireless Broadcast Systems Using Distributed Transmission/Directional Reception (DTDR)," U.S. Patent 5,345,599, Sept. 6, 1994.
[5] J. M. Cioffi and G. D. Forney, Jr., "Generalized Decision-Feedback Equalization for Packet Transmission with ISI and Gaussian Noise," Ch. 4, A. Paulraj, V. Roychowdhury, and C. Schaper, Eds., Communication, Computation, Control, and Signal Processing: A Tribute to Thomas Kailath, Kluwer, 1997.
[6] G. Ginis and J. M. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," IEEE JSAC, vol. 20, no. 5, June 2002, pp. 1085–1104.
[7] R. Cendrillon et al., "A Near-Optimal Linear Crosstalk Precoder for VDSL," IEEE Trans. Commun., May 2007.
[8] J. Cioffi and K. Song, "Level 3 DSM Results: Vectoring of Multiple DSLs," ANSI T1E1.4 contrib. 2002-059, Vancouver, Canada, Feb. 18, 2002.
[9] T. Starr, J. Cioffi, and P. Silverman, Understanding DSL, Prentice Hall, 1999.
[10] ATIS pre-published tech rep. ATIS-PP-0600024, "Multiple-Input Multiple-Output Crosstalk Channel Model," 2009.
[11] M. Mohseni, G. Ginis, and J. M. Cioffi, "Dynamic Spectrum Management for Mixtures of Vectored and Nonvectored DSL Systems," 44th Annual Conf. Info. Sciences and Sys., Princeton, NJ, Mar. 2010.

BIOGRAPHIES

VLADIMIR OKSMAN (vladimir.oksman@lantiq.com) received his M.S. and Ph.D. from the Leningrad Radio and Telecommunication State College, USSR, in 1976 and 1987. Currently, he is director of Technical Marketing at Lantiq. He is 34+ years in digital communications, developing the first long-distance digital transmission systems in USSR, early HDLS and ADSL, single-carrier and multicarrier VDSL systems, and home networking technologies. His recent efforts are devoted to vectored DSL and novel networking technologies for energy management. He actively participates in multiple standard development organizations.

HEINRICH SCHENK (heirich.schenk@lantiq.com) received Dipl. Ing. and Dr. Ing. degrees from Friedrich Alexander University, Erlangen, Germany, in 1974 and 1978. From 1979 to 1996 he worked for the Siemens Public Switching Division, Germany, as a senior engineer and head of the Development Division. Since 1997 he has worked in the field of concept and algorithm development in the Semiconductor Department of Siemens and later at Infineon. Presently he is a senior principal at Lantiq in Munich.

AXEL CLAUSEN (axel.clausen@lantiq.com) received his M.S.E. and Ph.D. in electrical engineering from Arizona State University in 1995 and 1999, respectively. He held different positions in system and concept engineering for broadband access, developing algorithms and system architecture for ADSL, VDSL, and home networking products. Currently, he is a principal system engineer in Lantiq, leading the development of system architecture and algorithms for vectored VDSL. He was actively involved in standardization of vectored DSL.

JOHN M. CIOFFI (jcioffi@assia-inc.com) is chairman and CEO of ASSIA Inc. and Hitachi Professor Emeritus, Stanford University. He received his B.S.E.E. in 1978 from the University of Illinois, his Ph.D.E.E. in 1984 from Stanford, and an honorary doctorate from the University of Edinburgh in 2010. He was the founder of Amati Communications Corp., 1991–1999, and the designer of the world's first ADSL modem and basic ADSL/VDSL patents. Affiliations and awards include the Boards of Alto Beam and ClariPhy; IEEE AG Bell (2010), Millennium (2000), Kobayashi (2001), and IEE Tomson (2001) medals; Marconi Fellow (2006); U.S. NAE (2001); FREng (2009); and ANSI Outstanding Achievement (1996).

MEHDI MOHSENI (mmohseni@assia-inc.com) joined ASSIA, Inc., in 2006 as a systems architect, focusing on DSL network optimization and improving management systems. His research interests incorporate multi-user information theory and optimization of wireless and broadband networks using convex optimization techniques. He holds a Bachelor's degree in electrical engineering from Sharif University of Technology, Tehran, Iran, as well as M.S. and Ph.D. degrees in electrical engineering from Stanford University.

GEORGE GINIS (gginis@assia-inc.com) is currently vice president of Expresse Engineering at ASSIA, Inc., overseeing development and deployment of the DSL Expresse management product. Between 2002 and 2005 he was a DSL systems engineer at Texas Instruments. He holds a Diploma in electrical and computer engineering from the National Technical University of Athens, and M.S. and Ph.D. degrees in electrical engineering from Stanford University.

CARL NUZMAN (nuzman@research.bell-labs.com) is a Distinguished Member of Technical Staff in the Enabling Computing Technologies Department, Bell Labs, Alcatel-Lucent. Recent research interests include vectored DSL systems, wireless power control, and optical networking. He holds a Ph.D. in electrical engineering from Princeton University, and B.S. degrees from the University of Maryland College Park.

JOCHEN MAES (jochen.maes@alcatel-lucent.com) received his M.S. and Ph.D. in physics in 2000 and 2004 from Katholieke Universiteit Leuven, Belgium, for his research on charge coupling in nano-structured semiconductor lasers of 1.3 μm and 1.55 μm wavelengths. Since 2006 he has been researching next-generation broadband access systems over copper at Bell Labs Alcatel-Lucent. He acts as a reviewer for several journals and conferences, and has published around 35 papers.

MIGUEL PEETERS (miguel.peeters@broadcom.com) is a graduate in engineering from the Universite Libre de Bruxelles and Ecole Centrale Paris. He joined Broadcom Corporation in 2000. For more than 10 years his activities have been related to the use of multicarrier modulations for DSL. He is a regular contributor to standard bodies like ITU-T, where he is editor of the G.997.1 Recommendation.

KEVIN FISHER (kfisher@ikanos.com) has more than 23 years of experience in the development of firmware, software, and complex silicon devices. His experience includes executive and senior-level engineering positions including CEO of Vector Silicon, VP of IC and DSP engineering at 2Wire, and department manager of advanced recording channels at Quantum. He earned his Ph.D. E.E. and M.S.E.E from Stanford University after a B.S. degree in computer engineering at the University of Illinois. He holds 20 U.S. patents.

PER-ERIK ERIKSSON (per-erik.s.eriksson@ericsson.com) is currently a senior research engineer at Ericsson Research working with next-generation DSL technologies involving VDSL2 vectoring, and technologies for high-frequency transmission over copper. He actively participates in the DSL standardization in ITU-T SG15/4. He holds an M.S. in electrical engineering from the Royal Institute of Technology, Stockholm, Sweden.