


VDSL: The Next Step in the DSL Progression



Krista S. Jacobsen
(jacobsen@ti.com)
Texas Instruments
Broadband Access Group
San Jose, CA

DSPS Fest
August 1999
Houston, TX

THE WORLD LEADER IN DSP SOLUTIONS



Table of Contents

| | | |
|-------|---|----|
| 1 | General overview..... | 5 |
| 2 | Introduction..... | 6 |
| 3 | VDSL fundamentals | 6 |
| 3.1 | VDSL deployment scenarios..... | 6 |
| 3.2 | Transmission directions..... | 7 |
| 3.3 | Transmission channels and twisted-pair lines..... | 7 |
| 3.4 | Power spectral density (PSD)..... | 10 |
| 3.5 | Noise..... | 10 |
| 3.5.1 | Crosstalk | 10 |
| 3.5.2 | Radio-frequency ingress | 12 |
| 3.5.3 | Impulse noise | 12 |
| 3.6 | Signal-to-noise ratio (SNR)..... | 13 |
| 3.7 | Noise margin | 16 |
| 4 | Key ANSI/ETSI VDSL system requirements | 16 |
| 4.1 | Data rates and ratios | 16 |
| 4.2 | Egress suppression | 17 |
| 4.3 | Transmit power and PSD | 18 |
| 5 | Additional system considerations | 19 |
| 5.1 | Spectral compatibility | 19 |
| 5.2 | Power consumption, size and complexity | 20 |
| 6 | VDSL solutions | 21 |
| 6.1 | Line code alternatives..... | 21 |
| 6.2 | DMT vs. CAP/QAM: Which approach meets the VDSL system requirements? | 23 |
| 6.3 | Duplexing alternatives..... | 26 |
| 6.4 | TDD vs. FDD: Which approach meets the VDSL system requirements? | 28 |
| 7 | Performance of DMT-based TDD VDSL..... | 31 |
| 8 | Summary..... | 33 |
| 9 | Acronyms..... | 34 |
| 10 | References..... | 35 |

List of Figures

| | | |
|------------|--|----|
| Figure 1. | FTTEx architecture | 6 |
| Figure 2. | FTTCab architecture | 7 |
| Figure 3. | Attenuation as a function of frequency of two 26 AWG (0.4 mm) lines and two 24 AWG (0.5 mm) lines | 8 |
| Figure 4. | Examples of loops with and without bridged taps | 9 |
| Figure 5. | Insertion losses of three lines, one without bridged taps, two with a single bridged tap | 9 |
| Figure 6. | Illustration of near-end crosstalk (NEXT) | 10 |
| Figure 7. | NEXT coupling due to 10 disturbers, illustrating dependence of coupling on frequency | 11 |
| Figure 8. | Illustration of far-end crosstalk (FEXT)..... | 12 |
| Figure 9. | FEXT coupling due to 10 disturbers, illustrating the dependence of coupling on frequency and line length (26 AWG)..... | 12 |
| Figure 10. | Impulse noise in the time and frequency domains | 13 |
| Figure 11. | Simplified view of the process of applying an FEC code and interleaving the result | 13 |
| Figure 12. | Applying a signal with flat PSD to a 1-km 24 AWG line..... | 14 |
| Figure 13. | SNR corresponding to example in Figure 12..... | 14 |
| Figure 14. | Example noise profile of 10 ADSL NEXT and -140 dBm/Hz AWGN | 15 |
| Figure 15. | Receiver SNR when a signal at -60 dBm/Hz is applied to Loop B and corrupted by the noise profile in Figure 14..... | 15 |
| Figure 16. | Illustrating the potential for ADSL and VDSL interactions in the FTTEx configuration | 19 |
| Figure 17. | Illustrating the potential for ADSL and VDSL interference in the FTTCab configuration | 20 |
| Figure 18. | Single-carrier (QAM) transmitter block diagram..... | 21 |
| Figure 19. | DFE block diagram | 22 |
| Figure 20. | DMT transmitter/receiver pair block diagram | 23 |
| Figure 21. | Typical placements of the downstream and upstream channels in FDD | 26 |
| Figure 22. | A single band supports both downstream and upstream transmission in TDD systems..... | 27 |
| Figure 23. | TDD superframes enable support of a variety of downstream-to-upstream data rate ratios..... | 28 |
| Figure 24. | NEXT caused by mixing symmetric and asymmetric FDD systems | 29 |
| Figure 25. | NEXT caused when symmetric and asymmetric TDD superframes are mixed..... | 29 |
| Figure 26. | Average maximum range of DMT-based TDD VDSL in symmetric mode with 6 dB downstream and upstream noise margins | 32 |
| Figure 27. | Average maximum range of DMT-based TDD VDSL when in 8:1 asymmetrical mode with 6 dB downstream and upstream noise margins..... | 32 |

List of Tables

| | | |
|----------|---|----|
| Table 1. | ETSI evaluation payload bit rate combinations | 16 |
| Table 2. | ANSI evaluation data rate combinations..... | 17 |
| Table 3. | Amateur radio bands recognized by ETSI | 17 |
| Table 4. | Amateur radio bands recognized by ANSI | 18 |

1 General overview

Very high-speed digital subscriber line (VDSL) technology promises to deliver information at speeds up to 52 megabits per second (Mbps) over ordinary copper telephone lines. These very high speeds push information access far beyond the limits of asymmetric digital subscriber line (ADSL) technology, which is now in the early stages of deployment. VDSL data rates, equivalent to fractional T3 or full T3 services, will increase Internet access speeds for businesses, provide fast links between local-area networks at separate sites, and bring video services to business and residential customers - all without requiring deployment of optical fiber to the home.

Realization of the potential of VDSL depends on the availability of enabling integrated circuit (IC) technology. Some vendors have already announced VDSL solutions, but not all of these new products meet the requirements specified by the ANSI and ETSI standards bodies. Equipment manufacturers who are searching for the right VDSL solution must understand the technical constraints imposed by the VDSL transmission environment and the ANSI and ETSI VDSL system requirements, and how competing solutions handle these constraints.

To operate successfully, VDSL equipment must overcome line attenuation, crosstalk, radio-frequency (RF) ingress and other interferences. Of particular importance is operation on existing unshielded twisted-pair lines. After all, if extensive re-wiring is necessary (in the home, for example) to enable VDSL, then perhaps the deployment of fiber to and within the home is a more cost-effective solution! VDSL modems must also sustain specified data rates over measured distances, suppress RF emissions, and coexist compatibly with frequency spectra of other services that may be present in the same cable bundle, such as ADSL, ISDN and HDSL, or on the same line, like plain old telephone service (POTS).

Discrete multi-tone (DMT) modulation achieves these VDSL operational requirements more effectively than quadrature amplitude modulation (QAM) or carrierless amplitude-phase (CAP) modulation. Because DMT provides near-optimal performance on all channels, it overcomes inherent line problems and operates successfully even when certain frequency regions are affected by severe noise. DMT is resilient to RF ingress, adapts to changing channel and noise conditions, supports all required data rates, suppresses RF egress in amateur radio bands, and provides spectral compatibility with other DSLs.

Since mid-1998, Texas Instruments has offered a VDSL solution that delivers the performance optimization and flexibility of DMT. TI's solution, which is based on high-performance application-specific integrated circuits (ASICs), enables equipment manufacturers to design low-power, low-complexity systems that meet the demanding ANSI and ETSI VDSL requirements. The TI approach gives service providers the flexibility in deployment and provisioning they desire. Software programmability allows the operator to configure remotely, for each line, the downstream and upstream data rates and their ratio. As a result, the same hardware supports both symmetric and asymmetric connections.

2 Introduction

Today, several IC vendors have announced solutions for very high-speed digital subscriber line (VDSL) systems. But equipment manufacturers who are undertaking VDSL designs need to understand that not all of these “solutions” are really solutions. That is, not all of these chipsets meet the requirements for VDSL, especially the requirements specified by ANSI and ETSI.

To aid in the evaluation of VDSL solutions, this paper discusses important technical constraints of the VDSL transmission environment, summarizes key ANSI and ETSI requirements for VDSL systems, and considers popular line code and duplexing schemes in respect to these factors.

3 VDSL fundamentals

This section begins by describing a number of VDSL deployment scenarios. Next, several fundamental VDSL concepts are presented, including transmission channels, twisted-pair lines, power spectral density, noise, signal-to-noise ratio, and noise margin. With an understanding of the ideas in this section, the reader will appreciate, given the requirements for VDSL systems, why TI’s VDSL chipsets perform better and consume less power than other alternatives.

3.1 VDSL deployment scenarios

VDSL is intended to support high bit-rate transmission on telephone lines shorter than about 5 kft (about 1.5 km) in length. In the public telephone network, there are two general loop architectures. In densely populated areas or cities, many customers are within a few kilofeet of the *central office (CO)* or *local exchange (LEx)*. In such cases, VDSL can be deployed directly from the CO or LEX. This configuration, known as *fiber-to-the-exchange (FTTEx)*, is illustrated in Figure 1. When fiber extends deeper into the network, VDSL is deployed from the *optical network unit (ONU)* in a configuration known as *fiber-to-the-cabinet (FTTCab)*. The FTTCab architecture is shown in Figure 2.




Figure 1. FTTEEx architecture

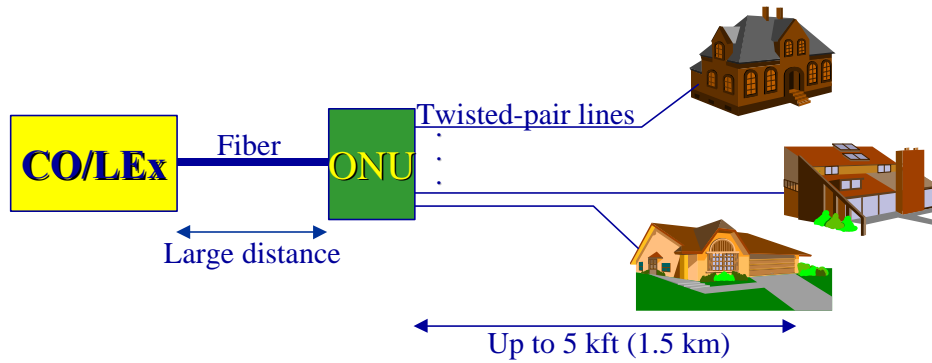


Figure 2. FTTCab architecture

3.2 Transmission directions

The transmission direction from the CO, LEx or ONU to the customer premise is called *downstream*. The direction from the customer to the CO, LEx or ONU is called *upstream*.

3.3 Transmission channels and twisted-pair lines

A transmission channel is a physical medium used to transfer information-bearing signals from one point to another. In telephony networks, the channels are twisted pair, which are manufactured by twisting together two insulated copper wires. Multiple pairs are then twisted together tightly and bound into cables. As described in the previous section, these cables emanate from a CO, LEx or ONU, and individual twisted pair are tapped from the cable as necessary to provide service to subscribers.

All physical channels attenuate signals. The amount by which a transmitted signal is attenuated when it reaches the receiver at the end of a twisted-pair line is a function of a number of variables, including frequency, line length and wire gauge. On all lines, attenuation increases with frequency. The rate at which attenuation increases with frequency is a function of line length and wire gauge (or the diameter of the copper wires). Signals on long lines composed of small-diameter wires are attenuated very rapidly with increasing frequency, whereas short lines made of larger-diameter wires cause a more gentle increase in attenuation with frequency. Figure 1 illustrates attenuation as a function of frequency of four lines: 300 and 1,000 meters of 26 AWG (or 0.4-mm) line, and 300 and 1,000 meters of 24 AWG (or 0.5-mm) lines. The attenuation curves are smooth because the lines are terminated in the appropriate characteristic impedance at both ends. A comparison of the four curves shows clearly the relationships between attenuation and line length and wire gauge.

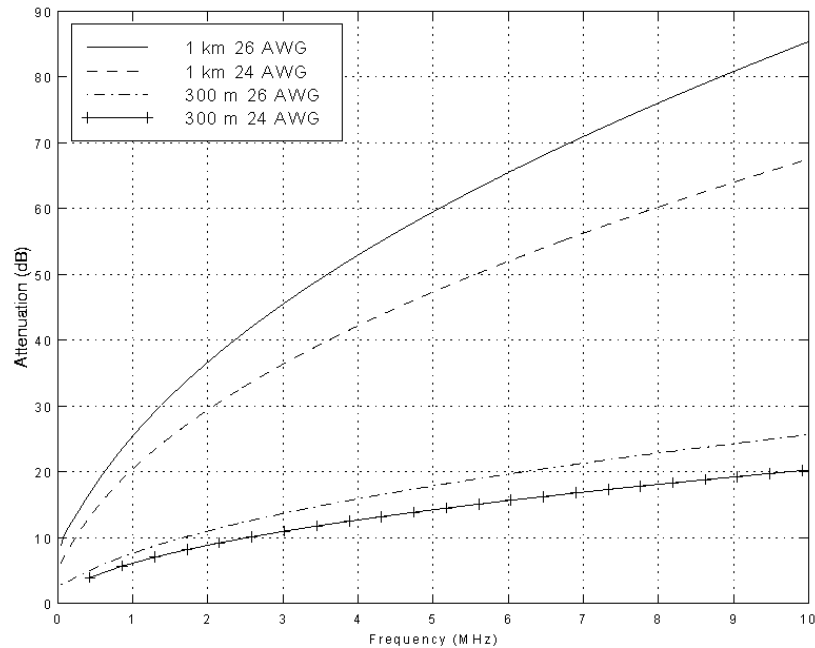


Figure 3. Attenuation as a function of frequency of two 26 AWG (0.4 mm) lines and two 24 AWG (0.5 mm) lines

Because the attenuation is smooth with frequency, signals transmitted on these types of lines are relatively simple to decode at the receiver. Thus, a single length of line terminated in the appropriate impedance at both ends is preferred for VDSL transmission. However, many twisted-pair lines do not exhibit such smooth attenuation. For example, *bridged tap* configurations are common in certain regions of the world, including the United States. In a bridged tap configuration, an unused twisted-pair line is connected in shunt to a main cable pair. The unused pair is left open-circuited across the main pair. Bridged taps were deployed originally to provide flexibility for future expansion in the plant. Today they are not installed in new telephone line deployments, although DSL systems must be capable of operating on older lines with existing bridged taps. On these lines, the unterminated stub of twisted-pair reflects signals that, at the receiver, add constructively in some frequency regions and destructively in others. Where destructive interference occurs, transmitted signals are attenuated significantly. In the frequency domain, destructive interference appears in the channel's frequency response (or insertion loss) as notches. The locations and depths of these notches depend on the lengths of unterminated stubs connected to the main line. In the time domain, bridged taps distort signals, stretching them in time and changing their amplitudes. VDSL transmission systems must be capable of operating with distorted signals.

Figure 4 illustrates three loop configurations, two of which feature bridged taps. Figure 5 shows the insertion losses of the loops A through C. The dashed curve contains three pronounced notches, from 8-12 dB deep, caused by the 100-meter unterminated stub in Loop B. The dash-dot curve demonstrates that the effect of a longer unterminated stub is passband ripple at low frequencies and a nearly-constant 3-dB loss at higher frequencies.




Figure 4. Examples of loops with and without bridged taps




Figure 5. Insertion losses of three lines, one without bridged taps, two with a single bridged tap

3.4 Power spectral density (PSD)

A power spectral density (PSD) describes how power is distributed in frequency. Signals and noises are both described by PSDs, which are often expressed in decibel milliwatts per Hertz (dBm/Hz).

3.5 Noise

Noise is all unwanted energy arriving at a receiver along with the desired signal. Noise is caused by a variety of sources and is, unfortunately, unavoidable. Together with channel attenuation, noise limits the rate at which data can be transported on a channel. This section describes the various noises encountered on twisted-pair lines.

3.5.1 Crosstalk

Individual wires that compose twisted-pair lines are insulated, and the twisting of these lines into cables limits electromagnetic interference to nearby lines. However, because the shielding between lines is not perfect, signals from one line can couple into other lines. As a result, a local receiver can detect signals transmitted on other lines, thus increasing the noise power and degrading the received signal quality on that line. The coupling of unwanted signals from one or more lines into another line is known as *crosstalk*, and in VDSL applications it can take two forms: *near-end crosstalk* and *far-end crosstalk*.

Near-end crosstalk (NEXT)

Near-end crosstalk (NEXT), illustrated in Figure 6, occurs when a local receiver detects signals transmitted on other lines by one or more local transmitters. In other words, signals coupling into the line are transmitted in the direction opposite to the received signal but in an overlapping frequency band. The level of NEXT detected at a local receiver is dependent primarily on the number of interferers and their proximity to the line of interest, the relative powers and spectral shapes of the interfering signals, and the frequency band over which NEXT occurs. In general, NEXT coupling between adjacent lines in a cable is worse than NEXT between lines spaced further apart. In addition, NEXT worsens if the transmit power on the interfering line(s) is increased: if the transmit power of all interferers is doubled, so is NEXT appearing at other receivers. Finally, NEXT is dependent on the frequency band over which it occurs. The coupling of signals is proportional to frequency raised to the power of 3/2. Figure 7 shows NEXT coupling due to 10 disturbers over the frequency range of 0 to 11 MHz. The dependence of coupling on frequency is evident in the figure.




Figure 6. Illustration of near-end crosstalk (NEXT)




Figure 7. NEXT coupling due to 10 disturbers, illustrating dependence of coupling on frequency

Far-end crosstalk (FEXT)

Far-end crosstalk (FEXT) occurs when a local receiver detects signals transmitted in its frequency band by one or more remote transmitters. In this case, interfering signals travel in the same direction as the received signal, as illustrated by Figure 8. As is the case for NEXT, the level of FEXT detected is dependent on the number of interferers and their proximity to the line of interest, the relative powers and spectral shapes of the interfering signals, and the frequency band over which FEXT occurs. However, in this case the frequency-dependence of FEXT coupling is to the power of 2 rather than to the $3/2$ power. Thus, FEXT is a stronger function of frequency than is NEXT. In addition, the detected FEXT level depends on the length of line over which coupling of unwanted signals occurs. In contrast to NEXT, which is essentially independent of line length¹, FEXT decreases with increasing line length because unwanted signals are attenuated by the length of line into which they couple. For this reason, FEXT is usually a minor² impairment on long lines such as ADSL, which are typically over 2 miles long. FEXT is more significant on shorter VDSL lines: it can dominate noise profiles. Figure 9 shows FEXT coupling due to 10 disturbers as a function of frequency for three line lengths. The solid curve is FEXT coupling assuming the length over which coupling occurs is 300 meters, and the dashed and dash-dot curves show coupling assuming lengths of 1 km and 1.5 km, respectively. The three curves show clearly that FEXT coupling decreases with increasing line length. They also illustrate the dependence of FEXT on line attenuation, which, as shown previously, increases with loop length and frequency. On longer loops, the line attenuation is severe enough to counteract the f^2 contribution to the FEXT coupling expression, resulting in coupling curves that, beyond some frequency, decrease rapidly with increasing frequency.

-
1. Studies have shown that on short loops, a few hundred feet long, for example, NEXT is dependent on loop length.
 2. "Minor" means other contributors to the noise profile are substantially more severe.
-




Figure 8. Illustration of far-end crosstalk (FEXT)




Figure 9. FEXT coupling due to 10 disturbers, illustrating the dependence of coupling on frequency and line length (26 AWG)

3.5.2 Radio-frequency ingress

Ingress noise appears at receivers when over-the-air signals in overlapping frequency bands couple into phone lines. Telephone lines, specifically overhead distribution cable and wires within the home, are particularly susceptible to ingress from AM radio signals due to the high power levels and density of AM radio stations in the over-the-air spectrum. AM interferers appear in the VDSL frequency spectrum as high-level noise spikes in the band between 525 kHz and 1.61 MHz.

3.5.3 Impulse noise

Impulse noise is a short-duration, high-power burst of energy that can, temporarily, overwhelm information-bearing signals. Impulse noise on twisted-pair lines can be caused by electronic and electro-mechanical devices such as telephones and other household appliances. In addition, power line discharges and lightning can also cause impulse noise.

In the frequency domain, during the time it exists, impulse noise is characterized by a power spectrum that is essentially flat and wide-band, as shown in Figure 10. Because its occurrences are impossible to predict, impulse noise can impair or disable poorly-designed systems.




Figure 10. Impulse noise in the time and frequency domains

To mitigate impulse noise, well-designed systems use forward error correcting (FEC) codes in conjunction with data interleaving, as shown in Figure 11. FEC codes add redundancy to the transmitted signal so that if a portion of the signal is corrupted by impulse noise, the remaining parts of the signal contain enough information to reconstruct the missing portion. FEC codes are applied byte-wise to blocks of data and are characterized by how many byte errors they can correct within a block of a specified length. If the actual number of byte errors exceeds the number the code can correct, then the resulting data stream contains some incorrect data. To maximize the effectiveness of FEC codes, interleaving can be used at the transmitter. After application of an FEC code, an interleaver rearranges the order of the coded bytes so that any impulse noise on the channel will corrupt a set of bytes that, when de-interleaved, are spread out in time. In effect, the interleaver reduces the time over which a single impulse harms the (de-interleaved) signal, thereby improving the chances the FEC code can correct the residual errors. The effectiveness of an interleaver depends on its depth, which is proportional to the amount of time by which consecutive coded bytes are separated after the interleaver has been applied. The penalty for interleaving is a receiver delay incurred while the FEC decoder waits to operate on the de-interleaved, but still coded, data stream.




Figure 11. Simplified view of the process of applying an FEC code and interleaving the result

3.6 Signal-to-noise ratio (SNR)

The signal-to-noise ratio (SNR) is the ratio of power of the information-bearing signal at the receiver to the power of the received noise. Essentially, the SNR describes the quality of the transmission channel. In the frequency domain, the SNR is computed by dividing the PSD of the information-bearing signal at the receiver by the PSD of the received noise. The PSD of the information-bearing signal is simply its transmit PSD multiplied by the insertion loss of the channel. Because the channel insertion loss and received noise profile always vary with frequency, and the transmit PSD may vary with frequency, the SNR is a function of frequency.

Together with the target detection error probability and channel bandwidth, the SNR determines the maximum rate at which information can be transmitted across a channel.

To illustrate SNR, Figure 12 shows the application of a signal with a flat -60 dBm/Hz transmit PSD to the 1-km 24 AWG line discussed in Section 3.1. A noise profile consisting only of additive white Gaussian noise (AWGN) at -140 dBm/Hz is present at the receiver. Figure 13 shows the resulting receiver SNR for this example.




Figure 12. Applying a signal with flat PSD to a 1-km 24 AWG line




Figure 13. SNR corresponding to example in Figure 12

A second example illustrates the impact of bridged taps and shaped noise on the SNR. In this case, a signal with a flat PSD at -60 dBm/Hz is applied to the line called Loop B in Section 3.1. The noise is NEXT from 10 ADSL lines in the same binder plus AWGN at -140 dBm/Hz. The composite PSD is shown in Figure 14. Figure 15 shows the receiver SNR for this example. The effects of bridged taps and ADSL NEXT are clear from the SNR plot.




Figure 14. Example noise profile of 10 ADSL NEXT and -140 dBm/Hz AWGN




Figure 15. Receiver SNR when a signal at -60 dBm/Hz is applied to Loop B and corrupted by the noise profile in Figure 14

3.7 Noise margin

As shown in Section 3.5, many twisted-pair impairments (such as impulse noise) are unpredictable and time-varying. As these impairments come and go, the receiver SNR varies with time. To account for unexpected increases in noise and the resultant decreases in receiver SNR, most systems do not operate at the maximum rate the channel can support. Instead, systems operate with a *noise margin*, which is the amount by which the SNR can degrade without increasing the receiver detection error probability. In essence, use of a noise margin “backs off” the system to prevent detection errors due to unforeseen increases in noise.

4 Key ANSI/ETSI VDSL system requirements

The requirements for VDSL systems have been established by telephone companies through their participation in standards groups ANSI T1E1.4 in the United States and ETSI TM6 in Europe. Although T1E1.4 and TM6 are independent groups whose specifications apply to disjoint geographical regions, the participants recognize the economic benefits of a common VDSL solution in the United States and Europe, and as a result they have attempted to generate consistent sets of VDSL system requirements. This section describes those requirements, highlighting differences when necessary. Additional details can be found in references [1] and [2].

4.1 Data rates and ratios

Both ANSI and ETSI require support of symmetric and asymmetric downstream-to-upstream VDSL data rate ratios. ETSI denotes modems that support asymmetric data rates as “Class I” and those that support symmetric rates as “Class II.” A compliant modem must support Class I or Class II operation, or both. For modem evaluation, ETSI specifies the payload bit rate combinations shown in Table 1.

Table 1. ETSI evaluation payload bit rate combinations

| Modem class | Downstream rate (kbps) | Upstream rate (kbps) |
|-------------|------------------------|----------------------|
| I | 6 x 1024 | 2 x 1024 |
| | 12 x 1024 | 2 x 1024 |
| | 24 x 1024 | 4 x 1024 |
| II | 6 x 1024 | 6 x 1024 |
| | 12 x 1024 | 12 x 1024 |
| | 24 x 1024 | 24 x 1024 |
| | 36 x 1024 | 36 x 1024 |

From Table 1, the required ETSI downstream-to-upstream data rate ratios are 6:1, 3:1, and 1:1. Some TM6 participants have recommended removing the 36.864 Mbps Class II requirement because of concerns that requiring accommodation of such a high symmetric data rate, which is possible only on a small percentage of short lines, would substantially increase modem cost and complexity.

ANSI also specifies bit rates for symmetric and asymmetric modes of operation. The data rate combinations used for modem evaluation are given in Table 2, which reveals the ANSI downstream-to-upstream data rate ratios are 8:1, 4:1, and 1:1. Note that ANSI and ETSI require different asymmetric payload bit rate ratios.

Table 2. ANSI evaluation data rate combinations

| Service type | Downstream rate (Mbps) | Upstream rate (Mbps) |
|---------------------|-------------------------------|-----------------------------|
| Asymmetric | 52 | 6.4 |
| | 34 or 38.2 | 4.3 |
| | 26 | 3.2 |
| | 19 | 2.3 |
| | 13 | 1.6 |
| | 6.5 | 1.6 or 0.8 |
| Symmetric | 34 | 34 |
| | 26 | 26 |
| | 19 | 19 |
| | 13 | 13 |
| | 6.5 | 6.5 |
| | 4.3 | 4.3 |
| | 2.3 | 2.3 |

4.2 Egress suppression

As mentioned previously, signals transmitted in the over-the-air spectrum can enter telephone lines and corrupt VDSL signals. The opposite effect occurs, too. Egress occurs when VDSL signals on twisted-pair lines leak into over-the-air bands and corrupt signals. Of particular concern to operators are emissions into the amateur radio bands. For this reason, both ANSI and ETSI require VDSL modems to restrict their transmit power spectral densities to levels no higher than -80 dBm/Hz within the amateur radio bands recognized in their respective geographical regions. The locations of the bands recognized by ETSI and ANSI are listed in Tables 3 and 4, respectively.

Table 3. Amateur radio bands recognized by ETSI

| Band start frequency (MHz) | Band stop frequency (MHz) |
|---------------------------------------|--------------------------------------|
| 1.810 | 2.000 |
| 3.500 | 3.800 |
| 7.000 | 7.100 |
| 10.100 | 10.150 |
| 14.000 | 14.350 |
| 18.068 | 18.168 |
| 21.000 | 21.450 |
| 28.000 | 29.100 |

Table 4. Amateur radio bands recognized by ANSI

| Band start frequency (MHz) | Band stop frequency (MHz) |
|---------------------------------------|--------------------------------------|
| 1.800 | 2.000 |
| 3.500 | 4.000 |
| 7.000 | 7.300 |
| 10.100 | 10.150 |
| 14.000 | 14.350 |
| 18.068 | 18.168 |
| 21.000 | 21.450 |
| 28.000 | 29.700 |

4.3 Transmit power and PSD

The *transmit PSD* describes how the power of an information-bearing signal is distributed in frequency when the signal is applied to the channel at the transmitter output. For example, if a total of 10 mW is applied to a line, and the transmitter applies that power evenly across a 1 MHz bandwidth, then the transmit PSD is a constant 10^{-5} mW/Hz or -50 dBm/Hz. A *transmit PSD mask* specifies the maximum allowable transmit PSD, which is, by definition, a function of frequency.

The maximum transmit power allowed by both ANSI and ETSI is 11.5 dBm. However, ANSI and ETSI differ on how that power may be distributed in frequency. Both groups have defined masks specifying the maximum allowable transmit PSD, and both require modems to be capable of reducing their transmit PSDs to -80 dBm/Hz within the amateur radio bands.

ANSI's approach to specifying the maximum transmit PSD involves three binary choices by the operator deploying a VDSL system. The choices are:

PSD enhancement (on/off):

When off, the transmit PSD is restricted to a maximum value of -60 dBm/Hz over the available frequency band. When on, the transmit PSD may be boosted above -60 dBm/Hz. The PSD boost is governed by a PSD mask and is subject to the total power constraint of 11.5 dBm.

ADSL compatibility (on/off):

When on, the transmit PSD below 1.104 MHz is restricted to -90 dBm/Hz. When off, the transmit PSD in this region is subject to the status of the PSD enhancement option. Section 5.1 illustrates why the ADSL-compatible option is necessary.

RF emission notching (on/off):

When on, the transmit PSD within the amateur radio bands is restricted to -80 dBm/Hz. When off, the PSD within the amateur bands is subject to the status of the PSD enhancement option.

It should be noted that members of ANSI T1E1.4 are presently working to define a spectral management standard that may affect the specified ANSI PSD masks.

ETSI defines a set of PSD masks, each of which applies to a particular deployment scenario. When VDSL is deployed from the ONU (in the FTTCab architecture), the same mask is applied to both the downstream and upstream directions. Separate upstream and downstream masks are defined for when VDSL is deployed from the CO or local exchange. Additional differences between the various masks are introduced in the frequency range from 0 to 276 kHz based on whether POTS or ISDN (or perhaps neither) also resides on the VDSL line. More aggressive PSD masks are defined for use in plants in which most or all cables are buried and emissions into the over-the-air spectrum are not anticipated. Finally, with all masks ETSI requires modems to be capable of reducing their transmit PSDs to no higher than -80 dBm/Hz within the amateur radio bands.

5 Additional system considerations

This section explains why VDSL implementations must ensure spectral compatibility with other DSL services in the same binder, and how spectral compatibility can be achieved. The section also comments on other factors of importance, including power consumption, size and complexity.

5.1 Spectral compatibility

If VDSL is to be viable, it must be spectrally compatible with other services that may reside in the same cable. The DSL most vulnerable to interference from VDSL is ADSL. In certain configurations, VDSL can harm ADSL performance unless precautions are taken. In some other configurations VDSL is vulnerable to interference from ADSL. This section examines how spectral incompatibilities between ADSL and VDSL can occur and determines how minimal interference from VDSL to ADSL can be ensured.

Fiber-to-the-exchange configuration

In the FTTEEx configuration, some twisted pair emanating from a CO or LEx may carry ADSL signals, which are presumed to be transmitted at -40 dBm/Hz PSD, while other lines deliver VDSL signals at -60 dBm/Hz transmit PSD to subscribers closer to the CO or LEx, as illustrated in Figure 16. In this configuration, the impact of VDSL on ADSL performance is negligible, although the reverse is not true. To illustrate, the upstream and downstream ADSL channels are considered separately.




Figure 16. Illustrating the potential for ADSL and VDSL interactions in the FTTEEx configuration

VDSL never affects upstream ADSL transmission in the FTTEEx configuration because the VDSL spectrum begins at 300 kHz, whereas the upstream ADSL channel ends at 138 kHz. Thus, the only interference from VDSL to the upstream ADSL band is out-of-band leakage, which should be -120 dBm/Hz at most in a well-designed VDSL system. Therefore, upstream ADSL transmission is unaffected by VDSL in the FTTEEx configuration. Note, however, that if the VDSL start frequency were lowered to less than 138 kHz, then VDSL could adversely affect ADSL performance. In this case downstream VDSL signals could couple into the upstream ADSL channel as NEXT. Looking at the effect of VDSL on ADSL, if VDSL uses the spectrum below 1.104 MHz to support upstream communications, then VDSL does not impact downstream ADSL at all. On the contrary, in this case downstream ADSL transmissions appear as high-level NEXT in the upstream VDSL channel.

The downstream ADSL channel is affected minimally if VDSL uses the spectrum below 1.104 MHz to support downstream transmission. In this case, VDSL injects FEXT into the downstream ADSL channel. However, because the VDSL PSD is -60 dBm/Hz at most, and VDSL signals couple into ADSL only for the first several hundred meters (usually no more than 1.5 km), this FEXT is far less severe than ADSL self-FEXT. It is well known that ADSL self-FEXT does not significantly impair the performance of ADSL due to the severe attenuation of long loops; thus, it can be concluded that FEXT from VDSL will degrade the ADSL performance even less in the FTTEEx configuration and is therefore not of concern. The effect of ADSL FEXT on the downstream VDSL channel cannot be ignored, however. This FEXT is significantly more severe than VDSL FEXT because the ADSL transmit PSD exceeds the VDSL transmit PSD by 20 dB.

Thus, in the FTTEEx configuration the effect of VDSL on ADSL performance is insignificant, but the reverse is not true. The degradation of VDSL due to ADSL in the FTTEEx configuration is evaluated in detail in [3]. Simulations show that in the FTTEEx configuration, the spectrum below about 2 MHz is virtually useless for supporting upstream VDSL transmission and only marginally more useful for supporting downstream VDSL transmission.

Fiber-to-the-cabinet configuration

Figure 17 illustrates the potential for interference between ADSL and VDSL in the FTTCab architecture. In this situation, an ADSL remote unit and a VDSL remote unit reside in the same binder near the customer premises. If the upstream VDSL spectrum overlaps the downstream ADSL band, then upstream VDSL signals couple into the ADSL downstream as NEXT. If the ADSL band is used by the VDSL system to support downstream transmission, then the VDSL downstream injects FEXT into the ADSL downstream. In either case, VDSL may be detrimental to ADSL performance.




Figure 17. Illustrating the potential for ADSL and VDSL interference in the FTTCab configuration

The effect of VDSL on ADSL performance is quantified in [3]. The paper shows that if upstream VDSL transmissions are supported using the spectrum overlapping the ADSL downstream band, then upstream VDSL signals couple into the ADSL downstream band as NEXT. Simulations show clearly that even a single VDSL operating at -60 dBm/Hz is severely detrimental to ADSL in the FTTCab configuration if the upstream VDSL band overlaps the downstream ADSL band. As an example, the 6-Mbps range of ADSL decreases from approximately 3.75 km to only 2.4 km when a single VDSL is allowed to use the downstream ADSL band to support upstream transmission. From another perspective, the bit rate that can be allocated to a subscriber at the unimpaired 6-Mbps range of 3.75 km decreases to only 3 Mbps when a single VDSL is activated and transmits upstream in the band below 1.104 MHz.

It has been suggested by some that downstream VDSL transmission below 1.104 MHz would not degrade ADSL performance because VDSL signals would appear only as FEXT at the ATU-R. However, simulations in [3] show that using the spectrum below 1.104 MHz to support downstream VDSL transmission adversely affects ADSL performance in the FTTCab configuration. Depending on the length of the VDSL interferer, the effect can be substantial, with shorter VDSL lines degrading ADSL performance more than longer VDSL lines. For example, when a single 1.5-km VDSL line resides in the FTTCab configuration with an ADSL line, the 6-Mbps ADSL range decreases from 3.75 km to 3.2 km. In addition, the bit rate at the unimpaired 6-Mbps range decreases to 4 Mbps. A more severe case is when a 300-meter VDSL injects FEXT into the ADSL line. In this case the 6-Mbps range decreases from 3.75 km to only 2.75 km, and the bit rate at the unimpaired 6-Mbps range decreases to only 3.5 Mbps.

To avoid such substantial degradations in ADSL performance, VDSL must be restricted from transmitting at -60 dBm/Hz either upstream or downstream in the band below 1.104 MHz in the FTTCab configuration when ADSL loops reside in the same binder. In the FTTEx configuration, however, VDSL can transmit below 1.104 MHz, although this spectrum is of little value for upstream or downstream transmission when ADSL is present.

5.2 Power consumption, size and complexity

As described previously, in many applications, VDSL modems will be deployed from the ONU, which is typically located in a small curbside cabinet with no temperature control mechanisms. Because cabinets do not provide cooling, VDSL power consumption must be very low - less than 1.5 Watts per transceiver, including line drivers. To fit in the ONU, VDSL line cards must also be small. The size of a VDSL line card is dependent on a number of factors, including the level to which components have been integrated and the complexity of the design.

6 VDSL solutions

VDSL solutions can be categorized by line code (single-carrier (CAP/QAM) or multi-carrier (DMT)) and by duplexing technique (frequency-division duplexed (FDD) or time-division duplexed (TDD)). This section examines these alternatives.

6.1 Line code alternatives

Although numerous techniques may be used to transmit digital signals on twisted-pair lines, given a finite-complexity constraint and the need to support high bit rates with low latency, there are two practical options: single-carrier modulation with equalization and multi-carrier modulation. This section examines the two approaches as candidates for VDSL.

Single-carrier modulation

The single-carrier modulation family includes both quadrature amplitude modulation (QAM) and carrierless amplitude-phase (CAP) modulation. In single-carrier modulation, a bitstream is encoded into symbols by mapping consecutive sets of b bits, where b is typically less than or equal to 8, into constellation points. These constellation points are then modulated, filtered, and transmitted within some predetermined channel bandwidth. Figure 18 shows a generic QAM transmitter for VDSL, where ω_c is the carrier frequency in radians. CAP transmitters are similar, but the carrier is eliminated and in-phase and quadrature filters are used in place of the lowpass filters.

The channel distorts the transmitted signal and causes successive symbols to interfere with each other, an effect known as *intersymbol interference (ISI)*. At the receiver, a finite-length equalizer is used to reduce ISI, thereby improving the performance of the system while maintaining a reasonable level of complexity. The ability of the equalizer to reconstruct the transmitted data stream depends on a number of factors, including the equalizer type and the length of its filter(s). In general, nonlinear equalizers outperform linear ones, and longer filters yield better performance than shorter filters.

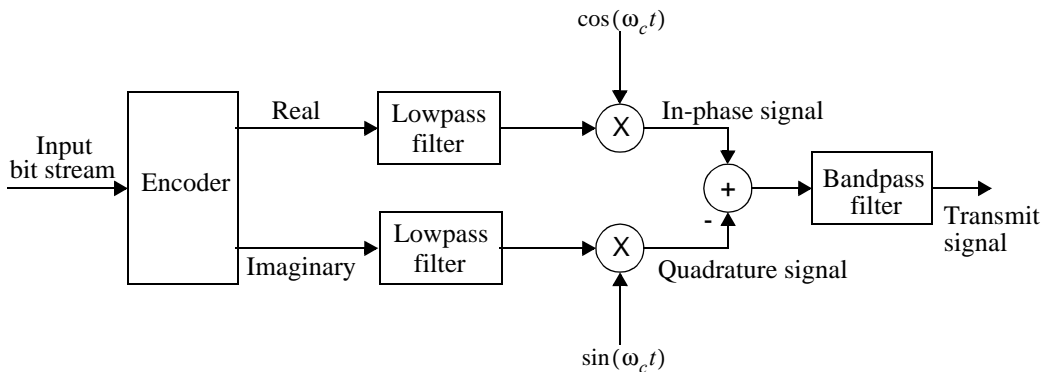


Figure 18. Single-carrier (QAM) transmitter block diagram

The best-performing equalizer is the *decision-feedback equalizer (DFE)*, a type of nonlinear equalizer that uses two filters and a decision device to reconstruct the desired signal. Figure 19 shows a simplified block diagram of the DFE. The slicer determines which of the constellation points is closest to the equalized (but noisy) received symbol. The slicer's decision is then fed back and subtracted from the signal emerging from the feedforward filter. In theory, the feedforward filter eliminates part of the ISI. The feedback filter then produces an estimate of the remaining ISI. When the feedback filter output is subtracted from the feedforward filter output, the remaining ISI should be eliminated.

Although the DFE provides the best performance of all equalizers, DFEs have some disadvantages. First, any errors the slicer makes are fed back and used to generate future estimates of the residual ISI. A single error can

corrupt a large number of future decisions, thus reducing the system's noise margin. As more and more incorrect decisions are made and fed back, the situation worsens. This process is known as *error propagation*. To mitigate error propagation with a DFE, a technique called *precoding* can be used. Precoding eliminates error propagation but requires a change to the transmitter that results in an increase in transmit power. Precoding also eliminates the possibility of blind training, a technique that allows the DFE to configure itself quickly during start-up without communicating with the transmitter.




Figure 19. DFE block diagram

Multi-carrier modulation

In multi-carrier modulation, a channel is partitioned into a set of orthogonal, independent subchannels, each of which supports a distinct carrier. *Discrete multi-tone (DMT)* modulation is a multi-carrier modulation that achieves excellent performance with finite complexity. Because it offers many advantages for transmission on twisted-pair lines, DMT has been standardized world-wide for ADSL. As this section shows, DMT is also an ideal choice for VDSL.

As in any multi-carrier modulation, a DMT transmitter partitions a channel's bandwidth into a large number of subchannels. Each subchannel is characterized by an SNR, which is measured whenever a connection is established and monitored thereafter. The source bitstream is encoded into a set of QAM subsymbols, each of which represents a number of bits determined by the SNR at the midpoint of its associated subchannel, the desired overall error probability, and the target bit rate. The set of subsymbols is then input as a block to a complex-to-real inverse discrete Fourier transform (IDFT), which is often implemented using the fast Fourier transform (FFT). Following the IDFT, a cyclic prefix is prepended to the output samples to mitigate intersymbol interference. The resulting time-domain samples are converted from digital to analog format and applied to the channel. At the receiver, after analog-to-digital conversion, the cyclic prefix is stripped, and the noisy samples are transformed back to the frequency domain by a DFT. Each output value is then scaled by a single complex number to compensate for the magnitude and phase of its subchannel's frequency response. The set of complex numbers, one per subchannel, is called the *frequency-domain equalizer (FEQ)*. After the FEQ, a memoryless (that is, symbol-by-symbol) detector decodes the resulting subsymbols. Thus, in contrast to CAP/QAM systems, DMT systems do not suffer from error propagation because each subsymbol is decoded independently of all other (previous, current, and future) subsymbols. Figure 20 shows a block diagram of a DMT transmitter and receiver pair.

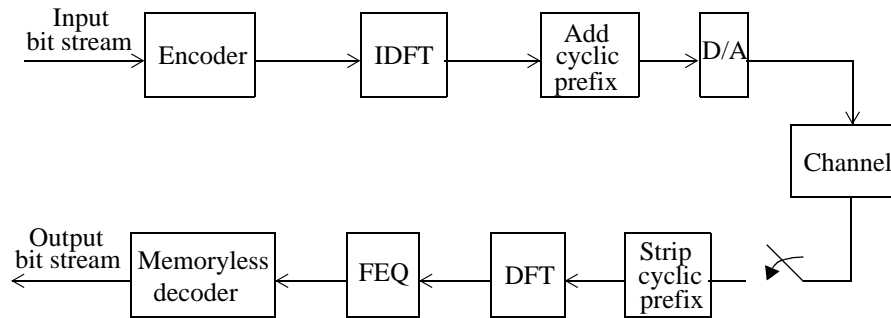


Figure 20. DMT transmitter/receiver pair block diagram

During steady-state operation, the subchannel SNRs are monitored in a data driven manner by the receiver. Upon detecting degradation in one or more subchannel SNRs, the receiver computes a modified bit distribution that better achieves the desired error performance. Depending on the SNR of a degraded subchannel, some or all of its bits may be moved to one or more other subchannels that can support additional bits (that is, subchannels with noise margins greater than zero). The required bit distribution change is reported to the transmitter, where it is implemented. This technique is known as *bit swapping*, and it allows the system to maintain near-optimal system performance as the channel and noise conditions change. Bit swapping algorithms are designed to provide the maximum noise margin at the desired bit rate and error probability.

6.2 DMT vs. CAP/QAM: Which approach meets the VDSL system requirements?

This section examines the advantages and disadvantages of the DMT and CAP/QAM approaches relative to the VDSL system requirements described earlier.

Operation on difficult channels

As mentioned previously, the performance of the single-carrier DFE receiver is dependent on the lengths of the feedforward and feedback filters. Unfortunately, equalizers with short (and therefore low-complexity) filters do not perform well on channels that have significant deviations in their insertion losses, such as telephone lines in bridged-tap configurations.

In contrast, DMT operates effectively on these difficult channels. Because the channel is partitioned into a large number of narrow subchannels, the SNR across any single subchannel is nearly flat. Subchannels overlapping low-SNR regions caused by bridged taps or high noise levels are allocated fewer bits of information than high-SNR subchannels. Furthermore, on long lines that are severely attenuated at high frequencies, DMT modems simply turn off subchannels that are too attenuated to support data. On any channel, therefore, a DMT system can send as much data as the channel allows using only those regions in which data should be transmitted.

Resilience to ingress

The resilience of CAP/QAM to radio-frequency ingress is closely tied to the length of the DFE's feedforward filter. The feedforward filter must generate notches at the locations of all radio-frequency interferers to reduce the probability these disturbers cause slicer errors. As the number of required notches grows, so must the length of the feedforward filter. A fixed-length filter can reject only a certain number of radio-frequency interferers at certain power levels.

In a DMT system, the effects of unpredictable noise sources such as radio-frequency ingress from over-the-air AM and amateur radio transmissions are mitigated in part because the channel is partitioned into narrow subchannels. An infinitely narrow interferer located at the center of the subchannel degrades only that subchannel. Of course, most interferers cannot be considered "infinitely narrow," nor do they usually fall centered within a subchannel. As a result, interferers can bleed into several subchannels. To mitigate the effects of radio-frequency ingress, a window

can be applied to the received signal before demodulation. TI has developed an additional, proprietary receiver technique that contains RF ingress to only a few subchannels to improve performance further.

Robustness to impulse noise

In the presence of certain types of impulse noise, DMT systems offer better immunity than CAP/QAM systems. An examination of TI's VDSL system illustrates why. The system partitions the band from 0 to 11.04 MHz into 256 subchannels, which means each subchannel is 43.125 kHz wide. Accounting for the cyclic prefix, the symbol rate is 40 kHz. Thus, each symbol spans 25 μ s. A CAP/QAM system using the same bandwidth has a symbol rate of 11.04 MHz and a symbol duration of 90.6 ns.

Now assume an impulse of duration 5 μ s hits a channel on which TI's VDSL system is operating. At most, this impulse corrupts one-fifth of a single DMT symbol. (Alternatively, it might corrupt the end of one symbol and the beginning of the next.) In the receiver, when the affected DMT symbol is transformed back to the frequency domain by the DFT, the impulse is spread out over the entire DMT symbol. In essence, this operation increases the effective duration of the impulse but substantially reduces its power. As long as the average power during the DMT symbol period does not consume the noise margin, the system does not suffer detection errors. Thus, the effects of most impulses that are shorter than the duration of a DMT symbol are mitigated automatically by the receiver.

In contrast, a 5 μ s impulse corrupts 55 consecutive symbols of the CAP/QAM system that operates in the same bandwidth as TI's VDSL system. Because CAP/QAM symbols are decoded one by one in the receiver, these systems offer no inherent protection against impulse noise. Instead, they must rely solely on forward error correcting codes and interleaving to protect data from degradations caused by impulse noise. Because impulse noise causes error propagation in the DFE, the effectiveness of the forward error correcting code (which operates on the data after it emerges from the DFE) is reduced.

DMT systems provide impulse noise immunity in another way, too. Although impulse noise is theoretically constant with frequency, in reality most impulses are spectrally shaped. When a DMT system is corrupted by a shaped impulse, not all subchannels are affected equally. Because forward error correcting codes are applied across the subchannels, the frequency selectivity of an impulse enables the code to correct more errors than it could if the impulse noise were indeed flat. CAP/QAM systems offer no such advantage when impulse noise is spectrally shaped.

Adaptivity to changing channel and noise conditions

CAP/QAM systems must rely on their noise margins and proper adaptation of the DFE filters to cope with changes in the channel insertion loss and noise. In contrast, DMT systems use bit swapping and adaptive FEQs to accommodate changes in a line's transfer function or noise that don't reduce its capacity substantially. Of course, if the channel capacity degrades to a point that the data rate used in an established connection can no longer be supported, both DMT and CAP/QAM fail.

Support of required data rates

Because CAP/QAM systems transmit data within a band that is substantially wider than a single DMT subchannel, the granularity in data rates a CAP/QAM system can support is coarse. If a required data rate falls between two possible rates, a CAP/QAM system must support the requirement using the higher available rate. As an example, if a system uses a bandwidth of 3 MHz to transmit data, then the smallest increase in data rate that can be supported is 3 Mbps³, which is provided when the number of bits per symbol is increased by 1 bit. With 1 bit per symbol, the system supports 3 Mbps; with 2 bits, the rate is 6 Mbps; and so on. If a data rate of 10 Mbps is required, then use of 4 bits per symbol yielding 12 Mbps is required. In the best case, 2 Mbps of available data rate is wasted. In the worst case, if the SNR is such that the maximum rate that can be supported by the channel is 10 Mbps, then a CAP/QAM system transmitting 12 Mbps will be degraded by errors because it is exceeding the capability of the channel. Under such conditions, error propagation is a severe problem that will likely disable the system.

3. This analysis neglects the use of excess bandwidth. In reality, to support data in multiples of 3 Mbps a single-carrier system would consume at least 3.3 MHz.

In contrast, DMT systems offer very fine granularity in data rates because the channel is partitioned into narrow subchannels, and bits are allocated independently to those subchannels. The granularity of data rates depends on the subchannel width. If a system uses subchannels that are 32 kHz wide (excluding the cyclic prefix), then the lowest possible bit rate granularity is 32 kbps. TI's VDSL solution provides 64 kbps data rate granularity.

Egress suppression

To provide adequate egress suppression (that is, to restrict the transmit PSD to -80 dBm/Hz within the amateur radio bands), the CAP/QAM transmitter must generate a notch for each amateur radio band the system overlaps. For example, Tables 3 and 4 indicate that a CAP/QAM system operating in the frequency range from 1 MHz to 11 MHz must generate four notches in the transmit signal's spectrum. At the receiver, notches must be generated in these same locations by the feedforward filter of the DFE to ensure energy (noise) received within the amateur radio bands does not corrupt the desired signal.

In CAP/QAM systems that operate in frequency bands overlapping the amateur radio bands, notches must be generated using filters. As a rule, the length of the filter grows with the number and depth of notches required. Because both the transmit filters and the DFE feedforward filter must generate notches over the amateur radio bands, and the feedforward filter must generate additional notches where radio-frequency ingress appears, the complexities of these filters can grow quickly. Vendors who provide shorter, lower-complexity filters risk egress into the amateur radio bands as well as performance degradations due to radio-frequency ingress.

An alternative CAP/QAM approach positions several CAP/QAM signals between the amateur radio bands, thus eliminating the need for notch filters to suppress egress. Although this approach, which is actually a primitive form of multicarrier, successfully eliminates egress into the amateur bands, separate transmitter and receiver hardware must be provided for each band. As the number of bands grows, so do the system complexity and cost.

In contrast to CAP/QAM systems, by dividing the channel into subchannels DMT systems are inherently well-equipped to meet strict egress requirements in the over-the-air amateur radio bands. Those subchannels overlapping the amateur radio bands can be disabled, which automatically reduces the transmit PSD within these bands to a level of -73 dBm/Hz. TI has developed and implemented an additional, low-complexity digital technique that reduces the PSD to the required -80 dBm/Hz level.

Spectral compatibility

CAP/QAM solutions are spectrally compatible with ADSL in the FTTCab configuration only if both the upstream and downstream bands are located above 1.1 MHz. For this reason, most CAP/QAM solutions use only the spectrum above 1 MHz, regardless of whether ADSL resides in the same cable. In situations when ADSL is not present, these systems cannot take advantage of the (typically high-SNR) bandwidth available below 1 MHz. Providing a separate bandwidth allocation for use when compatibility with ADSL is not required would increase the complexity and cost of a CAP/QAM modem.

In contrast to CAP/QAM, DMT provides tremendous flexibility in choosing the lower and upper frequency edges of the transmit band. Spectral compatibility with ADSL is achieved simply by reducing the transmitted PSD on subchannels below 1.104 MHz. When interference to ADSL lines in the same binder is not of concern, such as in the FTTE configuration, a DMT-based system can allocate more energy to those subchannels to support more bits and thus provide higher performance.

Power consumption

Proponents of CAP/QAM systems generally maintain that the power consumption of DMT greatly exceeds the power consumption of CAP/QAM. In reality, the power consumption of any solution is heavily dependent on its level of integration. Clearly, a solution implemented as an application-specific integrated circuit (ASIC) consumes less power than a solution implemented entirely with discrete components. Furthermore, if a solution uses a chipset that does not provide all the functionality required for VDSL modems (such as FEC, the ability to notch the amateur radio bands, providing additional notches to protect against radio-frequency ingress, etc.), additional components are required to meet the VDSL system requirements. Incorporating these components into a system can increase its power consumption substantially.

Furthermore, power consumption depends not only on the line code and its level of integration, but also on what type of duplexing is used. The impact of the duplexing choice on power consumption is addressed in an upcoming section.

Because power consumption is a complex mixture of a number of variables, sweeping statements about the relative power consumptions of CAP/QAM and DMT simply cannot be made.

The best line code for VDSL is DMT

The preceding comparisons show clearly that when the VDSL system requirements are considered, DMT is preferable to CAP/QAM. DMT provides:

- Near-optimal performance on all channels, including those with bridged taps and severe noise profiles
- Resilience to radio-frequency ingress
- Inherent robustness to impulse noise
- Adaptivity to changing channel and noise conditions
- Support of all required ANSI and ETSI data rates
- Egress suppression within the amateur radio bands
- Spectral compatibility with ADSL and other DSL

DMT is clearly the line code of choice for VDSL.

6.3 Duplexing alternatives

Duplexing defines how a system supports bi-directional data transmission. Two duplexing approaches have been proposed for VDSL. They are *frequency-division duplexing* and *time-division duplexing*.

Frequency-division duplexing (FDD)

Most frequency-division duplexed systems define two or more channels, at least one each for upstream and downstream transmission. These channels are disjoint in frequency; hence the name “frequency-division duplexing.” Critical to the performance of FDD systems are the bandwidths and placement in frequency of the upstream and downstream channels.

Figure 21 illustrates the simplest FDD case, which provides a single upstream band and a single downstream band. As the figure illustrates, the upstream band may reside above or below the downstream band. A viable VDSL system must operate on lines that range in length from 300 meters to 1.5 kilometers or perhaps even longer. As Section 3.1 showed, line attenuation increases more rapidly with frequency on longer loops. Thus, the useful frequency band of a line decreases as its length increases. To allow successful transmission, both the upstream and downstream VDSL channels must be located within the useful frequency band. If the useful frequency band falls below f_L , then either the upstream or downstream channel disappears, and bidirectional transmission is not possible.

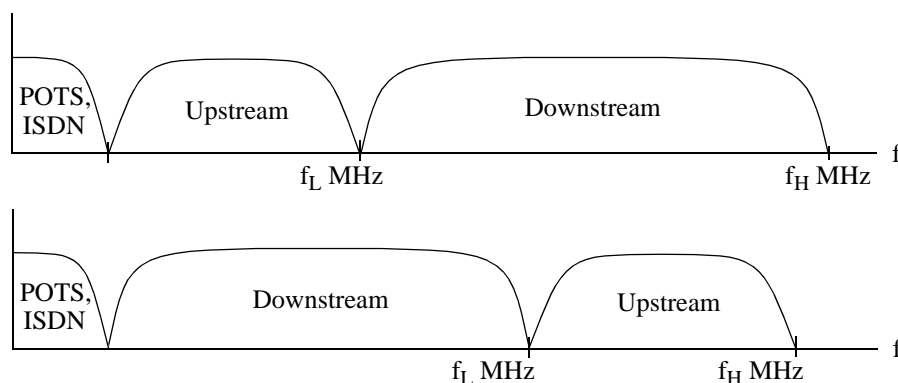


Figure 21. Typical placements of the downstream and upstream channels in FDD

Another primary consideration when designing an FDD system is the bandwidths of the upstream and downstream channels. The appropriate choices for the bandwidths depend on the desired data rates and downstream-to-upstream data rate ratio. The appropriate bandwidth allocation for asymmetric 8:1 data transport differs substantially from the appropriate bandwidth allocation to support symmetric data. Choosing the downstream and upstream channel bandwidths is complicated further by the variability of line length and, as a result, the variability of the SNR profile and useful frequency band. For example, the appropriate bandwidth allocation for 8:1 service on a 300-meter line is very different from the appropriate allocation for 8:1 service on a 1.5-km line.

To support a wide range of data rates and data rate ratios, and to support bi-directional transmission on the required range of loop lengths, FDD systems must provide variable-bandwidth downstream and upstream channels, which generally increases system complexity, particularly in the analog filters. The exception is an FDD implementation of DMT, which allows arbitrary downstream and upstream subchannel allocations by providing a full-bandwidth set of subchannels in each direction. Each subchannel is then used in either the downstream or upstream direction. Although this technique requires two full-sized DFTs in all modems, which increases the system's digital complexity, it eases analog requirements and provides tremendous flexibility in FDD bandwidth allocations.

Time-division duplexing (TDD)

In contrast to FDD VDSL solutions, which separate the upstream and downstream VDSL channels in frequency, time-division duplexed (TDD) systems support upstream and downstream transmissions within a single frequency band during different time periods. Figure 22 illustrates the single frequency band used by TDD systems. Use of the time-shared channel bandwidth is coordinated using *superframes*. A superframe consists of a downstream transmission period, a guard time, an upstream transmission period, and another guard time. The durations of the downstream and upstream transmission periods are integer numbers of DMT symbol periods. Superframes are denoted as A-Q-B-Q, where A and B are the number of symbol periods allocated for downstream and upstream transmission, respectively, and the Qs represent quiescent (guard) times that account for the channel's propagation delay and allow its echo response to decay between transmit and receive periods.



Figure 22. A single band supports both downstream and upstream transmission in TDD systems

In TI's TDD VDSL implementation, the duration of a superframe is 20 symbol periods (500 μ s). The sum of A and B is 18 symbol periods; the sum of the Qs is 2 symbol periods. The values of A and B are chosen by the operator to yield the desired downstream-to-upstream data rate ratio. For example, if the noise profiles in the upstream and downstream directions are assumed to be equivalent, setting A equal to B results in a configuration that supports symmetric transmission. Setting A = 16 and B = 2 yields an 8:1 downstream-to-upstream bit rate ratio. When A = 12 and B = 6, 2:1 transmission is supported. Figure 23 illustrates the superframes that support 8:1, 2:1, and 1:1 transmission.

The use of superframes enables TDD systems to compensate for differences in the downstream and upstream noise levels. For example, if the noise in the upstream direction is more severe than in the downstream direction, a TDD system can allocate additional symbols to the upstream direction to compensate. If symmetric transmission is required, then an 8-Q-10-Q superframe can be used instead of the nominal 9-Q-9-Q superframe, resulting in increased range at a given data rate.

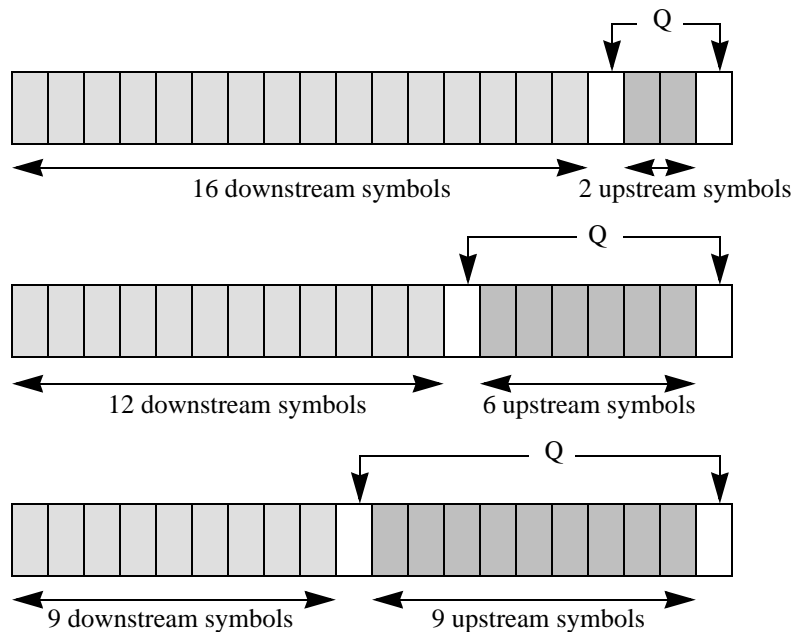


Figure 23. TDD superframes enable support of a variety of downstream-to-upstream data rate ratios

Use of TDD requires that modems on lines within a single binder group be synchronized to a common superframe clock so that all downstream transmissions occur simultaneously on all lines, and all upstream transmissions occur at approximately the same time on all lines. If a common superframe structure is not used, lines supporting TDD in a binder group can inject NEXT into one another, reducing the data rates they can support.

To ensure harmonious operation, all TDD modems at the CO/LEx or ONU must synchronize to a common superframe clock. There are a number of methods by which this clock can be provided; for example, it can be derived from the 8 kHz network clock, sourced by one of the TDD modems, or derived using GPS technology.

6.4 TDD vs. FDD: Which approach meets the VDSL system requirements?

This section examines the advantages and disadvantages of the TDD and FDD approaches relative to the VDSL system requirements.

Mixing symmetric and asymmetric services

VDSL differs from all other DSL because it supports both symmetric and asymmetric data rate ratios. Given that the FDD bandwidth allocations for symmetric and asymmetric services differ, as do the appropriate TDD superframe structures, the question of whether symmetric and asymmetric services can reside in the same binder arises naturally.

Unfortunately, when the optimal time/frequency allocations are used, spectral incompatibilities result regardless of whether a system is TDD or FDD. To illustrate, Figure 24 shows example spectral allocations for an FDD system. The upper allocation supports symmetric transmission, and the lower allocation supports 8:1 transmission. The shaded portion is the frequency band in which NEXT between lines occurs. Thus, mixing symmetric and asymmetric FDD VDSL systems causes NEXT in part of the frequency band, but all the time. In an attempt to provide spectral compatibility between symmetric and asymmetric services, suboptimal downstream and upstream channel allocations have been proposed. Unfortunately, there is no single spectral allocation for FDD that supports both symmetric and asymmetric services without a performance degradation to one or both types of service.

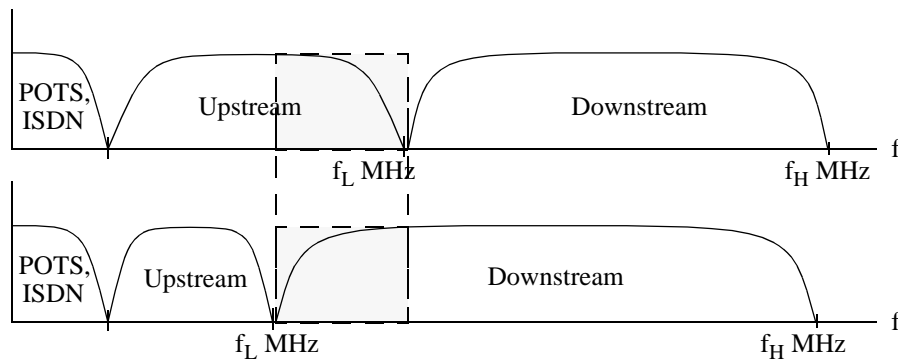


Figure 24. NEXT caused by mixing symmetric and asymmetric FDD systems

TDD systems suffer from a similar degradation when symmetric and asymmetric superframe structures are mixed in a binder. Figure 25 illustrates the worst case, which is when a line supporting 8:1 transmission with a 16-Q-2-Q superframe resides near a line supporting symmetric transmission with a 9-Q-9-Q superframe. Note that the 9-Q-9-Q superframe has been shifted in time by one symbol period to minimize overlap between the downstream symbols on the 8:1 line and the upstream symbols on the symmetric line. However, five symbols are still corrupted by NEXT.

Whereas NEXT in the FDD case spans only part of the bandwidth all the time, NEXT with TDD spans the entire bandwidth but only part of the time. The severity of the NEXT in either case is dependent on how different the mixed ratios are.

The point of this section is that spectral incompatibility between symmetric and asymmetric lines is not a deficiency of either FDD or TDD. Rather, it is a problem caused by the specification of both symmetric and asymmetric data rate ratios in VDSL.

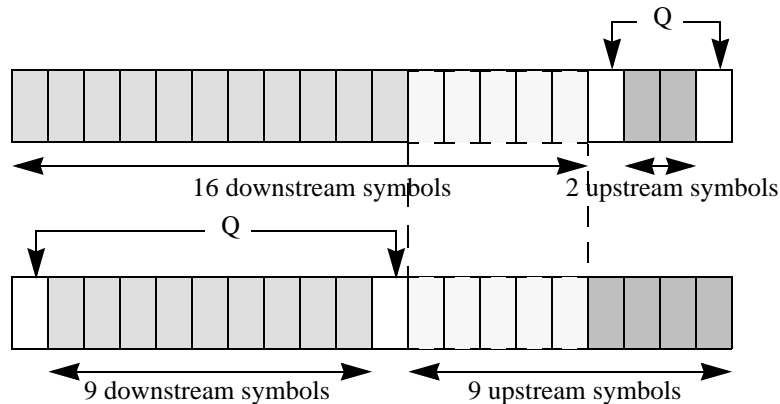


Figure 25. NEXT caused when symmetric and asymmetric TDD superframes are mixed

Support of required data rate ratios

The superframe structure used in TDD enables support of both symmetric and a wide range of asymmetric downstream-to-upstream bit rate ratios with a single transceiver. The desired bit rate ratio is determined for the most part by setting the software-programmable values of A and B to the appropriate values. Furthermore, if the downstream and upstream noise profiles differ substantially, the superframe structure can be modified to compensate for the difference.

Most FDD modems are typically less proficient at supporting multiple downstream-to-upstream data rate ratios.⁴ To support multiple ratios in a reasonable fashion, FDD modems must be capable of changing the bandwidths of the upstream and downstream channels, which usually requires nimble band-splitting analog filters. The complexity and power consumption of a transceiver are dependent on the complexities of the band-splitting filters, which is proportional to the number of bandwidth options the modem provides. As Section 4.1 showed, together ANSI and ETSI require support of the data rate ratios 1:1, 3:1, 4:1, 6:1, and 8:1. The complexity of a transceiver capable of supporting all these ratios on all possible line lengths would be prohibitive. To reduce complexity, many vendors choose to provide separate hardware for different downstream-to-upstream data rate ratios. However, use of separate hardware limits service alternatives.

To illustrate why hardware capable of supporting multiple data rate ratios is desirable, consider VDSL deployment in densely-populated regions and large cities. In such places, the lines of business and residential customers may reside in the same binder. Most operators agree that business customers will request symmetric service, whereas residential customers will want asymmetric service to support Internet access, video-on-demand, and the like. As the preceding section showed, mixing symmetric and asymmetric services causes spectral incompatibility for the ideal TDD and FDD time/frequency allocations. Thus, simultaneous support of symmetric and asymmetric data rate ratios is impractical unless a compromise (and suboptimal) time/frequency allocation is used. However, most business customers require service during weekdays, while most residential customers require service in the evening and on weekends. VDSL modems capable of providing multiple data rate ratios would enable operators to offer symmetric services during the day and asymmetric services at night. In this way, operators can provide the desired rates to business and residential customers whose lines happen to reside in the same binder.

Complexity, cost, and power consumption

As mentioned, the complexity of an FDD modem is dependent on the number of bandwidth allocation scenarios it supports. Modems that support all required VDSL data rate ratios on all lines could be extremely complex and costly, particularly if multiple band-splitting analog filters are used.

In contrast, TDD systems can provide reduced complexity, both in digital signal processing and in analog components. For DMT-based TDD modems, these complexity reductions result from sharing hardware common to both the transmitter and receiver. Sharing of hardware is possible because the transmit and receive functions of a DMT modem are essentially equivalent: both require the computation of a discrete Fourier transform (DFT), which is usually accomplished using the fast Fourier transform (FFT) algorithm. Because TDD is used, a modem can either transmit or receive at any particular time. As a result, hardware to compute only one FFT is required per modem. This FFT spans the entire system bandwidth and is active throughout the superframe except during the guard periods. Additional analog hardware savings are realized in TDD modems because the same band is used to transmit and receive. Whichever path is not in use can be turned off, which reduces power consumption. In contrast, modems using FDD always must provide power to both the transmit and receive paths.

It should be noted that use of TDD with CAP or QAM would not yield similar reductions in hardware complexity. Because single-carrier transmitters and receivers are dramatically different, separate transmitter and receiver hardware would need to be provided in each single-carrier modem, and no cost savings would be realized as a result of using TDD. In fact, it can be argued that a TDD CAP/QAM solution would be much more expensive than an FDD single-carrier solution because it would use a much wider bandwidth and would require more complex transmit and DFE filters, as well as a faster digital-to-analog and analog-to-digital converters.

The cost savings provided by a DMT-based TDD system translate into a low-cost, low-power-consumption solution for VDSL.

4. DMT-based FDD modems are an exception. Partitioning the bandwidth into subchannels provides significant flexibility in the downstream-to-upstream data rate ratio.

Synchronization requirements

As mentioned, TDD modems at the CO or ONU must be synchronized so the overall system performance is not compromised by NEXT from line to line. For this reason, a common superframe clock must be available to all VTU-Os. Unfortunately, operators view the distribution of this common clock as a difficult undertaking when unbundled loops are considered. Operators are uncomfortable with assuming the responsibility to provide a common, reliable clock to those who lease lines in their networks. They worry that a common clock failure would render them vulnerable to lawsuits from those companies or individuals who are leasing lines and relying on the common clock for their own VDSL systems. As a consequence, the FSAN (full-service access network) operators have recommended FDD as the VDSL duplexing scheme.

By default, the only practical duplexing scheme for VDSL is FDD

When the VDSL system requirements are considered, TDD is clearly preferable to FDD for VDSL, particularly when the underlying modulation is DMT. TDD enables provision of a flexible, low-power solution. However, the need to synchronize TDD modems presents a difficult logistical problem for service providers in an unbundled environment. Thus, by default, the duplexing scheme for VDSL must be FDD.

7 Performance of DMT-based TDD VDSL

The preceding sections of this paper showed that from the perspective of meeting the system requirements for VDSL, the combination of DMT modulation with time-division duplexing is the best solution. This section presents measured results that quantify the real-world performance of such systems.

Maximum range tests quantify the maximum range of a system when operating at specific downstream and upstream data rates with desired downstream and upstream noise margins. Figure 26 plots the average⁵ measured maximum ranges of a DMT-based TDD VDSL system operating at various symmetrical data rates while sustaining downstream and upstream noise margins of 6 dB and bit error rates no higher than 10^{-7} . The plot reveals that a symmetrical data rate of 13 Mbps can be supported on 0.5-mm (24 AWG) lines up to nearly 1.4 km in length and on 0.4-mm (26 AWG) lines up to 1.1 km long. A symmetrical rate of 26 Mbps can be sustained on 0.5-mm loops up to 850 meters in length and on 0.4-mm lines up to 700 meters long.

Figure 27 plots the average maximum range of a DMT-based TDD VDSL system when configured to support 8:1 asymmetrical transmission with 6 dB downstream and upstream noise margins and bit error rates no higher than 10^{-7} in both directions. The rate combination of 26 Mbps downstream and 3.2 Mbps upstream is sustainable on 0.5-mm loops up to 1.3 km long and on 0.4-mm lines up to 1.05 km long. Support of 52 Mbps downstream and 6.4 Mbps upstream is possible on 0.5-mm loops up to 700 meters in length and on 0.4-mm loops up to 600 meters long.

5. At least five measurements from different sites around the world were used to generate the average values.




Figure 26. Average maximum range of DMT-based TDD VDSL in symmetric mode with 6 dB downstream and upstream noise margins^a

a. Based on measured performance of TI Vivace SDMT VDSL modems.




Figure 27. Average maximum range of DMT-based TDD VDSL when in 8:1 asymmetrical mode with 6 dB downstream and upstream noise margins^a

a. Based on measured performance of TI Vivace SDMT VDSL modems.

8 Summary

Equipment manufacturers who are evaluating competing VDSL solutions need to be fully informed about the transmission environment and standards requirements, as well as how competing line code and duplexing schemes handle these factors. To aid in the evaluation, this paper began with an overview of VDSL lines and the impairments that complicate high-speed transmission. After a discussion of the requirements for VDSL systems, the paper then described the CAP/QAM and DMT line codes and evaluated their suitability for VDSL, concluding that DMT is preferable. The discussion then turned to the FDD and TDD duplexing schemes, with the conclusion that TDD provides greater flexibility and consumes less power than FDD, but is impractical in unbundled environments.

Texas Instruments currently provides a low-power VDSL solution that offers the performance optimization of DMT and the flexibility of TDD. Known as synchronized DMT (SDMT), TI's VDSL solution leverages high-performance ASICs. Lab and field evaluations have verified that the SDMT approach provides all the advantages of DMT and TDD described in this paper.

In response to operator recommendations that FDD be the VDSL duplexing scheme, TI is working to design a DMT-based FDD VDSL modem. Customers can be assured that the FDD design will provide the characteristics of all TI's xDSL solutions - best-in-class performance, robust operation, low power consumption and a highly integrated solution.

9 Acronyms

| | |
|--------|---|
| ADSL | Asymmetric digital subscriber line |
| ANSI | American National Standards Institute |
| ASIC | Application-specific integrated circuit |
| AWG | American wire gauge |
| AWGN | Additive white Gaussian noise |
| CAP | Carrierless amplitude/phase modulation |
| CO | Central office |
| DFE | Decision feedback equalizer |
| DFT | Discrete Fourier transform |
| DMT | Discrete multi-tone |
| DSL | Digital subscriber line |
| ETSI | European Telecommunications Standards Institute |
| FDD | Frequency-division duplexed |
| FEC | Forward error correcting (or correction) |
| FEQ | Frequency-domain equalizer |
| FEXT | Far-end crosstalk |
| FFT | Fast Fourier transform |
| FTTCab | Fiber-to-the-cabinet |
| FTTEx | Fiber-to-the-exchange |
| IC | Integrated circuit |
| IDFT | Inverse discrete Fourier transform |
| ISI | Intersymbol interference |
| LEx | Local exchange |
| NEXT | Near-end crosstalk |
| ONU | Optical network unit |
| PSD | Power spectral density |
| QAM | Quadrature amplitude modulation |
| RF | Radio-frequency |
| SDMT | Synchronized discrete multi-tone |
| SNR | Signal-to-noise ratio |
| TDD | Time-division duplexed |
| VDSL | Very high-speed digital subscriber line |

10 References

- [1] ANSI Draft Technical Document. "Very-high-speed Digital Subscriber Lines System Requirements." T1E1.4 contribution 98-043R6. December 1998.
- [2] ETSI Technical Specification TS 101 270-1 v1.1.1. "Transmission and multiplexing (TM); Access transmission systems on metallic access cables; Very high speed Digital Subscriber Line (VDSL); Part 1; Functional requirements." April 1998.
- [3] K.S. Jacobsen. "Interactions Between ADSL and VDSL in the Same Binder: Simulations and Test Results." ANSI T1E1.4 contribution number 99-298. Ottawa, Ontario. June 1999.
- [4] B. Wiese and J.A.C. Bingham. "Digital radio frequency cancellation for DMT VDSL." T1E1.4 contribution number 97-460. Sacramento, CA. December 1997.
- [5] VDSL Alliance, T. Pollet, ed. "Physical Medium Specific Specification for G.vdsl." ITU SG15 contribution number MA-050. March 1999.
- [6] J.A.C. Bingham. "Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come." *IEEE Communications Magazine*, vol. 28, no. 5, May 1990, pp. 5-14.
- [7] C. Modlin, et. al. "Complexity of VDSL." ITU contribution number MA-063. March 1999.
- [8] J.M. Cioffi, et. al. "Very-High-Speed Digital Subscriber Lines." *IEEE Communications Magazine*, vol. 37, no. 4, April 1999, pp. 72-79.
- [9] T. Starr, J.M. Cioffi, and P.J. Silverman. *Understanding Digital Subscriber Line Technology*. Prentice-Hall, Upper Saddle River, NJ, 1999.
- [10] VDSL Alliance website: www.vdslalliance.com.