A Discrete Multitone Transceiver System for HDSL Applications

Jacky S. Chow, *Student Member, IEEE*, Jerry C. Tu, *Student Member, IEEE*, and John M. Cioffi, *Senior Member, IEEE*

***Abstract*—A discrete multitone (DMT) transceiver design for high-bit-rate digital subscriber line (HDSL) access is presented and analyzed. The DMT transmitter and receiver structure and algorithms are detailed, and the computational requirements of DMT for HDSL are estimated. We find that at a sampling rate of 640 kHz, using an appropriate combination of a short FIR equalizer and a length-512 DMT system, 1.6 Mb/s data transmission is possible within the CSA at an error rate of 10^{-7} on a single twisted pair. We also show significant performance margin can be achieved when two coordinated twisted pairs are used to deliver a total data rate of 1.6 Mb/s. We conclude that, in terms of a performance-per-computation figure of merit, the DMT system is an excellent candidate for HDSL implementation.**

I. Introduction

MULTITONE modulation methods use an optimized frequency division allocation of energy and bits to maximize the reliably achievable data rates that can be transmitted over bandlimited communication channels. These systems are easily described because they use frequency division multiplexing to transport a single-input data stream on several carriers within the usable frequency band of the channel. The simplicity of multitone methods accounts for their use in some of the earliest data transmission modems, such as the Collin's Kineplex modems of 1957 [1]. While the simplicity and appeal of multitone modulation has existed for decades, it is only very recently that their optimal performance has been proven theoretically (see [2], [3]) and demonstrated in practice with recent products, such as Telebit's Trailblazer modems [4] or NEC's multitone voiceband and groupband modems [5]. These highest-performance products are sold at only a fraction of the cost of comparable performance single-channel or QAM modems.

The excellent high-performance/cost tradeoff of multitone modulation also makes it a strong candidate for the transceiver implementation for high-bit-rate digital subscriber lines (HDSL), a 1.6 Mb/s digital subscriber service on twisted-pair channels of up to two miles. This paper presents a multitone modulation structure that is of particularly high performance and low cost for the HDSL application. We generally find that, even on worst-case channels, a 1.6 Mb/s (1.536 Mb/s user data plus 64 kb/s control channel) data rate is achievable at an error rate of 10^{-7} on a single twisted pair (with 6 dB margin). Furthermore, we find that the described multitone system exhibits excellent margins (nearly 18 dB) for the so-called dual-duplex option of providing 1.6 Mb/s on two coordinated twisted pairs.[1]

The intersymbol interference, crosstalk, and attenuation of 24-gauge and/or 26-gauge twisted pair is particularly severe within the two-mile carrier serving area (CSA) requirements, as we discuss in Section II. Section II also briefly examines and reviews theoretically and practically achievable maximum data rates for HDSL. Equalization and coding methods have to be carefully applied in the HDSL application to render the high transported data rates reliable. In the approach presented, a particularly effective combination of a small amount of time-domain equalization and of a larger amount of multitone modulation is used to transmit the data reliably. The equalizer is used to shorten the effective response length of the data channel, while the multitone modulation is used to mitigate the remaining intersymbol interference and the effects of crosstalk, as is described in Section III. The multicarrier system is fully adaptive, employing both an adaptive algorithm for update of the adaptive filters and also a special channel identification procedure that is used to initialize and update the transmitter bit allocation to frequency-indexed subchannels. This relatively simple procedure is essential for maximizing performance. We also estimate the complexity in terms of the number of arithmetic operations in Section III and illustrate that the multitone transceiver can be implemented with programmable signal processors, although those processors may need to be custom designed for lowest cost.

The performance in achievable data rate of the multicarrier system is analyzed and characterized in Section IV as a function of input power, channel characteristic, number of tones or carriers, sampling rate, noise margins, and crosstalk.

Manuscript received November 14, 1990; revised May 17, 1991. This work was supported in part by a gift from Bell Communications and by a contract with the University Technology Transfer Institute.

The authors are with the Information Systems Laboratory, Department of Electrical Engineering, Stanford University, Stanford, CA 94305.

IEEE Log Number 9102134.

[1]We note that results in this paper use channel models with less attenuation than those of previous studies in [6]-[9], based on measured line results of several industrial sources.

0733-8716/91/0800-0895$01.00 © 1991 IEEE

II. Subscriber Loop Characterization and Modeling

In any study of HDSL transceivers, the channel models are critical to specification of performance. We outline in detail the HDSL channel models used throughout the study of DMT HDSL transceivers. We note that performance projections based on different channel models can lead to erroneous claims and comparisons. A detailed tutorial on HDSL technologies is presented by Lechleider in this issue [10].

A. CSA Loop Characteristics

The carrier serving area (CSA) is an identifiable subset of the current subscriber loop population in the U.S., and our HDSL study is restricted to loops that satisfy minimal CSA requirements. The CSA consists of mainly 24- and 26-gauge twisted pair, with lengths on 24-gauge wire to 12 kft and on 26-gauge wire to 9 kft. CSA loops can also contain bridge taps of limited length; a complete specification of the CSA appears in [11]. Impairments in the CSA data transmission channel include:

- intersymbol interference (ISI),
- crosstalk noise coupled from adjacent loops within the same cable bundle,
- impulsive noise (caused by switching transients, lightning, and other electrical machinery),
- echo noise resulting from the combination of imperfect hybrids and gauge changes on the line as well as signals reflected from bridged taps,
- electronics noise (such as quantization noise in analog-to-digital converters (ADC) and thermal noise in the analog portion of the transmitter and receiver), and
- inductive noise at 60 Hz and its harmonics caused by nearby power lines.

We consider a set of representative CSA channels, provided by Bell Communications Research [12], that have the respective loop cable arrangements illustrated in Fig. 1. We use a letter designation (A–H) on these channels for reference. These channels are generally representative of some of the worst, in terms of channel impulse response, expected to be found within the CSA limits. In general, longer wire length, smaller wire size, and presence of more and longer bridge taps results in poorer achievable data rates. We also note that on-premises wiring is often impossible to characterize statistically, so a 6 dB margin is imposed upon the performance of any HDSL transceiver to ensure that 10^{-7} error rate is maintained [13]. In all loop models in this paper, we include this 6 dB margin simply by increasing loop attenuation for the loops in Fig. 1 by 6 dB (see also Footnote 1).

B. Channel Modeling

A line simulation program, based on standard transmission line modeling and a standard set of twisted-pair characteristics, was used to generate an impulse response for each of the channels in Fig. 1. The effect of the coupling transformer at each end of the loop termination was

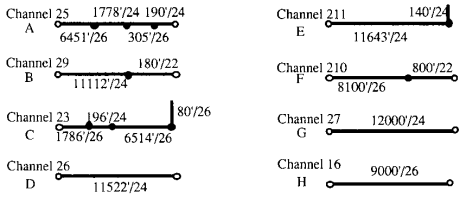

Fig. 1. CSA channels: length (in ft.)/gauge.

included by adding a double-pole double-zero highpass digital filter with a cutoff frequency of 300 Hz. These characteristics were found to be in near exact agreement with those provided by separate programs at various industrial groups working in the HDSL field [14], [12]. One such time-domain impulse response is shown in Fig. 2 for Channel H.

Crosstalk noise coupled from adjacent wire pairs in the cable bundle is the dominant noise impairment in CSA loops. Various characterizations have been proposed [15], [16], and we will use [15], which models crosstalk noise as near end crosstalk (NEXT) with the following transfer function:

$$|H_{\mathrm{NEXT}}(f)|^2 = K_{\mathrm{NEXT}} f^{3/2} \tag{1}$$

where K_{NEXT} is the coupling coefficient. Since far end crosstalk noise is attenuated by the channel, the resulting noise power is negligible compared to near end crosstalk. Hence, we will only consider NEXT. The coupling coefficient K_{NEXT} for a 50-pair cable was empirically determined to be approximately 10^{-13} [17].

To model the effects of residual echo after cancellation, we note that a typical CSA loop attenuation is on the order of 24 dB, and a good echo canceller can presumably reduce the echo to a level of 30–40 dB below the received signal. Assuming an average input power of 10 mW (10 dBm), this translates into -44 to -54 dBm of total echo noise power across the two-sided bandwidth of 640 kHz, which corresponds to a noise power spectral density of 6.22×10^{-11} to 6.22×10^{-12} mW/Hz. This residual echo noise is assumed to be additive white Gaussian noise (AWGN), and we fix this AWGN to have a power spectral density of -110 dBm/Hz for subsequent analysis and simulation. The electronics noise, as mentioned earlier, is typically on the order of -140 to -170 dBm/Hz [6] and is thus, by comparison to echo noise, negligible. Of course, crosstalk noise is larger than either echo or electronics noise.

While impulsive noise is more difficult to characterize, some [18] have attempted to model its effect. Impulsive noise usually consists of millivolt pulses (say, up to 40 mV) that last a few μs to 1000 μs.

Inductive noise can be modeled as a fixed noise level at 60 Hz, but the highpass filter simulating the transformers eliminates the inductive noise.

1) Pole-Zero Modeling: To facilitate simulation of CSA loops for performance analysis, we constructed a pole-zero (IIR) model for each channel impulse response

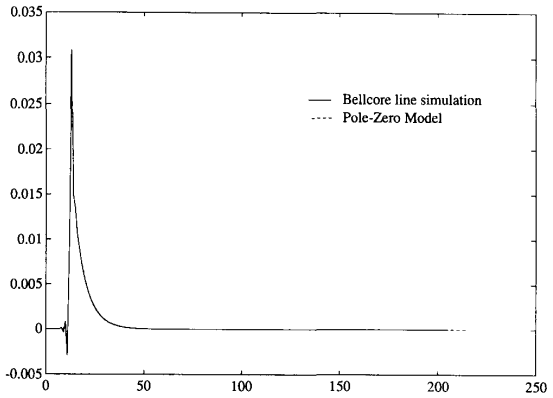

Fig. 2. Impulse response of a CSA loop.

at various sampling rates, using standard system identification tools. An impulse response obtained for Channel *H* is compared to the impulse response generated by a pole-zero model in Fig. 2. Since the pole-zero model closely approximates the true response, the use of pole-zero models for subsequent studies is justified.

The number of poles and zeroes varies with different CSA channels, but they all fall into the same range: 5–8 zeroes, and 6–12 poles. For future reference, assume the number of zeroes for a particular channel is equal to ν, and the number of poles is equal to $L - 1$. We also found that most of the CSA channels are not minimum phase, i.e., some of the zeroes from the pole-zero models lie outside the unit circle in the z-plane.

C. Capacity

The capacity of these loops with the aforementioned models is computed in [6], at the sampling rates of 400 and 800 kHz. As explained later in Section IV-A and [6], a sampling rate of 400 kHz is not sufficient to maximize the HDSL transceiver performance on most CSA channels shown in Fig. 1 as well as the capacity with a flat input spectrum (called "flat input capacity" in Table I) at 800 kHz sampling rate. The two columns under "Capacity" list the capacity using the optimal transmit power distribution [6], and represent the best data rates that can be achieved with an 800 kHz sampling rate. "Crosstalk Bound" assumes crosstalk noise is the only noise impairment (and, thus, ignores the contribution of AWGN), while "11 dBm Power" includes the effect of AWGN with 11 dBm transmit input power. The three columns under "Flat Input Capacity" do not optimize transmit power distribution; instead, a flat input power spectral density is assumed. The effects of near end crosstalk, 6 dB margin, and AWGN are included for the calculation, and the maximum achievable data rate based on the flat input power spectrum are shown for three different input powers: 1, 10, and 50 mW (0, 10, and 17 dBm, respectively). Notice that, with a sufficiently high input power (10 dBm or above), the penalty in achievable data rate

TABLE I
CAPACITY AND FLAT INPUT CAPACITY FOR CSA LOOPS, 800 KHZ RATE AND WITH 6 DB MARGIN (IN MB/S)

	Capacity with 6 dB Margin		Flat Input Capacity with 6 dB Margin		
Channel	Crosstalk Bound	11 dBm Power	0 dBm Power	10 dBm Power	17 dBm Power
A	2.42	2.30	1.87	2.29	2.40
B	2.39	2.27	1.85	2.27	2.37
C	2.32	2.20	1.78	2.19	2.30
D	2.31	2.19	1.77	2.19	2.29
E	2.26	2.14	1.72	2.13	2.24
F	2.29	2.17	1.75	2.16	2.27
G	2.19	2.07	1.65	2.06	2.16
H	2.14	2.02	1.60	2.01	2.11

(10 dBm or above), the penalty in achievable data rate without optimizing the transmit power distribution becomes negligible. Thus, we will assume a flat input power spectrum is used, and this will simplify analysis and simulation. Also, this allays any practical concern about loop-to-loop energy spectrum variation.

We can infer from this table that 1.6 Mb/s can be reliably transmitted over most CSA channels. However, an excellent transceiver design is required to accomplish this goal on either a single pair or on two coordinated pairs with a large performance margin such as 12 dB. Even with only a 6 dB performance margin on two pairs, 800 kb/s is 35–50% of capacity and more than a trivial receiver is required. On the other hand, a data rate of 160 kb/s on CSA channels (as in the case of basic rate ISDN) is well below the fundamental limits (although it may be a little more difficult on longer (18 kft) channels). Thus, one would not expect that a simple receiver used for basic rate ISDN would necessarily be good for use in HDSL.

III. The Discrete Multitone Transceiver

This investigation focuses on a discrete multitone (DMT) transceiver for HDSL. We briefly describe the modulation and demodulation processes and the major system components. We also accurately compute computational requirements for the DMT HDSL transceiver using Motorola's 56000 DSP as a representative host for implementation. A system using multiple Motorola 56000 DSP's has been used to implement the described DMT transceiver in real-time at Stanford University.

A. System Description

A simplified illustration of the DMT transmitter, receiver, and channel appears in Fig. 3. The basic concept is to use frequency division modulation to divide the transmission system into a set of frequency-indexed subchannels that appear to be modulated and demodulated independently. With a careful allocation of bits and transmit power to the subchannels, it is well known [2], [3] that such a system is capable of performing at the highest theoretical limits and that no other system can exceed its

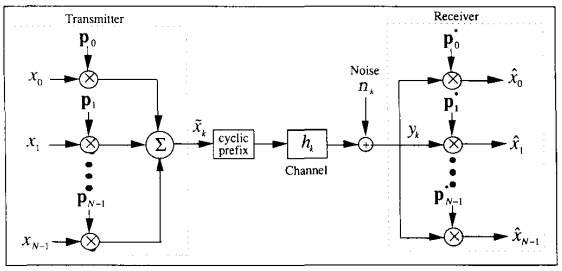

Fig. 3. Simplified illustration of DMT transceiver and channel.

performance. There are now at least three commercial 19.2 kb/s voiceband modems, marketed by three different companies, that use DMT to attain excellent performance. This section describes a DMT system for HDSL.

The system in Fig. 3 is presumed to be sampled (digitized) at a rate that is at least twice as high as any signal frequency, so that sampling theory arguments apply and the discrete-time system is equivalent to the actual continuous-time system. A *block* of input bits is divided into N *subsymbols*, x_n, $n = 0, \cdots, N-1$, that are independently modulated by N-dimensional sampled-sinusoid *modulating vectors* $\boldsymbol{p}_n$, and then summed to form an N-dimensional block of modulated *transmit signal samples* $\tilde{x}_k$, $k = 0, \cdots, N-1$. The nth modulating vector is given by:

$$\boldsymbol{p}_n = [p_{n,0} \quad p_{n,1} \quad \cdots \quad p_{n,N-1}] \tag{2}$$

where

$$p_{n,k} \triangleq \frac{1}{\sqrt{N}} e^{+j(2\pi/N)kn} \qquad k, n = 0, \cdots, N-1 \tag{3}$$

the inverse discrete Fourier transform (IDFT) vector. To ensure a real-valued transmit signal, $\tilde{x}_k = \Re(\tilde{x}_k)$, it is necessary that we impose the constraint on the subsymbols that ($*$ denotes complex conjugate, $\Re$ denotes real part)

$$x_n = x_{N-n}^* \qquad n = 1, \cdots, N-1. \tag{4}$$

It is convenient to think of DMT as consisting of $N/2$ quadrature amplitude modulation (QAM) channels. In practice, of course, the modulation is done digitally and no explicit use of any carrier modulation frequencies or modulators is required. The *channel output samples* y_k are found by the sum of the convolution of the modulated transmit samples, $\tilde{x}_k$, with the sampled channel *impulse response* h_k and the additive noise n_k,

$$y_k = h_k * \tilde{x}_k + n_k. \tag{5}$$

The impulse response h_k is assumed to be of finite length, and can be described by the $\nu + 1$ samples $h_0, \cdots, h_\nu$. Its *D-transform* is $h(D) = \sum_{k=0}^{\nu} h_k D^k$. In the receiver, the corresponding demodulation is accomplished by taking the inner product of the corresponding N channel outputs, $y_0, \cdots, y_{N-1}$, with $\boldsymbol{p}_n^*$. The noise at the channel output presently includes only additive white Gaussian noise; we will generalize to the crosstalking case as this development proceeds.

Because of the finite number of subchannels, a *cyclic prefix* (CP) of the transmitted signal samples $\tilde{x}_k$ is needed once for every transmitted block of bits. The CP prefixes the block of transmitted samples $\tilde{x}_0, \cdots, \tilde{x}_{N-1}$ by the (repeated) samples, $\tilde{x}_{N-\nu}, \cdots, \tilde{x}_{N-1}$. We denote the corresponding time indexes at the beginning of the block by $k = -\nu, -\nu+1, \cdots, -1$. When $\nu + 1$ is the length of the channel impulse response in sample periods, then the successive blocks of transmitted data do not interfere with each other, and one can write (see [3]) that

$$\hat{x}_n = |H_n| x_n + n_n \tag{6}$$

where

$$H_n = h_0 + h_1 e^{-j(2\pi/N)n} + h_2 e^{-j(2\pi/N)n} + \cdots + h_\nu e^{-j(2\pi/N)\nu n}. \tag{7}$$

In practice, the modulation and demodulation processing is implemented with fast Fourier transform (FFT) algorithms that require significantly less computations than would occur in the direct implementation shown in Fig. 3. We discuss specific computational requirements in Section III-C.

A trellis code [19] of coding gain γ can be applied to each subchannel so that the energy allocated to this channel will sustain a greater number of bits than in the absence of the applied trellis code. The number of bits can be approximated by

$$b_n = \log_2 \left(1 + \frac{\mathrm{SNR}_n}{\Gamma}\right) \tag{8}$$

where $\Gamma = 10$ dB $- \gamma$ for a bit error rate of 10^{-7} [20]. If no code is used, $\gamma = 0$ dB and $\Gamma = 10$ dB. The *signal-to-noise ratio* on the nth subchannel is denoted by SNR_n and is computed by

$$\mathrm{SNR}_n = \frac{\epsilon_x |H_n|^2}{\sigma^2} \tag{9}$$

where $\epsilon_x = E|x_n|^2$ is the signal energy of the nth subsymbol and σ^2 is the corresponding white noise energy. In practice with HDSL, the dominant noise is crosstalk noise, which is not white, so that $\sigma^2 \to \sigma_n^2$, a subchannel-dependent noise energy. When N is chosen sufficiently large, the noise on each of the subchannels will be approximately white within the occupied frequency band and will be uncorrelated with the noise on any of the other subchannels. In order to compute the bit allocation outlined by (8), the channel characteristic must be known at the transmitter, which is achieved as described in Section III-A-3. The energy allocation is approximated by a flat

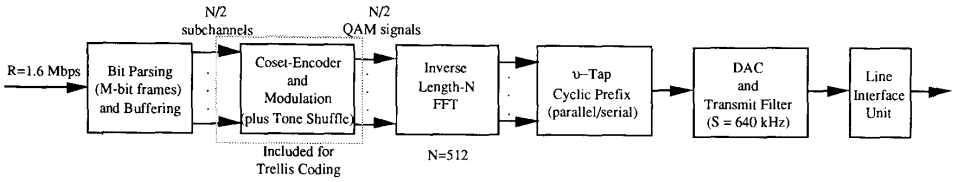

Fig. 4. Block diagram of a specific DMT HDSL transmitter.

transmit energy spectrum, so that ϵ_x is not a function of the frequency index n. Our studies have found that very little is lost in using this flat energy distribution in the HDSL application. The poorest subchannels carry only dummy data and are simply ignored by the receiver. This flat transmit spectrum has the effect of rendering the crosstalk power spectrum from any DMT HDSL transmitter constant in practice.

1) The DMT HDSL Transmitter: A specific embodiment of the HDSL transmitter is illustrated in block diagram form in Fig. 4. In the specific implementation of the system that we study, ν is always limited (through equalization, as discussed later) to a maximum of 8, and later in Section IV-A-3) and 4), we fix the sampling rate S to be equal to 640 kHz and $N = 512$.

For a single twisted pair to carry $R = 1.6$ Mb/s (the *single-duplex* option), the input data bit stream is parsed into M-bit blocks, where $M = (N + \nu)R/S = 1300$. These bits are then transformed into (a maximum of) $N/2 = 256$ QAM subsymbols that are then applied to a trellis encoder. The trellis coder sequentially processes the frequency-indexed subsymbols to avoid the large latency and memory requirement that would occur with multiple trellis encoders. This process requires that the subchannel indexes be "shuffled" to avoid any minor correlation between the noise on adjacent frequency subchannels. The output of the trellis encoder is then modulated, as described earlier, through the use of an $N = 512$ inverse FFT. Finally, an eight-sample cyclic prefix is placed in the beginning of the corresponding block of modulated transmit samples, and the extended block of $N + \nu$ samples is then applied to the channel through the digital-to-analog converter (DAC) and line-interface unit. The sampling rate is 640 kHz, which leaves the 256 subchannels effectively separated by 1.25 kHz. This corresponds to a *block symbol period* of $512 + 8 = 520$ samples or 812.5 μs. Eight samples is considerably shorter than most HDSL channel impulse responses at 640 kHz. An unusual and simple equalizer is used in the receiver to contain intersymbol interference to eight samples, or impulse-response length to nine samples.

For the so-called *dual-duplex* option with two twisted pairs, the 1300-bit blocks are parsed into the 512 QAM subchannels that exist on both of the channels. There are then two IFFT's, two DAC's, and two line-interface units. However, only one trellis code is necessary with this option with respect to the single-duplex option. This improves the performance margin but increases the cost of the dual-duplex option. If a sufficient margin can be obtained using the single-pair solution (called "single-duplex"), then it is preferable from a cost standpoint.

2) The DMT HDSL Receiver: A block diagram of the general DMT HDSL receiver appears in Fig. 5, and a specific receiver that corresponds to the specific transmitter of Fig. 4 is shown in Fig. 6. This receiver structure is a slight variation of the original receiver proposed in [7]. Both receiver structures achieve the same performance but the one in Fig. 6 results in lower complexity.

An actual HDSL channel response is of significant duration in practice, and we can model its polynomial-fraction D-transform

$$\tilde{h}(D) = \frac{a(D)}{p(D)}. \tag{10}$$

In practice, we identify this infinite-impulse-response model, $\tilde{h}(D)$, by choosing $a(D)$ and $p(D)$ to minimize the mean square error between this model and the channel output. Using this method, the crosstalk noise also affects the settings for $a(D)$ and $p(D)$. Due to the nonminimum phase nature of $a(D)$, the use of an IIR equalizer $1/\tilde{h}(D)$ to eliminate ISI results in unstable filter. A short, fixed L-tap feedforward equalizer is thus used to process the sampled channel output y_k sequentially. This equalizer does not remove intersymbol interference completely, but rather confines the impulse response so that its length is approximately $\nu + 1$ sample periods or less. By setting the equalizer to $p(D)$, we will have a minimum mean square error approximation of an additive white noise channel with impulse response characterized by $a(D)$, which has a response length of $\nu + 1$ sample periods or less. The last N samples of the $N + \nu$ samples that correspond to the transmit block are extracted from the equalizer output. An N-point FFT is performed on the equalizer output. The FFT outputs, z_n, $(n = 0, 1, \cdots, N - 1)$, are multiplied by N complex, 1-tap adaptive filters, w_n, $(n = 0, 1, \cdots, N - 1)$ so that a common decision device can be used to estimate the subsymbols on each of the subchannels. The initial tap setting for w_n

subchannels, then two IFFT's, two DAC's

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.