

Signaling Vulnerabilities in Wiretapping Systems

Many law enforcement wiretap systems are vulnerable to simple, unilateral countermeasures that exploit the unprotected in-band signals passed between the telephone network and the collection system. This article describes the problem as well as some remedies and workarounds.



Law enforcement agencies in the US and elsewhere use voice telephone interception systems to collect wiretap evidence and intelligence against criminal and national security subjects. Such systems provide a legal record of the digits dialed by the subject and, in some cases, the audio content of the calls themselves. Wiretapping is often credited as an essential tool in the investigation and prosecution of serious crime, especially when complex criminal enterprises and conspiracies are involved.

Unfortunately, however, many of the telephone interception technologies that law enforcement depends on for evidence collection are less reliable than previously thought. We found that the design and implementation of these systems often render them vulnerable to simple, unilateral countermeasures that allow wiretap subjects (or their correspondents) to prevent accurate and complete capture of call data and contents. These countermeasures exploit the in-band signals passed between the telephone network and the law enforcement agency.

In particular, the evidence collected by virtually all interception systems based on traditional technology, as well as at least some systems based on newer interfaces, can be manipulated by the subject with practical techniques and readily available hardware. We found one countermeasure, requiring only a standard PC, that prevents the accurate recording of dialed telephone numbers and line statuses. Perhaps more seriously, we also found simple countermeasures that effectively and selectively suppress the recording of call audio with only modest degradation of call quality.

Unlike traditional wiretap countermeasures (such as encryption), our techniques are entirely unilateral—they

don't require active cooperation between subjects and their associates—and they obscure not only the content, but also the metadata that indicates the presence of communication and its endpoints in a way that is sometimes difficult to detect. This has implications not only for the accuracy of the intelligence that can be obtained from these taps, but also for the acceptability and weight of legal evidence derived from it.

Our analysis is based entirely on information obtained from published sources and equipment purchased openly in the retail and surplus markets. Thus it is possible (perhaps even likely) that motivated wiretap targets such as those involved with organized crime have already discovered and actively employed them. We recommend that currently fielded telephone interception systems be evaluated with respect to these vulnerabilities and reconfigured or modified where possible to reduce their susceptibility. In addition, the possibility of these or similar countermeasures should be considered in analyzing previously collected wiretap evidence and intelligence.

Despite law enforcement's growing reliance on wiretaps, little attention has been paid in the open literature to their reliability. Indeed, this article could represent the first analysis of the security of modern telephone wiretap systems by the computing and communications research community. Drafts of this article have been made available to the law enforcement community.

Wiretapping and US law

Broadly speaking, the federal laws governing electronic surveillance for criminal (Federal Wiretap Act [Title III]¹) and national security (Foreign Intelligence Sur-

MICAH SHERR,
ERIC CRONIN,
SANDY CLARK,
AND MATT
BLAZE
*University of
Pennsylvania*

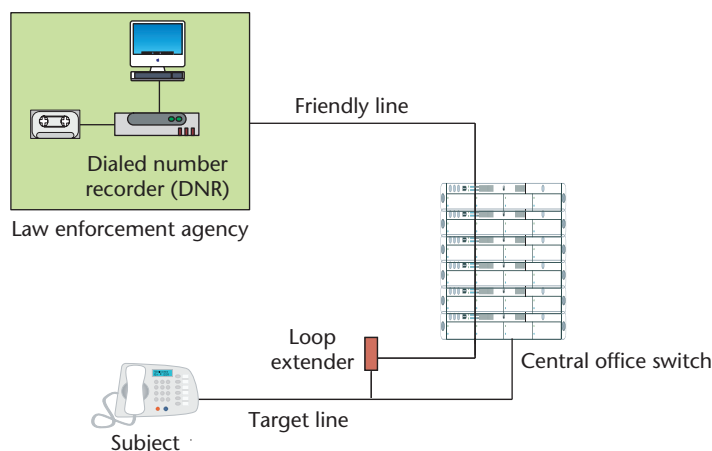


Figure 1. Loop-extender wiretap architecture. The target telephone line is tapped in the field with a loop-extender device, which relays the signals and content to the law enforcement agency over its own telephone line.

veillance Act [FISA]²) investigations authorize two categories of telephone wiretaps for use by US law enforcement agencies.

The first category, called a *dial number recorder* (DNR) or *pen register*, records the digits dialed and other outgoing signaling information. DNR taps, which provide traffic analysis information, but not the call's audio contents or speaker identity, must pass relatively modest judicial scrutiny to be authorized. A related investigative technique, called a *trap and trace* for historical reasons, provides analogous information about incoming calls.

The second category, *full audio interception* (sometimes called a Title III or FISA wiretap depending on its legal context), records not only the dialed digits and signaling, but also the actual call contents. Legal authorization for full audio interception taps entails a higher standard of proof and greater judicial scrutiny. Such taps are also more expensive (and labor intensive) for the law enforcement agency than DNR taps because they generally require continuous real-time monitoring by investigators.

Communication evidence isn't produced exclusively by wiretap interceptions, though: some investigative functions are served by examining telephone accounting and billing data that the carrier has collected. Law enforcement agencies occasionally subpoena telephone records and use them as a source of intelligence or evidence, but these aren't, strictly speaking, interceptions for the purposes of this article. First, they're inherently "retrospective," meaning they report on the subject's past telephone activity rather than on future activity. Second, they aren't ordinarily available to the law enforcement agency until sometime after the activity has occurred. DNR and full audio wiretaps, on the other hand, are

"prospective," meaning they report communication occurring immediately after their installation, typically in real or near-real time.

In practice, although the legal requirements for—and information collected by—these two kinds of legal wiretaps varies widely, the same equipment can be used to implement both (audio-capture features can be disabled for DNR-only taps). Two wiretapping technologies commonly available to law enforcement agencies are *loop-extender taps* and *CALEA taps*.

Loop-extender taps

The most basic and oldest wiretap technology involves a direct electrical connection between the subject's telephone line and a second line terminating at the law enforcement agency. Such a connection (literally a "tap") can be made anywhere along the length of the local loop serving the subject, in the telephone switching office, or on the subject's premises. In principle, no special hardware is required for such interceptions at the tap point; it's sufficient simply to "splice in" a pair of wires leading back to the law enforcement agency's facilities. To ensure proper isolation and level equalization of intercepted content, current law enforcement practice for such taps uses a small device, called a loop extender or dial-up slave, at the splice point. The device sends any audio on the subject's line to the law enforcement line, re-encodes the signals, and performs level equalization. DNR equipment at the law enforcement agency decodes the dialed digit and call activity signals and, when configured for a full audio interception, also records the calls' voice contents.

To tap a line with a loop extender, a voice-grade telephone line (either a dedicated leased line or regular dial-up line), controlled by law enforcement and terminating at the law enforcement agency, is provisioned in such a way that it shares at least one cable splice point with the subject's line. (In wiretap parlance, the subject line is called the *target line* and the law enforcement line is called the *friendly line*; see Figure 1.) The target line is physically tapped and connected to the friendly line through a loop-extender (or dial-up slave) device. (The terminology isn't completely standardized, but most vendors use the term "loop extender" to refer to a device that uses a leased line for the friendly line and the term "dial-up slave" when the friendly line is a regular telephone line. For simplicity, we use the term "loop extender" to refer to either arrangement.) Any detected signals (and audio content, when authorized) are decoded and logged by the law enforcement agency equipment at the other end of the friendly line.

Because loop extenders can intercept only wireline (plain old telephone service [POTS]) telephone lines, the technology has been largely supplanted in the US by the CALEA systems described in the next section. However,

analog loop-extender systems remain on the law enforcement market, and some agencies still rely on them for some or all of their interceptions.

CALEA taps

The second, newer, wiretap technology was designed to comply with the US 1994 Communications Assistance for Law Enforcement Act (CALEA),³ which mandates a standard interface between telephone service providers (including wireline and cellular services) and agencies that perform wiretaps. In CALEA taps, the telephone company (not the law enforcement agency) decodes the signaling information and, when a full audio intercept is authorized, separates out the call audio to its own channel. The law enforcement agency connects to the telephone company through a standard interface, defined in J-STD-025A,⁴ in which the signaling information (including dialed digits, on-/off-hook status, and so on) and call audio are sent to the agency over separate channels. Whereas CALEA applies only in the US, J-STD-025A-compliant switches and interception products are marketed in other countries as well.

Each law enforcement agency conducting a J-STD-025A interception leases one or more telephone lines between the agency facilities and the target's telephone switch (see Figure 2). The first of these lines carries a *call data channel* (CDC) that reports the signaling data (call times, numbers dialed, line status, and so on) associated with all lines monitored by the agency at that switch. Additional lines to the law enforcement agency carry *call content channels* (CCCs) that contain the live audio of any active monitored lines for which a full audio interception has been authorized. The CDC might carry call data for more than one active tap, and although a single CCC can carry only one call's audio at a time, a particular CCC can carry audio for different subjects at different times, when CCCs dynamically assigned as lines become active (with the assignment reported over the CDC).

The J-STD-025A standard specifies the messages sent on the CDC as well as several different delivery formats for CDCs and CCCs. The simplest (and, we understand, most widely deployed at this time) CDC and CCC arrangement is via standard analog (POTS) telephone lines or 56-Kbps ISDN bearer channels. However, the CCC and CDC can also be delivered with IP packets over a secure virtual private network (VPN).

Once a CDC (and, when needed, one or more CCCs) has been provisioned between a switch and a law enforcement agency, installing a tap on a new line is simply a matter of configuring the CALEA delivery system at the switch to report activity on the target line. Although telephone companies are free to implement the J-STD-025A interface any way they wish, most systems don't require a special physical connection to the individual target line's local loop.

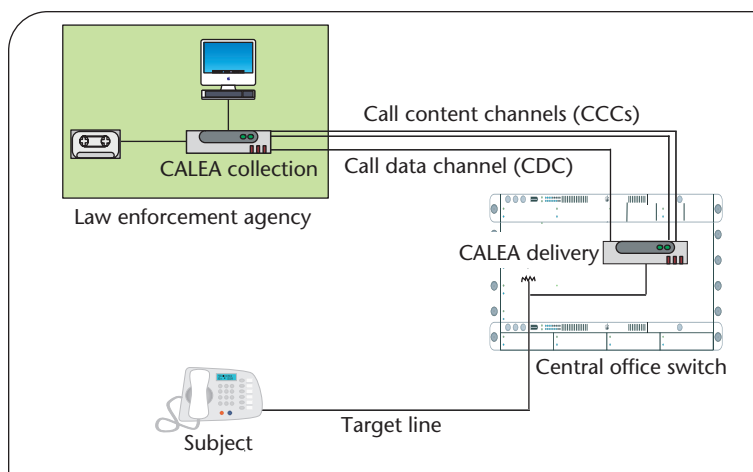


Figure 2. Communications Assistance for Law Enforcement Act (CALEA) wiretap architecture. The target line is tapped by the telephone switch itself, with signals and content relayed over separate lines.

Wiretap threat models

Perhaps surprisingly, there doesn't appear to be a widely agreed upon threat model against which law enforcement wiretap systems are measured.

Most scientific and engineering research in communication security views the eavesdropper as the adversary and is therefore concerned with guaranteeing that an abstract and powerful eavesdropping attack will fail. This perspective, while generally useful, doesn't readily apply when the threat is reversed to make the communicator the adversary.

The most mature area of existing work that examines the effectiveness of eavesdropping does so from the perspective of digital network intrusion detection. There, the focus is on thwarting eavesdropping countermeasures such as evasion and insertion.⁵⁻⁷ However, unlike network intrusion detection systems, telephone wiretaps aim to capture data about all communication, both normal and anomalous.

Detection

Perhaps the most prominently considered threat against eavesdropping systems is detection. Surreptitious interceptions are thought to produce better intelligence than those that aren't. A vibrant, if occasionally somewhat disreputable, technical surveillance and countermeasures industry fuels both sides of the bugging and bug-detecting arms race.

Wiretap systems that depend on direct metallic connection to the local loop (such as loop-extender systems) are potentially susceptible to detection by a range of means. A tapping device installed at or near a subject's premises might be noticed in a physical inspection. Depending on the circuitry used, taps that change the line's

transmission characteristics can sometimes be discovered electronically—for example, through sensitive loss measurements or time-domain reflectometry. Taps might also be exposed by penetrating a telephone company's information systems or facilities, such as via rogue insider access, computer compromise, or physical burglary. Loop extenders marketed to law enforcement agencies usually have a relatively innocuous physical appearance and high impedance circuitry, but depending on how and where they're installed, an expert might still be able to detect their presence.

J-STD-025A requires that switch-based CALEA monitoring be undetectable to the subject and that the computing systems provisioning and maintaining any interceptions be adequately secured. However, the standard doesn't prescribe specific security mechanisms, nor does it require special protection or authentication for CCC and CDC traffic or links.

Encryption and content obfuscation

The most well understood countermeasures against eavesdropping involve the use of cryptographic techniques, and modern cryptosystems are thought to provide very good end-to-end security when implemented properly. However, voice encryption isn't widely used by wiretap subjects. Furthermore, digital voice encryption systems for analog telephones aren't yet readily available on the commercial market, and require the participation of both parties to be effective. End-to-end encryption protects only the content, not the dialed numbers or other signaling (because the signal's endpoint is the phone network itself, not the called party).

Denial of service against CALEA CCCs

The law enforcement community and the telecommunications industry have already discovered at least one practical countermeasure against J-STD-025A CALEA call content collection. This countermeasure prevents the collection of subject call content on systems with dynamically assigned CCCs and exploits the fact that the number of different voice channels associated with a monitored line is potentially unbounded if the subject subscribes to a call-forwarding service. Essentially, the target and its correspondents "flood" a monitored line with unrelated calls that are forwarded elsewhere. (The number of forwarded calls is bounded only by the switch's call-forwarding limits.) Each additional call is assigned its own CCC, eventually leaving no CCCs open for monitoring significant calls.

Although CALEA was partly motivated by new services such as call forwarding, this countermeasure apparently wasn't considered in developing the original CALEA interfaces. The problem was first publicly suggested in recent patent disclosures^{8,9} for systems that allow the law enforcement agency to disconnect super-

fluous CCCs. (This capability isn't addressed in the J-STD-025A standard.) The published literature says little about whether wiretap subjects have actually employed CCC flooding countermeasures or whether currently fielded CALEA systems incorporate the defenses described in the patents.

CALEA systems might also fail if the telephone company provisions the tap to monitor the wrong target line. A recent report from the US Department of Justice (DOJ) found instances of recorded traffic from FISA taps that were later discovered to have originated from incorrect sources.¹⁰

Evasion, confusion, and the eavesdropper's dilemma

In a separate work,¹¹ we formalized the concepts of evasion and confusion as eavesdropping countermeasures and identified the "eavesdropper's dilemma" as a fundamental trade-off in certain interception architectures. Briefly, evasion occurs when a target can prevent legitimate traffic from reaching the interception system, and, conversely, confusion occurs when spurious traffic can be directed at it. If a system is susceptible to either countermeasure, the intercepted traffic's fidelity can be arbitrarily degraded, either by the target or, in some cases, by a third party. The architecture of many eavesdropping systems allows defense against evasion or confusion only at the expense of increased exposure to the other, hence the eavesdropper's dilemma.

An interception system is subject to the eavesdropper's dilemma whenever it has incomplete knowledge of how the network and receiver process traffic, or if it destroys information processed at low layers of the protocol stack.

Although we introduced confusion, evasion, and the eavesdropper's dilemma in the context of digital network interception, they can be readily applied against the analog law enforcement wiretap systems described in this article.

Signaling countermeasures against loop-extender taps

Loop-extender taps rely heavily on in-band signaling, with an architecture that makes them especially vulnerable to manipulation by the target. We found three kinds of practical countermeasures against systems that use these taps: the first masks the dialed digits of outgoing calls, the second obscures incoming caller ID signals, and the third (and perhaps most serious) disables audio monitoring and recording by the agency.

Dialed digit spoofing

A fundamental weakness in the loop-extender tap model arises from the way the tapping equipment decodes dialed digits and other audio signals. Although telephone number signals represent "digital" information, they're transmitted on telephone lines in analog form. The most

common dialing system uses audio dual-tone multifrequency (DTMF)^{12,13} signals. DTMF is also popularly known by its original AT&T trademark, TouchTone. Analog DTMF signals are decoded and converted to digital form at the telephone company switch.

DTMF digit signals consist of two audio frequency tones: a “low” tone corresponding to the horizontal “row” position of the digit on the keypad and a “high” tone corresponding to the “column” position. Although the familiar consumer telephone DTMF keypad has 12 digits (0 through 9, *, and #) arranged in four rows and three columns, the DTMF standard specifies a fourth column, giving four additional tone signals, usually called A, B, C, and D. The C button, for example, is conceptually located to the right of the 7, 8, and 9 buttons (see Figure 3; these fourth-column tones will be important to us later.)

Although most telephone instruments produce tone signals well within the “standard” acceptable range of properly operating DTMF decoders, tone signals at the edges of the standard parameters will be accepted by some decoders but not others. Many parameters affect whether a given decoder will recognize a given tone signal as a valid dialed digit, including the precise frequencies of the two tone components, their overall power level, the relative amplitude of the tones, signal duration, waveform distortion, external noise, and so on.

Two observations follow directly from the analog nature of DTMF signals and their decoding by telephone switches:

- Because DTMF transmission and decoding are analog processes, for any given parameter (frequency, amplitude, and so on), no two DTMF decoders will use precisely the same threshold to determine whether a given signal is accepted or rejected (the analog eavesdropper’s dilemma).
- By completing or not completing dialed calls, a telephone switch is an oracle for determining whether DTMF signals sent to it have parameters within its tolerances.

We found that by systematic DTMF dialing with selectively degraded parameters, analog telephone subscribers can discover the thresholds of their switches’ DTMF decoders efficiently and with sufficient accuracy and precision to construct signals that other decoders would likely treat differently.

In a loop-extender wiretap system, two different DTMF decoders process each dialed digit on the target line independently: one at the law enforcement agency (to determine the dialed number that is logged) and another at the telephone company switch (to determine the number the telephone network uses for actual call processing). This means that each tone parameter of each

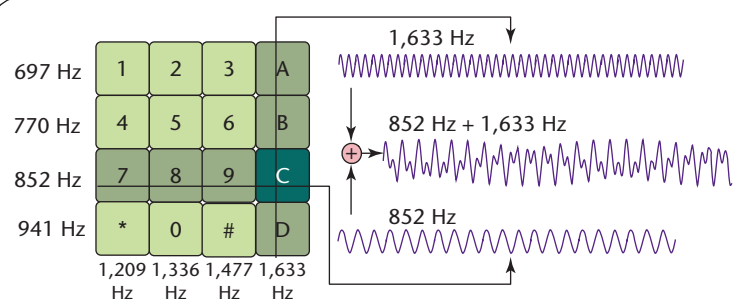


Figure 3. Dual-tone multifrequency (DTMF) keypad and waveforms of generated tone. Each row and column of the keypad generates a tone signal; when a key is pressed, the row and column tone are combined and sent on the line.

digit either has tone encodings that are accepted as valid by the switch but not the wiretapper, or encodings that are accepted as valid by the wiretapper but not the switch. Note that this property is inherent to any interception system in which a separate DTMF decoder is used at the tap; it doesn’t depend on the failure of any equipment to operate within standard specifications.

In our experiments, a simple, automated binary search (involving roughly 30 to 120 minutes of unattended experimental dialing and analysis with a laptop computer) could discover the precise threshold characteristics of a given telephone switch’s DTMF decoder with sufficient accuracy to distinguish it from other decoders. Spurious digit encodings can then be constructed that are just outside the accepted parameters of the switch’s decoder (such that they would have no effect on the actual number dialed), but that will be accepted with high probability by an external law enforcement tap. These signals thus attempt to *confuse* a tap. Conversely, non-standard digit encodings can be constructed that are just within the parameter range accepted by the switch, but that will be ignored with high probability by an external law enforcement tap. These signals thus attempt to *evade* a tap.

This probe discovers nothing about the limits of the wiretap’s DTMF decoder (or even if a wiretap is present), but it doesn’t need to. Because of the analog eavesdropper’s dilemma, and as our experiments confirm, a wiretap’s decoder will always be either more liberal or more conservative than the switch in handling signals at the edges of acceptance. When a wiretap is conservative, digit signals accepted by the switch evade detection by the tap. When a wiretap is liberal, it accepts extraneous confusion digit signals ignored by the switch.

Complete telephone numbers can be dialed with a combination of evasion digits and confusion noise, such that the complete, correct number will be received as intended at the switch, but where some or all of the digits

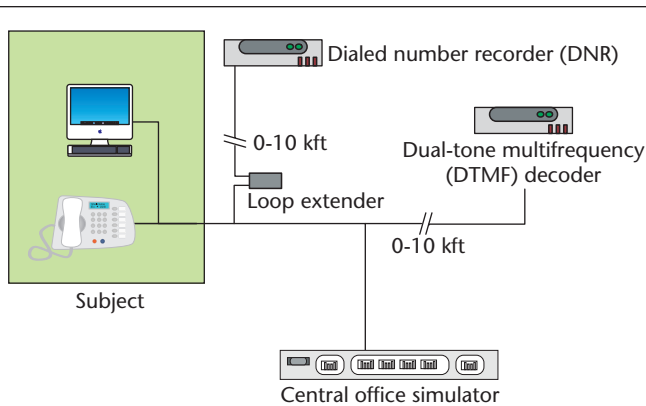


Figure 4. Experimental wiretapping configuration. A simulated central office with two telephones, simulated transmission losses, and connected to various wiretapping systems.

recorded by the tap are incorrect. (Depending on the exact hardware at the switch and the tap, the tap will predominantly tend to either ignore evasion digits or accept confusion digits. The subject can't predict whether confusion or evasion dominates without access to the tap hardware, but the tap logs incorrect numbers in either case.) The software to perform the probing and the dialing can be quite simple, requiring only modest computer, sound card, and modem hardware.

A wiretap can never reconstruct evaded digit signals that it ignores by being too conservative. The obvious defense against evasion is thus to use a more liberal decoder at the tap. Unfortunately for the wiretapper, an interception system that is too liberal becomes susceptible to confusion. Although an overly liberal decoder might seem to be a "lesser evil" than one that is overly conservative (because fewer potential digits are discarded), confusion can achieve perfect secrecy under ideal conditions.

Confusion and evasion dialing are effective against any eavesdropping system that performs its own DTMF decoding and commits to a single interpretation of the digits. In a full audio tap, it might be possible to later conduct offline forensic analyses of recorded tone signals and reconstruct evaded or confused digits. However, such an analysis would depend on the precise tolerances of the local central office (CO) and the transmission characteristics between the tap and the subject at the time of the call. Deriving these parameters requires active probing from the tap, similar to that performed by the subject. Current wiretapping technology has no mechanism for performing such probes, and so precludes such analyses.

Of course, a subject can't be sure whether an eavesdropper's DTMF decoder will be more conservative or liberal than that of the switch for any given parameter, and so can't be sure whether confusion or evasion will be more successful at masking dialed numbers. There's no need to

choose between the two, however; confusion and evasion dialing can be used in concert across an entire string of dialed digits. Here, the subject intersperses n random noise digits in random positions among the l "true" digits to be dialed. The true digits are signaled using evasion, and the noise digits are sent using confusion.

This combined dialing technique is effective to the extent that either confusion or evasion dialing succeeds often enough to mask the actual dialed number. Our experiments with standard law enforcement loop-extender tap hardware as well as with the taps we constructed based on laboratory and diagnostic DTMF signal analyzers suggest that these countermeasures are both practical and reliably effective at hiding the true dialed number.

Dialing experiments

To test the practical effectiveness of confusion and evasion dialing, we conducted experiments under a variety of simulated network conditions and with a range of tapping hardware. Our results support our analytical hypothesis that confusion and evasion are effective at preventing DTMF digit recording in taps (such as loop extenders) that rely on their own DTMF decoders.

Figure 4 depicts our experimental setup. A CO switch simulator interprets DTMF tones from the "target" line (the simulated subject's phone line), generates call progress, ringing, and caller ID signals, and switches calls. The subject line is equipped with both a normal telephone and a Pentium laptop computer with a modem, sound card, and telephone audio interface. The laptop modem is used to seize the line, while the sound card generates confusion and evasion dialing signals. The telephone set is used for the actual voice communication.

To account for different telephone line and CO conditions, we ran our experiments with several simulated and actual telephone switches, including Teltone TLS-5C and Ameritec AM-7 simulators and the Western Electric/Lucent 5-ESS switch that serves one of the authors' homes. (The Teltone and Ameritec devices simulate most aspects of a telephone switch, decoding and generating all loop signals and providing voice paths; we simulated the loss and distortion of the local loop with Telebyte model 453 variable-length 26-gauge cable simulators.) Results were similar under all setups; the data shown here used the Teltone simulator.

The primary tapping device in our tests was a Recall Technologies model NGNR-2000 (a current law enforcement loop-extender DNR and audio collection system). We also constructed our own taps using various laboratory and diagnostic DTMF decoders. In all setups, every tap was able to correctly reconstruct all dialed digits when no countermeasures were employed.

The first step in confusion and evasion dialing is determining the switch's limits. Although many parameters are defined for proper DTMF tones,^{13,14} we concentrated on

six: the amplitude of the high-tone component (recall that a DTMF signal consists of a high-frequency tone mixed with a low-frequency tone), the low tone's amplitude, and the positive and negative frequency skews of the high- and low-frequency tone components. Our probe software used a simple binary search for each parameter's limits. In each trial, it dialed a nonexistent number (say, 555-010X), with the parameter under test (say, the high-tone frequency) varied on the last digit. If the switch recognized the entire dialed number (as indicated by a call completion signal), then the parameter was within the switch's tolerance. However, if the switch didn't attempt to route the call (as indicated by a lack of response), then the parameter was outside the switch's tolerance. The binary search narrows a pair of differentiating values for each parameter and each digit. Signals with the parameter value within the switch's tolerance can be used for evasion, whereas tones outside the tolerance can be used for confusion. We performed this search for all six parameters and for each DTMF digit (0 through 9). As we mentioned earlier, the entire automated probing process took approximately 30 to 120 minutes to complete; in practice, a subject would repeat this search as needed to compensate for drift in the switch decoder hardware and changes in the transmission characteristics between the subscriber and the CO.

After probing the switch, we tested evasion and confusion dialing. Table 1 shows the effects of evasion dialing alone. Here, all digits were transmitted with evasion parameters, causing the switch to barely recognize the tones. As the table shows, although the switch received the full number and completed the call, only one of the wiretaps correctly captured all dialed digits.

In this case, evasion was moderately successful; it prevented most of our taps from capturing at least one digit of the dialed number. Adding confusion, however, makes the eavesdropper's job even harder. In our next experiment, we inserted confusion signals uniformly at random among the evasion-dialed digits. For readability, we show here an experimental run in which the target sends 20 confusion digits along with the 11-digit number; more confusion digits would ordinarily be sent in actual countermeasure use.

Table 2 shows the results of using both confusion and evasion. As before, the switch correctly processed the desired number. All tested DTMF decoders were susceptible to both confusion and evasion, and not only failed to record some evaded digits, but also accepted some confusion signals. Confusion was particularly effective against the Ameritec AM8a, which recorded a total of 22 dialed digits. The number of possible 11-digit reconstructions of this string is approximately

$$\binom{22}{11} = 705,432,$$

Table 1. Dual-tone multifrequency (DTMF) evasion of dialed number 1-987-654-3210.

DEVICE	INTERPRETATION IN THE PRESENCE OF EVASION
Recall model NGNR-2000 (DNR device)	18753210
Ameritec model AM8a	1976541
DSchmidt model DTMFLCD-2	1976543210
Harris model 25D	19876543210
Metro-Tel model TPM32MF	1976541
Metro-Tel model VNA70A	19765421

Table 2. DTMF evasion and confusion of dialed number 1-987-654-3210.*

DEVICE	INTERPRETATION IN THE PRESENCE OF EVASION AND CONFUSION
Recall model NGNR-2000 (DNR device)	149876465642392120
Ameritec model AM8a	1346676649919555432610
DSchmidt model DTMFLCD-2	1497645432120
Harris model 25D	139876419556432610
Metro-Tel model TPM32MF	1476411543210
Metro-Tel model VNA70A	14876411543210

* The dialed number is **17349668766**49916955**564239261200** (bold digits were sent using evasion parameters and the other digits were sent using confusion).

and we note that not a single one of these interpretations is the actual dialed number because some of the digits were evaded.

Incoming calling-number ID spoofing

Calling-number ID (CNID, sometimes referred to as caller ID) is an optional feature offered by local exchange carriers that allows a subscriber to screen incoming calls. If CNID service is enabled on the called party's line, the CO transmits the caller's telephone number and, if available, the name associated with that account. The CNID information is relayed using in-band signaling between the first and second ring signals. Special devices display the incoming calling number to the subscriber.

When CNID service is active on a target's line, any wiretap device can also decode and record the source of incoming calls. Note that if CNID isn't present, loop-extender taps must learn the caller's number through some other mechanism (such as billing records).

Because the CO transmits the CNID data, evasion isn't possible. However, a subject can trivially confuse the capture of the CNID transmission by injecting counterfeited signals on the line.

Table 3. Using confusion to forge calling-number ID.

DEVICE	NO CONFUSION (BASELINE)	WITH CONFUSION
Recall (model NGNR-2000) DNR device	(987) 654-3210 Tony Soprano	(215) 898-5000 Matt Blaze
RadioShack trim phone with caller ID (model 43-3909A)	987-654-3210 TONY SOPRANO	215-898-5000 MATT BLAZE
Tempo (model ID'R Plus)	(987) 654-3210 Tony Soprano	(215) 898-5000 Matt Blaze
US West Call Waiting ID (model CI-98)	987-654-3210 TONY SOPRANO	215-898-5000 MATT BLAZE

To confuse a wiretapper's recovery of incoming CNID, we simply replayed periodically (through the sound card) a forged CNID audio signal while the target's phone was on-hook. (We used a single static number for this purpose but a more sophisticated subject could generate new signals each time.) In every case in our experiments, the wiretap decoded the forged signal instead of the legitimate CNID transmission when incoming calls were received (see Table 3).

Line status spoofing and recording suppression

In-band signaling makes full audio loop-extender taps especially vulnerable to countermeasures. A subject can remotely disable any audio recording equipment (and cause the system to log line inactivity) for arbitrary periods during a call by spoofing the on-hook signal generated by a loop extender.

In loop-extender systems, all signaling data and audio are sent to the law enforcement agency over a single channel (the friendly line), entirely in the analog voice-band domain. Any call progress and line status signaling data from the target that the law enforcement system is to process or record must therefore be sent over the same channel that carries the target audio (in-band). This, as we shall see, is a rather fragile arrangement rich in potential for exploitation by the target.

An intercept-collection system must record several kinds of call-processing signals from the target's telephone line. Some of these signals (including DTMF-encoded dialed digits, incoming calling-number ID, and audible call progress signals such as dial tones and busy and audible ringing signals) are already encoded in the voice-band audio domain as part of the standard telephone interface. These signals simply pass through the dial-up slave's normal audio interface and can be decoded entirely by the law enforcement hardware (although, as noted earlier, not always correctly). Other telephone signals, however—most notably the on-/off-hook state, rotary "pulse" dialed digits, and incoming-call ringing signals—aren't encoded in the audio domain on the tar-

get line, depending instead on the (DC) voltage and current on the wire between the CO switch and the target's telephone instrument. Although the dial-up slave unit can detect these signals relatively easily by observing the line voltage, they can't simply be relayed back over the friendly line in the same form for processing by the agency (that line's voltage and current maintain the connection through the telephone network between the loop extender and the agency itself). These signals are therefore encoded as special audio tones superimposed on the friendly line audio, and recognized and decoded as such by the law enforcement equipment.

The most important signal that isn't already in the voice-band audio domain (and from whose state many of the other DC signals can be derived) is the on-/off-hook status. Ordinarily, when the line is on-hook, no target audio or other signals would be present on the line. This gives rise to a simple (perhaps too simple) audio encoding of line status: an idle tone is sent continuously on the friendly line whenever the target line is detected in the on-hook state and removed when the line goes to the off-hook state. Most loop-extender (and dial-up slave) systems marketed to law enforcement use this scheme.

In fact, not only do virtually all loop extenders indicate line status with an idle tone, they almost all use the same de facto standard idle-tone signal: the DTMF C digit (a two-frequency audio signal consisting of 852 Hz and 1,633 Hz; some literature mentions the use of the A tone for this purpose, but all current vendors of which we are aware use C). This is the only audio signal added to the friendly line by some dial-up slave and loop-extender products. (Other models also provide additional signals to indicate incoming ringing and periodic "keep alive" off-hook status signals, usually using fourth-column DTMF tones.) Because the loop extender sends the on-/off-hook status signal over the same channel that carries the target audio, legitimate indications of changes in target line status can't be distinguished from an identical-sounding signal generated by the target while a call is in progress.

New-call spoofing. Unexpected use of the idle signal can trigger bugs in some loop-extender equipment. At least one system that we tested (a Recall Technologies NGNR-2000) would become distressed if a DTMF C idle signal wasn't immediately followed by new call setup signals; this apparently caused the device to conclude that it lost the connection with the slave unit, disconnect the friendly line, and initiate a new connection. Under best-case conditions, it required more than 30 seconds to reestablish the connection to the loop extender or dial-up slave. It was very easy to exploit this vulnerability; we simply sent the DTMF C signal for 3 seconds. This would cause the collection function to stop recording audio for the 30 to 45 seconds required to establish the connection. Sending the DTMF C tone on the target line for 3 seconds every 30 seconds would allow no audio to be recorded by the wiretap with this hardware. We can't speculate on the performance of other DNR and recording devices in this regard without testing them, of course.

Even loop-extender systems that correctly process idle signals can be manipulated. The target (or the target's correspondent) can introduce a false new-call record by sending the 852-Hz + 1,633-Hz DTMF C signal on the line during a call for long enough for the wiretap to detect an on-hook condition and register (incorrectly) that the current call has ended. At this point, the target can send additional DTMF and audible ringing signals to simulate a new call being placed (presumably to a different number), all the while maintaining the connection with the original correspondent.

Of course, sending a brief burst of C tone doesn't by itself prevent the capture of the call content (except on buggy interception equipment such as the Recall), but it does allow the target to introduce spurious call records into the wiretap logs and to associate captured call content with false telephone numbers.

Recording suppression through C tone spoofing. The use of the in-band C tone idle signal has even more serious consequences for full audio wiretaps: the subject can suppress content recording for arbitrary periods.

Loop-extender systems turn off audio recording when a C tone signal is detected. Naturally, subjects can't easily converse while a spoofed C tone is sent at full volume over the target line (unless they employ special narrow-band filters to eliminate it). However, there is no need for the subject (or the correspondent) to send the tone at full volume.

We found that even under conditions very unfavorable to the target (in which the law enforcement equipment was attenuated by 10,000 more feet of 26-gauge cable than the total connection length between the target and the correspondent), it was possible to falsely indicate an on-hook condition and turn off the recording equipment with a continuous C tone sent at very low ampli-

tude. False signals of as little as -40 dBm total power on the target line were sufficient for this purpose. We found it readily possible to carry on an intelligible, even comfortable, conversation over this tone, with the audio completely evading wiretap collection because the recording was turned off and muted.

Note that the vulnerability to this countermeasure is a fundamental property of the in-band signaling architecture used between the loop extender and the interception recording system. It could be prevented only by the loop extender filtering out DTMF C tone signals from the target audio stream sent over the friendly line; we aren't aware of any loop extenders or dial-up slave products that perform such filtering, however.

You can find an MP3 audio example of a conversation evaded with C tone at www.crypto.com/papers/wiretapping/.

Signaling countermeasures against CALEA taps

At first blush, the J-STD-025A CALEA interfaces seem to effectively neutralize in-band signaling countermeasures; separate channels deliver the target's signaling (the CDC) and voice traffic (the CCC), and allow decoding of DTMF tones at the switch instead of at a second unit at the law enforcement agency. Because the telephone company is responsible for DTMF decoding before sending the data to the agency, it's likely that the reported digits are derived directly from the switch's call-processing system, and because the line status is reported over a separate signaling channel, such systems need not be vulnerable to in-band spoofing of the line status. Nevertheless, many CALEA implementations fall short of achieving the level of robustness that their architecture would appear to allow.

Many CALEA configurations may indeed be more reliable than traditional loop-extender systems with regard to susceptibility to confusion and evasion dialing. However, we note that CALEA and J-STD-025A specify only a standard interface between the telephone company and law enforcement; they don't require or assume

We found it readily possible to carry on an intelligible conversation over this tone, with the audio completely evading wiretap collection.

any particular implementation of these interfaces, and they don't require effectiveness or performance beyond that which pre-CALEA systems achieved. Therefore, although most CALEA-compliant telephone switches

presumably report the actual digits recorded by the DTMF decoder that processes the calls, there is no explicit requirement that they do so, and thus there is no guarantee that the dialed numbers reported to law en-

The problem of in-band signal abuse in particular has a long history in communication security.

forcement accurately reflect those processed by the switch. Moreover, “post-cut-through” digits reported on the CDC—those processed not by the switch but by a remote endpoint (such as a voicemail system)—can still be confused or evaded.

A more serious potential vulnerability in CALEA implementations is recording suppression via an in-band “continuity tone” signal that some collection system implementations recognize on the CCC. The processing of this signal renders such systems vulnerable to the same content evasion countermeasures that work against loop-extender systems.

Recording suppression in CALEA implementations

Although the J-STD-025A standard appears to eliminate the possibility of in-band signaling countermeasures by providing the law enforcement agency with call content and signaling in separate delivery channels, actual implementations sometimes blur this distinction. In particular, some CALEA implementations use the DTMF C tone signal to indicate that a CCC is in the idle state. Recall that this is the same signal used to indicate line status in loop-extender systems. The C tone is processed by some CALEA CCC collection systems in much the same way—as a signal to disable the recording equipment.

This mechanism might have been motivated by a desire for backward compatibility with loop-extender collection systems. Law enforcement agencies and telephone companies can construct any of a variety of CALEA collection system architectures. The CDC and CCC can be delivered to the agencies directly over separate telephone lines or an IP VPN, or they could employ the services of an intermediate CALEA service provider such as Pen-Link or Xlence. Some CALEA collection systems are designed to accept CDC and CCC channels directly, whereas others adapt “legacy” loop-extender recording systems. The US Federal Bureau of Investigation (FBI) and the DOJ explicitly requested that the continuity tone on idle CCC channels be a required CALEA feature in a list of proposed improvements to the original J-STD-025A specification.^{15,16} Although the US Federal Communications Commission (FCC) didn’t ultimately

adopt the continuity tone as a requirement,¹⁶ it has become a common optional feature cited in CALEA vendor literature and system patents.^{8,17–19} In particular, the C tone on the CCC is often specified in product literature as a mechanism to control the collection system’s audio recording equipment—that is, when the C tone is present, the CCC is assumed to be idle (regardless of the call status as reported on the CDC) and the collection system can automatically mute audio monitoring and stop the recording equipment. The C tone is also used internally by some switch-side CALEA delivery systems.

Just as with loop-extender systems, these configurations (sometimes called C tone supervision) make it possible for a subject to unilaterally disable content recording by sending a continuous C tone at an amplitude sufficiently high to trigger the recording-suppression mechanism but low enough to allow intelligible conversation. Because the same tone is also used to suppress recording in loop-extender systems, the target need not know whether CALEA or loop-extender taps are used by the agencies he or she wishes to evade.

Not all CALEA implementations support C tone supervision signals; it’s an optional feature not required by the standard. However, the C tone appears to be a relatively commonly available option among current CALEA systems that use analog or ISDN bearer channels for CCC delivery and in CALEA products designed for output to legacy collection equipment.

Discussion

The signaling protocol failures described in this article are of significance to a broad range of communities for several reasons. First, of course, is the immediate problem of conducting reliable lawful interception without evasion or manipulation by the subject. In fact, this is a subtler problem than it might first seem to be. Someone who believes that he or she is being wiretapped can reliably evade interception simply by refraining from using the suspected telephone line for incriminating conversations or by using a voice encryption system for such calls. However, history suggests that “telephone silence” isn’t a satisfying solution for many wiretap subjects because many criminal enterprises apparently rely extensively on telephone communication. Neither does end-to-end encryption provide widespread practical cover for many law enforcement targets. Encryption requires advance planning for the use of special hardware by both peers, and reliable voice telephone encryption systems still aren’t widely available on the commercial market. The signaling countermeasures in this article, on the other hand, not only require less sophisticated equipment than encryption (and only unilateral action), but also can be used to actively mislead an investigator with incorrect or incomplete interception records.

Whether signaling countermeasures are attractive to or likely to be employed by subjects depends on how they perceive the threat. Because wiretaps are usually secret, a subject can never be sure that he or she is actually being monitored or whether monitoring is conducted with a susceptible system. However, because of the relative lack of tapping-system diversity, a target could make an educated guess as to whether agencies could be expected to conduct an investigation are using vulnerable equipment. Federal and local agency procurement contracts for wiretap equipment are often publicly available. With a simple Internet search, we were able to discover vendors and even model numbers of the equipment that several agencies in various jurisdictions use.

There is evidence (albeit indirect and inconclusive) suggesting that sophisticated targets might sometimes employ signaling countermeasures. For example, law enforcement agents have noted (in trial testimony) unexplained audio gaps in wiretap recordings, with specific reference to C tone signals.²⁰

More broadly, the existence of signaling countermeasures suggests that the wiretap technology used by law enforcement should be critically evaluated against a wider range of threats than perhaps it has been. The scope of our analysis was deliberately restricted to information and materials we obtained from public sources and was limited to the narrow problem of confusion and evasion countermeasures. We made no attempt to be exhaustive or comprehensive, yet we quickly discovered practical attacks that seem rather obvious in hindsight. Some of the potential vulnerabilities in modern CALEA systems arise directly from features (for example, the CCC continuity tone) requested by the law enforcement community itself. It seems at least plausible that as-yet undiscovered weaknesses exist in the J-STD-025A specification (and the systems that implement and support it). A systematic effort to discover or rule out vulnerabilities would improve confidence in these systems.

Finally, the protocol failures and signaling weaknesses in voice wiretaps provide an overall case study in computer and communication security. Well-established principles of secure system design appear to have been violated in the loop-extender and CALEA tap architectures, with interfaces subject to multiple interpretations and complex interacting features and options that strain to maintain backward compatibility with “legacy” systems.

The problem of in-band signal abuse in particular has a long history in communication security, most famously exposed in the US and international long-distance telephone network of the 1960s and 1970s (see the sidebar). And yet, vulnerable systems that depend on unprotected in-band signals for critical functions continue to be de-

signed and fielded, suggesting that many security practitioners and system designers don’t adequately understand or appreciate the risks inherent in such designs.

There is, unfortunately, little room to make conventional analog loop-extender interception systems more robust against these countermeasures within their design constraints. Audio recording of dialing signals provide limited opportunities for subsequent forensic analysis of confused or evaded dialed digits (assuming the characteristics of the CO and transmission line can be accurately estimated later), although legal constraints generally preclude agencies from making such recordings on DNR-only taps.

Vulnerable CALEA systems, on the other hand, may be able to be made more robust against specific countermeasures with relatively minor configuration changes. In particular, the law enforcement equipment that processes the CCC should be configured not to shut off when a C tone is present on the channel. Instead, such systems should rely only on the CDC to determine when recording should commence or stop. Agencies should confirm their CALEA equipment’s behavior with their vendors.

Additionally, wiretap evidence, whether collected by loop-extender or CALEA systems, should be evaluated by investigators for signs of signaling countermeasures. In particular, records of dialed numbers and call times should be examined for discrepancies against telephone company call detail records. This reconciliation should be performed routinely and as soon as possible after the records become available.

We also strongly urge that J-STD-025A and other interception standards and practices undergo a critical security review against countermeasures such as those discussed here and, more generally, against a broader threat model. The relatively simple signaling countermeasures in this article became quickly apparent even from our somewhat cursory analysis. It appears that a systematic search for vulnerabilities under a threat model that includes subject-initiated countermeasures wasn’t part of the development process for either the J-STD-025A standard or many of the systems that implement it.

As wiretap systems become more homogeneous and standardized, the consequences of vulnerabilities become increasingly serious. Any weaknesses in J-STD-025A systems could have the unintended and somewhat ironic consequence of degrading law enforcement’s ability to conduct wiretaps on the advanced digital and mobile systems that CALEA envisioned. J-STD-025A standardizes the delivery of intercepted content to law enforcement across many different communications services, thus any countermeasures against these systems threaten law enforcement’s access to the entire spectrum

In-band signal abuse in the long-distance telephone network

In the late 1940s, the AT&T long-distance telephone network added features for “direct distance dialing” of long-distance subscriber calls without the need for manual operator assistance. These features required new protocols and mechanisms to support the automated signaling of trunk status and transmission of telephone numbers and other routing information between switching centers in different cities.

Although local subscriber dialing had been available for many years, the signals and protocols used in the local loop weren’t directly applicable to the long-distance trunk circuits used between switching centers. Subscriber signaling used (and still uses) a DC current loop, which, for various reasons, is inappropriate for circuits more than a few miles long or for wide-band interoffice trunks. The new long-distance system instead used audio signals in the voice band to signal line status and dialed numbers. A 2,600-Hz “idle” signal was placed on trunks when they were inactive; other tone signals (similar to dual-tone multifrequency [DTMF], but using different frequencies) were associated with individual number digits. To route a long-distance call, a switch would select an idle trunk to the next switch in the path, remove the idle tone, transmit the desired number, and connect the calling subscriber’s local loop to the trunk. The remote switch would then route the call to the next

switching center in the same way, until it finally reached the destination number’s local switch. Billing records for long-distance calls were maintained at the originating subscriber’s switch.

This arrangement had the significant advantage of allowing individual subscribers to use their existing equipment to perform long-distance dialing, because the new trunk signals were intended to be encoded and decoded by the internal network-switching equipment, not by end-user telephone instruments. The signals, although in the audible voice-band frequency range (hence “in-band”), were largely transparent to the end user. By the late 1950s, the system was fully deployed in most of the US.

Within just a few years, technically inclined telephone users had discovered ways to exploit the system to make fraudulent long-distance calls. The fraud technique used a specially constructed signal generator, which became popularly known as a *blue box*, to spoof the in-band long-distance signaling tones. A race condition in the way idle trunks were handled allowed an end user to briefly send the 2,600-Hz idle signal during a long-distance call, which would be misinterpreted by the remote trunk as signaling the end of the current call. The remote trunk then disconnected the call in progress and became ready to accept tone signals for a supposedly new call, which could be similarly spoofed by the caller. The

of intercepted communications. We suggest that the law enforcement community develop and articulate security and assurance requirements for interception systems against which existing and future standards and technologies will be measured. □

Acknowledgments

We are grateful for insights from Steve Bellovin, Bill Cheswick, Will Enck, Susan Landau, Patrick McDaniel, and Fernando Pereira. The Trustworthy Network Eavesdropping and Countermeasures project at the University of Pennsylvania is funded by the US National Science Foundation Cyber-Trust program under contract NSF-0524047.

References

1. *Federal Wiretap Act*, Title III, United States of America, 1968 (codified as amended in 18 US Code Section 2510–2522).
2. *Foreign Intelligence Surveillance Act of 1978*, Pub. L. No. 95–511, 92 Stat., United States of America, 1978 (codified as amended in 50 US Code Section 1801–1811, 1821–1829, 1841–1846, 1861–62).
3. *Communications Assistance for Law Enforcement Act*, Pub. L. No. 103–414, 108 Stat. 4279, United States of America, 1994 (codified as amended in 18 US Code and 47 US Code Section 229, 1001–1010, 1021).
4. *Lawfully Authorized Electronic Surveillance*, J-STD-025A, Am. Nat’l Standards Inst., 2003.
5. T. Ptacek and T. Newsham, *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection*, tech. report, Secure Networks, 1998.
6. U. Shankar and V. Paxson, “Active Mapping: Resisting NIDS Evasion without Altering Traffic,” *Proc. 2003 IEEE Symp. Security and Privacy*, IEEE CS Press, 2003, pp. 44–61.
7. V. Paxson, “Bro: A System for Detecting Network Intruders in Real-Time,” *Computer Networks*, vol. 31, nos. 23 and 24, 1999, pp. 2435–2463.
8. E. Kampmeier, D. Smith, and M. Smith, *Utilization of Communication Channels between a Central Office Switch and a Law Enforcement Agency*, US patent 6,728,338, Patent and Trademark Office, Nov. 2000 (issued Apr. 2004).
9. L. Prieur, *Automatic Monitoring Service for Telecommunications Networks*, US patent 6,470,075, Patent and Trademark Office, June 1999 (issued Oct. 2002).
10. Office of the Inspector General Audit Division, *Federal Bureau of Investigation’s Foreign Language Translation Program Follow-Up*, audit report 05–33, US Dept. of Justice, July 2005.
11. E. Cronin, M. Sherr, and M. Blaze, *The Eavesdropper’s Dilemma*, tech. report MS-CIS–05–24, Univ. of Penn-

accounting system at the caller's switch would continue to record the original call as if it were still in progress. It was therefore possible for subscribers to defeat the billing system by starting the process with a call to an inexpensive or toll-free long-distance number and spoofing the idle signal and routing tones for what would have otherwise been a more expensive call.

By 1971, the in-band signaling flaws came to national attention,^{1,2} bringing about an arms race of sorts between blue-box users (so-called phone phreaks) and telephone companies seeking to protect their billing revenue (and prevent unauthorized access to the internal signaling network in general). AT&T eventually developed an out-of-band "common channel" long-distance signaling architecture³ that defeated the blue box by eliminating in-band interoffice signaling. The new system wasn't fully deployed in the US until the mid 1980s, however.

The vulnerabilities in loop-extender and CALEA wiretap systems—and the methods for exploiting them—are strikingly reminiscent of the weaknesses of the 1960s telephone network. It's notable that even though CALEA primarily uses out-of-band signals, the C tone idle signal mechanism (when present) remains in-band and vulnerable to exploitation by the target.

Unfortunately, although the weaknesses themselves might be similar, the telephone industry's response to the problem isn't as directly applicable as might be hoped in forming a response to wiretap countermeasures. A contemporary article⁴ suggested a three-pronged approach to mitigating the effects of the blue box:

the first was user education on the ethical pitfalls of telephone fraud, the second was vigorous detection and prosecution of those committing fraud, and the third was migration to out-of-band signaling. Note that the first two prongs aimed for deterrence, not prevention. In retrospect, this strategy was at least partly effective—defrauding the telephone company is a serious crime. Fraud-detection technology was deployed selectively throughout the network to aid in identifying suspected blue-box users, who were subject to widely publicized criminal prosecutions. The prospect of detection and prosecution presumably dissuaded many otherwise law-abiding would-be blue boxers in the years before the vulnerability was fixed.

Detering exploitation of wiretap countermeasures seems to be a much more difficult problem than deterring toll fraud, not least because many of those most motivated to deploy such countermeasures are already criminals (and, in any case, the use of wiretap countermeasures is generally not, in and of itself, a crime).

References

1. R. Rosenbaum, "Secrets of the Little Blue Box," *Esquire*, Oct. 1971, pp. 117–125 and pp. 222–226.
2. R. Oklahoma, "Regulating the Phone Company in Your Home," *Ramparts*, vol. 10, June 1972, pp. 54–57.
3. A.E. Ritchie and J.Z. Menard, "Common Channel Interoffice Signalling: An Overview," *Bell Systems Tech. J.*, vol. 57, Feb. 1978, pp. 221–250.
4. M. Eleccion, "Beating the Blue-Box Bandits," *IEEE Spectrum*, vol. 9, Aug. 1972, pp. 52–58.

12. L. Schenker, "Pushbutton Calling with a Two-Group Voice Frequency Code," *Bell Systems Tech. J.*, vol. 39, Jan. 1960, pp. 239–255.
13. Recommendation Q.24, *Multifrequency Push-Button Signal Reception*, Int'l Telecommunications Union, 1988.
14. *Telcordia Notes on the Networks*, tech. report SR-2275 issue 4, Telcordia Technologies, Oct. 2000.
15. *Communications Assistance for Law Enforcement Act*, third report and order, docket no. 97–213, US Federal Communications Commission, 1999.
16. *Communications Assistance for Law Enforcement Act*, docket no. 97–213, US Federal Communications Commission, 2002.
17. *EWSD Integrated CALEA with Dial-Out Capability*, tech. bulletin 02PBCALEA01, Siemens, 2002.
18. R.M. Howell, *Method of Intercepting Telecommunications*, US patent 5,920,611, Patent and Trademark Office, Sept. 1996 (issued July 1999).
19. R.M. Howell, *Telecommunications Intercept System*, US patent 5,943,393, Patent and Trademark Office, Sept. 1996 (issued Aug. 1999).
20. *United States of America v. Ahmed Abdel Sattar, Lynne Stewart, and Mohammed Yousry*, trial transcripts, testimony of special agent Michael Elliot, New York, Southern Dist. Ct., Oct 2004, pp. 7392–7399.

Micah Sherr is a PhD candidate in computer and information sciences at the University of Pennsylvania. His research interests include network security, protocol design and analysis, network intrusion detection and prevention, and privacy and data confidentiality. Sherr has an MSE in computer and information science from the University of Pennsylvania. He is a member of the IEEE and Usenix. Contact him at msherr@cis.upenn.edu.

Eric Cronin is a PhD candidate in computer and information sciences at the University of Pennsylvania. His research interests include network security, privacy, and distributed systems. Cronin has an MS in computer science and engineering from the University of Michigan. He is a member of the ACM, the IEEE, and Usenix. Contact him at ecronin@cis.upenn.edu.

Sandy Clark is a visiting scholar in the Distributed Systems Lab at the University of Pennsylvania and a computing systems manager at Princeton University. Her research interests include human-scale and network security, privacy, and risk evaluation. Clark has a BA in German language and computer science from Brigham Young University. She is a member of the IEEE, Usenix, and SAGE. Contact her at clarks@seas.upenn.edu

Matt Blaze is an associate professor of computer and information sciences and director of the Trusted Network Eavesdropping and Countermeasures project at the University of Pennsylvania. His research interests include secure systems, cryptography and cryptographic protocols, and large-scale systems. Blaze has a PhD in computer science from Princeton University. He is a member of the ACM, IACR, and the IEEE, and is a director of the Usenix association. Contact him at mab@crypto.com.