

## Chapter 1

# ON THE RELIABILITY OF CURRENT GENERATION NETWORK EAVESDROPPING TOOLS

Eric Cronin, Micah Sherr, and Matt Blaze

**Abstract** This paper analyzes the problem of interception of Internet traffic from the eavesdropper’s point of view. We examine the reliability and accuracy of transcripts, and show that obtaining “high fidelity” transcripts is harder than previously assumed. Even in highly favorable situations, such as capturing unencrypted traffic using standard protocols, simple – and entirely unilateral – countermeasures are shown to be sufficient to prevent accurate traffic analysis in many Internet interception configurations. In particular, these countermeasures were successful against every available eavesdropping system we tested. Central to our approach is a new class of techniques that we call *confusion*, which, unlike cryptography or steganography, does not require cooperation by the communicating parties and, in some case, can be employed entirely by a third party not involved in the communication at all.

**Keywords:** eavesdropping, electronic interception, eavesdropping countermeasures

## 1. Introduction

The results of Internet interceptions are almost always accepted uncritically. While previous work has shown the potential for spurious errors [Bel00, BB00] or evasion [PN98, Pax99] to interfere with capture, there has been remarkably little exploration of the problems which face an eavesdropper who wishes to ensure the accuracy of their intercepts. We assert that the task of the eavesdropper is actually far more difficult than has previously been realized, and show that existing tools are insufficient to gauge the accuracy of captured traffic.

At least six properties of the Internet protocol stack and architecture make it difficult for an eavesdropper to accurately reconstruct communications from its intercepts: decentralized control and heterogeneous implementations; “best effort” (as opposed to reliable) message deliv-

ery that allows data to be re-ordered, duplicated or dropped in transit; shared state and context between communicating parties; dynamic (and often asymmetric) routing that can change during a flow's lifetime; lack of sender and receiver authentication; and ambiguities in protocols, implementations, and configurations.

These properties mean that a great deal of state information is involved in the correct interpretation of any given packet, and this state is spread across many places, including each of the communicating parties and the network itself. Without complete knowledge of this state, the mere presence of a packet somewhere on the network does not automatically imply that it will be accepted by the recipient given in its header, that it came from the supposed sender, or that it has not been (or will not be) altered, duplicated, or deleted somewhere along its path.

Any intercept system must take into account these properties (and all the corresponding state) in order to ensure not only that it is sufficiently *sensitive* (that it receives all data exchanged between the targets), but that it is also sufficiently *selective* (that it rejects spurious data that is not actually part of the targets' exchange) [CSB05a]. The figure of merit most often considered in judging intercept systems is sensitivity; adequate selectivity, on the other hand, is generally thought to be easily achieved by cursory examination of, e.g., packet headers. In fact, selectivity may be a far more difficult problem than most intercept systems recognize, especially in the presence of deliberate countermeasures.

Fortunately for the eavesdropper, on more benign networks at least, many of the factors that might introduce uncertainty about the veracity and interpretation of a given packet are relatively static, at least for the lifetime of a particular interception. For example, although routes can theoretically change midstream, in practice, they rarely do, and although routers and hosts are free to alter, reorder, delay, and duplicate packets, for the most part they refrain from doing so.

However, this lends a false sense of security to those producing eavesdropping tools. Depending on the network configuration, many ambiguities can be intentionally induced, either by one of the communicating parties or by a third party altogether. In fact, across much of the protocol stack, from the physical layer to the applications, it is surprisingly simple to introduce data that appears entirely valid but that might not be received and processed by the purported recipient. The Internet appears almost to have been designed to maximize uncertainty from the point of view of those eavesdropping on it.

In particular, we observe that a single party, which we call a *confuser*, can introduce traffic directed at an eavesdropper but that is never actually received (or if received, is rejected) by the ostensible recipi-

ent. Depending on the eavesdropper's configuration and its position in the network, this traffic can be made indistinguishable from legitimate traffic. In the presence of sufficient confusion, an eavesdropper may be able to be made arbitrarily uncertain as to whether a given intercepted message was real or spurious.

Let us introduce some terminology that will be used throughout the remainder of this paper. As is customary, *Alice* and *Bob* will represent our network communicators. Alice will often be a source while Bob will be a sink (although, of course, in most protocols the roles are symmetric and often alternating). *Eve* will be our eavesdropper. An interception system is vulnerable to *confusion* if it captures and records in its transcripts messages *purportedly* from Alice to Bob but that are rejected or otherwise not processed by Bob.

Although we do not advocate that confusion be used as a general confidentiality technique, we briefly note that confusion has some interesting qualities that make it particularly attractive as an eavesdropping countermeasure.

- While cryptography is typically used in a manner that ensures the confidentiality only of message payloads, confusion protects both a message's contents and metadata. It may therefore be advantageous to combine confusion with encryption to mask signaling information as well as content.
- Since confusion is transparent to Bob, it may be easily incorporated into existing protocols. Thus, it may be particularly useful when legacy applications and protocols cannot be easily upgraded or replaced.
- If the confuser is a third-party, then neither Alice nor Bob needs to be aware of the confusion. Unlike bilateral techniques in which it is obvious that Alice and Bob have colluded to disguise their messages, confusion allows Alice and Bob to deny that they even attempted to communicate privately.

## 2. Related Work

There has been little prior work investigating the general problem of traffic interception from the eavesdropper's point of view [Bel00, BB00, CSB05b]. However, considerable research has addressed the related (but not identical) topic of information privacy. Cryptography, steganography, subliminal or covert channels [Sim83], winnowing and chaffing [Riv98], quantum communication [BBB<sup>+</sup>90], and anonymous

communications [DMS04, RR98], for example, all focus on establishing confidential communication.

Work from the eavesdropper’s point of view has primarily been limited to the specialized subtopic of intrusion detection [SP03, PP03, Pax99]. In a network intrusion detection system (NIDS), the primary goal of the listener (eavesdropper) is real-time analysis of incoming traffic to recognize attack signatures and detect anomalies. These systems are deployed at the borders of controlled networks where it becomes much easier to make assumptions about the machines within the network that the system protects. Additionally, the communication patterns of an attacker are also unique compared to general bidirectional communications (hence the NIDS is able to flag suspicious traffic). However, unlike a NIDS, a general purpose eavesdropper must process all traffic, both normal and anomalous. Because of these differences, we may draw from work on NIDS, but their applicability is limited by the different constraints on topology and communication characteristics.

### 3. Confusion in the Internet Architecture

By design, the Internet is a very heterogeneous system. Machines of differing hardware and software configurations communicate and interoperate through the use of standard protocols. However, ambiguities in implementations, configurations, and protocol specifications create the opportunity for non-uniformity in the processing of specially crafted messages. Confusion exploits these inconsistencies by forcing the eavesdropper to consider multiple plausible interpretations of its transcripts. The IP and TCP specifications (which famously advise “be conservative in what you do, be liberal in what you accept from others. [Pos81]”) thus aggravate the problem of proper selectivity by recommending that implementations accept even outlier communications.

Below, we explore various vectors and techniques for injecting confusion in the Internet architecture. The confusion countermeasures are not intended to be exhaustive; rather, their purpose is to illustrate the ease and effectiveness at which reliable interception can be defeated.

#### 3.1 Physical Layer Confusion

At the physical layer, network devices convert analog signals into digital encodings. To allow interoperable devices, standards exist that define acceptable ranges for amplitudes, frequencies, voltages, and so forth [IEE85, IEE90, IEE97]. However, because transmission and decoding are analog processes, for any given parameter (frequency, amplitude, etc.), no two decoders will use precisely the same threshold to determine

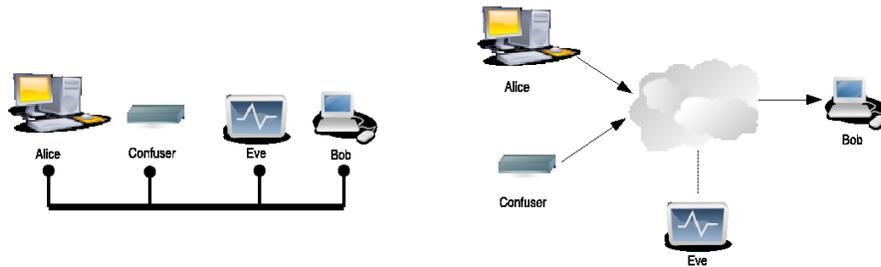


Figure 1. Left: An eavesdropping topology in which all parties communicate via the same shared bus. Right: A configuration in which Eve is located on the network between Alice and Bob.

whether a given signal is accepted or rejected. Thus, network devices, particularly commodity hardware, do not strictly abide by these standards and often interpret messages sent outside of the specified ranges.

Alice can exploit these differences to evade as well as confuse Eve. As depicted in Figure 1 (*left*), we assume a topology in which all parties share the same communication medium (e.g., a common bus or a wireless network). To *evade* Eve, Alice can transmit messages at a frequency, amplitude, or voltage that is imperceptible to Eve but acceptable by Bob. (Note that this type of physical evasion is more difficult when Alice, Bob, and Eve do not share a communication medium, as intermediary routers act as normalizers and reduce the likelihood of an effective evasion attack.) Generally, if Eve is less sensitive than Bob and the three parties share a communication medium, then Eve is susceptible to evasion.

Eve’s obvious counter-countermeasure (i.e., enhancing her sensitivity) has the unfortunate effect of increasing her vulnerability to confusion [CSB05a]. If Eve is *more* sensitive than Bob, evasion is not possible. However, a third-party confuser can now inject noise that is processed by Eve but ignored by Bob. As a result, Eve is forced to consider multiple interpretations, while Bob only sees the legitimate messages.

### 3.2 Link Layer Confusion

Confusion is possible at the link layer if the confuser and Eve share the same Ethernet. A typical example of such a topology is an unencrypted 802.11 network in which Eve “sniffs” wireless transmissions.

As we show empirically in Section 1.4, current eavesdropping systems suffer from inadequate selectivity. Although most eavesdropping systems are capable of recording traffic at the link layer, they often ignore Ethernet frames and instead process messages at either the network or transport layer. By crafting Ethernet frames with invalid MAC destina-

tion addresses, a confuser can inject noise that is processed by Eve but fails to be delivered to Bob [PN98]. Neither Bob nor the local gateway will process the noise since their operating systems silently discard Ethernet frames whose MAC addresses do not match that of the network interface.

This technique is obviously only effective when Eve has poor selectivity. If Eve examined the Ethernet frames, she would be capable of distinguishing the noise from the message text. Unlike other confusion countermeasures, the MAC technique is not indicative of a fundamental limitation of electronic eavesdropping. However, the significance of the approach is that it illustrates the dangers of inadequate selectivity: An eavesdropping system that fails to properly process Ethernet frames *is* inherently vulnerable to this form of confusion. Accordingly, an Internet eavesdropping system that observes traffic on a local Ethernet cannot claim to be reliable unless it both intercepts and processes link layer headers.

### 3.3 Network Layer Confusion

If Eve intercepts a packet on the path from Alice and Bob (see Figure 1, *right*), she must carefully examine the packet's IP header to form an opinion as to whether the packet is deliverable. There are several reasons that a packet may fail to be delivered: the packet's checksum may be incorrect, IP options may be specified that are unsupported by an intermediary router (e.g., source routing), the packet's size may exceed a hop's MTU, or the initial time-to-live (TTL) value may be insufficient to reach Bob [Pos81, PN98]. If the confuser has more knowledge about the network than Eve, he can inject noise that will be dropped either before reaching Bob or by Bob's IP implementation. If Eve processes all intercepted IP packets (which, as we show in Section 1.4, is the case with all tested eavesdropping systems), then she will interpret the noise along with the legitimate traffic.

As with the link layer techniques, the network layer confusion countermeasures highlight weaknesses in current eavesdropping systems. By enhancing Eve's selectivity, many of these countermeasures can be eliminated. However, an eavesdropper that either does not examine IP headers or lacks sufficient selectivity to determine whether packets are deliverable is inherently vulnerable to this type of confusion.

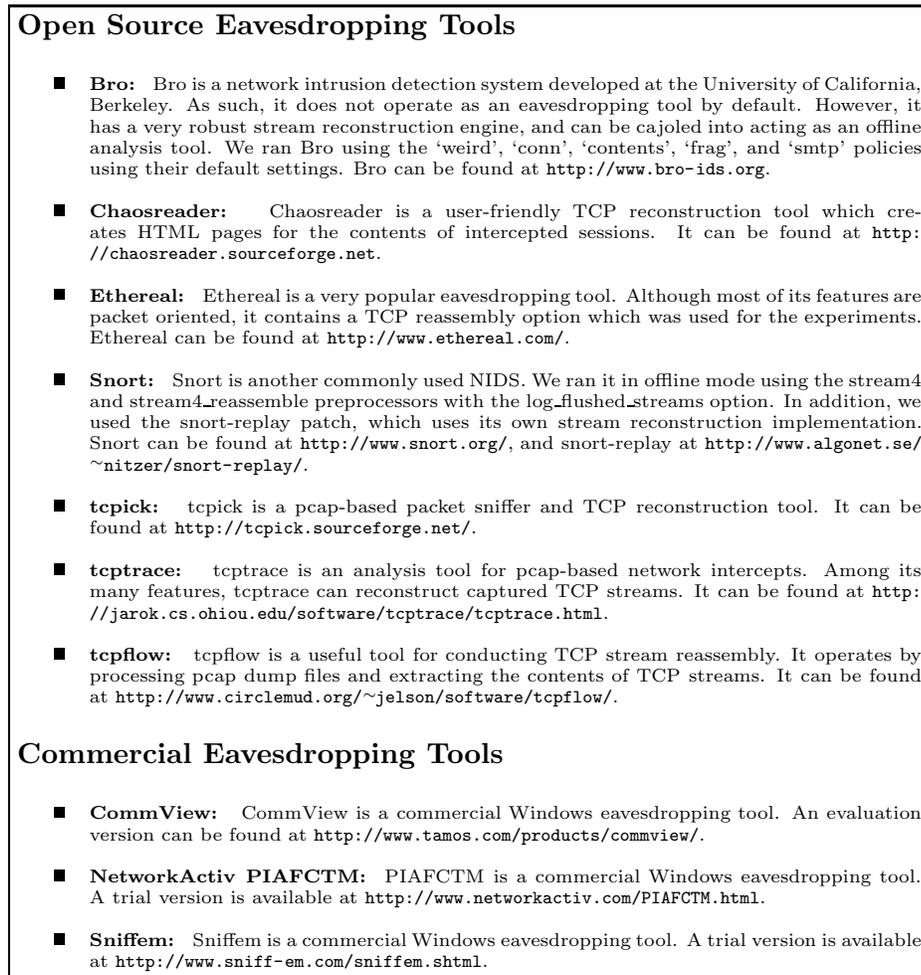


Figure 2. Eavesdropping software evaluated

## 4. Failure of Current Eavesdropping Systems

In this section we examine common tools for eavesdropping in several domains, and show how they fall vulnerable to simple, unilateral attacks. We look at both digital and analog examples.

### 4.1 Digital eavesdropping evasion

To demonstrate the susceptibility of current eavesdropping tools to confusion, we implemented the MAC and TTL confusion techniques described in Section 1.3 and originally introduced as NIDS attacks in [PN98].

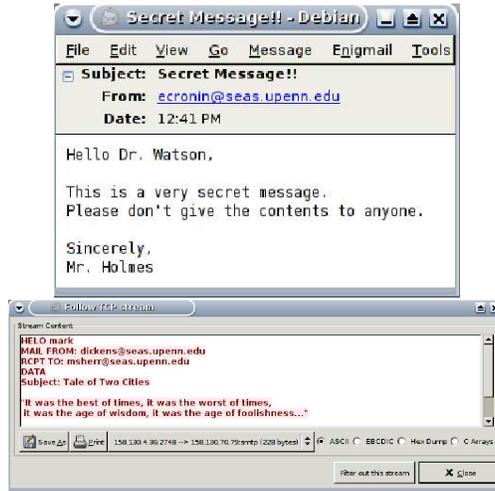


Figure 3. Top: Legitimate message received by the SMTP server (Bob) and the intended email recipient. Bottom: An eavesdropping system’s (Ethereal) reconstruction in which Eve not only fails to capture Alice’s message, she also perceives the covertedext as the legitimate message.

(Fragroute [Son99] also provides an implementation of the NIDS techniques, but it was found to be unsuitable for general purpose bidirectional communication.) The MAC approach relies on generating noise with invalid MAC destination addresses. While Eve will process the noise, the local gateway will not route such packets since it only accepts correctly addressed Ethernet frames. In the TTL technique, the confuser introduces noise with TTLs that are sufficient to reach Eve but not Bob. Note that both techniques can be trivially defeated by providing adequate selectivity. Here, our aim is not to introduce formidable countermeasures. Rather, we show that the current generation of eavesdropping tools are highly susceptible to even these weak forms of confusion.

In our experiments, Alice transmits an email via SMTP to our institution’s email server (Bob). To confuse Eve, Alice (functioning as the confuser) injects spurious noise using either the MAC or the TTL confusion techniques. To maximize confusion, Alice sends both the legitimate email and the noise in byte-sized packets (recall that since TCP is stream based, applications that rely on TCP are generally unaffected by the size of the transmitted packets). For every byte of legitimate text, Alice sends eight noise packets. Of the eight noise streams, the first is comprised of a “cover message”. This first stream, although composed of noise, constitutes a false but sensible message (a passage from

Dickens’ “A Tale of Two Cities” [Dic59]). The remaining seven streams of noise consist of random characters. In an attempt to cause Eve to interpret the false stream rather than her true message, Alice always sends the false stream first, followed by a random intermixing of the legitimate stream and the seven random noise streams. No modifications were made to the SMTP server (Bob).

We tested our link and network layer confusion tools against 11 eavesdropping systems, ranging from commercial applications to free open-source toolkits (descriptions of the eavesdropping systems are provided in Figure 2). Experiments were conducted on a testbed network in which Alice and Eve reside on the same local subnet. From this subnet, a minimum TTL of five is required to reach Bob. Both Alice and Eve are Pentium servers with 3COM Fast EtherLink XL 100MB/s network cards and are connected via a 100MB/s switch.

The performance of the eavesdroppers in the presence of confusion was startlingly lacking. Figure 4 describes Eve’s (in)ability to reliably reconstruct the email messages. Although all but one eavesdropping packages were able to correctly reconstruct Alice’s message in the absence of confusion, all tested systems failed to interpret her message once either of the two confusion techniques was applied. Anomalies were reported only by 18% of the eavesdroppers with the MAC-based approach and 27% of the systems when TTL confusion was used. Moreover, the cover message was perceived as the email in 45% of the cases when either technique was utilized (see Figure 3). In all cases, the email server (Bob) correctly received Alice’s communication and delivered the email to its intended recipient.

## 4.2 POTS Evasion and Confusion

Confusion and evasion can be practical threats to digital Internet eavesdropping, and indeed, such systems are the focus of this paper. However, similar techniques can also be applied to analog networks, especially when analog to digital conversion is performed (see Section 1.3.1). For a simple example, we consider voice telephone signaling between the subscriber and the switch.

Analog telephone service, also known as “plain old telephone service” (POTS), uses analog touch-tones to signal the caller’s desired number. The touch-tone system is an international standard known as DTMF [Int88], followed by both telecommunications equipment manufacturers and the telephone network to ensure compatibility. Each of the 16 DTMF digits is comprised of two base frequencies. The specification

lists acceptable ranges for tone duration, spacing, frequency, amplitude, and twist (relative amplitude of the high and low frequencies).

Constructing a decoder for DTMF signals which exactly follows these specifications turns out to be surprisingly difficult and expensive. Instead, most decoders relax the tolerances in one or more of the dimensions of the standard, accepting tones which are a little too loud or quiet, or a little out of pitch, etc. Since each decoder has a unique range of acceptable tones, it is possible to construct out-of-specification tones that some decoders will see but others will miss. Using this knowledge, it is possible to both evade and confuse an eavesdropper.

To test the practicality of this countermeasure, we constructed a simulated phone network in our lab. A Teltone TLS-5 Telephone Line Simulator was used as the central office, providing dialtone and four phone lines with unique phone numbers, as well as DTMF switching to connect calls between lines. Alice, the call initiator, was an Ameritec AM8a PCM/VF call analyzer with an ordinary POTS handset connected, Bob was a second POTS handset, and Eve was a Metrotel VNA 70a DTMF decoder.

Among its many features, the AM8a call analyzer used by Alice allows DTMF codes to be generated with very precise and configurable parameters. For our experiments, we focused on two variables: the frequency and amplitude of the higher of the two DTMF frequencies. Through binary search, settings which allowed for both evasion and confusion were discovered. Evasion can be applied by decreasing the amplitude of the higher frequency. At -39dBm, the central office still correctly decodes Alice's signal and completes the call, while the eavesdropper records nothing. Similarly, if instead the high frequency is increased by 3Hz the central office no longer recognizes Alice's touch-tones, but the eavesdropper records them as having been dialed. Using Alice's handset in coordination with the AM8a, the legitimate number can be dialed interspersed with out-of-range digits to provide confusion. In addition, although we did not test the scenario, by combining both techniques it is clear that Alice could drive Eve to a specific false phone number.

This experiment highlights the challenges which face an eavesdropper when positioned too close to the sender. Limited sensitivity and imperfect selectivity make it susceptible to both evasion and confusion countermeasures. While Eve may be certain that intercepts originate from Alice, she cannot be certain of where in the telephone network they terminate. A far more reliable form of dialed number recording is therefore achieved through analysis of call detail records generated by the switch itself, but this is, of course, not surreptitious with respect

to the operators of the switch. Confusion as a telephone wiretapping countermeasure is studied in more detail in [SCCB05].

## 5. Improving Eavesdropping Reliability

The experiments described in the previous section show how unilateral countermeasures can reduce the reliability of eavesdropping systems. In this section, we explore methods to improve eavesdropping tools' resilience to such countermeasures.

### 5.1 Enhancing Sensitivity

To reduce her susceptibility to evasion, Eve can improve her sensitivity. This implies recording at the lowest possible OSI layer, and recording everything available (even data that appears to be erroneous). Any action that could have been performed automatically by lower layers, such as discarding corrupt packets, can be carefully emulated by Eve in a more selective manner.

Unfortunately, this advice may be hard to follow. For example, many authorized uses of eavesdropping in the United States operate under strict limitations on what can be recorded to prevent traffic of those not under suspicion from being observed. In such environments the steps Eve can take to improve sensitivity are reduced.

### 5.2 Enhancing Confusion Detection and Eavesdropper Selectivity

In some situations, confusion may be made ineffective by deploying confusion-aware eavesdroppers. For example, the MAC confusion technique described in Section 1.3 can be defeated with improved software. By enhancing her sensitivity, Eve may be able to better identify and filter the noise, thereby improving her reliability. However, if Eve is careless in her selections and ignores packets with covert information, she provides Alice and Bob with an unmonitored communication channel.

### 5.3 Active Eavesdropping

Confusion is only possible when there is an asymmetry in knowledge between Eve and the confuser. To inject uncertainty in Eve's transcripts, the confuser exploits his knowledge (e.g., the network topology or Bob's TCP/IP stack configuration) to ensure that the noise will be removed or filtered before being processed by Bob. If Eve can also acquire this knowledge, then she can apply the same filter and can therefore trivially defeat confusion.

The intuitive solution to constructing a confusion-resistant eavesdropper is to make Eve active. In addition to passively observing traffic, an *active eavesdropper* attempts to learn more about the network and the communicating parties by sending out probes. For example, an active eavesdropper can counter the TTL confusion technique described in Section 1.3 by counting the number of network hops between itself and Bob. By acquiring additional knowledge, Eve can improve her selectivity and overall reliability.

Unfortunately, active eavesdropping is not always sufficient to ensure reliable reconstruction of the intercepted traffic. First, the probes used by an active Eve can themselves be subjected to a form of confusion. As a counter-counter-countermeasure, a confuser can inject a number of fake responses to Eve's probes. Returning to the TTL confusion example, a confuser can transmit fake ICMP TTL-exceeded messages to frustrate Eve's ability to discern the true TTL cutoff. Second, if Eve actively transmits probes, she may reveal her presence to Alice, Bob, and/or the confuser. Since eavesdropping is usually meant to be clandestine, active eavesdropping may be inappropriate for many situations.

#### 5.4 Improving Reliability through Eavesdropper Placement

The location of Eve in the network topology may affect her resilience to confusion. An intuitive approach is to position her in close proximity to Alice. The ability of distant third-party confusers to inject noise is thus diminished as Eve can better discern Alice's communications from those of a distant forger. Unfortunately, this strategy is ineffective when Alice functions as the confuser. Unless Eve can determine which of Alice's messages are authentic, her position does little to improve her reliability.

A better solution is to place Eve as close as possible to Bob (and henceforth as far as possible from any confusers). For example, the TTL confusion technique will be ineffective if Bob and Eve reside on the same local network. A disadvantage of this approach is that Eve can only make reliable claims about the messages received by Bob. Her distance from Alice may make the authenticity of intercepted messages harder to establish.

A more ideal strategy is to deploy a number of collaborating eavesdroppers throughout the network. By comparing messages intercepted near the sender versus the receiver, Eve may be able to remove likely noise and improve her reliability. We leave the analysis of colluding eavesdropping as a future research direction.

## 6. Conclusion

For electronic wiretapping systems to be reliable, they must exhibit correct behavior with regard to both sensitivity and selectivity. Since capturing traffic is a requisite of any monitoring system, considerable research has focused on preventing evasion attacks and otherwise improving sensitivity. However, little attention has been paid to enhancing selectivity or even recognizing the issue in the Internet context.

Traditional wisdom has held that eavesdropping is sufficiently reliable as long as the communicating parties do not participate in a bilateral effort to conceal their messages. We have demonstrated that even in the absence of cooperation between the communicating endpoints, reliable Internet eavesdropping is more difficult than simply capturing packets. If an eavesdropper cannot definitively and correctly select the pertinent messages from the captured traffic, the validity of the reconstructed conversation can be called into question. By injecting noise into the communication channel, unilateral or third-party confusion can make the selectivity process much more difficult, diminishing the reliability of electronic eavesdropping.

Whether eavesdropping can be performed reliably and confusion correctly detected and rejected on the Internet depends heavily on the specific interception topology and on the locations of potential sources of confusion traffic. Even in those configurations where confusion can theoretically be filtered out, the eavesdropping software itself may still be susceptible to confusion, and, in fact, current software appears to be especially vulnerable to even the simplest confusion techniques.

## Acknowledgments

The authors would like to thank Harry Hoffman for his assistance configuring Bro and Snort for the experiments in Section 1.4.1. This work was partially supported by the US National Science Foundation Cyber-Trust program under contract NSF-0524047.

Software	No Confusion (1B pkts)		MAC Confusion		TTL Confusion	
	Inter-pretation	Detected Anoma-lies	Inter-pretation	Detected Anoma-lies	Inter-pretation	Detected Anoma-lies
bro	Success	None reported	Failure (Cover-text)	Retrans. inconsis-tency	Failure (Cover-text)	Retrans. inconsis-tency
chaosreader	Success	None reported	Failure (Ran-dom noise)	None reported	Failure (Ran-dom noise)	None reported
CommView Eval. Ver-sion	Success	None reported	Failure (Cover-text)	None reported	Failure (Cover-text)	None reported
ethereal	Success	None reported	Failure (Cover-text)	None reported	Failure (Cover-text)	None reported
Network-Activ PI-AFCTM	Success	None reported	Failure (Cover-text)	None reported	Failure (Cover-text)	None reported
Sniffem	Failure (Ran-dom noise)	None reported	Failure (Ran-dom noise)	None reported	Failure (Ran-dom noise)	None reported
snort-replay	Success	None reported	Failure (Ran-dom noise)	None reported	Failure (Ran-dom noise)	None reported
snort-stream4	Success	None reported	Failure (Ran-dom Noise)	None reported	Failure (Ran-dom Noise)	TTL Exceeded
tcpick	Success	None reported	Failure (Cover-text)	None reported	Failure (Cover-text)	None reported
tcptrace	Success	None reported	Failure (Ran-dom noise)	TCP DUPs detected	Failure (Ran-dom noise)	TCP DUPs detected
tcpflow	Success	None reported	Failure (Ran-dom noise)	None reported	Failure (Ran-dom noise)	None reported

*Success* - The eavesdropping application correctly interpreted the messagetext.

*Failure (Coverttext)* - The eavesdropping application incorrectly interpreted the coverttext as the legitimate messagetext. See Figure 3.

*Failure (Random noise)* - No discernible English text could be obtained from the eavesdropper's interpretation.

*Figure 4.* Ineffectiveness of various eavesdropping software against confusion techniques.

## References

- [BB00] M. Blaze and S. M. Bellovin. Inside RISKS: Tapping, tapping on my network door. *Communications of the ACM*, 43(10), December 2000.
- [BBB<sup>+</sup>90] C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. A. Smolin. Experimental quantum cryptography. In *Advances in Cryptology - EUROCRYPT*, May 1990.
- [Bel00] S. M. Bellovin. Wiretapping the net. *The Bridge*, 20(2):21–26, 2000.
- [CSB05a] E. Cronin, M. Sherr, and M. Blaze. The eavesdropper’s dilemma. Technical Report MS-CIS-05-24, University of Pennsylvania, 2005.
- [CSB05b] E. Cronin, M. Sherr, and M. Blaze. Listen too closely and you may be confused. In *Proc. of 13th International Security Protocols Workshop* (to be published), 2005.
- [Dic59] C. Dickens. *A Tale of Two Cities*. Apr 1859.
- [DMS04] R. Dingleline, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. In *Proc. of the 13th Usenix Security Symposium*, pages 303–320, Aug 2004.
- [IEE85] IEEE. IEEE standards for local area networks: carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications. IEEE 802.3, 1985.
- [IEE90] IEEE. Information processing systems - local area networks - part 4: token-passing bus access method and physical layer specifications. IEEE 802.4, August 1990.
- [IEE97] IEEE. IEEE Std 802.11-1997 information technology-telecommunications and information exchange between

- systems-local and metropolitan area networks-specific requirements-part 11: Wireless lan medium access control (MAC) and physical layer (PHY) specifications. IEEE 802.11, November 1997.
- [Int88] International Telecommunication Union. Multifrequency push-button signal reception. Recommendation Q.24, Telecommunication Standardization Sector of ITU, 1988.
- [Pax99] V. Paxson. Bro: a system for detecting network intruders in real-time. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(23–24):2435–2463, 1999.
- [PN98] T. Ptacek and T. Newsham. Insertion, evasion, and denial of service: Eluding network intrusion detection. Technical report, Secure Networks, Inc., 1998.
- [Pos81] J. B. Postel. Internet protocol. RFC 791, Internet Engineering Task Force, September 1981.
- [PP03] R. Pang and V. Paxson. A high-level programming environment for packet trace anonymization and transformation. In *Proc. ACM SIGCOMM 2003*, Aug. 2003.
- [Riv98] R. Rivest. Chaffing and winnowing: Confidentiality without encryption. <http://theory.lcs.mit.edu/~rivest/chaffing.txt>, March 1998.
- [RR98] M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. In *ACM Transactions on Information and System Security*, 1998.
- [SCCB05] M. Sherr, E. Cronin, S. Clark, and M. Blaze. Signaling vulnerabilities in wiretapping systems. *IEEE Security and Privacy*, pages 24–36, Nov/Dec 2005.
- [Sim83] G.J. Simmons. The prisoners’ problem and the subliminal channel. In *Proc. of IEEE Workshop on Communications Security CRYPTO’83*, 1983.
- [Son99] D. Song. fragroute, 1999. <http://monkey.org/~dugsong/fragroute/>.
- [SP03] U. Shankar and V. Paxson. Active mapping: Resisting NIDS evasion without altering traffic. In *Proc. of the 2003 IEEE Symposium on Security and Privacy*, May 2003.