

Toward an Analytical Approach to Anonymous Wireless Networking

Parv Venkitasubramaniam, Ting He, Lang Tong, and Stephen B. Wicker, Cornell University

ABSTRACT

Communications in a wireless network are susceptible to unauthorized traffic analysis by eavesdroppers. Although cryptography can protect the contents of communication, the transmission times of packets alone can reveal significant *networking information* such as source-destination pairs and routes of traffic flow. This article focuses on analytical approaches to *anonymous* networking for the prevention of information retrieval through packet timing analysis. In particular, an analytical measure for route anonymity is proposed using information theoretic equivocation. Based on the metric, provably anonymous countermeasures are described for wireless ad hoc and sensor networks, where traffic is subjected to strict constraints on medium access and latency. A key objective is to bridge a long standing gap between the information theoretic approach to secrecy in communication, and a more pragmatic approach of Chaum Mixing used in anonymous systems on the Internet. The efficacy of the proposed approach is demonstrated using an orthogonal transmitter directed signaling network.

INTRODUCTION

As wireless networks increasingly dominate our means of communication, the need for security and privacy has gained significant prominence. Due to the unprotected communication medium, wireless networks are vulnerable to unauthorized retrieval of networking information. For example, by merely correlating packet transmission times of different nodes, a passive eavesdropper can decipher source-destination pairs and paths of traffic flow in a network. The retrieval of such information, known as *traffic analysis*, is a violation of user privacy. Furthermore, the retrieved information can also be used for jamming network traffic and launching a denial-of-service attack.

In emerging wireless technologies such as ad hoc and sensor networks, traffic analysis can be particularly damaging. For example, in sensor networks, even without knowledge of source-des-

tinuation pairs, the direction of data flow alone can reveal sensitive information such as event location or access points. Furthermore, the open deployment of nodes with limited or no physical protection makes the network susceptible to active means of traffic analysis such as node capturing or packet flooding.

The design of protective measures against traffic analysis depends on the total information available to the adversary. Apart from transmission times of nodes obtained through eavesdropping, an adversary may also possess knowledge of communication protocols and signaling strategies used in the network. While secrecy of messages is ensured through encryption, hiding the act of communication requires a fundamental redesign of networking protocols. The challenging task in designing protocols that are resilient to traffic analysis is to obfuscate the transmission patterns without violating networking constraints. In this regard the wireless medium has its share of advantages and disadvantages. On one hand, it is difficult for an eavesdropper to ascertain the sender-receiver pair of an encrypted wireless transmission, especially when many streams are multiplexed at a single node. On the other hand, the shared wireless medium is band-limited and susceptible to fading and interference, thereby constraining the network designer.

As a motivating example, consider a multi-access relay as shown in Fig. 1, where node B forwards packets received from nodes S_1 and S_2 (to D_1 and D_2 , respectively). Consider the relaying of delay-sensitive traffic, where every packet received by the relay needs to be forwarded within a fixed delay or otherwise dropped. Let the nodes transmit on orthogonal noninterfering channels, such that the maximum transmission rate of each node is bounded. The rates of packets that can be relayed successfully from the two sources lie in a pentagonal region, as shown in Fig. 1. Any point in this region can be achieved if the sources transmit at the corresponding rates, and the relay merely forwards each received packet as soon as it is received. It is easy to see that this forwarding strategy would result in the transmission schedules of the source nodes and relay

This work is supported in part by the National Science Foundation under awards CCF-0635070 and CCF-0728872, the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program DAAD19-01-2-0011, and the Team for Research in Ubiquitous Secure Technology (TRUST) sponsored by the National Science Foundation under award CCF-0424422.

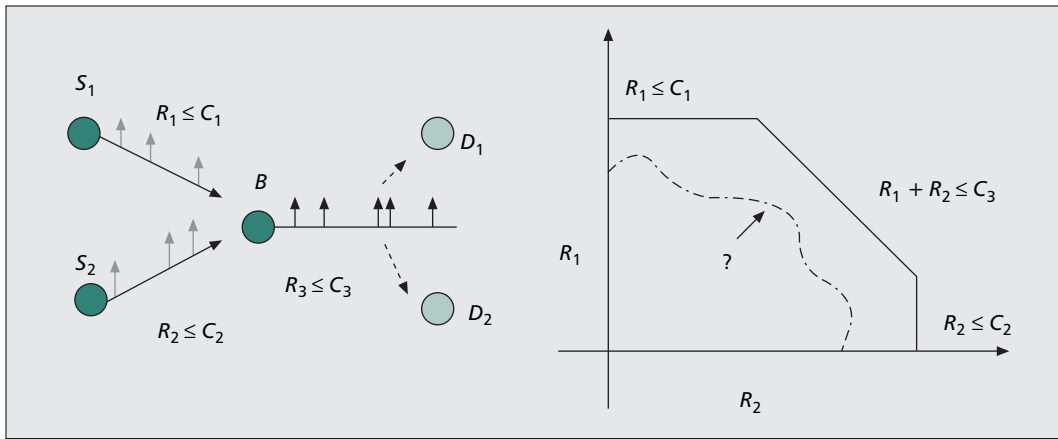


Figure 1. Achievable rate region for two-hop relay. Outer region: The transmission rates of nodes S_1 , S_2 are bounded by C_1 , C_2 , respectively. The sum rate of relayed packets from S_1 , S_2 are bounded by C_3 to maintain stability at the relay. Inner region: Unknown relay rate region when transmission schedules are independent.

being highly correlated; an eavesdropper would be able to detect the correlation and identify that B is an intermediate relay. Suppose the schedules of all nodes are forced to be statistically independent; an eavesdropper cannot detect a correlation across any pair of schedules, thus effectively hiding the relay operation. However, since each packet cannot be delayed indefinitely, enforcing statistically independent incoming and outgoing schedules would result in packets being dropped and therefore lower achievable rates of relayed packets.

The multiaccess relay of Fig. 1 represents a basic component in wireless networking, and hiding the relay operation is a key step toward designing anonymous network protocols. The example highlights that providing *anonymity* would result in reduced network performance. This raises some interesting and important questions. Can we measure the anonymity of routes in a wireless network? If so, how do we design protocols that guarantee a desired amount of anonymity? What are the trade-offs between achievable network performance and the level of anonymity provided?

In this article we present an analytical framework to address these questions. In particular, the proposed approach aims to combine the mathematical strength of information theoretic secrecy and practical algorithms used by anonymous systems on the Internet to develop rigorous analytical foundations for anonymous wireless networking.

The rest of this article is organized as follows. First, the concept of Chaum Mixing and its applicability in wireless networks are discussed. This is followed by the exposition of our analytical model for measuring anonymity. Based on the defined measure, the idea of covert relaying to guarantee anonymity at individual nodes is then described. Following this, the network level strategy is presented, where a connection to rate-distortion is used to provide a trade-off between achievable network performance and the desired level of anonymity. Finally, some concluding remarks and possible ramifications of this approach are discussed.

ANONYMITY IN INTERNET PROTOCOLS: MIXING

Anonymous protocols for Internet applications such as email or browsing have been designed based on the concept of Mixing [1], pioneered by Chaum. A Mix is a special node or server that collects packets from multiple users, and transmits them after re-encryption and length padding. This ensures that by comparing contents or packet lengths, it is impossible to match an incoming and outgoing packet at a Mix. Furthermore, packets are collected by the Mix from multiple sources and transmitted in batches, so that an eavesdropper is unable to identify the source of a packet using the arrival and departure times. A single Mix is liable to be compromised, and hence a (possibly random) sequence of Mixes are interposed between source and destination terminals to protect against active means of gaining information. The key idea is to perform encryption for multiple Mix servers in layers so that a compromised Mix can at most reveal the address of the subsequent Mix in the path. An eavesdropper can therefore not determine the source-destination pairs unless all the Mixes are compromised.

Subsequent to the original design proposed by Chaum, there have been several improvements to the batching strategies of Mixes to handle different types of traffic analysis attacks. While Mix-based solutions are useful in Internet applications such as anonymous remailers and Web browsing, a study of flow correlation attacks [2] showed that when long streams of packets with buffer or latency constraints are forwarded through Mixes, it is possible to correlate incoming and outgoing streams almost perfectly. Constraints on delay and buffer size are critical in network applications, such as media transmission, or in sensor networks, where battery powered sensor nodes cannot store packets indefinitely. A known solution for low-latency anonymous communication that caters to packet streams is the use of dummy transmissions to cover the actual flow of traffic. Systems that use this approach require users to maintain a con-

The key idea is to perform encryption for multiple Mix servers in layers so that a compromised Mix can at most reveal the address of the subsequent Mix in the path. An eavesdropper can therefore not determine the source-destination pairs unless all the Mixes are compromised.

In the event of active compromising of nodes, the model can be generalized by incorporating additional side information available to the adversary. One example of side information would be the partial knowledge of the routes as provided by a set of compromised nodes.

stant transmission rate of packets whether or not they have actual data to communicate. This ensures that to an external eavesdropper, the observed pattern of traffic is fixed irrespective of the routes of communication. However, the effects of dummy transmissions on network throughput or latency have not been analyzed or optimized. Our intent is to use ideas such as Mixing and dummy transmissions to develop strategies for an ad hoc wireless network setting, where the challenge lies in maximizing network performance under the constraints on bandwidth, latency, and anonymity.

AN INFORMATION THEORETIC MEASURE OF ANONYMITY

The first step toward an analytical model for this problem is to define a quantifiable measure of anonymity. In the context of Mix networks, anonymity has been measured using the *anonymity set* of an observed packet. The anonymity set is the set of all possible source-destination pairs for that packet. A weaker notion using the *sender anonymity set* was also provided by the General Packet Radio Service (GPRS) standard until very recently. The service allowed a mobile user to exchange data packets with a predefined host that can be addressed by the supported internetworking protocols. Since all user packets go through a GPRS gateway source node (GGSN), the sender's identity remains anonymous. Here, the anonymity is measured using the size of the sender anonymity set (set of all possible senders). Although the anonymity sets serve as good models in Internet applications, for networks such as wireless ad hoc or sensor networks, it is equally essential to protect the routes of data flow as it is to hide the source destination information. Furthermore, when streams of packets are transmitted, the anonymity has to be guaranteed, not only to individual packets, but to the information flows.

The analytical model we propose is motivated by the information theoretic approach to secrecy. Parallel to cryptography, secrecy in communication has been studied extensively from an information theoretic perspective using the idea of equivocation. Equivocation, proposed by Shannon [3], measures the uncertainty of a transmitted message with respect to an eavesdropper's observation. It has been used to measure secrecy in point-to-point communication channels such as wiretapped [4] and broadcast channels [5]; the goal was to maximize the reliable rate of communication to an intended receiver, while guaranteeing a desired level of secrecy (equivocation) with respect to the eavesdropper.

While previous applications of equivocation have been restricted to secrecy of transmitted messages, we use it to quantify the anonymity of the *routes* in a network. Consider a network where V represents the set of nodes. During any observation of the network traffic, some subset of nodes in V communicate using a fixed set of routes. This set of routes S , to which we refer as a *network session*, represents the information we wish to hide from the eavesdropper.

During any network session, each node transmits packets at arbitrary times, which can be represented as a point process, say Y_A for node A . In a wireless network, if packet headers are encrypted, detection of a transmitted packet may not directly reveal the identity of the transmitting or receiving node. However, eavesdroppers can be stationed such that the received power level can be used to estimate the location of transmitting nodes. Alternatively, if the underlying physical layer utilizes an orthogonal transmitter directed signaling model, the eavesdropper can identify the transmitting node using the knowledge of the spreading sequences. Since the locations of eavesdroppers cannot be determined, we assume every node is monitored; the eavesdropper's complete observation is therefore the set of point processes $Y = \{Y_A : A \in V\}$.

We define the anonymity in the network using the conditional uncertainty of the session with respect to eavesdroppers' observation. We model the network session as a random variable $S \sim p(S)$, and measure anonymity of the session using the normalized equivocation,

$$\alpha_s = \frac{H(S|Y)}{H(S)}.$$

The maximum value of α_s is 1, in which case the designed schedules provide *perfect anonymity*; upon observing Y , the eavesdropper obtains no additional information about the session S than that provided by the prior $p(S)$. For a general $\alpha_s < 1$, the physical interpretation of anonymity comes from Fano's Inequality [6]; equivocation provides a lower bound to the *error probability* of the eavesdropper in decoding the session correctly.

The fundamental design problem is, given a session S , what is the optimal scheduling strategy Y , such that anonymity $\alpha_s \geq \alpha$ is guaranteed. An optimal scheduling strategy is defined as the distribution $q(Y|S)$ that maximizes the desired performance metric such as network throughput, latency, or energy efficiency.

Note that this analytical model assumes that eavesdroppers do not use active means of gaining inference such as compromising nodes. In the event of active compromising of nodes, the model can be generalized by incorporating additional side information available to the adversary. One example of side information would be partial knowledge of the routes as provided by a set of compromised nodes.

PROVIDING ANONYMITY

The Mix network approach, as shown in [2], does not guarantee required anonymity under the constraints of wireless networks. An alternative solution is to fix a scheduling strategy for the nodes irrespective of the active routes. If a node does not have data packets to transmit, it adheres to the schedule by transmitting dummy packets. Fixed scheduling, analyzed in [7] for wireless networks, guarantees perfect anonymity to the routes at all times. However, the constant transmission of dummy packets makes it an inefficient alternative for large networks. Further-

more, the implementation of a fixed schedule requires synchronization across all nodes and a constant network topology, which is not practical in energy constrained ad hoc wireless networks.

We propose an approach that shares some traits with the ideas of Mixing and dummy transmissions, but guarantees any desired level of anonymity under strict network constraints on medium access and delay. Specifically, depending on the routes in each session, we divide the set of relays into two categories (Fig. 2), covert and visible, which are described below.

Covert relays: A relay B (as shown in Fig. 2) is covert if its outgoing transmission schedule is statistically independent of the transmission schedules of all nodes occurring previously in the paths that contain B . When a relay is covert, it would be impossible for an eavesdropper to correlate the transmission schedule of any preceding node in a path and that of relay B . In effect, the relaying operation of B is hidden. However, due to constraints on delay and bandwidth, the independence in transmission schedules would result in dropped packets or require dummy packet transmissions.

Visible relays: A visible relay B generates its schedule depending on the arrival times of packets at B . The relay forwards every packet received after a processing delay. For a visible relay, it is evident that the schedules of streams transmitted by a preceding node in the path and the relay would be highly correlated, and an eavesdropper can detect the relaying operation.

A covert relay that forwards packets from multiple sources plays the role of a Mix by completely decorrelating the incoming and outgoing schedules. In addition, due to encrypted packet headers, an eavesdropper would be unable to decipher which routes contain the covert relay, even when long streams of packets are transmitted. However, due to dummy transmissions and possible packet drops by the covert relay, the achievable relay rates would be lower than those of a visible relay.

The covert relaying approach entails two fundamental design issues: optimal scheduling strategies for a covert relay and optimal selection of covert relays. The strategy to select covert relays depends on the routes of a session and the desired level of anonymity. This is described later. In the remainder of this section we discuss scheduling strategies of covert relays in more detail.

COVERT RELAYING

One approach to generate independent schedules at a covert relay would be to derive a queuing discipline that forwards packets within the required delay constraints, and yet results in a statistically independent outgoing schedule. Such a strategy would, however, be vulnerable to active inference techniques such as packet insertion or flooding attacks [8]. We propose an alternative strategy, where a random transmission schedule is generated a priori by the relay, and a subset of arrived packets are chosen to be forwarded so that the delay constraints are satisfied with minimum packet drops.

In [9] we designed covert relays with strict delay constraints, where each received packet

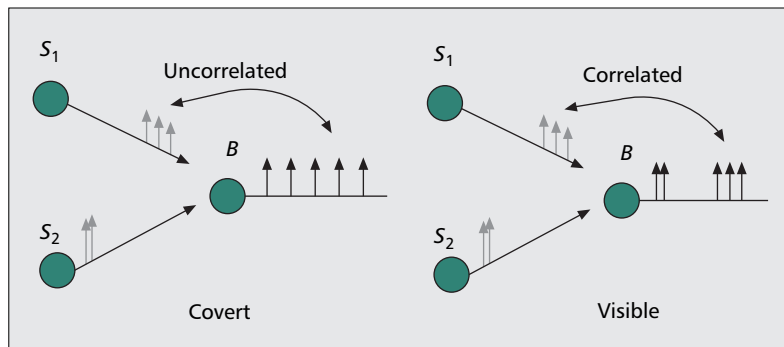


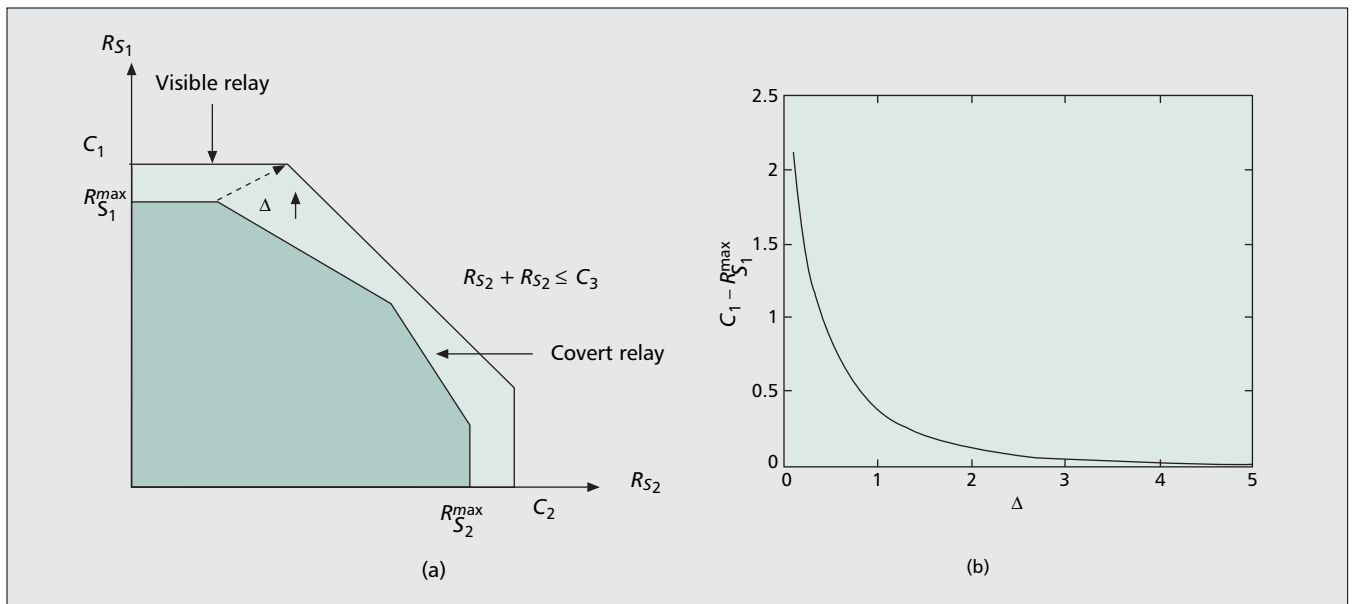
Figure 2. Covert and visible relaying.

needs to be relayed within Δ time units or dropped otherwise. For orthogonal transmitter directed signaling at the physical layer, we provided relaying strategies and characterized the set of achievable rates analytically for Poisson distributed schedules. The rate region for a 2×1 relay is illustrated in Fig. 3a. For any finite delay, the use of independent Poisson schedules results in a non-zero packet drop rate, as illustrated in the figure ($R_{S_1}^{\max} < C_1$). However, the rate region converges exponentially to that of a visible relay as $\Delta \rightarrow \infty$. This can be seen from Fig. 3b, which plots the difference in the maximum relay rate for S_1 between covert and visible relaying, as a function of Δ .

A key component in the optimal relaying strategies is the Bounded Greedy Match (BGM) algorithm, proposed in [10], in the context of chaff insertion for stepping stone attacks. For any pair of point processes, the BGM algorithm guarantees that packets are relayed within strict delay constraints with the fewest possible packet drops. For a multiaccess relay, we expanded this idea by assigning a priority level to each source node and choosing to relay packets depending on the priority levels. Furthermore, by employing time sharing of multiple strategies and forward error correction at each source, we characterized the achievable region (Fig. 3) in [9].

Poisson point processes are a particular scheduling distribution and generally suboptimal. An interesting open problem is to characterize the optimal distribution of transmission schedules that maximizes the achievable rates for a covert relay. The optimal distribution would vary depending on the medium access channel and the nature of delay constraints. For a transmitter directed signaling system with an average delay constraint, Poisson schedules have been shown to be optimal under certain medium access constraints [11]. Similarly, for a strict delay constraint, if the minimum allowed interpacket delay is much smaller than the delay bound, it can be shown that a periodic schedule is optimal.

Independent scheduling for covert relays is a specific technique for providing anonymity to routes in a network. Hiding the relaying operation of nodes is a broad idea and also an attractive solution because of its distributed nature. Each node decorrelates its incoming and outgoing schedules independent of the other nodes in the path. Furthermore, the relaying strategies we propose to maximize relay rates do not require



■ **Figure 3.** a) 2×1 relay rate region; b) difference between maximum relay rate for C_1 for a covert and visible relay, when $C_1 = C_2 = 3$, $C_3 = 4$.

the relay node to have a priori knowledge of the source nodes schedules; the decision to forward or drop a packet is made only after the packet has arrived at a relay.

THROUGHPUT-ANONYMITY TRADE-OFF

In general, every covert relay incurs a loss in performance, as either decreased rates of relayed packets or increased latency in relaying. It is therefore not beneficial to make all relays in a session covert, especially in a large network. Thus, the challenge is to optimize the choice of relays depending on the level of anonymity desired such that network performance is maximized. A key insight into this optimization comes from a connection to information theoretic rate distortion.

This connection is not based on the idea of covert relaying and can be explained using a general intuition. The objective of the rate distortion problem is to generate the fewest codewords for a set of source sequences, such that the distortion between each source sequence and the corresponding codeword is within a specified constraint. The idea is to divide the set of source sequences into the fewest bins and generate one codeword for each bin such that the distortion constraint is met. Alternatively, fixing the code rate fixes the total number of bins. Then the sequences are placed optimally within each bin such that the corresponding reconstruction sequences minimize the expected distortion. In the anonymous networking setup the key idea is to divide the set of all possible network sessions into bins such that for each bin there is a scheduling strategy that would make the sessions within that bin indistinguishable to an eavesdropper. The level of anonymity required determines the number of bins, and the optimal scheduling strategy plays the role of a codeword

by minimizing the performance loss across sessions within the bin.

In [11] we considered the special case of orthogonal transmitter directed signaling, and characterized the relationship between throughput and anonymity by proving an equivalence to a rate distortion function. Specifically, in any session S , the set of covert relays $B \setminus V$ was chosen randomly using a conditional distribution $q(B|S)$. Since the adversary is unable to correlate schedules at the covert relays, he/she only observes a portion of the session, denoted \hat{S} . The covert relays cause a reduction in throughput denoted $d(S, \hat{S})$, which is evaluated by adding the packet loss at each covert relay in the routes. For the transmitter directed model, we optimized the conditional distribution $q(B|S)$ in [11], and showed that the optimal throughput-anonymity trade-off is equivalent to a distortion rate function [6].

Theorem 1: (Throughput-Anonymity Trade-off) If $\lambda(0)$ is the throughput achievable with zero anonymity, throughput $\lambda(\alpha)$ is achievable with anonymity α if

$$\lambda(0) - \lambda(\alpha) \geq D(H(S)(1 - \alpha)),$$

where $D(R)$ is the information-theoretic distortion rate function:

$$D(R) = \min_{q(\hat{S}|S): I(S; \hat{S}) \leq R} E(d(S, \hat{S})). \quad (1)$$

Since \hat{S} denotes the session as observed by the eavesdropper, it is the equivalent of the reconstruction sequence in the rate-distortion optimization. Further, since both problems minimize mutual information, the level of anonymity α corresponds to the rate of compression. Therefore, characterizing the distortion-rate function is equivalent to characterizing the optimal throughput anonymity relation. The consequences of this equivalence, however, extend beyond the characterization of the optimal

throughput. Rate-distortion is a field that has been studied for many decades, and the numerous models and techniques developed therein could aid the design of strategies for anonymous networking. For example, applying the result of Theorem 1, the Blahut-Arimoto algorithm provides an efficient iterative technique to obtain the optimal conditional distribution of covert relays $q(B|S)$ and the throughput function $\lambda(\alpha)$.

EXAMPLE


Consider the example of a switching network, as shown in Fig. 4a. During any network session, each source S_i sends packets to a distinct destination D_j , and for each pair S_i, D_j there is a fixed path through the intermediate relays. There are 24 possible sessions (source-destination pairings) that are assumed equally probable.

For this setup, Fig. 4 plots the throughput anonymity region and the fraction of covert relays when all nodes have identical transmission rate constraints. As seen in Fig. 4b, the throughput is a convex function of the anonymity level, which is a property of the rate-distortion function. Intuitively, this can be attributed to the average nature of the metrics equivocation and throughput. Note that the maximum throughput is achievable with non-zero anonymity. This is a virtue of transmitter directed signaling; encrypted headers would ensure that final destination nodes cannot be inferred by the eavesdropper. Figure 4c shows the relation between average number of covert relays and level of anonymity. As can be seen, it is not necessary to make all relays covert to obtain maximum anonymity (in this case it is sufficient to make relays M_2, M_4 covert for $\alpha = 1$).

DISCUSSION AND CONCLUDING REMARKS

The model for network sessions used in this article assumes independent and static observations by eavesdroppers. This may not apply to the scenario where an eavesdropper monitors the network for long periods of time. In that case one would need a stochastic model to account for session changes, depending on when nodes start or stop communication. In the results presented thus far we have considered the special case of strict delay constrained traffic and optimized network throughput. Imposing a delay constraint at each relay is restrictive in a multihop network, and the ideal metric to consider is the end-to-end delay. This throws up some interesting questions: for a fixed end-to-end delay constraint, how does one allocate delays at each (covert) relay in the route? If the throughput requirement was fixed, how does the optimal end-to-end delay vary with anonymity? Using the analytical approach presented in this article, we believe these questions can be answered in a fundamental theoretical way.

From a practical standpoint, the translation of the theoretical results to implementable solutions also requires a decentralized relay selection strategy. One approach to address decentralization is to investigate message passing for distributed decision systems, where in every session the nodes exchange the minimum amount of



■ **Figure 4.** a) Switching network: sources $\{S_i\}$ transmit packets to destinations $\{D_j\}$ through relays $\{M_i\}$; b) throughput vs. anonymity; c) fraction of covert relays vs. anonymity.

To the best of our knowledge, the proposed metric is the first analytical measure designed to quantify the secrecy of routes in an eavesdropped wireless network. The preliminary results obtained so far clearly demonstrate the potential for analytical methods to address the scheduling design.

information so that the decision to remain covert or visible can be made optimally. However, in network applications where message exchanges across nodes may not be possible, each node would only have partial information about the session. This is true of Mix networks where layered encryption ensures that each Mix only has knowledge of the neighboring nodes in the routes. For such applications, a decentralized implementation would involve each relay independently making a decision to be covert based on the available local information. The performance of such a decentralized approach was shown to be characterizable using a constrained distortion-rate optimization [12].

This article primarily deals with a fixed set of routes between source-destination pairs. In many situations, using a fixed set of routes may not provide sufficient anonymity, especially when the adversary obtains side information by compromising nodes. In such situations anonymity can be improved by randomizing the routes between source-destination pairs. One approach would be to design a collection of routes for every session such that revealing an unknown subset of links provides insufficient information about the session. Conceptually, this problem is similar to coding for wiretap II channels [13], where channel codes are designed so that revealing an unknown subset of bits does not provide sufficient information about the transmitted message.

In conclusion, the main contribution in this article is an analytical approach to anonymous wireless networking. To the best of our knowledge, the proposed metric is the first analytical measure designed to quantify the secrecy of routes in an eavesdropped wireless network. The preliminary results obtained so far clearly demonstrate the potential for analytical methods to address the scheduling design. Furthermore, our results also present connections to classical information theoretic problems such as wiretapped channel communication and rate-distortion, which, although well studied in literature, now present novel applications.

REFERENCES

- [1] D. Chaum, "Untraceable Electronic Mail, Return Addresses and Digital Pseudonyms," *Commun. ACM*, vol. 24, no. 2, Feb. 1981, pp. 84–88.
- [2] Y. Zhu et al., "On Flow Correlation Attacks and Countermeasures in Mix Networks," *Proc. Privacy Enhancing Technologies Wksp.*, May 2004, Toronto, Canada.
- [3] C. E. Shannon, "Communication Theory of Secrecy Systems," *Bell Sys. Tech. J.*, vol. 28, 1949, pp. 646–715.
- [4] A. Wyner, "The Wiretap Channel," *Bell Sys. Tech. J.*, vol. 54, 1975, pp. 1355–87.
- [5] I. Csiszár and J. Körner, "Broadcast Channels with Confidential Messages," *IEEE Trans. Info. Theory*, vol. 24, no. 3, May, 1978, pp. 339–48.
- [6] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, 1991.
- [7] B. Radosavljevic and B. Hajek, "Hiding Traffic Flow in Communication Networks," *Proc. IEEE MILCOM.*, San Diego, CA, Oct. 1992.
- [8] A. Serjantov, R. Dingleline and P. F. Syverson, "From a Trickle to a Flood: Active Attacks on Several MIX Types," *Proc. 5th Int'l. Wksp. Info. Hiding*, Noordwijkerhout, The Netherlands, Oct. 2002.
- [9] P. Venkatasubramanian, T. He, and L. Tong, "Relay Secrecy in Wireless Networks with Eavesdroppers," *Proc. 2006 Allerton Conf. Commun., Control and Comp.*, Monticello, IL, Sept. 2006.

- [10] A. Blum, D. Song, and S. Venkataraman, "Detection of Interactive Stepping Stones: Algorithms and Confidence Bounds," *Proc. 2004 Conf. Recent Advance in Intrusion Detection*, Sophia Antipolis, France, Sept. 2004.
- [11] P. Venkatasubramanian, T. He and L. Tong, "Anonymous Networking amidst Eavesdroppers," to appear, *IEEE Trans. Info. Theory*, Special Issue on Information-Theoretic Security, 2008, <http://arxiv.org/abs/0710.4903>.
- [12] P. Venkatasubramanian and L. Tong, "A Rate-Distortion Approach to Anonymous Networking," *Proc. 2007 Allerton Conf. Commun., Control and Comp.*, Monticello, IL, Sept. 2007.
- [13] L. H. Ozarow and A. D. Wyner, "Wiretap Channel II," *Bell Labs. Tech. J.*, vol. 63, no. 10, Dec. 1984, pp. 2135–46.

BIOGRAPHIES

PARVATHINATHAN VENKATASUBRAMANIAM [S'03, M'07] (pv45@cornell.edu) received his B.Tech. in electrical engineering from the Indian Institute of Technology, Madras, in 1998, and his M.S and Ph.D. in electrical engineering from Cornell University, Ithaca, New York. He is currently a visiting postdoctoral researcher at the University of California, Berkeley. His research interests include network security, signal processing in wireless networks, and multiple access communication protocols. He received the 2004 Leonard G. Abraham Award from the IEEE Communication Society and a Best Student Paper Award at the 2006 IEEE ICASSP.

TING HE [S'04, M'07] (th255@cornell.edu) is a research staff member in the Networking Technologies group at IBM Research, Hawthorne, New York. She received a Ph.D. degree from the School of Electrical and Computer Engineering, Cornell University, in 2007 and a B.S. degree in computer science from Peking University, China, in 2003. Her research focuses on solving inference problems in networking environments, including sensor networks, wired and wireless ad hoc networks, and mobile ad hoc networks.

LANG TONG [F'05] (ltong@ece.cornell.edu) received a B.E. degree from Tsinghua University, Beijing, China, in 1985, and a Ph.D. degree in electrical engineering from the University of Notre Dame, Indiana, in 1991. He was a postdoctoral research affiliate at the Information Systems Laboratory, Stanford University, California, in 1991. He joined Cornell University in 1998, where he is now the Irwin and Joan Jacobs Professor in Engineering. Prior to joining Cornell University, he was on faculty at West Virginia University, Morgantown, and the University of Connecticut, Storrs. He was also the 2001 Cor Wit Visiting Professor at Delft University of Technology, The Netherlands. He received the Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 best paper award (with M. Dong) from the IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society (with P. Venkatasubramanian and S. Adireddy). He also received the Young Investigator Award from the Office of Naval Research and is the coauthor of five papers that won student paper awards.

STEPHEN B. WICKER (wicker@ee.cornell.edu) received a B.S.E.E. with high honors from the University of Virginia in 1982. He received an M.S.E.E. from Purdue University in 1983 and a Ph.D. in electrical engineering from the University of Southern California in 1987. He is a professor of electrical and computer engineering at Cornell University. He is the author of *Codes, Graphs, and Iterative Decoding* (Kluwer, 2002), *Turbo Coding* (Kluwer, 1999), *Error Control Systems for Digital Communication and Storage* (Prentice Hall, 1995), and *Reed-Solomon Codes and Their Applications* (IEEE Press, 1994). He has served as Associate Editor for Coding Theory and Techniques for *IEEE Transactions on Communications*, and two terms as a member of the Board of Governors of the IEEE Information Theory Society. He is currently an Associate Editor for **ACM Transactions on Sensor Networks**. He teaches and conducts research in wireless information networks, digital systems, self-configuring networks, and game theory. His current research focuses on the development of highly distributed adaptive networks. He was awarded the 1988 Cornell College of Engineering Michael Tien Teaching Award and the 2000 Cornell School of Electrical and Computer Engineering Teaching Award.