A Game-Theoretic Approach to Anonymous Networking

Parv Venkitasubramaniam, *Member, IEEE*, and Lang Tong, *Fellow, IEEE*

***Abstract*—Anonymous wireless networking is studied when an adversary monitors the transmission timing of an unknown subset of the network nodes. For a desired quality of service (QoS), as measured by network throughput, the problem of maximizing anonymity is investigated from a game-theoretic perspective. Quantifying anonymity using conditional entropy of the routes given the adversary's observation, the problem of optimizing anonymity is posed as a two-player zero-sum game between the network designer and the adversary: The task of the adversary is to choose a subset of nodes to monitor so that anonymity of routes is minimum, whereas the task of the network designer is to maximize anonymity by choosing a subset of nodes to evade flow detection by generating independent transmission schedules. In this two-player game, it is shown that a unique saddle-point equilibrium exists for a general category of finite networks. At the saddle point, the strategy of the network designer is to ensure that any subset of nodes monitored by the adversary reveals an identical amount of information about the routes. For a specific class of parallel relay networks, the theory is applied to study the optimal performance tradeoffs and equilibrium strategies. In particular, when the nodes employ transmitter-directed signaling, the tradeoff between throughput and anonymity is characterized analytically as a function of the network parameters and the fraction of nodes monitored. The results are applied to study the relationships between anonymity, the fraction of monitored relays, and the fraction of hidden relays in large networks.**

***Index Terms*—Anonymity, eavesdropper, saddle-point equilibrium, traffic analysis, wireless networks.**

I. Introduction

A. Motivation

THE PACKET transmission times[1] of nodes in a network can reveal significant information about the source–destination pairs and routes of traffic flow in the network [1], [2]. Equipped with such information, a malicious adversary can launch more powerful attacks such as wormhole, jamming, or denial of service. Anonymous networking is the act of communicating over a network without revealing the identities of source–destinations or the path of flow of packets.

The typical design of anonymous networking protocols models adversaries as omniscient and capable of monitoring every single transmission in the network perfectly. From a practical standpoint, this is far too conservative, and such universal information would be available only to the network owner or a centralized controller. In this paper, our goal is to study the problem of anonymity in networks under a more general adversary model, where an *unknown* subset of the nodes is monitored by the adversary. The subset of monitored nodes could depend on the physical location of the adversary or partial knowledge of network transmission protocols. It is also possible that in some public wireless networks, certain nodes may have weaker physical protection than others and are hence more vulnerable to transmission monitoring.

From a network design perspective, the goal is to design transmission and relaying strategies such that the desired level of network performance is guaranteed with maximum *anonymity of network routes*. Providing anonymity to the routes of data flow in a network requires modification of packet transmission schedules and additional transmissions of dummy packets to confuse an external observer. These modifications, however, reduce the achievable network performance, particularly in ad hoc wireless networks, where the scheduling needs to satisfy medium access constraints on the shared channel. Therefore, depending on the desired quality of service (QoS), it is necessary to pick the optimal set of nodes to modify transmission schedules so that anonymity is maximized without violating QoS requirements.

If the network designer were aware of which nodes of the network were being monitored by the adversary, the optimal set of nodes can be chosen such that minimum information is revealed through the monitored nodes. However, if the adversary is aware of the set of nodes that the network designer has chosen to protect, then he can alter his choice of nodes to monitor so that maximum information about the network routes is retrieved. This "interplay" between the network designer and the adversary is the main subject of this paper, and it is studied using a game-theoretic approach.

Since the set of monitored nodes is unknown to the network designer, a conservative approach would be to design the scheduling strategy assuming an omniscient adversary. However, since the power of the adversary, i.e., the maximum fraction of monitored nodes, is bounded, we would like to investigate if the strategies of the network designer and the adversary can be analyzed jointly to get a better tradeoff between anonymity and network performance compared to that under the omniscient assumption (see Fig. 1). To this end, we propose

Manuscript received February 09, 2010; revised November 22, 2010 and May 02, 2011; accepted September 01, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Sarkar. Date of publication January 12, 2012; date of current version June 12, 2012. This work was supported in part by the National Science Foundation under Awards CCF-0728872 and CNS-1117701 and the Army Research Office under MURI Awards W911NF-08-1-0238 and W911NF-10-1-0419.

P. Venkitasubramaniam is with the School of Electrical and Computer Engineering, Lehigh University, Bethlehem, PA 18015 USA (e-mail: parv.v@lehigh.edu).

L. Tong is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2011.2176511

[1]Transmission time in this paper refers to the time point of transmission, not the duration or latency.

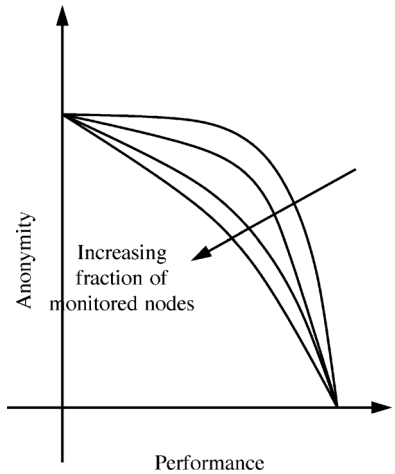

Fig. 1. Anonymity–performance tradeoff. As the fraction of monitored nodes gets smaller, we wish to compute the improvement in the performance tradeoffs.

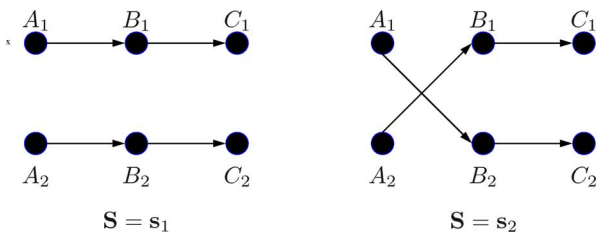

Fig. 2. 2-relay parallel network: two possible sessions, each containing two paths. $\mathbf{s}_1 = \{(A_1, B_1, C_1), (A_2, B_2, C_2)\}, \mathbf{s}_2 = \{(A_1, B_2, C_2), (A_2, B_1, C_1)\}$.

a two-player zero-sum game between the adversary and the network designer, where the payoff is anonymity, the action of the adversary is to choose which nodes to monitor to minimize payoff, and the action of the network designer is to choose which nodes of the network to "hide" from the adversary to maximize the payoff subject to the QoS constraint.

The game-theoretic perspective can be understood using an example of a 2-relay parallel network as shown in Fig. 2. During any period of observation of the adversary, we assume that the network operates in one of two configurations $\mathbf{s}_1$ or $\mathbf{s}_2$ (see Fig. 2) wherein

$$\mathbf{s}_1 = \{(A_1, B_1, C_1), (A_2, B_2, C_2)\},$$
$$\mathbf{s}_2 = \{(A_1, B_2, C_2), (A_2, B_1, C_1)\}$$

are the set of active routes in each configuration (henceforth referred to as a *network session*). The adversary's goal is to identify which of these sessions is currently active in the network by monitoring the transmission timing of the monitored nodes.

Consider a transmitter-directed signaling model, where each node transmits on a unique orthogonal channel such that transmissions of multiple nodes are noninterfering. Under this signaling scheme, merely detecting the transmission times of packets by a node will not reveal the identity of the intended receiver. Suppose in this setup, the adversary can only afford to monitor the transmissions of two nodes. An adversary would therefore have to detect correlations across transmission schedules of a source and a relay to identify the flow of traffic. For example, if B_1 forwarded packets as and when they arrived from the source, then during network session $\mathbf{s}_1$, the transmission schedules of A_1 and B_1 would be highly correlated, whereas during $\mathbf{s}_2$, the schedules of A_1 and B_1 would be statistically independent. An adversary who merely monitors nodes A_1 and B_1 would therefore be able to identify the network session perfectly by detecting the dependence between schedules. Suppose, instead, the relays B_1 and B_2 always use transmission schedules that are statistically independent of the arrival schedules from the sources. Then, no information about the session can be obtained by monitoring the transmission schedules of any pair of nodes. Using independent schedules, however, requires dummy transmissions by the relays, thus reducing the rate of data packets forwarded by each relay.

Consider a scenario when the throughput requirement mandates that at most one relay can generate independent schedules (using dummy transmissions). If only relay B_1 generates a transmission schedule that is statistically independent of that of A_1 and A_2, then the optimal strategy for the adversary would be to monitor (A_2, B_2) or (A_1, B_2), either of which would help him perfectly determine the session. However, given the knowledge that the adversary would monitor (A_1, B_2) or (A_2, B_2), the optimal strategy of the network designer would be to make the schedule of B_2 always independent thus maximizing anonymity.

A natural question that arises is the following: Is there a pair of strategies for the network designer and the adversary that neither has any incentive to modify? In other words, if formulated as a two-player zero-sum game between the adversary and the network designer with anonymity as the payoff, does a Nash equilibrium exist? As will be shown in Section III, a saddle-point equilibrium does exist in the class of mixed strategies. For this example, at the equilibrium point, the optimal strategy for the network designer is to choose one of the relays with probability $\frac{1}{2}$ to generate independent schedules, and the optimal strategy for the adversary is to monitor each source–relay pair with probability $\frac{1}{4}$. By definition, at this operating point, neither the network designer nor the adversary have any incentive to modify their strategies (see Theorem 3).

The example discussed above involves a simple scenario with only two possible network sessions, and the adversary has two kinds of observations: a pair of dependent or a pair of independent schedules. In a general multihop network, anonymity based on partial information about the session can be quantified using Shannon's equivocation [3], [4], and our goal in this work is to optimize the tradeoff between the desired network throughput and the achievable anonymity as a function of the adversary's monitoring capability.

B. Main Contributions

In this paper, we consider a game-theoretic formulation of anonymous networking in a general class of finite wireless networks when the number of nodes monitored by an adversary is bounded by a known constant. We pose the design problem as a two-player zero-sum game with equivocation (conditional entropy) of the network session as the payoff. The adversary's strategy is to pick a random subset of nodes to monitor, and the network designer's strategy is to pick a random subset of nodes to generate independent schedules, thus avoiding detection. For the class of finite multihop networks considered, we prove that

a saddle-point equilibrium always exists in the class of centralized strategies [2]. Note that since anonymity, as defined by conditional entropy, is a nonlinear function of the probabilities of mixing multiple strategies, the existence of Nash equilibria in classical two-player zero-sum games [5], where payoff of mixed strategies is the weighted sum of pure strategy payoffs, does not directly apply.

To demonstrate the applicability of the game-theoretic model, we consider a general class of parallel relay networks. For a symmetric relay model, we characterize analytically the throughput–anonymity tradeoff as a function of the adversary's power and, using the results on player strategies, derive the saddle-point strategies that are understandably symmetric. We then introduce asymmetry into the properties of the relay rate and the information model and, using the derived results on saddle-point strategies, demonstrate the gain of the game-theoretic approach over naive intuitive strategies. We also show that the game-theoretic approach can be used to study large parallel relay networks by characterizing the asymptotic relationships among anonymity, the fraction of monitored relays, and the fraction of covert relays.

C. Related Work

Anonymous communication over the Internet is fairly well studied, where many applications have been designed based on the concept of traffic mixes proposed by Chaum [6]. Mixes are routers or proxy servers that collect packets from multiple users and transmit them after reencryption and random delays so that incoming and outgoing packets cannot be matched by an external observer. While mix-based solutions have been used in applications such as anonymous e-mail or browsing, it has been shown that when long streams of packets with latency or buffer constraints are forwarded through mixes, it is possible to correlate incoming and outgoing streams almost perfectly [7].

In wireless networks, an alternative solution to mixing is the use of cover traffic [8], [9], which ensures that, irrespective of the active routes, the transmission schedules of all nodes are fixed *a priori*. If a node does not have any data packets, the transmission schedule is maintained by transmitting dummy packets. While the fixed scheduling strategy provides complete anonymity to the routes at all times, it was found to be inefficient [8] due to high rate of dummy transmissions, and the implementation required synchronization across all nodes, which is not practical in ad hoc wireless networks. In this paper, the technique used to provide anonymity is similar to that in [10], where a subset of relays (referred to as *covert relays*) generates independent transmission schedules using dummy transmissions.

The general adversary model considered here necessitates a game-theoretic formulation of the problem. Game theory [11] has been used in a wide range of multiagent problems from economics to networking. In the context of network security, earlier applications were focused on jamming. Basar considered the problem of jamming in Gaussian channels [12], where it was shown that the optimal jamming strategy is either a linear function of jammer's observation or an additive independent Gaussian noise. Borden *et al.* [13] considered the information-theoretic saddle points of the jamming game under hard/soft quantization schemes. More recent work along this line includes [14]–[16]. Game-theoretic models have also been used to model problems related to distributed intrusion detection [17], [18], where the goal is to design attacking and detection strategies with probability of detection as the payoff. In [19], game theory was used to study attacker and defense strategies on a graphical model of a network, where the attackers choose nodes to compromise, while the defender picks links to "clean up." To the best of our knowledge, ours is the first application of game theory to hide traffic flows in the presence of eavesdroppers. The work closest to ours in this regard is that of information concealing games using finite-dimensional data [20] where one of the players (the adversary) chooses a subset of available resources to hide, while the opponent (the network user) chooses a subset of resources based on the revealed observation to maximize his utility. The authors identify conditions under which Nash equilibria exist and provide approximation techniques to compute the equilibria. Conceptually, this problem has some similarities to our strategy of choosing covert relays, where the network designer chooses to hide a subset of relays, whereas the adversary chooses a subset of relays to monitor. In our model, the adversary's observation depends on the actions of both the players, which are decided *a priori*, and the payoff is a nonlinear function of the probabilities of mixing strategies, thus different from classical mixed strategy models [5].

Our mathematical model for anonymity is based on the framework proposed in [10], where conditional entropy of the network session was proposed as a metric for anonymity. Entropy and measures related to entropy such as K-L divergence have been proposed as payoffs in games of complexity [21]. Entropy in such contexts were however used as metrics of complexity rather than a measure of uncertainty.

II. System Model

Notation: Let the network be represented by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes in the network and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of directed links. (A, B) is an element of $\mathcal{E}$ if and only if node B can receive transmissions from node A. A sequence of nodes $P = (V_1, \ldots, V_n)$ is a *valid path* in $\mathcal{G}$ if $(V_i, V_{i+1}) \in \mathcal{E}, \forall i < n$. The set of all loopless paths is denoted by $\mathcal{P}(\mathcal{G})$.

A. Adversary Observation and Inference

During any network observation by the adversary, a subset of nodes communicate using a fixed set of paths. This set of paths $\mathbf{S} \in 2^{\mathcal{P}(\mathcal{G})}$ is referred to as a network *session*. The adversary's goal is to use his observation to identify the session. We model $\mathbf{S}$ as an i.i.d. random variable $\mathbf{S} \sim p(\mathbf{s})$. The prior $p(\mathbf{s})$ on sessions is assumed to be available to the adversary. The set of possible sessions $\mathcal{S}$ is given by $\mathcal{S} = \{\mathbf{s} \in \mathcal{P}(\mathcal{G}) : p(\mathbf{s}) > 0\}$. (See example sessions in Fig. 2.)

Transmitter-Directed Signaling: The adversary's observation would depend on the underlying physical-layer signaling model. In this paper, we consider orthogonal transmitter-directed signaling at the physical layer, where each node utilizes a unique orthogonal signaling scheme such that a transmission schedule

[2]Centralized strategies are strategies that require coordinated action across all nodes of the network. Such strategies can be implemented using a single central controller, the use of shared randomness across nodes, or limited message passing between nodes.

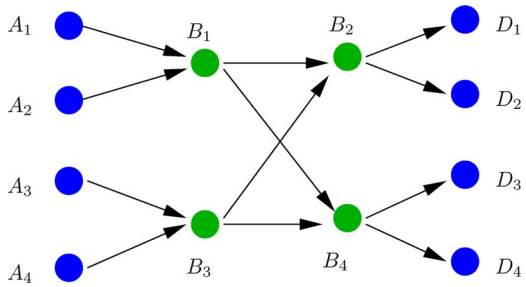

Fig. 3. Switching network. $\{A_i\}$ transmit to $\{D_i\}$ through relays $\{B_i\}$.

detected by the adversary would reveal only the transmitting node and not the intended receiving node.

Observable Session: The goal of the network designer is to modify transmission schedules of the nodes in every session such that the monitored nodes reveal as little information about the actual session as possible. For instance, if a subset of relays always generates independent transmission schedules, then it is not possible for the adversary to determine which paths pass through them. In effect, the set of (broken) paths observable would be a distorted version of the actual session. Let $\hat{\mathbf{S}}$ (henceforth referred to as *observable session*) denote the set of paths as observed by an omniscient adversary.

For example, consider the switching network in Fig. 3, where every session is defined by a unique pairing of sources and destinations (each A_i sends packets to a unique D_j through intermediate relays). In this network, consider a session $\mathbf{s}_1$ given by the set of paths

$$\mathbf{s}_1 = \left\{ \begin{array}{l} (A_1, B_1, B_4, D_3), (A_2, B_1, B_2, D_2), \\ (A_3, B_3, B_2, D_1), (A_4, B_3, B_4, D_4) \end{array} \right\}.$$

Suppose node B_1 generated an independent schedule regardless of the arrival times from A_1, A_2. Then, an omniscient adversary would not be able to identify the paths of the packet streams from A_1 and A_2 after they reach B_1. Therefore, the observable session would contain the set of paths

$$\hat{\mathbf{S}} = \left\{ \begin{array}{l} A_1, A_2, (B_1, B_4, D_3), (B_1, B_2, D_2), \\ (A_3, B_3, B_2, D_1), (A_4, B_3, B_4, D_4) \end{array} \right\}. \quad (1)$$

Adversary Observation: Under a general adversary model, packet transmission times of a subset of nodes are observed by the adversary. Specifically, the adversary randomly chooses any subset of nodes, denoted by $\mathbf{N}_\mathrm{a}$, to monitor. The maximum number of monitored nodes is denoted by k_a (also referred to as *power of the adversary*). We model $\mathbf{N}_\mathrm{a}$ as a random variable where the random distribution of $\mathbf{N}_\mathrm{a}$ is chosen by the adversary to maximize his payoff. Depending on the observable session $\hat{\mathbf{S}}$ and the set of monitored nodes $\mathbf{N}_\mathrm{a}$, the adversary's observation $\hat{\mathbf{S}}_\mathrm{a}$ would be a further distorted version of the underlying session $\mathbf{S}$. The adversary's net observation can be represented by a set of paths $\hat{\mathbf{S}}_\mathrm{a}$ and would be given by a deterministic function $f_\mathrm{a}(\hat{\mathbf{S}}, \mathbf{N}_\mathrm{a})$. (Note that $f_\mathrm{a}(\hat{\mathbf{S}}, \mathcal{V}) = \hat{\mathbf{S}}$.)

In the switching network example of Fig. 3, let B_1 be covert in session $\mathbf{s}_1$. Then, (1) provides the observable session (omniscient adversary). If the adversary monitors nodes A_1, A_3, B_1, and B_3, then

$$\hat{\mathbf{S}}_\mathrm{a} = \{A_1, B_1, (A_3, B_3)\}.$$

B. Performance Metrics: Anonymity and Throughput

The task of the network designer is to design the probability distribution of observable sessions, denoted by $q_\mathrm{n}(\hat{\mathbf{s}} \,|\, \mathbf{s})$, such that a desired QoS is achieved while the adversary obtains minimum information about the session $\mathbf{S}$ by observing $\hat{\mathbf{S}}_\mathrm{a}$. The task of the adversary, on the other hand, is to design the probabilities $q_\mathrm{a}(\mathbf{N}_\mathrm{a})$ of choosing nodes to monitor s.t. maximum information is obtained by observing $\hat{\mathbf{S}}_\mathrm{a}$.

Anonymity: We quantify anonymity using Shannon's equivocation [3], which measures the uncertainty of the underlying session given the adversary's observation.

Definition 1: We define the *anonymity* $A(q_\mathrm{n}, q_\mathrm{a})$ for a network strategy $q_\mathrm{n}(\hat{\mathbf{s}} \,|\, \mathbf{s})$ w.r.t. adversary strategy $q_\mathrm{a}(\mathbf{n}_\mathrm{a})$ as the normalized conditional entropy of the sessions given the adversary observation

$$A(q_\mathrm{n}, q_\mathrm{a}) \triangleq \frac{H(\mathbf{S} \,|\, \hat{\mathbf{S}}_\mathrm{a})}{H(\mathbf{S})}. \quad (2)$$

The normalization ensures that the anonymity is always between 0 and 1. The motivation behind the above definition comes from Fano's inequality, which lower-bounds the adversary's probability of error by the conditional entropy [22]. Note that previous entropy-based definitions of anonymity [4], [10] in the context of omniscient adversaries are special cases of Definition 1 (when $\mathbf{N}_\mathrm{a} \equiv \mathcal{V}$).

Throughput: Since distorting the observable session requires modification of transmission schedules, the latency and bandwidth constraints in the network would require transmission of dummy packets and result in a reduced rate of data packets delivered from the sources to destinations. Let $\Lambda(\mathbf{s}, \hat{\mathbf{s}})$ represent the sum-rate of packets deliverable from sources to destinations when the actual session is $\mathbf{s}$ and the observable session is $\hat{\mathbf{s}}$. Note that $\Lambda(\mathbf{s}, \hat{\mathbf{s}}) \leq \Lambda(\mathbf{s}, \mathbf{s})$.

Definition 2: The *throughput* $\Upsilon(q_\mathrm{n})$ of a scheduling strategy $q_\mathrm{n}(\hat{\mathbf{S}} \,|\, \mathbf{S})$ is defined as

$$\Upsilon(q_\mathrm{n}) = \mathbb{E}(\Lambda(\mathbf{S}, \hat{\mathbf{S}})) \quad (3)$$

where the expectation is over the joint probability density function (pdf) of $\mathbf{S}$ and $\hat{\mathbf{S}}$.

Anonymity and throughput are essentially two opposing paradigms in the design of the optimal scheduling strategy: Transmitting more dummy packets increases anonymity, whereas higher throughput necessitates fewer dummy transmissions. Unlike the omniscient adversary setup, since the power of the adversary is bounded, the uncertainty in the identities of the monitored nodes, i.e., the randomness in $\mathbf{N}_\mathrm{a}$, necessitates the game-theoretic formulation, as was illustrated in the example in Section I. In Section III, we formulate this problem as a two-player zero-sum game and establish the existence of a saddle-point equilibrium.

III. Two-Player Game Using Covert Relaying Strategy

Consider a two-player zero-sum game $\mathbb{G}_\mathrm{a}$, defined by a 3-tuple $(\mathcal{A}_\mathrm{n}, \mathcal{A}_\mathrm{a}, \phi)$, where $\mathcal{A}_\mathrm{n}$ and $\mathcal{A}_\mathrm{a}$ denote the action spaces of the network designer and the adversary, respectively, and $\phi : \mathcal{A}_\mathrm{n} \times \mathcal{A}_\mathrm{a} \mapsto [0, 1]$ is the payoff function for the network designer (the adversary's payoff is $-\phi(\cdot, \cdot)$).

A. Action Spaces

In its most general form, the action space for the network designer would include the set of all probability distributions $q_{\rm n}(\hat{\mathbf{S}} \,|\, \mathbf{S})$ over the space of all loopless paths $\mathcal{P}$. In this paper, we restrict the set of observable sessions to those achievable using the set of *covert relaying strategies*, where each relay node belongs to one of two categories: *covert* or *visible*.

Covert Relay: A covert relay B generates an outgoing transmission schedule that is statistically independent of the schedules of all nodes occurring previously in paths that contain B. Due to statistical independence, no adversary can detect the flow of traffic through a covert relay. Covert relaying is a modification to the transmission scheduling that provides anonymity and yet adheres to the medium access and delay constraints of the system.

Visible Relay: A visible relay B transmits every received packet immediately upon arrival, thereby ensuring all arriving packets are relayed successfully within the latency constraint. However, the traffic flow through the visible relay operating under this highly correlated schedule is easily detected by an eavesdropper. (A statistically consistent detector for this purpose has been designed in [23].)

In a given session $\mathbf{s}$, if the set of covert relays is $\mathbf{b}_{\rm n}$, then the observable session $\hat{\mathbf{s}}$ can be expressed as a deterministic function $f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n})$. For a transmitter-directed signaling model, $f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n})$ is a set of paths such that for every path in $\mathbf{s}$ that has k covert relays, $f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n})$ contains $k+1$ paths, each beginning at the source or a covert relay and terminating one relay before the subsequent covert relay or the destination. This is because covert relay schedules prevent the adversary from detecting any correlation between the schedule of a prior node in the path with that of the relay.

We model the set of covert relays in a session by a random variable $\mathbf{B}_{\rm n}$ with conditional distribution $\{q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s})\}$, and the class of covert relaying strategies is defined by the set of all probability distributions $\{q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s})\}$. Note that this is a restrictive action space where it may not be possible to realize all observable sessions in $2^{\mathcal{P}(\mathcal{G})}$ for any session $\mathbf{s}$.

As expected, maintaining independent schedules would require covert relays to drop packets or add dummy packets consequently reducing the rate of relayed data packets, whereas visible relays can relay every packet that arrives without any loss in rate. The loss in rate at a covert relay would be a function of the probability distributions of transmission schedules, delay and bandwidth constraints, and the relaying strategy. In a session $\mathbf{s}$, let $\Lambda'(\mathbf{s}, \mathbf{b}_{\rm n})$ denote the achievable sum-rate when the relays in the set $\mathbf{b}$ are covert. Note that since $\mathbf{s}, \mathbf{b}$ perfectly determine the observable session $\hat{\mathbf{s}}$

$$\Lambda'(\mathbf{s}, \mathbf{b}_{\rm n}) = \Lambda(\mathbf{s}, f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n})).$$

The characterization of the exact rate loss is not necessary for this exposition, and we will treat it as an abstract quantity. In the subsequent section, where we study parallel relay networks, we shall use specific scheduling and relaying strategies and provide an analytical characterization of the rate loss for that class of networks.

For a given strategy $q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s})$, the throughput Υ can be expressed as a linear function

$$\Upsilon(q_{\rm n}) = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}) \sum_{\mathbf{b} \in 2^{\mathcal{V}}} q_{\rm n}(\mathbf{b} \,|\, \mathbf{s}) \Lambda'(\mathbf{s}, \mathbf{b}).$$

By restricting ourselves to the class of covert relaying strategies, we define the action spaces for the network designer and the adversary in the game as follows.

The action of the network designer is to select the probability mass function $q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s})$ that chooses covert relays in each session $\mathbf{s}$. The key constraint in this design is the throughput requirement $(\Upsilon(q_{\rm n}) \geq \gamma)$. Accordingly

$$\mathcal{A}_{\rm n} = \begin{cases} \{q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s}) : \mathbf{s} \in \mathcal{S}, \mathbf{b}_{\rm n} \subset \mathcal{V}\} : \\ \Upsilon(q_{\rm n}) \geq \gamma \\ q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s}) \geq 0 \qquad \forall \mathbf{s}, \mathbf{b}_{\rm n} \\ \sum_{\mathbf{b}_{\rm n}} q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s}) = 1 \qquad \forall \mathbf{s}. \end{cases}$$

The action of the adversary is to design the probability distribution $q_{\rm a}(\mathbf{n}_{\rm a})$ of picking nodes to monitor during the session, subject to the constraint on the maximum number of monitored nodes $(\mathbf{n}_{\rm a} \in \mathcal{V}^{k_{\rm a}})$. Therefore

$$\mathcal{A}_{\rm a} = \begin{cases} \{q_{\rm a}(\mathbf{n}_{\rm a}) : \mathbf{n}_{\rm a} \in \mathcal{V}^{k_{\rm a}}\} \\ q_{\rm a}(\mathbf{n}_{\rm a}) \geq 0 \qquad \forall \mathbf{n}_{\rm a} \\ \sum_{\mathbf{n}_{\rm a}} q_{\rm a}(\mathbf{n}_{\rm a}) = 1. \end{cases}$$

B. Payoff and Saddle Point

For a given observable session $\hat{\mathbf{s}} = f_{\rm o}(\mathbf{s}, \mathbf{b})$, the adversary observation $\hat{\mathbf{s}}_{\rm a}$ would be restricted to the paths between monitored nodes in $\mathbf{n}_{\rm a}$. In other words

$$\hat{\mathbf{s}}_{\rm a} = f_{\rm a}(\hat{\mathbf{s}}, \mathbf{n}_{\rm a}) \stackrel{\triangle}{=} \{\mathbf{p} \bigcap \mathbf{n}_{\rm a} : \mathbf{p} \in \hat{\mathbf{s}}\}.$$

Define $\mathcal{F}_{\rm a} : 2^{\mathcal{P}(\mathcal{G})} \times 2^{\mathcal{V}} \mapsto 2^{\mathcal{S} \times 2^{\mathcal{V}}}$ to be the adversary's uncertainty set

$$\mathcal{F}_{\rm a}(\hat{\mathbf{s}}_{\rm a}, \mathbf{n}_{\rm a}) = \{(\mathbf{s}, \mathbf{b}) : f_{\rm a}(f_{\rm o}(\mathbf{s}, \mathbf{b}), \mathbf{n}_{\rm a}) = \hat{\mathbf{s}}_{\rm a}\}.$$

In other words, if the adversary monitors $\mathbf{n}_{\rm a}$, $\mathcal{F}_{\rm a}(\mathbf{p}, \mathbf{n}_{\rm a})$ is the set of possible pairs of session and covert relays that would lead to the observation $\mathbf{p}$ through the nodes $\mathbf{n}_{\rm a}$.

For a given pair of strategies $(q_{\rm n}, q_{\rm a}) \in (\mathcal{A}_{\rm n} \times \mathcal{A}_{\rm a})$, the payoff function $\phi(q_{\rm n}, q_{\rm a})$ is the anonymity which from Definition 1 is given by

$$\begin{aligned} \phi(q_{\rm n}, q_{\rm a}) &= \frac{H(\mathbf{S} \,|\, \hat{\mathbf{S}}_{\rm a})}{H(\mathbf{S})} \\ &= \frac{1}{H(\mathbf{S})} \sum_{\mathbf{n}_{\rm a} \in 2^{\mathcal{V}}} \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{b}_{\rm n} \in 2^{\mathcal{V}}} -q_{\rm a}(\mathbf{n}_{\rm a}) p(\mathbf{s}) \\ &\quad \times q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s}) \log q_{\rm ap}(\mathbf{s}, f_{\rm a}(f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n}), \mathbf{n}_{\rm a}), \mathbf{b}_{\rm a}) \end{aligned} \tag{4}$$

where

$$q_{\rm ap}(\mathbf{s}, \hat{\mathbf{s}}_{\rm a}, \mathbf{n}_{\rm a}) \stackrel{\triangle}{=} \frac{\sum_{\mathbf{b}: f_{\rm a}(f_{\rm o}(\mathbf{s}, \mathbf{b}), \mathbf{n}_{\rm a}) = \hat{\mathbf{s}}_{\rm a}} q_{\rm n}(\mathbf{b} \,|\, \mathbf{s}) p(\mathbf{s})}{\sum_{(\mathbf{s}', \mathbf{b}') \in \mathcal{F}_{\rm a}(\hat{\mathbf{s}}_{\rm a}, \mathbf{b})} q_{\rm n}(\mathbf{b}' \,|\, \mathbf{s}') p(\mathbf{s}')} \tag{5}$$

is the *aposterior* probability that the current session is $\mathbf{s}$ given the adversary observes $\hat{\mathbf{s}}_{\rm a}$ through the nodes $\mathbf{n}_{\rm a}$.

In a zero-sum game, we know that the interests of the network designer and the adversary are exactly opposite: While the network designer would prefer to make the monitored nodes covert, the adversary would prefer to monitor the visible nodes. We wish to determine if there is an operating point in the pair of action spaces, where neither the network nor the adversary has

any incentive to change their strategy—in other words, if this game has a saddle-point equilibrium.

Definition 3: A pair of strategies $(q_{\rm n}, q_{\rm a}) \in \mathcal{A}_{\rm n} \times \mathcal{A}_{\rm a}$ constitutes a *saddle-point equilibrium* if

$$\phi(q_{\rm n}, q_{\rm a}) = \sup_{q \in \mathcal{A}_{\rm n}} \phi(q, q_{\rm a}) = \inf_{q \in \mathcal{A}_{\rm a}} \phi(q_{\rm n}, q). \tag{6}$$

Note that, although it is well known that two-player zero-sum standard matrix games as defined in [5] always have a Nash equilibrium in the class of mixed strategies, the result does not extend to the game defined here. In fact, even if modeled as a continuous-kernel game as in [24], the existence of saddle-point equilibrium when action spaces are compact does not directly apply here. The reason being is that the payoff for a mixed strategy in such two-player games is a weighted sum of pure strategy payoffs; in our setup, the payoff is a nonlinear function of the pure strategy payoffs and the mixing probabilities [see (4)]. The existence of a saddle point in this game is shown in the following theorem.

Theorem 1: For the two-player zero-sum game $\mathbb{G}_{\rm a}$ defined by the action spaces $\mathcal{A}_{\rm n}, \mathcal{A}_{\rm a}$ and payoff function ϕ, there exists a saddle-point equilibrium.

Proof: Refer to the Appendix. □

The equilibrium condition guarantees that at the operating point, the adversary can use no other strategy to decrease the anonymity of the session. In addition to proving the existence of a saddle point, characterizing the optimal strategy for the adversary is also important, and particularly helpful in network scenarios where additional protection can be provided to nodes that are more likely to be monitored.

Note that the omniscient adversary setup is a specific instance of this game, when the adversary has exactly one action: monitor all nodes. The existence of an equilibrium is trivial, and the operating point is given by the rate distortion optimization [4]

$$\phi(\gamma) = H(\mathbf{S}) - \inf_{q_{\rm n}(\hat{\mathbf{S}} \,|\, \mathbf{S}) : \Upsilon(q_{\rm n}) \le \gamma} I(\mathbf{S}; \hat{\mathbf{S}}). \tag{7}$$

The uniqueness of the equilibrium follows from the zero-sum property of the game. Note that while the pair of strategies that achieves the saddle-point anonymity is not unique, the saddle-point anonymity in the two-player zero-sum game is indeed unique. This game is also an example of an incomplete information game [18], where the adversary does not have complete access to the session or the realization of the network designer's randomness, while the network designer does not have access to the realization of the adversary's randomness.

Although computing saddle-point strategies is hard since the action spaces are continuous, properties of player strategies can be derived by studying the conditions.

C. Insights Into Player Strategies

In this section, we investigate the properties of the saddle-point player strategies using the conditions for equilibrium.

Partial Information: For a given subset of nodes $\mathbf{b}$, we define the partial uncertainty from the adversary's perspective as

$$H_{\rm p}(\mathbf{b}) = \sum_{\mathbf{s}, \hat{\mathbf{s}}} p(\mathbf{s}) q_{\rm n}(\mathbf{b}_{\rm n} \,|\, \mathbf{s}) \log q_{\rm ap}(\mathbf{s}, f_{\rm a}(f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n}), \mathbf{n}_{\rm a}), \mathbf{b})$$

where $q_{\rm ap}$ is the *aposterior* probability defined in (5). The partial uncertainty represents the uncertainty of the session when the adversary monitors a particular subset of nodes.

Information Leakage Rate: For a given action by the network designer—making a set of relays $\mathbf{b}$ covert in a session $\mathbf{s}$—the rate of information leakage is defined as

$$\mathcal{L}(\mathbf{s}, \mathbf{b}) \triangleq \frac{d\phi(q_{\rm n}, q_{\rm a})}{dq_{\rm n}(\mathbf{s}, \mathbf{b})}. \tag{8}$$

Theorem 2: For the two-player zero-sum game $\mathbb{G}_{\rm a}$, at the saddle point $(q_{\rm n}^*, q_{\rm a}^*)$:

1) $\forall \mathbf{b}_{\rm a}^1, \mathbf{b}_{\rm a}^2$ s.t. $q_{\rm a}^*(\mathbf{b}_{\rm a}^1), q_{\rm a}^*(\mathbf{b}_{\rm a}^2) > 0$

$$H_{\rm p}\left(\mathbf{b}_{\rm a}^1\right) = H_{\rm p}\left(\mathbf{b}_{\rm a}^2\right).$$

2) $\forall \mathbf{s}$, if $\exists \mathbf{b}_1, \mathbf{b}_2$, s.t. $q_{\rm n}^*(\mathbf{s}, \mathbf{b}_1), q_{\rm n}^*(\mathbf{s}, \mathbf{b}_2) > 0$ and $\Lambda(\mathbf{s}, \mathbf{b}_1) = \Lambda(\mathbf{s}, \mathbf{b}_2)$, then

$$\mathcal{L}(\mathbf{s}, \mathbf{b}_1) = \mathcal{L}(\mathbf{s}, \mathbf{b}_2). \tag{9}$$

3) $\forall \mathbf{s}$, if $\exists \mathbf{b}_1, \mathbf{b}_2$, s.t. $q_{\rm n}^*(\mathbf{s}, \mathbf{b}_1), q_{\rm n}^*(\mathbf{s}, \mathbf{b}_2) > 0$ and $\Lambda(\mathbf{s}, \mathbf{b}_1) \neq \Lambda(\mathbf{s}, \mathbf{b}_2)$, then

$$\frac{\mathcal{L}(\mathbf{s}, \mathbf{b}_1) - \mathcal{L}(\mathbf{s}, \mathbf{b}_2)}{\Lambda(\mathbf{s}, \mathbf{b}_1) - \Lambda(\mathbf{s}, \mathbf{b}_2)} \text{ is a constant.} \tag{10}$$

Proof: Refer to the Appendix.

The above theorem states that, at the saddle point, for every subset of nodes monitored by the adversary (with nonzero probability), the partial uncertainty of the underlying session is identical. In other words, the set of covert relays is chosen such that any monitored subset reveals equal information about the session. At this operating point, from the perspective of the adversary, any probability distribution over these "degenerate" subsets would give rise to the same anonymity. There, however, exists a unique distribution to choose nodes to monitor, which, when employed, gives the network designer no incentive to deviate. At this point, the difference in information leakage rates for any pair of actions is proportional to the difference in throughput. In other words, the throughput cost per unit change in uncertainty is identical for every choice of covert relays (by the network designer).

Although the conditions in (9) and (10) appear complicated to analyze owing to *aposterior* probabilities, in many networks it is possible to utilize network structure and session models to analyze the condition and characterize the optimal throughput–anonymity tradeoffs.

In Section IV, we consider one such class of *parallel relay networks* to demonstrate the applicability of the game-theoretic approach. Specifically, we use the derived results on saddle-point strategies to study the optimal behavior of network nodes and the adversary and, in the process, demonstrate the performance improvement due to the game-theoretic approach over naive intuitive player strategies. We also apply the formulation to characterize fundamental asymptotic relationships between anonymity, throughput, and adversary capability in parallel relay networks. The asymptotic relationships are useful in the design of strategies in large networks where numerical computation becomes practically infeasible. In fact, we demonstrate

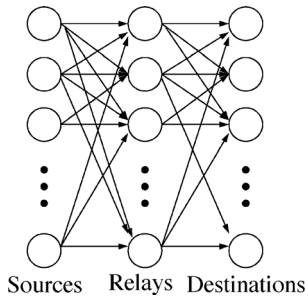

Fig. 4. Parallel relay network model.

that the maximum loss in using the asymptotic results on an n-node parallel relay network is bounded by $\frac{\log n}{n}$.

IV. Parallel Relay Networks

A. Network Model

Consider a *parallel relay network* as shown in Fig. 4, where the set of nodes $\mathcal{V}$ in the network can be divided into three subsets $\mathcal{V}^{\mathrm{s}}, \mathcal{V}^{\mathrm{r}}, \mathcal{V}^{\mathrm{d}}$ such that $\mathcal{V}^{\mathrm{s}} = \{A_1, \ldots, A_n\}$ is the set of source nodes, $\mathcal{V}^{\mathrm{d}} = \{D_1, \ldots, D_n\}$ is the set of destination nodes, and $\mathcal{V}^{\mathrm{r}} = \{R_1, \ldots, R_n\}$ is the set of intermediate relay nodes the network. The set of edges $\mathcal{E}$ can similarly be divided into two sets $\mathcal{E}_{\mathrm{s}}, \mathcal{E}_{\mathrm{r}}$, where $\mathcal{E}_{\mathrm{s}}$ denotes the set of edges between source nodes and relays, and $\mathcal{E}_{\mathrm{r}}$ is the set of edges between relays and the destinations.

We make the following two assumptions in the model.

1) *Full connectivity:* Every source is connected to every relay, and every relay is connected to every destination.
2) *Symmetry:* The probability of a source–relay–destination connection is identical across sources, relays, or destinations.

Note that these assumptions, while not critical to the analytical tractability, help to provide broader insights into optimal strategies for the network designer and the adversary.

Session Model: In each session, every source in $\mathcal{V}^{\mathrm{s}}$ picks a distinct destination in $\mathcal{V}^{\mathrm{d}}$ and a distinct intermediate relay in $\mathcal{V}^{\mathrm{r}}$ to forward its packets, such that all sources and relays are transmitting in every session. From a graph-theoretic perspective, each session corresponds to a unique pair of bipartite matchings from the sources to the relays and from the relays to the destinations.

Owing to the symmetry assumption, each session $\mathbf{s}$ has an identical prior probability

$$p(\mathbf{s}) = \frac{1}{n!n!}.$$

Medium Access Constraints: We consider a transmitter-directed signaling model, where every node (source or relay) has an independent transmission rate constraint. Let C^{s} denote the transmission rate constraint for any source, and let C^{r} denote the transmission rate constraint for any relay.

Transmission and Relaying Strategy: For purposes of analytical characterization, we consider independent Poisson schedules, where all source schedules and covert relay schedules are generated according to independent Poisson processes. The relaying strategy used by any covert relay is the Bounded Greedy Match algorithm [25], which was shown to maximize the sum-rate of relayed data packets.

Throughput: Given the transmission rates of the relay and the source nodes, [4, Theorem 1] characterizes the maximum achievable data rate when the BGM algorithm is used as the relaying strategy. Since all routes in the parallel relay network are 2-hop routes, the sum-rate $\Lambda(\mathbf{s}, \mathbf{b}_{\mathrm{n}})$ in a session $\mathbf{s}$ when relays in $\mathbf{b}_{\mathrm{n}}$ are covert is expressible as a sum of achievable rates for each source–destination pair

$$\Lambda(\mathbf{s}, \mathbf{b}) = (n - |\mathbf{b}|)\min(C^{\mathrm{s}}, C^{\mathrm{r}}) + |\mathbf{b}|\lambda^*(C^{\mathrm{s}}, C^{\mathrm{r}}),$$
$$\text{where } \lambda^*(a,b) = a\frac{be^{\Delta(b-a)} - b}{be^{\Delta(b-a)} - a}$$

is the maximum achievable rate for a covert relay using independent Poisson schedules under a strict delay constraint of Δ seconds per packet [4].

The throughput, as defined in Section II, is given by

$$\Upsilon(q_{\mathrm{n}}) = \sum_{\mathbf{s}} p(\mathbf{s}) \sum_{\mathbf{b}_{\mathrm{n}}} q_{\mathrm{n}}(\mathbf{b}_{\mathrm{n}} \,|\, \mathbf{s}) \Lambda(\mathbf{s}, \mathbf{b}_{\mathrm{n}}).$$

The maximum achievable throughput $\Upsilon_{\max}$ when all relays are visible is given by

$$\Upsilon_{\max} = n \min(C^{\mathrm{s}}, C^{\mathrm{r}}).$$

Note that sum-rate here is used as a specific scalar measure of performance to define the strategy space of the network nodes. In general, any function of capacity region can be used to define the strategy space of the network, and the results here can be extended to such models as well.

Adversary Model: The adversary monitors a subset of the nodes, which we denote by a pair of random variables $\mathbf{N}_{\mathrm{a}}^{\mathrm{s}}, \mathbf{N}_{\mathrm{a}}^{\mathrm{r}}$, where $\mathbf{N}_{\mathrm{a}}^{\mathrm{s}}$ and $\mathbf{N}_{\mathrm{a}}^{\mathrm{r}}$ denote the sources and relays that are monitored, respectively. For every monitored node, the adversary has perfect knowledge of the packet transmission times. We know that $|\mathbf{N}_{\mathrm{a}}^{\mathrm{s}}| + |\mathbf{N}_{\mathrm{a}}^{\mathrm{r}}| \leq k_{\mathrm{a}}$.

Given the bipartite session model, at every monitored relay, the schedule observed by the adversary is either correlated to that of a monitored source node or is independent of every monitored source node. In effect, the adversary observation $f_{\mathrm{a}}(f_{\mathrm{o}}(\mathbf{s}, \mathbf{b}_{\mathrm{n}}), \mathbf{n}_{\mathrm{a}}) = \mathbf{p}_{\mathrm{a}}^{\mathrm{s,r}} \cup \mathbf{p}_{\mathrm{a}}^{\mathrm{s}} \cup \mathbf{p}_{\mathrm{a}}^{\mathrm{r}}$, where we have the following.

1) $\mathbf{p}_{\mathrm{a}}^{\mathrm{s,r}}$ is a set of source–relay pairs with dependent schedules.
2) $\mathbf{p}_{\mathrm{a}}^{\mathrm{s}}$ is a set of source nodes whose schedules are not correlated with that of any monitored relay.
3) $\mathbf{p}_{\mathrm{a}}^{\mathrm{r}}$ is a set of relays whose schedules are not correlated with that of any monitored source.

For example, consider a session in a three-source parallel-relay network, where source A_i communicates with destination D_i through relay R_i. Let the network designer make relay R_1 covert and the adversary monitor the nodes A_1, A_2, R_1, R_2, and R_3. In this example, the adversary observation can be written as $\mathbf{p}_{\mathrm{a}}^{\mathrm{s,r}} \cup \mathbf{p}_{\mathrm{a}}^{\mathrm{s}} \cup \mathbf{p}_{\mathrm{a}}^{\mathrm{r}}$, where

$$\mathbf{p}_{\mathrm{a}}^{\mathrm{s,r}} = \{(A_2, R_2)\} \quad \mathbf{p}_{\mathrm{a}}^{\mathrm{s}} = \{A_1\} \quad \mathbf{p}_{\mathrm{a}}^{\mathrm{r}} = \{R_1, R_3\}.$$

Anonymity: By merely monitoring the transmissions of the nodes in the network, an adversary can at most identify every source–relay pair. Since the network utilizes transmitter-directed signaling, using transmission timing alone, it is impossible to determine any final destination. We therefore

measure anonymity using the set of source–relay pairs perfectly identifiable by the adversary. Let $\mathbf{S}'$ denote the set of source–relay pairs in the session. We can write

$$H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}) = H(\mathbf{S}' \mid \hat{\mathbf{S}}_{\mathrm{a}}) + H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}, \mathbf{S}').$$

Since $\mathbf{S}'$ contains all the source–relay pairings and $\hat{\mathbf{S}}_{\mathrm{a}}$ contains no information about destinations, $H(\mathbf{S} \mid \mathbf{S}', \hat{\mathbf{S}}_{\mathrm{a}}) = H(\mathbf{S} \mid \mathbf{S}')$, which is a constant irrespective of the set of monitored nodes. We therefore modify the payoff in the two-player game as

$$\phi = \frac{H(\mathbf{S}') \mid \hat{\mathbf{S}}_{\mathrm{a}})}{H(\mathbf{S}')}.$$

It is easy to see that the total anonymity as defined in Section II has a monotonic one-to-one relationship to the above definition.

Our goal is study the saddle-point strategies and throughput-anonymity tradeoffs of this network model by jointly optimizing the covert probability function $\{q_{\mathrm{n}}(\mathbf{b}_{\mathrm{n}} \mid \mathbf{s})\}$ and the adversary strategy $q_{\mathrm{a}}(\mathbf{n}_{\mathrm{a}})$ subject to the throughput constraint $\Upsilon(q_{\mathrm{n}}) \geq \gamma$ and the adversary power k_{a}. If $q_{\mathrm{n}}^*, q_{\mathrm{a}}^*$ denote the NE probability distributions of the network designer and adversary, respectively, then let

$$A^*(\gamma) = \phi(q_{\mathrm{n}}^*, q_{\mathrm{a}}^*)$$

represent the NE anonymity–throughput tradeoff.

Theorem 3: For an omniscient adversary, the NE throughput anonymity tradeoff is given by

$$A^*(\gamma) = \frac{(\Upsilon_{\max} - \gamma)}{n\epsilon}, \text{ where } \epsilon = \min(C^{\mathrm{s}}, C^{\mathrm{r}}) - \lambda^*(C^{\mathrm{s}}, C^{\mathrm{r}}).$$

Proof: Refer to the Appendix. □

The throughput–anonymity tradeoff under an omniscient adversary is linear, which is a consequence of the 2-hop nature and symmetry in the network model. The constant ϵ represents the per-node rate loss. As mentioned earlier, this operating point represents a trivial equilibrium. Against an omniscient adversary, the optimal strategy for the network designer is to make all relays covert together with probability

$$q_{\mathrm{n}}(\mathcal{V} \mid \mathbf{s}) = \frac{\Upsilon_{\max} - \gamma}{n\epsilon} \qquad \forall \mathbf{s}.$$

The general idea behind this strategy is as follows: If in a session, k relays are covert, then the anonymity from an omniscient adversary's perspective would be restricted to the k relays and cannot exceed $\log k!$. The corresponding loss in throughput for the network designer is $k\epsilon$. The optimal network design strategy would therefore correspond to minimizing the throughput cost per unit gain in anonymity.

B. General Adversary Model

Consider the simplest case of $k_{\mathrm{a}} = 2$. When $k_{\mathrm{a}} = 2$, the only way the adversary can obtain nonzero information is if one of the monitored nodes is a relay and the other is a source. Due to the symmetry assumption, intuition suggests that the optimal strategy for the adversary would be to monitor every source–relay pair with equal probability.

When $k_{\mathrm{a}} > 2$, there is an additional challenge in determining the ratio of relays and sources that should be monitored by the adversary. In general, the optimal ratio need not be fixed and could be a random quantity, as long as the total number of monitored nodes does not exceed k_{a}. However, optimizing the adversary and network strategies reveals that the optimal strategy would in fact have a fixed ratio. This is shown in the following theorem, which characterizes the equilibrium throughput–anonymity tradeoff for the general adversary.

Theorem 4: Let $p_{\mathrm{c}} = \frac{\Upsilon_{\max} - \gamma}{n\epsilon}$, $k = \lfloor \frac{k_{\mathrm{a}}}{2} \rfloor$, $k' = \lceil \frac{k_{\mathrm{a}}}{2} \rceil$, and

$$w(m) = \begin{cases} \frac{((n-k)!)^2}{(n-2k+m)!}, & k_{\mathrm{a}} \leq n \\ 0, & \text{o.w.} \end{cases}$$

Then, there exists a unique equilibrium throughput–anonymity tradeoff that is given by

$$\begin{aligned} A^*(\gamma) = &\left[p_{\mathrm{c}} + \frac{w(0)(1-p_{\mathrm{c}})}{n!}\right] \log(w(0)(1-p_{\mathrm{c}}) + n! p_{\mathrm{c}})) \\ &+ \frac{(n! - w(0))}{n!} p_{\mathrm{c}} \log p_{\mathrm{c}} \\ &+ \sum_{(k_{\mathrm{a}}-n-1)^+ + 1}^{k} \binom{k}{m} \binom{k'}{m} \frac{m!}{n!} (1-p_{\mathrm{c}}) w(m) \\ &\times \log(w(m)). \end{aligned}$$

Proof: Refer to the Appendix. □

The anonymity at the saddle point is composed of two components. The first term represents the uncertainty in determining which of the monitored relays is covert; since only a subset of sources are monitored, independence across schedules does not necessarily imply that the relay is covert. The remaining component of the anonymity is the uncertainty due to the unobserved nodes in the network. The quantity p_{c} represents the average probability with each a relay is covert, and this probability is influenced by the level of throughput required. The relationship is similar to the omniscient adversary case. As the network size increases, the first component converges to a constant, and the anonymity is dominated by the missing information from unobserved nodes (see Section V).

Saddle-Point Strategies: The optimal strategy for the adversary at the saddle point, as revealed in the proof, is to monitor an equal number of relays and sources such that each $\frac{k_{\mathrm{a}}}{2}$-size subsets of relays and sources are chosen uniformly randomly. When k_{a} is odd, the adversary monitors one additional relay. The intuitive argument for this strategy is as follows: If the number of sources monitored exceeded the number of monitored relays by 2 or more, then by removing one monitored source and adding a monitored relay, the number of possible routes that can be discovered only stands to increase.

The optimal strategy for the network designer is to make all the relays to be covert with probability

$$q_{\mathrm{n}}(\mathcal{V} \mid \mathbf{s}) = \frac{\Upsilon_{\max} - \gamma}{n\epsilon} \qquad \forall \mathbf{s}.$$

At first glance, this may be surprising since the adversary only monitors a subset of nodes in any session. However, if all relays were not covert, then the fraction of monitored relays that are visible provides more information per unit cost in throughput than that obtained from sessions when none of the relays are covert. Furthermore, uniform probabilities $q_{\mathrm{n}}(\mathbf{b}_{\mathrm{n}} \mid \mathbf{s})$ across sessions result in a uniform *aposterior* probability over all sessions, which maximizes entropy.

Fig. 5 plots the throughput-anonymity tradeoff for two parallel relay networks. The gain in anonymity due to the game-theoretic approach over the omniscient strategy is evident from the

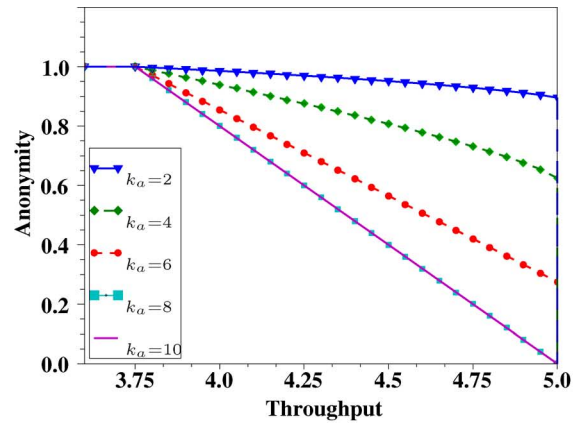

(a)

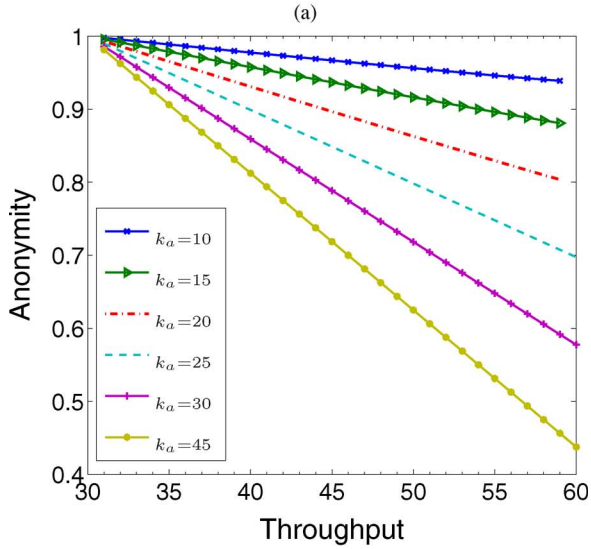

(b)

Fig. 5. Tradeoffs for parallel relay networks. (a) 5-relay parallel network: $C^{\mathrm{s}} = C^{\mathrm{r}} = 1, \Delta = 3$. (b) 60-relay parallel network: $C^{\mathrm{s}} = C^{\mathrm{r}} = 1, \Delta = 1$.

plots. Note that in the small network, while the tradeoff is linear for an omniscient adversary (Theorem 2), it is not so in general. For a large network, however, the tradeoffs are mostly linear, except for small values of k_{a}. This "asymptotic" linearity is shown analytically in Section V.

C. Asymmetric Networks

In the results thus far, the symmetry in the underlying network model resulted in symmetric strategies for the adversary and the network designer. When asymmetry is introduced in the networks, naive intuitions may not provide the saddle-point strategies. To understand the effect of asymmetry on the strategies, we consider two kinds of asymmetric networks: networks where the transmission capacities of the relays are unequal, and networks where the numbers of sources catered by the relays are unequal.

Asymmetry in Covert Relay Rates: Consider first the case of an n-parallel-relay network, where the transmission capacities of relays $B_1, \ldots, B_n$ are unequal. Specifically, there exist at least two relays B_i, B_j such that the loss in data rates $\epsilon_i \neq \epsilon_j$.

Theorem 5: For an n-relay parallel network, where an adversary monitors $k_{\mathrm{a}} = 2$ nodes, if rate losses due to covert relaying for the relays are given by $\epsilon_1, \ldots, \epsilon_n$, respectively, there exists a unique saddle point where we have the following.

1) $q_{\mathrm{n}}(B_i \,|\, \mathbf{s}) = \dfrac{\Upsilon_{\max} - \gamma}{\sum_i \epsilon_i} \quad \forall i \leq n.$
2) $q_{\mathrm{a}}(A_i, B_j) = \dfrac{\epsilon_j}{n \sum_i \epsilon_i}.$

Proof: Refer to the Appendix. □

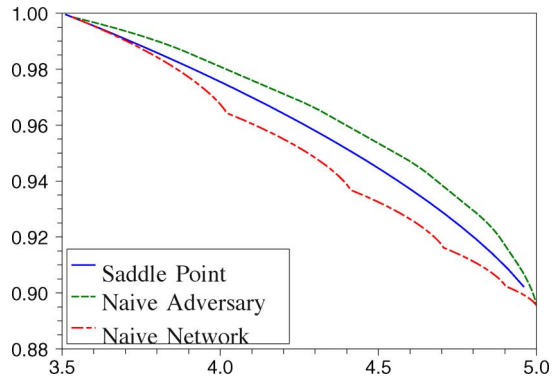

Fig. 6. Asymmetric rate loss model with $n = 5$ relays: comparison with naive strategies.

Interestingly, although the model is asymmetric, the covert relaying strategy is symmetric. This is because each relay, when visible, reveals an equal amount of information. Therefore, any asymmetry in the retrievable information from the two relays induced by the network strategy would automatically force the adversary to monitor the less protected (or more informative) relay exclusively. Such a pair of strategies cannot constitute a saddle point.

When the network design strategy is symmetric, the payoff is a constant regardless of the adversary's probability of monitoring each source–relay pair. However, there is only one strategy, at which point the optimal strategy for the network is symmetric, thus resulting in an equilibrium. In particular, the probability of monitoring a relay is proportional to the rate loss at the relay. As intuition would suggest, the more rate loss, the less likely a relay is to be covert and, consequently, a greater incentive for it to be monitored. In effect, at the saddle point, the adversary's strategy is to choose the probabilities of monitoring each relay so that the network is forced to make all relays covert with equal likelihood.

Under such an asymmetric model, if a network designer were to assume naively that the adversary's strategy were symmetric, then for a required level of throughput, the optimal strategy would be to make relays with lower throughput loss ϵ_i covert with higher probability so that the same level of throughput can be achieved with higher anonymity (w.r.t. the uniform adversary). However, the optimal adversary would then employ unequal probabilities of monitoring the relays, which would eventually result in lower-than-expected anonymity. The difference between the anonymity due to the naive networking strategy and the equilibrium strategy is shown in Fig. 6 and clearly demonstrates the benefit of using the game-theoretic approach. The figure also plots the tradeoff when the adversary employs the naive strategy of uniform monitoring, and the network designer optimizes the choice of covert relays assuming the uniform adversary.

Asymmetry in Relay Information: In the asymmetric model discussed above, the saddle-point strategy for the network

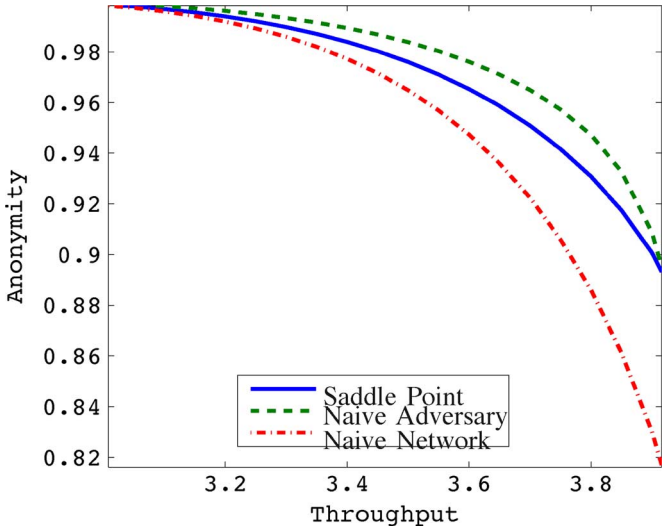

Fig. 7. Asymmetric relay information model with four sources and three relays: comparison with naive strategies.

designer was symmetric since each relay, when monitored, provided the same amount of information. We now consider a modification of the parallel network structure and introduce asymmetry in the amount of information provided by a relay. Specifically, let the number of relays be $n-k$, where k relays are multiplexing relays with two sources transmitting to each of them every session, and the remaining $n-2k$ relays are nonmultiplexing relays with exactly one source transmitting to each of them in every session. The capacities of relays are chosen such that each relay, when covert, incurs an identical throughput loss ϵ. We consider a two-player game where the adversary monitors at most two nodes.

Theorem 6: For an $n-1$ relay asymmetric parallel relay network, where an adversary monitors $k_{\mathrm{a}}=2$ nodes, there exists a unique saddle point, where the following applies.

1) The optimal strategy of the network is to make a nonmultiplexing relay covert with probability q_{r}^1 and a multiplexing relay covert with probability q_{r}^2, where

$$q_{\mathrm{r}}^1 \log\left(q_{\mathrm{r}}^1\right) - \left(q_{\mathrm{r}}^1+n-1\right)\log\left(q_{\mathrm{r}}^1+n-1\right) = 2q_{\mathrm{r}}^2\log\left(2q_{\mathrm{r}}^2\right) - \left(2q_{\mathrm{r}}^2+n-2\right)\log\left(2q_{\mathrm{r}}^1+n-2\right) - 2.$$

2) The optimal adversary strategy is to monitor a source-multiplexing relay pair with probability p_1 and a source nonmultiplexing relay pair with probability p_2 such that

$$\frac{p_1}{p_2} = \frac{(n-2k)\log\left(\frac{q_{\mathrm{r}}^1}{q_{\mathrm{r}}^1+n-1}\right)}{(k)\log\left(\frac{2q_{\mathrm{r}}^2}{2q_{\mathrm{r}}^2+n-2}\right)}.$$

Proof: Refer to the Appendix. □.

In this setup, the theorem states that the optimal strategy for the network designer is asymmetric as well. A naive adversary would choose to monitor nonmultiplexing relays with higher probability since they provide more information, whereas a naive network designer would choose to hide all relays with equal probability since all relays provide identical throughput loss. Fig. 7 plots the improvement in anonymity over naive strategies due to the game-theoretic approach.

The intuition behind the optimal strategies is similar to the asymmetric rate loss model. The more information provided by a relay, the more likely the adversary is to monitor that relay, and a greater incentive to make it covert. At the saddle point, the network increases the probability of nonmultiplexing relays being covert just enough so that the adversary obtains equal information from any relay.

D. Large Networks

In this section, we use the derived results to study equilibria in large networks. When the fraction of monitored nodes $\frac{k_{\mathrm{a}}}{2n}$ is a constant, the anonymity monotonically increases with n, but asymptotically converges toward a constant.

Theorem 7: If $\frac{k_{\mathrm{a}}}{2n}=\alpha$ is a constant, then the anonymity for a fixed throughput ratio $\gamma^*=\frac{\gamma}{\Upsilon_{\max}}$ converges as

$$\lim_{n\to\infty} A(\gamma^*) = 1-\alpha^2\frac{(\gamma^*-(1-\epsilon))^+}{\epsilon}.$$

Proof: Refer to the Appendix. □

According to the theorem, for a fixed throughput, the loss in anonymity is proportional to the square of the fraction of monitored relays. Put in another perspective, for a fixed number of monitored relays, the anonymity asymptotically converges to 1 as

$$A = 1 - O\left(\frac{1}{n^2}\right).$$

The intuition for this relationship can be understood by looking at the maximum throughput case: $\gamma^*=1$. At that operating point, $A(\gamma^*)=1-\alpha^2$. In the large n regime, the total uncertainty is approximately $n\log n$. Every monitored relay reduces uncertainty by $\log n$ if the corresponding source is also monitored. If the corresponding source is not among the monitored nodes, then the reduction in uncertainty is negligible. For every relay, the corresponding source would be monitored with approximate probability $\frac{k}{n}$. Since k relays are monitored, the net reduction in uncertainty is approximately $\frac{k^2}{n^2}$, thus resulting in the square law of the theorem.

Asymptotic relationships can be used to design approximate strategies for large networks. In particular, it would be useful to characterize the asymptotic relationship between the fraction of covert relays and the fraction of monitored relays. As the number of monitored relays increases, the fraction of relays that are covert per session would also increase. We can use Theorem 4 to obtain the asymptotic relationship. Specifically, for a fixed anonymity A, the fraction of covert relays per session β is given by

$$\beta = 1-\frac{1-A}{\alpha^2}.$$

Furthermore, if $\beta(n)$ is the exact fraction of covert relays required for a network of size n, it is easily shown that

$$\beta(n)-\beta = O\left(\frac{\log n}{n}\right).$$

This is of particular relevance to large wireless sensor networks where the number of covert relays (relays generating dummy transmissions) is directly related to energy overhead. Fig. 8 plots this relationship for finite networks in comparison with the asymptotic relationship.

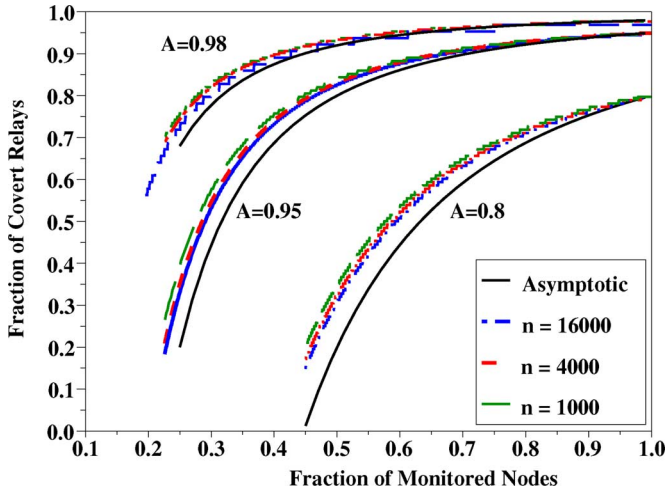

Fig. 8. Covert versus monitored relays. The three sets of curves are plotted for $A = 0.8, 0.95, 0.98$.

V. Concluding Remarks

In this paper, we considered the problem of providing anonymity to network communication when adversaries monitor or compromise an unknown subset of nodes in the network. We presented a game-theoretic formulation and proved the existence of saddle-point equilibria. Using the class of parallel relay networks, we demonstrated that this approach can be used to obtain optimal strategies for the network designer and the adversary, as well as provide insights into anonymity–throughput tradeoffs in large networks. The problem of computing the equilibria has not been dealt with in this paper, but efficient algorithms for this purpose would fortify the results here and are part of ongoing research. In this paper, we have used specific classes of networks and assumed knowledge of topology and sessions. A similar approach for random networks with random connections could shed valuable insights into scaling behavior of anonymous networking.

Appendix

A. Proof of Theorem 1

In order to prove the existence of a saddle point in the two-player game, it is sufficient to show the following.

1) $\mathcal{A}_{\rm n}$ and $\mathcal{A}_{\rm a}$ are closed convex and bounded sets.
2) The payoff is continuous in the domain $\mathcal{A}_{\rm n} \times \mathcal{A}_{\rm a}$.
3) For every $q_{\rm a} \in \mathcal{A}_{\rm a}$, $\phi(x, q_{\rm a})$ is concave in x.
4) For every $q_{\rm n} \in \mathcal{A}_{\rm n}$, $-\phi(q_{\rm n}, y)$ is concave in y.

If the two-player game satisfies the above conditions, then it constitutes a general two-player concave game, which was shown to have a guaranteed Nash equilibrium in [26].

1) *Convexity of action spaces:* The space $\mathcal{A}_{\rm a}$ is a finite-dimensional simplex, which is closed, bounded, and convex. $\mathcal{A}_{\rm n}$ is a subset of the simplex with the additional constraint

$$R(\mathbf{q}_{\rm a}) \geq r.$$

Since the constraint is not a strict inequality, the space is closed. $R(\,\cdot\,)$ is a linear function of $\mathbf{q}_{\rm a}$. Therefore, for any pair of probability vectors $\mathbf{q}_{\rm a}^1, \mathbf{q}_{\rm a}^2$

$$\alpha R\left(\mathbf{q}_{\rm a}^1\right) + (1-\alpha) R\left(\mathbf{q}_{\rm a}^2\right) = R\left(\alpha \mathbf{q}_{\rm a}^1 + (1-\alpha)\,\mathbf{q}_{\rm a}^2\right)$$

which proves the convexity of $\mathcal{A}_{\rm n}$.

2) Since the payoff is linear in $\mathbf{q}_{\rm a}$ and is an entropy function of $\mathbf{q}_{\rm n}$, the continuity of the payoff can be easily shown (the details are omitted here).
3) In order to show the concavity of ϕ w.r.t. to $\mathbf{q}_{\rm n}$, we need to show that for any $\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm n}^2 \in \mathcal{A}_{\rm n}$,$\mathbf{q}_{\rm a} \in \mathcal{A}_{\rm a}$

$$\alpha\phi\left(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}\right) + (1-\alpha)\phi\left(\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}\right) \leq \phi\left(\alpha\mathbf{q}_{\rm n}^1 + (1-\alpha)\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}\right).$$

Consider the following modification to the setup, where apart from the topology and set of network sessions, the network designer and the adversary are given access to a common Bernoulli random variable $Z \sim \mathcal{B}(\alpha)$. Consider any $\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm n}^2 \in \mathcal{A}_{\rm n}$. The network designer utilizes the following strategy: If the observed variable $Z = 1$, then the distribution $\mathbf{q}_{\rm n}^1$ is used to make relays covert, and if $Z = 0$, $\mathbf{q}_{\rm n}^2$ is used. Since Z is observed by the adversary as well, this strategy would amount the anonymity being equal to the conditional entropy $H(\mathbf{S} \mid \hat{\mathbf{S}}, Z)$.
Now, suppose the Bernoulli variable were only available to the network designer, and he utilizes the same strategy. Since the adversary has no knowledge of Z, his entropy would be $H(\mathbf{S} \mid \hat{\mathbf{S}})$, where the distribution of covert relays would be the effective distribution

$$\alpha\mathbf{q}_{\rm n}^1 + (1-\alpha)\mathbf{q}_{\rm n}^2.$$

Since conditioning reduces entropy, $H(\mathbf{S} \mid \hat{\mathbf{S}}, Z) \leq H(\mathbf{S} \mid \hat{\mathbf{S}}$, and therefore

$$\alpha\phi\left(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}\right) + (1-\alpha)\phi\left(\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}\right) \leq \phi\left(\alpha\mathbf{q}_{\rm n}^1 + (1-\alpha)\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}\right).$$

4) For any $\mathbf{q}_{\rm n}$, $\phi(\mathbf{q}_{\rm n}, \mathbf{q}_{\rm a})$ is a linear function of $\mathbf{q}_{\rm a}$, and therefore

$$\alpha\phi\left(\mathbf{q}_{\rm n}, \mathbf{q}_{\rm a}^1\right) + (1-\alpha)\phi\left(\mathbf{q}_{\rm n}, \mathbf{q}_{\rm a}^2\right) = \phi\left(\mathbf{q}_{\rm n}, \alpha\mathbf{q}_{\rm a}^1 + (1-\alpha)\mathbf{q}_{\rm a}^2\right)$$

which establishes the required concavity.

For uniqueness, consider two pairs of strategies $(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}^1)$ and $(\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}^2)$ that achieve saddle-point equilibrium. By the definition of saddle point, we know that

$$\begin{aligned}\phi\left(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}^1\right) &\leq \phi\left(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}^2\right) \leq \phi\left(\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}^2\right) \leq \phi\left(\mathbf{q}_{\rm n}^2, \mathbf{q}_{\rm a}^1\right) \\ &\leq \phi\left(\mathbf{q}_{\rm n}^1, \mathbf{q}_{\rm a}^1\right).\end{aligned}$$

The above sequence of inequalities establishes the uniqueness of the payoff. □

B. Proof of Theorem 2

According to the definition of payoff

$$\begin{aligned}\phi(q_{\rm n}, q_{\rm a}) = \frac{H(\mathbf{S} \mid \hat{\mathbf{S}}_{\rm a})}{H(\mathbf{S})} = \frac{1}{H(\mathbf{S})} \sum_{\mathbf{n}_{\rm a}} \sum_{\mathbf{s}, \mathbf{b}_{\rm n}} -q_{\rm a}(\mathbf{n}_{\rm a}) p(\mathbf{s}) \\ \times q_{\rm n}(\mathbf{s}, \mathbf{b}_{\rm n}) \log q_{\rm ap}(\mathbf{s}, f_{\rm a}(f_{\rm o}(\mathbf{s}, \mathbf{b}_{\rm n}), \mathbf{n}_{\rm a}), \mathbf{b}_{\rm a}). \quad (11)\end{aligned}$$

From the adversary's perspective, the goal is to choose $q_{\rm a}$ such that $\phi(q_{\rm n}, q_{\rm a})$ is minimized. Since $q_{\rm a}$ is a probability distribution, using Lagrange multipliers, we can write

$$L_{\rm a} = \phi(q_{\rm n}, q_{\rm a}) + \beta_{\rm a} \sum_{\mathbf{n}_{\rm a}} q_{\rm a}(\mathbf{n}_{\rm a}).$$

At the minimizing distribution, we know that

$$\frac{dL_{\mathrm{a}}}{dq_{\mathrm{a}}(\mathbf{n}_{\mathrm{a}})} = 0 \qquad \forall \mathbf{n}_{\mathrm{a}}.$$

Therefore, for any subset of nodes $\mathbf{n}_{\mathrm{a}}$ for which $q_{\mathrm{a}}(\mathbf{n}_{\mathrm{a}}) > 0$

$$H_{\mathrm{p}}\left(\mathbf{n}_{\mathrm{a}}^{1}\right) + \beta_{\mathrm{a}} \text{ is a constant}$$

which proves the first part of the theorem.

From the network designer's perspective, the goal is to design $q_{\mathrm{n}}(\mathbf{b}_{\mathrm{n}})$ such that $\phi(q_{\mathrm{n}}, q_{\mathrm{a}})$ is maximized, while maintaining a throughput γ. Again, using Lagrange multipliers, we can define

$$L_{\mathrm{n}} = \phi(q_{\mathrm{n}}, q_{\mathrm{a}}) + \beta_1 \sum_{\mathbf{s},\mathbf{b}} p(\mathbf{s}) q_{\mathrm{n}}(\mathbf{s},\mathbf{b}) \Lambda(\mathbf{s},\mathbf{b}) + \sum_{\mathbf{s}} p(\mathbf{s}) \sum_{\mathbf{b}} \beta_2(\mathbf{s}) q_{\mathrm{n}}(\mathbf{s},\mathbf{b}).$$

At the maximizing distribution, for every $q(\mathbf{s},\mathbf{b}) > 0$

$$\frac{dL_{\mathrm{n}}}{dq_{\mathrm{n}}(\mathbf{b}_{\mathrm{n}})} = 0.$$

$$\Rightarrow \sum_{\mathbf{n}_{\mathrm{a}}} q_{\mathrm{a}}(\mathbf{n}_{\mathrm{a}}) \left[p(\mathbf{s}) + p(\mathbf{s}) \log(q_{\mathrm{n}}(\mathbf{s},\mathbf{b}_{\mathrm{n}})) - p(\mathbf{s}) - p(\mathbf{s}) \log \left[\sum_{\mathbf{s}',\mathbf{b}'_{\mathrm{n}}} q_{\mathrm{n}}(\mathbf{s}',\mathbf{b}'_{\mathrm{n}}) p(\mathbf{s}') \right] \right] + \beta_1(\Lambda(\mathbf{s},\mathbf{b}_{\mathrm{n}})) + \beta_2(\mathbf{s}) = 0.$$

Equating the values of β_1, $\beta_2(\mathbf{b})$, the conditions are obtained. □

C. Proof of Theorem 3

Define $p_k = \sum_{\mathbf{s},\mathbf{b}:|\mathbf{b}|=k} p(\mathbf{s}) q_{\mathrm{n}}(\mathbf{b} \mid \mathbf{s})$. Due to the symmetric rates, the throughput achievable by a strategy q_{n} is

$$\Upsilon(q_{\mathrm{n}}) = \Upsilon_{\max} - \sum_{k} p_k k \epsilon$$

where $\epsilon = \min(C^{\mathrm{r}}, C^{\mathrm{s}}) - f(C^{\mathrm{r}}, C^{\mathrm{s}})$.

For a given strategy q_{n}, the anonymity for an omniscient adversary can be written as

$$H(\mathbf{S} \mid \mathbf{B}) = \sum_{\mathbf{b} \subset \mathcal{V}^{\mathrm{r}}} \left(\sum_{\mathbf{s}} p(\mathbf{s}) q_{\mathrm{n}}(\mathbf{b} \mid \mathbf{s}) \right) H(\mathbf{S} \mid \mathbf{B} = \mathbf{b}).$$

For a given realization of $\mathbf{B}$, the omniscient adversary can perfectly correlate the flows through all relays in $\mathcal{V}^{\mathrm{r}} \backslash \mathbf{B}$, therefore the information lost due to independent schedules can be upper-bounded by

$$H(\mathbf{S} \mid \mathbf{B} = \mathbf{b}) \leq \log(|\mathbf{b}|!).$$

$$\Rightarrow H(\mathbf{S} \mid \mathbf{B}) \leq \sum_{\mathbf{b}} \left(\sum_{\mathbf{s}} p(\mathbf{s}) q_{\mathrm{n}}(\mathbf{b} \mid \mathbf{s}) \right) \log(|\mathbf{b}|!) = \sum_{k} p_k \log(k!).$$

Consider maximizing $\sum_k p_k \log(k!)$ subject to

$$\sum_{k} p_k k \epsilon \leq \Upsilon_{\max} - \gamma.$$

If $\Upsilon_{\max} - \gamma \geq n\epsilon$, it is easy to see that $q_{\mathrm{n}} = 1$. When $\Upsilon_{\max} - \gamma \geq n\epsilon$, since $\frac{\log(k!)}{k}$ is increasing in k, the maximizing $\{p_k\}$ is given by

$$p_k = 0 \quad k < n \quad p_n = \frac{\Upsilon_{\max} - \gamma}{n\epsilon}.$$

Therefore, for any throughput t

$$H(\mathbf{S} \mid \mathbf{B}) \leq \frac{\Upsilon_{\max} - \gamma}{n\epsilon} \log(n!).$$

The above inequality is achievable by making all relays covert with probability p_n, and hence proves the theorem.

D. Proof of Theorem 4

Consider the following adversary strategy: During every session, the adversary picks $\frac{k_{\mathrm{a}}}{2}$ source–relay pairs with uniform probability. We characterize the optimal network strategy for this adversary and show that the adversary can do no better by changing his strategy, thus proving equilibrium.

For a given set of monitored nodes $\mathbf{B} \in (\mathcal{V}^{\mathrm{s}})^k \times (\mathcal{V}^{\mathrm{r}})^k$, let $X_{\mathbf{B}}$ be a random variable that denotes the set of communicating source–relay pairs within the set of monitored nodes. Then, for a given covert relaying strategy $q_{\mathrm{n}}()$, the anonymity for the specified adversary can be expressed as

$$H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}) = \sum_{\mathbf{b}} (H(X_{\mathbf{b}} \mid \hat{\mathbf{S}}_{\mathrm{a}}) + H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}, X_{\mathbf{B}_2}) = \sum_{\mathbf{b}} (H(X_{\mathbf{b}} \mid \hat{\mathbf{S}}_{\mathrm{a}}) + H(\mathbf{S} \mid X_{\mathbf{b}})$$

where the second equality is because, given the communications within the monitored nodes, the uncertainty of the rest of the network does not depend on the observation.

Furthermore, given the set of communicating pairs within the set of monitored nodes, the uncertainty in the unobserved portion of the network would be independent of any strategy, and therefore a constant.

Accordingly, consider maximizing $\sum H(X_{\mathbf{B}} \mid \hat{\mathbf{S}}_{\mathrm{a}})$ subject to the throughput constraint. This maximization is akin to the omniscient case; the uncertainty refers to the communications within the monitored nodes. The difference comes from the fact that since there are unobserved nodes in the network, some of the monitored sources or relays can communicate with nodes outside the set of monitored nodes. Nevertheless, it can be shown that the optimal network strategy is not affected by this modification. We prove this for $k_{\mathrm{a}} = 2$; the proof for general k_{a} is a straightforward generalization. Define

$$p^{c}(\mathbf{b}) = \sum_{\mathbf{S}} p(\mathbf{S}) \left(1 - \sum_{\mathbf{B}:\mathbf{b} \cap \mathbf{B} \neq \phi} q(\mathbf{B} \mid \mathbf{S}) \right).$$

In other words $p^{c}(\mathbf{b})$ is the probability that a flow through $\mathbf{b}$ is visible. Therefore

$$H(X_{\mathbf{b}} \mid \hat{\mathbf{S}}_{\mathrm{a}}) = h(p^{c}(\mathbf{b}))$$

where $h(p)$ is the binary entropy function. Due to the throughput requirement, we know that $\sum_{\mathbf{b}} p^c(\mathbf{b})$ is a constant. Since finite entropy is bounded by the size of the alphabet

$$\sum_{\mathbf{b}\in\mathcal{V}^{\mathrm{s}}\times\mathcal{V}^{\mathrm{r}}} H(X_{\mathbf{b}} \mid \hat{\mathbf{S}}_{\mathrm{a}}) \leq n^2 h\left(\frac{1}{n}\right)$$

where the equality is achieved when $\forall \mathbf{b}, p^c(\mathbf{b})$ is identical.

Furthermore, since $q(\mathbf{B} \mid \mathbf{S})$ is independent of $\mathbf{S}$

$$H(\mathbf{S} \mid X_{\mathbf{B}}) = \log((n-1)!)$$

which is independent of the covert relaying strategy.

The optimal covert relaying strategy is therefore symmetric across all relays and sessions. Using the two derived conditions, the maximizing anonymity is given by

$$H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}) = h\left(\frac{1}{n}\right) + \log((n-1)!).$$

For the derived covert relaying strategy, the anonymity w.r.t. to a general adversary can be written as

$$H(\mathbf{S} \mid \hat{\mathbf{S}}_{\mathrm{a}}) = \sum_{\mathbf{b}\in\mathcal{V}^{\mathrm{s}}\times\mathcal{V}^{\mathrm{r}}} q_{\mathrm{a}}(\mathbf{b})(H(X_{\mathbf{b}} \mid \hat{\mathbf{S}}_{\mathrm{a}}) + H(\mathbf{S} \mid X_{\mathbf{b}_2})$$

where $q_{\mathrm{a}}(\mathbf{b}_2)$ is the probability that the adversary monitors the source–relay pair $\mathbf{b}$. Due to the symmetry in covert relaying strategy, $H(X_{\mathbf{b}})$ and $H(\mathbf{S} \mid X_{\mathbf{b}})$ are identical across pairs $\mathbf{b}$. Therefore, for any probability mass function $\{q_{\mathrm{a}}(\,\cdot\,)\}$, the total information gained (or lost) would be no different for the adversary. In other words, there is no incentive for the adversary to deviate from the uniform monitoring strategy, and that pair of strategies is therefore a saddle point. □.

E. Proof of Theorem 5

Since uniform probability maximizes entropy, we can write

$$q_{\mathrm{n}}(B_1 \mid \mathbf{s}_1) = q_1 \qquad \forall \mathbf{s} \quad q_{\mathrm{n}}(B_2 \mid \mathbf{s}) = q_2 \qquad \forall \mathbf{s}.$$

Then, $\Upsilon_{\max} - \gamma = q_1\epsilon_1 + q_2\epsilon_2$. If the adversary monitors B_1 with probability p, then

$$\phi(p,(q_1,q_2)) = p\left[\frac{1}{2}\log\left(\frac{1+q_1}{q_1}\right) + \frac{1}{2}\log(1+q_1)\right] + (1-p)\left[\frac{1}{2}\log\left(\frac{1+q_1}{q_1}\right) + \frac{1}{2}\log(1+q_1)\right].$$

If $q_1 > q_2$, then $p = 0$ is optimal for the adversary. However, if $p = 0$, then the optimal network strategy is to make $q_1 = 0$, which is a contradiction. Hence

$$q_1 = q_2 = \frac{\Upsilon_{\max} - \gamma}{\epsilon_1 + \epsilon_2}.$$

If $p*$ is the saddle-point strategy for the adversary, then p^* must necessarily satisfy (from Theorem 2)

$$\frac{d}{dq_1}\phi\left(p^*, \left(q_1, \frac{\Upsilon_{\max} - \gamma - q_1\epsilon_1}{\epsilon_2}\right)\right) = 0$$

where $q_1 = \frac{\Upsilon_{\max}-\gamma}{\epsilon_1+\epsilon_2}$. It is easily verified that $p^* = \frac{\epsilon_1}{\epsilon_1+\epsilon_2}$ is the unique solution to the above equation.

F. Proof of Theorem 6

The adversary has two choices: either monitor a source and a nonmultiplexing relay, or a source and a multiplexing relay. Within the set of relays, condition 1 in Theorem 2 requires that the amount of information available through each relay is identical. In other words, within the set of multiplexing relays, the probability of covertness would be identical. Consequently, within the set of multiplexing relays, the probability of an adversary monitoring any particular multiplexing relay would be identical. Likewise, the argument applies to the set of nonmultiplexing relays as well. Therefore, if $q_{\mathrm{r}}^1, q_{\mathrm{r}}^2$ refers to the respective probabilities of monitoring a nonmultiplexing and multiplexing relay, and if p_1, p_2 refers to the respective probabilities of an adversary monitoring a nonmultiplexing and multiplexing relay, then

$$\phi = 1 - \frac{q_{\mathrm{a}}^1}{\log(S_T)}\left[q_{\mathrm{r}}^1 \log\left(q_{\mathrm{r}}^1\right) - \left(q_{\mathrm{r}}^1 + n - 1\right)\log\left(q_{\mathrm{r}}^1 + n - 1\right)\right] - \frac{q_{\mathrm{a}}^2}{\log(S_T)}\left[2q_{\mathrm{r}}^2\log\left(2q_{\mathrm{r}}^2\right) + 2 - \left(2q_{\mathrm{r}}^2 + n - 2\right)\log\left(2q_{\mathrm{r}}^2 + n - 2\right)\right]$$

where $S_T = \frac{n!}{2^k}$ is the total number of sessions. Applying the conditions in Theorem 2 to the expression above, the theorem is proved. Details are omitted due to paucity of space □.

G. Proof of Theorem 7

We know from Theorem 3 that the anonymity $A(\gamma)$ can be written as

$$A(\gamma) = \frac{A_1(\gamma) + A_2(\gamma)}{n\log n}$$

where

$$A_1(\gamma) = \left[p_{\mathrm{c}} + \frac{w(0)(1-p_{\mathrm{c}})}{n!}\right]\log(w(0)(1-p_{\mathrm{c}}) + n!p_{\mathrm{c}})) + \frac{(n! - w(0))}{n!}\log p_{\mathrm{c}}$$

$$A_2(\gamma) = \sum_{\max(1,2k-n)}^{k} \binom{k}{m}\binom{k'}{m}\frac{m!}{n!}(1-p_{\mathrm{c}})w(m)\log(w(m)).$$

Using Stirling's approximation for large n, we can write

$$\begin{aligned}\frac{w(0)}{n!} &= \frac{((n-\alpha n)!)^2}{n!(n-2\alpha n)!} \\ &\approx \frac{(n-\alpha n)^{2n-2\alpha n}\sqrt{(1-\alpha)^2 n^2 4\pi^2}e^{-2n+2\alpha n}}{n^n(n-2\alpha n)^{n-2\alpha n}\sqrt{(1-2\alpha)n^2 4\pi^2}e^{-2n+2\alpha n}} \\ &= \sqrt{\frac{(1-\alpha)^2}{1-2\alpha}}e^{n[(2-2\alpha)\log(1-\alpha)-(1-2\alpha)\log(1-2\alpha)]} \\ &\to 0 \quad \text{for any } \alpha \in (0,1). \\ &\text{Therefore } \lim_{n\to\infty} A_1(\gamma) = p_{\mathrm{c}}. \end{aligned} \tag{12}$$

Using Stirling's approximation on $\log w(m)$, for large n

$$\begin{aligned}\log w(m) &= 2\log((n-k)!) - \log((n-2k+m)!) \\ &= 2(n-k)\log(n-k) - (n-2k+m)\log(n-2k+m) \\ &\quad + \frac{1}{2}\log\left(\frac{(n-k)^2}{n-2k+m}\right) + O(1) \\ &= \left(n-m+\frac{1}{2}\right)\log n + O(1).\end{aligned}$$

Since $m \le \alpha n$, we can write

$$\begin{aligned}A_2(\gamma) &= \sum_{m=\max(1,2k-n)}^{k} \binom{k}{m}\binom{k'}{m}\frac{m!}{n!}(1-p_c)w(m)\log(w(m)) \\ &= \frac{1-p_c}{\binom{n}{k}} \sum_{m=\max(2k-n,1)}^{k} \binom{k}{m}\binom{n-k}{k-m}(n-m)\log n \\ &= \frac{(1-p_c)\log n}{\binom{n}{k}}\left[(n-k)\binom{n}{k} + \frac{k}{n}(n-k)\binom{n}{k}\right] \\ &= \frac{n^2-k^2}{n}(1-p_c)\log n. \qquad (13)\end{aligned}$$

Combining (12) and (13), the result is proven. □.

REFERENCES

[1] N. Matthewson and R. Dingledine, "Practical traffic analysis: Extending and resisting statistical disclosure," in *Proc. PET*, May 2004, vol. 3424/2005, pp. 784–786.

[2] T. He and L. Tong, "Detecting information flows: Improving Chaff tolerance by joint detection," in *Proc. Conf. Inf. Sci. Syst.*, Baltimore, MD, Mar. 2007, pp. 51–56.

[3] C. E. Shannon, "Communication theory of secrecy systems," *Bell Syst. Tech. J.*, vol. 28, no. 4, pp. 656–715, 1949.

[4] P. Venkitasubramaniam, T. He, and L. Tong, "Anonymous networking amidst eavesdroppers," *IEEE Trans. Inf. Theory*, vol. 54, no. 6, pp. 2770–2784, Jun. 2008.

[5] H. S. Kuhn, *Classics in Game Theory*. Princeton, NJ: Princeton Univ. Press, 1944.

[6] D. Chaum, "Untraceable electronic mail, return addresses and digital pseudonyms," *Commun. ACM*, vol. 24, pp. 84–88, Feb. 1981.

[7] Y. Zhu, X. Fu, B. Graham, R. Bettati, and W. Zhao, "On flow correlation attacks and countermeasures in mix networks," in *Proc. PET*, May 26–28, 2004, vol. 3424/2005, pp. 735–742.

[8] B. Radosavljevic and B. Hajek, "Hiding traffic flow in communication networks," in *Proc. IEEE Military Commun. Conf.*, 1992, vol. 3, pp. 1096–1100.

[9] R. E. Newman, I. S. Moskowitz, P. Syverson, and A. Serjantov, "Metrics for traffic analysis prevention," in *Proc. PET*, Apr. 2003, LNCS 2760, pp. 48–65.

[10] P. Venkitasubramaniam and L. Tong, "Throughput anonymity trade-off in wireless networks under latency constraints," in *Proc. IEEE INFOCOM*, Phoenix, AZ, Apr. 2008, pp. 241–245.

[11] J. F. Nash, "Equilibrium points in n-person games," in *Proc. Nat. Acad. Sci.*, Jan. 1950, vol. 36, pp. 48–49.

[12] T. Basar, "The Gaussian test channel with an intelligent jammer," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 1, pp. 152–157, Jan. 1983.

[13] J. M. Borden, D. M. Mason, and R. J. McEliece, "Some information theoretic saddle points," *SIAM J. Control Optimiz.*, vol. 23, pp. 129–143, Jan. 1985.

[14] M. M'edard, "Capacity of correlated jamming channels," in *Proc. 35th Annu. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, Sep. 1997, vol. 35, pp. 1043–1052.

[15] A. Kashyap, T. Basar, and R. Srikant, "Correlated jamming on MIMO Gaussian fading channels," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 2119–2123, Sep. 2004.

[16] A. Kashyap, T. Basar, and R. Srikant, "Mutual information games in multiuser channels with correlated jamming," *IEEE Trans. Inf. Theory*, vol. 55, no. 10, pp. 4598–4607, Oct. 2009.

[17] T. Alpcan and T. Basar, "A game-theoretic analysis of intrusion detection in access control systems," in *Proc. IEEE Conf. Decision Control*, Paradise Island, Bahamas, Dec. 2004, vol. 2, pp. 1568–1573.

[18] Y. Liu, C. Comaniciu, and H. Man, "Modeling misbehaviour in adhoc networks: A game-theoretic approach to intrusion detection," *Int. J. Security Netw.*, vol. 1, no. 3–4, pp. 243–254, 2006.

[19] K. Lye and J. M. Wing, "Game strategies in network security," *Int. J. Inf. Security*, vol. 4, pp. 71–86, Feb. 2005.

[20] S. Sarkar, E. Altman, R. El-Azouzi, and Y. Hayel, "Information concealing games in communication networks," in *Proc. IEEE INFOCOM*, Phoenix, AZ, Apr. 2008, pp. 2119–2127.

[21] F. Topsoe, "Entropy and equilibrium via games of complexity," *Physica A, Statist. Mech. Appl.*, vol. 340, pp. 11–31, Sep. 2004.

[22] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[23] T. He and L. Tong, "Detecting information flows: Fundamental limits and optimal algorithms," *IEEE Trans. Inf. Theory*, 2007, submitted for publication.

[24] G. Owen, *Game Theory*. New York: Academic, 1995.

[25] A. Blum, D. Song, and S. Venkataraman, "Detection of interactive stepping stones: Algorithms and confidence bounds," in *Proc. RAID*, Sophia Antipolis, France, Sep. 2004, vol. 3224/2004, pp. 258–277.

[26] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave N-person games," *Econometerica*, vol. 33, pp. 520–534, Jul. 1965.

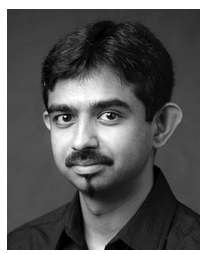

Parv Venkitasubramaniam (S'03–M'07) received the B.Tech. degree from the Indian Institute of Technology, Madras, India, in 1998, and the M.S. and Ph.D. degrees from Cornell University, Ithaca, NY, in 2005 and 2008, respectively, all in electrical engineering.

He is presently a P. C. Rossin Assistant Professor with the Electrical and Computer Engineering Department, Lehigh University, Bethlehem, PA. His research interests include security and anonymity in networks, information theory, distributed signal processing, and smart energy distribution.

Dr. Venkitasubramaniam received the 2004 Leonard G. Abraham Award from the IEEE Communication Society and a Best Student Paper Award at the 2006 IEEE ICASSP.

Lang Tong (S'87–M'91–SM'01–F'05) received the B.E. degree from Tsinghua University, Beijing, China, in 1985, and the M.S. and Ph.D. degrees in electrical engineering from the University of Notre Dame, Notre Dame, IN, in 1987 and 1991, respectively.

He is the Irwin and Joan Jacobs Professor in Engineering with Cornell University, Ithaca, NY. He was a Postdoctoral Research Affiliate with the Information Systems Laboratory, Stanford University, Stanford, CA, in 1991. He was the 2001 Cor Wit Visiting Professor with the Delft University of Technology, Delft, The Netherlands, and had held visiting positions with Stanford University and the University of California, Berkeley. His research is in the general area of statistical signal processing, wireless communications and networking, and information theory.

Prof. Tong has served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the IEEE TRANSACTIONS ON INFORMATION THEORY, and IEEE SIGNAL PROCESSING LETTERS. He received the 1993 Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 Best Paper Award (with Min Dong) from the IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society (with Parv Venkitasubramaniam and Srihari Adireddy). He is also a coauthor of seven student paper awards. He received the Young Investigator Award from the Office of Naval Research. He was named as a 2009–2010 Distinguished Lecturer by the IEEE Signal Processing Society.