A Framework for Energy-efficient Adaptive Jamming of Adversarial Communications

Jiasi Chen, Soumya Sen, Mung Chiang
Princeton University
Princeton, NJ, USA

David J. Dorsey
Lockheed Martin ATL
Cherry Hill, NJ, USA

***Abstract*—This work proposes a framework for jamming wireless networks that incorporates probabilistic models of internal states and observable characteristics of link protocols, where protocols are divided into two general classes: random access (RA) or channelized access (CA). Without exact knowledge of network parameters and internal state, the proposed intelligent jammer optimizes its strategy to be energy efficient while achieving the target throughput. Probabilistic models for jamming FDMA and CSMA-based protocols are described for illustration of the framework: A frequency-hopping voice network is analyzed to determine the optimal jam strategy for proactive frequency jammer; and a CSMA packet protocol is analyzed for varying packet arrival rates at the nodes. Since RA protocols display observable reaction to channel conditions, we propose a feedback-control loop that uses observable feedback to infer network parameters. Both protocols are evaluated through simulation for their energy-throughput tradeoff compared to a naive jammer.**

I. Introduction

In this paper, we design jamming strategies for FDMA and CSMA-based wireless networks. Our goal is to develop an energy-efficient jamming strategy that is based on practical system capabilities, and robust to protocol modifications and extensions. The proposed jamming framework is posed as an optimization problem using probabilistic models of observed and unobserved states of the target protocol. Fig. 1 shows the setup of our network: a single jammer is attempting to jam a network of adversary nodes.

The ability to model the internal states of adversarial communications is important because one of the most difficult aspects of jamming an adversary network is estimating the effectiveness of the attack. The challenge is due to the fact that the target network parameters that we are attempting to affect (e.g., SINR at the target receiver, throughput of the network, etc.) are not observable in general. If we have models that link observable features to unobserved states, we can provide a feedback loop for optimal jamming. In cases where there are no easily observable features to provide feedback, we must rely on models that estimate the internal state.

Our framework is not designed to exploit particular vulnerabilities of a specific protocol implementation, so it may be applied to a larger set of existing protocols, unlike existing approaches. Brown et al., for example, advocate sensing the network to estimate protocol specifics, then using this information to jam the network [1]. Because our framework does not rely on specific protocol features, the jamming strategies do not require demodulating, decoding, or decrypting messages.

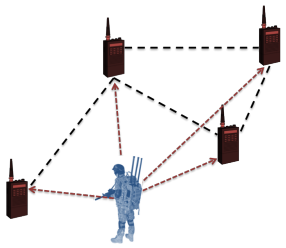

Fig. 1: A single jammer jams communications of an enemy network.

Therefore, they are robust to modifications of the protocol, and do not require sophisticated sensing techniques or special hardware for implementation. Our framework focuses on general access scheme models, and can extended to combinations of the general access scheme models.

The 5-layer network model provides a large design space to design optimal jam strategies. Several possible control knobs are shown in Fig. 2. If we consider each layer of the communication stack to be executing a protocol that optimizes a cost function (e.g., minimize bit error rate), then an optimal jamming strategy should mirror the intent of the target network. The control knob should then be selected to reflect the control at the target stack layer. For example, a distributed *ad hoc* network routing protocol may use local information to select the best set of links; the jamming strategy should mirror this objective by jamming a set of links that maximize packet error [2]. In this effort, we focus on the link layer. Here, the target protocols coordinate the use of a finite resource (time and frequency) in order to minimize collisions or interference among themselves. Thus, our effort focuses on when and at what frequency the jammer should transmit.

At the link layer, there are several types of jammers that may be considered: constant, periodic, deceptive, and reactive [3]. In this work, we develop strategies for an intelligent reactive jammer, since this can best mirror the target protocols while remaining robust to protocol modifications, and offers the greatest potential for energy savings. For our purposes, link layer access schemes are divided into two classes: channelized access (CA) and random access (RA). CA methods include TDMA, FDMA, and CDMA. RA methods include contention-based protocols such as ALOHA and CSMA. In this work, we will analyze frequency-hopping FDMA from the CA class and CSMA from the RA class. Frequency hopping is often

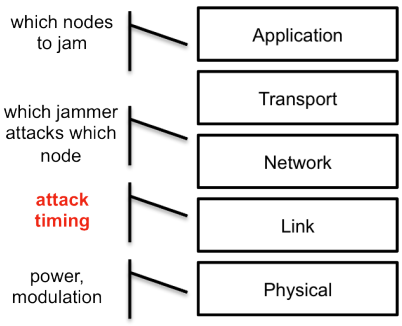

Fig. 2: Control layers.

proposed as a strategy to avoid jammers, and CSMA is arguably one of the most popular wireless access methods.

In the following sections, Section II provides the system model and parameters; Section III and IV analyzes FDMA and CSMA protocols, respectively, and provides analytical and numerical results; and Section V concludes this work.

II. System Model for Adaptive Jamming

A. Energy-constrained Jammer

We consider a single jammer that jams a wireless network, attempting to reduce the network's throughput below a given target. Clearly, if the jammer has unlimited resources, it can jam constantly and easily meet the throughput target. In practical situations, though, the jammer is energy-constrained. Therefore, we formulate the jammer's problem as:

$$\begin{aligned} &\textit{minimize} && \text{jammer power} \\ &\textit{subject to} && \text{network throughput} \leq \text{throughput target} \end{aligned} \tag{1}$$

Our system operates based on information that is easily obtainable in real implementation: the transmit probability and the frequency band occupation of the links.[1]

B. System Parameters

1) Unobservable:

- *Throughput*: The jammer tries to reduce the throughput of the network below target, but it cannot perform any direct over-the-air measurement of the network throughput since packet retransmissions and collisions are not observed. We assume sources and sinks are one-hop apart so that network throughput can be calculated from link throuhgput.
- *Packet arrival rate*: The packet arrival process at the transmitting node cannot be directly observed from the transmit probability. This is due to the nature of RA protocols: the absence of a tranmsission could be due to no packet arrivals, or to the node being in a backoff state.
- *Packet retransmision*: The model should account for retransmissions induced by both collisions and jamming.
- *Collisions*: Simultaneous transmissions from different links can collide with each other and reduce system throughput, thereby reduce the need for the jammer to be active. This collision probability is not easily estimable by the jammer and depends on factors such as the density of nodes in the network and the packet transmission probability of the links.

2) Observable:

- *Link activity*: The jammer can detect when a link is active or inactive.
- *Transmit probability*: The jammer can observe how often links transmit data. For RA protocols, the packet transmission probability of a node depends not only on the effectiveness of jamming and collisions, but also on the packet arrival rate at that node. This probability can be learned from empirical observations of the network over time, and periodically updated.
- *Frequency band occupation*: The jammer can detect on which frequency bands the links are communicating. From empirical observation of the network over time, this can be translated into a probability of a link occupying a given frequency band.

In order to estimate system parameters such as transmnit probability and network topology, we must first isolate and track individual transmissions and separate their transmitters. This is not the focus of this work, but we will briefly outline possible techniques based on work in [4].

We assume that the individual transmitters are distinguishable either through their power profiles or using a direction-finding algorithm. The network topology and link transmit probability parameters can be inferred by tracking the rise and fall times for each packet detection in the network (each associated with a transmitter), and comparing the distributions of durations and inter-arrivals between them to determine a casual relationship between the inter-arrival processes. In the case of CSMA, the inter-arrival distributions between pairs of transmitters that are in direct communication will exhibit lower entropy than transmitter pairs that are not in direct communication. Then, using the Kullback-Liebler distance between each pair of inter-arrival distributions, one can obtain an estimate of the transmitters that constitute links in the network. By isolating these links and observing the activity on each, the probability that a transmission will occur on this link can be estimated. Later through simulation, we will examine the effect of observation error on jammer performance.

Furthermore, it is necessary to classify the target network's channel access scheme. Such a coarse classification can be also accomplished using simple observable features of the communication (frequency, bandwidth, power, on/off times), and computing distributions for inter-arrival times, durations, and frequency occupancy. A representative framework for classifying protocols using these observable parameters and their relationships is discussed in [4].

[1] In this work, a "link" refers to an active communication session between a transmitter node and a receiver node.

III. FDMA Frequency-Hopping Voice Networks

A frequency-hopping model for voice communications is considered. Frequency hopping is a mechanism that helps the network avoid frequency bands with high interference, such as the interference caused by a jammer. This section is devoted to the analysis of such frequency-hopping networks and the optimal strategy of the attacking jammer. We develop a new strategy of pro-active jamming based on the jammer's ability to transmit on multiple frequency bands simultaneously.

The target network operates as follows: A centralized controller assigns active links to a periodic time slot. The transmission frequency of each time slot changes according to a pre-determined schedule. Multiple links, and therefore frequency bands, are active simultaneously during a single time slot. Each link in the network hops randomly to a new frequency band in the next time slot. Collisions may occur between links. This is similar to the standard TDMA/FDMA approach which is used in GSM, for example.

The state of the network is which frequency bands are active per time slot. If the jammer knows the internal state of the network for all time, the problem (1) is trivial since the jammer can simply jam T_o fraction of the active frequency bands in each time slot. The situation is more interesting when the future states of the network are random. In this case, our jammer also adopts a probabilistic strategy.

A. Previous works

Previous work on frequency hopping networks can be divided into two areas: reactive frequency hopping, and proactive frequency hopping. In reactive frequency hopping, the network hops to a new frequency only when it detects an attack. This is the approach proposed by Xu et al. as a general technique to counter jamming denial-of-service attacks [5]. However, jamming detection and hop synchronization are key implementation challenges. For these reasons, proactive frequency hopping has been more often studied and is the network strategy we assume in this work.

In proactive frequency hopping, the network hops frequencies according to a pre-defined sequence known by the transmitter and receiver. The most common jam strategy is repeater jamming, where the jammer simply detects and jams busy frequency bands. Torrieri and Navda both analysed the optimal hop rate to avoid a repeater jammer, based on physical distances and detection and switching delay [6], [7]. Our jamming strategy in turn counters this defensive fast-hopping network by proactively jamming the network.

Pelechrinis et al. considered a game-theoretic model where the jammer and a single link probabilistically transmit on the set of available frequencies[8]. The problem was posed as an optimization game and the equilibrium strategy was found. In our work, we also use a probabilistic model of the jammer and frequency band occupation. Unlike [8], we consider a multi-link network and account for the collisions and interactions between links. We assume the strategy of the network does not change and solve for the jammer's optimal strategy. The previous works all consider narrow-band jamming where the jammer can only transmit on one frequency. In this work, we consider the case of a wideband jammer who can transmit on multiple frequency bands simultaneously.

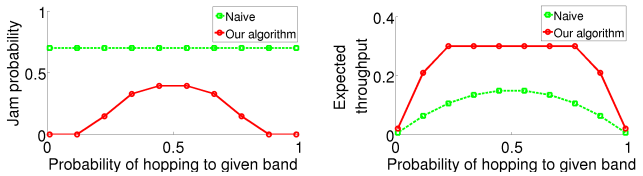

(a) (b)

Fig. 3: Jammer performance in simple case of two frequency bands and two links with identical hop probabilities.

B. Jammer attack strategy

In this section, we derive the optimization problem that the jammer solves in order to find its optimal jamming strategy. Let there be B frequency bands available to the N active links. $P \in \mathbb{R}^{B\times N}$ has entries p_{ij}, the probability that active link j transmits on frequency band i. Let $\mathbf{p}_j$ be the *link band probability*. The jammer determines the probability of jamming each frequency band in order to decrease the throughput of the network. We call this variable $\mathbf{0} \leq \mathbf{q} \leq \mathbf{1} \in \mathbb{R}^B$. T_o is the target throughput. Assuming the power expended by the jammer is proportional to the number of frequency bands jammed, the expected power is $\mathbf{1}^\intercal\mathbf{q}$. The problem (1) is then to minimize $\mathbf{1}^\intercal\mathbf{q}$, subject to throughput constraint $\mathbf{E}[T] \leq T_o$.

We will derive the expression for $\mathbf{E}[T]$ to see that the problem is linear. Let X_l be a binary random variable that is 1 when the transmission on link l is successful, and 0 otherwise. The expected value of X_l and network throughput are:

$$\begin{aligned}\mathbf{E}[X_l] &= \sum_{i=1}^{B} p_{il}(1-q_i)\prod_{j\neq l}(1-p_{ij}) \\ \mathbf{E}[T] &= \frac{1}{N}\mathbf{E}\left[\sum_{l=1}^{N} X_l\right] = \frac{1}{N}\sum_{i=1}^{B}(1-q_i)\sum_{l=1}^{N} p_{il}\prod_{j\neq l}(1-p_{ij})\end{aligned} \tag{2}$$

So the jammer's optimization problem is:

$$\begin{aligned}\underset{\mathbf{q}}{\text{minimize}} \quad & \mathbf{1}^\intercal\mathbf{q} \\ \text{subject to} \quad & \mathbf{c}^\intercal\mathbf{q} \leq T_o + \mathbf{1}^\intercal\mathbf{c} \\ & c_i = -\frac{1}{N}\sum_{l=1}^{N} p_{il}\prod_{j\neq l}(1-p_{ij}) \\ & 0 \leq \mathbf{q} \leq 1\end{aligned} \tag{3}$$

C. Numerical simulations

We compare the performance of our jammer against a naive jammer who chooses $\mathbf{q} = (1-T_o)\mathbf{1}$. This naive jammer always achieves the throughput target, but does not take advantage of the knowledge of each link's frequency band probabilities.

First, a simple example of the optimal jamming strategy is presented. There are two frequency bands and two links with identical band probabilities and $T_o = 0.3$. In Fig. 3a, when the probability of being on a given band is low for both links, the jammer is inactive. This is because the probability of being on

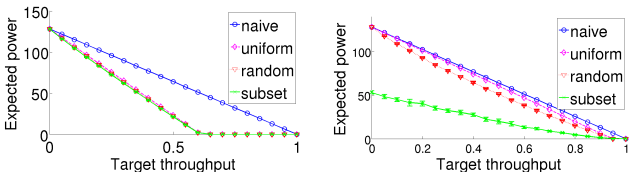

Fig. 4: Throughput-power tradeoff for frequency-hopping network, under various simulation environments.

the other frequency band is high, implying collisions and low throughput which will satisfy the target throughput constraint. As the band probability grows, expected throughput of the network increases, so the jammer must take action. When band probability is high enough, collisions are likely on this link and the jammer no longer needs to be active. The naive jammer, on the other hand, has constant jam probability.

The change of throughput with band probability is shown in Fig. 3b. When the jammer is inactive, the expected throughput varies with the band probability. The jammer activates to cap the throughput of the network at the target throughput of 0.3. The throughput with the naive jammer is lower than the target throughput across all probabilities, but does not take advantage of possible energy savings by lowering the jam probability.

Now we move onto simulations under more realistic environments. In the *uniform* environment, the probability of choosing a frequency band is uniform; in the *random* environment, the probabilities of choosing a frequency band are an arbitrary distribution; in the *subset* environment, each link chooses uniformly from 64 random frequency bands. There are 64 links and 128 possibly frequency bands. Each data point is an average over 10 trials, and the error bars show the standard deviation.

Our algorithm has two major advantages over the naive algorithm: a) knowledge of collisions that decrease network throughput, and b) ability to optimize which frequency bands it transmits over. The first advantage is shown in Fig. 4a through the power-target throughput tradeoff curve. There are 128 bands, 64 active links, and band subset of size 64. When the jammer is inactive, the target throughput is about 0.6, indicating that there are naturally occurring collisions that decrease network throughput for this number of bands and links. Our algorithm uses less energy while meeting the same target throughput as the naive jammer.

The jammer's second advantage of choosing the optimal jam probabilities is illustrated in Fig. 4b. In this setup, there are 128 bands, 8 active links, and band subset of size 8. The jammer is inactive only when the target throughput close to 1, indicating that natural collision probability is low and the performance of the algorithm is due to its selection of the optimal jam probabilities. Our algorithm achieves a better energy target-throughput tradeoff than the naive solution. This is particularly evident in the subset environment, when great energy savings are possible by not transmitting on empty frequency bands.

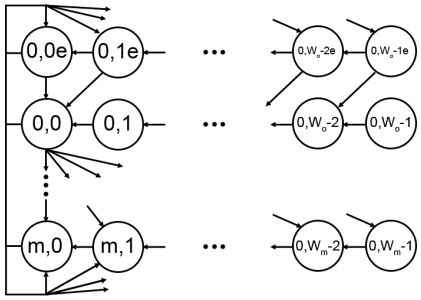

Fig. 5: Markov model of CSMA exponential backoff. Horizontal transitions represent the backoff counter, and vertical transitions represent the number of retransmissions.

IV. CSMA Packet Networks

In this section, we explore the jamming strategy for a packet-based CSMA network with exponential backoff. Unlike the FA protocols, the network in RA protocols reacts to collisions in the environment, including collisions caused by the jammer. This makes the internal state of the network difficult to infer. We propose a feedback-control loop for the jammer to utilize available feedback from the network to determine the network state and the optimal jam strategy.

The CSMA protocol is as follows. A transmitter node first senses the medium is clear before transmitting. A transmission is successful when the transmitter receives an ACK from the receiver. If the transmission is unsuccessful, the transmitter chooses a random backoff time up to a limit, before continuing to sense the medium. Each time the transmission is unsuccessful, the backoff limit increases exponentially.

Bianchi derived an analytic expression for the throughput of a CSMA network [9]. Chinta et al. developed an adaptive jamming strategy based on this formulation [10]. However, this work inherited a key assumption of Bianchi: infinite packet arrival rate at the nodes. Malone et al. analyzed a CSMA network with finite packet arrival rates but did not consider jamming [12]. In this work, in keeping with the theme of practicality, we extend Malone et al. and develop an adaptive jamming strategy that accounts for finite backlog at the nodes.

A. Jammer attack strategy

A Markov model for a CSMA network with homogeneous packet arrival at each node can be constructed, as shown in Fig. 5 [12]. It is hard for the jammer to know the value of each transmitter node's backoff counter, so the internal state is unknown. To address this, the jammer can instead consider the stationary distribution of the Markov chain. A system of nonlinear equations must be solved to obtain the stationary distribution, the transmission probability of each link τ_i, and the collision probability p_i, when a jammer is not present:

$$\begin{aligned} \tau_i &= f(p_i, \lambda_i) \qquad & i = 1, 2, \ldots, N \\ p_i &= 1 - \prod_{j \neq i} (1 - \tau_i) \qquad & i = 1, 2, \ldots, N \end{aligned} \tag{4}$$

where f is given in [12].

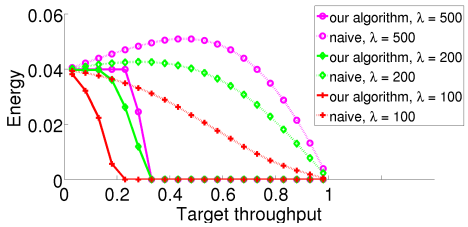

Fig. 6: Throughput-power tradeoff comparison for CSMA network.

The jammer must determine the probability of jamming a transmission q, given that at least one tranmsission occurred [10]. So when the jammer is present, the new probability of collision p_i' depends both on p_i and the jam probability q:

$$\tau_i = f(p_i + q - p_i q, \lambda_i) \quad i = 1, 2, \ldots, N \tag{5}$$

Then the throughput T can be calculated as a function of $\tau_i, p_i', \lambda_i, q$ similar to [12].

The jammer is reactive and only jams when it hears a transmission, either a data packet or an ACK [10], [11]. Thus, the power used by the jammer is proportional to the probability that at least one node transmits, $\left(1 - \prod_i (1 - \tau_i)\right) q$. The problem (1) faced by the jammer is to solve:

$$\begin{aligned} \underset{q}{\text{minimize}} \quad & \left(1 - \prod_i (1 - \tau_i(q))\right) q \\ \text{subject to} \quad & T(q) \leq T_o \end{aligned} \tag{6}$$

where T is obtained by solving (5) and plugging into the formula for T [12]. This problem is nonlinear and must be solved numerically.

B. Numerical simulations

We first show how the network reacts to jam probability and packet arrival rate. Fig. 7a shows the relationship between throughput, jam probability, and packet arrival rate for 10-link network. Fig. 7b shows how power used by the jammer changes with jam probability and packet arrival rate. This figure suggests some non-intuitive behaviour: when q is high, power expenditure is low. This apparent contradiction can be resolved by realizing that when q is high, the network is forced into a state of long backoff, so there are fewer packets for the jammer to jam [10].

We next compare the power-target throughput performance of our algorithm to a naive jamming algorithm. The naive jammer simply jams with $q = (1 - T_o)$. Fig. 10 shows the energy savings of our algorithm for various packet arrivals rates for a 5-link network. The naive jammer's energy usage emulates Fig. 7b. Our intelligent jammer, on the other hand, realizes when the the network's collisions themselves are sufficient to meet the target throughput, and jamming is unnecessary. This transition between high and low energy expenditure is sharp for high arrival rates and smooth for low arrival rates. This can be explained by Fig. 7b: the flatness of the power surface for low arrival rates means that there are more opportunities for the jammer to choose low q and save energy.

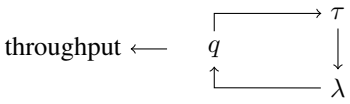

Fig. 9: Relationship between jam probability, transmit probability, and arrival rate in CSMA jammer feedback loop.

C. Feedback-control loop

A key question that arises from the preceding discussion is how to estimate the arrival rate of packets at the nodes. In both RA and CA protocols, the jammer has incomplete knowledge of the throughput of the enemy network. However, in the FDMA protocol, the arrival rate of packets at the node can be directly measured by observing how many times a nodes transmits in the absence of jamming. In a CSMA protocol, though, the jammer does not know the arrival rate. This is because a silent CSMA link could be due to two reasons: (a) the node having no packets in its transmit buffer, or (b) the node being in a backoff state.

The unobservability of the packet arrival rate poses a problem to the previous jamming optimization (6), which can be numerically solved only given known λ. In this section, we discuss how λ can be estimated from observable features of the network such as τ. Since the jammer has incomplete information about the arrival rate, it must adapts to any available information about the network to choose the its jamming strategy. The actions of the jammer in turn create new observations of the network, which force the jammer to update its strategy. Concretely, the jammer chooses q that realizes the throughput target. This q results in some observable transmit probability τ. From τ, the arrival rate λ can be estimated, and a new q chosen. This operational loop is shown in Fig. 9, and a toy example given in Fig. 8:

1) Given the target throughput, find the feasible region of q, as illustrated by the horizontal bars in the left subfigure.
2) From the set of feasible q, find the q that minimizes energy consumed, as shown in the middle subfigure.
3) The network reacts to the system by changing τ, which is observed by the jammer. The jammer uses observed τ and Fig. 7c to estimate the packet arrival rates λ of the network. This is shown in the right subfigure.
4) Using the updated λ, the cycle repeats.

We evaluated the performance of the CSMA jammer feedback loop. First, we examined how mis-estimation by the jammer of node transmit probability affected its arrival rate estimate. Fig. 10a shows how the jammer's arrival rate estimate changes over time based on network feedback. In each iteration, the jammer mis-observes the transmit probability by a maximum of $\pm 20\%$. We see that although the arrival rate estimate never converges to the true arrival rate exactly, it is fairly close. Moreover, not knowing the true arrival rate does not greatly affect jammer performance because the optimal jam strategy is robust to changes in arrival rate. We also swept over different values of transmit probability

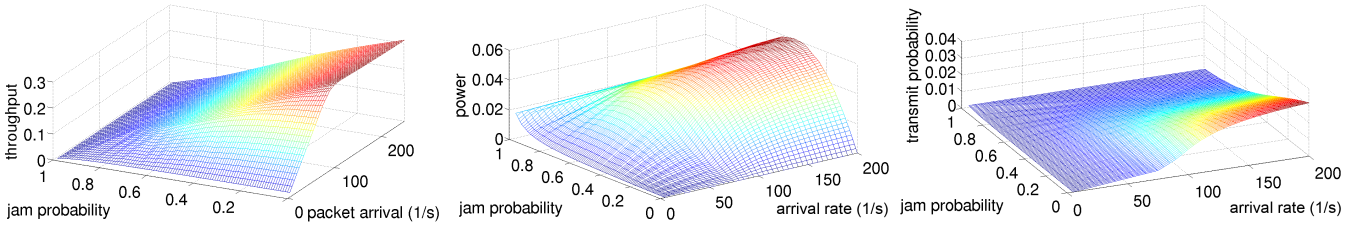

(a) Throughput (b) Power (c) Transmit probability

Fig. 7: Effect of jam probability and packet arrival rate on CSMA network's throughput, power, and transmit probability.

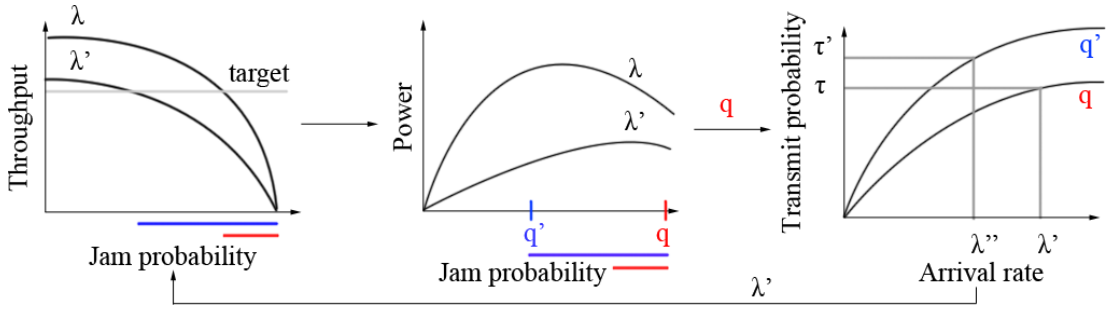

Fig. 8: Example of CSMA jammer feedback loop.

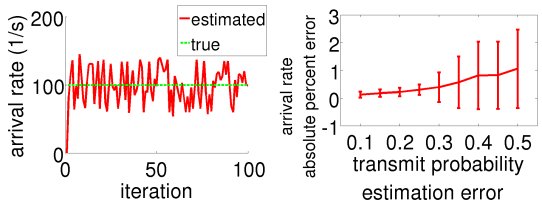

(a) (b)

Fig. 10: Efficacy of feedback loop. (a) Estimated arrival rate with $\pm 20\%$ transmit probability observation error. (b) Arrival rate estimation error increases with transmit probability observation error.

estimation error. As shown in Fig. 10b, both average and standard deviation of arrival rate estimation error increase with transmit probability estimation error. For example, in this 5-node scenario, our results suggest that transmit probability observation error below 0.3 is necessary for accurate arrival rate estimation.

V. Conclusions

In this work, we analyzed frequency-hopping voice networks and CSMA packet-based networks to develop jam strategies for an energy-efficient jammer. We recast the general problem of the jammer for each protocol model. In both cases, the internal state of the network (active frequency bands and backoff counter, respectively) are unknown. To address this in the first case, we optimized for the expected state of the network. In the second case, we solved the optimization problem with respect to the stationary distribution of the states. In both cases, the energy efficiency of our algorithm was demonstrated through simulation versus a naive jammer.

The next step is implementation and evaluation of our system in hardware. Real experiments would address how often probability distributions of link transmission and frequency band occupation should be learned and updated. In addition, there remain many more protocol classes to be analyzed, and future work will also address these.

References

[1] Brown TX, James JE, Sethi A, "Jamming and sensing of encrypted wireless ad hoc networks", *MobiHoc*, 2006.

[2] Tague P, Slater D, Noubir G, Poovendran R "Quantifying the Impact of Efficient Cross-Layer Jamming Attacks via Network Traffic Flows", *IEEE/ACM Trans. on Networks*, 2009.

[3] Xu W, Trappe W, Zhang Y, Wood T, "The feasibility of launching and detecting jamming attacks in wireless networks", *MobiHoc*, 2005.

[4] Rosenbluth D, Dorsey D, Tilghman P, Byron J, Whitaker ET, Trewhitt E, Sardana V, Olivieri M, "Cognitive Representations of Communications for Electronic Warfare", *MILCOM*, 2012.

[5] Xu W, Wood T, Trappe W, Yanyong Z, "Channel surfing and spatial retreats: defenses against wireless denial of service", *Proc. ACM Workshop on Wireless security*, 2004.

[6] Torrieri DH, "Fundamental limitations on repeater jamming of frequency-hopping communications", in *IEEE J. Sel. Areas Comm.*, 1989.

[7] Navda V, Bohra A, Ganguly S, Rubenstein D, "Using Channel Hopping to Increase 802.11 Resilience to Jamming Attacks", *INFOCOM*, 2007.

[8] Pelechrinis K, Koufogiannakis C, Krishnamurthy SV, "On the Efficacy of Frequency Hopping in Coping with Jamming Attacks in 802.11 Networks", in *IEEE Trans. Wireless Comm.*, 2010.

[9] Bianchi G, "Performance analysis of the IEEE 802.11 distributed coordination function", *IEEE J. Sel. Areas Comm.*, 18(3), 2000.

[10] Chinta T, Wong TF, Shea JM, "Energy-efficient jamming attack in IEEE 802.11 MAC", *MILCOM*, 2009.

[11] Zhang Z, Wu J, Deng J, Qiu M, "Jamming ACK Attack to Wireless Networks and a Mitigation Approach", *IEEE GLOBECOM*, 2008.

[12] Malone D, Duffy K, Leith DJ, "Modeling the 802.11 distributed coordination function in non-saturated heterogeneous conditions", *IEEE Trans. Networking*, 15(1), 159-172, 2007.