Quantifying the Impact of Efficient Cross-Layer Jamming Attacks via Network Traffic Flows

Patrick Tague, David Slater, Guevara Noubir, and Radha Poovendran

***Abstract*—We investigate a class of coordinated jamming attacks in which multiple jammers collaboratively apply knowledge about the network layer functionality to efficiently reduce the throughput of network traffic. We show how a constrained optimization framework can be used to characterize coordinated jamming attacks and allow the impact of the attack to be quantified from the perspective of the network. Using this network-centric interpretation of jamming attacks, a network designer can attain a greater understanding of the potential threat of jamming. To illustrate our approach, we propose and evaluate a variety of metrics to model the attack impact, serving both as adversarial objective functions and as network evaluation metrics, and present a simulation study to quantify and compare attack performance.**

***Index Terms*—Jamming, Network flow, Optimization, Resource allocation**

I. Introduction

The open, shared nature of a broadcast communication channel introduces vulnerabilities to wireless networks such as denial-of-service (DoS) attacks [2]. Jamming is a particularly debilitating DoS attack where an adversary transmits interfering signals over the shared medium to block valid communications, for instance by transmitting wide-band noise or high-power narrow-band pulses [3], [4]. Communication systems attempting to perform jamming mitigation typically employ spread-spectrum techniques, forcing the adversary to exhaust significantly more resources by increasing the jamming bandwidth or power required to perform the attack [3]–[5]. Such techniques significantly increase the resource cost required for the jamming attack, leading to effective anti-jamming against resource-constrained jammers.

By incorporating cross-layer information and network communication into the jamming attack, a resource-constrained adversary can significantly increase the efficiency of the attack by targeting specific communication channels, helping to counteract the effect of the anti-jamming systems. Recent work has shown that intelligent jammers can exploit the structure of wireless link layer and MAC protocols [6]–[8] and link layer error correction protocols [9] to jamming attacks that require significantly less energy than jamming continuously or randomly.

P. Tague, D. Slater, and R. Poovendran are with the Network Security Lab (NSL), Electrical Engineering Department, University of Washington, Seattle, Washington. Email: {tague,dmslater,rp3}@u.washington.edu.

G. Noubir is with the College of Computer and Information Science, Northeastern University, Boston, MA, USA. Email: noubir@ccs.neu.edu.

A preliminary version of this material appeared at the 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'08) [1].

In this work, we propose to additionally incorporate information from the network layer into the jamming attack, leading to a further reduction in the required energy resources. By noting that network flows traverse multiple links, the jamming adversary can effectively block an entire network flow [10] by jamming only the link where minimal energy is required. Additionally, since the probability of correct packet decoding is a function of the interference power at the receiver, the adversary can adjust its transmission power to moderate the probability of successfully jamming a packet [4]. This jamming power regulation can be applied independently for each network flow and allows jammers to balance the resource expenditure over multiple flows in trade for a decreased probability of jamming success.

An adversary with a network of jammers can optimize the jamming attack by combining the network-layer information with transmission power regulation, and balancing the jamming workload across the jamming network. Furthermore, to counteract the probabilistic success of jamming packets, the jamming workload can be allocated such that packets unsuccessfully jammed by upstream jammers can be targeting by downstream jammers, given sufficient coverage of the network. Hence, the adversary can optimize a global utility function such as the expected flow rate reduction, energy expenditure, or jamming network lifetime through intelligent assignment of jamming workload and transmission power levels to the jamming networks. We denote this cross-layer DoS attack as a *flow-jamming attack*.

In this article, we quantify the effect of flow-jamming attacks on network performance, and identify crucial concepts which may then be incorporated into network protocol design. We make the following contributions toward this problem.

- We formulate flow-jamming attacks as constrained optimization problems which jointly optimize over jamming transmission power levels and jamming workload allocation to a network of jammers distributed throughout a wireless network.
- We propose a variety of metrics to evaluate the effect of flow-jamming attacks on network traffic and the fractional resource expenditure of the jamming adversary, which double as objective functions for optimization.
- We introduce convex and linear programming relaxations to the optimization framework by decomposing the optimization into two independent optimization problems, thereby enabling efficient computation.
- We propose a cooperative distributed algorithm for flow-jamming attacks for a decentralized jamming network and compare the performance to the centralized approach.

The following outline summarizes the remainder of this article. In Section II, we describe our wireless network model and necessary assumptions. Likewise, in Section III we formulate our adversary model and assumptions. Then in Section IV, we propose evaluation metrics and formulate constraints for the flow-jamming attack framework. We formulate optimal centralized attacks as constrained optimization problems in Section V. We then propose a distributed algorithm for flow-jamming attacks in the absence of centralized adversarial control in Section VI. In Section VII, we evaluate and compare the performance of the centralized and distributed flow-jamming attacks using the proposed metrics. Finally, we summarize our results in Section VIII.

II. Network Model

In this work, we utilize cross-layer information from the network (layer 3), data-link/MAC (layer 2), and physical channel (layer 1). We state our models and assumptions governing these layers, from the top down. A summary of the notation and metrics defined throughout this article is given in Table I.

TABLE I
A SUMMARY OF NOTATION AND METRICS IS PROVIDED.

Symbol	Definition
$\mathcal{N}$	Set of wireless network nodes
$\mathcal{F}$	Collection of network flows
r_f	Packet rate of flow $f \in \mathcal{F}$
$\mathcal{J}$	Set of jammers
E_j	Energy budget for jammer $j \in \mathcal{J}$
d_{jf}	Distance from jammer $j \in \mathcal{J}$ to flow $f \in \mathcal{F}$
P_{jf}	Transmission power used by $j \in \mathcal{J}$ targeting $f \in \mathcal{F}$
$q(d_{jf}, P_{jf})$	Packet error rate (PER) in $f \in \mathcal{F}$ due to $j \in \mathcal{J}$
x_{jf}	Jammer-to-flow assignment for $j \in \mathcal{J}$ and $f \in \mathcal{F}$
$\boldsymbol{P}$	Vector of variables P_{jf}
$\boldsymbol{x}$	Vector of variables x_{jf}
$\pi_f(j)$	Order of jammer j along flow f
$\lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j)$	Resource expenditure of jammer j
$\lambda(\boldsymbol{x}, \boldsymbol{P})$	Resource expenditure metric (18)
$I(\boldsymbol{x}, \boldsymbol{P})$	Jamming impact metric (17)
$G(\boldsymbol{x}, \boldsymbol{P})$	Jamming gain metric (19)
$V(\boldsymbol{x}, \boldsymbol{P})$	Resource variation metric (20)
$\Phi(\boldsymbol{x}, \boldsymbol{P})$	Demand penalty function (21)

A. Network Layer

Let $\mathcal{N}$ denote the set of nodes which make up the wireless network. Traffic flows $f \in \mathcal{F}$ are established throughout the network between source-destination ordered pairs (s, d) via a suitable routing algorithm. For each (s, d) pair, we define an associated traffic flow f with rate r_f packets per second, which we assume is fixed over the duration of interest to the adversary. Without loss of generality, we assume that each flow consists of a single path from the source s to the destination d, as any multiple-path flow can be decomposed into a set of corresponding single-path flows.

B. Data-Link/MAC Layer

We are primarily interested in characterizing the worst-case effect of flow-jamming on the performance of networking protocols, without coupling its effect with that of typical link-layer errors. Therefore, we assume a number of network idealities, which if removed would only increase the impact of the jamming attack, by requiring additional network transmissions. Firstly, we assume that packetization or framing at the data-link layer [11] occurs without error, i.e. that there are no framing errors at each receiving node. Similarly, we assume that the packet transmissions of all flows in $\mathcal{F}$ do not lead to collisions at the MAC layer. Collision-free scheduling can be achieved, for example, by requiring that neighboring nodes transmit on orthogonal communication channels using any of the class of OFDMA protocols [12], which includes TDMA, FDMA, and CDMA as special cases. A network attempting to mitigate jamming is likely to utilize OFDMA or other spread-spectrum techniques to accomplish this goal.

C. Physical Layer

In this section, our goal is to derive a physical layer model that can incorporate a broad range of environments and modulation/coding techniques. We are primarily interested in analyzing the physical layer packet error rate (PER), the probability that a packet is received in error, as the primary purpose of a jamming attack is to drastically increase this quantity.

For each transmitter and receiver, we assume the following physical communication model [4]. Given that the transmitter T and receiver R are separated by a distance d and T transmits with constant power P_T, the received signal power P_R at node R is given by

$$P_R = \rho P_T d^{-\alpha}. \tag{1}$$

The constant ρ in (1) incorporates the antenna gains G_{TR} of T in the direction of R and G_{RT} of R in the direction of T, the transmission wavelength λ, and the constant loss factor L, assumed to be independent of transmit power P_T. The constant α is the path-loss exponent, assumed to be $\alpha \geq 2$, that captures the decay in signal power with distance. Typical values of the path-loss exponent vary in the range $2 \leq \alpha \leq 4$ for open, outdoor environments and in the range $4 \leq \alpha \leq 7$ for constricted, indoor environments.

The PER of the physical layer communication protocol can be analyzed with respect to the signal-to-interference-and-noise ratio (SINR) s given by

$$s = \frac{P_R}{I + N} \tag{2}$$

where I is the interference power and N is the noise power at the receiver. In the absence of jamming, we assume the interference power I is zero and refer to the corresponding quantity P_R/N as the signal-to-noise ratio (SNR). We note that this assumption is equivalent to incorporating ambient interference into the quantity N. Under this model, we assume that each transmitter in the network adjusts its transmit power P_T as a function of the fixed distance d to a sufficient level to maintain an SNR of γ at the receiver. The SINR at the receiver can thus be expressed as

$$s = \frac{\gamma}{1 + I/N}. \tag{3}$$

The probability of packet error (PER) can be computed as a function $q(s)$ of the SINR s. We note that the exact form of the function $q(s)$ depends on the modulation and coding schemes. Examples of PER function $q(s)$ based on Gaussian noise and interference include

$$q(s) = \beta e^{-\xi s} \tag{4}$$

and

$$q(s) = \beta \operatorname{erfc}\left(\sqrt{\xi s}\right), \tag{5}$$

where $\operatorname{erfc}(\cdot)$ is the complementary error function for the Gaussian distribution [4], ξ is a constant depending on the modulation and coding schemes, and β is a constant maximum PER. As a particular example, the PER for packets of length L bits using uncoded BPSK or QPSK modulation under a Gaussian noise model has the form

$$q(s) = 1 - \left(1 - b \operatorname{erfc}\left(\sqrt{\xi s}\right)\right)^L, \tag{6}$$

which is well approximated by the PER function in (5) with $\beta = bL$ for reasonable values of parameters b and L.

Since the constant ξ in (4) and (5) is a scalar coefficient of s, it can be chosen to fit a reference PER. For example, if interference power $I = mP_R$ is sufficient to cause a PER of p, then using (4), we have

$$p = \beta e^{-\xi s} = \beta e^{-\frac{\xi\gamma}{1+I/N}} = \beta e^{-\frac{\xi\gamma}{1+m\gamma}}, \tag{7}$$

and the constant ξ can be parameterized by m and p as

$$\xi = \frac{1+m\gamma}{\gamma} \ln\left(\frac{\beta}{p}\right). \tag{8}$$

III. Adversary Model

Let $\mathcal{J}$ denote a set of jamming nodes which make up the adversarial wireless network. We assume that each jammer $j \in \mathcal{J}$ is able to sense transmissions and infer network flow topology and rates within a particular sensing region around the jammer's location. In addition, we assume that each jammer $j \in \mathcal{J}$ may exchange information with a subset $\mathcal{J}_j \subseteq \mathcal{J}$ of neighboring jammers without interfering with the network of nodes $\mathcal{N}$. This can be achieved by selecting channels orthogonal to those used by the communicating network.

We suppose that each jammer j is constrained by an energy budget E_j, which denotes a finite supply of energy allotted for a particular time interval of attack, since a jamming network unconstrained in terms of energy would be able to brute force jam all nearby flows without the need for any network information. With a constrained energy budget, on the other hand, the jamming adversary will seek an optimal allocation of resources with respect to the jammers' energy budgets and their affect on the underlying communicating network topology, by incorporating cross-layer information and distributing the jamming workload. We thus define the *jammer-to-flow assignment* variable $x_{jf} \in [0,1]$ as the fraction of packets in flow f which jammer j will attempt to jam. Since we assumed that the underlying link layer does not admit collisions, we further assume that it is not possible to simultaneously jam multiple packets from non-interfering links with a single narrow-band jamming transmission. Once the assignment variables x_{jf} have been determined, we assume the jammers can coordinate their jamming transmissions such that neighboring jammers do not simultaneously attempt to jam the same packet in a common flow.

We further suppose that each jammer j can select the jamming transmission power P_{jf} used to jam each packet in flow f, which is a primary parameter in determining the energy exhausted by j. The choice of transmission power P_{jf} determines the PER $q(s)$ (e.g. (4) or (5)) through the received interference power I, which is related according to (1). Since I is inversely proportional to the distance from the jammer to the receiver, we assume that the each jammer chooses to jam a multi-hop flow at the receiver closest to itself, thus maximizing the effect of jamming. This minimum distance is denoted by d_{jf}. Using (1), the interference power I can thus be expressed as $\rho P_{jf} d_{jf}^{-\alpha}$. Combined with (3), this expression yields the SINR s as a function of P_{jf} and d_{jf}, given by

$$s = \frac{\gamma}{1 + \rho P_{jf} d_{jf}^{-\alpha}/N}. \tag{9}$$

The PER $q(s)$ can now be expressed by the function $q(d_{jf}, P_{jf})$ via (9), assuming the jammer can estimate the constants ρ, α, and N. For instance, the PER model in (4) with $\beta = 1$ is given by

$$q(d_{jf}, P_{jf}) = e^{-\frac{\xi\gamma}{1+\rho P_{jf} d_{jf}^{-\alpha}/N}}. \tag{10}$$

This function of d_{jf} and P_{jf} captures the intuitive behavior that the probability of jamming success (PER) increases as a function of transmission power P_{jf} and decreases as a function of distance d_{jf}. Furthermore, given a fixed distance d_{jf}, the function $q(d_{jf}, P_{jf})$ behaves according to a sigmoid, or "s-shaped", function [13], which is useful for optimization techniques.

Given that jammer j is jamming with power P_{jf}, the average energy per packet which is exhausted is equal to a constant cost c times the jamming power P_{jf}, where c may depend on parameters such as modulation and coding schemes, the duty cycle of the jammer, and the packet length. The energy expended on a particular flow is then this quantity times the flow rate r_f and the flow assignment x_{jf}.

For a given set of deployed jammers with energy budgets E_j and set of network flows with rates r_f, the flow-jamming attack is uniquely specified by the jammer-to-flow assignment variables x_{jf} and the jamming transmission powers P_{jf}. Hence, the resource allocation problem of interest is for jammers to collaboratively and optimally determine the assignment variables x_{jf} and transmission powers P_{jf}. An example of a network and jammer topology to illustrate the model is given in Figure 1.

The ability for the jammers to collaboratively optimize depends on their ability to exchange locally sensed information about the network flows as well as their own resource constraints. In the case that the set of neighboring jammers $\mathcal{J}_j$ exchanging information is equal to the entire set $\mathcal{J}$, the optimization problem is essentially a centralized optimal resource allocation problem [14]. Such centralized solutions serve as

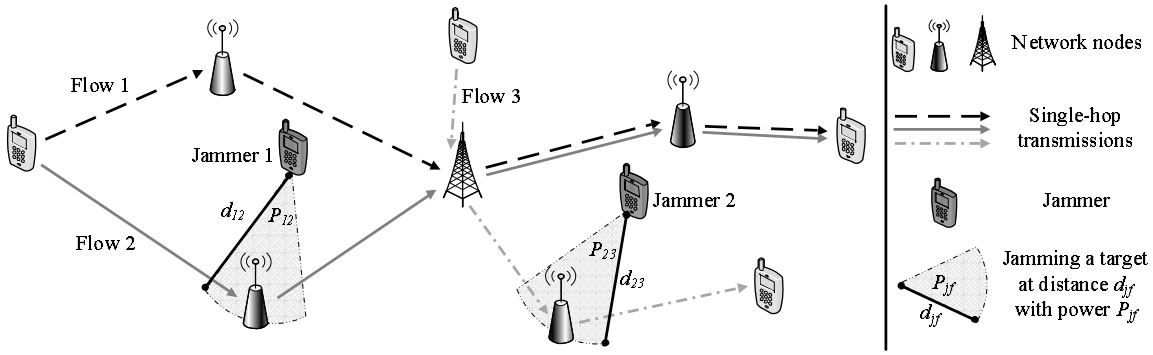

Fig. 1. An example network and jammer topology is illustrated with three network flows and two jammers. Sample jamming options are indicated by the corresponding minimum distance d_{jf} and power P_{jf}.

a performance baseline for comparison and are addressed in Section V for various evaluation metrics introduced in Section IV-B. Noting that extensive jammer communication overhead is counter-productive in practice, we develop a distributed algorithm for flow-jamming in Section VI.

IV. Attack Metrics and Constraints

In order to define the feasible set of allocations for the jamming network, we present a set of constraints which must be satisfied by the flow-jamming attack with respect to the jammer resources and traffic flows. Additionally, we define a set of metrics regarding the effects of the attack in terms of network throughput and jammer resource expenditure, which serve as objective functions for evaluation and optimization of feasible attacks.

We let $\boldsymbol{x}$ and $\boldsymbol{P}$ respectively denote the vectors of jammer-to-flow assignment variables x_{jf} and jamming transmission powers P_{jf}. When convenient, we refer to the sub-vectors of variables $\boldsymbol{x}$ and $\boldsymbol{P}$ corresponding to a single flow f as $\boldsymbol{x}_f$ and $\boldsymbol{P}_f$ and to a single jammer j as $\boldsymbol{x}_j$ and $\boldsymbol{P}_j$. We let $\mathbf{1}$ denote a vector of ones such that $\mathbf{1}^T\boldsymbol{x}$ is equivalent to the ℓ_1 vector sum norm [15] for non-negative $\boldsymbol{x}$.

A. Attack Constraints

We formulate the base constraints on flow-jamming attacks with respect to the jammer-to-flow assignment vector $\boldsymbol{x}$ and power vector $\boldsymbol{P}$. The first pair of constraints follow the definitions of x_{jf} and P_{jf} and restrict the variables to their corresponding domains as

$$0 \leq x_{jf} \leq 1, \tag{11}$$

$$P_{\min} \leq P_{jf} \leq P_{\max} \tag{12}$$

for all $j \in \mathcal{J}$ and $f \in \mathcal{F}$, where $P_{\min}$ and $P_{\max}$ are respectively the minimum and maximum jamming transmission powers.

For a given jammer-to-flow assignment vector $\boldsymbol{x}$ and power vector $\boldsymbol{P}$, it is necessary that the resource of each jammer $j \in \mathcal{J}$ does not exceed their energy budget E_j. Since the jammers may be heterogeneous in terms of their energy budgets, we normalize the resource expenditure to the fraction of their available energy that is exhausted. The resource expenditure is then given by

$$\lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j) = \frac{c}{E_j} \sum_{f \in \mathcal{F}} P_{jf} r_f x_{jf}. \tag{13}$$

The *supply constraint* yielded by these fractional resource expenditures is given by

$$0 \leq \lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j) \leq 1 \tag{14}$$

for all $j \in \mathcal{J}$.

The jamming allocation must additionally satisfy a *flow constraint*, as jammers cannot jam more flow than exists in the network. Since jamming success is probabilistic, we formulate the constraint in the average case by interpreting $q(d_{jf}, P_{jf})$ as the average fraction of jamming attempts which are successful. Since the flow available to downstream jammers is thus dependent on upstream and adjacent jammers, we define the *order* $\pi_f(j)$ of each jammer j along the flow f such that $\pi_f(j_i) < \pi_f(j_k)$ implies that jammer j_i will jam f upstream of (at a node closer to the source than) jammer j_k. Similarly, $\pi_f(j_i) = \pi_f(j_k)$ implies that jammers j_i and j_k will jam the same node in f. We let $\pi_f^{-1}(m)$ denote the set of jammers j with $\pi_f(j) = m$. The desired flow constraint is thus given by

$$\sum_{j \in \pi_f^{-1}(m)} x_{jf} + \sum_{j \in \bigcup_{i<m} \pi_f^{-1}(i)} q(d_{jf}, P_{jf}) x_{jf} \leq 1 \tag{15}$$

for all $f \in \mathcal{F}$ and for each order m. We note that if all jamming success probabilities $q(d_{jf}, P_{jf})$ are forced to 1, then the set of flow constraints given by (15) for a particular flow f and for all m reduces to the single linear constraint $\mathbf{1}^T\boldsymbol{x}_f \leq 1$. Solving for $\boldsymbol{x}$ in this special case is equivalent to finding the optimal partition of the jammer-to-flow assignment.

Finally, we allow for an optional constraint based on the jammer's desire to jam a certain portion of the network traffic. This *demand constraint* on each flow f imposes a lower bound z_f on the expected fraction of throughput reduction as

$$\sum_{j \in \mathcal{J}} q(d_{jf}, P_{jf}) x_{jf} \geq z_f \tag{16}$$

for all $f \in \mathcal{F}$. While the previous constraints always allow a feasible solution, adding the demand constraint further restricts the domain, possibly resulting in an infeasible set of

constraints. We again note that a special case arises when the probabilities $q(d_{jf}, P_{jf})$ are forced to 1. The jammer then has the ability to impose a lower bound of $z_f = 1$ for all flows f, which due to (15) is satisfied only when all flows are completely jammed.

B. Attack Evaluation Metrics

With constraints defining the feasible set, it is natural to define metrics that analyze the effectiveness of a given solution with respect to jammer resource expenditure and network throughput, thus providing objective functions for optimization. If the jamming network is locally optimizing using a distributed protocol, it may result in a solution that violates the overall flow or demand constraints, resulting in undefined metrics. We consider this problem of over-allocation alongside our distributed algorithm in Section VI.

To begin, we evaluate the extent to which the jammers in $\mathcal{J}$ are able to reduce network throughput by defining the *jamming impact* $I(\boldsymbol{x}, \boldsymbol{P})$, the average fractional throughput reduction, as

$$I(\boldsymbol{x}, \boldsymbol{P}) = \frac{1}{|\mathcal{F}|} \sum_{\substack{j \in \mathcal{J} \\ f \in \mathcal{F}}} q(d_{jf}, P_{jf}) x_{jf}. \tag{17}$$

The average resource expenditure for the set $\mathcal{J}$ of jammers is defined as the fraction $\lambda(\boldsymbol{x}, \boldsymbol{P})$ given by

$$\lambda(\boldsymbol{x}, \boldsymbol{P}) = \frac{1}{|\mathcal{J}|} \sum_{j \in \mathcal{J}} \lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j). \tag{18}$$

By combining (17) and (18), we can evaluate the overall jamming impact per unit of resource expenditure. Thus, we define the *jamming gain* $G(\boldsymbol{x}, \boldsymbol{P})$ equal to the ratio of jamming impact to resource expenditure as

$$G(\boldsymbol{x}, \boldsymbol{P}) = \frac{I(\boldsymbol{x}, \boldsymbol{P})}{\lambda(\boldsymbol{x}, \boldsymbol{P})}. \tag{19}$$

The jamming gain serves as a basis for optimizing the resource efficiency of the attack, instead of optimizing for resource expenditure or impact alone.

The above metrics reflect the average behavior over the set $\mathcal{J}$ of jammers. However, in order to maximize the lifetime of the jamming network [16], we are interested in measuring the ability to fairly distribute the resource expenditure among the jammers. We thus define the *jamming resource variation* $V(\boldsymbol{x}, \boldsymbol{P})$ equal to the relative difference between the maximum and minimum resource expenditure as

$$V(\boldsymbol{x}, \boldsymbol{P}) = 1 - \frac{\min_j \lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j)}{\max_j \lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j)}, \tag{20}$$

falling in the range $[0, 1]$. Large variation indicates that some jammers are fully exhausting their energy budgets while others have unallocated energy resources due to network geometry. Small variation implies a balance of relative energy expenditure and allows for prolonged minimum jammer lifetime and flow-jamming attack duration.

As an alternative to imposing the demand constraint (16) for all $f \in \mathcal{F}$, we can incorporate a penalty function $\Phi(\boldsymbol{x}, \boldsymbol{P})$ into the metrics which measures the degree of violation of the demand constraint. This also ensures the existence of a feasible solution. The penalty function is thus defined as

$$\Phi(\boldsymbol{x}, \boldsymbol{P}) = \sum_{f \in \mathcal{F}} \left(z_f - \sum_{j \in \mathcal{J}} q(d_{jf}, P_{jf}) x_{jf} \right)^+, \tag{21}$$

where $(y)^+ = \max(0, y)$. The demand constraint (16) can be written equivalently in terms of the penalty function as $\Phi(\boldsymbol{x}, \boldsymbol{P}) = 0$.

V. Flow-Jamming Attack Formulations

In order to quantify the effect of worst-case flow-jamming attacks on network traffic, we formulate these attacks as constrained optimization problems subject to the constraints given in Section IV-A, using the evaluation metrics presented in Section IV-B as objective functions. The general formulation yields a non-convex optimization problem, which demonstrates the upper bound on the effectiveness of the attack. We then demonstrate a two-step convex optimization relaxation, which yields a real-time approximation of the desired solution by first solving for transmission powers $\boldsymbol{P}$ and then optimizing the assignment $\boldsymbol{x}$ with respect to the chosen $\boldsymbol{P}$. We discuss the trade-offs between optimality of solutions and associated computational overhead in Section VII.

A. Optimal Flow-Jamming Attacks

Suppose the adversary is interested in maximizing an objective function $g(\boldsymbol{x}, \boldsymbol{P})$ that measures the performance of the jamming attack (e.g. gain, negative resource expenditure, etc.) and minimizing the penalty function $\Phi(\boldsymbol{x}, \boldsymbol{P})$. In order to determine the optimal solution $(\boldsymbol{x}^*, \boldsymbol{P}^*)$ subject to the constraints outlined in Section IV-A, the adversary must solve the constrained optimization problem

$$\boxed{\begin{array}{c} (\boldsymbol{x}^*, \boldsymbol{P}^*) = \underset{\boldsymbol{x}, \boldsymbol{P}}{\arg\max}\, g(\boldsymbol{x}, \boldsymbol{P}) - \Delta\Phi(\boldsymbol{x}, \boldsymbol{P}) \\ \text{s.t. } (11), (12), (14), \text{ and } (15) \end{array}} \tag{22}$$

where $\Delta \geq 0$ is a weight on the penalty $\Phi(\boldsymbol{x}, \boldsymbol{P})$ for violating the demand constraint (16). The objective and penalty functions are optimized simultaneously, with Δ determining the relative importance of the two functions. A large Δ value would seek primarily to satisfy the demand constraint, using $g(\boldsymbol{x}, \boldsymbol{P})$ as a secondary criterion, with the opposite true for small Δ values. We present the following objective functions $g(\boldsymbol{x}, \boldsymbol{P})$ of interest using the metrics defined in Section IV-B.

Case 1 - Maximum Impact Attack: To maximize the overall impact without a demand penalty, the adversary can set $g(\boldsymbol{x}, \boldsymbol{P}) = I(\boldsymbol{x}, \boldsymbol{P})$ with $\Delta = 0$. Alternatively, to impose a demand penalty, the adversary can set $\Delta > 0$ and $z_f > 0$ for at least one flow f, noting that the penalty is assessed individually for each flow f, while the overall impact is averaged over all $f \in \mathcal{F}$.

Case 2 - Minimum Resource Attack: To minimize the resource expenditure $\lambda(\boldsymbol{x}, \boldsymbol{P})$ required to meet the demand constraint (16), the adversary can set $g(\boldsymbol{x}, \boldsymbol{P}) = -\lambda(\boldsymbol{x}, \boldsymbol{P})$ and set $z_f > 0$ for at least one flow f to prevent the trivial solution $\boldsymbol{x} = 0$. For sufficiently large Δ, the optimization of

this problem will focus on meeting the demand constraint with the minimum resource expenditure.

Case 3 - Maximum Gain Attack: The goals of maximizing impact and minimizing resource expenditure can be combined in maximizing the jamming gain $g(\boldsymbol{x}, \boldsymbol{P}) = G(\boldsymbol{x}, \boldsymbol{P})$. Since maximizing the gain does not necessarily lead to a high jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$, demand penalties with $\Delta > 0$ and $z_f > 0$ for at least one flow f can be imposed.

Case 4 - Minimum Variation Attack: To balance the resource expenditure over the set of jammers $\mathcal{J}$, the adversary can set $g(\boldsymbol{x}, \boldsymbol{P}) = V(\boldsymbol{x}, \boldsymbol{P})$ and set $z_f > 0$ for at least one flow f to prevent the trivial solution $\boldsymbol{x} = 0$. The optimal solution will attempt to meet the demand constraint with the optimal balance of resources over the jammers.

B. Computational Considerations and Attack Approximations

By inspection of the constraints and objective functions, we see that the flow constraint, demand constraint, and evaluation metrics are dependent on $q(d_{jf}, P_{jf})$, which is sigmoidal in P_{jf}. Hence, the general optimization formulation in (22) is a non-convex optimization problem [17], and the determination of the optimal solution $(\boldsymbol{x}^*, \boldsymbol{P}^*)$ must thus rely on heuristic methods. The implication is that only local optimums are guaranteed, and the computation time for finding these points may prohibit real-time jamming attacks.

Given the complications involved in solving non-convex optimization problems, we present an alternative convex optimization formulation. The goal of our approach is to decompose the joint optimization of variables $\boldsymbol{x}$ and $\boldsymbol{P}$ into two independent optimization problems. Fixing $\boldsymbol{P}$ yields linearity in the constraints (14), (15), and (16), and convexity in the objective and penalty functions (17), (18), and (21), with respect to the optimization variable $\boldsymbol{x}$. Hence, the two-stage decomposition is given by

$$\boxed{\begin{aligned} P_{jf}^* &= \arg\max_{P_{jf}} h(P_{jf}) \\ &\quad \text{s.t. } (12) \\ \\ \boldsymbol{x}^* &= \arg\max_{\boldsymbol{x}} g(\boldsymbol{x}, \boldsymbol{P}^*) - \Delta\Phi(\boldsymbol{x}, \boldsymbol{P}^*) \\ &\quad \text{s.t. } (11), (14), \text{ and } (15), \end{aligned}} \tag{23}$$

which can be efficiently solved when the objective function $h(P_{jf})$ is convex in P_{jf} and $g(\boldsymbol{x}, \boldsymbol{P}^*)$ is convex in $\boldsymbol{x}$. For $\boldsymbol{x}$, the metric of jamming gain $G(\boldsymbol{x}, \boldsymbol{P})$ is given by ratio of two linear functions in $\boldsymbol{x}$, and can thus be formulated through linear-fractional transformation as a convex optimization problem [17]. The metric of resource variation in (20) is generally non-convex, and motivates the need for a similar convex objective.

Attacks aiming to minimize the resource variation $V(\boldsymbol{x}, \boldsymbol{P})$ can be approximated by an alternative formulation which simultaneously minimizes the maximum resource expenditure and maximizes the minimum resource expenditure for each jammer j, effectively tightening the upper and lower bounds on each $\lambda_j(\boldsymbol{x}, \boldsymbol{P})$. Hence, instead of minimizing the variation $V(\boldsymbol{x}, \boldsymbol{P})$, we can instead introduce a variable upper bound λ_U and lower bound λ_L and minimize the difference $\lambda_U - \lambda_L$ with the modified supply constraint

$$\lambda_L \leq \lambda_j(\boldsymbol{x}_j, \boldsymbol{P}_j) \leq \lambda_U, \tag{24}$$

for all $j \in \mathcal{J}$, with the bounding constraints $0 \leq \lambda_L \leq 1$ and $0 \leq \lambda_U \leq 1$.

Given that the attack formulation in (23) for the given value of $\boldsymbol{P}^*$ is a convex optimization problem, the remaining piece is to define the objective function $h(P_{jf})$ in order to aptly fix the value $\boldsymbol{P}^*$. In what follows, we discuss two heuristic methods that can be used to define this objective function independently of the optimization variable $\boldsymbol{x}$.

1) Decomposition for High-Gain Attack: The first approach seeks to find $\boldsymbol{P}$ that will maximize the gain of the flow-jamming attack independently of $\boldsymbol{x}$. For a single jammer j and flow f, the contribution I_{jf} toward the impact $I(\boldsymbol{x}, \boldsymbol{P})$ of the flow-jamming attack can be expressed as

$$I_{jf} = x_{jf} q(d_{jf}, P_{jf}), \tag{25}$$

where d_{jf} is constant. Likewise, the fraction λ_{jf} of resources used by jammer j to jam packets in flow f can be written as

$$\lambda_{jf} = c P_{jf} r_f x_{jf} / E_j. \tag{26}$$

Thus, the individual gain of the single jammer-to-flow assignment can be measured by the ratio I_{jf}/λ_{jf}, which does not depend on x_{jf}. We thus define the objective function $h(P_{jf})$ as

$$h(P_{jf}) = \frac{E_j q(d_{jf}, P_{jf})}{c r_f P_{jf}}. \tag{27}$$

The maximum value of the objective function $h(P_{jf})$ for each (j, f) pair can thus be individually chosen independent of the assignment variable x_{jf}. Due to the sigmoidal shape of the PER function $q(d_{jf}, P_{jf})$, the optimal power P_{jf}^* can be determined analytically [18]. Using these values for $\boldsymbol{P}^*$ yields the optimal gain for a single jammer and flow, though it is not necessarily optimal when multiple jammers and flows are considered. It does, however, allow the assignment variables x_{jf} to be efficiently optimized using (23).

2) Decomposition for High-Impact Deterministic Attack: A second approach is to fix $\boldsymbol{P}$ to maximize the impact of the attack for each node-flow pair. If $P_{\max}$ is sufficiently large, the PER $q(d_{jf}, P_{jf})$ can be made close enough to 1 that the margin $1 - q(d_{jf}, P_{jf})$ is negligible. In this case, the choice of objective function $h(P_{jf}) = P_{jf}$ yields the optimal value $P_{jf}^* = P_{\max}$, allowing for a deterministic flow-jamming attack in which all probabilities $q(d_{jf}, P_{jf})$ are assumed to be 1. In this case, as mentioned in Section IV-A, the set of non-convex flow constraints (15) for flow f and all orders m reduces to the single linear constraint $\mathbf{1}^T \boldsymbol{x}_f \leq 1$. Additionally, this deterministic attack allows the adversary to achieve $z_f = 1$ in the demand constraint (16) given sufficient resources. This greedy demand constraint simplifies the penalty function, as the the argument y to the penalty function $(y)^+$ is always non-negative, implying that $(z)^+ = z$, which results in a linear program. Furthermore, a sufficiently high Δ forces the adversary to maximize the jamming impact for all jammers, yielding a corresponding heuristic for high-impact attacks.

VI. Distributed Flow-Jamming Attacks

As the size of the jamming network increases, the communication and computational overhead of calculating and coordinating the centralized attack algorithm becomes prohibitively expensive. Therefore, in this section we present a distributed flow-jamming attack in which each jammer j computes the jamming transmission power P_{jf} and assignment variable x_{jf} using only local information. Since the optimality of the attack is dependent on the amount of information available, which in turn is dependent on the size chosen for the local region, there is an inherent trade-off between communication overhead and attack performance. However, increasing the required amount of communication may reduce the responsiveness of the jamming attack. In order to demonstrate the feasibility of this attack for resource-constrained jammer networks, we present the case where each jammer j computes only the local variables $\boldsymbol{x}_j$ and $\boldsymbol{P}_j$ using messages received from jammers in the neighborhood $\mathcal{J}_j = \bigcup_f \pi_f^{-1}(\pi_f(j))$ of jammers targeting the same nodes as j.

Prior to the jamming attack commencing, each jammer senses its local region, determining the initial traffic flow rates r_f as well as the jamming costs of nodes incident to those flows of interest $\mathcal{F}_j \subseteq \mathcal{F}$, and exchanges this information with its neighborhood. After the attack has begun, a downstream jammer j may sense a packet rate at the nearest node in flow f which is less than the initial packet rate r_f. We thus assume that each jammer j occasionally senses this *residual packet rate*, and we let $r_{jf}^{(t)}$ denote the sampled value of this rate at time t. We assume that two jammers j_1 and j_2 with $\pi_f(j_1) = \pi_f(j_2)$ will sense the same rates $r_{j_1 f}^{(t)} = r_{j_2 f}^{(t)}$ at each time t. We additionally introduce the projected flow rate $\hat{r}_{jf}$ which acts as an estimate of the value of the residual flow $r_{jf}^{(t)}$ after jamming, which is determined from past allocations. Note that initially, $r_f = r_{jf}^{(0)} = \hat{r}_{jf}$.

In this distributed attack formulation, each jammer j individually solves a constrained optimization problem involving only their local view. Since this step involves no coordination with neighboring nodes, this may result in jamming allocation conflicts, where neighboring jammers assignments $\boldsymbol{x}$ sum to more than unity, resulting in sub-optimal performance. To compensate for this, each jammer j exchanges $\boldsymbol{P}_j$ and $\boldsymbol{x}_j$ with its neighbors, and an additional optimization problem is solved for each set of jamming allocation conflicts.

In order to formulate the constrained optimization problem, it is first necessary to reformulate the objective function $g(\boldsymbol{x}, \boldsymbol{P})$, the penalty function $\Phi(\boldsymbol{x}, \boldsymbol{P})$, and the constraints in the scope of a single jammer. Since resource variation is not well-defined in a single jammer context, we focus only on the metrics of jamming impact and jamming gain. The objective function $g(\boldsymbol{x}, \boldsymbol{P})$ and penalty function $\Phi(\boldsymbol{x}, \boldsymbol{P})$ described in Section V-A can be similarly re-defined as $g_j(\boldsymbol{x}, \boldsymbol{P})$ and $\Phi_j(\boldsymbol{x}, \boldsymbol{P})$ for each jammer j with respect to the local flows $\mathcal{F}_j$ and itself. Since only the jammer j computes only its own allocation, the allocation constraint (11) and flow constraint (15) can be combined by incorporating the projected flow rate $\hat{r}_{jf}$ for each flow f as

$$0 \leq x_{jf} \leq \frac{\hat{r}_{jf}}{r_f}. \tag{28}$$

The supply constraint (14) remains unchanged. Use of the simplified formulation yields the (non-convex) optimization problem

$$\boxed{\begin{aligned}(\boldsymbol{x}_j^*, \boldsymbol{P}_j^*) = \underset{\boldsymbol{x}_j, \boldsymbol{P}_j}{\arg\max}\, & g_j(\boldsymbol{x}_j, \boldsymbol{P}_j) - \Delta\Phi_j(\boldsymbol{x}_j, \boldsymbol{P}_j) \\ \text{s.t. } & (12), (14), \text{ and } (28).\end{aligned}} \tag{29}$$

We again consider the alternative convex formulation in Section V-B allowing each jammer j to solve the two-stage decomposition

$$\boxed{\begin{aligned}P_{jf}^* = \underset{P_{jf}}{\arg\max}\, & h(P_{jf}) \\ \text{s.t. } & (12) \\ \\ \boldsymbol{x}_j^* = \underset{\boldsymbol{x}_j}{\arg\max}\, & g_j(\boldsymbol{x}_j, \boldsymbol{P}_j^*) - \Delta\Phi_j(\boldsymbol{x}_j, \boldsymbol{P}_j^*) \\ \text{s.t. } & (14) \text{ and } (28).\end{aligned}} \tag{30}$$

After the optimization has been solved for all jammers in $\mathcal{J}$, each jammer j shares its allocations with its neighbors $\mathcal{J}_j$. For each flow $f \in \mathcal{F}_j$, if the subset $\pi_f^{-1}(\pi_f(j)) \subseteq \mathcal{J}_j$ of jammers violate the local flow constraints

$$\sum_{k \in \pi_f^{-1}(\pi_f(j))} x_{kf} \leq \frac{r_{jf}^{(t)}}{r_f} \tag{31}$$

at the corresponding target node, the conflict is resolved as follows. For each $k \in \pi_f^{-1}(\pi_f(j))$, let $e_{kf} = cP_{kf}^* r_{kf}^{(t)} x_{kf}^*$ denote the energy allocated to flow f. Each conflicting jammer k then simultaneously uses the exchanged information to solve the sub-problem

$$\boxed{\begin{aligned}(\boldsymbol{x}_f^*, \boldsymbol{P}_f^*) = \underset{\boldsymbol{x}_f, \boldsymbol{P}_f}{\arg\max}\, & \sum_{k \in \pi_f^{-1}(\pi_f(j))} q(d_{kf}, P_{kf}) x_{kf} \\ \text{s.t. } & (12), (31) \\ & cP_{kf} r_{kf}^{(t)} x_{kf} \leq e_{kf} \text{ for each } k \in \pi_f^{-1}(\pi_f(j)) \\ & 0 \leq x_f\end{aligned}} \tag{32}$$

after which jamming commences. At the next update time t after δ seconds, jammers sense the current residual flow rate $r_{jf}^{(t)}$ and use the previously sensed rate $r_{jf}^{(t-\delta)}$ to compute the projected flow for the next set of optimizations as

$$\hat{r}_{jf} = \begin{cases} r_{jf}^{(t)} + r_{jf}^{(t-\delta)} \displaystyle\sum_{k \in \mathcal{J}_j \setminus \{j\}} x_{kf} & \text{if } r_{jf}^{(t)} \leq r_{jf}^{(t-\delta)} \\ r_{jf}^{(t)} \left(1 - \displaystyle\sum_{k \in \mathcal{J}_j \setminus \{j\}} x_{kf}\right) & \text{else.} \end{cases} \tag{33}$$

Thus, each jammer j performs the distributed flow-jamming attack via the following iterative algorithm.

1) At time $t = 0$, initialize and exchange the distances d_{jf} with jammers in $\mathcal{J}_j$. Set the projected flow rates to $\hat{r}_{jf} = r_{jf}^{(0)} = r_f$.

2) Exchange the projected flow rates $\hat{r}_{jf}$ with neighbors and solve the optimization problem in (29) or (30) using the local information.
3) Exchange the resulting values $\boldsymbol{x}_j^*$ and $\boldsymbol{P}_j^*$ with jammers in $\mathcal{J}_j$.
4) Resolve any allocation conflict by solving (32).
5) Jam using the computed parameters $\boldsymbol{x}_j^*$ and $\boldsymbol{P}_j^*$.
6) At the next update time t, sense the new residual flow rate $r_{jf}^{(t)}$ and compute the new projected flow rate using (33), then return to step 2.

We note that certain special cases of the objective functions $h(P_{jf})$ and $g_j(\boldsymbol{x}, \boldsymbol{P})$ yielding linear programming or convex optimization problems may allow for Lagrangian dual decomposition methods to be used [19], though we do not address these special cases in this work. A comparison of the proposed distributed algorithm to the centralized formulation is given in the next section.

VII. Performance Evaluation

In this section, we evaluate the performance of the flow-jamming attacks using the metrics presented in Section IV-B, the centralized optimization problem in Section V, and the distributed optimization formulation in Section VI. We first describe the simulation setup and then illustrate and compare the performance of the centralized and distributed flow-jamming attacks using various objective functions.

A. Simulation Setup

We use the following setup to obtain our simulation results. A network comprising the set $\mathcal{N}$ of nodes is randomly deployed over a given area, and a link is formed between any pair of nodes within a fixed communication range. Each network flow $f \in \mathcal{F}$ is formed between a randomly selected source s and destination d using a randomized geometric routing algorithm that chooses the next hop toward d from the set of neighbors that are closer to d in terms of either distance or hop count. As discussed in Section II-C, we assume that each transmitting node sets its transmission power to maintain an SNR of γ at the receiving node.

A network comprising the set $\mathcal{J}$ of jammers is deployed over the same area as the nodes in $\mathcal{N}$, and the neighboring subsets $\mathcal{J}_j$ and $\mathcal{F}_j$ for each jammer j are determined by fixed communication and sensing ranges. The path-loss constant ρ and exponent α were set using measurements made in an open environment. We further assume that the interference model is given by the PER function in (10), with the constant ξ chosen using the reference parameters m and p in (8). Table II summarizes the default parameter values used in our simulation study, noting that specified parameters are varied in certain cases.

B. Comparison of Optimization Formulations

We first evaluate the performance of the centralized attack formulation in (22) using the objective functions in Cases 1-4 in Section V-A. We illustrate the results of a single network and jammer deployment with each of the four attacks performed using the same data set. Figure 2 illustrates the four values for each metric $I(\boldsymbol{x}, \boldsymbol{P})$, $\lambda(\boldsymbol{x}, \boldsymbol{P})$, $G(\boldsymbol{x}, \boldsymbol{P})$ and $V(\boldsymbol{x}, \boldsymbol{P})$, with the results of each metric normalized with respect to the largest result for ease of comparison.

As seen in Figure 2, each attack optimizes the corresponding metric in trade for weaker performance in terms of other evaluation metric. The maximum impact attack in Case 1 yields the highest impact $I(\boldsymbol{x}, \boldsymbol{P})$ but also has the highest resource expenditure $\lambda(\boldsymbol{x}, \boldsymbol{P})$. The minimum resource attack in Case 2 yields the lowest resource expenditure $\lambda(\boldsymbol{x}, \boldsymbol{P})$ but also has the lowest impact $I(\boldsymbol{x}, \boldsymbol{P})$. The maximum gain attack in Case 3 balances the trade-off between impact and resource expenditure, yielding an increase in both metrics over the minimum resource attack in Case 2. We note that all of Cases 1-3 achieve their maximum values by allowing a subset of the jammers to do a majority of the work, leading to a high resource variation $V(\boldsymbol{x}, \boldsymbol{P})$. The minimum variation

TABLE II
THE PARAMETERS IN OUR SIMULATION STUDY ARE SUMMARIZED.

Parameter	Value
Network area	$100\ m \times 100\ m$
Number of nodes	$\lvert\mathcal{N}\rvert = 200$
Network radio range	$20\ m$
Number of network flows	$\lvert\mathcal{F}\rvert = 20$
Flow rates	$r_f = 1500\ pkts/sec$
Signal-to-noise ratio	$\gamma = 5$
Path-loss constant	$\rho = 2.5 \times 10^{-4}$
Path-loss exponent	$\alpha = 2.7$
Receiver noise	$N = 1\ nW$
PER parameters	$m = 10$, $p = 0.9$
Number of jammers	$\lvert\mathcal{J}\rvert = 10$
Jammer radio and sensing ranges	$50\ m$
Jammer energy supply	$E_j = 10\ mJ$
Jammer cost coefficient	$c = 10^{-6}$
Minimum jamming power	$P_{\min} = 0\ mW$
Maximum jamming power	$P_{\max} = 500\ mW$
Jamming demand	$z_f = 1/2$
Penalty coefficient	$\Delta = 10^4$

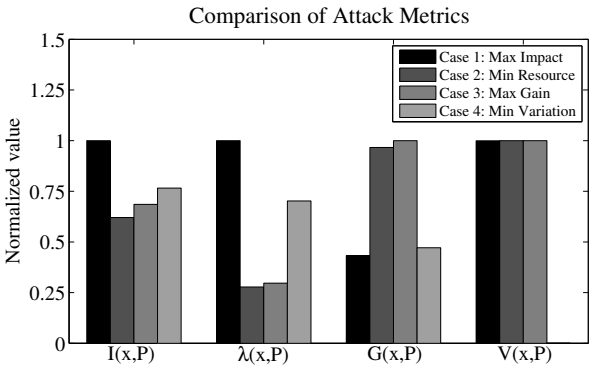

Fig. 2. The centralized flow-jamming attacks in Cases 1-4 presented in Section V-A are simulated. The metrics of jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$, resource expenditure $\lambda(\boldsymbol{x}, \boldsymbol{P})$, gain $G(\boldsymbol{x}, \boldsymbol{P})$, and resource variation $V(\boldsymbol{x}, \boldsymbol{P})$ are illustrated for each attack. The value of each metric is normalized by the maximum for each group.

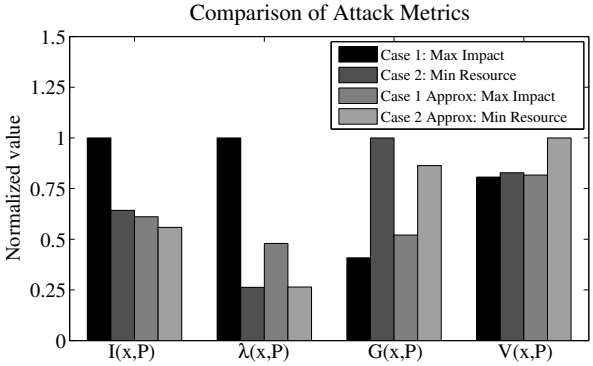

Fig. 3. The centralized flow-jamming attacks in Cases 1-2 presented in Section V-A are simulated using both the non-convex formulations and the convex approximations, requiring respective computational run-time of 3178 seconds, 2784 seconds, 0.7 seconds, and 0.4 seconds. The metrics of jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$, resource expenditure $\lambda(\boldsymbol{x}, \boldsymbol{P})$, gain $G(\boldsymbol{x}, \boldsymbol{P})$, and resource variation $V(\boldsymbol{x}, \boldsymbol{P})$ are illustrated for each attack. The value of each metric is normalized by the maximum for each group.

attack in Case 4 balances this workload and in a further trade-off between impact and resource expenditure, reducing the jamming gain but increasing the impact.

We next compare the non-convex formulation with the performance of the centralized convex attack formulation in (23) using the two-stage decomposition with the objective functions in Cases 1-2. Figure 3 illustrates the normalized results of these cases. As can be seen, the attacks using convex formulations yield reasonable approximations of the objective attained by the centralized solution. We note that the computational run-times of 3178 seconds for Case 1 and 2784 seconds for Case 2 are several orders of magnitude greater than the run-times of 0.7 seconds and 0.4 seconds required for the convex approximations. The simulated results and corresponding run-times for the non-convex formulations were obtained using the `glcSolve` solver in the TOMLAB Optimization Environment [20], and the results and run-times for the convex were obtained using the `CVX` package for solving convex optimization problems [21]. These computation times demonstrate that real-time jamming attacks using the non-convex attack formulations are likely impractical. We thus focus our attention on the use of convex approximations for the remainder of this simulation study.

C. Effect of Parameter Variation

In order to quantify the effect of parameter variation on flow-jamming attack performance, we next simulate attacks using the convex approximation and distributed formulations over a range of parameters. We illustrate the effect of varying the number of jammers, the total jamming energy, the number of traffic flows, and the path-loss exponent α on the resulting jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$.

To evaluate the effect of the size of the adversarial network on jamming efficiency, we vary the number of jammers $|\mathcal{J}|$ while maintaining a fixed total jamming energy $\sum_{j\in\mathcal{J}} E_j$. Figure 4 illustrates the jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ as a function

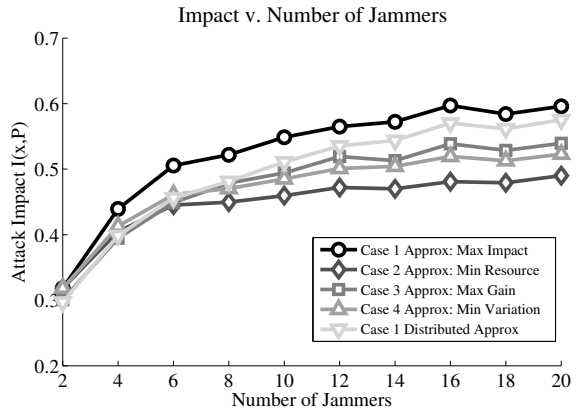

Fig. 4. The jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ resulting from each of the centralized convex and distributed flow-jamming attacks is simulated for various numbers of jammers $|\mathcal{J}|$, keeping the total jamming energy $\sum_{j\in\mathcal{J}} E_j$ constant.

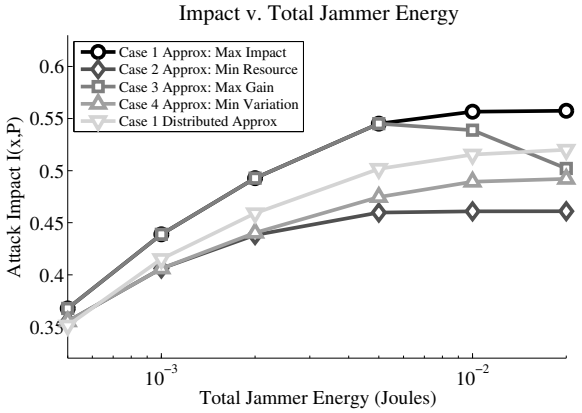

Fig. 5. The jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ resulting from each of the centralized convex and distributed flow-jamming attacks is simulated for various values of the total jamming energy $\sum_{j\in\mathcal{J}} E_j$, keeping all other network and jamming parameters constant.

of the number of jammers for attacks using the convex approximations of Cases 1-4 and the distributed algorithm in Section VI. As seen in Figure 4, dispersing the jamming energy over a larger adversarial network increases the jamming impact via close jammer proximity.

We demonstrate the effect of the jammers' resource constraints on jamming impact by varying the total jamming energy $\sum_{j\in\mathcal{J}} E_j$ for a fixed number of jammers. Figure 5 illustrates the jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ as a function of the total jamming energy for attacks using the convex approximations of Cases 1-4 and the distributed algorithm in Section VI. As seen in Figure 5, the resulting jamming impact increases with the total available energy, with diminishing returns as the energy increases.

To show the effect of network flow topology, we evaluate the effect of varying the number of network traffic flows $|\mathcal{F}|$, keeping the total traffic rate fixed. Figure 6 illustrates the jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ as a function of the number of traffic flows for attacks using the convex approximations of Cases 1-4 and the distributed algorithm in Section VI. As seen in Figure 6, the jamming impact initially increases with the number of flows due to the decreasing minimum distances

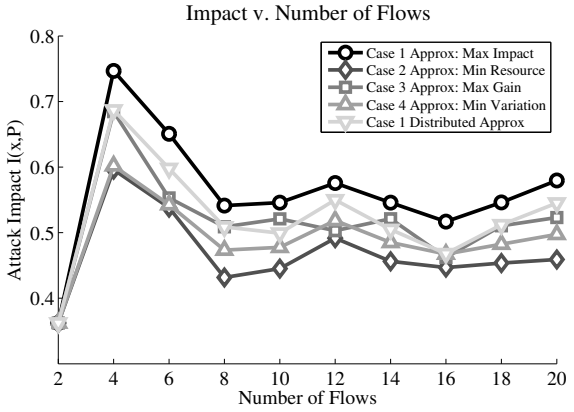

Fig. 6. The jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ resulting from each of the centralized convex and distributed flow-jamming attacks is simulated for various numbers of network traffic flows $|\mathcal{F}|$, keeping the total flow rate $\sum_{f \in \mathcal{F}} r_f$ constant.

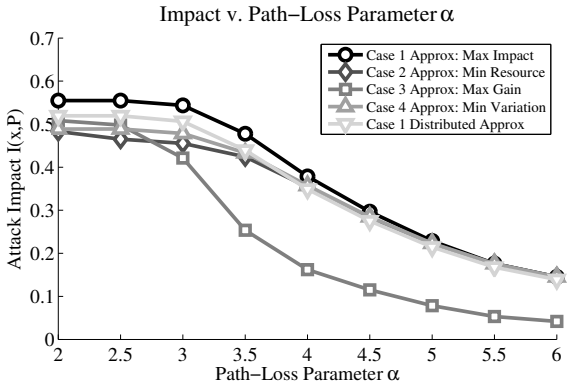

Fig. 7. The jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ resulting from each of the centralized convex and distributed flow-jamming attacks is simulated for various values of the path-loss exponent α, keeping all other network and jamming parameters fixed.

between jammers and flows. However, as the number of flows increases beyond a threshold, the jamming impact decreases due to the traffic being spread evenly among areas not covered by the adversarial network.

To show the effect of the physical medium, we evaluate the effect of varying the path-loss parameter α. Figure 7 illustrates the jamming impact $I(\boldsymbol{x}, \boldsymbol{P})$ as a function of the path-loss exponent for attacks using the convex approximations of Cases 1-4 and the distributed algorithm in Section VI. As seen in Figure 7, the jamming impact decreases with increasing path-loss exponent due to the increase in transmission power required to maintain the interference power level. In this case, it is also necessary for the network nodes to increase their transmission power to maintain connectivity. We additionally note that the performance of the distributed algorithm converges to that of the centralized convex approximation as the path-loss parameter α increases, demonstrating the highly localized effect of jamming in the lossy environment.

VIII. Conclusion

In this work, we showed that cross-layer information and optimization techniques allow a resource-constrained adversary to intelligently allocate jamming resources over an adversarial network. We introduced the class of *flow-jamming attacks* which seek to efficiently reduce network throughput, demonstrating the potential effects of jamming by even the most resource-starved adversarial networks. We presented a set of metrics to quantify the effect of flow-jamming on network flows and jammer resource expenditure and developed a constrained, non-convex optimization problem to model these attacks. Furthermore, in order to demonstrate the feasibility of these attacks in real-time, we proposed convex relaxations and a distributed version of jamming attacks in this framework. We demonstrated the potential impact of cross-layer flow-jamming attacks through a simulation study and showed that the convex formulations closely approximate the non-convex optimization problems. Future work can leverage the understanding of jamming expounded by this work to design network protocols that are robust to cross-layer jamming attacks.

References

[1] P. Tague, D. Slater, G. Noubir, and R. Poovendran, "Linear programming models for jamming attacks on network traffic flows," in *Proc. 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'08)*, Berlin, Germany, Apr. 2008.

[2] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons, Inc., 2001.

[3] R. A. Poisel, *Modern Communication Jamming Principles and Techniques*. Artech House, 2004.

[4] D. J. Torrieri, *Principles of Secure Communication Systems*, 2nd ed. Boston: Artech House, 1992.

[5] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Prentice Hall, 2001.

[6] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *Proc. USENIX Security Symposium*, Washington, DC, Aug. 2003, pp. 15–28.

[7] D. J. Thuente and M. Acharya, "Intelligent jamming in wireless networks with applications to 802.11b and other networks," in *Proc. 25th IEEE Communications Society Military Communications Conference (MILCOM'06)*, Washington, DC, Oct. 2006, pp. 1–7.

[8] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 10, pp. 54–62, Oct. 2002.

[9] G. Lin and G. Noubir, "On link layer denial of service in data wireless LANs," *Wireless Communications and Mobile Computing*, vol. 5, no. 3, pp. 273–284, May 2005.

[10] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[11] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Prentice Hall, 1992.

[12] K. Fazel and S. Kaiser, *Multi-Carrier and Spread Spectrum Systems*. Wiley, 2003.

[13] F. Meshkati, H. V. Poor, S. C. Schwartz, and N. B. Mandayam, "An energy-efficient approach to power control and receiver design in wireless data networks," *IEEE Transactions on Communications*, vol. 53, no. 11, pp. 1885–1894, Nov. 2005.

[14] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Dover, 1998.

[15] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge: Cambridge University Press, 1985.

[16] J.-H. Chang and L. Tassiulas, "Energy conserving routing in wireless ad-hoc networks," in *Proc. 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'00)*, Tel Aviv, Israel, Mar. 2000, pp. 22–31.

[17] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, 2004.

[18] V. Rodriguez, "Maximizing the ratio of a generalized sigmoid to its argument," Polytechnic University, Tech. Rep. WICAT Tech. Rep. 02-010, May 2002.

[19] D. P. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, Aug. 2006.

[20] "The TOMLAB optimization environment," http://tomopt.com/tomlab/.

[21] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming (web page and software)," Feb. 2009, http://stanford.edu/~boyd/cvx.