Jamming-Aware Traffic Allocation for Multiple-Path Routing Using Portfolio Selection

Patrick Tague, Sidharth Nabar, James A. Ritcey, and Radha Poovendran

***Abstract*—Multiple-path source routing protocols allow a data source node to distribute the total traffic among available paths. In this article, we consider the problem of jamming-aware source routing in which the source node performs traffic allocation based on empirical jamming statistics at individual network nodes. We formulate this traffic allocation as a lossy network flow optimization problem using portfolio selection theory from financial statistics. We show that in multi-source networks, this centralized optimization problem can be solved using a distributed algorithm based on decomposition in network utility maximization (NUM). We demonstrate the network's ability to estimate the impact of jamming and incorporate these estimates into the traffic allocation problem. Finally, we simulate the achievable throughput using our proposed traffic allocation method in several scenarios.**

***Index Terms*—Jamming, Multiple path routing, Portfolio selection theory, Optimization, Network utility maximization**

I. Introduction

Jamming point-to-point transmissions in a wireless mesh network [1] or underwater acoustic network [2] can have debilitating effects on data transport through the network. The effects of jamming at the physical layer resonate through the protocol stack, providing an effective denial-of-service (DoS) attack [3] on end-to-end data communication. The simplest methods to defend a network against jamming attacks comprise physical layer solutions such as spread-spectrum or beamforming, forcing the jammers to expend a greater resource to reach the same goal. However, recent work has demonstrated that intelligent jammers can incorporate cross-layer protocol information into jamming attacks, reducing resource expenditure by several orders of magnitude by targeting certain link layer and MAC implementations [4]–[6] as well as link layer error detection and correction protocols [7]. Hence, more sophisticated anti-jamming methods and defensive measures must be incorporated into higher-layer protocols, for example channel surfing [8] or routing around jammed regions of the network [6].

The majority of anti-jamming techniques make use of diversity. For example, anti-jamming protocols may employ multiple frequency bands, different MAC channels, or multiple routing paths. Such diversity techniques help to curb the effects of the jamming attack by requiring the jammer to act on multiple resources simultaneously. In this paper, we consider the anti-jamming diversity based on the use of multiple routing paths. Using multiple-path variants of source routing protocols such as Dynamic Source Routing (DSR) [9] or Ad-Hoc On-Demand Distance Vector (AODV) [10], for example the MP-DSR protocol [11], each source node can request several routing paths to the destination node for concurrent use. To make effective use of this routing diversity, however, each source node must be able to make an intelligent allocation of traffic across the available paths while considering the potential effect of jamming on the resulting data throughput.

In order to characterize the effect of jamming on throughput, each source must collect information on the impact of the jamming attack in various parts of the network. However, the extent of jamming at each network node depends on a number of unknown parameters, including the strategy used by the individual jammers and the relative location of the jammers with respect to each transmitter-receiver pair. Hence, *the impact of jamming is probabilistic from the perspective of the network*[1], and the characterization of the jamming impact is further complicated by the fact that the jammers' strategies may be dynamic and *the jammers themselves may be mobile*[2].

In order to capture the non-deterministic and dynamic effects of the jamming attack, we model the packet error rate at each network node as a random process. At a given time, the randomness in the packet error rate is due to the uncertainty in the jamming parameters, while the time-variability in the packet error rate is due to the jamming dynamics and mobility. Since the effect of jamming at each node is probabilistic, the end-to-end throughput achieved by each source-destination pair will also be non-deterministic and, hence, must be studied using a stochastic framework.

In this article, we thus investigate the ability of network nodes to characterize the jamming impact and the ability of multiple source nodes to compensate for jamming in

This work was supported in part by the following grants: ARO PECASE, W911NF-05-1-0491; ONR, N000-07-1-0600; ONR, N00014-07-1-0600; ARL CTA, DAAD19-01-2-0011; and ARO MURI, W911NF-07-1-0287. This document was prepared through collaborative participation in the Communications and Networks Consortium sponsored by the US Army Research Laboratory under the Collaborative Technology Alliance Program, DAAD19-01-2-0011. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the US Government.

P. Tague, S. Nabar, J. A. Ritcey, and R. Poovendran are with the Network Security Lab (NSL), Electrical Engineering Department, University of Washington, Seattle, Washington. Email: {tague,snabar,jar7,rp3}@u.washington.edu. P. Tague is currently with Carnegie Mellon University, Silicon Valley Campus, Moffett Field, California.

A preliminary version of this material appeared at IEEE PIMRC 2008.

[1]We assume that the network does not rely on a jamming detection, localization, or tracking infrastructure.

[2]We note that factors other than jamming that similarly impact throughput can be included as well. We focus on jamming in this work as it is likely the prominent source of packet loss.

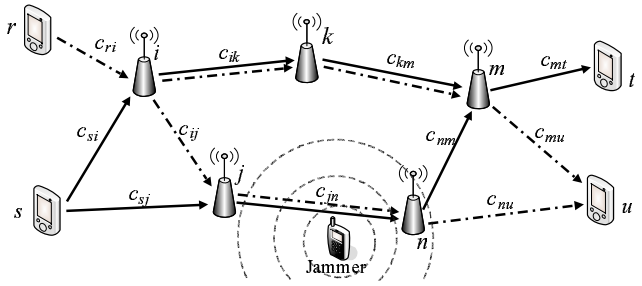

Fig. 1. An example network with sources $\mathcal{S} = \{r, s\}$ is illustrated. Each unicast link $(i, j) \in \mathcal{E}$ is labeled with the corresponding link capacity.

the allocation of traffic across multiple routing paths. Our contributions to this problem are as follow:

- We formulate the problem of allocating traffic across multiple routing paths in the presence of jamming as a lossy network flow optimization problem. We map the optimization problem to that of asset allocation using portfolio selection theory [12], [13].
- We formulate the centralized traffic allocation problem for multiple source nodes as a convex optimization problem.
- We show that the multi-source multiple-path optimal traffic allocation can be computed at the source nodes using a distributed algorithm based on decomposition in network utility maximization (NUM) [14].
- We propose methods which allow individual network nodes to locally characterize the jamming impact and aggregate this information for the source nodes.
- We demonstrate that the use of portfolio selection theory allows the data sources to balance the expected data throughput with the uncertainty in achievable traffic rates.

The remainder of this article is organized as follows. In Section II, we state the network model and assumptions about the jamming attack. To motivate our formulation, in Section III, we present methods that allow nodes to characterize the local jamming impact. These concepts are required to understand the traffic allocation optimization and the mapping of this problem to Portfolio selection. In Section IV, we formulate the optimal multiple path traffic allocation problem for multi-source networks. In Section V, we evaluate the performance of the optimal traffic allocation formulation. We summarize our contributions in Section VI.

II. System Model and Assumptions

The wireless network of interest can be represented by a directed graph $G = (\mathcal{N}, \mathcal{E})$. The vertex set $\mathcal{N}$ represents the network nodes, and an ordered pair (i, j) of nodes is in the edge set $\mathcal{E}$ if and only if node j can receive packets directly from node i. We assume that all communication is unicast over the directed edges in $\mathcal{E}$, i.e. each packet transmitted by node $i \in \mathcal{N}$ is intended for a unique node $j \in \mathcal{N}$ with $(i, j) \in \mathcal{E}$. The maximum achievable data rate, or capacity, of each unicast link $(i, j) \in \mathcal{E}$ in the absence of jamming is denoted by the predetermined constant rate c_{ij} in units of packets per second[3].

Each source node s in a subset $\mathcal{S} \subseteq \mathcal{N}$ generates data for a single destination node $d_s \in \mathcal{N}$. We assume that each source node s constructs multiple routing paths to d_s using a route request process similar to those of the DSR [9] or AODV [10] protocols. We let $\mathcal{P}_s = \{p_{s1}, \ldots, p_{sL_s}\}$ denote the collection of L_s loop-free routing paths for source s, noting that these paths need not be disjoint as in MP-DSR [11]. Representing each path $p_{s\ell}$ by a subset of directed link set $\mathcal{E}$, the sub-network of interest to source s is given by the directed subgraph

$$G_s = \left(\mathcal{N}_s = \bigcup_{\ell=1}^{L_s} \{j : (i, j) \in p_{s\ell}\}, \mathcal{E}_s = \bigcup_{\ell=1}^{L_s} p_{s\ell} \right)$$

of the graph G.

Figure 1 illustrates an example network with sources $\mathcal{S} = \{r, s\}$. The subgraph G_r consists of the two routing paths

$$p_{r1} = \{(r, i), (i, k), (k, m), (m, u)\}$$
$$p_{r2} = \{(r, i), (i, j), (j, n), (n, u)\},$$

and the subgraph G_s consists of the two routing paths

$$p_{s1} = \{(s, i), (i, k), (k, m), (m, t)\}$$
$$p_{s2} = \{(s, j), (j, n), (n, m), (m, t)\}.$$

In this article, we assume that the source nodes in $\mathcal{S}$ have no prior knowledge about the jamming attack being performed. That is, we make no assumption about the jammer's goals, method of attack, or mobility patterns. We assume that the number of jammers and their locations are unknown to the network nodes. Instead of relying on direct knowledge of the jammers, we suppose that the network nodes characterize the jamming impact in terms of the empirical packet delivery rate. Network nodes can then relay the relevant information to the source nodes in order to assist in optimal traffic allocation. Each time a new routing path is requested or an existing routing path is updated, the responding nodes along the path will relay the necessary parameters to the source node as part of the reply message for the routing path. Using the information from the routing reply, each source node s is thus provided with additional information about the jamming impact on the individual nodes.

III. Characterizing the Impact of Jamming

In this section, we propose techniques for the network nodes to estimate and characterize the impact of jamming and for a source node to incorporate these estimates into its traffic allocation. In order for a source node s to incorporate the jamming impact in the traffic allocation problem, the effect of jamming on transmissions over each link $(i, j) \in \mathcal{E}_s$ must be estimated and relayed to s. However, to capture the jammer mobility and the dynamic effects of the jamming attack, the local estimates need to be continually updated. We begin with an example to illustrate the possible effects of jammer mobility on the traffic allocation problem and motivate the use of continually updated local estimates.

[3]We assume that this capacity is an available constant which corresponds to the maximum packet rate for reliable transport over each wireless link. We do not address the analysis or estimation of this link capacity parameter.

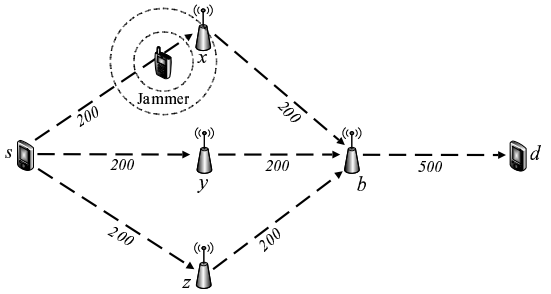

Fig. 2. An example network that illustrates a single-source network with three routing paths. Each unicast link (i, j) is labeled with the corresponding link capacity c_{ij} in units of packets per second. The proximity of the jammer to nodes x and y impedes packet delivery over the corresponding paths, and the jammer mobility affects the allocation of traffic to the three paths as a function of time.

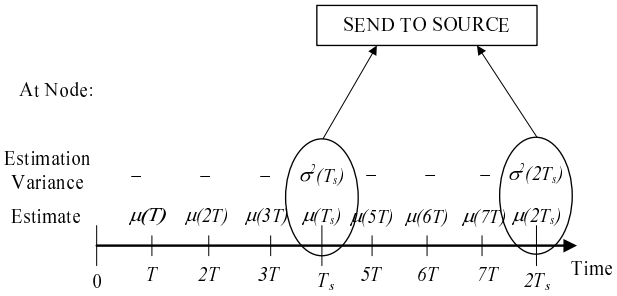

Fig. 3. The estimation update process is illustrated for a single link. The estimate $\mu_{ij}(t)$ is updated every T seconds, and the estimation variance $\sigma^2_{ij}(t)$ is computed only every T_s seconds. Both values are relayed to relevant source nodes every T_s seconds.

A. Illustrating the Effect of Jammer Mobility on Network Throughput

Figure 2 illustrates a single-source network with three routing paths $p_1 = \{(s,x),(x,b),(b,d)\}$, $p_2 = \{(s,y),(y,b),(b,d)\}$ and $p_3 = \{(s,z),(z,b),(b,d)\}$. The label on each edge (i,j) is the link capacity c_{ij} indicating the maximum number of packets per second ($pkts/s$) which can be transported over the wireless link. In this example, we assume that the source is generating data at a rate of 300 $pkts/s$. In the absence of jamming, the source can continuously send 100 $pkts/s$ over each of the three paths, yielding a throughput rate equal to the source generation rate of 300 $pkts/s$. If a jammer near node x is transmitting at high power, the probability of successful packet reception, referred to as the *packet success rate*, over the link (s,x) drops to nearly zero, and the traffic flow to node d reduces to 200 $pkts/s$. If the source node becomes aware of this effect, the allocation of traffic can be changed to 150 $pkts/s$ on each of paths p_2 and p_3, thus recovering from the jamming attack at node x. However, this one-time re-allocation by the source node s does not adapt to the potential mobility of the jammer. If the jammer moves to node y, the packet success rate over (s,x) returns to one and that over (s,y) drops to zero, reducing the throughput to node d to 150 $pkts/s$, which is less than the 200 $pkts/s$ that would be achieved using the original allocation of 100 $pkts/s$ over each of the three paths. Hence, each node must relay an estimate of its packet success rate to the source node s and the source must use this information to reallocate traffic in a timely fashion if the effect of the attack is to be mitigated. The relay of information from the nodes can be done periodically or at the instants when the packet success rates change significantly. These updates must be performed at a rate comparable to the rate of the jammer movement to provide an effective defense against the mobile jamming attack.

Next, suppose the jammer continually changes position between nodes x and y, causing the packet success rates over links (s,x) and (s,y) to oscillate between zero and one. This behavior introduces a high degree of variability into the observed packet success rates, leading to a less certain estimate of the future success rates over the links (s,x) and (s,y). However, since the packet success rate over link (s,z) has historically been more steady, it may be a more reliable option. Hence, the source s can choose to fill p_3 to its capacity and partition the remaining 100 $pkts/s$ equally over p_1 and p_2. This solution takes into account the historic variability in the packet success rates due to jamming mobility. In the following section, we build on this example, providing a set of parameters to be estimated by network nodes and methods for the sources to aggregate this information and characterize the available paths on the basis of expected throughput.

B. Estimating Local Packet Success Rates

We let $x_{ij}(t)$ denote the packet success rate over link $(i,j) \in \mathcal{E}$ at time t, noting that $x_{ij}(t)$ can be computed analytically as a function of the transmitted signal power of node i, the signal power of the jammers, their relative distances from node j, and the path loss behavior of the wireless medium. In reality, however, the locations of mobile jammers are often unknown, and, hence, the use of such an analytical model is not applicable. Due to the uncertainty in the jamming impact, we model the packet success rate $x_{ij}(t)$ as a random process and allow the network nodes to collect empirical data in order to characterize the process. We suppose that each node j maintains an estimate $\mu_{ij}(t)$ of the packet success rate $x_{ij}(t)$ as well as a variance parameter $\sigma^2_{ij}(t)$ to characterize the estimate uncertainty and process variability[4].

We propose the use of a recursive update mechanism allowing each node j to periodically update the estimate $\mu_{ij}(t)$ as a function of time. As illustrated in Figure 3, we suppose that each node j updates the estimate $\mu_{ij}(t)$ after each *update period* of T seconds and relays the estimate to each relevant source node s after each *update relay period* of $T_s \gg T$ seconds. The shorter update period of T seconds allows each node j to characterize the variation in $x_{ij}(t)$ over the update relay period of T_s seconds, a key factor in $\sigma^2_{ij}(t)$.

We propose the use of the observed packet delivery ratio (PDR) to compute the estimate $\mu_{ij}(t)$. While the PDR incorporates additional factors such as congestion, it has been shown by extensive experimentation [8] that such factors

[4]At a time instant t, the estimate $\mu_{ij}(t)$ and estimation variance $\sigma^2_{ij}(t)$ define a random variable describing the current view of the packet success rate. This random variable can be appropriately modeled as a beta random variable [15], though the results of this article do not require such an assumption.

do not affect the PDR in a similar manner. Furthermore, we propose to average the empirical PDR values over time to smooth out the relatively short-term variations due to noise or fading. During the update period represented by the time interval $[t-T,t]$, each node j can record the number $r_{ij}([t-T,t])$ of packets received over link (i,j) and the number $v_{ij}([t-T,t]) \leq r_{ij}([t-T,t])$ of valid packets which pass an error detection check[5]. The PDR over link (i,j) for the update period $[t-T,t]$, denoted $PDR_{ij}([t-T,t])$, is thus equal to the ratio

$$PDR_{ij}([t-T,t]) = \frac{v_{ij}([t-T,t])}{r_{ij}([t-T,t])}. \quad (1)$$

This PDR can be used to update the estimate $\mu_{ij}(t)$ at the end of the update period. In order to prevent significant variation in the estimate $\mu_{ij}(t)$ and to include memory of the jamming attack history, we suggest using an exponential weighted moving average (EWMA) [16] to update the estimate $\mu_{ij}(t)$ as a function of the previous estimate $\mu_{ij}(t-T)$ as

$$\mu_{ij}(t) = \alpha\mu_{ij}(t-T) + (1-\alpha)PDR_{ij}([t-T,t]), \quad (2)$$

where $\alpha \in [0,1]$ is a constant weight indicating the relative preference between current and historic samples.

We use a similar EWMA process to update the variance $\sigma^2_{ij}(t)$ at the end of each update relay period of T_s seconds. Since this variance is intended to capture the variation in the packet success rate over the last T_s seconds, we consider the sample variance $V_{ij}([t-T_s,t])$ of the set of packet delivery ratios computed using (1) during the interval $[t-T_s,t]$ as

$$V_{ij}([t-T_s,t]) = Var\left\{PDR_{ij}([t-kT,t-kT+T]) : k=0,\ldots,\lceil T_s/T\rceil - 1\right\}. \quad (3)$$

The estimation variance $\sigma^2_{ij}(t)$ is thus defined as a function of the previous variance $\sigma^2_{ij}(t-T_s)$ as

$$\sigma^2_{ij}(t) = \beta\sigma^2_{ij}(t-T_s) + (1-\beta)V_{ij}([t-T_s,t]), \quad (4)$$

where $\beta \in [0,1]$ is a constant weight similar to α in (2).

The EWMA method is widely used in sequential estimation processes, including estimation of the round-trip time (RTT) in TCP [17]. We note that the parameters α in (2) and β in (4) allow for design of the degree of historical content included in the parameter estimate updates, and these parameters can themselves be functions $\alpha(t)$ and $\beta(t)$ of time. For example, decreasing the parameter α allows the mean $\mu_{ij}(t)$ to change more rapidly with the PDR due to jammer mobility, and decreasing the parameter β allows the variance $\sigma^2_{ij}(t)$ to give more preference to variation in the most recent update relay period over historical variations. We further note that the update period T and update relay period T_s between subsequent updates of the parameter estimates have significant influence on the quality of the estimate. In particular, if the update period T_s is too large, the relayed estimates $\mu_{ij}(t)$ and $\sigma_{ij}(t)$ will be outdated before the subsequent update at time $t+T_s$. Furthermore, if the update period T at each node is too large, the dynamics of the jamming attack may be averaged out over the large number of samples $r_{ij}([t-T,t])$. The update periods T and T_s must thus be short enough to capture the dynamics of the jamming attack. However, decreasing the update period T_s between successive updates to the source node necessarily increases the communication overhead of the network. Hence, there exists a trade-off between performance and overhead in the choice of the update period T_s. We note that the design of the update relay period T_s depends on assumed path-loss and jammer mobility models. The application-specific tuning of the update relay period T_s is not further herein.

Using the above formulation, each time a new routing path is requested or an existing routing path is updated, the nodes along the path will include the estimates $\mu_{ij}(t)$ and $\sigma^2_{ij}(t)$ as part of the reply message. In what follows, we show how the source node s uses these estimates to compute the end-to-end packet success rates over each path.

C. Estimating End-to-End Packet Success Rates

Given the packet success rate estimates $\mu_{ij}(t)$ and $\sigma^2_{ij}(t)$ for the links (i,j) in a routing path $p_{s\ell}$, the source s needs to estimate the effective end-to-end packet success rate to determine the optimal traffic allocation. Assuming the total time required to transport packets from each source s to the corresponding destination d_s is negligible compared to the update relay period T_s, we drop the time index and address the end-to-end packet success rates in terms of the estimates μ_{ij} and σ^2_{ij}. The end-to-end packet success rate $y_{s\ell}$ for path $p_{s\ell}$ can be expressed as the product

$$y_{s\ell} = \prod_{(i,j)\in p_{s\ell}} x_{ij}, \quad (5)$$

which is itself a random variable[6] due to the randomness in each x_{ij}. We let $\gamma_{s\ell}$ denote the expected value of $y_{s\ell}$ and $\omega_{s\ell m}$ denote the covariance of $y_{s\ell}$ and y_{sm} for paths $p_{s\ell}, p_{sm} \in \mathcal{P}_s$. Due to the computational burden associated with in-network inference of correlation between estimated random variables, we let the source node s assume the packet success rates x_{ij} as mutually independent, even though they are likely correlated. We maintain this independence assumption throughout this work, yielding a feasible approximation to the complex reality of correlated random variables, and the case of in-network inference of the relevant correlation is left as future work. Under this independence assumption, the mean $\gamma_{s\ell}$ of $y_{s\ell}$ given in (5) is equal to the product of estimates μ_{ij} as

$$\gamma_{s\ell} = \prod_{(i,j)\in p_{s\ell}} \mu_{ij}, \quad (6)$$

and the covariance $\omega_{s\ell m} = E[y_{s\ell}y_{sm}] - E[y_{s\ell}]E[y_{sm}]$ is similarly given by

$$\omega_{s\ell m} = \prod_{(i,j)\in p_{s\ell}\oplus p_{sm}} \mu_{ij} \prod_{(i,j)\in p_{s\ell}\cap p_{sm}} \left(\sigma^2_{ij} + \mu^2_{ij}\right) - \gamma_{s\ell}\gamma_{sm}. \quad (7)$$

[5]In the case of jamming attacks which prevent the receiving node j from detecting transmissions by node i, additional header information can be periodically exchanged between nodes i and j to achieve the convey the total number of transmissions, yielding the same overall effect.

[6]If the x_{ij} are modeled as beta random variables, the product $y_{s\ell}$ is well-approximated by a beta random variable [18].

In (7), $\oplus$ denotes the exclusive-OR set operator such that an element is in $A \oplus B$ if it is in either A or B but not both. The covariance formula in (7) reflects the fact that the end-to-end packet success rates $y_{s\ell}$ and y_{sm} of paths $p_{s\ell}$ and p_{sm} with shared links are correlated even when the rates x_{ij} are independent. We note that the variance $\omega_{s\ell}^2$ of the end-to-end rate $y_{s\ell}$ can be computed using (7) with $\ell = m$.

Let $\boldsymbol{\gamma}_s$ denote the $L_s \times 1$ vector of estimated end-to-end packet success rates $\gamma_{s\ell}$ computed using (6), and let $\boldsymbol{\Omega}_s$ denote the $L_s \times L_s$ covariance matrix with (ℓ, m) entry $\omega_{s\ell m}$ computed using (7). The estimate pair $(\boldsymbol{\gamma}_s, \boldsymbol{\Omega}_s)$ provides the sufficient statistical characterization of the end-to-end packet success rates for source s to allocate traffic to the paths in $\mathcal{P}_s$. Furthermore, the off-diagonal elements in $\boldsymbol{\Omega}_s$ denote the extent of mutual overlap between the paths in $\mathcal{P}_s$.

IV. Optimal Jamming-Aware Traffic Allocation

In this section, we present an optimization framework for jamming-aware traffic allocation to multiple routing paths in $\mathcal{P}_s$ for each source node $s \in \mathcal{S}$. We develop a set of constraints imposed on traffic allocation solutions and then formulate a utility function for optimal traffic allocation by mapping the problem to that of portfolio selection in finance. Letting $\phi_{s\ell}$ denote the traffic rate allocated to path $p_{s\ell}$ by the source node s, the problem of interest is thus for each source s to determine the optimal $L_s \times 1$ rate allocation vector $\boldsymbol{\phi}_s$ subject to network flow capacity constraints using the available statistics $\boldsymbol{\gamma}_s$ and $\boldsymbol{\Omega}_s$ of the end-to-end packet success rates under jamming.

A. Traffic Allocation Constraints

In order to define a set of constraints for the multiple-path traffic allocation problem, we must consider the source data rate constraints, the link capacity constraints, and the reduction of traffic flow due to jamming at intermediate nodes. The traffic rate allocation vector $\boldsymbol{\phi}_s$ is trivially constrained to the non-negative orthant, i.e. $\boldsymbol{\phi}_s \geq \mathbf{0}$, as traffic rates are non-negative. Assuming data generation at source s is limited to a maximum data rate R_s, the rate allocation vector is also constrained as $\mathbf{1}^T \boldsymbol{\phi}_s \leq R_s$. These constraints define the convex space $\boldsymbol{\Phi}_s$ of feasible allocation vectors $\boldsymbol{\phi}_s$ characterizing rate allocation solutions for source s.

Due to jamming at nodes along the path, the traffic rate is potentially reduced at each receiving node as packets are lost. Hence, while the initial rate of $\phi_{s\ell}$ is allocated to the path, the residual traffic rate forwarded by node i along the path $p_{s\ell}$ may be less than $\phi_{s\ell}$. Letting $p_{s\ell}^{(i)}$ denote the sub-path of $p_{s\ell}$ from source s to the intermediate node i, the residual traffic rate forwarded by node i is given by $y_{s\ell}^{(i)} \phi_{s\ell}$, where $y_{s\ell}^{(i)}$ is computed using (5) with $p_{s\ell}$ replaced by the sub-path $p_{s\ell}^{(i)}$.

The capacity constraint on the total traffic traversing a link (i, j) thus imposes the stochastic constraint

$$\sum_{s \in \mathcal{S}} \sum_{\ell : (i,j) \in p_{s\ell}} \phi_{s\ell} y_{s\ell}^{(i)} \leq c_{ij} \tag{8}$$

on the feasible allocation vectors $\boldsymbol{\phi}_s$. To compensate for the randomness in the capacity constraint in (8), we replace the residual packet success rate $y_{s\ell}^{(i)}$ with a function of its expected value and variance. The mean $\gamma_{s\ell}^{(i)}$ and variance $(\omega_{s\ell}^{(i)})^2$ of $y_{s\ell}^{(i)}$ can be computed using (6) and (7), respectively, with $p_{s\ell}$ replaced by the sub-path $p_{s\ell}^{(i)}$. We thus replace $y_{s\ell}^{(i)}$ in (8) with the statistic $\gamma_{s\ell}^{(i)} + \delta \omega_{s\ell}^{(i)}$, where $\delta \geq 0$ is a constant which can be tuned based on tolerance to delay resulting from capacity violations[7]. We let $\boldsymbol{W}_s$ denote the $|\mathcal{E}| \times L_s$ *weighted link-path incidence matrix* for source s with rows indexed by links (i, j) and columns indexed by paths $p_{s\ell}$. The element $w((i,j), p_{s\ell})$ in row (i, j) and column $p_{s\ell}$ of $\boldsymbol{W}_s$ is thus given by

$$w\left((i,j), p_{s\ell}\right) = \begin{cases} \min\left\{1, \gamma_{s\ell}^{(i)} + \delta \omega_{s\ell}^{(i)}\right\}, & \text{if } (i,j) \in p_{s\ell} \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

Letting $\boldsymbol{c}$ denote the $|\mathcal{E}| \times 1$ vector of link capacities c_{ij} for $(i, j) \in \mathcal{E}$, the link capacity constraint in (8) including expected packet loss due to jamming can be expressed by the vector inequality

$$\sum_{s \in \mathcal{S}} \boldsymbol{W}_s \boldsymbol{\phi}_s \leq \boldsymbol{c}, \tag{10}$$

which is a linear constraint in the variable $\boldsymbol{\phi}_s$. We note that this statistical constraint formulation generalizes the standard network flow capacity constraint corresponding to the case of $x_{ij} = 1$ for all $(i, j) \in \mathcal{E}$ in which the incidence matrix $\boldsymbol{W}_s$ is deterministic and binary.

B. Optimal Traffic Allocation Using Portfolio Selection Theory

In order to determine the optimal allocation of traffic to the paths in $\mathcal{P}_s$, each source s chooses a utility function $U_s(\boldsymbol{\phi}_s)$ that evaluates the total data rate, or throughput, successfully delivered to the destination node d_s. In defining our utility function $U_s(\boldsymbol{\phi}_s)$, we present an analogy between traffic allocation to routing paths and allocation of funds to correlated assets in finance.

In Markowitz's portfolio selection theory [12], [13], an investor is interested in allocating funds to a set of financial assets that have uncertain future performance. The expected performance of each investment at the time of the initial allocation is expressed in terms of return and risk. The return on the asset corresponds to the value of the asset and measures the growth of the investment. The risk of the asset corresponds to the variance in the value of the asset and measures the degree of variation or uncertainty in the investment's growth.

We describe the desired analogy by mapping this allocation of funds to financial assets to the allocation of traffic to routing paths. We relate the expected investment return on the financial portfolio to the estimated end-to-end success rates $\boldsymbol{\gamma}_s$ and the investment risk of the portfolio to the estimated success rate covariance matrix $\boldsymbol{\Omega}_s$. We note that the correlation between related assets in the financial portfolio corresponds to the correlation between non-disjoint routing paths. The analogy between financial portfolio selection and the allocation of traffic to routing paths is summarized below.

[7]The case of $\delta = 0$ corresponds to the average-case constraint and will lead to increased queueing delay whenever $y_{s\ell}^{(i)} > \gamma_{s\ell}^{(i)}$. Increasing the value of δ improves the robustness to variations around the mean but decreases the amount of traffic which can be allocated to the corresponding path.

Portfolio Selection	Traffic Allocation
Funds to be invested	Source data rate R_s
Financial assets	Routing paths $\mathcal{P}_s$
Expected Asset return	Expected Packet success rate $\gamma_{s\ell}$
Investment portfolio	Traffic allocation $\boldsymbol{\phi}_s$
Portfolio return	Mean throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$
Portfolio risk	Estimation variance $\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s$

As in Markowitz's theory, we define a constant *risk-aversion factor* $k_s \geq 0$ for source $s \in \mathcal{S}$ to indicate the preference for source s to allocate resources to less risky paths with lower throughput variance. This risk-aversion constant weighs the trade-off between expected throughput and estimation variance. We note that each source s can choose a different risk-aversion factor, and a source may vary the risk-aversion factor k_s with time or for different types of data. For a given traffic rate allocation vector $\boldsymbol{\phi}_s$, the expected total throughput for source s is equal to the vector inner product $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$. The corresponding variance in the throughput for source s due to the uncertainty in the estimate $\boldsymbol{\gamma}_s$ is equal to the quadratic term $\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s$. Based on the above analogy making use of portfolio selection theory, we define the utility function $U_s(\boldsymbol{\phi}_s)$ at source s as the weighted sum

$$U_s(\boldsymbol{\phi}_s) = \boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s - k_s \boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s. \tag{11}$$

Setting the risk-aversion factor k_s to zero indicates that the source s is willing to put up with any amount of uncertainty in the estimate $\boldsymbol{\gamma}_s$ of the end-to-end success rates to maximize the expected throughput. The role of the risk-aversion factor is thus to impose a penalty on the objective function proportional to the uncertainty in the estimation process, potentially narrowing the gap between expected throughput and achieved throughput. The cases of $k_s = 0$ and $k_s > 0$ are compared in detail in Section V.

Combining the utility function in (11) with the set of constraints defined in Section IV-A yields the following jamming-aware traffic allocation optimization problem which aims to find the globally optimal traffic allocation over the set $\mathcal{S}$ of sources.

Optimal Jamming-Aware Traffic Allocation

$$\begin{aligned} \boldsymbol{\phi}^* = \quad & \underset{\{\boldsymbol{\phi}_s\}}{\arg\max} \sum_{s \in \mathcal{S}} \boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s - k_s \boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s \\ \text{s.t.} \quad & \sum_{s \in \mathcal{S}} \boldsymbol{W}_s \boldsymbol{\phi}_s \leq \boldsymbol{c} \\ & \mathbf{1}^T \boldsymbol{\phi}_s \leq R_s \text{ for all } s \in \mathcal{S}, \\ & \mathbf{0} \leq \boldsymbol{\phi}_s \text{ for all } s \in \mathcal{S}. \end{aligned} \tag{12}$$

Since the use of centralized protocols for source routing may be undesirable due to excessive communication overhead in large-scale wireless networks, we seek a distributed formulation for the optimal traffic allocation problem in (12).

C. Optimal Distributed Traffic Allocation using NUM

In the distributed formulation of the algorithm, each source s determines its own traffic allocation $\boldsymbol{\phi}_s$, ideally with minimal message passing between sources. By inspection, we see that the optimal jamming-aware flow allocation problem in (12) is similar to the network utility maximization (NUM) formulation of the basic maximum network flow problem [14]. We thus develop a distributed traffic allocation algorithm using Lagrangian dual decomposition techniques [14] for NUM.

The dual decomposition technique is derived by decoupling the capacity constraint in (10) and introducing the *link prices* λ_{ij} corresponding to each link (i, j). Letting $\boldsymbol{\lambda}$ denote the $|\mathcal{E}| \times 1$ vector of link prices λ_{ij}, the Lagrangian $L(\boldsymbol{\phi}, \boldsymbol{\lambda})$ of the optimization problem in (12) is given by

$$L(\boldsymbol{\phi}, \boldsymbol{\lambda}) = \sum_{s \in \mathcal{S}} \boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s - k_s \boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s + \boldsymbol{\lambda}^T \left(\boldsymbol{c} - \sum_{s \in \mathcal{S}} \boldsymbol{W}_s \boldsymbol{\phi}_s \right). \tag{13}$$

The distributed optimization problem is solved iteratively using the Lagrangian dual method as follows. For a given set of link prices $\boldsymbol{\lambda}_n$ at iteration n, each source s solves the local optimization problem

$$\boldsymbol{\phi}_{s,n}^* = \underset{\boldsymbol{\phi}_s \in \boldsymbol{\Phi}_s}{\arg\max} \left(\boldsymbol{\gamma}_s^T - \boldsymbol{\lambda}_n^T \boldsymbol{W}_s \right) \boldsymbol{\phi}_s - k_s \boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s. \tag{14}$$

The link prices $\boldsymbol{\lambda}_{n+1}$ are then updated using a gradient descent iteration as

$$\boldsymbol{\lambda}_{n+1} = \left(\boldsymbol{\lambda}_n - a \left(\boldsymbol{c} - \sum_{s \in \mathcal{S}} \boldsymbol{W}_s \boldsymbol{\phi}_{s,n}^* \right) \right)^+, \tag{15}$$

where $a > 0$ is a constant step size and $(\boldsymbol{v})^+ = \max(\mathbf{0}, \boldsymbol{v})$ is the element-wise projection into the non-negative orthant. In order to perform the local update in (15), sources must exchange information about the result of the local optimization step. Since updating the link prices $\boldsymbol{\lambda}$ depends only on the expected link usage, sources must only exchange the $|\mathcal{E}| \times 1$ link usage vectors $\boldsymbol{u}_{s,n} = \boldsymbol{W}_s \boldsymbol{\phi}_{s,n}^*$ to ensure that the link prices are consistently updated across all sources. The iterative optimization step can be repeated until the allocation vectors $\boldsymbol{\phi}_s$ converge[8] for all sources $s \in \mathcal{S}$, i.e. when $\|\boldsymbol{\phi}_{s,n}^* - \boldsymbol{\phi}_{s,n-1}^*\| \leq \epsilon$ for all s with a given $\epsilon > 0$. The above approach yields the following distributed algorithm for optimal jamming-aware flow allocation.

Distributed Jamming-Aware Traffic Allocation

Initialize $n = 1$ with initial link prices $\boldsymbol{\lambda}_1$.

1. Each source s independently computes
   $\boldsymbol{\phi}_{s,n}^* = \underset{\boldsymbol{\phi}_s \in \boldsymbol{\Phi}_s}{\arg\max} \left(\boldsymbol{\gamma}_s^T - \boldsymbol{\lambda}_n^T \boldsymbol{W}_s \right) \boldsymbol{\phi}_s - k_s \boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s.$
2. Sources exchange the link usage vectors
   $\boldsymbol{u}_{s,n} = \boldsymbol{W}_s \boldsymbol{\phi}_{s,n}^*$.
3. Each source locally updates link prices as
   $\boldsymbol{\lambda}_{n+1} = \left(\boldsymbol{\lambda}_n - a \left(\boldsymbol{c} - \sum_{s \in \mathcal{S}} \boldsymbol{u}_{s,n} \right) \right)^+.$
4. If $\|\boldsymbol{\phi}_{s,n}^* - \boldsymbol{\phi}_{s,n-1}^*\| > \epsilon$ for any s, increment n and go to step 1.

Given the centralized optimization problem in (12) and the above distributed formulation for jamming-aware traffic allocation, a set of sources with estimated parameters $\boldsymbol{\gamma}_s$ and $\boldsymbol{\Omega}_s$ can proactively compensate for the presence of jamming on network traffic flow.

[8] In order to prevent premature termination at a local minimum, sources could additionally exchange a flag f_s indicating whether or not local convergence has been attained such that all sources continue to iterate until all convergence flags have been set.

D. Computational Complexity

We note that both the centralized optimization problem in (12) and the local optimization step in the distributed algorithm are quadratic programming optimization problems with linear constraints [13]. The computational time required for solving these problems using numerical methods for quadratic programming is a polynomial function of the number of optimization variables and the number of constraints.

In the centralized problem, there are $\sum_{s\in\mathcal{S}}|\mathcal{P}_s|$ optimization variables corresponding to the number of paths available to each of the sources. The number of constraints in the centralized problem is equal to the total number of links $|\bigcup_{s\in\mathcal{S}}\mathcal{E}_s|$, corresponding to the number of link capacity constraints. In the distributed algorithm, each source iteratively solves a local optimization problem, leading to $|\mathcal{S}|$ decoupled optimization problems. Each of these problems has $|\mathcal{P}_s|$ optimization variables and $|\mathcal{E}_s|$ constraints. Hence, as the number of sources in the network increases, the distributed algorithm may be advantageous in terms of total computation time. In what follows, we provide a detailed performance evaluation of the methods proposed in this article.

V. Performance Evaluation

In this section, we simulate various aspects of the proposed techniques for estimation of jamming impact and jamming-aware traffic allocation. We first describe the simulation setup, including descriptions of the assumed models for routing path construction, jammer mobility, packet success rates, and estimate updates. We then simulate the process of computing the estimation statistics $\mu_{ij}(t)$ and $\sigma^2_{ij}(t)$ for a single link (i,j). Next, we illustrate the effects of the estimation process on the throughput optimization, both in terms of optimization objective functions and the resulting simulated throughput. Finally, we simulate a small-scale network similar to that in Figure 2 while varying network and protocol parameters in order to observe performance trends.

A. Simulation Setup

The simulation results presented herein are obtained using the following simulation setup. A network of nodes is deployed randomly over an area, and links are formed between pairs of nodes within a fixed communication range. The set $\mathcal{S}$ of source nodes is chosen randomly, and the destination node d_s corresponding to each source $s\in\mathcal{S}$ is randomly chosen from within the connected component containing s. Each routing path in the set $\mathcal{P}_s$ is chosen using a randomized geometric routing algorithm which chooses the next hop toward the destination d_s from the set of neighboring nodes that are closer to d_s in terms of either distance or hop-count. Nodes transmit using fixed power P_t.

We simulate the case of continuous jamming at a fixed power P_j using omnidirectional antennas. The mobility of each jammer j consists of repeatedly choosing a random direction $\theta_j\in[0,2\pi)$ and a random speed $v_j\in[0,V_{\max}]$ and moving for a random amount of time $\tau_j>0$ at the chosen direction and speed. At each instant in time, the packet error rate is a function of the transmission powers P_t and P_j, the distance d_{tr} from the transmitter to the receiver, and the distances d_{jr} from each jammer to the receiver. The packet error rate is set equal to $e^{-\xi s}$ where s is the signal to interference and noise ratio (SINR) $s=S/(I+N)$. The SINR is computed as a function of the received signal power $S=\rho P_t d_{tr}^{-\nu}$ from the transmitter, the received interference power $I=\rho\sum_j P_j d_{jr}^{-\nu}$ from the jammers, and the noise N at the receiver. The constant $\xi>0$ determines the relationship between the SINR and the packet error rate, and the constants $\rho>0$ and $\nu\geq 2$ characterize the path-loss of the wireless medium. In our simulation study, we choose parameters based on IEEE 802.15.4 and the CC2420 transceiver, and these parameters are summarized in Table I.

TABLE I
Summary of simulation parameters.

Parameter	Value
Network area	$500\ m\times 500\ m$
Radio range	$100\ m$
Number of sources	$\lvert\mathcal{S}\rvert=1$
Number of nodes	$\lvert\mathcal{N}\rvert=200$
Maximum source data rate	$R_s=200\ pkts/s$
Maximum number of paths	$\lvert\mathcal{P}_s\rvert\leq 5$
Transmission power	$P_t=1\ mW$ $(0\ dBm)$
Link capacity	$c_{ij}=500\ pkts/s$
Jamming transmission power	$P_j=1\ mW$ $(0\ dBm)$
Maximum jammer mobility speed	$V_{\max}=5\ m/s$
Packet error rate parameter	$\xi=1.16$
Path-loss constant	$\rho=2.5\times 10^{-4}$
Path-loss exponent	$\nu=2.7$
Receiver noise	$N=10^{-10}\ mW$ $(-100\ dBm)$
EWMA coefficients	$\alpha=0.7$, $\beta=0.3$
Update period	$T=0.05\ s$
Update relay period	$T_s=2\ s$

We are interested in comparing the performance of several methods of traffic allocation using the given network and jamming models. We define the following cases of interest.

Case I - Ignoring jamming: Each source s chooses the allocation vector $\boldsymbol{\phi}_s$ using the standard maximum-flow formulation corresponding to $\mu_{ij}=1$ and $\sigma^2_{ij}=0$ for all links (i,j). This case is included in order to observe the improvement that can be obtained by incorporating the jamming statistics.

Case II - Maximum throughput: The allocation vectors $\boldsymbol{\phi}_s$ are chosen using the jamming-aware optimization problem in (12) with risk-aversion constant $k_s=0$. This case incorporates the estimates μ_{ij}, updated every T_s seconds, in the allocation.

Case III - Minimum risk-return: Similar to Case II with $k_s>0$. This case incorporates the estimates μ_{ij} and uncertainty parameters σ^2_{ij} to balance the mean throughput with the estimation variance.

Case IV - Oracle model: Each source s continuously optimizes the allocation vector $\boldsymbol{\phi}_s$ using the true values of the packet success rates x_{ij}. This impractical case is included in order to illustrate the effect of the estimation process.

Our simulations are performed using a packet simulator which generates and allocates packets to paths in a fixed network according to the current value of the allocation vector $\boldsymbol{\phi}_s$. Each trial of the simulation compares several of the above

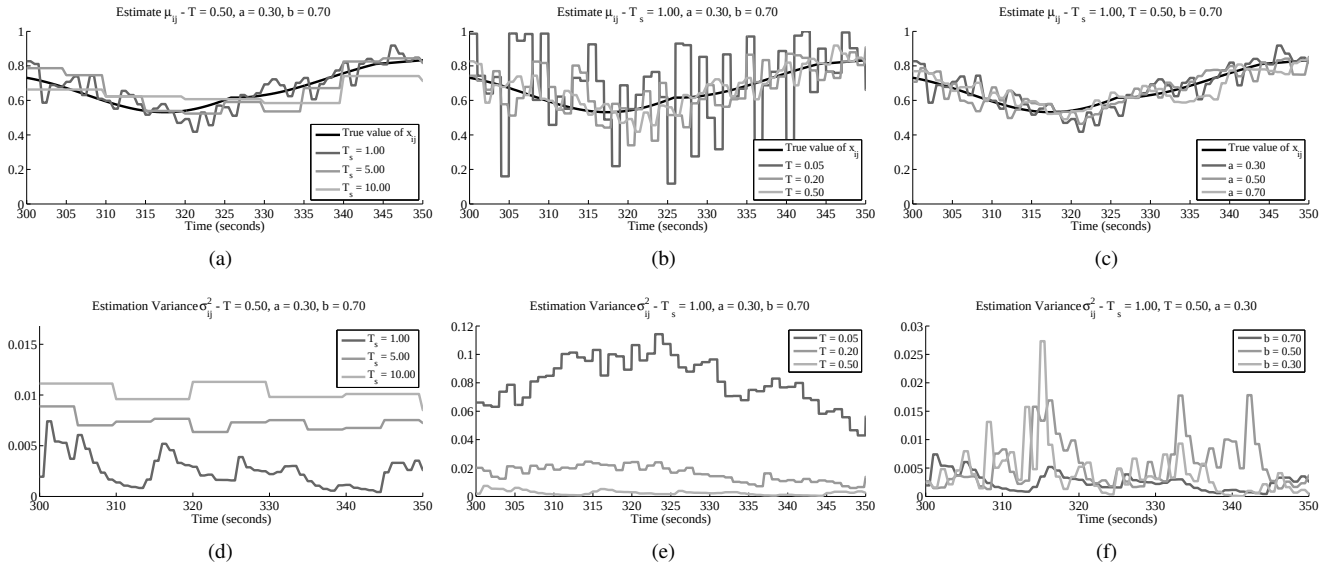

Fig. 4. The estimate $\mu_{ij}(t)$ and estimation variance $\sigma^2_{ij}(t)$ are simulated and for various values of the update relay period T_s in (a) and (d), the update period T in (b) and (e), and the EWMA coefficients in (c) and (f).

cases using the same jammer mobility patterns.

B. Simulation of Estimation Process

We first simulate the process of computing the estimate $\mu_{ij}(t)$ and the variance $\sigma^2_{ij}(t)$ over a single link (i, j). Figure 4 shows the true packet success rate $x_{ij}(t)$ with the estimate $\mu_{ij}(t)$ and the estimation variance $\sigma^2_{ij}(t)$ for various parameter values. By inspection of Figure 4, we see that a shorter update relay period T_s and a longer update period T yield a more consistent estimate $\mu_{ij}(t)$ with less variation around the true value of $x_{ij}(t)$. In addition, a smaller value of α allows the estimate $\mu_{ij}(t)$ to reflect rapid changes in $x_{ij}(t)$, while a larger value of α smooths the estimate $\mu_{ij}(t)$ over the sampled PDRs. We similarly see that a shorter update relay period T_s and a longer update period T yield a lower estimation variance $\sigma^2_{ij}(t)$. In addition, a smaller value of the EWMA coefficient β allows the estimation variance $\sigma^2_{ij}(t)$ to primarily reflect recent variations in the sampled PDRs, while a larger value of β incorporates PDR history to a greater degree.

C. Network Simulation

We next simulate the jamming-aware traffic allocation using the estimated parameters $\mu_{ij}(t)$ and $\sigma^2_{ij}(t)$ as described in Section V-A. To observe the effects of the jamming-aware formulation and the estimation process, we first compare the optimal expected throughput and the actual achieved throughput for Case I, Case II, and Case IV in Figure 5. Figure 5(a) illustrates the expected throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$ and throughput variance $\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s$ over time, and Figure 5(b) illustrates the resulting throughput $\boldsymbol{y}_s^T \boldsymbol{\phi}_s$ over time. By inspection, we see that all of Cases II, III, and IV consistently outperform Case I, showing the benefit of incorporating any type of jamming statistics into the allocation problem. The effect of the estimation error in Case II is seen in the difference between the expected throughput in Figure 5(a) and the achieved throughput in Figure 5(b).

To observe the effect of the risk-aversion constant k_s, we next compare the optimal expected throughput and the actual achieved throughput for Case II with $k_s = 0$ to that of Case III with $k_s > 0$ in Figure 5. Figure 5(c) illustrates the expected throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$ and throughput variance $\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s$ over time, and Figure 5(d) illustrates the resulting throughput $\boldsymbol{y}_s^T \boldsymbol{\phi}_s$ over time. By inspection, we see that Case III exhibits a significant reduction in the throughput variance compared to that of Case II, resulting in achievable throughput much closer to the expected throughput. This reduction in variance in Case III sometimes comes in trade for a reduction in both expected and achieved throughput compared to that of Case II. However, due to the higher variance in Case II, Case III can sometimes achieve higher throughput than Case II, for example over the interval 375-390 seconds in Figure 5(d). The most important feature of Case III is that the achieved throughput in Figure 5(d) closely matches the expected throughput in Figure 5(c).

The choice of k_s in the multi-path traffic allocation is similar to the choice of the risk-aversion parameter in financial portfolio selection [12]. As shown, this parameter introduces a trade-off between the expected throughput and the associated uncertainty. Hence the design of this parameter is a problem of interest in many scenarios, including when timely packet delivery is required (e.g. delivery of control messages) or when packet losses can be tolerated (e.g. streaming video).

D. Simulation of Parameter Dependence

We next evaluate the effect of varying network and protocol parameters in order to observe the performance trends using the jamming-aware traffic allocation formulation. In particular, we are interested in the effect of the update relay period T_s and the maximum number of routing paths $|\mathcal{P}_s|$ on the performance of the flow allocation algorithm. In order to compare trials with different update times or numbers of paths, we average the simulated results over each simulation run, yielding a single

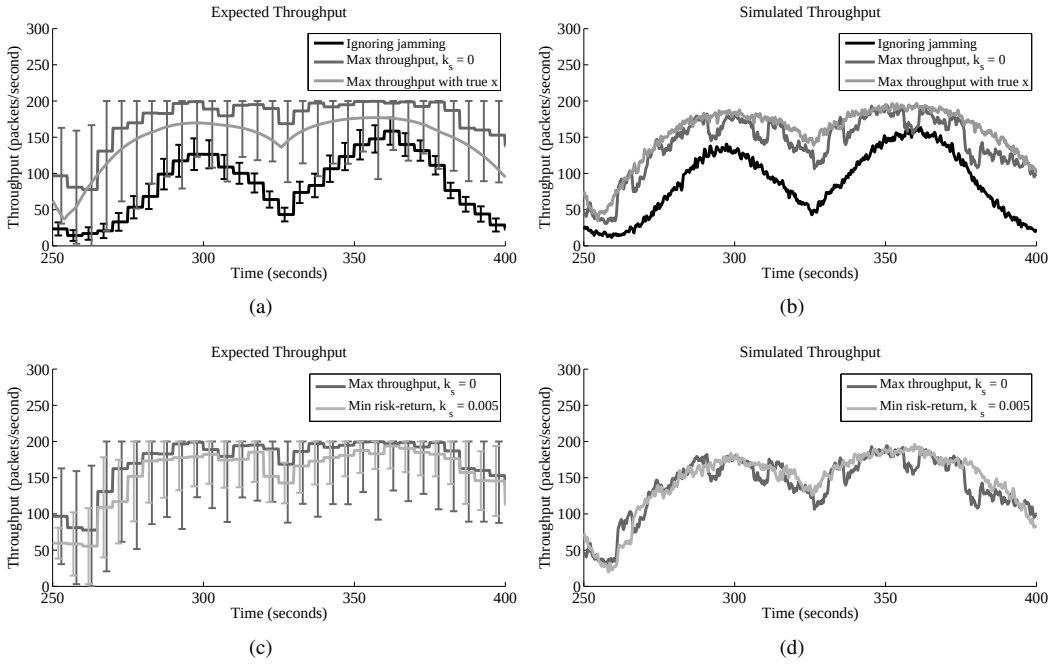

Fig. 5. In (a) and (b), Case I, Case II, and Case IV are compared in terms of the optimal expected throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$ and the actual achieved throughput $\boldsymbol{y}_s^T \boldsymbol{\phi}_s$. In (c) and (d), Case II is similarly compared with Case III. The error bars in (a) and (c) indicate one standard deviation $\sqrt{\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s}$ above and below the mean, limited by the maximum rate of 200 $pkts/s$.

value for each trial. In addition to comparing the expected throughput for various parameter values, we compute the Sharpe ratio [19], given by the ratio of the expected throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$ to the standard deviation $\sqrt{\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s}$, measuring the throughput-per-unit-risk achievable by the different methods. To ensure that the observed trends are due to the intended parameter variation, we simulate a simple network topology similar to that given in Figure 2. Figure 6 illustrates the trends in expected throughput, throughput variance, and Sharpe ratio as the update relay period T_s and the number of routing paths $|\mathcal{P}_s|$ increase. Since increased update times lead to increased variance, as previously seen in Figure 4(d), the Sharpe ratio decreases with increasing T_s. Figure 6(c) illustrates the improvement in throughput due to increased routing diversity.

VI. Conclusion

In this article, we studied the problem of traffic allocation in multiple-path routing algorithms in the presence of jammers whose effect can only be characterized statistically. We have presented methods for each network node to probabilistically characterize the local impact of a dynamic jamming attack and for data sources to incorporate this information into the routing algorithm. We formulated multiple-path traffic allocation in multi-source networks as a lossy network flow optimization problem using an objective function based on portfolio selection theory from finance. We showed that this centralized optimization problem can be solved using a distributed algorithm based on decomposition in network utility maximization (NUM). We presented simulation results to illustrate the impact of jamming dynamics and mobility on network throughput and to demonstrate the efficacy of our traffic allocation algorithm. We have thus shown that multiple-path source routing algorithms can optimize the throughput performance by effectively incorporating the empirical jamming impact into the allocation of traffic to the set of paths.

References

[1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," *Computer Networks*, vol. 47, no. 4, pp. 445–487, Mar. 2005.

[2] E. M. Sozer, M. Stojanovic, and J. G. Proakis, "Underwater acoustic networks," *IEEE Journal of Oceanic Engineering*, vol. 25, no. 1, pp. 72–83, Jan. 2000.

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

[4] 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.

[5] 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.

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

[7] 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.

[8] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: Attack and defense strategies," *IEEE Network*, vol. 20, no. 3, pp. 41–47, May/Jun. 2006.

[9] D. B. Johnson, D. A. Maltz, and J. Broch, *DSR: The Dynamic Source Routing Protocol for Multihop Wireless Ad Hoc Networks*. Addison-Wesley, 2001, ch. 5, pp. 139–172.

[10] E. M. Royer and C. E. Perkins, "Ad hoc on-demand distance vector routing," in *Proc. 2nd IEEE Workshop on mobile Computing Systems and Applications (WMCSA'99)*, New Orleans, LA, USA, Feb. 1999, pp. 90–100.

[11] R. Leung, J. Liu, E. Poon, A.-L. C. Chan, and B. Li, "MP-DSR: A QoS-aware multi-path dynamic source routing protocol for wireless ad-hoc networks," in *Proc. 26th Annual IEEE Conference on Local Computer Networks (LCN'01)*, Tampa, FL, USA, Nov. 2001, pp. 132–141.

[12] H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–92, Mar. 1952.

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

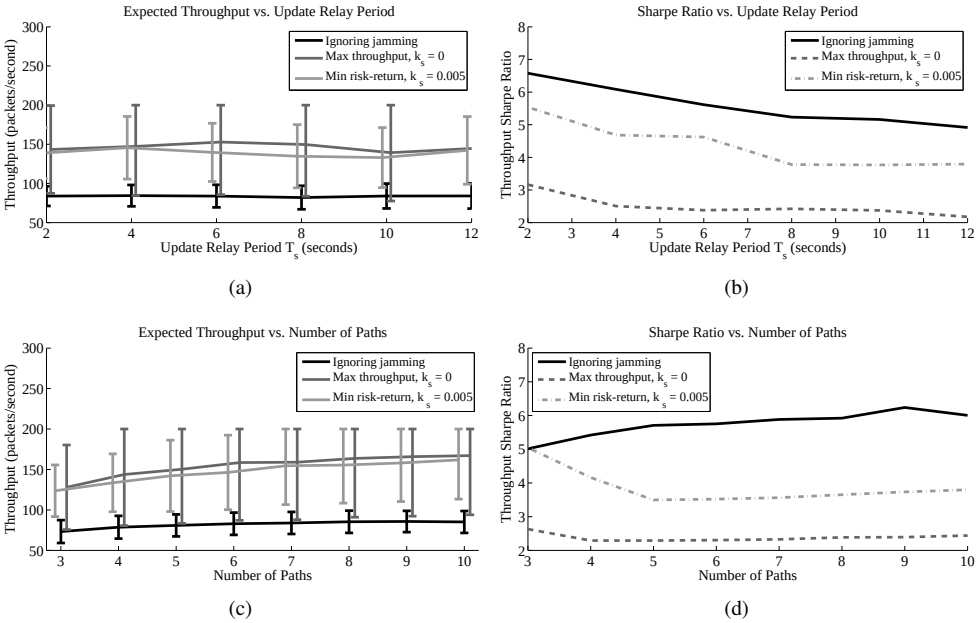

Fig. 6. The expected throughput is computed for Cases I, II, and III with varying update relay period T_s in (a) and (b) and number of routing paths $|\mathcal{P}_s|$ in (c) and (d). In (a) and (c), the expected throughput $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s$ is illustrated with error bars to indicate one standard deviation $\sqrt{\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s}$ around the mean, limited by the maximum rate of 200 $pkts/s$. In (b) and (d), the Sharpe ratio $\boldsymbol{\gamma}_s^T \boldsymbol{\phi}_s / \sqrt{\boldsymbol{\phi}_s^T \boldsymbol{\Omega}_s \boldsymbol{\phi}_s}$ is illustrated.

[14] 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.

[15] M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions*, 3rd ed. New York: John Wiley & Sons, Inc., 2000.

[16] S. W. Roberts, "Control chart tests based on geometric moving averages," *Technometrics*, vol. 42, no. 1, pp. 97–101, Feb. 2000.

[17] V. Paxson and M. Allman, "Computing TCP's retransmission timer," RFC 2988, Nov. 2000, http://www.ietf.org/rfc/rfc2988.txt.

[18] I. R. James, "Products of independent beta variables with applications to connor and mosimann's generalized dirichlet distribution," *Journal of the American Statistical Association*, vol. 67, no. 340, pp. 910–912, Dec. 1972.

[19] W. F. Sharpe, *Investors and Markets: Portfolio Choices, Asset Prices, and Investment Advice*. Princeton University Press, 2007.

Patrick Tague received BS degrees in Computer Engineering and Mathematics from the University of Minnesota in 2003 and MS and PhD degrees in Electrical Engineering from the University of Washington in 2007 and 2009, respectively, as a member of the Network Security Lab (NSL). He is an Assistant Research Professor at Carnegie Mellon University, Silicon Valley Campus. His research interests include security and reliability of wireless ad-hoc and sensor networks. Dr. Tague received the Yang Research Award in 2009 for outstanding graduate research in the Electrical Engineering Department at the University of Washington and the Outstanding Graduate Research Award from the UW Center for Information Assurance and Cybersecurity in 2009.

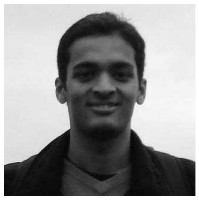

Sidharth Nabar received the Bachelors degree in Electrical and Electronics Engineering from National Institute of Technology Surathkal, India in 2007 and the MS degree in Electrical Engineering from University of Washington, Seattle, in 2009. He is a PhD candidate in the department of Electrical Engineering, University of Washington and is a member of the Network Security Lab (NSL). His current research interests include medical sensor networks and security in wireless sensor networks.

James A. Ritcey received the BSE degree from Duke University, the MSEE degree from Syracuse University, and the PhD degree in Electrical Engineering (communication theory and systems) from the University of California, San Diego in 1985. Since 1985, he has been with the Department of Electrical Engineering at the University of Washington, where he now holds the rank of Professor. From 1976 to 1981 he was with the General Electric Company, and graduated from GEs Advanced Course in Engineering. His research interests include communications and statistical signal processing for radar and underwater acoustics. He has published extensively in these areas. Professor Ritcey served as the General Chair of the 1995 International Conference on Communications in Seattle. He also served as Technical Program Chair of the 1992 and General Chair of the 1994 Asilomar Conference on Signals, Systems, and Computers and is currently a member of the Steering Committee.

Radha Poovendran received the PhD degree in Electrical Engineering from the University of Maryland, College Park, in 1999. He is a Professor and founding Director of the Network Security Lab (NSL), Electrical Engineering Department, University of Washington, Seattle. His research interests are in the areas of applied cryptography for multiuser environment, wireless networking, and applications of information theory to security. He is a coeditor of the book *Secure Localization and Time Synchronization in Wireless Ad Hoc and Sensor Networks* (Springer-Verlag, 2007). Dr. Poovendran was a recipient of the NSA Lucite Rising Star Award and Faculty Early Career Awards, including the National Science Foundation CAREER Award in 2001, the Army Research Office YIP Award in 2002, the Office of Naval Research YIP Award in 2004, PECASE in 2005 for his research contributions to multiuser security, and a Graduate Mentor Recognition Award from the University of California San Diego in 2006. He co-chaired the first ACM Conference on Wireless Network Security (WiSec) in 2008.