ABSTRACT

Wireless mesh networks (WMNs) are interconnected systems of wireless access points (APs) that provide untethered network connectivity for a group of users who require data, voice, and/or video communication. The wireless access medium of a WMN makes it vulnerable to electromagnetic attack and interference. We apply the game-theoretic defender-attacker-defender (**DAD**) optimization modeling technique to the design, attack, and operation of a WMN. We present a sampling-based algorithm and associated decision-support tool that can quickly prescribe WMN topologies to minimize the worst possible disruption an adversary can inflict through deliberate interference, i.e., jamming. Our model considers radio operating characteristics, the relative importance of client coverage and network throughput, and the effects of radio propagation over terrain. We implement our solution technique in a decision-support tool that runs on a laptop, does not require commercial solvers or other add-ins, and can use terrain information freely downloaded from the Internet. To our knowledge, we are the first to apply the **DAD** framework to the problem of WMN topology design.

INTRODUCTION

Wireless mesh networks (WMNs) are interconnected systems of wireless access points (APs) that provide untethered network connectivity for a group of users who require data, voice, and/or video communication. Each AP has two radio devices: the first connects to local client devices, such as laptops and smartphones; the second connects to other APs to create a backhaul network. To function, APs require only a local power source, such as a battery or portable generator. This makes WMNs well-suited to operations in austere environments such as combat and humanitarian assistance and/or disaster relief (HA/DR) operations. The wireless access medium of a WMN makes it particularly vulnerable to attack and exploitation (Mpitziopoulos et al., 2009; Pelechrinis et al., 2011). Such actions may include passive eavesdropping and packet capture, spoofing trusted identities to gain unauthorized access to the network, injecting malicious code, or denial of service (DoS) attacks (Xu et al., 2004). During physical-layer noise jamming DoS attacks that we consider, an attacker constantly broadcasts noise on the same radio frequency (RF) channel(s) used by the WMN in an attempt to overpower the friendly signal, degrading or denying use of the channel(s) (Pelechrinis et al., 2011; Poisel, 2011; Vakin et al., 2001; Xu et al., 2005). Powerful commercial and military jamming systems are readily available, but this type of attack can be conducted with inexpensive equipment and little technological prowess, and can be very challenging to defend against (Mpitziopoulos et al., 2009; Xu et al., 2004; *The Economist*, 2011; IET, 2013; Wood et al., 2003). Even unintentional interference can be as harmful as an intentional attack (see, e.g., Cox (2007)). Hence, jamming is of increasing concern in both civilian and military operating environments (Caro, 2007).

Designers of WMNs employ various strategies to defend against such threats, including frequency hopping and spread spectrum techniques, filtering noisy connections, adjusting the signal-to-noise ratio threshold, and various other security protocols (Poisel, 2011; Ståhlberg, 2000; Zhang et al., 2008).

We describe a method for quickly designing WMN physical topologies (i.e., the placement of APs) that are inherently robust to the effects of deliberate jamming or other electromagnetic interference (EMI) emanating from point sources (i.e., jammers). Our method considers constraints on network service and the effects of radio propagation over terrain. Although we focus on intentional noise jamming, our technique can be generalized to any form of WMN interference in which network performance is a function of the physical distance between interference sources, WMN APs, and client devices.

There has been much recent research in defending WMNs from jamming attacks. Xu et al. (2005) find that devices that constantly jam (which we assume) are more prone to detection; they develop algorithms to improve the classification rate of jamming attacks. Wood et al. (2003) describe a method of mapping areas affected by physical-layer jamming to avoid placing sensors in these denied areas. Ståhlberg

Fast Design of Wireless Mesh Networks to Defend Against Worst-Case Jamming

Paul J. Nicholas

Johns Hopkins University Applied Physics Laboratory, paul.nicholas@jhuapl.edu

David L. Alderson

Naval Postgraduate School, dlalders@nps.edu

APPLICATIONS AREAS: Command and Control, Land and Expeditionary Warfare, Modeling, Simulation and Wargaming, Computing Advances in Military OR

OR METHODS: Nonlinear Programming, Multi-objective Optimization, Network Methods

(2000) and Lazos and Krunz (2011) each recommend several methods of increasing the robustness of wireless networks to attack, including the use of directional antennae and frequency hopping, but neither specifically consider defensive placement of APs. Xu (2007) examines the effectiveness of adjusting transmission power to avoid jamming, but she assumes jammers will operate at a transmission power less than that of the APs (we make no such assumption). Xu et al. (2004) and Ma et al. (2005) examine spatial retreats, i.e., moving APs physically away from the sources of interference, as a form of defense against a jamming attack. However, neither consider jammers that could then move and attack the newly configured network. As Mpitziopoulos et al. (2009) observe, this type of defense is ineffective against an adversary that can again move jammers.

The conflicting interests of a network designer and attacker in respectively maximizing and minimizing network performance make this problem a natural candidate for the use of game theory. Thamilarasu and Sridhar (2009) consider the use of game theory in modeling optimal jamming attack and detection strategies. However, they do not consider the actions taken by a network designer or defender, and both they and Srivastava et al. (2005) consider only *strategic-form games* (wherein players move simultaneously), vice *extensive form games* (wherein players move sequentially) that we consider (see Fudenberg and Tirole (1991) and Myerson (1991) for reviews of game theory).

Our game-theoretic approach to building a robust WMN topology is similar to a spatial retreat in that the only defensive method we consider to minimize the effects of jamming is to place an AP elsewhere. However, unlike any of the previous work that focuses on static or random jamming, we consider WMN network design in the presence of an *intelligent adversary* who observes the WMN and then places the jammer(s) to maximally disrupt network performance.

As noted in Wood et al. (2003), overcoming the effects of jamming can quickly escalate into a costly game of one-upmanship, where the network designer and adversary are constantly trying to outmaneuver each other. We adopt the game-theoretic defender-attacker-defender (**DAD**) methodology of Alderson et al. (2014, 2011), and Brown et al. (2006) to model the design, attack, and operation of a wireless mesh network. The application of **DAD** to our model can identify WMN topologies that minimize the worst possible damage an adversary can inflict, avoiding such endless competition. To our knowledge, we are the first to apply the **DAD** framework to the design of WMNs that are robust to jamming.

Next, we describe our jammer-cognizant model of WMN performance and our application of the **DAD** framework. In the following section, we describe our solution method and then we explore the behavior of our model and algorithm under various conditions. We conclude by describing areas of future research.

PROBLEM FORMULATION

Overview

Building on the notation of Nicholas and Alderson (2012, 2015), we define N to be the set of all AP nodes, indexed by $i = 1, 2, \ldots, n$, where $n = |N|$. We define M to be the set of all jammer nodes, indexed by $k = 1, 2, \ldots, m$, where $m = |M|$. Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n) \in \Lambda$ represent the locations of the APs, and let $\chi = (\chi_1, \chi_2, \ldots, \chi_m) \in \mathcal{X}$ represent the locations of the jammers. We define the *operating area* as the topographic area where an AP i or jammer k may be physically located. We assume that APs and jammers, once placed, remain stationary. We divide the operating area into a set of discrete coverage regions R, indexed by $r = 1, 2, \ldots, |R|$. Although our formulation allows the use of any discretization scheme, our implementation assumes rectangular regions arranged in a grid (see Figure 1). Each coverage region has an associated elevation that we assume is uniform throughout the region. This assumption is not true in practice, but is consistent with much of the available elevation data.

We measure the performance of a WMN in two dimensions: its ability to provide adequate client coverage (i.e., sufficient signal strength to users) in coverage regions, and the throughput of the backhaul network. Even in the absence of jammers, the designer of a WMN

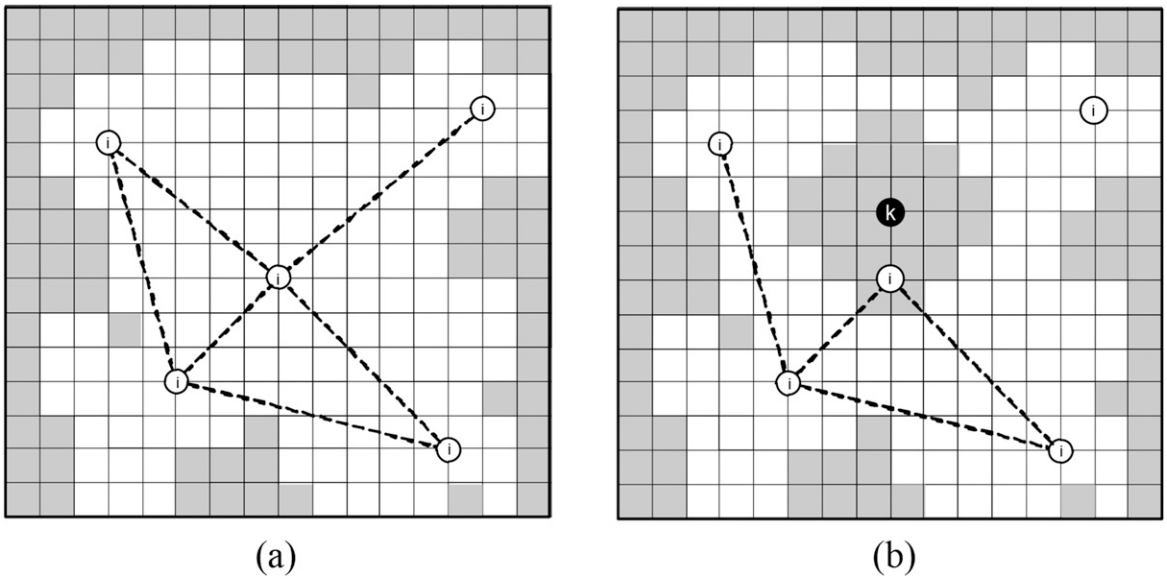

Figure 1. A representative discretized operating area and wireless mesh network. (a) White circles denote the location of access points, shaded regions denote the areas with insufficient client coverage, and dashed lines denote links in the backhaul network. (b) The placement of a jammer, denoted by a black circle, decreases client coverage and disrupts backhaul network connectivity.

faces a fundamental challenge in the placement of APs. The need to provide client coverage to a large number of regions suggests spreading out the APs. However, because wireless transmission capacity degrades with increasing distance between APs, the desire for high backhaul network throughput creates a drive to keep the APs close together. A ''good'' WMN topology balances these two conflicting objectives.

Figure 1(a) depicts a typical WMN in the absence of jamming. AP nodes are illustrated as white circles, and shaded grid elements represent coverage regions that receive insufficient client coverage from the APs. The coverage obtained at each grid location depends on several factors including the local terrain, AP and client radio characteristics, as well as any interference. The dashed lines represent the backhaul network used to communicate between AP nodes.

Figure 1(b) depicts a WMN in the presence of jamming. Here, a single jammer node (illustrated as a black circle) may have two active transmitters, one interfering with nearby APs and client devices (resulting in a greater number of shaded regions) and the other interfering with nearby AP backhaul network radios (effectively degrading or eliminating backhaul links).

Without loss of generality, we assume APs are not subject to self-jamming or interference from other APs. We assume jammers are conducting broadband noise or barrage jamming, where the jammer transmission power is spread across the entire targeted channel (Mpitziopoulos et al., 2009; Poisel, 2011). Such jammers provide lower power spectral densities (i.e., power per Hertz) than an equivalent tone or single-channel jammer because transmission power is spread over a larger frequency range (Poisel, 2011). Following Ståhlberg (2000), we assume the effects of multiple jammers are perfectly additive at the respective receivers. Hence, barrage jamming cannot be overcome by simply placing an additional, redundant AP, as all APs are subject to the same interference.

DAD Model

We apply the **DAD** methodology of Alderson et al. (2014, 2011) and Brown et al. (2006) to model the design, attack, and operation of a WMN. In our version of this three-stage, sequential Stackelberg game (von Stackelberg,

1952), the defender-as-designer, or simply designer **D**, first places n APs in the operating area. In the second stage, the attacker **A**, cognizant of the AP topology, places m jammers to disrupt client coverage and total delivered flow. In the final stage, the defender-as-operator, or simply operator $\underline{\mathbf{D}}$ calculates client coverage and sends traffic flow across the network (in reality, the operator is a routing algorithm computed by the APs).

The optimal solution to our **DAD** problem identifies the locations of APs to create a WMN that is the most robust to the worst possible jamming attack. Such an attack could represent the actions of a rational human opponent, or the worst-case positioning of unintentional interference sources such as civilian radios, other RF devices, or high-voltage electrical devices. We describe each stage of **DAD**, beginning at the innermost stage.

The Operator's Problem: Client Coverage and Network Throughput. We model the operator's problem as a modification of the simultaneous routing, resource allocation, and coverage (SRRA+C) problem (Nicholas and Alderson, 2012).

Given fixed AP locations $\hat{\lambda}$ and fixed jammer locations $\hat{\chi}$, the operator $\underline{\mathbf{D}}$ computes the shortfall in client coverage (i.e., the number of regions whose delivered signal strength falls short of the required level) denoted $Z_{coverage}(\hat{\lambda}, \hat{\chi})$ and then selects variables F that determine network flow, denoted $Z_{flow}(\hat{\lambda}, \hat{\chi}, F)$. The objective is to minimize the combination of this shortfall and negative network flow (i.e., maximize positive network flow) by choice of flow variables $F \in \mathcal{F}$.

Explicitly, the operator's problem is:

$$Z_{\underline{\mathbf{D}}}(\hat{\lambda}, \hat{\chi}) = \min_{F \in \mathcal{F}} \big(Z_{coverage}(\hat{\lambda}, \hat{\chi}) - w Z_{flow}(\hat{\lambda}, \hat{\chi}, F) \big). \quad (1)$$

where w is a positive scalar representing the relative importance of network flow. We use $w = 1$, meaning coverage shortfall and network flow are equally weighted.

The feasible region $\mathcal{F}$ for network flow variables includes constraints for the balance of flow in the backhaul network, as well as relationships between transmission power, transmission capacity, and coverage. See the appendix for a complete specification of the operator's problem.

The Attacker's Problem: Placing Jammers. The attacker **A**, given fixed AP node locations $\hat{\lambda}$, wishes to maximize coverage shortfall and minimize delivered network flow by placing jammer nodes at locations χ:

$$Z_{\mathbf{A}\underline{\mathbf{D}}}(\hat{\lambda}) = \max_{\chi \in \mathcal{X}} \min_{F \in \mathcal{F}} \big(Z_{coverage}(\hat{\lambda}, \chi) - w Z_{flow}(\hat{\lambda}, \chi, F) \big). \quad (2)$$

Here, the feasible region $\mathcal{X}$ for attacks is constrained by the operating area and a limit on the number of jammers.

Note this is a two-stage problem (**A**$\underline{\mathbf{D}}$), as the operator's problem (1) is solved after the attacker chooses $\chi \in \mathcal{X}$. Shankar (2008) similarly considers the deliberate placement of jammers by an intelligent adversary to maximally disrupt WMN operations. He models network flow using the simultaneous routing and resource allocation (SRRA) problem (Xiao et al., 2004) but does not consider client coverage. He exhaustively enumerates a fixed number of candidate locations for jammers, whereas we consider a continuous space for jammer placement. We compare these search methods in a later section.

The Designer's (Restricted) Problem: Placing Access Points. The network designer **D** wishes to minimize coverage shortfall and maximize delivered network flow by placing AP nodes at locations $\lambda \in \Lambda$.

Consider the simplified situation where the designer knows in advance the fixed jammer node locations $\hat{\chi}$. In this restricted case, the placement problem is:

$$Z_{\mathbf{D}\underline{\mathbf{D}}}(\hat{\chi}) = \min_{\lambda \in \Lambda} \min_{F \in \mathcal{F}} \big(Z_{coverage}(\lambda, \hat{\chi}) - w Z_{flow}(\lambda, \hat{\chi}, F) \big). \quad (3)$$

Note the original SRRA+C problem (Nicholas and Alderson, 2012) is special case of the designer's problem (3) with no jammer nodes. However, the real challenge to the network designer is that she does not know in advance the location of the jammer nodes.

Overall SRRA+C DA$\underline{\mathbf{D}}$ Problem. By nesting the problems of the operator, attacker, and

designer, we obtain the overall SRRA+C **DAD** formulation:

$$Z_{\mathbf{DA\underline{D}}} = \min_{\lambda\in\Lambda}\max_{\chi\in\mathcal{X}}\min_{F\in\mathcal{F}}\big(Z_{coverage}(\lambda,\chi) - wZ_{flow}(\lambda,\chi,F)\big). \quad (4)$$

The designer **D** first chooses AP locations λ, which the attacker **A** then aims to maximally disrupt by placing jammers at locations χ. Given AP and jammer locations, the operator **D** calculates client coverage and determines how best to route traffic. By allowing the designer to move first in this sequential Stackelberg game, we assume the designer is operating in an area that will subsequently be subject to jamming. Had we allowed the attacker to move first (i.e., **AD**$\underline{\mathbf{D}}$), we would assume the designer is being forced to operate in an area already being jammed.

The solution to the SRRA+C **DAD** problem indicates where the network designer should place APs to minimize the worst-case disruption possible by EMI. That is, when solved to optimality, the obtained AP network topology is completely immune to greater degradation, as the attacker cannot possibly do more damage without additional resources. Note the converse is not true. Because we assume the designer (with perfect information of the worst possible attack) places his APs first, it is possible (indeed, likely) that the designer could improve upon this design given fixed jammers. Likewise, if we allow the attacker to move first (**AD**$\underline{\mathbf{D}}$), it is likely he could improve upon his attack given fixed APs. In other words, we find a Stackelberg equilibrium but not a Nash equilibrium (Fudenberg and Tirole, 1991; Cruz, 1975), as the designer could likely unilaterally improve his strategy after the attacker's move.

SOLUTION METHOD

Solving the Operator's Problem

We solve the operator's problem (1) by calculating its two components separately. Calculating client coverage $Z_{coverage}$ is a straightforward series of calculations based on input data. Solving the SRRA problem (Xiao et al., 2004) to calculate the value of network flow Z_{flow} is more challenging. Xiao et al. (2004) observe that the SRRA problem has special structure that allows it to be solved using dual decomposition. We use the same approach to solve the problem using the subgradient method (Bertsekas, 1999), stopping after a given number of iterations. See Nicholas and Alderson (2012, 2015) and Nicholas (2009) for further details on calculating client coverage and our SRRA solution technique.

Solving the Attacker's and Designer's Problems

The attacker's and designer's problems (like the SRRA+C problem) are nondifferentiable, nonconvex, nonlinear optimization problems. The difficulty of finding exact solutions to such problems increases the desirability of using heuristic computational techniques such as genetic or simulated annealing algorithms, and sampling algorithms such as mesh adaptive direct search (Audet and Dennis, 2006). In our previous work, we use the DIviding RECTangles (DIRECT) algorithm of Jones et al. (1993) to sample the SRRA+C solution space (i.e., the designer's problem with no jammers) to quickly find solutions. This same approach will work for the attacker's problem $Z_{\mathbf{A\underline{D}}}$ (2), given fixed AP nodes, and for the designer's problem $Z_{\mathbf{D\underline{D}}}$ (3), given fixed jammers.

DIRECT is a sampling optimization algorithm based on Lipschitzian optimization (Horst and Hoang, 1996). The algorithm iteratively samples from the continuous, hyper-rectangular solution space, where the number of dimensions is $2m$ (attacker's problem) or $2(n-1)$ (designer's problem), the length of each dimension is proportional to the operating area length or width, and a single point in the solution space represents the locations of all nodes being placed. The algorithm progressively samples from and divides the solution space into smaller hyper-rectangles. At each step, it chooses to explore a particular sub-hyper-rectangle based on both the solution value of the center point and the total volume of the given partition, where larger volumes are more desirable because they indicate greater unexplored territory and hence greater potential for an improved incumbent solution. DIRECT is guaranteed to eventually find

the globally optimal solution to a continuous problem, though this convergence may be slow due to the effects of the curse of dimensionality (Bellman, 1961; Hastie et al., 2009). In practice, we find reasonably good solutions to networks of six APs and three jammers after relatively few (e.g., 10–15) DIRECT iterations.

Solving the DAD Problem

To solve the SRRA+C **DAD** problem, we cannot simply use one large instance of DIRECT to search concurrently for good AP locations λ and jammer locations χ, as the attacker and designer are playing against each other and have opposing (i.e., maximization and minimization) goals. Instead, we follow Alderson et al. (2011) and decompose the **DAD** problem into a designer **D** master problem with separate attacker **A** subproblems. We solve using nested instances of DIRECT, each with its own objective function. We present pseudocode in Algorithm DIRECT for SRRA+C **DAD**. The master problem uses DIRECT to choose AP locations λ_u (step 6) for each iteration $u = 1, 2, \ldots,$ *max_master_iterations*. For those given AP locations, another instance of DIRECT is initialized to solve the associated subproblem, choosing jammer locations χ_v (step 10) for each iteration $v = 1, 2, \ldots,$ *max_sub_iterations*. Given AP locations λ_u and jammer locations χ_v, the overall objective value is then obtained by solving the operator's problem (1) (step 11). The best attack (i.e., the highest overall objective value) is stored as the incumbent (steps 12–15). After *max_sub_iterations*, the subproblem returns jammer locations $\tilde{\chi}$ yielding the best attack found. The master problem continues searching for the best AP locations λ^* to minimize the damage caused by the worst attack found, storing the best incumbent design (steps 18–22), until *max_master_iterations*.

For AP locations λ_u and given enough iterations, DIRECT will eventually find a solution within an arbitrary distance of the optimal jamming attack. In practice, we are constrained by time and the computational limits of our computer implementation. Note that we cannot precisely calculate the optimality gap of any particular solution, but this is ameliorated by the practical need to identify ''good'' solutions quickly for time-sensitive network design in support of HA/DR or combat operations.

Algorithm DIRECT for SRRA+C DAD

1. **begin**
2. Store map data
3. Initialize $u \leftarrow 1$
4. **Master problem (*Designer*)**
5. **while** ($u <$ *max_master_iterations*) **do**
6. Calculate AP locations λ_u using DIRECT
7. Initialize $v \leftarrow 1$
8. ***Subproblem* (*Attacker*)**
9. **while** ($v <$ *max_sub_iterations*) **do**
10. Calculate jammer locations χ_v using DIRECT
11. Solve operator's problem $Z_{\underline{\mathbf{D}}}$ for λ_u and χ_v
12. **if** $Z_{\underline{\mathbf{D}}}(\lambda_u, \chi_v) > Z_{\underline{\mathbf{D}}}(\lambda_u, \tilde{\chi})$//If best attack yet, store as incumbent
13. $\tilde{\chi} \leftarrow \chi_v$
14. $Z_{\underline{\mathbf{D}}}(\lambda_u, \tilde{\chi}) \leftarrow Z_{\underline{\mathbf{D}}}(\lambda_u, \chi_v)$
15. **endif;**
16. $v \leftarrow v+1$
17. **end;**
18. **if** $Z_{\underline{\mathbf{D}}}(\lambda_u, \tilde{\chi}) < Z_{\underline{\mathbf{D}}}(\lambda^*, \chi^*)$//If best design yet, store as incumbent
19. $\lambda^* \leftarrow \lambda_u$
20. $\chi^* \leftarrow \tilde{\chi}$
21. $Z_{\underline{\mathbf{D}}}(\lambda^*, \chi^*) \leftarrow Z_{\underline{\mathbf{D}}}(\lambda_u, \tilde{\chi})$
22. **endif;**
23. $u \leftarrow u+1$
24. **end;**
25. **end;**
26. Return AP locations λ^*, jammer locations χ^*, and operator's solution $Z_{\underline{\mathbf{D}}}(\lambda^*, \chi^*)$

COMPUTATIONAL RESULTS

Building on the software we initially developed for solving the SRRA+C problem (Nicholas, 2009), we implement our algorithm for this **DAD** problem using Microsoft Visual C++. Our decision-support tool runs on a laptop, does not require commercial solvers or other add-ins, and can use terrain information freely downloaded from the Internet.

Many factors affect the shape of the SRRA+C solution space, including the technical

characteristics of the APs and jammers, their relative numbers and signal strengths, the type and strength of jamming, the amount of overlap in client coverage, the effects of terrain on EM propagation, and the assignment of traffic destination nodes. An exhaustive exploration of these factors is beyond the scope of this paper. In the following analyses, we follow Wood et al. (2007) and assume each radio in each AP and the associated radio in each jammer are identical, transmitting with the same output power and similar antennae. We model our AP and jammer radio characteristics on the Cisco Aironet 1550 WMN AP, and our client devices on a generic internal laptop 802.11n wireless interface card. We begin with a simple analysis on flat ''tabletop'' terrain to gain intuition on optimal jamming and defense strategies, and then consider a realistic case study using actual terrain data. We also provide a brief performance analysis of our method.

Behavior on Artificial Flat Terrain

We explore the attacker's problem (2) by finding the optimal single jammer attack against a network of two fixed APs. Consider a 1 square kilometer operating area (gridded into 100 $\times$ 100 regions) with flat terrain, with an AP placed near the top and bottom of the region (Figure 2(a)). With no jammer present, these two APs (depicted as black circles) will provide the client coverage shown in white and deliver network traffic to each other (the solid line) at a maximum rate of 419 kbps. Darker areas indicate areas of increasing client coverage shortfall.

We enumerate solutions by placing the jammer in each coverage region and solving the operator's problem (1). In a single-channel jamming attack, the optimal attack is to simply place the single jammer directly on top of either AP, depicted as an X on the bottom AP in Figure 2(b). This direct-AP attack eliminates client coverage by the bottom AP, and reduces network traffic flow between the APs to essentially zero. Shankar (2008) uses the direct-AP attack with barrage jammers, but he does not consider client coverage. In our model, we consider client coverage and often find that the optimal barrage jamming attack (as opposed to a single-channel attack) is to place jammers in a between-AP attack, such as in Figure 2(c). In such a location, the jammer is able to concurrently reduce the client coverage provided by both APs and reduce the delivered network traffic to essentially zero. In the rest of this paper, we consider only barrage jammers.

The between-AP attack may at first seem counterintuitive, as the center of operating area in Figure 2(a)) receives less client coverage than the area immediately surrounding each AP; it may seem this center area has ''less to lose'' than an attack directly on each AP. However, recall our formulation penalizes the degree of coverage

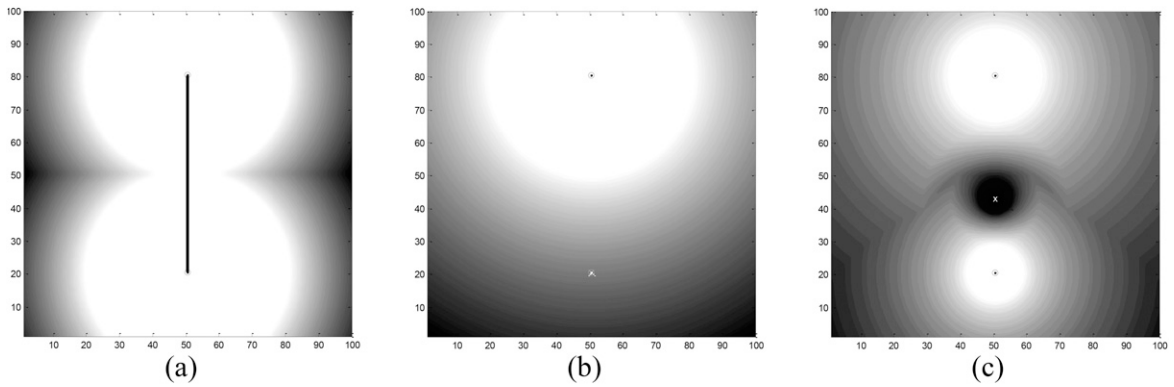

Figure 2. Optimal solution for two APs on flat terrain. White areas indicate sufficient client coverage where client devices are able to connect to APs. Darker areas indicate progressively worse client coverage shortfall. (a) Client coverage provided by two APs (black circles) with a single backhaul network link (solid line) in the absence of jammers. (b) The optimal single-channel jamming attack targets one of the two nodes (at the bottom). (c) A barrage jamming attack places the jammer (the X) in between the two APs.

shortfall. By placing the jammer in between each AP, the jammer maximizes this penalty by making already-deficient client coverage that much worse. Additionally, network flow is maximally disrupted in a between-AP attack because this reduces delivered flow to both APs concurrently.

Case Study in Fort Ord, California

While simple rules-of-thumb such as ''place jammers between APs'' may be useful when designing WMNs for flat surfaces, the effects of terrain greatly complicate the problem. We conduct a case study on Fort Ord, California terrain using our algorithm to consider these effects. Our operating area covers 116 acres, gridded into a $73 \times 73 = 5{,}329$ coverage regions. The area has gently rolling hills, a large parking lot, a stadium, and several roads. We use elevation data from the National Elevation Dataset (NED) (USGS, 2013). Figure 3(a) is an elevation contour plot, and Figure 3(b) is a Google Maps (2013) image of the area.

Enumeration Analysis. To demonstrate the nonlinearities of operations on real terrain, we first consider the effect of placing one jammer among four fixed APs arranged in a square about 160 meters across (open circles in Figure 3). We enumerate the placement of the jammer at each of the 5,329 regions r. The shading at each point represents client coverage (Figure 3(c)) and network flow (Figure 3(d)), where dark represents less desirable service (i.e., more effective jamming). Unlike the results on flat terrain, these results are highly nonlinear and cannot be prescribed using simple rules-of-thumb. For example, placing the jammer at the lower-left AP location provides only moderate client coverage jamming, but provides the most effective network flow jamming among the four AP locations.

DIRECT Analysis. Next, we use DIRECT to examine the solutions to the unjammed, undefended, and defended networks consisting of four, five, and six APs in Figures 4, 5, and 6, respectively. We present the results for each case to include the best unjammed solution found (i.e., the designer's problem without jammers); the worst possible jamming attack against this undefended solution (i.e., the solution to the attacker's problem (2)); and the best placement of APs that minimizes the effects of the worst-case jamming attack found (i.e., the solution to the **DAD** problem (4). We run DIRECT until the solution objective values have not changed significantly for more than 10 function evaluations, or 20 master and subproblem iterations of DIRECT (whichever occurs first).

The results for the case of four APs and one jammer are depicted in Figure 4. Figure 4(a) is the unjammed solution. In the worst-case attack against the undefended network (Figure 4(b)), the attacker places a single jammer between two APs and causes considerable damage to the network. In the interference-robust design (Figure 4(c)), the designer chooses locations for the APs that reduce the damage done to client coverage even when the attacker again chooses a between-AP attack. Note that because the jammers have the same operating characteristics as

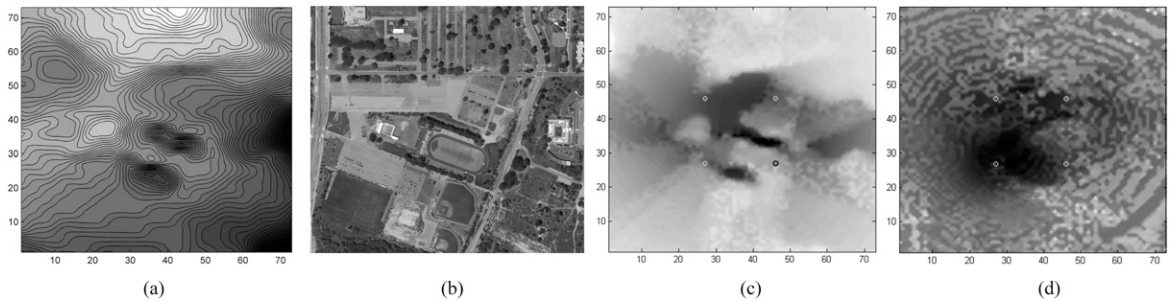

Figure 3. Case study of the 116-acre operating area on Ft Ord, California: (a) elevation contour map, (b) Google Maps image, (c) contour plot of client coverage values, and (d) network flow values. The shade at each location in (c) and (d) indicates the overall client coverage value or network flow value when a jammer is placed at that location (more effective jamming attacks are indicated by darker shading).

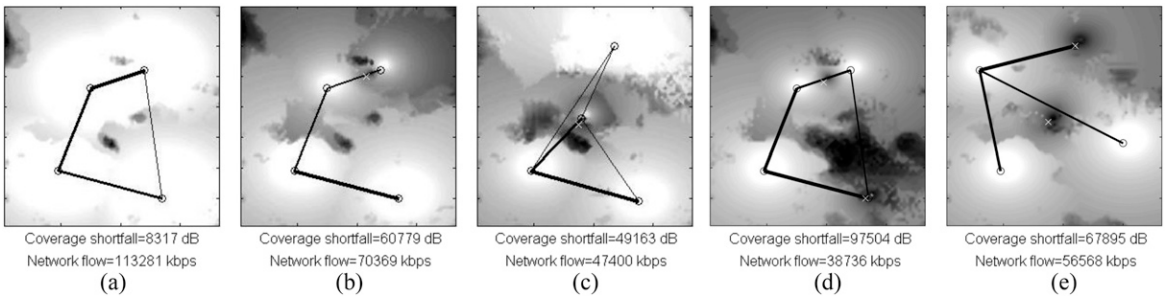

Figure 4. Analysis on Ft. Ord terrain with four APs. The shade at each location indicates the overall client coverage value when a jammer is placed at that location (more effective jamming attacks are indicated by darker shading). Solid lines indicate the backhaul network formed among APs: (a) selected placement of APs in the absence of any jammers; (b) selected attack of single jammer against fixed AP locations; (c) selected placement of APs anticipating a single jammer, along with selected attack for that jammer; (d) selected attack of two jammers against fixed AP locations; and (e) selected placement of APs anticipating two jammers, along with selected attack for those jammers. With one jammer, the attacker chooses a between-AP attack; with two jammers, the attacker chooses a between-AP attack and a direct-AP attack.

the APs, the only way to completely eliminate client coverage is to place a jammer directly on top of an AP. With two jammers (Figure 4(d) and 4(e)), the attacker places the first jammer in a position near that chosen in the one jammer scenario, and the second jammer in a direct-AP attack.

In Figure 5, we consider networks of five APs. Figure 5(a) is the unjammed solution, along with the solution in the presence of two (Figures 5(b) and 5(c)) and three (Figures 5(d) and 5(e)) jammers. The worst-case jamming attack against each undefended solution chooses direct-AP attacks for each jammer. These results illustrate a tension when placing jammers: as a jammer gets nearer an AP, it more effectively jams that AP but less effectively jams distant APs. As the ratio of jammers to APs increases, the direct-AP attack becomes more attractive because this tension slackens: distant APs are more likely to already be effectively jammed. In the interference-robust **DAD** solution, the designer places the APs farther apart.

In Figure 6, we consider networks of six APs. Figure 6(a) is the unjammed solution, along with the solution in the presence of two (Figures 6(b) and 6(c)) and three (Figures 6(d) and 6(e)) jammers. With six APs and two jammers (Figure 6), the worst-case attack against the undefended network is again a direct AP attack. In the interference-robust **DAD** solution, the worst-case attack nearly severs the top two

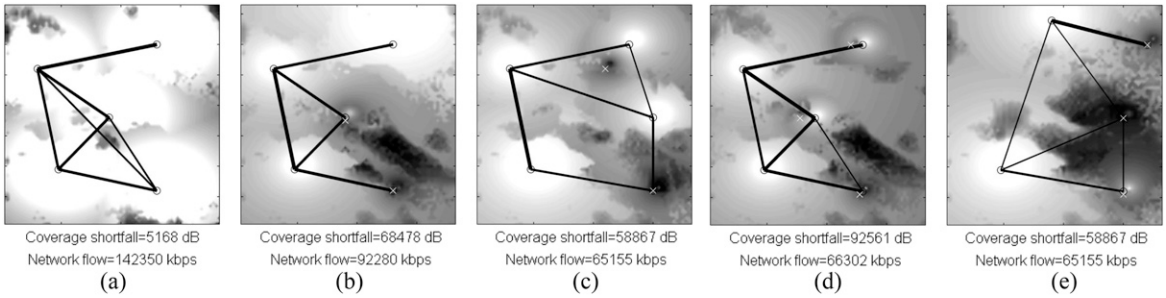

Figure 5. Analysis on Ft Ord terrain with five APs: (a) selected placement of APs in the absence of any jammers; (b) selected attack of two jammers against fixed AP locations; (c) selected placement of APs anticipating two jammers, along with selected attack for that jammer; (d) selected attack of three jammers against fixed AP locations; and (e) selected placement of APs anticipating three jammers, along with selected attack for those jammers. In both, the defender chooses to move APs farther apart to lessen the damage done by jamming.

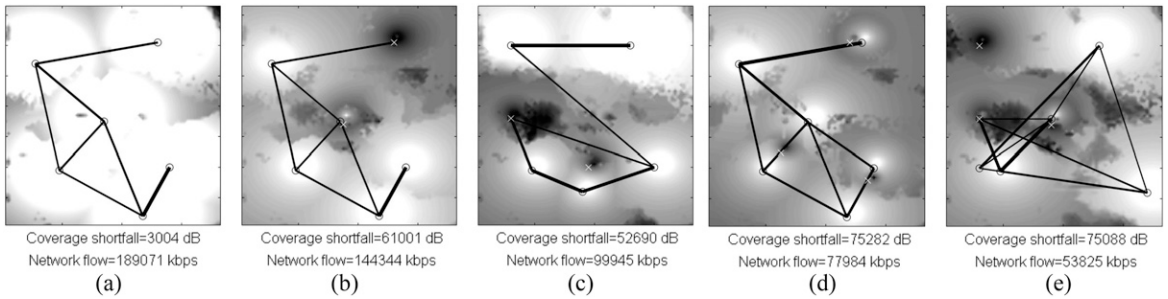

Figure 6. Analysis on Ft Ord terrain with six APs: (a) selected placement of APs in the absence of jammers; (b) selected attack of two jammers against fixed AP locations; (c) selected placement of APs anticipating two jammers, along with selected attack for that jammer; (d) selected attack of three jammers against fixed AP locations; and (e) selected placement of APs anticipating three jammers, along with selected attack for those jammers. In both, direct-AP attacks essentially deny client coverage in the middle-left section of the operating area.

APs from the rest of the network: traffic between these two APs is less than 8 kbps. In Figures 6(d) and 6(e), we consider six APs and three jammers. In this case, the attacker essentially denies use of the middle-left portion of the operating area.

Cost and Benefit of Anticipating Attacks. The results of this case study suggest that network planning that does not anticipate jamming can yield designs that are vulnerable. However, how many attacks should a planner anticipate, and is there a ''cost'' of planning for too many attacks?

In Table 1, we observe the importance of correctly anticipating the number of enemy jammers. For a given number of APs, we show the resulting objective value obtained for specific numbers of planned and actual jammers. Looking across each row, we observe in all cases that performance gets worse with an increasing number of jammers, as one would expect.

The upper-right triangle of Table 1 reflects the consequences of underestimating the number of actual jammers, whereas the lower-left triangle of Table 1 reflects the consequences of

Table 1. Importance of correctly anticipating the number of enemy jammers.

		Actual # jammers							
#APs	**Planned # jammers**	**0**		**1**		**2**		**3**	
4	0	**8,331**		60,805	(630%)	97,579	(1071%)	129,617	(1456%)
	1	10,217	(23%)	**49,230**		91,633	(86%)	125,166	(154%)
	2	9,852	(18%)	40,501	(−18%)	**67,960**		104,044	(53%)
	3	13,789	(66%)	47,852	(−3%)	82,475	(21%)	**97,676**	
5	0	**5,199**		34,731	(568%)	68,595	(1219%)	92,655	(1682%)
	1	7,416	(43%)	**30,736**		49,846	(62%)	81,114	(164%)
	2	7,092	(36%)	32,888	(7%)	**46,775**		76,855	(64%)
	3	11,696	(125%)	43,934	(43%)	56,718	(21%)	**61,161**	
6	0	**3,066**		28,211	(820%)	61,162	(1895%)	107,770	(3415%)
	1	5,408	(76%)	**25,522**		57,305	(125%)	105,660	(314%)
	2	5,804	(89%)	27,989	(10%)	**52,861**		76,737	(45%)
	3	10,433	(240%)	36,556	(43%)	53,664	(2%)	**75,328**	

Diagonal values (in bold) are objective values corresponding to instances where the designer correctly anticipates the number of enemy jammers. Values in the upper right triangle reflect the consequences of underestimating the number of actual jammers (values in parentheses represent horizontal percentage difference with diagonal). Values in the lower-left triangle reflect the consequences of overestimating the number of actual jammers (values in parentheses represent vertical percentage difference with diagonal).

overestimating the number of actual jammers. In general, we observe that the ''cost'' of underestimating jammers (i.e., the difference between the diagonal and upper horizontal off-diagonal value) is much larger than the cost of overestimating the actual number of jammers (i.e., the difference between the diagonal and lower vertical off-diagonal value). Consider the case of five planned APs. The lower vertical off-diagonal values range from being 7–125 percent larger than the diagonal, while the upper horizontal off-diagonal values range from 62–1,682 percent larger.

Table 1 also contains some unexpected results. In particular, one might expect that with increasing vertical off-diagonal distance in the lower-right portion of the table, values are strictly increasing (i.e., the cost of overestimating the number of jammers increases with additional jammers). However, for the situation of four planned APs, this is not the case. We observe some instances where overestimating by more jammers actually reduces the off-diagonal cost (in the case of two planned jammer and zero actual jammers), as well as instances where overestimating the number of jammers results in objective values that are lower than the diagonal value (in the case of two or three planned jammers and only one actual jammer). We speculate that these values are attributed to uneven optimality gaps across the individual runs for this smallest problem instance (where small changes in discrete values are likely to have the biggest effect). However, understanding this more comprehensively is a topic for further investigation.

Finally, it is not lost on us that the results of Table 1 represent the payoffs of a two-person game between designer and attacker, where the defender chooses the number of jammers for which to plan and the attacker chooses the number of jammers to employ. Recommendations following many attacker-defender models ultimately require answering the question ''For how many attacks should we prepare?'' and the values in Table 1 provide a quantitative means to analyze this question. Using game theory to formally assess the tradeoffs between overpreparing and underpreparing in attacker-defender games is a topic for further exploration.

Performance Analysis

Since our method is intended to support hasty network design, we desire good solutions relatively quickly. Using our tool, we compare the performance of our algorithm to exhaustive enumeration. Our algorithm places $n-1$ APs (as we assume the location of a headquarters node is known a priori) and m jammers in a continuous space. We discretize this space by limiting feasible AP and jammer locations to the same grid used to define the set of coverage regions R within the operating area. Thus, the number of possible AP topologies is $(|R|_{n-1})$, and for each AP topology there are $(|R|_m)$ possible jammer topologies, yielding $(|R|_{n-1})(|R|_m)$ solutions to this discretized variant of the SRRA+C **DAD** problem. The exponential increase in the number of solutions as n, m, and $|R|$ grow restricts the use of this enumeration method to small problems, but we provide a few examples to demonstrate the superior performance of the DIRECT algorithm.

We cannot validly compare the **DAD** solutions found using the nested DIRECT algorithm and exhaustive enumeration by merely determining which produces a lower overall objective value. If this was our goal, we could simply set DIRECT to run with very few subproblem (i.e., attacker) iterations and thus limit it from finding particularly

Table 2. DIRECT jamming attacks on designs obtained using discrete enumeration.

			Objective runtime			Objective runtime	
N	m	Function evaluations	Value	(hr:min:sec)	Function evaluations	Value	(hr:min:sec)
2	1	4,950	642.73	0:00:49	123	**703.78**	0:00:0.7
3	1	495,000	472.32	1:40:56	101	**502.74**	0:00:01
2	2	495,000	5,268.58	0:41:37	167	**6,087.93**	0:00:01
3	2	24,502,500	933.43	53:54:27	103	**1,019.68**	0:00:01

Table 3. Enumerated jamming attacks on designs obtained from DIRECT.

			Objective runtime			Objective runtime	
N	m	Function evaluations	Value	(hr:min:sec)	Function evaluations	Value	(hr:min:sec)
2	1	11,249	**703.53**	0:00:41	100	701.52	0:00:01
3	1	5,591	**561.98**	0:00:50	100	539.94	0:00:01
2	2	6,447	**3,909.00**	0:00:27	4,950	2,003.40	0:00:27
3	2	33,963	**1,387.69**	0:09:22	4,950	945.59	0:01:02

good jamming attacks. However, our goal is to find those WMN designs that are most robust to jamming attacks. To demonstrate the performance of DIRECT in finding such designs, we use the algorithm to attack the best (i.e., most interference-robust) AP design found using enumeration.

Consider a small flat operating area discretized into $10 \times 10 = 100$ coverage regions. We first enumerate all possible discrete solutions for WMNs consisting of two APs and one jammer, three APs and one jammer, and two APs and two jammers. We then use DIRECT to attack the fixed AP topology of the best **DA**$\underline{\mathbf{D}}$ solution found. We present the results in Table 2. Bold values indicate cases where DIRECT obtains a more effective jamming attack than enumeration. DIRECT does this in each case considered, and does so in less than a second of processing time. In Table 3, we present results of the opposite approach: we use DIRECT to find **DA**$\underline{\mathbf{D}}$ solutions and then use enumeration to attack the best fixed AP topology. In no case does the enumeration method yield an attack more damaging than that found using DIRECT.

Finally, we analyze the performance of our algorithm using the 50-node network considered in Shankar (2008) and Xiao et al. (2004), discretized into the same $10 \times 10 = 100$ flat coverage regions (see Figure 7). We denote five nodes as destinations for traffic (indicated by large black circles), but unlike previous studies, we allow any node to serve as a source for network traffic. (This follows from our assumption that each AP will service client devices in the surrounding coverage regions.) Shankar (2008) uses enumeration to calculate the damage incurred by iteratively placing jammers at locations defined by a fixed grid. We compare attacks generated using his enumeration method and DIRECT, and present the results in Table 4. In both cases, DIRECT is able to find a more effective attack than the enumeration method, and does so in considerably less time.

Note that while these results demonstrate instances where DIRECT is more effective and efficient than discrete enumeration, this is not a fair comparison. DIRECT is a continuous algorithm and is guaranteed to eventually sample within an arbitrary distance of any point in the solution space, whereas discrete enumeration is limited to placing nodes at fixed, finite locations. Hence, as the number of iterations goes to infinity, DIRECT is guaranteed to eventually find a solution at least as good as discrete enumeration. Future research could compare the use of DIRECT to other algorithms, such as genetic or simulated annealing algorithms (see, e.g., Serafino et al. (2011)).

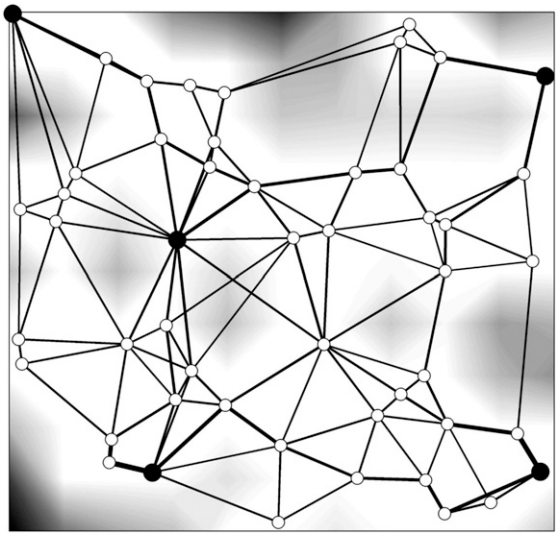

Figure 7. SRRA+C analysis of the 50-node network considered by Xiao et al. (2004) without jammers. The large black nodes denote traffic destinations. Client coverage shortfall is indicated by shaded areas. Line thickness is proportional to the traffic flow along each respective link.

Table 4. Comparison of enumerated and DIRECT attacks on a 50-node network.

	DAD solved using enumeration			DIRECT attack		
M	Function evaluations	Overall value	Runtime (hr:min:sec)	Function evaluations	Overall value	Runtime (hr:min:sec)
1	100	376.16	0:51:55	7	**421.96**	0:03:36
2	4,950	871.15	55:02:7	25	**963.45**	0:25:45

CONCLUSION

Our model of WMN performance is based on arguably the most fundamental factor in wireless communications: the transmission and reception of EM energy over terrain (Molisch, 2010). Using the game-theoretic **DAD** framework and based on our SRRA+C model, we develop a method for quickly designing WMN topologies that are robust to the effects of EMI.

The SRRA+C **DAD** formulation may be useful in modeling the interactions of other, similar systems where areas (whether physical or logical) need to be serviced by a fixed number of interconnected entities and need to be robust to worst-case disruption. For instance, the formulation could be applied to a logistics network or facility location problem (see, e.g., Church et al. (2004)), where warehouses (i.e., APs) need to distribute goods to customers in known locations (i.e., client coverage areas) in the presence of road construction or traffic jams (i.e., jammers). Another application area may be electrical distribution systems, where substations (i.e., APs) need to service client areas despite blown transformers, fallen trees, and intentional attacks (i.e., jammers).

Future research could incorporate the allocation of EM spectrum, the effects of RF phase, or the use of directional antennae (see, e.g., Ståhlberg, 2000). The modular nature of our formulation allows us to use essentially any WMN model, including high-fidelity simulations like OPNET (Riverbed Technology, 2014), but increased fidelity will incur increased runtimes and possibly less tractability.

Finally, the results here assume that the number of APs is fixed and decided up front. From Figures 4–6, we observe that the optimal design for four, five, or six APs places them in different locations. In practice, it might be important to design a topology that can be built incrementally (e.g., deploy four APs, then add a fifth, then add a sixth). Such a design could support a land force that is evolving or perhaps moving over terrain. As envisioned here, APs are not permanently fixed and could be repositioned as needs evolve. In general, it is not always possible to design a network where all intermediate topologies are optimal (e.g., Nehme and Morton (2010)); however, the formulation presented here could be extended to include a constraint that requires near-optimal nested intermediate designs. This and other extensions serve as potentially important topics for future study.

ACKNOWLEDGEMENTS

The authors thank several colleagues for constructive comments on this and earlier versions of this paper. Alderson was supported by the Office of Naval Research and the Defense Threat Reduction Agency. The results in this paper were drawn from a recent technical report (Nicholas and Alderson (2015)). Our invention is the subject of US patent award 9,788,213.

REFERENCES

Alderson, D. L., Brown, G. G., and Carlyle, W. M. 2014. Assessing and Improving Operational Resilience of Critical Infrastructures and Other Systems. *Tutorials in Operations Research.* INFORMS, 180–215.

Alderson, D. L., Brown, G. G., Carlyle, W. M., and Wood, R. K. 2011. Solving Defender-Attacker-Defender Models for Infrastructure Defense. *Proceedings Operations Research, Computing and Homeland Defense, 12th*

INFORMS Computing Society Conference, Wood, K. and Dell, R., eds., 28–49.

Audet, C., and Dennis Jr, J. E. 2006. Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM Journal on Optimization*, Vol 17, No 1, 188–217.

Bellman, R. E. 1961. *Adaptive Control Processes: A Guided Tour.* Volume 4. Princeton University Press.

Bertsekas, D. 1999. *Nonlinear Programming.* Athena Scientific, 185–186.

Brown, G., Carlyle, M., Salmeron, J., and Wood, K. 2006. Defending Critical Infrastructure, *Interfaces*, Vol 36, No 6, 530–544.

Caro, D. 2007. Users Fear Wireless Networks for Control, *InTech Magazine*, May 1, 2007. http://lists.jammed.com/ISN/2007/05/0122.html, retrieved May 7, 2013.

Church, R., Scaparra, M., and Middleton, R. 2004. Identifying Critical Infrastructure: The Median and Covering Facility Interdiction Problems, *Annals of the Association of American Geographers*, Vol 94, No 3, 491–502.

Cox, J. 2007. Xbox Accused of Jamming WLANs, *Techworld.* http://news.techworld.com/mobile-wireless/10941/xbox-accused-of-jamming-wlans.

Cruz, J. B. 1975. Survey of Nash and Stackelberg Equilibrium Strategies in Dynamic Games, *Annals of Economic and Social Measurement*, Vol 4, No 2, 339–344.

Fudenberg, D., and Tirole, J. 1991. *Game Theory.* MIT Press.

Google Maps. 2013. http://maps.google.com/, retrieved May 26, 2013.

Hastie, T., Tibshirani, R., and Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer Series in Statistics.

Horst, R., and Hoang, T. 1996. *Global Optimization.* Springer.

Institute for Engineering and Technology (IET). 2013. Jamming and Radio Interference: Understanding the Impact. *IET Sector Insights.* http://www.theiet.org/sectors/information-communications/, retrieved on May 28, 2013.

Jones, D. R., Perttunen, C. D., and Stuckman, B. E. 1993. Lipschitzian Optimization Without the Lipschitz Constant, *Journal of Optimization Theory and Applications*, Vol 79, No 1, 157–181.

Lazos, L., and Krunz, M. 2011. Selective Jamming/Dropping Insider Attacks in Wireless Mesh Networks, *IEEE Network*, Vol 25, No 1, 30–34.

Ma, K., Zhang, Y., and Trappe, W. 2005. Mobile Network Management and Robust Spatial Retreats via Network Dynamics, *Proceedings IEEE Conference Mobile Ad Hoc and Sensor Systems*, 235–242.

Molisch, A. 2011. *Wireless Communications.* Wiley, United Kingdom.

Mpitziopoulos, A., Gavalas, D., Konstantopoulos, C., and Pantziou, G. 2009. A Survey on Jamming Attacks and Countermeasures in WSNs, *IEEE Communications Surveys & Tutorials*, Vol 11, No 4, 42–56.

Myerson, R. 1991. *Game Theory: Analysis of Conflict.* Harvard University Press.

Nehme, M., and Morton, D. 2010. Efficient Nested Solutions of the Bipartite Network Interdiction Problem, *Proceedings of the IIE Annual Conference.*

Nicholas, P. J. 2009. Optimal Transmitter Placement in Wireless Mesh Networks. Master's thesis, Operations Research Department, Naval Postgraduate School.

Nicholas, P. J., and Alderson, D. 2012. Fast, Effective Transmitter Placement in Wireless Mesh Networks, *Military Operations Research*, Vol 17, No 4, 69–84.

Nicholas, P. J., and Alderson, D. L. 2015. Designing Interference-Robust Wireless Mesh Networks Using a Defender-Attacker-Defender Model. Technical report NPS-OR-15-002, Naval Postgraduate School.

Pelechrinis, K., Iliofotou, M., and Krishnamurthy, S. 2011. Denial of Service Attacks in Wireless Networks: The Case of Jammers, *IEEE Communications Surveys & Tutorials*, Vol 13, No 2, 245–257.

Poisel, R. 2011. *Modern Communications Jamming: Principles and Techniques.* Artech House.

Riverbed Technology. 2013. OPNET Modeler Suite. http://www.riverbed.com/products-solutions/products/network-planning-simulation/Network-Simulation.html, retrieved May 28, 2013.

di Serafino, D., Liuzzi, G., Piccialli, V., Riccio, F., and Toraldo, G. 2011. A Modified DIviding RECTangles Algorithm for a Problem in Astrophysics, *Journal of Optimization Theory and Applications*, Vol 151, No 1, 175–190.

Shankar, A. 2008. Optimal Jammer Placement to Interdict Wireless Network Services. Master's thesis, Operations Research Department, Naval Postgraduate School.

Shannon, C. 1949. Communication in the Presence of Noise, *Proceedings of the IRE*, Vol 37, No 1, 10–21.

Srivastava, V., Neel, J., Mackenzie, A., Menon, R., Dasilva, L. Hicks, J., Reed, J., and Gilles, R. 2005. Using Game Theory to Analyze Wireless Ad Hoc Networks, *IEEE Communications Surveys*, Vol 7, No 4, 46–56.

von Stackelberg, H. 1952. *The Theory of the Market Economy.* William Hodge.

Ståhlberg, M. 2000. Radio Jamming Attacks against Two Popular Mobile Networks. Helsinki University of Technology Seminar on Network Security.

Thamilarasu, G., and Sridhar, R. 2009. Game Theoretic Modeling of Jamming Attacks in Ad Hoc Networks, *Proceedings 18th International Conf. Computer Communications and Networks*, IEEE, 1–6.

The Economist. 2011. GPS Jamming: No Jam Tomorrow.

United States Geological Survey (USGS). 2013. National Elevation Dataset. http://ned.usgs.gov/, retrieved May 26, 2013.

Vakin, S., Shustov, L., and Dunwell, R. 2001. *Fundamentals of Electronic Warfare.* Norwood, MA: Artech House, 61.

Wood, A. D., Stankovic, J. A., and Son, S. H. 2003. JAM: A Jammed-Area Mapping Service for Sensor Networks, *Proceedings IEEE Real-time System Symposium*, 286–297.

Wood, A. D., Stankovic, J. A., and Zhou, G. 2007. DEEJAM: Defeating Energy-Efficient Jamming in IEEE 802.15.4-based Wireless Networks, *Proceedings 4th Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks*, 60–69.

Xiao, L., Johansson, M., and Boyd, S. 2004. Simultaneous Routing and Resource Allocation Via Dual Decomposition, *IEEE Transactions on Communications*, Vol 52, No 7, 1136–1144.

Xu, W. 2007. On Adjusting Power to Defend Wireless Networks from Jamming, *Proceedings MobiQuitous 2007, 4th Int. Conf. Mobile and Ubiquitous Systems*, 1–6.

Xu, W., Trappe, W., Zhang, Y., and Wood, T. 2005. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks, *Proceedings 6th ACM International Symposium Mobile Ad-hoc Networking and Computing*, 46–57.

Xu, W. Wood, T., Trappe, W., and Zhang, Y. 2004. Channel Surfing and Spatial Retreats: Defenses against Wireless Denial of Service, *Proceedings 3rd ACM Workshop on Wireless Security*, 80–89.

Zhang, Y., Zheng, J., and Hu, H., eds. 2008. *Security in Wireless Mesh Networks.* CRC Press.

APPENDIX: THE OPERATOR'S PROBLEM

Building on Xiao et al. (2004), the following is a formulation of the operator's problem as a nonlinear optimization problem. See Nicholas and Alderson (2015) for a complete derivation.

Index Use

$i \in N$	AP node (*alias j*)
$k \in M$	jammer node
$(i, j) \in A$	directed arc (*link*)
$d \in D$ & N	destination node

Input Data

$\widehat{\lambda}_i$	locations of AP nodes, $\hat{\lambda} = \{\widehat{\lambda}_i, i \in N\}$
$\widehat{\chi}_k$	locations of jammer nodes, $\hat{\chi} = \{\widehat{\chi}_k, k \in M\}$
p_i	maximum total transmission power per AP node, $i \in N$ [watts]
b	channel bandwidth [Hertz]

Calculated Data

$gain_{ij}$	product of antilog gain terms from $i \in N$ to $j \in$ N

$loss_{ij}$	product of antilog loss terms from $i \in N$ to $j \in$ N
$interference_j$	total received EMI and background noise power at $j \in N$ [watts]
$Z_{coverage}(\hat{\lambda}, \hat{\chi})$	total coverage shortfall of given AP locations $\hat{\lambda}$ and jammer locations $\hat{\chi}$

Decision Variables

P_{ij}	total transmission power along arc $(i, j) \in A$ [watts]
S_i^d	total traffic flow from origin $i \in N$ to destination $d \in D$ [bps]
T_{ij}	total traffic flow along arc $(i, j) \in A$ [bps]
X_{ij}^d	traffic flow along arc $(i, j) \in A$ to destination $d \in D$ [bps]

Formulation

$$\min_{P,S,T,X} \left(Z_{coverage}(\hat{\lambda}, \hat{\chi}) - w \sum_d \sum_{i \neq d} \log_2 \left(S_i^d \right) \right) \tag{5}$$

$$\sum_{j:(i,j)\in A} X_{ij}^d - \sum_{j:(j,i)\in A} X_{ji}^d = S_i^d \quad \forall i \in N, \forall d \in D, i \neq d \tag{6}$$

$$T_{ij} = \sum_d X_{ij}^d \quad \forall (i,j) \in A \tag{7}$$

$$T_{ij} - b \log_2 \left(1 + \frac{gain_{ij}}{interference_j loss_{ij}} P_{ij} \right) \leq 0 \quad \forall (i,j) \in A \tag{8}$$

$$\sum_{j:(i,j)\in A} P_{ij} \leq p_i \quad \forall i \in N \tag{9}$$

$$P_{ij} \geq 0 \quad \forall (i,j) \in A \tag{10}$$

$$S_i^d \geq 0 \quad i \neq d \tag{11}$$

$$T_{ij} \geq 0 \quad \forall (i,j) \in A \tag{12}$$

$$X_{ij}^d \geq 0 \quad \forall (i,j) \in A, \forall d \in D \tag{13}$$

The objective function (5) maximizes delivered network flow, where w is a constant indicating the relative importance of network flow, and coverage shortfall $Z_{coverage}(\hat{\lambda}, \hat{\chi})$ is calculated based on AP locations $\hat{\lambda}$ and jammer locations $\hat{\chi}$. Constraints (6) ensure balance of flow at each node. Constraints (7) define total flow along each arc as the sum of individual flows. Constraints (8) define arc capacity based on the Shannon limit (Shannon, 1949). Constraints (9) restrict total transmission power at each AP. Constraints (10–13) enforce nonnegativity. If we let F denote the tuple of decision variables (P, S, T, X), and use $F \in \mathcal{F}$ to mean explicitly that constraints (6–13) are satisfied, then this formulation is equivalent to (1).