Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2012

Fast, Effective Transmitter Placement in Wireless Mesh Networks

Nicholas, P.; Alderson, D.L.

Nicholas, P., Alderson, D.L., 2012, Fast, Effective Transmitter Placement in Wireless Mesh Networks, Military Operations Research, Vol. 17, No. 4, pp. 69-84.
<https://hdl.handle.net/10945/36932>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

ABSTRACT

Wireless mesh networks are systems of interconnected wireless access points that provide digital services to client devices via radio transmission. We consider the challenges of a communications planner who must quickly design a wireless mesh network, as might be expected during combat operations or in support of humanitarian assistance and disaster relief operations. We seek a network that maximizes client coverage area subject to constraints on network service, the technical characteristics of the available access points, and radio propagation over terrain. We create a nondifferentiable, nonconvex, nonlinear optimization problem and use a sampling algorithm to quickly find good solutions. We validate our formulation and solutions via numerical experiments and several field tests, and we demonstrate that our technique can generate network topologies capable of functioning in real-world scenarios.

We construct a corresponding decision support tool that allows a communications planner to design working wireless mesh network topologies quickly, with no guesswork, and requiring very little expertise. The tool runs on a laptop, supports virtually any type of access point, uses terrain information freely downloadable from the Internet, and does not require any additional software or solver licenses.

INTRODUCTION

Description of Problem

A wireless mesh network (WMN) is a communications network of fixed *access points* (APs) that exchange electronic messages via radio transmission to and from *client devices* (such as computers, sensors, or mobile devices). The fixed position of its APs differentiates a WMN from a so-called ad hoc network, where the APs can be constantly moving (Zhang et al. 2006, p. 565).

Military and civilian organizations can benefit from the inherent advantages of WMNs. During combat operations, WMNs can quickly and securely relay time-critical information such as intelligence reports, tactical orders, and sensor readings to separated small units. During humanitarian

assistance and disaster relief (HA/DR) operations, WMNs can provide maps, floorplans, video surveillance, emergency aid requests, and other critical information to first responders.

In the system under consideration, each AP uses two separate and configurable radio devices to create a WMN with two levels of connectivity. The first level of connectivity supports client-to-AP communication within a *client coverage area*, whereas the second level consists of a *backhaul radio network* that routes traffic between APs. Client devices may roam within the coverage area and communicate with one another or to an outside network, such as the Internet, through a gateway (e.g., via a satellite uplink). We assume that wireless APs alone provide traffic routing services, thus forming an *infrastructure mesh* type of WMN (see Nicholas 2009, pp. 2–4). Additionally, we assume client devices communicate with a single AP at a time, and hence do not serve as intermediate relay points (Zhang et al. 2006, pp. 564–567).

The physical topology of a WMN, as defined by the locations of the wireless APs, is critical to its performance. We consider the challenge of a communications officer or a *network designer* who must choose the locations and configurations of wireless APs to provide service to clients in desired coverage areas while meeting restrictions on the quantity, placement, and characteristics of the APs and also satisfying any requirements for coverage, bandwidth, and other service standards. The designer must also consider the effects of terrain and other aspects of the operating environment on radio wave propagation. Because combat and HA/DR operations are time-sensitive, the designer must build the WMN quickly and with as little guesswork as possible.

We focus on designing a WMN topology that maximizes client coverage while considering network flow requirements. There is a fundamental tension when designing such a WMN: maximizing client coverage tends to place APs far apart, whereas the requirements for backhaul network throughput tend to keep APs relatively close. The challenge is to balance this trade-off for a fixed number of APs while respecting their technological details and the physics of radio propagation over terrain.

We evaluate the performance of a given WMN topology in two steps. First, we

Fast, Effective Transmitter Placement in Wireless Mesh Networks

Paul J. Nicholas

Marine Corps Combat
Development Center
paul.nicholas@usmc.mil

Dr. David L. Alderson

Naval Postgraduate School
dlalders@nps.edu

APPLICATION AREAS:
Battle Management/
Command and Control
OR METHOD: Multi-
objective Optimization,
Network Methods,
Nonlinear
Programming,
Simulation

calculate the *value of client coverage* as dictated by specific AP locations and their configuration, ground terrain, and environmental data. We then calculate the *value of network flow* according to the Simultaneous Routing and Resource Allocation (SRRA) techniques of Xiao et al. (2004). We combine these two subproblems to formulate SRRA+C, a nondifferentiable, nonconvex, nonlinear optimization problem. We implement the Dividing RECTangles (DIRECT) sampling algorithm of Jones et al. (1993) to quickly find good solutions to SRRA+C.

To validate the quality of the network topologies generated by our technique, we conduct numerical experiments and field tests using commercial equipment. As an aid to the network designer, we create a customized graphical decision support tool to solve the SRRA+C problem. The standalone tool reads and graphically displays digital terrain elevation information, obtains its best solution to the SRRA+C problem, and displays the resulting network and client coverage regions. We demonstrate that our technique can quickly create WMN topologies that function in realistic scenarios, and that our decision support tool can assist communications officers or network designers in building WMNs in support of combat or HA/DR operations.

Previous Work

He et al. (2004) use the DIRECT sampling algorithm of Jones et al. (1993) to find good AP placements in indoor wireless networks, using very accurate (and computationally expensive) ray-tracing techniques that predict radio propagation. We build on this general idea to calculate the value of client coverage, and use the DIRECT algorithm similarly for outdoor environments.

Xiao et al. (2004) solve the SRRA problem via dual decomposition to identify optimal traffic routes and allocation of AP transmission power within a wireless network. As noted above, we use SRRA to calculate the value of network flow.

Shankar (2008) uses the SRRA formulation to determine network flow among prepositioned wireless nodes, and then adopts the attacker-defender techniques of Brown et al.

(2006) to calculate optimal jammer locations that maximally disrupt network flow. Our work complements Shankar's research, providing the initial design of a WMN.

To our knowledge, ours is the first technique for designing WMN topologies that maximize client coverage while considering backhaul network service requirements on real terrain. See Nicholas (2009, pp. 6–9) for a detailed literature review.

This paper is organized as follows. In the next section, we describe each element of the SRRA+C formulation along with our technique to solve it. We then consider several notional and real network design problems, contrast our solution to what we obtain from brute force enumeration, and briefly summarize the results of our field tests. We conclude with suggestions for follow-on research.

MODEL FORMULATION

Our goal is to position APs in locations that maximize client coverage, subject to restrictions on network service, AP characteristics and placement, and radio propagation over terrain. As noted earlier, there is an inherent tension between maximizing client coverage and network traffic flow. The crux of the SRRA+C formulation is to capture and quantify this tension.

We represent each AP as a *node* in a network. Let N denote the set of all AP nodes, indexed by $i = 1, 2, \dots, n$, where $n = |N|$. Let $\lambda = (\lambda_1, \lambda_2 = \lambda_n)$ represent the locations of the nodes. Although the mathematics support any coordinate system to represent location, in our implementation each node location λ_i is itself a two-dimensional coordinate representing the northing and easting for node i . We assume that node locations are *fixed* in the sense that nodes will remain in position once placed. In what follows, we use the terms AP and node interchangeably.

We define the *operating region* as the topographic area where an AP may be physically located. We partition the operating region into a set of discrete *regions* R , indexed by $r = 1, 2, \dots, |R|$. Although our formulation allows the use of any discretization scheme,

our implementation assumes rectangular regions arranged in a grid (see Figure 1).

Each AP has two radio devices. The first provides client coverage in the immediate vicinity of the AP; the specific coverage obtained (e.g., shaded area in Figure 1) depends on several factors including the local terrain. The second radio creates a backhaul network between nodes (e.g., dashed lines in Figure 1). Let $A \subseteq N \times N$ denote the set of directed backhaul arcs between nodes, with flow along arc $(i, j) \in A$ representing the directed transmission from node i to node j . In general, the wireless nature of this network means that transmission along any of the backhaul arcs is possible, although distance, terrain, and background interference will dictate which of these arcs has the transmission capacity to be useful. Following Xiao et al. (2004), we assume APs are not subject to self-jamming or interference from other APs.

We identify a single node $d \in N$ as the *headquarters (HQ) node* (see Figure 1), and assume

the network designer predefines the location of this node. (In what follows, we will assume that node d is always the first node in the set N .) We assume the vast majority of client network traffic will be directed to or received from the HQ node, as this location will connect to the Internet or other outside network, as well as host email, domain, and storage servers. Hence, we optimize our network for traffic flow from client service areas to the HQ node, though our formulation is more general and allows any number of APs to serve as the destination for network traffic.

Calculating Client Coverage

The client coverage provided by a particular WMN topology is a function of its AP locations. We calculate the *received signal strength* (RSS) ρ_{ir} (measured in dBm) from the transmitter at node i to the center of region r using the standard link budget formula (Olexa 2005, p. 79):

Figure 1. Discretized operating region and wireless mesh network (WMN). Circles denote the location of access points (APs), shaded regions denote the areas with sufficient client coverage (i.e., zero coverage shortfall), and dashed lines denote the backhaul network. In this example, there is a single headquarters (HQ) node (labeled d) and four APs (labeled 2, 3, 4, 5). The locations of the APs are represented by the vector $\lambda = (\lambda_d, \lambda_2, \lambda_3, \lambda_4, \lambda_5)$.

$$\rho_{ir} = power_i + gain_i - loss_i - loss_{path} - loss_{misc} + gain_r - loss_r \quad (1)$$

where $power_i$ is transmitted power (in dBm) from node i ; $gain_i$ and $gain_r$ are respectively the gains (in dBi) of the transmitter at node i and receiver in region r ; $loss_i$ and $loss_r$ are respectively the losses (in dB) from cables, connectors, etc., of the transmitter and receiver; $loss_{path}$ is total path loss (in dB) from λ_i to the center of region r ; and $loss_{misc}$ is miscellaneous loss (in dB), such as fade margin. All of the terms in Equation (1) are input data, determined by the equipment and environment, except for the total path loss $loss_{path}$, which depends on the position of the transmitter and receiver in question.

Our formulation allows any model for computing $loss_{path}$, including the Irregular Terrain Model (ITM) (Longley and Rice 1968) and Hata-COST 231 (COST 1999). Our preferred model is the Terrain Integrated Rough Earth Model (TIREM) of Alion Science & Technology Corporation (<http://www.alionscience.com>). This model samples terrain elevation to compute path loss, and considers the effects of free space loss, diffraction around obstacles, atmospheric absorption and reflection, and other factors.

We adopt and modify He et al.'s (2004) concept of power coverage to quantify the value of client coverage. We specify for each region $r \in R$ a minimum *coverage threshold* τ_r in dBm; any received signal ρ_{ir} above this threshold from AP i to region r qualifies as adequate client coverage. The difference between ρ_{ir} and τ_r represents a quantity we define as *coverage shortfall* at region r from node i . Summarizing, for each location $r \in R$ and each node $i \in N$, we have

$$(\text{Coverage Shortfall})_{ir} \equiv (\tau_r - \rho_{ir})_+$$

where $()_+$ denotes the projection onto the non-negative real line. Because a positive difference represents inadequate client coverage, we wish to minimize this quantity. We need consider only the minimum coverage shortfall from each node i , as we assume each client device can connect to only one AP. We sum over all $r \in R$ to calculate *total coverage shortfall*, denoted $Z_{coverage}$:

$$Z_{coverage}(\lambda) = (\text{Total Coverage Shortfall}) \equiv \sum_{r \in R} \min_{i \in N} (\tau_r - \rho_{ir})_+ \quad (2)$$

The total coverage shortfall is a function of node locations λ . By allowing only positive terms, we disallow the benefit of transmitting excessive power to any given coverage location. This method also limits the amount any coverage location can be penalized: no more than τ_r . Note it is possible to increase the relative importance of any region r by multiplying the region's coverage shortfall by a positive scalar.

Calculating Network Flow

Given fixed AP locations $\lambda_i, i \in N$, we use the Shannon capacity formula (1949) to calculate the transmission capacity along each arc $(i, j) \in A$. This formula establishes a theoretical upper bound on transmission capacity in bits per second (bps). Following Xiao et al. (2004), the capacity from node i to node j is:

$$(\text{Capacity})_{ij} = \text{bandwidth} \log_2 \left(1 + \frac{\text{gain}_{ij}}{\text{noise}_{ij} \text{loss}_{ij}} P_{ij} \right) \quad \forall (i, j) \in A \quad (3)$$

where *bandwidth* is channel bandwidth in Hertz; gain_{ij} is the sum of the antilog gain terms (gain_i and gain_j from Equation 1); noise_{ij} is the background noise power in watts from node i to node j ; and loss_{ij} is the sum of the antilog loss terms (loss_i , loss_j , loss_{path} , and loss_{misc} from Equation 1). Again, these input data depend on the placement of APs, which we assume are at known, fixed locations. In theory, we assume that all backhaul arcs are possible (i.e., $A = N \times N$).

We assume each AP has limited total transmission power denoted p_i (in watts), and we define P_{ij} to be the amount of p_i used to transmit from i to j . Thus, each AP is additionally constrained by

$$\sum_{j: (i,j) \in A} P_{ij} \leq p_i \quad (4)$$

Here, P_{ij} is a decision variable representing the AP-to-AP transmission power from node i to node j , whereas the transmission power for AP-to-client communication $power_i$ is a (constant) input parameter.

We measure each individual traffic flow in bps. We adopt the approach of Xiao et al. (2004) to quantify the value of total network flow according to a log-utility function that places a zero value on unit flow, positive values on flows greater than one, and negative values on flows less than one. We note that a zero flow has an infinite penalty and therefore there is strong incentive to ensure that each origin-destination pair receives some flow. Defining S_i^d to be the total flow originating at node i and destined for node d , we have

$$(\text{Utility of Total Network Flow}) \equiv \sum_d \sum_{i \neq d} \log_2(S_i^d). \quad (5)$$

Collectively, we obtain our version of the Xiao et al. (2004) SRRA problem for given AP locations λ as shown in Figure 2.

Given AP locations λ , this is a multicommodity network flow problem. The objective function (S0) maximizes the total utility of traffic flow between each origin-destination pair. Constraints (S1) ensure balance of flow at each node. Constraints (S2) define the total flow along any arc as the sum of all traffic flows along that arc. Constraints (S3) ensure that total flow along any arc is less than or equal to its capacity. Constraints (S4) enforce a budget on the total transmission power at each AP. Constraints (S5)-(S8) ensure nonnegativity.

Xiao et al. (2004) observe that the SRRA problem can be solved via dual decomposition because of its layered structure. Specifically, by introducing the Lagrange multipliers $\alpha_{ij} \forall (i, j) \in A$ for the capacity constraint (S3), we obtain the partial Lagrangian

<u>Index Use</u>			
$i \in N$	node (<i>alias j</i>)		
$(i, j) \in A$	directed arc (<i>link</i>)		
$d \in D \subseteq N$	destination nodes		
<u>Input Data</u>			
λ_i	the locations of nodes, $\lambda = \{\lambda_i, i \in N\}$	[none]	
p_i	maximum total transmission power per node, $i \in N$	[watts]	
<i>bandwidth</i>	channel bandwidth	[hertz]	
<u>Calculated Data</u>			
<i>gain_{ij}</i>	sum of antilog gain terms from $i \in N$ to $j \in N$	[none]	
<i>loss_{ij}</i>	sum of antilog loss terms from $i \in N$ to $j \in N$	[none]	
<i>noise_{ij}</i>	background noise power from $i \in N$ to $j \in N$	[watts]	
<u>Decision Variables</u>			
S_i^d	total flow of traffic from origin $i \in N$ to destination $d \in D$	[bps]	
F_{ij}^d	traffic flow along arc $(i, j) \in A$ to destination $d \in D$	[bps]	
T_{ij}	total flow along arc $(i, j) \in A$	[bps]	
P_{ij}	total transmission power along arc $(i, j) \in A$	[watts]	
<u>Formulation</u>			
$Z_{\text{flow}}(\lambda) = \max_{S, F, T, P} \sum_d \sum_{i \neq d} \log_2(S_i^d)$			(S0)
s.t. $\sum_{i: (j, i) \in A} F_{ji}^d - \sum_{i: (i, j) \in A} F_{ij}^d = S_j^d$			$\forall j \in N, \forall d \in D$ (S1)
$T_{ij} = \sum_d F_{ij}^d$			$\forall (i, j) \in A$ (S2)
$T_{ij} - \text{bandwidth} \log_2 \left(1 + \frac{\text{gain}_{ij}}{\text{noise}_{ij}} P_{ij} \right) \leq 0$			$\forall (i, j) \in A$ (S3)
$\sum_{j: (i, j) \in A} P_{ij} \leq p_i$			$\forall i \in N$ (S4)
$S_i^d \geq 0$			$i \neq d$ (S5)
$F_{ij}^d \geq 0$			$\forall (i, j) \in A, \forall d \in D$ (S6)
$T_{ij} \geq 0$			$\forall (i, j) \in A$ (S7)
$P_{ij} \geq 0$			$\forall (i, j) \in A$ (S8)

Figure 2. Formulation SRRA.

$$L(S, F, T, P, \alpha) = \sum_d \sum_{i \neq d} \log_2(S_i^d) - \sum_{(i,j) \in A} \alpha_{ij} \left(T_{ij} - \text{bandwidth} \log_2 \left(1 + \frac{\text{gain}_{ij}}{\text{noise}_{ij} \text{loss}_{ij}} P_{ij} \right) \right). \quad (6)$$

The objective function of the dual problem is:

$$V(\alpha) = \max_{S, F, T, P} L(S, F, T, P, \alpha). \quad (7)$$

One immediate observation is that this dual function can be evaluated separately in the network flow variables S , F , and T and the communications variable P . Thus, the problem of evaluating the dual function separates into a *network flow (net) subproblem* and a *resource allocation (comm) subproblem*, that is:

$$V(\alpha) = V_{\text{net}}(\alpha) + V_{\text{comm}}(\alpha) \quad (8)$$

where

$$\begin{aligned} V_{\text{net}}(\alpha) &= \max_{S, F, T} \sum_d \sum_{i \neq d} \log_2(S_i^d) - \sum_{(i,j) \in A} \alpha_{ij} T_{ij} \\ \text{s.t.} \quad &\sum_{i:(j,i) \in A} F_{ji}^d - \sum_{i:(i,j) \in A} F_{ij}^d = S_j^d \quad \forall j \in N, \forall d \in D \\ &T_{ij} = \sum_d F_{ij}^d \quad \forall (i,j) \in A \\ &S_i^d \geq 0 \quad i \neq d \\ &F_{ij}^d \geq 0 \quad \forall (i,j) \in A, \forall d \in D \\ &T_{ij} \geq 0 \quad \forall (i,j) \in A \end{aligned}$$

$$\begin{aligned} V_{\text{comm}}(\alpha) &= \max_P \sum_{(i,j) \in A} \alpha_{ij} \text{bandwidth} \log_2 \left(1 + \frac{\text{gain}_{ij}}{\text{noise}_{ij} \text{loss}_{ij}} P_{ij} \right) \\ \text{s.t.} \quad &\sum_{j:(i,j) \in A} P_{ij} \leq p_i \quad \forall i \in N \\ &P_{ij} \geq 0 \quad \forall (i,j) \in A. \end{aligned}$$

The Lagrange dual problem is thus:

$$\begin{aligned} \min_{\alpha} V(\alpha) &= V_{\text{net}}(\alpha) + V_{\text{comm}}(\alpha) \\ \text{s.t.} \quad &\alpha \geq 0. \end{aligned} \quad (9)$$

The dual function $V(\alpha)$ is always convex. Xiao et al. (2004) assume that a feasible solution (S, F, T, P) exists such that the nonlinear capacity constraints hold with strict inequality (known as *Slater's condition*, see Boyd and Vandenberghe 2004, Sec. 5.2), and therefore conclude strong duality holds and the optimal value of this dual problem is equal to the optimal value of the primal problem. However, they also note that the objective function of the primal problem is not strictly concave in the variables F and T , and thus the dual function is only piecewise differentiable. As a result, the dual problem is a nondifferentiable convex optimization problem, to which they apply the subgradient method to obtain a solution.

We also apply the subgradient method to solve this problem, running it only a fixed number of iterations (typically 500) to obtain an approximate solution. Each iteration of the subgradient method might not necessarily improve the dual objective value, but each iteration reduces the optimality gap (Bertsekas 1999, p. 621).

As noted by Xiao et al. (2004), both subproblems are convex optimization problems with special structure lending themselves to very efficient computational techniques. We solve the network flow subproblem as a multicommodity network flow problem, and use an algorithm described by Luss and Gupta (1975) to solve the resource allocation subproblem. Further details of the subgradient algorithm and its implementation to solve this problem are available in Nicholas (2009, pp. 35–45). We use the approximate solution to the SRRA problem as a quantification of the value of network flow.

SRRA + C Formulation

The overall SRRA+C problem minimizes $Z(\lambda)$ by choice of AP locations $\lambda_i, \forall i \in N$. We calculate the overall performance of a WMN as:

$$Z(\lambda) \equiv Z_{\text{coverage}}(\lambda) - wZ_{\text{flow}}(\lambda). \quad (10)$$

That is, the overall objective function is a linear combination of client coverage (calculated as client coverage shortfall) and network flow. The objective function is in units of dBm, although the objective value has no direct practical interpretation. Rather, the objective value serves as a relative method of comparing different network topologies.

We use w as a positive scalar (in units of dBm/log₂ bps) representing the relative weight placed on network flow. Larger values of w increase the value of network flow, and in general increase the appeal of more compact network topologies. We typically use $w = 1$, meaning that we weight these terms equally. See Nicholas (2009, pp.78–80) for a detailed sensitivity analysis of w .

Solving the SRRA + C Problem

One method of solving the SRRA + C problem is simply to restrict APs to a discrete set of fixed locations (e.g., the center of each location $r \in R$) and then try all possible network solutions via complete enumeration. This requires that we calculate the overall objective value for each unique topology and keep track of the best one(s). Recall that there are $|R|$ operating regions and n APs, with the position of the first AP (i.e., the HQ node) already fixed. Hence, the total number of unique topologies is $\binom{|R|}{n-1}$, and the solution space grows exponentially in both $|R|$ and n . Clearly, we need a more efficient

method of solving SRRA + C if it is to be used for fast design of WMNs. We show how to use the Diving RECTangles (DIRECT) algorithm of Jones et al. (1993) to meet this need.

Let $\lambda^k = (\lambda_1^k, \lambda_2^k, \dots, \lambda_n^k)$ denote the positions of the APs in iteration k of the algorithm. We follow the algorithm in Figure 3 to iteratively select λ^k for increasing k until we reach a specified stopping criterion (number of iterations).

This iterative scheme is a modified version of DIRECT, which is a sampling optimization algorithm based on Lipschitzian optimization (Horst and Hoang 1996, pp. 43–46) that iteratively divides the multidimensional solution space of AP locations into smaller hyper-rectangles (hence the name Diving RECTangles). Unlike Lipschitzian optimization, DIRECT does not require a priori specification of the Lipschitz constant, nor knowledge of the objective function gradient, which makes it appealing in solving this non-differentiable, non-convex, nonlinear problem.

In this implementation, DIRECT samples candidate solutions within a unit hypercube defined by the multidimensional SRRA + C solution space. The dimensionality of this space

Input: Full SRRA problem data (number and operating characteristics of APs, HQ node location, elevation and coverage requirements for each $r \in R$) and desired number of iterations $max_iterations$.

Output: Best estimate of optimal AP locations $\lambda^* = (\lambda_1^*, \lambda_2^*, \dots, \lambda_n^*)$ and overall objective value, $Z(\lambda^*)$.

```

begin
  Store map data
  Initialize starting AP locations  $\lambda^1$  at center of unit hypercube; calculate  $Z(\lambda^1)$ 
  Initialize  $k \leftarrow 2$ 
  while ( $k < max\_iterations$ ) do
    Select candidate AP locations  $\lambda^k$  based on previous objective values  $Z(\lambda^1), \dots, Z(\lambda^{k-1})$ 
    and hyper-rectangles for previous solutions  $\lambda^1, \dots, \lambda^{k-1}$ 
    Calculate coverage shortfall value  $Z_{coverage}(\lambda^k)$  according to Equation (2)
    Solve formulation SRRA and calculate network flow values  $Z_{flow}(\lambda^k)$ 
    Calculate overall objective value  $Z(\lambda^k)$  according to Equation (10)
    if ( $Z(\lambda^k) < Z(\lambda^*)$ )
       $\lambda^* \leftarrow \lambda^k$ 
       $Z(\lambda^*) \leftarrow Z(\lambda^k)$ 
    endif;
     $k \leftarrow k + 1$ 
  end;
  Return best AP locations  $\lambda^*$  and overall objective value  $Z(\lambda^*)$ 
end;

```

Figure 3. Algorithm DIRECT for SRRA + C.

is based on the number of APs n and the number of variables necessary to define the location of each AP (in our case, a pair of variables for northing and easting). The vector λ^k representing the locations of all APs at iteration k is a single point in this multidimensional solution space.

The initial solution λ^1 is chosen as the exact center of the solution space. During each iteration, DIRECT divides the solution space into smaller hyper-rectangles, based on the unexplored territory in the solution space and previously calculated objective function values. From any incumbent solution, the algorithm samples other locations in the solution space along each dimension, computing the corresponding objective function value for each. The largest hyper-rectangles are placed around those sampled points with more desirable objective values because they have more unexplored territory and hence greater potential for improvement. During each iteration, the next candidate solution λ^k depends on the previously stored solution values and the size of their associated hyper-rectangles. See Nicholas (2009, pp. 49–63) for more details on our implementation of DIRECT.

The DIRECT algorithm is guaranteed to converge eventually to a global optimum if the objective function is continuous (Jones et al. 1993). To meet this requirement, we use bilinear interpolation on the discrete elevation points within our map data to smooth the terrain surface.

One drawback of the DIRECT algorithm is that the optimality gap at any iteration is not known. Given the Lipschitz constant, we can calculate a lower bound, but in our application finding this constant would be as difficult as solving the problem to optimality. Another drawback is that the rate of convergence cannot be calculated exactly. That is, while DIRECT is guaranteed to converge, we do not know a priori how long this will take. In general, this is a serious concern, and it could mean that our algorithm runs too slowly to be of practical help. Fortunately, we observe in practice that it quickly yields solutions that are very good when compared to those obtained by simple techniques like exhaustive enumeration over a finite grid of candidate solutions. Though there

is no guarantee our technique will always provide good solutions, it is a significant advance over the kind of trial and error currently used by network designers in the field.

As a crude means of estimating an optimality gap for a given solution, we can calculate a lower bound to the SRRA+C objective value by considering an idealized network with zero coverage shortfall and no propagation loss between APs. Such an idealized network does not reflect the main tension in our model: the need to spread APs to get good coverage versus the need to keep APs close together to get good backhaul capacity. The laws of physics and technical limitations of the APs prevent us from being able to build this idealized network in practice, so this lower bound is weak, but the value provides a point of reference against which to gauge the progress of the DIRECT algorithm.

SRRA + C Decision Support Tool

We implement the algorithm DIRECT for SRRA+C in a decision support tool built using Microsoft Visual C++. The standalone program displays terrain data, which can be of any scale and any grid-based format, such as Universal Transverse Mercator (UTM) (Defense Mapping Agency 1989). The user can input all required variables, including drawing the desired coverage region(s) directly on a graphical display of the terrain. The program supports three different modes of calculating path loss (inverse-square, Hata COST-231, and TIREM). After solving the problem, the final AP locations, connections between APs, and coverage areas are depicted atop the map. Note our implementation does not require any third-party solvers or additional software.

ANALYSIS AND RESULTS

As case studies for our analysis, we consider two different operating regions in Fort Ord, California. These regions are respectively 45 acres and 145 acres, and the terrain consists of gently rolling hills, grass fields, pavement, and some buildings. We use terrain information from the United States Geological Survey (USGS) via MapMart (<http://www.mapmart.com>).

Case Study 1: 3-node network on 45 acres in Fort Ord, California

We begin with a 45-acre region (see Figure 4), which we divide into a grid of size $|R| = 65 \times 33 = 2,145$ rectangles. We place the HQ AP directly in the middle of the operating region and solve SRRA+C for a network of three APs using complete enumeration. Enumeration guarantees that we will find the optimal solution to this restriction of the original problem (in discrete space), but the method is extremely resource-intensive as the number of unique solutions grows exponentially in both locations $|R|$ and number of APs n . Recall the HQ node location is determined in advance, so building a network of n nodes means that we need to locate $n-1$ APs. We enumerate all $\binom{2145}{3-1} = 2,299,400$ unique solutions in eight hours, 38 minutes, and 55 seconds. Figure 5 ranks the solutions by their objective value.

Figure 5 shows that a small number of solutions are significantly better than the others. Specifically, about 0.29% are within 50% of the best solution obtained, 1,834.8. Again, this solution value has no absolute physical interpretation and serves only to allow relative comparison. However, this shows that the best solutions are rare among the possible solutions.

We then run DIRECT on this problem (see Figure 6) and stop it at each iteration 1, 2, ..., 10

to record the current overall objective value. These values are compared to our lower bound for this network (dashed line) and the best value found by enumeration (solid line).

In only 2.05 seconds, the DIRECT algorithm obtains a *better* solution than was found by enumeration in 8 hours and 39 minutes. This is possible because enumeration considers only discrete locations within the operating region, while DIRECT places APs continuously at any location in the region. Because it is continuous, DIRECT is *guaranteed* to find a solution at least as good as enumeration as the number of iterations approaches infinity. In this limit, DIRECT will almost certainly find a better one, because it considers an increasingly dense subset of solutions and will eventually sample to within an arbitrary distance of any point (to include the global optimum) within the solution space.

In practice, networks with more APs have more dimensions and a corresponding greater number of sub-hyper-rectangles. Hence for larger networks the DIRECT algorithm must be run for more iterations to find good solutions. There is no clear way to avoid this when comparing networks of varying numbers of APs. However, by running it until we reach the computational limit of our implementation of the algorithm, we provide the best answer possible. In this case, running DIRECT for up to 15 iterations does not result in a better solution.

Figure 4. Elevation profile of 45-acre Ford Ord region.

Figure 5. Rank-ordered solutions for a 3-node network in 45-acre Fort Ord region obtained using enumeration.

Case Study 2: 5-node network on 145 acres in Fort Ord, CA

We next examine a network of five APs over a 145-acre section of Fort Ord. This network is too large to consider using enumeration, so we solve only using DIRECT. We again stop the algorithm at each iteration 1, 2, ..., 30 to record the current overall objective value. Figure 7 presents the results.

The best solution value is computed by DIRECT after 30 iterations (796.02) and is within an order of magnitude of the lower bound (71.413). Although we cannot certify optimality, this boundary places a limit on how far we are from the global optimum. Here, we are willing to exchange optimality for speed: we obtain this solution after 22 iterations of DIRECT in slightly more than 17 seconds; to solve this problem using our enumeration technique would take

Figure 6. Comparison of solutions obtained by enumeration and DIRECT on actual terrain.

Figure 7. Example comparing DIRECT solutions to a lower bound.

more than 10^4 years. Of course, making this tradeoff makes sense only if the solutions we obtain are useful.

SRRA + C Validation: Does It Work in Practice?

The validity of solutions to the SRRA+C formulation, specifically our evaluation of network performance, rests on the accuracy of radio wave propagation and network throughput predictions. To validate our use of TIREM to predict received signal strength over real terrain, we conducted several field experiments in collaboration with the Hastily-Formed Networks (HFN) Research Group at the Naval Postgraduate School.

First, we conducted two point-to-point network tests in Ford Ord over generally flat terrain consisting of pavement and packed gravel, with no trees or other obstructions to the line of sight (LOS) path. We deployed Cisco AP1000-series Aironet WMN access points (Cisco 2009) broadcasting at the 5.8 GHz and 2.4 GHz operating frequencies (Figure 8 left and right, respectively) and used a laptop computer with an internal Intel wireless transceiver to measure the received signal strength. We compared these observations with TIREM predictions at ranges from 0 to 465 meters.

We observe that the measured values are reasonably close to the predicted ones, with the exception of measurements at 2.4 GHz from 300 meters and beyond. This is likely due to the

presence of nearby buildings at that end of our testing range: the radio waves may have reflected off the buildings and provided a stronger signal than would have otherwise been received. This suggests a potential weakness in the use of TIREM for urban environments. Overall, however, these two tests demonstrate that TIREM is capable of making very reasonable received signal strength predictions using our testing equipment in a real-world environment. See Nicholas (2009, pp. 69–75) for detailed testing results. In particular, these results are relatively accurate over ranges (i.e., up to 250 m) consistent with having five APs located uniformly across a 145-acre area.

However, the most important form of design validation is whether you can actually build and operate the system of interest. To determine if the topologies obtained from SRRA+C actually function in realistic environments, we conducted another network field test in the 45-acre rectangular region in Fort Ord, mentioned earlier. Elevation here ranges from 143 to 226 feet above sea level. The left side of Figure 9 is a Google Maps image of the operating region.

We used the SRRA+C Decision Support Tool to generate network topologies of three, four, and five APs. We deployed the same Cisco AP1000-series Aironet WMN APs, which use the 802.11b/g protocol (2.4 GHz) to provide client coverage and the 802.11a protocol (5.8 GHz) to provide the backhaul network. We positioned each AP using a Global Positioning System (GPS) device, and placed atop a two meter mast. The right side of Figure 9 depicts our setup.

FAST, EFFECTIVE TRANSMITTER PLACEMENT IN WIRELESS MESH NETWORKS

Figure 8. Validation of TIREM as a model for received signal strength. At 5.8 GHz (left) and 2.4 GHz (right), we observe that our implementation of TIREM reasonably predicts actual received signal strength in point-to-point field tests.

When required, we made small adjustments (within five meters) to AP position to avoid transmitting directly into trees, etc. We used small portable generators and car batteries as power sources.

Figure 10 displays the topologies prescribed by the SRRA+C Decision Support Tool for the 4-node and 5-node networks. After building each network, we used a laptop to measure the actual throughput from each AP to the HQ node (d). The results appear alongside each node in Figure 10.

The network topologies in Figure 10 worked similarly to our theoretical predictions. The backhaul network traffic from each AP to the HQ node followed exactly the routes predicted by SRRA+C. The actual throughput from each

AP to the HQ node was within an order of magnitude of our predictions, but typically exceeded them. While the throughput values predicted by SRRA+C were based on the assumption that all APs were transmitting concurrently, in our experiments we were limited to measuring the throughput from each AP individually. Because a lone transmitting node does not have to compete for backhaul network resources, it is natural that observed throughput should exceed prediction. Given our simplifying assumptions about the way the APs work (e.g., ignoring the signal processing technologies that try to compensate for environmental conditions), the level of correspondence between the predicted and observed values exceeded our expectations.

Figure 9. Field testing in Fort Ord, California. Left: Aerial view of 45-acre Fort Ord operating region [image courtesy Google Maps (<http://maps.google.com>)]. Right: Wireless mesh access point, mast, and portable generator

Figure 10. Network topologies and their measured performance during field tests. Circles denote position of each AP during field tests of 4-node (left) and 5-node (right) networks. For each node, we report the actual (and predicted) throughput from that node to the HQ node, denoted d . Arrows represent the direction of actual point-to-point wireless transmissions along the backhaul network, which exactly matched the predicted flows. For the 5-node network, SRRA+C predicted that node 5 should split its transmission in two directions; we observed that this node alternatively sent flow in two directions.

Not all of our field tests were successful. In the three-AP scenario (not shown), a node was unable to connect with the HQ node and the throughputs were lower than expected. We attribute this to the presence of intervening foliage and man-made obstructions, which are not considered by the TIREM model. Nonetheless, our field testing results show that SRRA+C solved with DIRECT can quickly provide working network designs with no guesswork. See Nicholas (2009, pp. 92–102) for a detailed description of these tests and results.

How Many APs Do We Need?

The SRRA+C problem allows us to determine where to place a known number of APs within a given operating region. But an equally important question is: *How many APs do we need?*

To address this question, we again consider the 45-acre section of Fort Ord and we use DIRECT to solve for networks of 2, 3, ..., 8 APs using weight $w = 1$, running the algorithm until we reach the computational limit of 32 hyper-rectangle divisions (the maximum possible using double-precision floating point numbers). We plot the best solution obtained for each network as a function of the two competing terms in the overall objective function (Equation 10). The information presented in such a plot serves two purposes. First, it provides informa-

tion on the relative “goodness” of a particular solution by comparing it to the bounds of network flow and coverage shortfall. It also provides a quantification of the value of additional APs and the appropriate number of APs for a particular scenario.

Figure 11 presents the best solutions to networks having two to eight APs in the 45-acre Fort Ord region and illustrates the relative contribution of client coverage and network flow to the overall objective value. We observe a decreasing gain in coverage with more than four APs, which we interpret to mean that four APs are sufficient to provide coverage to the majority of that region. As expected, additional APs provide greater network flow.

Observe the main improvement from the optimal 2-node network to the optimal 3-node network is a decrease in coverage shortfall (and in fact there is a slight decrease in delivered network throughput). The optimal 4-node network provides significantly higher delivered network throughput while also reducing the coverage shortfall. For networks of size five and larger, additional APs primarily increase delivered network throughput. Interestingly, it appears that only after achieving adequate client coverage does it make sense to use additional APs to enable greater network flow.

This chart helps answer the questions “How good is our network?” and “How much better can we do?” This information is of

Figure 11. Example bi-objective plot of SRRA+C for the 45-acre Fort Ord region. Each point represents our best solution to $Z(\lambda)$ for the specified number of APs, solved with equal objective weight given to the coverage and flow terms ($w = 1$). Beyond four APs, we observe decreasing benefit in coverage shortfall and increasing benefit in delivered network flow.

considerable value to a decision-maker, and the speed of our SRRA+C algorithm makes this analysis possible.

CONCLUSION AND RECOMMENDATIONS FOR FUTURE WORK

In this paper, we present a technique for designing WMNs that maximizes client coverage area while considering network service and radio propagation over terrain. Our SRRA+C formulation does not require information about device-specific characteristics such as radio wave modulation scheme or network routing protocol. Although these simplifications may reduce the predictive power of the formulation, it is possible to add such considerations. Further, this general approach makes it very easy to quickly model networks of diverse devices and capabilities.

Using our decision support tool, we compare the performance of enumeration and the DIRECT algorithm in solving the SRRA+C prob-

lem. We show that the DIRECT algorithm can find good solutions much faster than complete enumeration, and is capable of finding better solutions than enumeration because DIRECT considers a continuous solution space. DIRECT is guaranteed to find the optimal solution to the SRRA+C problem as the number of iterations goes to infinity, and each iteration of DIRECT provides a solution at least as good as the previous.

Although we cannot provide certificates that guarantee the optimality of any SRRA+C solution, proof of optimality in this domain is not necessary and in fact wastes critical time. In practice, operators need good working WMN topologies quickly, and our technique provides exactly this.

Our techniques and associated decision support tool can be used by HA/DR personnel and combat communications planners to quickly design WMNs to support their specific operations. The software runs on a laptop computer, does not require any additional software or solver licenses, accepts map data in a generic file format that is widely available on the Internet, and creates network topologies for virtually

any type of terrain and mesh AP device. This technique requires very little technical expertise and no guesswork.

Future work could investigate the use of higher-fidelity radio propagation and network traffic models, and compare the benefit of increased accuracy to any additional computational workload. In particular, we anticipate the need to investigate radio propagation models that consider vegetation and man-made obstructions, which are critical in calculating radio propagation in urban environments.

We have shown that complete enumeration has limited usefulness as a solution technique for comparison to DIRECT. Future work could compare DIRECT to other sampling algorithms or heuristic approaches. Future work could also explore the use of DIRECT in a parallel- or multiple-processor architecture as DIRECT naturally lends itself to this type of computation (see, e.g., He et al. 2004).

We are currently exploring the use of SRRA+C with DIRECT to build larger networks (of 100 APs and more) and networks with more than one traffic destination. Assessing the impact of disruptions (accidental or intentional) on the design of WMNs (e.g., Grotschel et al. 1995, and Shankar 2008) is a topic that also warrants additional study.

ACKNOWLEDGEMENTS

The authors are indebted to the Naval Postgraduate School (NPS) Hastily-Formed Networks Research Group, specifically R. Meyer and B. Steckler, for technical support and assistance during field testing. The authors also thank several colleagues for constructive comments on this and earlier versions of this paper. Alderson was supported under Office of Naval Research Grant Award Numbers N0001411WX20791 and N0001412WX20058. Our invention has been submitted for a US patent, application serial number 12/772,434.

REFERENCES

Bertsekas, D. 1999. *Nonlinear Programming*, Athena Scientific.

- Brown, G., Carlyle, M., Salmeron, J., Wood, K. 2006. "Defending Critical Infrastructure," *Interfaces*, Vol. 36, N. 6, pp. 530–544.
- Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*, Cambridge University Press.
- Cisco Systems, Inc. 2009. "Cisco 100 Series Lightweight Access Points Data Sheet," http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps6306/ps6315/product_data_sheet0900aecd8025708a_ps6306_Products_Data_Sheet.html, retrieved Nov. 1, 2012.
- Defense Mapping Agency. 1989. "DMATM 8358.2, The Universal Grids: Universal Transverse Mercator (UTM) and Universal Polar Stereographic (UPS)."
- COST (European Cooperation in the Field of Scientific and Technical Research). 1999. "Digital Mobile Radio Towards Future Generation Systems: COST 231 Final Report."
- Grotschel, M., Monma, C.L., and Stoer, M. 1995. "Design of Survivable Networks," *Network Models*. M.O. Ball, T.L. Magnanti, C.L. Monma, and G.L. Nemhauser, eds. *Handbooks in Operations Research and Management Science*, Vol. 7. Elsevier North-Holland, pp. 617–669.
- He, J., Verstak, A., Watson, L., Stinson, C., Ramakrishnan, N., Shaffer, C., Rappaport, T., Anderson, C., Bae, K., Jiang, J., and Tranter, W. 2004. "Globally Optimal Transmitter Placement for Indoor Wireless Mesh Networks," *IEEE Transactions on Wireless Communications*, Vol. 3, No. 6, pp. 1906–1911.
- Horst, R., and Hoang, T. 1996. *Global Optimization*. Springer.
- 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, pp. 157–181.
- Longley, A.G., and Rice, P.L. 1968. *Prediction of Tropospheric Radio Transmission Loss Over Irregular Terrain. A Computer Method-1968*, Institute for Telecommunications Sciences.
- Luss, H., and Gupta, S.K. 1975. "Allocation of Effort Resources Among Competing Activities." *Operations Research*, Vol. 23, No. 2, pp. 360–366.

- Nicholas, P. 2009. "Optimal Transmitter Placement in Wireless Mesh Networks." MS thesis, Operations Research Department, Naval Postgraduate School.
- Olexa, R. 2005. *Implementing 802.11, 802.16, and 802.20 Wireless Networks: Planning, Troubleshooting and Operations*, Elsevier.
- Shankar, A. 2008. "Optimal Jammer Placement to Interdict Wireless Network Services." MS thesis, Operations Research Department, Naval Postgraduate School.
- Shannon, C. 1949. "Communication in the Presence of Noise," *Proceedings of the IRE*, Vol. 37, pp. 10–21.
- Xiao, L., Johansson, M., and Boyd, S. 2004. "Simultaneous Routing and Resource Allocation Via Dual Decomposition," *IEEE Transactions on Communications*, Vol. 52, pp. 1136–1144.
- Zhang, Y., Luo, J., and Hu, H. 2006. *Wireless Mesh Networking: Architectures, Protocols, and Standards*, Auerbach Publications.