

Mobility Control for Throughput Maximization in Ad Hoc Networks

Tamer Nadeem

Siemens Corporate Research

755 College Road East

Princeton, NJ 08540

Email: tamer.nadeem@siemens.com

Srinivasan Parthasarathy

Department of Computer Science

University of Maryland

College Park, MD 20742

Email: sri@cs.umd.edu

Abstract—Physical topology of an ad hoc wireless network imposes fundamental limits on its throughput capacity. In this work, we present a network design algorithm for configuring node locations in an 802.11 ad hoc network with the goal of improving the throughput capacity of the network. Given a network configuration (i.e., a mapping of network nodes to physical locations), we establish an inverse proportional relationship between the interference degree of the network and the guaranteed link-throughput. Motivated by this observation, we present an algorithm which progressively guides the network towards better configurations with lower interference degrees, which in turn increases the network’s throughput capacity. Packet level simulations using the NS-2 simulator shows that the reconfigured topologies obtained by our algorithm consistently outperform the original network topologies for throughput and delay related metrics. In particular, for many cases, we observed over 100% increase in the total network throughput and the minimum guaranteed throughput, over 60% increase in throughput fairness and over 50% reduction in the mean-service delay of packets. Our algorithm admits a simple distributed implementation and can be viewed as a distributed mobility control primitive for improving the throughput performance of mobile ad hoc networks. Alternately, our techniques can also be employed by a centralized designer during network creation time to obtain a network configuration with a high throughput capacity. To the best of our knowledge, ours is the first work which explores the use of guided network configuration strategies for improving the throughput capacity in ad hoc wireless networks.

I. INTRODUCTION

Bandwidth is a premium resource in wireless networks. Wireless networks are typically composed of nodes and links whose peak data-rates are orders of magnitude lower, and whose channel error-rates are several times higher than their wired counterparts. In addition, communication in wireless networks is fundamentally different from wired networks due to the broadcast nature of the wireless medium. When two proximate wireless nodes attempt transmission at the same time, one or both of the transmissions may fail due to interference. Consequently, the total available network bandwidth

is far less than the sum of its parts (link / node bandwidths), and optimally utilizing the scarce bandwidth resources in a wireless network is the key to delivering acceptable throughput performance to end-users.

Several existing techniques for wireless throughput maximization rely on resource augmentation. For example, multi-channel multi-radio (MCMR) wireless networks [17], [23] equip nodes with multiple wireless interfaces, each of which operates on a non-interfering channel; this increases spatial reuse of the communication channel and the network’s throughput capacity. The use of directional antennas [31] localizes wireless signal propagation to a smaller region, limits interference, and thus improves throughput performance. Hybrid networks [2] employ high-speed wired backbones for connecting different parts of a wireless network and augment the available wireless bandwidth. While all these techniques are known to be effective in improving the throughput capacity, the additional investment required for augmenting network resources could be prohibitively expensive and limit their applicability in cost-constrained scenarios.

In this work, we focus on *optimizing the structure (topology) of the network during the time of deployment using interference-aware network design, and adaptive (re)configuration of the network during its operation using controlled mobility* with the goal of maximizing the throughput capacity of the network. Some prime candidates which can benefit from our work include sensor network deployments for long term habitat and environmental monitoring (see for instance, the Great Duck Island (GDI) sensor network [27]), self-organizing network of mobile robots that are equipped with various sensing capabilities for military surveillance [6], [11], and wireless mesh networks for connecting remote hosts in rural communities to the wired Internet [9]. In general, our techniques are applicable to long-term ad hoc wireless deployments that exhibit static traffic patterns, or communication patterns which change relatively infrequently over time. In such settings where the traffic is regular enough, we envision that the benefits of adaptive network configuration in order to best serve the communication demands of the

network far out-weigh its costs (such as the energy expended in network (re)configuration, or the bandwidth utilized by the (re)configuration scheme). Our techniques can also inter-operate with resource augmentation schemes described earlier to jointly enhance the network capacity.

We note that mobility control and network design have both been explored in the context of low-energy routing, improving geographic coverage and connectivity, and improving end-to-end reliability for packet transmissions in ad hoc / sensor networks [13], [32], [33], [22], [5], [19], [20], [29]. However, to the best of our knowledge, ours is the first work to explore the use of mobility and location control to optimize the throughput capacity of wireless networks. In particular, we view the following as the main contributions of our work:

A. Our Contributions

- 1) We characterize the impact of network topology on per-link throughputs in an 802.11 network. Specifically, given a *multi-hop* 802.11 network and a set of concurrently active links, we show that the minimum guaranteed link-throughput under saturation conditions is inversely proportional to the maximum interference degree of the network. This generalizes the well-known analysis of Bianchi [7] for the case of *single-hop* 802.11 networks, where the communication graph of the nodes is a clique.
- 2) We model the problem of minimizing the maximum interference degree (and hence optimizing link throughputs) in an 802.11 multi-hop network as a Mixed Integer Quadratic Program (MIQP). The significantly non-convex nature of the MIQP does not allow the use of known optimization techniques for obtaining exact (or approximate) solutions to the MIQP; however, we use the MIQP to derive key insights with important implications for the design of distributed protocols for interference minimization.
- 3) Motivated by the relationship between throughput and interference degree, we present a simple distributed mobility control algorithm which guides nodes towards regions of low-interference in the network, using purely local-control strategies. This ensures that the network progressively acquires a higher throughput capacity without violating network connectivity constraints or geographic locations constraints which must be respected by the nodes in the network. The mechanisms we propose in this paper aim to mitigate the impact of interference through intelligent mobility control, and hence they can be used to optimize many related objectives (such as minimizing packet losses due to interference, fairness, etc.).

We evaluate the performance of our algorithm using packet-level simulations with the NS-2 network simulator for a variety of traffic scenarios. The reconfigured network topologies obtained by our algorithm is seen to consistently outperform the original topology for a variety of throughput related metrics as well the mean-service delay. Our simulations also reveal traffic regimes under which the mobility control algorithm is most effective in increasing the network capacity.

The rest of the paper is organized as follows. We survey related work in Section II. In Section III, we discuss the models, metrics, definitions, and notation that are used in the rest of the paper. In Section IV, we analyze the impact of the maximum interference degree on the guaranteed link-throughput of the network under saturation conditions and establish an inverse relationship between the two. This naturally suggests the use network design strategies which minimize the interference degree in order to improve throughput capacity. We then show in Section V that the problem of minimizing interference degree in ad hoc networks can be modeled as a mathematical program. We derive important insights from this model, and design a mobility control algorithm in Section VI which readily admits a simple distributed implementation, and discuss further optimizations to our algorithm in Section VII. In Section VIII, we deal with the performance evaluation of our algorithm using NS-2 simulations.

II. RELATED WORK

The idea of using controlled mobility in an ad hoc network for improving its performance was first pioneered by Goldenberg *et al.* [13]. In this work, the authors study the problem of minimizing the total energy for routing a given set of multi-hop flows. The authors present a distributed algorithm with the guarantee that the network converges to the optimal energy configuration. Our work is similar in spirit to this work in the following sense. We seek an optimal configuration which maximizes the throughput metrics of interest to us for the given set of flows. However, the underlying combinatorial structure in our problem is considerably more complex as we show in Section IV with important implications for the design of distributed protocols. In particular, unlike the energy minimization problem, we show that local-gradient based distributed algorithms are *not* guaranteed to achieve the optimal throughput configuration.

Mechanisms for deploying and repositioning mobile gateways, sinks, or base stations for minimizing energy consumption and prolonging network lifetime has also been addressed by several researchers [24], [12], [3], [4], [5], [19], [20], [29]. For instance, authors in [24], [12] propose integer programming and ray-tracing models respectively to compute the minimum number of base stations and their locations, for coverage in indoor WLAN networks. A simple analytical formulation is proposed in [3], [4] to allow the gateways in sensor networks to react and alter its positions in an attempt to maximize the efficiency of inter-domain communications. The proposed mechanism moves the gateway to the weighted geographic centroid of its sensors by considering the location and traffic generated by those sensors. Gateways reposition periodically to respond to the change in node locations and traffic pattern in order to gradually get closer to nodes with high traffic and maximize network lifetime. Unlike the assumption of the unconstrained gateway's capability to move, authors in [5], [19], [20] propose mechanisms which also consider topology and connectivity constraints in controlling gateway mobility. A mechanism for controlling the mobility of relaying sensors instead of controlling the sink in order to prolong network lifetime is investigated in [29].

Several authors attempt to exploit mobility of one or more designated nodes in the network to improve the success probability of end-to-end packet receptions in sparse ad hoc networks that suffer from frequent network partitions [32], [33], [16]. Authors in [32], [33] introduce a pro-active Message Ferrying (MF) approach which utilizes a set of special mobile nodes called “message ferries”, to provide communication service for nodes in the deployment area. Message ferries move around the deployment area and take responsibility for carrying data between nodes. With knowledge about ferry trajectories, nodes can adapt their own trajectories to meet the ferries and transmit or receive messages in order to communicate with distant nodes that are out of range or not currently connected. A similar approach for sensor networks was proposed in [16]. As a specific mobile node moves in close proximity to sensors, data is transferred to the mobile node for later depositing at the destination of the data. These mechanisms, unlike our work, target delay-tolerant applications and are *not* geared towards exploiting controlled mobility. A related approach to recovering from ad hoc network partitions is by deploying mobile forwarding nodes as in [22]. The mobile forwarding nodes detect network partitions and adaptively change their location for connecting partitions.

The dynamic aspects of the coverage in mobile sensor networks has been explored recently in [14], [30], [25], [18]. For instance, authors in [25] propose Dynamic Coverage Maintenance (DCM) schemes which exploit the limited mobility of the sensor nodes, and compensate the loss of coverage through controlled and careful repositioning of sensors with minimum expenditure of energy. The authors propose a set of distributed DCM schemes which can be executed on individual sensor nodes having knowledge of only their local neighborhood topology.

While all of the above literature point to effective use of mobility (both controlled and uncontrolled), the main point of departure for our work in comparison with the above literature is in our use of novel mechanisms which exploit controlled mobility for maximizing the throughput capacity of the network without violating connectivity or geographic constraints.

III. BACKGROUND

We now present the models, metrics and the notation used in the rest of the paper.

Network Model: We model an 802.11 ad hoc wireless network using a connected undirected graph $G = (V, E)$. All nodes in V lie on the two dimensional XY plane. Nodes employ a common transmission power which results in a uniform transmission range r_{tx} . Nodes u and v can form a link $(u, v) \in E$ only if their distance $d(u, v) \leq r_{tx}$. Note that the distance constraint is a *necessary condition but not a sufficient one*. Links between sufficiently proximate nodes may not exist due to high error levels in the communication channel, physical obstructions, etc. Further, ad hoc networks often employ virtual backbone mechanisms (see [10], [26] for instance) which expose a limited subset of network edges to

higher-level routing protocols. Such mechanisms significantly reduce route-message exchanges and other routing protocol overheads while preserving network connectivity and the capacity of the underlying physical network. In general, let E_{data} denote the set of ordered node-pairs (u, v) which require one-hop data communication between themselves; let E_{conn} denote the set of links that are needed to maintain network connectivity (for e.g., E_{conn} could be selected by a virtual backbone mechanism); we let $E = E_{data} \cup E_{conn}$ denote the set of all links in the network.

We now briefly review the key ingredients of Distributed Coordination Function (DCF), which is the fundamental mechanism used by 802.11 for accessing the wireless channel. A node with a new packet to transmit monitors the channel for a fixed period of time called the Distributed Inter-frame Space - DIFS. If the channel is idle during this duration, the node transmits. Otherwise, the node monitors the channel until it is found idle for a time of length DIFS. At this point, the node waits for a time generated by the random exponential backoff scheme before transmitting; the backoff is also employed between consecutive transmissions by the node. The exponential scheme works as follows: at each packet transmission, the random wait time is chosen uniform in the range $(0, \dots, w - 1)$ (in units of time slots). In the beginning, each node sets its *contention window* w to the minimum value W . After each unsuccessful attempt, the value w is doubled up to a maximum value of size $2^{b-1}W$, where b is the *number of backoff stages*. Nodes maintain a backoff time counter which is decremented as long as the channel is sensed *idle, frozen* when the channel is detected busy, and reactivated when the channel is sensed idle again for a period of time equal to DIFS.

The sensing mechanism at a node u deems the channel busy if the total received signal strength is above a certain threshold, called the carrier-sense threshold: this in turn is the case if there exists another node v which is currently transmitting, and which is within the *carrier-sense range*, $r_{cs}(> r_{tx})$, of u . Given a network configuration Γ , we define the interference graph $H_{\Gamma}(V, F)$ as follows: let edge $(u, v) \in F$ if both u and v are active nodes (i.e., end-point of some link in E_{data}) and if $d(u, v) \leq r_{cs}$. For all $u \in V$, let $\delta(u)$ denote its *interference degree* (i.e., the degree of u in H). We define $\delta_{\max} = \max_{u \in V} \delta(u)$. In Section IV, we discuss the crucial role played by the interference degree δ in limiting the throughput of an 802.11 ad hoc network.

Mobility model: We assume a *controlled node-mobility model* as follows: each node $u \in V$ has two dimensional XY coordinates $(x(u), y(u))$. Nodes can employ a mobility control algorithm which allow them to *choose their locations* subject to the following constraints: **(P1)** network connectivity constraint for each link $(u, v) \in E$; this requires the end-points u and v to choose locations such that they can communicate directly between themselves; in particular, the mutual distance $d(u, v)$ can be at most r_{tx} ; **(P2)** physical boundary constraints for each node u of the form $x(u) \in [x_{low}(u), x_{high}(u)]$ and $y(u) \in [y_{low}(u), y_{high}(u)]$; for e.g., nodes in a mesh/sensor network may be required to serve/sense a fixed target region; this also models stationary nodes; **(P3)** network design constraints which decide the time at which the mobility control

strategy can be executed; we assume that in mesh/sensor networks, node locations need to be decided during the time of network creation; in a mobile ad hoc network, we assume nodes are allowed to employ a (distributed) mobility control strategy which allow nodes to choose and change their locations when the network is operational.

Metrics: Recall that E_{data} denotes the set of links which involve direct data communication across themselves. For each link $e \in E_{data}$, we assume that the transmitter always has a non-zero number of packets to send to the receiver, i.e., we assume that the network is *saturated*. Let $\tau(e)$ denote the *saturation throughput* of link e , i.e., the throughput achieved by the 802.11 MAC protocol across link e in network saturation conditions. Define the *network configuration* to be a mapping of each node to a fixed location, such that the node locations satisfy the network connectivity **(P1)** and boundary constraints **(P2)**.

Given a fixed network configuration Γ , we are interested in three throughput metrics: **1)** the guaranteed saturation throughput $\eta(\Gamma) \doteq \min_{e \in E_{data}} \tau(e)$; this is the throughput which is *guaranteed* to be achieved by *every link* in E_{data} ; for multi-hop connections, the end-to-end throughput of the multi-hop flow is equal to that of the throughput of its bottleneck link; thus, maximizing η directly translates to maximizing the *guaranteed end-to-end throughput* for a given set of multi-hop connections. **2)** The total throughput $\mu(\Gamma) \doteq \sum_{e \in E_{data}} \tau(e)$ which is a natural measure of network utilization. **3)** The fairness index $\lambda(\Gamma) = \frac{(\sum_{e \in E_{data}} \tau(e))^2}{|E_{data}| \sum_{e \in E_{data}} \tau(e)^2}$; while there is no fairness metric which is considered “universally good”, the above metric introduced by Jain, Chiu and Hawe [15] encapsulates the fact that a fairness value close to 1.0 corresponds to equitable channel sharing and all throughput values being nearly equal, while a fairness value close to 0.0 corresponds to disparate throughput values due to inequitable sharing of the channel. For clarity of exposition, whenever it is clear from the context, we will omit the argument Γ from η , μ , and λ in the rest of the paper.

Our approach in brief: The objective of our work is to achieve the optimal network configuration which maximizes the value of the three metrics of interest to us (possibly, in a distributed manner). The basic idea underlying our algorithm is as follows: each transmitter/receiver (transceiver) estimates the amount of interference it experiences currently and sequentially attempts to move to a new location with lesser interference, subject to network connectivity and node boundary constraints. We discuss how interference estimation as well as sequential mobility can be achieved in three increasingly complex environments: where network designer has complete knowledge of the network (mesh/sensor networks), where nodes have accurate information about their interference neighborhood (MANETs/ sensor networks), and where nodes know only their single-hop communication neighborhood (MANETs/sensor networks).

IV. IMPACT OF INTERFERENCE DEGREE

We now analyze the crucial role played by the interference degree δ_{max} on the guaranteed saturation throughput η for

multi-hop networks. A well-understood and important special case which was first studied by Bianchi [7] is as follows: consider a set C of m nodes, all of which are within the transmission range, r_{tx} , (and hence r_{cs}) of each other and where each node in C acts as a transmitter to some fixed intended receiver. Bianchi showed that the total achievable saturation throughput μ_{cliq} for this setting can be approximated as follows:

$$\mu_{cliq} = \frac{\mathbf{E}[P]}{T_s + \sigma K + T_c(K(e^{\frac{1}{K}} - 1) - 1)} \quad (1)$$

Here, $\mathbf{E}[P]$ is the expected size of the payload, T_s is the average time taken for a successful packet transmission, T_c is the average length of time the channel is sensed busy by each node during a collision, σ is the duration of an empty slot, $T_c^* = \frac{T_c}{\sigma}$, and $K = \sqrt{\frac{T_c^*}{2}}$. Observe that these values and hence μ_{cliq} are *fixed constants* independent of the number of nodes m . By symmetry, it follows that the guaranteed per-link throughput $\eta_{cliq} = \mu_{cliq}/m$.

Bianchi’s analysis relies on the fact that every sender contends with every other sender in the network and due to *network symmetry within a clique*, all nodes would share the maximum achievable network throughput equitably. However, symmetry is generally not a valid assumption in the case of a multi-hop network: nodes differ in the structure of their interference neighborhood and may experience differing amounts of channel contention. This naturally motivates the following question: how does the saturation throughput of a node depend upon its interference degree in an arbitrary multi-hop wireless network? In order to better understand the scenario, we first present the results of a simulation study: fix a transmitter node u and its intended recipient v , where v is within the transmission range of u . Let \mathcal{D}_{cs} , the disk of radius r_{cs} centered at u , denote the interference neighborhood of u . We place a set of δ nodes uniformly at random within \mathcal{D}_{cs} ; each of these δ nodes act as a sender to a unique neighbor within its transmission range. We measure the effect of u ’s interference degree δ as a function of the throughput observed on the link (u, v) in this scenario. The results of this scenario are plotted in Figure 1. It is immediately obvious from the hyperbolic curve that the saturation throughput $\tau((u, v))$ is inversely proportional to the interference degree δ of u .

We establish this inverse relationship between the interference degree and saturation throughput in an *arbitrary multi-hop network* using Theorem 1. The theorem states that in a network of maximum degree δ_{max} , the guaranteed saturation throughput η is upper-bounded as follows: $\eta \leq \frac{6\mu_{cliq}}{\delta_{max} + 1}$. The proof of Theorem 1 uses a careful packing argument to show that in any given 802.11 ad hoc network of degree $\delta_{max} + 1$, there is always a large clique C' of size $\frac{\delta_{max} + 1}{6}$. This fact combined with Bianchi’s analysis implies that the guaranteed saturation throughput for any link with a sender in C' is at most $\eta \leq \frac{6\mu_{cliq}}{\delta_{max} + 1}$ (this is an upperbound and *not* an exact estimate). *For ease of analysis*, the statement and proof of theorem 1 assumes that the set of edges in E_{data} form a matching (i.e., each node communicates with at most one of its neighbors). We discuss later, how our analysis can be easily

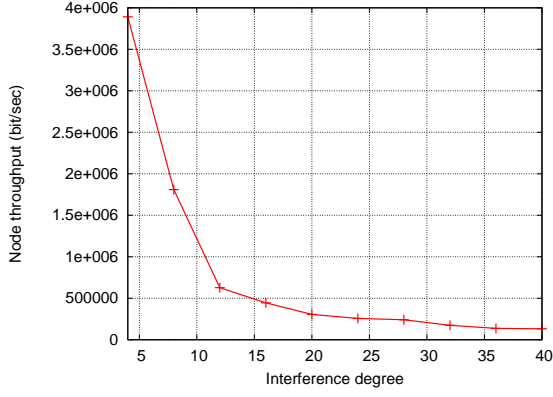


Figure 1. The saturation throughput for a node is inversely proportional to its interference degree

generalized to the case where E_{data} is an arbitrary subset of E .

Theorem 1: For any given network $G = (V, E)$ whose interference degree is δ_{\max} , the guaranteed saturation throughput $\eta \leq \frac{6\mu_{cliq}}{\delta_{\max}+1}$.

Proof: Let $u \in V$ be a node of interference degree δ_{\max} . Consider a disk $\mathcal{D}_{cs}(u)$ of radius r_{cs} centered at u . Clearly all nodes which interfere with u , lie within the disk $\mathcal{D}_{cs}(u)$. Consider a sector of $\mathcal{D}_{cs}(u)$ which has a radius r_{cs} , and which subtends an angle $\frac{\pi}{3}$ at u : we first observe that any two nodes within this sector is of distance at most r_{cs} . Hence, all nodes within this sector interfere with each other and form a clique of interfering nodes (see Figure 3). Next, we note that $\mathcal{D}_{cs}(u)$ can be partitioned into at most 6 such disjoint sectors; hence, it follows that there is a clique of size at least $\frac{\delta_{\max}+1}{6}$ in G which contains u . Crucially, all the clique nodes share the channel with each other, since all of them are within the carrier-sense range of each other. It now follows from the analysis of Bianchi (Eq. 1) that the guaranteed saturation throughput of links whose transmitters belong to this clique is at most $\eta \leq \frac{\mu_{cliq}}{\frac{\delta_{\max}+1}{6}} = \frac{6\mu_{cliq}}{\delta_{\max}+1}$. ■

We now discuss how Theorem 1 can be generalized to the case when E_{data} is an arbitrary subset of E . Let $\alpha(u)$ denote the data-neighbor degree of a node u . This is simply the number of nodes in the network to which u transmits data (i.e., the number of edges in E_{data} which emanate from node u). Observe that Theorem 1 imposes an upper-bound of $\frac{6\mu_{cliq}}{\delta(u)+1}$ on the total throughput of node u . It follows that there exists a link incident on u , whose saturation throughput is upper-bounded by $\frac{6\mu_{cliq}}{\alpha(u)\cdot(\delta(u)+1)}$. In general, the guaranteed saturation throughput η can now be upper-bounded by $\min_u \frac{6\mu_{cliq}}{\alpha(u)\cdot(\delta(u)+1)}$, which provides the required generalization of Theorem 1.

Next, we observe that our analysis yields an optimistic upper-bound for η (rather than a tight one) due to two reasons. Firstly, for a given node u , our analysis ignores the effect of the *hidden* nodes which do not lie within the carrier-sense range of u , but lie within the carrier-sense range of u 's intended receiver(s) (and hence interfere with them). The packet losses at the receiver(s) caused due to collisions with hidden nodes would lead to a further throughput reduction which is not

captured by Theorem 1. We note that incorporating the effect of hidden nodes does not change this result qualitatively, but only leads to an increased value of the constant factor involved (this constant is 6 in Theorem 1). Secondly, Eq. 1 holds only when the number of backoff stages b and the length of the smallest backoff window W for the 802.11 protocol at each node has been optimized w.r.t. its interference degree [7]. However, the current 802.11 DCF standard specifies fixed values of b and W which are specific to the physical layer but independent of the node's interference degree [7]. In the latter case, the total saturation throughput μ'_{cliq} can be evaluated numerically by solving for the values of t and p in the following set of equations and by substituting the value of t in Eq. 2. Here, t represents the transmission probability at a given time slot and p denotes the conditional probability that a collision occurs given that a transmission occurred at a particular time slot (see Bianchi [7] for details).

$$\begin{aligned} t &= \frac{2 \cdot (1-p)}{(1-2p) \cdot (W+1) + pW \cdot (1-(2p)^b)} \\ p &= 1 - (1-t)^{(m-1)} \end{aligned}$$

$$\mu'_{cliq} = \frac{\mathbf{E}[P]}{T_s - T_c + \frac{T_c^* - (1-t)^m \cdot (T_c^* - 1)}{mt \cdot (1-t)^{(m-1)}}} \quad (2)$$

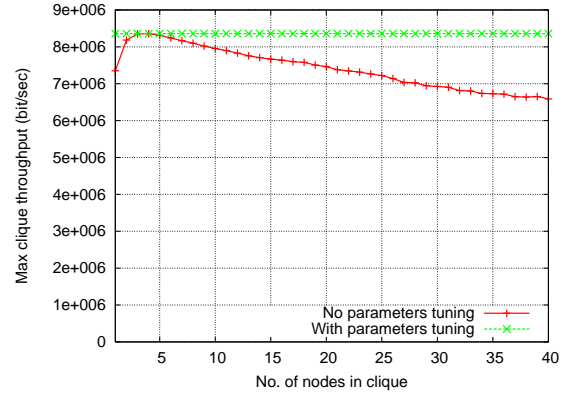


Figure 2. Variation of μ'_{cliq} w.r.t. to m

Figure 2 plots μ'_{cliq} for a range of values of m ; the figure illustrates that $\mu'_{cliq} \leq \mu_{cliq}$ and more significantly, μ'_{cliq} is a decreasing function of m for fixed values of b and W . This implies that in a multi-hop network, the impact of the interference degree δ on the saturation throughput η for a high-degree node is more severe than that of a low-degree node.

To conclude the arguments in this section, Theorem 1 and Eq. 2 establish the crucial connection between the network degree δ_{\max} and the guaranteed saturation throughput η in a *multi-hop* ad hoc network. More importantly, they suggest the use of network design primitives which minimize the interference degree of the network, for maximizing the guaranteed saturation throughput across the active links in the network; this observation forms the basis of the mobility control models and algorithms developed in the remainder of this paper.




Figure 3. Interference neighborhood: the disk of radius r_{cs} with 6 disjoint 60° sectors represents the interference neighborhood of the central node. Nodes within a sector interfere with each other; at least one of the sectors contain $1/6^{th}$ of the interfering nodes.

V. A QUADRATIC PROGRAMMING FORMULATION

We now present the Mixed Integer Quadratic Programming (MIQP) formulation for the minimum-interference network design problem. The objective of the MIQP, presented in Figure 4, is to minimize the network interference degree δ_{\max} subject to network connectivity, interference and node location constraints. Recall that $x(u), y(u)$ denote the X and Y coordinates of node u . In the MIQP, S is the set of transmitting nodes in the network. The indicator variable $z_{u,v}$ denotes if transmitters u and v lie within the carrier-sense range of each other: if $z_{u,v} = 1$, then it contributes to the interference degrees of both u and v (Eq. 6); if $z_{u,v} = 0$, then the two transmitters are required to be more than a distance of r_{cs} apart (Eq. 5). Constraints (4) state that for all links in E , the two end-points must be within a distance of r_{tx} . Eq. (10) capture the location constraints associated with each node. Among all network configurations that respect the location and connectivity constraints, the optimal configuration for the MIQP is the one which has the least interference degree.

$$\begin{aligned}
 & \min \delta_{\max} & (3) \\
 \forall (u, v) \in E : & (x(u) - x(v))^2 + (y(u) - y(v))^2 \leq r_{tx}^2 & (4) \\
 \forall ((u, v) \in S \times S \text{ and } (u \neq v)) : & & \\
 (x(u) - x(v))^2 + (y(u) - y(v))^2 & > (1 - z_{u,v})r_{cs}^2 & (5) \\
 \forall (v \in S) : & \delta(v) = \sum_{u \in S} z_{u,v} & (6) \\
 \forall ((u, v) \in S \times S \text{ and } (u \neq v)) : & z_{u,v} \in \{0, 1\} & (7) \\
 \forall (v \in V) : & \delta(v) \leq \delta_{\max} & (8) \\
 \forall (v \in V) : & x(v) \in [x_{low}(v), x_{high}(v)] & (9) \\
 & \& \quad y(v) \in [y_{low}(v), y_{high}(v)] & (10)
 \end{aligned}$$

Figure 4. The network design problem: δ_{\max} is the interference degree of the network which is to be minimized subject to network connectivity constraints (4) and node location constraints (10). The indicator variable $z_{u,v}$ denotes if u and v interfere with each other: if they do not interfere, i.e., $z_{u,v} = 0$, then the interference constraints (5) require that $d(u, v) > r_{cs}$.

In general, solving a constrained quadratic program is hard even when the constraints and the objective function are continuous: efficient *convex-optimization* techniques exist for

special cases where the objective is to minimize a convex function over a set of convex constraints. In our MIQP, the set of constraints (5) span a *non-convex* subspace which renders convex-optimization techniques infeasible for obtaining an optimal solution. Intuitively, the basic difficulty in solving our MIQP optimally is the existence of local-minima which prevent local-gradient based solution strategies from converging to an optimal solution. We emphasize that the existence of local optima for our MIQP is not an artifact of our modeling technique but is an inherent feature of the optimization problem itself (see Figure 5 for an illustration).




Figure 5. Existence of local minima: the six big circles denote interference-regions of six nodes located at their centers (and not shown in the figure). The middle node in the box wants to minimize its interference degree by choosing a location within the dotted rectangular box; the optimal position is the left-extreme (denoted by solid circle around a node) while a greedy local-gradient approach would choose the right-most position which is sub-optimal (denoted by dotted circle).

From the perspective of protocol design, the above discussion motivates the need for a *local-control heuristic which admits a simple distributed implementation and which works well in practice* rather than computationally expensive optimal solutions. We focus on the design of such a heuristic in the next section.

VI. DISTRIBUTED MOBILITY CONTROL

We now present the distributed implementation of our mobility control algorithm for ad hoc networks. Briefly, our algorithm consists of two subroutines, namely *location management* and *token management*. The algorithm works in an iterative fashion: within a single iteration, the token management subroutine ensures that a special control packet (the token) is passed around exactly once to all nodes within the network. The node which currently holds the token is the one which executes the location management subroutine: this subroutine is responsible for relocating the node to a new location where it experiences lesser interference without violating its connectivity and boundary constraints. Together, both these subroutines ensure that nodes change their locations in a coordinated manner such that the network migrates to progressively better configurations with each iteration.

At the very outset, we emphasize that the description below is aimed at providing a specific implementation of the ideas which underlie the mobility control algorithm in a *distributed* setting. Several elements in the following description can be simplified (or even omitted) in the case of centralized setting: for e.g., in the case of a mesh network designer who

has complete knowledge about the conditions in her network (we assume such a centralized setting in our simulations). In particular, the token management subroutine whose objective is to prevent simultaneous uncoordinated relocation of nodes, can be totally omitted in the centralized version. We now present the details of the two subroutines below.

A. Location Management

Consider a fixed node u whose current position is given by the coordinates $x(u), y(u)$. Let B_u denote the square whose length is $2r_{tx}$ and which is centered at u . Let A denote the uniform partitioning of B_u into $m \times m$ grids such that each grid is a square of length $\frac{2r_{tx}}{m}$ (see Figure 6). Let $A[i]$ denote the midpoint of the i^{th} grid in a fixed linear ordering of the m^2 grids (the ordering will be specified later). Algorithm 1 presents the pseudocode for the location management subroutine for node u . As a precondition, the algorithm requires u to possess the token before the execution of the subroutine. Lines 1 and 2 initialize the variables $final(u)$ and $\delta_{final}(u)$ to u 's current position and current interference degree respectively. Line 3 allows node u to explore each grid location sequentially: if grid i lies outside the boundaries specified by the node's location constraints (line 4), then the node does not visit this location (or equivalently, does not consider this location further after visiting it). Otherwise, u visits $A[i]$ (line 7) and checks if it is still connected to all its neighbors in the network. If this is not the case, then grid i is marked as *unsafe* (line 10); otherwise, if i is safe and has the least interference degree amongst all the safe grids visited so far (line 15), then $A[i]$ is updated as the final target location for u (line 16) and the final interference degree is the one which u experiences at $A[i]$ (line 17). After exploring all the m^2 grids in this manner, node u finally moves to the safe grid which has the least interference degree (line 20).

A node can use simple local-control techniques for computing its interference degree at a specific location in the network. For instance, a node can observe the number of unique signal-strength values measured by itself at its current location which are greater than the carrier-sense threshold; the node can then estimate this value as its interference degree. This is not an exact value but only an approximation due to the assumption that two nodes which interfere with a fixed node u , are never at the same distance from u (and hence, their signals have different signal-strength values at u). In Section VII, we discuss how other distributed approximation techniques could be used in the absence of accurate mechanisms to estimate received signal-strength at a node.

We now specify the linear ordering of the grids as explored by node u . The natural greedy strategy for node u is to start from the center of the box b_u (the initial location of node u) and spiral outwards towards the perimeter of b_u (see Figure 6). This ensures that the node travels the minimum possible distance when moving from one grid to the next. In particular, we note that the nodes only need to know their current location *relative* to their initial coordinates and can be easily implemented *without* resorting to a full-fledged location determination system.

Algorithm 1 Location Management sub-routine for node u

Require: Node u possesses the token currently

- 1: $final(u) \leftarrow$ current position of u
- 2: $\delta_{final}(u) \leftarrow$ current number of interfering nodes for u
- 3: **for** ($i = 1; i \leq m^2; i = i + 1$) **do**
- 4: **if** ($A[i] \notin [x(u)_{low}, x(u)_{high}] \times [y(u)_{low}, y(u)_{high}]$) **then**
- 5: **break**
- 6: **end if**
- 7: visit $A[i]$
- 8: **for all** $v \in N(u)$ **do**
- 9: **if** $d(u, v) > r_{tx}$ **then**
- 10: status[i] = unsafe
- 11: **break**
- 12: **end if**
- 13: **end for**
- 14: $\delta_{current}(u) \leftarrow$ current number of interfering nodes
- 15: **if** (status[i] \neq unsafe) and ($\delta_{current}(u) < \delta_{final}(u)$) **then**
- 16: $final(u) = A[i]$
- 17: $\delta_{final}(u) = \delta_{current}(u)$
- 18: **end if**
- 19: **end for**
- 20: Move to $final(u)$

B. Token Management

Concurrent and uncoordinated execution of the location management algorithm in a distributed setting could lead to two undesirable consequences. First, two end-points of a link could simultaneously choose to move in opposite directions leading to a disconnected network. Second, a large set of nodes which do not mutually interfere with each other to start with could simultaneously choose to move to the same region: this would lead to an increased interference degree for all these nodes in their new location. Both these issues are illustrated in Figures 7 and 8 respectively.

The token management subroutine avoids these undesirable scenarios by a simple distributed locking mechanism. The subroutine *sequentially* passes a special control packet called the *token* around the network. As noted earlier, only the node which is currently in possession of the token can execute the location management scheme. At the termination of the location management subroutine, the node which currently holds the token passes it to a designated neighbor, which then initiates its location management algorithm. A single *iteration* of the subroutine involves all nodes possessing the token (and executing the location management subroutine) exactly once. Repeated iterations of this algorithm drives the network towards progressively better configurations and the algorithm can be terminated after a pre-specified number of iterations.

The sequential token passing can be achieved within a single iteration of the algorithm through the following mechanism. Using a standard distributed spanning tree algorithm (see Lynch [21] for a specific implementation), we compute a spanning tree for the entire network. Token passing can now be easily accomplished through a distributed traversal (such as




Figure 6. Grid exploration in location management: the solid square of length $2r_{tx}$ is the box b_u explored by center node. The dotted square represents the boundary constraints for u . Box b_u is partitioned into 11×11 grids. Grid numbers denote the spiral order used by node u to visit the grids (numbers shown only up to 14).




Figure 7. Uncoordinated location management could lead to loss of connectivity: the figure illustrates the initial and final position of two nodes (small filled circles and small unfilled circles respectively). The big circles represent the initial and final transmission ranges: the nodes are initially within each other's transmission range but lose connectivity; note that the nodes would still be connected if one of them remained at its initial location.

in/pre/post-order traversal) of the spanning tree. The existence of a spanning tree is guaranteed since we assume that the underlying graph G is connected (this is the only reason why the assumption is warranted). In particular, in the context of mesh/sensor networks where centralized network design is possible, we can eliminate token management since there is no longer a need for nodes to coordinate their location management in a distributed fashion.

VII. DISCUSSION

This section deals with the various optimizations that are possible and the modifications that are needed to the mobility control scheme depending upon the amount of information available at the nodes. The location management relies on the ability of the nodes to estimate the interference degree




Figure 8. Uncoordinated location management could lead to high interference: figure illustrates several interfering nodes (small filled circles) and their interference ranges (big circles). The central region does not belong to any node's interference range. Uncoordinated movement could lead to all nodes moving towards the center; this increases all their interference degree.

by observing the received signal-strengths. In the absence of a reliable signal-strength detection mechanism, a node may use the Kalman filter based channel estimation techniques of Bianchi and Tinnirello [8] or the Bayesian estimation techniques of Vercauteren, Toledo and Wang [28] for computing its interference degree: both these techniques rely on local variables observable at a node, such as the conditional collision probability p for a packet being transmitted, and the probability t of the channel being busy to estimate the number of nodes that compete for the channel in an 802.11 network. We note that both these techniques also yield only approximations of the actual interference-degree since they do not take the effect of hidden nodes into account.

The location management scheme can be optimized further when nodes possess a reliable location detection mechanism. In particular, if nodes know their current location and that of their one-hop communication neighbors, a node can easily pre-compute the set of grids which are feasible under the connectivity and boundary constraints and visit only those locations (potentially saving the time and cost of visiting all the grid locations). Further, if nodes know the coordinates of all other nodes within their carrier-sense range, then nodes can pre-compute the optimal location which will be chosen by the location management scheme. This enables the nodes to relocate to their new location without exploring any of their surrounding grids. The latter can be easily determined in dense ad hoc networks with location information, since the carrier-sense range is typically a small constant factor times the transmission range, and nodes which interfere with each other will be two or three communication hops away from each other.

Another optimization which could potentially speed-up our algorithm is distributing multiple tokens concurrently to well-

separated nodes so that they may execute their location management subroutines concurrently. However, two key issues need to be resolved before multiple tokens can be concurrently active in the network: the first is the need for a reliable mechanism which ensures that tokens are not held by two nodes which are within close proximity to each other (this would again lead to the uncoordinated location management problems discussed in Section VI which the token management algorithm had set out to solve in the first place); second is the need for a strategy to terminate the algorithm after a fixed number of iterations. Currently, since there is only a single token which is sequentially circulated across all the nodes, a designated node can be chosen to destroy the token after the fixed number of rounds. This needs to be generalized for the case when multiple tokens are active concurrently. We note that while allowing nodes to concurrently change their locations in a completely distributed manner could optimize the convergence of the algorithm, a complete exploration of techniques which can achieve it is beyond the scope of this paper and is deferred to future work.

VIII. PERFORMANCE EVALUATION

This section deals with the impact of our mobility control algorithm on the throughput performance of the network. Specifically, our simulations are aimed at evaluating the effect of mobility control on total throughput μ , guaranteed throughput η , fairness λ , and maximum and average network degrees under a variety of saturated network conditions. We now discuss the basic setup of our simulations followed by our experiments, observations, and conclusions.

A. Simulation Setup

We use the *ns-2* simulator [1] enhanced with the CMU-wireless extensions for our simulations. The underlying MAC layer mechanism is the IEEE 802.11b protocol with the four way handshake (RTS/CTS/DATA/ACK) scheme. Each node uses a fixed data rate of 11 Mbps and has a transmission range of 250m and a carrier sense range of 550m. We evaluate network performance under three distinct settings: a linear network of nodes placed on a line-segment of length 4000m with one-hop data connections (i.e., each flow's source-destination pair forms the end-points of some link), a two dimensional network of nodes placed on a square 1000m \times 1000m area with one-hop connections, and two dimensional network of nodes placed on a square grid (of varying dimensions) with multi-hop connections (i.e., connections have intermediate relay nodes). Each connection is a saturated flow of UDP packets of 1000 bytes in size.

The independent variables which are varied across the different runs of the experiment are the total number of nodes in the network and the communication degree of each node (the number of neighbors for which each node acts as a transmitter). The original network topology, which is part of the input to the mobility control algorithm, is created by placing the nodes uniformly at random within the network area. For one-hop connections, we choose the source-destination pairs as follows: for each source u , we pick a random neighbor v

within the transmission neighborhood of u and create a one-hop connection between them; if the communication degree (specified as part of the input for each run) is one, then all nodes in the network are picked as sources. If it is greater than one, then additional sources are picked uniformly at random throughout the network until the average number of connections per node equals the communication degree. For multi-hop connections we experimented with three specific example scenarios (which will be discussed later). The choice of the network topology and traffic conditions are motivated by the need to evaluate our algorithm under a wide range of scenarios and also to keep the simulations reasonably short.

The mobility control algorithm is executed under the following network connectivity and boundary constraints: we require only the end-points of communicating links to be within the transmission ranges of each other in the final topology; We impose two types of boundary constraints: for all the single-hop connection scenarios and for two of the multi-hop scenarios, nodes are only required to be within the network boundary (i.e., inside the 4000m line-segment for the linear network and within the 1000m \times 1000m square for the two dimensional network); for one of the multi-hop scenario, we consider an example where connection end-points are assumed to be fixed and only the relay nodes are allowed to relocate. We view our simulation scenarios as examples of a centralized network designer who optimizes network performance for arbitrary (but known) network topology and traffic conditions using our mobility control algorithm.

In each run, we execute the mobility control algorithm subject to the connectivity and boundary constraints for 100 rounds and obtain the final topology. We measure the metrics of interest to us in both the initial and final topologies by running the simulator for a (simulation) time period of 100 seconds. We now discuss our results.

B. Results

Figures 9 and 10 present the variation of the total network throughput vs. the number of network nodes for linear and two-dimensional networks respectively; the 'degree' parameter for each curve refers to the per-node communication degree. We discuss two observations which hold for both one and two dimensions. **1)** The total network throughput for the original network topology does *not* seem to vary much with the per-node communication degree. This can be attributed to the fact that, under saturation conditions, the total available throughput within a given region of the network is fixed (and independent of the number of connections which contend for the channel within this region). Note that this is also consistent with that of our analysis in Section IV. **2)** The total throughput for the final network topology output by our mobility control algorithm depends on the per-node communication degree. In particular, while the mobility control algorithm always seems to yield a final configuration with a better total throughput than the original configuration, the extent of improvement is greater when the per-node communication degree is smaller. This observation can be explained by the fact that each communicating link imposes a connectivity constraint on the mobility-control

algorithm: hence large communication degrees impose heavy constraints on the choice of network topologies which in turn restricts the total throughput of the final topology.

Observe that the percentage of throughput increase for two-dimensional networks is higher than that of linear networks for any fixed communication degree. We believe this can be explained by the fact that the mobility control algorithm has two degrees of freedom while moving a node to a new location (nodes can move along both X and Y axes) as opposed to a linear network where nodes can move only within the line-segment: the algorithm thus has a greater chance of discovering a topology which reduces node degree in two dimensions rather than one dimension. Finally, we note that the throughput improvements in the mobility controlled topologies can be directly attributed to the reduced node degrees. Let $\delta(v)$ denote the interference degree of node v . Figure 11 plots the sum $\sum_{v \in V} \frac{1}{\delta(v)}$ for the same settings which were studied for the two-dimensional network in Figure 10. Recall our arguments in Theorem 1 that the saturation throughput achievable by a node is inversely proportional to its interference degree. The plots in Figure 11 along with 10 empirically validates these analytical arguments.




Figure 9. Total throughput for networks with different degree versus number of nodes for original generated topologies and controlled topologies

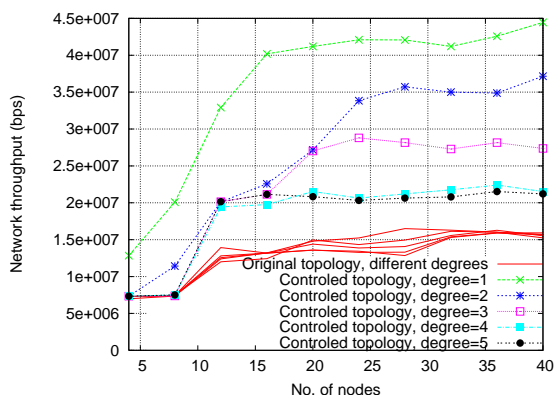


Figure 10. Total throughput for networks with different degree versus number of nodes for original generated topologies and controlled topologies

Figure 12 plots the effect of our algorithm on the maximum interference degree of the two dimensional network. As

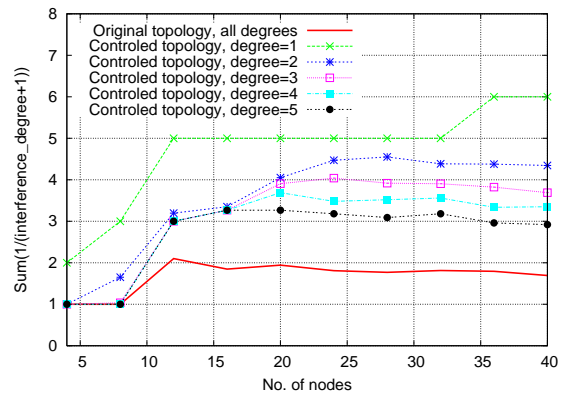


Figure 11. The total sum of the inverse of node degrees for original and controlled topologies

expected, the maximum interference degree increases with the number of nodes in the network, as all of them occupy a fixed network area. More significantly, the mobility control algorithm is more effective in reducing the maximum interference degree, when the communication degree is lower (and hence, network connectivity constraints are fewer in number). This effect is also reflected in the measurement of the minimum link throughput for these scenarios as seen in Figure 13. We remark that high communication degree implies that more connections share the throughput available at any given node; together, they account for the low values of guaranteed per-link throughput in scenarios with high communication degree, even after our mobility control algorithm is executed.

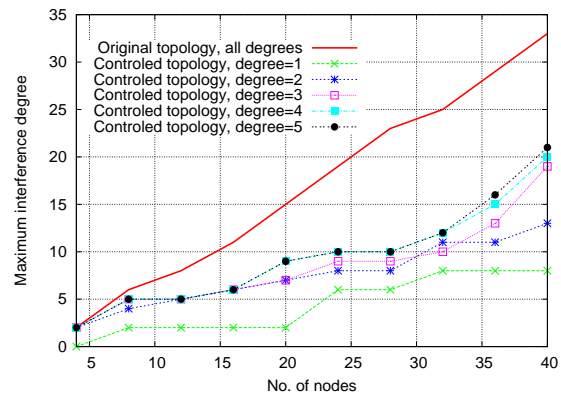


Figure 12. Maximum interference degree for networks with different degree versus number of nodes for original generated topologies and controlled topologies

Figure 14 plots the fairness experienced by the packets for the original and mobility controlled two-dimensional topologies. Recall that the fairness index λ was defined as the ratio $\frac{(\sum_{e \in E_{data}} \tau(e))^2}{|E_{data}| \sum_{e \in E_{data}} \tau(e)^2}$, where E_{data} is the set of active links, and $\tau(e)$ denotes the saturation throughput of link e . A fairness index closer to 1.0 implies that links have similar throughput values and share the wireless channel equitably. For the case where the communication degree is one, observe how the mobility controlled topology has a fairness value of 0.9 or

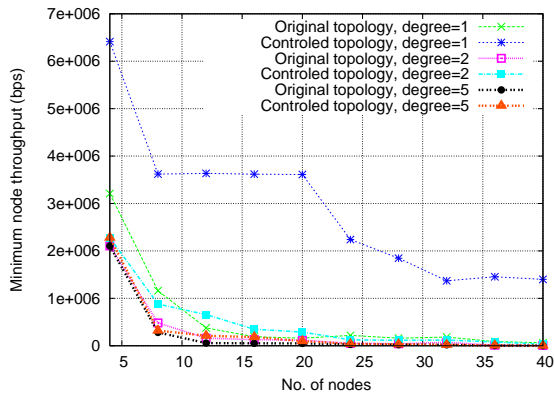


Figure 13. Minimum flow throughput for networks with different degree versus number of nodes for original generated topologies and controlled topologies

above which improves the fairness of the original topology by above 50%. However, as with the throughput metrics studied before, the effectiveness of mobility control for improving fairness, drops with increasing communication degree. Indeed, as the final point on the curve for degree two indicates, the mobility controlled fairness could be actually lesser than the original fairness when the communication degree increases, *even if the minimum link throughput and the total network throughput increase*. This is explained by the fact that our algorithm has a significantly higher effect on some of the links than others, which makes the throughput values of the links more disparate, which in turn leads to a lower fairness value.

Figure 15 plots the effect of our algorithm on the mean service delay over all packets transmitted: unlike the minimum throughput metric, and the fairness metric, our algorithm can be seen to reduce the packet delays under all scenarios: this is the direct consequence of the lesser interference degree which further decreases chances of packet collisions, and the number of retransmissions. In particular, we note that the mobility controlled topology experiences a 50% or higher reduction in the mean delay, even when the number of network nodes and the communication degree increase.

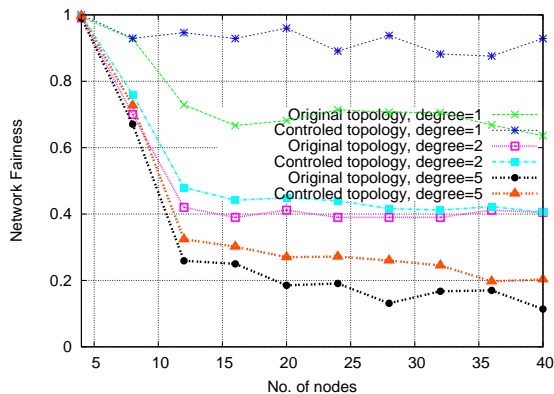


Figure 14. Network fairness for networks with different degree versus number of nodes for original generated topologies and controlled topologies

Figures 16, 17, and 18 visually capture the effect of our

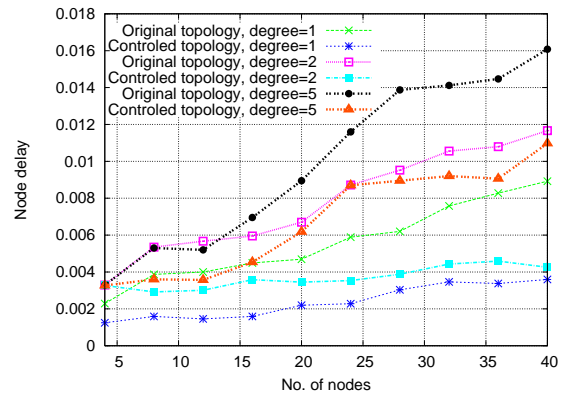


Figure 15. Packet delay for networks with different degree versus number of nodes for original generated topologies and controlled topologies

algorithm in three multi-hop scenarios. Figures 16 and 17 depict 3 multi-hop flows each of which consists of 5 nodes. The source and destination nodes of each flow are located at the end-points of the flow. In Figure 16, the distance between two successive nodes in a flow is 200m (which is close to the maximum transmission range of 250m). In Figure 17, the successive nodes in a flow are separated by distance of 100m. The three flows do not share any common node, although the central relay node for all three flows share the same location in the original topology (denoted by darker nodes and links). The two topologies mainly differ in the structure of their interference graph: in Figure 17, each node interferes with every other node in the network, while this is not true in the case of Figure 16. Figure 18 depicts two multi-hop flows: in the original topology, one of the two flows consist of three nodes with the distance between successive nodes being 200m, while the second flow consist of 7 nodes separated by 165 meters. In this scenario, the source and destination nodes for both the flows are constrained to be fixed in their given locations (and are not allowed to be relocated by our algorithm). In all the three figures, the network can be easily seen to have achieved the optimal configuration possible subject to the given set of connectivity and location constraints. Table VIII-B presents the interference degree and the throughput data for these scenarios for the original and the controlled network configurations.

Topology	Original topology		Controlled topology	
	Interference degree	Network Throughput	Interference degree	Network Throughput
1	12	1813280	2	4962240
2	13	3624640	2	5131040
3	8	2045040	4	3551680

Table I

THE MAXIMUM INTERFERENCE DEGREE AND THE NETWORK TOTAL THROUGHPUT FOR ORIGINAL AND CONTROLLED TOPOLOGIES

To summarize this section, we note that the mobility control algorithm is seen to be highly effective in traffic regimes with low communication degree: nodes experience significant reduction in their interference degree which leads to a vastly




Figure 16. Network topology #1: original topology is tightly clustered while our mobility control mechanism makes it sparse

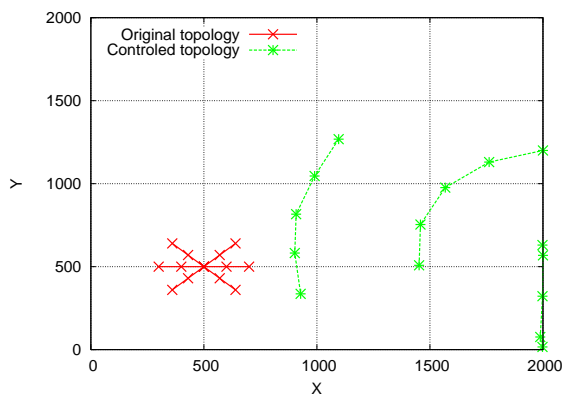


Figure 17. Network topology #2: original topology is tightly clustered while our mobility control mechanism makes it sparse

improved performance for all three throughput metrics (total throughput, minimum throughput, and fairness) as well as a significant decrease in the mean service delay of the packets. However, as expected, the additional connectivity constraints imposed on the mobility algorithm by increased communication degrees, degrades the effectiveness of the algorithm. This leads to decreased throughput and fairness related gains in the controlled topology. We believe that increased communication degree places fundamental constraints on the performance of any mobility control (or centralized network design) algorithm and needs to be rectified using orthogonal techniques (such as load balanced multi-hop routing techniques which reduces the per-node communication degree).

IX. CONCLUSIONS

As ad hoc network deployments become increasingly common and as they start assuming crucial roles in military surveillance, disaster relief, sensing, monitoring, and community Internet access through mesh networking, it becomes essential to manage them optimally to extract the best possible system performance. In this work, we explore one such significant question in this domain, namely, intelligent mobility control schemes for maximizing the throughput capacity of ad hoc wireless networks. We first rigorously establish the impact




Figure 18. Network topology #3: original topology is tightly clustered while our mobility control mechanism makes it sparse

of network topology on network throughput. We then formulate the problem of optimizing the network design to maximize network capacity, using a mixed integer quadratic program. Although the structure of this program makes it computationally infeasible to obtain an optimal solution, we derive key insights from this program to design a simple distributed mobility control algorithm which progressively takes the network to higher capacity configurations. We demonstrate through extensive packet-level simulations that the final configurations yielded by our algorithm enjoy significantly superior throughput and delay characteristics compared to the initial configurations. In the future, we intend to build upon this study by developing provably-good approximation algorithms for wireless network design, as well as exploring controlled mobility strategy for maximizing end-to-end throughput related objectives.

REFERENCES

- [1] The Network Simulator ns-2, <http://www.isi.edu/nsnam/ns/>.
- [2] Ashish Agarwal and P. R. Kumar. Capacity bounds for ad hoc and hybrid wireless networks. *SIGCOMM Comput. Commun. Rev.*, 34(3):71–81, 2004.
- [3] Mohiuddin Ahmed, Son Dao, and Randy Katz. Positioning range extension gateways in mobile ad hoc wireless networks to improve connectivity and throughput. In *Proceedings of IEEE Military Communications Conference (MILCOM 2001)*, October, 2001.
- [4] Mohiuddin Ahmed, Srikanth Krishnamurthy, Randy Katz, and Son Dao. Trajectory control of mobile gateways for range extension in ad hoc networks. *Computer Networks*, 39(6):809–825, August, 2002.
- [5] Kemal Akkaya, Mohamed Younis, and Meenakshi Bangad. Sink repositioning for enhanced performance in wireless sensor networks. *Computer Networks*, 49:512–434, 2005.
- [6] B. Badrinath, M. Srivastava, K. Mills, J. Scholtz, and K. Sollins. Special issue on smart spaces and environments. *IEEE Personal Communications*, Oct 2000.
- [7] Giuseppe Bianchi. Performance analysis of the ieee 802.11 distributed coordinated function. *IEEE Journal on Selected Areas in Communications*, 18(3), March 2000.
- [8] Giuseppe Bianchi and Ilenia Tinnirello. Kalman filter estimation of the number of competing terminals in an ieee 802.11 network. In *INFOCOM*, 2003.
- [9] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 31–42, New York, NY, USA, 2005. ACM Press.

- [10] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. SPAN: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 85–96. ACM Press, 2001.
- [11] Deborah Estrin, Ramesh Govindan, and John Heidemann. Embedding the Internet. *Communications of the ACM*, 43(5):39–41, May 2000. (special issue guest editors).
- [12] Thom Fruehwirth and Pascal Brisset. Optimal placement of base stations in wireless indoor telecommunication. *Lecture Notes in Computer Science*, 1520, 1998.
- [13] David Kiyoshi Goldenberg, Jie Lin, A. Stephen Morse, Brad E. Rosen, and Y. Richard Yang. Towards mobility as a network control primitive. In *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 163–174, New York, NY, USA, 2004. ACM Press.
- [14] Guohong Cao Guiling Wang and Tom La Porta. Movement-assisted sensor deployment. In *Proceedings of The 23rd Conference of the IEEE Communications Society (INFOCOM 2004)*, March, 2004.
- [15] Raj Jain, Dah-Ming Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. *CoRR*, cs.NI/9809099, 1998.
- [16] S. Jain, R. Rahul C. Shahv, W. Brunette, G Borriello, and S Roy. Exploiting mobility for energy efficient data collection in wireless sensor networks. *to appear in the ACM/Kluwer Mobile Networks and Applications Journal*, 2005.
- [17] Murali Kodialam and Thyaga Nandagopal. Characterizing the capacity region in multi-radio multi-channel wireless mesh networks. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 73–87, New York, NY, USA, 2005. ACM Press.
- [18] B. Liu, P. Brass, O. Dousse, P. Nain, and D. Towsley. Mobility Improves Coverage of Sensor Networks. In *Proceedings of MobiHoc'05*, Urbana-Champaign, Illinois, USA, May 2005.
- [19] Jun Luo and Jean-Pierre Hubaux. Joint mobility and routing for lifetime elongation in wireless sensor networks. In *Proceedings of The 24th Conference of the IEEE Communications Society (INFOCOM 2005)*, March, 2005.
- [20] Jun Luo, Jacques Panchard, Michal Piorkowski, Matthias Grossglauser, and Jean-Pierre Hubaux. Mobiroute: Routing towards a mobile sink for improving lifetime in sensor networks. Technical Report LCA-REPORT-2005-008, Ecole Polytechnique Federale De Lausanne (EPFL), 2005.
- [21] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.
- [22] C. Ou, K-F. Ssu, and H. Jiau. Connecting Network Partitions with Location-Assisted Forwarding Nodes in Mobile Ad Hoc Environments. In *Proceedings of the 10th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC'04)*, 2004.
- [23] Ashish Raniwala, Kartik Gopalan, and Tzi cker Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(2):50–65, 2004.
- [24] Ricardo Rodrigues, Geraldo Mateus, and Antonio Loureiro. Optimal base station placement and fixed channel assignment applied to wireless local area network projects. In *Proceedings of IEEE International Conference on Networks*, 1999.
- [25] A. Sekhar, B. Manoj, and C. Murthy. Dynamic Coverage Maintenance Algorithms for Sensor Networks with Limited Mobility. In *Third IEEE International Conference on Pervasive Computing and Communications (PERCOM'05)*, Hawaii, USA, March 2005.
- [26] Raghupathy Sivakumar, Bevan Das, and Vaduvur Bharghavan. Spine routing in ad hoc networks. *Cluster Computing*, 1(2):237–248, 1998.
- [27] Robert Szweczyk, Eric Osterweil, Joseph Polastre, Michael Hamilton, Alan Mainwaring, and Deborah Estrin. Habitat monitoring with sensor networks. *Commun. ACM*, 47(6):34–40, 2004.
- [28] T. Vercauteren, A.L. Toledo, and Xiaodong Wang. Online bayesian estimation of hidden markov models with unknown transition matrix and applications to ieee 802.11 networks. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 18–23, march 2005.
- [29] Wei Wang and Vikram Srinivasanand Kee-Chaing Chua. Using mobile relays to prolong the lifetime of wireless sensor networks. In *Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 270 – 283, 2005.
- [30] Jie Wu and Shuhui Yang. SMART: A scan-based movement-assited sensor deployment method in wireless sensor networks. In *Proceedings of The 24th Conference of the IEEE Communications Society (INFOCOM 2005)*, March, 2005.
- [31] Su Yi, Yong Pei, and Shivkumar Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 108–116, New York, NY, USA, 2003. ACM Press.
- [32] W. Zhao, M. Ammar, and E. Zegura. A Message Ferrying Approach for Data Delivery in Sparse Mobile Ad Hoc Networks. In *Proceedings of MobiHoc'04*, Roppongi, Japan, May 2004.
- [33] W. Zhao, M. Ammar, and E. Zegura. Controlling the Mobility of Multiple Data Transport Ferries in a Delay Tolerant Network Connecting Network Partitions with Location-Assisted Forwarding Nodes in Mobile Ad Hoc Environments. In *Proceedings of the 10th IEEE INFOCOM 2005*, Miami, USA, March 2005.