Performance Improvements in Ad Hoc Networks Through Mobility Groups and Channel Diversity

Invited Paper

Albert F. Harris III
Department of Computer Science
University of Illinois at Urbana–Champaign
aharris@cs.uiuc.edu

Marco Miozzo, Michele Rossi and Michele Zorzi
Department of Information Engineering
University of Padova
{marco.miozzo,rossi,zorzi}@dei.unipd.it

ABSTRACT

The physical characteristics of device mobility in many different circumstances often suggest the logical grouping of nodes that move together. The use of mobility group protocols, treating groups of devices as ad hoc networks with a node elected as the gateway, can lead to significant performance gains. However, the addition of ad hoc networks can lead to starvation and fairness problems. Multi-channel diversity schemes can be used to mitigate these effects. The main contribution of this work is the design and analysis of a combined routing group formation and multi-channel diversity algorithm. We show that, used independently, our routing group protocol and the multi-channel diversity scheme do not perform well in one or more of the metrics tested. However, when combined, the schemes complement each other, and effectively increase performance across all metrics. Our analysis is performed through extensive simulation using a recently proposed group mobility model and shows a number of new results for routing group protocols and multi-channel diversity schemes.

Categories and Subject Descriptors: C.2.1 [Network Architecture and Design]: Wireless Communication.

General Terms: Design, Performance

Keywords: Wireless network, group mobility, multi-channel

1. INTRODUCTION

The need for protocols that maximize throughput and energy efficiency continues to grow with the number of mobile computing devices with wireless networking capabilities. Mobility patterns greatly affect the performance of these protocols. There are many scenarios in which the physical mobility patterns of users allow the formation of logical groups (*e.g.*, a user carrying multiple devices). These logical groups of devices based on mobility patterns, called mobility groups, can be leveraged to build new routing protocols. Additionally, such devices are commonly equipped with radios that may provide a number of orthogonal channels on which communication can take place. This may facilitate increased throughput and fairness among nodes. Recent research has focused either on the use of novel schemes to make use of mobility groups [1, 2] or on protocols to effectively leverage multi-channel diversity within a single technology [3, 4]. The main contributions of our work deal with leveraging both routing groups and multi-channel diversity to achieve significant gains in a number of metrics, including throughput and fairness.

Physical mobility patterns that lend themselves to mobility group creation can be exploited by using short-range, low-power, low-bandwidth devices (*e.g.*, IEEE 802.15.4 (Zigbee) [5]) to create and maintain ad hoc mobility group networks and using high bandwidth interfaces (*e.g.*, IEEE 802.11 [6]) to transmit the data. These ad hoc networks can elect member nodes to act as gateways to the Internet on behalf of the rest of the mobility group [1, 2]. Consider a tour bus carrying a number of people around a city, each with a computing device with both Zigbee and IEEE 802.11 interfaces. All of the people on the bus will be moving as a group. Therefore, any information that needs to be streamed to each individual in the bus (*e.g.*, information about attractions in the city), could be sent to a single node on the bus and then travel via local links to all the members of the tour. Additionally, anyone needing to communicate with nodes outside the bus could route their traffic through the gateway node on the bus. We refer to the group of users on the bus as a routing group (RG) and the gateway node as the group leader. The Zigbee interface could be used to create the RG and handle control messages while data could be sent via the IEEE 802.11 interfaces. More generally, often times the physical movements of nodes will be correlated (*e.g.*, people moving along a sidewalk or cars moving down a highway at similar speeds). Therefore, a mechanism to build these RGs dynamically allows their use without the need of manual configuration by the users. In our previous work [1, 7], we have explored some of the advantages of RG protocols, which include increased connectivity range and throughput. In the present work, we expand this analysis to include a number of other metrics.

The introduction of ad hoc networks in the above example through the use of RGs leads to new problems, one of them being that distributed CSMA-based random access algorithms (*e.g.*, IEEE 802.11 DCF [6]) are well known to result in unfairness and flow starvation [3]. This is due to the fact that transmitters in multihop ad hoc networks are not all within range of each other and therefore have different views of the channel state. Fairness and throughput improvements can be achieved through multi-channel diversity algorithms [3]. Such algorithms exploit the use of multiple orthogonal channels to allow multiple nodes to transmit at the same time. The key challenge in using such protocols is that many of them require interfaces that have unrealistic abilities such as the capability of listening on all channels simultaneously [8, 9, 10, 11] or

WICON 2007 Ocotober 22-24, 2007 Austin, Texas USA

require clock synchronization [12, 13, 14, 15]. We use a scheme based on the Asynchronous Multi-channel Coordination Protocol (AMCP) [3] to perform asynchronous channel selection, allowing nodes to reap the benefits of multi-channel diversity using current wireless interface capabilities. We observe that our analysis complements the results in [3] as we look at the performance of the protocol under mobility by also considering the effect of routing mechanisms. We show that in some cases in the absence of RGs, multi-channel diversity may actually lead to performance degradation.

The main contribution of this work is the combination of RG algorithms with multi-channel diversity. We demonstrate, through extensive simulation, that the combination of these two techniques leads to large increases in throughput and fairness, while decreasing delay, average queue size, and energy consumption. We compare the performance of dynamic source routing (DSR) with and without RGs. For each protocol, we evaluate its performance with and without multi-channel diversity. This locates precisely the benefits of each technique. Our analysis shows that each individual technique leads to degradations in performance under certain conditions; however, these degradations are mitigated by their combined use. RGs allow significant performance gains in the face of group mobility, while multi-channel diversity provides significant gains in the face of multihop wireless networks. Our results demonstrate that in order to gain significant benefits across all of the metrics, a combination of RGs and multi-channel diversity is very effective.

The rest of this paper is as follows. Section 2 presents the state of the art in terms of mobile grouped devices and multi-channel diversity algorithms and motivates the need for a combination of these techniques. Section 3 presents our approach to combined RG and multi-channel diversity. We divide the presentation into two parts: Section 3.1 describes our RG algorithm and Section 3.2 presents the multi-channel diversity algorithm. Section 4 reports our results. It begins with a description of our methodology, including the simulator, the mobility model, and the network scenario. The section ends with the presentation of the results and a discussion of their implications. Finally, Section 5 presents some conclusions and future directions.

2. MOBILE PERFORMANCE WITH GROUP MOBILITY AND MULTI-CHANNEL DIVERSITY

With an increased awareness of the possibility of making use of group mobility patterns and the number of interfaces with multiple orthogonal channels comes the possibility of leveraging these characteristics to increase performance in terms of throughput, fairness, and energy consumption. There are two individual research lines that have been followed in the literature, one for the use of mobility groups, and one for the use of multiple channels.

The recognition of the fact that physical mobility patterns often lead to logical groups of devices led researchers to consider grouping nodes into ad hoc networks, with one or more nodes acting as a gateway to the Internet. Such group mobility solutions began with early work on personal area networks [2]. If a user had a number of devices, then it was likely that they would move as a unit. Therefore, algorithms were developed to treat the PANs as a single device as seen from the outside. In such networks, external routing was achieved by sending all data to a gateway node, or group leader. The leader could be selected according to various criteria (*e.g.*, current connectivity [2]). Methods to make sure that the PAN maintained a global address and that, as the PAN moved, packets would continue to reach group members, followed along the lines of dynamic DNS [16], NAT [17], and MobileIP [18]. Internal routing in the PANs could be achieved through the use of any suitable ad hoc routing mechanism [19, 20].

However, some logical mobility groups can be formed even if the devices are not all owned by the same user (as in the tour bus example from the introduction as well as many others). Therefore, mechanisms to automatically configure routing groups, including group leader elections (*e.g.*, via cluster-head algorithms [21]), are needed. To this end, in our previous work, we defined a novel online group formation algorithm [1]. In the present work, we perform a more complete analysis of the effects of this algorithm on several different metrics. Additionally we analyze the combination of this algorithm with a multi-channel diversity protocol. In fact, while increased connectivity is certainly a benefit of using RG solutions, decreased node fairness and overall throughput can result from the use of ad hoc networks in the RGs. Therefore, to solve this problem, we advocate the use of multi-channel diversity protocols.

Previous work has shown that multi-channel diversity has the potential both to increase throughput and fairness and to prevent starvation [8, 9, 10, 11, 12, 13, 14, 15, 22]. The proposed protocols can be divided into two types: scheduled access protocols [12, 13, 14, 15] and contention-based access protocols [3, 8, 9, 10, 11]. Typical scheduled access protocols assign nodes to a slotted periodic frame synchronized to a global clock. Control messages are sent during a control frame and are used to reserve data-frame slots which do not create conflicts. Such solutions are hard to implement in ad hoc settings however, since mobility makes it difficult for nodes to keep track of the position of the control portion of the frame [3].

Contention-based access protocols do not assume general time-slot synchronization. Some early solutions assumed the ability to receive packets on all channels simultaneously [9, 10] or used a separate transceiver for the control channel [8, 11]. MMAC [4] employs a single transceiver but uses synchronized control frames and therefore has the problems inherent to scheduled access protocols. There are a few solutions that focus on the use of a single transceiver and do not require scheduling (*e.g.*, SSCH [22] and AMCP [3]). In this work, we use a protocol similar to the very recent AMCP [3] to provide multi-channel diversity, as it can be easily implemented and provides significant performance gains. However, we perform a more complete analysis of the protocol, demonstrating that, under certain traffic loads, the protocol degrades performance if used alone. We show that these effects can be mitigated through the use of RGs.

Combining RG protocols with multi-channel diversity techniques to maximize the benefits of both is a subject that, to our knowledge, has not yet been analyzed. In this paper, we show that, due to their highly complementary nature, the use of combined RG formation and multi-channel diversity leads to significant performance gains across all of the metrics we analyze.

3. A COMBINED APPROACH

In order to exploit routing groups and the inclusion of interfaces with multi-channel capabilities to increase throughput and fairness in mobile ad hoc network scenarios, we propose a two-fold approach. First, a good RG scheme is used to exploit the inherent correlation of physical mobility patterns. Second, a multi-channel diversity protocol is used to increase the throughput and fairness of nodes within interference range of each other. In the following subsections, both our RG algorithm and the multi-channel diversity algorithm used in our experiments are described in detail.

3.1 Routing Group Scheme

Our online RG creation algorithm [1] facilitates the use of RGs based on mobility pattern similarities among different users without the need for explicit user configuration. At a high level, each node asynchronously sends periodic HELLO messages, using a low-power interface. Other nodes track these HELLO messages, using them to make RG decisions based on the perceived stability of the links.

For the algorithm, HELLO messages (*i.e.*, stability updates) are sent by each node every $T_H + \delta T$ seconds, where T_H is a configurable lapse of time, which is constant during the operation of the algorithm, and δT is a random quantity used to decrease the probability of collision and to avoid the synchronization of HELLO messages. Two parameters control the sensitivity of the scheme to changes in stability: T_{SCAN} and W. Each node updates its stability belief every T_{SCAN} seconds based on the number of HELLO messages received in the interval. W is the size of the stability measurement buffer, which controls the amount of history used in making stability calculations.

Specifically, to make a decision as to whether a device should be included in a RG, each node, i, tracks three quantities: a link stability vector $\mathbf{V}_i$, a HELLO message counter $\mathbf{R}_i^k$, where k covers the last W stability measurements, and a stability value S_i.

V_{ij}, the jth entry of vector $\mathbf{V}_i$, represents the stability belief at node i with respect to device j and is calculated as follows:

$$V_{ij} = \sum_{k=0}^{W-1} \left[\frac{R_{ij}^k}{\eta W \lfloor T_{SCAN}/T_H \rfloor} \right], \qquad (1)$$

where R_{ij}^k is the jth entry of vector $\mathbf{R}_i^k$, containing the number of HELLO messages received from node j during the kth T_{SCAN} interval. Note that $\lfloor T_{SCAN}/T_H \rfloor$ is the number of HELLO messages that a device can send in T_{SCAN} seconds. Equation (1) returns one if the number of HELLO messages received by node i from node j in the last WT_{SCAN} seconds equals $\eta W \lfloor T_{SCAN}/T_H \rfloor$. The parameter $\eta \in (0, 1]$ is used to tune the algorithm. After having calculated V_{ij} for every active neighbor j, each element of the vector $\mathbf{R}_i^k$ is deterministically shifted one position back for each k and all entries in $\mathbf{R}_i^0$ are reinitialized to zero.

S_i is node i's stability estimate and is calculated as follows:

$$S_i = \sum_{j \in \mathcal{U}_i} I\{V_{ij}\}, \qquad (2)$$

where $\mathcal{U}_i$ is the set of the neighbors of node i from which at least one HELLO message was received during the last WT_{SCAN} seconds. $I\{V_{ij}\}$ returns one if $V_{ij} \geq 1$ and zero otherwise. It can easily be seen, due to Equation (1), that S_i corresponds to the number of neighbors from which user i has received at least $100\eta\%$ of the maximum number of receivable HELLO messages during the last WT_{SCAN} seconds. Observe that in the limiting case where $\eta = 1$, S_i equals the number of neighboring devices from which user i has received 100% of the expected HELLO messages during the last WT_{SCAN} seconds. Stability values S_i are piggybacked on the HELLO messages. Upon receiving a HELLO message from a node j, the stability measure S_j therein is locally stored as $\tilde{S}_j$.

RG membership decisions are made every T_{SCAN} seconds soon after calculating the stability value S_i (see Equation (2)). A node i only takes part in an RG if its own stability value S_i is greater than or equal to some γ (*i.e.*, during the last T_{SCAN} seconds, this node has received a sufficient number of HELLO messages from at least γ other neighbors). Moreover, each node i keeps track of members of its RG via a group vector ($\mathbf{G}_i$). If $S_i \geq \gamma$, node i checks the stability value V_{ij} of each $j \in \mathcal{U}_i$: if $V_{ij} \geq 1$ and $\tilde{S}_j \geq \gamma$, node i sets $G_{ij} = W$. If either of these conditions is not met, node i updates G_{ij} according to $G_{ij} \leftarrow (G_{ij} - 1)$. If G_{ij} reaches zero, node j is removed from the RG. Note that in making RG decisions the node considers the most recently received stability estimates of its neighbors $\tilde{S}_j$. These values may be slightly stale due the inherent protocol latencies. The exact effects of the parameters on RG formation are described in detail in [1]. Furthermore, even though each node has its own view of link stabilities, group membership is consistent across nodes, since they communicate group membership status via HELLO messages.

3.2 Multi-Channel Scheme

Our multi-channel diversity algorithm was built using the Asynchronous Multi-channel Coordination Protocol (AMCP) as a base. In this section, we describe our protocol, highlighting the differences. We need a multi-channel protocol that does not assume the use of independent radios for channel negotiation, but still avoids the problems discussed in Section 2. We assume that each node's IEEE 802.11 interface can select one of N channels, of which 1 is a control channel and the other $N - 1$ are data channels. The multi-channel scheme in this paper follows the basic design of AMCP, using the same two data structures. First, each node maintains a local N-entry *Channel Table*. This structure contains an *available bit* and an *availability timer* for each data channel supported by the wireless technology. The *availability timer* indicates the amount of time a given data channel will be unavailable, based on the data transmission duration. When a node joins the network, all data channels are set to unavailable. Each node also maintains a *channel preference* variable that stores the number of the data channel that the node prefers to compete for, with zero indicating no preference.

Following AMCP, the protocol used in this paper has five steps, which occur sequentially, and one error recovery step that occurs any time there is an error (*i.e.*, a CTS timeout in Step 3 or an ACK timeout in Step 4). Each of these steps is described in detail next.

Step 1: Channel Selection. A node ready to transmit first selects a data channel for which to contend. To do this, it selects the data channel indicated by the *channel preference* variable, if it is non-zero and that channel is available, otherwise it selects an available data channel at random. If no data channel is available, the node waits until an *availability timer* expires and selects the corresponding channel.

Step 2: Channel Contention. The node inserts the selected channel from step one into the RTS packet and contends for the control channel using IEEE 802.11 DCF CSMA/CA [6]. In the CTS we include two NAV intervals. The first NAV interval expires at the end of the CTS transmission, rather than at the end of the DATA/ACK as in standard IEEE 802.11, and is used to schedule the transmissions on the control channel. If a CTS is not received, a timeout occurs and the control channel becomes free. The second NAV field in the RTS stores the time of a total transmission, including DATA/ACK, and this value is used to set the corresponding channel *availability timer*, in case the channel is successfully contended for.

Step 3: Channel negotiation. When the destination node receives an RTS packet, it checks the value of the *availability bit* corresponding to the channel being contended for. If the channel is available, the destination node responds with a confirming CTS packet containing the channel number. It then switches to that channel and waits for a DATA packet. If the channel is not available, then the destination node sends a rejecting CTS containing a list of its own available channels and remains on the control channel. If no CTS is received the node enters the **error recovery** step.

If the sending node receives a confirming CTS, it switches to the selected data channel and transmits the DATA packet. If it receives a rejecting CTS, it randomly selects a channel from the list contained in the rejecting CTS that is also marked available in its own data structure and uses that channel in another contention step.

Step 4: Data transmission. Once the destination node receives the DATA packet, it sends an ACK and switches back to the control channel. The sending node waits on the data channel for this ACK and then also transitions back to the control channel. If no ACK is received, it transitions back to the control channel and enters the **error recovery** step.

Step 5: Setting channel availability. After both sender and receiver return to the control channel, they set their *channel preference* to the data channel just used. They also reset the availability timers to their maximum values for all data channels except the one they just used, in case a contention that was won for some other channel was missed.

Error recovery step. When an error event occurs (*e.g.*, a CTS timeout), a backoff procedure is started. When the backoff timer expires, the node sets its *channel preference* variable to zero and starts availability timers for each channel.

For our experiments, we set the time to switch between channels to zero; however, this has little effect on the performance, since there are many other delays that would be larger anyway (*e.g.*, SIFS, backoff, etc.), see [3].

4. PERFORMANCE OF RGS AND MULTI-CHANNEL DIVERSITY ALGORITHMS

To demonstrate that the combined use of RGs and multi-channel diversity protocols leads to better overall performance in scenarios where group mobility is present, we performed a large number of simulation experiments. Section 4.1 describes the ANEMURAS network simulator used in this work. Section 4.2 presents our novel group mobility model. Section 4.3 presents the metrics used to quantify the performance of the protocols tested, and Section 4.4 presents the considered network scenario. Finally, Section 4.5 and Section 4.6 present and discuss our findings.

4.1 Simulator

We ran a large number of simulations using ANEMURAS [1], an event-driven network simulator for heterogeneous wireless systems designed as part of the Ambient Networks project [23]. ANEMURAS has been specifically designed to model a multi-technology wireless communication scenario, where both mobile users and fixed access points coexist and communicate through the wireless medium. Node mobility, wireless channel variability, and inter-user/inter-system interference have been explicitly modeled in the simulator.

4.1.1 Channel Model

The wireless channel is modeled accounting for path loss, shadowing, and multipath fading phenomena using their product as the link gain associated with each transmission. Path loss is implemented according to the well-known Hata model. Shadowing is modeled according to the Gudmunson model and multi-path fading is modeled for each link through a Jakes simulator with a programmable number of oscillators [24].

4.1.2 Physical Layer Model

The physical layer model takes as input the channel gain matrix created by the channel model and the transmission powers selected by each user and returns the signal to interference plus noise ratio (SINR) metric for each link. The simulator implements a physical layer module for IEEE 802.11 and IEEE 802.15.4. Errors in the transmission stream are tracked at the bit level and coding is accounted for through pre-computed coding gain curves. The receiving model for each physical layer module accounts for possible interfering transmissions during the reception of a packet; therefore, the *capture effect* is accounted for at the physical layer.

4.1.3 MAC Layer Model

The MAC layer model includes modules for both IEEE 802.11b and IEEE 802.15.4. The IEEE 802.11b module implements the distributed coordination function (DCF) algorithm and both basic access mode and RTS/CTS for collision avoidance [6], along with the multi-channel diversity protocol described in Section 3.2. The IEEE 802.15.4 module used for RG formation implements the beaconless, peer-to-peer mode [5].

4.1.4 Routing Layer

In our experiments, we compare two routing algorithms (*i.e.*, standard dynamic source routing (DSR) [19] and DSR augmented with RGs). The algorithm used to choose the group leader follows previous work by Basagni [21], who presented a distributed algorithm that partitions the nodes into clusters and uses a weight-based mechanism to choose cluster-heads. This mechanism provides sufficient flexibility in the selection criteria to allow the selection of one or more group leaders per routing group. These group leaders are responsible for route discovery; therefore, only a single route discovery is performed for all group members. This, as will be shown later, greatly improves performance. We use a practical implementation of this protocol, sending group leader selection metrics on top of our HELLO message structure. Furthermore, our implementation removes the assumption in the original work about having a perfect channel. While RG awareness could be implemented in many ad hoc routing protocols, the goal of this study is to examine the benefits of such awareness; therefore, we chose a single routing protocol to augment though the benefits for other protocols (*e.g.*, AODV [20]) can be expected to be similar.

Essentially, the DSR augmented with RGs functions as follows. The RG algorithm is run to elect group leaders. From then on, only the group leaders perform route requests and replies. When a routing group member has a packet to send, it notifies the routing group leader, who in turn performs route discovery and sends the packet on behalf of the member. RG members do not perform any path discovery algorithm, simply relying on the RG leaders for this process.

4.2 Group Mobility Model

In previous work [7], we defined a novel group mobility model, which we use here. Consider a set of nodes moving around in given area. One of these nodes is defined as the *leader* node. This node's movement can be defined according to any method (*e.g.*, random waypoint). Each of the other nodes (called *followers*) has its movement governed by two terms: drift movement and random movement. The drift movement is governed by an attraction to the leader and the random movement is superimposed on the drift movement, essentially adding some noise. We call the group of leader and follower nodes a *mobility group*.

To model group mobility, we treat each node as a particle of unit mass, allowing the amount of attraction between nodes, which causes them to move in groups, to be modeled using standard Newtonian physics. The attraction to the leader node for each follower is controlled by a *force field* defined as follows:

$$\overrightarrow{F_a}(d) = \begin{cases} \beta \, \dfrac{C_\ell C_f}{d^\alpha} \, \overrightarrow{u_a}, & d \geq d_{min} \\ \overrightarrow{F_a}(d_{min}), & d < d_{min} \end{cases}, \tag{3}$$

where $C_\ell, C_f \in \Re^+$ are the node charges, d is the distance separating the two nodes, $\overrightarrow{u}_a$ is the unit vector aligned with the axis connecting the two nodes, and $\alpha, \beta \in \Re$ are parameters used to tune the force field. α is left constant, where β is independently drawn at each time step from a Gaussian distribution with mean μ_β and variance σ_β^2. σ_β^2 is used to control the dispersion of the followers around the leader. $\overrightarrow{F}_a(d_{min})$ is used to account for a minimum distance at which to limit the strength of the attraction.

The change in velocity of a node towards the leader, $\Delta\overrightarrow{v_a}$, within an interval of time Δt is then calculated as follows:

$$\Delta\overrightarrow{v_a} = \Delta t \overrightarrow{F_a}(d). \tag{4}$$

This model is very general, with no need to define reference points for followers [7] and can also be easily modified to take into account other factors, such as the presence of obstacles, streets, etc.

4.3 Metrics

To analyze the performance gains due to the combined use of our routing group algorithm and multi-channel diversity, we evaluate a number of different metrics. First, we evaluate the effect of the combined protocol on delay and queue length. As specified later in Section 4.4, we consider an uplink communication scenario, where all mobile users are interested in transmitting their data to a static, randomly placed, access point. The delay is defined as the time needed for the data packets to reach the access point through (possibly) multi-hop communication. Delay is critical for a number of real-time applications, including voice applications. Second, we analyze the throughput, another important metric for several multimedia applications. Third, we analyze network fairness. For this we use the well-known Jain's fairness index [24], which is defined as follows:

$$\text{fairness} = \left(\sum_{i=0}^{n} x_i\right)^2 \Big/ \left(n \sum_{i=0}^{n} x_i^2\right), \tag{5}$$

where i is the flow index, n is the total number of flows in the network, and $x_i = A_i/O_i$ represents the actual allocation (A_i) over the optimal allocation (O_i) for flow i. This fairness index has values between zero and one, one being perfect fairness, and is not overly sensitive to atypical network flow patterns. The final metric examined is the energy goodput, defined as the energy consumed per good packet received.

4.4 Network Scenario

All simulations occur on a 200 m by 200 m square, with one randomly placed static access point and 30 mobile nodes. Twenty of these mobile nodes are equally distributed among four RGs, which move according to the group mobility model described in Section 4.2. The remaining ten nodes are independent nodes, not belonging to any RG. Each node is equipped with an IEEE 802.11b radio and an IEEE 802.15.4 radio. The former radio is used to deal with all data traffic and the latter is used only for the RG formation algorithm (*i.e.*, the transmission of the HELLO messages).

Data flows are sent using different power levels for intra-group and inter-group communication, using the lowest power possible while still maintaining reliability for intra-group communication. This has the effect of limiting the interference due to intra-group communication. For inter-group communication we consider standard transmission power levels, as specified at the end of the Section.

This scenario allows us to extensively test the protocols over a range of network loads and situations that present the possibility of forming mobility groups. For the simulations, we model pedestrian behavior for node mobility, setting user speeds in the range of 0.5 m/s to 1.2 m/s. The attraction within groups is set so that the average distance between any two users within the same group is around 15 m (*i.e.*, the coverage range used by IEEE 802.15.4 devices [5]). To this end, the mobility model parameters are set as follows: average speed $= 2$ m/s, $\alpha = -1$, $C_\ell = C_f = 0.8$. The RG algorithm used $W = 4$, $\eta = 0.7$, $T_H = 12$ seconds and $T_{SCAN} = 30$ seconds. These parameters were found to maximize group stability for pedestrian traffic. Finally, the IEEE 802.11 interface has a transmit power of 2.1 W to 2.2 W and a receive power of 1.3 W, both based on the Cisco Aironet 350 series [25]. We also use a constant packet size of 512 bytes.

We tested our protocols with a number of different scenarios, but due to space limitations, only this scenario is presented. Results from other scenarios led to the same conclusions. We chose to present this scenario because it most accurately models airports, museums, and other places that our solution targets. This simulation set up is large enough that single hop communication with an IEEE 802.11 access point is not always possible and nodes may have to route through routing groups. The number of users was chosen to be close to the saturation point of the network so that the various protocols would be stressed. A number of independent users are simulated in the scenario to provide background contention for the routing groups. Additionally, we present a scenario with multiple independent routing groups to show they can coexist effectively in the same region.

4.5 Simulation Results

We use DSR [19] as our base routing algorithm, and a version of it augmented with the RG algorithm (denoted by RG in the following). In each of the graphs in this section, we use the following convention: *protocol*(i,j), where *protocol* is either DSR or RG, i is the number of nodes actively transmitting data in each mobility group, and j is the number of data channels used by the multi-channel diversity algorithm (when $j = 1$ there is no multi-channel algorithm running). We tested DSR with and without the multi-channel algorithm and RG with and without the multi-channel algorithm. Each figure presents the average of 200 simulation rounds each lasting 600 seconds, which has been verified to yield tight confidence intervals. We present each metric as the traffic generation rate of the active nodes in the groups varies for different numbers of active nodes. This allows the traffic load of the network to be varied along two dimensions. The greater the number of active nodes, the more contention between users, the greater the traffic generation rate, and the greater the per-flow demand. Each metric is then averaged across the active nodes.

Delay is a critical metric for any real-time application (*e.g.,* VoIP). The use of routing groups, even without multi-channel diversity, significantly reduces the average delay experienced by nodes (see Figures 1 to 3). This is due to the following two main reasons. First, the use of RGs allows for a reduction in contention for the base station. Second, as RG leaders are in charge of transmitting route discovery messages for all group members, the number of control packets is substantially reduced with respect to the standard DSR case. A reduction in the number of route discoveries and route replies means an increased capacity which positively influences all performance metrics. Multi-channel diversity, on the

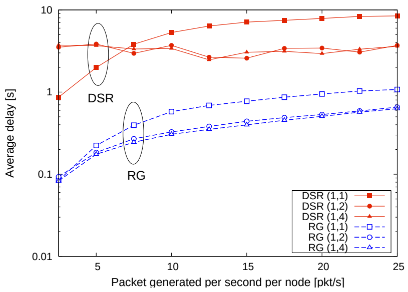

Figure 1: Delay vs. traffic rate (one active transmitter per group)

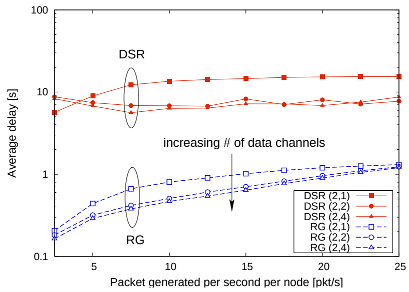

Figure 2: Delay vs. traffic rate (two active transmitters per group)

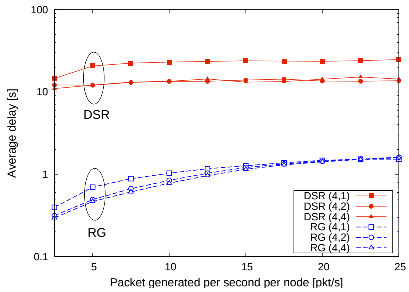

Figure 3: Delay vs. traffic rate (four active transmitters per group)

other hand, increases the delay experienced for low packet generation rates. This is due to the overhead for the availability timers not being justified by the network load. However, when routing groups are combined with multi-channel diversity, this effect is mitigated and the use of extra channels always results in lower delays.

A metric closely related to delay is the average queue size: larger queues imply longer delays (see Figures 4 to 6). The first phenomenon to note is that, for DSR, the addition of multiple channels at low packet generation rates for one, two, and four active senders performs worse than the case without multiple channels. This is due to the need for nodes to wait for the availability timers to expire even when a channel is free if they did not hear the contention for that channel, which is an introduced overhead without a benefit outweighing it (corresponding to the increase in delay in the previous figures). However, once traffic rates saturate the one channel case, the multi-channel algorithm begins to outperform the single channel case. However, when the RG protocol is added, this effect is effectively mitigated.

For our RG protocol with a single channel, the queues grow dramatically as the traffic rate increases. This is primarily due to the group leader becoming a bottleneck. However, when combined with multi-channel diversity, the queue sizes are kept very low and in all cases lower than in DSR with or without multi-channel diversity. In Figure 6, at packet generation rates of greater than 20 packets per second per node, RG with multi-channel diversity is outperformed by DSR. This is due to the fact that the group leader has become a bottleneck. This suggests that for high network traffic, RG should elect multiple group leaders, which would solve the problem while adding only minimal complexity to the algorithm.

For multimedia and bulk data transfer, the average throughput per node, in terms of packets per second, is a critical metric (see Figures 7 to 9). In this case, for DSR, extra channels always increase the average throughput; however, for RGs, the throughput performance with a single channel outperforms the case with multiple channels, until the group leader becomes a bottleneck. The bottleneck effect is mitigated by the addition of multi-channel diversity until very large network loads (see Figure 9), where DSR with multi-channel diversity overtakes RG with multiple channels at a packet generation rate of about 17 packets per second.

Jain's fairness index will show whether certain flows are getting starved (see Figures 10 to 12). DSR's fairness stays roughly constant as traffic increases; however, the fairness is rather low, again, with the single channel case outperforming the multi-channel cases. This is not intuitive, since one of the reasons to use multi-channel protocols is to avoid starvation. Our results show that while no nodes starve, the average fairness is slightly reduced by the use of multi-channel diversity alone. This is a new observation about the multi-channel diversity scheme, resulting from the overhead of the availability channels forcing nodes into constantly contending for their preferred channels. However, when routing groups are leveraged, multi-channel diversity then increases fairness. Furthermore, multi-channel diversity mitigates the effect of the group leader becoming a bottleneck and decreasing fairness.

To evaluate the energy goodput, we consider the total energy consumed by the network, including idle energy, receive, and transmit, for both data and control messages. Lower energy goodput is desirable as it shows that more data can be successfully transmitted for some fixed amount of energy. This is a critical metric for any mobile system as battery capacity is still a constraint (see Figures 13 to 15). In general, multi-channel diversity improves the energy goodput for both DSR and RG. This is because nodes do not have to contend for the same data channel and therefore have few backoff periods. As in the previous results, however, the performance of RG degrades when the group leader becomes the bottleneck, in the worst case (see Figure 15) being outperformed by DSR with multi-channel diversity.

4.6 Discussion of the Results

The previous section's results all demonstrate that, in our scenario, a combined routing group, multi-channel diversity approach yields significant improvements over the performance of either of these two techniques in isolation. In fact, the two techniques prove to be very complementary in their effects.

Our RG algorithm is very effective in decreasing delay, while further delay gains are achieved through the use of multi-channel diversity. However, queue lengths are greatly reduced through the use of multi-channel diversity. Significant throughput gains can be achieved through the use of RGs; however, these throughput gains decrease rapidly after the group leader becomes a bottleneck. Fairness also decreases for this same reason. The use of multi-channel diversity can significantly move this point to larger traffic loads. Furthermore, the group leader bottleneck problem can be easily addressed by either limiting the size of the RG (*e.g.*, where size is measured in terms of traffic generation), or using a traffic rate threshold to trigger the election of new group leaders for a single RG.

In terms of energy goodput, multi-channel diversity greatly improves performance in the absence of our RG algorithm. However, the RG algorithm also improves the energy goodput, and the combination of the two performs the best.

Finally, we note that significant improvements are gained through the use of only two data channels. This is an important result since in this case the protocols presented in this work could be implemented in the current IEEE 802.11 frequency allocation.

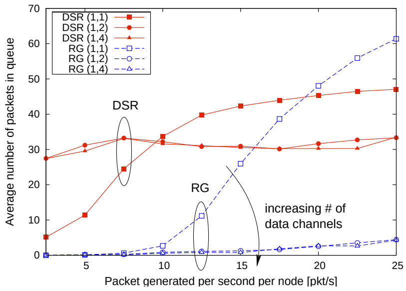

Figure 4: Queue size vs. traffic rate (one active transmitter per group)

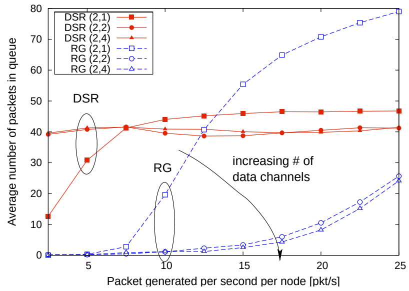

Figure 5: Queue size vs. traffic rate (two active transmitters per group)

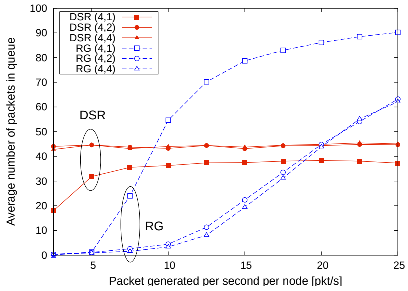

Figure 6: Queue size vs. traffic rate (four active transmitters per group)

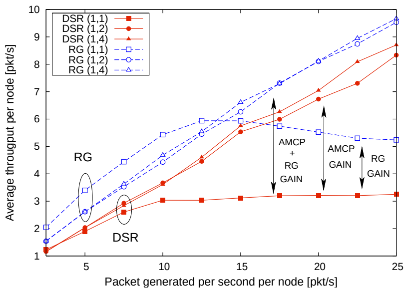

Figure 7: Throughput vs. traffic rate (one active transmitter per group)

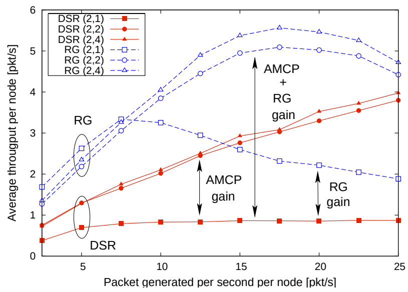

Figure 8: Throughput vs. traffic rate (two active transmitters per group)

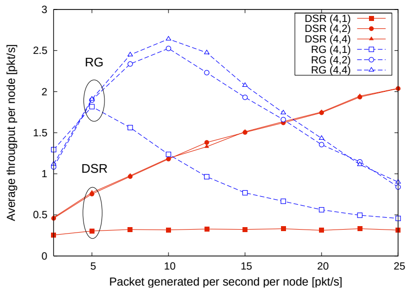

Figure 9: Throughput vs. traffic rate (four active transmitters per group)

5. CONCLUSION AND FUTURE DIRECTIONS

This paper has explored the use of combined routing group algorithms and multi-channel diversity schemes. We present our novel routing group formation and maintenance algorithm, which exploits multi-radio capabilities to efficiently manage dynamic, online routing group formation. As the main contribution of our work, we demonstrate that in scenarios where the characteristics of the physical mobility patterns lend themselves to group formation, the use of combined multi-channel diversity schemes and routing group algorithms yields significant performance improvements in terms of delay, throughput, fairness, and energy efficiency. We demonstrate this through an extensive set of simulations, using a recent model for group mobility. Our results also show that using only three orthogonal channels, significant gains are achieved; therefore, current IEEE 802.11 technology provides a sufficient number of orthogonal channels for our protocols.

Future work in this area includes examining the inclusion of other technologies, such as UMTS, and analyzing their effects on the performance of our combined algorithms. A hierarchical approach, including the PANs as nodes within routing groups, is also another interesting direction. In both cases, we expect the same type of gains as seen in this work. Furthermore, adding more advanced features to the multi-channel diversity algorithm could significantly improve performance as well.

6. REFERENCES

[1] L. Badia, N. Bui, M. Miozzo, M. Rossi, and M. Zorzi, "On the Exploitation of User Aggregation Strategies in Heterogeneous Wireless Networks," in *IEEE CAMAD*, Trento, Italy, June 2006.

[2] R. Kravets, C. Carter, and L. Magalhaes, "A Cooperative Approach to User Mobility," *ACM Computer Communications Review*, vol. 31, no. 5, pp. 57–69, October 2001.

[3] J. Shi, T. Salonidis, and E. Knightly, "Starvation Mitigation Through Multi-Channel Coordination in CSMA Multi-hop Wireless Networks," in *ACM MobiHoc*, Florence, Italy, May 2006.

[4] J. So and N. Vaidya, "Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Cahnnel Hidden Terminals Using A Single Transceiver," in *ACM MobiHoc*, Roppongi Hills, Tokyo, Japan, May 2004.

[5] IEEE 802 LAN/MAN Standards Committee, "Wireles Medium Access and Physical Layer Specification for Low Rate Wireless Personal Area Networks," IEEE Standard 802.15.4, 2006.

[6] ——, "Wireless LAN Medium Access Control MAC and Physical Layer (PHY) Specifications," IEEE Standard 802.11, 1999.

[7] M. Rossi, L. Badia, N. Bui, and M. Zorzi, "On Group Mobility Patterns and Their Explotation to Logically Aggregate Terminals in Wireless Networks," in *IEEE VTC Fall*, Dallas, TX, US, September 2005.

[8] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A Multi-Radio Unification Protocol For IEEE 802.11 Wireless Networks," in *IEEE BROADNETS*, San José, CA, US, October 2004.

[9] N. Jain and S. Das, "A Multichannel CSMA MAC Protocol with Receiver-based Channel Selection for Multihop Wireless Networks," in *IEEE ICCCN*, Scottsdale AZ, US, October 2001.

[10] A. Nasipuri and S. Das, "A Multichannel CSMA MAC Protocol for Multihop Wireless Networks," in *IEEE WCNC*, New Orleans, LA, US, September 1999.

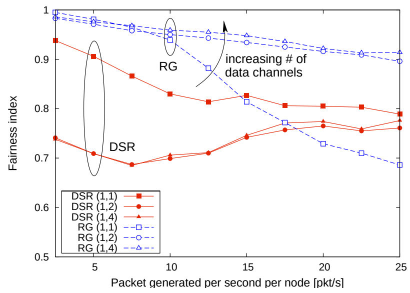

Figure 10: Fairness vs. traffic rate (one active transmitter per group)

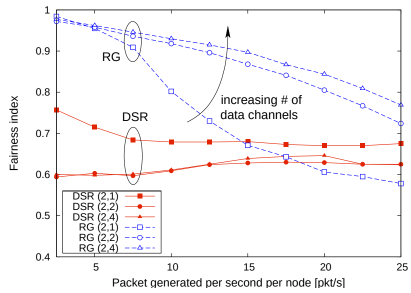

Figure 11: Fairness vs. traffic rate (two active transmitters per group)

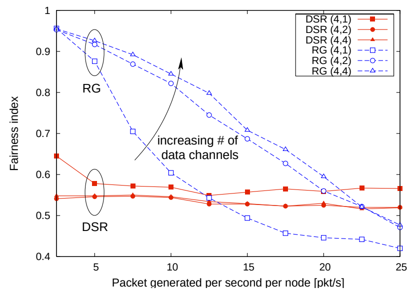

Figure 12: Fairness vs. traffic rate (four active transmitters per group)

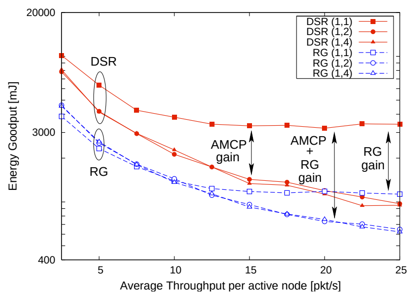

Figure 13: Energy consumption vs. traffic rate (one active transmitter per group)

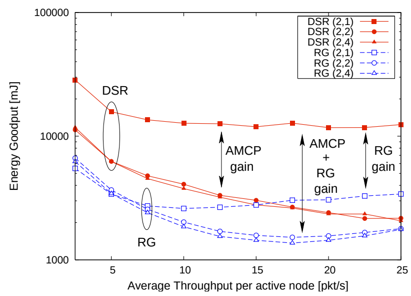

Figure 14: Energy consumption vs. traffic rate (two active transmitters per group)

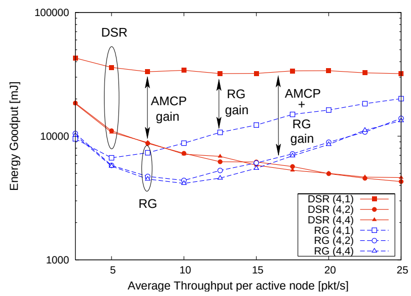

Figure 15: Energy consumption vs. traffic rate (four active transmitters per group)

[11] S. Wu, C. Lin, Y. Tseng, and J. Sheu, "A New Multi-Channel MAC Protocol with On-Demand Channel Assignment for Mobile Ad Hoc Networks," in *IEEE I-SPAN*, Dallas/Richardson, TX, US, December 2000.

[12] I. Chlamtac and A. Lerner, "Fair Algorithms for Maximal Link Activation in Multihop Radio Networks," *IEEE Transactions on Communications*, vol. 35, no. 7, pp. 739–746, July 1987.

[13] A. Ephremides and T. Truong, "Scheduling Broadcasts in Multihop Radio Networks," *IEEE Transactions on Communications*, vol. 38, no. 4, pp. 456–460, April 1990.

[14] C. Lin, "On-demand QoS Routing in Multihop Mobile Networks," in *IEEE INFOCOM*, Anchorage, Alaska, April 2001.

[15] C. Zhu and M. Corson, "QoS Routing for Mobile ad hoc Networks," in *IEEE INFOCOM*, New York, NY, US, June 2002.

[16] P. Vixie, S. Thomson, Y. Rekhter, and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)," Network Working Group, RFC 2136, April 1997.

[17] P. Srisuresh and M. Holdrege, "The IP Network Address Translator (NAT) Terminology and Considerations," Network Working Group, RFC 2663, August 1999.

[18] C. Perkins, "IP Mobility Support for IPv4," Network Working Group, RFC 3344, August 2002.

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

[20] C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *IEEE WMCSA*, New Orleans, LA, US, February 1999.

[21] S. Basagni, "Distributed and Mobility-adaptive Clustering for Multimedia Support in Multi-hop Wireless Networks," in *IEEE VTC*, Amsterdam, The Netherlands, September 1999.

[22] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted Seeded Channel Hopping for Capacity Improvements in IEEE 802.11 Ad-hoc Wireless Networks," in *ACM MobiCom*, Philadelphia, PA, US, Sep.-Oct. 2004.

[23] The Ambient Networks Project, http://www.ambient-networks.org/.

[24] B. Sklar, *Digital Communications: Fundamentals and Applications*, 2nd ed. Prentice Hall PTR, 2001.

[25] Cisco, "Cisco Aironet 350 Client Data Sheet," http://www.cisco.com/.