ICAM: Integrated Cellular and Ad-Hoc Multicast

Haiyun Luo, Li Erran Li, Ram Ramjee, Randeep Bhatia, Paul Sanjoy

June 2004
UCLA-CSD-TR040026

Computer Science Department
University of California, Los Angeles
Los Angeles, CA 90095-1596

Abstract

In third generation (3G) wireless data networks, multicast throughput decreases with the increase in multicast group size, since a conservative strategy for the base station is to use the *lowest* data rate of all the receivers so that the receiver with the *worst* downlink channel condition can decode the transmission correctly. This paper proposes ICAM, Integrated Cellular and Ad-Hoc Multicast, to increase 3G multicast throughput through opportunistic use of ad-hoc relays. In ICAM, 3G base station delivers packets to *proxy* mobile devices with better 3G channel quality. The proxy then forwards the packets to the receivers through an IEEE 802.11-based ad-hoc network. In this paper, we first propose a localized greedy algorithm that discovers for each multicast receiver the proxy with the highest 3G downlink channel rate. We discover that due to capacity limitations and interference of the ad-hoc relay network, maximizing the 3G downlink data rate of each multicast receiver's proxy does not lead to maximum throughput for the multicast group. We then show that the optimal ICAM problem is NP-hard, and derive a polynomial-time 4-approximation algorithm for the construction of the multicast forest. This bound holds when the underlying wireless MAC supports broadcast or unicast, single rate or multiple rates ($4(1+\epsilon)$ approximation scheme for the latter), and even when there are multiple simultaneous multicast sessions. Through both analysis and simulations we show that our algorithms achieve throughput gains up to 840% for 3G downlink multicast with modest overhead on the 3G uplink.

1 Introduction

Third-generation (3G) CDMA wide-area wireless networks have experienced significant growth recently. As of Feb 29, 2004, the number of CDMA20001X subscribers worldwide has increased by more than 100% last year and exceeded 85 million. At the same time the number of subscribers of CDMA20001xEV-DO, also known as HDR (High Data Rate), has exceeded 5.9 million [1]. As the user population base builds up, group communications such as on-demand video streaming, group messaging and gaming through hand-held wireless devices have been spurring the development of multicast functions in the 3G network infrastructure. 3G standard bodies 3GPP and 3GPP2 have been actively standardizing multicast services [25, 26, 27].

Existing multicast strategy in 3G networks suffers in terms of decreased downlink channel utilization and throughput as the size of the multicast group increases. In order for the multicast receiver with the *worst* downlink channel condition to correctly decode data frames from the 3G downlink transmission, a conservative strategy for the 3G base station is to use the *lowest* data rate among all the receivers in the multicast group. Due to path loss and fast fading characteristics of the wireless medium, the likelihood that at least one multicast receiver experiences bad downlink channel condition increases as the multicast group size increases, resulting in decreased downlink channel utilization and throughput. For example, as we will see in Section 3, our simulation results indicate that although the average downlink data rate is as high as 600Kbps for a single client, the throughput for a multicast group of five users decreases to around 80Kbps and drops close to the lowest achievable rate of 38.4Kbps with ten or more users. This phenomenon is in stark contrast to the scenarios of unicast traffic where increasing the number of users *increases* the downlink channel utilization using Proportional Fairness Scheduling [3]. While it is possible to use sophisticated coding such as Reed-Solomon codes combined with Turbo codes to support an increased data rate of about 200Kbps for broadcast services [4], it still results in significantly lower throughput compared to the unicast case.

One approach to increasing 3G throughput is through the use of ad-hoc relays. In this model, mobile devices are assumed to have both 3G and IEEE 802.11 interfaces. The mobile receiver first discovers a *proxy* client (e.g., another mobile device located in the same cell) with superior downlink channel condition and higher data rate connection with the 3G base station. On behalf of the receiver, the proxy client receives data packets from the base station at higher data rate. The proxy then forwards the packets using the 802.11-based ad-hoc network to the mobile receiver. While this model has been shown to significantly improve 3G cell throughput for *unicast* traffic in the unified cellular and ad-hoc network (UCAN) architecture [19], extending this model for *multicast* traffic is not trivial since multicast traffic can easily overload an IEEE 802.11-based

ad-hoc network, limiting the achievable throughput gains. We refer to our problem as *ICAM* for Integrated Cellular and Ad-Hoc Multicast.

Thus, *in order to maximize throughput in ICAM, it is not sufficient to choose the best proxy with the highest data rate connection. The data rate of the proxy must be balanced by the data rate achievable over the interference-prone multicast ad-hoc network.* Therefore, it is important that the choice of the proxies and the algorithm for the construction of the multicast forest be performed jointly with the algorithm explicitly aware of capacity limitations and interference possibilities in the ad-hoc network.

To this end, we characterize the interferences of multihop wireless network with a general model that uses a graph-theoretic representation called the *interference graph*. With the interference graph and the network topology around each multicast receiver, we derive a polynomial-time 4 approximation algorithm for the construction of the optimal throughput multicast forest, with consideration of both the proxies' downlink data rates and potential bottlenecks along the ad-hoc relay paths.

While the near optimal algorithm achieves high throughput, it also incurs high overhead in the presence of users with high mobility since the topological information of the mobile nodes have to be updated constantly to the base station. Therefore, we also propose a greedy algorithm that computes the best proxy for each receiver and attempts to merge multiple receivers to proxies, where possible, to limit the impact of interference. We leverage the 3G base station's explicit coordination and the omni-present 3G downlink/uplink to achieve ad-hoc relay network efficiency and reliability in the presence of ad-hoc network topological changes due to node mobility.

The contribution of this paper is three-fold. First, to the best of our knowledge, we are the first to propose multicast ad-hoc relay protocols to improve the throughput efficiency of 3G multicast. Second, we propose a polynomial-time approximation algorithm that outputs near optimal multicast relay strategy. Our algorithm is based on a very general interference model and has an approximation factor of four. To our knowledge, this is also the first multicast routing design that explicitly considers multihop wireless interferences in ad-hoc networks. We assume a very general network connectivity model and do not make the assumption that the ad-hoc network is a unit disk graph. Our algorithm and its bounds are equally applicable when the underlying wireless media access control supports broadcast or unicast, single rate or multiple rates ($4(1+\epsilon)$ approximation scheme for the latter) and even when there are multiple simultaneous multicast sessions. We also propose a greedy algorithm that discovers proxies and constructs an ad-hoc multicast relay network to support efficient 3G multicast with highly mobile nodes. Finally, we evaluate the performance of the near optimal and greedy algorithms using extensive simulations,

showing throughput gains of up to 840%.

The rest of the paper is organized as follows. In Section 2, we review related work. In Section 3, we present the motivation for ad-hoc relaying in order to improve the efficiency of the 3G multicast problem. In Section 4, we present the models and assumptions for ICAM. In Section 5, we present the greedy algorithm. In Section 6, we propose an algorithm that achieves near optimal end-to-end throughput. In Section 7, we present extensive simulation results of our implementation of the multicast relay protocols. We discuss related issues in Section 8. Finally Section 9 concludes this paper.

2 Background and Related Work

In this section we briefly review the hybrid wireless network architectures to which our work is applicable. We then survey the related work on multicast in ad hoc networks.

2.1 Hybrid Wireless Network Architectures

Various architectures that combine 3G networks with ad hoc networks have been proposed. We first discuss 3G networks and then ad hoc networks.

The 3G network that our work is applicable is 1xEV-DO (Evolution-Data Only), also known as HDR (High Data Rate). HDR is an integral part of the CDMA2000 family of 3G standards. Designed for bursty packet data applications, it provides a peak data rate of 2.4Mbps and an average data rate of 600Kbps within one 1.25MHz CDMA carrier. Users share the HDR downlink using time multiplexing with time slots of 1.67ms each. At any time instant, data frames are transmitted to one specific client, and the data rate is determined by the client's channel condition. While HDR has the potential to provide "anywhere" "always-on" wide-area wireless Internet access, its peak downlink data rate of 2.4Mbps is relatively low compared with the data rate of 11Mbps of IEEE 802.11b links.

The ad hoc networks can be formed by infrastructure relays or mobile devices of wireless subscribers. For example, in the iCAR [29] architecture, relay nodes are stationary special-purpose devices. Infrastructure relays can also be mobile, e.g. devices mounted on top of bus, cabs, etc. In the UCAN [19], relays are mobile devices themselves. UCAN requires that each mobile device be equipped with these two types of wireless interfaces. Fortunately, given the popularity of the IEEE 802.11b (Wi-Fi) interface, it is already being embedded in every mobile device and thus the device only needs a 3G HDR interface card to operate in UCAN. It can be a portable computer with both 3G wireless modem and IEEE 802.11b PCMCIA card, or a PDA

Figure 1: UCAN Architecture

with both interfaces integrated in a single card [2].

The UCAN architecture is based on the key idea of opportunistic use of the license-free IEEE 802.11 interfaces to improve the proprietary 3G cell throughput. Figure 1 shows the UCAN network architecture. For those mobile devices associated with the HDR base station, some of them may be actively receiving data packets from the Internet via the HDR downlink, while others may have their HDR interfaces in the dormant mode. Associated clients monitor the pilot bursts of the HDR downlink to estimate their current downlink channel conditions. At the same time, these devices turn on their IEEE 802.11b interfaces in ad-hoc mode, and run UCAN protocols. If a destination client experiences low HDR downlink channel rate (e.g., 38.6Kbps), instead of transmitting directly to the destination, the HDR base station transmits the data frames to another client (proxy client) with a better channel rate (up to 2.4Mbps). These frames are further relayed through IP tunneling via intermediate relay clients to the destination, using the high-bandwidth IEEE 802.11b links.

2.2 Related Work

Multicast over ad-hoc networks has been intensively studied in recent years. The proposed multicast routing protocols can be classified into two categories. One category is tree-based, including Reservation-Based Multicast (RBM) [10], Lightweight Adaptive Multicast (LAM) [16], Ad hoc Multicast Routing Protocol (AMRoute) [8], Ad hoc Multicast Routing protocol utilizing Increasing id-numberS (AMIS) [28], and multicast extension of Ad hoc On-demand Distance Vector (MAODV) [22]. They all build a shared or core-based tree to deliver multicast data, but differ in detailed mechanisms for tree construction, maintenance, and adaptation to the network topological dynamics. The other category is mesh-based, including Core-Assisted Mesh Protocol (CAMP) [11], and On-demand Multicast Routing Protocol (ODMRP) [17]. They enhance the connectivity by building a mesh with multiple forwarding paths, therefore improve the resilience

as the network topology changes.

However, all the aforementioned multicast protocols focus on addressing node mobility induced topological changes, and none of them explicitly considers the multihop wireless channel interferences. Therefore, they suffer from traffic concentration and link-layer contention, especially when the multicast group size is large. Besides, it is unclear how they perform compared to the optimal case. In this paper, we quantify the impact of channel interferences on the end-to-end throughput. Our design carefully engineers the distribution of the multicast relay traffic to avoid hot-spot congestion. Besides, by leveraging the explicit coordination function of the 3G HDR base station and the omni-present uplink/downlink, the implementation of our optimal multicast algorithms is significantly simpler and more reliable even in the presence of high topological dynamics.

3 Motivation

With increasing use of high bandwidth data applications in 3G wireless networks, especially with large number of users receiving the same high data rate services, efficient information distribution is essential. 3G standard bodies recognize this need and has been standardizing multicast and broadcast services [26, 25]. The standards describe how to set up radio resources to enable point to multipoint transmission. The standards also describe other related functionality such as protocols to setup multicast in the current 3G architecture, protocols that enable charging and security of multicast sessions.

However, 3G multicast is inherently limited by the worst channel rate among group members. To state it more precisely, assume that there are n clients covered by a 3G base station. Let l of these clients (denoted by the set R) belong to a multicast group, henceforth called *multicast receivers*. If multicast receiver $v \in R$ has an instantaneous downlink channel rate of r_v^i at time slot i, then the data rate for the multicast at time slot i is

$$\min_{v \in R} r_v^i$$

and the average throughput between time slots i and j, $j \geq i$ is

$$\sum_{t=i}^{j} \left(\min_{v \in R} r_v^t \right) / (j - i + 1)$$

Due to path loss and fast fading characteristics of the wireless medium, the likelihood that at least one multicast receiver experiences bad downlink channel condition increases as the multicast group size increases. Therefore, increasing the number of receivers results in lower average multicast throughput. To quantify this effect, we simulated a network setting where all multicast

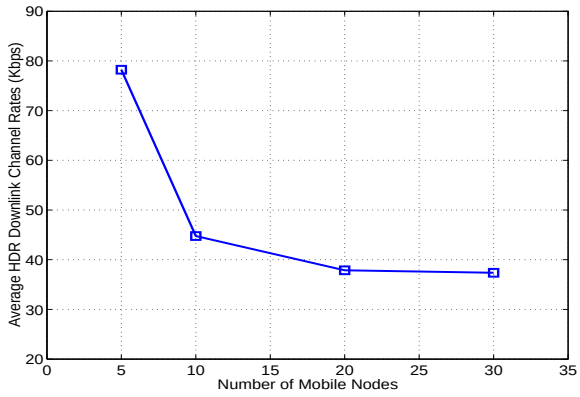

Figure 2: Multicast Throughput

receivers are static and randomly distributed over a 600×600m^2 HDR cell. The HDR downlink channel included both the slow fading (depends on user location) and fast fading components. As we can see from Figure 2 the average throughput decreases dramatically as the size of the multicast group grows. Although the average HDR downlink channel rate is as high as 600Kbps for a single client, the throughput for a multicast group of five users decreases to around 80Kbps and drops close to the lowest achievable rate of 38.4Kbps with ten or more users. While sophisticated coding can improve this throughput to about 200Kbps [4], it still falls significantly short of the achievable unicast throughput; this is discussed more in Section 8.

The inefficiency of 3G multicast motivates us to use relays to improve its throughput. Specifically, for each multicast receiver v with low average downlink channel rate r_v, find a proxy client with higher average downlink channel rate $p(v)$ and an ad-hoc relay route from the proxy to the receiver. We present a more precise definition of the problem that includes the capacity constraints and interference issues in the ad-hoc network in Section 6.1.

4 Models and Assumptions

We use the same model as that in [19] for a mobile user's HDR downlink channel condition with both slow and fast fading considered. Slow fading is modeled as a function of the client's distance from the HDR base station. Fast fading is modeled by Jakes' Rayleigh fading [15]. The combined E_c/N_t for both slow and fast fading is then mapped to a table of supported data rate with 1% error [6]. Figure 3 presents a snapshot of HDR downlink instantaneous channel rates, and the average rate over a long time period for clients with different distances from the base station.

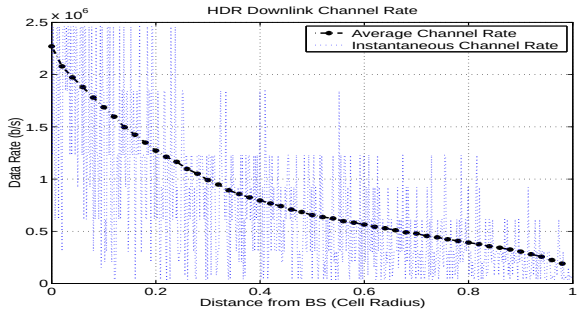

Figure 3: HDR Downlink Instantaneous and Average Channel Rate

4.1 IEEE 802.11 Unicast vs Broadcast

The media access control of IEEE 802.11 [13] based ad-hoc networks defines two options to access the wireless media: unicast and broadcast. Unicast is intended for point-to-point transmission. Besides carrier sensing, RTS/CTS/ACK are used to protect the data transmission. Broadcast can reach all nodes that are within distance R_t away. However, its reliability is lower than that of unicast since RTS/CTS/ACK is not used. Also, due to the lack of ACK messaging, IEEE 802.11 broadcast provides no per-hop reliability. The lack of ACK messaging also makes it necessary to introduce heart-beat messages to detect link breakages due to mobility, increasing the overhead. Furthermore, without the RTS/CTS handshake broadcast can only uses the data rates that are defined in the basic set (i.e., 1Mbps and 2Mbps), while the data rate of unicast can be up to 11Mbps. Although unicast introduces extra control overhead (RTS/CTS/ACK), Figure 4 shows that using 11Mbps unicast the throughputs of CBR/UDP traffic are still 100% higher than the throughputs using 2Mbps broadcast. Thus, though our proposed algorithms work with both unicast and broadcast transmissions, but we use unicast for relaying 802.11 traffic in our simulation evaluation.

4.2 Interference model

We assume a general proximity-based interference model for the IEEE 802.11 based multihop wireless network. In this model, the transmission of a node u does not cause interference to the transmission of a node x, if their distance is greater than R_I where R_I is the maximal interference range. We assume that R_I is $q \times R_t$ where R_t is the transmission range of each node and $q \geq 1$. Our interference model subsumes the protocol model and physical model [14].

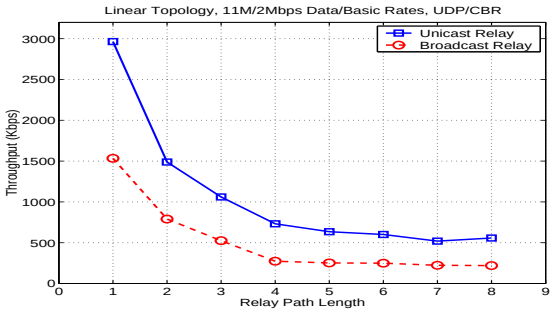

Figure 4: The decrease of throughput with respect to hop length

4.3 Hop limit for proxy discovery

We assume that the distance between a proxy and any of its receiver is upper bounded by a small number h. That is the proxy of a given receiver could be 0 to h hops away for some small value h. There are several reasons for h being small. First, due to interferences, wireless channel error, and lack of scheduling, IEEE 802.11 ad-hoc network throughput decreases very fast as the number of hops increase. (See Figure 4). Similar observations have also been presented in the study of the capacity of IEEE 802.11 ad-hoc networks in [18]. Thus, if the ad-hoc throughput decreases to the minimal HDR downlink channel rate in less than h number of hops, then any proxy of a multicast receiver has to be within h hops; otherwise the 802.11 ad-hoc network would become the bottleneck contradicting the need for relays to increase 3G multicast traffic. Second, paths of length exceeding a certain number of hops are not desirable because of the increased probability of route breakages due to mobility, higher latencies, high overhead of proxy discovery and high routing update overheads over the HDR uplink. Our simulation study shows that, for a 500m radius cell, using a proxy beyond a $h = 4$ hop neighborhood of the multicast receiver does not result in increasing gains and in fact, results in slight decrease in gains in some cases.

4.4 Minimal Separation and Location

Our near optimal algorithm in section 6 makes two more assumptions although our greedy heuristics in section 5 works without them. We assume a minimal separation of distance sR_t between any pair of transmitters where $0 < s$. This assumption is natural since the two transmitters can not be co-located in space. We also assume the base station knows the location of each node. Clearly, this is not an issue when the relays are fixed and part of the infrastructure. For the case where relays are mobile nodes, location information have to be obtained in other ways such as through the base station's estimate using signal strength and angle of arrival as part of the E911 service, or through GPS and other localization mechanisms [12] applied on the mobile nodes. If

explicit location information is not available, the base station can compute an embedding of the connectivity graph of the ad-hoc networks [20] and our algorithm works on the embedding rather than the real coordinates.

5 Greedy Algorithm

In this section, we present our greedy algorithm that chooses proxies and establishes multicast routing entries for the distribution of packets from the proxies to the multicast receivers. We then point out the unique issue of multihop wireless channel interferences that limit the performance gain of the greedy ad-hoc relay, and motivate our analysis and design of the near optimal throughput multicast relay as presented in the next section.

Our greedy ad-hoc relay protocol consists of two components: *multi-path greedy proxy discovery* and *opportunistic relay path merging*. In multi-path greedy proxy discovery, each multicast receiver establishes multiple greedy relay paths at the base station. With the partial topology around each multicast receiver available, the base station chooses proxies according to both their instantaneous downlink channel rates and the total number of multicast receivers that the proxy can reach, thereby opportunistically merging relay paths to different multicast receivers.

5.1 Multi-path greedy proxy discovery

The proxy discovery is initiated from a multicast receiver by broadcasting a RTREQ message within a certain range. The RTREQ message carries the multicast receiver's average HDR downlink channel rate, the multicast group ID, and a sequence number that is incremented every time the multicast receiver initiates a new round of proxy discovery.

The processing of a RTREQ message is shown in Figure 5. Whenever it receives a RTREQ message, a client compares the sequence number with the largest RTREQ sequence number it has seen for the multicast receiver. It drops the RTREQ message if the sequence number is smaller, or if the sequence numbers are equal but the hop number is no smaller. The client then compares its own HDR downlink channel rate with that included in the RTREQ message. It further processes the RTREQ message only if its own HDR downlink channel rate is higher. The client then writes its own channel rate into RTREQ, and forwards a copy of the RTREQ message to the HDR base station. Finally the client decrements the RTREQ message's TTL. If the TTL is still positive, the client attaches its identifier into the relay path in the RTREQ message, and further broadcasts the updated RTREQ.

The RTREQ message is propagated *greedily* along paths of clients with increasing HDR down-

```
recvRTREQ(pkt, w)//u gets RTREQ from w
1.  (seqNo,hop) = getSeqNoHopCnt(pkt.src,pkt.mcastid);
2.  if (pkt.seqNo > seqNo) OR
3.     ((pkt.seqNo == seqNo)AND(pkt.hop < hop)) then
4.        createRouteEntry(pkt.src,pkt.mcastid,w);
5.        if (pkt.channelRate < channelRate) then
6.           pkt.channelRate = channelRate;
              // u declares to BS that it can be a proxy
7.           declareProxyToBS(u,pkt.src,pkt.mcastid);
8.           if (pkt.TTL > 0) then
9.              pkt.TTL = pkt.TTL-1;
10.             attachToPath(pkt.path,u);
                 // u broadcasts RTREQ to its neighbors
11.             broadcast(pkt);
12.          endif
13.       endif
14. endif
```

Figure 5: Multi-path greedy proxy discovery at node u

link rates. We choose to only allow these nodes with higher average HDR downlink channel rates to report to the base station since other nodes with lower HDR downlink channel rates are unlikely to have high instantaneous HDR downlink channel rates to serve as proxies for the multicast receivers.

Note that, there is no route reply messages from proxies back to the multicast receiver that sends out the RTREQ messages. Candidate proxies send path information to the base station. The reason is that, the base station has more topology and channel information than the multicast receiver and can make better decisions in selecting proxies. However, it is the multicast receiver that initiates the proxy discovery process. If no proxy information has been established yet, the base station will simply default to 3G multicast.

5.2 Opportunistic relay path merging

When it receives a RTREQ message, the HDR base station extracts the relay path from the RTREQ messages and updates the ad-hoc network topology that is constructed and maintained using RTREQ messages. The topology represents all the greedy relay paths for each multicast

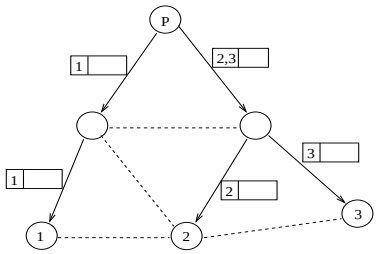

Figure 6: Ad-hoc relay from proxy P to multicast receiver v_1, v_2, and v_3 using IEEE 802.11 unicast

receiver in the the multicast group. With such topology available, the base station ranks each relay client according to the total number of multicast receivers reachable within certain range, e.g., 3 hops. Higher ranked relay clients, i.e., the clients that are connected to more multicast receivers within certain range, will be chosen as proxies with higher priority. This way, the base station merges the paths to different multicast receivers opportunistically to save the relay overhead on the IEEE 802.11 ad-hoc network. By merging, we mean that, for each common link among different receiver relay paths, we only send one copy of the packet. The data rate of the HDR downlink broadcast is then set to the minimum of all proxies' HDR downlink rate: $\min_{v \in R} p(v)^t$.

5.3 Ad-hoc relay of multicast packets

Once a proxy client receives a packet from the HDR downlink broadcast, it sends out the packet to all the multicast receivers in its routing table that is established during the RTREQ message propagation. To avoid sending multiple copies of the same packet to the same next hop client, the proxy client adds a new "Destination Header" specifying all the intended destination multicast receivers for each packet transmitted to the next hop relay client. This destination header will split as the packet propagates down the multicast tree. See Figure 6 for an illustration.

Although the destination header eliminates redundant transmissions of the same packet along a single multicast tree, a multicast receiver may still receive multiple copies of the same packet from different proxies due to overhearing. The multicast receiver can send out a "prune" message hop-by-hop upstream to eliminate the relay from certain proxies, e.g., the proxies from which it receives delayed packets compared with another proxy. Under high dynamics of the ad-hoc topology due to node mobility, the multicast receiver may simply choose to keep multiple proxies to increase the relay resilience.

Client	Greedy Relay	Optimal Relay
v_2	544.2	1046.9
v_3	547.2	1041.9

Table 1: Throughput of two relay strategies for the example in Figure 7 (Kbps)

5.4 Relay path maintenance

A relay path breaks when the proxy, relay or multicast receiver moves out of range. When the next-hop relay client is out of reach, the IEEE 802.11b MAC layer calls a callback function to inform the relay client of such failures[1]. The relay client then reports this routing failure to the HDR base station using the HDR uplink. The routing failure message deletes the broken wireless link from the topology maintained at the base station, and initiates the re-computation of the proxies. Besides, the relay client also sends out "prune" messages one-hop upstream to notify its upstream node the unreachability of the multicast receivers included in the destination header of the multicast packet. Similar approaches apply when an existing multicast node leaves or a new node joins the multicast group.

5.5 Impact of wireless interference

As we shall see later in Section 7, in many situations the greedy ad-hoc relay significantly improves the multicast throughput by as much as 400∼600%. However, we now discuss why, in some situations, the greedy ad-hoc relay may not perform well.

The primary goal of the greedy ad-hoc relay strategy is to choose the best proxy. The opportunistic merging tries to minimize the number of forwarding hops on the ad-hoc relay path but this is secondary to proxy discovery. However, since the offered load to the ad-hoc network equals the data rate of the 3G downlink at the proxy "magnified" by a factor of the number of multicast receivers who are using proxies, it could turn out in some situations that the 3G downlink channel rate of the best proxy may be higher than the capacity of the ad-hoc relay path, i.e., the ad-hoc relay path becomes the bottleneck. We use a simple example shown in figure 7 to illustrate this problem.

In figure 7 two multicast receivers, i.e., clients v_2 and v_3, belong to a multicast group. For each multicast receivers there are two alternative relay paths, as shown in the figure with solid and dot lines. Clients v_{14} and v_{15} are located closer to the base station, and their average HDR downlink

[1]In the case of unicast, link failure is detected by failure to receive CTSs or ACKs; in the case of broadcast failure is detected by lack of heart-beat messages.

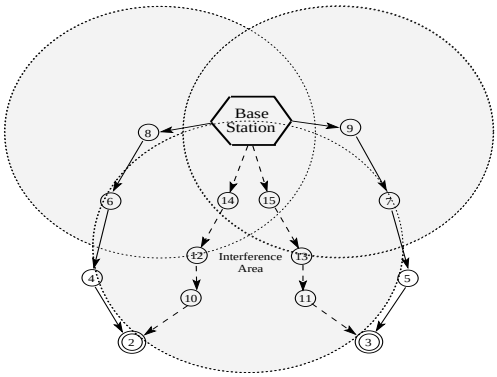

Figure 7: Greedy Ad-Hoc Relay v.s. Optimal Relay to Multicast Receiver v_2 and v_3

channel rates are higher than that of clients v_8 and v_9. The greedy ad-hoc relay will discover client v_{14} and v_{15}, in pursuit of the best HDR downlink proxy and highest HDR downlink data rate. However, because the two relay paths, i.e., $v_{14} \rightarrow v_{12} \rightarrow v_{10} \rightarrow v_2$ and $v_{15} \rightarrow v_{13} \rightarrow v_{11} \rightarrow v_3$ interfere with each other, at any given time only one path can transmit and receive packets. On the other hand, although clients v_8 and v_9 have slightly lower HDR downlink channel rates than those of clients v_{14} and v_{15}, relaying through $v_8 \rightarrow v_6 \rightarrow v_4 \rightarrow v_2$ and $v_9 \rightarrow v_7 \rightarrow v_5 \rightarrow v_3$ results in higher end-to-end multicast throughput because these two paths are out of their interference range and they can transmit and receive concurrently. The throughput of these two different relay strategies are shown in table 1. The throughput of greedy relay is only half of that of the optimal relay strategy.

The reason for greedy ad-hoc relay's sub-optimal performance is that *greedily optimizing the throughput for each multicast receiver individually does not yield globally optimal throughput due to interferences in the ad-hoc network*. Therefore, in addition to discovering a good proxy, explicit consideration of the wireless channel interferences between different relay paths is necessary for maximizing relay throughput. This is discussed in more detail in the next section.

6 Near Optimal Algorithm

In this section, we first present our notations and a precise definition of the problem. We then present a 4-approximation algorithm that runs in polynomial time. For ease of presentation, we first describe our algorithm assuming broadcast transmission and single-rate ad hoc networks in Section 6.1 and 6.2. We then show how our algorithm can be extended to multi-rate ad hoc networks in Section 6.3 . All our results easily carry over to the case of unicast transmissions with minor modifications. We make clear these modifications during our description of the algorithm in

$G = (V, E)$	connectivity graph of the 802.11 ad hoc network
$I = (V, A)$	interference graph of the 802.11 ad hoc network
R_t	transmission range of each node in V
R_I	the maximum interference range, $R_I = qR_t$
r_v	the 3G downlink channel rate of $v \in V$
$p(v)$	the 3G downlink channel rate of the proxy of node $v \in V$
r_v^a	the data rate of v received from the ad hoc relay network $r_v^a = \min\{f/k(G'), p(v)\}$
$r(G)$	the optimal multicast rate
f	the channel rate of 802.11 ad hoc network
$k(G')$	minimal number of colors needed to color G', a subgraph of I
k_C	minimal number of colors needed to color the best relay network for receivers in cell $C \in \Gamma$
R	the set of multicast receivers in a multicast group
R_{3G}	a subset of R that receive the multicast session directly from 3G and are not selected in ad hoc relay subnetwork
R_C	a subset of R that are in cell $C \in \Gamma$
Γ	the number of grid cells under a base station
$G'_C = (V'_C, E'_C)$	the optimal multicast relay subnetwork for cell $C \in \Gamma$
$G_O = (V_O, E_O)$	the ad hoc relay network output by algorithm ALGO

Table 2: A summary of notations used in Section 6.

Section 6.2 and 6.3. We further describe some techniques that result in improved performance of our algorithm in practice. Finally, we discuss how our algorithm can deal with multiple multicast groups.

6.1 Problem Statement

In this section, we introduce notations and formally state our problem. Our notations are summarized in Table 2. We are given an ad hoc network $G = (V, E)$ where V is the set of (n) nodes and E the set of (m) links. If there exists $(u, v) \in E$, then node u and v are at most R_t apart. Note that we do not make the assumption that G is a unit disk graph [9]. The set $R \subseteq V$ is the set of receivers of a given 3G multicast group. Every node $v \in V$ is associated with a rate $r_v \geq 0$, which is node v's average downlink rate from the $3G$ base station. A receiver $v \in R$ may receive data in one of the two ways: either directly from the $3G$ base station–at a rate of r_v or via multicast by a

proxy on the ad hoc network. The ad hoc relay subnetwork for multicast is formed by a subset of the nodes $V' \subseteq V$ and their incident links. Thus the ad hoc relay subnetwork for a given multicast group is a graph $G' = (V', E')$ such that $(u, v) \in E'$ iff $u, v \in V'$ and $(u, v) \in E$. For example, one relay subnetwork for multicast in Figure 7 consists of the set of nodes $\{v_8, v_6, v_4, v_2, v_9, v_7, v_5, v_3\}$ and the set of links $\{(v_8, v_6), (v_6, v_4), (v_4, v_2), (v_9, v_7), (v_7, v_5), (v_5, v_3)\}$. Node v_2 and v_3 are receivers. Node v_2's proxy is v_8 and node v_3's proxy is v_9. Note that G' can thus be specified by V' since E' is determined by V'. Within the ad hoc relay network the communication takes place over a (forest) collection of rooted and directed trees F spanning the nodes in V'. For example, in Figure 7, there are two trees, namely the two relay paths, one for each receiver. Within each ad hoc rooted tree $T \in F$ data (received from the $3G$ base station) is multicasted (using the nodes of T) by its root (proxy) to all the receivers in T. For ease of description, we first assume that all links in the ad hoc network support a common rate f as is common in 802.11 networks. We relax this assumption in Section 6.3. Note that not all receivers may be in V'. However such receivers (denoted by set R_{3G}) must receive data from $3G$ directly for a rate r_v for node v. The rate at which data is received by the rest of the receivers (in $R - R_{3G}$) is the better of r_v and r_v^a for receiver v, where the latter is the rate at which data is received by v via the ad hoc relay subnetwork. For a given G' we denote $\max\{r_v, r_v^a\}$ by $r_v(G')$, the data rate that receiver v can receive. Note that, we do not place any restrictions on receivers, that is, a receiver whether it receives directly from 3G or through proxy relay, it can be on the relay path or be a proxy of another receiver.

We denote by $I = (V, A)$ the interference graph for the ad hoc network. Thus, two ad hoc nodes u and v interfere with each other iff there is a link (u, v) in A, implying that any two interfering nodes cannot broadcast simultaneously[2]. As we stated earlier, we make the natural assumption that $(u, v) \notin A$ if u and v are at least qR_t apart for some fixed constant q. For a given ad hoc network $G' = (V', E')$ let $k(G')$ denote the minimum number of colors required to color the nodes[3] in V' such that two nodes $u, v \in V'$ have the same color iff there is no link (u, v) in A. Thus, the best multicast rate that can be achieved in the ad hoc network G' is at most $f/k(G')$. Note that, to be precise, $k(G')$ is determined by the optimal coloring of only the non-leaf nodes (receivers) of the trees of F since the leaf nodes do not participate in any transmissions. Our results, although applicable to this more precise model, are more involved and hence for ease of presentation we will use the model where all leaf receivers, except R_{3G}, are colored.

As stated earlier, we make the assumption that in any solution G' the best proxy for a receiver

[2]For unicast transmissions, each node in the interference graph will represent an edge in G.

[3]For unicast transmissions, $k(G')$ denotes the minimum number of colors required to color the edges in E'.

v is some node no more than h hops from v for some small value h. ¿From our simulation results in Section 7, for a 500m radius cell, we find that searching for a proxy beyond a $h = 4$ hop neighborhood does not result in increasing gains and thus, a practical value of h is 4. Consider a receiver $v \in R - R_{3G}$. Let $p(v)$ be the rate of the $3G$ proxy for v. Then $r_v^a = \min\{f/k(G'), p(v)\}$. Hence the multicast rate for v in G' is:

$$r_v(G') = \max\{r_v, \min\{f/k(G'), p(v)\}\}$$

Denote $r(G') = \min_{v \in R} r_v(G')$. The ICAM problem is to compute a G' such that the multicast rate $r(G')$ for its associated ad hoc subnetwork G' is maximized.

We show that ICAM problem with broadcast transmission[4] is NP-complete using a reduction from the problem of coloring civilized unit disk graphs. A graph is a (r, s)-civilized graph if its vertices can be mapped to points in d-dimensional space so that the length of each edge is $\leq r$ and the distance between any two points is $\geq s$. The MINCOLOR problem for unit disk graphs was shown to be NP-complete in [9]. An examination of this proof shows that the graph resulting from the reduction is a (1,1)-civilized graph. Hence the MINCOLOR problem remains NP-complete for (r, s)-civilized graphs. Given a civilized unit disk graph $G = (V, E)$ we create an instance of the ICAM problem as follows. We set $R_t = 1$ and G as the ad hoc network. We set $p = 1$ and the interference graph $I = G$. The set of receivers $R = V$. That is every node is a receiver. All receivers v have $r_v = 0$ except those defined by the following algorithm, which have r_v set to a very large value M. The algorithm picks the set of receivers v(denoted by the set R') with $r_v = M$ as follows. Initially $R' = V$. We say a node v covers node u if v and u are at most hR_t apart. We assume without loss of generality that every node u covers at least one other node $v \in V$. As long as there is a vertex v in R' such that all the nodes that are covered by v (including v) are also covered by some other node in $R' - \{v\}$ then R' is set to $R' = R' - \{v\}$. Note that when this procedure stops, then for every node $v \in R'$ there exists a node $u \in V - R'$ (u must be different from v) that is covered by v but not by any other node in R'.

We now show that any optimal solution to the ICAM problem must pick the subset $V' = V$ (i.e. all the nodes in the ad hoc network) as the ad hoc relay subnetwork and R' as the set of proxies or in other words for this choice, the multicast rate $r(G')$ for the associated ad hoc network G' is maximized. Note that $V - R' \subseteq V'$ since otherwise the optimal multicast rate is zero (recall only nodes in R' have non-zero 3G channel rate). However, by the choice of R' all nodes in R' must also be in V' because otherwise some receiver's optimal proxy v will also have rate $r_v = 0$ and hence the optimal multicast rate will be zero. Finally, since M is chosen to be a large number

[4]We omit the proof for unicast transmission

the optimal multicast rate is f/k where k is the minimum number of colors needed to color the graph induced by the vertices in $V' = V$. Or in other words the optimal solution must color G with the least number of colors. Note that our reduction also works if the network connectivity graph is not a (r, s)-civilized graph.

6.2 Approximation Algorithm

For broadcast transmissions, given the nodes $V' \subseteq V$ of the ad hoc relay subnetwork G' the multicast rate achieved by G' is independent of the actual multicast trees[5] in G', and depends only on the set V'.

computeMcastForest()

1. Divide the 3G base station coverage area into a grid of cell size $(2h + q + \epsilon)R_t \times (2h + q + \epsilon)R_t$
2. for each grid cell $C \in \Gamma$ //Γ is the set of grid cells
3. $\quad V'_C$ = computeOptCellForest(C);
4. // Merge the solution of each cell
5. $V_O = \cup_{C \in \Gamma} V'_C$
6. $R_{3G} = R \cap (V - V_O)$

Figure 8: Algorithm ALGO for Computing Multicast Relay Forest

As illustrated in Figure 8, the approximation algorithm ALGO for the problem ICAM works by dividing the coverage area of a 3G base station into a 2-dimensional grid of cell sizes $(2h + q + \epsilon)R_t \times (2h + q + \epsilon)R_t$ (ϵ is any value greater than 0). The rationale for this choice of the grid size is described later. We denote the set of grid cells Γ. Given an instance of the problem ICAM, ALGO independently computes a solution for each cell of the grid Γ that contains at least one receiver. More specifically ALGO computes the best solution for the given problem instance when restricted only to the receivers $R_C \subseteq R$ in any cell $C \in \Gamma$. Next ALGO merges these solutions for all cells to compute a feasible solution to the original instance of the problem.

Let C be a cell of the grid with at least one receiver ($|R_C| > 0$). Let $V_C \subseteq V$ be the set of all nodes that are at most h hops from at least one receiver in R_C. Note that in any optimal solution the set of proxies for the receivers in R_C and any intermediate nodes necessary for multicasting to these receivers must be in V_C. Hence multicast rate in the optimal solution to the ICAM problem when restricted to the subgraph of G induced by the nodes V_C and with the receivers R_C must be greater than the multicast rate in the optimal solution of the original problem.

[5]For unicast transmissions, we also need to pick the actual multicast tree.

computeOptCellForest(C)

1. Enumerate all subsets V'_C of V_C // V_C is the set of
   // nodes that are within h hops from receivers in cell C
2. for each $v \in V'_C$
3. p(v) = findBestProxy(v);
4. rm = computeMinProxyRate();
5. $k(G'_C)$ = minColor(V'_C);
6. $rate$ = min(rm, $f/k(G'_C)$);
7. Output the subset V'_C with maximal $rate$

Figure 9: Algorithm for Computing Optimal Multicast Relay Forest in one Grid Cell

Algorithm ALGO computes the solution for a cell C as follows (see Figure 9). It enumerates all subsets V'_C of nodes in V_C. For a given subset V'_C, let G'_C be its associated ad hoc relay subnetwork (as defined earlier in Section 6.1). ALGO computes the minimum number of colors $k(G'_C)$ needed to color the vertexes of G'_C based on the interference graph I. We will show later that this can be done efficiently. Algorithm ALGO then computes the best proxy $p(v)$ in G'_C for every receiver $v \in V'_C$, as described before in Section 6.1. Note that all this information is sufficient to compute $r_v(G'_C)$ for every receiver $v \in V'_C$. The $r_v(G'_C)$ of all receivers $v \in R_C$ that are not in V'_C is $r_v(G'_C) = r_v$. In other words, these receivers are considered to be not part of the ad hoc relay subnetwork. By taking the minimum of all the $r_v(G'_C)$ values for all receivers $v \in R_C$, ALGO is able to compute the multicast rate $r(G'_C)$ for the ad hoc relay subnetwork associated with the subset V'_C for the receivers in R_C. Finally by selecting the best subset $V'_C \subseteq V_C$, whose associated ad hoc relay subnetwork has the highest rate, algorithm ALGO is able to compute the optimal solution for the receivers R_C in cell C. We will show later that all this can be done in constant time.

Having computed the best ad hoc network induced by the subset V'_C for the receivers R_C in every cell C ($|R_C| > 0$), the algorithm ALGO can output the union of the subnetworks computed for each grid cell, i.e. $V_O = \bigcup_{C \in \Gamma} V'_C \subseteq V$ as the solution for the original problem instance. The set of receivers R_{3G} that receive directly from 3G base station is $R \cap (V - V_O)$. Note that this is a feasible solution for the original problem instance.

We now show that the algorithm ALGO runs in polynomial time and is a 4-approximation.

Lemma 1 *The number of nodes in any set $V_C \subseteq V$ for a cell C is bounded by a constant.*

Proof: Note that the nodes in V_C are all contained in a bounding box of size $(4h+q)R_t \times (4h+q)R_t$

(for receivers at the edge of a cell in the grid, there can be a proxy at most hR_t outside the cell resulting in a $2hR_t$ increase to the size of the edge of the cell). Also note that by the civilized graph assumption, i.e minimal separation distance of sR_t, $0 < s < 1$, between any two nodes, there are at most $((4h+q)/s)^2$ nodes in any bounding box of this size. Hence $|V_C| \leq ((4h+q)/s)^2$.

Lemma 2 *Algorithm ALGO runs in polynomial time.*

Proof: Note that there are polynomial number (at most $|R|$) of cells C with $|R_C| > 0$. By Lemma 1 the set V_C has a constant number of nodes for each cell C. Thus there are a constant number ($2^{|V_C|}$) of subsets of V_C. Each of these subsets has a constant number of nodes and hence its optimal coloring can be determined in constant time (by using brute force or by using the linear algorithm for coloring graphs with bounded tree-width [7]). Thus overall the algorithm runs in polynomial time. Note, however that even though the running time is polynomial it may still be quite large (in the worst case). Later we describe some pruning techniques that can be incorporated in the algorithm to make it quite practical.

Now we show that algorithm ALGO finds a solution to the ICAM problem, whose multicast rate is at least one fourth the multicast rate of any optimal solution. Let $r(G'_C)$ be the multicast rate of the best ad hoc relay subnetwork among those induced by the subsets $V'_C \subseteq V_C$ for the receivers R_C in any cell C ($|R_C| > 0$). Let $r(G)$ be the optimal multicast rate for the original problem instance. Note that for all C ($|R_C| > 0$), we have $r(G'_C) \geq r(G)$. Let k_C be the minimum number of colors used for coloring the best ad hoc relay subnetwork computed by the algorithm for cell C.

Lemma 3 *The optimal multicast rate for the original problem instance is at most* $\frac{f}{max_{C \in \Gamma} k_C}$.

Proof: Let C be the cell for which the maximum $max_{C \in \Gamma} k_C$ is attained. Note that if $k_C = 0$ then the result trivially holds. Otherwise there must exist a receiver $v \in R_C$ which receives a better data rate via the ad hoc relay subnetwork than its $3G$ rate r_v in the best ad hoc relay subnetwork G'_C computed by the algorithm for cell C. For this receiver v we must thus have that its data rate in G'_C is $\min\{f/k_C, p(v)\}$, where $p(v)$ is the best proxy for receiver v in G'_C. Thus

$$r(G) \leq r(G'_C) \leq \min\{f/k_C, p(v)\} \leq f/k_C.$$

Recall that $V'_C \subseteq V_C \subseteq V$ denotes the nodes in the best ad hoc relay subnetwork selected by Algorithm ALGO for cell C. In addition the algorithm ALGO outputs the ad hoc relay subnetwork associated with the set $V_O = \bigcup_C V'_C \subseteq V$ as the solution for the original problem

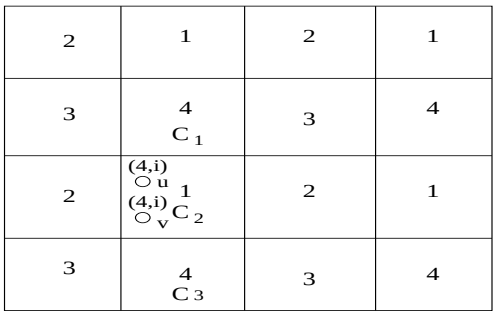

Figure 10: Grid can be colored with 4 colors

instance and then applies a merging step. For ease of description we however assume V_O as the solution output by the algorithm in the ensuing proof. Note that the merging step can only make the solution better.

Lemma 4 *The graph $G_O = (V_O, E_O)$ induced by the node set V_O can be colored with at most $4max_{C\in\Gamma}k_C$ colors*

Proof: Consider two cells C_1 and C_2 of the grid. Note that by the choice of the cell sizes $((2h+q+\epsilon)R_t \times (2h+q+\epsilon)R_t)$ a node in V_{C_1} can interfere with a node in V_{C_2} iff C_1 and C_2 are neighboring cells. Note that the cells of the grid can be colored with 4 colors such that no two adjacent cells get the same color (see Figure 10). Denote this coloring by the function $f(C)$. Thus $f(C)$ is the color for cell C in this coloring. Consider the optimal coloring (which uses k_C colors) of the best ad hoc relay subnetwork G'_C computed by the algorithm for receivers in cell C. Without loss of generality the colors used for this coloring are $1, 2, \ldots k_C$. Consider a new coloring of the nodes V'_C of G'_C where any node v with color i is given the color $(i, f(C))$. Note that this is a proper coloring of the nodes of V_O that induces the graph $G_O = (V_O, E_O)$. This is because nodes that receive the same color are strictly more than qR_t apart. For example, in Figure 10, node u and v may receive the same color $(4, i)$ from the solution of cell C_1 and C_3 respectively. However, u is at most hR_t away from the boundary of cell C_1 and v is at most hR_t away from the boundary of cell C_3. Since grid cell size is $(2h+q+\epsilon)R_t$, u and v must be strictly more than qR_t apart. By construction this coloring uses at most $4max_{C\in\Gamma}k_C$ colors.

Theorem 5 *Algorithm ALGO is a 4-approximation for the problem ICAM*

Proof: Let $k = max_{C\in\Gamma}k_C$. Consider a receiver $v \in V_O$, which is in cell C of the grid. Let $r(G)$ be the optimal multicast rate. Let G'_C denote the best ad hoc relay subnetwork computed by the algorithm for receivers in cell C. Recall that the rate for receiver v in G'_C is

$$r_v(G'_C) = \max\{r_v, \min\{f/k_C, p(v)\}\}$$

where $p(v)$ is the best proxy for node v in V_C and hence in V_O. Note that as shown before $r(G) \leq r_v(G'_C)$. The rate of v in G_O is

$$r_v(G_O) = \max\{r_v, \min\{f/4k, p(v)\}\}.$$

Note that if

$$r(G) \leq r_v(G'_C) = r_v$$

then

$$r_v(G_O) = r_v \geq r(G).$$

Otherwise

$$r(G) \leq \min\{f/k_C, p(v)\} \leq p(v).$$

Also by Lemma 3, $r(G) \leq f/k$. Thus

$$r(G)/4 \leq \min\{f/4k, p(v)\} \leq r_v(G_O).$$

For any other receiver $v \in R_C$ its rate is r_v in both solutions hence

$$r(G) \leq r_v(G'_C) = r_v(G_O).$$

Thus for any receiver $v \in R$ we have

$$r(G)/4 \leq r_v(G_O),$$

thus implying the claimed bound.

6.3 Extending to Multi-rate Ad Hoc Networks

In this section we present at a high level the extensions required for algorithm ALGO to work in multi-rate ad-hoc networks. For ease of presentation we describe these extension assuming broadcast transmissions. Again, our results easily carry over to networks (e.g. current 802.11 networks) where unicast transmission is used. As the single-rate algorithm in Section 6.2, the multi-rate algorithm ALGO works by finding an (approximately) optimal solution for each cell C of the grid Γ by enumeration, and then obtains the overall solution by combining these solutions for all the cells. Within each cell algorithm ALGO enumerates all potential multicast trees and for each tree finds the (approximately) best possible broadcast schedule. It then outputs that tree for the cell for which the achieved multicast rate is maximized. Note that given a multicast tree T for a given cell the best broadcast schedule may assign different number of broadcast slots (colors) to the nodes of T. More precisely let $f_v(T)$ be the rate when node v broadcasts to its

downstream receivers in T ($f_v(T)$ is determined by the lowest link rate for all its receivers) then in the best schedule a node with relatively low value of $f_v(T)$ must broadcast more often to ensure that the multicast rate achieved is high. Algorithm ALGO is modified to enumerate the possible slot assignments for the nodes of T. The main technical challenge is that the number of such possible slot assignments can be infinitely large. Hence algorithm ALGO must carefully select a small set of possible slot assignments, such that the best slot assignment within this set is close to the overall optimal slot assignment in terms of the achieved multicast rate. The size of the search space used by algorithm ALGO can be controlled by a user specified error tolerance ϵ: the smaller ϵ is, the larger the search space and the better the guarantees on the rate of the computed solution. More precisely the algorithm enumerates all those possible slot assignments in which the number of slots assigned to a node v with the largest $f_v(T)$ ranges from 1 to $\lceil(1+1/L)/\epsilon\rceil L$ where L is the least common multiplier of all link rates in the ad-hoc network (e.g. $L = 22$ in the standard 802.11b network). For a given number of slot requirements $k_v(T)$ of each node v of T, algorithm ALGO finds an minimum coloring of the interference graph induced by the set of nodes in T, in which $k_v(T)$ distinct colors are assigned to each node v. Note that if $K(T)$ denotes the total number of colors used by this optimal solution, then the multicast rate achievable in this slot assignment is computed as $r(T) = \min_{v \in T} f_v(T)k_v(T)/K(T)$. Since assigning more colors to a node u such that $f_u(T)k_u(T)/K(T) > r(T)$ does not improve the multicast rate, we assume $f_u(T)k_u(T)/K(T) = f_v(T)k_v(T)$ for the rest of our discussion. Algorithm ALGO approximates the optimal multicast rate for a given tree T to within a factor of $1+\epsilon$ and runs in time which is exponential in $1/\epsilon$, resulting in a Polynomial Time Approximation Scheme (PTAS). Let $r(T)$ be the maximum rate (achievable by T without considering 3G rate) given by our algorithm, and r^* be the corresponding optimal rate. Let k^* be the total number of colors used by the optimal. We now give the formal proof.

Lemma 6 *Algorithm ALGO outputs $r(T)$ such that $r^* \leq (1+\epsilon)r(T)$.*

Proof: Let $m = \lceil(1+1/L)/\epsilon\rceil$. Let k^*_{max} be the number of colors assigned by the optimal solution to a node with the largest rate. Let k^* be the total number of colors used by theis optimal solution. Here we assume that k^*_{max} cannot be reduced further without effecting the optimal multicast rate. In other words the rate r^* achieved by this solution satisfies:

$$\frac{(k^*_{max}-1)f_{max}}{k^*} < r^* \leq \frac{(k^*_{max})f_{max}}{k^*}. \quad (1)$$

Note that this can be easily ensured by reducing k^*_{max} until this inequality is satisfied. There are two cases.

- $k^*_{max} \leq mL$. Since our search will cover the optimal solution, our algorithm will achieve the optimal multicast rate.

- $k^*_{max} > mL$. Consider a restriction of the problem in which the node with the maximum rate f_{max} in the tree T is limited to use exactly L number of colors (slots). Note that under this restriction any other node v with rate f_v must be assigned $\lceil \frac{Lf_{max}}{f_v} \rceil$ colors. Note that by construction $\frac{L}{f_v}$ is integral, implying that

$$\lceil \frac{Lf_{max}}{f_v} \rceil = \frac{Lf_{max}}{f_v}.$$

Consider an optimal solution to this restricted problem in which the achieved multicast rate is maximized. Let $K(T)$ be the total number of colors (slots) used by this optimal solution. Then the rate $r(T)$ achieved for this solution is exactly $r(T) = Lf_{max}/K(T)$, since this is the rate achieved at every node of T. Note that since our algorithm in its enumeration considers this restricted problem, thus our algorithm outputs a solution of rate at least $r(T)$. We now show that the optimal rate for the original problem (unrestricted version) is at most $(1+\epsilon)r(T)$, thus establishing the bound.

Note that since $k^*_{max} > mL$ there must exist a $p > m$ such that, $pL \leq (k^*_{max} - 1)$ and $(p+1)L > (k^*_{max} - 1)$. This is because $k^*_{max} = pL$ is equivalent to $k^*_{max} = L$ since scaling does not change the solution. We claim that, $k^* \geq pK(T)$. This is because as defined in (1) the rate r^* is determined by the number of colors (k^*_{max}) assigned to a node with the largest rate f_{max}. Hence any other node v with rate f_v in the optimal tree T^* must be assigned at least

$$\frac{(k^*_{max} - 1)f_{max}}{f_v} \geq \frac{pLf_{max}}{f_v}$$

colors by the optimal solution. Hence the optimal coloring is still valid if we decrease the number of colors assigned to each node v to $\frac{pLf_{max}}{f_v}$. Then by scaling by factor p both the total number of colors used and the color assignment to each node we get a solution that assigns $\frac{Lf_{max}}{f_v}$ colors to every node but uses less than $K(T)$ total number of colors a contradiction due to the optimality of $K(T)$.

Therefore, we have,

$$\begin{aligned} r^* \leq k^*_{max} f_{max}/k^* &\leq k^*_{max} f_{max}/(pK) \\ &< ((p+1)L+1)f_{max}/(pK) \\ &\leq ((p+1)L+1)/(pL)r(T) \\ &\leq ((m+1)L+1)/(mL)r(T) \\ &\leq (1+\epsilon)r(T) \end{aligned}$$

Algorithm ALGO merges the optimal solutions of all the cells C as described for the single-rate case. The proof for the overall performance bound of the algorithm is along the line of the single-rate case with the following major modifications. Let K_C and $k_C(v)$ be the total number of colors and the number of colors assigned to node v respectively by (approximately) optimal solution for grid C. Let $K = LCM(\{K_C | C \in \Gamma\})$, where LCM denotes least common multiplier. We extend the set of colors assigned to the optimal multicast tree for cell C to a total of K colors by replacing each color i ($1 \leq i \leq K_C$) by the K/K_C colors $i, i+K_C, i+2K_C, \ldots i+(K/K_C-1)K_C$. Note that now every node v has $k_C(v)K/K_C$ colors and the new coloring uses exactly K total number of colors and is a valid coloring for the solution for cell C. Note that as shown in Lemma 4 this implies that the merged solution can be colored with at most $4K$ colors. Note that in this coloring of the merged solution a node v for cell C gets at least $(k_C(v)K/K_C)/4K = (k_C(v)/K_C)/4$ fraction of the total number of colors, which is at least one fourth of the fraction of the total number of colors it gets in the (approximately) optimal solution for cell C. This observation combined with the proof of Lemma 5 establishes the following theorem:

Theorem 7 *Algorithm ALGO is a $4(1+\epsilon)$ polynomial time approximation scheme for the problem ICAM for multi-rate ad hoc networks.*

Proof: Recall $K = LCM(\{K_C | C \in \Gamma\})$, i.e. K is the least common multiplier among the set of total number of colors for coloring grid cells. If a node gets color number j, it will get all colors such that $K - iK_c = j$ for $i > 0$. That is, if a node v gets k_v colors in cell C, it will get k_vK/K_c number of colors in our merged solution G_O. Since G_O uses $4k$ colors, the rate v gets is $(f_vk_vK/K_c)/4K = f_vk_v/(4K_c)$. Consider a receiver $v \in V_O$, which is in cell C of the grid. Let $r(G)$ be the optimal multicast rate. Let G'_C denote the best ad hoc relay subnetwork computed by the algorithm for receivers in cell C. By Lemma 6, the rate for receiver v in G'_C is

$$r_v(G'_C) \leq \max\{r_v, \min\{(1+\epsilon)f_vk_v/K_C, p(v)\}\}$$

where $p(v)$ is the best proxy for node v in V_C and hence in V_O. Note that as shown before $r(G) \leq r_v(G'_C)$. The rate of v in G_O is

$$r_v(G_O) = \max\{r_v, \min\{f_vk_v/(4K_c), p(v)\}\}.$$

Note that if

$$r(G) \leq r_v(G'_C) = r_v$$

then

$$r_v(G_O) = r_v \geq r(G).$$

Otherwise

$$r(G) \leq \min\{(1+\epsilon)f_v k_v/K_C, p(v)\} \leq p(v).$$

Note that $r(G) \leq (1+\epsilon)f_v k_v/K_C$. Thus

$$r(G)/4 \leq \min\{(1+\epsilon)f_v k_v/(4K_C), p(v)\} \leq (1+\epsilon)r_v(G_O).$$

For any other receiver $v \in R_C$ its rate is r_v in both solutions hence

$$r(G) \leq r_v(G'_C) = r_v(G_O).$$

Thus for any receiver $v \in R$ we have

$$r(G)/(4(1+\epsilon)) \leq r_v(G_O),$$

thus implying the claimed bound.

6.4 Practical considerations

We first discuss techniques to speed up the computation for one grid cell. When we enumerate relay network $\hat{G_C}$ for cell C, we keep the best $r^* = r(\hat{G*_C})$ so far. The set of receivers R^C_{3G} not in $\hat{G_C}$ (not participating in ad hoc relay) receives data directly from 3G. Since the number of nodes in a level of the tree decreases toward the root, a tree can only have $N = h \times |R_C - R^C_{3G}|$ number of nodes. Recall that R_C is the set of receivers in cell C. If the set cardinality of the enumerated set $\hat{V_C}$ minus R_C is greater N, then we discard the enumeration of all $\hat{V_C}$'s superset. For an enumerated relay network $\hat{G*_C}$, if the minimal rate among R^C_{3G} is smaller than r^*, then we do not process this set further. If the minimal rate among proxies of $R_C - R^C_{3G}$ is smaller than r^*, then we stop processing this set.

After we get $G_O = (V_O, E_O)$ from algorithm ALGO, one can apply the following optimization to make the solution more efficient in practice. The intuition is that, when we consider one cell C, some receivers in C can achieve a better rate using proxies. However, when we combine the solution, the whole session rate can be lower. Therefore, some of these receivers may not need proxies any more and can receive directly from 3G. In these cases, we can prune proxies not needed and links needed only for these proxies. Detailed description is as follows. First compute the achievable rate $r(G')$ of multicast for all receivers. Then identify all receivers that can receive directly from 3G base station, i.e. v such that $r_v \geq r(G')$. Denote this set J. Next prune all the proxies that have no receivers within h hops in $R - R_{3G} - J$. Finally, prune all relay nodes which do not connect any proxy with receivers. This optimization however has no implication on the worst case performance bound for the algorithm ALGO.

We can also adopt certain techniques that affects the approximation ratio. Instead of using optimal coloring algorithm, we can use a simple greedy heuristics [5] which has an approximation ratio of 3. We can also use smaller grid cells to trade off running time for worst case bound.

6.5 Multiple Multicast Groups

While the algorithm ALGO provided the near-optimal multicast relay strategy for a single multicast group, it can be applied independently to multiple multicast groups as well because of the following reason. Recall that HDR transmits frames in the downlink in a time-multiplexed fashion with a time slot duration of 1.67ms. Thus, the 802.11 ad-hoc network has 1.67ms to sink the traffic to the receivers before it receives the next packet from the 3G base station. We claim that this is sufficient time for the 802.11 network to deliver the packet to the receiver (barring failure scenarios such as route breakages etc.) for the following reasons. a) the proxy is at most h hops away from the receiver for some small value h (e.g. $h = 4$) and b) the 802.11 network is not the bottleneck (otherwise, the optimal algorithm would determine that it is best not to use any relays and let the 3G base station deliver the packets directly to the receivers). Our simulation results in Section 7.3.4 for the multiple multicast groups support this claim.

7 Performance Evaluation

In this section, we evaluate the performance of the greedy ad-hoc relay and the near-optimal ad-hoc relay in improving the multicast throughout. We first present the models, metrics and methodology for our evaluation. We then compare the performance of the greedy ad-hoc relay and near-optimal ad-hoc relay algorithms in networks with stationary nodes. Finally we present the simulation results of the greedy ad-hoc relay protocol in mobile scenarios, investigating the impact of the multicast group size, node density, node mobility, and multiple multicast groups.

7.1 Model, Metrics and Methodology

We implement the near-optimal ad-hoc relay and greedy ad-hoc relay in the *ns-2* simulator. The HDR downlink channel is modeled according to the published experimental data in [3, 6] (see Section 3). We use the IEEE 802.11b implementation in ns-2 version *2.1b9a* where 11Mbps data rate is supported at 115-meter communication range. The channel rate in our ad hoc network is set to 11Mbps. The radio propagation model for IEEE 802.11b is the Two-Ray Ground reflection model [21]. The node mobility is set according to the improved random waypoint model [30]. A mobile node starts at a random location, waits for a certain pause time, and randomly chooses

Throughput	No-Relay	Greedy	Near-Optimal
Max (Kbps)	103	760	740
Min (Kbps)	54	615	683
Avg (Kbps)	80	678	719
Avg Gain	-	785%	840%

Table 3: Throughput comparison over 20 random stationary scenarios

	Greedy	Near-Optimal
Avg HDR Downlink Rate	990Kbps	834Kbps
Goodput	0.685	0.862
802.11 Energy Consumption	67.4	42.5

Table 4: Average HDR Downlink Channel Rate and Goodput over 20 random stationary scenarios

a new destination and moves with a random speed chosen from a non-zero minimum to the maximum speed. We set the pause time to be 3 seconds, and vary the speed in the range of $[0.5, 1] \times \text{Speed}_{\text{max}}$ to investigate the impact of node mobility.

We simulate UDP/CBR multicast traffic. The packet size is set to 512 bytes. The total load of the CBR flows is set to 1.01×2.457Mbps. That is, the packet queue at the HDR base station is always backlogged.

We use four metrics to evaluate the performance of our relay protocols. We use the **average HDR downlink data rate** to evaluate the performance of our proxy discovery since the higher the downlink data rate, the better the proxy selection. We use **goodput**, i.e., the ratio of the number of packets a multicast receiver receives over the number of packets multicasted over the HDR downlink, as an indicator of relay loss, to evaluate the relay capability of the ad-hoc relay paths. We compare the **average throughput gains** to evaluate the overall performance of our proxy discovery and ad-hoc routing. Finally we show the **routing overhead on HDR uplink**.

7.2 Comparison between Greedy and Near-optimal Ad-hoc Relay

In this section we compare the performance of the greedy and near-optimal ad-hoc relay in stationary scenarios. While we show a hand-crafted example earlier (see Figure 7 and Table 1), where the optimal algorithm outperforms the greedy algorithm by 92%, we now study more realistic random scenarios. We place 30 nodes randomly in a square cell of $600 \times 600\text{m}^2$ with the HDR base station located in the center. We randomly place 5 multicast receivers in the cell. All

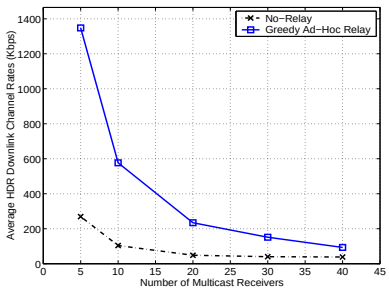

Figure 11: Mcast Group Size Impact: Average HDR downlink channel rate

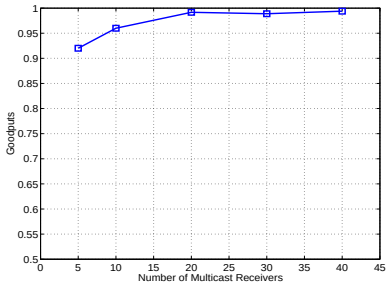

Figure 12: Mcast Group Impact: Goodput

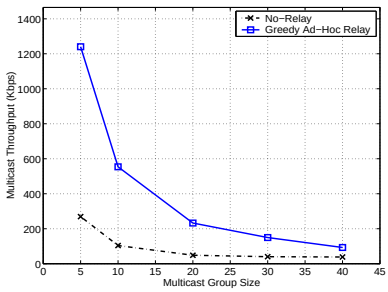

Figure 13: Mcast Group Size Impact: Throughput

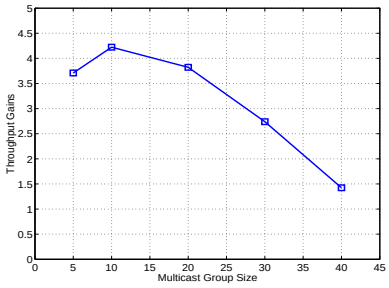

Figure 14: Mcast Group Size Impact: Throughput Gain

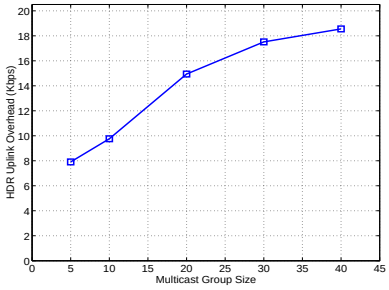

Figure 15: Mcast Group Impact: HDR Uplink Overhead

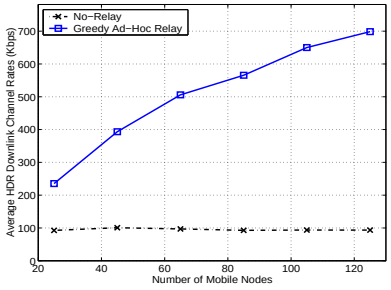

Figure 16: Node Density Impact: Average HDR downlink channel rate

results are the average over 20 random topologies generated by the `setdest` tool [30].

Table 3 shows the maximum, minimum and average end-to-end throughput over 20 random topologies for both greedy and near-optimal ad-hoc relay. Greedy ad-hoc relay protocol achieves throughput gains of 572∼897% with an average throughput gain of 785%. Near-optimal relay protocol further increases the average throughput by 55% and achieves an average throughput gain of as high as 840%.

Table 4 shows the average HDR downlink channel rate, goodput, and IEEE 802.11 interface energy consumption for the simulated scenarios. As we can see greedy proxy discovery finds proxies that are 18% better than the near-optimal algorithm in terms of the HDR downlink channel utilization. However, because of its explicit consideration of the multihop wireless channel interference, near-optimal algorithm achieves a 25.8% higher goodput than that of the greedy ad-hoc relay by carefully engineering the distribution of ad-hoc relay traffic. Since high goodput also means low packet loss along the ad-hoc relay paths and each packet loss involves four to seven IEEE 802.11 link layer re-transmissions, the near-optimal algorithm consumes 36.9% less energy

than the greedy algorithm on proxy and relay clients' IEEE 802.11 interfaces.

7.3 Mobile Scenarios

In this section we study the performance of the greedy ad-hoc relay protocol in mobile scenarios with different multicast group sizes, node densities, mobility speeds, and multiple multicast groups. Due to the computation cost of the near-optimal relay algorithm we do not apply it in scenarios with frequent topological changes. We also increase the size of the cell to $886 \times 886 m^2$ (approximating a 500-meter radius circular cell with the HDR base station located in the center) in order to accommodate more mobile nodes. We place mobile nodes randomly inside the cell.

7.3.1 Multicast group size

We first investigate the impact of the multicast group size. We randomly place 65 mobile nodes in an HDR cell, since a maximum number of 60 connected mobile users are allowed in an HDR sector [3]. We set the maximum moving speed as 15m/s. Figure 11-15 show the simulation results. As we can see from Figure 11 the greedy proxy discovery protocol is able to locate proxies with high HDR downlink channel rate. In scenarios with five multicast receivers the greedy proxy discovery improves the HDR downlink channel utilization from 270Kbps to 1.37Mbps, a gain of 407%. Note that the average downlink rate for the no-relay case has increased to 270Kbps from 80Kbps in the scenarios of Section 7.2 since we have increased the node density (we study the impact of varying node density on performance later).

Figure 11 also shows that as the multicast group size increases, the HDR downlink channel utilization decreases. The reason is that given the same node density, large number of multicast receivers increases the probability that at least one multicast receiver cannot locate a good proxy with high HDR downlink rate. Figure 12 shows that the goodput of the ad-hoc relay paths increases as the multicast group size increases, due to decreased offered load, i.e., the average HDR downlink channel rate. The combination of the HDR downlink channel utilization and the goodput leads to a throughput gain of 147$\sim$420%, as shown in Figure 13 and 14. Note that the throughput gain with the multicast group size of 5 nodes is slightly larger than the throughput gain with group size 10. This is due to the smaller number of proxy candidates, i.e., mobile nodes with HDR downlink data rate $\geq$1.4Mbps, and the lower goodput in the scenarios of multicast group size 5.

The relay path discovery and maintenance consumes the bandwidth resources of the HDR uplink. Figure 15 shows that the overhead increases as the multicast group size increases. In all scenarios, the greedy relay protocol consumes less than 10% of the HDR uplink's 153.6Kbps

bandwidth [3].

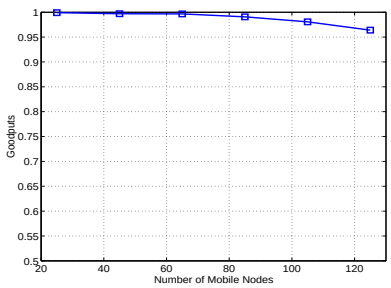

Figure 17: Node Density Impact: Goodput

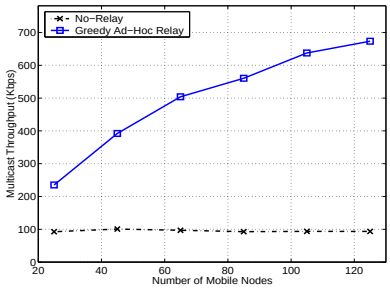

Figure 18: Node Density Impact: Throughput

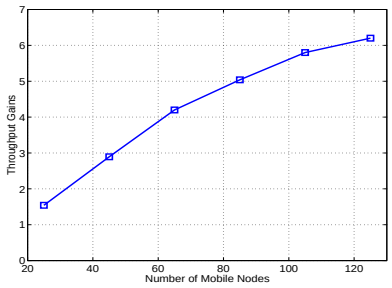

Figure 19: Node Density Impact: Throughput Gain

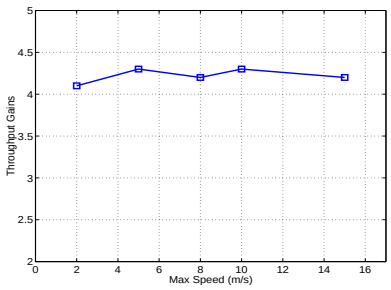

Figure 20: Mobility Impact: Throughput Gain

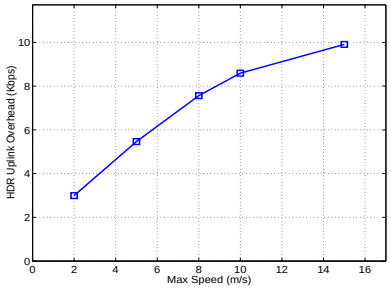

Figure 21: Mobility Impact: HDR Uplink Overhead

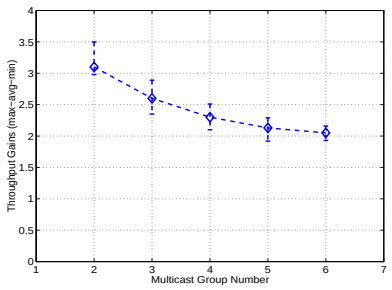

Figure 22: Multiple Multicast Groups: Throughput Gain

7.3.2 Node density

In order to study the impact of the node density in the ad-hoc relay network, we fix the multicast group size as 10 and the maximum node moving speed at 5m/s. We then change the total number of mobile nodes in the HDR cell, including those 10 multicast receivers, from 25 to 125. The results are shown in Figure 16-19. From Figure 16 we can see that the average HDR downlink channel rate increases as the ad-hoc network density increases. The reason is that high network density improves connectivity. It increases the chance for a multicast receiver to locate a proxy with high HDR downlink channel rate. The goodput, shown in Figure 17, decreases slightly as the network density increases, mainly due to the increased offered load (the higher average HDR downlink channel rate). As Figure 18 and 19 show, greedy ad-hoc relay achieves an end-to-end multicast throughput gain of 621% at the network density of an average of 5.3 neighbors per node.

Number of Groups	2	3	4	5	6
No-relay (Kbps)	118	136	145	150	157
Relay (Kbps)	485	489	478	469	478
Throughput Gain	311%	259%	228%	213%	205%

Table 5: Aggregate throughput of multiple multicast groups

7.3.3 Node mobility

To study the impact of node mobility, we use an HDR cell with 65 mobile nodes, including 10 multicast receivers. We set the maximum moving speed of the mobile nodes (including the multicast receivers) from 2 to 15m/s. As Figure 20 shows, the throughput gain of the greedy ad-hoc relay remains robust at around 410%, demonstrating effectiveness of the greedy relay path maintenance. However, the overhead on the HDR uplink (Figure 21) increases due to the increased number of link breakage reports sent by the relay clients.

7.3.4 Multiple multicast groups

We finally show the throughput gains for multiple multicast groups receiving packets from the base station. We simulate 2∼6 multicast groups, with 10 multicast receivers in each group, in an HDR cell of 65 mobile nodes including the multicast receivers. Proportional fairness scheduling [3, 19] is used to schedule among different multicast groups. We set the maximum moving speed at 5m/s. Note that, with 6 multicast groups instead of one, HDR scheduler is able to better exploit user diversity using proportional fairness (schedule the group with the best instantaneous rate among the six groups, subject to fairness), thus increasing the base throughput (without ad-hoc relay) to 157Kbps from 118Kbps (see Table 5). Thus, even though the throughput gain decreases from 311% to 205% as the number of groups increase in Figure 22, the absolute throughput achieved using the relay model remains almost the same. This is further supported by the fact that the goodput of the ad-hoc relay remains almost the same as that in the scenarios with one single multicast group (see Figure 17), supporting our earlier claim in Section 6.5 that since all multicast groups take turns sharing the HDR downlink bandwidth, there is not much additional contention in the ad-hoc network introduced by multiple multicast groups.

8 Discussion

In this section, we discuss other relevant issues.

Coding gains: While a conservative strategy for the 3G base station is to clearly send the packets at the data rate corresponding to the receiver with the worst downlink channel condition so that no packets are dropped, it is possible to do better by the use of coding. For example, in [4], the authors employ Reed-Solomon codes at the MAC layer with Turbo codes at the physical layer and achieve broadcast throughput of up to 204Kbps for stations with one antenna. However, this is still significantly lower than the multicast throughput achieved using a relay architecture such as ours. While the relay architecture will no longer provide gains of up to 800% over such a coded channel, we note that it still delivers throughput of 800Kbps for small multicast groups. Furthermore, the relay architecture will also benefit from improved coding efficiencies and we are currently investigating the impact of coding on ICAM.

Security: 3G multicast content is encrypted using a unique and frequently changing *short-term key* (SK). The Mobile Equipment part of the 3G mobile station decrypts the content using SK, which is derived from a *broadcast access key* (BAK). Only users subscribed to the service can obtain BAK. Since data is decrypted by frequently changing SK, even if a mobile station leaks the key to other relay nodes, they can not gain a substantial amount of the session packets. As the mobile station uses secure hardware, a user who owns a mobile station can not compromise the hardware and obtain the key that generates SK. Therefore, ad-hoc relay does not compromise security.

Incentive issues: If the relays are part of the infrastructure, then there are no incentive issues. If they are not, incentives must be provided to encourage nodes participating in relaying traffic. BS sends all the data and knows the multicast forest. Therefore, the BS can keep proper accounting for charging receivers and crediting proxies and other relay nodes. We now describe possible cheating behaviors and show how they can be prevented. (1) Nodes can not falsely claim their rate and be proxies. If they get chosen, they would not be able to correctly decode the packet. Since credits are based on the number of successfully forwarded packets, they will not get credits. (2) Receivers can not lie about received packets. If they do, 3G BS will try to send them directly which will cost more. (3) Forwarding nodes have no incentive to drop packets since they only get credits on successfully forwarded packets. (4) Relay nodes have less incentive to add other nodes in the relay path. This is because relay nodes do not know how BS chooses routes. BS is likely to choose a shorter route. For authenticating each relay path and preventing other collusion scenarios such as free riding, we can easily adapt the charging and rewarding schemes from [23] to our context.

Further improvement of the greedy ad-hoc relay: for simplicity, we adopted the simple greedy heuristics that perform well in our simulated scenarios. However, the greedy heuristics can

be further improved with explicit consideration of the wireless interferences, at the cost of higher computation complexity. One improvement can be as follows. Let $\bar{G} = (\bar{V}, \bar{E})$ be the partial topology discovered by proxy discovery. We do a binary search to find the highest multicast rate $r(G')$ in the interval $[\min_{v \in R} r_v, r_{max}]$ where r_{max} is 2.4Mbps in our architecture. For each guessed rate $\hat{r}$, let $Y = \{v | r_v \geq \hat{r}\}$ and let $R_{3G} = Y \cap R$. If all receivers in $R - R_{3G}$ are within h hops away from at least one node in Y, then we have found the best rate $r(G')$. After we find the best rate, we construct a rooted Steiner tree with bounded tree depth problem to minimize the number of transmissions. Create a new root node u. Create a zero-cost edge from u to any v such that $r_v \geq r(G')$. Let the set of terminals to connect in the Steiner tree problem be $R - R_{3G}$. Let the depth of the tree be bounded by h. Let all edge cost in $\bar{E}$ be 1. For the case of broadcast transmissions, the Steiner tree problem is node-weighted rather than edge weighted. The Steiner tree problems are well-known NP-hard problems but approximation algorithms can be used to obtain the solution.

Multi-channel 802.11: Another possibility for avoiding contention in the 802.11 ad-hoc network is through the use of multiple channels (802.11b has three orthogonal channels). While common techniques for choosing a channel in a multi-channel 802.11-based ad-hoc network is based on local knowledge [24], the use of a central 3G base station can definitely help in designing a global channel assignment algorithm. Another simpler possibility is to use standard graph coloring algorithms and allocate these orthogonal channels statically such that nodes belonging to a particular sector (some 3G cells can have up to 6 sectors) and cell or assigning a particular channel. We are examining this as part of the future work.

9 Conclusion

In this paper, we proposed the ICAM architecture for improving 3G multicast throughput using ad-hoc relays. 3G multicast throughput is limited by the receiver with the worst channel rate. By finding proxies for receivers with poor channel quality and relaying multicast packets through an IEEE 802.11-based ad-hoc network, we showed that the throughput of multicast sessions can be significantly improved. We designed protocols and algorithms to enable such ad-hoc relays. We presented two novel algorithms to determine the set of proxies: one based on greedy heuristics and the other, a near-optimal centralized algorithm that runs in polynomial time. Our near optimal approximation algorithm assumes a very general interference model and does not assume unit disk graph as the connectivity model. The bound holds when the underlying wireless media access control supports broadcast or unicast, single rate or multiple rates and even when there

are multiple simultaneous multicast sessions. Through extensive simulation, we showed that both our algorithms improve the average 3G multicast throughput by up to 840% with the near optimal algorithm outperforming the greedy heuristics by 6-92% in static scenarios while the greedy algorithm performing very well in relatively high mobility scenarios with throughput gains of as much as 410%.

References

[1] 3G Today. `http://www.3gtoday.com/`.

[2] GTRAN Dual-Mode 802.11/CDMA Wireless Modem. `http://www.gtranwireless.com/newsevents/pressreleases_20020422.htm`.

[3] 1xEV: 1x EVolution IS-856 TIA/EIA Standard - Airlink Overview. QUALCOMM Inc. White Paper, Nov. 2001.

[4] P. Agashe, R. Rezaiifar, and P. Bender. CDMA2000 High Rate Broadcast Packet Data Air Interface Design. *IEEE Communications Magazine*, pages 83–89, Feb. 2004.

[5] Y. Bejerano and R. Bhatia. Mifi: managed wifi for qos assuranc, fairness and high throughput of current ieee 802.11 networks with multiple access points. In *IEEE INFOCOM*, 2004.

[6] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi. CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data Service for Nomadic Users. *IEEE Communications Magazine*, 38:70–77, Jul. 2000.

[7] H. Bodlaender. Dynamic programming on graphs of bounded treewidth. In *Proceedings of 15th International Colloquium on Automata, Languages and Programming*, pages 631–643. Lecture Notes in Computer Science, 317, Springer-Verlag, Berlin, 1988.

[8] E. Bommaiah, M. Liu, A. McAuley, and R. Talpade. AMRoute: Ad hoc Multicast Routing protocol. Internet Draft, draft-talpade-manet-amroute-00.txt, 1998.

[9] B. N. Clark, C. J. Colbourn, and D. Johnson. Unit disk graphs. *Discrete Mathematics*, 86:165–177, 1990.

[10] M. S. Corson and S. G. Batsell. A Reservation-Based Multicast (RBM) Routing Protocol for Mobile Networks: Initial Route Construciton Phase. *Wireless Networks*, 1(4):427–450, Dec. 1995.

[11] J. J. Garcia-Luna-Aceves and E. L. Madruga. The Core-assisted Mesh Protocol. *IEEE Journal on Selected Areas in Communications*, 17(8):1380–1384, Auguest 1999.

[12] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher. Range-free localization schemes for large scale sensor networks. In *ACM MOBICOM*, 2003.

[13] IEEE Computer Society. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. IEEE standard 802.11, 1999.

[14] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *ACM MOBICOM*, pages 66–80, 2003.

[15] W. C. Jakes. *Microwave Mobile Communication*. Wiley, 1974.

[16] L. Ji and M. S. Corson. A Lightweight Adaptive Multicast Algorithm. In *IEEE GLOBECOM*, 1998.

[17] S.-J. Lee, W. Su, and M. Gerla. On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks. *Mobile Networks and Applications*, 7:441–453, 2002.

[18] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris. Capacity of Ad Hoc Wireless Networks. In *ACM MOBICOM*, pages 61–69, 2001.

[19] H. Luo, R. Ramjee, P. Sinha, L. E. Li, and S. Lu. UCAN: A Unified Cellular and Ad-Hoc Network Architecture. In *ACM MOBICOM*, pages 353–367, 2003.

[20] A. Rao, C. Papadimitriou, S. Shenker, and I. Stoica. Geographic routing without location information. In *ACM MOBICOM*, pages 96–108, 2003.

[21] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.

[22] E. M. Royer and E. E. Perkins. Multicast Operation of the Ad-hoc On-demand Distance Vector Routing Protocol. In *ACM MOBICOM*, pages 207–218, 1999.

[23] N. B. Salem, L. Buttyan, J. P. Hubaux, and M. Jakobsson. A Charging and Rewarding Scheme for Packet Forwarding in Multi-hop Cellular Networks. In *ACM MOBIHOC*, pages 13–24, 2003.

[24] J. So and N. Vaidya. A multi-channel mac protocol for ad hoc wireless networks. In *UIUC Technical Report*, Jan 2003.

[25] The 3rd Generation Partnership Project 2(3GPP2). roadcast/Multicast Services Stage 1 . 3GPP2 TS-3GB-S.R0030-0v1.0, 2003.

[26] The 3rd Generation Partnership Project (3GPP). Multimedia Broadcast/Multicast Service (MBMS); Stage 2. 3GPP TR 23.846, 2003.

[27] J. Wang, R. Sinnarajah, T. Chen, Y. Wei, and E. Tiedemann. Broadcast and Multicast Services in cdma2000. *IEEE Communications Magazine*, Feb. 2004.

[28] C. Wu, Y. C. Tay, and C.-K. Toh. Ad hoc Multicast Routing Protocol Utilizing Increasing ID-numberS (AMRIS) Functional Specification. Internet Draft, draft-ietf-manet-amris-spec-00.txt, 1998.

[29] H. Wu, C. Qiao, S. De, and O. Tonguz. An integrated Cellular and Ad hoc Relaying system: iCAR. *IEEE Journal on Selected Areas in Communications*, 19(10):2105–2115, Oct. 2001.

[30] J. Yoon, M. Liu, and B. Noble. Random waypoint considered harmful. In *IEEE INFOCOM*, 2003.