ICAM: Integrated Cellular and Ad-Hoc Multicast

Randeep Bhatia†, Li (Erran) Li†, Haiyun Luo*, Ram Ramjee†
†Bell Labs, Lucent Technologies.
Emails: {randeep,erranlli,ramjee}@bell-labs.com
*Department of Computer Science, UIUC. Email: haiyun@cs.uiuc.edu

***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, a 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.**

Keywords Ad-hoc networks, Cellular networks, Multicast, Network architecture, Routing

I. Introduction

Third-generation (3G) CDMA wide-area wireless networks have experienced significant growth recently. As of April 30, 2004, the number of CDMA20001X subscribers worldwide has increased by more than 100% last year and exceeded 100 millions. At the same time the number of subscribers of CDMA2000 1xEV-DO, also known as HDR (High Data Rate), has exceeded 6.7 million [1]. As the user population 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.

Existing multicast strategy in 3G networks suffers in terms of decreased downlink channel utilization 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 3G downlink, 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 of the wireless medium, the likelihood that at least one receiver experiences bad downlink channel condition increases as the multicast group size increases, resulting in decreased throughput for the multicast group. For example, our simulation (Section III) shows that although the average downlink data rate is as high as 600Kbps for a single receiver, the throughput for a multicast group of five users decreases to around 80Kbps. With ten or more receivers the throughput drops close to the lowest achievable rate of 38.4Kbps. This phenomenon is in stark contrast to the unicast scenarios where increasing the number of users *increases* the downlink channel utilization using Proportional Fairness Scheduling [2].

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 better 3G downlink channel condition. 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 through the IEEE 802.11-based ad-hoc network to the mobile receiver. While this model has been shown to significantly improve throughput for 3G *unicast* traffic [3], 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 proxy with the highest 3G data rate connection. The 3G data rate of the proxy must be balanced by the throughput achievable over the interference-prone ad-hoc relay network*. Therefore, it is important that the choice of the proxies and the construction of the multicast forest be performed jointly with explicit awareness of the capacity limitations of the ad-hoc relay network. To this end, we characterize the interference of multihop wireless network using a general 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.

The contribution of this paper is three-fold. First, to the best of our knowledge, we are the first to propose ad-hoc multicast 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 interference in ad-hoc networks. Our algorithm and its bounds are equally applicable 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. 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. Section II reviews related work. Section III presents the motivation for ICAM model. Section IV lays down our models and assumptions. Section V presents the greedy algorithm. In Section VI, we propose an algorithm that achieves near optimal end-to-end throughput. We evaluate the performance of our implementation of the multicast relay protocols through extensive simulations in Section VII. Finally Section VIII concludes this paper.

II. Background and Related Work

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 at the highest power, 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 [4] 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 [3], relays are mobile devices themselves. UCAN requires that each mobile device be equipped with both 3G and IEEE 802.11 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 with both interfaces integrated in a single card.

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) [5], Lightweight Adaptive Multicast (LAM) [6], Ad hoc Multicast Routing Protocol (AMRoute) [7], Ad hoc Multicast Routing protocol utilizing Increasing id-numberS (AMIS) [8], and multicast extension of Ad hoc On-demand Distance Vector (MAODV) [9]. 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) [10], and On-demand Multicast Routing Protocol (ODMRP) [11]. 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 interference. 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 interference on the end-to-end throughput with proven throughput bounds. Our design carefully engineers the distribution of the multicast traffic to avoid hot-spot congestion. Moreover, by leveraging the coordination of the 3G base station, the implementation of all the aforementioned ad-hoc multicast protocols in ICAM can be significantly simpler and more responsive to highly dynamic network topologies.

III. Motivation

3G multicast is inherently limited by the worst channel rate of the multicast group. More precisely, assume that there are n clients covered by a 3G base station, and l of these clients (denoted by the set R) belong to a multicast group. 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 $\min_{v \in R} r_v^i$. Due to path loss and fast wireless channel fading, 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 throughput for the multicast group.

To quantify this effect, we simulated a 30-node network setting where all multicast receivers are static and randomly distributed over a $600 \times 600 \text{m}^2$ HDR cell. The HDR downlink channel included both slow fading (depends on user location) and fast fading components. As we can see from Figure 1 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 [12], it still falls significantly short of the achievable unicast throughput.

3G multicast's inefficiency motivates us to use relays to improve throughput. Specifically, for each multicast receiver v with 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 formal

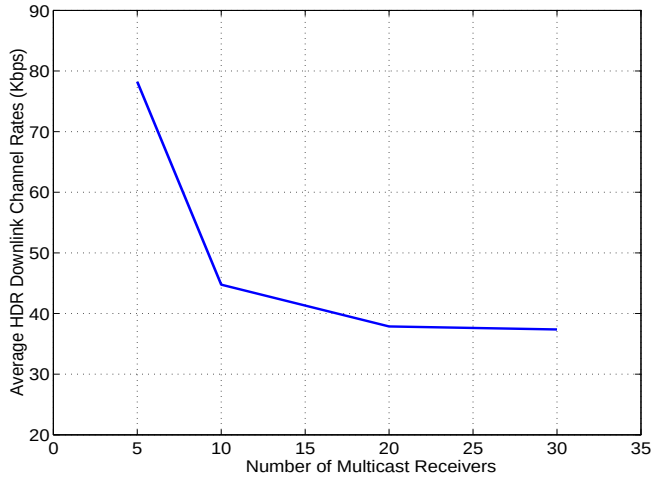

Fig. 1. Multicast Throughput

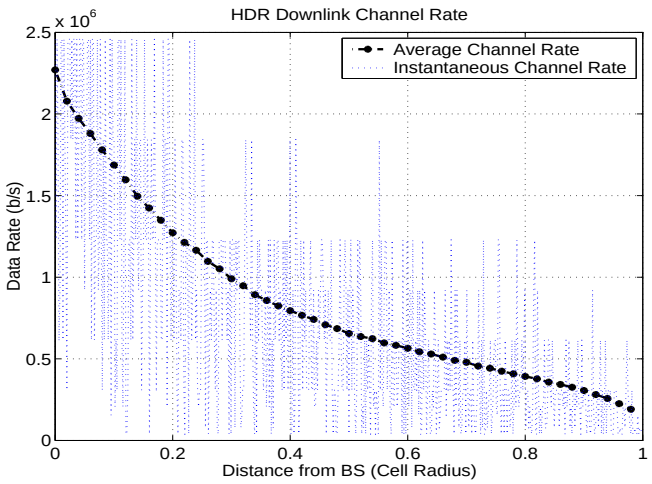

Fig. 2. HDR Downlink Instantaneous and Average Channel Rate

definition of the problem that includes the capacity constraints and interference of the ad-hoc network in Section VI-A.

IV. Models and Assumptions

The HDR downlink channel of a mobile client is modeled with both slow and fast fading. 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 [13]. 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 [14]. Figure 2 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.

A. 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 $q \geq 1$.

B. Hop limit for proxy discovery

We assume that the hop distance between a proxy and any of its receivers is upper bounded by a small number h. There are several reasons for h being small. First, due to interference, wireless channel error, and lack of scheduling, IEEE 802.11 ad-hoc network throughput decreases very fast as the number of hops increases [15]. Thus, if the ad-hoc throughput decreases to the minimal HDR downlink channel rate in less than h number of hops, the IEEE 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 breakage due to mobility, latency, overhead of proxy discovery, and routing update overhead 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 further increase throughput.

C. Minimal Separation and Location

Our near optimal algorithm (Section VI) makes two more assumptions. 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 as 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 [16]. If explicit location information is not available, the base station can compute an embedding of the connectivity graph of the ad-hoc networks [17] and our algorithm works on the embedding rather than the real coordinates.

V. Greedy Algorithm

In this section, we present our greedy algorithm that discovers proxies and establishes multicast routing entries for the distribution of packets from proxies to multicast receivers. We then introduce the unique issue of multihop wireless channel interference that limits the performance gain of the greedy ad-hoc relay. It motivates our analysis and design of the near optimal throughput multicast relay as presented in the next section.

A. 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 increased by one every time the multicast receiver initiates a new round of proxy discovery.

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 decreases the TTL of RTREQ by one. 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 downlink 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.

Note that there is no route reply messages from proxies back to the multicast receiver. Candidate proxies send the entire relay path (in RTREQ) to the base station. The base station collects and maintains topology and channel information for relay, and makes decisions in selecting proxies. If no proxy information has been established yet, the base station will simply default to 3G multicast.

B. Opportunistic relay path merging and packet relay

The HDR base station extracts the relay path from RTREQ messages and constructs/updates the partial ad-hoc network topology around the multicast receiver. The partial topology includes all the greedy relay paths for a multicast receiver. With such topology available, at each time slot t HDR base station calculates for each multicast receiver v the highest HDR downlink data rate possible through its potential proxy clients, i.e., $p(v)^t$. The base station then sets the actual HDR downlink broadcast data rate to $\min_{v \in R} p(v)^t$ for correct reception at all multicast receivers or their proxies.

Note that for each multicast receiver there may exist multiple potential proxy clients whose HDR downlink channel rates are higher than $\min_{v \in R} p(v)^t$. HDR base station ranks each relay client according to the total number of multicast receivers reachable from the relay client in h hops. Relay clients that are connected to more multicast receivers 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, since for each common link among different receiver relay paths, we only send one copy of the packet.

C. Relay path maintenance

A relay path breaks when the proxy, relay, or the multicast receiver moves out of range. When the next-hop relay client is not reachable, the IEEE 802.11b MAC layer calls a callback function to inform the relay client of such failures. 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 the lack of heart-beat messages [18]. 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.

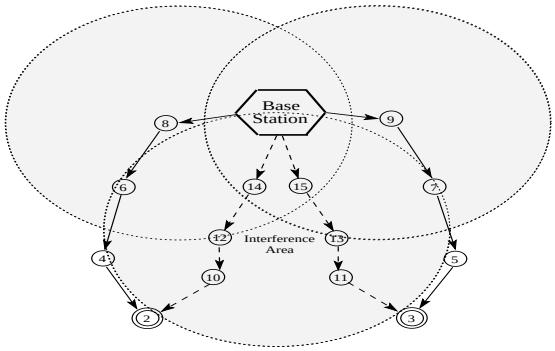

Fig. 3. Greedy Ad-Hoc Relay v.s. Optimal Relay

Client	Greedy Relay	Optimal Relay
v_2	544.2	1046.9
v_3	547.2	1041.9

TABLE I: THROUGHPUT OF TWO RELAY STRATEGIES (KBPS)

D. Impact of wireless interference

As we shall see later in Section VII, 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 does not perform well.

The primary goal of the greedy ad-hoc relay strategy is to choose the proxy with the best 3G downlink channel rate. However, since the offered load to the ad-hoc network equals to the data rate of the 3G downlink "magnified" by a factor of the number of multicast receivers, it turns out that in some situations the offered load can be higher than the capacity of the ad-hoc relay paths. We use a simple example shown in Figure 3 to illustrate this problem.

In Figure 3 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 channel rates are higher than those of clients v_8 and v_9. The greedy algorithm will discover client v_{14} and v_{15}, in pursuit of the proxy with 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

$G = (V, E)$	connectivity graph of the ad-hoc network
$I = (V, A)$	interference graph of the ad-hoc network
R_t	transmission range of each node $v \in V$
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 through 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 receives in cell $C \in \Gamma$
R	the set of multicast receivers in the multicast group
R_{3G}	a subset of R that receive the multicast data directly from 3G base station 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 II: NOTATION SUMMARY

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 $v_8 \to v_6 \to v_4 \to v_2$ and $v_9 \to v_7 \to v_5 \to v_3$ results in higher 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 I. In this scenario greedy algorithm achieves only half the throughput of the optimal relay strategy.

The reason for greedy ad-hoc relay's sub-optimal performance is that *greedily maximizing the throughput for each individual multicast receiver does not yield globally optimal throughput due to interference in the ad-hoc network*. Explicit consideration of the wireless interference among different relay paths is necessary.

VI. NEAR OPTIMAL ALGORITHM

In this section, we present a formal definition ICAM problem and a 4-approximation algorithm that runs in polynomial time. For ease of presentation, we describe our algorithm (Section VI-A and VI-B) assuming broadcast transmission and single-rate ad-hoc networks. We then extend our algorithm to multi-rate ad-hoc networks in Section VI-C. 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. Finally, we discuss how our algorithm can deal with multiple multicast groups.

A. Problem Statement

We summarize the notations in Table II. Ad-hoc network is represented as a graph $G = (V, E)$ where V is the set of (n) nodes and E the set of (m) links. If $(u, v) \in E$, then node u and v are at most R_t apart. The set $R \subseteq V$ is the set of receivers of a given 3G multicast group. The average 3G downlink rate of node $v \in V$ is $r_v \geq 0$. A receiver $v \in R$ may receive data either directly from the $3G$ base station – at rate r_v, or through a proxy. The ad-hoc relay subnetwork for a given multicast group is graph $G' = (V', E')$, where $(u, v) \in E'$ iff $u, v \in V'$ and $(u, v) \in E$. For example, one relay subnetwork for multicast in Figure 3 consists of node set $V' = \{v_8, v_6, v_4, v_2, v_9, v_7, v_5, v_3\}$ and link set $E' = \{(v_8, v_6), (v_6, v_4), (v_4, v_2), (v_9, v_7), (v_7, v_5), (v_5, v_3)\}$.

Relay subnetwork G' is composed of a (forest) collection of directed trees F spanning node set V'. Along each tree $T \in F$, data (received from the $3G$ base station) is multicasted from the root (proxy) to all the receivers in T. Receivers in $R - R_{3G}$ receive data at the rate of $\max\{r_v, r_v^a\}$, where r_v^a is the rate at which data is received through the ad-hoc relay subnetwork G'. For a given G' we denote $\max\{r_v, r_v^a\}$ as $r_v(G')$.

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 $(u, v) \in A$[1]. $(u, v) \notin A$ if u and v are at least qR_t apart for some fixed constant q. Given $G' = (V', E')$, let $k(G')$ denote the minimum number of colors required to color the nodes[2] in V' such that two nodes $u, v \in V'$ have the same color iff $(u, v) \notin A$. Therefore, the best multicast rate that can be achieved in G' is at most $f/k(G')$. To be precise, only non-leaf nodes need to be colored since the leaf nodes (receivers) do not participate in 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 in Section IV-B, we make the assumption that the best proxy for receiver v is no more than h hops away for some small value h, e.g., $h = 4$. For 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 G' such that the multicast rate $r(G')$ for its associated ad-hoc subnetwork G' is maximized. The ICAM problem is NP-hard. Due to lack of space we leave the detailed reduction in [19].

B. Approximation Algorithm

The basic idea behind our polynomial-time approximation algorithm is to leverage the property that for each multicast receiver the best proxy can be at most h-hops away (see Section IV-B). We perform localized proxy search and construct the optimal multicast trees for the receivers within certain locality. All local optimal multicast trees are then combined to form the global multicast forest.

Figure 4 illustrates our approximation algorithm ALGO for ICAM . It first divides the coverage area of a 3G base station into a 2-dimensional grid of $(2h + q + \epsilon)R_t \times (2h + q + \epsilon)R_t$ cells, where $\epsilon > 0$. The rationale for this choice of the grid

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

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

$computeMcastForest()$
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: the set of grid cells
3. V'_C = computeOptCellForest(C);
4. // Merge the solution of each cell
5. $V_O = \cup_{C\in\Gamma} V'_C$
6. $R_{3G} = R \cap (V - V_O)$

Fig. 4. Algorithm ALGO for Computing Multicast Relay Forest

size is described later. ALGO then computes a solution for each cell of the grid Γ that contains at least one receiver. In other words, ALGO computes the optimal solution for ICAM when restricted only to the receivers $R_C \subseteq R$ in a cell $C \in \Gamma$. Finally ALGO merges these solutions for all cells to compute a feasible solution to the original instance of the problem.

Let C be a cell with at least one receiver ($|R_C| > 0$), and $V_C \subseteq V$ 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 relay nodes must be in V_C. Algorithm ALGO computes the optimal solution for a cell C as follows (see Figure 5). 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 VI-A). 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 VI-A. 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 $r_v(G'_C) = r_v$. Taking the minimum of $r_v(G'_C)$ for all receivers $v \in R_C$, ALGO is able to compute the multicast rate $r(G'_C)$ for the ad-hoc relay subnetwork for the receivers in R_C. ALGO then selects the subset $V'_C \subseteq V_C$, whose associated ad-hoc relay subnetwork has the highest rate, to generate the optimal multicast strategy for the receivers R_C in cell C. We will show later that all this can be done in constant time. Finally ALGO outputs 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 $R \cap (V - V_O)$.

$computeOptCellForest(C)$
1. Enumerate all subsets V'_C of V_C // V_C is the set of // nodes that are within h hops from receivers in 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 $rate$

Fig. 5. Algorithm for Computing Optimal Multicast Relay Forest in one Grid Cell

We now show that ALGO runs in polynomial time and achieves an approximation factor of 4.

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 brute force or the linear algorithm for coloring graphs with bounded tree-width [20]). Thus, overall the algorithm runs in polynomial time. We remark that, in our practical setting, the number of 3G subscribers $|V_c|$ in any given time within a grid cell is a small number. Therefore, the time to compute an optimal coloring for each grid cell is managible in practice. ■

Let $r(G'_C)$ be the optimal multicast rate of the 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 achieves a higher throughput via the ad-hoc relay subnetwork G'_C than via its 3G downlink channel. 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 $r(G) \leq r(G'_C) \leq \min\{f/k_C, p(v)\} \leq f/k_C$. ■

We assume V_O as the solution output by ALGO 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 $4\max_{C\in\Gamma} k_C$ 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 (see Figure 6). Denote $f(C)$ as the color for cell C. 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

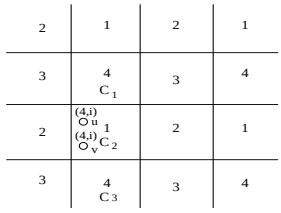

Fig. 6. Grid can be colored with 4 colors

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 $G_O = (V_O, E_O)$, since nodes that receive the same color are strictly more than qR_t apart. By construction this coloring uses at most $4 \max_{C \in \Gamma} k_C$ colors. ∎

For example, in Figure 6, node u and v may receive the same color $(4, i)$ from the solution of cell C_1 and C_3 respectively. u is at most hR_t away from the boundary of cell C_1, while 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.

We finally are ready to prove the approximation bound for ALGO as follows.

Theorem 5: ALGO is a 4-approximation for ICAM

Proof: Let $k = \max_{C \in \Gamma} k_C$. Consider a receiver $v \in V_O$ in cell C. Let $r(G)$ be the optimal multicast rate. Let G'_C denote the best ad-hoc relay subnetwork computed by ALGO 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

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

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 $r(G)/4 \leq r_v(G_O)$. ∎

C. 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. We leave formal proofs in [19] due to the lack of space. For ease of presentation we describe these extension assuming broadcast transmissions, but our results easily carry over to networks (e.g. current 802.11 networks) where unicast transmission is used.

Note that given a multicast tree T for a given cell C the best broadcast schedule may assign different number of broadcast slots (colors) to the nodes of T. Let $f_v(T)$ be the rate when node v broadcasts to its downstream receivers in T. Then a node with relatively low $f_v(T)$ must broadcast more often to ensure that the end-to-end throughput. ALGO is modified to enumerate the possible slot assignments for the nodes in T. Since the number of such possible slot assignments can be infinitely large, ALGO must carefully select a small set of possible slot assignments. We control the size of the search space by a user specified error tolerance ϵ: the smaller ϵ is, the larger the search space and the better the throughput of the computed solution. ALGO then 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 $\lceil (L+1)/\epsilon \rceil$ 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). Given $k_v(T)$ for v of T, ALGO finds an minimum coloring of the interference graph induced by 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 throughput achievable in this slot assignment is computed as $\min_{v \in T} f_v(T) k_v(T)/K(T)$. 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).

Algorithm ALGO merges the solutions of all the cells C similarly to the single-rate case. The proof for the overall performance bound is along the line of the single-rate case with the following major modifications. Let K_C be the total number of colors and $k_C(v)$ be the number of colors assigned to node v 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 $(1 \leq i \leq K_C)$ with the K/K_C colors $i, i + K_C, i + 2K_C, \ldots i + (K/K_C - 1)K_C$. Now every node v has $k_C(v)K/K_C$ colors. The new coloring uses exactly K total number of colors and is a valid coloring for the solution for cell C. As shown in Lemma 4 this implies that the merged solution can be colored with at most $4K$ colors. Note that node v for cell C has at least $(k_C(v)K/K_C)/4K = (k_C(v)/K_C)/4$ fraction of the total number of colors, 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 6: ALGO is a $4(1 + \epsilon)$ polynomial time approximation scheme for ICAM with multi-rate ad-hoc networks.

D. Multiple Multicast Groups

While ALGO provides the near-optimal multicast relay strategy for a single multicast group, it can be applied independently to multiple multicast groups. Recall that HDR transmits frames in the downlink in a time-multiplexed fashion with time slots of 1.67ms each. 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

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

TABLE III: THROUGHPUT COMPARISON

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

TABLE IV: HDR DOWNLINK CHANNEL RATE, GOODPUT, AND IEEE 802.11 INTERFACE ENERGY CONSUMPTION

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 VII-C.4 for the multiple multicast groups further confirm this claim.

VII. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the greedy and the near-optimal relay. We first present the metrics and methodology. We then compare the performance of the greedy and near-optimal algorithms in networks with stationary nodes. Finally we show the performance of the greedy algorithm in mobile scenarios, investigating the impact of the multicast group size, node density, node mobility, and multiple multicast groups.

A. Metrics and Methodology

We implement the greedy and near-optimal relay in the *ns-2* simulator. The HDR downlink channel is modeled according to the published experimental data in [2], [14] (see Section III). We use the IEEE 802.11b implementation in ns-2 version *2.1b9a* where 11Mbps data rate is supported at 115-meter communication range. Due to the unreliable nature of 802.11 MAC broadcast, we use 802.11 unicast for wireless media access. 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. The node mobility is set according to the improved random waypoint model [21]. We set the pause time between consecutive moves to 3 seconds, and vary the speed in the range of $[0.5, 1]\times \text{Speed}_{\max}$. 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. **Average HDR downlink data rate** evaluates the performance of our proxy discovery. 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 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**.

B. Greedy v.s. 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 3 and Table I), 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 600×600m^2 with the HDR base station located in the center. We randomly place 5 multicast receivers in the cell. All results are the average over 20 random topologies generated by the `setdest` tool [21].

Table III 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 IV 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.

C. Mobile Scenarios

In this section we study the performance of the greedy algorithm in mobile scenarios with different multicast group sizes, node densities, mobility speeds, and multiple multicast groups. As a future work, we discuss in Section VIII on how to speed up the computation of the near-optimal algorithm for highly mobile scenarios. We increase the size of the cell to 886×886m^2 (approximating a 500-meter radius circular cell with the HDR base station located in the center) in order to accommodate more mobile nodes.

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 [2]. We set the maximum moving speed as 15m/s. Figure 7-11 show the simulation results. As we can see from Figure 7,

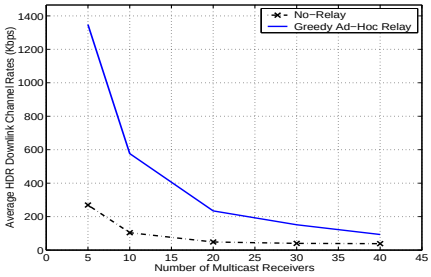

Fig. 7. Group Size Impact: Average HDR downlink channel rate

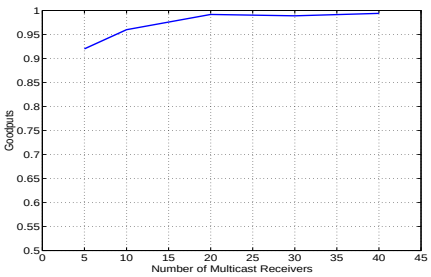

Fig. 8. Group Size Impact: Goodput

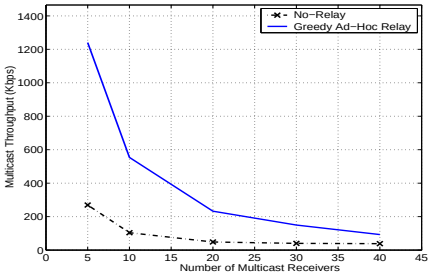

Fig. 9. Group Size Impact: Throughput

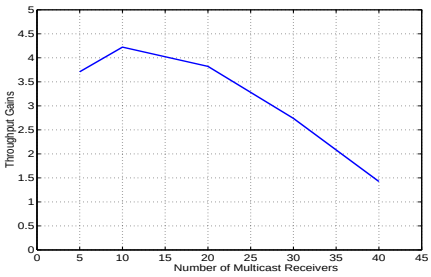

Fig. 10. Group Size Impact: Throughput Gain

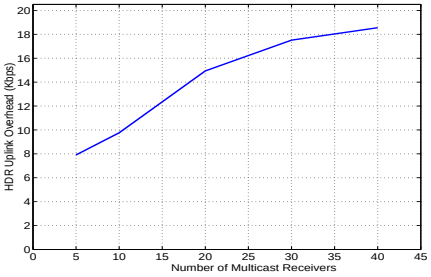

Fig. 11. Group Size Impact: HDR Uplink Overhead

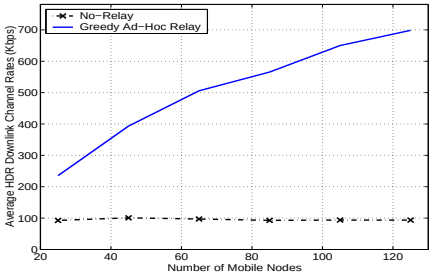

Fig. 12. Node Density Impact: Average HDR downlink channel rate

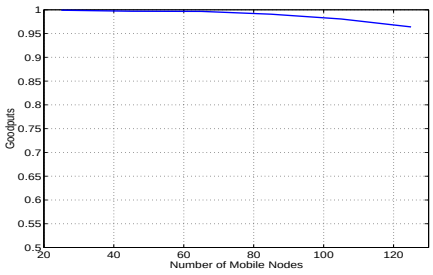

Fig. 13. Node Density Impact: Goodput

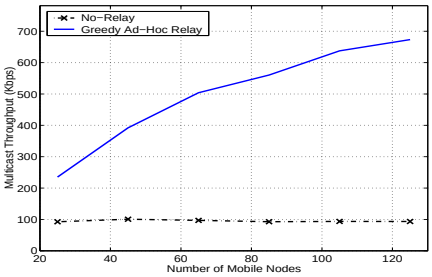

Fig. 14. Node Density Impact: Throughput

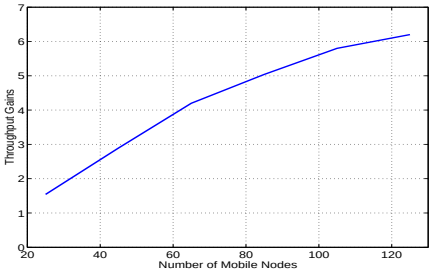

Fig. 15. Node Density Impact: Throughput Gain

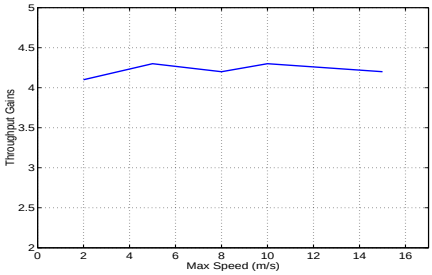

Fig. 16. Mobility Impact: Throughput Gain

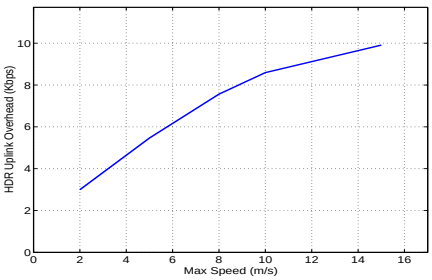

Fig. 17. Mobility Impact: Uplink Overhead

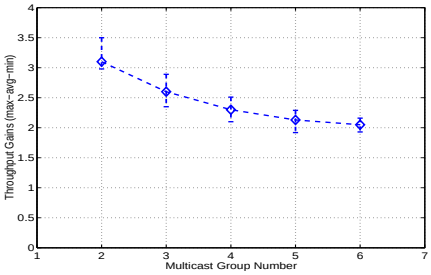

Fig. 18. Multiple Groups: Throughput Gain

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 VII-B, since we have increased the node density (see density impact next).

Figure 7 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 8 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∼420%, as shown in Figure 9 and 10. Note that the throughput gain with the multicast group size of 5 nodes is slightly smaller than the throughput gain with multicast 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,

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 V: THROUGHPUT OF MULTIPLE MULTICAST GROUPS

and the lower goodput in the scenarios of multicast group size 5.

Figure 11 shows that the overhead on HDR uplink increases as the multicast group size increases. In all scenarios, the greedy algorithm consumes less than 10% of the HDR uplink's 153.6Kbps bandwidth [2].

2) Node density: In order to study the impact of the node density of 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 12-15. From Figure 12 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 13, decreases slightly as the network density increases, mainly due to the increased offered load (the higher average HDR downlink channel rate). As Figure 14 and 15 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.

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 16 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 17) increases due to the increased number of link breakage reports.

4) Multiple multicast groups: We finally show the throughput gains for multiple multicast groups simultaneously receiving packets from the HDR 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. We set the maximum moving speed at 5m/s. Note that, with six multicast groups instead of one, HDR proportional fairness scheduler [2], [3] is able to exploit user diversity with 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 V). Therefore, even though the throughput gain decreases from 311% to 205% as the number of groups increase in Figure 18, the absolute throughput achieved using ad-hoc relay remains almost the same as shown in Table V. This result confirms our earlier claim in Section VI-D that since all multicast groups take turns sharing the HDR downlink bandwidth, there is not additional channel contention in the ad-hoc network introduced by multiple multicast groups.

VIII. 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 presented two novel algorithms to determine the set of proxies and the relay paths: one based on greedy heuristic and the other a near-optimal algorithm that runs in polynomial time. For our near-optimal algorithm, we prove the performance bound that holds when the underlying wireless MAC supports broadcast or unicast, single rate or multiple rates, and even when there are multiple simultaneous multicast sessions. Through extensive simulations, we showed that the near-optimal algorithm improves the average 3G multicast throughput by up to 840% and outperforms the greedy heuristics up to 92% in static scenarios, while the greedy algorithm achieves throughput gains of as much as 410% in relatively high mobility scenarios.

ACKNOWLEDGMENT

We thank Dr. Sanjoy Paul for initial discussions on the problem.

REFERENCES

[1] "3G Today." `http://www.3gtoday.com/`.
[2] "1xEV: 1x EVolution IS-856 TIA/EIA Standard - Airlink Overview." QUALCOMM Inc. White Paper, Nov. 2001.
[3] H. Luo, R. Ramjee, P. Sinha, L. E. Li, and S. Lu, "UCAN: A Unified Cellular and Ad-Hoc Network Architecture," in *Proceedings of ACM MOBICOM*, pp. 353–367, 2003.
[4] H. Wu, C. Qiao, S. De, and O. Tonguz, "An integrated Cellular and Ad hoc Relaying system: iCAR," *Proceedings of IEEE Journal on Selected Areas in Communications*, vol. 19, pp. 2105–2115, Oct. 2001.
[5] M. S. Corson and S. G. Batsell, "A Reservation-Based Multicast (RBM) Routing Protocol for Mobile Networks: Initial Route Construciton Phase," *Wireless Networks*, vol. 1, pp. 427–450, Dec. 1995.
[6] L. Ji and M. S. Corson, "A Lightweight Adaptive Multicast Algorithm," in *Proceedings of IEEE GLOBECOM*, pp. 1036–1042, 1998.
[7] E. Bommaiah, M. Liu, A. McAuley, and R. Talpade, "AMRoute: Ad hoc Multicast Routing protocol." Internet Draft, draft-talpade-manet-amroute-00.txt, 1998.
[8] 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.
[9] E. M. Royer and E. E. Perkins, "Multicast Operation of the Ad-hoc On-demand Distance Vector Routing Protocol," in *Proceedings of ACM MOBICOM*, pp. 207–218, 1999.
[10] J. J. Garcia-Luna-Aceves and E. L. Madruga, "The Core-assisted Mesh Protocol," *IEEE Journal on Selected Areas in Communications*, vol. 17, pp. 1380–1384, Auguest 1999.
[11] S.-J. Lee, W. Su, and M. Gerla, "On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks," *Mobile Networks and Applications*, vol. 7, pp. 441–453, 2002.
[12] P. Agashe, R. Rezaiifar, and P. Bender, "CDMA2000 High Rate Broadcast Packet Data Air Interface Design," *IEEE Communications Magazine*, pp. 83–89, Feb. 2004.
[13] W. C. Jakes, *Microwave Mobile Communication*. Wiley, 1974.
[14] 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*, vol. 38, pp. 70–77, Jul. 2000.

[15] J. Li, C. Blake, D. S. J. De Couto, H. Lee, and R. Morris, "Capacity of Ad Hoc Wireless Networks," in *Proceedings of ACM MOBICOM*, pp. 61–69, 2001.
[16] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher, "Range-free localization schemes for large scale sensor networks," in *Proceedings of ACM MOBICOM*, pp. 81–95, 2003.
[17] A. Rao, C. Papadimitriou, S. Shenker, and I. Stoica, "Geographic routing without location information," in *Proceedings of ACM MOBICOM*, pp. 96–108, 2003.
[18] S. Marti, T. Giuli, K. Lai, and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," in *Proceedings of ACM MOBICOM*, pp. 255–265, 2000.
[19] R. Bhatia, L. E. Li, H. Luo, and R. Ramjee, "ICAM: Integrated Cellular and Ad-Hoc Multicast," tech. rep., UIUC Computer Science, November 2004.
[20] H. Bodlaender, "Dynamic programming on graphs of bounded treewidth," in *Proceedings of 15th International Colloquium on Automata, Languages and Programming*, pp. 631–643, Lecture Notes in Computer Science, 317, Springer-Verlag, Berlin, 1988.
[21] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in *Proceedings of IEEE INFOCOM*, pp. 1312–1321, 2003.