Clustering Overhead for Hierarchical Routing in Mobile Ad hoc Networks

John Sucec and Ivan Marsic
Rutgers University

***Abstract*--Numerous clustering algorithms have been proposed that can support routing in mobile ad hoc networks (MANETs). However, there is very little formal analysis that considers the communication overhead incurred by these procedures. Further, there is no published investigation of the overhead associated with the recursive application of clustering algorithms to support hierarchical routing.**

This paper provides a theoretical upper bound on the communication overhead incurred by a particular clustering algorithm for hierarchical routing in MANETs. It is demonstrated that, given reasonable assumptions, the average clustering overhead generated per node per second is only *polylogarithmic* in the node count. To derive this result, novel techniques to assess cluster maintenance overhead are employed.

***Index terms*--Mobile ad hoc network, hierarchical routing**

I. INTRODUCTION

Mobile ad hoc networks (MANETs) are comprised of mobile nodes that perform multiple hop datagram forwarding over wireless links. The mobility of network nodes combined with the transient nature of wireless links results in a rapidly changing network topology. The dynamic nature of the network environment arguably makes the task of routing in MANETs far more difficult than in wired networks. Further, it is commonly assumed for MANETs that the wireless links tend to be relatively low capacity fixed-sized links (i.e., no hierarchy in the physical topology of the network). This means that neither traffic aggregation nor summation of routing information can be achieved through hierarchically proportioned physical links. Thus, not only is maintaining and acquiring routing information in MANETs difficult to achieve but so is achieving this in a manner that scales well with increasing network size.

This paper addresses the scalability, with respect to increasing node count, of *hierarchical routing* in MANETs. The performance metric under consideration is the *control overhead per node* (ψ) required by hierarchical routing. This assessment considers only the overhead due to the maintenance of routing tables and hierarchical clustering. The overhead due to location (or address) management is considered elsewhere. The objective here is to express ψ as a function of $|V|$, where V is the set of network nodes. The finding of this paper is that with some reasonable constraints on the network, $\psi = O(\log^2|V|)$ bits per second per node.

Numerous papers have been published on hierarchical routing. Among these include [1], [7], [8], [9], [12], [13] and [16]. Although a number of these papers provide detailed assessments of hierarchical routing performance, only [8] and [16] attempt to quantify analytically control packet overhead. In [16], control packet overhead required for constructing routing tables in a *two*-level hierarchically organized network is considered. However, the analysis was performed to address primarily the overhead of updates due to link cost changes and did not address the control packet overhead incurred by node mobility. In [8], the control packet overhead required for routing table maintenance is also considered but for a three-level hierarchical network. However, the assessment in [8] does not consider the case of $\Theta(\log|V|)$ hierarchical levels or attempt to explicitly bound ψ as a function of $|V|$.

Scalability performance metrics, considered elsewhere, include the ratio of hierarchical path length to least-hop path length and routing table storage overhead. Although these metrics are important, control packet overhead is of chief interest here. The justification for focusing on ψ is as follows. One, control packet overhead is arguably more critical than routing table size because scarce wireless link capacity poses a more severe performance limit than the available memory in today's computers. Two, whereas extensive earlier work exists that analyzes hierarchical path lengths (e.g., [17]), little analysis has been published that assesses ψ.

The remainder of this work is organized as follows. In Section II the essential features of the network environment are described. Section III provides an overview of hierarchical routing and hierarchical cluster formation techniques. Section IV presents a detailed analysis of the overhead incurred by hierarchical routing operations. Conclusions on the results of Section IV are provided in Section V.

II. NETWORK ENVIRONMENT

The underlying physical topology of a MANET is represented here by a *connected*, undirected graph, $G = (V,E)$, where V is the set of nodes and E is the set of bi-directional links. It is assumed that, at any time, nodes are situated randomly throughout a fixed size area (A) in accordance with a two-dimensional uniform random variable distribution. For

This work was supported in part by U.S. Army CECOM contract number DAAB07-00-D-G505.

the purpose of analyzing the frequency of location update events, the random waypoint model for node mobility, employed in [2], with zero pause time is assumed here.

Each node is equipped with a single network interface card (NIC) having a transmission radius of R_{TX} m. If the distance separating a pair of nodes is less than R_{TX}, then a bi-directional link exists between them and they are considered to be neighbors of one another. Otherwise, the nodes are not connected. Each NIC employs carrier sense multiple access with collision avoidance (CSMA/CA) so that each node operates in a shared *broadcast* media with its neighbors.

The scalability of a routing protocol may be assessed in terms of a number of distinct criteria. Among these include scalability with respect to increasing node count ($|V|$), increasing average node density (nodes per unit area) and increasing average node speed (m/s). In order to isolate the performance of hierarchical routing with respect to increasing $|V|$, it is assumed that average node density and average node speed are held constant.

Lastly, it is shown in [10] that the average hop count on the shortest path between an arbitrary pair of nodes in two-dimensional network (e.g., Fig. 1) is $\Theta\left(\sqrt{|V|}\right)$. As noted in [14], to maintain connectivity in random graphs, R_{TX} must be $\Theta\left(\sqrt{\log|V|}\right)$. Thus, for random graphs average hop count is actually $\Theta\left(\sqrt{|V|/\log|V|}\right)$. However, the $\log|V|$ term that appears in the expression for average hop count will be ignored here for the sake of compactness of notation and, therefore, R_{TX} is assumed constant and the $\Theta\left(\sqrt{|V|}\right)$ result given in [10] is employed here, instead.

III. Hierarchical Routing Overview

A. *Hierarchical Principles*

Fig. 1 illustrates the fundamental concept of a clustered hierarchy. All network nodes (i.e., V) are level-0 clusters. Level-0 clusters organize themselves into level-1 clusters, via some clusterhead election process such as one of the methods described in [3]. The level-1 clusterheads, in turn, organize themselves into level-2 clusters. That is, a level-k node which is elected as the clusterhead for a level-k cluster becomes a level-(k+1) node. This clustering procedure is performed recursively until the desired number of cluster levels have been constructed.

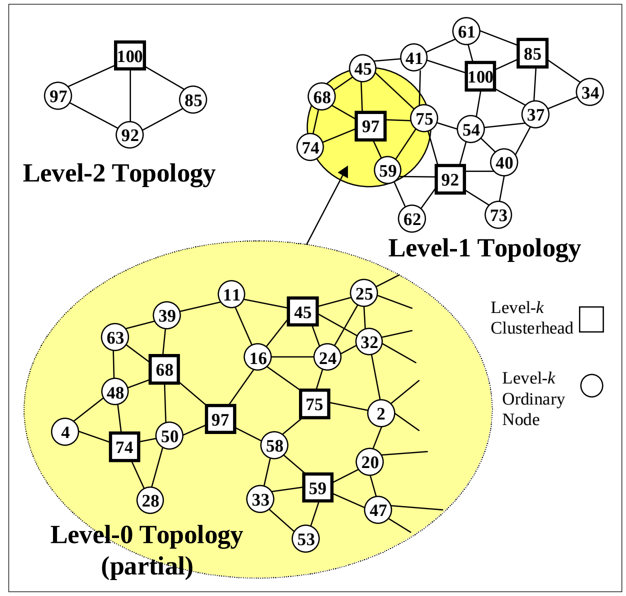

Fig. 1. Example of 3-level hierarchy.

Hierarchical routing has long been known to be instrumental in affording scalability in computer networks. Reference [1] shows that through hierarchical clustering of network nodes, the average size of the routing table maintained at each node is $L\times|V|^{1/L}$. Here, L is the number of levels in the clustered hierarchy and $|V|^{1/L}$ corresponds to the *arity* of the hierarchical tree. That is, the number of level-(k–1) clusters divided by the number of level-k clusters, $k \in \{1,2,\ldots,L\}$. Letting some constant $\beta = |V|^{1/L}$ be the arity, the average table size at each node is $O\left(\beta\cdot\log_\beta|V|\right)=O\left(\log_\beta|V|\right)$. (Here forward, the base of the logarithm terms is assumed to be some constant and is omitted from expressions.) On the other hand, letting $L = \log|V|$ results in *increased average path length*. For example in torus networks, it is shown in [1] that using the minimum size routing table incurs an increase in average path length by as much as a factor of 6. This effectively reduces network throughput by a factor of 6. However, this factor is invariant in the node count. Therefore, in terms of scalability, it is arguably more crucial to reduce overheads that grow with increasing network size. Hence, the interest in hierarchical routing overhead.

The principles of hierarchical routing have seen application in military-based packet radio networks, such as the Survivable Packet Radio Network (SURAN) described in [7] and [8]. More recently, the Hierarchical State Routing (HSR) protocol proposed in [9,12] and multimedia support for mobile wireless networks (MMWN) proposed in [13] represent hierarchical approaches designed to support group mobility and multimedia, respectively, in the MANET environment.

The analysis of this paper assumes *strict hierarchical routing*, based on the description provided in [11], to be in effect. HSR and MMWN are examples of strict hierarchical routing implementations. An implementation recommended for SURAN in [8] also falls into this category. Two important concepts concerning packet forwarding in hierarchical networks is that packet forwarding decisions are made solely on the hierarchical address of the destination node and every node has a $\Theta(\log|V|)$ hierarchical map for the clusters of the network hierarchy to which it belongs. This means that forwarding of user packets need *not* be directed through clusterheads and are forwarded via clusterhead and/or non-clusterhead nodes along the shortest hierarchical path to the destination, thereby, preventing the creation of hot spots. Detailed knowledge of how packet forwarding is implemented in hierarchical networks is not essential for understanding this paper and, therefore, further discussion of this topic is omitted here.

B. Clustering Techniques

A number of clustering schemes have been proposed in previous literature (e.g., [3], [4], [5] and [6]). Of particular interest here are the max-min D-hop clustering strategy of [3] and the linked cluster algorithm (LCA) of [4]. Each of these approaches is an ID-based clustering technique. The max-min D-hop strategy is shown to converge in $O(D)$ rounds and generates only $O(D)$ messages per node. It represents, therefore, a scalable clustering procedure. The 1-hop clustering case is equivalent to an *asynchronous* version of the LCA. It is an asynchronous version of the LCA that is assumed to be in effect for election of level-k clusterheads, $k \in \{1,2,\ldots,L\}$, known here as asynchronous LCA (ALCA).

Based on the analysis of [3], formation of 1-hop clusters requires only 2 rounds of communication. This fact is significant for bounding the overhead required for cluster maintenance. That is, once level-k clusterheads are initially elected by the ALCA, $k \in \{1,2,\ldots,L\}$, if a single perturbation to network topology (e.g., clusterhead death) triggers a clusterhead reorganization, then reorganization will incur only two rounds of messaging. This means that the messaging required to react to a single topology perturbation is confined to a two-hop radius about the location of the event. Thus, the impact of a level-k topology change on ALCA cluster maintenance has only a local effect on the level-k topology, unlike some of the other clustering approaches where a single level-k perturbation can subsequently effect the entire level-k cluster topology.

To better understand the ALCA, the ALCA election process is described briefly. Essentially, a level-k node v_k is elected as a level-k clusterhead by a neighbor u_k if its node ID v_k is the largest among all nodes in the closed neighborhood of u_k (i.e., the union of u_k and its level-k neighbors). For example, in the level-0 topology of Fig. 1, node 97 is elected to serve as a clusterhead because it is the largest node in its neighborhood. As another example, node 68 is also elected to serve as a clusterhead because it has the largest node ID in the level-0 neighborhood of node 63, even though 68 is *not* the largest node in its own level-0 neighborhood. The recursive application of this election process is illustrated in Fig. 1 by the level-1 and level-2 topologies. Thus, yielding a 3-level clustered hierarchy for this example network.

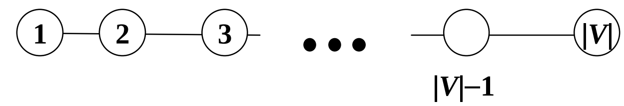

Fig. 2. Worst case node arrangement for the ALCA.

A brief discussion of hierarchical cluster links is now given. As illustrated by the level-1 and level-2 topologies, level-(k+1) cluster links may be due to one of three possible level-k cluster link effects. For example, node 100 has level-2 cluster links between itself and each of the other three level-2 nodes. Each cluster link, however, is due to a different topology effect. The cluster link between node 100 and 97 is due to a pair of level-1 gateway nodes (e.g., 54 and 75). The cluster link between nodes 100 and 92 is due to a single level-1 gateway node (i.e., 54). Lastly, it is possible that the level-k clusterheads may be level-(k−1) neighbors as is the case for nodes 100 and 85. In this case, no level-1 gateway node is required.

As an aside, it is noted that an advantage of the clustering approach of [6] over the ALCA is that it handles better the pathological network topology of Fig. 2. That is, it will organize the network of Fig. 2 into $\lceil |V|/2 \rceil$ level-1 clusters. The ALCA, on the other hand, would form $|V|-1$ level-1 clusters. Simulation results, however, indicate that recursive application of the ALCA results in hierarchical trees with a suitable average arity for random two-dimensional graphs. Although the approach of [6] may have additional advantages, the $\Theta(1)$ bound on cluster formation time is crucial for the analysis of this paper, making the ALCA a logical choice for clustering overhead assessment.

IV. Communication Overhead

Communication overhead in hierarchically organized networks may result from the following phenomenon:

- Hello protocol (ψ_{HELLO})
- Level-k cluster formation and cluster maintenance messaging, $k \in \{1,2,\ldots,L\}$ (ψ_{CL})
- Flooding of the L-level hierarchy to cluster members (ψ_{FLOOD})
- Addressing information required in datagram headers (ψ_{DATA})
- Location registration events (ψ_{REG})
  - Due to changes in the clustered hierarchy
  - Due to node mobility between clusters
- Handoff or transfer of location management data (ψ_{XFER})
  - Due to changes in the clusterhead hierarchy
  - Due to node mobility between clusters
- Location query events (ψ_{QRY})

Total communication overhead *per node* in hierarchically organized networks is the sum of the above contributing elements. The last three of the above factors, ψ_{REG}, ψ_{XFER}

and ψ_{QRY}, correspond to overhead associated with location management (or, equivalently, address management) required for hierarchical routing. Assessing the communication overhead due to location management is beyond the scope of this paper and will not be considered here. However, an evaluation of location registration overhead is given in [15] and a result of $\Theta(\log|V|)$ packet transmissions per node is derived for hierarchical location management when $L = \Theta(\log|V|)$. A theoretical handoff framework is assessed, in [18], for which it is found that handoff incurs average overhead per node that is polylogarithmic in the node count. Also discussed in [15], is the overhead due to location queries. ψ_{QRY} depends on the query frequency per node (f_{QRY}) and the average hop distance from a querying node to the location management server of the target node. Assuming an average hop distance of $\Theta\left(\sqrt{|V|}\right)$ between an arbitrary pair of nodes and assuming $f_{QRY} = \Theta(1)$, then $\psi_{QRY} = \Theta\left(\sqrt{|V|}\right)$.

This paper considers the overhead required for correct packet forwarding and the construction and maintenance of packet forwarding tables (i.e., the overhead incurred by routing). This corresponds to an assessment of the first four of the above factors, ψ_{HELLO}, ψ_{CL}, ψ_{FLOOD} and ψ_{DATA}. For this purpose, the following claims are made:

Claim 1. The Hello packet transmission count (ψ_{HELLO}) is $\Theta(1)$ per node. Justification for this claim is provided in Section IV-A.

Claim 2. The packet transmission count due to $\Theta(\log|V|)$ levels of cluster formation ($\psi_{CL\text{-}F}$) is $O(\log|V|)$ per node. Justification for this claim is provided in Section IV-C.

Claim 3. The packet packet transmission count due to cluster maintenance ($\psi_{CL\text{-}M}$) is $O(\log|V|)$ per node. Justification for this claim is provided in Section IV-D.

Claim 4. The packet transmission count due to flooding of the hierarchical map (ψ_{FLOOD}) is $\Theta(\log|V|)$ per node. Justification for this claim is provided in Section IV-E.

Claim 5. The overhead due to hierarchical addressing (ψ_{DATA}) is $\Theta(\log|V|)$ per *datagram*. Justification for this claim is provided in Section IV-F.

A. Hello Protocol Overhead

The Hello protocol is employed for nodes to learn and verify adjacencies. Discovery of adjacencies can be facilitated by periodic broadcast of a single Hello message over the shared CSMA/CA transmission media. This is sufficient for a node v to announce itself to all nodes within R_{TX} of it. Once an adjacency is discovered, robust exchange of neighborhood data can be facilitated by periodic unicast communication of Hello messages between the two neighbors via some collision avoidance handshaking sequence. This is done, on average, with n_0 neighbors of v, where n_0 is the average of level-0 neighbors of a node (i.e., the number nodes within R_{TX} of v). Since it has been assumed in Section II that average node density is constant with respect to increasing $|V|$, $n_0 = \Theta(1)$. Including also the periodic broadcast of the Hello message, the total number of Hello message transmissions per Hello interval is $1+n_0 = \Theta(1)$.

The frequency of Hello messages is proportional to the average node movement rate μ and inversely proportional to the transmission radius R_{TX}. That is, $f_{HELLO} \propto \mu/R_{TX}$. Clearly, f_{HELLO} can be bounded from above by some constant that accounts for all realistic values of μ and, therefore, is independent of $|V|$. Thus, $f_{HELLO} = \Theta(1)$. Combining this fact with $1+n_0 = \Theta(1)$ transmissions per Hello interval means ψ_{HELLO} is $\Theta(1)$, as per Claim 1.

B. Definitions and Clustering Assumptions

The following additional definitions are useful here.

- $n_k \equiv$ The average number of level-k neighbors for a level-k node, $k \in \{0,1,\ldots,L\}$.
- $n_{max} \equiv \max\{n_0,n_1,\ldots,n_L\}$
- $f_k \equiv$ The average frequency at which level-k cluster link state changes occur, $k \in \{0,1,\ldots,L\}$. In other words, the frequency at which either a pair of level-k one-hop neighbors become two-hop neighbors (cluster link deleted) or a pair of level-k two-hop neighbors become one-hop neighbors (cluster link created).
- $\beta_k \equiv$ Average *arity* at level-k of the hierarchical tree, $k \in \{1,2,\ldots,L\}$. That is, the average number of level-(k–1) nodes per level-k cluster.
- $D_k \equiv$ The average distance, in terms of level-0 hops, separating level-k nodes (i.e., level-(k–1) clusterheads), $k \in \{1,2,\ldots L\}$.
- $p_k \equiv$ The probability that an arbitrary level-(k–1) cluster link state change effects the state of a level-k cluster link, $p_0 = 1$ and $0 < p_k \le 1$ for $k \in \{1,2,\ldots,L\}$.
- $l_k \equiv$ Average number of level-k cluster links effected given that a level-(k–1) cluster link state change impacts level-k, $l_0 = 1$.

Prior to assessing the overhead due to hierarchical clustering, some conditions on the network environment, in addition to those of Section 2, are stated. These assumptions are made to simplify the analysis.

i. Clusterhead election is performed recursively to construct successively higher level clusters until a level-k cluster is formed such that its clusterhead is the only level-k clusterhead in the network. This is denoted as level-L of the clustered hierarchy.

ii. $f_0 = \Theta\left(\frac{\mu}{R_{TX}} \cdot |V|\right) = \Theta(|V|)$

iii. $D_k = \Theta\left(\prod_{j=1}^{k}\sqrt{\beta_j}\right)$, $D_1 = \Theta(1)$.

iv. $n_k = \Theta(1)$, $\forall\, k \in \{1,2,\ldots,L\}$.

v. The pathological case illustrated by Fig. 2, where the ALCA yields a very small $\beta_k > 1$, either does not occur or is extremely rare so that its effect is neglibible, on average.

vi. $l_k = \Theta(1)$, $k \in \{1,2,\ldots,L\}$.

Justification of (iii) is based on the observation that the geographical area covered by a level-k cluster will be greater than that of a level-$(k-1)$ cluster by a factor of β_k, on average. Thus, the average geographical distance separating a pair of level-k clusterheads will be greater than that separating a pair of level-$(k-1)$ clusterheads by a factor $\sqrt{\beta_k}$. Assumption (iii) then follows readily from this.

Some justification of (iv) and (v) is provided in Fig. 3. Here, the average degree (n_k) and arity (β_k) for randomly generated networks consisting of 2800 nodes is shown. Each trial resulted in the ALCA generating 5 levels of clustered hierarchy. From Fig. 3, it is apparent that $\Theta(n_k) = \Theta(n_0)$, $k \in \{1,2,3,4(=L)\}$. Applying the earlier assumption that average node density is invariant in the node count means $\Theta(n_0) = \Theta(1)$ and, therefore, (iv) appears to hold. Further it is apparent that the ALCA, on average, generates hierarchical trees of sufficient arity at each level in the tree. This result tends to validate (v).

Asssumption (vi) is based on the fact that a single level-$(k-1)$ cluster link state change will have only local impact on the ALCA. Further discussion of this is provided in Section IV-D.

C. Cluster Formation

The formation of level-k clusters, $k \in \{1,2,\ldots,L\}$, involves the recursive application of the ALCA. At each level-k node, 2 rounds of communication must be performed with its level-k neighbors to elect a level-k clusterhead (i.e., a level-$(k+1)$ node). The communication with level-k neighbors is via unicast over a path consisting of $D_k = \Theta\left(\prod_{j=1}^{k}\sqrt{\beta_j}\right)$ level-0 node hops. That is, the length of paths connecting level-k nodes is $\sqrt{\beta_k}$ times longer than paths connecting level-$(k-1)$ nodes, on average. This would suggest that the communication overhead due to level-k cluster formation increases with k. However, the increase in path length between level-k nodes is offset by a decrease in the number of nodes at each successively higher level in the hierarchy. Specifically, the number of level-k nodes is less than the number of level-$(k-1)$ nodes by a factor of β_k. Thus, it is fairly straightforward to show, as follows, that cluster formation overhead is $O(|V|)$ for each level in the hierarchy.

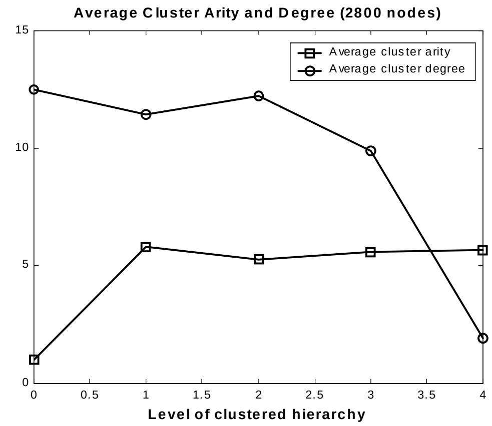

Fig. 3: Average degree and arity for a 2800 node network

Evaluation of level-1 cluster formation is omitted here as it follows the procedure of [3], requiring two rounds of communication between neighbors and $\Theta(|V|)$ network-wide communication overhead. Level-2 cluster formation is now evaluated. First, it is recalled that $D_1 = \Theta(1)$ represents the average number of hops separating adjacent level-1 nodes. Therefore, each level-1 node must communicate 2 rounds of cluster formation messaging with on average $n_1 \le n_{max} = \Theta(1)$ neighboring level-1 nodes over level-0 paths consisting of, on average, of D_1 hops. Since there are $|V|/|\beta_1|$ level-1 clusters in the network, the aggregate communication overhead due to level-2 cluster formation is (2 rounds of messaging) $\times$ (n_1 unicast sessions) $\times$ (D_1 transmissions per message per unicast session) $\times$ ($|V|/\beta_1$ nodes) $= 2 \cdot n_1 \cdot D_1 \cdot |V|/\beta_1 = \Theta(|V|)$ packet transmissions.

Next, level-3 cluster formation is considered. Here, the average number of level-0 hops separating level-2 nodes is on average larger than that separating level-1 nodes by a factor of $\sqrt{\beta_2}$ (i.e., $\sqrt{\beta_2} \cdot D_1$). Thus, the overhead due to a level-2 unicast communication session is greater than that incurred by a level-1 unicast session by a similar factor. However, the number of level-2 nodes is smaller than the number of level-1 nodes by a factor of β_2. Recalling that $n_2 \le n_{max} = \Theta(1)$ is the average number neighbors for each level-2 node and that there are $|V|/(\beta_2 \cdot \beta_1)$ level-2 clusters, the aggregate level-3 cluster formation packet transmission count is $2 \cdot \sqrt{\beta_2} \cdot D_1 \cdot n_2 \cdot |V|/(\beta_2 \cdot \beta_1) = 2 \cdot D_1 \cdot n_2 \cdot |V|/\left(\sqrt{\beta_2} \cdot \beta_1\right) < 2 \cdot D_1 \cdot n_{max} \cdot |V|/\beta_1 = \Theta(|V|)$. Thus, the additional number of packet transmissions due to increased path length for level-3 cluster formation (versus that of level-2 cluster formation) is more than offset by the reduced number of level-2 nodes involved in the process (versus the number of level-1 nodes involved in level-2 cluster formation).

Applying this analysis to level-k cluster formation, $k \in \{4,5,\ldots,L\}$, yields a similar upper bound. Since there are $\Theta(\log|V|)$ cluster levels, the aggregate number of packet transmissions due to ALCA cluster formation is

$O(|V| \cdot \log|V|)$. Dividing by $|V|$ yields the *per node* cluster formation overhead, $\psi_{CL\text{-}F} = O(\log|V|)$, as per Claim 2.

D. Cluster Maintenance

The assessment of level-k cluster maintenance follows logic similar to that given for level-k cluster formation. Here, however, the analysis is performed with respect to a baseline link state change frequency f_0 corresponding to the frequency of level-0 (i.e., node level) cluster link state change events. The maintenance of level-1 and higher level clusters is based on recursive application of the ALCA when impacted by level-0 cluster link state changes.

An important concept for the ensuing analysis is that of *diminishing probability of larger link state change effect radius*. That is, although a level-0 cluster link state change will necessarily result in messaging between level-0 nodes in the vicinity of the change, it will impact level-1 nodes with some probability $p_1 < 1$. Further, the level-0 link state change events also impact level-2 nodes with probability $p_2 \cdot p_1 < p_1 < 1$.

A second important concept is a property of the ALCA. Recalling from Section III-B, cluster organization incurs only 2 rounds of messaging. That is, only the level-(k–1) neighborhood of the endpoints for the updated level-(k–1) cluster link will be effected. This means that the effect of a level-(k–1) cluster link state change does not propagate throughout the level-k network topology. Further, since $n_{max} = \Theta(1)$, this neighborhood is also $\Theta(1)$ and $l_k = \Theta(1)$, as assumed in Section IV.

A third concept of importance here is that as in cluster formation, the average hop distance level-k control messaging must traverse increases by a factor of $\sqrt{\beta_k}$ over that required at level-(k–1). Thus, a critical issue to assess is whether p_k becomes sufficiently small with increasing k to offset the combined effects of l_k and increased path length between level-k nodes.

Simulation results, reported later in this section, address whether p_k offsets the effects of l_k and $\sqrt{\beta_k}$. However, justification for this to be true is provided now as follows.

a) Considering the level-1 topology of Fig. 1 makes an appeal to intuition. There are 33 level-1 cluster links. However, it is possible to delete up to 17 of these links without affecting the level-2 topology. This example illustrates that level-(k–1) cluster link changes which do *not* impact level-k cluster links are likely to be common.

b) Aggregating β_k level-(k–1) clusters per level-k cluster means that the number of level-k cluster links is less than the number of level-(k–1) links by a factor of approximately β_k.

c) As the average separation between level-k clusterheads increases with increasing k, so does the distance required for the two clusterheads to move relative to one another to trigger a cluster link state change between the two clusters.

To assess the aggregate overhead due to ALCA cluster maintenance, a single level-0 cluster link state change (either a link is added or a link is broken) and the subsequent effect on level-0 nodes and higher level nodes is considered. Beginning with the effect on level-0 clusters, with probability $p_0 = 1$ (because the cluster link state change has occurred at the level-0 topology) notification of this change will be announced by the two nodes forming the endpoints of the added/deleted link. After two rounds of messaging, each of the two nodes forming the endpoint of the added/deleted edge will either declare itself a level-0 clusterhead (or relinquish its role as a clusterhead), trigger a neighbor to serve as its new level-0 clusterhead (or to relinquish its role as a clusterhead) or incur no change to the clusterhead status of its neighborhood. The average number of packet transmissions is computed as follows: (2 level-0 nodes affected) $\times$ (n_0 neighbors/node) $\times$ (2 messages/neighbor) $\times$ (1 hop/message) $\times$ (l_0 = 1 level-0 link) $\times\, p_0 = 4 \times n_0$.

Now, with probability p_1, this level-0 cluster link change will also have effect on the local level-1 topology. The average number of level-0 hop across which ALCA cluster maintenance will have to traverse is D_1. The expected number of level-1 packet transmissions is computed as follows: (2 level-1 nodes affected) $\times$ (n_1 neigbors/node) $\times$ (2 messages/neighbor) $\times$ (D_1 hops/message) $\times$ (l_1 level-1 links) $\times\, p_1 = 4 \times n_1 \times l_1 \times p_1 \times D_1$.

Next, with probability p_2, this level-1 cluster link state change will also effect the local level-2 topology. The average number of level-0 hops across which ACLA cluster maintenance messaging will have to traverse will be larger than that for the level-1 case by a factor of $\sqrt{\beta_2}$. Thus, the expected number of level-2 packet transmissions is computed as follows: (2 level-2 nodes affected) $\times$ (n_2 neighbors/node) $\times$ (2 messages/neighbor) $\times$ ($\sqrt{\beta_2} \times D_1$ hops/message) $\times\, l_1 \times p_1 \times l_2 \times p_2 = 4 \times \sqrt{\beta_2} \times n_2 \times l_1 \times p_1 \times l_2 \times p_2 \times D_1$. Letting $\Psi_{CL\text{-}M}$ be the aggregate control packet overhead due to level-k cluster maintenance, this quantity may be expressed as follows:

$$\Psi_{CL-M} = 4 \cdot \sum_{k=0}^{L} n_k \cdot f_k \cdot D_k \qquad (1)$$

Noting that $f_k = f_{k-1} \cdot l_k \cdot p_k$ and $D_k = \sqrt{\beta_k} \cdot D_{k-1}$ allows $\Psi_{CL\text{-}M}$ to be expressed in terms of f_0 and D_1.

$$\Psi_{CL-M} = 4 \cdot f_0 \cdot \left(n_0 + \frac{D_1}{\sqrt{\beta_1}} \cdot \sum_{k=1}^{L} \left(n_k \cdot \prod_{j=1}^{k} \sqrt{\beta_j} \cdot l_j \cdot p_j \right) \right) \qquad (2)$$

Noting that $n_{\max} = \Theta(1)$, $D_1 = \Theta(1)$ and $\beta_1 > 1$, the expression for $\Psi_{\text{CL-M}}$ given by (2) may be upper bounded as follows:

$$\Psi_{CL-M} < 4 \cdot D_1 \cdot f_0 \cdot n_{\max} \cdot \left(1 + \sum_{k=1}^{L}\left(\prod_{j=1}^{k}\sqrt{\beta_j} \cdot l_j \cdot p_j\right)\right) \quad (3)$$

If $\rho = \max\left\{\sqrt{\beta_k} \cdot l_k \cdot p_k\right\} < 1 \;\; \forall \; k \in \{1,2,\ldots,L\}$, then a more definitive upper bound for $\Psi_{\text{CL-M}}$ may be formulated:

$$\Psi_{CL-M} < 4 \cdot D_1 \cdot f_0 \cdot n_{\max} \cdot \sum_{k=0}^{L} \rho^k \quad (4a)$$

$$\rightarrow \Psi_{CL-M} < 4 \cdot D_1 \cdot f_0 \cdot n_{\max} \cdot (L+1) \quad (4b)$$

$$\rightarrow \Psi_{CL-M} = O(f_0 \cdot L) = O(|V| \cdot \log|V|) \quad (5)$$

The $O(|V| \cdot \log|V|)$ figure concluding (5) is based on the fact that $D_1 = \Theta(1)$, $n_{\max} = \Theta(1)$, $L = \Theta(\log|V|)$ and $f_0 = \Theta(|V|)$.

The above analysis assumes that $\rho < 1$, based on reasons (a)-(c) given earlier. This phenomenon has been verified in simulation of hierarchical networks for the ALCA. Fig. 4 compares the reduction in cluster link state change frequency, $1/(l_k \cdot p_k)$, with the increased average path length separating clusterheads. Although in the simulations, $1/(l_1 \cdot p_1) = f_0 / f_1 < \sqrt{\beta_1}$, $1/(l_k \cdot p_k) = f_{k-1} / f_k > \sqrt{\beta_k} \;\; \forall \; k \in \{2,3,\ldots,L\}$ thereby allowing the summation of (3) to still converge. Fig. 5 shows a plot of the $l_k \cdot p_k \cdot \sqrt{\beta_k}$ product (ρ_k).

The net result of this analysis is that $\Psi_{\text{CL-M}}$ is $O(|V| \cdot \log|V|)$. Dividing this result by the node count yields $\psi_{\text{CL-M}} = O(\log|V|)$ per node, as per Claim 3.

Lastly, concerning Fig. 4, the decrease in cluster link state change frequency, when comparing level-0 cluster link state change frequency to that of level-1 cluster link states, is relatively modest (approximately 1.5). It is conjectured that this is due to the nature of physical (i.e., level-0) cluster link state changes and virtual (i.e., level-k, $k \geq 1$) cluster link state changes. Whereas level-0 cluster link states depend only on whether node pairs are within some distance of one another, a level-k node's link states depend not only on the geographic distance separating it from other nodes but also on the value of IDs of its neighbors relative to its own ID.

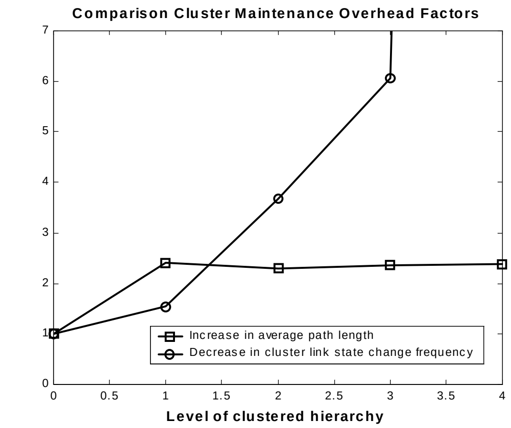

Fig. 4. Effect of clustering in a 2800 node network. μ = 10m/s, network radius = 3705m and R_{TX} = 250m.

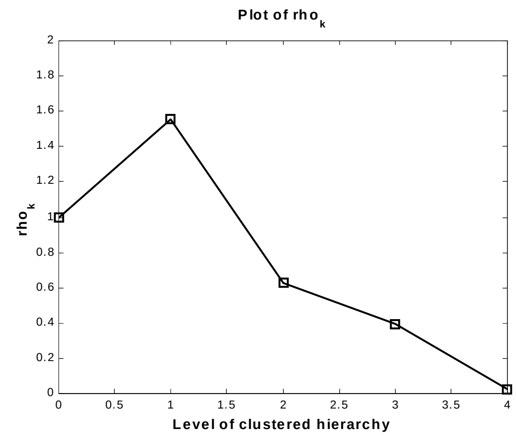

Fig. 5. ρ_k for the network scenario of Fig. 4.

E. Flooding Overhead

In order for a node v to unambiguously perform packet forwarding based on strict hierarchical routing, v must know the topology for each cluster to which it belongs. That is, for each level-k cluster to which it belongs, v must know the node IDs of the cluster, the connectivity among the cluster nodes and to which clusters (if any) each cluster node serves as a cluster gateway. Of these three data items, the intra-cluster connectivity matrix contributes the most to the level-k routing table size—a contribution that is quadratic in the cluster size.

To assess routing table size, $\beta_{\max}$ is defined as $\max\{\beta_1,\beta_2,\ldots,\beta_{L+1}\}$. That is, $\beta_{\max}$ is the maximum average number of level-(k-1) clusters per level-k cluster among all levels of the clustered hierarchy. Since $\beta_k \leq 1+n_{k-1}$, it follows that $\beta_{\max} \leq 1+n_{\max} = \Theta(1)$. Thus, the maximum table size for the hierarchical map of the branch in the network layer hierarchy to which a node v belongs is $O(L \cdot |\beta_{\max}|^2)$. Since $L = \Theta(\log|V|)$ and $|\beta_{\max}|^2 = \Theta(1)$, this maximum table size is $\Theta(\log|V|)$.

Now, to efficiently distribute the hierarchical routing table relevant to each network node, a recursive flooding procedure is employed for each level in the hierarchy. First, dissemination of the level-L cluster topology is initiated by the level-L clusterhead (i.e., the highest node in the network hierarchy). The topology consists of not more than $|\beta_{max}|$ nodes and, therefore, requires a message size that does not exceed $\Theta(|\beta_{max}|^2) = \Theta(1)$ bytes and incurs $\Theta(|V|)$ packet transmissions.

Unlike the level-L topology which must be disseminated throughout the network, distribution of a level-(L–1) cluster topology (and all other level-(L–1) clusters) must be confined to that specific level-(L–1) cluster. That is, the topology for a level-(L–1) cluster is flooded only to members of that level-(L–1) cluster. The flooding is initiated by the level-(L–1) clusterhead (a level-L node). Since there are β_L level-(L–1) clusters, on average, the flooding overhead incurred for a specific level-(L–1) cluster is $\Theta\left(|V|/\beta_L\right)$. Also, there are on average β_L such level-(L–1) clusters. Thus, the aggregate number of packet transmissions required to flood level-(L–1) cluster topologies is $\Theta(|V|)$. Here again, the maximum message size is $\Theta(|\beta_{max}|^2) = \Theta(1)$ bytes. The analysis given here for disseminating level-(L–1) cluster topology is applied recursively for level-(L–2) and subsequent lower levels of the clustered hierarchy until level-1 cluster topology distribution overhead has been assessed. Therefore, at each level in the network hierarchy, the aggregate overhead to distribute to each node $v \in V$ the cluster topology for the level-k cluster to which it belongs is $\Theta\left(|\beta_{max}|^2 \cdot |V|\right) = \Theta(|V|)$, $k \in \{1,2,\ldots,L\}$.

Since $L = \Theta(\log|V|)$, the aggregate overhead incurred by flooding an L-level hierarchical map to each node is $\Theta\left(|V| \cdot \log|V|\right)$ and $\psi_{FLOOD} = \Theta(\log|V|)$, as per Claim 4. With the hierarchical map, each node can unambiguously forward packets toward an arbitrary destination address.

F. Hierarchical Addressing

To facilitate unambiguous packet forwarding via hierarchical routing, each datagram header must contain the hierarchical address of the target node t. The hierarchical address consists of the concatenation of the L cluster IDs of the clusters to which t belongs as well as the ID of t, itself. Thus, the hierarchical address consists of $(L+1)\times b$ bits, where b is the number of bits in a node ID. Thus, the hierarchical address adds $\Theta(L) = \Theta(\log|V|)$ bits of overhead to the header of every datagram.

The $\Theta(\log|V|)$ addressing overhead impacts the overhead associated with cluster maintenance as unicast communication is employed to exchange cluster link state information between adjacent clusterheads. Thus, the length of cluster link state packets is $\Theta(\log|V|)$. Combining this with $\psi_{CL\text{-}M} = O(\log|V|)$ given in Section IV-D, means that clustering overhead is $O(\log^2|V|)$.

Additionally, because the $\Theta(\log|V|)$ hierarchical address appears in every datagram header it also effects the network capacity available for actual user traffic. As shown in [10] the maximum aggregate throughput in packet radio networks $\Theta\left(\sqrt{|V|}\right)$. In [14] it is shown that the aggregate throughput is only $\Theta\left(\sqrt{|V|/\log|V|}\right)$ when the uniform node distribution is *random*. Thus, the presence of hierarchical addressing chokes the available throughput per node by an additional $\log|V|$ factor.

V. Conclusions

This paper has argued that the number of packet transmissions per node required for a particular MANET hierarchical clustering scheme (the ALCA) is $\Theta(\log|V|)$. Combining this with $\Theta(\log|V|)$ datagram header bits required for hierarchical addressing results in per node overhead of $O(\log^2|V|)$ bits per second. This is an important result because it implies that the sizing of network links need only be $O(\log^2|V|)$ in order to accommodate the growth in control traffic incurred by hierarchical routing. Thus, given the clustering procedure under consideration here, hierarchical routing actually *scales very well* with respect to increasing node count. In contrast, for example, non-hierarchical link state routing incurs aggregate link state packet overhead that is $\Theta(|V|^2)$. This means the sizing of network links must be $\Theta(|V|)$ in order to accommodate the growth in traffic due to flooding of link state packets. Although $O(\log^2|V|)$ overhead per node for hierarchical clustering may be an intuitively sensible figure, there is no previous work in published literature that justifies this figure or even states this figure for multi-hop, mobile packet radio networks. Determining whether a similar bound holds for other clustering approaches represents a direction for future work.

The effect of $\Theta(\log|V|)$ levels of hierarchical addressing is that the length of datagram headers must be $\Theta(L) = \Theta(\log|V|)$. This means that the average throughput available for each network node is throttled by a factor of $\Theta(\log|V|)$ compared to what theoretically can be achieved via non-hierarchical routing. However, non-hierarchical link state routing incurs overhead that is $\Theta(|V|)$ per node. Thus, although hierarchical addressing may constrict network throughput, hierarchical routing is still clearly more scalable in comparison. For example, in the 2800 node simulation of Fig. 3, $L = 4$ which means that the hierarchical address consists of the concatenation of 5 node IDs. Assuming a 64-bit NIC number is used as the node ID, means that a 5-level hierarchical address incurs 32 additional bytes of datagram header content that do not occur for a non-hierarchical address. An extra 32 *bytes* in each datagram header is likely to be substantially less than the $\Theta(2800)$ *packet transmissions per node* that occur at a frequency of f_0 in order to propagate link state packet updates.

Of course, the $O(\log^2|V|)$ result derived here is not the complete picture in assessing the scalability of topology-based hierarchical routing. First, there are issues of location management (i.e., address management) as considered in [15] and [18]. Second, there is the matter of increased average path length due to the sub-optimal packet forwarding paths that are followed as a result of the summarized topology information in hierarchical routing. Although the work of [1] indicates that the average hierarchical path length is worse than the optimal path length by only a constant, simulations to determine the actual average hierarchical path length for networks consisting of $\Theta(\log|V|)$ cluster levels represent a logical next step in comparing routing performance. Lastly, the effect of the scaling constant for the $O(\log^2|V|)$ result derived here has yet to be determined. A large scaling constant would mean that the scalability advantage of hierarchical routing with $\Theta(\log|V|)$ levels of clustered hierarchy can not be realized until $|V|$ becomes very large. A small scaling constant, on the other hand, would mean that the hierarchical routing techniques discussed here might be appropriate even for "small" networks.

References

[1] L. Kleinrock and F. Kamoun, "Hierarchical routing for large networks," *Computer Networks*, Vol. 1, No. 3, pp. 155-174.

[2] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," *Proc. Mobicom '98*, Dallas, TX, pp. 85-97.

[3] A. D. Amis, R. Prakash, T. H.P. Vuong and D. T. Huynh, "Max-min D-cluster formation in wireless ad hoc networks," *Proc. IEEE Infocom 2000*, pp. 32-41.

[4] D. J. Baker and A. Ephremides, "The architectural organization of a mobile radio network via a distributed algorithm," *IEEE Trans. Commun.*, Vol. COM-29, No. 11, November 1981, pp. 1694-1701.

[5] D. J. Baker, A. Ephremides and J. A. Flynn, "The design and simulation of a mobile radio network with distributed control," *IEEE J. Select. Areas Commun.*, Vol. SAC-2, No. 1, January 1984, pp. 226-237.

[6] C. R. Lin and M. Gerla, "Adaptive clustering for mobile wireless networks," *IEEE J. Select. Areas Commun.*, Vol. 15, No. 7, September 1997, pp. 1265-1275.

[7] J. Westcott and G. Lauer, "Hierarchical routing for very large networks," *Proc. IEEE MILCOM '84*, 21-24 October 1984, pp. 214-218.

[8] G. S. Lauer, "Hierarchical routing design for SURAN," *Proc. IEEE ICC '86*, 22-25 June 1986, pp. 93-102.

[9] G. Pei, M. Gerla, X. Hong and C.-C. Chiang, "A wireless hierarchical routing protocol with group mobility," *Proc. IEEE WCNC '99*, September 1999, pp. 1538-1542.

[10] L. Kleinrock and J. Sylvester, "Optimum transmission radii for packet radio networks or why six is a magic number," *Proc. IEEE National Telecommunications Conference*, Birmingham, AL, December 1978, pp. 4.3.1-4.3.5.

[11] M. Steenstrup, "Cluster-based networks," *Ad hoc networking*, Editor C. E. Perkins, Ch. 4, pp. 75-138, Addison Wesley, 2001.

[12] A. Iwata, C.-C. Chiang, G. Pei, M. Gerla and T.-W. Chen, "Scalable routing strategies for ad hoc wireless networks," *IEEE J. Select. Areas Commun.*, Vol. 17, No. 8, August 1999, pp. 1369-1379.

[13] R. Ramanathan and M. Steenstrup, "Hierarchically-organized, multihop mobile wireless networks for quality-of-service support," *ACM/Baltzer Mobile Networks and Applications*, Vol. 3 (1998), pp. 101-119.

[14] P. Gupta P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Information Theory*, Vol. 46, No. 2, March 2000, pp. 388-404.

[15] J. Sucec and I. Marsic, "Location management for link state routing in mobile ad hoc networks," *Proc. IEEE WCNC 2002*, 18-21 March 2002.

[16] C. V. Rammamoorthy and W.-T. Tsai, "An adaptive hierarchical routing algorithm," *Proc. IEEE CSAC*, Fall 1983, pp. 93-104.

[17] J. J. Garcia Luna Aceves and N. Shacham, "Analysis of routing strategies for packet radio networks," *Proc. IEEE Infocom '85*, pp. 292-302.

[18] J. Sucec and I. Marsic, "Location management handoff overhead in hierarchically organized mobile ad hoc networks," *Proc. Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, 19 April 2002.