

Efficient Paging in Large-Scale Urban Mesh Networks

Stephan Bohacek
Dept. of Electrical and
Computer Engineering
University of Delaware
Newark, DE, USA
bohacek@udel.edu

Vinay Sridhara
Dept. of Electrical and
Computer Engineering
University of Delaware
Newark, DE, USA
vsridahar@udel.edu

Lei Lou
Dept. of Computer and
Information Science
University of Delaware
Newark, DE, USA
llou@udel.edu

ABSTRACT

The task of forming a route can be broken into two parts, finding any route to the desired host and refining that route so it is suitable for the connection. We refer to this first task as paging and it consists of simply delivering a single packet to the destination. Traditional MANET routing protocols perform this task via flooding. However, when the number of nodes ranges into the thousands, flooding becomes highly inefficient. In the case of cellular networks, paging is performed with one or more base stations broadcasting a page message. However, in large-scale urban networks, nodes are not necessarily in direct contact with the infrastructure, hence such techniques are not applicable. In this paper, the topology and architecture of a large-scale urban mesh network is examined and a scheme is developed for paging. This scheme is designed for and takes advantage of the topology, mobility, and architecture found in urban mesh networks. The resulting scheme is highly efficient in that it generates little overhead and is scalable.

Categories and Subject Descriptors: C.2.2 [Computer-communication networks]: Network Protocols

General Terms: Performance

Keywords: Large-scale urban mesh networks, Mobility management, Paging

1. INTRODUCTION

Large-scale urban mesh networks (LUMNets) are poised to provide ubiquitous internet access to mobile users. In early 2006 the city of Philadelphia [16] finalized a partnership to begin the construction of an mesh network that will cover the 135 sq. mile city. The City of San Francisco has recently received proposals for a LUMNet in that city [9]. Taipei is establishing wireless access in 5% of the city, which includes the city's most-populated area, and within a

year expects to cover around 90% of the city [20]. Several other cities such as Minneapolis [14], Anaheim, California, and Tempe, Arizona are planning similar LUMNet deployments. These deployments are meant to enhance city and emergency services communication as well as to provide city-wide, low-cost, ubiquitous Internet access for residents and visitors.

While these networks are being deployed, there are several challenges that remain unsolved. Perhaps the most pressing issues are related to scalability and coverage. Consider, for example, that during the lunch time, downtown Manhattan has approximately 10,000 people outdoors in 1 km² [17]. Similarly, office usage surveys [22] and maps of downtown areas [1] [7] indicate that an urban core may have office worker densities that exceed 40,000 people per km².

Coverage is a critical challenge facing the usability of these networks. In the case of Philadelphia, the initial plan was that client nodes would be only one hop from the infrastructure nodes (either wired or wireless infrastructure). Due to the administrative and financial difficulties with the city owned networking including infrastructure nodes that are inside of private buildings, the infrastructure nodes are restricted to be outdoors. However, the 802.11 physical layer is not able to easily communicate from indoors to outdoors. As a result, it is not possible for the planned outdoor mesh network to provide one hop connectivity to all locations inside buildings.

Recognizing the limitation that an outdoor only network places on the utility of the network, the planners in Philadelphia considered the possibility that consumer owned devices could act as relays and therefore extend the coverage. In terms of topology, such LUMNets then resemble hybrid networks in that there is an infrastructure component and an MANET component. While there has been some attention to developing protocols for hybrid networks, protocols for LUMNets has not been well explored.

This paper examines paging in LUMNets. By paging, we mean the task of finding a client node in the LUMNet, or more specifically, delivering a single packet to a client node. Paging is the first step in forming a connection with a client and is an important part of mobility management. In the case of cellular networks, it has been found that if a client node is idle (i.e., a connection is not currently established), then it is more efficient for the client to *not* alert the base station every time it moves out of range of one base station and into range of another. The details of when a client alerts the base station that it has moved is the focus of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WMASH'06, September 25, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-484-7/06/0009 ...\$5.00.




Figure 1: Performance of Paging via Flooding in LUMNets in an 9-block region of Chicago. Scenarios with 1000, 2000, 3000, and 4000 client nodes were examined. In each scenario, ten node pages (i.e., node searchers) were performed, and the average overhead required to deliver a single packet to a destination is shown above. The paging is performance via flooding. Note that when there are more than 1000 nodes, the route search may fail, and hence, on average, the source initiates more than one route requests.

considerable research [13], [18], [3], [11], [4], [5]. However, in all cases, when a new connection is to be established, the infrastructure must find the client node, that is, the client must be paged.

Since the clients in the LUMNet architecture considered here are not able to directly communicate with the infrastructure, the paging schemes developed for cellular networks cannot be applied. On the other hand, since these LUMNets have components that are similar to ad hoc networks, techniques developed for such networks can be considered. In on-demand routing protocols for MANETs (e.g., DSR [12], AODV [15]), paging is integrated into the route request phase. Such protocols use flooding to find the destination. A critical drawback with these protocols is that flooding generates excessive overhead. For example, consider the small scale example shown in Figure 1. Here, a 9-block region of Chicago is examined with between 1000 to 4000 clients and 65 outdoor base stations. The modeling and propagation for this simulation is based on the UDeLModels [2] and QualNet was used for packet simulator. The simulation results show that a tremendous amount of overhead is generated to find a single node in even this small topology (note that the overhead appears to grow exponentially with the population size). Clearly, it is not practical to use such flooding-based schemes in LUMNets. We conclude that a new approach to find nodes in LUMNets is required.

In this paper, a paging scheme suitable for LUMNet is developed. The scheme has some proactive features and reactive features. Specifically, in a way that is similar to cellular phone mobility management, the infrastructure proactively maintains an estimate of the location of the client, and when a client is to be paged, the paging message is flooded in a specific region. A key ingredient of paging in LUMNets is that urban propagation and mobility naturally partition the client nodes and hence the infrastructure only needs to track which cell of the partition the client is located. The result is




Figure 2: (a) The figure illustrates the disparity in coverage of indoor and outdoor locations due to infrastructure nodes placed outdoors at various pathloss thresholds. (b) The figure illustrates the coverage on different floors. The results obtained are from a simulation carried out for a section of downtown Chicago using raytracing mechanism. 27 infrastructure nodes were placed outdoors on lamp posts (3.5m high) with a separation of 50m between them.

that the location of the clients can be tracked in a scalable fashion.

This paper proceeds as follows. Since the paging scheme takes advantage of the topologies that naturally occur in LUMNets, the next section details several important aspects of these topologies. Section 3 then describes the paging problem and Section 4 presents the scalable paging technique. And Section 5 presents some simulation-based performance analysis. These simulations are based on the UDeLModels, which provide realistic propagation in mobility in urban areas [2]. And finally, concluding remarks can be found in Section 6. Note that due to space limitation, analytical performance analysis is not included, but can be found in the full version of the paper [6].

2. THE TOPOLOGY OF LUMNETS

2.1 Coverage of the Infrastructure

The propagation of wireless signals in urban areas results in a specific structure on the LUMNet. Specifically, the outdoor communications occurs over good propagation environments, while indoor communications suffers from poor propagation environments. Thus, the outdoor area is well connected. For example, a relatively small number of hops are required to travel great distances. As a result, LUMNets are expected to provide very good coverage outdoors. On the other hand, since the infrastructure is restricted to be outdoor, client nodes within buildings are not well connected. They are mostly only connected to client nodes on the same floor or, at most, a few floors away. For example, consider Figure 2. This figure is based on UDeL Models' 3-D ray-tracing propagation computational tools applied to a nine block region of downtown Chicago where the infrastructure nodes were approximately 50 m apart (See [2] for details). The left-hand frame shows the fraction of all outdoor and indoor location that are reachable by the infrastructure at different channel gains. As expected, outdoor locations are well connected and can support high bit-rate communication (i.e., they are connected with low path loss), while




Figure 3: The figure shows the coverage inside of buildings due to the infrastructure nodes placed outdoors with a separation of 50m between them. It can be observed as the floors number increases the coverage becomes almost nil.

even if low bit-rate communication is used (e.g., a 70 dB path loss threshold is used), only a small fraction of indoor locations are reachable from the infrastructure. The right-hand frame of Figure 2 provides further insight into the problem with indoor coverage. This figure shows how the coverage varies as a function of the floor (here it is assumed that a path loss of 70 or less will support communication). Clearly, while good coverage is provided at the lower floors, the outdoor radios are not able to penetrate the upper floors. If the mesh network is composed of roof top infrastructure, then the figure would be flipped, with good coverage on the upper floors, and poor coverage on the lower floors. While one solution is to abandon the hope of ubiquitous communication and insist that indoor communication is provided by the individuals with administrative control of the buildings we seek protocols that support the goal of city-wide ubiquitous communication.

To understand the topology further, consider Figure 3. This figure indicates the coverage of in an urban area. The coverage of a building is indicated by the color. Specifically, the darker the color, the higher the fraction of the locations on that floor of the building is covered by the infrastructure. In Figure 3, the infrastructure is designed so that each location outdoors is covered by at least one infrastructure node. However, as clearly shown, in Figures 3 and 2 a significant part of the indoor area is not within range to the infrastructure. Specifically, the white colored floors on the upper floors indicate that no locations on these floors can communicate directly with the infrastructure.

2.2 Topologies with Mobile Relaying

If client nodes are permitted to act as relays, then the coverage of the LUMNet is greatly extended. Indeed, if the subscriber density is high enough, then all regions of the city will be able to communicate with the infrastructure. However, due to the structure of propagation (i.e., good outdoor propagation, and bad indoor propagation), the topology of the LUMNet takes a particular form with a highly connected core that is composed of outdoor nodes and weakly connected tendrils that are composed of client nodes within buildings. Figure 4 shows an example of a tendrill that resulted from a ray-tracing experiment.




Figure 4: The figure depicts the structure of multi-hop communication in a building. This structure resembles that of a tendrill.

We use the following nomenclature to describe the clients within the LUMNet topology (also, see Figure 5). If a client node is directly reachable by the infrastructure, then the client is referred to as a *direct client* (DC). While most, or even all, outdoor nodes will be DCs, some DCs will be indoors as well. We refer to tendrill clients (TC) as client nodes that are not directly reachable from the infrastructure, but are reachable when client nodes are used as relays. Of the DCs, some can act as a gateway to the tendrill in that they have infrastructure nodes as neighbors and TCs as neighbors. Such DCs are referred to as tendrill gateways (TGWs). Of course, as client nodes move, their status may change.

Besides the client nodes, there are two types of infrastructure nodes, namely base stations that act as gateways between the wireless network and wired Internet, and fixed wireless relays. We make no distinction between these types of infrastructure nodes, and simply refer to them as WFIN (Wi-Fi infrastructure nodes).

3. PAGING IN WIRELESS NETWORKS

As mentioned in Section 1, this paper focuses on paging, or the problem of finding a destination. It is important to note that we only consider the problem of the infrastructure finding a client, and not the problem of the client finding the infrastructure. However, since the infrastructure is large (i.e., it is made up of many nodes) and a particular client is small (i.e., it is only one node), this second problem is considerably more simple than the one examined here.

Paging is an integral part of mobility management and there has been extensive work on paging, especially in cellular phone networks [13], [18], [3], [11], [4], [5]. Paging has also been integrating into mobility management for data networks (e.g., [23], [10]). However, in these cases, the clients are within one hop of the infrastructure, and hence these techniques cannot be applied for paging TCs. Nonetheless, the basic ideas are common to all paging algorithms, including the one developed here. Specifically, each client node will occasionally register with the infrastructure. This reg-

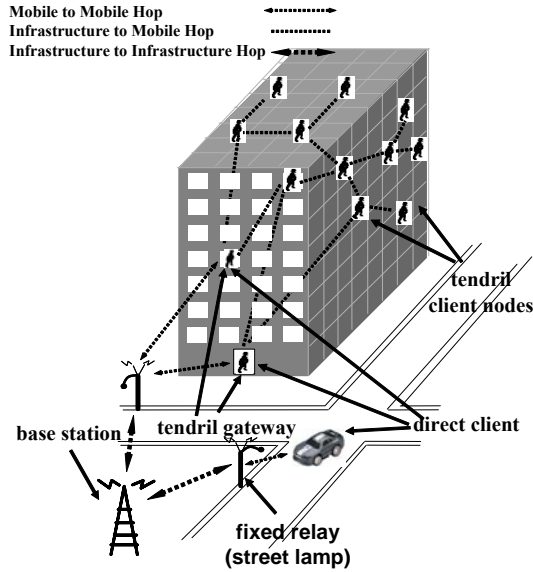


Figure 5: A conceptual visualization of a mesh network. The picture above highlights the different elements of multihop mesh networks.

istration will contain various information and allows the infrastructure to learn the current location of the client. One approach is for clients to periodically register with the infrastructure, while another is for nodes to register when an event occurs. For example, a client might register when it can no longer "hear" its last WFIN or it can not hear any node from a set of WFINs that was determined during the last registration. Alternatively, a client might register when it has changed WFINs a specific number of times [3]. These techniques assume that the client uses WFINs to determine whether and how far it has moved. Hence, the client must be able to directly communicate with the WFINs.

The basic goal of client registration is to reduce the size of the region that must be searched when the client is paged. For example, when a DC registers, it is clear which WFIN it is reachable through, and hence, the DC can easily be found by the WFIN where the client most recently registered broadcasting a paging packet. However, when a TC registers, the information about which WFIN the client is reachable from is not sufficient to reduce the region searched. To see this, one must consider the structure of the propagation discussed above, i.e., the good¹ propagation outdoors and bad propagation indoors. This structure means that distance parts of the city can be reached in a few hops, while the upper floors of a building requires a large number of hops. Consider, for example, a TC that has recently registered with a WFIN and has provided the further information that it is 10 hops from the WFIN (such values are common when the TC is on the upper floors of a building). Thus, if the WFIN desires to reach the TC, it could, perhaps, broadcast a paging message that is to be relayed no

¹Due to the straight streets and buildings lining the streets, it is common that when propagating down a street, the signal strength decays like $d^{-\alpha}$, where $\alpha < 2$. This effect is known as propagation in urban canyons.




Figure 6: Heterogenous Propagation. The 5th and 95th percentile of the distances from the WFIN to nodes that are exactly 1, 2, 3, and 4 hops away from the WFIN.

more than $10 + K$ times, where K allows for a few extra hops in case the TC has moved. However, due to the good propagation conditions outdoors, such a message would be propagated to the far reaches of the city. Furthermore, the message be relayed into all buildings near the WFIN. Consequently, the fixed radius flood would utilize bandwidth in a large part of the city.

This effect can also be seen in Figure 6. This figure shows the 5th and 95th percentile of the distance from the WFIN to nodes that are exactly 1, 2, 3, and 4 hops away from the WFIN. Note that as the number of hops increases, the 5th percentile does not change. These nodes are within buildings. However, the 95th percentile increases with the number of hops. As a result, specifying that a flood only spread K hops is ineffective in reducing the geographic area the flood covers.

4. SCALABLE PAGING FOR LUMNETS

4.1 Overview

The basic idea of the scalable paging approach developed here is that the WFINs should limit the search of the TC to the tendril where TC is located. Thus, the infrastructure must keep track of which tendril a node resides in and TCs within the tendril must also know which tendril they reside in. If these objectives are met, then the WFIN simply broadcasts a page message with the correct tendril identification and all nodes in the identified tendril forward the page message according to an efficient flooding scheme (in this study, Scalable Broadcast Algorithm (SBA) was used [8]). Hence, the overhead for the search is the overhead to efficiently flood a single tendril. A tendril may be a single building as shown in Figure 4, or may have some other structure. Regardless, the overhead required to efficiently flood a tendril is far less than the overhead required to efficiently flood the entire network. Furthermore, the overhead is independent of the size of the network and only depends on the size of the tendrils within the network.

A tendril is identified by the TGW nodes that provide connectivity between the tendril and the infrastructure. In order for the TC to know which TGWs serve the tendril,

the TGW floods beacons through the tendrils (here, and in all situations where flooding is used, it is assumed that some sort of efficient flooding scheme is used). By examining how many hops a beacon has traveled, and by comparing this hop count to the hops traveled by beacons from other recently heard TGWs, the TC decides whether to propagate the TGW beacon. In this way, TCs can determine a set of TGWs that are between it and the WFIN. When a TC sends a registration packet to the infrastructure, it includes a list of all TGWs that it has recently heard and that provide high quality communication between the WFIN and the TC (these TGWs are called *active TGWs*). The set of TCs that share the same active TGWs constitute a tendrils. Thus, as long as the tendrils information is up-to-date in both the infrastructure and the tendrils nodes, then the region of the network that must propagate the page message is greatly limited.

An important challenge is keeping the tendrils information up-to-date without requiring excessive overhead due to registrations. The next section describes a method that reduces the overhead due to registrations by allowing the registration of a single TC to be gratuitously applied to all the TCs within the tendrils. One complication arises when a client moves into a new tendrils. In this case, the TC's registrations may be gratuitously applied to all the TCs within last tendrils it was in, and the registrations by TCs within the new tendrils will not be gratuitously applied to TC that just entered the tendrils. This complication can be eliminated by a client rapidly detecting that it has changed tendrils. The algorithm for this detection is developed in Section 4.3.

4.2 Similarity and Gratuitous Registration

While restricting the page message to remain within a specific tendrils greatly reduces the overhead involved in paging, it requires that tendrils clients keep the WFINs informed about the TGW that they are reachable through. This requirement can lead to excessive registration overhead when TGWs move or go off/online. For example, when a DC moves close to a building and is able to communicate with nodes within the building, it becomes a TGW for the tendrils within the building. Similarly, when a node inside the building moves toward the edge of the building, it may be able to communicate with the infrastructure and become a TGW. When either of these events occurs, there is a risk that every node within the tendrils will try to register with the information that it is reachable through the new TGW.

To reduce the number of registrations when the TGW change, the similarity of nearby nodes is exploited, where two nodes are *similar* if they share the same set of active TGWs (active TGW will be defined shortly). The idea is that when a registration from a tendrils client is received, the WFIN applies any changes in the registration to all nodes that are similar to the registering node. Thus, one registration will suffice for the entire set of similar tendrils clients. To implement this idea, the tendrils nodes must be able to determine whether a node has registered and whether the WFIN will apply this registration to it. Before discussing the scheme that achieves this goal, we discuss active TGWs.

4.2.1 Active TGW

A TC's *active TGWs* are those TGWs that offer good connectivity between the tendrils and infrastructure. Specif-

ically, an active TGW must have offered connectivity from itself to the tendrils client for a sufficiently long time and with high enough quality and offer connectivity between itself and the infrastructure for a sufficiently long time and with high enough quality. Here we specify that communication quality is the fraction of successful transmissions. While not investigated here, quality could also include other metrics such as SNR and interference.

To determine which TGWs are active TGWs, the TGW beacon contains a beacon number, the beaconing period, and a measure of the quality of communication from the TGW's "best" WFIN to the TGW. The "best" WFIN is the one that has high enough quality of communication to the TGW for the longest duration. By including a beacon number, the TC is able to determine if any beacons have been missed. The beaconing period allows the tendrils client to determine whether and how many beacons have been missed since the last one was received. Together, these allow the TC to determine the probability of packet delivery from the TGW to the TC. The WFIN beacons also include beacon numbers and beaconing period and allow the TGW to compute the communication quality from the WFINs. It is possible to include other measures within the TGW beacon. For example, as the beacon is being flooded through the tendrils, each node could append the received signal strength. On the other hand, the nodes could append whether the node is a fully mobile node (e.g., mobile VoIP hand-set), or semi-permanent (e.g., a desktop or consumer owned device that is specifically designed to act as a relay). Such techniques are not investigated here.

As TGW beacons are received, the TC is able to determine whether the TGW is stable (i.e. whether the number of beacon received exceeds a threshold) and whether the communication quality from the TGW is sufficient (e.g., the probability of a beacon being successfully received exceeds a threshold). If both of these requirements are met, then the TGW becomes an active TGW. On the other hand, a TGW is removed from the list of active TGWs when the communication quality falls below a threshold or when a sufficient number of consecutive TGW beacons have not been received. It is important to note that it takes time for a TGW to become an active TGW. Thus, a mobile node that happens to momentarily be near a building will not become an active TGW. Furthermore, hysteresis is included so that there is limited "flap" in the list of active TGWs.

4.2.2 Gratuitous Registration

In order to support similarity-based implicit registration, three lists of active TGWs are maintained for each TC. First, the TC maintains a list of currently active TGWs. Second, the WFIN that the TC registers with maintains an estimate of the TC's active TGWs. Third, the TC maintains an estimate of the WFIN's estimate of its list of active TGWs. In general, when the TC detects that the WFIN's estimate of active TGW differs from the TC actual list of active TGWs, the TC sends a registration. After the WFIN receives a registration, the WFIN's estimate of the TC's list of active TGWs matches the TGWs list. And when the acknowledgement of the registration is received by the TC, all three lists are identical.

The following scheme is used to reduce registration overhead. This scheme is also depicted in Figures 7 and 8. When

a TC receives a TGW beacon from a TGW, it determines its list of active TGWs (perhaps adding the TGW whose beacon just arrived and perhaps deleting any TGW from which beacons have not recently arrived). The TC then compares this updated list of active TGWs to the TC's estimate of the WFIN's list of active TGW for this node. If these lists are different, then the node will perform two tasks; first it will prepare to register with the WFIN, and second, it appends its list of active TGWs and its estimate of the WFIN's estimate of its list of TGWs on the TGW beacon and propagates the TGW beacon. As is the case with all flooding, the propagation of the TGW beacon follows an efficient flooding algorithm.

Now consider the propagation of the TGW beacon that has been appended with the list of active TGWs. We will refer to the TC that appended the TGW information as the *registering TC* and any TC that receives the appended TGW beacon as a *receiving TC*. The appended active TGW information includes both the registering TC's updated list of active TGWs and its estimate of the WFIN's list of active TGW for the registering client. Upon receiving this appended TGW beacon, the receiving TC first confirms that this beacon has not travel too many hops as compared to other TGW beacons that arrive from other active TGWs. If it has, the TGW beacon is dropped. Otherwise, the receiving TC compares its estimate of the WFIN's list of active TGW to the estimate that is included in the beacon. If these two lists match, then the registering TC and the receiving TC are similar. In this case, when the WFIN receives the registration from the registering TC, it will apply the updated active TGW list to all similar clients, which includes the receiving TC. Thus, if the registering TC and the receiving TC are similar, then the receiving TC updates its estimate of the WFIN's list of active TGW for it to match the list of active TGWs included in the beacon. The receiving TC then updates its actual list of active TGW according to the received TGW beacon (i.e., to reflect the new active TGW represented by the reception of the TGW beacon and/or to reflect the TGW deleted due to missed beacons). In the typical case, the actual list and the newly updated estimate of the WFIN's list match, and hence the receiving node does not prepare to send a registration packet. Rather, it updates the number of hops that the TGW has traveled and propagates the TGW beacon (following the rules of the efficient flooding algorithm). On the other hand, if the receiving TC is not similar to the registering TC or if the receiving TC's estimate of the WFIN's list of active TGWs for it does not match the receiving TC actual list of active TGWs, then the receiving TC prepares to register, and append its estimate of the WFIN's list of active TGW for the it along with its actual list of active TGWs. This extended TGW beacons is then propagated. In this way, a single client registration will serve to register a large number of clients greatly reducing the number of registrations. But, if this registration does not suffice for some nodes, then these nodes will register as well.

The number of registrations is further reduced as follows. When a registering node prepares to register, it first listens for a random duration for other registrations. If it overhears another node transmit a registration and these two nodes are similar (i.e., they have their estimate of the WFIN's estimate of active TGWs is the same), then the node that

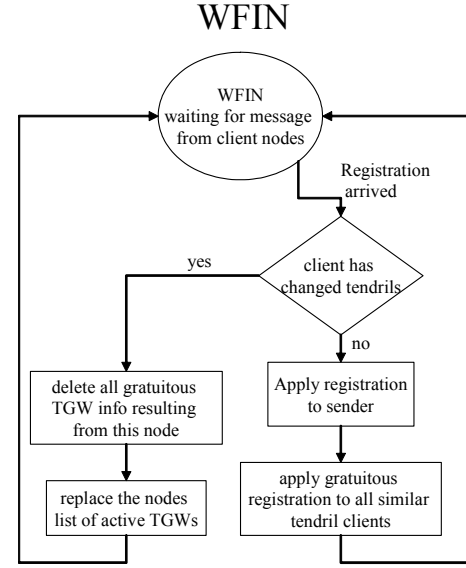


Figure 7: Flowchart describing the functionality of a WFIN on reception of a registration from a client.

has not yet transmitted its registration will terminate the registration process.

While this scheme reduces the number of registration messages, it increases the impact of a dropped or lost registration packet. Hence, a node that transmits a registration packet will repeat the transmission periodically until an acknowledgement from the WFIN is received.

4.3 Detecting a Change of Tendril and Appropriation of Neighbors' Active TGW Lists

4.3.1 Detecting a Change of Tendril

One problem with the scheme above arises when a client node changes tendrils. For example, a TC may be similar to a number of TCs within a tendril. Thus, any registration sent by this TC will be applied to these other TC. However, when this TC moves to another tendril, it will detect new TGWs and hence will eventually register with this new TGW information. Upon receiving such registrations, the WFIN will update the list of active TGWs for all the TCs that are similar to the registering node, which includes TC in the old tendril. Similarly, and more importantly, when a node changes tendrils, it will eventually remove TGWs from its active TGW list. When this TC sends registrations, the registrations will result in the incorrect deletion of the TGW from similar nodes. While the spurious assignment of additional TGWs to the WFIN's estimate of a TC's list of active TGWs will result in page messages being propagated through additional tendrils, the spurious deletion of TGWs could result in page messages not being propagated through the tendril where the TC resides.

In order to reduce such spurious updating of active TGW list, the WFIN will mark all TGW information it has applied to a client as gratuitous and it will record which node or nodes are responsible for the gratuitous registrations. Furthermore, each TC will attempt to detect when it changes

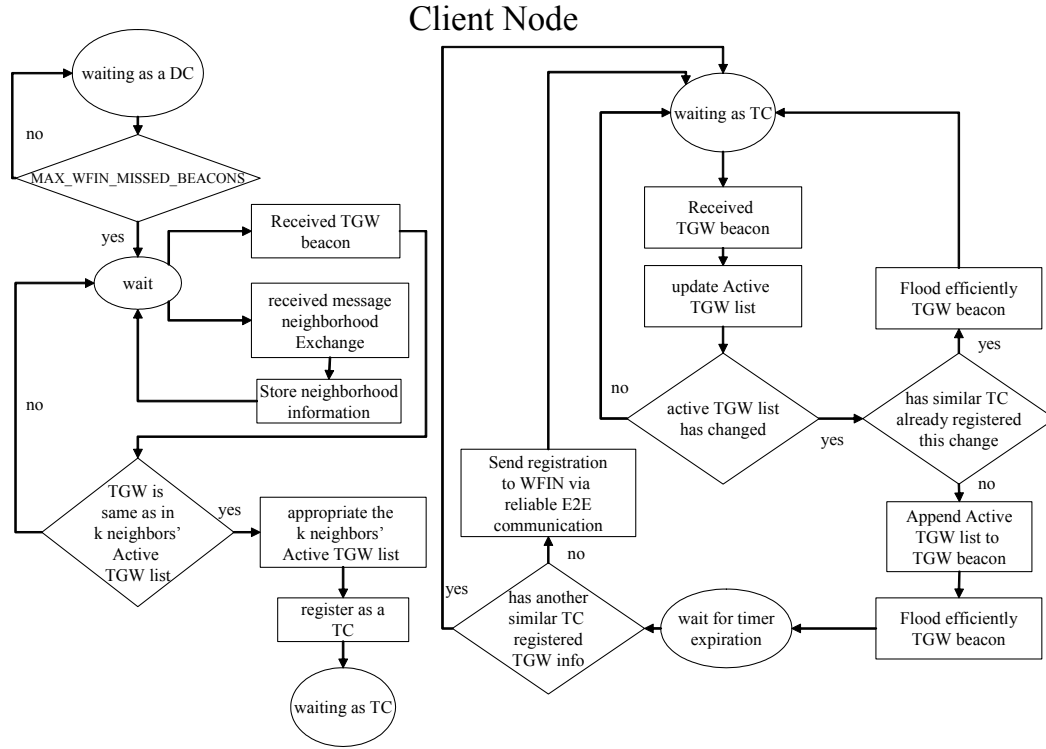


Figure 8: The figure shows the functionality of a client node as a flowchart. The left hand flowchart depicts the functionality of a Tendril Gateway or Direct Client. The right hand flowchart shows the functionality of a Tendril Client

tendrils. When a change of tendril is detected, the node will send a message to the WFIN, which will delete any gratuitous TGW information that resulted from this node.

Furthermore, when a client node detects that it has changed tendrils, it will also appropriate the list of active TGW from its new neighbors. This list of TGWs is sent in a registration message to the WFIN. The WFIN then not only deletes any gratuitous TGW information due to this node, but also replaces the node's list of active TGWs with this new list. In this way, the client instantly becomes similar to the other nodes within the new tendril. These steps used to detect a change of tendril are described next. This algorithm is depicted in Figures 7 and 8. More information about this scheme is provided in the analysis provided in the full version of this paper [6].

As mentioned above, in support of efficient flooding, TCs will periodically broadcast neighborhood information. For example, it is typical to exchange 1-hop neighbor information so that each client is able to construct its 2-hop neighborhood. Piggy-backed on the neighborhood broadcasts, TCs include their list of active TGWs. Thus, a node can easily determine whether it is similar to its neighbors. If a client is not similar to its neighbors, then it is likely that one or more nodes moved. But a difference in neighboring TCs list of active TGWs only indicates that some client has moved, it does not indicate which client moved. However, once a TGW beacon is flooded through the tendril, the TGW beacon will clarify that some active TGWs are correct and others are incorrect. When a sufficient number

of TGW beacons confirm that which of the list of active TGWs is correct, the TC or TCs with an incorrect list of active TGWs will appropriate the list of active TGW from the neighbors with correct a list. This TC (or TCs) will send a special registration message indicating that the TC has changed tendrils and the new list of active TGWs will be added.

Note that the appropriating of neighbor's list of active TGWs greatly speeds up the generation of the list of active TGWs. More specifically, as mentioned above, to reduce the number of registrations and changes in active TGWs, it is important that the active TGW lists not include transitory TGWs (i.e., client nodes that are TGW for only a short time). This can easily be accomplished by only making a TGW active once a TGW when a large number of TGW beacons have arrived. A drawback of this approach is that it increases the amount of time that a TC must remain in a new tendril before it makes the TGWs that serve the tendril into active TGWs. However, by appropriating the list of active TGWs from its neighbors, a TC can determine its active TGW quickly and register with the WFIN, but still reduce the impact of transitory TGWs.

4.3.2 Detecting Entry into a Tendril

When a client moves within a large building it may move to a location where there is a different set of active TGWs (i.e., two tendrils are overlapped). The scheme discussed above will detect such changes. A second way in which a client changes tendrils is when it moves from being an DC to a TC. It is important to distinguish the change of tendril

problem and the change from a client changing from being a DC to a TC. The first scenario typically occurs when tendrils overlap. Hence, while the TC is in the area of overlap, it is reachable via both tendrils (i.e., both sets of TGWs). However, when a client becomes a TC, it means that it is no longer a DC and hence is no longer reachable by the infrastructure. Reachability will only be restored when the TC registers. The duration until the TC becomes reachable is an important part of the performance of the paging scheme. As will be shown in Section 5.1, there is a trade-off between the average duration of the unreachable period and the overhead.

The algorithm to detect when a client has entered a tendril is similar to the scheme to detect that a TC has changed tendrils. We briefly describe the algorithm here. Further details are included in [6] where the trade-off between the duration of unreachability and overhead is analyzed. Also, Figures 7 and 8 depict the essential parts of this algorithm. Briefly, a client becomes aware that it is no longer a DC when a sufficient number of WFIN beacons have been missed, we denote this threshold as $MAX_IN_MISSED_BEACONS$. The client then changes its status as a DC and becomes a TC. This event triggers the broadcast of its neighborhood information to support efficient flooding. This broadcast induces the client's new neighbors to broadcast their neighborhood information. Appended on to the exchanged neighborhood information are the TCs' list of active TGWs. The new TC records this list of TGWs. When a TGW beacon is flooded, the new TC is able to confirm which of its neighbors' active TGWs are valid. The new TC then appropriates the active TGW from these TCs (i.e., the new TC replaces its list of active TGWs with its new neighbors' list) and registers with the WFIN.

4.4 Other Issues

While there are many minor aspects to the paging scheme, some of the more important ones are as follows. While the registration process discussed above is sufficient for a node to always be reachable, the tendril nodes will periodically send registrations to the WFIN. These registrations are useful to confirm gratuitous registrations and provide an upper bound on the worst-case duration that a client is not able to be paged due to incorrect or corrupted registrations (e.g., infrastructure node failure/rebooting). However, rebooting can be accommodated with specialized beacons that request registrations at a particular hop distance. Hence, the WFIN could request registrations in an expanding ring fashion.

In spite of the mechanisms described above, in many cases, when a new TGW becomes an active TGW, several nodes in the tendril will register. For example, if two nodes are within range of a newly active TGW but not within range of each other, they will both register that they are reachable via the newly active TGW. While this slightly increases the overhead due to registration, it also provides robustness in that it reduces the impact of any single TC's registration.

Registrations only occur shortly after a TGW beacon is received from an active TGW. The registration follows the reverse path that the beacon took. Thus, the path is likely to exist. Furthermore, when a TC or DC forwards the registration, it appends its own registration information. Moreover, if a TC desires to make a connection with the WFIN or some node not in the mesh network, the TC can initiate

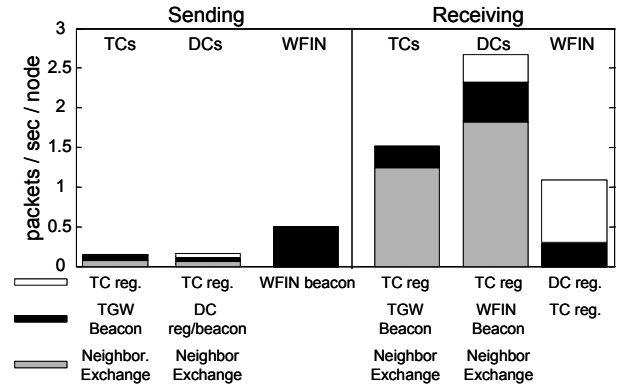


Figure 9: Rates and types of packets that are sent and received. These packet types are those necessary to maintain the client information.

a connection to the base station by using the reverse path followed by the most recently received TGW beacon.

The TGWs do not send TGW beacons into the tendril. Rather, they periodically register with the WFIN. TCs overhear the registration and flood it into the tendril as a TGW beacon.

There are several aspects of this scheme that are not discussed, as they can be achieved using techniques commonly used in mobility management and wireless networking. For example, the infrastructure must maintain and distribute information about which WFIN each client is reachable from. Also, deciding when a client is a DC, and allowing a TGW to become active more quickly when the TC and its neighbors do not have any active TGWs can be attacked using standard techniques that employ thresholds etc.

5. SIMULATION RESULTS

While the full version of the paper includes analytic performance analysis [6], here simulation is used. These simulations used the UDelModels [2] and QualNet [19]. The UDelModels are a set of simulation tools that provide realistic mobility and propagation in urban areas. Several different scenarios were examined, namely, a 1-building/1-WFIN scenario, a 2-building/1-WFIN scenario, a 4-building/2-WFIN scenario, and an 8-building/3-WFIN scenario. All buildings had 10 floors and the node population was approximately $100 \times$ the number of buildings. Note that 100 nodes / building is considerably fewer than typical occupation in a 10 story building. However, as the node density increases, efficient flooding plays an increasingly important role in the performance. While efficient flooding is well studied [21], there has been little work devoted to the specifics of flooding in an urban environment. This paper devoted to paging, and hence, to avoid the impact of the efficient flooding protocol, we restrict our attention to relatively sparse populations.

5.1 Client Information Maintenance

The performance of the scheme clearly depends on the values of parameters. Specifically, we define $T_{WFIN Beacon}$ to be the period between WFIN beacon transmissions, define T_{NB} to be the period between tendril client beacon transmissions (which includes neighborhood information), and define




Figure 10: Mean time for a client node to register after entering a tendril. In these simulations, $MAX_IN_MISSED_BEACONS=6$ and $k=3$. The values of the other parameters are as shown. The two building scenario was used in these simulations.

$T_{TGWBeacon}$ to be the period between TGW beacon transmissions. Figure 9 shows the overhead when $T_{WFINBeacon} = 2.0$ sec., $T_{NB} = 15$ sec., $T_{TGWBeacon} = 20$ sec. These results were generated from the two building scenario, but provide a good approximation into the other scenarios. As expected, little overhead results from the TC registration, but significant overhead is due to the TGW beacons and neighborhood exchanges.

Section 4.3.2 discussed how a client may briefly become unreachable when it changes from a DC to a TC. However, by increasing the rate that beacons are sent, the duration of unreachability can be reduced. Figure 10 shows the relationship between several parameters and the mean time for a node to become reachable (i.e., the mean time to register) after entering into a tendril from being a direct client (more details of this process are provided in the full version [6]).

The parameter values should be selected to achieve a trade-off between the overhead and the average duration until a node registers after entering a tendril. It is important to note that nodes infrequently move from being a direct client to being a TC (indeed, once a node becomes a TC, it must remain a TC until either it has verified that it is no longer reachable from its TGW and is reachable from the WFIN, or the node has a strong connection with the WFIN for a sufficiently long time). Thus, during the course of the day, a node is only unreachable for a small fraction of time. Hence, in order to reduce overhead, it may be desirable to allow extended periods where the node is unreachable. Nonetheless, we select $T_{NB} = 15$ sec., $T_{WFINBeacon} = 0.5$ sec., $T_{TGWBeacon} = 20$ sec., and $MAX_IN_MISSED_BEACONS=6$.

5.2 Route Search/Paging Overhead

The main objective of the above algorithms is to find clients such that the overhead remains bounded as the size of the network increases. Figure 11 shows the overhead for performing a single page/route search using the algorithm presented above as compared to the route search technique




Figure 11: Comparison of overhead for performing a single page using the efficient paging algorithm and route search technique used in AODV. Each building has 10 floors and is populated with 100 nodes.

used in AODV. In the case of AODV, it is assumed that nodes have no prior information as to the location of the desired node. Hence, AODV will flood the network when performing route search.

As shown in Figure 11, the overhead for a single route search grows linearly with the number of buildings when AODV is used, while it remains constant when the scalable paging described here is used. Note that in the case of one building, AODV and the scalable paging perform similarly. This is expected since this paging algorithm aims to flood a single building. (As mentioned above, the impact of efficient flooding is minimal since the node density is low).

6. CONCLUSION

While large-scale urban mesh networks (LUMNets) are being deployed around the world, there remain substantial challenges related to coverage and scalability. Coverage can be enhanced if mobile nodes are permitted to act as relays. In this case, traditional mobility management techniques are not applicable. In particular, paging or searching for a mobile node in such networks has not been addressed. Flooding, which is used in MANET routing protocols cannot be used since it does not scale to the sizes of the LUMNets being developed. This paper presents a scalable technique to page nodes in such LUMNets. The scheme leverages the topology that arises in urban areas. Specifically, the fact that wireless propagation is good outdoors, but greatly restricted indoors produces a topology with a highly connected core that is composed of outdoor nodes, and weakly connected tendrils made up of indoor nodes. The scheme presented restricts the flooding to the tendril where the mobile node resides.

7. REFERENCES

- [1] Geographic Information System (GIS). <http://www.gis.com>.
- [2] BOHACEK S., SRIDHARA, V., AND JONGHYUN, K. UDel Models. <http://udelmodels.eecis.udel.edu>.
- [3] AKYILDIZ, I., HO, J., AND LIN, Y.-B. Movement-based location update and selective paging for PCSnetworks. *IEEE/ACM Transactions on Networking* 4 (1996), 629–638.

- [4] BAR-NOY, A., AND KESSLER, I. Tracking Mobile Users in Wireless Networks. In *IEEE Transactions on Information Theory* (1993), vol. 39, pp. 1877–1886.
- [5] BAR-NOY, A., KESSLER, I., AND NAGHSHINEH, M. Topology-based Tracking Strategies for Personal Communication Networks. In *Mobile Networks and Applications* (August 1996), vol. 1, pp. 49–96.
- [6] BOHACEK, S., SRIDHARA, V., AND LOU, L. Efficient paging in large-scale urban mesh networks. Tech. Rep. Available at <http://udelmodels.eecis.udel.edu>, University of Delaware, 2006.
- [7] US CENSUS BUREAU. Topologically integrated geographic encoding and referencing (TIGER). Available at <http://www.census.gov/geo/www/tiger/>.
- [8] CHANGSHA, W. P., AND LU, X.-C. On the reduction of broadcast redundancy in mobile ad hoc networks. In *Proceedings of the 1st ACM International Symposium on Mobile Ad Hoc Networking and Computing* (2000).
- [9] CITY AND COUNTY OF SAN FRANCISCO. Request for proposals 2005-19 - TechConnect: Community wireless broadband network. Available at <http://www.sfgov.org/>, 2005.
- [10] DAS, S. K., AND SEN, S. K. A New Location Update Strategy for Cellular Networks and Its Implementation Using a Genetic Algorithm. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 1997)* (September 1997), pp. 185–194.
- [11] HAC, A., AND ZHOU, Z. Locating Strategies for Personal Communication Networks: A Novel Tracking Strategy. In *IEEE Journal on Selected Areas in Communication* (Oct. 1997), vol. 15, pp. 1425–36.
- [12] JOHNSON, D., AND MALTZ, D. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, Imielinski and Korth, Eds. kluwer academic publishers, 1996, pp. 153–181.
- [13] KIM, S. J., AND LEE, C. Y. Modeling and Analysis of the Dynamic Location Registration and Paging in Microcellular Systems. In *IEEE Transactions on Vehicular Technology* (Feb 1996), vol. 45, pp. 82–90.
- [14] THE CITY OF MINNEAPOLIS. Request for proposal to provide broadband IP data access services, 2005.
- [15] PERKINS, C. E., AND ROYER, E. M. Ad hoc on-demand distance vector routing. *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications* (1999), 90–100.
- [16] PHILADELPHIA MAYOR’S OFFICE. Government leaders teleconference - wireless: The 21st Century technology, 2004.
- [17] POZZI, F., AND SMALL, C. Vegetation and population density in urban and suburban areas in the u.s.a. *Presented at the Third International Symposium of Remote Sensing of Urban Areas* (June 2002).
- [18] ROSE, C., AND YATES, R. Minimizing the Average Cost of Paging Under Delay Constraints. In *Wireless Networks* (1995).
- [19] SCALABLE NETWORK TECHNOLOGIES. The QualNet simulator <http://www.qualnet.com/>.
- [20] TAIPEI TIMES. Ma says taipei to be wireless in two days, December 19, 2005.
- [21] WILLIAMS, B., AND CAMP, T. Comparison of broadcasting techniques for mobile ad hoc networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)* (2002), pp. 194–205.
- [22] WIST, M. Office space - how much is enough? Tech. rep., Gerald Eve, (2001).
- [23] ZHANG, X., CASTELLANOS, J. G., AND CAMPBELL, A. T. P-MIP: Paging in mobile IP. In *Proceedings of the 4th ACM International Workshop on Wireless Mobile Multimedia* (2001).