

# Wireless Ad Hoc Discovery of Parking Meters<sup>1</sup>

P. Basu and T.D.C. Little

Department of Electrical and Computer Engineering  
Boston University, Boston, Massachusetts 02215, USA

(617) 353-9877

{*basu, tdcl*}@bu.edu

MCL Technical Report No. 01-08-2004

**Abstract**— Locating a suitable parking spot is a common challenge faced by millions of city-dwellers every day. As common is the revenue generation by fee and fine collection in these municipalities. Wireless ad hoc networking technologies offer a new and efficient means to both simplify the process of parking and find collection as well as extending the convenience for drivers. In this paper we describe a multi-hop wireless parking meter network (PMNET) that, when coupled with a GPS receiver, allows a user (driver) to quickly locate and navigate to an available parking space. Our solution is achieved by equipping existing parking meters with wireless radio frequency (RF) transceivers and auxiliary hardware and software. We believe that this is a compelling application that applies wireless ad hoc networking and low-power, short range RF technologies. The attractiveness of the proposal stems from the fact that such a network of nodes can function without any *fixed* wired or wireless infrastructure such as cellular or satellite networks. In this work, we model a PMNET as a special class of ad hoc networks characterized by a combination of static, immobile nodes (parking meters) and mobile nodes (vehicles). We propose scalable techniques for satisfying a mobile user's query in a distributed fashion.

---

<sup>1</sup>*Proc. MobiSys 2004 Workshop on Applications of Mobile Embedded Systems (WAMES 2004)*, Boston, MA, USA, June 2004. This work is supported by the NSF under grant No. ANI-0073843. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

# 1 Motivation

It is not uncommon to see a frustrated group of friends trying to locate parking in a crowded city plaza (such as Manhattan in New York City or Harvard Square in Cambridge, Massachusetts) before heading to their destinations<sup>2</sup>. Now imagine a scenario in which the occupants have a feature that responds to a query of “Find me the closest open spots within 100 yards of Fenway Park” or “Find me a nearby place where two spots are open” or “Find me a 50c spot near Frank’s restaurant which allows parking for more than 2 hours,” and the results are presented on an in-car display almost immediately. The user might select and reserve a parking spot of choice, and pay the parking fee electronically before driving to the spot. Finally, the use of GPS allows a personal navigator to show on-screen driving directions to the selected spot.

A search through US Patent Office’s website indicates several patented techniques for detecting a vacant spot in a parking area using image processing methods and then sharing that information via a centralized wired or wireless network. Our solution proposes to utilize wireless ad hoc networks in order to facilitate the same ends. Details of our scheme have been published elsewhere [2]. In related work, Subramanian and Katz have briefly mentioned “parking lot networks” as being an example of generic self-organizing systems [3].

## 2 Architecture of a Parking Meter Network

Existing parking meter function needs to be augmented in order to be able to support the application described in Section 1. Parking meters (or *pm-nodes*) will require a low-cost, low-power embedded processor, programmable memory and operating system. They require an inexpensive infra-red (IR) sensor to detect occupancy; and a short-range, low-power RF transceiver to communicate with neighboring meters and nearby mobile vehicles. Since the transmission range of Bluetooth’s [5] radios ( $\approx 10 m$ ) may be too limiting to maintain a connected ad hoc network, we believe that Millennial Net’s i-Bean [4] is a better candidate technology since it can support transmission ranges upto 40 *m* and has nice energy saving features. We leave the detailed investigation of how to leverage i-Bean technology for this application as future work.

---

<sup>2</sup>In a recent book [1], Calvin Trillin eloquently captures the psyche of a New Yorker, Tepper who enjoys sitting in his car all day reading the newspaper, just for the opportunities to turn other would-be parkers away. The most closely held secret in the story concerns the easy availability of parking meters on Madison Avenue after 7 PM.

A stationary pm-node has a fixed geographical location attribute which can be exploited during the discovery process. Instead of putting GPS receivers into every parking meter, meter deployment can include the use of a portable GPS receiver to permanently burn the geographical coordinates of each pm-node into its memory. A locally unique ID along with the GPS location attribute can serve as a globally unique ID (GUID). These attributes and some others such as “street location” are *non-volatile* attributes. Other attributes that are likely to be stored at a pm-node at any time instant include the size of the spot, its current availability, the fee for use, and the time limit. These are referred to a *volatile* attributes. Wirelessly enabling these devices results in their connection to a virtual network of devices thus allowing them to function as resources that can be discovered remotely, instead of physically.

In addition to supporting drivers seeking parking, PMNETs can be used by the municipality to simplify and optimize revenue generation from fees or penalties. Arguably a game-theoretic analysis of the distinct goals of individuals vs. municipalities can yield a variety of approaches to parking rules. Some representative queries that might be posed by the fee collector or driver are:

Q0: Which streets in the locality have vacant spots right now?

Q1: Which meters are about to expire within 200 yards of my location (GPS coordinate specified)?

Q2: Is there any parking spot scheduled to become available soon on Main St. until midnight?

Q3: Locate all vehicles that reside in expired spots.

A query processing software module running on a user’s unit structures a query in a suitable query language before it is dumped in the network via pm-nodes that are RF-reachable from the user’s current location. Responses to a query can be sorted by user defined criterion before presentation to the user. The process can be visualized as querying a *physically distributed* database formed by pm-nodes possessing dynamically changing attributes. In this paper, we do not address the details of structuring a user query or a response since our focus is on algorithms for discovery and scalable updates which are described in the next section.

### 3 Ad hoc Parking Spot Discovery

Although the pm-nodes are themselves static, the content served by them changes dynamically. In addition to the non-volatile attributes such as geographic location, it is this dynamic content (e.g., is the spot available now?) which is often of great interest to the users. Most mobile user queries require both static and dynamic attributes for resolution.

A *query processing* module receives structured queries from mobile users, determines if a query matches the current values of its relevant attributes and takes appropriate action. If a query can be satisfied locally, the pm-node responds to the user using the underlying unicast routing protocol described in Sec. 3.3. In the basic version of the system, if a pm-node is unable to satisfy the query itself, it *rebroadcasts* the query to its neighbors. In a more sophisticated version of the system as proposed in Sec. 3.1, query processing is performed in a hierarchical fashion on pm-nodes which act as ClusterHeads (CH) for a set of other pm-nodes.

The *status update processing* module receives status updates (values of volatile attributes) from other pm-nodes and updates the local data structures accordingly. These data structures need to be queried repeatedly while responding to user queries. We advocate a hybrid approach (partly proactive and partly reactive) for facilitating querying and status update so as to minimize bandwidth usage as well as user latency.

#### 3.1 Cluster-based Status Updates

A hybrid cluster-based approach which utilizes the natural geographical clustering in PMNETs is ideal for facilitating quick status updates and speeding up the discovery process. Non-volatile attributes such as street name can be used to *cluster* a set of pm-nodes together. The radio transmission range is assumed to be large enough so that the pm-nodes on opposite sides of a street can communicate with each other.

**Election of Clusterheads:** In our approach, exactly one of all pm-nodes with the same *street* attribute serves as a clusterhead (CH) of the set. A CH can be elected to its role during installation or can automatically be selected by a simple leader election process executed by all pm-nodes<sup>3</sup>.

---

<sup>3</sup>Each pm-node broadcasts to its neighbors a  $(ID, street, val)$  tuple where *val* corresponds to a measure of its computing resources. For example, *val* can be a function of the amount of CPU and memory resources, as well as the remaining battery power (if applicable). Every pm-node then periodically neighbor-broadcasts

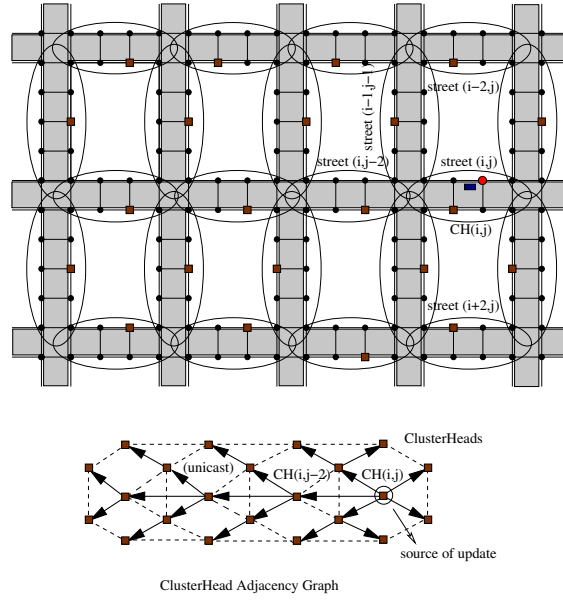


Figure 1: Scalable Status Update Dissemination

**Dissemination of Status Updates:** A non-CH pm-node of street  $X$  upon hearing about the election of a CH in a neighboring street  $Y$  informs its own CH about it. Thus every CH comes to know about the CH's in the neighboring streets. This is useful for efficient dissemination of current status information as the CH nodes in the entire locality form an overlay graph and the dissemination of status updates is performed only from the affected CH to the rest of the CH's by means of unicast. Specifically, each pm-node reports a change in its status (vacant  $\rightarrow$  full, or vice versa) to its CH which then floods this incremental information to other CH's in the PMNET. As a result of this, a user node only needs to query a nearby CH about the attributes/status of pm-nodes in other streets while discovering a spot.

Fig. 1 illustrates the process of performing a status update with a specific example. When a car leaves a parking spot on street  $(i,j)$ , the corresponding pm-node (marked by a red filled circle) informs its clusterhead,  $CH(i,j)$  about this change. Now,  $CH(i,j)$  needs to send a status update to all other CH's in the locality. This can be accomplished in a number of ways. In the first approach, a CH attempts to send reliable updates to its neighboring CH's over unicast (utilizing the underlying unicast routing substrate described in Sec. 3.3)

---

the maximum *val* from the same *street* that it has heard so far. The pm-node with maximum *val* is chosen as the CH of the particular street by every pm-node therein. Ties are broken by *IDs*. This process consumes little bandwidth as only local communication happens, and that too after installation, failure, or repair of a pm-node. It can also be triggered when the existing CH's battery is about to run out, i.e., if the pm-nodes run on battery power.

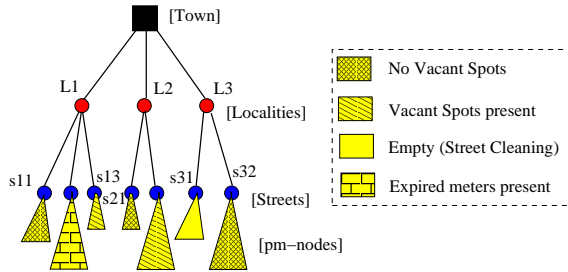


Figure 2: Hierarchical Name Trees

which then re-forward the updates further downstream. To determine who to forward the updates to, a Clusterhead Adjacency Graph (CHAG) is constructed from the topology of street intersections<sup>4</sup>. Now, unicast forwarding needs to be done only along the edges of a minimum spanning tree of CHAG rooted at  $CH(i,j)$ <sup>5</sup>.

There are several trade-offs between latency and control overhead that can be exploited in this system. A CH can choose to disseminate only a salient part of the status update to other CH's, e.g., if a spot in street S becomes vacant,  $CH(S)$  can inform other CH's of only the fact that there is an additional vacant parking spot available, instead of furnishing the details of the spot. This can help in quick resolution of queries like [Q0]. However, when a user in street X issues a more involved query,  $CH(X)$  may not be able to answer that query immediately from this delta-information, but it can query the relevant CHs about the details of the parking spots at the current time. This approach is advantageous as a CH needs to store information only at a coarse level of granularity. The obvious penalty is that of additional time taken to discover a spot. The best approach in a particular scenario would depend on factors such as the predominant nature of queries and their frequency, the frequency of changes, size of the PMNET etc.

### 3.2 Hierarchical Name-trees

Hierarchical name-trees (similar to MIT's Intentional Naming System [6]) can be useful for answering queries such as [Q2]. A pm-node's unique GUID, non-volatile attributes, and useful associated information such as nearby landmarks are programmed set at the time of installation. When CH's exchange information about pm-nodes, they can create tree-

<sup>4</sup>Two CH's are adjacent in a CHAG if and only if their corresponding streets intersect with each other.

<sup>5</sup>In the second approach, the dissemination problem is treated as a multicast routing problem where all the CH's are sinks in a multicast group and  $CH(i,j)$  is the source. This approach however assumes that all pm-nodes support multicast. Moreover, it is easier to guarantee reliability in updates in this approach in contrast with doing that in the pure multicast based solution.

like data structures with actual pm-node records as leaf nodes, and “towns”, “streets” and “localities” as *virtual* nodes in the hierarchy. Fig. 2 depicts a possible instance of a *name-tree* which is maintained at every CH pm-node in a street. The triangular blocks hanging from the “Street” nodes depict the individual pm-nodes. The relative sizes of triangles indicate the numbers of pm-nodes existing in each street. The various visual patterns inside them characterize the status of parking on a particular street as a whole at a particular instant of time. Such hierarchical presentation of information is possible because of the hierarchical geographic clustering of the pm-nodes.

When the a pm-node’s status changes, it informs its CH, which then reflects the change in its name-tree. The latter also *floods* the update to the other CH’s using the mechanisms described in Sec. 3.1. When a query such as [Q2] is issued by the user, it gets routed to the nearest CH pm-node and the query processing module on that pm-node searches under the “Street” metanode in its name-tree. After finding the “Main St.” node, it searches for suitable parking meters (children of that node in the attribute tree) which can satisfy the request. Note that the status updates need to occur on a detailed/individual pm-node basis and not just at an aggregated level such as “Main St. has a 2 vacant spots now.” If aggregates updates are being used, then the current CH has to query the CH(Main St.) pm-node to get more details.

The name-tree shown in Fig. 2 can also be visually presented to a user of the system when he/she submits a general query “Display the overall status of parking in the area”. It can be an invaluable tool for the fee collectors who then would not have to go to each individual parking meter and verify if the meter has expired or not. Ideally, a fee collector can on his/her PDA, zoom in on the streets which have expired meters and navigate to the meters individually. This can save significant effort and resources of the municipality.

### 3.3 A Hybrid Unicast Routing Architecture

After discovering a set of candidate spots, a user may initiate communication with some of them to check the validity of the status or to reserve a spot, if possible. Unicast routing is necessary to facilitate this. Because pm-nodes are static and possess a fixed GPS location attribute, the topology of the static part of PMNET changes infrequently. Therefore, a dynamic MANET routing protocol is unnecessary in its full generality for routing between two pm-nodes.

We believe that position based routing schemes [7] are ideal in static scenarios since they

do not need to exchange packets or store full routing table information. In position based routing, packets are routed hop-by-hop in a greedy fashion by forwarding to the neighboring pm-node which is *closest* to the destination location coordinate<sup>6</sup>.

Now, since users of a PMNET are usually mobile, their positions change continuously until they find a parking spot. Hence, the location coordinates of a user at the time of issuing the query may not be exactly the same when a pm-node issues a response. Therefore, static position based routing needs to be augmented slightly in order to facilitate routing in the reverse path (from pm-node to U). Routing table entries with reverse paths to U are setup at the intermediate pm-nodes during the forward routing (greedy location-based forwarding) phase, thus enabling reverse routing of messages to the user at a slightly later time instant.

However, even the above protocol augmentation is inadequate for handling routing from a pm-node to a user or between two users at a much later instant of time since the reverse routing table entries may have become stale owing to the movement of the user from his/her earlier location. The former mode of communication may be needed when a particular pm-node wants to inform a user about competitors for a parking spot or if it wants to warn the user about the expiring time. The latter mode can be used by anonymous users for negotiations about a spot. In order to solve this problem, we believe that the use of a cluster-based MANET routing protocol (piggybacked onto the status updates) is appropriate.

### 3.4 Competition between Vehicles

When multiple users are looking for parking in a crowded area, the scenario becomes much more interesting and a competition between drivers begins for a small number of available spots. In a baseline version of the system, vehicles which do not ultimately get the spot, query the system again from their new locations hoping to find a new spot. However, the system can be improved further by including more state information; e.g., if vehicles  $v_1$  and  $v_2$  are competing for the same spot  $P$ , the corresponding pm-node can inform one about the other's current location. If  $v_1$  is physically much farther from  $P$  than  $v_2$ , it cancels its query and searches for another spot, therefore saving time by not driving to  $P$ . More complex stochastic allocation problems can be formulated in this context but we leave that

---

<sup>6</sup>In regions of the network where such a greedy path does not exist (i.e., the only path requires that one move temporarily farther away from the destination), position based routing schemes recover by forwarding in *perimeter mode*, in which a packet traverses successively closer faces of a planar subgraph of the full radio network connectivity graph, until reaching a node closer to the destination, where greedy forwarding resumes. An example of a position based routing protocol which behaves in this manner is Greedy Perimeter Stateless Routing [8].



discussion for the paucity of space.

### 3.5 Fee Structure and User Behavior

Although we have not addressed fee models or rules in this paper, we fully expect that driver behavior is coupled with the fee structure and the competing goals of the municipality in system design. For example, if a driver violates the time limit, the system might automatically debit his/her account instead of issuing a ticket. Our proposal is amenable to such variants, including an increase in the parking fee in a non-linear fashion after the parking meter expires. For example, the 2 hour fee could be \$2, but if one stays for 3 hours, it could increase to \$10, and then to \$30 after 4 hours. The fee structure could thus determine user behavior and vice versa.

## 4 Conclusion

In this paper, we proposed an application for a wireless parking meter network in which drivers can quickly locate and navigate to available parking spaces. Enabling existing parking meters with IR occupancy sensors and low power RF transceivers can yield a self-organizing and self-managing distributed system while relying on techniques in mobile ad hoc networking. We described in detail algorithms for efficient discovery of vacant spots as well as scalable dissemination of status updates in order to reduce both discovery latency and the bandwidth usage. We believe that the proposed network is both economically and technologically feasible, although the details of a complete solution are not shown here. For example, the issues of energy management, failure and partition tolerance, and security are significant for the proper functioning of the proposed system and were not discussed in this paper. We plan to investigate these in the future.

## References

- [1] C. Trillin, "Tepper Isn't Going Out," *New York: Random House*, 213 pp.
- [2] P. Basu and T. D. C. Little, "Networked Parking Spaces: Architecture and Applications," *IEEE Vehicular Technology Conference (VTC)*, Vancouver, Canada, September 2002.

- [3] L. Subramanian and R. Katz, “An Architecture for Building Self-Configurable Systems,” *Proc. MobiHoc 2000*, Boston, MA, August 2000.
- [4] Millenial Net’s i-Bean. <http://www.millenial.net>
- [5] Bluetooth SIG. <http://www.bluetooth.com>
- [6] W. Adjie-Winoto, E. Schwartz, H. Balakrishnan, and J. Lilley, “The design and implementation of an intentional naming system,” *Proc. ACM SOSP 1999*, Kiawah Island, SC, December 1999.
- [7] M. Mauve, J. Widner, and H. Hartenstein, “A Survey on Position-Based Routing in Mobile Ad-Hoc Networks,” *IEEE Network*, November 2001.
- [8] B. Karp and H.T. Kung, “Greedy Perimeter Stateless Routing for Wireless Networks,” *Proc. ACM MobiCom 2000*, Boston, MA, August, 2000, pp. 243–254.