

# LIRA: Lightweight, Region-aware Load Shedding in Mobile CQ Systems

Buğra Gedik<sup>◇♣</sup>

Ling Liu<sup>◇</sup>

Kun-Lung Wu<sup>♣</sup>

Philip S. Yu<sup>♣</sup>

<sup>◇</sup> CERCs, College of Computing, Georgia Tech

<sup>♣</sup> Thomas J. Watson Research Center, IBM Research

{bgedik,lingliu}@cc.gatech.edu, {klwu,psyu}@us.ibm.com

## Abstract

To provide high-quality results for location-based, continual queries (CQs) in a mobile system, the query processor usually demands receiving frequent position updates from the mobile nodes. However, processing frequent updates often causes the query processor to become overloaded, under which updates must be dropped randomly, bringing down the query-result accuracy and negating the benefits of frequent updates. In this paper, we develop LIRA – a lightweight, region-aware load-shedding technique for preventively reducing the position-update load of a query processor, while maintaining high-quality query results. Instead of receiving too many updates and then randomly dropping some of them, LIRA uses a region-aware partitioning mechanism to identify the most beneficial shedding regions to cut down the position updates sent by the mobile nodes within those regions. Based on the densities of mobile nodes and queries in a region, LIRA judiciously applies different amounts of update reduction for different regions, aiming to minimize the negative impacts of load shedding on query-result accuracy. Experimental results show that LIRA is vastly superior to random update dropping and clearly outperforms other alternatives that do not possess region-aware load-shedding capabilities. Moreover, due to its lightweight nature, LIRA introduces very little overhead.

## 1 Introduction

The proliferation of mobile devices and advances in wireless communications are creating an increasing interest in rich, value-added location-based services, which are expected to form an important part of the future computing environments that will seamlessly integrate into our lives [17]. A recent example from the industry is the Google Ride Finder [6] service, which provides mobile users with the capability to employ CQs to monitor nearby taxi services.

Mobile CQ systems serve as an enabling technology for location monitoring applications. Scalable CQ middleware for location monitoring has been an active area of research, attested by several recent works, such as SINA [12], SRB [8], MAI [5], and others [3, 1, 13]. In almost all of these systems, position updates and query re-evaluations are two predominant and costly components of processing location-based CQs, consuming CPU, disk, and wireless network resources.

To produce high-quality query results, the query processor usually demands receiving frequent position updates from the mobile nodes. However, receiving and processing frequent updates often causes the query processor to become over-

loaded. As a result, the update problem in mobile CQ systems has received significant attention from the research community, producing several spatial index structures for efficiently integrating position updates into the system [18, 9, 11, 21]. Although indexes can speed up the processing of position updates, they do not solve the fundamental problem of overload. When overloads happen, the position updates will clog the system buffers and cause updates to be dropped randomly, which (as we show in this paper) is an ineffective way to handle overload. Surprisingly, none of the prior works have addressed the problem of effective *update load shedding*. Hence, there is a cogent need for developing intelligent update load-shedding techniques for mobile CQ systems. The load shedding algorithms should prevent overloads by reducing the number of position updates received by the query processor, while minimizing the side-effects on query-result accuracy.

In this paper, we develop a lightweight load-shedding technique for reducing the update load in mobile CQ systems, called LIRA. The main idea behind LIRA is that, given an update budget (which is either calculated automatically by LIRA or specified as a system-level *throttle fraction* parameter), LIRA creates a partitioning of the monitoring space into a set of *shedding regions* and associates an *update throttler* with each shedding region, where these update throttlers define the amount of load shedding to be performed for each region in accordance with the overall update budget. Generally, update load shedding decreases the quality of query results. In LIRA both the partitioning and the settings of the update throttlers are performed with the objective of minimizing the negative impact of load shedding on the accuracy of query results.

LIRA has four unique properties. First, the partitioning scheme employed by LIRA is *region-aware*. Contiguous geographical areas that have similar characteristics in terms of densities of mobile nodes and queries are grouped into the same load shedding regions. Second, the update throttlers are set based on the following principle: the regions where update load shedding may cut down a large number of updates while reducing the query-result accuracy minimally, are subjected to larger amounts of load shedding. Third, LIRA provides an adjustable bound on the maximum difference between the update throttlers of different shedding regions, ensuring that all mobile nodes are tracked by the system, albeit with varying accuracies. This feature extends the applicability of LIRA to mobile CQ systems with snapshot and historical query support. Last but not the least, LIRA introduces very little overhead and can be employed in conjunction with CQ systems that use update-efficient indexes, like TPR-tree [16].

We evaluate our load shedding approach using realistic location data synthetically generated using existing road maps and real-world traffic volume data. We devise a set of evaluation metrics to assess the effectiveness of LIRA and empirically show that LIRA is vastly superior to update dropping and clearly outperforms other alternatives that do not provide full-scale, region-aware load shedding capabilities.

## 2 Overview

In this section we describe the fundamental concepts underlying the LIRA load shedder, introduce some of the notations used in the paper, and present the system architecture.

### 2.1 Design Ideas

There are two major types of load shedding techniques for reducing the number of position updates processed by the query processor: *server-actuated* and *source-actuated*. In server-actuated load shedding, the position updates are dropped by the CQ server in order to match the update arrival rate with the server service rate. This has two disadvantages. First, the dropped updates are unnecessarily transferred from the mobile nodes to the CQ server, wasting the wireless network bandwidth. Second, the excessive updates still have to be received by the server, and thus contribute to the processing load. On the other hand, the source-actuated approach requires some coordination between the server and the mobile nodes, since the load shedding decisions are made by the server. With a lightweight process for coordination, LIRA uses the source-actuated approach for load shedding.

#### 2.1.1 Adjusting the Inaccuracy Threshold

A commonly used mechanism for actuating the position update reduction on the mobile nodes is motion modeling, also known as *dead reckoning*. Motion modeling uses approximation for location update prediction. Instead of reporting their position updates each time they move, mobile nodes only report the parameters of a model approximating their motion when the model parameters change significantly, which is decided based on an *inaccuracy threshold*  $\Delta$  and the last reported model parameters. When the predicted position of a mobile node deviates from its actual position by more than  $\Delta$ , the new motion parameters are reported, i.e., a position update is sent by the mobile node. In LIRA this inaccuracy threshold  $\Delta$  is used as a control knob to adjust the number of position updates sent by a mobile node.

A popular motion model is piece-wise linear approximation of the mobile node movement [20], whereas more advanced models also exist [2]. Without loss of generality, we adopt linear motion modeling in LIRA, because the particular motion model used is not of importance for this paper. Note that many of the existing mobile CQ systems have built-in support for linear motion modeling [16, 18, 5].

A straightforward but naive way of shedding update load is to have all mobile nodes use a single system-controlled inaccuracy threshold. Let  $\Delta_{\perp}$  be the minimum value that the inaccuracy threshold can take, which de-

finer the ideal resolution of position updates. Let  $\Delta_{\dashv}$  be the maximum value that the inaccuracy threshold can take, which defines the lowest resolution of position updates required to achieve reasonable query-result accuracy.

The inaccuracy threshold  $\Delta$  can be set to a value within  $[\Delta_{\perp}, \Delta_{\dashv}]$  to adjust the update expenditure of the system. By increasing  $\Delta$  from  $\Delta_{\perp}$  to  $\Delta_{\dashv}$ , the number of updates will decrease even though this reduction is not linear as shown in Fig. 1<sup>1</sup>.

It plots  $f(\Delta)$ , called the *update reduction factor*. For a given inaccuracy threshold  $\Delta \in [\Delta_{\perp}, \Delta_{\dashv}]$ ,  $f(\Delta)$  gives the number of position updates received relative to the case of  $\Delta = \Delta_{\perp}$ . As observed from Fig. 1, the rate of reduction in the update expenditure is more pronounced while  $\Delta$  is increased within the proximity of  $\Delta_{\perp} = 5$  meters, whereas it reduces to a fixed slope (linear decrease in the number of updates) as  $\Delta$  gets closer to its maximum value of  $\Delta_{\dashv} = 100$  meters.

A key observation we make in this paper is that different regions exhibit different characteristics in terms of the densities of mobile nodes and queries and can benefit from differing amounts of load shedding. This observation suggests that a uniform  $\Delta$  approach is significantly suboptimal. To understand this better, we plot the desirability of load shedding for regions with varying characteristics in Table 1.

		$m$	
		low	high
$n$	low	<	×
	high	✓	>

Table 1: Region characteristics and preference of load shedding

Let  $n$  be the number of mobile nodes and  $m$  be the number of queries within a region. When  $n$  is low and  $m$  is high for a region, load shedding should be avoided as much as possible (upper right quadrant in Table 1). This is because a small number of nodes that generate few updates are used for answering a large number of queries for this region. This implies that increasing  $\Delta$  here will significantly impact the overall query accuracy, while bringing only a small reduction in the number of position updates sent to the server. In contrast, load shedding is very desirable when  $n$  is high and  $m$  is low for a region (lower left quadrant in Table 1). In this case, a large number of nodes that generate many updates are used for answering few queries. It implies that increasing  $\Delta$  will minimally impact the overall query accuracy, while bringing a large reduction in the number of position updates sent to the server. Interestingly, the ratio  $m/n$  does not completely characterize the preference of one region over another for increasing the inaccuracy threshold  $\Delta$ . This is because the overall inaccuracy introduced in the mobile node positions increases linearly with increasing  $\Delta$ , whereas the amount of update reductions in-




Fig. 1: Update reduction factor

<sup>1</sup>The experimental setup and default parameters used to generate the graphs in Fig. 1 are given in Section 4.2

creases non-linearly as  $\Delta$  increases. This is why regions with small  $m$  and  $n$  are less attractive for load shedding compared to the regions with large  $m$  and  $n$  (the symbols  $<$  and  $>$  indicate this in Table 1), but compared to the scenario of high  $m$  and low  $n$  both are better choices.

### 2.1.2 Region-aware Load Shedding

This insight leads us to a region-aware approach to update load shedding. In LIRA we partition the geographical area of interest into  $l$  shedding regions, denoted by  $A_i, i \in [1..l]$ . Furthermore, we associate an inaccuracy threshold with each shedding region  $A_i$ , denoted by  $\Delta_i$ . We call  $\Delta_i$  the update throttler of the region  $A_i$ . A simple way of determining the shedding regions is to partition the entire geographical space of interest into  $l$  regions evenly. However, such even partitioning of the space is unlikely to produce an effective solution, since the level of heterogeneity (in terms of the number of mobile nodes and queries) inside two given equally-sized regions may differ significantly. Intuitively, a region where further partitioning generates sub-regions of similar characteristics in terms of the densities of mobile nodes and queries does not provide any gain with regard to reducing the number of position updates while minimizing the query-result inaccuracy. Thus the design of the LIRA load shedder should address the following two challenges: (1) How to partition the geographical space of interest into a set of shedding regions effectively? (2) How to set the update throttler for each shedding region to minimize the inaccuracy introduced in query results while meeting our update budget constraint?

In LIRA we introduce the concept of throttle fraction to define the position update budget of the system, denoted by  $z \in [0, 1]$ . For instance,  $z = 0.75$  means that the number of updates should be reduced by a quarter, compared to the case of using a common  $\Delta = \Delta_{\perp}$ . The throttle fraction can be calculated automatically by the server in reaction to overload situations by observing the size of the system message queue (see Section 3.4). Alternatively, when the server is not overloaded but the wireless communication load of receiving updates are putting a heavy burden on the network, the throttle fraction can be manually set as a system-level parameter.

## 2.2 System Architecture

Fig. 2 illustrates the system architecture of LIRA, which consists of three layers. The first layer is formed by the mobile CQ server. The server has three main responsibilities. First, it sets the throttle fraction  $z$  to define the position update budget of the system. Second, it is responsible for calculating the shedding regions and the associated update throttlers for a given update budget. Third, it is responsible for reporting to each base station in the second layer, the subset of shedding regions and update throttlers corresponding to the base station's coverage area.

The set of base stations that cover the space of interest form the second layer. The base stations are assumed to be connected to the mobile CQ server via the wired network. They provide wireless networking services to the mobile nodes.

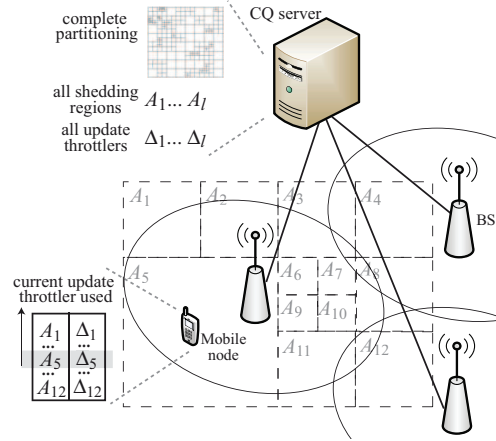


Fig. 2: System Architecture

The base stations are responsible for broadcasting the subset of load shedding regions and update throttlers corresponding to their coverage area to mobile nodes, when the server reports a change in the partitioning or the update throttler values. The base stations are also responsible for sending the shedding regions and update throttlers to a mobile node entering into a new base station's coverage area during a hand-off.

The set of mobile nodes form the third layer of our system. The mobile nodes are responsible for reporting their positions to the mobile CQ server using dead reckoning. However, the inaccuracy threshold used by a mobile node is dependent on the region in which it resides. As a result, the mobile nodes store a subset of shedding regions and update throttlers corresponding to the coverage area of their current base station. As they move from one shedding region to another within their base station's coverage area, the mobile nodes use the update throttler corresponding to their current shedding region as the inaccuracy threshold. This is performed locally. When the mobile nodes switch base stations, they change the subset of shedding regions and update throttlers they store based on the information they receive from the new base station.

### Factors Affecting the Number of Shedding Regions

From Fig. 2, one may observe an interesting trade-off in setting the number of shedding regions  $l$ . On the one hand, the larger the number of shedding regions, the more fine grained the partitioning, leading to more fully exploiting the potential heterogeneity existent in terms of densities of mobile nodes and queries. On the other hand, as the number of shedding regions increase, the average number of update throttlers and shedding regions per base station coverage area grows, increasing the cost for each mobile node in terms of both computation and communication. Thus a careful setting of  $l$  is critical for the overall system scalability in terms of both service quality and wireless communication bandwidth.

## 3 The LIRA Load Shedder

In this section we describe the main technical components of the LIRA load shedder, encompassing the three major server-side functionalities: (1) partitioning the geographical space of interest into  $l$  shedding regions for a given  $l$ , per-

formed by the GRIDREDUCE algorithm, (2) determining the update throttler for each of the  $l$  shedding regions, performed by GREEDYINCREMENT algorithm, and (3) setting the throttle fraction  $z$  to adjust the system-wide position update budget, performed by THROTTLOOP algorithm. These three algorithms work in cooperation to perform the load shedding. In particular, the THROTTLOOP algorithm monitors the performance of the system under the current workload and resource availability to decide the throttle fraction  $z$ . Given  $z$  computed by THROTTLOOP and the number  $l$  of shedding regions specified as a system-supplied parameter, the GRIDREDUCE algorithm creates a partitioning of the entire geographical space of interest and computes the set of  $l$  shedding regions, i.e.,  $A_i$  ( $i \in [1..l]$ ). Finally, given  $z$ ,  $l$ , and  $A_i$ 's, the GREEDYINCREMENT algorithm determines the update throttlers for the  $l$  shedding regions, i.e.,  $\Delta_i$  ( $i \in [1..l]$ ).

### 3.1 Problem Formulation

The problem is to find a partitioning of  $A_i, i \in [1..l]$ , and an associated set of update throttlers  $\Delta_i, i \in [1..l]$ , such that certain constraints are met (e.g., the update budget is respected) and an objective function is optimized (i.e., inaccuracy in query results is minimized). We start with formulating the two constraints. Let  $n_i$  denote the number of mobile nodes within shedding region  $A_i$ . The following two constraints should hold:

$$\begin{aligned} \sum_{i \in [1..l]} n_i \cdot f(\Delta_i) &\leq z \cdot n \cdot f(\Delta_{\vdash}) \\ \forall_{i \in [1..l]} \Delta_{\vdash} &\leq \Delta_i \leq \Delta_{\dashv} \end{aligned}$$

The first constraint, which we call the *update budget constraint*, states that the number of updates received under the region-aware load shedding approach should not exceed  $z$  (throttle fraction) times the number of updates that would have been received if there were no load shedding applied, i.e., we were to use a uniform inaccuracy threshold of  $\Delta_{\vdash}$  for all nodes. Note that we have  $f(\Delta_{\vdash}) = 1$ . The second constraint defines the domain of update throttlers ( $\Delta_i$ 's).

We now formulate the objective function of the problem we want to minimize, that is the inaccuracy in query results. For the purpose of this problem formalization we define the inaccuracy introduced by using an update throttler value of  $\Delta_i$  for a given region  $A_i$  as the number of queries in the region  $A_i$ , denoted by  $m_i$ , times the inaccuracy threshold  $\Delta_i$ , that is  $m_i \cdot \Delta_i$ . When computing  $m_i$ , queries partially intersecting the shedding region  $A_i$  are fractionally counted. The objective function that we want to minimize can be formulated as:

$$InAcc(\{A_i\}, \{\Delta_i\}) = \sum_{i \in [1..l]} m_i \cdot \Delta_i$$

Note that  $m_i$  and  $n_i$  are functions of the partitioning  $\{A_i\}$ . We now discuss an important extensions to this basic problem, in order to provide a system-level control over the difference in the inaccuracy thresholds used in different regions.

#### 3.1.1 The Fairness Threshold

We introduce a parameter called the fairness threshold, denoted by  $\Delta_{\Leftrightarrow}$ . In the original problem formulation, the shedding regions that do not contain any queries (i.e.,  $\{A_i : m_i = 0\}$ ) may be overly penalized by setting their update throttlers to maximum inaccuracy value of  $\Delta_{\dashv}$ , since the update reduction for those regions does not impact the query results. However, for mobile CQ systems supporting historic and ad-hoc queries this may be undesirable, thus  $\Delta_{\Leftrightarrow}$  can be used to reduce this effect. Formally, we have the following additional constraint on the update throttlers:

$$\forall_{i,j \in [1..l]} |\Delta_i - \Delta_j| \leq \Delta_{\Leftrightarrow}$$

One extreme case of  $\Delta_{\Leftrightarrow} = \Delta_{\dashv} - \Delta_{\vdash}$  represents the original formulation, whereas the other extreme case of  $\Delta_{\Leftrightarrow} = 0$  represents the uniform  $\Delta$  scenario.

We consider the partitioning and the setting of update throttlers as separate problems. In what follows, we first provide a heuristic-based partitioning algorithm for constructing the shedding regions and then give an optimal (under certain conditions) algorithm for setting the update throttlers for a given partitioning of the space. It is worth mentioning that the problem of setting the update throttlers is not a linear program, since the update reduction function  $f$  is not linear and as a result the update budget constraints are not linear.

### 3.2 GRIDREDUCE: Partitioning the Space

The goal of the GRIDREDUCE algorithm is to partition the geographical space of interest into  $l$  shedding regions, such that this partitioning produces lower query-result inaccuracy. For each shedding region  $A_i$  generated, the algorithm also determines the number of nodes  $n_i$  and the number of queries  $m_i$  for that region. This information is later used by GREEDYINCREMENT to set the update throttlers.

The GRIDREDUCE algorithm works in two stages and uses a *statistics grid* as the base data structure to guide its decisions. The statistics grid serves as a uniform, maximum fine-grained partitioning of the space of interest. In the first stage of the algorithm, which follows a bottom-up process, we create a region hierarchy over the statistics grid and aggregate the query and mobile node statistics for the higher-level regions in this hierarchy. This region hierarchy serves as a template from which a non-uniform partitioning of the space can be constructed. The second stage follows a top-down process and creates the final set of  $l$  shedding regions, starting from the highest region in the hierarchy (the whole space). The main idea is to selectively pick and drill down on a region using the hierarchy constructed in the first stage. The region to drill down is determined based on the expected gain in the query-result accuracy, called the *accuracy gain* (see Section 3.2.3), which is computed using the aggregated region statistics.

#### 3.2.1 The Statistics Grid

The statistics grid is an  $\alpha \times \alpha$  evenly spaced grid over the geographical space, where  $\alpha$  is the number of grid cells on each side of the space. We describe the relationship between  $\alpha$  and

$l$  later in this section. For each grid cell  $c_{i,j}$  the statistics grid stores the average number of mobile nodes  $n_{i,j}$  and queries  $m_{i,j}$  for that grid cell. The only data structure maintained over time by the LIRA load shedder is this grid.

The maintenance of the grid can be performed in a number of ways. For instance, if the mobile CQ server uses a grid-based index on mobile node positions [10, 12] the statistics grid can be trivially supported as a part of the grid index. Alternatively, the grid can be explicitly maintained by processing position updates. Note that it takes constant time to process an update for maintaining the grid. Moreover, all of the updates need not be processed, since the statistics can easily be approximated using sampling. In an off-line alternative, the average number of mobile nodes can be pre-computed for different times of the day based on historic data, in which case the maintenance cost is close to zero. In all three alternatives, maintenance of the statistics grid is a lightweight operation. The partitioning generated by the GRIDREDUCE algorithm using an  $\alpha \times \alpha$  grid is called an  $(\alpha, l)$ -partitioning.

### 3.2.2 Stage I: Building the Region Hierarchy

In the first stage (see lines 1- 8 in Algorithm 1), we build a complete quad-tree over the grid. Each tree node corresponds to a different region in the space, where regions get larger as we move closer to the root node which represents the whole space. Each level of the quad-tree is a uniform, non-overlapping partitioning of the entire space. Through a post-order traversal of the tree, we aggregate the statistics associated with the grid cells for each node of the tree, i.e., we compute the number of mobile nodes and number of queries for each tree node's region. The first stage of the algorithm takes  $\mathcal{O}(\alpha^2)$  time and consumes  $\mathcal{O}(\alpha^2)$  space.

### 3.2.3 Stage II: Drilling Down in the Hierarchy

In the second stage of the algorithm (see lines 9- 21 in Algorithm 1) we start with the root node of the tree, i.e., the entire space. At each step, we pick a visited tree node (initially only the root) and replace it with its 4 child nodes in the quad-tree. This process continues until we reach  $l$  visited tree nodes (corresponding to  $l$  shedding regions), assuming  $l \bmod 3 = 1$ . The crux of this stage lies in how we choose a region to further partition during each step. For this purpose we maintain a max-heap of all visited tree nodes based on their accuracy gains, a metric we introduce below, and at each step we pick the node with the highest accuracy gain.

Given a tree node, the accuracy gain is a measure of the expected reduction in the query-result inaccuracy, achieved by partitioning the node's region into 4 sub-regions corresponding to its child nodes. For a tree node  $t$ , the accuracy gain  $V[t]$  is calculated as follows. Let  $E[t]$  be the average result inaccuracy if we only had one shedding region that is  $t$ 's region. Formally, we have  $E[t] \leftarrow \min_{\Delta} (m[t] \cdot \Delta)$ , s.t.  $f(\Delta) \leq z \cdot f(\Delta_{\vdash})$ . Now let  $E_p[t]$  be the average result inaccuracy if we had 4 shedding regions that correspond to the regions of  $t$ 's child nodes  $t_i, i \in [1..4]$ . Formally, we have  $E_p[t] \leftarrow \min_{\{\Delta_i\}} \sum_{i=1}^4 \Delta_i \cdot m[t_i]$  subject to the constraint

Algorithm 1:  $(l, \alpha)$ -partitioning of the space

```

GRIDREDUCE( $\alpha, l, z$ )
1) Construct  $\log_2 \alpha + 1$ -level quadrant tree over the  $\alpha \times \alpha$  grid
2) foreach tree node  $t$  in post-order
3)   if  $t$  is a leaf node, {initialize # objs. and # qrys.}
4)      $c_{i,j}$ : corresponding grid cell of  $t$ 
5)      $n[t] \leftarrow n_{i,j}, m[t] \leftarrow m_{i,j}$ 
6)   else { $t$  is not a leaf, aggregate # objs. and # qrys.}
7)      $t_i$ :  $i^{\text{th}}$  children of  $t, i \in [1..4]$ 
8)      $n[t] \leftarrow \sum_{i=1}^4 n[t_i], m[t] \leftarrow \sum_{i=1}^4 m[t_i]$ 
9)  $H$ : empty max. heap of nodes, based on  $V$  (accuracy gain) values
10)  $L$ : empty list of tree nodes,  $i \leftarrow 0$ 
11)  $t \leftarrow$  root of the tree,  $H.INSERT(t)$ 
12) while  $L.SIZE() + H.SIZE() < l$  { $l$  regions not reached}
13)    $t \leftarrow H.POPMAX()$  {region to partition}
14)   if  $t$  is not a leaf {further partitioning possible}
15)     for  $i = 1$  to 4 {partition the region}
16)        $g \leftarrow \text{CALCERRGAIN}(t_i), H.INSERT((t_i, g))$ 
17)     else { $t$  is a leaf node} {no further partitioning}
18)      $L.INSERT(t)$  {store the region in  $L$ }
19) foreach  $t \in L \cup H$  {process the regions}
20)    $n_i \leftarrow n[t], m_i \leftarrow m[t]$  {set the region statistics}
21)    $A_i \leftarrow$  Area of  $t$ 's quadrant,  $i \leftarrow i + 1$  {set the area}
CALCERRGAIN( $t$ )
1)  $E \leftarrow \min_{\Delta} (m[t] \cdot \Delta)$ , s.t.  $f(\Delta) \leq z \cdot f(\Delta_{\vdash})$ 
2)  $E_p \leftarrow \min_{\{\Delta_i\}} \sum_{i=1}^4 \Delta_i \cdot m[t_i]$ ,
   s.t.  $\sum_{i=1}^4 n[t_i] \cdot f(\Delta_i) \leq z \cdot n[t] \cdot f(\Delta_{\vdash})$ 
3)  $V[t] \leftarrow E - E_p$  {accuracy gain is the difference in error}

```

$\sum_{i=1}^4 n[t_i] \cdot f(\Delta_i) \leq z \cdot n[t] \cdot f(\Delta_{\vdash})$ . Then the difference  $E[t] - E_p[t]$  gives us the accuracy gain  $V[t]$ .

The computation of  $E[t]$  and  $E_p[t]$ , and thus the accuracy gain  $V[t]$ , requires solving the problem of update throttler setting for a fixed  $l$  of shedding regions. Concretely, computation of  $E[t]$  requires to solve for node  $t$  with  $l = 1$  and computation of  $E_p[t]$  requires to solve for the 4 child nodes of  $t$  with  $l = 4$ . As we will show in Section 3.3, this general problem can be solved in loglinear time on  $l$ . As a result, the accuracy gain is computed in constant time for a tree node  $t$ . The second stage of the GRIDREDUCE algorithm takes  $\mathcal{O}(l \cdot \log l)$  time and consumes  $\mathcal{O}(l)$  space, bringing the combined time complexity to  $\mathcal{O}(l \cdot \log l + \alpha^2)$  and space complexity to  $\mathcal{O}(\alpha^2 + l)$ .

### 3.2.4 An Example Partitioning

Fig. 3 depicts an example  $(\alpha, l)$ -partitioning. The mobile node distribution (generated from a road map) is shown on the left, whereas the query distribution is shown on the center, and the final  $(\alpha, l)$ -partitioning is shown on the right. It is important to note that the regions are not being further partitioned when the further partitioning will not benefit the query-result accuracy. Here are the two interesting examples: the shedding regions marked with  $\times$  and  $*$  in Fig. 3, which we denote by  $A_{\times}$  and  $A_{*}$ . We see that  $A_{\times}$  is larger than some of the nearby regions. This is because the number of queries is zero for  $A_{\times}$  and as a result further partitioning is not needed.  $A_{*}$  is also larger than some of the nearby regions, but in contrast to  $A_{\times}$  the number of queries is large for  $A_{*}$ . However, what matters is the heterogeneity of the region in terms of the number of mobile nodes and queries within. In the case of  $A_{*}$  further partitioning of the region results in sub-regions of similar characteristics, implying that partitioning is unnecessary.




Fig. 3: An example  $(\alpha, l)$ -partitioning

### 3.2.5 The Relationship Between $l$ and $\alpha$

To find a pragmatic way of configuring the statistics grid parameter  $\alpha$ , we first observe the relationship between  $l$  and  $\alpha$ . Assume that the partitioning is performed such that all the shedding regions are evenly sized. This will yield a grid partitioning with  $\sqrt{l}$  number of cells on each side, which we refer to as the  $l$ -partitioning. Our aim is to have a statistics grid that is fine grained enough to provide us with an  $(\alpha, l)$ -partitioning whose non-uniformly sized shedding regions are sufficiently flexible in terms of the size of their area compared to the case of  $l$ -partitioning in which all regions are equal-sized. The side length of the minimum possible shedding region in  $(\alpha, l)$ -partitioning is proportional to  $1/\alpha$  (the shedding region is equal to a cell of the statistics grid), whereas the side length of a region in  $l$ -partitioning is proportional to  $1/\sqrt{l}$ . To achieve around  $x^2$  times difference in the areas of minimum possible shedding regions of  $l$ -partitioning and  $(\alpha, l)$ -partitioning, we should determine  $\alpha$  using the formula  $\alpha = 2^{\lceil \log_2(x \cdot \sqrt{l}) \rceil}$ . Having  $x = 10$  provides around 100 times difference in size. In our experimental studies we have found that this setting gives effective results.

### 3.3 GREEDYINCREMENT: Setting the $\Delta_i$ 's

The goal of the GREEDYINCREMENT algorithm is to find the optimal setting of the update throttlers associated with the  $l$  shedding regions produced by the GRIDREDUCE algorithm<sup>2</sup>, so that the inaccuracy in query results is minimized (while respecting the fairness thresholds). We first consider this problem without the fairness threshold constraints. The main idea is to increase the update throttlers in order to match the update budget. The update throttlers that bring a larger reduction in the update expenditure of the system in return for a smaller reduction in the result accuracy are preferred for increment.

#### 3.3.1 The Greedy Steps

As the name suggests, the algorithm is a greedy one. It starts by setting all  $\Delta_i$ 's to  $\Delta_{\downarrow}$ , the current update expenditure  $U$  to  $n \cdot f(\Delta_{\downarrow})$  and the update budget  $U_{\downarrow}$  to  $z \cdot U$ . Note that the initial setting is an infeasible solution since the update expenditure is higher than the update budget, that is  $U > U_{\downarrow}$ . At each greedy step one of the update throttlers is selected based on the *update gain*, a criterion to be defined in the next subsection, and is increased by  $c_{\Delta}$ , called the *increment* (or by a smaller value in the case that we undershoot the update budget). When

<sup>2</sup>Note that GREEDYINCREMENT is also used within GRIDREDUCE during computation of the accuracy gains, albeit in a much smaller scale.

$\Delta_i$  is incremented by  $c_{\Delta}$ , the current update expenditure is decreased by  $n_i \cdot (f(\Delta_i) - f(\Delta_i + c_{\Delta}))$ . This process continues until the current update expenditure decreases to match the update budget (i.e.,  $U = U_{\downarrow}$ ), or all the update throttlers reach their maximum bound (i.e.,  $\Delta = \Delta_{\uparrow}$ ). The former condition implies that the update expenditure is reduced to a value equal to throttle fraction times the maximum update expenditure associated with the case of  $\forall_{i \in [1..l]} \Delta_i = \Delta_{\uparrow}$ . This means that the update constraint is satisfied. On the other hand, the latter condition implies that the update budget can not be met for the given throttle fraction  $z$  and the update throttler range  $[\Delta_{\uparrow}, \Delta_{\downarrow}]$ , leading to the solution  $\forall_{i \in [1..l]} \Delta_i = \Delta_{\downarrow}$ .

#### 3.3.2 Update Gain Calculation

The key point of GREEDYINCREMENT is the selection of the update throttler to use at each greedy step. We pick the update throttler that has the highest update gain. The update gain is defined as the ratio of the decrease in update expenditure to the additional inaccuracy introduced in the query results. We denote the rate of decrease in the update expenditure at a point  $\Delta$  by  $r(\Delta)$ , and define it as the negative of the update reduction function  $f$ 's derivative at point  $\Delta$ . Formally:

$$r(\Delta) = - \left. \frac{d(f(x))}{dx} \right|_{x=\Delta}$$

Based on this definition, making a  $dx$  increase in  $\Delta_i$  will reduce the update expenditure by  $n_i \cdot r(\Delta_i) \cdot dx$ , and will decrease the query-result inaccuracy by  $m_i \cdot dx$ . Thus, the update gain for the update throttler  $\Delta_i$ , denoted by  $S_i$ , is:

$$S_i(\Delta) = (n_i/m_i) \cdot r(\Delta)$$

In each step of the GREEDYINCREMENT algorithm, an update throttler  $\Delta_j$  is selected such that we have  $j = \text{argmax}_{i \in [1..l]} S_i(\Delta_i)$ . If the update gain for  $\Delta_j$  is larger than the update gain for  $\Delta_k$ , then increasing  $\Delta_j$  provides better update reduction compared to  $\Delta_k$  for the same amount of increase in query-result inaccuracy.

#### 3.3.3 Optimality and Setting of the Increment $c_{\Delta}$

To provide an optimality guarantee and to guide the setting of  $c_{\Delta}$ , we approximate the update reduction function  $f$  by a non-increasing, piece-wise linear function of  $\kappa$  segments, each of size  $(\Delta_{\downarrow} - \Delta_{\uparrow})/\kappa$ . This enables us to prove the following:

**Theorem 3.1.** *For  $c_{\Delta} = (\Delta_{\downarrow} - \Delta_{\uparrow})/\kappa$ , the GREEDYINCREMENT algorithm is optimal for the non-increasing piece-wise linear approximation of the update reduction function  $f$  with  $\kappa$  segments of size  $c_{\Delta}$  each. **Proof:** See technical report [4].*

The time complexity of the GREEDYINCREMENT algorithm is given by  $\mathcal{O}(\kappa \cdot l \cdot \log l)$  or by  $\mathcal{O}(l \cdot \log l)$  if  $\kappa$  is constant. The space complexity is  $\mathcal{O}(l)$ . See [4] for details.

#### 3.3.4 Supporting the Fairness Threshold

In order to support the fairness constraints dictated by the fairness threshold  $\Delta_{\leftrightarrow}$ , we make the following changes to the base algorithm. At each greedy step, the update throttler with

the highest update gain, say  $\Delta_i$ , is incremented by *at most*  $c_\Delta$ , making sure that it does not go beyond a value that will violate the fairness constraint. Concretely, if the minimum update throttler we have is  $\Delta_{\triangleright} = \min_{j \in [1..l]} \Delta_j$ , then  $\Delta_i$  is not increased beyond  $\Delta_{\triangleright} + \Delta_{\Leftrightarrow}$ . When an update throttler  $\Delta_i$  reaches the limit, that is we have  $\Delta_i = \Delta_{\triangleright} + \Delta_{\Leftrightarrow}$ , then it is moved to a *blocked list* and is not considered for the following steps of the algorithm until it is removed from the blocked list. Whenever the minimum update throttler  $\Delta_{\triangleright}$  is changed, the set of update throttlers in the blocked list that are no more on the limit are removed and are included in the following steps of the algorithm. The pseudo code of GREEDYINCREMENT can be found in our technical report [4].

### 3.4 THROTTLOOP: Setting the Throttle Fraction

The throttle fraction  $z$  can be adaptively adjusted by the LIRA load shedder, when it is not set as a fixed system-level parameter to retain only a pre-defined fraction of position updates. The adjustment of the throttle fraction is performed by the THROTTLOOP, which observes the position update queue and periodically decides the fraction of position updates that should be retained (throttle fraction  $z$ ). The aim is to reduce the system load so that the rate at which the position updates are received ( $\lambda$ ) and the rate at which these updates are processed ( $\mu$ ) are balanced to prevent dropping updates from the input queue. The *utilization* of the system, denoted by  $\rho$ , is given by  $\lambda/\mu$ . Let us denote the maximum size of the input queue by  $B$ . Assuming an  $M/M/1$  queuing model, we should have the following relationship between  $\rho$  and  $B$  to make sure that the average queue length is no more than the maximum queue size [15]:  $\rho = 1 - 1/B$ . If the utilization is larger than  $1 - B^{-1}$ , it represents an overload situation and thus the throttle fraction  $z$  should be decreased. On the other hand, if the utilization is smaller than  $1 - B^{-1}$ , it implies that the system is not fully utilized and the throttle fraction  $z$  should be increased. This understanding leads to the following procedure that describes the operation of THROTTLOOP:

$$\begin{aligned} \text{Initially: } & i \leftarrow 0, z^{(i)} \leftarrow 1 \\ \text{Periodically: } & u \leftarrow \rho/(1 - B^{-1}), i \leftarrow i + 1 \\ & z^{(i)} \leftarrow \min(1, z^{(i-1)}/u) \end{aligned}$$

## 4 Experimental Evaluation

In this section we present experimental results on the effectiveness of the LIRA load shedder in cutting the cost of receiving and processing position updates in mobile CQ systems, while minimally affecting the accuracy of the query results. Before describing the experimental setup, we first define a set of evaluation metrics to assess the effectiveness of LIRA.

### 4.1 Evaluation Metrics

We define two sets of evaluation metrics. The first set of evaluation metrics is used to measure the accuracy of the query results under load shedding and the second set of metrics deals with the cost of performing load shedding.

#### 4.1.1 Query-result Accuracy

*Mean Containment Error*, denoted by  $E_{rr}^C$ , defines the average containment error in query results. Containment error for a query result is defined as the ratio of the number of missing and extra items in the result to the correct result set size. Let  $Q$  denote the set of queries,  $R(q)$  denote the result set for a query  $q \in Q$  under load shedding, and  $R^*(q)$  denote the correct result set under  $\forall_{i \in [1..l]} \Delta_i = \Delta_{\vdash}$ . Then:

$$E_{rr}^C = \sum_{q \in Q} \frac{|R^*(q) \setminus R(q)| + |R(q) \setminus R^*(q)|}{|Q| \cdot |R^*(q)|}$$

*Mean Position Error*, denoted by  $E_{rr}^P$ , defines the average position error in query results. Position error for a query result is defined as the average error in the positions of mobile nodes in the query result compared to the correct positions. Let  $p(o)$  denote the position of a mobile node  $o$  in a query result  $q$  under load shedding and  $p^*(o)$  denote the correct position of  $o$  under  $\forall_{i \in [1..l]} \Delta_i = \Delta_{\vdash}$ . We have:

$$E_{rr}^P = \sum_{q \in Q} \sum_{o \in q} \frac{|p(o) - p^*(o)|}{|Q| \cdot |R(q)|}$$

*Standard Deviation of Containment Error*,  $D_{ev}^C$ , and *Coefficient of Variance of Containment Error*,  $C_{ov}^C$ , are fairness metrics that measure the variation among the query results in terms of containment error. We have  $C_{ov}^C = D_{ev}^C/E_{rr}^C$ .

#### 4.1.2 Cost of Load Shedding

To evaluate the cost incurred by load shedding, we measure *i*) the time it takes to execute the adaptation step that involves running the THROTTLOOP, GRIDREDUCE, and GREEDYINCREMENT algorithms and *ii*) the number of shedding regions that should be known by a mobile node on average. The former metric measures the cost of load shedding from the perspective of the server, whereas the latter measures it from the perspective of the mobile node as well as the wireless network.

### 4.2 Experimental Setup

The experiments were performed using an hour long car (mobile node) position trace<sup>3</sup> generated from real-world road networks available from the National Mapping Division of the United States Geological Survey (USGS) [19] and traffic volume data taken from [7]. We used a map from the Chamblee region of the state of Georgia in the USA (which covers a rich mixture of expressways, arterial roads, and collector roads) to generate the trace used in this paper. The map covers a region of  $\approx 200\text{km}^2$ . The trace is generated by simulating the cars going on roads in accordance with the traffic volume data.

The queries used in the experiments are range CQs. The side length for the range queries are randomly selected from the interval  $[w/2, w]$  where  $w$  is called the *side length parameter*. We use three different distributions for the locations of the queries, namely *Proportional*, *Inverse*, and *Random*. When the query distribution is Proportional, the locations of

<sup>3</sup>The trace generator is available at <http://www.prism.gatech.edu/~gtg470c/research/research.html#kanom>

Parameter	Description	Default Value
$l$	number of shedding regions	250
$\alpha$	statistics grid side cell count	128
$z$	throttle fraction	0.5
$\Delta_{\vdash}$	minimum inaccuracy threshold	5 meters
$\Delta_{\dashv}$	maximum inaccuracy threshold	100 meters
$c_{\Delta}$	increment	1 meter
$\Delta_{\Leftrightarrow}$	fairness threshold	50 meters
$m/n$	# of queries to # of nodes ratio	0.01
$w$	query side length	1000 meters

Table 2: Experimental parameters

the queries follow the mobile node distribution. Similarly, they follow the inverse of the mobile node distribution when the query distribution is Inverse, and are randomly distributed when the query distribution is Random. Due to limited space, in this paper we mainly present our results on the Proportional query distribution. The results for the Inverse and Random distributions are very similar. More details can be found in [4].

In the experiments presented in this paper we compare our LIRA load shedder with the following alternatives:

- **Random Drop:** The excessive position updates are not admitted to the input FIFO queue and are dropped.
- **Uniform  $\Delta$ :** A uniform inaccuracy threshold  $\Delta$  is used to retain only throttle fraction times the original number of location updates. The THROTLLOOP algorithm is still used, but the approach is not region-aware and thus space partitioning and update throttle settings are not performed.

- **Lira-Grid:** A downgraded version of the LIRA load shedder, lacking the GRIDREDUCE algorithm which determines the shedding regions based on  $(l, \alpha)$ -partitioning. Instead, it uses equally-sized shedding regions based on an  $l$ -partitioning, yet still employs GREEDYINCREMENT.

Table 2 presents the set of experimental parameters used and the default values they take when not stated otherwise. As we show in this section, the default setting  $l = 250$  of the number of shedding regions provides sufficient granularity in partitioning (for a region of size  $\approx 200\text{km}^2$ ) to improve the query-result accuracy significantly, while putting very little load on the mobile nodes and the wireless network.

All experiments are performed on an IBM PC with 512MB main memory and 2.4Ghz Intel Pentium4 processor, using Java with Sun JDK 1.5.

### 4.3 Experimental Results

We present the set of experimental results in two groups. The first group of results are on the query-result accuracy and highlight the superiority of LIRA compared to competing approaches for shedding position update load. The second group of results are on the additional cost brought by the LIRA load shedder, and show that the overhead is minimal.

#### 4.3.1 Query-result Accuracy

We study the impact of several system and workload parameters on the query-result accuracy and the relative advantage of LIRA over competing approaches.

**Impact of the Throttle Fraction:** The graphs in Figs. 4 and 5 plot the mean position error  $E_{rr}^P$  and mean

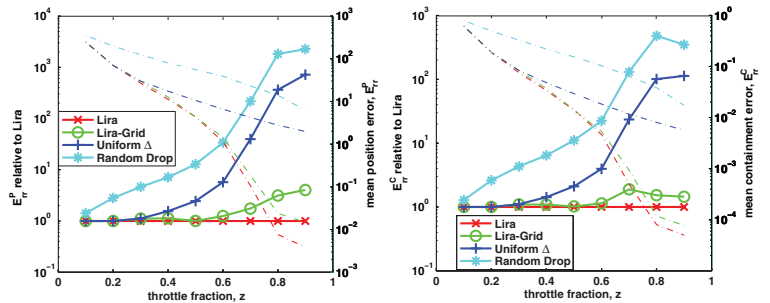


Fig. 4: Position Error vs. throttle fraction

Fig. 5: Containment Error vs. throttle fraction

containment error  $E_{rr}^C$  as a function of the throttle fraction  $z$ , for the Proportional query distribution. The left  $y$ -axis is used to show the relative values (solid lines) with respect to the error of LIRA and the right  $y$ -axis is used to show the absolute errors (dashed lines). Both  $y$ -axes are in logarithmic scale. We make three observations from the figure.

First, the LIRA load shedder outperforms all other approaches throughout the entire throttle fraction range. Random Drop performs the worst, followed by Uniform  $\Delta$  and Lira-Grid. At  $z = 0.75$ , Random Drop has 300 times the mean position error of LIRA, Uniform  $\Delta$  has 40 times that of LIRA, and Lira-Grid has 2 times that of LIRA. At  $z = 0.5$ , Random Drop, Uniform  $\Delta$ , and Lira-Grid has 10, 2, and 1.08 times the  $E_{rr}^P$  of LIRA. The results for the mean containment error  $E_{rr}^C$  are similar. Second, we observe that as the throttle fraction  $z$  gets smaller, the relative errors approach to 1, while at the same time the absolute errors increase and finally merge. The increasing errors are the result of decreasing update budget, whereas the relative errors decrease to 1 due to the maximum inaccuracy bound  $\Delta_{\dashv}$ . When the update budget gets smaller than the minimum update expenditure of the system achieved at  $\forall_{i \in [1..l]} \Delta_i = \Delta_{\dashv}$ , all of the three approaches that use inaccuracy thresholds converge at this same solution. For this experimental setting, this convergence occurs around  $z = 0.25$ . Last, we observe very high (in the order of  $10^3$ 's) relative errors for Random Drop and Uniform  $\Delta$  as  $z$  gets closer to 1. This seems surprising at first, as for the case of  $z = 1$  (not shown in the figures) all approaches have zero error. However, a slight decrease in the throttle fraction, that is when we have  $z = 1 - \epsilon$ , introduces some error in the query results for the case of Random Drop and Uniform  $\Delta$ , whereas it introduces almost no error in the case of LIRA. This is because LIRA cuts the required fraction of position updates from the regions that do not contain any queries. Close to none error of LIRA near  $z = 1$  boosts the relative error results for Random Drop and Uniform  $\Delta$ .

**Impact of the Number of Shedding Regions:** The graphs in Fig. 6 plot the relative mean containment error  $E_{rr}^C$  of Lira-Grid with respect to LIRA as a function of the number of shedding regions  $l$ , for different query distributions. The throttle fraction is set as  $z = 0.5$ . We observe that Lira-Grid has up to 35% higher containment error in query results compared to LIRA. The improvement provided by LIRA is more pronounced when Inverse query distribution is used and is small-



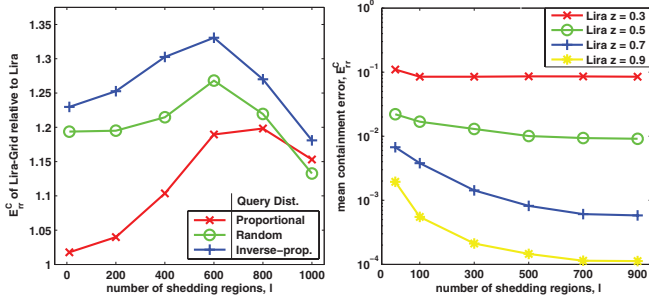


Fig. 6:  $E_{rr}^C$  of Lira-Grid w.r.t. to LIRA vs. # of shedding regions




Fig. 7: Containment Error of LIRA vs. # of shedding regions

est for the case of Proportional query distribution. As  $l$  increases, the flexibility provided by having a larger number of shedding regions improves the error incurred by LIRA at a better rate than Lira-Grid, since LIRA utilizes an intelligent space partitioning algorithm. However, when  $l$  gets too large the grid partitioning of Lira-Grid achieves enough granularity to catch Lira in terms of the query-result inaccuracy, as observed from the figure. This is because after a certain level of granularity is reached, more fine-grained partitioning is of no use, since the accuracy gain is close to zero for all of the shedding regions. The graphs in Fig. 7 attest to this latter intuition. They plot the mean containment error  $E_{rr}^C$  of LIRA as a function of the number of shedding regions, for different throttle fractions. We see that the error reduction rate decreases with increasing  $l$  and the errors stabilize. The reduction in error is more pronounced for larger  $z$  values. Note that the default setting of  $l = 250$  for the number of shedding regions is rather conservative based on Fig. 7, yet it still performs significantly better than the competing approaches as illustrated by Fig. 5. This conservative setting of  $l$  also results in a lightweight load shedding solution, as we illustrate later in this section.

**Impact of the Fairness Threshold:** The graphs in Fig. 8 plot the standard deviation of containment error  $D_{ev}^C$  (on the left  $y$ -axis corresponding to solid lines) and coefficient of variance of containment error  $C_{ov}^C$  (on the right  $y$ -axis corresponding to dashed lines) for LIRA and Uniform  $\Delta$  as a function of the fairness threshold  $\Delta_{\leftrightarrow}$ . Note that Uniform  $\Delta$  does not use a fairness threshold, thus the evaluation metrics stay constant. The surprising observation from the figure is that, with increasing fairness threshold the standard deviation in containment error decreases for LIRA and at all times stays smaller than the  $D_{ev}^C$  of Uniform  $\Delta$ . Even though larger  $\Delta_{\leftrightarrow}$  values imply less fairness, the resulting relaxed constraints in setting the update throttlers enable smaller containment errors and thus the standard deviation also gets smaller. If we look at the coefficient of variance of containment error, which is a better measure of fairness, we see that increasing  $\Delta_{\leftrightarrow}$  increases  $C_{ov}^C$  in LIRA and Uniform  $\Delta$  is more fair compared to LIRA. To put this into simple terms, we can say that on average the difference in errors of two query results will be smaller for LIRA compared to Uniform  $\Delta$ , yet when judged based on the relative average query errors of LIRA and Uniform  $\Delta$  respectively, the error in query results is more fair among different queries in the case of Uniform  $\Delta$ .

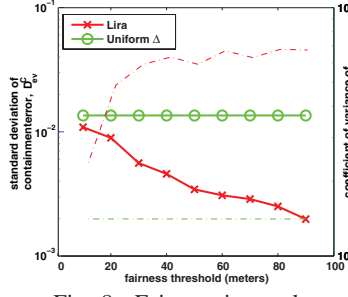


Fig. 8: Fairness in result accuracy for LIRA and Uniform  $\Delta$

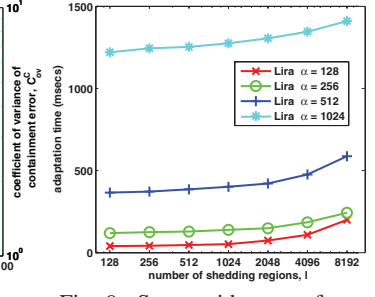


Fig. 9: Server side cost of configuring LIRA

### 4.3.2 Cost of Load Shedding

The cost of load shedding consists of *i*) configuring the parameters of LIRA on the server side, which includes setting the throttle fraction, shedding regions, and update throttlers, *ii*) broadcasting the subset of shedding regions and update throttlers that correspond to the coverage area of each base station, and *iii*) installing the new set of shedding regions and update throttlers on the mobile node side.

**Server Side Cost:** The graphs in Fig. 9 plot the time it takes to execute the THROTLLOOP, GRIDREDUCE, and GREEDYINCREMENT algorithms as a function of the number of shedding regions  $l$ , for different numbers of cells ( $\alpha^2$ ) for the statistics grid. For the default parameters of  $l = 250$  and  $\alpha = 128$ , the configuration of LIRA takes around 40 msecs. This will enable frequent adaptation, even though for most applications that involve monitoring cars or pedestrians it is unlikely that the update load will fluctuate with a period less than tens of minutes. Even for an adaptation period of 10 minutes, the configuration of LIRA will take only  $6.6 \cdot 10^{-5}$  fraction of the adaptation period. Note that these values are for a region of size  $200\text{km}^2$ . If we have a 16 times larger region of size  $3200\text{km}^2$  ( $\approx 10$  times the size of Atlanta, the capital city of the state of Georgia, USA), then we should have  $l = 16 \cdot 250 = 4000$ , and from  $\alpha = 2^{\lceil \log_2(10 \cdot \sqrt{l}) \rceil}$  we should have  $\alpha = 512$ . For this setting the configuration of LIRA takes 500 msecs. This corresponds to  $8 \cdot 10^{-4}$  fraction of a 10 minute adaptation period. These numbers show that LIRA is lightweight and introduces little overhead on the server side.

**Messaging Cost:** Table 3 shows the average number of shedding regions that should be known to a base station as a function of the base station coverage area radius. However, in reality base stations have smaller coverage regions at places where the number of users is large (urban areas) and larger coverage regions at places where the number of users is small (suburban areas) [14]. This nature of base stations match perfectly with LIRA's space partitioning scheme, since the number of partitions are usually larger for dense areas and the small base station coverage areas help decreasing the average number of shedding regions known to a mobile node. Following this logic, we have used a node density dependent base station placement scheme and found that on the average each node and thus each base station should know around 41 shedding regions. Assuming a shedding region (which is square

base station radius (km)	1.0	2.0	3.0	4.0	5.0
# of $\Delta_i$ 's per node	3.1	12.5	28.2	50.2	78.5

41  $\Delta_i$ 's on average, takes  $41 \cdot (3 + 1) \cdot 4$  bytes = 656 bytes.

Table 3: Number of shedding regions per base station

in shape) is represented by 3 floats and an update throttler is represented by a single 4 byte float, the size of the broadcast data sent by a base station to all nodes in its coverage area to install the shedding regions and update throttlers is around  $41 \cdot (3 + 1) \cdot 4$  bytes = 656 bytes on average. To assess the messaging cost of LIRA, compare this number to 1472 bytes, which is the maximum payload available to an UDP packet over Ethernet with a typical MTU of 1500 bytes. When LIRA reconfigures the load shedding parameters, the new information is installed on all mobile nodes by using an average of one wireless broadcast packet per base station.

**Mobile Node Side Cost:** Since the total number of shedding regions known to a mobile node at any time is only around 41, LIRA does not put a major burden on mobile nodes in terms of memory consumption or processing load. By employing a tiny  $5 \times 5$  grid index on the mobile node side, the shedding region that contains the current position of the mobile node can be found quickly. As a result, LIRA will work on computationally weak mobile nodes without any problem.

## 5 Related Work

To the best of our knowledge, this is the first work on position update load shedding in mobile CQ systems. Several works have appeared in the literature on handling the position updates efficiently in mobile CQ systems [18, 9, 11, 21] or using motion modeling to reduce the number of position updates received [20, 2]. The first set of works do not directly address the update load shedding problem, but instead aim at decreasing the IO and CPU cost of integrating the position updates into spatial index structures. This does not involve suppressing or dropping the position updates from the mobile nodes, which is inevitable when the current resources of the system are not sufficient to handle the update load. Our work is complementary in nature to this line of previous work. The second set of previous work use motion modeling to cut the update load, and ensure that the resulting position updates do not have inaccuracy beyond a pre-specified threshold. A key difference is that, our work is driven by the update budget enforced by the load on the system. We adjust the inaccuracy thresholds to reduce the update expenditure of the system to meet the update budget. In other words, our work utilizes the previous work on motion modeling at the mobile node side for actuating the position update suppressing. However, the core of our solution is to find a partitioning and a set of inaccuracy thresholds to associate with each partition, so that the position updates received from the mobile nodes can answer the queries installed in the system accurately.

There have also been a number of distributed solutions to evaluate CQs in mobile systems [1, 8, 3]. In these systems, the position updates are only received if they affect a query result. Even though these systems do not provide any load shedding

capability, their update load is expected to be significantly lower compared to solutions that track all mobile nodes. However, these solutions cannot support historic queries, since the location updates are not received from all objects. The ad-hoc snapshot queries are also expensive to evaluate. Interestingly, LIRA can be configured to mimic the behavior of these systems by setting the maximum inaccuracy bound to a large value. Moreover, our system has the additional advantage of not being tied to any specific query processing technique and has very little overhead.

## 6 Conclusion

We presented LIRA, a position update load shedder for mobile CQ systems. The primary feature of LIRA is its region-awareness, which enables it to partition the space into a set of shedding regions and apply differing amounts of update throttling for different shedding regions. We developed a heuristic algorithm to discover a partitioning of the space that leads to reduced error in query results, and an optimal algorithm that sets the update throttlers associated with each shedding region to minimize the query-result inaccuracy. We showed that the LIRA load shedder is significantly superior to random update dropping and uniform inaccuracy threshold schemes. Moreover, LIRA is lightweight by design and can be used in conjunction with many of the existing update indexing and mobile CQ processing techniques.

## References

- [1] Y. Cai and K. A. Hua. An adaptive query management technique for efficient real-time monitoring of spatial regions in mobile database systems. In *IEEE International Performance Computing and Communications Conference*, 2002.
- [2] A. Civilis, C. S. Jensen, and S. Pakalnis. Techniques for efficient road-network-based tracking of moving objects. *IEEE TKDE*, 17(5), 2005.
- [3] B. Gedik and L. Liu. Distributed processing of continuously moving queries on moving objects in a mobile system. In *EDBT*, 2004.
- [4] B. Gedik, L. Liu, K.-L. Wu, and P. S. Yu. LIRA: Lightweight, region-aware load shedding in mobile CQ systems. Technical Report CERCs-06-07, GaTech, 2006.
- [5] B. Gedik, K.-L. Wu, P. S. Yu, and L. Liu. Processing moving queries over moving objects using motion adaptive indexes. *IEEE TKDE*, 18(5), 2006.
- [6] Google RideFinder home page. <http://labs.google.com/ridefinder>, Febuary 2006.
- [7] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *ACM MobiSys*, 2003.
- [8] H. Hu, J. Xu, and D. Lee. A generic framework for monitoring continuous spatial queries over moving objects. In *ACM SIGMOD*, 2005.
- [9] C. S. Jensen, D. Lin, and B. C. Ooi. Query and update efficient B+-tree based indexing of moving objects. In *VLDB*, 2004.
- [10] D. V. Kalashnikov, S. Prabhakar, S. Hambrusch, and W. Aref. Efficient evaluation of continuous range queries on moving objects. In *DEXA*, 2002.
- [11] M.-L. Lee, W. Hsu, C. S. Jensen, B. Cui, and K. L. Teo. Supporting frequent updates in r-trees: A bottom-up approach. In *VLDB*, 2003.
- [12] M. F. Mokbel, X. Xiong, and W. G. Aref. SINA: Scalable incremental processing of continuous queries in spatio-temporal databases. In *ACM SIGMOD*, 2004.
- [13] S. Prabhakar, Y. Xia, D. V. Kalashnikov, W. G. Aref, and S. E. Hambrusch. Query indexing and velocity constrained indexing: Scalable techniques for continuous queries on moving objects. *IEEE TC*, 51(10), 2002.
- [14] Qualcomm. Wireless access solutions using 1xEV-DO. [www.qualcomm.com/technology/1xev-do/webpapers/wp\\_wirelessaccess.pdf](http://www.qualcomm.com/technology/1xev-do/webpapers/wp_wirelessaccess.pdf), 2005.
- [15] S. Ross. *A First Course in Probability*. Prentice Hall, 2005.
- [16] S. Saltinis, C. S. Jensen, S. T. Leutenegger, and M. A. Lopez. Indexing the positions of continuously moving objects. In *ACM SIGMOD*, 2000.
- [17] C. Science and T. Board. *IT Roadmap to a Geospatial Future*. The National Academics Press, November 2003.
- [18] Y. Tao, D. Papadias, and J. Sun. The TPR\*-Tree: An optimized spatio-temporal access method for predictive queries. In *VLDB*, 2003.
- [19] U.S. Department of the Interior. U.S. geological survey web page. <http://www.usgs.gov/>, November 2003.
- [20] O. Wolfson, P. Sistla, S. Chamberlain, and Y. Yesha. Updating and querying databases that track mobile units. *Distributed and Parallel Databases*, 7(3), 1999.
- [21] X. Xiong and W. G. Aref. R-trees with update memos. In *IEEE ICDE*, 2006.