

MobiQual: QoS-aware Load Shedding in Mobile CQ Systems

Buğra Gedik ^{#1}, Kun-Lung Wu ^{#2}, Philip S. Yu ^{#3}, Ling Liu ^{*4}

[#]Thomas J. Watson Research Center, IBM Research, 19 Skyline Dr, Hawthorne, NY 10532

¹bgedik@us.ibm.com, ²klwu@us.ibm.com, ³psyu@us.ibm.com

^{*}College of Computing, Georgia Institute of Technology, 801 Atlantic Drive, Atlanta, GA 30332

⁴lingliu@cc.gatech.edu

Abstract—Freshness and accuracy are two key measures of quality of service (QoS) in location-based, mobile continual queries (CQs). However, it is often difficult to provide both fresh and accurate CQ results due to (a) limited resources in computing and communication and (b) fast-changing load conditions caused by continuous mobile node movement. Thus a key challenge for a mobile CQ system is: How do we achieve the highest possible quality of the query results, in both freshness and accuracy, with currently available resources under changing load conditions? In this paper, we formulate this problem as a load shedding one, and develop MobiQual – a QoS-aware framework for performing both update load shedding and query load shedding. The design of MobiQual highlights three important features. (1) *Differentiated load shedding*: Different amounts of query and update load shedding are applied to different groups of queries and mobile nodes, respectively. (2) *Per-query QoS specifications*: The overall freshness and accuracy of the query results are maximized with individualized QoS specifications. (3) *Low-cost adaptation*: MobiQual dynamically adapts, with a minimal overhead, to changing load conditions and available resources. We show that, through a careful combination of update and query load shedding, the MobiQual approach leads to much higher freshness and accuracy in the query results in all cases, compared to existing approaches.

I. INTRODUCTION

With the ever increasing accessibility of wireless communications and the proliferation of mobile devices, we are experiencing a world where we can stay connected while on-the-go. Combined with the availability of low-cost positioning devices (such as GPS sensors), this has created a new class of applications and business opportunities in the area of mobile location-based services (LBSs). Examples include location-aware information delivery and resource management, such as transportation services (NextBus bus locator [1], Google ride finder [2]), fleet management, mobile games, and battlefield coordination.

A key challenge for LBSs is a scalable location monitoring system capable of handling large number of mobile nodes and processing complex queries over their positions. Although several mobile continual query (CQ) systems have been proposed to handle long-running location monitoring tasks in a scalable manner [3], [4], [5], [6], [7], the focus of these works is primarily on efficient indexing and query processing techniques, not on accuracy or freshness of the query results.

Accuracy (inaccuracy) is defined based on the amount of mobile node position errors found in the query results *at the time of* query re-evaluation. This accuracy measure is strongly tied to the frequency of position updates received from the mobile nodes. Although one can also use a higher level concept to measure accuracy, such as the amount of containment errors found in the query results¹, including both false positives (inclusion errors) and false negatives (exclusion errors), we argue that using position update errors for accuracy measure will provide higher level of precision. This is primarily because by utilizing the amount of node position errors as the accuracy measure, one can easily bound the inaccuracy by a threshold-based position reporting scheme [8], [9].

Freshness (staleness), on the other hand, refers to the age of the query results *since* the last query re-evaluation. It is dependent on the frequency of query re-evaluations performed at the server. As mobile nodes continue to move, there are further deviations in mobile node positions after the last query re-evaluation. However, such deviations are not attributed to inaccuracy. Hence, freshness can be seen as a metric capturing the post-query-re-evaluation deviations in mobile node positions. It is important to note that higher freshness does not necessarily imply higher accuracy and vice versa.

To obtain fresher query results, the CQ server must re-evaluate the continual queries more frequently, requiring more computing resources. Similarly, to attain more accurate query results, the CQ server must receive and process position updates from the mobile nodes in a higher rate, demanding communication as well as computing resources. However, it is almost impossible for a mobile CQ system to achieve 100% fresh and accurate results due to continuously changing positions of mobile nodes. A key challenge therefore is: How do we achieve the highest possible quality of the query results in both freshness and accuracy, in the presence of changing available resources and changing workloads of location updates and location queries?

In this paper, we propose MobiQual – a resource-adaptive and QoS-aware load shedding framework for mobile CQ systems. MobiQual is capable of providing high-quality query

¹A bound on the amount of containment errors can be approximated by a bound on the position errors, if the distribution of the mobile nodes around the query boundaries is at hand or can be approximated.

results by dynamically determining the appropriate amount of update load shedding and query load shedding to be performed according to the application-level QoS specifications of the queries. An obvious advantage of combining query load shedding and update load shedding within the same framework is to empower MobiQual with *differentiated load shedding* capability, that is configuring *query re-evaluation periods* and *update inaccuracy thresholds* for achieving high overall QoS with respect to both freshness and accuracy.

We have conducted detailed experimental studies on the effectiveness of MobiQual. Our results show that (1) a careful combination of location update load shedding and location query load shedding can significantly outperform the approaches that are based on query-only or update-only load shedding; and (2) MobiQual provides higher quality guarantees compared to the approaches that lack the supports of QoS awareness and differentiated load shedding.

Previous work on mobile CQ systems have focused on four major themes with respect to scalability and performance. These are: *i*) indexing schemes to process position updates more efficiently [10], [5], [11], [12]; *ii*) query processing techniques to evaluate continual queries more efficiently [6], [3], [13], [14]; *iii*) motion modeling techniques to reduce the number of position updates received from the mobile nodes, while keeping the position accuracy high [8], [9]; and *iv*) distributed mobile CQ systems that achieve scalability by performing query-aware update filtering on the mobile node side to receive updates that only relate to the current set of queries installed in the system [15], [4], [16]. Most of these works, with the exception of the works listed under item *iv*, are mostly orthogonal to our work and can be incorporated into MobiQual easily. For instance, MobiQual can use a TPR-tree [10] as its underlying index structure on the server side, can make use of advanced motion modeling techniques [8] on the mobile node side, and can employ incremental query processing techniques [3] for query re-evaluation. Unlike the set of works listed under item *iv*, MobiQual receives updates from all the nodes, so that ad-hoc and historical queries can also be supported. However, MobiQual prefers to shed position updates from regions that have minimal impact on the currently installed queries, thus achieving best of both worlds.

In our previous work [17], we introduced the Lira system for performing update load shedding in mobile CQ systems. However, the Lira system not only lacks query load shedding capabilities, but also ignores any QoS-aware mechanisms as a part of the load shedding framework. Thus the Lira results are limited. To the best of our knowledge, none of the existing work has exploited the potential of combining update load shedding and query load shedding to maximize the application-level freshness and accuracy of mobile queries.

II. DESIGN OVERVIEW

A. Load Shedding in Mobile CQ Systems

In a mobile CQ system, the CQ server receives position updates from the mobile nodes through a set of base stations




Fig. 1: Mobile CQ system and load shedding

(see Figure 1) and periodically evaluates the installed continual queries (such as continual range or nearest neighbor queries) over the last known positions of the mobile nodes. Since the mobile node positions change continuously, motion modeling [8], [9] is often used to reduce the number of updates sent by the mobile nodes. The server can predict the locations of the mobile nodes through the use of motion models, albeit with increasing errors. Mobile nodes generally use a threshold to reduce the amount of updates to be sent to the server and to limit the inaccuracy of the query results at the server side below the threshold. Smaller thresholds result in smaller errors and higher accuracy, at the expense of a higher load on the CQ server. This is because a larger number of position updates must be processed by the server, for instance, to maintain an index [10], [5]. When the position update rates are high, the amount of position updates is huge and the server may randomly drop some of the updates if resources are limited. This can cause unbounded inaccuracy in the query results. In MobiQual, we use accuracy-conscious update load shedding to regulate the load incurred on the CQ server due to position update processing by dynamically configuring the inaccuracy thresholds at the mobile nodes.

Another major load for the CQ server is to keep the query results up to date by periodically executing the CQs over the mobile node positions. More frequent query re-evaluations translate into increased freshness in the query results, also at the expense of a higher server load. Given limited server resources, when the rate of query re-evaluations is high, the amount of queries to be re-evaluated is vast and the server may randomly drop some of the re-evaluations, causing stale query results (low freshness). In MobiQual, we utilize freshness-conscious query load shedding to control the load incurred on the CQ server due to query re-evaluations by configuring the query re-evaluation periods.

In general, the total load due to evaluating queries and processing position updates dominates the performance and scalability of the CQ server. Furthermore, the time-varying processing demands of a mobile CQ system entails that update and query load shedding should be dynamically balanced and adaptively performed in order to match the current workload with the server's capacity.

B. The MobiQual Approach

The MobiQual system aims at performing dynamic load shedding to maximize the overall quality of the query results, based on per-query QoS specifications and subject to

processing capacity constraints. The QoS specifications are defined based on two factors: accuracy and freshness. In MobiQual, the QoS specifications are used to decide on not only how to spread out the impact of load shedding among different queries, but also how to find a balance between query load shedding and update load shedding. The main idea is to apply differentiated load shedding to adjust the accuracy and freshness of queries. Namely, load shedding on position updates and query re-evaluations is done in such a way that the freshness and accuracy of queries are non-uniformly impacted.

From the perspective of query load shedding, we make two observations to show that nonuniform freshness in the query results can increase the overall QoS of the mobile CQ system: (1) Different queries have different costs in terms of the amount of load they incur. (2) Different queries have different tolerance to staleness in the query results. Thus it is more effective to shed load (by sacrificing certain amount of freshness) on a costly query than an inexpensive one. This is especially beneficial if the costly query happens to be less stringent on freshness, based on its QoS specification. Bearing these observations in mind, in MobiQual we employ *QoS-aware query load shedding*: We use query re-evaluation periods as control knobs to perform query load shedding, where the same amount of increase in query re-evaluation periods for different queries brings differing amounts of load reduction and QoS degradation with respect to freshness.

Similar to query load shedding, we make two observations regarding update load shedding to show that nonuniform result accuracy can increase the overall QoS. First, different geographical regions have different numbers of mobile nodes and queries. Second, different queries have different tolerance to position errors in the query results. This means that shedding more updates from a region with a higher density of mobile nodes and a lower density of queries can bring a higher reduction on the update load and yet have a smaller impact on the overall query result accuracy. This is especially true if the queries within the region have less stringent QoS specifications in terms of accuracy. Thus, in MobiQual we employ *QoS-aware update load shedding*: We use inaccuracy thresholds from motion modeling as control knobs to adjust the amount of update load shedding to be performed, where the same amount of increase in inaccuracy thresholds for different geographical regions brings differing amounts of load reduction and QoS degradation with respect to accuracy.

MobiQual dynamically maintains a *throttle fraction*, which defines the amount of load that should be retained. It performs both update load shedding and query load shedding to control the load of the system according to this throttle fraction, while maximizing the overall quality of the query results. MobiQual not only strikes a balance between freshness and accuracy by employing both query and update load-shedding, but also improves the overall quality of the results by utilizing per-query QoS specifications to capture each query's different tolerance to staleness and inaccuracy.

C. Notation and Fundamentals

The set of continual queries installed in the system is denoted by Q . For each query $q \in Q$, there is an associated QoS specification S_q . The QoS function $S_q(\tau_q, \epsilon_q)$ takes a value in $[0, 1]$, where 1 represents perfect quality in terms of freshness and position error, and 0 represents the worst. τ_q and ϵ_q are used to denote the degree of staleness and inaccuracy in the query results, respectively. τ_q corresponds to the query re-evaluation period for q , whereas ϵ_q corresponds to the average of the inaccuracy thresholds used in motion modeling for the mobile nodes within the query result of q . At any given time, the result of query q can be at most τ_q seconds old and at the time of query evaluation the position of a mobile node in the query result can deviate from its actual position by ϵ_q meters on average. The mobile CQ system supports a minimum staleness value of τ_- and a minimum position error of ϵ_- . For any query q , we have $S_q(\tau_-, \epsilon_-) = 1$. Similarly, we introduce a maximum staleness value, denoted by τ_+ , and a maximum position error, denoted by ϵ_+ . The staleness in the query results cannot exceed the maximum threshold value of τ_+ , at which point the results are assumed to be useless. Also the position error is bounded by ϵ_+ . In summary, we have $\tau_q \in [\tau_-, \tau_+]$ and $\epsilon_q \in [\epsilon_-, \epsilon_+]$. The minimum and maximum staleness and position error thresholds are system parameters.

Since a scalable mobile CQ system should be able to handle tens of thousands of queries and hundreds of thousands of mobile nodes, it is inefficient, even if it is possible, to adjust and dynamically maintain the re-evaluation periods for queries and inaccuracy thresholds for mobile nodes individually. In MobiQual, we divide the set of m queries into k groups, denoted by $C_j, j \in [1..k]$, where the number of queries in C_j is denoted by m_j and $\sum_{j=1}^k m_j = m$. The queries within the same group C_j share the same re-evaluation period P_j , i.e. we have $\forall q_u \in C_j, \tau_u = P_j$. We denote the one-time cost of processing the set of queries in C_j as $f_c(C_j)$, which is simply the sum of one-time processing costs of individual queries. The usage of the cost model in MobiQual does not require absolute values of query costs and can work with relative values for cost-based analysis. A key question for query load shedding is how to divide the queries into k query groups and how to compute the re-evaluation period P_j for each query group C_j ($j \in [1..k]$).

Similarly, given a total number of n mobile nodes, we partition the geographical area of interest into l regions, denoted by $A_i, i \in [1..l]$, where the number of mobile nodes in A_i is denoted by n_i and $\sum_{i=1}^l n_i = n$. The mobile nodes within the same region A_i use the same inaccuracy threshold Δ_i . A query q_u whose result lies completely within region A_i will have $\epsilon_u = \Delta_i$. For queries whose results contain mobile nodes from different regions, ϵ_u is given by a weighted average of Δ_i values of the involved regions.

We denote the fraction of updates received from a region A_i , when using an inaccuracy threshold Δ_i , as $f_r(\Delta_i)$. f_r is relative to the ideal case where all Δ_i 's are equal to the minimum position error ϵ_- . Thus we have $f_r(\epsilon_-) = 1 > f_r(\epsilon_+)$. f_r is

a non-increasing continuous function with a positive second derivative. More detailed characterization of such functions exist for specific motion modeling and prediction schemes [8], [17]. A key challenge for update load shedding is how to partition the geographical area of interest into l regions and how to compute the inaccuracy threshold Δ_i for each region A_i ($i \in [1..l]$).

D. Trade-offs in Setting k and l

In general, the larger the number of query groups (k) we have, the higher the quality of the query results is in terms of freshness, as it enables performing differentiated load shedding with finer granularity. The only restriction in setting the value of k is the computational cost (which forms a major part of the adaptation cost) of finding an effective setting for the re-evaluation periods $P_j, j \in [1..k]$. Similar trade-off is observed in setting the number of regions (l) and thus the number of inaccuracy thresholds, with one exception. Since the changes in inaccuracy thresholds have to be communicated back to the mobile nodes through control messages (broadcasts from base stations), there is a second dimension to this trade-off: The larger the l value is, the higher the control cost of the adaptation step will be. In Section VII, we experimentally evaluate the benefit/cost trade-off in setting k to show that with lightweight adaptation we can achieve high quality query results. The details of setting l can be found in [17].

E. Solution Outline

There are three functional components in the MobiQual system: *reduction*, *aggregation*, and *adaptation*.

– **Reduction** includes the algorithm for grouping the queries into k clusters and the algorithm for partitioning the geographical space of interest into l regions. The query groups are incrementally updated when queries are installed or removed from the system. The space partitioning is re-computed prior to the periodic adaptation.

– **Aggregation** involves computing aggregate-QoS functions for each query group and region. The aggregated QoS functions for each query group represent the freshness aspect of the quality. The aggregated QoS functions for each region represent the accuracy aspect of the quality. We argue that the separation of these two aspects is essential to the development of a fast algorithm for configuring the re-evaluation periods and the inaccuracy thresholds to perform adaptation. QoS-aggregation is repeated only when there is a change in the query grouping or the space partitioning.

– **Adaptation** is performed periodically to determine: (i) the throttle fraction $z \in [0, 1]$, which defines the amount of load that can be retained relative to the load of providing perfect quality (i.e., $\forall_{j \in [1..k]} P_j = \tau_{\perp}$ and $\forall_{i \in [1..l]} \Delta_i = \epsilon_{\perp}$); (ii) the setting of re-evaluation periods $P_j, j \in [1..k]$; and (iii) the setting of inaccuracy thresholds $\Delta_i, i \in [1..l]$. The latter two are performed with the aim of maximizing the overall QoS. The computation of the throttle fraction is performed by monitoring the performance of the system and adjusting z in a feedback loop (see our technical report for details [18]).

In the remaining sections, we first present the aggregation of QoS functions, assuming that the query grouping and space partitioning are performed (Section III). We then present the formulation of the QoS-aware query load shedding problem and present the *quality loss based clustering* (QLBC) algorithm for clustering the queries into k groups (Section IV). Then we formalize the QoS-aware update load shedding problem and provide a brief description of the QoS-aware space partitioning algorithm for dividing the geographical space of interest into l regions (Section V). Finally, we present the formulation of the problem of combining query load shedding with update load shedding, and present the *minimum quality loss per cost step* (MQLS) algorithm for performing the adaptation step (Section VI).

III. AGGREGATING THE QOS FUNCTIONS

The aim of QoS aggregation is to associate an aggregate function $\mathcal{V}_j^*(P_j)$ for each query group C_j , and an aggregate function $\mathcal{U}_i^*(\Delta_i)$ for each region A_i , such that the overall QoS of the system, denoted by Ψ , is maximized. We define:

$$\Psi = \frac{1}{m} \sum_{q \in Q} \mathcal{S}_q(\tau_q, \epsilon_q) \quad (1)$$

$\mathcal{S}_q(\tau_q, \epsilon_q)$ denotes the QoS specification for query q and can be defined as follows:

$$\mathcal{S}_q(\tau_q, \epsilon_q) = \alpha_q \cdot \mathcal{V}_q(\tau_q) + (1 - \alpha_q) \cdot \mathcal{U}_q(\epsilon_q) \quad (2)$$

In other words, $\mathcal{S}_q(\tau_q, \epsilon_q)$ is a linear combination of the freshness QoS function $\mathcal{V}_q(\tau_q)$ and the accuracy one $\mathcal{U}_q(\epsilon_q)$. The parameter $\alpha_q \in [0, 1]$, called *freshness weight*, is used to adjust the relative importance of the two components, freshness and accuracy. $\mathcal{V}_q(\tau_q)$ and $\mathcal{U}_q(\epsilon_q)$ are non-increasing positive functions, where $\mathcal{V}_q(\tau_{\perp}) = 1$ and $\mathcal{U}_q(\epsilon_{\perp}) = 1$.

Since the query groups are non-overlapping, we have:

$$\mathcal{V}_j^*(P_j) = \sum_{q \in C_j} \alpha_q \cdot \mathcal{V}_q(P_j) \quad (3)$$

We approximate the \mathcal{V}_q functions using piece-wise linear functions of κ equal-sized segments along the input domain $[\tau_{\perp}, \tau_{\perp}]$. This enables us to represent the aggregate QoS functions (\mathcal{V}_j^* 's) as piece-wise linear functions of κ segments as well.

Recall that the set of queries that intersect a region A_i can overlap with the set of queries that intersect a different region, since a query q can intersect more than one region. Let $m_q(i)$ denote the fraction of q 's query region that lies within A_i and Q denotes the set of queries in the system. Then, we have:

$$\mathcal{U}_i^*(\Delta_i) = \sum_{q \in Q, \text{ s.t. } m_q(i) > 0} (1 - \alpha_q) \cdot m_q(i) \cdot \mathcal{U}_q(\Delta_i) \quad (4)$$

The equality in Equation 4 holds when (a) \mathcal{U}_q 's are linear functions,² or (b) \mathcal{U}_q 's are piece-wise linear functions and there are no queries crossing the region borders. However, it is still

²For Equation 4 to hold, we should be able to write $\mathcal{U}_q(\epsilon_q) = \mathcal{U}_q(\sum_{i, m_q(i) > 0} m_q(i) \cdot \Delta_i) = \sum_{i, m_q(i) > 0} m_q(i) \cdot \mathcal{U}_q(\Delta_i)$. This can be done if and only if \mathcal{U}_q is a linear function.

a good approximation for the general case of piece-wise linear functions if the crossings are not frequent. Because the size of a region is significantly larger than that of a query, query crossings are indeed infrequent. Like \mathcal{V}_j^* 's, we also represent \mathcal{U}_i^* 's as piece-wise linear functions with κ segments. Based on this analysis, the Equation 1 can be written in the following form:

$$\Psi = \frac{1}{m} \left(\sum_{j=1}^k \mathcal{V}_j^*(P_j) + \sum_{i=1}^l \mathcal{U}_i^*(\Delta_i) \right) \quad (5)$$

Note that, for a given $j \in [1..k]$, \mathcal{V}_j^* is independent of Δ_i 's ($i \in [1..l]$). Similarly, for a given $i \in [1..l]$, \mathcal{U}_i^* is independent of P_j 's ($j \in [1..k]$). This separation allows us to operate at the granularity of query groups for configuring query load shedding and at the granularity of regions for configuring the update load shedding.

It is critical to note that the queries within C_j may intersect a number of different areas, and similarly queries within area A_i may be contained in a number of different query groups. As a result, if \mathcal{U}_i^* 's were not independent of P_j 's, altering the re-evaluation period P_j for queries within C_j may have altered more than m_j different aggregate QoS functions belonging to different regions. A similar argument is valid for altering the inaccuracy threshold Δ_i for A_i when \mathcal{V}_j^* 's are not independent of Δ_i 's. Thus, without a clear separation of re-evaluation periods and inaccuracy thresholds in aggregated QoS functions, we end up creating a significant overhead for system optimization. This will defy reduction by making \mathcal{V}_j^* 's and \mathcal{U}_i^* 's dependent on a large number of parameters, making their computation costly.

One downside of representing a query's QoS specification as a linear combination of a freshness-related QoS function and an accuracy-related QoS function is the loss of certain amount of expressiveness, compared to the case of an unrestricted QoS function of two parameters. Yet, the presented model still manages to capture a wide spectrum of QoS specifications, ranging from staleness insensitive ($\alpha_q = 0$) to inaccuracy insensitive ($\alpha_q = 1$) scenarios.

We first discuss query load shedding and update load shedding separately in the next two sections, and then present our final solution for combining them.

IV. QOS-AWARE QUERY LOAD SHEDDING

We now focus on the QoS-aware query load shedding problem, by only considering the freshness aspect of the quality and the cost of query re-evaluation.

A. Formalization of the Problem

The aim of the query load shedding problem is to maximize the first component of the overall quality from Equation 5, denoted by Ψ_v . Given k query groups, recalling that $\mathcal{V}_j^*(P_j)$ denote the aggregation function for the query group C_j , and P_j denote the setting of the re-evaluation period for C_j , we define Ψ_v as follows:

$$\Psi_v = \sum_{j=1}^k \mathcal{V}_j^*(P_j) \quad (6)$$

Assume that the throttle fraction z is given, which defines the fraction of query load to keep. (The details for computation of z are described in [18]). Under this assumption, the one-time re-evaluation cost of queries within C_j is given by $f_c(C_j)$ and since these queries are re-evaluated every P_j seconds, the overall cost is given by $f_c(C_j)/P_j$. As a result, the load under a given set of re-evaluation periods $\{P_j\}$ is $\sum_{j=1}^k f_c(C_j)/P_j$, which should be less than or equal to the throttle fraction times the load of the ideal case of $\forall_{q \in Q}, \tau_q = \tau_+$, which is given by $z \cdot \sum_{q \in Q} f_c(\{q\})/\tau_+$. In summary, the query load shedding algorithm should respect z as the query re-evaluation budget, while maximizing the freshness in the query results. This can be modeled by the following processing constraint:

$$\sum_{j=1}^k f_c(C_j)/P_j \leq z \cdot \sum_{q \in Q} f_c(\{q\})/\tau_+ \\ \forall_{j \in [1..k]}, \tau_+ \leq P_j \leq \tau_+$$

The second constraint defines the scope of the re-evaluation period P_j ($j \in [1..k]$). The key problem here is to define the set of query groups ($C_j, j \in [1..k]$), so as to maximize Ψ_v . This is performed by the QLBC algorithm described next.

B. Measuring Quality Loss Per Unit Cost

The first question for clustering queries is to find which metric should be used as a distance measure to define similar queries. One intuitive observation is that two queries are similar for the purpose of query load shedding if the amount of reduction in quality per unit decrease in cost is similar for the two queries. We call this measure *quality loss per unit cost* (*qlpc*) metric. Let $G(q, z)$ denote the quality loss per unit cost for a given query q and a given throttle fraction z . We define $G(q, z)$ using the following formula:

$$G(q, z) = \frac{\alpha_q \cdot \left. \frac{d(\mathcal{V}_q(\tau))}{d\tau} \right|_{\tau=\tau_+/z}}{f_c(\{q\}) \cdot \left. \frac{d(1/\tau)}{d\tau} \right|_{\tau=\tau_+/z}} \quad (7)$$

Note that $\mathcal{V}_q(\tau)$ is the freshness-related QoS function associated with q , whereas $f_c(\{q\})/\tau$ is the cost function of q . Setting the re-evaluation period to $\tau = \tau_+/z$ reduces the overall cost of re-evaluating q to z times the cost for the ideal case of $\tau = \tau_+$. Since queries within the same group will share the same re-evaluation period, Equation 7 captures the quality loss per unit cost for a $d\tau$ increase in the re-evaluation period.

Clearly, a query q with a small $G(q, z)$ value is a good choice for shedding the query load, as it brings a small loss in QoS for a large amount of decrease in load. Therefore, if two queries have similar \mathcal{V}_q functions, then the one with the larger evaluation cost $f_c(\{q\})$ will be preferred for load shedding. However, if two queries have similar $f_c(\{q\})$ values, then the query with the smaller (absolute) derivative of its \mathcal{V}_q function will be preferred for load shedding. Note that the derivative of the QoS function \mathcal{V}_q is constant over each linear segment and thus Equation 7 can be simplified as follows, where $\mathcal{V}_q^a(i)$

denotes the slope of the i th ($i \in [1..\kappa]$) linear segment of \mathcal{V}_q :

$$G(q, z) = \frac{\alpha_q \cdot \mathcal{V}_q^a \left(\left\lceil \kappa \cdot \frac{\tau_+ / z - \tau_-}{\tau_+ - \tau_-} \right\rceil \right)}{-f_c(\{q\}) \cdot (z/\tau_+)^2} \quad (8)$$

C. Grouping Queries with QLBC

MobiQual uses the quality loss per unit cost (*qlpc*) metric to define the similarity of queries and the distance function used for clustering queries into desirable query groups in terms of load shedding effectiveness. We call this algorithm the *Quality loss based clustering* algorithm, QLBC for short.

It is obvious that putting queries that have diverse $G(q, z)$ values into the same group is very ineffective, because queries with larger $G(q, z)$ values are not good candidates for query load shedding compared to others. Hence there will be less overall benefit from increasing the common re-evaluation period. The QLBC algorithm finds the similarity between two queries q_1 and q_2 in two steps. First, it models the quality loss per unit cost of each query at different z values using a *qlpc vector*, where each element of the vector corresponds to the $G(q, z)$ value at a different load shedding level z (κ different levels equally spaced between 0 and 1). Second, the QLBC algorithm uses the Euclidean distance between the *qlpc* vectors of queries to define the similarity of queries. This similarity, denoted by $D(q_1, q_2)$, is defined as:

$$D(q_1, q_2) = \sum_{\iota \in ([1..\kappa] - 0.5)/\kappa} (G(q_1, \iota) - G(q_2, \iota))^2 \quad (9)$$

The QLBC algorithm uses k -means clustering [19] to form the final k set of query groups, based on Equation 9.

V. QoS-AWARE UPDATE LOAD SHEDDING

In this section we describe the QoS-aware update load shedding problem, by only considering the accuracy aspect of the quality and the cost of position update processing.

A. Formalization of the Problem

The goal of the update load shedding problem is to maximize the second component of the overall quality from Equation 5, denoted by Ψ_u . Given l regions of the geographical space of interest, recalling that $\mathcal{U}_i^*(\Delta_i)$ denote the aggregation function for region A_i , and Δ_i denote the inaccuracy threshold associated with A_i , we define Ψ_u as follows:

$$\Psi_u = \sum_{i=1}^l \mathcal{U}_i^*(\Delta_i) \quad (10)$$

Assume that the throttle fraction z is given, which defines the fraction of update load to keep. The number of updates and thus the relative cost of update processing for a given region A_i are proportional to $n_i \cdot f_r(\Delta_i)$. As a result, the load under a given set of inaccuracy thresholds $\{\Delta_i\}$ can be computed by $\sum_{i=1}^l n_i \cdot f_r(\Delta_i)$. This load should be less than or equal to the throttle fraction times the load of the ideal case of $\forall_{i \in [1..l]}, \Delta_i = \epsilon_+$, which is given by $z \cdot n \cdot f_r(\epsilon_+)$. Thus,

the following processing constraints must hold for the update load shedding problem:

$$\sum_{i=1}^l n_i \cdot f_r(\Delta_i) \leq z \cdot n \cdot f_r(\epsilon_+)$$

$$\forall_{i \in [1..l]}, \epsilon_+ \leq \Delta_i \leq \epsilon_-$$

The second constraint defines the domain of the inaccuracy threshold Δ_i ($i \in [1..l]$). The key question here is how to partition the space of interest into a number of regions such that the overall quality Ψ_u is maximized. To perform this we use an extension of the previously proposed GRIDREDUCE algorithm [17]. Our extensions are targeted toward building QoS-awareness into the algorithm.

B. The Space Partitioning Algorithm

The goal of the space partitioning algorithm is to partition the geographical space of interest into l shedding regions, such that this partitioning produces query results of higher accuracy. Concretely, the algorithm first builds a partition hierarchy over the space by constructing a quad-tree, where each tree node corresponds to a different region in the space. The partition hierarchy contains a single region at the top level and becomes more fine-grained as we go down in the hierarchy. This is achieved by dividing each partition into four quadrant partitions at the next lower level. Each level of the quad-tree is a uniform and non-overlapping partitioning of the entire space. In order to capture the QoS characteristics of each partition, we aggregate the QoS functions for each partition in the hierarchy through a post-order traversal of the quad-tree. Once the hierarchy is constructed and the QoS functions are aggregated for each partition, the selection of the l regions follows a top-down process. The algorithm starts from the top-most partition in the hierarchy. At each step one partition is picked and is replaced by its four quadrants. This process continues until l regions of possibly different sizes are reached. The criterion used to decide which partition to pick at each step forms the crux of the algorithm. With the QoS functions at hand, we compute how beneficial it is to divide a partition into its quadrants by solving the small scale version of the QoS-aware update load shedding problem we formalized in the previous subsection by restricting it to the four quadrants at hand. The partition that provides the highest gain in terms of the QoS measure Ψ_u is picked for further partitioning.

VI. PUTTING IT ALL TOGETHER: MOBIQUAL SOLUTION

In this section we first formalize the problem of combining QoS-aware update load shedding and QoS-aware query load shedding. Then we present a fast greedy algorithm called the *Minimum quality loss per cost step* (MQLS) to configure the re-evaluation periods $P_j, j \in [1..k]$ and the result inaccuracy thresholds $\Delta_i, i \in [1..l]$ within the same framework, aiming at achieving high overall QoS and better satisfying the freshness and accuracy requirements of mobile location queries.

A. Problem Formalization

The objective of the combined load shedding problem is to maximize the overall quality $\Psi = \frac{1}{m}(\Psi_v + \Psi_u)$ given in Equation 5. We now restate the processing constraint by combining the load due to query re-evaluation and update processing.

Let z_v denote the fraction of the query load retained for a given set of re-evaluation periods $\{P_j\}$. We have: $z_v = \frac{\sum_{j=1}^k f_c(C_j)/P_j}{\sum_{q \in Q} f_c(\{q\})/\tau_-}$. Similarly, let z_u denote the fraction of the update load retained for a given set of inaccuracy thresholds $\{\Delta_i\}$. We have: $z_u = \frac{\sum_{i=1}^l n_i \cdot f_r(\Delta_i)}{n \cdot f_r(\epsilon_-)}$. With these definitions, we can state the processing constraint as follows:

$$z_v + z_u \cdot \gamma \leq z \cdot (1 + \gamma) \quad (11)$$

The parameter γ in Equation 11 represents the cost of performing update processing with the setting of $\forall_i, \Delta_i = \epsilon_-$ compared to the cost of performing query re-evaluation with the setting of $\forall_j, P_j = \tau_-$. In other words, for the ideal case the query re-evaluations costs 1 unit, whereas the update processing costs $\gamma \in (0, \infty]$ units. Note that γ is *not* a system specified parameter and is learned adaptively as follows. Let U be the observed cost of update processing and V be the observed cost of query re-evaluation during the last adaptation period. Then we have $\gamma = \frac{U/z_u}{V/z_v}$. This assumes that the workload does not significantly change within the time frame of the adaptation period. Recall that the load shedding parameters are configured after each adaptation period, thus yielding new values for z_u and z_v (by way of changing P_j 's and Δ_i 's). Thus the combined load shedding problem is formalized as follows:

| |
|---|
| <p>maximize $\Psi = \frac{1}{m} \left(\sum_{j=1}^k \mathcal{V}_j^*(P_j) + \sum_{i=1}^l \mathcal{U}_i^*(\Delta_i) \right)$ subject to $\frac{\sum_{j=1}^k f_c(C_j)/P_j}{\sum_{j=1}^k f_c(C_j)/\tau_-} + \gamma \cdot \frac{\sum_{i=1}^l n_i \cdot f_r(\Delta_i)}{n \cdot f_r(\epsilon_-)} \leq z \cdot (1 + \gamma)$ $\forall_{j \in [1..k]}, \tau_- \leq P_j \leq \tau_+, \forall_{i \in [1..l]}, \epsilon_- \leq \Delta_i \leq \epsilon_+$</p> |
|---|

Note that this is a non-linear program, since the constraints have $1/P_j$ terms and are not linear. We now describe MQLS – a fast, greedy algorithm for setting the re-evaluation periods and inaccuracy thresholds to solve the above stated QoS-aware load shedding problem.

B. The MQLS Algorithm

The basic principle of the MQLS algorithm is to start with the ideal case of $\forall_j, P_j = \tau_-$ and $\forall_i, \Delta_i = \epsilon_-$ and incrementally reduce the load to z times that of the ideal case by repetitively increasing the re-evaluation period or the inaccuracy threshold that gives the smallest quality loss per unit cost reduction. The algorithm is greedy in nature, since it takes the minimum quality loss per cost step. Concretely, we partition the domain of re-evaluation periods and inaccuracy thresholds into β segments, such that we increase the P_j 's and Δ_i 's in increments of size $c_v = (\tau_+ - \tau_-)/\beta$ and $c_u = (\epsilon_+ - \epsilon_-)/\beta$, respectively. The MQLS algorithm maintains a

min. heap that stores a *qlpc* (quality loss per unit cost³) value for each re-evaluation period and each inaccuracy threshold. The *qlpc* value of a re-evaluation period (or an inaccuracy threshold) gives the quality loss per unit cost for increasing it by c_v units (or c_u units). The *qlpc* value is denoted by S_j^v for query group C_j and S_i^u for shedding region A_i . We have:

$$S_j^v = \sum_{q \in Q} f_c(\{q\}) \cdot \frac{\mathcal{V}_j^*(P_j + c_v) - \mathcal{V}_j^*(P_j)}{f_c(C_j) \cdot \left(\frac{1}{P_j + c_v} - \frac{1}{P_j} \right)} \quad (12)$$

$$S_i^u = \gamma \cdot n \cdot f_r(\epsilon_-) \cdot \frac{\mathcal{U}_i^*(\Delta_i + c_u) - \mathcal{U}_i^*(\Delta_i)}{n_i \cdot (f_r(\Delta_i + c_u) - f_r(\Delta_i))} \quad (13)$$

The nominators of the second components in the above equations represent the changes in the quality due to the increment, whereas the denominators represent the changes in the cost. Note that the first components of the above equations are used to normalize the costs in the denominators, so that S_j^v 's and S_i^u 's can be compared.

When the MQLS algorithm starts, the current load expenditure of the system, which is the sum of the load due to update and query load shedding appropriately weighted by γ , is above our load budget imposed by the throttle fraction z . The algorithm iteratively pops the topmost element of the min. heap and depending on whether we have a re-evaluation period or inaccuracy threshold makes the increment using either c_v or c_u . The *qlpc* value of the popped element is updated based on Equation 12 (or Equation 13) and is put back into the heap unless no further increments are possible. The algorithm runs until the load expenditure of the system is within the budget or all the re-evaluation periods and inaccuracy thresholds hit their maximum value. In the latter case the load cannot be shed to meet the processing constraint and random dropping of incoming updates as well as delay in query re-evaluations will unavoidably take place.

The total number of greedy steps the algorithm can take is given by $\beta \cdot (l + k)$, which happens when all re-evaluation periods and inaccuracy thresholds have to be increased to their maximum values. Each greedy step takes $\mathcal{O}(\log(l + k))$ time, since the min. heap has $l + k$ elements and the heap operations used take logarithmic time on the heap size. The final time complexity of the MQLS algorithm directly follows as $\mathcal{O}(\beta \cdot (l + k) \cdot \log(l + k))$ and the space complexity as $\mathcal{O}(l + k)$.

VII. EXPERIMENTAL EVALUATION

In this section we compare the performance of the MobiQual system to a number of other alternatives. These are:

- **Query-only load shedding:** QoS-aware differentiated load shedding with respect to re-evaluation periods only (see Section IV) and uses a fixed inaccuracy threshold of ϵ_- .
- **Update-only load shedding:** QoS-aware differentiated load shedding with respect to inaccuracy thresholds only (see Section V) and can be seen as the QoS-aware extension of the Lira approach [17]. Thus we name it as Lira⁺.

³This is *qlpc* for a query group or for a region, and not for a query as it was first introduced in Section IV-B. The core concept is the same.




Fig. 2: Example QoS functions, with different mid-point QoS values ($y_{0.5}$)

– **Single Δ -P**: Combined QoS-aware query and update load shedding, but without query grouping (QLBC algorithm from Section IV-C) and space partitioning (extended GRIDREDUCE algorithm from Section V-B). It represents a special case of the MobiQual system with $k = l = 1$.

We evaluate the MobiQual system using four main evaluation metrics. These include:

- i) The overall quality metric Ψ , as defined by Equation 5.
- ii) The mean period delay D , which is defined as the average difference between the ideal case period τ_{\perp} and the assigned period of queries, $\tau_q = P_j$ for $q \in C_j$. The mean period delay is formulated as:
$$D = \frac{1}{m} \sum_{q \in Q} (\tau_q - \tau_{\perp})$$
- iii) The mean position error R , which is defined as the average error in the positions of the mobile nodes within query results, relative to the error for the ideal case of $\forall i \in [1..l] \Delta_i = \epsilon_{\perp}$. It is formulated as:
$$R = \frac{1}{m} \sum_{q \in Q} (\epsilon_q - \epsilon_{\perp})$$
- iv) The running time of the adaptation step, which includes configuring a new set of re-evaluation periods and inaccuracy thresholds using the MQLS algorithm.

A. Experimental Setup

To create the mobile node movement trace used in the experiments, we used a real-world road map from the Chamblee region of the state of Georgia, USA. The trace covers a region of around 200km². We used real-world traffic volume data at the granularity of specific road types (such as expressway, arterial, collector), taken from [20], to simulate cars going on roads. The trace contains around 15K mobile nodes. The default re-evaluation period range used for the experiments is $[\tau_{\perp}, \tau_{\rightarrow}] = [1, 10]$ seconds, whereas the inaccuracy threshold range used is $[\Delta_{\perp}, \Delta_{\rightarrow}] = [5, 100]$ meters. The number of regions used for partitioning is set as $l = 250$ (see [17]). The increments used by the MQLS algorithm are determined using $\beta = 100$, i.e., the maximum number of increments possible is 100 for each re-evaluation period and inaccuracy threshold. The queries used in the experiments are range queries. The query distribution is proportional to the object distribution. Inverse and random distributions were also used, with similar results. The query side lengths were randomly chosen from the range $[0, 1000]$ meters.

A number of system and workload parameters were varied in the course of the experiments to understand their impact on the query result quality and running-time performance of the MobiQual system. These include the number of query groups used, i.e., the k parameter used by the QLBC algorithm

(default value: 16), the number of queries to number of objects ratio (default value: 0.01), the emulated capacity of the system (default: $z = 0.5$), and the QoS functions specified by the queries. Figure 2 gives the general template of the QoS functions that were used for both V_q (freshness) and U_q (accuracy) components of a query q 's QoS specification function S_q , whereas the α value that adjusts the relative importance of the freshness and accuracy components of quality were chosen at random from the range $[0, 1]$. The QoS functions were approximated by 10 linear segments and a parameter called *mid-point QoS threshold* was used to pick a random V_q or U_q component from the set of available functions, a subset of which is shown in Figure 2. Any given V_q or U_q is chosen by randomly picking a number, say $y_{0.5}$, between 0 and the mid-point quality threshold, and determining the QoS function whose value for the mid-point of its domain is equal to $y_{0.5}$ (and matching the template given in Figure 2).

B. Experimental Results

We divide the experimental results into three parts. The first part deals with the impact of the amount of load to be shed on the query result quality. The second part deals with the performance of MobiQual under different query loads and the impact of the number of query groups on the query result quality as well as on the time it takes to perform the adaptation step. The third part deals with the impact of the QoS specifications on the performance of MobiQual.




Fig. 3: Overall result quality as a function of throttle fraction

1) *Impact of the Throttle Fraction*: The graphs in Figure 3 plot the overall quality of the query results as a function of the throttle fraction (i.e., at different load shedding levels) for the competing approaches. At a given load shedding level, if $1 - z$ fraction of the load cannot be shed by a load shedding algorithm, then the QoS value is not plotted for that z value, and for smaller z values thereof. For instance, we observe from Figure 3 that, for the default settings, the query-only approach can only support load shedding for $z \geq 0.7$ and Lira⁺ for $z \geq 0.5$, whereas MobiQual and Single Δ -P can support $z \geq 0.2$. MobiQual significantly outperforms update-only and query-only load shedding schemes, as it is observed from the rapidly declining QoS values of the latter two approaches with decreasing z . Furthermore, MobiQual outperforms Single Δ -P, for a wide range of z values. While shedding 60% ($z = 0.4$) of the load, MobiQual is able to keep the QoS around $\Psi = 0.9$, whereas this value is only around 0.75 for the Single Δ -P approach. Similarly, MobiQual manages to sustain a QoS value of $\Psi = 0.7$ for 70% load shedding, compared

to a mere 0.4 for Single Δ -P. The two approaches both hit the $\Psi = 0$ boundary when MobiQual is forced to set all query re-evaluation periods and inaccuracy thresholds to their maximum value, at which point there is no difference between the two. The superior performance of MobiQual compared to Single Δ -P illustrates the strength of the differentiated load shedding concept, whereas the poor performances of update-only and query-only load shedding attest to the importance of performing combined query and update load shedding.




Fig. 5: Mean position error as a function of the throttle fraction

The graphs in Figures 4 and 5 plot the mean period delay and mean position error as a function of the throttle fraction for competing approaches, respectively. Note that the query-only approach has zero mean position error (as observed from Figure 5), whereas the update-only approach has zero mean period delay (as observed from Figure 4). However, since a good overall quality requires balancing freshness and accuracy, these two approaches do not provide good overall QoS as observed from Figure 3. The mean period delay of Single Δ -P stays slightly above that of MobiQual for $z > 0.3$. After this point Single Δ -P registers lower mean period delays. This is because further increasing the single re-evaluation period has diminishing benefit in terms of the qlpc metric, since Single Δ -P cannot provide differentiated load shedding. In contrast, the MobiQual approach can locate queries that can tolerate further staleness with less impact on the QoS value, due to the QLBC algorithm, and thus can increase the re-evaluation periods further for such queries. Even though this results in higher mean period delay compared to Single Δ -P, it translates into a higher overall QoS due to a better balance between query and update load shedding. It is observed from Figure 5 that MobiQual consistently outperforms Single Δ -P in terms of the mean position error. This is not because MobiQual sheds less update load, but it is because MobiQual sheds the update load from regions that has lesser impact on the query results, due to the QoS-aware partitioning algorithm it employs.

2) *Impact of the # of Queries and Query Groups:* The graphs in Figure 6 plot the overall QoS of the query results as a function the number of queries to number of mobile nodes ratio, for MobiQual vs. query-only and update-only (Lira+) load shedding. The throttle fraction is set to 0.75 for this experiment. It is interesting to observe that as the number of queries increase, the update-only load shedding loses its advantage over query-only load shedding. This is because with increasing number of queries, the dominant cost




Fig. 6: Overall result quality with changing query workload




Fig. 7: The impact of number of query groups on result quality

becomes the query re-evaluation, since the full update load does not depend on the number of queries. This shows the importance of performing combined query and update load shedding, which is effective independent of the number of queries or the number of mobile nodes, as evidenced by the superior performance of MobiQual compared to query-only and update-only approaches with changing number of queries to number of mobile nodes ratio (see Figure 6).

An important parameter that impacts the performance of the MobiQual system is the number of query groups, k . As discussed in Section II-D, in general the higher the number of query groups the more fine grained is the differentiated load shedding. The only limiting factor in increasing the value of k is the time it takes to execute the adaptation step, as the computational complexity of the MQLS algorithm is dependent on k . However, increasing the number of k has diminishing return in terms of the overall QoS, as shown by Figure 7, since the query groups become more and more homogeneous in terms of the QoS functions of the queries contained within. The graphs in Figure 7 plot the overall QoS as a function of the throttle fraction (x -axis is in logarithmic scale) for different k values. This experiment is run for 1000 continual queries. We clearly see from the figure that the gain in QoS when going from $k = 8$ to $k = 16$ is significantly lower than the gain in QoS when going from $k = 1$ to $k = 2$. This shows that having query groups smaller than 50-60 queries does not bring much gain in overall query result quality. Even though small query groups are unnecessary, the MQLS algorithm can support large k values with low overhead. Figure 8 shows that for $k = 16$ and $z = 0.5$ the adaptation step takes around 110 milliseconds. In a mobile CQ system, the change in the workload in terms of the number of CQs and mobile nodes is not spontaneous, and significant shifts in the workload is likely to happen within minutes. Thus the time it takes to run the adaptation step in order to configure the new set of re-evaluation periods and inaccuracy thresholds is relatively small compared to the adaptation period, resulting in a very lightweight load shedding scheme.

3) *Impact of the QoS Specifications:* The graphs in Figures 9 and 10 plot the overall query result quality as a function of the mid-point QoS threshold used for the freshness component of the QoS specifications and the accuracy component of it, respectively. Decreasing values along the x -axis represent QoS specifications with increasingly stringent freshness components for Figure 9 and increasingly stringent accuracy components for Figure 10. A high throttle fraction value of




Fig. 8: Adaptation time with different number of query groups




Fig. 9: Query result quality for varying freshness QoS specs.




Fig. 10: Query result quality for varying accuracy QoS specs.

0.75 was used to make sure that all competing approaches can shed the required fraction of the load. Note that update-only load shedding (Lira⁺) is indifferent to the freshness components of the QoS specifications, whereas query-only load shedding is indifferent to the accuracy components. As a result, the lines for update-only and query-only load shedding are flat in Figures 9 and 10, respectively. We observe from Figure 9 that MobiQual is very robust to changes in the freshness components of the QoS specifications and shows a smaller decrease in overall QoS with increasing intolerance to staleness in QoS specifications, compared to alternative approaches. It provides up to 12% better QoS compared to query-only load shedding and 5% better compared to Single Δ -P. These values are valid for shedding 25% percent of the load. The improvement provided by MobiQual over the closest competitor reaches 80% when shedding 70% of the load, and for a mid-point freshness QoS threshold of 0.75 (see previous Figure 3). The results from Figure 10 for the mid-point accuracy QoS threshold are very similar.

VIII. CONCLUSION

In this paper we presented MobiQual, a load shedding system aimed at providing high quality query results in mobile CQ systems. MobiQual has three unique properties. First, it uses per-query QoS specifications that characterize the tolerance of queries to staleness and inaccuracy in the query results, to maximize the overall QoS of the system. Second, it effectively combines query load shedding and update load shedding within the same framework, through the use of differentiated load shedding concept. Finally, the load shedding mechanisms used by MobiQual are lightweight, enabling quick adaption to changes in the workload. Through experimental studies, we showed that MobiQual significantly outperforms approaches that are based on query-only or update-only load shedding, as well as approaches that do combined query and update load shedding but lack the differentiated load shedding elements of the MobiQual solution.

REFERENCES

- [1] "NextBus," <http://www.nextbus.com/>, January 2004.
- [2] "Google RideFinder home page," <http://labs.google.com/ridefinder>, Febuary 2006.
- [3] B. Gedik, K.-L. Wu, P. S. Yu, and L. Liu, "Processing moving queries over moving objects using motion adaptive indexes," *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 5, 2006.

- [4] H. Hu, J. Xu, and D. Lee, "A generic framework for monitoring continuous spatial queries over moving objects," in *Proceedings of ACM SIGMOD*, 2005.
- [5] C. S. Jensen, D. Lin, and B. C. Ooi, "Query and update efficient B+-tree based indexing of moving objects," in *Proceedings of VLDB*, 2004.
- [6] M. F. Mokbel, X. Xiong, and W. G. Aref, "SINA: Scalable incremental processing of continuous queries in spatio-temporal databases," in *Proceedings of ACM SIGMOD*, 2004.
- [7] K.-L. Wu, S.-K. Chen, and P. S. Yu, "Efficient processing of continual range queries for location-aware services," *Information Systems Frontiers*, vol. 7, 2005.
- [8] O. Wolfson, P. Sistla, S. Chamberlain, and Y. Yesha, "Updating and querying databases that track mobile units," *Distributed and Parallel Databases*, vol. 7, no. 3, 1999.
- [9] A. Civilis, C. S. Jensen, and S. Pakalnis, "Techniques for efficient road-network-based tracking of moving objects," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 5, 2005.
- [10] S. Saltenis, C. S. Jensen, S. T. Leutenegger, and M. A. Lopez, "Indexing the positions of continuously moving objects," in *Proceedings of ACM SIGMOD*, 2000.
- [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 *Proceedings of VLDB*, 2003.
- [12] X. Xiong and W. G. Aref, "R-trees with update memos," in *Proceedings of IEEE ICDE*, 2006.
- [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 Transactions on Computers*, vol. 51, no. 10, 2002.
- [14] K.-L. Wu, S.-K. Chen, and P. S. Yu, "Incremental processing of continual range queries over moving objects," *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 11, 2006.
- [15] Y. Cai and K. A. Hua, "Real-time processing of range-monitoring queries in heterogeneous mobile databases," *IEEE Transactions on Mobile Computing*, vol. 5, no. 7, 2006.
- [16] B. Gedik and L. Liu, "Mobieyes: A distributed location monitoring service using moving location queries," *IEEE Transactions on Mobile Computing*, vol. 5, no. 10, 2006.
- [17] B. Gedik, L. Liu, K.-L. Wu, and P. S. Yu, "Lira: Lightweight, region-aware load shedding in mobile CQ systems," in *Proceedings of IEEE ICDE*, 2007.
- [18] —, "Region-aware load shedding in mobile continual query systems," GaTech, Tech. Rep. CERCS-06-07, 2006.
- [19] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*. Morgan Kaufmann, August 2000.
- [20] M. Gruteser and D. Grunwald, "Anonymous usage of location-based services through spatial and temporal cloaking," in *Proceedings of ACM MobiSys*, 2003.