

The Pennsylvania State University
The Graduate School

ADVANCED PATH QUERIES IN ROAD NETWORKS

A Thesis in
Computer Science and Engineering
by
Yuan Tian

© 2013 Yuan Tian

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

December 2013

The thesis of Yuan Tian was reviewed and approved* by the following:

Wang-Chien Lee
Associate Professor of Computer Science and Engineering
Thesis Advisor

Sencun Zhu
Associate Professor of Computer Science and Engineering

Raj Acharya
Department Head and Professor of Computer Science and Engineering

*Signatures are on file in the Graduate School.

Abstract

On a road network, the minimum cost path (or min-cost path for short) from a source location to a destination is a path with the smallest travel cost among all possible paths. Despite that min-cost path queries on static networks have been well studied, the problem of monitoring min-cost paths on a road network in presence of updates is not fully explored. In this thesis, we present *PathMon*, an efficient system for monitoring min-cost paths in dynamic road networks.

PathMon addresses two important issues of the min-cost path monitoring problem, namely, (i) *path invalidation* that identifies min-cost paths returned to path queries affected by network changes, and (ii) *path update* that replaces invalid paths with new ones for those affected path queries. For (i), we introduce the notion of *query scope*, based on which a *query scope index* (QSI) is developed to identify affected path queries. For (ii), we devise a *partial path computation algorithm* (PPCA) to quickly recompute the updated paths.

We have conducted a comprehensive performance evaluation by simulation, on various road networks and simulation settings. In these evaluations, we demonstrate that the proposed techniques, QSI and PPCA, are very effective on the path invalidation and path update issues. The query processing performance of QSI and PPCA outperformed the compared algorithms.

Table of Contents

List of Figures	vi
List of Tables	vii
Acknowledgments	viii
Chapter 1	
Introduction	1
Chapter 2	
Preliminaries	3
2.1 Road Network Data Modeling	3
2.1.1 Path Queries in Location Based Services	5
2.2 Problem Formulation	6
2.3 Research Issues and Solution Overview	7
Chapter 3	
Theoretical Analysis	9
3.1 Min-cost path search properties	9
3.2 Path invalidation with query scopes	11
3.3 Path update with processing states	13
3.4 PathMon Framework	17
Chapter 4	
PathMon: Data Structures and Implementation	19
4.1 Path Invalidation	19
4.1.1 Query Scope Representation	19
4.1.2 Query Scope Index	20

4.2	Partial Path Re-computation	24
Chapter 5		
	Performance Evaluation	30
5.1	Experiment Setup	30
5.2	Effectiveness of Query Scope Representations	31
5.3	Effectiveness of Checkpoint Selection	32
5.4	Path Monitoring Performance	33
Chapter 6		
	Conclusion	37
	Bibliography	39

List of Figures

2.1	Representing a spatial network in graph	4
2.2	Example road network	7
3.1	Illustration of the min-cost path spanning tree	11
3.2	Example of processing state	15
3.3	The PathMon framework	18
4.1	Illustration of partial materialization	21
4.2	The query scope index structure	22
4.3	Illustration of PPCA	27
5.1	Re-computation ratio	32
5.2	Checkpoint selection strategy comparison (CA)	33
5.3	Performance under various path length	34
5.4	Performance under various # of updated edges	34
5.5	Performance under A* algorithm	35
5.6	Performance for various cities ($L_p = 0.25D_N$, $ E_u = 0.05\% E $)	35

List of Tables

3.1 The trace of a search for $P(s, d)$ 10

5.1 The evaluation parameters 31

Acknowledgments

I want to take this opportunity to thank my adviser, Professor Wang-Chien Lee, for the inspiring guidance and strict training over the past few years. I would not have grown so much without his advice. I also deeply appreciate him for supporting my career decision and encouraging me to pursue what I am passionate about. All I have to give back, is only a thankful heart, a promise of life-time friendship, and an endless fight through the path to career success.

Dr. Ken C. K. Lee, my research mentor and trustful friend, has contributed a lot into these research. He is not only a great mentor, who led me step-by-step in doing research, but also a personal model, with his smartness, self-discipline, patience and kindness.

My parents, with all these years I am away from home, have done everything they can to support me and help me. Their unconditional love and trust are the source of my energy.

Last but not the least, thanks to my husband, Xingjie, for being an awesome friend and co-worker, and a perfect partner in life. His unusual ideas inspired me many times when I am blocked in the research. In the down times, it is always him to cheer me up and regain the courage to move forward.

Chapter 1

Introduction

How can I get out from the traffic jam and arrive at home as soon as possible? How do I get to the Empire State Building in Manhattan from State College on a Monday afternoon? What is the least expensive route from Napa Valley to San Francisco? If you have ever asked yourself any of these questions or something similar, you were probably able to get the answers just by opening your web browser or picking up your mobile phone. Thanks to the increasing availability of map data, GPS-enabled devices and the rapidly developing spatial data management technology, the location based services (LBSs) is about to boom big time. Major investors of LBS include leading IT companies such as Google, Yahoo!, Microsoft, and Apple, GPS device providers like Garmin, Tomtom and Lowrance.

Path search is one of the most popular and essential location-based services in our daily life. Given a starting location and a destination, there can be multiple possible paths. Among all these paths, users may have different preferences in choosing an optimal one. How to help users define the optimal path and efficiently process the queries are the core of path search problem, and here we look into the aspect of *minimum cost paths*.

Minimum cost path in dynamic road networks. If we consider the metrics in travel along a path as the cost of this path, such as distance, driving time, toll charge, etc., then the one with the smallest cost is usually seen as desirable to users. However, as the road network is dynamic, changes of the underlying network may cause the cost of any min-cost path between a given source and a destination no longer the smallest. Once a min-cost path query is issued, the user may want to stay

updated with the latest optimal route. In this case, a min-cost path monitoring system is demanded. Therefore, we present a min-cost path monitoring framework called *PathMon*, to monitor the changes of dynamic road networks and determine the most current min-cost paths between specified sources and destinations.

There are many real world applications that can apply our research solutions. For example, trucking and courier companies expect their vehicles to stay on min-cost paths to keep their operation costs low. Emergency service vehicles need to take the fastest routes to arrive at the destinations for life-critical missions. The min-cost path monitoring service will help efficiently finding the paths for these users in real time. In addition, during rush hours, if the routing system is aware of multiple min-cost path queries, it may be able to strategically adjust the routes to alleviate traffic load on different roads.

This thesis will present the definition/modeling, problem analysis and insights into this interesting path query problem. The rest of this proposal is organized as follows. In Chapter 2, we introduce the necessary background of LBSs and road networks. The theoretical foundation of the minimum cost path monitoring techniques are discussed in Chapter 3 and Chapter 4 presents the implementation details and data structures. Chapter 5 provide a comprehensive evaluation of our solution against other approaches. Chapter 6 concludes our work.

Preliminaries

2.1 Road Network Data Modeling

There are two major types of data models used in LBSs. One is based on the assumption that the space is free to movements, so that the distances between spatial objects can be measured simply using the Euclidean distance [2, 3, 4, 5]. We call this the *Euclidean space model*. Public squares, grasslands and water fields are typical examples. The other model, which assumes that the spatial object movements are constrained by predefined trajectories, is called the *spatial network model*, because the predefined trajectories form a network in the space [6, 7, 8, 9, 10, 11, 12, 1]. Road networks and airlines are the most commonly seen representatives of the spatial networks. The distance between spatial objects is measured using the *network distance*, i.e., the total cost of the minimum cost path between the objects. Many online map services are based on this model. Because computing the network distance is an expensive operation, the data management and processing techniques for the spatial network model are more complicated, but they can be easily generalized to the Euclidean space model. Hence, in our study, we mainly focus on the spatial network model.

Generally, a graph model is used to represent a spatial network, in which the nodes are road ends, intersections or significant turning points, and the edges are the road segments connecting nodes. We denote such a graph as $\mathcal{G} = \langle N, E \rangle$, where $N = \{n_i\}$ and $E = \{\langle n_i, n_j \rangle | n_i, n_j \in N\}$. Figure 2.1 gives an example of the graph representation of a real spatial network. We remark that here the

edges may be directed or undirected, so that the one-way constraints and turn constraints are satisfied. Also, there are *spatial objects* distributed in the network, representing a wide variety of entities like people, vehicles, buildings, parks and so on. The geo-tagged multimedia data such as texts, photos and videos can be considered as spatial objects, too. We denote the spatial objects as $\mathcal{O} = \{O_i\}$. The nodes, edges and spatial objects are usually associated with a set of features, such as geo-locations, road travel costs, shapes, address names, etc. We define a feature function $f(x)$ on the nodes, edges and objects to extract the features of a given target ϵ ($\epsilon \in N$ or $\epsilon \in E$ or $\epsilon \in \mathcal{O}$). The graph model naturally captures the connectivity properties of a spatial network, and with the features, the information associated with the graph can be very rich. Among all possible features, the location is one of the most interested. A location may be the coordinates in a Euclidean space, i.e. $f(\epsilon).location = (x, y)$, but in the spatial network, we usually attach the location with one or more edges, so that it can be written as $f(\epsilon).location = (\langle n_i, n_j \rangle, n_i, \sigma)$. Here, the edge $\langle n_i, n_j \rangle$ is the road from which the location can be directly reached (without going through any other roads), and σ ($0 \leq \sigma \leq f(\langle n_i, n_j \rangle).length$) is the physical length of the road segment from the node n_i to the location of ϵ .

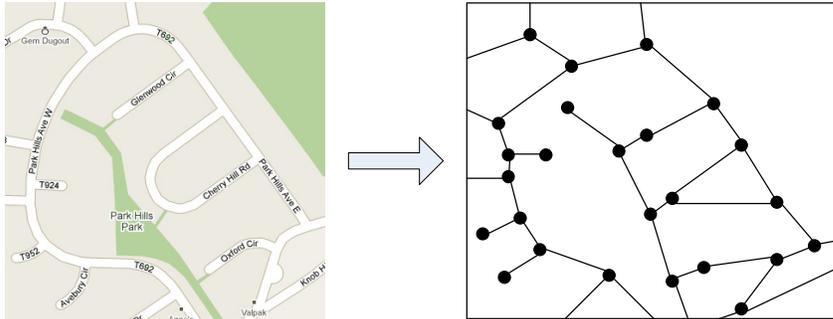


Figure 2.1. Representing a spatial network in graph

The dynamics of spatial networks to be dealt by LBSs come from three cases: 1) the change of road costs, 2) the movement of spatial objects and 3) the movement of reference points (i.e. locations at which queries are issued). Several works have been done on modeling the dynamic behavior of spatial network road costs. Some models such as stochastic models, speed patterns, time-delay functions consider the change of road networks to be predictable. In details, stochastic models [13]

handles the network changes in accordance with a distributed Markov chain. Speed pattern [14, 15] specifies the road costs at different time for different vehicle types. Time-delay functions [16] are used to represent the travel costs on individual road segments. Based on the predicted network costs, shortest paths can be determined. On the other hand, some works consider the dynamic networks as unpredictable. One of the most popular model is the quasi-dynamic model, where the edge costs are assumed to change quickly and then remain invariant until next changes [17, 18]. To model the movement of a spatial object or a reference point, a simple way is to use a discrete sequence of locations along with time stamps. If the trajectory is known before hand, we may also use a continuous curve to approximate it.

2.1.1 Path Queries in Location Based Services

The queries in LBS systems can be roughly categorized into two classes, namely, *path search* and *object search*. Path search aims at computing a set of paths from one location to another (or others), such that the aggregated path feature satisfy certain conditions. Minimum cost path query [19, 20, 21, 22, 8, 11] is the most frequently used path search function, because a minimum cost path is not only interesting to users, but also a fundamental component in a lot of other queries where network distances are concerned. Object search, on the other hand, aims at finding a set of spatial objects whose features satisfy certain conditions. In this thesis, we focus on path queries.

Min-cost path queries. Min-cost path queries for static network are well studied. Single-source min-cost path search algorithms include Bellman-Ford algorithm [23, 24], Dijkstra’s algorithm [19] and A* algorithm [20], while all-pair shortest path algorithms include Floyd-Warshall algorithm [25] and Seidel’s approach [26]. Since network traversal is a very expensive operation, especially in a large network, HEPV [27], HiTi [11] and a grid-based partition method in [28] were proposed to pre-compute and materialize some shortest paths to facilitate online path searches. Single source min-cost path update algorithms have been studied in [29, 30]. The basic idea is to propagate the changes from an updated edge to a set of vertices succeeding it in the shortest path tree. Then, all edges connecting to the affected vertices are examined to see whether they can form a path with smaller cost. [31] has studied the min-cost path monitoring problem. The main contribution of this

work is to quickly identify min-cost paths affected by a network cost change with elliptic bounding areas.

2.2 Problem Formulation

We model an ever-changing network as a set of time-stamped graphs $\{G_t | G_t = \langle N_t, E_t, \mathcal{C}_t \rangle\}$, where N_t is a set of nodes representing the end points of roads and road intersections present at time t ; $E_t = \{\langle n_i, n_j \rangle | n_i, n_j \in N_t\}$ is a set of edges standing for the road segments and $\mathcal{C}_t = \{C_{\langle n_i, n_j \rangle, t}\}$ is a set of positive values suggesting the cost of the edges at time t . Here, the cost may be the physical edge length, expected travel time, toll charge or others defined by applications.

Upon G_t , a path from a source node s to a destination node d ($s, d \in N_t$), denoted as $p(s, d)$, is a sequence of interleaving edges and nodes: $p(s, d) = (n_{i_0}, \langle n_{i_0}, n_{i_1} \rangle, n_{i_1}, \langle n_{i_1}, n_{i_2} \rangle, \dots, n_{i_{k-1}}, \langle n_{i_{k-1}}, n_{i_k} \rangle, n_{i_k})$ where $n_{i_0} = s$ and $n_{i_k} = d$. Because in general, an edge can be uniquely identified with its end nodes, we only use the nodes to represent a path, i.e. $p(s, d) = (n_{i_0}, n_{i_1}, \dots, n_{i_k})$. The total cost of a path $p(s, d)$ at time t , denoted by $C_{p(s,d),t}$, is equal to $\sum_{j=1}^k C_{\langle n_{i_{j-1}}, n_{i_j} \rangle, t}$. Among all possible paths formed between s and d , the min-cost path in G_t denoted by $P_t(s, d)$ is one with the smallest total cost. Hereafter, we denote the smallest path cost from s to d in G_t by $\|s, d\|_t$, which is also referred to as the *network distance* from s to d . We will use network distance and minimum path cost interchangeably.

Now, given a series of time-stamped graphs (i.e., $G_{t_0}, G_{t_1}, G_{t_2}, \dots$) that represent a network at different times, a continuous min-cost path query $q(s, d)$, as formalized in Definition 1, provides a series of min-cost paths, i.e. $P_{t_i}(s_t, d)$, corresponding to each G_{t_i} .

Definition 1. Continuous min-cost path query. *Given a dynamic road network G_t , a moving source node s (with the location at time t denoted as s_t) and a fixed destination node d , a continuous min-cost path query $q(s, d)$ returns a series of the latest min-cost paths $P_t(s_t, d)$ on G_t from the source node s_t to d , invoked by road network change at time t .¹ \square*

¹We assume that a mobile user always follows the path returned by the query. Hence, if road costs in the network remains unchanged, the previously returned path does not need updates.

2.3 Research Issues and Solution Overview

In a large-scale path monitoring system where a massive number of continuous min-cost path queries are running, an efficient mechanism to identify updated min-cost paths whenever network changes occur is needed. An intuitive approach to process continuous min-cost path queries is to reevaluate min-cost paths from scratch whenever a network update occurs. This is apparently inefficient since not all paths are necessarily *affected* by a network change. Moreover, path searches from scratch is computationally expensive. In fact, re-computation of some non-affected paths can be avoided and those affected queries can be partially recomputed.

Running example. In a simple road network as shown in Figure 2.2(a), a continuous min-cost path query is issued to monitor a min-cost path from node s to d . At time t_0 , a result path $P_{t_0}(s, d) = (s, e, f, d)$ is indicated in a dashed line with path cost 7. At time t_1 ($t_1 > t_0$), the cost of edge $\langle b, c \rangle$ changes from 4 to 1. However, it does not lead to a new min-cost path, i.e. $P_{t_0}(s, d) = P_{t_1}(s, d)$. In this case, $\langle b, c \rangle$ is a *non-affecting edge* with respect to $P_{t_0}(s, d)$.

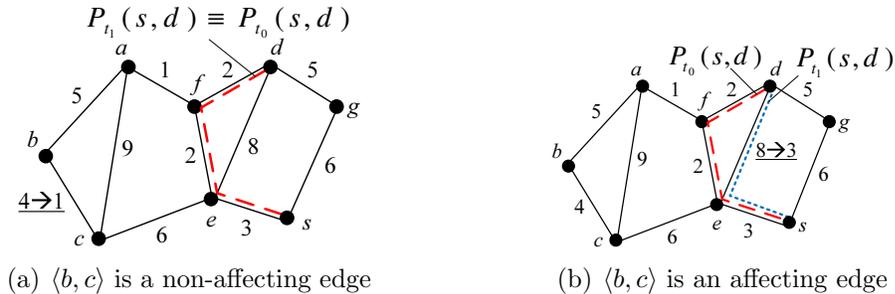


Figure 2.2. Example road network

Besides, at time t_2 ($t_2 > t_1$), the cost of $\langle e, d \rangle$ decreases from 8 down to 3 (as shown in Figure 2.2(b)). Now, $P_{t_2}(s, d)$ becomes (s, e, d) . In this case, $\langle e, d \rangle$ is an affecting edge. \square

From the example, we obtain two observations. First, with respect to a determined min-cost path, not all edge changes result in a new min-cost path. Therefore, those edges are non-affecting edges. Whenever an edge is identified to be a non-affecting edge to a min-cost path, it is safe to keep the original min-cost path. On the other hand, by detecting changes happened to the affecting edges, we can quickly identify which min-cost path queries are impacted by the changes. Second,

although an existing min-cost path may be affected by an edge change, we expect a large portion of the existing min-cost path remaining in the new min-cost path. Recall in the above example that the edge (s, e) remains in the new path $P_{t_2}(s, d)$. As such, it is possible to reuse some common parts of an existing path and only to compute a new portion in the new min-cost path.

In light of these observations, we approach the min-cost path monitoring problem by addressing two important issues, namely, (i) *path invalidation* and (ii) *path update*, and develop a novel, efficient and scalable path monitoring framework, called *PathMon*. PathMon contains two important components, namely, the *query scope index (QSI)* and the *partial path computation algorithm (PPCA)*. *QSI* keeps track of *query scopes*, i.e., subnetworks that are involved in derivation of answer paths for individual queries. In presence of edge cost updates within some query scopes, corresponding queries are detected to be affected (and the originally calculated min-cost path possibly becomes invalid). *PPCA* utilizes the intermediate processing states captured in previous path search processes to facilitate partial re-computation of affected paths. Since paths computed by a deterministic search algorithm on a slightly changed network would access almost the same set of road network data in the same order, some path re-computation can be alleviated by reusing some processing states of a previous search.

Theoretical Analysis

To address the path invalidation and path update issues, we need to look back to the fundamental operation, i.e. the path search algorithm, which provides critical information and useful properties for deriving solutions.

3.1 Min-cost path search properties

Without loss of generality, we consider Dijkstra's algorithm here due to its generality and popularity, and more importantly, its insightful properties. Please note that although the following discussions only focus on Dijkstra's algorithm, our analysis and theorems can be easily generalized to other path search algorithms.

Dijkstra's algorithm employs the best-first strategy in network traversal. Starting from the destination d , in each step, the algorithm picks a node v with the minimum accumulative cost to d (denoted as c_v) among all unexplored nodes to visit.¹ Once v is being visited, the accumulative cost to d is determined as the network distance $\|v, d\|$. Then, each unvisited neighbor w of v is assigned with the accumulative distance $c_w = \min\{c_w, \|v, d\| + C_{\langle w, v \rangle}\}$. This process repeats until the source node s is visited, and the answer path can be constructed by tracing back from s to d . A min-cost path spanning tree T rooted at d is formed during the algorithm, in which the path from any tree node n to d is the min-cost path between n and d .

¹Since the source node moves as the driver travels in the road network, we use the fixed destination as the starting node of the search.

We illustrate the search process of Dijkstra’s algorithm with the same sample road network as given in Figure 2.2(a). Two main data structures are utilized to support the search process. One is a visited node set V that prevents re-visiting of nodes and maintains the predecessor information for back tracking. The other is a priority queue H that sorts the unvisited nodes in non-descending order of accumulative costs to the destination.² The trace of the search is given in Table 3.1.

Step	Latest visited node (node id, predecessor, accumulative cost)	content of the heap (after exploring edges)
0	\emptyset	$((d, \perp, 0))$
1	$(d, \perp, 0)$	$((f, d, 2), (g, d, 5), (e, d, 8))$
2	$(f, d, 2)$	$((a, f, 3), (e, f, 4), (g, d, 5), (e, d, 8))$
3	$(a, f, 3)$	$((e, f, 4), (g, d, 5), (b, a, 8), (e, d, 8), (c, a, 12))$
4	$(e, f, 4)$	$((g, d, 5), (s, e, 7), (b, a, 8), (e, d, 8), (c, e, 10), (c, a, 12))$
5	$(g, d, 5)$	$((s, e, 7), (e, d, 8), (b, a, 10), (c, e, 10), (s, g, 11), (c, a, 12))$
6	$(s, e, 7)$	search terminates.

Table 3.1. The trace of a search for $P(s, d)$

The min-cost path spanning tree formed by Dijkstra’s algorithm is shown in Figure 3.1 in bold black lines (with the circled numbers indicating the order of node being visited). This spanning tree (denoted as T) has the following properties.

Property 1. *The min-cost spanning tree properties.*

1. Recursiveness. Let $\hat{T}(u)$ denote the subtree of T rooted at node u , then $\hat{T}(u)$ is also a min-cost path spanning tree from all nodes in $\hat{T}(u)$ to u , i.e. $\forall v \in \hat{T}(u)$, $P(v, u) \subseteq \hat{T}(u)$.
2. Distance priority. $\forall v \in \hat{T}(u)$, $\|v, d\| \geq \|u, d\|$.
3. Distance bound. $\forall v \in T$, $\|v, d\| \leq \|s, d\|$. On the other hand, $\forall v' \notin T$, $\|v', d\| \geq \|s, d\|$.

PROOF. Recursiveness and distance priority are quite straightforward and we omit the proof. Here we prove the distance bound property. The node visit order of Dijkstra’s algorithm is non-descending order of accumulative cost to the destination d . When the source node s is being explored, all the nodes whose distance to d is smaller than s must have been visited, so they are all covered by

²When two nodes have the same accumulative costs to the destination, we use the node IDs as the tie breaker.

T . Meanwhile, all the unvisited nodes must have their distance to d greater than or equal to $\|s, d\|$. ■

As exemplified in Figure 3.1, we have $e \in \hat{T}(f)$, thus $\|e, d\| = 4 > \|f, d\| = 2$, and $P(e, f) = \{e, f\} \subset \hat{T}(f)$. Also, all the nodes covered by T have smaller or equal distances to d than s .

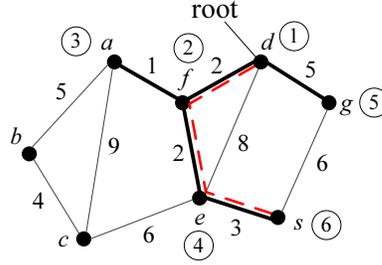


Figure 3.1. Illustration of the min-cost path spanning tree

3.2 Path invalidation with query scopes

Since s is a leaf node of T , if the min-cost path $P(s, d)$ changes, the structure of T must change. However, only monitoring the edges on the path is not enough. To show the relationship between the tree and the path in presence of changes, we consider the following cases (assuming $P_{t_0}(s, d) = (s, e, f, d)$, $t_1 > t_0$):

1. The updated edge is already in the tree T : e.g., $\langle s, e \rangle \in T$. If $C_{\langle s, e \rangle, t_1}$ increases to 10, $P_{t_0}(s, d) = (s, e, f, d)$ will be invalid, and the new path will be $P_{t_1}(s, d) = (s, g, d)$.
2. The updated edge is not in T , but both of its end nodes are covered by T : e.g., $s, g \in T$ but $\langle s, g \rangle \notin T$. If $C_{\langle s, g \rangle, t_1}$ decrease to 1, then correspondingly we will have $P_{t_1}(s, d) = (s, g, d)$, and the original path will become invalid.
3. Some updated edges have only one end node in T : e.g., $e, a \in T$, $c \notin T$, so $\langle e, c \rangle \notin T$, $\langle c, a \rangle \notin T$. If $C_{\langle e, c \rangle, t_1}$ and $C_{\langle c, a \rangle, t_1}$ both decrease to 0.1, $P_{t_0}(s, d)$ will change to $P_{t_1}(s, d) = (s, e, c, a, f, d)$.
4. None of the updated edges has any end node covered by T : e.g., $c, b \notin T$. If $C_{\langle c, b \rangle, t_1}$ changes, no matter what the new value is, the path will not be affected.

From the above discussion, we can see that if an edge has at least one end node covered by T , it may result in a change of the min-cost path. In order to catch all possible affecting edge changes, we need to monitor all edges linking to the nodes in T . In other words, the key issue in path monitoring is to detect changes involving the min-cost path spanning tree. Based on the spanning tree, we define the *query scope* of a path query as in Definition 2.

Definition 2. Query scope. *Given a min-cost path query q with the answer path $P_t(s, d)$ and the min-cost path spanning tree $T_t(q)$, the query scope $QS_t(q)$ is a set of nodes covered by $T_t(q)$. Only the edges with at least one end node in $QS_t(q)$ may be an affecting edge of q with respect to $P_t(s, d)$. \square*

According to the definition of query scope and the properties of the min-cost path spanning tree, we obtain the following property and condition for eliminating irrelevant network updates.

Property 2. Distance bound of query scope. *For a query $q(s, d)$, at any time t , its query scope $QS_t(q)$ satisfies $\forall v \in QS_t(q), \|v, d\|_t \leq \|s, d\|_t$. Meanwhile, $\forall v' \notin QS_t(q), \|v', d\|_t \geq \|s, d\|_t$.*

Lemma 1. *If at time $t_1 > t_0$ no edge with at least one end node in $QS_{t_0}(q)$ is updated, i.e. $\forall \langle x, y \rangle$ whose cost is changed, $x \notin QS_{t_0}(q) \vee y \notin QS_{t_0}(q)$, the min-cost path $P_t(s, d)$ will not be affected, i.e. $P_{t_1}(s, d) \equiv P_{t_0}(s, d)$. \square*

PROOF. Assume that $\forall \langle x, y \rangle$ updated at time t_1 ($t_1 > t_0$), $x \notin QS_{t_0}(q)$, $y \notin QS_{t_0}(q)$, and $P_{t_1}(s, d) \not\equiv P_{t_0}(s, d)$. According to Property 2, we have $\|x, d\|_{t_0} \geq \|s, d\|_{t_0}$, $\|y, d\|_{t_0} \geq \|s, d\|_{t_0}$. Since $P_{t_0}(s, d)$ is entirely covered by $QS_{t_0}(q)$, its cost does not change at t_1 . Then we have $C_{P_{t_1}(s, d)} < C_{P_{t_0}(s, d)}$, and $P_{t_1}(s, d)$ must include at least one updated edge (otherwise at time t_0 it already has smaller cost than $P_{t_0}(s, d)$, which is a contradiction).

Let u be one end node of a certain updated edge on $P_{t_1}(s, d)$, and $P_{t_1}(u, d) = (u, n_{i_1}, n_{i_2}, \dots, n_{i_k}, d)$ is part of $P_{t_1}(s, d)$. $\|d, d\|_{t_0} = 0 < \|s, d\|_{t_0}$, so all edges linking to d are not updated, i.e. $C_{\langle n_{i_k}, d \rangle, t_1} = C_{\langle n_{i_k}, d \rangle, t_0}$. $\|n_{i_k}, d\|_{t_1} = C_{\langle n_{i_k}, d \rangle, t_1} = C_{\langle n_{i_k}, d \rangle, t_0} \geq \|n_{i_k}, d\|_{t_0}$, and since $\|n_{i_k}, d\|_{t_1} < \|s, d\|_{t_0}$ (n_{i_k} is on the new min-cost path), we have $\|n_{i_k}, d\|_{t_0} < \|s, d\|_{t_0}$.

By induction, we have that all the nodes on $P_{t_1}(u, d)$, including u , have

$$\|n_i, d\|_{t_1} = \sum_{j=i}^{k-1} C_{\langle n_j, n_{j+1} \rangle, t_1} = \sum_{j=i}^{k-1} C_{\langle n_j, n_{j+1} \rangle, t_0} \geq \|n_i, d\|_{t_0}, \text{ and}$$

$$\|n_i, d\|_{t_0} \leq \|n_i, d\|_{t_1} < \|s, d\|_{t_0} \quad (n_i = u, n_{i_1}, n_{i_2}, \dots, n_{i_k}, d).$$

Thus, we have $\|u, d\|_{t_0} < \|s, d\|_{t_0}$, this contradicts with the assumption that the end node u of any updated edge have $\|u, d\|_{t_0} > \|s, d\|_{t_0}$.

Thus, $P_{t_1}(s, d) \equiv P_{t_0}(s, d)$. ■

Further, when no edge inside the query scope is updated, it can be proved that not only the original min-cost path does not change, but the entire min-cost path spanning tree remains unchanged.

Lemma 2. $T_{t_1}(q) \equiv T_{t_0}(q)$ if no edge with at least one end node in $QS_{t_0}(q)$ is updated at time t_1 . Hence, the query scope $QS_{t_1}(q) \equiv QS_{t_0}(q)$. □

PROOF. By treating each node in $T_{t_0}(q)$ as s and apply Lemma 1, we see that no path in $T_{t_0}(q)$ has been changed at time t_1 . So, $T_{t_1}(q) \equiv T_{t_0}(q)$. ■

Due to the distance bound property, we can express Lemma 1 and 2 using the distance criterion, as stated in Corollary 1.

Corollary 1. For a query $q(s, d)$, if at time $t_1 > t_0$, $\forall \langle x, y \rangle$ that is updated, $\|x, d\|_{t_0} > \|s, d\|_{t_0}$ and $\|y, d\|_{t_0} > \|s, d\|_{t_0}$, then $P_{t_1}(s, d) \equiv P_{t_0}(s, d)$, $QS_{t_1}(q) \equiv QS_{t_0}(q)$. □

Given the query scope $QS_t(q) = \{v | v \in T_t(q)\}$ and the above properties, we can derive our path invalidation algorithm based on detecting the changes in the query scopes.

3.3 Path update with processing states

Next, let us consider the path update problem.

As we have discussed, to guarantee the correctness of the min-cost path, the whole min-cost path spanning tree needs to be always up-to-date. According to the definition of a min-cost path spanning tree T (rooted at the destination node d), $\forall v \in T$, $P(v, d) \subseteq T$. If after some network updates, some min-cost paths in T remain unchanged, they may be reused as a part of the new spanning tree, i.e. if $\exists v \in T_{t_0}$, $P_{t_0}(v, d) \equiv P_{t_1}(v, d)$ and $\|v, d\|_{t_1} < \|s, d\|_{t_1}$, then $P_{t_1}(v, d) \subseteq T_{t_1}$. In

the path search algorithm, this min-cost path information from each node in the spanning tree to the root is stored in the visited node set V and priority queue H . From the visited entries in V , we can depict the current structure of the spanning tree. Meanwhile, H provides a guidance for the future steps without needing to retrieve the neighbors of nodes in V repeatedly.

From the trace in Table 3.1, we can also observe that, the node to be visited in a specific step is determined by the content of V and H at that step, and then V and H are updated according to the neighboring nodes and edges of the visited node. For instance, when $V = \{(d, \perp, 0)\}$, $H = ((f, d, 2), (g, d, 5), (e, d, 8))$, the node to be visited next must be f because it has the minimum accumulative distance to d in H and is not in V . This process is deterministic, i.e. if the network remains unchanged, the algorithm will always follow the same steps and find out the same min-cost path, and the final layout of the spanning tree will also be the same. Let us denote the content of V and H in the step that node n is visited as V^n and H^n , then (V^n, H^n) represents the state of the search process. We define (V^n, H^n) as the *processing state* at node n (see Definition 3), thus the entire search process can be viewed as a sequence of processing states, i.e. $S(n_1), S(n_2), \dots, S(n_k)$ where $S(n_i) = (V^{n_i}, H^{n_i})$, $i = 1, 2, \dots, k$.

Definition 3. Processing state. *Given a min-cost path query $q(s, d)$, the processing state $S(n)$ at a node n contains a set of visited nodes V^n and the current content of priority queue H^n just after n is visited. \square*

We exemplify the processing state in Figure 3.2. In the search for $P(s, d)$, the processing state $S(f)$ at node f contains the visited set $V^f = \{(d, \perp, 0), (f, d, 2)\}$ from which we can recover the spanning tree at the time f is visited, and the priority queue H^f with all nodes (a , e and g) linked to d and f but not yet explored. Since $(a, f, 3)$ is the first entry in H^f , the processing state $S(f)$ determines that the node a is going to be visited next.

Treating a min-cost path search as a process of growing a spanning tree, a processing state is generated each time a new node is added to the tree. We have the following observations. Let G and G' represent a network before and that after an edge change, respectively. Given the same path query evaluated on G and G' , we can obtain two sequences of states \mathcal{S} and \mathcal{S}' , respectively. If both \mathcal{S} and \mathcal{S}' are

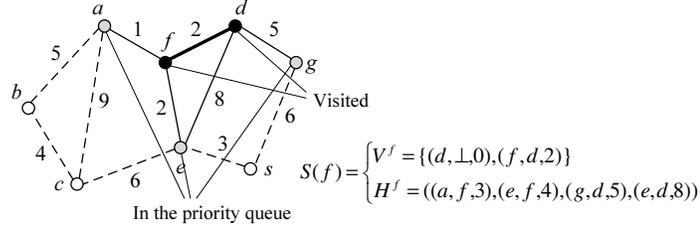


Figure 3.2. Example of processing state

exactly the same, it indicates the spanning trees T (obtained in G) and T' (obtained in G') are identical (thus the min-cost paths must be identical). We refer two sequences of states to be the same if they consist of the same number of states and each corresponding pair of states have the same sets of visited nodes and contents of priority queues. On the other hand, if \mathcal{S} and \mathcal{S}' are different, the spanning tree on G' may possibly become different from one obtained in G (even though the min-cost paths may or may not be different). While determining \mathcal{S}' , which equals to a full path evaluation, is computationally expensive, it is economic to reuse some of the states in \mathcal{S} if they are previously determined on G and maintained in the system. Logically, if not too many edges have been updated in G , some paths in T will still appear in T' . With the stored states of \mathcal{S} , we can directly rebuild the part of T that is unchanged in T' and avoid re-computation from scratch. Based on this idea, we can restart a path search at a state $S(n_k) \in \mathcal{S}$ which guarantees that $S(n_1), \dots, S(n_k)$ are identical to their counterparts in \mathcal{S}' . We explain this in Lemma 3.

Lemma 3. *Let $S(n)$ be the processing state at node n ($\|n, d\|_{t_0} \leq \|s, d\|_{t_0}$) in a search process for $P_{t_0}(s, d)$ before the network changes. If $\forall \langle x, y \rangle$ whose cost is updated, $\|x, d\|_{t_0} > \|n, d\|_{t_0}$ and $\|y, d\|_{t_0} > \|n, d\|_{t_0}$, and $S'(n)$ is the processing state at n in a search process after the changes, then $S'(n)$ is identical to $S(n)$, denoted as $S'(n) \equiv S(n)$. \square*

PROOF. From Lemma 2 we know that $\|x, d\|_{t_0} > \|n, d\|_{t_0} \Rightarrow \|x, d\|_{t_1} > \|n, d\|_{t_1}$, and similarly $\|y, d\|_{t_1} > \|n, d\|_{t_1}$. $\langle x, y \rangle$ can be accessed only when x or y is visited. According to the node ordering in H , i.e. non-descending order of the accumulative cost to d , n has already been visited before the visit to x or y . Therefore, $S'(n)$ has already been determined by the search process before

accessing any changed edge $\langle x, y \rangle$. Since the search process has only accessed unchanged edges by the time of generating $S'(n)$, due to the determinism of the search algorithm, $S'(n) \equiv S(n)$. \blacksquare

To illustrate Lemma 3, we consider the change of edge $\langle s, e \rangle$ as shown in Figure 3.2. Since $\min\{\|s, d\|, \|e, d\|\} = 4$ and $\|f, d\| = 2$, $S'(f)$ will remain the same as $S(f)$. In fact, because d, f and a are all nearer to d than $\langle s, e \rangle$, $S(d)$, $S(f)$ and $S(a)$ are not affected. Finally, as $\|a, d\| = 3 > \|f, d\| > \|d, d\|$, $S(a)$ should be the latest state not affected by the edge changed. Thus, the search can restart with the visited nodes and the priority queue of $S(a)$. This is analogy to “cut” the spanning tree to the shape stored in $S(a)$ and let it grow again.

Besides, if a state $S'(n) \in \mathcal{S}'$, which is obtained after incorporating the changed edges inside the (potentially enlarged) query scope, is identical to the state $S(n) \in \mathcal{S}$, it can be certain that the rest of states in \mathcal{S} and \mathcal{S}' are exactly the same. Thus, the computation of \mathcal{S}' can be terminated earlier by sharing the remaining states after $S(n)$ from \mathcal{S} . Denote $\{\langle u, v \rangle^+\}$ as the edges whose cost are increased, and $\sum \Delta C_{\langle u, v \rangle^+}$ gives the upper bound of query scope enlargement (so that $\|s, d\|_{t_1} \leq \|s, d\|_{t_0} + \sum \Delta C_{\langle u, v \rangle^+}$). The termination condition is stated in Lemma 4.

Lemma 4. *Let $S(n)$ be the processing state at node n ($\|n, d\|_{t_0} \leq \|s, d\|_{t_0}$) in a search process for $P_{t_0}(s, d)$ before the network changes. Given that all updated edge $\langle x, y \rangle$ satisfies either 1) $\|x, d\|_{t_1} > \|s, d\|_{t_0} + \sum \Delta C_{\langle u, v \rangle^+}$ and $\|y, d\|_{t_1} > \|y, d\|_{t_0} + \sum \Delta C_{\langle u, v \rangle^+}$, or 2) $\|x, d\|_{t_1} < \|n, d\|_{t_1}$ and $\|y, d\|_{t_1} < \|n, d\|_{t_1}$, and $S'(n)$ is the processing state at n in a search process after the changes, if $S'(n) \equiv S(n)$, then for any node m such that $\|n, d\|_{t_0} < \|m, d\|_{t_0} \leq \|s, d\|_{t_0}$, $S'(m) \equiv S(m)$. \square*

PROOF. In the Dijkstra’s algorithm, a processing state is only determined by the processing state generated in the previous expansion iteration and the costs of unexplored edges, but not directly related to earlier processing states. The updated edges $\langle x, y \rangle$ with $\|x, d\|_{t_1} > \|s, d\|_{t_0} + \sum \Delta C_{\langle u, v \rangle^+} \geq \|s, d\|_{t_1}$ and $\|y, d\|_{t_1} > \|s, d\|_{t_0} + \sum \Delta C_{\langle u, v \rangle^+} \geq \|s, d\|_{t_1}$ will not influence the search (Lemma 3); The updated edges with the network distances from both end points to d smaller than $\|n, d\|_{t_1}$ all explored, so the unexplored edges all are unchanged. Given a processing state $S'(n)$ that is identical to $S(n)$ and a set of unchanged edges to be accessed afterwards, the same search steps can be deterministically derived. In

each step, the node being visited and its associated processing state will then be identical to those in the previous search process. ■

Our partial path re-computation algorithm can be derived based on the above analysis. By ensuring two conditions: i) restarting from a processing state which was generated in a previous search and unaffected by updates; ii) terminating at another processing state that guarantees the equality between the following search steps and existing search results; we can effectively save computation cost on unnecessary network traversals while preserve the algorithm correctness.

3.4 PathMon Framework

The PathMon framework, as shown in Figure 3.3, includes three major components: a) the network data index, b) the query index, and c) the path search / update processing engine. The data index maintains the network topological information, including nodes, edges and edge costs. In the implementation, we index those nodes and their edges as adjacent lists based on the CCAM storage scheme designed for road networks [42]. The data index is mainly used to support the evaluation of path queries and path updates. On the other hand, the query index keeps track of the registered continuous min-cost path queries. Upon network changes, the query index is looked up to determine affected path queries. Based on the idea of query scopes, we devise a *query scope index* as the query index. We leave the detailed discussion to Section 4.1. Last, the path search/update processing engine, triggered by events, is responsible for search or update paths. For example, when an edge is updated, the engine accesses the query index to identify affected queries and then accesses the data index to recompute the corresponding result paths. The engine also maintains the indices upon query arrival/removal and network updates. To reduce the cost of re-computation, we explore the partial path computation algorithm and detail it in Section 4.2.

In the PathMon framework, each continuous min-cost path query goes through a life cycle of three stages: (i) query initialization, (ii) path update and (iii) query removal. Query initialization starts when a new path query is registered to the path monitoring server (①). Then the server computes the answer path based on the current state of the network (②) and returns it to the user (③). Meanwhile,

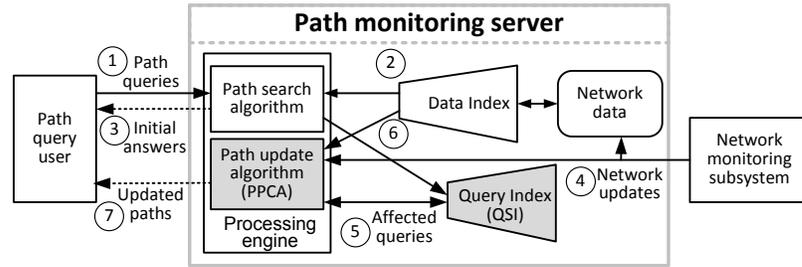


Figure 3.3. The PathMon framework

the query index is updated to accommodate this newly registered continuous query. Path update takes place when network updates from network monitoring subsystem are received (④). While network status and the data index are updated, the path update algorithm is triggered to detect affected queries from the query index (⑤) and accesses the data index for updated paths (⑥). Then, updated path is delivered to the user (⑦). In case that a registered query is not affected by the change, no path update is needed and the path previously delivered to the user is asserted to be valid. Query removal occurs when a query is terminated by the user. Terminated queries are removed from the query index.

PathMon: Data Structures and Implementation

4.1 Path Invalidation

In Chapter 3, we have analyzed the properties of min-cost path queries, and presented the concept of *query scope*. In this section we study different query scope representations, and present the design of query scope index (QSI) and describe how to efficiently look up affected queries via QSI.

4.1.1 Query Scope Representation

According to Definition 2, a min-cost path query $q(s, d)$ has a query scope $QS_t(q)$ that contains all nodes in the spanning tree $T_t(q)$. In our study, we consider three alternative representations for query scopes in terms of representation precision and storage overhead, as described in the following:

1. **All-node materialization.** This approach records exactly every node inside a query scope. It is the most precise one (the entire query scope is covered, and no irrelevant nodes are involved), but consumes the most storage space.
2. **Anchor node materialization.** Observing that nearby nodes (in terms of network proximity) are very likely to be accessed together by the path

search algorithm, we may pick only one node among some nearby nodes as the “anchor” to refer a query (e.g. the query identifier) and the distance from the anchor to the destination. Other nodes can obtain the information by traversing neighboring subnetworks to visit nearby anchors, and then decide whether it is in a query scope. With only the anchor nodes recorded, the anchor node materialization reduces the storage cost, but upon checking the association between nodes and query scopes, this approach incurs additional computational and I/O cost due to extra network traversals.

3. **Grid-based materialization.** Rather than individual nodes, we use affecting edges to determine the bounding area for a query scope. Here, we first partitioning the areas covered by the road network into a grid. A grid cell is preserved as a part of the query scope if it covers any affecting edges. Though this representation is simple and takes less storage than the above two, this geo-spatial partition suffers from false hits since some edges within affected grid cells may not necessarily be affecting edges.

In PathMon, we adopt the anchor node materialization for the query scope as the system default, because it generalizes the all-node materialization and fully exploits the network connectivity. By carefully selecting anchor nodes, the storage overhead for keeping the query scopes is effectively reduced while not incurring significant extra traversal cost. We also implemented the grid-based materialization for comparison purpose.

4.1.2 Query Scope Index

To alleviate the affected path detection, we store the query scopes identified based on Definition 2 in the *query scope index* (in short, the *QSI*). Let $Q = \{q\}$ denote a set of continuous min-cost path queries registered in the PathMon system. QSI provides a reversed one-to-many mapping between queries and nodes covered by their query scopes, i.e., $m : N \rightarrow Q$. This mapping facilitates a quick lookup of affected queries from end nodes of updated edges.

In the all-node based query scope materialization, this mapping can be implemented by simply marking every node that is visited in a search process with the query ID. However, due to the high storage space overhead and update cost, we

further consider the anchor node materialization. Since the path search algorithm generates a min-cost path spanning tree covering all nodes inside the query scope, the anchor node assignment can be naturally integrated with the search process with a minor computational overhead. The basic idea is to evenly spread a query (and its registered information) to nodes within its query scope based on their min-cost path to the destination node. A system parameter α is a preset number of maximum traversal steps between a new anchor node and its preceding anchor node on the same path of the spanning tree, so that it can control the number of anchor nodes. Thus, a node is assigned as an anchor node if the number of nodes on its min-cost path to the destination is a multiply of α . If α is small, more covered nodes will be assigned as anchor nodes, which improves search efficiency by consuming more storage. Otherwise, if α is large, storage overhead is reduced with a deterioration of search efficiency.

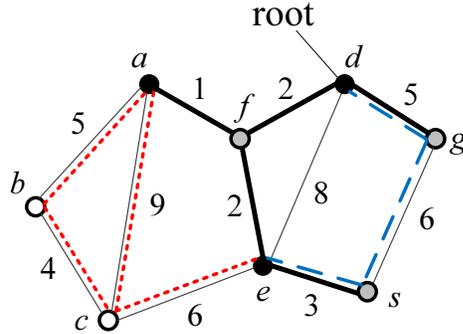


Figure 4.1. Illustration of partial materialization

Fig 4.1 illustrates the anchor node selection and affected path search process. Here the min-cost path from s to d , along with the min-cost path spanning tree rooted at d , are shown in bold lines. Assume $\alpha = 2$, and the dark nodes are selected as anchor nodes. The nodes in grey are non-anchors but covered by the query scope, while other white nodes lay outside the query scope (and their distance to d is greater than $\|s, d\|_{t_0}$). When an edge cost $C_{\langle c, b \rangle}$ changes, an expansion is performed from both c and b to visit the nodes that are reachable within $\alpha - 1$ hops (indicated by the dotted, red lines). During the expansion from c , anchor nodes e and a are found, and $\|c, d\|$ can be calculated by $\|c, d\| = \min\{\|c, e\| + \|e, d\|, \|c, a\| + \|a, d\|\}$. Similarly, for b , we have $\|b, d\| = \|b, a\| + \|a, d\|$. The calculation results show that $\|b, d\| > \|s, d\|$ and $\|c, d\| > \|s, d\|$, so we know that the edge $\langle c, b \rangle$ has no end

node inside the query scope, and this edge update could not possibly change the min-cost path. On the other hand, consider another update of edge $\langle s, g \rangle$, since d is an anchor node visited in the expansion of $\alpha - 1$ hops from g (indicated by the dashed, blue lines), $\|g, d\| = \|g, d\| + \|d, d\| < \|s, d\|$, we know that the edge $\langle s, g \rangle$ has at least one end node covered by the query scope. The query $q(s, d)$ is then identified as an affected query.

In our implementation of PathMon server, QSI is realized by two sets of lists $\{QSI.qlist(n) | n \in N\}$ and $\{QSI.anchors(q) | q \in Q\}$, where $QSI.qlist(n)$ is a list of query IDs for an anchor node n and $QSI.anchors(q)$ is a list of anchor node IDs for a query q . Figure 4.2 illustrates the two basic data structures of QSI.

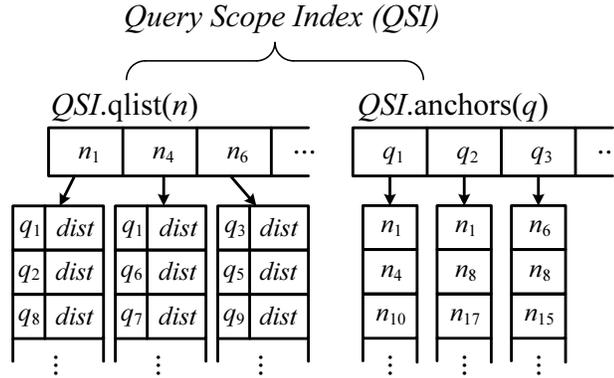


Figure 4.2. The query scope index structure

With $QSI.qlist(n)$ we can easily retrieve the queries that are mapped with a node n , and with $QSI.anchors(q)$ we can quickly obtain all anchor nodes that maintain information regarding to a query q . QSI is updated upon registration of a new query, as listed in Algorithm 1. This is integrated with the path search algorithm, which iteratively examines nodes in non-descending order of the network distance to the destination d of the new query q . The search starts from the destination d . A node n positioned at x hops away (along the min-cost path) from d is an anchor node if $x \bmod \alpha = 0$ (line 6). The position of d is 0 (thus d is always an anchor node). The algorithm for finding queries affected by an edge $\langle n, n' \rangle$ is shown in Algorithm 2. It performs a traversal started from the end nodes of updated edges for at most $\alpha - 1$ nodes apart, while collecting queries that are associated with the traversed anchor nodes.

When a query q^* is terminated, we first obtain all anchor nodes of q^* from

Algorithm 1 RegisterQuery($QSI, q(s, d), t_0$)

```

1: Visited node set  $V \leftarrow \emptyset$ 
2: heap  $H.insert(d, \perp, 0, 0)$ 
3: while  $H \neq \emptyset$  do
4:    $(n, pre_n, c_n, x_n) \leftarrow H.getMin()$ 
5:   if  $n \in V$  then
6:     break
7:   end if
8:   if  $x_n \bmod \alpha = 0$  then
9:      $QSI.qlist(n).add(q, c_n)$  {assign  $n$  as an anchor node}
10:     $QSI.anchors(q).add(n)$  {map  $n$  with  $q$ }
11:   end if
12:    $V.insert(n)$   $\{||n, d|| = x_n\}$ 
13:   if  $n = s$  then break
14:   for each  $n'$  such that  $\langle n', n \rangle \in E$  do
15:     if  $n' \notin V$  then
16:        $H.insert((n', n, c_n + C_{\langle n', n \rangle, t_0}, x_n + 1))$ 
17:     end if
18:   end for
19: end while

```

Algorithm 2 SearchQuery($QSI, \langle n, n' \rangle$)

```

1:  $SQ \leftarrow \emptyset$  { $SQ$  is the set of affected queries}
2: Expand from  $n$  and  $n'$  for  $\alpha - 1$  hops using Dijkstra's algorithm
3: for each visited node  $m$  do
4:   if  $QSI.qlist(m) \neq \emptyset$  then
5:     for each  $q \in QSI.qlist(m)$  do
6:       if  $||n, m||_{t_1} + ||m, d_q||_{t_0} \leq ||s_q, d_q||_{t_0}$  or  $||n', m||_{t_1} + ||m, d_q||_{t_0} \leq ||s_q, d_q||_{t_0}$ 
7:         then
8:            $SQ.insert(q)$ 
9:         end if
10:      end for
11:   end if
12: end for
13: return  $SQ$ 

```

$QSI.anchors(q^*)$, and for each anchor node m , remove the entry of q^* from $QSI.qlist(m)$. After that, the entire list of $QSI.anchors(q^*)$ is deleted.

4.2 Partial Path Re-computation

In the discussions earlier, we have already presented the main idea of the path update algorithm, i.e. path update can start at a certain processing state of the previous search and safely terminate at another processing state existing in the previous search after incorporating all updates in the query scope range. As such, part of the path re-computation mandated for affected path queries can be eliminated. This section mainly focuses on the issues of implementing and optimizing the path re-computation algorithm.

The key idea of our partial path re-computation is to reuse the processing states explored in previous search processes. For this purpose, in each path search process, we select a set of *checkpoint nodes* at which the processing states are stored. Since it is expensive to maintain the processing states V^n, H^n for all visited nodes, we select a few nodes as the checkpoint nodes. Below we present five checkpoint node selection strategies. The performances of these strategies will be studied in Chapter 5.

1. Random selection (RAND). Randomly pick a given number of nodes as the checkpoint.
2. Associated with Anchor nodes (ANCH). Each anchor node also is responsible for carrying its corresponding processing state. Thus the restart point can be found along with the path invalidation process.
3. Fixed step interval (FSI). Store a processing state every fixed number of steps in the search. Thus the restarting processing state is guaranteed to be within a pre-set number of steps from visiting the updated edge.
4. Fixed cost interval (FCI). The checkpoint nodes are evenly distributed according to the distance to the destination.
5. Adaptive-to-updates (ATU). Initially, use one of the above strategies to select the checkpoint nodes. As the network updates happen, dynamically add checkpoints near the frequent-updated areas and remove the checkpoints where updates rarely occur. If the update distribution is skewed in the network, i.e. a few “hot-spots” are constantly updated while other places remain almost static, this approach is expected to achieve better performance.

Accordingly, we extend QSI to QSIU (namely, *QSI for Update*) in order to facilitate a quick lookup of preserved state information when affected queries are detected. Thus, while QSI is mainly developed for path invalidation, QSIU also supports efficient path updates by the Partial Path Computation Algorithm (PPCA). In addition to QSI, a *processing state repository* $QSIU.procs(q, n)$ is introduced. Given a query q and a node n , $QSIU.procs(q, n)$ returns the processing state in the path search process for q at node n . If n is not a checkpoint node of q , a NULL value is returned. Meanwhile, we maintain a list $QSIU.check(q)$ to record all the checkpoint nodes for the query q to support query removal.

A processing state $S(n)$ at a checkpoint node n consists of 1) a visited node set V^n , and 2) a priority queue H^n . Since there are multiple queries registered in QSIU, we use a subscript q to indicate that $S_q(n)$ is the processing state of query q at node n , i.e. $S_q(n) = (V_q^n, H_q^n)$. $S_q(n)$ is available in the path search process of query q . When a node is selected as a checkpoint node, we maintain the visited node set and the priority queue of current state by adding $S_q(n)$ into $QSIU.procs(q, n)$.

QSIU is not only used for identifying the affected paths, but also collecting important information in the search process for use in PPCA, i.e., the *affecting range*. For a query $q(s, d)$, the affecting range is the range of the distances from the updated edges located within the query scope to the destination. Let the affecting range be denoted as $[X_q, Y_q]$, $0 \leq X_q \leq Y_q$. All the updated edges $\langle n, n' \rangle$ in the network satisfy that $X_q \leq \|n, d\| \leq Y_q \vee X_q \leq \|n', d\| \leq Y_q$, or $\|n, d\| > \|s, d\| \vee \|n', d\| > \|s, d\|$. X_q provides the initialization information to PPCA, and Y_q is used to check whether all affecting updated edges have been visited and thus the re-computation algorithm may terminate. Particularly, in our algorithm we set X_q equal to the nearest distance from an updated edge to the destination, and the upper bound Y_q of the possible (updated) affecting range is determined as

$$Y_q = \max_{\Delta C_{\langle n, n' \rangle} \neq 0} \{ \max\{\|n, d\|, \|n', d\|\} \} + \sum_{\Delta C_{\langle n, n' \rangle} > 0} \Delta C_{\langle n, n' \rangle}$$

in which each $\langle n, n' \rangle$ is an affecting edge for q . Here, we conservatively incorporate all edge cost increments into the upper bound estimation. We will discuss the

usage of the affecting range $[X_q, Y_q]$ and the starting node later. Algorithm 3 lists the query identification and re-computation information collecting process with QSIU. Given a set of network updates, the algorithm returns all affected query IDs and their corresponding $[X_q, Y_q]$.

Algorithm 3 SearchQuery($QSIU, t_1$)

```

1:  $SQ \leftarrow \emptyset$ 
2: for all  $q, X_q \leftarrow \infty, Y_q \leftarrow 0, \Delta Y_q \leftarrow 0$ 
3: for all  $\langle n, n' \rangle$  that is updated do
4:    $\Delta C \leftarrow C_{\langle n, n' \rangle, t_1} - C_{\langle n, n' \rangle, t_0}$ 
5:   Expand  $\langle n, n' \rangle$  for  $\alpha - 1$  hops using Dijkstra's algorithm
6:   for each anchor node  $a$  visited in the expansion do
7:     for each  $q \in QSIU.qlist(a)$  do
8:        $c_n \leftarrow ||n, a|| + ||a, d_q||$ 
9:        $c_{n'} \leftarrow ||n', a|| + ||a, d_q||$ 
10:      if  $c_n \leq ||s_q, d_q||_{t_0}$  or  $c_{n'} \leq ||s_q, d_q||_{t_0}$  then
11:         $SQ.insert(q)$ 
12:        if  $\Delta C > 0$  then  $\Delta Y_q \leftarrow \Delta C + \Delta Y_q$ 
13:        if  $c_n < X_q$  or  $c_{n'} < X_q$  then  $X_q \leftarrow \min\{c_n, c_{n'}\}$ 
14:        if  $c_n > Y_q$  or  $c_{n'} > Y_q$  then  $Y_q \leftarrow \max\{c_n, c_{n'}\}$ 
15:      end if
16:    end for
17:  end for
18: end for
19: for all  $q$  do  $Y_q \leftarrow Y_q + \Delta Y_q$ 
20: return  $SQ, \{[X_q, Y_q]\}$ 

```

It is noticeable that in Algorithm 3, the cost c_n computed from $||n, a|| + ||a, d_q||$ may deviate from the actual $||n, d_q||$ at time t_1 (after the network change) if a is not on the min-cost path from n to d_q or $||a, d_q||$ itself is updated. However, the correctness of the algorithm can still be guaranteed because the smallest $||n, d||$ value must be consistent with the current network state. This is quite clear with Lemma 2 and 3.

PPCA initializes the re-computation with a “stable” processing state that does not change upon edge updates, in order to eliminate the steps of path computation from destination node (the very beginning state) to this state. According to Lemma 3, this initial state is positioned in maximum number of steps from the destination node, without being affected by edge changes. Thereafter, the algorithm traverses the changed portion of road network, factoring in the updates to obtain

new processing states for affected nodes, until either of the following termination conditions is satisfied: (1) *early termination* - after all the changed edges within the query scope are visited, there is a processing state at a certain node remaining the same as it was in the previous search algorithm; (2) *normal termination* - the source node s_t is visited. If the algorithm meets condition (1), based on Lemma 4, the rest of states in the search process remains unchanged and thus can be skipped. If condition (2) is meet, the algorithm terminates naturally.

Fig 4.3 gives an illustration of PPCA using Dijkstra’s algorithm. The lines marked with a “▲” represent the updated edges. Other black lines are the edges traversed in a re-computation algorithm. The re-computation starts from a checkpoint node m that holds the latest processing state with smaller distance to d than any updated edge (so that $X_q = ||m, d||$), and terminates at another checkpoint node m' where $||m', d|| > Y_q$.

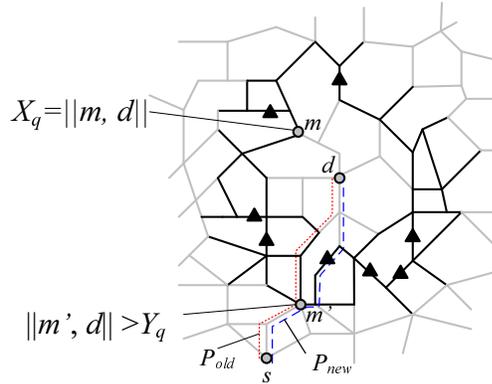


Figure 4.3. Illustration of PPCA

With the information collected in QSIU search algorithm, the re-computation algorithm PPCA is performed in four steps, elaborated below. Algorithm 4 lists the pseudo code for this update process.

1. Given the lower bound of affecting range X_q , find the latest stable processing state $S_q(m)$ ($||m, d|| < X_q$), initialize the search algorithm with processing state $S_q(m)$ (lines 1-7).
2. Re-run the path search algorithm from the initialized state to update the min-cost path. In this step, update of the QSIU is performed seamlessly. If a node n is selected as a checkpoint node, the algorithm associates the new

Algorithm 4 UpdatePath($QSIU, q(s_t, d), s_t, X_q, Y_q$)

```

1:  $Dist_{max} \leftarrow \infty, S_{start} \leftarrow \text{NULL}$ 
2: for all  $m \in QSIU.check(q)$  do
3:   if  $\|m, d\| < X_q$  AND  $\|m, d\| > Dist_{max}$  then
4:      $S_{start} \leftarrow QSIU.procs(q, m)$ 
5:      $Dist_{max} \leftarrow \|m, d\|$ 
6:   end if
7: end for
8: Visited node set  $V \leftarrow S_{start}.V$ , priority queue  $H \leftarrow S_{start}.H$ 
9: while  $H \neq \emptyset$  do
10:   $(n, pre_n, c_n, x_n) \leftarrow H.getMin()$ 
11:  if  $n \notin V$  then
12:     $V.insert(n, pre_n, c_n)$ 
13:    if  $x_n \bmod \alpha = 0$  then
14:       $QSIU.anchors(q).add(n)$ 
15:       $QSIU.qlist(n).add(q)$ 
16:    end if
17:    if  $n$  is selected as a checkpoint then
18:       $QSIU.procs(q, n).insert(V, H)$  {add a new processing state}
19:       $QSIU.check(q).add(n)$ 
20:    else if  $QSIU.procs(q, n) \neq \text{NULL}$  then { $n$  is an expired checkpoint}
21:       $QSIU.procs.remove(q, n)$  {remove an obsolete processing state}
22:       $QSIU.check(q).remove(n)$ 
23:    end if
24:    if  $(n = s_t)$  OR  $(c_n > Y_q$  AND  $S_q(n) \equiv QSIU.procs(q, n))$  then
25:      break
26:    end if
27:    for each  $n'$  such that  $\langle n', n \rangle \in E$  do
28:       $H.insert((n', n, c_n + C_{\langle n', n \rangle}, x_n + 1))$ 
29:    end for
30:  end if
31: end while
32: Trace back from  $s_t$  to construct the path  $P(s_t, d)$ 

```

processing state $S'_q(n)$ to n by inserting new entries (if n was not a checkpoint node of q) or updating existing entries (if n was originally a checkpoint node of q) in $QSIU.procs(n)$ (lines 17-19). If a previously selected checkpoint n is no longer a checkpoint node in the new search, its associated processing state expires and is removed from $QSIU.procs(n)$ (lines 20-22). The anchor nodes of q are also updated with similar operations (lines 13-16).

3. If s is visited, the search terminates anyway. Otherwise, after all updated

edges affecting q have been explored (we may check this by comparing the current expansion distance with Y_q), the path re-computation terminates if there exists an anchor node m' such that $S_q(m') \equiv S'_q(m')$ (m' is then called the termination node) (lines 24-25).

4. Finally, starting from s , trace back the precedences to construct the entire path (line 30).

While retaining the same methodology, an *incremental representation* of processing states is adopted in our PPCA implementation to further reduce the storage overhead, i.e. we differentiate the current processing state (V^n, H^n) from what has been stored at the last checkpoint node (V^m, H^m) , and only keep the difference $(\Delta V^n, \Delta H^n)$. When (V^n, H^n) is to be retrieved from the processing state repository of QSIU, the incremental processing states $\{(\Delta V^m, \Delta H^m)\}$ corresponding to all the checkpoint nodes m visited before n are retrieved and then merged with $(\Delta V^n, \Delta H^n)$ to recover the complete (V^n, H^n) . If the node n currently being visited is selected as a checkpoint node, we differentiate the current processing state (V^n, H^n) from that has been stored at the last checkpoint node (V^m, H^m) , and only keep the difference, i.e. $(\Delta V^n, \Delta H^n)$. When (V^n, H^n) is to be retrieved from the processing state repository of QSIU, first all the incremental processing states $\{(\Delta V^m, \Delta H^m)\}$, where $\{m\}$ is the set of checkpoint nodes visited before n , are retrieved and then they are merged with $(\Delta V^n, \Delta H^n)$ to recover the complete (V^n, H^n) .

Performance Evaluation

In this chapter, we evaluate the proposed QSI and PPCA, i.e., the core components of PathMon, via an extensive set of experiments by simulation. Besides, we compare our approaches with the most related approach, namely, EBM [31], and a naive approach that reevaluates all path queries upon any network change. In what follows, we first describe the experiment setups. Then, we examine the effectiveness of QSI and different checkpoint selection heuristics. Finally, we evaluate the overall performance of PathMon and compare it with EBM and the naive re-computing approach.

5.1 Experiment Setup

We implemented the aforementioned algorithms with GNU C++. For QSI (denoted by QSI), we implemented both anchor node materialization and grid based materialization and compare their effectiveness in terms of the ratio of identified affected queries. For PPCA and QSIU (denoted by QSIU+PPCA), we only used anchor node materialization. For anchor node materialization, we defaulted α to 10. For EBM (denoted by EBM), an elliptic affecting area is formulated for every min-cost path. Here, we assume V_{max} of EBM to be 1.5. The naive approach (labeled as Naive) reevaluates all path queries when updates occur in the network. We use a grid file with 50×50 grid cells to support grid based materialization and as a grid index used by EBM. Here, except for PPCA that recomputes path partially, all others, including QSI, EBM and Naive, use Dijkstra’s shortest path algorithm

and/or A* algorithm to compute paths. Here, A* is applicable for situations where edge costs are related to geometrical distance.

We ran the simulation for each setting for 30 times on Linux 2.6.9 servers with Intel Xeon 3.2GHz CPU and 4GB RAM and report the average results. In each setting, we issued 1,000 continuous min-cost path queries. Source nodes of the queries are randomly picked, whereas destination nodes are selected about 0.15 through 0.35 of the network diameter away from respective source nodes. We summarize all the experiment parameters in Table 5.1.

Parameter	Value	Default
Network	CA (21,048 nodes, 21,693 edges), NY (46,254 nodes, 63,552 edges), WM (61,333 nodes, 80,589 edges)	CA
No. of queries	1,000	1,000
No. of changed roads ($ E_u $)	0.03%, 0.05%, 0.08%, 0.1%	0.05%
Length of paths (L_p)	0.15, 0.20, 0.25, 0.30, 0.35 of network diameter	0.25
Path search alg.	Dijkstra's, A*	Dijkstra's

Table 5.1. The evaluation parameters

5.2 Effectiveness of Query Scope Representations

We expect the representation of query scopes to have an impact on performance of path monitoring system. In the naive approach, the query scope contains the entire network, while in EBM, the query scope is an ellipse approximated by grid cells. Our QSI and QSIU+PPCA are based on the “network connectivity” and thus consist of nodes covered by the query scope. In the first experiment, we compare the path invalidation ability of these approaches in terms of number of affected paths identified.

The result is plot in Figure 5.1. As shown, Naive incurs all path recomputed (i.e., 100%) due to blind path reevaluation. EBM saves some path re-computation for those edges not covered by the elliptic area. QSI, exploiting query scopes to precisely determine non-affecting edges, has a significantly lower re-computation

ratio. QSIU+PPCA uses the same path invalidation method as QSI so they have the identical re-computation ratio. Meanwhile, the grid-based representation of QSI (denoted as Grid) has less false hits than EBM because it is more precise. Comparing with QSI, however, some redundant re-computations still exist in Grid.

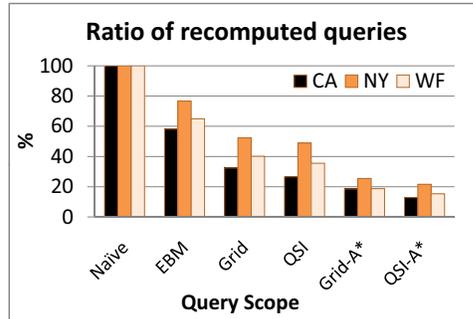


Figure 5.1. Re-computation ratio

In the scenario where certain geometric assumptions are applicable in the path cost estimation, path search algorithms utilizing the estimation heuristics can achieve a much smaller query scope size. For example, the QSI-A* and Grid-A* curves represent the ratio of recomputed queries identified by the query scope built based on the A* algorithm. While Grid-A* incurs more re-computations than QSI-A*, they both have quite low ratio compared with the ones based on Dijkstra’s algorithm. This is because, in the heuristic-based algorithms, both the connectivity and physical constraints of the road network are considered, resulting in effective elimination of the false hits.

5.3 Effectiveness of Checkpoint Selection

In QSIU+PPCA, five checkpoint placement strategies have been presented, namely, random (RAND), anchor-node-association (ANCH), fixed-step-interval (FSI), fixed-cost-interval (FCI) and adaptive-to-update (ATU). Here, we study the effectiveness of these strategies.

In this set of experiments, the total number of checkpoints selected by the strategies is set to $|N|/100$ where $|N|$ is the number of nodes in the network. The results are plotted in Fig 5.2. In terms of I/O cost, RAND, FSI and FCI incur almost the same number of page accesses, because the checkpoint selected by these

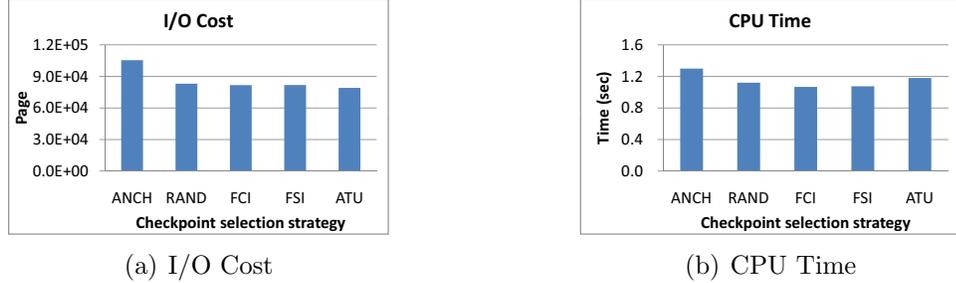


Figure 5.2. Checkpoint selection strategy comparison (CA)

strategies tend to be evenly distributed in the network. The running times of these three strategies are also very close. Meanwhile, ATU achieves slightly better I/O efficiency due to adaptively selecting checkpoints that are close to the places where updates happen frequently. Nevertheless, due to the extra computational costs needed to re-allocate checkpoints, ATU takes longer CPU time to perform the re-computation. Last, ANCH incurs both the most I/O cost and CPU time among all five strategies, indicating that the anchor node assignment is not suitable for selecting checkpoints.

5.4 Path Monitoring Performance

Next, we evaluate PathMon and compare its performance against other techniques in terms of CPU time and I/O costs, the two common metrics considered for system performance. CPU time measures the time the system spends to update path queries. I/O cost measures the number of disk pages accessed to detect/update affected path queries.

Impact of path length. The current min-cost path length directly affects the size of the query scope. Here, we evaluate its impact on the performance of examined techniques by varying the path length from 0.15 to 0.35 with a step of 0.05 of the network diameter. The results are reported in Figure 5.3. As the path length increases, the ratio of recomputed queries, processing time and I/O cost of all approaches increase accordingly. Due to effective elimination of redundant path update, QSIU+PPCA reduces both I/O cost and CPU time as shown in Figure 5.3(a) and Figure 5.3(b), respectively. Also, we can see that Naive is the worst, while QSIU+PPCA significantly outperforms EBM. With fewer paths recomputed, QSI

performs better than EBM but still incurs about 50% more I/O cost and CPU time than QSIU+PPCA.

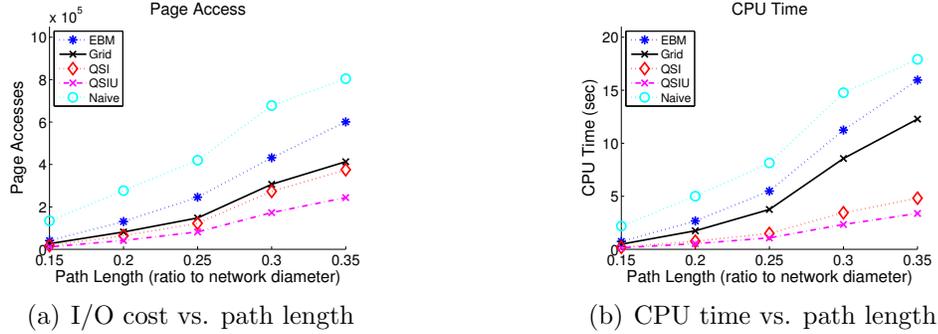


Figure 5.3. Performance under various path length

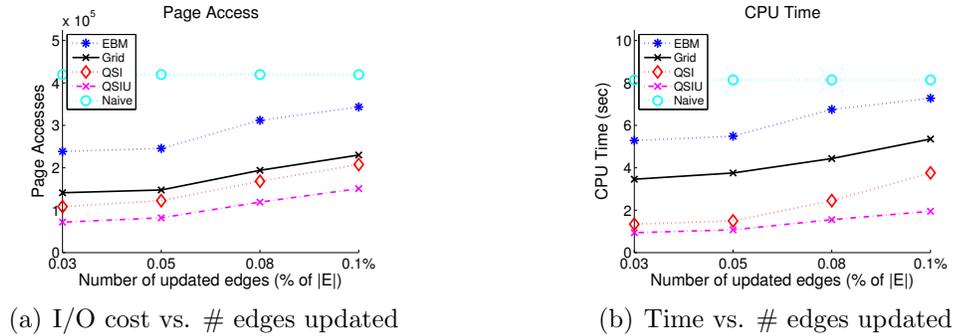


Figure 5.4. Performance under various # of updated edges

Impact of number of updated edges. If a network is highly dynamic, a large number of edges would be updated simultaneously. In this experiment, we model the degree of network dynamics by varying the number of updated edges up to 0.01% of the total number of edges. Please note that the query processing time is very short for our proposed algorithm (i.e. a few seconds for QSIU+PPCA), so the approach is capable to deal with high update frequency. Again, I/O cost and CPU time are significantly saved by QSIU+PPCA as shown in Figure 5.4(a) and Figure 5.4(b), respectively. On the other hand, this experiment shows a trend that both performance of EBM and QSIU+PPCA degrade as the number of updated edges increase. Naive will eventually prevail, which indicates a room for future research in PathMon.

Impact of path search algorithm. The path search algorithm can significantly affect the query scope, and potentially lead to more efficient path monitoring.

Here we show that, besides Dijkstra’s algorithm, our proposed algorithms can be applied to other best-first path search algorithms (e.g., A* algorithm) and also achieve superior performance.

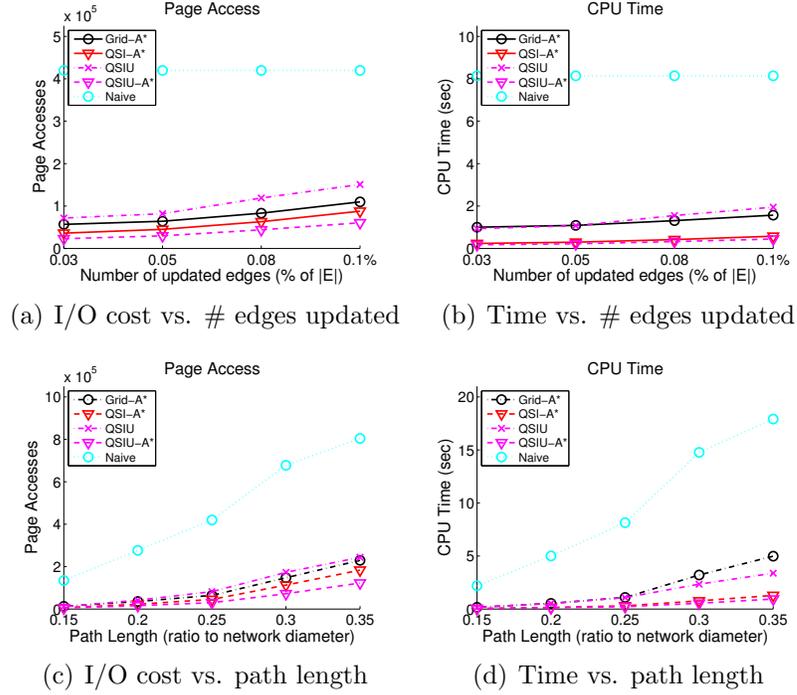


Figure 5.5. Performance under A* algorithm

We use Naive and QSIU as the baseline in our comparison. As shown in Fig 5.5, when A* algorithm is applied, both the I/O and time performance improve significantly for different query scope representations (i.e. the grid-based and the anchor node-based representation). Furthermore, we can see that the QSIU-A* approach outperforms QSI-A* and Grid-A*.

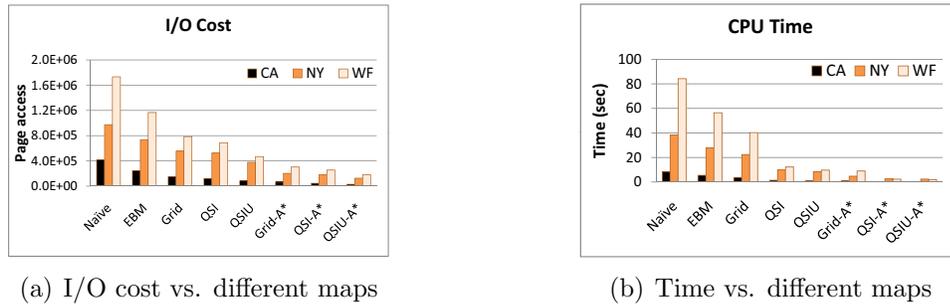


Figure 5.6. Performance for various cities ($L_p = 0.25D_N$, $|E_u| = 0.05\%|E|$)

Evaluation on various maps. To validate the efficiency of our approach, a set of experiments on various cities has been conducted. We report the results in Fig 5.6. For all algorithms, as the network size increases, both I/O cost and CPU time grow. QSIU is superior among all Dijkstra's algorithm based approaches, and QSIU-A* performs the best in the A* algorithm based approaches.

Conclusion

Path queries are essential features in today's location-based services, and they bring great convenience to people's daily lives. In this thesis, we dived into the minimum cost path monitoring query, which helps users to always stay updated about the route with lowest travel cost to the desired destination. The min-cost path monitoring queries, if built into existing location-based services, will be come powerful tool to facilitate path search. However, processing a large number of path monitoring queries is not trival; it presents several challenges. In this thesis, we proposed novel and efficient approaches to address the technical challenges.

We studied two research issues of the problem, namely, path invalidation and path update, and developed novel techniques based on our problem analysis and our understanding of the min-cost path search properties. Based on the result of our analysis, we introduced an index structure called query scope index (QSI) to record query scopes, which consist of nodes in the min-cost path spanning trees and in turns cover all affecting edges. QSI is efficient to identify affected path queries. Also, we devised a partial path re-computation algorithm (PPCA) to partially recompute a new min-cost path when the old min-cost path becomes invalid, by resuming a search at certain processing states. To facilitate PPCA, another index called QSIU that extends QSI is devised to maintain processing states. Our PathMon system seamlessly integrates QSI, QSIU and PPCA. PathMon is a centralized system to support monitoring of massive continuous min-cost path queries. Upon network changes, the PathMon server identifies the queries affected by the changes from QSI/QSIU, and quickly recomputes the new min-cost paths using

PPCA. Through comprehensive experiments, we examined the effectiveness of different designs for QSI and PPCA. Additionally, we studied the impact of expected query path length, network update rates and network size to the performance of the PathMon system.

Bibliography

- [1] LEE., K. C. K., W.-C. LEE, and B. ZHENG (2009) “Fast object search on road networks,” in *Proceedings of EDBT’09*, pp. 1018–1029.
- [2] TSUNG CHANG, K. (2007) *Introduction to Geographic Information Systems (4th Edition)*, McGraw Hill.
- [3] GUTTMAN, A. (1984) “R-trees: a dynamic index structure for spatial searching,” in *Proceedings of the 1984 ACM SIGMOD international conference on Management of data (SIGMOD’84)*, Boston, Massachusetts, pp. 47–57.
- [4] BECKMANN, N., H.-P. KRIEGEL, R. SCHNEIDER, and B. SEEGER (1990) “The R*-tree: an efficient and robust access method for points and rectangles,” *SIGMOD Record*, **19**(2), pp. 322–331.
- [5] CHENG, R., D. V. KALASHNIKOV, and S. PRABHAKAR (2004) “Querying Imprecise Data in Moving Object Environments,” *IEEE Transaction on Knowledge and Data Engineering*, **16**(9), pp. 1112–1127.
- [6] PAPADIAS, D., J. ZHANG, N. MAMOULIS, and Y. TAO (2003) “Query Processing in Spatial Network Databases,” in *Proceedings of 29th International Conference on Very Large Data Bases (VLDB), Berlin, Germany, Sep. 9-12*, pp. 802–813.
- [7] WALTER, V. and D. FRITSCH (1994) “Modelling and Storage of Road Network Data,” in *Proceedings of Advanced Geographic Data Modelling 1994 (AGDM), Delft, The Netherlands*.
- [8] KLUNDER, G. A. and H. N. POST (2006) “The Shortest Path Problem on Large-Scale Real-Road Networks,” *Networks*, **48**(4), pp. 182–194.
- [9] JENSEN, C. S., J. KOLÁRVR, T. B. PEDERSEN, and I. TIMKO (2003) “Nearest Neighbor Queries in Road Networks,” in *Proceedings of the 11th ACM International Symposium on Geographic Information Systems (ACM-GIS), New Orleans, Louisiana, USA*, pp. 1–8.

- [10] SANKARANARAYANAN, J., H. ALBORZI, and H. SAMET (2006) “Enabling Query Processing on Spatial Networks,” in *Proceedings of the 22nd International Conference on Data Engineering (ICDE), Atlanta, GA, USA*, p. 163.
- [11] JUNG, S. and S. PRAMANIK (2002) “An Efficient Path Computation Model for Hierarchically Structured Topographical Road Maps,” *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, **14**(5), pp. 1029–1046.
- [12] HU, H., D. L. LEE, and J. XU (2006) “Fast Nearest Neighbor Search on Road Networks,” in *Proceedings of the International Conference on Extending Database Technology (EDBT)*, pp. 186–203.
- [13] ORDA, A., R. ROM, and M. SIDI (1993) “Minimum Delay Routing in Stochastic Networks,” *IEEE/ACM Transaction on Networking*, **1**(2).
- [14] KANOULAS, E., Y. DU, T. XIA, and D. ZHANG (2006) “Finding Fastest Paths on A Road Network with Speed Patterns,” in *Proceedings of the 22nd International Conference on Data Engineering (ICDE), Atlanta, GA, USA*.
- [15] GONZALEZ, H., J. HAN, X. LI, M. MYSLINSKA, and J. P. SONDAG (2007) “Adaptive Fastest Path Computation on a Road Network: A Traffic Mining Approach,” in *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB), University of Vienna, Austria, Sep. 23-27*, pp. 794–805.
- [16] DING, B., J. X. YU, and L. QIN (2008) “Finding Time-Dependent Shortest Paths over Large Graphs,” in *Proceedings of the International Conference on Extending Database Technology (EDBT)*, Nantes, France, pp. 205–216.
- [17] MCQUILLAN, J. M., I. RICHER, and E. C. ROSEN (1980) “The new routing algorithm for the ARPANet,” *IEEE Transaction on Communications*, **28**(5), pp. 711–719.
- [18] SEGALI, A. (1983) “Distributed network protocols,” *IEEE Transaction on Information Theory*, **29**.
- [19] DIJKSTRA, E. (1959) “A Note on Two Problems in Connection with Graphs,” *Numerical Mathematics*.
- [20] DECHTER, R. and J. PEARL (1985) “Generalized Best-First Search Strategies and the Optimality of A*,” *Journal of ACM*, **32**(3), pp. 505–536.
- [21] DEO, N. and C. PANG (1984) “Shortest path algorithms: Taxonomy and Annotations,” *Networks*, **14**.
- [22] LUBY, M. and P. L. RAGDE (1989) “A Bidirectional Shortest-path Algorithm with Good Average-case Behavior,” *Algorithmica*, **4**, pp. 551–567.

- [23] BELLMAN, R. (1958) “On a Routing Problem,” *Quarterly of Applied Mathematics*, **16**(1), pp. 87–90.
- [24] BERTSEKAS, D. P. (1993) “A Simple and Fast Label Correcting Algorithm for Shortest Paths,” *Networks*, **23**, pp. 703–709.
- [25] FLOYD, R. W. (1962) “Algorithm 97: Shortest Path,” *Communications of the ACM*, **5**(6).
- [26] SEIDEL, R. (1992) “On the All-Pairs-Shortest-Path Problem,” in *Proceedings of the 24th annual ACM symposium on Theory of Computing (STOC)*, Victoria, British Columbia, Canada, pp. 745–749.
- [27] JING, N., Y.-W. HUANG, and E. A. RUNDENSTEINER (1998) “Hierarchical Encoded Path Views for Path Query Processing: An Optimal Model and Its Performance Evaluation,” *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, **10**(3), pp. 409–432.
- [28] MÖHRING, R. H., H. SCHILLING, B. SCHÜTZ, D. WAGNER, and T. WILLHALM (2006) “Partitioning Graphs to Speedup Dijkstras Algorithm,” in *ACM Journal of Experimental Algorithmics*, vol. 11, pp. 1–29.
- [29] FRIGIONI, D., A. MARCHETTI-SPACCAMELA, and U. NANNI (1996) “Fully Dynamic Output Bounded Single Source Shortest Path Problem,” in *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, pp. 212–221.
- [30] NARVÁEZ, P., K.-Y. SIU, and H.-Y. TZENG (2000) “New Dynamic Algorithms for Shortest Path Tree Computation,” *IEEE/ACM Transaction on Networking*, **8**(6), pp. 734–746.
- [31] LEE, C.-C., Y.-H. WU, and A. L. P. CHEN (2007) “Continuous Evaluation of Fastest Path Queries on Road Networks,” in *Proceedings of the 10th International Symposium on Advances in Spatial and Temporal Databases (SSTD)*, Boston, MA, USA, July 16-18, pp. 20–37.
- [32] CARRAWAY, R. L., T. L. MORIN, and H. MOSKOWITZ (1990) “Generalized Dynamic Programming for Multicriteria Optimization,” *European Journal of Operational Research*, **44**, pp. 95–104.
- [33] MODESTI, P. and A. SCIOMACHEN (1998) “A utility measure for finding multiobjective shortest paths in urban multimodal transportation networks,” *European Journal of Operational Research*, **111**, pp. 495–508.
- [34] MARTINS, E. Q. V. (1984) “On a Multicriteria Shortest Path Problem,” *European Journal of Operational Research*, **16**, pp. 236–245.

- [35] BRUMBAUGH-SMITH, J. and D. SHIER (1989) “An empirical investigation of some bicriterion shortest path algorithms,” *European Journal of Operational Research*, **43**(2), pp. 216–224.
- [36] GUERRIERO1, F. and R. MUSMANNO (2001) “Label Correcting Methods to Solve Multicriteria Shortest Path Problems,” *Journal of Optimization Theory and Applications*, **111**, pp. 589–613.
- [37] CLIMACO, J. C. N. and E. Q. V. MARTINS (1982) “A bicriterion shortest path algorithm,” *European Journal of Operational Research*, **11**(4), pp. 399–404.
- [38] AZEVEDO, J. A. and E. Q. V. MARTINS (1991) “An algorithm for the multiobjective shortest path problem on acyclic networks,” *Investigacao Operacional*, **11**, pp. 52–69.
- [39] CURRENT, J. R., C. REVELLE, and J. L. COHON (1990) “An interactive approach to identify the best compromise solution for two objective shortest path problems,” *Computers and Operations Research*, **17**(2), pp. 187–198.
- [40] COUTINHO-RODRIGUES, J. M., J. C. N. CLÍMACO, and J. R. CURRENT (1999) “An interactive bi-objective shortest path approach: searching for unsupported nondominated solutions,” *Computers and Operations Research*, **26**(8), pp. 789–798.
- [41] GRANAT, J. and F. GUERRIERO (2003) “The interactive analysis of the multicriteria shortest path problem by the reference point method,” *European Journal of Operational Research*, **151**, pp. 103–118.
- [42] SHEKHAR, S. and D.-R. LIU (1997) “CCAM: A Connectivity-Clustered Access Method for Networks and Network Computations,” *Transaction on Knowledge and Data Engineering*, **9**(1), pp. 102–119.
- [43] CHOMICKI, J., P. GODFREY, J. GRYZ, and D. LIANG (2003) “Skyline with Presorting,” in *Proceedings of ICDE’03*.
- [44] JING, N., Y.-W. HUANG, and E. A. RUNDENSTEINER (1998) “Hierarchical Encoded Path Views for Path Query Processing: An Optimal Model and Its Performance Evaluation,” *IEEE Transaction on Knowledge and Data Engineering*, **10**(3), pp. 409–432.
- [45] “Feifei Li’s collection, <http://cs-people.bu.edu/lifeifei/tpq.html>,” .

Vita

Yuan Tian

Ms. Yuan Tian obtained her Bachelors degree in Electrical Engineering from Tsinghua University, Beijing, China in Summer 2007, and started her graduate study in the Pennsylvania State University in Fall 2007. In her research studies, Yuan mainly focused on three area: path queries in road networks, social network analysis, and distributed uncertain databases. Her research works on road networks have been published in SIG-SPATIAL GIS, and the papers on social network analysis are published in CIKM.

Yuan has always been passionate about developing real systems, and she joined Facebook Inc as a software engineer in Summer 2010. At Facebook, she continue to receive professional training in both software development and product thinking. She was an active contributor to several important projects at Facebook. Currently, Yuan is a senior engineer in an e-commerce analytics startup.