

Local Network Voronoi Diagrams

Sarana Nutanong[†], Egemen Tanin^{†‡}, Mohammed Eunus Ali[†], Lars Kulik^{†‡}

[†]Department of Computer Science and Software Engineering
University of Melbourne, Victoria, Australia

[‡]NICTA Victoria Research Laboratory, Australia
{sarana,egemen,eunus,lars}@csse.unimelb.edu.au

ABSTRACT

Continuous queries in road networks have gained significant research interests due to advances in GIS and mobile computing. Consider the following scenario: “A driver uses a networked GPS navigator to monitor five nearest gas stations in a road network.” The main challenge of processing such a moving query is how to efficiently monitor network distances of the k nearest and possible resultant objects. To enable result monitoring in real-time, researchers have devised techniques which utilize precomputed distances and results, e.g., the *network Voronoi diagram (NVD)*. However, the main drawback of preprocessing is that it requires access to all data objects and network nodes, which means that it is not suitable for large datasets in many real life situations. The best existing method to monitor k NN results without precomputation relies on executions of snapshot queries at network nodes encountered by the query point. This method results in repetitive distance evaluation over the same or similar sets of nodes. In this paper, we propose a method called the *local network Voronoi diagram (LNVD)* to compute query answers for a small area around the query point. As a result, our method requires neither precomputation nor distance evaluation at every intersection. According to our extensive analysis and experimental results, our method significantly outperforms the best existing method in terms of data access and computation costs.

Keywords

Nearest neighbor, Spatial databases, Query processing.

1. INTRODUCTION

Advances in GIS and mobile computing have led to research interests in continuous spatial queries in road networks. An application scenario is: “a driver who uses a GPS navigator to continually query the nearest gas station with respect to the moving query location.” The main challenge of performing such a task in real time is that the user’s movement may result in changes in the shortest paths from the user to the surrounding nodes. These changes cause reevaluation of network distances, which is a costly operation in terms of graph traversal [7, 14].

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

ACM GIS '10, November 2-5 2010, San Jose, CA, USA
Copyright 2010 ACM 978-1-4503-0428-3/10/11 ...\$10.00.

To facilitate monitoring of the nearest neighbor (NN) with respect to a moving query point, research in this area largely focuses on techniques which utilize precomputed network distances and/or partial results. One common example of such techniques is the *network Voronoi diagram (NVD)* [10, 11]. The NVD of a set \mathcal{D} of data objects consists of $|\mathcal{D}|$ cells (i.e., collections of road segments) such that every location in a cell shares the same nearest object in \mathcal{D} . Figure 1 shows the NVD of four objects a , b , c and d on a road network. As shown in the figure, the road network is decomposed into four cells $VC(a)$, $VC(b)$, $VC(c)$ and $VC(d)$ corresponding to a , b , c and d , respectively. Assume that a user at the location q is searching for the nearest neighbor. Using the NVD technique, we can see that q is in the Voronoi cell $VC(c)$ (the grey region), hence c is the NN of q . The NN remains unchanged as long as q remains inside the same region.




Figure 1: Network Voronoi diagram of a dataset $\{a, b, c, d\}$

A major drawback of this technique is that it requires access to all objects and network nodes in the data space. In a large network, the practicality of this technique lies in an assumption that construction of an NVD can be done offline. The offline-construction assumption may not hold in navigation applications where the dataset \mathcal{D} contains different kinds of objects and users submit object preferences during runtime. For example, a driver searching nearby gas stations that offer bio-diesel and accept petrol discount cards would require an NVD to be built from that specific subset of \mathcal{D} . Another case is when a navigation application uses the travel time derived from the current traffic condition as the proximity measure. In such an application, changes in the traffic condition will invalidate precomputed NVDs. When offline-construction is inapplicable, the best existing method [1] is to execute snapshot k NN queries at network nodes encountered by the query point. As mentioned earlier, a drawback of this method is the high graph traversal cost.

To alleviate both drawbacks, we propose a technique called the *local network Voronoi diagram (LNVD)*. Specifically, an LNVD is constructed by computing the NVD of a subspace around the query point. The main novelty of our approach is a method to determine

upon which part of the LNVD can be relied to produce accurate results. No graph traversal operation is required as long as the query point remain in this part. Consequently, our method requires neither offline-construction nor repetitive graph traversal operations.

The contributions of this paper can be summarized as follows.

- We propose a technique, called LNVD, to locally compute an NVD. Our technique is the first to address the problem of locally computing a network Voronoi diagram.
- We generalize the LNVD technique to the order- k LNVD (k LNVD) to support monitoring of k nearest objects and provide proof of correctness for the technique.
- We conduct a complexity analysis and show the construction time with respect to the value of k and the number n of nodes in the k LNVD. We also show that k LNVD has a lower overall cost than the best competitor [1].
- We conduct extensive performance evaluations and our results show our proposed technique significantly outperforms the best competitor.

2. PROBLEM FORMULATION

We study the problem of road-network nearest neighbor monitoring for a moving query point and a static dataset. A road network can be represented as a graph G , i.e., a set N of nodes (junctions) and a set E of edges (road segments). The weight of an edge $\text{EDGE}(n_i, n_j)$ denotes the cost of travelling from node n_i to node n_j , e.g., the distance or time. Edge weights can only assume non-negative values. For any given two locations v_1 and v_2 on the graph G , $\text{DIST}(v_1, v_2)$ denotes the cost of traversing $\text{PATH}(v_1, v_2)$, the path connecting v_1 to v_2 that minimizes the travelling cost. Since in practice a road network may consist of both one-way and two-way roads, G is modeled as a directed graph. That is, an edge in E can be unidirectional or bidirectional. In summary, G satisfies the following quasi-metric conditions: (i) *non-negativity*, (ii) *identity of indiscernibles*, and (iii) *triangle inequality*.

The k nearest neighbor (k NN) query [6] finds the nearest k data objects with respect to a given query point. A formal definition of the k NN query is given as follows.

DEFINITION 1 (k NEAREST NEIGHBOR (k NN) QUERY).

Given a set \mathcal{D} of objects and a query point q , the k NN query finds a list \mathcal{A} of objects such that:

- \mathcal{A} contains k objects from \mathcal{D} ;
- for any object $x \in \mathcal{A}$ and object $y \in (\mathcal{D} - \mathcal{A})$, $\text{DIST}(q, x) \leq \text{DIST}(q, y)$;
- for any adjacent entries x_i and x_{i+1} in the list \mathcal{A} , $\text{DIST}(q, x_i) \leq \text{DIST}(q, x_{i+1})$.

The moving k NN (Mk NN) query is defined as follows.

DEFINITION 2 (MOVING k NN (Mk NN) QUERY). Given a

set \mathcal{D} of objects and a moving query point q , the Mk NN query finds a k NN result for every position of q .

3. RELATED WORK

Papadias et al. [12] proposed two approaches to processing spatial-network k NN queries, the *incremental network expansion (INE)* and *incremental Euclidean restriction (IER)*. INE- k NN is based on the Dijkstra's algorithm [4], i.e., incrementally expanding the search space along network edges until k nearest objects are found. IER- k NN finds a number k' of nearest neighbors in the Euclidean distance where k' is greater than or equal to k and calculates the network distances of these k' objects to obtain the actual spatial-network k NNs. Furthermore, an optimization in terms of

data access and distance calculations to IER- k NN was proposed by Deng et al. [3]. In the same manner as the A* algorithm [13], their method uses the Euclidean distance measure to derive an upper bound value for pruning purposes.

Cho and Chung [1] proposed a continuous k NN technique that performs snapshot k NN queries at the intersections on the query path. Specifically, using one of the aforementioned snapshot k NN techniques [3, 12], continuous query answers can be produced by exploiting the fact that the movement of a user travelling in a spatial network is constrained by network connectivities. They showed that k NN results between any two nodes can be inferred from those of the nodes. The drawback of this approach is that it can incur repetitive distance calculations over similar sets of nodes.

Mouratidis et al. [9] studied a k NN monitoring problem in a dynamic environment where locations of the query point and data objects as well as edge weights can change. They proposed a method to handle this problem by keeping part of the k NN answer though incremental maintenance of the shortest path tree (SPT). Incremental SPT maintenance is suitable for changing edge weights and moving data objects, since such updates trigger only small changes to the SPT. However, using this method with moving query points still incurs frequent reevaluation of the k NN answer. This is because as the query point moves from the original location q to q' , a vast majority of the shortest path tree (with the root at q) can become invalid with respect to the new location q' .

Both Dijkstra's and the A* algorithms incur expensive graph traversal operations, this cost is accentuated with continuous queries. This is because, the shortest paths from a query point q to its surround nodes change as q moves from one node to another [14]. To address this issue, algorithms that utilize precomputed distances or partial results have been proposed. One classic example is the *network Voronoi diagram (NVD)* [10, 11]. In a spatial network, the NVD of a set \mathcal{D} of data objects consists of regions (i.e., collections of road segments) and each object in \mathcal{D} is associated with a region/cell in which it dominates.

Kolahdouzan and Shahabi proposed a technique called the *Voronoi-based network nearest neighbor (VN^3)* search [7] based in the first-order network Voronoi diagram. The technique can provide Mk NN results via a simple lookup operation, while processing Mk NN queries with a higher value of k can be done by examining neighboring regions. In addition, they also proposed continuous k NN algorithms [8] based on the VN^3 technique. An important drawback of these approaches [7, 10, 11] is that data objects are used in the precomputation which means that changes in the dataset or edge weights can render the NVD obsolete.

To remedy the dataset update problem, data objects are decoupled from the precomputation process. This decoupling is done by precomputing shortest paths between network nodes and storing the shortest path information in a hierarchical or grid-based data structure. Samet et al. [14] formulated a k NN algorithm that utilizes precomputed shortest path information stored in quadtrees. Specifically, for each node x , the shortest path to every other node in the network is computed. Each edge $\text{EDGE}(n_x, n_y)$ (from n_x to an immediate neighbor n_y) is associated with destination nodes n_z such that $\text{EDGE}(n_x, n_y)$ forms part of the shortest path $\text{PATH}(x, z)$. This creates a shortest path map of all nodes reachable by n_x , where each destination node is categorized according to the edges to the immediate neighbors. For efficiency purposes, collections of destination nodes are stored as quadtree regions where each region contains only one node type.

Demiryurek et al. [2] proposed a grid-based solution to storing shortest path information. Each $\text{EDGE}(n_x, n_y)$ stores a boolean vector $\langle b_1, \dots, b_c \rangle$ where b_i indicates whether $\text{EDGE}(n_x, n_y)$ forms

part of the shortest path from x to any node in Grid i . In addition to a k NN algorithm, they also proposed k NN monitoring techniques which utilize this shortest path information. Although these techniques [2, 14] are unaffected by dataset updates, other updates such as changes in edge weights may affect shortest paths and invalidate precomputed shortest path information.

4. PRELIMINARIES

In this section, we describe the *shortest path tree (SPT)* and *network Voronoi diagram (NVD)* concepts [11] which form a basis of our proposed method. The simplest form of an SPT is one with a single reference point, which can be obtained by executing the Dijkstra’s algorithm [4]. Figure 2 illustrates how $\text{PATH}(n_{12}, c)$ is calculated by incrementally discovering surrounding nodes according to their distances to c . In this example, n_9 is discovered first with a distance of 3 units. Next, n_8 , n_5 and n_7 are discovered with distances of 5 (via c), 6 (via n_9) and 6 units (via n_8), respectively. The search halts when n_{12} is discovered with a distance of 8 units via n_9 .

Figure 2 also shows that each of the nodes involved in the search is associated with a label (p, d, n_h) where p denotes the reference point, d denotes the shortest distance, and n_h denotes the node from which the shortest distance is derived. The shortest path from n_{12} to c can be obtained by recursively traversing the next hop n_h in the label (p, d, n_h) until 9 is reached. In this case, we obtain (n_{12}, n_9, c) as the shortest path. As displayed in Figure 2, a *shortest path tree* (where the reference point c is the root) is formed by linking the next hop n_h of all nodes involved in the calculation.




Figure 2: Calculation of $\text{DIST}(n_{12}, c)$ via the Dijkstra’s algorithm where each node (involved in the computation) is associated with a label (q, d, n_h) .

The same incremental expansion principle can also be used to construct an NVD of a set \mathcal{D} of data objects. The first step is to build an SPT with objects in \mathcal{D} as reference points and incrementally expand the search space around each of these objects until all nodes are visited. By ensuring that nodes are visited according to the minimum distance to the objects in \mathcal{D} , each node is associated with its nearest object.

After the SPT of \mathcal{D} is constructed, NVD cell boundaries are computed based on the information regarding the distance and NN obtained during the SPT construction. Specifically, a boundary point is computed for each edge where the two end nodes have different nearest objects. After the Voronoi cells boundaries are identified, the NN with respect to any query location q in the data space can be obtained by identifying the cell in which q resides. The main drawbacks of this method is that it requires: (i) global access to the data, where costs could be prohibitive; (ii) a preconstructed Voronoi diagram for each dataset. In the next section, we will show how such drawbacks can be mitigated by our proposed technique.

5. LOCALIZING VORONOI DIAGRAMS

This section presents our proposed technique, the *local network Voronoi diagram (LNVD)*, and how it can be used to process moving nearest neighbor queries. As exemplified in Figure 3, an LNVD is a Voronoi diagram of a subspace. In other words, we delimit the area of interest to road segments (edges), intersections (nodes) and data objects within a specific search region. As a result, construction of an LNVD requires only local information and by limiting the search region to a small, manageable size, an LNVD can be constructed on the fly. In this example, we set our search region to the region within a range of 14 units from the query point q . This search region (highlighted in grey) is denoted as $\text{SR}(q, 14)$. Based on information within $\text{SR}(q, 14)$, an LNVD is constructed. Since $\text{SR}(q, 14)$ contains three objects a , b and c , our LNVD, in this case, has three local Voronoi cells where each is associated with its nearest neighbor (NN). The correctness of these NNs are guaranteed even under partial information.

An LNVD comprises local Voronoi cells of data objects within the search region. Each local Voronoi cell $\text{LVC}(p)$ can be bounded by two bisector types:

- (i) *internal*: a boundary between $\text{LVC}(p)$ and one of its immediate neighbors known.
- (ii) *external*: a boundary between a cell of a known neighbor and a region in which the search region does not provide further NN information.

As illustrated in Figure 3, the cell $\text{LVC}(c)$ is bounded by three internal bisectors and three external bisectors. The LNVD in this example is the road segments covered by the three cells $\text{LVC}(a)$, $\text{LVC}(b)$ and $\text{LVC}(c)$. As long as the query point q remains inside the LNVD, the NN with respect to q can be obtained by identifying the cell in which q resides.




Figure 3: Local network Voronoi diagram with respect to a search region $\text{SR}(q, 14)$ and objects a , b and c

5.1 Data Retrieval

The first step to constructing an LNVD is to retrieve the data objects in the search region $\text{SR}(q, r)$. We utilize the *incremental-network-expansion (INE) range query* [12] to retrieve objects and to calculate the network distances of nodes within the search region. For example, $\text{SR}(q, 14)$ in Figure 3 comprises nodes and edges in the grey region whose distances from q are less than or equal to 14 units. The boundaries that delimit this search space are shown as grey crosses in the figure.

5.2 Reliability of Local Information

Since we consider only a subset of the whole data space, the Voronoi diagram constructed based on this local information cannot be complete. In this subsection, we present a method to determine the reliability of a local NN result. The condition of an object being reliable with respect to a node n_i is formally defined as follows.

DEFINITION 3 (RELIABILITY OF A DATA OBJECT). Given an object p inside $SR(q, r)$ and s outside $SR(q, r)$, p is reliable with respect to n_i if it is guaranteed that s is not closer to n_i than p . That is, $DIST(n_i, p)$ is less than or equal to a lower bound of $DIST(n_i, s)$.

We now describe how to derive $INFIMUM(DIST(n_i, s))$, the greatest lower bound of $DIST(n_i, s)$. Since we only know that s is outside $SR(q, r)$, $INFIMUM(DIST(n_i, s))$ is the MINIMUM DISTANCE (MINDIST) from n_i to the boundary set \mathcal{B} of $SR(q, r)$. That is,

$$INFIMUM(DIST(n_i, s)) = MINDIST(n_i, \mathcal{B}).$$

As a result, we need to rule out the region containing locations v such that v is closer to \mathcal{B} than any object. This is because, an object outside $SR(q, r)$ can be the NN with respect to v . In other words, $SR(q, r)$ does not pertain enough information to produce the NN with respect to v . As shown in Figure 3, $MINDIST(n_{12}, \mathcal{B})$ is 4 units while c as the nearest object in $SR(q, r)$ produces a distance of 8 units. Hence, n_{12} does not have a guaranteed NN. A method to identify regions that have no guaranteed NN is presented in the next subsection.

5.3 Construction

The intuition behind our LNVD construction method is to identify the following two region types:

- *local Voronoi cell (LVC)*, a region with a guaranteed NN;
- *boundary-dominated cell (BDC)*, a region in which the NN cannot be inferred from information in $SR(q, r)$.

We propose a method to incrementally label nodes based on their distances to its nearest object and the boundaries of $SR(q, r)$. Specifically, we consider the boundary set \mathcal{B} as one data object whose distance from a node n_i is given by $MINDIST(n_i, \mathcal{B})$. By associating the closest object (or the boundary set \mathcal{B}) to each node in $SR(q, r)$, we can identify edges $EDGE(n_i, n_j)$ such that n_i and n_j are associated with two different nearest objects (denoted as p_i and p_j , respectively). The bisector between the cells of p_i and p_j on $EDGE(n_i, n_j)$ is the location equidistant to p_i and p_j .

Algorithm 1 provides the steps of our proposed LNVD construction method. We first define the network data structure used in the algorithm. A network/graph is represented using the adjacency-list format, i.e., a list of nodes where each node entry contains information regarding its adjacent nodes. The data structure NODE is defined as follows.

DEFINITION 4 (NODE). The structure of a node n_i contains the following attributes:

- ID: the node identification.
- Adjacency list (AdjList): a list of edges to/from immediate neighbors and associated weights.
- Label: a label of n_i contains three attributes (p, d, n_h) where
 - p is the associated NN,
 - d represents $DIST(n_i, p)$, and
 - n_h represents the next hop in order to reach p .
- Type: a node type is “Labelable” by default and becomes “Permanent” after it is labelled.

Another important data structure used in the construction process is the *priority queue*, which is defined as follows.

DEFINITION 5 (PRIORITY QUEUE). A *priority queue* is a container of priority queue entries. Each priority queue entry comprises four attributes (n, p, d, n_h) . The first attribute n is a reference of the node to which the entry corresponds. The other three

attributes (p, d, n_h) form a labeling candidate for n . Entries in a priority queue are organized in such a way that an entry (p, d, n_h) with the smallest labelling distance d is always the head entry.

Algorithm 1: Construct-LNVD(\mathcal{D}, G, q, r)

input : Dataset \mathcal{D} , Graph G , Query Point q , Distance r
output: Graph G with labels and LVC boundaries

- 1 Initialize Priority Queue PQ ;
- 2 $\mathcal{A} \leftarrow \text{INE-RangeQuery}(\mathcal{D}, G, q, r)$;
- 3 $\mathcal{B} \leftarrow \{v : \text{DIST}(q, v) = r\}$;
- 4 **for each** (object p in \mathcal{A}) **do**
- 5 Node $n_p \leftarrow$ Create a network node from p ;
- 6 $G.\text{Insert}(n_p)$;
- 7 $PQ.\text{Insert}(PQEntry(n_p, p, 0, -))$;
- 8 **for each** (node n_b adjacent to \mathcal{B}) **do**
- 9 $PQ.\text{Insert}(PQEntry(n_b, \mathcal{B}, r - \text{DIST}(q, n_b), \mathcal{B}))$;
- 10 **while** PQ is **not** empty **do**
- 11 $PQEntry(n, p, d, n_h) \leftarrow PQ.\text{DequeueHead}()$;
- 12 **if** $n.$ Type is Labelable **then**
- 13 $n.$ Label $\leftarrow (p, d, n_h)$;
- 14 $n.$ Type \leftarrow Permanent;
- 15 **for each** Adjacent node n_a of n in $SR(q, r)$ **and** n_a is Labelable **do**
- 16 Distance $d_a \leftarrow d + \text{weight of } \text{EDGE}(n_a, n)$;
- 17 $PQ.\text{Insert}(PQEntry(n_a, p, d_a, n))$;
- 18 **for each** bidirectional edge $EDGE(n_i, n_j)$ in $SR(q, r)$ **do**
- 19 Check for external and internal bisectors on $EDGE(n_i, n_j)$;
- 20 **return** G ;

The LNVD construction steps in Algorithm 1 can be described as follows. The algorithm accepts four input parameters: the dataset \mathcal{D} , the graph G , the query point q and the range r . The construction process is divided into the following three stages.

Initialization (Lines 1 to 9): The first step is initialization of a priority queue PQ . Next we execute a range query to retrieve objects whose distances from q is less than or equal to r , and to calculate network distances of nodes within the range. Upon completion of the range query execution, we also identify the boundary set \mathcal{B} of $SR(q, r)$, i.e., a set of locations v such that $DIST(q, v)$ is equal to r . (A method to determine an appropriate size for $SR(q, r)$ is given in Section 8.1.) For each of the retrieved objects p , we create a node entry n_p , insert it into G and modify the affected edges accordingly. Next, a priority queue entry $(n_p, p, 0, -)$ is created and inserted into PQ . Note that the fourth attribute, next hop, is null (“-”) since n_p is already at p . For each node n_b adjacent to \mathcal{B} , we create an entry $(n_b, \mathcal{B}, MINDIST(n_b, \mathcal{B}), \mathcal{B})$, where $MINDIST(n_b, \mathcal{B})$ is equal to $(r - \text{DIST}(q, n_b))$. The entry is then inserted into PQ .

In the context of our running example (Figure 3), after the initialization steps, PQ has

- initial object entries of $(a, a, 0, -)$, $(b, b, 0, -)$ and $(c, c, 0, -)$;
- initial boundary entries of $(n_2, \mathcal{B}, 3, \mathcal{B})$, $(n_1, \mathcal{B}, 4, \mathcal{B})$, $(n_{10}, \mathcal{B}, 4, \mathcal{B})$, and $(n_{12}, \mathcal{B}, 4, \mathcal{B})$.

SPT construction (Lines 10 to 17): Our SPT construction process is conducted in a best-first manner. In particular, at each iteration of the *while loop*, the head entry (n, p, d, n_h) is retrieved from PQ . Then, we check whether the corresponding node n can still be labelled. If so, the node label is set to (p, d, n_h) and the node type is changed to permanent. Next, for each node n_a in $SR(q, r)$ that can reach n by a single hop, we calculate a labelling distance d_a for (n_a) as the sum of d and the weight of $EDGE(n_a, n)$. The

while loop iterates until PQ is exhausted. An end result of this step is an SPT of objects in $SR(q, r)$ as generators. Each node in the SPT is associated with a reliable NN or the boundary set \mathcal{B} as displayed in Figure 4.

Based on the example given in Figure 3, Figure 4 shows the order in which nodes are labelled during the SPT construction. The SPT is generated from $\{a, b, c\}$ and \mathcal{B} as reference locations. The construction process starts from nodes with smallest distances to these reference locations and incrementally explore neighboring nodes until every node $SR(q, r)$ is associated with its nearest object.

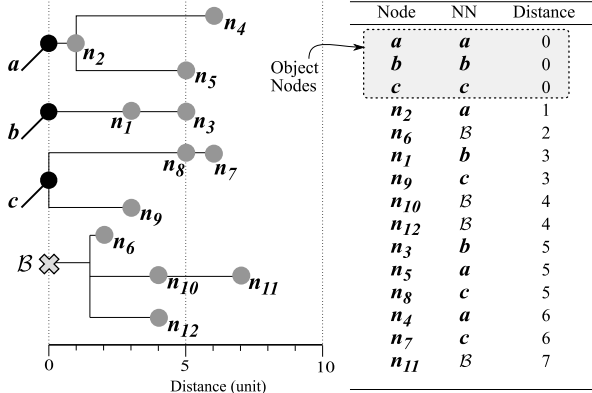


Figure 4: Shortest path trees generated from $\{a, b, c\}$ and \mathcal{B}

Bisector calculation (Lines 18 to 19): After having created an SPT, each node is associated with its reliable NN or the boundary set \mathcal{B} . The next step is to compute external and internal bisectors by checking each bidirectional edge $EDGE(n_i, n_j)$ in $SR(q, r)$.¹

- *Internal bisector.* An internal bisector is on an edge $EDGE(n_i, n_j)$ such that n_i and n_j are associated with two different NNs (assuming no boundaries on $EDGE(n_i, n_j)$). Let p_i and p_j be the NNs of n_i and n_j respectively. The internal bisector on $EDGE(n_i, n_j)$ is the location v such that

$$DIST(v, p_i) = DIST(v, p_j).$$

Consider $EDGE(n_3, n_7)$ in Figure 3 as an example. The NN of n_3 is b with the distance of 5 units and the NN of n_7 is c with the distance of 6 units. The bisector between b and a on $EDGE(n_3, n_7)$ is the location 1 units from n_3 and 2 units from n_7 . The location is 7 units to both b and a .

- *External bisector.* An external bisector is on an edge $EDGE(n_i, n_j)$ such that n_i is associated with an NN p_i and n_j is associated with \mathcal{B} .² The external bisector on $EDGE(n_i, n_j)$ is the location v such that

$$DIST(v, p_i) = MINDIST(v, \mathcal{B}).$$

Consider $EDGE(n_5, n_6)$ in Figure 3 as an example. The NN of n_5 is a with the distance of 4 units while n_6 is associated with \mathcal{B} with $MINDIST(n_6, \mathcal{B})$ of 2 units. The bisector between a and \mathcal{B} on $EDGE(n_5, n_6)$ is the location 2 units from n_5 and 4 units from n_6 . The location is 6 units to both n_5 and \mathcal{B} .

¹If $EDGE(n_i, n_j)$ is unidirectional ($n_i \rightarrow n_j$), the NN of any location v between n_i and n_j is the same as that of n_j , since we cannot travel to n_i without passing n_j . This principle also applies to the order- k LNVD described in the next section.

²For the purpose of bisector calculations, we consider each boundary in \mathcal{B} as a node with the $MINDIST$ to \mathcal{B} of 0 units.

Note that boundary calculation on one edge is independent from others. Hence one may choose to defer the computation until an edge is visited by the query point to avoid unnecessary boundary computation.

5.4 Query Processing

Processing a moving nearest neighbor query using the LNVD technique can be done as follows. Assume that the initial location of the query point is q . We can use Algorithm 1 to construct an LNVD with respect to a search region $SR(q, r)$. After an LNVD is constructed, it can be used to provide the following information:

- the nearest neighbor with respect to any location q' within the LNVD by identifying the local Voronoi cell in which q' resides;
- the shortest path to the nearest neighbor using the shortest path tree.

If q' is outside the LNVD (i.e., having crossed an external bisector), we can again use Algorithm 1 to construct an LNVD with respect to the current query location q' . It is also possible to derive a technique to make use of existing information to reduce the construction cost of the new LNVD. Formulation of such a technique and its computational benefits (in comparison to constructing a new one from scratch) can be investigated as future work.

5.5 Proof of Correctness

This section shows that Algorithm 1 produces accurate results. First, we show that network distances calculated by the algorithm is correct in Lemma 1.

LEMMA 1 (CORRECTNESS OF DISTANCE CALCULATIONS). *For any location v in $SR(q, r)$, only nodes in $SR(q, r)$ are required to calculate (i) the distance from v to the set \mathcal{B} of $SR(q, r)$ boundaries, and (ii) the distance from v to its guaranteed NN p (if exists).*

PROOF. First, we show that a path from v to a boundary in \mathcal{B} that includes a location v' outside $SR(q, r)$ can never be the path that minimizes the distance from v to \mathcal{B} . Since $SR(q, r)$ is fully enclosed by \mathcal{B} , in order to travel from v to v' , one has to encounter a boundary x in \mathcal{B} . Since, $DIST(x, v')$ cannot assume a negative value, $DIST(v, v')$ can never be smaller than or equal to $MINDIST(v, \mathcal{B})$ and the calculation of $MINDIST(v, \mathcal{B})$ need not include any node outside $SR(q, r)$.

Second, we show that for each location v with a reliable NN of p , $PATH(v, p)$ does not include any location v' outside $SR(q, r)$. According to Definition 3, for p to be reliable with respect to v , the following condition has to be satisfied:

$$DIST(v, p) \leq MINDIST(v, \mathcal{B}).$$

Since v' is outside the boundaries of $SR(q, r)$,

$$DIST(v, p) \leq MINDIST(v, \mathcal{B}) < DIST(v, v').$$

Hence, $PATH(v, p)$ can never include v' .

As a result, we need to consider only nodes and edges in $SR(q, r)$ in order to calculate distances from a location v to its guaranteed NN or the boundary set \mathcal{B} . \square

After having established that distances calculated by Algorithm 1 is correct, we are now ready to show that the NN associated to each local Voronoi cell is also correct in Lemma 2.

LEMMA 2 (CORRECTNESS OF ALGORITHM 1). *Given an object p inside a search space $SR(q, r)$, for any location v in $LVC(p)$, p is the nearest neighbor of v .*

PROOF. Since $LVC(p)$ is bounded by internal and external bisectors, any location v in $LVC(p)$ satisfies the following two conditions. First, for all data objects p' inside $SR(q, r)$,

$$DIST(v, p) \leq DIST(v, p').$$

Second, p is reliable with respect to v . That is, for all data objects s outside $SR(q, r)$,

$$DIST(v, p) \leq DIST(v, s).$$

Since each data object is either inside or outside $SR(q, r)$, p is guaranteed to be the nearest neighbor with respect to any given location v in $LVC(p)$. \square

6. EXTENDING TO ORDER-K

The earlier described LNVD concept can be generalized to the order- k LNVD (k LNVD). An k LNVD with respect to $SR(q, r)$ consists of local Voronoi cells (LVCs), where each cell corresponds to a specific k NN result. As displayed in Figure 5, an k LNVD (where k is 2) consists of four LVCs corresponding to the following four rank-sensitive k NN answers: $\langle a, c \rangle$, $\langle b, c \rangle$, $\langle c, a \rangle$ and $\langle c, b \rangle$. Similar to an order-1 LNVD, the boundaries of each LVC in an k LNVD may comprise internal and external bisectors which are defined as follows.

- *Internal bisector (white cross)*. A location where the query answer changes from one guaranteed k NN to another.
- *External bisector (black cross)*. A location where the boundary set \mathcal{B} has the same distance as the k th NN. That is, the k th NN is about to become unreliable.




Figure 5: Order-2 local network Voronoi diagram with respect to a search region $SR(q, 14)$ and objects a, b and c

6.1 Algorithm

Algorithm 2 provides the steps to construct a k LNVD. The algorithm is based on the same best-first graph traversal principle in Algorithm 1. In the same fashion as the order-1 LNVD where each network node is associated with its NN, each network node in a k LNVD is associated with its k NNs. Specifically, for each node n , $n.Label$ is replaced by $n.LabelList$ (a list of labels). Hence, the labelling of n is done by inserting a label into $n.LabelList$ and n becomes permanent upon completion of k labels or inclusion of a boundary label.

The algorithm shares the same initialization steps (Lines 1 to 9) as Algorithm 1. For the SPT construction (Lines 10 to 18), we also utilize the best-first graph traversal principle using the priority queue PQ . At each iteration, the head entry (n, p, d, n_h) is dequeued from PQ . The entry is added into $n.LabelList$ if (i) $n.Type$ is labelable, and (ii) there is no label with p as the associated NN

Algorithm 2: Construct- k LNVD(\mathcal{D}, G, k, q, r)

input : Dataset \mathcal{D} , Graph G , k , Query Point q , Distance r
output: Graph G with labels and LVC boundaries

- 1 Initialize Priority Queue PQ ;
- 2 $\mathcal{A} \leftarrow INE\text{-RangeQuery}(\mathcal{D}, G, q, r)$;
- 3 $\mathcal{B} \leftarrow \{v : DIST(q, v) = r\}$;
- 4 **for each** (object p in \mathcal{A}) **do**
- 5 Node $n_p \leftarrow$ Create a network node from p ;
- 6 $G.Insert(n_p)$;
- 7 $PQ.Insert(PQEntry(n_p, p, 0, -))$;
- 8 **for each** (node n_b adjacent to \mathcal{B}) **do**
- 9 $PQ.Insert(PQEntry(n_b, \mathcal{B}, r - DIST(q, n_b), -))$;
- 10 **while** PQ is not empty **do**
- 11 $PQEntry(n, p, d, n_h) \leftarrow PQ.DequeueHead()$;
- 12 **if** $n.Type$ is Labelable **and** no existing label with p **then**
- 13 $n.LabelList.Add(Label(p, d, n_h))$;
- 14 **if** $n.LabelList.Length = k$ **or** Object p is the set \mathcal{B} of $SR(q, r)$ boundaries **then**
- 15 $n.Type \leftarrow$ Permanent;
- 16 **for each** Adjacent node n_a of n in $SR(q, r)$ **and** n_a is Labelable **do**
- 17 Distance $d_a \leftarrow d +$ weight of $EDGE(n_a, n)$;
- 18 $PQ.Insert(PQEntry(n_a, p, d_a, n))$;
- 19 **for each** bidirectional edge $EDGE(n_i, n_j)$ in $SR(q, r)$ **do**
- 20 Check for external and internal bisectors on $EDGE(n_i, n_j)$;
- 21 **return** G ;

in $n.LabelList$. The node type is changed from labelable to permanent if there are k labels or an entry with \mathcal{B} is inserted. For each adjacent node n_a in $SR(q, r)$, a priority queue entry is created and is inserted into PQ . The *while* loop repeats until PQ is exhausted. As a result, each node n_i is associated with at most k unique reliable NNs, which are ranked according to the distances from n_i . For each node n_i in $SR(q, r)$, the number of associated NNs depends on the number of reliable NNs it has, which is determined by $MINDIST(n_i, \mathcal{B})$. For examples,

- n_{10} has no NN associated because \mathcal{B} is closer to n_{10} than any object in $SR(q, r)$;
- n_1 is associated with only b , because b is the only object with a distance smaller than $MINDIST(n_1, \mathcal{B})$;
- n_5 has two NNs, a and c , because both $DIST(n_5, a)$ and $DIST(n_5, c)$ are smaller than $MINDIST(n_5, \mathcal{B})$.

Internal and external bisectors can be categorized by categorizing the bidirectional edges $EDGE(n_i, n_j)$ into three types:

- *Both n_i and n_j have data objects as k NNs.* In this case, we can use the *split-point calculation* method [8] to identify locations that the k NN answer changes. All bisectors calculated in this case are internal, since there is no boundary set \mathcal{B} involved in the calculations. As exemplified in Figure 6, the bisector $BT(a, c)$ separates the two LVCs, $LVC(\langle a, c \rangle)$ and $LVC(\langle c, a \rangle)$.




Figure 6: Both end nodes have 2 reliable data objects

- *One of the end nodes (namely n_j) has the boundary set \mathcal{B} in its label list.* We still use the split-point calculation method [8] to identify internal bisectors based on the NNs from the two end nodes. The only difference is that the calculation terminates when \mathcal{B} replaces the k th nearest ob-

ject. That is, an external bisector x is found. By this means we guarantee that for any location v between n_i and x , $\text{MINDIST}(v, \mathcal{B})$ is always greater than or equal to the k th NN with respect to v . Hence, all of the k NNs are reliable with respect to v (Definition 3). As exemplified in Figure 7, n_7 is associated with 2 reliable objects c and b , while n_3 only has b as its only one reliable object. The edge $\text{EDGE}(n_7, n_3)$ has one internal bisector and one external bisector. The internal bisector $\text{BT}(c, b)$ separates $\text{LVC}(\langle c, b \rangle)$ and $\text{LVC}(\langle b, c \rangle)$. The external bisector separates $\text{LVC}(\langle b, c \rangle)$ from the exterior of the LNVD.




Figure 7: Only one end node has 2 reliable data objects

- Both end nodes have the boundary set \mathcal{B} in their label lists. Such edges are ignored since neither end nodes have reliable k NNs.

As a result, each $\text{EDGE}(n_i, n_j)$ is decomposed into local Voronoi cells (LVCs) where each corresponds to a specific rank-sensitive k NN answer.

6.2 Proof of Correctness

Since network distances in Algorithm 2 are calculated in the same way as Algorithm 1, their correctness is also supported by Lemma 1. For the k NN results, we can extend Lemma 2 and its associated proof as follows.

LEMMA 3 (CORRECTNESS OF ALGORITHM 2). *Given an LVC with an associated k NN answer of $\langle p_1, \dots, p_k \rangle$, for any location v the LVC, $\langle p_1, \dots, p_k \rangle$ are the rank-sensitive k NN answer with respect to v .*

PROOF. The function of internal and external bisectors of k LVND is similar to those of order-1 LNVD, so we extend the proof for Lemma 2 as follows. For any location v in an LVC of a k NN list $\langle p_1, \dots, p_k \rangle$, internal and external bisectors ensures the following.

- First, internal bisectors ensure that the k NN list is sorted according to the distances to v , and for any object p' in $\text{SR}(q, r)$ but not in the k NN list, p' is not closer to v than any object in the k NN list.
- Second, external bisectors ensure that all of objects in the k NN list are reliable with respect to v . That is, there is no object outside $\text{SR}(q, r)$ closer to v than any object in $\langle p_1, \dots, p_k \rangle$.

As a result, $\langle p_1, \dots, p_k \rangle$ is guaranteed as the rank-sensitive k NN answer for all locations v in $\text{LVC}(p)$. \square

7. COMPLEXITY ANALYSIS

In this section, we analyze the time complexity of the k LVND construction algorithm and the best competitor by Cho and Chung [1]. For conciseness, we call their method “CC- k NN”. We assume that the number of intersections is much greater than the number of data objects as it is the case in existing literature [11, 12].

7.1 Proposed Method: Order- k LNVD

We now derive the complexity of Algorithm 2 as follows. The total cost can be considered as the sum of the costs of the following steps: (i) execution of a range query and insertion of initial entries into the priority queue; (ii) the main k LVND construction loop.

In Sections 5 and 6, for exposition purposes, a search region is expressed as $\text{SR}(q, r)$, where q denotes the search region’s centroid and r denotes the search range. In other words, q determines the location and r determines the size of the search region. In our analysis, we use the number n of nodes inside the search region instead of r to represent the size of the search region. This is because, n has a more easy to analyze, observable effect on the processing cost than r . Note that the same INE range query concept can still be applied to discover n nearest nodes around q . That is, the search terminates when n nearest nodes around the query point is discovered instead of when the distance r is reached. In this case, the value of r is determined by n . That is, r is the distance from q to the farthest node of the n nodes in the search space.

We first derive the range query cost. Assume that the Dijkstra’s algorithm is used in the retrieval process to compute the distances and shortest paths of nodes in the search region. We obtain $\mathcal{O}(e + n \log n)$ as the cost of this operation, where e denotes the number of edges in the search region [5]. In a planar graph, e is proportional to the product of n and the degree of the graph, which is a small constant. Therefore, we have $\mathcal{O}(n \log n)$ as the range query cost in our setting. The next operations are insertions of the generators in \mathcal{A} and nodes adjacent to the boundary set \mathcal{B} into the priority queue. The cost of each insertion operation (using the Fibonacci heap) is constant [5]. Therefore the cost of these insertion operations are $\mathcal{O}(|\mathcal{A}| + |\mathcal{B}|)$. Both $|\mathcal{A}|$ and $|\mathcal{B}|$ are much smaller than n , so it is dominated by the range query cost. Therefore, the cost of the first step is $\mathcal{O}(n \log n)$.

Next, the cost of constructing a k SPT can be derived from the number of iterations of the *while loop* (Lines 10 to 18) and the cost of priority queue operations (i.e., dequeue and insertion) for each iteration. The number of *while loop* iterations is the maximum number of priority queue entries, which can be derived as a product of three parameters: the number n of nodes, the number k of labels per node and the maximum number of priority queue entries that can be generated from each label.³ The last parameter is equal to the degree d of the graph, i.e., the number of adjacent nodes. We have $(n \cdot k \cdot d)$ as the number of priority queue entries and the number of iterations.

For each iteration, the cost of a dequeue operation is logarithmic with respect to the priority queue size. We also use $(n \cdot k \cdot d)$ as the priority queue size. Hence, the total cost of the dequeue operations is $\mathcal{O}(n \cdot k \cdot d \cdot \log(n \cdot k \cdot d))$. For the priority queue insertions, only at most $(n \cdot k)$ of priority queue entries would satisfy the condition in Line 12 and the *for loop* (Lines 16 to 18) iterates at most d times. Therefore we have $\mathcal{O}(n \cdot k \cdot d)$ as the total cost of the insertion operations, which is dominated by the dequeue operations. In a planar graph, the degree d is a small constant, so the total cost of the second step is $\mathcal{O}(kn \log kn)$.

The cost of bisector calculations is not included in the construction cost. This is because calculations of bisectors on an edge $\text{EDGE}(n_i, n_j)$ can be deferred until $\text{EDGE}(n_i, n_j)$ is visited by the query point. That is, only edges along the query trajectory have their bisectors calculated. As a result, the construction cost involves only two cost components: (i) $\mathcal{O}(n \log n)$ for the initialization, and (ii) $\mathcal{O}(kn \log kn)$. Since the latter dominates the former, we have

$$\mathcal{O}(kn \log kn)$$

as the total construction cost.

The value of k is the number of nearest neighbors reported to the user. The number n of nodes determines the search region size,

³Under the assumption that $(|\mathcal{A}| + |\mathcal{B}|)$ is much smaller than n , the $(|\mathcal{A}| + |\mathcal{B}|)$ initial priority queue entries can be neglected.

which has the two following effects. First, n has a positive correlation with the construction cost. Second, a greater value of n also means that we have a larger k LNVD, which in turn provides more k NN answers per each construction. Studies on these effects of the value of n are given in Section 8.1.

7.2 Best Existing Competitor

We now consider the competitor, CC - k NN [1], which executes the k NN query at each network node. Assume that the Dijkstra’s algorithm [4] is used to discover the k NNs at each instance of the k NN query. The cost for each query instance can be given as $\mathcal{O}(n' \log n')$, where n' denotes the number of nodes visited by the Dijkstra’s algorithm in order to obtain the k NNs. The number n' of visited nodes can be estimated as a product of the number k of objects and the relative density ρ of nodes with respect to objects (i.e., *the number of nodes divided by the number of objects*). Figure 8 provides an empirical support for this estimate ($k \cdot \rho$) by comparing it with measured values of n' . The test was conducted on a road network dataset of North America with 175,813 nodes.

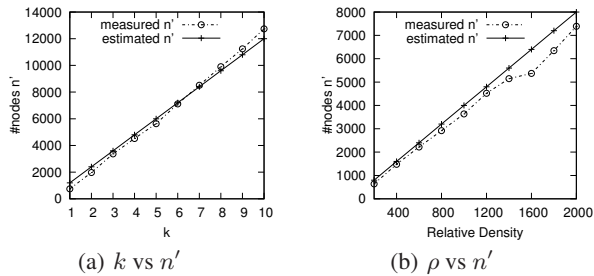


Figure 8: The number n' of nodes involved in a k NN evaluation in comparison to the estimate ($k \cdot \rho$) (Each result is the average from 20 query instances.)

We can therefore express the cost of each k NN retrieval as

$$\mathcal{O}(k\rho \log k\rho).$$

That is, the k NN cost is loglinear with respect to the product of k and ρ . This cost is incurred every time the query point encounters a new node.

7.3 Summary

For the k LVND, since the search space must include at least k objects, $(k \cdot n)$ is greater than $(k \cdot \rho)$. As a result, the k LVND cost of $\mathcal{O}(kn \log kn)$ is greater than the k NN cost. However, the k NN cost of $\mathcal{O}(k\rho \log k\rho)$ is incurred every time the query point encounters a new node, which can be disadvantageous in a setting with a high node density. In the next section, through experimental studies, we show that the order- k LNVD has a lower cost than the competitor, CC - k NN, when taking into account the fact that an LNVD contains k NN answers of multiple nodes.

8. EXPERIMENTAL STUDIES

In this section, we evaluate the performance of our proposed method in comparison to the best competitor, CC - k NN [1]. The experiments were conducted on a machine with an 2.66GHz Intel Core 2 Duo CPU, 4GB of main memory, and Linux 2.6.32 kernel as the operating system. All techniques were implemented on Java. We used a road network dataset of North America with 175,813 nodes and a diameter of 18,579 units. The experiments consist of two studies. In the first study, we present an empirical method to determine the size of a search region for the LNVD construction algorithm (Algorithm 2). In the second study, we compare the performance of our proposed method to that of the competitive one.

8.1 Determination of the Search Region Size

In the first experimental study, we present an empirical method to derive an appropriate size of the search region (SR) with different values of k and ρ . As mentioned earlier, the SR size can be varied by adjusting the number n of nodes it contains. The value of n is important to the performance of our k LNVD method. A greater n value provides a larger k LNVD but incurs a greater construction cost. If the n value is too small on the other hand, k LNVD may not contain enough complete nodes (nodes with a complete k NN answer) and results in frequent k LVND constructions. We use m to denote the number of complete nodes in a k LNVD.

In this experimental study we observe the effect of n on the processing cost t and the payoff, i.e., the number m of complete nodes. The processing cost (denoted as t) is the product of the following two measurements:

- the number of *while-loop* iterations in Algorithm 2, and
- the logarithm of the maximum priority queue size.

Our main objective is to find an SR size n that yields a reasonable cost-payoff ratio (t/m) for given values of k and ρ .

An appropriate search region size n depends on the k values, and the relative density ρ of nodes with respect to objects. Specifically, a greater value of k requires a greater value of n to counteract a stricter k NN completion requirement for each node, and a greater value of ρ requires a greater value of n to compensate for the reduced number of objects per node. As a result, we use $(n/\rho/k)$ as a normalized SR size to neutralize the effects of k and ρ . We varied the value of $(n/\rho/k)$ from 1 to 6 in this study. Experiments were conducted on k of 4 and 8 objects and ρ of 400 and 800 nodes per object.

Figure 9(a) shows the effect of $(n/\rho/k)$ on the processing cost t . We can see that t has a positive correlation with $(n/\rho/k)$, since an increase in n results in a greater number of iterations and a greater priority queue size. It can be seen that an increase in $(n/\rho/k)$ results in both an increase in the number m of complete nodes and the processing cost t . In order to find an appropriated value of n (with respect to fixed values of k and ρ), we display the processing cost per complete node, i.e., (t/m) , in Figure 9(b). For all experimented values of k and ρ , (t/m) drastically reduced as $(n/\rho/k)$ increased from 1 to 2. An increase in the value of $(n/\rho/k)$ beyond 3 hardly improved the processing cost per complete node. In the interests of keeping the LNVD in a small and manageable size, we use 3 as our value of $(n/\rho/k)$. That is, we set the SR size such that the number n of nodes is

$$n = 3 \cdot \rho \cdot k. \quad (1)$$

Another measure we use to display the behavior of Algorithm 2 is the completion ratio ($m : n$), i.e., the number of complete nodes divided by the number n of nodes inside the SR. Figure 9(c) displays the effect of $(n/\rho/k)$ on the completion ratio. The figure shows that the completion ratio increases in a diminishing rate as $(n/\rho/k)$ increases. The results from different values of k and ρ produce approximately the same curve. As a result, Equation 1 effectively sets the completion ratio to approximately 35%

Since we set the number n of nodes inside SR to $(3 \cdot \rho \cdot k)$, we can express the k LNVD construction cost as

$$\mathcal{O}(\rho k^2 \log \rho k^2).$$

As earlier discussed, for this value of n , the number m of complete nodes is 35% of n , i.e., $m = (0.35 \cdot 3 \cdot \rho \cdot k)$. Hence, we can express the cost per complete node as

$$\mathcal{O}(k \log \rho k^2).$$

Since only some of the m complete nodes in SR are visited by the

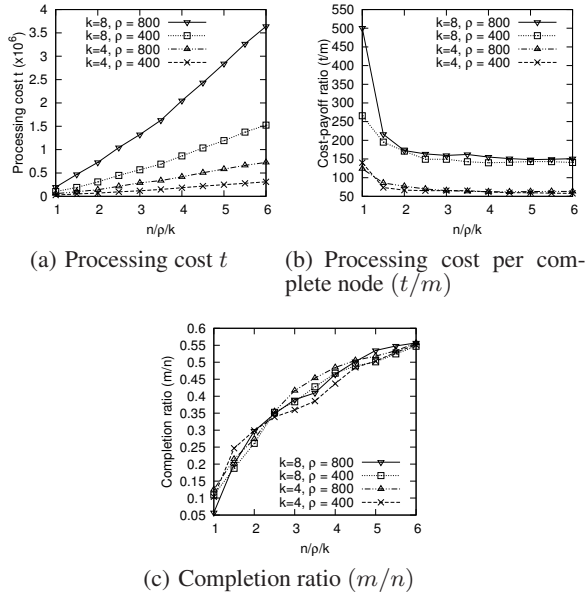


Figure 9: k LNVD construction costs and payoff

query point, a more accurate comparison can be done by measuring the performance of the two methods when used to process an actual moving k NN query. Specifically, the comparative cost of the two methods depends on the number n_q of nodes encountered by the query point per each time a k LNVD is constructed. For example, in a planar graph, we can assume that n_q is proportional to \sqrt{m} , which results in a cost per visited node of

$$\mathcal{O}\left(k\sqrt{\rho k} \log \rho k^2\right).$$

In comparison to the CC- k NN's processing cost of $\mathcal{O}(\rho k \log \rho k)$, we can see that k LNVD is more suitable to a setting of a small k value and high node density ρ . We argue that a setting of large ρ and small k is a valid assumption in actual application environments through the following reasons. First, a mobile user who is driving or walking can pay attention to only a small number k of resultant objects. For example, most navigation devices display only 5 to 10 NNs at a time. Second, a high number of nodes is required to produce an accurate model of a road network. For examples, eight nodes are required to model an intersection where right turns (assuming a left-hand traffic) are not permitted, and four nodes are required to model a road u-turn. To provide an insight into the behaviors of the two methods in a more realistic query processing environment, we report results of our experimental studies in the next section.

8.2 Performance Comparison

In this section, we study the performance of our proposed method in comparison to CC- k NN [1]. We used two trajectory types: *single destination (SD)* and *multiple destination (MD)*.

- An SD trajectory represents a user travelling along a shortest path from a start location and a destination, e.g., a user travelling interstate.
- An MD trajectory represents a case of a user changing their destination every 300 units, e.g., a delivery truck driver dropping off goods at multiple locations.

We generated 20 trajectories of each type and experimental results are reported as average of these trajectories. The trajectory length is fixed to 3,000 units, i.e., 1.61% of the network diameter. Experi-

ments were conducted to study the effect of two parameters:

- *The k value*: ranged from 2 to 10 with a default of 6.
- *The density ρ of nodes with respect to objects*: ranged from 400 to 2,000 with a default of 1,200.

The aim of this set of experiments is to test the validity of the cost per visited node of the k LNVD, $\mathcal{O}(k\sqrt{\rho k} \log \rho k^2)$, and the cost of CC- k NN, $\mathcal{O}(\rho k \log \rho k)$, derived in Section 8.1 (also recall the analysis in Section 7.) The processing cost can be given as the product of the number of iterations and the cost per iteration. The cost per iteration is the logarithm of priority queue size. For k LNVD, $\mathcal{O}(k\sqrt{\rho k})$ represents the number of iterations and $\mathcal{O}(\rho k^2)$ represents the priority queue size. For CC- k NN, both number of iterations and priority queue size are $\mathcal{O}(\rho k)$. As we vary the parameters k and ρ , we expect the behavior of both algorithms to conform with this analysis.

For each parameter, the two methods are compared using the following cost measures.

- *Total time*: the time taken to process a trajectory.
- *Processing cost*: the product of: (i) the number of iterations of the *while-loop* in Algorithm 2, and (ii) the logarithm of the maximum priority queue size.
- *Number of queries*: The number of times the NN query or the k LNVD construction method is executed.⁴

Total time. Figure 10 shows that the total times of both methods have positive correlations with ρ and k , which conform with the cost analysis given in Section 8.1. The figure also shows that SD trajectories are more expensive to process than MD, since MD trajectories are more likely to remain in the same area than SD. Our proposed method, k LNVD, consistently outperforms CC- k NN by approximately one order of magnitude. The *breakdown* of the total time is given in the next paragraph.

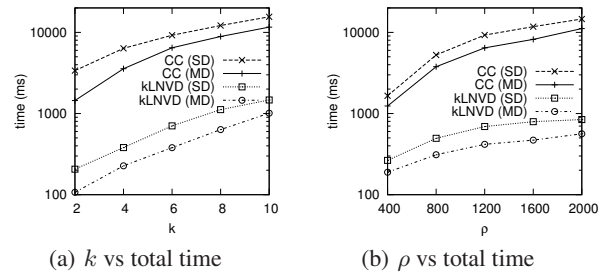


Figure 10: Effects of k and ρ on the total time

Processing cost. We define the processing cost as the product of the number of *while-loop* iterations and the cost per iteration (logarithm of the priority queue size). Figure 11 shows the processing cost and these two components with respect to different values of k and ρ . As suggested by the analysis, Figure 11(a) shows that an increase in k provides a more drastic effect on the number of iterations than CC- k NN and Figure 11(b) shows that an increase in ρ produces a drastic increase in the number of iterations of CC- k NN and a moderate increase for k LNVD. The number of iterations of k LNVD in both figures are significant lower than that of CC- k NN.

In terms of the priority queue size (Figures 11(c) and 11(d)), CC- k NN has a significantly smaller priority queue than k LNVD. This is because k LNVD construction involves more nodes and labels per node than CC- k NN. As suggested by the cost analysis, the priority queue size of k LNVD drastically increases as k increases, while an increase in ρ produces an increase in the same rate as CC- k NN.

The overall processing cost (Figures 11(e) and 11(f)) is approx-

⁴This measure is indicative of the number of communications when data objects are retrieved from a remote location.

imately the product of the number of iterations (Figures 11(a) and 11(b)) and the logarithm of the priority queue size (Figures 11(c) and 11(d)). Hence, the number of iterations has a greater effect on the processing cost than the priority queue size. Figures 11(e) and 11(f) also shows that k LNVD significantly outperforms CC- k NN.

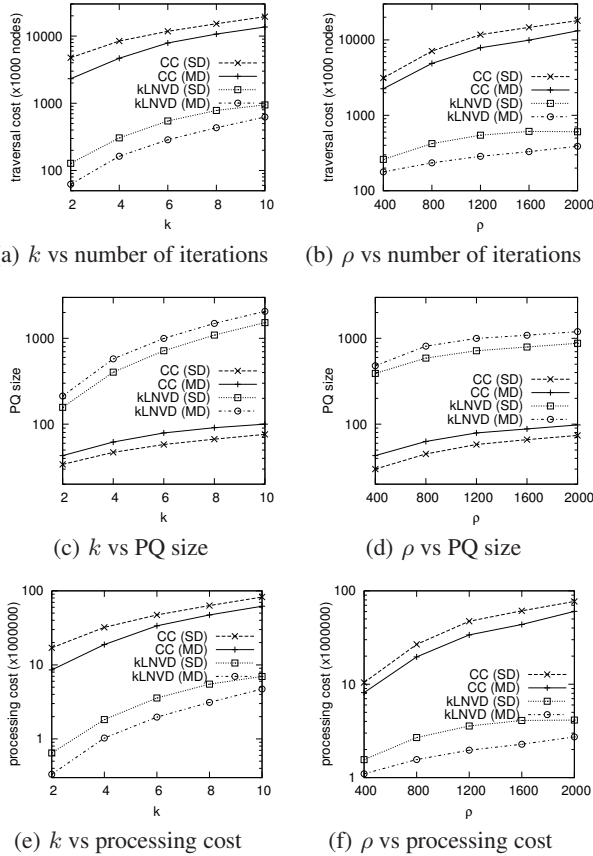


Figure 11: Effects of k and ρ on the processing cost

Number of retrievals. The number of retrievals represents the number of queries executed to retrieve objects from the database, i.e., the number of k LNVD constructions, and the number of visited nodes (discounting the repeats) for CC- k NN. Figure 12 shows that the number of retrievals for k LNVD reduces as k and ρ increase, since the number m of complete nodes is given as $(0.35 \cdot 3 \cdot \rho \cdot k)$. The number of retrievals for CC- k NN is determined by the number of visited nodes. Hence, the parameters ρ and k have no effect on this measure. We can also see that k LNVD continues to outperform CC- k NN significantly.

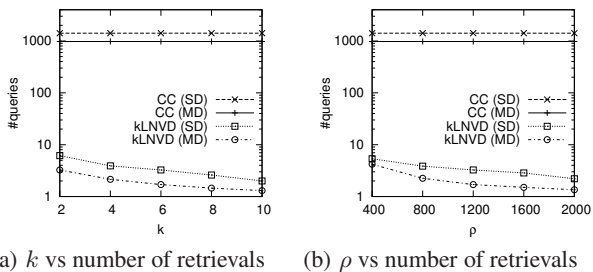


Figure 12: Effects of k and ρ on the number of retrievals

8.3 Summary

We show that the number n of nodes in SR should be set to $(3 \cdot \rho \cdot k)$ to obtain a reasonable completion ratio (Figure 9(c)) and to keep the k LNVD size small and manageable. We continue to show that by setting n to this value, k LNVD is more scalable than CC- k NN as ρ increases and outperforms CC- k NN for all experimental values of k . We take advantage of the fact that the best-first priority queue processing cost is linear with respect to the number of iterations (traversal operations) and logarithmic with respect to the number of entries. Hence, by having a larger priority queue but less traversal operations, the overall processing cost is improved.

9. CONCLUSIONS

We have proposed a method to process *moving k nearest neighbor (Mk NN)* queries which requires neither preprocessing nor repetitive distance evaluation. The complexity analysis and experimental results have shown that our proposed method, *order- k local network Voronoi diagram (k LNVD)* scales better than CC- k NN [1] as ρ increases. Specifically, k LNVD significantly outperforms CC- k NN in terms of the total time, traversal cost, processing cost and the number of data retrievals. As future work, we plan to apply the local computation of Voronoi diagrams to different domains, e.g., the Euclidean and other normed vector spaces.

Acknowledgement. This work is partially supported under the Australian Research Council's Discovery funding scheme (project number DP0880215).

10. REFERENCES

- [1] H.-J. Cho and C.-W. Chung. An efficient and scalable approach to CNN queries in a road network. In *VLDB*, pages 865–876, 2005.
- [2] U. Demiryurek, F. B. Kashani, and C. Shahabi. Efficient continuous nearest neighbor query in spatial networks using Euclidean restriction. In *SSTD*, pages 25–43, 2009.
- [3] K. Deng, X. Zhou, H. T. Shen, S. W. Sadiq, and X. Li. Instance optimal query processing in spatial networks. *VLDB J.*, 18(3):675–693, 2009.
- [4] E. W. Dijkstra. A note on two problems in connection with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [5] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987.
- [6] G. R. Hjaltason and H. Samet. Distance browsing in spatial databases. *ACM Trans. Database Syst.*, 24(2):265–318, 1999.
- [7] M. R. Kolahdouzan and C. Shahabi. Voronoi-based k nearest neighbor search for spatial network databases. In *VLDB*, pages 840–851, 2004.
- [8] M. R. Kolahdouzan and C. Shahabi. Alternative solutions for continuous k nearest neighbor queries in spatial network databases. *Geoinformatica*, 9(4):321–341, 2005.
- [9] K. Mouratidis, M. L. Yiu, D. Papadias, and N. Mamoulis. Continuous nearest neighbor monitoring in road networks. In *VLDB*, pages 43–54, 2006.
- [10] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley, Chichester, second edition, 2000.
- [11] A. Okabe, T. Satoh, T. Furuta, A. Suzuki, and K. Okano. Generalized network Voronoi diagrams: Concepts, computational methods, and applications. *International Journal of Geographical Information Science*, 22(9):965–994, 2008.
- [12] D. Papadias, J. Zhang, N. Mamoulis, and Y. Tao. Query processing in spatial network databases. In *VLDB*, pages 802–813, 2003.
- [13] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, second edition, 2003.
- [14] H. Samet, J. Sankaranarayanan, and H. Alborzi. Scalable network distance browsing in spatial databases. In *SIGMOD*, pages 43–54, 2008.