

Traffic Data Management
for Advanced Driver
Information Systems

CTS
TE
228.3
.S53
1995



Technical Report Documentation Page

1. Report No. MN/RC - 95/22	2.	3. Recipient's Accession No.	
4. Title and Subtitle TRAFFIC DATA MANAGEMENT FOR ADVANCED DRIVER INFORMATION SYSTEMS		5. Report Date May 1995	
		6.	
7. Author(s) Shashi Shekhar		8. Performing Organization Report No.	
9. Performing Organization Name and Address Computer Science Department University of Minnesota Minneapolis MN 55455		10. Project/Task/Work Unit No.	
		11. Contract (C) or Grant (G) No. (C) 71784 TOC# 133 (G)	
12. Sponsoring Organization Name and Address Minnesota Department of Transportation 395 John Ireland Boulevard St. Paul Minnesota, 55155		13. Type of Report and Period Covered Final Report 1994-95	
		14. Sponsoring Agency Code	
15. Supplementary Notes			
16. Abstract (Limit: 200 words) <p>Advanced Traveler Information Systems (ATIS) offer the potential to help a driver find the quickest and safest route to a destination. An effective navigation system requires effective route planning services, which need to provide three facilities: route computation, route evaluation, and route display. This project focuses on route planning algorithms for ATIS. The cost models and performance studies in this report show that single-pair algorithms can outperform traditional algorithms in many situations.</p>			
17. Document Analysis/Descriptors Advanced Traveler Information Systems (ATIS) Route Computation		18. Availability Statement Advanced Driver Information Systems (ADIS) Routing Algorithm	
19. Security Class (this report) Unclassified		20. Security Class (this page) Unclassified	21. No. of Pages 40
		22. Price	

Traffic Data Management for Advanced Driver Information Systems

Final Report

Prepared by

Shashi Shekhar

Computer Science Department
University of Minnesota
Minneapolis MN 55455

May 1995

Published by

Minnesota Department of Transportation
Office of Research Administration
200 Ford Building Mail Stop 330
117 University Avenue
St Paul Minnesota 55155

The contents of this report reflect the views of the author who are responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the views or policies of the Minnesota Department of Transportation at the time of publication. This report does not constitute a standard, specification, or regulation.

The author and the Minnesota Department of Transportation do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to this report.

Acknowledgments

This work was supported by the Minnesota Department of Transportation, the Center for Transportation Studies at the University of Minnesota, the Federal Highway Authority, and the University of Minnesota Undergraduate Research Opportunities Program. We would like to thank Andrew Fetterer, Mark Coyle, Ashim Kohli, Jeff Johnson, and Maryam Kalantar for their contributions. Finally, we would like to thank the project review committee members, namely Marthand Nookala, Ron Dahl, and Jeff Southward, for their insightful comments.

Table of Contents

1. Introduction	1
1.1. Related Work and Our Contributions	2
1.2. Problem Definition and Scope of Work	4
1.3. Outline of the Paper	5
2. Single Pair Algorithms	7
2.1. Iterative Algorithm	7
2.2. Dijkstra's Algorithm	8
2.3. Best First A* Algorithms	9
3. Database Implementation	11
3.1. Cost Analysis	11
3.2. Experimental Analysis of Database Performance	16
4. Main Memory Implementation	29
4.1. Data Structure and User Interface	29
4.2. Results	30
5. Conclusions and Future Work	35
References	37

List of Figures

Figure 1.	Iterative Algorithm	7
Figure 2.	Dijkstra's Algorithm	8
Figure 3.	A* Algorithm	9
Figure 4.	Synthetic graphs and node pairs	17
Figure 5.	Effect of graph size on execution time	18
Figure 6.	Effect of Path Length on Execution Time	19
Figure 7.	Effect of Edge Cost Models on Execution Time	21
Figure 8.	Minneapolis Road Map	22
Figure 9.	Minneapolis Road Map Results	23
Figure 10.	Effect of Graph Size on Execution Time of Versions	25
Figure 11.	Effect of Varying Edge Cost on Execution Time of Versions	26
Figure 12.	Effect of Path on Execution Time of Versions	27
Figure 13.	Sample Menus for the Graphical User Interface	30
Figure 14.	Result of Travel Distance Experiment	31
Figure 15.	Result of Travel Time Experiment	32

List of Tables

Table 1.	Notations Used in Cost Analysis	12
Table 2.	Cost of Iterative BFS Algorithm	13
Table 3.	Cost Model for Dijkstra's and A* (version 3)	14
Table 4A.	Parameter Values	15
Table 4B.	Estimated costs, 30x30 Graph, 20% Variance on edge cost	16
Table 5.	Effect of Graph Size on Number of Iterations	19
Table 6.	Effect of Path Length on Number of Iterations	20
Table 7.	Effect of Edge Cost Models on Iterations	21
Table 8.	Effect of Path Length and Orientation on Iterations	23
Table 9.	Effect of Travel Time Heuristic on Iterations	32
Table 10.	Effect of Travel Distance Heuristic on Iterations	33
Table 11.	Percent Speedup from Taking First Path over Optimal Path	33

Executive Summary

This project focuses on route planning algorithms of Advanced Driver Information Systems, which are also known as Advanced Traveler Information Systems (ATIS). Route planning has been approached by graph-theoretic algorithms for all-pair (transitive closure) and single-source (partial transitive) path computations. These algorithms compute many more paths beyond the single-pair path that is of interest to ATIS, and hence may not be satisfactory for ATIS due to the dynamic nature of edge costs (travel-time). For single-pair path computation, we explore specialized algorithms such as A*, which are designed to reduce irrelevant computation and to quickly discover the shortest paths. Our cost models and performance studies show that single-pair algorithms can outperform traditional algorithms in many situations.

1. Introduction

Advanced Traveler Information Systems (ATIS) assist travelers with planning, perception, analysis and decision making to improve the convenience, safety and efficiency of travel [1]. ATIS is one facet of the Intelligent Highway Vehicle System (IVHS), which is currently being developed to improve the safety and efficiency of automobile travel. Route planning is an essential component of ATIS. It aids travelers in choosing the optimal path to their destinations in terms of travel distance, travel time and other criteria. Estimates made by government agencies show that approximately 6% of all driving time in the U.S. is due to incorrect choice of routes [2]. An effective navigation system with static route selection, coupled with real-time traffic information, is crucial in eliminating unnecessary travel time. Reducing vehicles' exposure to congestion also reduces their exposure to accidents, reduces pollution, and allows efficient calculation of routes. Furthermore, the commercial vehicle sector will derive additional advantages, including lower transportation costs and less delay at checkpoints.

Route planning services need to provide three facilities: route computation, route evaluation and route display. The goal of route planning is to locate a connected sequence of road segments from a current location to a destination. Route computation may be based on criteria such as the shortest travel distance or travel time. Route computation is also useful for travel during rush hour, travel in unfamiliar areas, and/or travel to an unfamiliar destination. The goal of route evaluation is to find the attributes of a given route between two points. These attributes may include travel time and traffic congestion information, and thus route evaluation is also useful for selecting travel time by a familiar path. The goal of route display is to effectively communicate the optimal route to the traveler for navigation. In this paper, we focus on route computation and also examine algorithms for real-time route computation on maps stored in a database.

1.1. Related Work and Our Contributions

The single-pair path computation problem is a special case of single-source path computation and all-pair path computation. Traditional research in database query languages[3-5], transitive closure [6-17], and recursive query processing[18-21] has approached single-pair path computation as a special case of more general problems. For example, partial transitive closure computation[22] and transitive closures [8] have been used for single-pair path computations. Previous evaluation of the transitive closure algorithms examined the iterative, logarithmic, Warren's, Depth first search (DFS), hybrid, and spanning-tree-based algorithms. [6, 10, 12, 13]. This study was based on implementing algorithms in Fortran, Quel* and Ingres. The study did not examine A* and other estimator-based algorithms, and did not examine the effect of path length and edge costs on the relative performance of search algorithms. In applications such as ATIS, the desired route will vary in length and rarely will traversal of the full map be necessary. Iterative algorithms need to examine all nodes in the graph to find the shortest path between any pair of nodes. They must carry out the same amount of work irrespective of the path length. There is a need for a more thorough study of algorithms for single-pair path computations.

The Advanced Mobile Traffic Information and Communication System (AMTICS)[23] is an integrated traffic information and navigation system. This system displays traffic information gathered by traffic control and surveillance centers by means of a screen in each vehicle. The traffic information is processed by a computer and broadcast to vehicles on a radio data communication system called *teleterminal*. Other wide area radio-frequency data communication systems have also been considered[24]. The on-board equipment consists of a display, a CD-ROM reader for retrieving map information stored on CDs and a microprocessor to calculate the vehicle's position and to superimpose it on the display. The system is able to display the vehicle's current position, route, traffic congestion, regulation, road work, and parking. Some other systems that are commercially available and provide comparable features include the *Etak Navigator*, the *Bosch Travelpilot*, and the *ALI-SCOUT* infrastructure-supported route guidance system[25, 26].

Some of the work done in storage and display of precomputed shorted routes includes the TRRL Navigator[27]. The data linking the major intersections in the road network with the shortest route is stored and used to guide the driver from intersection to intersection to the destination. The system responds to vehicle requests by means of either a simple route listing on a

TV screen or a digitally coded cassette to be read by an in-car unit with an LCD display. Successive instructions are displayed as needed at the push of a button. A further development of the system will be able to detect the position of the vehicle by means of inductive coupling loop detectors linked to road-side units and thereby enabling automatic dispatch of successive instructions.

One of the requirements of a good ADIS is an easy user interface. Some research effort has been directed in this area. Standard text based queries such as name or subject as well as geographic queries such as proximity to a location have been studied[28]. Some studies have been done on the spatial data infrastructure, the map navigation software, and the design and validation of headup displays for navigation[29]. The database used by pathfinding is a topological connectivity network where streets between intersections are edges and intersections are nodes on the adjacency graph. Edges have properties such as driving restrictions, distance, road classification, carrying capacity etc[26].

Work on in-vehicle real time traffic information systems includes the pathfinder project[30], the Detroit transportation center transit information project[31]. The pathfinder is an information and navigation system designed to test the feasibility of using the latest technological devices to aid motorists in avoiding adverse traffic conditions[32, 33]. Sensor data management and rapid-detection-and-control systems for incident management have been studied in [34-36] and as part of the Traffic Reporter project[37].

In this paper, we evaluate three path-planning algorithms for single-pair path computation. These algorithms are the iterative breadth-first search, Dijkstra's single-source path-planning algorithm, and the A* single-path planning algorithm. These algorithms represent three classes of algorithms, namely transitive closure, partial transitive closure, and single-pair path planning. The algorithms have been chosen for their simplicity and for their representation of the essential features of the algorithms in their class. The iterative algorithm does not utilize path-length information to reduce the work required for single-pair path computation, as is typical of most transitive-closure algorithms. Dijkstra's algorithm has limited lookahead of one edge during searching, and therefore, like most partial-transitive closure algorithms, it is not able to focus the search in the direction of the destination. A* represents the single-pair path-planning algorithms which use heuristic lookaheads to focus the search.

We also evaluate the performance of the algorithms on graphs representing the roadmap of Minneapolis. In order to get an insight into their relative performance, we also use synthetic grid maps as a benchmark computation. We examine the effect of two parameters, namely path length and edge-cost distribution, on the performance of the algorithms. Finally, we examine the effect of implementation decisions on the performance of the A* algorithm. The main hypothesis is that estimator functions can improve the average-case performance of single-pair path computation when the length of the path is small compared to the diameter of the graph. We examine this hypothesis via experimental studies and analytical cost modeling (studies include a grid graph as well).

1.2. Problem Definition and Scope of Work

A (directed) graph $G = (N, E, C)$ consists of a node set N , a cost set C , and an edge set E . The edge set E is a subset of the cross product $N * N$. Each element (u, v) in E is an edge that joins node u to node v . Each edge (u, v) is associated with a cost $C(u, v)$. Cost $C(u, v)$ takes values from the set of real numbers. A node v is a neighbor of node u if edge (u, v) is in E . The degree of a node is the number of neighboring nodes. A path in a graph from a source node s to a destination node d is a sequence of nodes $(v_0, v_1, v_2, \dots, v_k)$ where $s = v_0$, $d = v_k$, and the edges (v_0, v_1) , (v_1, v_2) , ..., (v_{k-1}, v_k) are present in E . The cost of the path is the sum of the cost of the edges, i.e. $\sum_{i=1}^k C(v_{i-1}, v_i)$. An optimal path from node u to node v is the path with the smallest cost. The algorithms used two measures of cost: shortest travel distance and shortest travel time. The cost for the shortest distance paths is the euclidean distance between the nodes, and the cost for the shortest travel time paths will be the average travel time across the edge.

We investigated the suitability of applying traditional path computations algorithms to Advanced Driver Information Systems. The algorithms were tested in a database and in main memory. We developed a graphical user interface (GUI) to aid in algorithm evaluation and to verify experimental accuracy.

1.3. Outline of the Paper

Section 2 defines the problems and introduces the terminology and notation used in the paper. Section 3 lists the candidate algorithms for single-pair path computation. Section 4 provides analytical cost modeling of the algorithms. Section 5 describes the experiments and observations to evaluate the effectiveness of estimator functions and implementation decisions. Section 6 presents the conclusions and future work.

2. Single Pair Algorithms

2.1. Iterative Algorithm

This is one of the oldest algorithms for transitive closure, all-pair path computation, and graph traversal[38], and is also known as the breadth-first search algorithm. The iterative algorithm can be described by the psuedo-code shown in Figure 1.

```
Procedure Iterative( N, E, s, d);
begin
foreach u in N do { C(s,u) = ∞; C(u,u) = 0; path(u,v) := null;}
frontierSet := [ s ];
while not_empty(frontierSet) do
{ foreach u in frontierSet do
  { frontierSet := frontierSet - [u];
    fetch( u.adjacencyList);
    foreach <v, C(u,v)> in u.adjacencyList
      { if C(s,v) > C(s,u) + C(u,v) then
        { C(s,v) := C(s,u) + C(u,v); path(s,v) := path(s,u) + edge(u,v);
          if not_in(v, frontierSet) then frontierSet := frontierSet + [v];
        }
      }
    }
}
end;
```

Figure 1. Iterative Algorithm

Lemma 1: Procedure Iterative(N,E,s,d) finds the shortest path between nodes s and d, if the edge costs C(u,v) are non-negative.

The optimality property, coupled with the termination condition of the procedure, guarantees discovery of the shortest path between s and d. The procedure terminates when the frontier-Set becomes empty: i.e. the ExploredSet contains all the nodes which have a path from s. This completes the proof.

We note that the iterative algorithm cannot be terminated before exploring the entire graph to find the shortest path between two nodes. It cannot reduce its work for paths with only a few edges, because its execution time is not sensitive to path length.

2.2. Dijkstra's Algorithm

Dijkstra's algorithm has been influential in path computation research[22], and has been applied to transportation planning. It provides competent worst-case I/O cost for partial transitive closure and single-pair path-computation problem[22]. The psuedo-code in Figure 2 describes the algorithm.

```
Procedure Dijkstra( N, E, s, d);
begin
foreach u in N do { C(s,u) = ∞; C(u,u) = 0; path(u,v) := null;}
frontierSet := [ s ]; exploredSet := emptySet;
while not_empty(frontierSet) do
{ select u from frontierSet with minimum C(s, u);
  frontierSet := frontierSet - [u]; exploredSet := exploredSet + [u];
  if (u = d) then terminate
  else { fetch( u.adjacencyList);
        foreach <v, C(u,v)> in u.adjacencyList
          if C(s,v) > C(s,u) + C(u,v) then
            { C(s,v) := C(s,u) + C(u,v); path(s,v) := path(s,u) + (u,v);
              if not_in(v, frontierSet U exploredSet) then
                frontierSet := frontierSet + [v];
            }
          }
        }
}
end;
```

Figure 2: Dijkstra's algorithm

Lemma 2: Assuming that the graph does not contain negative edge costs, then as soon as the destination node d in the frontierSet is selected and removed, its path(s,d) contains the shortest path from source node s to d .

The proof for this lemma is provided in [22]. We note that the procedure terminates after the iteration which selects destination node d as the best node in the frontierSet. The procedure can terminate quickly if the shortest path from s to d has few edges, since in many cases it does not have to examine all the nodes to discover the shortest path.

2.3. Best First A* Algorithms

The best-first search has been a framework for heuristics which speed up algorithms by using semantic information about a domain and has been explored in database contexts for single-pair path computation[39]. A* is a special case of the best-first search algorithms. It uses an estimator function $f(u,d)$ to estimate the cost of the shortest path between node u and d . A* has been quite influential due to its optimality properties[40]. The basic steps of these algorithms are described in Figure 3. A best-first search without estimator functions is not very different from Dijkstra's algorithm.

```
Procedure A*( N, E, s, d, f);
begin
foreach u in N do { C(s,u) = ∞; C(u,u) = 0; path(u,v) := null;}
frontierSet := [ s ]; exploredSet := emptySet;
while not_empty(frontierSet) do
{ select u from frontierSet with minimum ( C(s,u) + f(u,d) ); **
  frontierSet := frontierSet - [u]; exploredSet := exploredSet + [u];
  if (u = d) then terminate
  else { fetch (u.adjacencyList);
        foreach <v, C(u,v)> in u.adjacencyList
        { if C(s,v) > C(s,u) + C(u,v) then
          { C(s,v) := C(s,u) + C(u,v); path(s,v) := path(s,u) + (u,v);
            if not_in(v, frontierSet) then frontierSet := frontierSet + [v];
          }
        }
      }
}
end;
```

Figure 3: A* Algorithm

Lemma 3: Assuming that the graph does not contain negative edge costs and that the estimator function never overestimates the cost of the shortest path(u,d), then as soon as the destination node d ($f(d,d) = 0$) in the frontierSet is selected and removed by A, its path(s,d) contains the shortest path from source node s to d .*

The proof for Lemma 3 is provided in[41].

We note that the procedure terminates after the iteration which selects destination node d as the best node in the frontierSet. The procedure can terminate quickly if the shortest path from s to d has few edges. It does not have to examine all the nodes to discover the shortest path, in many cases. Furthermore, the estimator can provide extra information to focus the search on the shortest path to the destination, reducing the number of nodes to be examined.

3. Database Implementation

3.1. Cost Analysis

We discuss the cost model and derive the cost formulas for the algorithms in this section. The analysis is based on the following representation of the data-structures used by various algorithms. Directed graphs are represented as pairs of relations: edge (S) and node (R). The edge relation S is a read-only relation and it stores the edges of the graph along with their costs. Its fields include: Begin-node, End-node, and Edge-cost. The node relation R stores the internal data-structures of various routing algorithms. Its fields include: node-id, x-coordinate, y-coordinate, status, a path to the source node, and path-cost. The fields Begin-node and End-node of S contain the values of the node-ids from the node relation R of the graph to represent an edge in the graph, and the edge-cost contains the cost of traversing this edge. The neighbors of a node v , i.e. its adjacency list, can be found by retrieving the edges with $S.Begin-node = v$. An undirected graph can be represented by storing two directed-edge entries in S for each undirected edge in the graph. The relation S has a primary index (random hash) on the field S.Begin-node. Each tuple in R represents a node with status and path attributes. The status field is used to represent node lists such as frontierSet and exploredSet. Nodes with status = open represent the frontierSet. Nodes with status = closed represent the exploredSet. Node(s) with status = current represent the current node(s) being explored. Nodes with status = null are ones that are not open, "closed" or current. The path field in R points to a neighboring node on the best path to the source node. The complete path to the source node can be constructed by traversing this pointer, starting at the destination node. The path-cost field for a node n represents the total cost of the path from n to the source node. The relation R has a primary index (ISAM) on the node-id. Furthermore, we will choose between four join algorithms for the joins computed by the algorithms.

Deciding how to best manage duplicates in the frontierSet is important. It can be done in three ways: avoiding duplicates, removing duplicates, or allowing duplicates. Allowing duplicates leads to redundant iterations of the algorithm. Duplicates can be avoided by making sure that the status of the node is null before adding it to the frontierSet. Duplicates can also be eliminated after insertion in frontierSet by duplication-elimination algorithms, but we prefer

duplicate avoidance for its cost effectiveness.

Although the algebraic cost formulas do not account for system overhead, they still provide us with an insight into the relative performance of the algorithms. In order to account for other detailed costs, we implemented a query optimizer simulation in C which was able to choose between several Select and Join strategies and produce very accurate approximations of the I/O cost of the algorithms, based on varying parameters. Using the simulation, we were able to verify the validity of the experimental results obtained by measuring the execution-times of the implementations of various algorithms on the INGRES DBMS. We use the notation in Table 1 to derive the cost models.

Symbol	Meaning
S	Edge relation = (Begin-node, End-node, Edge-cost)
R	Temporary Node relation = (Node-id, x-coordinate, y-coordinate, status, path, path-cost)
JOIN	Paths to neighbors of current node(s)
L	Path Length, Number of edges in path
I	I/O cost of creating a temporary relation
I_l	ISAM index level
s_r	Selection cardinality of nodes in R
A	Average Number of nodes in the adjacency list (neighbors)
n	Number of nodes in the Graph
S	Number of tuples in S
R	Number of tuples in R
C	Number of nodes in R that are marked current
JS	Join selectivity = $\frac{ S \text{ join } R }{ S * R } = \frac{ JOIN }{ S * R }$
D_t	Cost for deleting all tuples in a relation t
B	Disk block size (in bytes)
T_s	Tuple size of the source relation (edge-relation) S
T_r	Tuple size of the resultant relation (node-relation) R
Bf_s	Blocking factor for S = B / T_s
Bf_r	Blocking factor for R = B / T_r
Bf_{rs}	Blocking factor for R x S = $B / (T_r + T_s)$
B_s	$ S / Bf_s$ Number of blocks of source relation S
B_r	$ R / Bf_r$ Number of blocks of resultant relation R
B_{join}	Number of blocks of JOIN
B_c	$ C / Bf_r$ Number of blocks of current-nodes
t_{read}	Time for reading one block from disk
t_{write}	Time for writing one block to disk
t_{update}	Time for updating one tuple ($t_{read} + t_{write}$)
C_j	Cost of step j

Table 1: Notations Used in Cost Analysis

Cost Model of the Iterative Algorithm

The Iterative algorithm can be broken into eight steps for cost modeling. The steps and their associated costs are shown in table 2.

Cost	Steps of Iterative Algorithm
$C_1 = I$	-1. Creating the resultant relation R (node-relation)
$C_2 = B_s * t_{read} + B_r * t_{write}$	-2. Initializing R with all nodes in S
$C_3 = 2 * (B_r * \log(B_r) + B_r) * t_{update}$	-3. Indexing and Sorting the node-relation by node-name
$C_4 = (I_L + S_r) * t_{update} + B_r * t_{read}$	-4. Mark start node in R as "current" and count current-nodes
For each iteration i:	
$C_5^i = B_r * t_{read}$	-5. Fetch all current-nodes from R
$C_6^i = F(B_c, B_s, B_{join})$	-6. Perform a Join to get the neighbors of all current-nodes
$C_7^i = 2 * B_r * t_{update}$	-7. Update status and path of nodes in R (mark explored nodes as "closed")
$C_8^i = B_r * t_{read}$	-8. Scan R to count the number of current-nodes
Cost of iteration i = $\Gamma_i = C_5^i + C_6^i + C_7^i + C_8^i$	
Total Cost $T = C_1 + C_2 + C_3 + C_4 + \sum_{i=1}^{B(L)} \Gamma_i$	

Table 2: Cost of Iterative BFS Algorithm

Each iteration of the iterative algorithm is composed of steps 5,6,7 and 8. Step 6 performs a join which helps in fetching the adjacency lists of all the current nodes. The cost of this join is described by a function $F(B_1, B_2, B_3)$. The input parameters for F are the sizes (in blocks) of the relations to be joined and the size of the resultant relation. The value of this function depends on the join strategy that is chosen to carry out the join. The function uses the input parameters to choose the cheapest join strategy from among four viable choices: (1) Hash Join, (2) Nested-Loop Join, (3) SortMerge Join, and (4) Primary Key Join. In the case of the Iterative algorithm, the size of the resultant relation is estimated as $B_{join} = (JS * |C| * |S|) / Bf_{rs}$. Step 8 scans the relation R to determine the size of the current node-list. The algorithm terminates if the current node-list is empty. The number of iterations taken by the BFS algorithm can be expressed as a function $B(L)$, which is dependent on several factors such as the start node and the graph diameter. It is difficult to predict the number of current nodes at any point in time because they are dependent

on the amount of backtracking performed by this algorithm. This backtracking is a dynamic factor dependent on the edge costs. On the average, the number of current nodes ($|C|$) per iteration before Step 6 is estimated to be $|R|/B(L)$, if there is no backtracking at all. The average number of tuples (paths) after the Join in Step 6 is estimated to be $|S|/B(L)$. The join selectivity (JS) is estimated to be $\frac{|S|/B(L)}{|R|/B(L) * |S|} = \frac{1}{|R|}$. The total cost can be approximated as

$T = C_1 + C_2 + C_3 + C_4 + B(L) * \Gamma_{average}$, where $\Gamma_{average}$ is the average cost per iteration.

Cost Models for Dijkstra's and A* Algorithms

Cost	Steps of Dijkstra's / A* Algorithms
$C_1 = I$	-1. Creating the resultant relation R
$C_2 = B_s * t_{read} + B_r * t_{write}$	-2. Initializing the node-relation with all nodes in S
$C_3 = 2 * (B_r * \log(B_r) + B_r) * t_{update}$	-3. Indexing and Sorting the node-relation by node-id
$C_4 = (I_L + S_r) * t_{update} + B_r * t_{read}$	-4. Mark start node in R as "current" and count current-nodes
For each iteration i:	
$C_5^i = B_r * t_{read}$	-5. Find the minimum cost of nodes in the frontierSet
$C_6^i = B_r * t_{update}$	-6. Fetch the minimum cost node X from the frontierSet and mark it "current"
$C_7^i = F(B_c, B_s, B_{join})$	-7. Perform a Join to get the neighbors of node X
$C_8^i = (B_r * t_{read} * B_{join} + B_{join} * (t_{read} + t_{write}))$	-8. Scan R and the result of Step 7 and update the status and node-costs in R
$C_9^i = B_r * t_{update}$	-9. Update status of nodes in R by marking explored nodes as "closed"
$C_{10}^i = B_r * t_{read}$	-10. Scan R to count the number of current-nodes
Cost of iteration $i = \Gamma_i = C_5^i + C_6^i + C_7^i + C_8^i + C_9^i + C_{10}^i$	
Total Cost $T = C_1 + C_2 + C_3 + C_4 + \sum_{i=1}^{Z(n,L)} \Gamma_i$	

Table 3: Cost Model for Dijkstra's and A* (version 3)

Dijkstra's algorithm and A* (version 3) can be decomposed into 10 steps for cost analysis. The steps and their associated costs are shown in Table 3. The number of blocks occupied by the paths in JOIN is expressed as $B_{join} = (JS * 1 * |S|) / B_{fs} = |A| / B_{fs}$.

The number of iterations taken by any algorithm is denoted by a function $Z(n,L)$. In general, $Z(n,L)$ for a given algorithm depends on the number of nodes in the graph and the path length from the source to the destination. For long paths, the function $Z(n,L)$ could equal the number of nodes in the graph (n). For short paths, $Z(n,L)$ may be much smaller than the total number of nodes.

In each iteration, A* (version 3) and Dijkstra's perform similar computations. The main difference appears in the selection of the minimum-cost node to expand at each iteration. Dijkstra's minimizes only the actual cost for this selection, whereas A* (version 3) selects the node with the minimum actual + heuristic cost. This makes c_5 and c_6 depend on the algorithm.

Since it is difficult to algebraically predict the number of iterations, we extract it from the trace of the actual execution of the algorithms. The average join selectivity in Step 7 for both these algorithms is $JS = |A|/|S|$, because there can only be one current node at each iteration and it can only have $|A|$ neighbors. The average number of neighbors for each node in a grid graph is four (4). The total cost for A* (version 3) and Dijkstra's can be approximately $T = C_1 + C_2 + C_3 + C_4 + B(L) * \Gamma_{average}$, where $\Gamma_{average}$ is the average cost per iteration.

Example

Using the algebraic cost formulas described in Tables 2 and 3, we can compute the algorithm costs for different path lengths. For illustration, we assume that all the algorithms choose the nested-join approach for Step 7 of the algebraic cost models. Under this assumption, the join function can be expressed as $F(B_1, B_2, B_3) = B_1 * t_{read} + (B_1 * B_2) * t_{read} + B_3 * t_{write}$. Let us also assume certain average values for the size of the JOIN relation, the number of current node(s) $|C|$, the join selectivities for the different algorithms, and also for several other parameters shown in Table 4A.

Parameter	Value	Parameter	Value
I	0.5 units	T_s	32 bytes
I_t	3	T_r	16 bytes
S_r	1	Bf_s	128 records/block
$ A $	4	Bf_r	256 records/block
$ S $	3480 edges	Bf_{rs}	86 records/block
$ R $	900 nodes	t_{read}	0.035 units
D_t	0.5 units	t_{write}	0.05 units
B	4096 bytes	t_{update}	0.085 units

Table 4A: Parameter Values

Let us assume the ratio *Number of nodes in the graph* $|R|$ / *Number of iterations* N of the algorithm as the average number of current-nodes per iteration of the iterative algorithm. The join selectivity for the same algorithm is $JS = 1/|R|$. Therefore, B_{join} for the iterative would be $= |S| / (B(L) * Bf_{rs})$. Since in each iteration of the A* (version 3) and Dijkstra's algorithms, there exists exactly one current node; therefore, the average join selectivity would be $JS =$ *Average number of neighbors of a node* $|A|$ / *Total number of edges* $|S|$. The B_{join} for these algorithms would be $= |A| / Bf_{rs}$. Let us take the number of iterations (refer to Table 5) from the execution trace of the QUEL programs to predict the execution time.

Algorithm / Path	Horizontal	Semi-Diagonal	Diagonal
Dijkstra	1055.6	1656.8	1941.2
A* (version 3)	66.7	881.2	1809.8
Iterative	176.9	176.9	176.9

Table 4B: Estimated costs, 30x30 Graph, 20% Variance on edge cost

Table 4B shows an estimate of the costs incurred by each algorithm on a 30x30 graph. We note the similarity of these estimates to the actual experimental results shown in Figure 6.

3.2. Experimental Analysis of Database Performance

We evaluated the algorithms in two stages. First, the algorithms implemented in EQUER were run on the graphs, and we obtained measurements of processing time. The experiments were repeated a number of times to arrive at average execution times. We also used Ingres in single-user mode to reduce overhead, due to contention for resources. We then designed a simulation using our algebraic cost models to verify that the time measurement was a reasonable comparator. The simulation took the number of iterations from the execution trace of the EQUER programs to predict the execution-time. With our algebraic cost models and simulation, we were able to predict actual execution time within ten percent. We report the results from the EQUER implementation in this section. We also evaluate three different versions of the A* algorithm to determine the effect of different estimator functions and implementation decisions.

Comparative Performance Evaluation in a Synthetic Grid

Grid Benchmark Computation: The benchmark path computation is based on an undirected graph with cycles for single-pair path computations. The node pairs and graphs for benchmark computations are shown in Figure 4. The synthetic graph represents two-dimensional grids with 4 neighbor nodes. The grid includes $k \times k$ nodes, with k nodes along each row and each column, and with an edge connecting adjacent nodes along rows and columns. We chose three node pairs for path computation: diagonally opposite nodes, linearly opposite nodes, and a random-node pair.

Two-dimensional grid graphs are typical of the navigation problems in an unstructured environment. These have been used in prior studies of single-pair path computation[39]. The choice of a synthetic graph also simplifies interpretation of results.

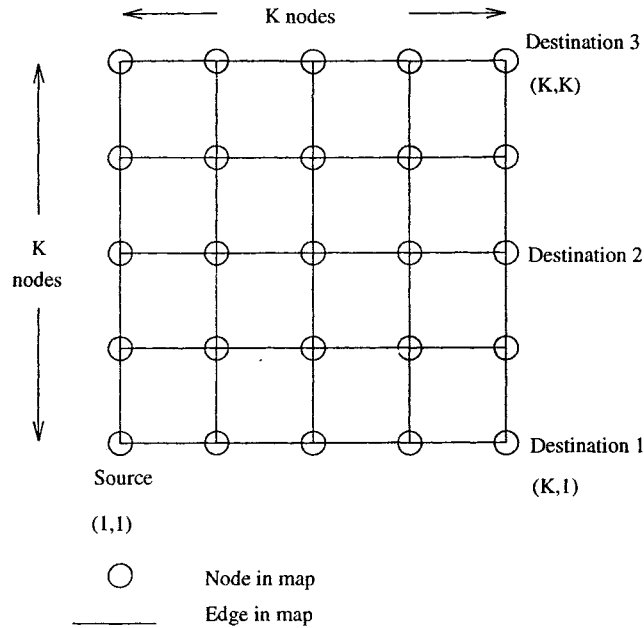


Figure 4: Synthetic graphs and node pairs

Candidate Algorithms: We chose three algorithms: Iterative, Dijkstra’s, and A* (version 3).

Variable Parameters: The parameters for the first set of experiments included graph size, path length, and edge cost models. Experiments were carried out for the following graph sizes: 10*10, 20*20, and 30*30 node grids. We used three cost models for the edges: a uniform cost, a uniform cost with a small variance, and a skewed cost. The uniform-cost model assigns unit cost (i.e. 1) to each edge in the graph. A uniform cost with a 20% variation assigns a cost of $1 + 0.2 * U[0,1]$, where $U[0,1]$ is a random number uniformly distributed between 0 and 1. This cost model will change the degree of backtracking required in the execution of estimator-based algorithms such as A* (version 3). Finally, the skewed-cost model assigns a small cost to the edges $[(1, i), (1, i+1)]$ on the bottom of the grid and the edges $[(k, i), (k, i+1)]$ on the right side of the grid. This model eliminates backtracking from estimator-based A* (version 3), creating the best case for that version.

We have chosen problem instances of grid graphs and parameters, such as graph sizes and path length, based on the literature[13, 39]. The parameter choice of path length is motivated by our application. Other parameters including edge-cost models are exploratory, under the hypothesis that these will impact the performance of estimator-based path computation algorithms significantly.

Effect of Graph Size

The execution time for all the figures in this section represents the sum of the user time and the system time. The algorithms compute the path between diagonally opposite nodes (i.e. the longest path) to compare the worst-case performance of the various algorithms. The performance of A* (version 3), Dijkstra's, and the Iterative BFS is shown by the graph in Figure 5. Table 5 gives the number of iterations for each of the algorithms for different graph sizes.

The graph size significantly affects the performance of A* (version 3) and Dijkstra's algorithms for computing the longest diagonal path. The number of iterations and execution times for A* (version 3) and Dijkstra's grow linearly with the number of nodes in the graph. The iterative algorithm is affected to a lesser extent. Its execution time and the number of iterations to be performed for computing a diagonal path grows sublinearly with the number of nodes in the graph.




Figure 5: Effect of graph size on execution time

Algorithm / Graph Size	10 x 10	20 x 20	30 x 30
Dijkstra	99	399	899
A* (version 3)	85	360	838
Iterative	19	39	59

Table 5: Effect of Graph Size on Number of Iterations
(20% Edge Cost Variance, Diagonal Path)

Effect of Path Length

The effect of path length on the performance of the three algorithms is shown in Figure 6 and Table 6. Figure 6 shows the execution time of the algorithms for different path-length values. Table 6 gives the number of iterations for each of the algorithms for the different path lengths. As we can see in figure 6, A* (version 3) outperforms Iterative and Dijkstra’s algorithm for horizontal paths. For the remaining two path lengths, the Iterative algorithm performs better.

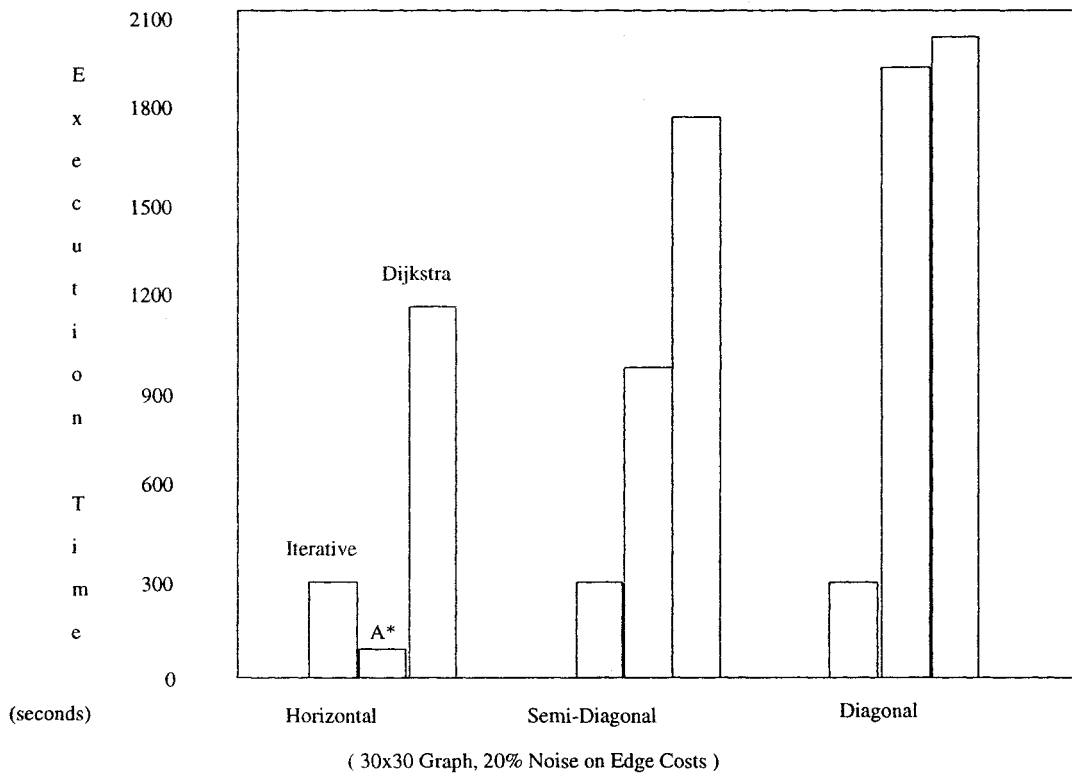


Figure 6: Effect of path length on execution time

Algorithm / Path	Horizontal	Semi-Diagonal	Diagonal
Dijkstra	488	767	899
A* (version 3)	29	407	838
Iterative	59	59	59

Table 6: Effect of Path Length on Number of Iterations
(20% Edge Cost Variance, 30x30 Graph)

Effect of Edge Cost Models

The effect of edge-cost models on the performance of the three algorithms is shown in Figure 7 and Table 7. Figure 7 shows the execution time of the algorithms for three cost models, a uniform cost, edge costs with a 20% variance, and skewed edge-costs. Table 7 gives the number of iterations for each of the algorithms using the different edge cost models.

The execution cost incurred by Iterative BFS depends on the edge-cost model, even though the BFS algorithm does not take advantage of the cost information on the edges. This is due to the possibility of reopening a node and revising the path. Since the remaining two algorithms use edge costs to drive the search, it is expected that they are affected by variations in the edge-cost model. In the 20%-edge-cost variance model, A* (version 3) requires more backtracking (360 iterations vs 85 iterations for uniform cost), and therefore has higher execution times. When the edge costs are skewed (i.e. the edges from the source to the destination have a much lower edge cost), A* (version 3) and Dijkstra’s algorithm perform very well. Backtracking is eliminated as the sub-path generated at every iteration is optimal.

Comparative Performance on the Minneapolis Road Map

The Minneapolis road map data consisted of 1089 nodes and 3300 edges that represented highway and freeway segments for a 20-square-mile section of the Minneapolis area. The data about each segment includes the x and y position of the two nodes, average speed for the segment, average occupancy, and road type. For these preliminary experiments, we used only the distance between the edges as the edge cost. The edges that connected freeway segments were one-way, making the resulting graph directed, as opposed to undirected. The road segments were obtained by digital imaging of the area, followed by regional identification. The map is shown in Figure 8. The more dense center region is the downtown Minneapolis region. In this region, the highways and freeways are not parallel to the x or y-axis. The outlying areas do show a more gridlike pattern of roads, except where lakes interrupt in the lower left corner and where the Mississippi river flows north to southeast in the upper right quadrant of the map.




Figure 7: Effect of edge cost models on execution time

Algorithm / Cost	Uniform Cost	20% Variance	Skewed
Dijkstra	399	399	48
A* (version 3)	189	360	38
Iterative	39	39	56

Table 7: Effect of Edge Cost Models on Iterations
(20x20 Graph, 20% Edge Cost Variance, Diagonal Path)

Table 8 and Figure 9 show both the number of iterations and the execution times for the various algorithms for 4 different paths. The paths that we considered included two diagonal paths (from A to B and from C to D), and two short paths (from D to G and E to F). The results are consistent with our predictions from the synthetic grid and the algebraic simulations. The graph size is slightly larger (1089 nodes) compared to the 30x30 synthetic grid (900 nodes). The iterative algorithm must perform approximately the same number of iterations, regardless of the path length. The cost per iteration will change depending on the amount of backtracking that must be done. In the diagonal case, the path from point A to point B is against the slope of the downtown area, resulting in more backtracking. The lower cost for the C to D path may result from less backtracking, as the downtown area grid is almost parallel to the path. The iterative

algorithm is superior to the estimator-based algorithms when the number of iterations is high. With a smaller number of iterations (i.e. shorter path length), the estimator-based algorithms clearly outperform the iterative algorithm. The path from D to G required only 17 iterations for the optimal A* algorithm, resulting in a cost that is 95% less than that of the iterative algorithm. Similar results were obtained for the second short path from E to F.





Figure 8: Minneapolis Road Map

Algorithm / Path	A to B	C to D	G to D	E to F
Iterative	55	51	55	41
A* (version 3)	453	266	17	64
Dijkstra	1058	1006	105	307

Table 8: Effect of Path length and orientation on Iterations
(Data from Minneapolis)



(Minneapolis Road Map Experiment)

Figure 9: Minneapolis Road Map Results

Analysis of Design Decisions in Implementing A*

We examined two design issues in implementing the estimator-based A* algorithms: frontierSet management and the choice of estimator function.

FrontierSet operations account for most of the work in each iteration of estimator-based path-computation algorithms. We examine two implementations of the frontierSet: as an independent relation, and as an attribute in the nodes relation. The frontierSet can be managed as

an independent relation. New reachable nodes can be added by an insert operations, and deleted from the frontierSet with a delete operation. The best node can be selected by a scanning of the frontierSet. This implementation requires adjustment of the index on the part of relation R, representing the frontierSet. Alternatively, the frontierSet can be implemented by adding an attribute status to each node in the node relation. The status attribute can take four values: open, closed, current, and null. Reachable nodes can be modified with a replace operation, which changes the status of newly reachable nodes to open. An explored node can be deleted by a replace operation which changes the status attribute to closed. Selection of the best node in the frontierSet could be implemented by a scan. It has been argued that the latter implementation is cost effective, due to the smaller overhead of index maintenance[39]. We use the QUEL[39] command REPLACE instead of APPEND and DELETE in implementing all the algorithms in Section 2.

Estimator functions are used to select the best node on the frontierSet to be explored in the current iteration. A perfect estimator function helps the algorithm to discover the shortest path by exploring the minimum number of nodes in the graph[41]. We examine two estimator functions: euclidean distance and manhattan distance. Consider two nodes located at coordinates (x_1, y_1) and (x_2, y_2) , respectively. The euclidean distance between the nodes is defined to be $\sqrt{(x_1-x_2)^2 + (y_1-y_2)^2}$. Euclidean distance always underestimates the cost of the shortest path between nodes. The manhattan distance between them is defined to be $(|x_1 - x_2| + |y_1 - y_2|)$. Manhattan distance is a perfect estimate of the length of the shortest path between nodes in grid graphs with a uniform cost model. The manhattan distance estimator function is not perfect on the Minneapolis data set, however, due to the non-uniform costs between edges. In fact, the manhattan distance on the Minneapolis data set is not always an underestimate; thus, for this graph, use of the manhattan distance does not guarantee an optimal solution.

We examined three implementations of estimator based-algorithms in this experiment: A* version 1, A* version 2, and A* version 3. Version 1 implements the frontierSet as a separate relation and uses the euclidean estimator function. Version 2 implements the frontierSet via a status attribute in the node relation and uses the euclidean estimator function. Version 3 implements the frontierSet via a status attribute in the node relation and uses the manhattan estimator function.




Figure 10: Effect of graph size on Execution Time of Versions

Effect of FrontierSet Management

The effect of frontierSet management can be examined by comparing the performances of version 1 and version 2. Version 1 and version 2 show quite different performances for different graph sizes, as shown in Figure 10. As the graph size increases, the performance of A* version 1 becomes worse than version 2. This is attributed to the fact that the APPEND and DELETE operations cost more than the REPLACE operation in Ingres [39]. Another reason for this difference in performance is that the A* version 2 updates the cost and flag of a neighboring node only if the original cost of that neighbor is greater than the cost of traveling through the "current" node. This step further combines the APPEND and DELETE in A* version 1 to a REPLACE in version 2.

Figure 11 shows that A* version 1 performs worse than version 2 for a uniform-cost grid. However, it does better than version 2 for a skewed graph, because of the higher initialization costs for version 2. A* version 1 expands nodes and appends them to the resultant relation as it goes along, unlike version 2, which begins by loading all the neighbors into the resultant relation. Both versions perform worst on the graph with a 20% variation in edge cost.




Figure 11: Effect of varying edge cost on Execution Time of Versions

We can see from Figure 12 that the execution time for all of the versions increases as the path length increases. A* version 1 starts out much better than version 2, but for longer paths it falls behind. The poor performance of versions 2 in the straight-line path could be attributed to higher initialization costs.

Effect of Estimator Functions

The effect of estimator functions can be studied by comparing their performance in reducing execution time for most cases of path computation. Therefore, choosing a good estimator is of the utmost importance. The effect of estimator functions on performance can be studied by comparing the performances of version 2 and version 3 of A* in Figures 10, 11, and 12. Figure 10 shows that the execution cost of version 3 does not grow as rapidly as the execution cost of version 2. For the 30x30 grid, version 3 performs ten times better than version 2. A* version 2 has a comparatively higher cost of execution in the the graph with the 20% edge-cost variation (Figure 11), whereas version 3 has a comparatively higher cost in a uniform-cost graph. Similarly, in Figure 12 we observe that execution time for version 3 is almost linear to the increase in

path length, whereas the execution time for version 2 increases very rapidly when compared to the increase in path length. Manhattan distance also outperforms euclidean distance for grid graphs.

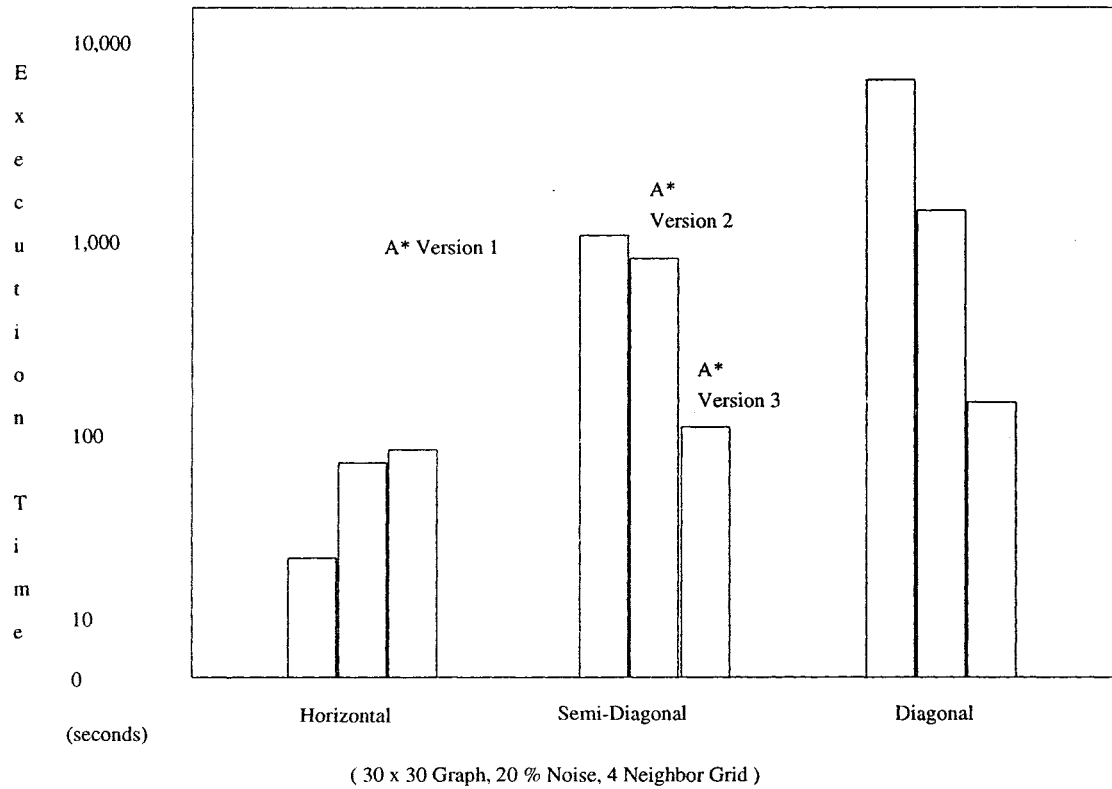


Figure 12: Effect of Path on Execution Time of Versions

4. Main Memory Implementation

4.1. Data Structure and User Interface

Main memory route computation necessitates the selection of efficient data structures that store the set of vertices and the frontierSet. The set of vertices is stored in a special balanced binary tree called a red-black tree, which ensures a worst case performance of $O(\lg n)$ for lookup operations. A *red-black tree* [38] is a balanced binary tree that satisfies the following properties: (1) that every node is either red or black, (2) that every leaf is black, (3) that if a node is red, then both its children are black, (4) and that every simple path from a node to a descendant leaf contains the same number of black nodes. Other balanced binary trees such as 2-3 trees, AVL-trees, and B-trees produce a more balanced tree, but their implementation is more complex. It is also important to efficiently store the vertices in the frontierSet. The frontierSet is the set of vertices sorted by the likelihood of leading to the destination vertex. A *heap* is an array-based data structure that can be viewed as a binary tree. For a tree to be a heap, it must satisfy the heap property: for all nodes i other than the root, $A[\text{parent}(i)] \geq A[i]$. Since a heap is based on a tree, heap operations such as insert and delete will take at most $O(\lg n)$. The heap contains only a single instance of the node. If the node is already in the heap, the node's values are updated and the heap is adjusted. Heap adjustment also takes $\lg(N)$ time. Implementing a more efficient heap data structure such as a binomial heap or a Fibonacci heap is not practical since the heap does not grow large.

The graphical user interface (GUI) displays a map. Travelers can query the map for roadway information such as road name and average road speed. Travelers are able to set up a favorite point alias to frequented map locations such as their home and work place, thus enhancing usability. The GUI allows users to add incidents to the map. An incident can model roadway hazards such as bad weather, accidents, and congestion. Travelers are able to choose between shortest distance routing and shortest time routing. Also, users can choose sub-optimal routes in order to reduce computation time. Figure 13 shows the menu structure for the graphical user interface.




Figure 13: Sample Menus for Graphical User Interface

Routing information must be presented to the traveler in a meaningful but concise format. We chose three familiar formats for route presentation, each with a separate level of detail and concentration. Routes are first presented to the traveler through a highlighted path on a map, distinguishing the route from the rest of the road. Routing information is also presented as a list of directions that provide a greater level of detail for travelers who are unfamiliar with the road network. In order to minimize interactions with the system, the list of directions is read to the traveler.

4.2. Results

Storage Considerations

Two factors are used to determine where the map information will be stored: the amount of memory and the size of the map. For city-sized maps, it is advantageous to use main memory. We developed the main memory route finding software and GUI on a SUN/4 with 24MB of main memory running the SunOS 4.1.3. The route computation software consists of 7000 lines of C++ code. The GUI is 1000 lines of Tcl/Tk code.

Implementation

For this experiment, we examined the performance of two algorithms, A* and Dijkstra's using two metric functions. We repeated the experiment a number of times to arrive at an average execution time. We also kept track of the number of nodes that had been explored and the

¹ The code was compiled with the publicly available GNU C++ compiler, g++ version 2.4. Tcl/Tk is a publicly available, multiplatform graphical user interface toolkit. We used the Wish version 3.3 Tk interpreter.

number of nodes that had been discovered. The number of nodes that have been explored can be thought of as the number of iterations that the algorithm takes, which is a measure of the effectiveness of the metric function.

Performance on Minneapolis Road Map

The Minneapolis road map data set consists of 2716 edges and 946 that represent highway and freeway segments for a 20 square mile section of the Minneapolis area. Each segment includes information on the x and y positions of the two nodes, road type, and average speed. The map is shown in in Figure 8. We used two edge costs: the distance between the nodes and the travel time between the nodes. The travel time is taken from the historical travel time for that line segment.




Figure 14: Results of Travel Distance Experiment




Figure 15: Results of Travel Time Experiment

Figures 14 and 15 show the execution time and number of iterations using both algorithms for four different paths and using both edge cost models. The paths that we considered include two diagonal paths (from A to B and from C to D) and two short paths (from D to G and E to F). When using distance as the edge cost, the A* algorithm is up to 50% faster than the Dijkstra algorithm for long paths. When using the travel-time heuristic function, the performance gap between the two algorithms is less clear. This difference may be attributed to the heuristic function underestimating the remaining travel time.

Algorithm/Path	A to B	C to D	G to D	E to F
A*	884	487	89	44
Dijkstra	933	861	114	97

Table 9: Effect of Travel Time on Iterations

Algorithm/Path	A to B	C to D	G to D	E to F
A*	417	603	43	28
Dijkstra	928	921	97	115

Table 10: Effect of Travel Distance on Iterations

We also examined the tradeoff between path quality and computation time. The first model halts execution when the goal node is removed from the heap, i.e. the standard halting condition. The second model halts its search when the goal node is first discovered. The test paths consisted of paths from all the labeled nodes on the Minneapolis map (Figure 8) to the other labeled nodes. For each pair of nodes, A* distance, A* travel time, Dijkstra Distance, and Dijkstra time routing strategies were tested. Table 11 shows the improvement for the A* algorithm with shortest distance paths, shortest travel-time paths, short paths (B to G, E to F) and long paths (A to B, C to D). The average computational speed up for distance paths when using the first path (model 2) is 2 percent over using the guaranteed optimal path (model 1). However, when using time based paths, the average speedup is just under 10 percent, with a maximum speedup of just under 30 percent.

When short paths (B to G, E to F) are considered, using both travel time and travel distance, the average speedup of method 2 was under two percent. Considering only long paths, the average speedup was just over 3 percent.

In this experiment, path A to D, using Dijkstra distance routing was the only path to have its final-path cost differ from its first-path cost. Since Dijkstra's algorithm has lesser computational cost than A*, the path degradation may be ignored. It is still important to note that path quality is not guaranteed to be optimal if the first path cost is used. Initial results indicate that using a path not guaranteed to be optimal would be necessary only for time based routing.

Case / Method	Average over all Paths		Average of 2 Criteria for Path Quality	
	Distance Path	Time Path	Short Path	Long Paths
Max % Speedup	10.7	28.8	6.0	9.6
Average % Speedup	2.0	9.8	1.8	3.1

Table 11: Percent Speedup from Taking First Path over Optimal Path

5. Conclusions and Future Work

We have carried out algebraic and implementation based evaluations of the performance of three algorithms for computing the shortest path between a pair of points (source, destination)

Conclusions

- A* with effective estimator functions seems to be the preferred algorithm in most situations.
- In real applications such as ATIS, the tradeoff between optimality and speed may allow for sub-optimal algorithms to speed the processing.
- A pure SQL approach is good for storing large maps, but has unacceptable performance for route computation.
- A main memory approach has acceptable performance, but is limited in the size of the maps that it can store.

Future Work

- Storage and access methods need to be evaluated.
- We must explore buffering strategies for efficient query processing.
- Different heuristics for A* will be explored.
- Routing algorithms and access methods will be integrated into a network engine.

References

1. J. H. Rillings and R. J. Betsold, Advanced Driver Information Systems, *Trans. on Vehicular Technology* 40(1)IEEE, (February 1991).
2. G. F. King and T. M. Mast, Excess Travel: Causes, Extent and Consequences, *Transport. Res. Rec., 1111, TRB, Nat. Res. Council*, (1987).
3. J. Eder, Extending SQL with General Transitive Closure and Extreme Value Selections, *Trans. on Knowledge and Data Engineering* 2(4)IEEE, (1990.).
4. M. Mannino and L. D. Shapiro, Extensions to Query Languages for Graph Traversal Problems, *Trans. on Knowledge and Data Eng.* 2(3)IEEE, (Sept. 1990).
5. R. Agrawal, Alpha: An Extension of Relational Algebra to Express a Class of Recursive Queries, *Trans. on Software Eng.* 14(7)IEEE, (July 1988).
6. H. Lu, K. Mikkilineni, and J. P. Richardson, Design and Evaluation of Algorithms to Compute the Transitive Closure of a Database Relation, *Proc. Intl. Conf. on Data Eng., IEEE*, (1987).
7. J. Han, G. Qadah, and C. Chaou, Processing and Evaluation of Transitive Closure Queries, *Conf. on Extending Database Technology*, EDBT, (1988).
8. H. V. Jagadish, R. Agrawal, and L. Ness, A Study of Transitive Closure As a Recursion Mechanism, *Proc. Conf. on Management of Data*, ACM, (1987).
9. H. Jagadish, A Compressed Transitive Closure Technique for Effective Fixed-Point Query Processing, *Proc. Intl. Conf. on Expert Database Systems*, Benjamin Cummings, (1989).
10. Shaul Dar and H. V. Jagadish, A Spanning Tree Transitive Closure Algorithm, *Proc. Intl. Conf. on Data Eng., IEEE*, (1992).
11. R. Agrawal, A. Borgida, and H. V. jagadish, Efficient Management of Transitive Relationships in Large Data Bases, *Proc. Conf. on Management of Data*, ACM SIGMOD, (1989).
12. R. Agrawal and H. V. jagadish, Hybrid Transitive Closure Algorithms, *Proc. Int. Conf. Very Large Data Bases*, VLDB, (1990).
13. Y. Ioannidis, On the Computation of the Transitive Closure of Relational Operators, *Proc. Intl. Conf. on Very Large Data Bases*, VLDB, (1987).
14. P. Valduriez and S. Khoshafian, Transitive Closure of Transitively Closed Relations, *Proc. Intl. Conf. on Expert Database Systems*, Benjamin Cummings, (1989).
15. I. F. Cruz and T. S. Novell, Aggregate Closure: An Extension of Transitive Closure, *Intl. Conf. on Data Engineering*, IEEE, (1989).
16. Y. E. Ioannidis and R. Ramakrishnan, An Efficient Transitive Closure Algorithm, *Intl. Conf. on Very Large Data Bases*, VLDB, (1988).
17. B. Jiang, A Suitable Algorithm for Computing Partial Transitive Closures in Databases, *Intl. Conf. on Data Engineering*, IEEE, (1990).
18. F. Banchillon and R. Ramakrishnan, An Amateur's Introduction to Recursive Query Processing Strategies, *Conf. on Management of Data*, ACM SIGMOD, (1986).
19. F. Banchillon, Naive Evaluation of Recursively Defined Relations, *On Knowledge Base Management Systems - Integrating Database and AI Systems*, Springer Verlag (Ed. Brodie, Mylopoulos), (1985).
20. P. Valduriez and H. Boral, Evaluation of Recursive Queries Using Join Indices, *Intl. Conf. Expert Database Systems*, (1986).
21. Y. Kusumi, S. Nishio, and T. Hasegawa, File Access Level Optimization Using page Access Graph on Recursive Query Evaluation, *Proc. Conf. on Extending Database Technology*, EDTB, (1988).

22. B. Jiang, I/O Efficiency of Shortest Path Algorithms: An Analysis, *Proc. Intl. Conf. on Data Eng.*, IEEE, (1992).
23. O. Hiroyuki, Advanced Mobile Traffic Information and Communication System (AMITCS), *Proceedings of the 4th World Congress of the International Road Safety Organization, Tokyo, Japan*, (June 5-7, 1990).
24. A.M. Kirson, RF Data Communications Considerations in ADIS, *Transactions on Vehicular Technology* **40** pp. 51-55 IEEE, (1991).
25. R.L. French, Advanced Driver Information Systems, *Transactions on Vehicular Technology* **40** pp. 30-55 IEEE, (1991).
26. T. Kuznetsov, High Performance Routing for IVHS, *IVHS America*, pp. 540-543 (1993).
27. A. Goode, Information Systems Under Development To Make Driving Easier, *World Highways* **34**(6) pp. 10-11 (1983).
28. A. Norman, W.B. Zavoli, and M. Heideman, *Vehicle Information Systems and Electronic Display Technology. Integrating Business Listings With Digital Maps For Use In Vehicles*, Society of Automotive Engineers (1991).
29. *VNIS '91: Vehicle Navigation and Information Systems Conference proceedings.* 1991.
30. A. Clelland, M. Wendtland, S. Tedesco, F. Mammano, and G. Endo, Evaluation Results and Lessons Learned from The Pathfinder Operational Test, *IVHS America*, pp. 448-456 (1993).
31. J. D'Ignazio, *Detroit Transportation Center Transit Information. Research in Progress*, Federal Transit Administration (1993).
32. R. Sumner, *Pathfinder Project*, Farradyne Systems, Inc. (1988).
33. M.G. Sheldrick, Driving While Automated. Planning Smart Highways for Tomorrow's Smart Cars, *Scientific American* **263**(1) p. 2p (1990).
34. P.A. McOwen and J. Collura, Technology Considerations for a fault-tolerant, Windowed video computer control system for fast-response incident management, *IVHA America*, (1992).
35. R.W. Behnke, *German "SMART BUS" Systems*, Aegis Transportation Information Systems, Inc. (1993).
36. T.B. Sheridan, *Human Factors of Driver-Vehicle Interaction in the IVHS environment*, Massachusetts Institute of Technology; Center for Transportation Studies (1991).
37. M. Haselkorn, W. Barfiels, J. Spyridakis, L. Conquest, D. Dailey, P. Crosby, B. Goble, and M. Garner, *Real-Time Motorist Information for reducing urban freeway congestion*, Washington University (1992).
38. T. H. Cormen, C. E. Leisserson, and R. L. Rivest, *Introduction to Algorithms*, MIT Press, (1992).
39. R. Kung, E. Hanson, Y. Ioannidis, T. Sellis, L. Shapiro, and M. Stonebraker, Heuristic Search in Data Base Systems, *Proc. Expert Database Systems*, Benjamin Cummings Publications, (1986).
40. R. Detcher and J. Pearl, *Generalized best-first strategies and the optimality of A**, UCLA-ENG-8219, University of California, Los Angeles ().
41. D. Galperin, On the optimality of A*, *Artificial Intelligence* **8**(1) pp. 69-76 (1977).

Date: Thu, 18 Apr 1996 11:39:11 PST
From: RUTZBM@baron.wosc.osshe.edu
To: Sarah J Faust <faus0005@gold.tc.umn.edu>
Subject: Re:

Sarah,
Hiya! Thanks for the poem you wrote last week, it was very nice, but sad. I like those poems that don't rhyme, they're just like an eclectic collection of phrases that convey your thoughts. Well done. To answer your questions:

-Kristan and I don't have a pet, Kristan is allergic to cat hair. her parents got the coolest cat ever. She is so friendly and likes to play, Kristan even likes the cat, but every time we go to Bend for the weekend she gets all sneezy and her eyes water all weekend. This could put the cabosh on ever getting a pet, but ya never know. I think that's why babies don't have full body hair, because God knew a lot of adults are allergic to small furry critters and wouldn't it suck if you were allergic to your own child? (kind of a deep thought by Jack Handy)

-Our address is 4000 Witham Hill Dr. #16 Corvallis, OR 97330
-I don't remember any more questions, but if I think of one I'll answer it later.

The thing about Corvallis (on KQ) was on the Oregon State campus. Some guys in the dorms were doing some racist stuff. I'm surprised that it made the news out there because it wasn't that big of a story here. It lasted a couple of days and they talked about it because there were student protests and stuff. Still, it wasn't anything you couldn't witness on an average day in front of Morrill Hall. You know, some group of students is always pissed about something. If I had a dime for every protest I saw while I was at the U of M, I wouldn't have even needed to go to college, I could have retired. Those protesters would love it out here. There's a person chained to a tree every day at some logging sight. The logging issue really is big out here- it's on the news almost daily.

Have you guys talked to the Borgy's lately? We haven't talked to them since the day Brennan was born - We have been planning on calling them for the last 4 days but haven't gotten around to it. We don't get home until 7:00 or 8:00 every night which is like 10:00 or 11:00 their time and we don't want to rouse little Brennan from his sleep, not to mention Chad and Misti. Send me info if you have it. We are expecting pictures in the mail from them soon. Well, I should cruise. Hope to hear from ya soon. Who knows, maybe you are at work right now reading this. I'll be here for a short while longer so write back if you get this. Do you know how to have a conversation on e-mail? That'd be cool, heh, heh.

Love, Chris

