UCI-ITS-LI-WP-02-1

The Relative Performance of Heuristics for the Dynamic Traveling Salesman Problem

UCI-ITS-LI-WP-02-1

Amelia C. Regan [1]
Jiri Herrmann [2]
Xiangwen Lu [2]

[1] Department of Civil and Environmental Engineering, Graduate School of Management and Institute of Transportation Studies, University of California, Irvine
aregan@uci.edu

[2] Institute of Transportation Studies, University of California, Irvine
jherrmann@uci.edu, xlu@uci.edu

February 2002

Institute of Transportation Studies
University of California, Irvine
Irvine, CA 92697-3600, U.S.A.
http://www.its.uci.edu

The Relative Performance of Heuristics for the Dynamic Traveling Salesman Problem

Amelia Regan

Civil and Environmental Engineering and Graduate School of Management

Jiri Herrmann

Institute of Transportation Studies

and

Xiangwen Lu

Institute of Transportation Studies

University of California, Irvine

Irvine, CA, 92697

ABSTRACT

In this working paper we investigate the convergence of an asymptotically optimal algorithm for the M/G/1 queueing model with switchover costs, when applied to the related dynamic traveling salesman problem. In this paper we present simulation based analysis of heuristics for the dynamic traveling salesman problem in which a mobile server provides service to customers whose positions are known. Service requests are generated according to a Poisson process which is uniform across customer locations. In the general case we assume that the mean service time is known and its variance is bounded. Service time is independent of customer location. We first examine a special case of the problem in which the optimal traveling salesperson problem (TSP) tour and minimum spanning tree involves only links of equal length and then discuss the case for a general graph. The goal of this work is to examine the relative performance of algorithms intended to minimize the average waiting time of each customer. We show that the simple cyclic polling algorithm in which customers are visited in order of the optimal TSP tour has surprising robust performance.

February, 2002

Key Words: Dynamic and stochastic routing and scheduling, traveling salesperson problem, TSP, Dynamic TSP

Working Paper, early version appeared in the Proceedings of the 2002 meeting of the Transportation Research Board

INTRODUCTION

The dynamic traveling salesman problem concerns the development of a routing policy for a single mobile server providing service to customers whose positions are known. Service requests are generated according to a Poisson process which is uniform across customer locations. In the general case we assume that the mean service time is known and its variance is bounded. Service time is independent of customer location. This problem, called the Dynamic Traveling Salesman Problem (DTSP), was first introduced by Psaraftis (*1*). Bertsimas and van Ryzin (*2*) studied a similar problem, the Dynamic Repairman Problem (DTRP), in which customer locations are either uniformly distributed in a bounded area in the Euclidean plane or distributed according to a distribution with probability density function f(X).

In an earlier paper, Lu, Irani and Regan (*3*) examined a special case of the DTSP in which the optimal TSP tour and minimum spanning tree across customer locations involves only links of equal length (see Figure 1 for some examples). In this paper we also begin with this special case. For this special case, the researchers show that the average waiting time when the server follows the a priori tour generated by the well known "cyclic polling" algorithm is bounded from above by approximately $\frac{2?\ ?_1}{1?\ ?_1}$ times the average waiting time of the optimal algorithm where $?_1$, the fraction of time spent in on-site service time for a single node, is less than $\frac{1}{n}$ (in this case the server utilization rate, $?\ ?\ n?_1$). Note that when n is large $\frac{2?\ ?_1}{1?\ ?_1}$ is close to 2, and that when n is small $\frac{2?\ ?_1}{1?\ ?_1}$ is bounded by 3. They also identify circumstances under which this bound is tight. This implies that under certain conditions the cyclic polling algorithm is close to optimal. In this paper we examine the performance of various heuristics for the DTSP using simulation based analysis. We compare the performance of a set of heuristics for both special and general networks. We show that cyclic polling has very good, and quite possibly optimal performance as $?$, the fraction of time the server spends in on-site service time approaches 1.

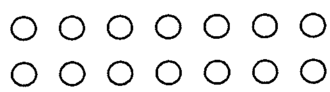

Example network with two rows

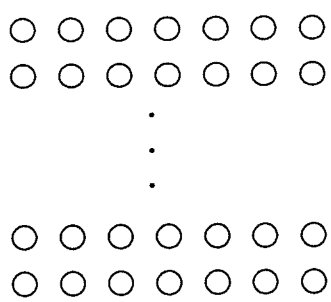

Example network with 2*K rows (K = 1,2,3…)

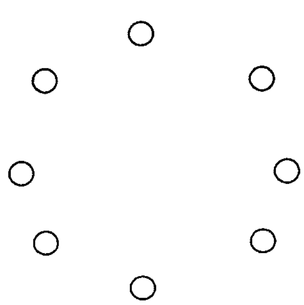

Figure 1. Example networks

RELEVANT LITERATURE

As mentioned in the introduction, the DTSP is closely related to the DTRP (see for example, Bertsimas and van Ryzin (*2*) and Psaraftis (*1*)). We also draw on related research in queueing theory in which cyclic polling algorithms are examined. A cyclic polling algorithm for a network of M/G/1 queues assumes that the server travels along a pre-defined tour, visiting each queue once in the tour. If the queue is empty, then the server simply moves on to the next queue. Much of the research on cyclic polling systems has as its focus the analysis of the waiting time or queue length under various service policies. An *exhaustive* policy is one in which the server will provide service to all waiting customers and to all customers who arrive during service to waiting customers. Under a *gated* policy, on the other hand, the server will provide service to waiting customers and then move on to the next queue, ignoring arrivals during the service period. In a *patient* system the server will spend some time waiting (either at its current location or some other location) if the system is idle, rather than continuing to cycle. Cooper, Niu and Srinivasan (*4*) developed an explicit expression for the average waiting time under gated or exhaustive policies. Srinivasan, Niu and Cooper (*5*) extended those results to describe the relationship between the waiting time distributions in systems with zero and non-zero switchover costs when a gated or exhaustive service discipline is enforced. Eisenberg (*6*) analyzed the polling system in which the server comes to a stop when the system is empty rather than continuing to cycle. That work examines a variety of stopping and starting rules. Later, Srinivasan and Gupta (*7*) consider the circumstances under which the server should be patient. We mention this literature here because the DTSP is closely related to the M/G/1 queueing model with switchover costs. Lu, Regan and Irani (*8*) examine the M/G/1 queueing model with switchover costs. In that paper the researchers develop a lower bound for the waiting time in these systems under any arbitrary algorithm, including those that are optimal.

In this paper we examine both cases where the switchover costs are constant, and those in which the switchover costs are proportional to the distance between nodes (or queues). We examine cyclic polling under exhaustive and gated service rules as well as systems with impatient and patient servers. We examine these relative to a longest queue first policy which is known to be optimal for the case with constant switchover costs. Our simulation shows that the a priori cyclic polling tour is surprisingly robust, relative to longest queue first policies and, as the fraction of time the server is busy increases, in networks in which the distance between nodes is not constant, cyclic polling actually outperforms longest queue first.

THE SIMULATION MODEL

The simulation model was developed in C++. The run lengths for the simulation are very long. It is well known that even simple queuing systems require very long periods to reach steady state behavior. Empirical analysis of our system showed that this was much more true than we could have imagined. However, we found that including a warm up period, however long, had no impact on the long run averages of the performance measures of the system. This is because even for rather high values of ? , the system returns periodically to an empty state. We present results obtained from serving one million customers. Because our simulation run times are so long, we do not present confidence intervals for the performance measures of interest, instead, we present these performance measures directly. The two performance measures of interest are the average wait time for service and the average number of customers in the system. We compare the following policies: cyclic polling with exhaustive service, cyclic polling with gated service and longest queue first with exhaustive service. We present the simulation results regarding the relative performance of these three heuristics for two types of six node graphs. The number six was selected arbitrarily. These results assume that the service times are exponential. This is without loss of generality. The first type of graph is shown in figure 1, except that we assume that longest queue first pays the same switching cost, no matter which node it is switching to. This is typical of many telecommunications and computer networks. The total distance that cyclic polling server must traverse in a single cycle is six. The second type of graph is one in which the locations of nodes are generated uniformly in a unit square. Then, the relative positions of the nodes are maintained but their locations are shifted so that the length of the optimal TSP tour on these nodes is equal to exactly six time(space) units. The simulations are then run on twenty of these random networks and the values obtained for the twenty (one million customer) runs are averaged.

SIMULATION RESULTS

We begin by presenting a comparison of closed form results for cyclic polling systems in which switchover costs are constant verses our simulation results. We show this in figure 2. The cyclic polling systems examined here assume that the servers are "impatient". This means that they keep moving when the system is empty. One can see that the results of our simulation model are virtually identical to those predicted by the closed form solution for the number of customers in the system under exhaustive cyclic polling.

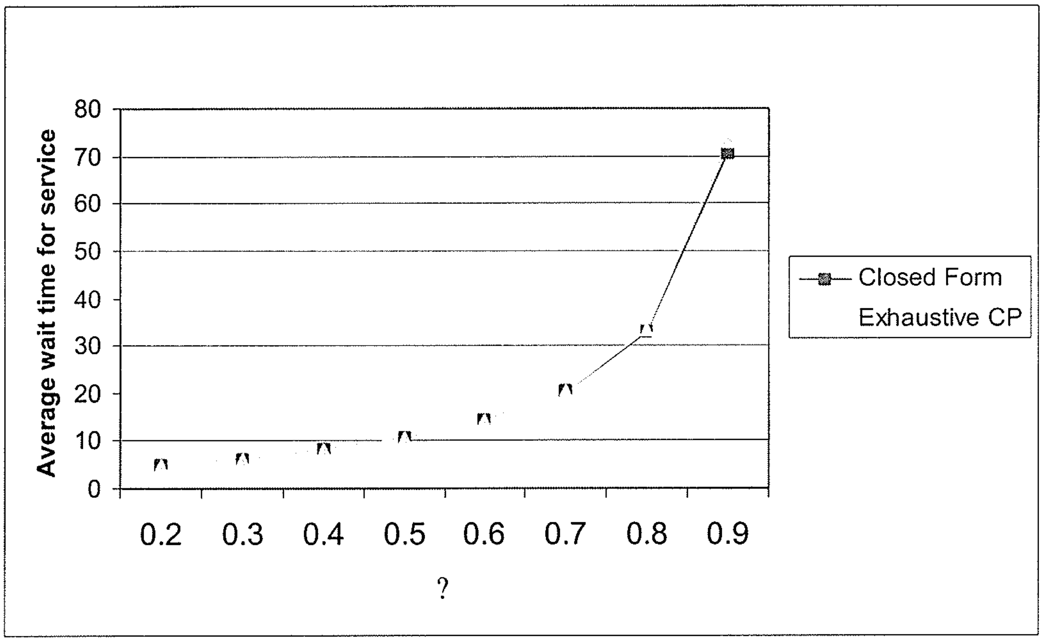

Figure 2. Closed form solution vs. simulation results

We next examine the relative performance of three heuristics for the DTSP. As predicted, longest queue first out-performs the others when the switching costs are constant. Figure 3 presents those results. However, we find that for high values of ?, the performance of cyclic polling is very close to that of longest queue first. Figure 4 presents their relative performance for values of ? between 0.95 and 0.99.

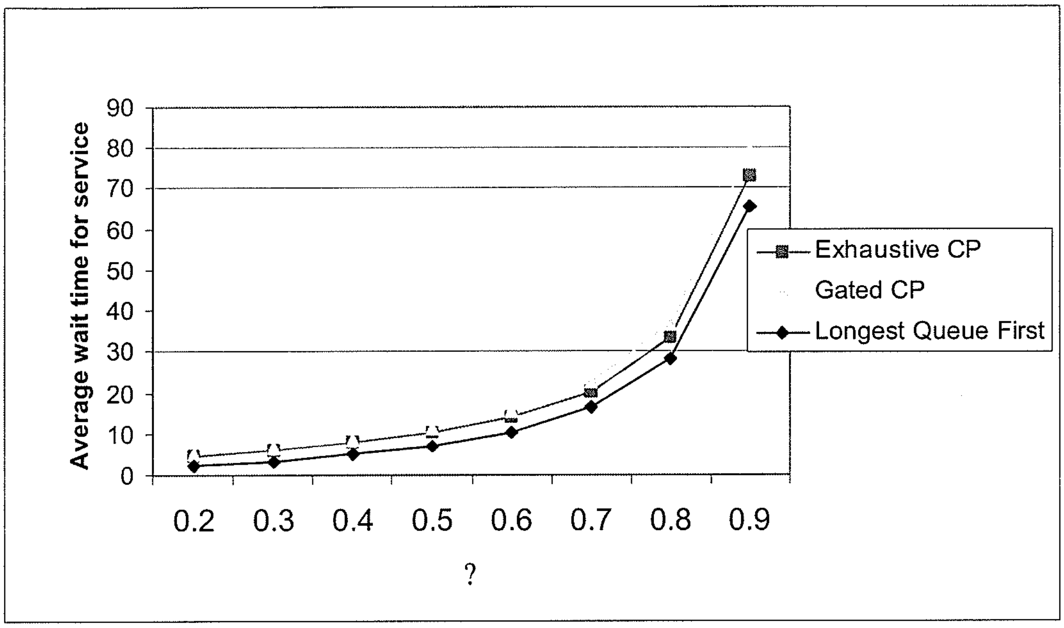

Figure 3. A comparison of three heuristics when the switching costs are constant

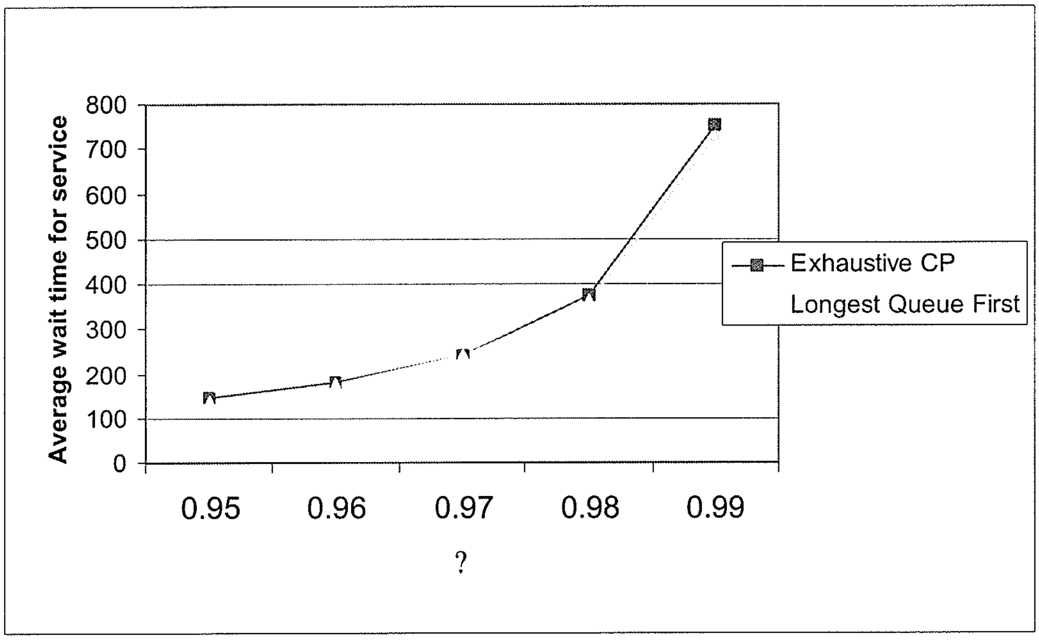

Figure 4. The relative performance of CP and longest queue first for ? approaching 1

Next we compare the performance of the same three heuristics for randomly generated networks. In this case, the cyclic polling tour follows the optimal TSP tour across the customers. The travel time in this case is proportional to distance. Figure 5 shows some examples of these networks. The cyclic polling solution, which corresponds to following the a priori generated TSP tour outperforms longest queue first for large values of ?. Figure 6 presents those results. Of particular interest is the value of ? ?after which cyclic polling outperforms longest queue first. Figure 7 presents more detailed simulation (smaller step size for ?) for the region from ? = 0.75 to 0.94. We can observe that cyclic polling consistently outperforms longest queue first for ? > 0.85. The point at which cyclic polling out performs longest queue first depends upon size of the problem investigated. For 11 node networks generated in the same way the cross-over point is 0.75. We are currently performing additional simulation experiments to try to estimate a function that describes the relationship between the number of nodes in the network and the point at which cyclic polling dominates longest queue first.

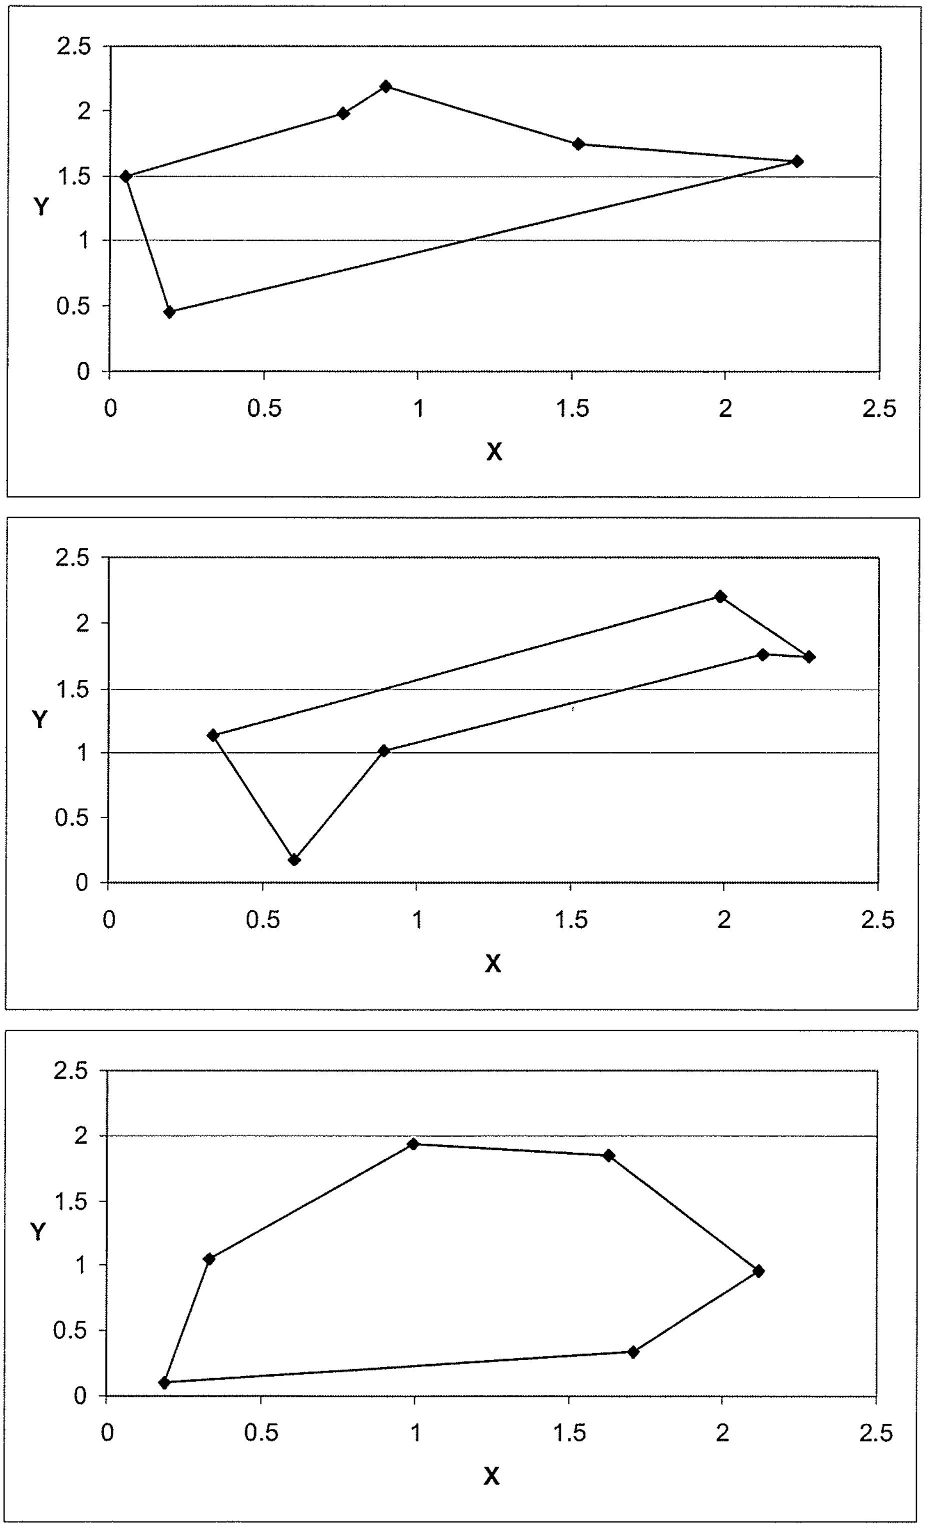

Figure 5. Randomly generated six node networks with L[TSP] = 6 units

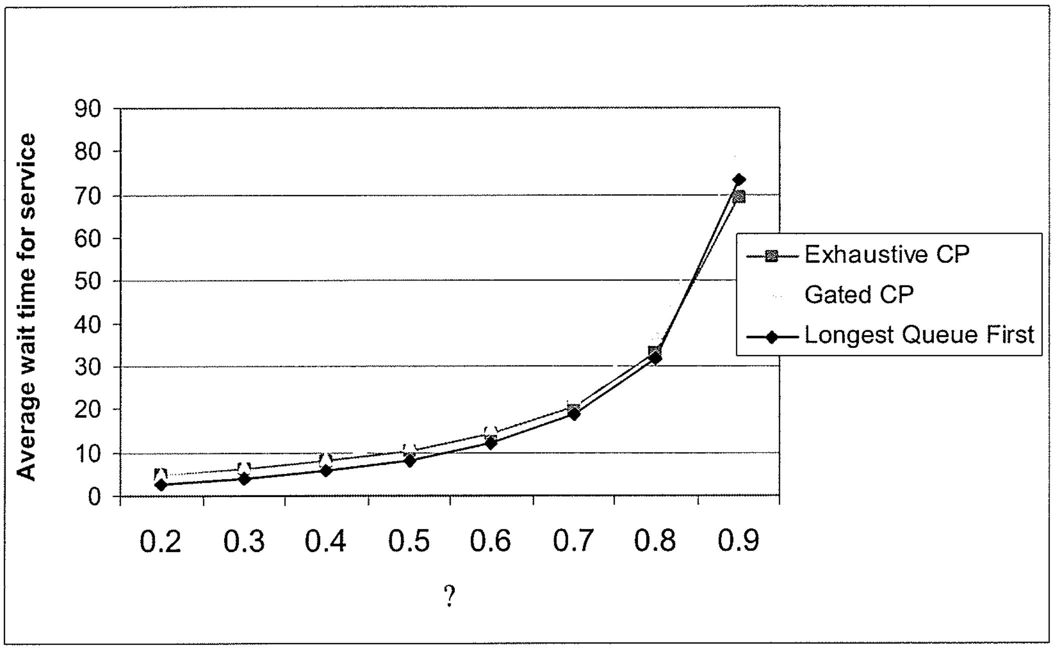

Figure 6. A comparison of three heuristics when the switching costs are proportional to distance (randomly generated 6 node networks)

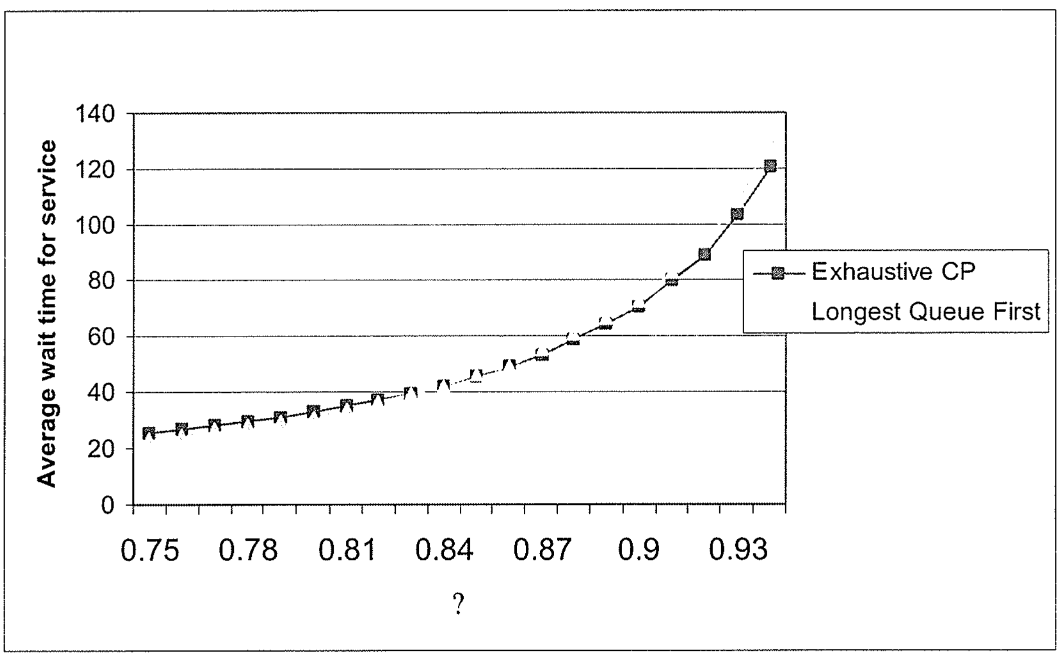

Figure 7. An examination of the point at which cyclic polling outperforms longest queue first when travel time is proportional to distance (randomly generated 6 node networks)

Now we examine the benefits of allowing a "patient" server to wait if there are no customers in the system. We assume that the server waits at its current location until a service request arrives in the system. Then it begins to move again. While the gains over the case with the impatient server are small, they are non-zero. The intuition here is that since arrivals occur according to the same Poisson process at each node, the next arrival is equally likely to occur at any node. It will definitely not occur between nodes, which is where an impatient server will be when the next customer arrives.

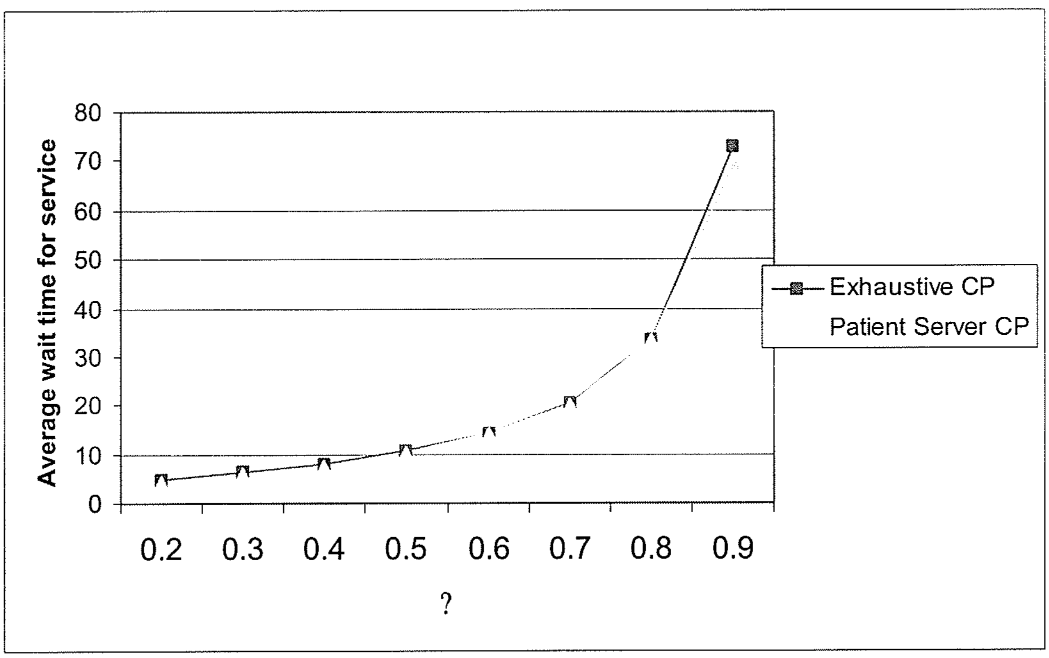

Figure 8. The benefit of allowing the server to be patient when no customers are in the system.

Finally, we mention that in our earlier work we derived two lower bounds for the wait time in the system in an M/G/1 system with constant switching costs (*2*). The true bound is the maximum of these two. Without providing the details of the bounds here, we present an empirical analysis of the tightness of these bounds. The closed form solution for the waiting time under cyclic polling provides an upper bound for the optimal. We can see that while this is the best analytic bound to date for this system, that in some cases it is quite loose. As discussed in (*2*) and as shown in Table 1., the bound becomes tighter as *?* approaches 1.0.

Table 1. Tightness of our lower bound on the average waiting time for service

?	Lower Bound	Closed Form CP	Ratio
0.2	1.25	4.88	3.90
0.3	2.14	6.21	2.90
0.4	3.33	8.00	2.40
0.5	5.00	10.50	2.10
0.6	7.50	14.25	1.90
0.7	11.67	20.50	1.76
0.8	20.00	33.00	1.65
0.9	45.00	70.50	1.57
0.99	495.00	745.50	1.51
0.999	4995.00	7495.50	1.50

CONCLUSION

In this working paper, using a simulation framework, we examine the relative performance of heuristics for the dynamic traveling salesman problem. Because closed form solutions and bounds on system performance are known for the special case in which the cost (distance, time) of switching between nodes is constant, we examine these systems first. We find that as hypothesized by other researchers, that a longest queue first policy performs better than a cyclic polling strategy. However, when we examine these policies on graphs in which the cost of switching between nodes varies, and is proportional to distance, we find that cyclic polling, thought to provide very robust solutions for this problem when the ? , the fraction of time spent in on-site service, approaches 1.0, performs very well. In our simulation, we find that when ? ? 0.85 , the average performance of cyclic polling on randomly generated networks is better than that of the longest queue first rule. Of interest in future research is the identification of network topologies which favor one method over another. This is the subject of on-going research.

We observe that in general graphs, the cyclic polling tour corresponds exactly to the a priori TSP tour over the nodes. Its robust performance agrees with similar research on an equally important stochastic optimization problem, namely the probabilistic traveling salesman problem. In the probabilistic TSP, the goal is to find an a priori solution with the least expected cost, in a network in which node locations are known, but in which customers require a visit in each problem instance according to some known probability. Researchers have shown that when this probability is high across all nodes in the network (which means that the system has high demand) that the TSP tour across the nodes provides good solutions to the PTSP, in which the order of the TSP is followed, but missing nodes are skipped (*9,10*).

ACKNOWLEDGEMENT

This research was partially supported by the University of California Transportation Center (UCTC) and by the US National Science Foundation (NSF) under Grant No. CMS-9875675. The authors gratefully acknowledge this generous support. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the UCTC or the NSF.

REFERENCES

(1) Psaraftis, H N, "Dynamic vehicle routing problems," in *Vehicle Routing: Methods and Studies*, B.L. Golden and A.A. Assad (eds), 223-248, Elsevier Science Publishers (1988).

(2) Bertsimas, D. and G. van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean Plane," *Operations Research*, 39, pp. 601-615 (1991).

(3) Lu, X, S Irani and A.C. Regan, The Dynamic Traveling Salesman Problem: An examination of alternative heuristics, *Transportation Science*, under review (2001).

(4) Cooper, B.R., S.C. Niu and M.M. Srinivasan, "A decomposition theorem for polling models: The switchover times are effectively additive," *Operations Research*, 44, 629-633 (1996).

(5) Srinivasan, M.M., S.C. Niu and B.R. Cooper, "Relating polling models with zero and nonzero switchover times," *Queueing Systems*, 19, 149-168 (1995).

(6) Eisenberg, M.,"The polling system with a stopping server," *Queueing Systems,* 18, 387-431 (1994).

(7) Srinivasan, M.M. and D. Gupta, "When should a roving server be patient?," *Management Science*, 42, 437-451 (1996).

(8) Lu, X., A.C. Regan and S. Irani, "The M/G/1 queue with switchover costs: An examination of alternative heuristics", *Queueing Systems*, under review (2001).

(9) Jaillet, P. , "A Priori Solution of a Traveling Salesman Problem in which a Random Subset of the Customers Are Visited", *Operations Research*, 36, (6), Nov.-Dec. 1988. p.929-936.

(10) Bertsimas, D.J., P. Jaillet and A.R. Odoni, A Priori Optimization, *Operations Research*, 38(6), 1019-1033 (1990).