

ANALYSIS OF AN $O(N^2)$ HEURISTIC FOR THE SINGLE VEHICLE MANY-TO-MANY EUCLIDEAN DIAL-A-RIDE PROBLEM

HARILAOS N. PSARAFTIS

Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

(Received 2 February 1981; in revised form 3 May 1982)

Abstract—We develop an $O(N^2)$ heuristic to solve the single vehicle many-to-many Euclidean Dial-A-Ride problem. The heuristic is based on the Minimum Spanning Tree of the nodes of the problem. The algorithm's worst case performance is four times the length of the optimal Dial-A-Ride tour. An analysis of the algorithm's average performance reveals that in terms of sizes of single-vehicle problems that are likely to be encountered in the real world (up to 100 nodes) and in terms of computational complexity, the $O(N^2)$ heuristic performs equally well, or, in many cases, better than heuristics described earlier by Stein for the same problem. The performance of the heuristic exhibits statistical stability over a broad range of problem sizes.

1. INTRODUCTION

The number of papers on routing and scheduling of Dial-A-Ride transportation systems has grown quite rapidly during the past few years. For instance, Wilson *et al.* (1976, 1977) have developed routing algorithms for Dial-A-Ride Systems operating in Rochester, N.Y.; Hendrickson (1978) and Daganzo (1978) have developed approximate models to evaluate the performance of such systems; Stein (1978a, b) has presented a probabilistic analysis of the problem; Sexton and Bodin (1979, 1980) have developed approximate algorithms based on Benders decomposition which they have applied to the subscriber Dial-A-Ride system of Baltimore, Md; Psaraftis (1980) has developed an exact approach, based on Dynamic Programming for solving the single vehicle problem; Gavish and Srikanth (1979) have developed mathematical formulations of the problem; and finally many other researchers have investigated routing problems connected or potentially connected in one form or another with the Dial-A-Ride problem: Christofides *et al.* (1980), Jarvis and Ratliff (1980), Baker (1981) and Tharakan and Psaraftis (1981). For a comprehensive survey of this and other vehicle routing problems see Bodin *et al.* (1981).

The Dial-A-Ride Problem (henceforth abbreviated DARP) is a very difficult combinatorial problem. In the DARP's generic version, a vehicle, initially located at a point A , is called to service N customers, each of whom wishes to travel from a distinct origin to a distinct destination, and then return to A so that the total length of the route is minimized.

It should be mentioned that this paper refers to the *static, immediate-request* version of the DARP. In that version, all N customers request *immediate* (as soon as possible) service and no new customer requests are considered until all of the above N customers are serviced. In the immediate-request version, no customer desired pickup or delivery times are considered. In that respect, it should be emphasized that the algorithm developed in this paper makes no attempt to satisfy time constraints, and hence is not guaranteed to perform well in case those constraints are present. The reader is referred to Sexton and Bodin (1979, 1980) and to Jaw *et al.* (1981) for the time-constrained problem.

It can be seen that the DARP is a constrained version of the classical Traveling Salesman Problem (TSP), the constraints regarding what we call *route legitimacy*: each customer's origin must precede that customer's destination on the route. Due to the inherent difficulty of the DARP, it is no surprise that all known exact solution algorithms for it are exponential. For instance, using Dynamic Programming (DP) to solve the DARP requires an $O(N^2 3^N)$ time (Psaraftis, 1980), a fact which limits the tractable problem size to no more than 8–10 customers (or 17–21 points).

In this paper, we present an $O(N^2)$ heuristic algorithm to solve the Euclidean version of the DARP as described above. The algorithm's worst-case performance is 300% above the optimum

and there exists a pathological case in the Euclidean plane in which this upper bound is reached asymptotically.

An extensive analysis of the $O(N^2)$ algorithm's average performance is subsequently carried out. The analysis consists of two parts: First, simulation runs in which the algorithm is applied to random problem instances are performed. Various statistics concerning the length of the Dial-A-Ride tour as well as each customer's waiting and riding times are also produced. It is seen that the heuristic exhibits a fairly robust "square root" behavior, according to which the length of the tour and the average waiting and riding times are proportional to the square root of the number of customers requesting service, with the constants of proportionality being fairly independent of the size of the problem. Similar behaviors have been found by Stein (1978a, b) for the same problem and by a number of other researchers for various other combinatorial problems (e.g. Beardwood *et al.* (1959) for the TSP). Thus, while the heuristic itself is intended to be used in an operational setting, its statistical stability can be also useful in the preliminary design of Dial-A-Ride systems, in terms of being able to predict the system's performance under given demand conditions.

The second part of the analysis attempts to answer the question of how much this heuristic deviates from optimality on the average. In that respect, a comparison with the algorithms suggested by Stein (1978a, b) leads to several interesting observations.

At the end of the paper we suggest various other heuristics that can be developed using the same philosophy as in the $O(N^2)$ heuristic. We also discuss various extensions of this work, mainly toward the multi-vehicle and dynamic cases.

2. DESCRIPTION OF THE $O(N^2)$ ALGORITHM

The heuristic is rather simple: It is based on the Minimum Spanning Tree (MST) that is defined on the N origins and the N destinations of the problem. From the MST in question, an initial Traveling Salesman tour T_0 through the above $2N$ nodes can be constructed. The heuristic of this paper produces a legitimate Dial-A-Ride tour by traversing T_0 in such a way so that no destination is visited before the corresponding customer has been picked up. The generic version of the heuristic is as follows (Version 0):

Step 1: Without distinguishing origins from destinations, construct a Traveling Salesman tour T_0 through the $2N$ points based on their Minimum Spanning Tree (MST).

Step 2: Choose any customer origin on T_0 as the first pick-up point P_1 on the Dial-A-Ride route. From this point, move on T_0 clockwise until all points are visited and then return to A . While doing this, do not visit any point that has been previously visited, or any destination whose origin has not been previously visited. Call this Dial-A-Ride tour T_1 .

Step 3 (optional): Improve upon T_1 by performing a sequence of local interchanges, (see details below).

Step 4 (optional): Repeat Step 2 (and optionally Step 3) but moving counterclockwise. Choose the tour in which T_1 has minimum length.

Step 5 (optional): Repeat Step 2 (and optionally Steps 3 and 4) N times, each time choosing a different customer origin as P_1 . Choose the tour in which T_1 has minimum length.

It should be noted that this algorithm belongs to a general class of DARP heuristics, whose Step 1 produces (by any method) an initial Traveling Salesman tour T_0 through $2N$ points, and whose steps 2–5 remain unchanged. In that respect, any TSP heuristic could be used in Step 1. For instance, the heuristic of Christofides (1976) or the interchange heuristics of Lin (1965) and Lin and Kernighan (1973). A comprehensive survey of such heuristics is presented in Golden, Bodin *et al.* (1980). For the purposes of this paper we chose to examine only the MST heuristic for the TSP mainly because of its simplicity.

Obtaining a Traveling Salesman tour T_0 from the MST of $2N$ points is straightforward. In its simplest form, T_0 can be obtained just by duplicating the MST. In this way, every node is visited at least once and thus T_0 is a feasible initial Traveling Salesman tour. An example is shown in Fig. 1(a), which shows the duplicated MST for a three-customer problem, where the origin of customer i is depicted by $+i$ and his/her destination by $-i$ ($i = 1, 2, 3$). Figure 1(b) shows one way to traverse T_0 . In this paper we will call such a traversal "circumnavigation", for obvious reasons.

Of course, *shortcuts* can (and should) be used to reduce the length of a tour obtained in

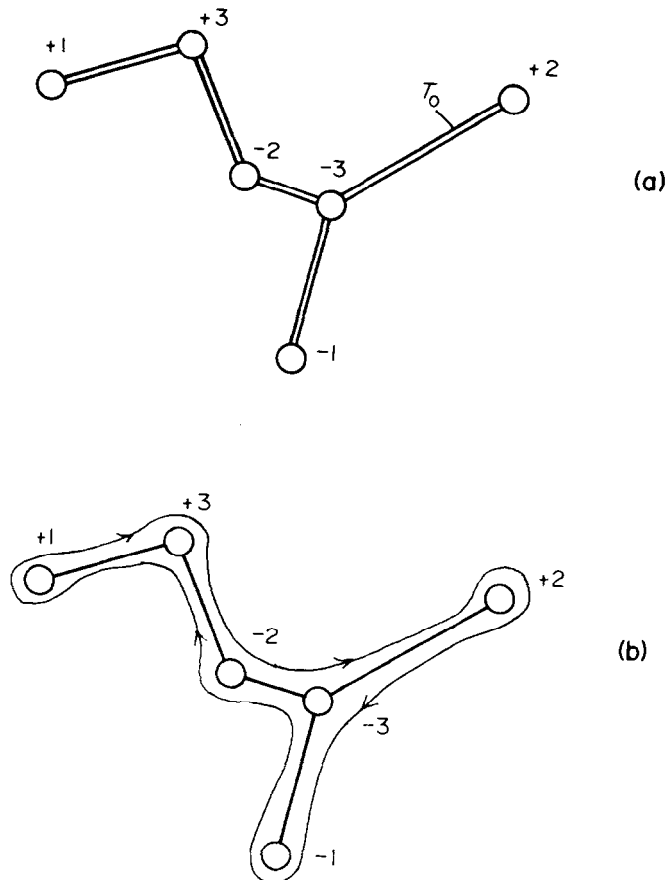


Fig. 1. (a) Duplicate MST. (b) Circumnavigate MST.

such a naive way, for nodes whose MST degree is greater than one will be visited twice or more if the above procedure is used. There are *two* alternative steps in the DARP algorithm where the shortcut operation can be performed. The performance of the heuristic has been seen to depend on which of these steps is chosen.

The obvious step to apply shortcuts is Step 1, the formation of T_0 itself. Doing this will produce a T_0 in which each point is visited exactly once. The alternative is to *omit* this operation in Step 1 and proceed immediately to Step 2 with a T_0 which is *exactly* the duplicated MST. Note that Step 2 incorporates a shortcut operation in itself, by not visiting any previously visited points. In both cases, shortcuts preserve the legitimacy (feasibility) of the DARP tour.

Surprisingly, despite the fact that *shorter* Traveling Salesman tours T_0 will be produced if the first alternative is used, all indications are to the effect that the heuristic performs better if the *second* shortcut alternative is followed. This conjecture stems from the following two points: First, the algorithm's worst-case performance was observed only in cases where shortcuts were used in Step 1 (see Psaraftis, 1981a). And second, it seems that a tour T_0 which visits some points *more than once* is likely to provide an increased flexibility in Step 2 of the algorithm, and hence, a better chance for a shorter Dial-A-Ride tour. In that respect, executing Step 2 with a T_0 that visits points more than once is conjectured to be more useful than doing so with a shorter T_0 in which each point is visited exactly once.

Version 0 of the algorithm will therefore be thought of as the version where T_0 is produced just by MST duplication and where shortcuts are sought in Step 2. The alternative version of the algorithm is called Version 1; in Version 1 T_0 is produced by MST duplication plus shortcuts.

Other versions of the algorithm are Versions 2 and 3, which differ from Versions 0 and 1

(respectively) in that the MST, as well as T_0 , include the starting point A . In those versions, P_1 is the first customer origin that is encountered starting from A and moving clockwise on T_0 (Step 2). Steps 3 and 4 are identical to those of Versions 0 and 1 and Step 5 is omitted since there is no flexibility on the choice of P_1 . It has been observed that this lack of flexibility is a slight disadvantage in the average performance of Versions 2 and 3 as opposed to the average performance of Versions 0 and 1.

In all versions, a word of caution concerns the application of Step 3 (sequence of local interchanges). Figure 2 shows the operation involved in a local interchange, where the route of Fig. 2(a) is preferable to the one of Fig. 2(b) if $d_{ij} + d_{km} - d_{ik} - d_{jm} < 0$, and if k is not the destination of the customer whose origin is j . Only if both of these conditions are met (the first one regarding route improvement and the second route legitimacy) can we interchange j with k on the Dial-A-Ride route. It should be noted that i, j, k and m are adjacent nodes (Fig. 2), and therefore the above interchange scheme is a "very local" operation, in the sense that it is not generally likely to produce dramatic tour improvements. The concept of k -interchange, introduced by Lin (1965) and Lin and Kernighan (1973) in their interchange heuristics for the TSP is generally a more powerful tool, but is likely to require significant modifications for application to the DARP in order to satisfy the route legitimacy constraints. This has been carried out recently by the author in developing 2- and 3- interchange procedures for the DARP. Such procedures are believed to enhance the local search after an initial Dial-A-Ride tour is produced (Psaraftis, 1981b).

An example on the application of Version 0 of the heuristic is shown in Fig. 3. Figure 3(a) depicts the situation after construction of the MST for ten origins (+) and ten destinations (-). The starting point is at A . Figure 3(b) shows the best Dial-A-Ride route obtained by this heuristic after execution of Steps 1, 2, 4 and 5. In other words, no local interchange has been performed, and it can be immediately observed that there is room for improvement. (Notice sequences $-8, -3, +5, -4$ and $+5, -4, -10, -6$). Figure 3(c) displays the Dial-A-Ride route after all steps of the heuristic have been executed.

The computational complexity of this heuristic is $O(N^2)$. The MST algorithm is, in itself, an $O(N^2)$ algorithm (Prim, 1957). Despite the fact that it is also possible to find the MST in the Euclidean plane in $O(N \log N)$ time (Shamos and Hoey, 1975), executing Steps 3 and 5 will still require $O(N^2)$ time.

As stated before, this algorithm has a worst-case performance of 300% above the optimum and there exists a pathological case in the Euclidean plane where this limit is asymptotically reached. The reader is referred to (Psaraftis, 1981a) for more details.




Fig. 2. An elementary local interchange. (a) is preferable to (b) if $d_{ij} + d_{km} - d_{ik} - d_{jm} < 0$ and if $k \neq j$.






Fig. 3. Application of $O(N^2)$ heuristic (Version 0). (a) Construct MST. (b) Best T_1 without interchange.

Fig. 3. (c) Best T_1 with interchange.

3. AVERAGE PERFORMANCE

An issue which is likely to be of interest to transportation planners much more than the algorithm's worst-case performance is how the heuristic performs "on the average". It is known that an algorithm's average performance is something which may bear very little or no connection with the algorithm's worst-case performance. Typical demonstrations of this can be found in the average performance of the heuristic of Christofides (1976) for the TSP, which is reported to be of the order of 10% from the optimum (Golden *et al.* 1980), in spite of a 50% worst-case performance, and in the apparently better average performance of the Lin and Kernighan (1973) heuristic, for which worst-case analysis has revealed that in certain pathological cases the heuristic's relative error can be arbitrarily high (Papadimitriou and Steiglitz, 1978).

A general approach to evaluating an algorithm's average performance is to attempt to derive a relationship that yields the average of the algorithm's optimal value (here, the length of the Dial-A-Ride tour) for a random problem instance. Examples of such approaches are the classical paper of Beardwood *et al.* (1959) for the asymptotic behavior of the optimal length of the TSP, the partitioning algorithm of Karp (1977) for the same problem and the probabilistic analysis of Stein (1978a, b) for the DARP. Trying to proceed according to the above method for the $O(N^2)$ heuristic represents an extremely complicated task, for that method lends itself only to heuristics whose structure is particularly simple, even naive. The reason seems to be that any non-trivial heuristic generally has the property of introducing strong dependencies between its various stages, thus rendering its probabilistic analysis intractable.

In this paper our approach is as follows:

First, we perform a series of simulation runs of the algorithm (Version 0) for random problem instances. Second, we address the question of how much the heuristic deviates from optimality on the average by performing a comparison with the heuristic algorithms suggested by Stein. We now describe these results in detail.

3.1 Description of the simulation runs

Each particular run consists of the execution of the algorithm for a problem of N origins and N destinations uniformly distributed on the unit square. It should be mentioned here that although the uniformity assumption may be appropriate for the general many-to-many case, it is not necessarily realistic in cases where many customers travel to a common destination, such as a railroad station or an airport. However, the assumption was made mainly for purposes of simplicity and compatibility with other analyses of the many-to-many case, and in that respect, causes no loss of generality in our study. Any other user-specified distribution could be tested.

In all runs the vehicle starts from (and finally returns to) the center of the square. Unit speed is assumed for the vehicle. For each value of N several runs have been made, each using a different random number generator seed. Various statistics are produced, whose minimum, maximum and sample average values are tabulated in Table 1. The description of those statistics follows:

(a) *Normalized length of the Dial-A-Ride tour.* L , the length of the Dial-A-Ride tour is normalized by dividing it by $\sqrt{2N}$. This normalization scheme is consistent with the results of Beardwood *et al.* (1959) for the TSP and of Stein (1978a, b) for the DARP. Referring to Table 1, one can see that $L/\sqrt{2N}$ tends to converge to a constant which is fairly independent of N . This constant seems to be in the neighborhood of 1.15.

(b) *Normalized average waiting and riding times per customer.* For each customer i , the waiting time WT_i is counted from the moment the vehicle departs from A until that customer is picked up. Similarly, the riding time RT_i is counted from the moment the customer is picked up until that customer is delivered. For a specific random instance of the problem, the quantities

$$\overline{WT} = \frac{1}{N} \sum_{i=1}^N WT_i$$

and

$$\overline{RT} = \frac{1}{N} \sum_{i=1}^N RT_i$$

are unbiased estimators of the average per customer waiting and riding times respectively. Table 1 shows that these quantities, normalized by $\sqrt{2N}$, seem to converge to constants which again are fairly independent of N . These constants seem to be in the neighborhood of 0.40 and 0.35 respectively.

(c) *Normalized standard deviation of waiting and riding times.* The quantities

$$\sigma_{WT}^2 = \frac{1}{N-1} \sum_{i=1}^N (WT_i - \overline{WT})^2$$

and

$$\sigma_{RT}^2 = \frac{1}{N-1} \sum_{i=1}^N (RT_i - \overline{RT})^2$$

are unbiased estimators of the variances of the waiting and riding times respectively. Table 1 shows that σ_{WT} and σ_{RT} , normalized again by $\sqrt{2N}$, seem to converge to values which are fairly independent of N (approx. 0.25 and 0.23 respectively).

(d) *Histograms of normalized waiting and riding times.* For a series of 10 runs, each for $N = 50$, we have prepared histograms of the normalized waiting and riding times (sample size of each = 500 customers). These histograms are shown in Figs. 4(a) and (b). They constitute an approximation to the probability distribution of the waiting and riding times of a random customer in the system.

A number of interesting observations concerning the average performance of this heuristic can be made:

(1) This simple heuristic is robust in the sense that it exhibits a fairly consistent "square

Table 1. Various statistics from simulation runs, $N = 10-50$ customers (11-101 points). The minimum, maximum and (sample) average value of each statistic are presented

N	10	20	30	40	50
# of runs	4	4	5	4	10
<hr/>					
$L/\sqrt{2N}$					
Minimum	1.01	1.01	1.07	1.08	1.07
Maximum	1.21	1.15	1.24	1.18	1.21
Average	1.09	1.04	1.13	1.13	1.15
<hr/>					
$\overline{WT}/\sqrt{2N}$					
Minimum	0.38	0.38	0.39	0.36	0.34
Maximum	0.46	0.47	0.47	0.40	0.46
Average	0.42	0.42	0.43	0.38	0.40
<hr/>					
$\overline{RT}/\sqrt{2N}$					
Minimum	0.28	0.23	0.33	0.32	0.30
Maximum	0.38	0.30	0.44	0.42	0.38
Average	0.33	0.25	0.37	0.38	0.35
<hr/>					
$\sigma_{WT}/\sqrt{2N}$					
Minimum	0.25	0.23	0.22	0.23	0.22
Maximum	0.31	0.30	0.28	0.26	0.29
Average	0.28	0.27	0.25	0.24	0.25
<hr/>					
$\sigma_{RT}/\sqrt{2N}$					
Minimum	0.14	0.18	0.20	0.20	0.21
Maximum	0.27	0.23	0.27	0.29	0.27
Average	0.19	0.20	0.23	0.24	0.23

root" behavior over a broad range of number of customers. According to this behavior, the average tour length, as well as the mean and standard deviation of each customer's waiting and riding time are proportional to the square root of the number of customers requesting service. This statistical stability of the heuristic can prove useful for preliminary system design considerations, when one would like to have a "quick-and-dirty" method for predicting system performance.

(2) It is also interesting to compare the customer average riding time with the average time it would take a customer to go directly from origin to destination. From Table 1 we can see that the former is growing roughly as $0.35\sqrt{2N} = 0.49\sqrt{N}$. The latter is known (Larson and Odoni, 1980) to be equal to 0.52. It is interesting and perhaps surprising to observe that these two quantities take approximately the same value for a value of N as low as 1. This means that the average riding time the heuristic would give for the trivial case $N = 1$ (0.52) remains in the same neighborhood (when normalized by \sqrt{N}) with the value the heuristic would give for large values of N (0.49). This argument adds to the robustness of this heuristic.

3.2 Deviation from optimality

In his probabilistic analysis of the DARP, Stein (1978a) showed that for $N \rightarrow \infty$, and for pickup and delivery points uniformly distributed in a unit area, the optimal length L^* of a Dial-A-Ride tour converges with probability 1.0 to $(4/3)b\sqrt{2N}$ where b is the asymptotic constant for the TSP (Beardwood, Halton and Hammersley, 1959). For $b = 0.765$ (the speculated, approximate value of b) this means that $L^*/\sqrt{2N}$ converges with probability 1.0 to the value of 1.02. Stein has actually described a simple DARP heuristic whose asymptotic error is only 6% above the optimum (Stein 1978b), plus another, more sophisticated heuristic which is asymptotically optimal (Stein 1978a). Since Table 1 shows that the heuristic of this paper produces Dial-A-Ride tours for which $L/\sqrt{2N} \approx 1.15$, that is, about 13% above the asymptotic value of 1.02, a natural question to ask is the following: "What is the value of the $O(N^2)$ heuristic if it produces tours which on the average are 13% longer than the optimum and potentially can be much longer than that?"





Fig. 4. Histograms of (a) $WT_i/\sqrt{2N}$ and (b) $RT_i/\sqrt{2N}$ from simulation runs. Sample size: $N = 50$, 10 runs (500 customers).

Addressing this very important issue requires some investigation into Stein's analysis. Stein's simple heuristic (henceforth referred to as Stein's Heuristic #1), optimally solves two TSP's, one on the N origins and one on the N destinations and then traverses the two tours sequentially. The result is Dial-A-Ride tour whose length converges with probability 1.0 to $2b\sqrt{N}$ (two TSP's of N nodes each). This length is asymptotically 6% longer than the optimum length of $(4/3)b\sqrt{(2N)}$ (Stein, 1978b). It should be noted that since this algorithm assumes that the two TSP's are solved *exactly*, it requires a computational effort which is exponential with respect to N . Thus, if a Dynamic Programming (DP) algorithm is used to solve each TSP (Held and Karp, 1962), Stein's Heuristic #1 has an $O(N^2 2^N)$ complexity. This implies that Stein's Heuristic #1 is tractable only for problems involving no more than $N \approx 15$ customers and that the cost of any attempt to check its rate of convergence to the 6% figure is prohibitive.

Table 2 presents a comparison of our $O(N^2)$ heuristic with the exact $O(N^{2.3N})$ algorithm of Psaraftis (1980) and with Stein's Heuristic #1 for a series of simulation runs for small problem sizes. Such a comparison is necessarily limited because of the computational difficulties that the $O(N^{2.3N})$ algorithm encounters if N is in the range of 8–10 customers or more. It should be noted that the figures reported for Stein's Heuristic #1 are *lower bounds* for what that heuristic would actually give, since we modified the heuristic so as to produce *optimal* Dial-A-Ride tours subject to the constraint that the set of all origins should precede the set of all destinations, rather than connect two TSP tours arbitrarily.† We observe that indeed the size of problems examined is too small for Stein's Heuristic #1 to converge to its 6% asymptotic value. We also observe that the $O(N^2)$ heuristic produces shorter tours than Stein's Heuristic #1 for the majority of cases.

For greater problem sizes, Stein's Heuristic #1 is very quickly out-performed by the $O(N^2)$ heuristic in terms of computational effort. A more sophisticated heuristic which seems suitable to compare our heuristic to if N is large is the asymptotically optimal heuristic suggested by Stein and referred to as Stein's Heuristic #2. In that heuristic, the area of service is partitioned into M regions of equal size. On the "first pass" through the regions the vehicle collects in

†Such a modified algorithm can be readily developed via a straightforward screening procedure in the state space of the exact DP algorithm of Psaraftis (1980).

Table 2. $L/\sqrt{(2N)}$ for three DARP algorithms applied to small problem instances. Each row represents a particular problem sampled from a uniform distribution on the unit square. Figures in parentheses are percentages above optimality. Stein's Heuristic #1 is modified so as to produce improved DARP tours (see text)

	Exact DP Algorithm	$O(N^2)$ Heuristic	Stein's Heuristic # 1 (modified)
N = 3	0.70	0.70 (0%)	0.70 (0%)
	1.15	1.15 (0%)	1.38 (20%)
	0.85	0.85 (0%)	0.85 (0%)
N = 4	0.80	0.80 (0%)	0.87 (9%)
	0.82	0.82 (0%)	0.82 (0%)
	1.13	1.14 (1%)	1.26 (12%)
N = 5	0.71	0.71 (0%)	0.90 (27%)
	1.29	1.37 (6%)	1.31 (2%)
	0.85	0.85 (0%)	0.97 (14%)
N = 6	0.86	0.90 (5%)	0.91 (6%)
	1.30	1.45 (12%)	1.42 (9%)
	1.07	1.16 (8%)	1.18 (10%)
N = 7	0.84	0.88 (5%)	0.97 (15%)
	1.17	1.22 (4%)	1.42 (21%)
	0.80	1.01 (26%)	1.01 (26%)
N = 8	0.84	0.93 (11%)	0.91 (8%)
	0.97	1.06 (9%)	1.06 (9%)
	1.20	1.30 (8%)	1.39 (16%)

region i ($i = 1, 2, \dots, M$) all origins in that region as well as delivers all destinations from regions $1, 2, \dots, i - 1$. On the “second pass” through the regions the vehicle visits the remaining destinations in each region. If an exact TSP algorithm is used for the routing in each region, this heuristic produces Dial-A-Ride tours for which $L/\sqrt{(2N)}$ converges to 1.02 with probability 1.0, hence the heuristic is asymptotically optimal.

It is interesting to shed some light on the assumptions that Stein used to establish this result, as well as to discuss the practical implications of these assumptions on the rate of convergence of $L/\sqrt{(2N)}$ towards the asymptotic value of 1.02 and on the algorithm’s computational complexity. The issue of convergence is of particular importance if one asks for what finite values of N does Stein’s Heuristic #2 start to behave as an optimal algorithm; this is an issue that has not been investigated in the literature to date.

An important assumption in Stein’s analysis is that not only $N \rightarrow \infty$ but also that the number of regions M , as well as the number of customers N/M within each region go to infinity, for only if the latter is true can one use the Beardwood, Halton and Hammersley (BHH) results for the lengths of the TSP tours that are defined in each region. If an exact, DP algorithm is used for those TSP’s, then Stein’s Heuristic #2 exhibits a running time of $O(M(N/M)^2 2^{N/M}) = O(N^2 2^{N/M}/M)$. This time can be polynomial with respect to N for certain rates of growth of M with respect to N , provided that both M and N/M go to infinity. One such type of growth (but not the only one possible) is $M = N/\lg N$ where \lg denotes the logarithm of base 2. If this is true, then the running time becomes $O(N^2 \lg N)$, indeed polynomial with respect to N . However, for $N = 50$ customers (the highest value for which we made simulation runs), this would imply approx. 9 regions, and between 5 and 6 customers per region, numbers which are probably too small for the BHH results to converge (Eilon *et al.*, 1971, p. 170). $N = 1000$ would change those figures to 15 regions and between 6 and 7 customers per region, and $N = 1000$ to 100 regions and around 10 customers per region. If 20 customers per region are considered as adequate for the BHH results to converge in each individual region, then this would imply about 52,500 regions and slightly over one million customers.

Of course, other rates of growth of M can produce a polynomial running time as well. A slower rate of growth, say, $M = (1/3)N/\lg N$ would imply 5 regions and a total of about 100 customers for 20 customers per region. However the running time of such an algorithm would be $O(N^4 \lg N)$, time that is certain to be prohibitive if N is large. A faster rate of growth would lower the algorithm’s complexity but slow down its rate of convergence even more.

The practical implication of the above analysis is the following: If one is to maintain computational tractability in Stein’s Heuristic #2 (in the sense of a low power polynomial complexity), then the range of convergence of that heuristic to an asymptotically optimal behavior is of the order of thousands (if not millions) of customers. For problems of smaller size, the assumptions underlying the performance of this heuristic become less and less valid and the performance itself more and more uncertain. A similar observation was made by Golden *et al.* (1980) for the performance of an algorithm methodologically very similar to Stein’s, Karp’s partitioning algorithm for the TSP (1977).

Based on the above, the idea of comparing the simulation results of the $O(N^2)$ heuristic that were presented earlier in this section for $10 \leq N \leq 50$ directly with the asymptotic value of 1.02 does not seem to have a very strong practical justification. An approach that we pursued instead was the following:

We ran the $O(N^2)$ heuristic and Stein’s Heuristic #2 for four cases of $N = 50$ customers (101 points) uniformly distributed on the unit square. We chose $M = 9$ subdivisions by assuming a $M = N/\lg N$ growth which guarantees an $O(N^2 \lg N)$ running time for Stein’s procedure. The 9 subdivisions were assumed to be $1/3 \times 1/3$ squares, boxed in a 3×3 fashion inside the unit square. The TSP’s defined in each subdivision were solved exactly and interconnected by visual inspection arbitrarily, as in Stein’s procedure, usually by deleting their longest links. The results appear in Table 3. We observe that the $O(N^2)$ heuristic performs better in three out of four cases and that Stein’s procedure itself consistently deviates from its asymptotic value of 1.02. Thus, this investigation seems to indicate that, for the sizes of single-vehicle problems that are likely to be encountered in the real world, the $O(N^2)$ heuristic performs equally well as, or, in many cases, better than Stein’s procedures with respect to both computational complexity and tour length despite the apparent superior theoretical performance of the latter for $N \rightarrow \infty$.

Table 3. $L/\sqrt{2N}$ for the $O(N^2)$ Heuristic and Stein's asymptotically optimal Heuristic #2. Each row represents a different problem ($N = 50$ customers). 9 subdivisions were used for Stein's procedure (see text). Figures in parentheses are percentages above the asymptotic value of 1.02

$O(N^2)$ Heuristic	Stein's Heuristic # 2
1.17 (15%)	1.22 (20%)
1.31 (28%)	1.23 (21%)
1.12 (10%)	1.33 (30%)
1.22 (20%)	1.31 (28%)

We conclude this unusually lengthy discussion with a rather philosophical note: It is clear that the preceding analysis is, to some extent, unfair for Stein's procedures, since the range of problem sizes at which we tested those procedures (N up to 50 customers) turned out to be quite different from the range they were designed for ($N = \text{"large"}$). This is so because this analysis has produced rather strong evidence that "large" for the DARP probably means "much larger than 50", a fact that was not made clear in Stein's work, or anywhere else in the literature to date. In all fairness to the above author, we should mention that Stein's intent was not to produce algorithms for the problem in question, but to provide insight and understanding into the structure of algorithms for the problem, and argue that a certain class of simple algorithms works well (Stein, 1981). By the same token, we also feel it is equally misleading to test the performance of the $O(N^2)$ heuristic (or any other DARP algorithm) for those "small and intermediate" problem sizes, against limits that can be reached only asymptotically, and, most likely, only for very large problem sizes. It is therefore our belief that the analysis of this section has also identified some possible pitfalls in comparing algorithms designed to work under different settings and shed some light on other important points regarding the evaluation of an algorithm's average performance.

4. DIRECTIONS FOR FURTHER RESEARCH

As mentioned in Section 2, many other DARP heuristics which are structurally similar to ours can be developed. Practically any heuristic that produces a Traveling Salesman tour T_0 can replace Step 1 of our algorithm. As an example, we can use the heuristic of Christofides (1976) for an initial Traveling Salesman Tour. If this is the case, it is easy to show that the worst-case performance of the DARP heuristic would drop to a 200% error ratio, and its complexity would increase to $O(N^3)$. No computational experience with any of those alternative heuristics exists to date (see Psaraftis, 1981a).

The development of sophisticated k -interchange procedures to make successive improvements in the Dial-A-Ride tour is a potentially more promising direction. However, computational experience with these algorithms is still limited (Psaraftis, 1981b).

Another direction that has been pursued by the author and his colleagues concerns the multi-vehicle case (Jaw *et al.*, 1981). In this case, the $O(N^2)$ heuristic of this paper has been incorporated into a multi-vehicle, advance reservation algorithm. Grouping customers into clusters, satisfying pickup or delivery time constraints and maximizing vehicle productivity are the algorithm's concerns.

A final and very important direction concerns developing a dynamic version of the algorithm. In the dynamic version, customer requests arrive and are considered dynamically in time. In (Psaraftis, 1980) it was shown that the conversion of the static version of the algorithm to its equivalent dynamic version is relatively straightforward. However, the results of the last section will no longer be true if the algorithms operate in a dynamic environment. For instance, the average rate of customer requests will now play a major role in determining the probability distribution of the random customer's waiting and riding time. Queuing considerations will become important, in a yet unspecified fashion. An analysis of that aspect of the problem is very important for the design of a real-time Dial-A-Ride dispatching system.

Acknowledgements—Work on this paper was supported in part by a grant to MIT from the University Research Program of the Urban Mass Transportation Administration. I would like to thank Amedeo Odoni, Nigel Wilson and Christos Papadimitriou for comments and assistance, as well as David Stein for some very useful feedback on Section 3.2. Finally, I am indebted to Larry Bodin, whose particularly insightful comments on two earlier drafts gave me the opportunity to develop a much improved paper.

REFERENCES

- Baker E. (1981) An algorithm for vehicle routing with time window constraints, working paper, University of Miami, Coral Gables.
- Beardwood J., Halton S. H. and Hammersley J. M. (1959) The shortest path through many points. *Proc. Camb. Phil. Soc.* **55**, 229–327.
- Bodin L. D., Golden B., Assad A. and Ball M. (1981) The state of the art in the routing and scheduling of vehicles and crews. *Working paper MS/S #81-035*, University of Maryland, September.
- Bodin L. D. and Sexton T. R. (1980) The subscriber Dial-A-Ride problem: The Baltimore, Maryland Benchmark. *ORSA/TIMS Meeting*, Colorado Springs, November.
- Christofides N. (1976) Worst case analysis of a new heuristic for the traveling salesman problem. *Management Sciences Res. Rep. No. 388*, Carnegie-Mellon University.
- Christofides N., Mingozzi A. and Toth P. (1980) Exact algorithms for the TSP with additional constraints. *Euro-IV Cong.*, Cambridge, England, July.
- Daganzo C. (1978) An approximate analytic model of many-to-many demand responsive transportation systems. *Trans. Res.* **12**, 325–333.
- Eilon S., Christofides N. and Watson-Gandy C. D. T. (1971) *Distribution Management: Mathematical Modeling and Practical Analysis*. Hafner, New York.
- Frederickson U. N., Hecht M. S. and Kim C. E. (1977) Approximation algorithms for some routing problems. Department of Computer Science, University of Maryland.
- Gavish B. and Srikanth K. K. (1979) Mathematical formulations for the Dial-A-Ride problem. *Working paper*, Graduate School of Management, University of Rochester, March.
- Golden B., Bodin L. D., Doyle T. and Stewart W. Jr. (1980) Approximate traveling salesman algorithms. *Oper. Res.* **28**, 694–711.
- Held M. and Karp R. M. (1962) A dynamic programming approach to sequencing problems. *SIAM* **19**, 196–210.
- Hendrickson C. T. (1978) Approximate, analytic performance models of integrated transit system components. Ph.D. thesis, Dept. of Civil Engineering, M.I.T.
- Jarvis J. J. and Ratliff H. D. (1980) Interactive heuristics for large scale transportation models. *D.O.T. Contract TSC-1618*, Report of School of Industrial and Systems Engineering, Georgia Institute of Technology, August.
- Jaw J. J., Odoni A. R., Psaraftis H. N. and Wilson N. H. M. (1981) A heuristic algorithm for the multi-vehicle many-to-many advance-request Dial-A-Ride problem (Version 2). *Working paper*, Massachusetts Institute of Technology, August.
- Karp R. M. (1977) Probabilistic analysis of partitioning algorithms for the traveling salesman problem. *Math. Oper. Res.* **2**, 209–224.
- Larson R. C. and Odoni A. R. (1980) *Urban Operations Research*. Prentice Hall, Englewood Cliffs, New Jersey.
- Lin S. (1965) Computer solutions to the travelling salesman problem. *Bell System Tech. J.* **44**, 2245–2269.
- Lin S. and Kernighan B. W. (1973) An effective heuristic algorithm for the traveling salesman problem. *Oper. Res.* **21**, 498–516.
- Papadimitriou C. H. and Steiglitz K. (1978) Some examples of difficult traveling salesman problems. *Oper. Res.* **26**, 434–443.
- Prim R. C. (1957) Shortest connection networks and some generalizations. *Bell System Tech. J.* **36**, 1389–1401.
- Psaraftis, H. N. (1980) A dynamic programming solution to the single vehicle many-to-many immediate request Dial-A-Ride problem. *Transpn Sci.* **14**, 130–154.
- Psaraftis H. N. (1981a) Two heuristic algorithms for the single vehicle many-to-many Dial-A-Ride problems. *Working Paper OE-UMTA-81-1*, MIT, January.
- Psaraftis H. N. (1981b) *k*-Interchange procedures for local search in a precedence-constrained routing problem. *Working paper OE-OR-81-1*, MIT, October. (to appear in the *European J. Oper. Res.*).
- Sexton T. R. (1979) The single vehicle many-to-many routing and scheduling problem. Ph.D. thesis, State University of New York at Stony Brook.
- Shamos M. I. and Hoey D. (1975) Closest point problems. *Proc. 6th Annual Symp. on Foundations of Computer Science*.
- Stein D. M. (1978a) An asymptotic, probabilistic analysis of a routing problem. *Math. Oper. Res.* **3**, 89–101.
- Stein D. M. (1978b) Scheduling Dial-A-Ride transportation systems. *Trans. Sci.* **12**, 232–249.
- Stein D. M. (1981) Personal communication.
- Tharakan G. G. and Psaraftis H. N. (1981) An exact algorithm for the exponential disutility Dial-A-Ride problem. Working paper, MIT, May.
- Wilson N. H. M. and Weissberg H. (1976) Advanced Dial-A-Ride algorithms research project: Final Report, Dept. of Civil Engineering, M.I.T. Report R76-20.
- Wilson N. H. M. and Colvin N. J. (1977) Computer control of the Rochester Dial-A-Ride system. Dept of Civil Engineering, M.I.T. Report R77-31.