New Competitive Ratios for Generalized Online Routing

Patrick Jaillet * and Michael R. Wagner †

August 2006

Abstract

We consider online routing optimization problems where the objective is to minimize the time needed to visit a set of locations under various constraints; the problems are online because the set of locations are revealed incrementally over time. We make no probabilistic assumptions whatsoever about the problem data. We consider two main problems: (1) the online Traveling Salesman Problem (TSP) with precedence and capacity constraints and (2) the online TSP with m salesmen. For both problems we propose online algorithms, each with a competitive ratio of 2; for the m-salesmen problem, we show our result is best-possible. We also consider polynomial-time online algorithms as well as various generalizations of our results.

1 Introduction

The Traveling Salesman Problem (TSP) is a very important problem in Operations Research; TSP solutions are valuable in their own right as well as in the solution of more complicated problems. In a common version of the TSP, we are given a metric space and a set of points in the space, representing cities. Given an origin city, the task is to find a tour of minimum total length, beginning and ending at the origin, that visits each city at least once. Assuming a constant speed, we can interpret this objective as minimizing the *time* required to complete a tour. We may also incorporate release dates, where a city must be visited on or after its release date; in this case the problem is known as the "TSP with release dates."

Additional constraints can be added to the above salesman problems. We will consider several in this extended abstract. The salesman can be considered a vehicle/server that transports packages and/or people. We can introduce precedence constraints where some cities must be visited before others. Precedence constraints are appropriate, for example, if packages/people have to be picked up at one location and delivered to another location. It is also natural to introduce a capacity for

*Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, jaillet@mit.edu.

†Department of Management and Finance, California State University East Bay, Hayward, CA 94542, michael.r.wagner@gmail.com.

the server; in other words, a server can visit only a subset of all cities in a given tour and must traverse multiple tours. Finally, we consider the case where we have multiple servers to manage.

In this extended abstract we are concerned with *online* versions of the above mentioned routing optimization problems. In our framework, the problem data is revealed dynamically over time, independent of the server's location, at release dates. We make no probabilistic assumptions whatsoever about the problem data. The assumption that problem instances are completely known a priori is unrealistic in many applications. Taxi services, buses and courier services, for example, require an online model in which locations to be visited are revealed over time, while the server is en route serving previously released requests. The focus of our extended abstract is on studying algorithms for a variety of online routing problems. They are evaluated using the competitive ratio criteria, which is defined as the worst case ratio of the online algorithm's cost to the cost of an optimal offline algorithm, where all data is known a priori. We also say that an algorithm is c-competitive if the competitive ratio of the algorithm is at most c. We call an online algorithm (or competitive ratio) *best-possible* if there does not exist another online algorithm with strictly smaller competitive ratio. We provide online algorithms for new online routing problems and we derive new competitive ratio bounds. A number of our competitive ratio results are best-possible.

1.1 Literature Review

Research concerning online versions of the TSP have been introduced relatively recently. Most related to our extended abstract is the stream of works which started with the paper by Ausiello, Feuerstein, Leonardi, Stougie, Talamo [3]. In this paper, the authors studied the online TSP that is a special case of the problems we consider here; they analyzed the problem on the real line and on general metric spaces, developing online algorithms for both cases and achieving a best-possible online algorithm for general metric spaces, with a competitive ratio of 2. These authors also provide a polynomial-time online algorithm, for general metric spaces, which is 3-competitive. Subsequently, Ausiello, Demange, Laura, Paschos [2] gave a polynomial-time algorithm, for general metric spaces, which is 2.78-competitive. Jaillet and Wagner [6] gave a $(2 + \epsilon)$-competitive polynomial-time algorithm for Euclidean spaces with dimension $d \geq 2$, for any $\epsilon > 0$.

1.2 Our Contributions

We first consider single server online routing problems with precedence and capacity constraints. We give an online algorithm that is 2-competitive; the power of this statement is that adding precedence and capacity constraints to the online TSP does not increase the competitive ratio. Considering polynomial-time algorithms, a modification to our algorithm is 2ρ-competitive, where ρ is the approximation ratio of a simpler offline problem. Next, we study multiple server routing problems (without precedence and capacity constraints) and show similar results to those just mentioned. We design a new algorithm with a competitive ratio of 2, a result that is best-possible. A modification of our algorithm again results in a polynomial-time online algorithm that is 2ρ-competitive. We note that adding servers to the problem statement does not increase nor decrease the competitive ratio with respect to the online TSP.

2 Single Server Routing Problems

We first consider routing problems where a single server must service a sequence of requests. The data for our problems is a set of points $(\mathbf{l}_i, r_i, \mathbf{d}_i),\ i = 1, \ldots, n$, where n is the number of requests and $k(i)$ is the number of cities in request i: $\mathbf{l}_i = (l_i^1, l_i^2, \ldots, l_i^{k(i)})$ and $\mathbf{d}_i = (d_i^1, d_i^2, \ldots, d_i^{k(i)})$. The quantity $l_i^j \in \mathcal{M}$, $\mathcal{M}$ an arbitrary metric space, is the location of the j^{th} city in the i^{th} request. The quantity $r_i \in \mathbb{R}_+$ is the i^{th} request's release date; i.e., r_i is the first time after which that cities in request i will accept service. We assume, without loss of generality, that $r_1 \leq r_2 \leq \cdots \leq r_n$. The quantity $d_i^j \in \mathbb{R}_+$ is the demand of city l_i^j. The server has a capacity[1] Q; the sum of city demands visited on any given tour can be at most Q; we assume $d_i^j \leq Q$ for all i, j. Precedence constraints exist *within* a request; i.e., for a fixed i, arbitrary precedence constraints of the form $l_i^j \preceq l_i^k$ (l_i^j must be visited before l_i^k) for any $j \neq k$. The service requirement at a city is zero. Unless stated otherwise, the server travels at unit speed or is idle. The problem begins at time 0, and the server is initially at a designated origin o of the metric space. The objective is to minimize the time required to visit all cities and have the server return to the origin. We also let $\mathcal{N} = \{1, \ldots, n\}$.

From the online perspective, the total number of requests, represented by the parameter n, is not known, and request i only becomes known at time r_i. $Z_n^A(Q)$ denotes the cost of online algorithm A on an instance of n cities with server capacity Q and $Z_n^*(Q)$ is the corresponding optimal offline cost where all data is known a priori. $Z_n^{r=0}(Q)$ is the optimal cost when all release dates are equal to zero; clearly, $Z_n^{r=0}(Q) \leq Z_n^*(Q)$. The problem instance underlying $Z_n^{r=0}(Q)$, $Z_n^A(Q)$ and $Z_n^*(Q)$ will be clear from context. At times, the n term will be suppressed. Finally, define L_{TSP} as the optimal TSP tour length through all cities in an instance; i.e., $L_{TSP} = Z^{r=0}(\infty)$.

We measure the performance of online algorithms using the competitive ratio. The competitive ratio is defined as the worst-case ratio, over all problem instances, of online to offline costs: $\max_{instances} Z^A(Q)/Z^*(Q)$. An online algorithm is also said to be c-competitive if its competitive ratio is at most c. An online algorithm is said to be best-possible if there does not exist another online algorithm with a strictly smaller competitive ratio.

We give an online algorithm that generalizes PAH, which was given by Ausiello et al. [3]; we denote our algorithm Plan-At-Home-Generalized (PAH-G). Note that the competitive ratio of the original PAH is 2.

Algorithm 1 : PAH-G

(1) Whenever the server is at the origin, it calculates and implements a ρ-approximate solution to $Z^{r=0}(Q)$ over all requests whose release dates have passed but have not yet been served

[1] It is possible to generalize our capacity model to allow positive and negative demands as well as different types of products being transfered. However, we study the current problem to limit the complexity of the analysis.

completely.

(2) If at time r_i, for some i, a new request is presented, the server takes one of two actions depending on the server's current position p and the farthest location in the current request l_i^*:

$$l_i^* = \arg \max_{\{l_i^j \mid 1 \le j \le k(i)\}} d(o, l_i^j) :$$

(2a) If $d(l_i^*, o) > d(p, o)$, the server goes back to the origin where it appears in a Case (1) situation.

(2b) If $d(l_i^*, o) \le d(p, o)$, the server ignores request i until it completes the route it is currently traversing, where again Case (1) is encountered.

Theorem 1 *Algorithm PAH-G is 2ρ-competitive.*

As an example, if we consider the online capacitated TSP without precedence constraints, we can apply the Iterated Tour Partition (ITP) heuristics given by Altinkemer and Gavish [1] and Haimovich and Rinnooy Kan [5]. If $d_i^j = 1$ for all i, j, there exists a ITP heuristic with approximation ratio $\rho \le (5/2 + 3/2Q)$. If demands are arbitrary, there exists a ITP heuristic with approximation ratio $\rho \le (7/2 - 3/Q)$.

This result shows interesting properties. First, it is possible to relate the competitive ratio of PAH-G to the approximation ratio of a simpler but related optimization problem $Z^{r=0}(Q)$. Also, if we have access to exact algorithms for $Z^{r=0}(Q)$, adding capacity and precedence constraints results in no increase in the competitive ratio, with respect to the online TSP.

3 Multiple Server Routing Problems

We now consider routing problems with m identical servers. We do not consider capacity or precedence constraints. The data for our multiple server problems is closely related to that of the single server problems: the data is a set of points (l_i, r_i), $i = 1, \dots, n$ where $l_i \in \mathcal{M}$ ($r_i \in \mathbb{R}^+$) is the location (release date) of the i-th request. We again assume, without loss of generality, that $r_1 \le r_2 \le \cdots \le r_n$. The service requirement at a city is again zero. Unless stated otherwise, the servers travel at unit speed or are idle. The problem begins at time 0, and all servers are initially at a designated origin o of the metric space. The objective is to minimize the time required to visit all cities and have all servers return to the origin.

$Z_n^A(m)$ denotes the cost of online algorithm A on an instance of n cities with m identical servers and $Z_n^*(m)$ is the corresponding optimal offline cost where all data are known a priori; we assume $n \ge m$. $Z_n^{r=0}(m)$ is the optimal cost when all release dates are equal to zero; clearly, $Z_n^{r=0}(m) \le Z_n^*(m)$. Note that $Z_n^{r=0}(m)$ is equivalent to the problem of finding a set of m tours, that collectively visit all locations, such that the maximum tour length is minimized; e.g., see [4]. The problem instance underlying $Z_n^{r=0}(m)$, $Z_n^A(m)$ and $Z_n^*(m)$ will be clear from context. Finally, note that

$L_{TSP} = Z^{r=0}(1)$. The competitive ratio and competitiveness are defined similarly to the single server case.

We again give an online algorithm that generalizes PAH, which was given by Ausiello et al. [3]; we denote our algorithm Plan-At-Home-m-Servers (PAH-m).

Algorithm 2 : PAH-m

(1) Whenever all servers are at the origin, they calculate and implement a ρ-approximate solution to $Z^{r=0}(m)$ over all requests whose release dates have passed but have not yet been served.

(2) If at time r_i, for some i, a new request is presented, the servers take one of two actions depending on the request's location l_i and the farthest server's current position p^* (ties broken arbitrarily):

$$p^* = \arg \max_{\{p_i \mid 1 \le i \le m\}} d(o, p_i) :$$

(2a) If $d(l_i, o) > d(p^*, o)$, all servers go back to the origin where they appear in a Case (1) situation.

(2b) If $d(l_i, o) \le d(p^*, o)$, all servers **except** *p^* return to the origin; server p^* ignores request i until it completes the route it is currently traversing, where again Case (1) is encountered.*

Theorem 2 *Algorithm PAH-m is 2ρ-competitive.*

As an example, we can apply the approximation algorithm for $Z^{r=0}(m)$ given by Frederickson, Hecht and Kim [4] that has an approximation ratio $\rho \le 5/2 - 1/m$.

Theorem 3 *If we use an exact algorithm in step(1) for calculating an optimal offline $Z^{r=0}(m)$, the competitive ratio of PAH-m is 2 and this result is best-possible.*

Again, it is possible to relate the competitive ratio to the approximation ratio of a simpler but related optimization problem $Z^{r=0}(m)$. Also, if we have access to exact offline algorithms for $Z^{r=0}(m)$, adding extra vehicles results in no change (increase or decrease) in the competitive ratio, with respect to the the online TSP.

4 Conclusion

The focus of this extended abstract has been on generalizations of the online traveling salesman problems to allow for precedence constraints, capacity constraints, and multiple vehicles. We derived competitive ratio results for these online problems, several being best-possible.

We conclude by mentioning two areas of additional current research we are pursuing. The introduction of accept/reject decisions adds a rich dimension to the problems considered in this extended abstract – the online algorithm has the ability to accept or reject a given request. This introduces additional difficulties however; under this framework is it easy to create problem instances with unbounded competitive ratios. Therefore, new measures of online routing algorithms under accept/reject decisions must be designed and utilized. Another area of research is to investigate the value of varying degrees of information about the problem data. For example, if there were a service time at each location, how much would it be worth to know the exact service time, a distribution for the service time, etc. A rich variety of problems are available for investigation.

References

[1] K. Altinkemer and B. Gavish. Heuristics for delivery problems with constant error estimates. *Transportation Science*, 24:295–297, 1990.

[2] G. Ausiello, M. Demange, L. Laura, and V. Paschos. Algorithms for the on-line quota traveling salesman problem. *Information Processing Letters*, 92(2):89–94, 2004.

[3] G. Ausiello, E. Feuerstein, S. Leonardi, L. Stougie, and M. Talamo. Algorithms for the on-line travelling salesman. *Algorithmica*, 29(4):560–581, 2001.

[4] G. Frederickson, M. Hecht, and C. Kim. Approximation algorithms for some routing problems. *SIAM Journal on Computing*, 7(2):178–193, 1978.

[5] M. Haimovich and A.H.G. Rinnooy Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10:527–542, 1985.

[6] P. Jaillet and M. Wagner. Online routing problems: value of advanced information as improved competitive ratios. *Transportation Science*, 40(2), 2006.