The On-line Asymmetric Traveling Salesman Problem [1]

Giorgio Ausiello [a], Vincenzo Bonifaci [a,b,*], Luigi Laura [a]

[a] *University of Rome "La Sapienza", Department of Computer and Systems Science, Via Salaria, 113 - 00198 Rome, Italy.*

[b] *Eindhoven University of Technology, Department of Mathematics and Computer Science, Den Dolech 2 - Postbus 513, 5600 MB Eindhoven, The Netherlands.*

Abstract

We consider two on-line versions of the asymmetric traveling salesman problem with triangle inequality. For the *homing* version, in which the salesman is required to return in the city where it started from, we give a $\frac{3+\sqrt{5}}{2}$-competitive algorithm and prove that this is best possible. For the *nomadic* version, the on-line analogue of the shortest asymmetric hamiltonian path problem, we show that the competitive ratio of any on-line algorithm depends on the amount of asymmetry of the space in which the salesman moves. We also give bounds on the competitive ratio of on-line algorithms that are *zealous*, that is, in which the salesman cannot stay idle when some city can be served.

Key words: On-line algorithms, competitive analysis, real time vehicle routing, asymmetric traveling salesman problem

---

1 Introduction

In the classical traveling salesman problem, a set of cities has to be visited in a single tour with the objective of minimizing the total length of the tour. This is one of the most studied problems in combinatorial optimization, together with

---

* Corresponding author.

Email addresses: `ausiello@dis.uniroma1.it` (Giorgio Ausiello), `bonifaci@dis.uniroma1.it` (Vincenzo Bonifaci), `laura@dis.uniroma1.it` (Luigi Laura).

[1] A preliminary version appeared in *Proceedings of the 9th Int. Workshop on Algorithms and Data Structures*, Waterloo, Canada, 2005, pp. 306–317.

its dozens of variations [11,16]. In the asymmetric version of the problem, the distance from one point to another in a given space can be different from the inverse distance. This variation, known as the Asymmetric Traveling Salesman Problem (ATSP) arises in many applications; for example, one can think of a delivery vehicle traveling through one-way streets in a city, or of gasoline costs when traveling through mountain roads.

The ATSP has been much studied from the point of view of approximation algorithms. However, if the condition is that every city or place has to be visited *exactly* once, the problem is **NPO**-complete and thus essentially no approximation is possible in polynomial time, unless **P**=**NP** [19]. Instead, in the case where every city or place given in the input has to be visited *at least* once or, equivalently, the distance function satisfies the triangular inequality, approximation algorithms exist having a reasonable approximation factor. In particular, the best algorithms known have an approximation ratio of $O(\log n)$ [10,13]. The problem is also known to be **APX**-hard [18]. The question of the existence of an algorithm with a constant approximation ratio for the asymmetric case is still open after more than two decades.

Here we are interested in the on-line version of the ATSP, named OL-ATSP. The on-line versions of a number of vehicle routing problems, including the standard TSP, the traveling repairman problem, the quota TSP and dial-a-ride problems have been studied recently [2,3,4,8,14,15,17]. In the on-line TSP and ATSP, the places to visit in the space are requested over time and a server (the salesman or vehicle) has to decide in what order to serve them, without knowing the entire sequence of requests beforehand. The objective is to minimize the completion time of the server. To analyze our algorithms, we use the established framework of *competitive analysis* [5,9,20], where the cost of the algorithm being studied is compared to that of an ideal optimum off-line server, knowing in advance the entire sequence of requests (notice, however, that even the off-line server cannot serve a request before it is released). The ratio between the on-line and the off-line costs is called the *competitive ratio* of the algorithm and is a measure of the loss of efficacy due to the absence of information on the future. Our paper is the first to address the on-line ATSP from the point of view of competitive analysis. Previous work, both theoretical and experimental, has focused on the off-line version [7,10,13].

Our results are summarized in Table 1, where they are also compared with the known results for the symmetric case. As we will see, the asymmetric TSP is substantially harder than the normal TSP even when considered from an on-line point of view; in other words, OL-ATSP is not a trivial extension of OL-TSP. In fact, as Table 1 shows, most bounds on the competitive ratio are strictly higher than the corresponding bounds for OL-TSP, and in particular in the nomadic case there cannot be on-line algorithms with a constant competitive ratio.

Although the algorithmic techniques we adopt in the asymmetric case come essentially from the symmetric case, they require some adjustment in order to attain useful competitive ratios. On the other hand, it is worth noting that the lower bound techniques are quite different from the previously known ones and we hope they can be of some use in future work.

We should also mention that we present our algorithms in simplified versions that compute optimal traveling salesman tours or paths. Thus, they do not run in polynomial time unless **P**=**NP**. However, if one is interested in polynomial running time, it is possible to compute approximately optimal tours instead, the competitive ratio degrading by a factor that is essentially the approximation ratio of the subroutine being used. For example, as a consequence of our results, an $O(1)$ approximation algorithm for the ATSP would automatically imply an $O(1)$-competitive polynomial time algorithm for the OL-ATSP. We further discuss this issue in Section 5.

The rest of this paper is organized as follows. After the necessary definitions and the discussion of the model, we study in Section 3 the homing case of the problem, in which the server is required to finish its tour in the same place where it started; we give a $\frac{3+\sqrt{5}}{2}$-competitive algorithm and show that this is also best possible. In Section 4, we address the nomadic version, also known as the wandering traveling salesman problem [12], in which the server is not required to finish its tour at the origin. For this case we show that in general an on-line algorithm with a competitive ratio independent of the space cannot exist; indeed, we show that the competitive ratio has to be a function of the amount of asymmetry of the space. In Section 5 we explain how our algorithms can be combined with polynomial time approximation algorithms in order to obtain polynomial time online algorithms. In the last section, we give our conclusions and discuss some open problems.

2 The model

An input for the OL-ATSP consists of a space M from the class $\mathcal{M}$ defined below, a distinguished point $O \in M$, called the origin, and a sequence of requests $r_i = (t_i, x_i)$ where x_i is a point of M and $t_i \in \mathbb{R}_+$ is the time when the request is presented. The sequence is ordered so that $i < j$ implies $t_i \leq t_j$.

The server is located at the origin O at time 0 and the distances are scaled so that, without loss of generality, the server can move at most at unit speed.

We will consider two versions of the problem. In the *nomadic* version, the server can end its route anywhere in the space; the objective is just to minimize the makespan, that is, the time required to serve all presented requests. In the

Problem	**Best Lower Bound**	**Best Upper Bound**	**References**
Homing OL-TSP	2	2	[1,3]
Homing OL-ATSP	$(3+\sqrt{5})/2$	$(3+\sqrt{5})/2$	
Homing OL-TSP (zealous)	2	2	[3]
Homing OL-ATSP (zealous)	3	3	
Nomadic OL-TSP	2.03	$1+\sqrt{2}$	[17]
Nomadic OL-ATSP	$\sqrt{K}$	$1+\sqrt{K+1}$	
Nomadic OL-TSP (zealous)	2.05	2.5	[3,17]
Nomadic OL-ATSP (zealous)	$(K+1)/2$	$K+2$	

Table 1
The competitive ratio of symmetric and asymmetric routing problems. Refer to Section 2 for the definition of K.

homing version, the objective is to minimize the time required to serve all presented requests and return to the origin.

An *on-line* algorithm for the OL-ATSP has to determine the behavior of the server at a certain moment t as a function only of the requests (t_i, x_i) such that $t_i \leq t$. Thus, an on-line algorithm does not have knowledge about the number of requests or about the time when the last request is released. We will use T to denote some tour or route over a subset of the requests; in this case, $|T|$ will be the length of that tour.

We will use Z^{OL} to denote the completion time of the solution produced by a generic on-line algorithm OL, while Z^* will be the completion time of the optimal off-line solution. An on-line algorithm OL is *c-competitive* if, for any sequence of requests, $Z^{\mathrm{OL}} \leq cZ^*$.

Finally, we would like to clarify the conditions that the space M should satisfy. Usually, in the context of the on-line TSP, continuous path-metric spaces are considered [3]. However, here the main issue is precisely asymmetry, so we have to drop the requisite that for every x and y, $d(x, y) = d(y, x)$. We review here the definitions. A set M, equipped with a distance function $d : M^2 \to \mathbb{R}_+$, is called a *quasi-metric space* if, for all $x, y, z \in M$:

(i) $d(x, y) = 0$ if and only if $x = y$;
(ii) $d(x, y) \leq d(x, z) + d(z, y)$.

We call a space M an *admissible space* if M is a quasi-metric and, for any $x, y \in M$, there is a function $f : [0, 1] \to M$ such that $f(0) = x$, $f(1) = y$ and

f is continuous, in the following sense: $d(f(a), f(b)) = (b - a)d(x, y)$ for any $0 \leq a \leq b \leq 1$. This function represents a *shortest path* from x to y. Notice that every admissible space is connected.

We will use $\mathcal{M}$ to denote the class of admissible spaces. Notice that the discrete metric induced by a weighted graph is not admissible if we take M to be the set of vertices. However, we can always make such a space admissible by adding (an infinity of) extra points "along the arcs".

In particular, to see how a directed graph with positive weights on the arcs can define an admissible space, consider the all-pairs shortest paths matrix of the graph. This defines a finite quasi-metric. Now we add, for every arc $a = (x, y)$ of the graph, an infinity of points π^a_γ, indexed by a parameter $\gamma \in (0, 1)$. Let π^a_0 and π^a_1 denote x and y respectively. We extend the distance function d so that:

$$d(\pi^a_\gamma, \pi^a_{\gamma'}) = (\gamma' - \gamma)d(x, y) \text{ for all } 0 \leq \gamma < \gamma' \leq 1.$$

It can be verified that π^a represents a shortest path from x to y. For $\gamma \notin \{0, 1\}$, the distance from a point π^a_γ to a point z not in π^a is defined as $d(\pi^a_\gamma, z) = d(\pi^a_\gamma, y) + d(y, z)$; that is, the shortest path from π^a_γ to z passes through y. Vice versa, the distance from z to π^a_γ is defined as $d(z, \pi^a_\gamma) = d(z, x) + d(x, \pi^a_\gamma)$. Finally,

$$d(\pi^a_{\gamma'}, \pi^a_\gamma) = (1 - (\gamma' - \gamma))d(x, y) + d(y, x) \text{ for all } 0 \leq \gamma < \gamma' \leq 1.$$

We say that such a space is *induced* by the original directed weighted graph. We remark that this model, while still including the originally proposed one [3] as a special case, can also capture the situation in which the server is not allowed to do U-turns.

Finally, it will be useful to have a measure of the amount of asymmetry of a space. Define as the *maximum asymmetry* of a space $M \in \mathcal{M}$ the value

$$K(M) = \sup_{x,y \in M} \frac{d(x, y)}{d(y, x)}.$$

We will say that a space M has *bounded asymmetry* when $K(M) < \infty$.

3 **Homing** OL-ATSP

In this section we consider the homing version of the on-line ATSP, in which the objective is to minimize the completion time required to serve all presented requests and return to the origin. We establish a lower bound of about 2.618 and a matching upper bound. Note that in the case of symmetric on-line TSP, the corresponding bounds are both equal to 2 [3,14].

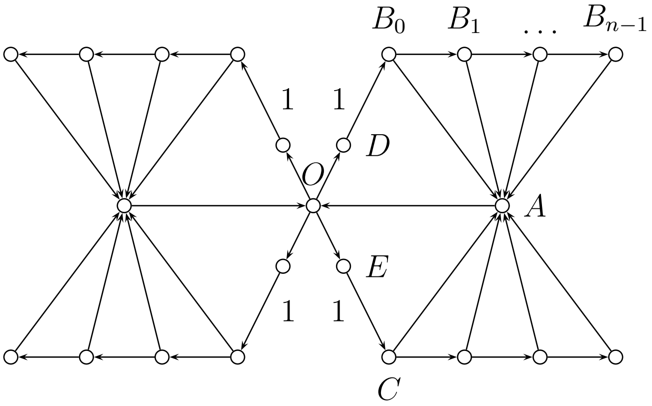

Figure 1. The graph used in the proof of Theorem 3.1.

Let ϕ denote the golden ratio, that is, the unique positive solution to $x = 1 + 1/x$. In closed form, $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$.

Theorem 3.1 *The competitive ratio of any on-line algorithm for homing* OL-ATSP *is at least* $1 + \phi$.

Proof. Fix any $\epsilon > 0$. The space used in the proof is the one induced by the graph depicted in Figure 1. The graph has $7+4n$ nodes, where $n = 1 + \lceil \frac{\phi-1}{\epsilon} \rceil$, and the length of every arc is ϵ, except for those labeled otherwise. Observe that the space is symmetric with respect to an imaginary vertical axis passing through O. Thus, we can assume without loss of generality that, at time 1, no request being released yet, the on-line server is in the left half of the space. Then at time 1 a request is given in point A, in the other half. Now let t be the first time at which the on-line server reaches point D or E.

If $t \geq \phi$, no further request is given. In this case $Z^{\text{OL}} \geq t + 1 + 2\epsilon$ while $Z^* \leq 1+3\epsilon$ so that, when ϵ approaches zero, Z^{OL}/Z^* approaches $1+t \geq 1+\phi$.

Otherwise, if $t \in [1, \phi]$, at time t, we can assume that the on-line server has just reached E (again, by symmetry). At this time, the adversary gives a request in B_i, where $i = \lceil \frac{t-1}{\epsilon} \rceil$. Now the on-line server has to traverse the entire arc EC before it can go serve B_i, thus

$$Z^{\text{OL}} \geq t + 1 + 3\epsilon + 1 + \epsilon \left\lceil \frac{t-1}{\epsilon} \right\rceil + 2\epsilon \geq 2t + 1 + 5\epsilon.$$

Instead, the adversary server will have moved from O to B_i in time at most $t+2\epsilon$ and then served B_i and A, achieving the optimal cost $Z^* \leq t+4\epsilon$. Thus, when ϵ approaches zero, Z^{OL}/Z^* approaches $2 + \frac{1}{t} \geq 1 + \phi$. $\square$

To prove a matching upper bound on the competitive ratio, we use a variation of algorithm SMARTSTART, first introduced by Ascheuer et al. [1].

Algorithm SMARTSTART(α)
The algorithm keeps track, at every time t, of the length of an optimal tour $T^*(t)$ over the unserved requests, starting at and returning to the origin. At the first instant t' such that $t' \geq \alpha|T^*(t')|$, the server starts following at full speed the currently optimal tour, ignoring temporarily every new request. When the server is back at the origin, it stops and returns monitoring the value $|T^*(t)|$, starting as before when necessary. As we will soon see, the best value of α is $\alpha^* = \phi$.

Theorem 3.2 SMARTSTART(ϕ) *is* $(1+\phi)$*-competitive for homing* OL-ATSP.

Proof. We distinguish two cases depending on whether the last request arrives while the server is waiting at the origin or not.

In the first case, let t be the release time of the last request. If the server starts immediately at time t, it will follow a tour of length $|T^*(t)| \leq t/\alpha$, ending at time at most $(1 + 1/\alpha)t$, while the adversary pays at least t, so the competitive ratio is at most $1 + 1/\alpha$. Otherwise, the server will start at a time $t' > t$ such that $t' = \alpha|T^*(t)|$ (since T^* does not change after time t) and pay $(1 + \alpha)|T^*(t)|$, so the competitive ratio is at most $1 + \alpha$.

In the second case, let $T^*(t)$ be the tour that the server is following while the last request arrives; that is, we take t to be the starting time of that tour. Let $T'(t)$ be an optimal tour over the requests released *after* time t. If the server has time to wait at the origin when it finishes following $T^*(t)$, the analysis is the same as in the first case. Otherwise, the completion time of SMARTSTART is $t + |T^*(t)| + |T'(t)|$. Since SMARTSTART has started following $T^*(t)$ at time t, we have $t \geq \alpha|T^*(t)|$. Then

$$t + |T^*(t)| \leq (1 + 1/\alpha)t.$$

Also, if $r_f = (t_f, x_f)$ is the first request served by the adversary having release time at least t, we have that $|T'(t)| \leq d(O, x_f) + Z^* - t$ (recall that Z^* is the off-line cost), since a possibility for T' is to go to x_f and then do the same as the adversary (subtracting t from the cost since we are computing a length, not a completion time, and on the other hand the adversary will not serve r_f at a time earlier than t).

By putting everything together, we have that SMARTSTART pays at most

$$(1 + 1/\alpha)t + d(O, x_f) + Z^* - t$$

and since two obvious lower bounds on Z^* are t and $d(O, x_f)$, this is easily seen to be at most $(2 + 1/\alpha)Z^*$.

Now $\max\{1 + \alpha, 2 + \frac{1}{\alpha}\}$ is minimum when $\alpha = \alpha^* = \phi$. For this value of the parameter the competitive ratio is $1 + \phi$. □

3.1 Zealous algorithms

In the previous section we have seen that the optimum performance is achieved by an algorithm that, before starting to serve requests, waits until a convenient starting time is reached. In this section we consider instead the performance that can be achieved by *zealous* algorithms [4]. A zealous algorithm does not change the direction of its server unless a new request becomes known, or the server is at the origin or at a request that has just been served; furthermore, a zealous algorithm moves its server always at full (that is, unit) speed when there are unserved requests.

We show that, for zealous algorithms, the competitive ratio has to be at least 3 and, on the other hand, we give a matching upper bound.

Theorem 3.3 *The competitive ratio of any zealous on-line algorithm for homing* OL-ATSP *is at least* 3*.*

Proof. We use the same space used in the lower bound for general algorithms (Figure 1). At time 1, the server has to be at the origin and the adversary gives a request in A. Thus, at time $1+\epsilon$ the server will have reached wlog E (by symmetry) and the adversary gives a request in B_0. The completion time of the on-line algorithm is at least $3+6\epsilon$, while $Z^* \leq 1+3\epsilon$. The result follows by taking a sufficiently small ϵ. □

The following algorithm is best possible among the zealous algorithms for homing OL-ATSP.

Algorithm PLAN AT HOME

When the server is at the origin and there are unserved requests, the algorithm computes an optimal tour over the set of unserved requests and the server starts following it, ignoring temporarily every new request, until it finishes its tour at the origin. Then it waits at the origin as before.

Theorem 3.4 PLAN AT HOME *is zealous and* 3*-competitive for homing* OL-ATSP*.*

Proof. Let t be the release time of the last request. If $p(t)$ is the position of PLAN AT HOME at time t and T is the tour it was following at that time, we have that PLAN AT HOME finishes following T at time $t' \leq t+|T|$. At that time, it will eventually start again following a tour over the requests which remain unserved at time t'. Let us call T' this other tour. The total cost payed by PLAN AT HOME will be then at most $t+|T|+|T'|$. But $t \leq Z^*$, since even the off-line adversary cannot serve the last request before it is released, and on the other hand both T and T' have length at most Z^*, since the

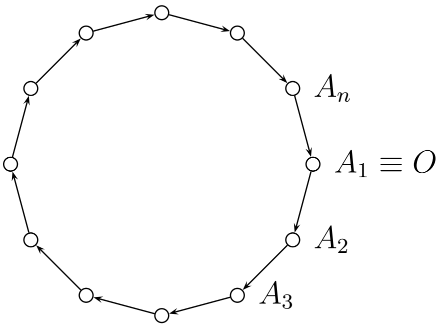

Figure 2. The graph used in the proof of Theorem 4.1.

off-line adversary has to serve all of the requests served in T and T'. Thus, $t + |T| + |T'| \leq 3Z^*$. □

4 **Nomadic** OL-ATSP

In this section we consider the nomadic version of the on-line ATSP, in which the server can end its route anywhere in the space. We show that no on-line algorithm can have a constant competitive ratio (that is, independent of the underlying space). Then we show, for spaces with a maximum asymmetry K, a lower bound $\sqrt{K}$ and an upper bound $1 + \sqrt{K+1}$. Note that in the case of symmetric nomadic on-line TSP, the best lower and upper bounds are 2.03 and $1 + \sqrt{2}$, respectively [17].

Theorem 4.1 *For every $L > 0$, there is a space $M \in \mathcal{M}$ such that the competitive ratio of any on-line algorithm for nomadic* OL-ATSP *on M is at least L.*

Proof. Let $\epsilon \leq 1/(2L+1)$, and consider the space induced by a directed cycle on $n = \lceil \frac{L}{\epsilon} \rceil$ nodes, where every arc has length ϵ (Figure 2). At time 0 a request is given in node A_3. Let t be the first time the on-line algorithm reaches node A_2.

Now if $t \geq 1$, the adversary does not release any other request so that $Z^* = 2\epsilon$, $Z^{\text{OL}} \geq 1 + \epsilon$ and $Z^{\text{OL}}/Z^* \geq \frac{1}{2\epsilon} + \frac{1}{2} \geq L$.

Otherwise, if $t \leq 1$, at time t the adversary releases a request at the origin. It is easily seen that $Z^* \leq t + 2\epsilon$ and $Z^{\text{OL}} \geq t + \epsilon(\lceil \frac{L}{\epsilon} \rceil - 1) \geq t + 2\epsilon + L - 3\epsilon$ so that

$$Z^{\mathrm{OL}}/Z^* \geq 1 + \frac{L - 3\epsilon}{t + 2\epsilon} \geq 1 + \frac{L - 3\epsilon}{1 + 2\epsilon} \geq L.$$

□

Corollary 4.2 *There is no on-line algorithm for nomadic* OL-ATSP *on all spaces $M \in \mathcal{M}$ with a constant competitive ratio.*

We also observe that the same lower bound can be used when the objective function is the sum of completion times.

Thus, we cannot hope for an on-line algorithm which is competitive for all spaces in $\mathcal{M}$. Indeed, we will now show that the amount of asymmetry of a space is related to the competitive ratio of any on-line algorithm for that space.

Theorem 4.3 *For every $K \geq 1$, there is a space $M \in \mathcal{M}$ with maximum asymmetry K such that any on-line algorithm for nomadic* OL-ATSP *on M has competitive ratio at least $\sqrt{K}$.*

Proof. Consider a set of points $M = \{x_\gamma : \gamma \in [0, 1]\}$ with a distance function

$$d(x_\gamma, x_{\gamma'}) = \begin{cases} \gamma' - \gamma & \text{if } \gamma \leq \gamma' \\ K(\gamma - \gamma') & \text{if } \gamma \geq \gamma'. \end{cases}$$

The origin is x_0. The adversary releases a request at time 1 in point x_1. Let t be the time the on-line algorithm serves this request. If $t \geq \sqrt{K}$, no more requests are released and $Z^{\mathrm{OL}} \geq \sqrt{K}$, $Z^* = 1$, $Z^{\mathrm{OL}}/Z^* \geq \sqrt{K}$.

Otherwise, if $t \leq \sqrt{K}$, at time t a request is given at the origin. Now $Z^{\mathrm{OL}} \geq t + K$, $Z^* \leq t + 1$ and

$$Z^{\mathrm{OL}}/Z^* \geq \frac{t + K}{t + 1} = 1 + \frac{K - 1}{t + 1} \geq 1 + \frac{K - 1}{\sqrt{K} + 1} = \sqrt{K}.$$

□

A natural algorithm, on the lines of the best known algorithm for the symmetric version of the problem [17], gives a competitive ratio which is asymptotically the same as that of this lower bound.

Algorithm RETURN HOME(β)
At any moment at which a new request is released, the server returns to the origin via the shortest path. Once at the origin at time t, it computes an optimal route T over all requests presented up to time t and then starts following T at the highest possible speed ensuring that $d(O, s(t')) \leq \beta t'$ at any time t' (here $s(t')$ is the position of the server at time t').

Theorem 4.4 *For every space $M \in \mathcal{M}$ with maximum asymmetry K, there is a value of β such that* RETURN HOME(β) *is $(1 + \sqrt{K+1})$-competitive on M.*

Proof. There are two cases to consider. In the first case RETURN HOME does not need to reduce its speed after the last request is released. In this case, if t is the release time of the last request, we have

$$Z^{\mathrm{RH}} \leq t + K\beta t + |T| \leq Z^* + K\beta Z^* + Z^* = (2 + K\beta)Z^*.$$

In the second case, let t be the last time RETURN HOME is moving at reduced speed. At that time, RETURN HOME has to be serving some request; let x be the location of that request. Since RETURN HOME is moving at reduced speed we must have $d(O, x) = \beta t$; afterwards RETURN HOME will follow the remaining part T_x of the route at full speed. Thus

$$Z^{\mathrm{RH}} \leq t + |T_x| = (1/\beta)d(O, x) + |T_x|.$$

On the other hand, $Z^* \geq |T| \geq d(O, x) + |T_x|$. Thus, in this case, the competitive ratio is at most $1/\beta$.

Obviously, we can choose β in order to minimize $\max\{2 + K\beta, 1/\beta\}$. This gives a value of $\beta^* = \frac{\sqrt{K+1}-1}{K}$, for which we obtain the competitive ratio of the theorem. □

4.1 Zealous algorithms

Also in the case of the nomadic version of the on-line ATSP, we wish to consider the performance of zealous algorithms. Of course, no zealous algorithm will be competitive for spaces with unbounded asymmetry. Here we show that the gap between non-zealous and zealous algorithms is much higher than in the homing case, the competitive ratio increasing from $\Theta(\sqrt{K})$ to $\Theta(K)$.

Theorem 4.5 *For every $K \geq 1$, there is a space $M \in \mathcal{M}$ with maximum asymmetry K such that the competitive ratio of any zealous on-line algorithm for nomadic* OL-ATSP *on M is at least $\frac{1}{2}(K + 1)$.*

Proof. We use the same space used in the proof of Theorem 4.3. At time 0, the adversary releases a request in point x_1. The on-line server will be at point x_1 exactly at time 1. Then, at time 1, the adversary releases a request in point x_0. It is easy to see that $Z^{\mathrm{OL}} \geq 1 + K$, while $Z^* = 2$. □

We finally observe that RETURN HOME(1) is a zealous algorithm for nomadic OL-ATSP and, by the proof of Theorem 4.4, it has competitive ratio $K + 2$.

Problem	**Algorithm**	**Competitive ratio**
Homing OL-ATSP	SMARTSTART(α^*_ρ)	$(1+2\rho+\sqrt{1+4\rho})/2$
	PLAN AT HOME	$1+2\rho$
Nomadic OL-ATSP	RETURN HOME(β^*_ρ)	$(1+\rho+\sqrt{(1+\rho)^2+4K})/2$
	RETURN HOME(1)	$1+\rho+K$

Table 2
The competitive ratio as a function of ρ and K.

5 Polynomial time algorithms

None of the algorithms that we have proposed in the previous sections runs in polynomial time, since all of them need to compute optimal tours on some subsets of the requests. On the other hand, a polynomial time on-line algorithm with a constant competitive ratio could be used as an approximation algorithm for the ATSP, and thus we do not expect to find one easily. However, our algorithms use off-line optimization as a black box and thus can use approximation algorithms as subroutines in order to give polynomial time on-line algorithms, the competitive ratio depending of course on the approximation ratio. In particular, in the homing version we need to solve instances of the off-line ATSP. The best polynomial time algorithm known for this problem has an approximation ratio of $0.842 \log n$ [13]. For the nomadic version, the corresponding off-line problem is the shortest asymmetric hamiltonian path, which also admits $O(\log n)$ approximation in polynomial time [6].

We do not repeat here the proofs of our theorems taking into account the approximation ratio of the off-line solvers, since they are quite straightforward. However, we give the competitive ratio of our algorithms as a function of ρ, the approximation ratio, and K, the maximum asymmetry of the space, in Table 2. Note that, with respect to the values in Table 1, the competitive ratio becomes worse by a factor that is strictly less than the approximation ratio. In the case of SMARTSTART and RETURN HOME, this is also due to the fact that the algorithms can adapt to the approximation ratio by suitably choosing the parameters α and β. For SMARTSTART, the optimal choice is

$$\alpha^*_\rho = \frac{1}{2\rho}\Big(1+\sqrt{1+4\rho}\Big),$$

while for RETURN HOME it is

$$\beta^*_\rho = \frac{1}{2K}\Big[\sqrt{(1+\rho)^2+4K}-(1+\rho)\Big].$$

6 Conclusions

We have examined several on-line variations of the asymmetric traveling salesman problem. Our results confirm that these asymmetric variations are indeed strictly harder than their symmetric counterparts.

The main conclusion is that, as usual in on-line vehicle routing when minimizing the completion time, waiting can improve the competitive ratio substantially. This is particularly evident in the case of nomadic ATSP on spaces with bounded asymmetry, where zealous algorithms have competitive ratio $\Omega(K)$ while RETURN HOME is $O(\sqrt{K})$-competitive.

We expect the competitive ratio of the homing OL-ATSP to be somewhat lower than $1+\phi$ when the space has bounded asymmetry. Also, since the proof that no on-line algorithm can have a constant competitive ratio in the nomadic case also applies when the objective function is the sum of completion times (the *traveling repairman* problem [15]), it would be interesting to investigate this last problem in spaces with bounded asymmetry.

Finally, we remark that the existence of polynomial time $O(1)$-competitive algorithms for the homing version is indissolubly tied to the existence of an $O(1)$-approximation algorithm for the off-line ATSP.

Acknowledgments

The second author would like to thank Luca Allulli for some helpful discussion.

References

[1] N. Ascheuer, S. O. Krumke, and J. Rambau. Online dial-a-ride problems: Minimizing the completion time. In H. Reichel and S. Tison, editors, *Proc. 17th Symp. on Theoretical Aspects of Computer Science*, volume 1770 of *Lecture Notes in Computer Science*, pages 639–650. Springer-Verlag, 2000.

[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] M. Blom, S. O. Krumke, W. E. de Paepe, and L. Stougie. The online TSP against fair adversaries. *INFORMS Journal on Computing*, 13(2):138–148, 2001.

[5] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.

[6] C. Chekuri and M. Pál. An $O(\log n)$ approximation ratio for the asymmetric traveling salesman path problem. Manuscript, available at `http://cm.bell-labs.com/cm/cs/who/chekuri/pub.html`, 2005.

[7] J. Cirasella, D. S. Johnson, L. A. McGeoch, and W. Zhang. The asymmetric traveling salesman problem: Algorithms, instance generators, and tests. In *Proc. 3rd Int. Workshop on Algorithm Engineering and Experimentation*, pages 32–59, 2001.

[8] E. Feuerstein and L. Stougie. On-line single-server dial-a-ride problems. *Theoretical Computer Science*, 268(1):91–105, 2001.

[9] A. Fiat and G. J. Woeginger, editors. *Online Algorithms: The State of the Art.* Springer-Verlag, 1998.

[10] A. M. Frieze, G. Galbiati, and F. Maffioli. On the worst-case performance of some algorithms for the asymmetric traveling salesman problem. *Networks*, 12(1):23–39, 1982.

[11] G. Gutin and A. P. Punnen, editors. *The Traveling Salesman Problem and its Variations.* Kluwer, Dordrecht, The Nederlands, 2002.

[12] M. Jünger, G. Reinelt, and G. Rinaldi. The traveling salesman problem. In M. O. Ball, T. Magnanti, C. L. Monma, and G. Nemhauser, editors, *Network Models, Handbook on Operations Research and Management Science*, volume 7, pages 225–230. Elsevier, 1995.

[13] H. Kaplan, M. Lewenstein, N. Shafrir, and M. Sviridenko. Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. In *Proc. 44th Symp. on Foundations of Computer Science*, pages 56–66, 2003.

[14] S. O. Krumke. Online optimization: Competitive analysis and beyond. Habilitation Thesis, Technical University of Berlin, 2001.

[15] S. O. Krumke, W. E. de Paepe, D. Poensgen, and L. Stougie. News from the online traveling repairman. In J. Sgall, A. Pultr, and P. Kolman, editors, *Proc. 26th Symp. on Mathematical Foundations of Computer Science*, volume 2136 of *Lecture Notes in Computer Science*, pages 487–499. Springer-Verlag, 2001.

[16] E. L. Lawler, J. K. Lenstra, A. Rinnooy Kan, and D. B. Shmoys. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization.* Wiley, Chichester, England, 1985.

[17] M. Lipmann. *On-Line Routing.* PhD thesis, Eindhoven University of Technology, 2003.

[18] C. H. Papadimitriou and M. Yannakakis. The traveling salesman problem with distances one and two. *Mathematics of Operations Research*, 18(1):1–11, 1993.

[19] S. Sahni and T. F. Gonzalez. P-complete approximation problems. *Journal of the ACM*, 23(3):555–565, 1976.

[20] D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.