A Constant Approximation Algorithm for the Uniform A Priori Capacitated Vehicle Routing Problem with Unit Demands

Finn Fernstrøm[a], Teresa Anna Steiner[a,*]

[a]*DTU Compute, Richard Petersens Plads, Bygning 324, 2800 Kgs. Lyngby*

Abstract

In this paper we introduce an a priori variant of the capacitated vehicle routing problem and provide a constant factor approximation algorithm, when the demands of customers are independent and identically distributed Bernoulli experiments. In the capacitated vehicle routing problem (CVRP) a vehicle, starting at a depot, must visit a set of customers to deliver the requested quantity of some item. The vehicle has capacity k, which is the maximum number of items that the vehicle can carry at any time, but it can always return to the depot to restock. The objective is to find the shortest tour subject to these constraints. In the a priori CVRP with unit demands the vehicle has to visit a set of active customers, drawn from some distribution. Every active customer has a demand of one. The goal is to find a master tour, which is a feasible solution to the deterministic CVRP, i.e. where all customers are active. Then, given a set of active customers, the tour is shortcut to only visit those customers. The cost of the tour is the expected cost with respect to the distribution of active customers. We consider the model, where every customer is independently active with the same probability.

Let N be the number of customers. We provide an algorithm, which takes as input an a priori TSP solution with approximation factor γ, and gives a solution to the a priori CVRP with unit demands, whose cost is at most $(1+k/N+\gamma)$ times the value of the optimal solution. Specifically, this gives an expected $5.5-$approximation by using the $3.5-$approximation to the a priori TSP from Van Ee and Sitters (2018) and a deterministic $8.5-$approximation by using the deterministic $6.5-$approximation from Van Zuylen (2011).

*Corresponding author

Email addresses: finnfernstroem@gmail.com (Finn Fernstrøm), terst@dtu.dk (Teresa Anna Steiner)

Keywords: Vehicle Routing Problem, A Priori Optimization, Approximation Algorithms

1. Introduction

Vehicle routing problems are well studied problems, where the goal is to find the shortest delivery routes from a set of depots to a set of customers [1, 2, 7, 8, 10, 13, 14]. Each customer has a certain demand of goods which have to be delivered. Additionally, there may be certain constraints like vehicle capacity (i.e., a maximum amount of goods that a vehicle can carry at any time), or time constraints [13]. Traditionally, these problems are deterministic, meaning that all of the customers have to be served and they have a predetermined demand. In some settings stochastic variations of the problems are more realistic: Not every customer needs to be served every day, and demands will typically vary over time. Also dynamic variations exist, where for example new customers may appear during routes [14]. An overview of stochastic vehicle routing problems is given by Gendrau et al. [8]. One way of approaching stochastic settings is in considering *a priori* versions of these problems. The general idea is to find a *master tour*, that is, a tour for all possible customers that can easily be modified to serve the demand on a given day. The cost is then defined to be the expected cost with respect to the distribution of demands.

We consider an a priori variant of the capacitated vehicle routing problem (CVRP), where the demands are all independent and identically distributed Bernoulli experiments. In the deterministic variation of CVRP we are given a single depot, a set of customers, which all have a demand of one, and a vehicle with capacity k. The goal is to find a tour (consisting of a set of subtours) of minimum total cost that visits every customer and returns to the depot after serving at most k customers. In the *a priori* variant a set of *active* customers is drawn from some probability distribution. The goal is to find as master tour a feasible solution to the deterministic problem. The tour is then shortcut given the demand, which in this case is the active set of customers. The objective is to minimize the expected cost after shortcutting. This is the same model used for the a priori TSP in Shmoys and Talwar [16] and for the a priori travelling repairman problem in Van Ee and Sitters [17]. Note, however, that this definition of an a priori strategy is slightly different from the a priori strategies, which have typically been considered for the CVRP (for example [3, 6, 9]). Here they follow a TSP tour of all customers, and then go back to the depot whenever necessary. In contrast to this we require that the master tour itself is already a feasible solution to the deterministic CVRP (i.e., it includes the returns to the depot).

This means that for our approach we can have a fleet of vehicles, where each vehicle visits (a subset of) the same customers every time. This can be important for some applications, such as nurses visiting their patients.

Related work. The idea of using an a priori strategy was first introduced for the travelling salesman problem by Jaillet in his PhD thesis [11] (later published in an article [12]). Shmoys and Talwar [16] provide a randomized algorithm with an approximation factor of 4 for the a priori TSP and a deterministic algorithm with an approximation factor of 8. In their work they consider the independent activation model, which means that every customer is active with a probability that is independent from the probabilities of the other customers. These bounds have been improved to a 3.5 factor for the randomized algorithm by Van Ee and Sitters [17] and a deterministic $6.5-$approximation by Van Zuylen [18]. In an earlier result, Schalekamp and Shmoys [15] give an $O(\log(N))-$ approximation for arbitrary distributions. Van Ee and Sitters [17] prove a constant approximation ratio for the a priori travelling repairman problem under the assumption that every node is independently active with the same probability. This is the same assumption we impose on our problem.

The deterministic vehicle routing problem and its variations are well studied [1, 2, 7, 10]. Two variations of the problem are the split-delivery CVRP and the unsplit-delivery CVRP. In the former variation the vehicle can deliver the demands over multiple visits, whereas it has to deliver the full demand in one visit in the latter. For split-delivery CVRP Haimovich and Rinnooy Kan [10] use an algorithm called the cyclic heuristic, and Altinkemer and Gavish [2] use a slight variation of it. Both take as input a TSP solution with approximation factor α and achieve an approximation ratio of $1 + \alpha(1 - 1/k)$ in the case where all the demands are one. This is essentially equivalent to the split-delivery CVRP with arbitrary demands, as shown by Haimovich and Rinnooy Kan [10]. For the non-split delivery CVRP Altinkemer and Gavish [1] give an algorithm with an approximation ratio of $2+\alpha(1-2/k)$, which has subsequently been improved to $2 + \alpha(1 - 2/k) - 1/(3k^3)$ by Bompadre et al. [7].

The a priori CVRP with stochastic demands was first studied by Bertsimas [5]. The problem considered by Bertsimas (and others) differs from ours in that, as mentioned, they fix an order of visited customers, but do not have a master tour which is a feasible solution. Bertsimas introduces two stochastic variants for the split-delivery variation: In the first strategy every customer is visited, and the demand only becomes clear once the vehicle arrives at the customer. As such, every customer needs to be visited, even if it turns out that they have zero demand. The other strategy, which is the strategy we apply to our problem, is defined such that the demands are fixed as soon as the vehicle first leaves the depot, and customers with zero demand are skipped. For the first strategy Bertsimas obtains an approximation ratio of $1 + \alpha + O(1/N)$, when the demands of the customers are identically distributed. Berman and Das [4] generalize this result to arbitrarily distributed demands and give a randomized algorithm with an approximation ratio of $1 + \alpha$ and a deterministic algorithm with an approximation ratio of $1 + 2\alpha$. Gupta et al. [9] give a similar algorithm, which also achieves an approximation ratio of $1 + \alpha$ for the split-delivery variation and additionally provide a $(2+\alpha)-$approximation for the unsplit-delivery variation. They also show that the cyclic heuristic is a deterministic $(1 + 2\alpha)-$approximation algorithm.

A common approach to CVRP (used for example in the cyclic heuristic) is to modify a solution to the TSP, such that it becomes a valid CVRP solution. We will use the same idea for the a priori variation of the problem.

Our results. We give an approximation algorithm for the a priori CVRP with unit demands. As in Shmoys and Talwar [16] we consider the independent activation model. Furthermore, we restrict ourselves to the case where every customer is independently active with the same probability. We prove the following theorem:

Theorem 1. *Given a $\gamma-$approximation algorithm for the a priori TSP, there exists a $(1 + k/N + \gamma)-$approximation algorithm for the a priori vehicle routing problem with unit demands and vehicle capacity k, where all nodes are independently active with the same probability.*

The algorithm first computes a solution to the a priori TSP, then uses a similar idea as Altinkemer and Gavish [2] to transform this into a solution to the a priori CVRP, but now optimizes with respect to the expected cost over the distribution of active customers. To prove our constant performance guarantees we also bound the value of an optimal a priori CVRP solution by the value of an optimal a priori TSP solution. Furthermore, we use the fact that for every subtour in the optimal solution, the active customer that is furthest away from the depot has to be served, and thus the length of the subtour is at least twice the distance from the depot to this customer. We bound this by two times the average distance to the active customers, which we then rewrite to a form that enables us to prove a constant factor approximation.

2. Problem Description

In the capacitated vehicle routing problem with unit demands we are given a complete undirected graph $G = (V, E)$ with N vertices (customers) and edge weights

$d_{u,v}$ denoting the distance between vertex u and vertex v. Additionally, we have a source (depot) $s \notin V$, and for every vertex $v \in V$ the cost of travelling to and from the source is $d_v = d_{s,v} = d_{v,s}$. Additionally, we assume the distances form a metric on $V \cup \{s\}$. The goal is to find, for a given capacity $k \in \mathbb{N}$, a minimal cost tour that visits every vertex in V exactly once, and goes back to the source after visiting at most k vertices.

In the a priori capacitated vehicle routing problem with unit demands we consider a probability distribution Π over the powerset of V. The goal is to find a master tour (that is, a valid solution for the CVRP) minimizing the expected cost of the tour after shortcutting to an active set A drawn from Π. From here on, whenever we talk about the cost of an a priori CVRP solution, we mean the expected cost with respect to Π.

In this paper we consider the independent activation model, where the probability of a vertex being active is independent of the probabilities of the other vertices. Furthermore, we restrict ourselves to the case, where all vertices $v \in V$ are active with the same probability p.

3. A Constant Approximation Algorithm for the Uniform A Priori CVRP with Unit Demands

In this section we give a $(1+k/N+\gamma)-$approximation algorithm to the uniform a priori CVRP with unit demands, where γ is the approximation ratio of an algorithm to the a priori TSP. Our idea is based on the work by Altinkemer and Gavish [2] for the deterministic CVRP with unit demands. They consider k different ways of splitting the TSP solution into subtours of length at most k, then choose the one with minimum cost. Inspired by this, we propose a similar algorithm, which constructs N sets of subtours, and chooses the optimal one with respect to the expected cost.

More formally, we start from a TSP solution of V, that is, a TSP solution of all nodes excluding the source. We then create N sets of subtours $\mathcal{S}_0, \ldots, \mathcal{S}_{N-1}$ in the following way: We enumerate the vertices in V as $v_0, \ldots, v_{N-1}$ in the order of the TSP tour. Then, $\mathcal{S}_i$ will be the tour split into $\lceil \frac{N}{k} \rceil$ subtours, starting at vertex i, where at the beginning and end of each subtour we return to the source. That is, for every $\mathcal{S}_i$ we have $\lceil \frac{N}{k} \rceil - 1$ subtours of length k of the form $(s,\ v_{zk+i \mod N}, v_{zk+i+1 \mod N}\ \ldots,\ v_{(z+1)k+i-1 \mod N},\ s)$, for $0 \leq z < \lceil \frac{N}{k} \rceil - 1$. The last subtour is given by $(s,\ v_{(\lceil \frac{N}{k} \rceil -1)k+i \mod N},\ v_{(\lceil \frac{N}{k} \rceil -1)k+i+1 \mod N}\ \ldots,\ v_{i-1 \mod N},\ s)$. Out of $\mathcal{S}_0, \ldots, \mathcal{S}_{N-1}$, we choose the $\mathcal{S}_i$ with minimum expected cost. This is the master tour returned by the algorithm. Note that in the case of $N \bmod k = 0$, some of the sets $\mathcal{S}_i$ will be identical. This has no influence on the analysis.

Clearly, this algorithm returns a valid solution with respect to the problem description. We show how to compute the expected cost for a set of subtours $\mathcal{S}_i$ in polynomial time. The expected cost of this solution is the expected cost of the set of subtours when shortcut to an active set A. Denote the resulting set of edges by $\mathcal{S}_i(A)$. Then the cost of $\mathcal{S}_i$ is given by:

$$E_A\left[\sum_{u,v\in V\cup\{s\}} d_{u,v}\mathbb{1}\left((u,v)\in\mathcal{S}_i(A)\right)\right] \tag{1}$$

$$=\sum_{u,v\in V\cup\{s\}} d_{u,v}P\left((u,v)\in\mathcal{S}_i(A)\right). \tag{2}$$

Here, $\mathbb{1}(X)$ is the indicator function that returns 1 if X is true and 0 otherwise. If u and v are elements of different subtours, then probability in (2) is zero. Otherwise, if u is the depot, the probability is $p(1-p)^l$, where l is the number of elements before v in the subtour in $\mathcal{S}_i$. That is, the probability that v is active and no elements before v are active. Similarly, if v is the depot, the probability is $p(1-p)^l$, where l is the number of elements after u in the subtour in $\mathcal{S}_i$. If both u and v are in V and in the same subtour, then the probability of the edge (u,v) being included in the shortcut solution is $p^2(1-p)^l$, where l is the number of vertices between u and v in $\mathcal{S}_i$. Thus, we can quickly compute $P\left((u,v)\in\mathcal{S}_i(A)\right)$ for any pair (u,v) and plug it into (2) to compute the expected cost of $\mathcal{S}_i$.

We now prove that the algorithm gives an expected $(1+k/N+\gamma)-$approximation to the a priori CVRP. We first prove two lower bounds for the expected cost of the optimal solution, then prove an upper bound for the expected cost of the solution returned by the algorithm. Finally, we combine these to obtain the claimed bound.

In the following lemma we show how to bound the optimal solution for the a priori CVRP in terms of the value of the a priori TSP solution.

Lemma 3.1. *Let C_{APTSP} and OPT_{APTSP} be the costs of a $\gamma-$approximation algorithm and an optimal solution to the a priori TSP on the set of customers V, respectively. Let OPT be the cost of an optimal solution to the uniform a priori CVRP with unit demands. Then*

$$\gamma OPT \geq \gamma OPT_{APTSP} \geq C_{APTSP}\,.$$

Proof. By definition $\gamma OPT_{APTSP} \geq C_{APTSP}$. At the same time we have $OPT \geq OPT_{APTSP}$, since we can get a feasible solution to the a priori TSP by shortcutting any CVRP solution such that we do not visit the source. Combining these two facts gives the lemma. □

In the next lemma we prove a lower bound, which is based on the observation that in each subtour we have to visit the active vertex that is furthest away from the depot.

Lemma 3.2. $OPT \geq \frac{2\left(1-(1-p)^k\right)}{k} \sum_{v \in V} d_v$.

Proof. Denote by $\mathcal{S}_{OPT}$ the set of subtours in the optimal solution. Then, by the triangle inequality, the following holds:

$$\begin{aligned} OPT &\geq 2E_A\left[\sum_{\tau \in \mathcal{S}_{OPT}} \operatorname{avg}_{v \in \tau \cap A} d_v\right] \\ &= 2 \sum_{\tau \in \mathcal{S}_{OPT}} E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v\right] \\ &= 2 \sum_{\tau \in \mathcal{S}_{OPT}} E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v | \tau \cap A \neq \emptyset\right] P(\tau \cap A \neq \emptyset). \end{aligned}$$

To compute $E_A(\operatorname{avg}_{v \in \tau \cap A} d_v | \tau \cap A \neq \emptyset)$ partition the outcome space by the size of $\tau \cap A$. That is,

$$E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v | \tau \cap A \neq \emptyset\right] = \sum_{i=1}^{|\tau|} E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v || \tau \cap A| = i\right] P(|\tau \cap A| = i).$$

Now, for a fixed i, each set A such that $|\tau \cap A| = i$ is equally likely. As such

$$E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v || \tau \cap A| = i\right] = \operatorname{avg}_{v \in \tau} d_v,$$

for all i. We get

$$\begin{aligned} OPT &\geq 2 \sum_{\tau \in \mathcal{S}_{OPT}} E_A\left[\operatorname{avg}_{v \in \tau \cap A} d_v | \tau \cap A \neq \emptyset\right] P(\tau \cap A \neq \emptyset) \\ &= 2 \sum_{\tau \in \mathcal{S}_{OPT}} (\operatorname{avg}_{v \in \tau} d_v)(1 - (1-p)^{|\tau|}) \\ &= 2 \sum_{\tau \in \mathcal{S}_{OPT}} \left(\frac{1}{|\tau|} \sum_{v \in \tau} d_v\right)(1 - (1-p)^{|\tau|}). \end{aligned}$$

The last expression decreases as a function of $|\tau|$ (see Appendix A).

Since we require the optimal solution to be feasible for the deterministic problem, each subtour τ can have at most k customers. Therefore,

$$\begin{aligned} OPT &\geq 2 \sum_{\tau \in \mathcal{S}_{OPT}} \sum_{v \in \tau} d_v \frac{1-(1-p)^{|\tau|}}{|\tau|} \\ &\geq 2 \sum_{\tau \in \mathcal{S}_{OPT}} \sum_{v \in \tau} d_v \frac{1-(1-p)^{k}}{k} \\ &= \frac{2\left(1-(1-p)^k\right)}{k} \sum_{v \in V} d_v \ . \end{aligned}$$

□

The following lemma gives an upper bound for the expected cost of the algorithm. The proof is similar to the one by Altinkemer and Gavish [2], however we now consider the expected costs of the sets of subtours $\mathcal{S}_0, ..., \mathcal{S}_{N-1}$, where the expectation is taken over Π.

Lemma 3.3. *Denote by C_{APTSP} the expected cost of the a priori TSP solution used in the algorithm. Then*

$$E_A\left[ALG(A)\right] \leq 2\left(\frac{1}{k}+\frac{1}{N}\right)\left(1-(1-p)^k\right)\sum_{v \in V} d_v + C_{APTSP}.$$

Proof. We can decompose the cost of every subtour into the cost of the edge that leaves the source (outgoing edge), the cost of the edge that returns to the source (incoming edge) and the rest of the subtour.

To determine the total cost of the algorithm we consider the expected contribution of each vertex separately, and use the sum of the expected costs of all possible sets of subtours $\mathcal{S}_0, \ldots, \mathcal{S}_{N-1}$ to bound the cost of the best set of subtours.

Each vertex will be the j'th vertex in a subtour in at most $\lceil \frac{N}{k} \rceil$ of the subtours $\mathcal{S}_0, \ldots, \mathcal{S}_{N-1}$. Now consider the following: If a vertex v is the j'th first vertex in a subtour, the chance of this vertex being the first *active* vertex in the subtour is $(1-p)^{j-1}p$. Summing up over all the sets of subtours it thus contributes at most

$$\left\lceil \frac{N}{k} \right\rceil \sum_{j=0}^{k-1} (1-p)^j p d_v \leq \left(\frac{N}{k}+1\right) \sum_{j=0}^{k-1} (1-p)^j p d_v$$

to the total cost of the outgoing edges. By the same argument, each vertex v contributes no more than $\left(\frac{N}{k}+1\right)\sum_{j=0}^{k-1}(1-p)^j p d_v$ to the total cost of the incoming edges. For every set of subtours, shortcut to A, the cost of the edges which are neither outgoing nor incoming can be bounded by the a priori TSP solution. For all sets, we thus bound the rest of the cost by $N \cdot C_{APTSP}$.

Hence, we get:

$$\sum_{i=0}^{N-1} E_A\left[\text{cost}(\mathcal{S}_i(A))\right] \leq 2\sum_{v\in V}\left(\frac{N}{k}+1\right)\sum_{j=0}^{k-1}(1-p)^j p d_v + N\cdot C_{APTSP} \tag{3}$$

$$= 2\left(\frac{N}{k}+1\right)\sum_{v\in V}\frac{1-(1-p)^k}{1-(1-p)} p d_v + N\cdot C_{APTSP} \tag{4}$$

$$= 2\left(\frac{N}{k}+1\right)\left(1-(1-p)^k\right)\sum_{v\in V} d_v + N\cdot C_{APTSP}. \tag{5}$$

Here (4) follows from (3) by using the closed form of the geometric series.

We choose the best set of subtours, so N times the cost of this solution must be smaller than sum of the costs of the N sets of subtours. Thus:

$$N\cdot E_A\left[ALG(A)\right] \leq \sum_{i=0}^{N-1} E_A\left[\text{cost}(\mathcal{S}_i(A))\right]$$

$$\leq 2\left(\frac{N}{k}+1\right)\left(1-(1-p)^k\right)\sum_{v\in V} d_v + N\cdot C_{APTSP}.$$

Hence:

$$E_A\left[ALG(A)\right] \leq 2\left(\frac{1}{k}+\frac{1}{N}\right)\left(1-(1-p)^k\right)\sum_{v\in V} d_v + C_{APTSP}.$$

□

We are now ready to prove the main result of this paper:

Theorem 1. *Given a $\gamma-$approximation algorithm for the a priori TSP, there exists a $(1+k/N+\gamma)-$approximation algorithm for the a priori vehicle routing problem with unit demands and vehicle capacity k, where all nodes are independently active with the same probability.*

Proof. We bound the ratio $\frac{E_A[ALG(A)]}{OPT}$ using Lemma 3.1, 3.2 and 3.3:

$$\frac{E_A\left[ALG(A)\right]}{OPT} \leq \frac{2\left(\frac{1}{k}+\frac{1}{N}\right)\left(1-(1-p)^k\right)\sum_{v\in V} d_v}{\frac{2\left(1-(1-p)^k\right)}{k}\sum_{v\in V} d_v} + \frac{C_{APTSP}}{OPT} \leq 1 + k/N + \gamma.$$

□

We can use the expected 3.5−approximation algorithm for the a priori TSP from Van Ee and Sitters [17] and the deterministic 6.5−approximation algorithm from Van Zuylen [18] to get the following corollary:

Corollary 3.4. *There exists a randomized* 5.5− *approximation algorithm and a deterministic* 8.5− *approximation algorithm for the a priori capacitated vehicle routing problem with unit demands, when nodes are independently active with the same probability.*

Proof. This follows immediately from Theorem 1 by plugging in the randomized 3.5−approximation algorithm from Van Ee and Sitters [17] and the deterministic 6.5−approximation algorithm from Van Zuylen [18], respectively. □

4. Conclusion

We have provided an algorithm, which, given a γ−approximation algorithm for a priori TSP, gives an approximation factor of $1 + k/N + \gamma$ for the a priori capacitated vehicle routing problem in the independent activation model with the additional constraints that the customers have unit demands and are active with the same probability. We can use the randomized algorithm for the a priori TSP from Van Ee and Sitters [17] to get an expected 5.5−approximation or we can use the determistic algorithm from Van Zuylen [18] to get a deterministic 8.5-approximation. It remains an open question whether the approach in this paper can be used to get a constant factor approximation algorithm for the general independent activation model.

Acknowledgements. We thank Inge Li Gørtz for calling this problem to our attention and for helpful discussions throughout the progress, as well as the anonymous reviewers for constructive feedback.

References

[1] K. Altinkemer and B. Gavish. Heuristics for unequal weight delivery problems with a fixed error guarantee. *Oper. Res. Lett*, 6(4):149 – 158, 1987.

[2] K. Altinkemer and B. Gavish. Heuristics for delivery problems with constant error guarantees. *Transp. Sci.*, 24(4):294–297, 1990.

[3] E. Ando, B. Bhattacharya, Y. Hu, T. Kameda, and Q. Shi. Selecting good a priori sequences for vehicle routing problem with stochastic demand. In *Proc. 8th ICTAC*, pages 45–61, 2011.

[4] P. Berman and S. K. Das. On the vehicle routing problem. In *Proc. 9th WADS*, pages 360–371, 2005.

[5] D. Bertsimas. *Probabilistic combinatorial optimization problems.* PhD thesis, Massachusetts Institute of Technology, 1988.

[6] D. J. Bertsimas. A vehicle routing problem with stochastic demand. *Oper. Res.*, 40(3):574–585, 1992.

[7] A. Bompadre, M. Dror, and J. B. Orlin. Improved bounds for vehicle routing solutions. *Discrete Optim.*, 3(4):299 – 316, 2006.

[8] M. Gendreau, G. Laporte, and R. Seguin. Stochastic vehicle routing. *Eur. J. Oper. Res*, 88(1):3–12, 1996.

[9] A. Gupta, V. Nagarajan, and R. Ravi. Approximation algorithms for VRP with stochastic demands. *Oper. Res.*, 60(1):123–127, 2012.

[10] M. Haimovich and A. H. G. Rinnooy Kan. Bounds and heuristics for capacitated routing problems. *Math. Oper. Res.*, 10(4):527 – 542, 1985.

[11] P. Jaillet. *Probabilistic traveling salesman problems.* PhD thesis, Massachusetts Institute of Technology, 1985.

[12] P. Jaillet. A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Oper. Res.*, 36(6):929–936, 1988.

[13] G. Laporte. The vehicle routing problem: An overview of exact and approximate algorithms. *Eur. J. Oper. Res*, 59(3):345–358, 1992.

[14] V. Pillac, M. Gendreau, C. Guéret, and A. L. Medaglia. A review of dynamic vehicle routing problems. *Eur. J. Oper. Res*, 225(1):1 – 11, 2013.

[15] F. Schalekamp and D. B. Shmoys. Algorithms for the universal and a priori TSP. *Oper. Res. Lett*, 36(1):1–3, 2008.

[16] D. Shmoys and K. Talwar. A constant approximation algorithm for the a priori traveling salesman problem. In *Proc. 13th IPCO*, pages 331–343, 2008.

[17] M. van Ee and R. Sitters. The a priori traveling repairman problem. *Algorithmica*, 80(10):2818–2833, 2018.

[18] A. Van Zuylen. Deterministic sampling algorithms for network design. *Algorithmica*, 60(1):110–151, 2011.

Appendix A. Proof that $\frac{1-(1-p)^x}{x}$ is a decreasing function of x

To prove this we differentiate the function wrt. x and show that the result is negative:

$$\frac{d}{dx}\left(\frac{1-(1-p)^x}{x}\right) = -\frac{(1-p)^x\ln(1-p)}{x} - \frac{1-(1-p)^x}{x^2} = -\frac{(1-p)^x(\ln(1-p)x-1)+1}{x^2}$$

We show that the numerator is always positive:

$$\begin{aligned} &(1-p)^x(\ln(1-p)x-1)+1 \geq 0 \\ \Leftrightarrow &\frac{1}{(1-p)^x} \geq -\ln(1-p)x-1 \\ \Leftrightarrow &1+\frac{1}{(1-p)^x} \geq -\ln((1-p)^x) \end{aligned}$$

This is true since for all $z > 0$ it holds that $\frac{z+1}{z} \geq -\ln(z)$