

On the greedy walk problem

CHARLES BORDENAVE*, SERGEY FOSS† AND GÜNTER LAST‡

April 20, 2011

Abstract

This note introduces a greedy walk on Poisson and Binomial processes which is a close relative to the well known greedy server model. Some open problems are presented.

1 Introduction

Assume that customers drop on a circle at the times of a homogeneous Poisson process. Each customer chooses its position uniformly, independently of everything else. There is one server on the circle that moves with constant speed towards the nearest customer (ignoring new arrivals). Whenever a customer is encountered, it is removed from the system. It takes $c \geq 0$ units of time to serve a customer. This is the well known greedy server system introduced in [4]. A number of discrete-space variations of it have been studied in [9, 5, 6], and there are many open problems around – see, e.g. a paper [7] in this issue of QUESTA. Here is one of the key problems: it is conjectured that the system is stable if the intensity λ of N is strictly less than c^{-1} . In particular, if it takes no time to serve the customers, then it is believed that the system is stable whatever the value of λ is.

This note introduces the *greedy walk*, a version of the model for a possibly infinite state space and without service. We formulate a number of open problems and conjectures.

2 The greedy walk

Consider a homogeneous Poisson process (Poisson point field) N in d -dimensional Euclidean space \mathbb{R}^d . A “vacuum cleaner” (server) starts from the origin and removes points one by one. (There are no external arrivals.) The first point $X_1 \in N$ to be removed is the nearest neighbour (w.r.t. Euclidean distance) of the origin. The second point $X_2 \in N$ to be removed is the nearest neighbour of X_1 . The $(n + 1)$ -st point X_{n+1} to be removed is the nearest neighbour of X_n among all remaining points $N \setminus \{X_1, \dots, X_n\}$. Hence the

*Institut de Mathématiques de Toulouse, CNRS & Université de Toulouse, 31062 Toulouse, France

†School of Mathematics and Computer Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, Scotland, UK and Institute of Mathematics, Novosibirsk, Russia

‡Institut für Stochastik, Karlsruhe Institute of Technology, 76128 Karlsruhe, Germany

vacuum cleaner, following a greedy policy, always walks to the nearest remaining point of N . We call this process the *greedy walk on N* .

Let $N_g := \{X_1, X_2, \dots\}$ denote the point process of all removed points.

Problem 2.1. Is it true that $N_g \neq N$ a.s.?

A simple Borel-Cantelli argument shows that the answer is positive in one dimension. Eventually the server will move in one of the two possible directions. In two and higher dimensions the behaviour of the server is much less obvious. The authors' conjecture is that the answer to Problem 2.1 is positive in all dimensions $d \geq 3$.

It should be noted that N_g is not stationary. The point process $N_g + x$ has the same distribution as the process of removed points when the server does not start in the origin but in the point $x \in \mathbb{R}^d$.

Let $W_r := [0, r]^d$ denote a cube with side length $r \geq 0$ and $N_g(W_r)$ the number of removed points in W_r .

Problem 2.2. Assume that the answer to Problem 2.1 is positive. Is it true that the limit

$$\mu := \lim_{r \rightarrow \infty} r^{-d} \mathbb{E} N_g(W_r) \quad (2.1)$$

exists almost surely?

Because of the absence of stationarity a solution to Problem 2.2 is not obvious. In one dimension the answer is positive and μ is half the intensity of N . This motivates the next problem. Let $\lambda := \mathbb{E} N(W_1)$ denote the intensity of N .

Problem 2.3. Assume that the answers to both Problems 2.1 and 2.2 are positive. Is it then true that $\mu \in (0, \lambda)$?

The fraction μ/λ is a performance parameter of the system. It can be interpreted as the probability that a “randomly chosen” point of N will be removed. Even if the answers to Problems 2.1 and 2.2 are positive, it is still possible that $\mu = \lambda$.

We conjecture that $\mu > 0$ for $d = 2$ and that $\mu = 0$ in all higher dimensions.

Rahul Roy has mentioned that the following problem had been proposed by Harry Kesten.

Problem 2.4. Does the vacuum cleaner cross any fixed hyperplane infinitely often?

3 Non-homogeneous case

In this section we consider a Poisson process N on \mathbb{R}^d with a non-constant intensity function $\lambda : \mathbb{R}^d \rightarrow [0, \infty)$ of the form

$$\lambda(x) = (\max(|x|, 1))^{-\alpha},$$

where $\alpha \in \mathbb{R}$ is a given parameter. The case $\alpha = 0$ corresponds to the homogeneous Poisson process.

Consider the greedy walk algorithm from the previous section and denote again by N_g the point process of all removed points. For a cube W_r with side length r , let $N(W_r)$ be the number of points of process N located in W_r , and $N_g(W_r)$ the number of removed points in W_r . In this more general setting, we may formulate a number of new problems.

Problem 3.1. What is the range of parameters α for which

1. $N_g \neq N$ a.s.?
2. the limit $\lim_{r \rightarrow \infty} N_g(W_r)/N(W_r)$ exist ?
3. $\limsup_{r \rightarrow \infty} N_g(W_r)/N(W_r) < 1$?
4. $\liminf_{r \rightarrow \infty} N_g(W_r)/N(W_r) > 0$?
5. $\lim_{r \rightarrow \infty} N_g(W_r)/N(W_r) = 0$?

Clearly, $N = N_g$ is finite a.s. if $\alpha > 2$. We also believe that statement 5 of Problem 3.1 holds for any negative α (therefore, statements 1-3 of the problem hold too).

Problem 2.4 also raises a natural question in the non-homogeneous case.

Problem 3.2. For which parameters α , the vacuum cleaner crosses any fixed hyperplane infinitely often ?

4 The greedy walk on a finite number of points

We may change the perspective and consider a metric space \mathbb{X} with bounded diameter Δ . Let $N_n = \{Y_0, \dots, Y_n\}$ be a set of $n + 1$ points in \mathbb{X} . At time 0, a server starts at Y_0 and performs the greedy walk at unit speed until time L_n where it has removed all points in N_n . Obviously, L_n is the length of the path performed by the vacuum cleaner. Hence, if L_n^* is the length of the shortest (traveling salesperson) tour on N_n , we have

$$L_n \geq L_n^* - \Delta.$$

A reverse inequality is known: for some constant $c > 0$ and all $n \geq 1$,

$$L_n \leq (c \ln n) L_n^*$$

where c is a universal constant, see [8].

Without extra assumptions, this logarithmic factor is optimal up to a constant factor. Assume now that \mathbb{X} is a bounded Borel set of \mathbb{R}^d , $d \geq 2$, with positive Lebesgue measure. Then, the above upper bound could be significantly improved if $N_n = \{Y_0, \dots, Y_n\}$ is a random set of $n + 1$ independent points uniformly distributed on \mathbb{X} . Note that the distribution of N_n is that of a homogeneous Poisson process on \mathbb{X} conditioned to have $n + 1$ points. A celebrated theorem of Beardwood, Halton, and Hammersley [1] asserts that there exists a constant $\alpha(\mathbb{X}) > 0$ such that almost surely

$$\lim_{n \rightarrow \infty} n^{\frac{1}{d}-1} L_n^* = \alpha(\mathbb{X}),$$

and $\alpha(\mathbb{X}) = \alpha([0, 1]^d) \text{vol}(\mathbb{X})^{\frac{1}{d}}$. The intuition is that there are $n + 1$ edges on the tour and that each edge has length of order $n^{-1/d}$ (which is the order of magnitude of the distance of the closest points to a given point). We refer to the monographs [11, 12], for more background on the probabilistic theory of Euclidean combinatorial optimization.

For the greedy path, $n^{1-\frac{1}{d}}$ is also the correct order of magnitude. It was proved by Bentley and Saxe [2] and Steele [10] that for any point set N_n ,

$$L_n \leq Cn^{1-\frac{1}{d}},$$

where C is a constant depending on the diameter of \mathbb{X} . On his webpage, David Aldous raises the two-dimensional version (in terms of averages) of following natural question:

Problem 4.1. Is it true that the limit

$$\beta(\mathbb{X}) := \lim_{n \rightarrow \infty} n^{\frac{1}{d}-1} L_n$$

exists almost surely, as a number in $[\alpha(\mathbb{X}), \infty]$?

Numerical experiments support a positive answer to Problem 4.1 as well as the equation $\beta(\mathbb{X}) = \beta([0, 1]^d) \text{vol}(\mathbb{X})^{1-\frac{1}{d}}$.

The greedy walk on N_n is naturally related to the greedy walk on the homogeneous Poisson point process N of intensity $\lambda = \text{vol}(\mathbb{X})^{-1}$ introduced in section 2. Indeed, if x lies in the interior of \mathbb{X} , then $n^{\frac{1}{d}}(N_n - x)$ converges weakly to the point process N in the topology of vague convergence of locally finite measures. This poses the following problem:

Problem 4.2. Consider again the greedy walk on a homogeneous Poisson process N of unit intensity. Is it true that the limit

$$\beta_g := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_1^n |X_i - X_{i-1}|$$

exists a.s.?

If the answer to Problem 2.1 is negative and if the answer to Problem 4.1 is positive, one could expect the answer to Problem 4.2 is also positive and, moreover, the constants β_g and $\beta([0, 1]^d)$ should coincide.

5 A space-time model

This section introduces a greedy walk model based on a space-time Poisson process. It closely related to Problem 4 in [7], see also [3].

Let \mathbb{X} be a metric space with metric ρ and let Π be a Poisson point process on $\mathbb{X} \times [0, \infty)$. Assume that the intensity measure of Π is the product of a locally finite measure $\Lambda \neq 0$ on \mathbb{X} and Lebesgue measure on $[0, \infty)$. A point $(x, s) \in \Pi$ is interpreted as a customer arriving at time s at location x . Let

$$N^t := \Pi(\cdot \times [0, t])$$

be the point process of all customers arrived by time $t \geq 0$ and let $N := \Pi(\cdot \times [0, \infty))$. It is quite possible (for instance if \mathbb{X} is discrete) that N has multiple points, that is $N(\{x\}) \geq 2$ for some $x \in \mathbb{X}$. A vacuum cleaner is moving in \mathbb{X} at constant speed $v < \infty$ and removes

the points of N one by one. (In contrast with the model discussed in [7], there are no service times here, so removals are instantaneous). After each removal the vacuum cleaner picks the point nearest to its current position. If there is no point available, the vacuum cleaner waits for the next arrival.

To give a more accurate description of the process, assume first that Λ is finite. At time 0 the server is positioned at a fixed location $x_0 \in \mathbb{X}$. Let $(X_1, \tau_1) \in \Pi$ be the point of Π with the smallest arrival time. Then X_1 is the first point to be served. It is removed from the system at time $T_1 := \tau_1 + v^{-1}\rho(x_0, X_1)$. Let τ_2 be the smallest $t \geq T_1$ such that $N^t - \delta_{X_1} \neq 0$. (It is possible that $\tau_2 = T_1$.) The next point X_2 to be removed from the system at time $T_2 := \tau_2 + v^{-1}\rho(X_1, X_2)$, is the nearest neighbour of X_1 in the support of $N^{\tau_2} - \delta_{X_1}$. If there is more than one nearest neighbour, one of them is chosen at random. It is then clear how to proceed, in order to obtain the whole sequence (X_n, T_n) , $n \geq 1$, of removed points and associated removal times. If Λ has atoms, it is possible that $(X_n, T_n) = (X_{n+1}, T_{n+1})$ for some $n \geq 1$. If Λ is infinite, there is no smallest arrival time. Then let (X_1, τ_1) be the point $(x, s) \in \Pi$ minimizing $\rho(x_0, x) + s$. The sequence $((X_n, T_n))$ can then be defined as before. (After the removal of a point there will always be a point of N available.)

Let $\Pi_g := \sum_{n \geq 1} \delta_{(X_n, T_n)}$ and $N_g^t := \Pi_g(\cdot \times [0, t])$. The next problem is the greedy server problem in the absence of service times, formulated for a general phase space. Let $M := \{T_n : N^{T_n} = N_g^{T_n}\}$ denote the point process of all times of removals leaving the system empty.

Problem 5.1. Assume that $\Lambda(\mathbb{X}) < \infty$. Is it true that

$$\liminf_{t \rightarrow \infty} t^{-1} M([0, t]) > 0$$

holds almost surely?

For a finite set \mathbb{X} the problem was solved in [5]. If $\Lambda(\mathbb{X}) = \infty$ (and $\mathbb{X} = \mathbb{R}^d$) the behaviour of the server might resemble a random walk. This suggests the next problem.

Problem 5.2. Assume that $\Lambda(\mathbb{X}) = \infty$. Does the server eventually leave any bounded set?

Both problems do relate to Problem 4 in [7].

Remark 5.3. It is easy to extend the model of this section so as to cover also the static model of Section 2. To do so, the intensity measure of Π can be taken as the sum of $\Lambda_0 \otimes \delta_0$ and $\Lambda \otimes L_1$, where Λ_0 is a locally finite measure and L_1 is Lebesgue measure on $(0, \infty)$, and where Λ might be the zero measure. Then Λ_0 is the intensity measure of a Poisson process of customers present in the system at time 0. If Λ equals zero, then speed v does not play a role.

References

- [1] BEARDWOOD, J. HALTON, J. H. AND HAMMERSLEY, J. M. (1959). The shortest path through many points. *Proc. Cambridge Philos. Soc.*, **55**, 299–327.

- [2] BENTLEY, J.L. AND SAXE, J.B. (1980). An analysis of two heuristics for the Euclidean traveling salesman problem. *Proc. 18th Allerton conference on Communication, Control and Computing* **6**, 41–49. University of Illinois, Urbana-Champaign
- [3] Bertsimas, D.J. and Garrett van Ryzin, G. van. (1991). A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane. *Operations Research* **39**, 601–615.
- [4] COFFMAN JR., E.G. AND GILBERT, E.N. (1987). Polling and greedy servers on a line. *Queueing Systems* **2**, 115–145.
- [5] FOSS, S. AND LAST, G. (1996). Stability of polling systems with exhaustive service policies and state dependent routing. *Annals of Applied Probability* **6**, 116–137.
- [6] FOSS, S. AND LAST, G. (1998). On the stability of greedy polling systems with general service policies. *Probab. Eng. Inform. Sci.* **12**, 49–68.
- [7] ROJAS-NANDAYAPA, L., FOSS, S., AND KROESE, D. (2011). Stability and Performance of Greedy Server Systems: A Review and Open Problems. *Queueing Systems*, this volume.
- [8] ROSENKRANTZ, D. STEARNS, R. AND LEWIS, P. (1977). An Analysis of Several Heuristics for the Traveling Salesman Problem. *SIAM Journal on Computing* **6(5)**, 563–581.
- [9] SCHAßBERGER, R. (1995). Stability of polling networks with state-dependent server routing. *Probab. Eng. Inform. Sci.*, **9**, 539–550.
- [10] STEELE, J. M. (1989). Cost of sequential connection for points in space. *Oper. Res. Lett.* **8**, no. 3, 137–142.
- [11] STEELE, J. M. (1997). Probability theory and combinatorial optimization. *CBMS-NSF Regional Conference Series in Applied Mathematics* **69**. SIAM, Philadelphia.
- [12] YUKICH, J. (1998). Probability theory of classical Euclidean optimization problems. *Lecture Notes in Mathematics* **1675**. Springer-Verlag, Berlin.