Helsinki University of Technology Institute of Mathematics Research Reports
Espoo 2009 A579

STABILITY OF A SPATIAL POLLING SYSTEM WITH GREEDY MYOPIC SERVICE

Lasse Leskelä **Falk Unger**

TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY
TECHNISCHE UNIVERSITÄT HELSINKI
UNIVERSITE DE TECHNOLOGIE D'HELSINKI

Helsinki University of Technology Institute of Mathematics Research Reports
Espoo 2009 A579

STABILITY OF A SPATIAL POLLING SYSTEM WITH GREEDY MYOPIC SERVICE

Lasse Leskelä Falk Unger

Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Mathematics and Systems Analysis

Lasse Leskelä, Falk Unger: *Stability of a spatial polling system with greedy myopic service*; Helsinki University of Technology Institute of Mathematics Research Reports A579 (2009).

Abstract: *This paper studies a spatial queueing system on a circle, polled at random locations by a myopic server that can only observe customers in a bounded neighborhood. The server operates according to a greedy policy, always serving the nearest customer in its neighborhood, and leaving the system unchanged at polling instants where the neighborhood is empty. This system is modeled as a measure-valued random process, which is shown to be positive recurrent under a natural stability condition that does not depend on the server's scan radius. When the interpolling times are light-tailed, the stable system is shown to be geometrically ergodic. We also briefly discuss how the stationary mean number of customers behaves in light and heavy traffic.*

AMS subject classifications: 60K25

Keywords: spatial queueing system, dynamic traveling repairman, quadratic Lyapunov functional, spatial-temporal point process, spatial birth-and-death process

Correspondence

Lasse Leskelä
Helsinki University of Technology
PO Box 1100
02015 TKK
Finland
lasse.leskela@iki.fi
http://www.iki.fi/lsl/

Falk Unger
UC Berkeley
665 Soda Hall
CA 94720
USA
funger@eecs.berkeley.edu

Received 2009-09-11

ISBN 978-952-248-077-4 (print) ISSN 0784-3143 (print)
ISBN 978-952-248-081-1 (PDF) ISSN 1797-5867 (PDF)

Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Mathematics and Systems Analysis
P.O. Box 1100, FI-02015 TKK, Finland
email: math@tkk.fi http://math.tkk.fi/

Stability of a spatial polling system with greedy myopic service

Lasse Leskelä* Falk Unger†

September 11, 2009

Abstract

This paper studies a spatial queueing system on a circle, polled at random locations by a myopic server that can only observe customers in a bounded neighborhood. The server operates according to a greedy policy, always serving the nearest customer in its neighborhood, and leaving the system unchanged at polling instants where the neighborhood is empty. This system is modeled as a measure-valued random process, which is shown to be positive recurrent under a natural stability condition that does not depend on the server's scan radius. When the interpolling times are light-tailed, the stable system is shown to be geometrically ergodic. We also briefly discuss how the stationary mean number of customers behaves in light and heavy traffic.

Keywords: spatial queueing system, dynamic traveling repairman, quadratic Lyapunov functional, spatial–temporal point process, spatial birth-and-death process

AMS 2000 Subject Classification: 60K25

1 Introduction

This paper studies a spatial queueing system where customers arrive to random locations in space that are a priori unknown to the server. The server operates sequentially in time by scanning a bounded neighborhood of

*Postal address: Helsinki University of Technology, PO Box 1100, 02015 TKK, Finland. URL: www.iki.fi/lsl/. Email: lasse.leskela@iki.fi

†Postal address: UC Berkeley, 665 Soda Hall, CA 94720, USA. Email: funger@eecs.berkeley.edu

a randomly chosen location, and serving the nearest customer it observes. After a service completion or an unsuccessful scan, a new scan is performed. This kind of models may be encountered in applications, such as disk storage or wireless sensor networks, where the time required for the server to find a customer has a major impact on the customer's sojourn time in the system.

To simplify the analysis, we assume that the service times are negligible compared to the scanning times. Observe that in the opposite case where the service times dominate, the number of customers in the system behaves like the standard single-server queue. A further simplification is to assume that the customers arriving during a scanning period are ignored from the ongoing scan. We expect that this assumption is valid when the scan radius is small. The resulting system can be rephrased as follows: A server polls the system sequentially at random locations, and either serves the nearest observed customer immediately, or leaves the system unchanged if there are no customers within the scan radius. Instead of a single polling server, we may alternatively think that there is an infinite stream of servers, each of which can only perform one scan during its operational lifetime.

The service policy is greedy in the sense that the server always aims to serve the nearest customer. We call this model *greedy polling server*, to distinguish it from spatial queueing systems where the server travels in space towards a nearest customer (*greedy traveling server*). Although a natural choice for many applications, the greedy policy is hard to analyze mathematically, as confirmed by earlier studies on traveling servers [1, 2, 4, 8, 12, 13]. Coffman and Gilbert [4] conjectured that the greedy traveling server on a circle is stable when the traffic intensity is less than one, regardless of the server's traveling speed. Kroese and Schmidt [12] proved this statement for several nongreedy policies, Altman and Levy [1] proved it for a so-called gated-greedy policy, and Foss and Last [8] proved it for greedy traveling servers on a finite space, see also Meester and Quant [16]. Despite these affirmative results for closely related systems, the stability of the greedy traveling server on a circle still remains an open problem. Analytical results on the distribution of customers related to greedy traveling servers are scarce. Coffman and Gilbert [4] showed that the greedy traveling server on a circle closely resembles a cyclic policy in heavy traffic (see also Litvak and Adan [15] for a similar result). Bertsimas and van Ryzin [2] found two lower bounds for the mean sojourn time in the system, which are valid for all travel policies. Kroese and Schmidt [13] derived second-order approximations for the number of customers and workload in light traffic.

The analytical challenges related to greedy traveling servers suggest that greedy polling servers may be difficult to analyze as well. This is why we set a modest research goal for this paper, namely, to characterize the stability of the system. This problem may be viewed as determining the system's throughput capacity, defined as the maximal sustainable arrival rate for which the number of customers in the system remains stable (stochastically bounded). Foss [7] recently presented an open problem, attributing it to V. Anantharam, conjecturing that the greedy polling server on a circle is stable if the arrival rate is less than the polling rate, regardless of the server's scan radius.

In this paper we prove Anantharam's conjecture [7, Section 3.2] by showing that the greedy polling server on a circle is stable if and only if the polling rate exceeds the arrival rate. The proof is based on presenting the system as a measure-valued Markov process, and developing a novel quadratic Lyapunov functional on the space of finite counting measures for which the measure-valued process has negative mean drift for large customer configurations. Besides incrementing the collection of known provable facts on spatial queueing systems with greedy service, our analytical results may be interesting in other application areas. For instance, the greedy polling server can be viewed as a spatial birth-and-death process. Spatial birth-and-death processes have usually been studied in the case where all individuals have a constant death rate, see for example Ferrari, Fernández, and Garcia [6]; and Garcia and Kurtz [9], who give sufficient conditions for stability in terms of the birth rates. The greedy polling server differs from the above birth-and-death processes in that the death rates of individuals are governed by the Voronoi tessellation generated by the customer locations. Borovkov and Odell [3] have recently studied a class of spatial–temporal point processes based on Voronoi cells, where the number of individuals is assumed constant over time. We expect that the quadratic Lyapunov functional presented in this paper may turn out useful in studying the ergodicity of more general spatial birth-and-death processes.

During the final writing stage of this article, we came across an interesting recent work of Robert [18], who considers the same problem from a different point of view. Using entirely different techniques (a stochastically monotone construction of a stationary solution), he proves a weaker form of stability, stating that the system has a limiting distribution for which the number of customers is finite almost surely. Our results based on Foster–Lyapunov drift criteria allow to prove stronger forms of stability, such as positive Harris recurrence and geometric ergodicity, depending on the tail behavior of the interpolling time distribution.

The rest of the paper is organized as follows. In Section 2 we describe the system as a measure-valued Markov process and derive formulas for its transition operators. Section 3 shows the positive recurrence of the system, and Section 4 is devoted to geometric ergodicity. In Section 5 we illustrate the system dynamics with numerical simulations complemented with heuristical ideas on the behavior of the system in light traffic and heavy traffic. Section 6 concludes the paper.

2 System description

2.1 Notation

The server operates on the circle $S = \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 = \ell^2\}$ of circumference $\ell > 0$, where the distance $d(x, y)$ between points x and y is defined as the length of the shortest arc connecting them. The state space of the system is the set $M_+(S)$ of finite counting measures ζ on S, so that $\zeta(B)$ denotes the number of customers located at $B \subset S$. The elements of $M_+(S)$ will also be called configurations, and the zero counting measure is called the empty configuration. The total number of customers in configuration ζ is denoted by $||\zeta||_{\mathrm{v}} = \zeta(S)$; this quantity also equals the total variation of ζ. We equip the space $M_+(S)$ with the sigma-algebra generated by the maps $\zeta \mapsto \zeta(B)$, with B being a Borel set in S. For real functions f on S we denote

$$\int_S f(x)\,\zeta(dx) = \sum_{x\in\zeta} f(x)\,\zeta(\{x\}),$$

where $x \in \zeta$ is shorthand for $\zeta(\{x\}) > 0$. For more background, see for instance Daley and Vere-Jones [5].

2.2 Population process in continuous time

Customers arrive to the circle S at uniformly distributed random locations. Assuming that the arrival locations and interarrival times are independent, the spatial–temporal arrival process can be described as a Poisson random measure on $\mathbb{R}_+ \times S$ with intensity measure $\lambda\, dt\, m(dx)$, where $\lambda > 0$ denotes the mean arrival rate, dt the Lebesgue measure on $\mathbb{R}_+$, and $m(dx)$ the uniform distribution (Haar measure) on S. The server polls the circle sequentially in time by scanning a neighborhood of radius $r > 0$ of a randomly chosen location, and either serving the nearest observed customer immediately, or leaving the system unchanged if the scanned neighborhood is empty. Hence the probability that a customer located at $x \in \zeta$ is served during a polling event is equal to

$$m(B_r \cap \Gamma_\zeta(x)),$$

where $B_r(x)$ denotes the open r-ball centered at x, and

$$\Gamma_\zeta(x) = \{y \in S : d(x,y) < d(x',y) \ \forall x' \in \zeta\}$$

denotes the Voronoi cell of a point x with respect to configuration ζ. (If there are many customers located at x, each of them is served with equal probability.) Assuming that the interpolling times are independent and exponential, say with parameter μ, and independent of the arrival process, the customer population in the system is described by a Markovian spatial birth-and-death process [9] with generator

$$\begin{aligned} Af(\zeta) = \lambda \int_S (f(\zeta + \delta_x) - f(\zeta))\, m(dx) \\ + \mu \int_S (f(\zeta - \delta_x) - f(\zeta))\, m(B_r(x) \cap \Gamma_\zeta(x))\, \zeta(dx), \end{aligned}$$

where δ_x denotes the Dirac measure at x.

2.3 Population process at polling instants

The stability regions of many ordinary queueing systems are insensitive to the shape of the service time distribution. We will show in Section 3 that analogously, the shape of the interpolling distribution does not affect the stability of our model. To show this claim, we will from now on assume that the interpolling times follow a general distribution G on $\mathbb{R}_+$. To keep the presentation concise, we will restrict the study into the discrete-time population process W obtained by sampling the system at polling instants, so that $W_t(B)$ denotes the number of customers in $B \subset S$ just after the t-th polling instant, $t \in \mathbb{Z}_+$ (we assume that also the initial state W_0 is observed just after a polling instant). We expect that the stability results for the discrete-time process extend to continuous time following similar techniques as in Kroese and Schmidt [12].

The population process W is a discrete-time Markov process on $M_+(S)$. Given an initial state $\zeta \in M_+(S)$, we denote the one-step transition operator of W by

$$Af(\zeta) = \mathrm{E}_\zeta\, f(W_1) = \mathrm{E}\,(f(W_1) \,|\, W_0 = \zeta)\,,$$

where f is a bounded or positive measurable function on $M_+(S)$. The associated probability kernel will be denoted by

$$P(\zeta, B) = A1_B(\zeta) = \mathrm{P}_\zeta(W_1 \in B), \quad B \subset M_+(S).$$

Because the polling locations are independent of the arrivals, we can decompose the transition operator according to

$$A = A_{\mathrm{a}} \circ A_{\mathrm{p}},$$

where the arrival operator A_{a} and the polling operator A_{p} are defined as follows. The operator A_{a} acts on bounded or positive measurable functions by

$$A_{\mathrm{a}} f(\zeta) = \sum_{n \geq 0} A_0^n f(\zeta)\, G_\lambda(n), \tag{2.1}$$

where

$$A_0 f(\zeta) = \int_S f(\zeta + \delta_x)\, m(dx)$$

corresponds to adding one new customer to a uniform random location, and

$$G_\lambda(n) = \int_{\mathbb{R}_+} e^{-\lambda s} \frac{(\lambda s)^n}{n!}\, G(ds)$$

is the probability that n customers arrive during an interpolling time. The operator A_{p} is defined by

$$A_{\mathrm{p}} f(\zeta) = f(x)(1 - k_r(\zeta)) + \sum_{x \in \zeta} f(\zeta - \delta_x)\, m(B_r(x) \cap \Gamma_\zeta(x)), \tag{2.2}$$

where

$$k_r(\zeta) = m(\cup_{x \in \zeta} B_r(x)) \tag{2.3}$$

is the probability that the server finds a customer during a scan targeted into configuration ζ.

2.4 Irreducibility and aperiodicity

The following result shows that the Markov process W describing the customer population in the system may become empty with probabilities uniformly bounded away from zero. As a consequence, the process satisfies the irreducibility and aperiodicity properties summarized below (see Meyn and Tweedie [17] for details).

Lemma 2.1. *The Markov process W is ϕ-irreducible and strongly aperiodic, where ϕ is the Dirac measure on $M_+(S)$ assigning unit mass to the zero counting measure on S. Moreover, the level sets of the form $C_n = \{\zeta \in M_+(S) : ||\zeta||_{\mathrm{v}} \le n\}$, $n \ge 0$, are small for W.*

Proof. Consider an initial configuration ζ with $||\zeta||_{\mathrm{v}} = k \le n$ customers. Then the probability that the system is empty after n polling instants is greater than the probability that no customers arrive during the first n interpolling times and the server finds a customer at each of the k first polling events. Because the probability of finding a customer in a nonempty system is at least $m(B_r)$, it follows that

$$P^n(\zeta, \{\zeta_0\}) \ge \epsilon^n m(B_r)^k \ge \epsilon^n m(B_r)^n,$$

where ζ_0 denotes the empty configuration and $\epsilon = \int_{\mathbb{R}_+} e^{-\lambda s} G(ds)$ is the probability that no customers arrive during an interpolling time. Denoting $\mu = \epsilon^n m(B_r)^n \phi$, we may thus conclude that

$$\inf_{\zeta \in C_n} P^n(\zeta, B) \ge \mu(B)$$

for all measurable $B \subset M_+(S)$. Thus the level set C_n is small [17, Section 5.2]. Further, the inequality $P(\zeta_0, \{\zeta_0\}) \ge \epsilon$ implies that W is strongly aperiodic [17, Section 5.4]. □

3 Positive recurrence

This section is devoted to deriving the main stability results (Theorems 3.8 and 3.11), which together imply that the system is positive recurrent if and only if the arrival rate of customers is strictly less than the polling rate, regardless of the server's scan radius. We start in Section 3.1 by discussing why it is not sufficient to analyze the mean drift of the population size, and introduce in Section 3.2 a quadratic functional for which the mean drift analysis works. Section 3.3 discusses a key interpolation inequality that is applied in Section 3.4 to show that the mean drift with respect to the quadratic functional is negative for large configurations. Section 3.5 summarizes the behavior of the unstable system.

3.1 Mean drift with respect to population size

A common method to prove the stochastic stability of a queueing system is show that the mean drift of the system with respect to the number of customers is strictly negative for large configurations. To see why this approach is not sufficient for the greedy polling system in this paper, denote the number of customers in configuration ζ by $h(\zeta) = ||\zeta||_{\mathrm{v}}$, and recall that the mean drift of the system with respect to h is defined by

$$Dh(\zeta) = Ah(\zeta) - h(\zeta),$$

where A is the one-step transition operator of the system. Recall the decomposition of $A = A_{\mathrm{a}} \circ A_{\mathrm{p}}$ in Section 2.3, and observe that

$$A_{\mathrm{a}}h(\zeta) = h(\zeta) + \lambda s_1,$$

where $s_1 = \int s\, G(ds)$ is the mean interpolling time, and

$$A_{\mathrm{p}}h(\zeta) = h(\zeta) - k_r(\zeta),$$

where k_r is the probability of a successful scan defined by (2.3). As a consequence,

$$Dh(\zeta) = \lambda s_1 - A_{\mathrm{a}}k_r(\zeta).$$

Consider a configuration $\zeta = n\delta_x$, where n customers are located in a single point $x \in S$. Then $k_r(\zeta) = m(B_r)$, so by conditioning on whether customers arrive or not during a polling instant, we find that

$$\begin{aligned} A_{\mathrm{a}}k_r(\zeta) &\le k_r(\zeta)G_\lambda(0) + (1 - G_\lambda(0)) \\ &= 1 - G_\lambda(0)(1 - m(B_r)), \end{aligned}$$

where $G_\lambda(0) = \int e^{-\lambda s}\, G(ds)$. Hence

$$Dh(\zeta) \ge \lambda s_1 - 1 + G_\lambda(0)(1 - m(B_r)).$$

Because the right side above does not depend on n, we see that $Dh(\zeta)$ can be strictly positive for arbitrarily large configurations, if $\lambda s_1 > 1 - G_\lambda(0)(1 - m(B_r))$. Hence the mean drift with respect population size cannot be used to show that the system is stable whenever $\lambda s_1 < 1$.

3.2 Quadratic energy functional

Let $M(S)$ be the space of signed counting measures on S, that is, measures of the form $\zeta = \sum_{i=1}^n z_i \delta_{x_i}$, where $z_i \in \mathbb{Z}$ and $x_i \in S$. The space $M(S)$ is a normed vector space with the total variation norm $||\zeta||_{\mathrm{v}} = \sum_i |z_i|$, and the subspace of finite positive counting measures $M_+(S)$ is a convex cone in $M(S)$.

Given $0 < a \le \ell/2$, define for $\zeta, \eta \in M(S)$,

$$\langle \zeta, \eta \rangle_{\mathrm{a}} = \int_S \int_S (a - d(x,y))_+ \, \zeta(dx) \, \eta(dy),$$

and denote

$$||\zeta||_{\mathrm{a}} = \sqrt{\langle \zeta, \zeta \rangle_{\mathrm{a}}}.$$

Lemma 3.1. *The map $(\zeta, \eta) \mapsto \langle \zeta, \eta \rangle_{\mathrm{a}}$ is symmetric, bilinear, and positive semidefinite. The map $\zeta \mapsto ||\zeta||_{\mathrm{a}}$ is a seminorm on $M(S)$ satisfying*

$$||\zeta||_{\mathrm{a}} \le ||\zeta||_{\mathrm{v}} \tag{3.1}$$

for all $\zeta \in M(S)$.

Proof. Clearly, $\langle \zeta, \eta \rangle_{\mathrm{a}}$ is bilinear and symmetric, so we only need to show positive semidefiniteness. We shall prove this using a probabilistic argument, representing the bilinear functional as an expectation of a random quantity. Observe that $(a - d(x,y))_+ = m(B_{a/2}(x) \cap B_{a/2}(y))$, where m denotes the uniform probability distribution on S. Hence by changing the order of integration we see that

$$\begin{aligned}
\langle \zeta, \eta \rangle_{\mathrm{a}} &= \int_S \int_S m(B_{a/2}(x) \cap B_{a/2}(y)) \, \zeta(dx) \, \eta(dy) \\
&= \int_S \int_S \int_S 1(x \in B_{a/2}(z)) 1(y \in B_{a/2}(z)) \, m(dz) \, \zeta(dx) \, \eta(dy) \\
&= \int_S \int_S \int_S 1(z \in B_{a/2}(x)) 1(z \in B_{a/2}(y)) \, \zeta(dx) \, \eta(dy) \, m(dz).
\end{aligned}$$

Hence

$$\langle \zeta, \eta \rangle_{\mathrm{a}} = \mathrm{E} \, \zeta(B_{a/2}(U)) \eta(B_{a/2}(U)),$$

where is U is a random variable uniformly distributed on S. Especially,

$$\langle \zeta, \zeta \rangle_{\mathrm{a}} = \mathrm{E} \, \zeta(B_{a/2}(U))^2, \tag{3.2}$$

which shows that $\langle \zeta, \zeta \rangle_{\mathrm{a}} \ge 0$ for all $\zeta \in M(S)$. As a consequence, $\zeta \mapsto ||\zeta||_{\mathrm{a}}$ is a seminorm (see for example Rudin [19, 4.2]). The representation (3.2) further shows the validity of (3.1). ∎

Remark 3.2. When a is chosen so that ℓ/a is not an integer, one could perhaps strengthen the statement of Lemma 3.1 by showing that $\langle \zeta, \eta \rangle_{\mathrm{a}}$ is an inner product on $M(S)$. However, in the sequel we will only need the fact that $||\zeta||_{\mathrm{a}}$ is a seminorm.

Lemma 3.3. *For any integer n and any configuration $\zeta \in M_+(S)$, there exists a closed ball B in S with diameter n^{-1} such that*

$$\zeta(B) \geq n^{-1}||\zeta||_{\mathrm{v}}. \tag{3.3}$$

Proof. Given an integer n, cover the unit circle with closed balls $B_1, \dots, B_n$, each having diameter n^{-1}. Then

$$||\zeta||_{\mathrm{v}} \leq \sum_{i=1}^{n} \zeta(B_i) \leq n \max_i \zeta(B_i),$$

which shows that $\max_i \zeta(B_i) \geq n^{-1}||\zeta||_{\mathrm{v}}$. □

Lemma 3.4. *For any $a > 0$ for all $\zeta \in M_+(S)$,*

$$\frac{\sqrt{a/2}}{1 + 2/a}||\zeta||_{\mathrm{v}} \leq ||\zeta||_{\mathrm{a}} \leq \sqrt{a}||\zeta||_{\mathrm{v}}. \tag{3.4}$$

Proof. The upper bound in (3.4) follows directly by observing that $(a - d(x,y))_+ \leq a$ for all x and y. To prove the corresponding lower bound, let n be an integer such that $2/a \leq n \leq 2/a + 1$, and use Lemma 3.3 to choose a closed ball B with diameter n^{-1} such that (3.3) holds. Because $(a - d(x,y))_+ \geq a - n^{-1} \geq a/2$ for all $x, y \in B$, it follows that

$$||\zeta||_{\mathrm{a}}^2 \geq \int_B \int_B (a - d(x,y))_+ \, \zeta(dx) \, \zeta(dy) \geq \frac{a}{2} \zeta(B)^2.$$

Because $n \leq 2/a + 1$, the lower bound now follows using (3.3). □

3.3 Interpolation inequality

This section is devoted to proving a key inequality (Lemma 3.7), which is needed for analyzing the mean drift of the system with respect to the seminorm $||\zeta||_a$. For a point x on the circle S and a positive real number a, we denote by $x + a$ the point on the circle obtained by traveling distance a from x anticlockwise on the circle, and $x - a$ the corresponding point obtained by traveling in the clockwise direction. Moreover, we denote by $[u, v]$ the closed arc formed by drawing a line from u to v moving anticlockwise on the circle, and by (u, v) the interior of $[u, v]$.

Lemma 3.5. *Let $0 < a \leq \ell/2$. Then for all $x \in S$ and for all $\zeta \in M_+(S)$ having atoms at $x - a$, x, and $x + a$,*

$$\sum_{y \in \zeta} (a - d(x,y))_+ m(\Gamma_\zeta(y)) = a^2. \tag{3.5}$$

Remark 3.6. Equation (3.5) may interpreted in terms nearest-neighbor interpolation as follows. Given a bounded measurable function f on S, define its nearest-neighbor interpolant (with respect to configuration ζ) by

$$I_\zeta f(z) = \sum_{y\in\zeta} f(y)1(z \in \Gamma_\zeta(y)).$$

Then

$$\int_S I_\zeta f(z)\, m(dz) = \sum_{y\in\zeta} f(y)\, m(\Gamma_\zeta(y)).$$

Given ζ and x as in Lemma 3.5, define $f(z) = (a - d(x,z))_+$. Then the left side in (3.5) equals $\int_S I_\zeta f(z)\, m(dz)$, and a short calculation shows that $\int_S f(z)\, m(dz) = a^2$. Hence (3.5) can be rewritten as

$$\int_S I_\zeta f(z)\, m(dz) = \int_S f(z)\, m(dz),$$

stating that the interpolation error made in replacing f by $I_\zeta f$ is zero (see Figure 1).

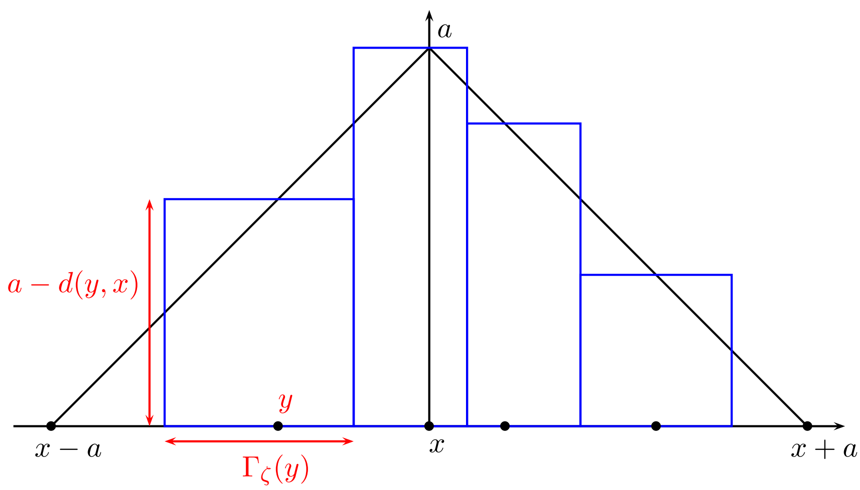

Figure 1: Nearest-neighbor interpolation of $z \mapsto (a - d(x,z))_+$.

Proof of Lemma 3.5. Assume first that x is the only atom of ζ in $(x-a, x+a)$. Then

$$m(\Gamma_\zeta(x)) = m([x - a/2, x + a/2]) = a,$$

which implies (3.5).

Assume next that (3.5) holds for some configuration ζ having atoms at $x-a$, x, and $x+a$. We shall show that the same is true also for $\zeta' = \zeta + \delta_z$, where $z \in S$ is arbitrary. We only need to consider the case where $z \in (x-a, x+a)$ such that $z \notin \zeta$, because otherwise the sum on the left side of (3.5) remains unaffected. Assume without loss of generality that $z \in (x-a, x)$, the other case being symmetric, and let $y_1, y_2 \in [x-a, x]$ be the points of ζ lying nearest to z in the anticlockwise and clockwise direction, respectively. Then

$$\begin{aligned}
m(\Gamma_{\zeta'}(z)) &= \frac{1}{2}(d(y_1, z) + d(z, y_2)), \\
m(\Gamma_{\zeta'}(y_1)) &= m(\Gamma_\zeta(y_1)) - \frac{1}{2}d(z, y_2), \\
m(\Gamma_{\zeta'}(y_2)) &= m(\Gamma_\zeta(y_2)) - \frac{1}{2}d(y_1, z).
\end{aligned}$$

Denote by $g(x, \zeta)$ the left side of (3.5). Because the other atoms of ζ remain unchanged, it follows that

$$\begin{aligned}
g(x, \zeta') - g(x, \zeta) = \frac{1}{2}&(a - d(z, x))(d(y_1, z) + d(z, y_2)) \\
&- \frac{1}{2}(a - d(y_1, x))d(z, y_2) - \frac{1}{2}(a - d(y_2, x))d(y_1, z).
\end{aligned}$$

Because $d(y_1, x) = d(y_1, z) + d(z, x)$ and $d(y_2, x) = d(z, x) - d(y_2, z)$, the terms on the right side of the above equation cancel out each other, so we may conclude that $g(x, \zeta') = g(x, \zeta)$. Hence (3.5) holds also for $\zeta' = \zeta + \delta_z$, and the proof is complete by induction. □

Lemma 3.7. *Assume that $0 < a \le \min(\ell/2, 2r)$. Then for all $\zeta \in M_+(S)$ and for all $x \in \zeta$,*

$$\sum_{y \in \zeta} (a - d(x, y))_+ m(B_r(y) \cap \Gamma_\zeta(y)) \ge a^2. \tag{3.6}$$

Proof. Denote the left side of (3.6) by $g(x, \zeta)$. Because $a \le \frac{\ell}{2}$, the sum on the left of (3.6) runs over the locations $y \in \zeta$ such that $y \in [x-a, x+a]$.

Let $\zeta' = \zeta + \delta_{x-a} + \delta_{x+a}$. We shall first show that

$$g(x, \zeta) \ge g(x, \zeta'). \tag{3.7}$$

Denote by y_l and y_r the atoms of ζ in $[x-a, x+a]$ located nearest to $x-a$ and $x+a$, respectively. Then (see Figure 2) $\Gamma_{\zeta'}(y) = \Gamma_\zeta(y)$ for all $y \in \zeta \cap (y_l, y_r)$, and moreover,

$$\Gamma_{\zeta'}(y_l) \subset \Gamma_\zeta(y_l) \quad \text{and} \quad \Gamma_{\zeta'}(y_r) \subset \Gamma_\zeta(y_r).$$

Because the points $y = x-a$ and $y = x+a$ contribute nothing to the sum on the left side of (3.6), we may conclude that (3.7) holds.

Further, because $a \le 2r$, it follows that $\Gamma_{\zeta'}(y) \cap B_r(y) = \Gamma_{\zeta'}(y)$ for all $y \in \zeta' \cap (x-a, x+a)$. Hence Lemma 3.5 shows that

$$g(x, \zeta') = \sum_{y \in \zeta'} (a - d(x,y))_+ m(\Gamma_\zeta(y)) = a^2.$$

In light of (3.7), this implies the claim. $\square$

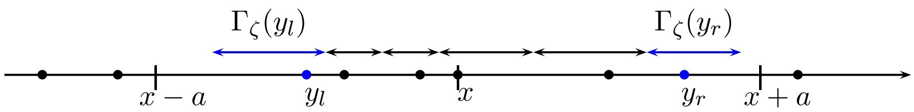

Figure 2: Voronoi cells of y_l and y_r.

3.4 Positive Harris recurrence

The following result shows that the system is stable under the natural condition $\lambda s_1 < 1$. Recall that s_1 and s_2 denote the first and the second moments of the interpolling time distribution, respectively. For a converse to Theorem 3.8, see Theorem 3.11 in Section 3.5.

Theorem 3.8. *Assume $\lambda s_1 < 1$ and $s_2 < \infty$. Then for all values of the scan radius $r > 0$ and the circle circumference $\ell > 0$, the population process W is positive Harris recurrent with stationary distribution π such that*

$$\int ||\zeta||_{\mathrm{v}} \, \pi(d\zeta) < \infty.$$

Moreover, for all initial states $\zeta \in M_+(S)$,

$$\sup_{g : |g| \le 1 + ||\cdot||_{\mathrm{v}}} \left| \mathrm{E}_\zeta \, g(W_t) - \int g \, d\pi \right| \to 0 \qquad \textit{as } t \to \infty.$$

Remark 3.9. Meyn and Tweedie call such a process f-ergodic with $f(\zeta) = 1 + ||\zeta||_{\mathrm{v}}$.

The proof of Theorem 3.8 is based on the following Foster–Lyapunov bound on the mean drift of the system with respect to the functional $h(\zeta) = \langle \zeta, \zeta \rangle_{\mathrm{a}}$.

Lemma 3.10. *Assume that $s_1, s_2 < \infty$, and let $0 < a \le \min(\ell/2, 2r)$. Then the mean drift of the system with respect to $h(\zeta) = \langle \zeta, \zeta \rangle_{\mathrm{a}}$ satisfies*

$$Dh(\zeta) \le -c_1 ||\zeta||_{\mathrm{v}} + c_2$$

for all $\zeta \in M_+(S)$, where

$$\begin{aligned} c_1 &= 2a^2(1 - \lambda s_1), \\ c_2 &= a(1 + \lambda s_1) + a^2(\lambda^2 s_2 - 2\lambda s_1). \end{aligned}$$

Proof. Recall that $Dh(\zeta) = Ah(\zeta) - h(\zeta)$, where $A = A_{\mathrm{a}} \circ A_{\mathrm{p}}$, and the transition operators A_{a} and A_{p} are defined by (2.1) and (2.2). Observe that the transition operator corresponding to n arrivals satisfies

$$\begin{aligned} A_0^n h(\zeta) &= \mathrm{E}\, h(\zeta + \sum_{i=1}^n \delta_{X_i}) \\ &= h(\zeta) + 2\,\mathrm{E}\langle \zeta, \sum_{i=1}^n \delta_{X_i} \rangle_{\mathrm{a}} + \mathrm{E}\langle \sum_{i=1}^n \delta_{X_i}, \sum_{j=1}^n \delta_{X_j} \rangle_{\mathrm{a}} \\ &= h(\zeta) + 2n\,\mathrm{E}\langle \zeta, \delta_{X_1} \rangle_{\mathrm{a}} + n\,\mathrm{E}\langle \delta_{X_1}, \delta_{X_1} \rangle_{\mathrm{a}} + (n^2 - n)\,\mathrm{E}\langle \delta_{X_1}, \delta_{X_2} \rangle_{\mathrm{a}}, \end{aligned}$$

where X_i are independent and uniformly distributed on S. Because the uniform distribution m on S is shift-invariant,

$$\int_S (a - d(x,y))_+ \, m(dy) = a^2$$

for all $x \in S$, so that

$$\mathrm{E}\langle \zeta, \delta_{X_1} \rangle_{\mathrm{a}} = \int_S \sum_{y \in \zeta} (a - d(x,y))_+ \zeta(\{y\}) \, m(dx) = a^2 ||\zeta||_{\mathrm{v}},$$

and

$$\mathrm{E}\langle \delta_{X_1}, \delta_{X_2} \rangle_{\mathrm{a}} = \int_S \int_S (a - d(x,y))_+ \, m(dx)\, m(dy) = a^2.$$

Hence

$$A_0^n h(\zeta) = h(\zeta) + 2na^2 ||\zeta||_{\mathrm{v}} + na + (n^2 - n)a^2.$$

Because

$$\int_{\mathbb{R}_+} \sum_{n \geq 0} (n^2 - n)\, e^{-\lambda s} \frac{(\lambda s)^n}{n!}\, G(ds) = \lambda^2 \int s^2\, G(ds),$$

we find using (2.1) that

$$A_{\mathrm{a}} h(\zeta) = h(\zeta) + 2\lambda s_1 a^2 ||\zeta||_{\mathrm{v}} + \lambda s_1 a + \lambda^2 s_2 a^2,$$

where $s_k = \int s^k\, G(ds)$ denotes the k-th moment of the interpolling distribution.

To calculate $A_{\mathrm{p}} h$, observe that

$$h(\zeta - \delta_x) - h(\zeta) = -2\langle \zeta, \delta_x \rangle_{\mathrm{a}} + a,$$

which shows that

$$A_{\mathrm{p}} h(\zeta) = h(\zeta) + a k_r(\zeta) - 2 \sum_{x \in \zeta} \langle \zeta, \delta_x \rangle_{\mathrm{a}}\, m(B_r(x) \cap \Gamma_\zeta(x)),$$

where $k_r(\zeta)$ is the probability of a successful scan, defined by (2.3). Lemma 3.7 implies that for all $\zeta \in M_+(S)$ and all $x \in \zeta$,

$$\sum_{y \in \zeta} (a - d(x,y))_+\, m(B_r(y) \cap \Gamma_\zeta(y)) \geq a^2,$$

so especially,

$$\sum_{y \in \zeta} \langle \zeta, \delta_y \rangle_{\mathrm{a}}\, m(B_r(y) \cap \Gamma_\zeta(y)) \geq a^2 ||\zeta||_{\mathrm{v}}.$$

Hence

$$A_{\mathrm{p}} h(\zeta) \leq h(\zeta) + a - 2a^2 ||\zeta||_{\mathrm{v}}.$$

Because

$$A_{\mathrm{a}} h(\zeta) = h(\zeta) + 2\lambda s_1 a^2 ||\zeta||_{\mathrm{v}} + \lambda s_1 a + \lambda^2 s_2 a^2$$

and

$$A_{\mathrm{a}}(||\cdot||_{\mathrm{v}})(\zeta) = ||\zeta||_{\mathrm{v}} + \lambda s_1,$$

we find that

$$A h(\zeta) \leq h(\zeta) + 2\lambda s_1 a^2 ||\zeta||_{\mathrm{v}} + \lambda s_1 a + \lambda^2 s_2 a^2 + a - 2a^2(\lambda s_1 + ||\zeta||_{\mathrm{v}}).$$

□

Proof of Theorem 3.8. By Lemma 3.10, the mean drift of the system with respect to $v(\zeta) = ||\zeta||_{\mathrm{a}}^2$ for $0 < a \le \min(\ell/2, 2r)$ satisfies

$$Dv(\zeta) \le -c_1 ||\zeta||_{\mathrm{v}} + c_2,$$

for $c_1 > 0$ and c_2. Define $f(\zeta) = 1 + c||\zeta||_{\mathrm{v}}$ where $c = c_1/2$, and let n be an integer such that $n \ge (1 + c_2)/(c_1/2)$. Then

$$Dv(\zeta) \le -f(\zeta)$$

for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}} > n$. Moreover,

$$Dv(\zeta) + f(\zeta) \le 1 + c_2$$

for all $\zeta \in M_+(S)$. Lemma 2.1 shows that the system is ϕ-irreducible and aperiodic, and the level set $C_n = \{\zeta \in M_+(S) : ||\zeta||_{\mathrm{v}} \le n\}$ is small and thus petite. Hence the f-norm ergodic theorem of Meyn and Tweedie [17, Theorem 14.0.1] shows the claim. □

3.5 Instability

Theorem 3.11. *If $\lambda s_1 \ge 1$, then the system is not positive recurrent, and if $\lambda s_1 > 1$, then $||W_t||_{\mathrm{v}} \to \infty$ almost surely as $t \to \infty$, regardless of the initial state.*

Proof. Denote the population process of the system by W, and let W' be the population process of a modified system with scan radius $r' = \ell/2$, so that the server finds a customer at polling events whenever the system is nonempty. Then $||W'||_{\mathrm{v}}$ equals the number of customers in a standard single-server queue with rate-λ Poisson arrivals and service times distributed according to G, observed just after service completions. It is not hard to see that there exists a coupling of W and W' so that $||W_t||_{\mathrm{v}} \ge ||W'_t||_{\mathrm{v}}$ for all $t \in \mathbb{Z}_+$ almost surely, whenever $||W_0||_{\mathrm{v}} \ge ||W'_0||_{\mathrm{v}}$ (see for instance Leskelä [14, Theorem 4.8]). When $\lambda s_1 \ge 1$, it is well-known that the mean return time to zero of $||W'||_{\mathrm{v}}$ is infinite [11]. The coupling then implies that the same is true for the process $||W||_{\mathrm{v}}$, which shows that W' is not positive recurrent. The second claim follows using the same coupling, because $||W'_t||_{\mathrm{v}} \to \infty$ almost surely when $\lambda s_1 > 1$. □

4 Geometric ergodicity

The standard single-server M/G/1 queue is known to be geometrically ergodic, when the tail of the service time distribution is light enough [10, Theorem 3.2]. An analogous result is true for the population process W, as the next result shows.

Theorem 4.1. *Assume that $\lambda s_1 < 1$ and the interpolling time distribution satisfies $\int e^{\theta s} G(ds) < \infty$ for some $\theta > 0$. Then the system is geometrically ergodic in the sense that there exist constants $\alpha > 0$, $\beta > 0$, and $c < \infty$ such that*

$$\sum_{t=0}^{\infty} e^{\alpha t} \sup_{g: |g| \le e^{\beta ||\cdot||_{\mathrm{a}}}} \left| \mathrm{E}_\zeta\, g(W_t) - \int g \, d\pi \right| \le c e^{\beta ||\zeta||_{\mathrm{a}}}$$

for all initial states $\zeta \in M_+(S)$.

Before proceeding with the proof of Theorem 4.1, we will show that the seminorm $||\zeta||_{\mathrm{a}}$ satisfies a similar Foster–Lyapunov bound (Lemma 4.3) as the function $\langle \zeta, \zeta \rangle_{\mathrm{a}} = ||\zeta||_{\mathrm{a}}^2$ in Lemma 3.10. Using the Foster–Lyapunov bound for $||\zeta||_{\mathrm{a}}$, we will then proceed along similar lines as in the proof of [17, Theorem 16.3.1] to bound the mean drift of the system with respect to the function $e^{\beta ||\zeta||_{\mathrm{a}}}$ for some $\beta > 0$ (Lemma 4.4), which is key to proving Theorem 4.1.

Lemma 4.2. *The function $x \mapsto \sqrt{1-x}$ satisfies*

$$\sqrt{1-x} = 1 - \frac{1}{2}x + R(x),$$

where $|R(x)| \le 2^{-3/2} x^2$ for all $x \in [-\frac{1}{2}, \frac{1}{2}]$.

Proof. Taylor's first order approximation shows that for all $x \in [-\frac{1}{2}, \frac{1}{2}]$, there exists $s \in [0,1]$ such that

$$R(x) = \frac{1}{8}(1 - sx)^{-3/2} x^2.$$

Because $(1-sx)^{-3/2} \le 2^{3/2}$ for all $|x| \le \frac{1}{2}$ and $s \in [0,1]$, the claim follows. □

Lemma 4.3. *Assume that $\lambda s_1 < 1$ and $s_2 < \infty$, and let $0 < a \le \min(\ell/2, 2r)$. Then there exist $\alpha > 0$, $b > 0$, and an integer n such that the mean drift of the system with respect to the seminorm $v(\zeta) = ||\zeta||_{\mathrm{a}}$ satisfies*

$$Dv(\zeta) \le -\alpha \tag{4.1}$$

for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}} > n$, and

$$Dv(\zeta) \le b \tag{4.2}$$

for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}} \le n$.

Proof. Jensen's inequality shows that $Av \leq (Av^2)^{1/2}$ holds pointwise on $M_+(S)$, so the mean drift with respect to v is bounded by

$$Dv \leq (Av^2)^{1/2} - v = (v^2 + Dv^2)^{1/2} - v. \tag{4.3}$$

Because $\lambda s_1 < 1$, we see using Lemma 3.10 that there exist $c > 0$ such that $Dv^2(\zeta) \leq -c||\zeta||_{\mathrm{v}}$ whenever $||\zeta||_{\mathrm{v}}$ is large enough. Because $||\zeta||_{\mathrm{a}} \leq ||\zeta||_{\mathrm{v}}$ (Lemma 3.1), we see that

$$Dv \leq (v^2 - cv)^{1/2} - v = v\left((1 - cv^{-1})^{1/2} - 1\right)$$

for all $||\zeta||_{\mathrm{v}}$ large enough. Lemma 3.4 shows that $v(\zeta) \to \infty$ as $||\zeta||_{\mathrm{v}} \to \infty$ in $M_+(S)$. Thus, for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}}$ is large enough, $cv^{-1} \leq \frac{1}{2}$, and using Lemma 4.2 we see that

$$\begin{aligned} Dv &\leq v\left(-\frac{1}{2}cv^{-1} + R(cv^{-1})\right) \\ &\leq v\left(-\frac{1}{2}cv^{-1} + 2^{-3/2}c^2v^{-2}\right) \\ &= -\frac{1}{2}c + 2^{-3/2}c^2v^{-1}. \end{aligned}$$

This shows the validity of (4.1) for a suitable chosen n.

Lemma 3.10 also shows that $Dv^2 \leq c_2$ for all $\zeta \in M_+(S)$. Because $v(\zeta) \leq ||\zeta||_{\mathrm{v}}$ (Lemma 3.1), inequality (4.3) shows that (4.2) holds with $b = (n^2 + c_2)^{1/2}$. □

Lemma 4.4. *Assume that $\lambda s_1 < 1$ and $\int_{\mathbb{R}_+} e^{\theta s}\, G(ds) < \infty$ for some $\theta > 0$, and let $0 < a \leq \min(\ell/2, 2r)$. Then there exist $\alpha > 0$, $\beta > 0$, $b > 0$, and an integer n such that the mean drift of the system with respect to $v_\beta(\zeta) = \exp(\beta||\zeta||_{\mathrm{a}})$ satisfies*

$$Dv_\beta(\zeta) \leq -\alpha v_\beta(\zeta) \tag{4.4}$$

for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}} > n$, and

$$Dv_\beta(\zeta) + \alpha v_\beta(\zeta) \leq b \tag{4.5}$$

for all $\zeta \in M_+(S)$ such that $||\zeta||_{\mathrm{v}} \leq n$.

Proof. Define $v_\beta(\zeta) = e^{\beta||\zeta||_{\mathrm{a}}}$ for some $\beta > 0$, to be chosen later. Then the mean drift with respect to v_β equals

$$Dv_\beta(\zeta) = v_\beta(\zeta)\, \mathrm{E}_\zeta\left\{e^{\beta(||W_1||_{\mathrm{a}} - ||W_0||_{\mathrm{a}})} - 1\right\}.$$

Using a first order Taylor series approximation we see that

$$e^{\beta t} = 1 + \beta t + R(t),$$

where the error term is bounded by $|R(t)| \le \frac{1}{2}\beta^2 t^2 e^{\beta|t|}$ for all $t \in \mathbb{R}$. Because $\frac{1}{2}t^2 \le s^{-2}e^{s|t|}$ for all t and all $s > 0$, it follows by setting $s = \beta^{1/3}$ that

$$|R(t)| \le \beta^{4/3} e^{(\beta+\beta^{1/3})|t|}.$$

This bound implies that

$$v_\beta(\zeta)^{-1} D v_\beta(\zeta) \le \beta D v(\zeta) + \beta^{4/3}\, \mathrm{E}_\zeta\, e^{(\beta+\beta^{1/3})|\, ||W_1||_{\mathrm{a}} - ||W_0||_{\mathrm{a}}|},$$

where we denote $v(\zeta) = ||\zeta||_{\mathrm{a}}$.

Let us next bound the exponential term. Because $||\zeta||_{\mathrm{a}}$ is a seminorm in the space of signed counting measures on S (Lemma 3.1), the triangle inequality shows that

$$|\, ||W_1||_{\mathrm{a}} - ||W_0||_{\mathrm{a}}| \le ||W_1 - W_0||_{\mathrm{a}}.$$

Let us write

$$W_1 - W_0 = \eta_a - \eta_p,$$

where η_a is a random counting measure describing the arrivals during an interpolling time, and η_p is a random counting measure describing the number of served customers during the first polling instant ($\eta_p = 0$ if the server sees no customers and $\eta_p = \delta_x$ for some $x \in W_0 + \eta_a$ otherwise). Because $||\cdot||_{\mathrm{a}} \le ||\cdot||_{\mathrm{v}}$ and $||\eta_p||_{\mathrm{v}} \le 1$, we find that

$$||W_1 - W_0||_{\mathrm{a}} \le ||\eta_a||_{\mathrm{v}} + 1.$$

Observe next that

$$\begin{aligned} \mathrm{E}_\zeta\, e^{(\beta+\beta^{1/3})||\eta_a||_{\mathrm{v}}} &= \int_{\mathbb{R}_+} \sum_{n=0}^{\infty} e^{(\beta+\beta^{1/3})n}\, e^{-\lambda s} \frac{(\lambda s)^n}{n!}\, G(ds) \\ &= \int_{\mathbb{R}_+} e^{\lambda(e^{(\beta+\beta^{1/3})}-1)s}\, G(ds). \end{aligned}$$

Choose now β small enough such that $\lambda(e^{(\beta+\beta^{1/3})} - 1) \le \theta$ and $e^{\beta+\beta^{1/3}} \le 2$. Then

$$v_\beta(\zeta)^{-1} D v_\beta(\zeta) \le \beta D v(\zeta) + 2\beta^{4/3}\hat{G}(\theta), \tag{4.6}$$

where $\hat{G}(\theta) = \int e^{\theta s} G(ds)$. By Lemma 4.3, $Dv(\zeta)$ is strictly negative for $||\zeta||_{\mathrm{v}}$ large enough. Hence there exist $\alpha > 0, \beta > 0$ and n such that (4.4) holds for $||\zeta||_{\mathrm{v}} > n$.

Inequality (4.6) further shows that

$$Dv_\beta(\zeta) + \alpha v_\beta(\zeta) \le \Big(\beta Dv(\zeta) + 2\beta^{4/3}\hat{G}(\theta) + \alpha\Big)\, v_\beta(\zeta).$$

for all $\zeta \in M_+(S)$. Because $||\zeta||_{\mathrm{a}} \le ||\zeta||_{\mathrm{v}}$ by Lemma 3.1, we have the bound $v_\beta(\zeta) \le e^{\beta n}$ for $||\zeta||_{\mathrm{v}} \le n$. Inequality (4.2) in Lemma 4.3 thus shows that (4.5) holds for some b large enough. $\square$

Proof of Theorem 4.1. By Lemma 2.1, the system is ϕ-irreducible and aperiodic. By Lemma 4.4, the geometric drift condition is satisfied for the level set $C_n = \{\zeta \in M_+(S) : ||\zeta||_{\mathrm{v}} \le n\}$ for some n large enough. The set C_n is small and thus petite by Lemma 2.1. Hence the claim follows by the geometric ergodic theorem of Meyn and Tweedie [17, Theorem 15.0.1] $\square$

5 Population size in steady state

Having seen that the population process W is positive recurrent for $\lambda s_1 < 1$ and $s_2 < \infty$, it would be interesting to find out an explicit expression for the stationary distribution of W, or at least for the stationary distribution of the number of customers in the system. Following [12], we could try the Laplace functional approach. Denote the stationary distribution of W by π, and denote the Laplace function of the stationary number of customers by

$$L(\theta) = \int e^{-\theta ||\zeta||_{\mathrm{v}}}\, \pi(d\zeta), \quad \theta > 0.$$

Then a straightforward calculation shows that

$$L(\theta) = \hat{G}(\lambda(1 - e^{-\theta})) \int e^{-\theta ||\zeta||_{\mathrm{v}}} (1 + (e^{\theta} - 1) k_r(\zeta))\, \pi(\zeta),$$

where $\hat{G}$ denotes the Laplace transform of G, and $k_r(\zeta)$ is defined by (2.3). However, solving $L(\theta)$ from the above equation appears intractable due to the nonlinearity of k_r.

To gain insight on the behavior of the number of customers in the system, we have numerically simulated the system for a choice of parameter combinations with exponential interpolling times. Figure 3 displays simulated paths of the population size for $\lambda = 0.1$ (left) and $\lambda = 0.9$ (right),

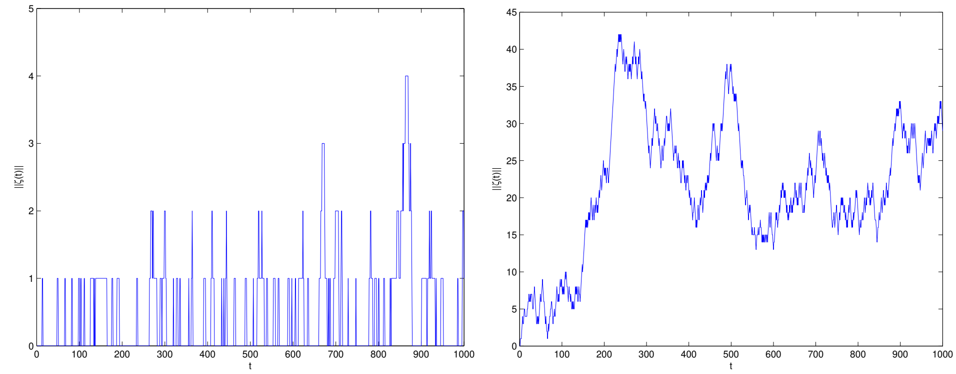

Figure 3: Population size as a function of time in light traffic ($\lambda = 0.1$, left) and heavy traffic ($\lambda = 0.9$, right).

where $s_1 = 1$, $r = 0.1$, and $\ell = 1$. The simulations suggest that the system in light traffic is empty a large proportion of time, whereas even for moderately heavy traffic ($\lambda s_1 = 0.9$), the empty system appears to be a rare event.

In light traffic, we can heuristically argue as follows. Assuming there is one customer present in the system, and no new customers arrive, finding the customer requires a geometrically distributed number of polling instants with parameter $m(B_r)$. Hence, assuming that no new customers arrive, the time required to serve the single customer has mean $s_1/m(B_r)$. Thus we might expect that the mean number of customers is roughly similar to the stationary probability that a renewal on/off process with mean on-time $1/\lambda$ and mean off-time $s_1/m(B_r)$ is in on-state, so that

$$\mathrm{E}\,||W||_{\mathrm{v}} \approx \frac{s_1/m(B_r)}{1/\lambda + s_1/m(B_r)} \approx \frac{\lambda s_1}{m(B_r)}. \tag{5.1}$$

In heavy traffic, many customers arrive during each interpolling time. Because the customers arrive to independent uniform locations, we can heuristically argue that with high probability there will be lots of customers relatively uniformly spread over the circle, so that the probability of finding a customer during a scan is close to one. Hence we might expect that in heavy traffic, the mean number of customers is not too far from the number of customers in a standard M/G/1 queue, so that

$$\mathrm{E}\,||W||_{\mathrm{v}} \approx \lambda + \frac{\lambda^2 s_2}{2(1-\lambda)}, \quad \lambda s_1 \approx 1. \tag{5.2}$$

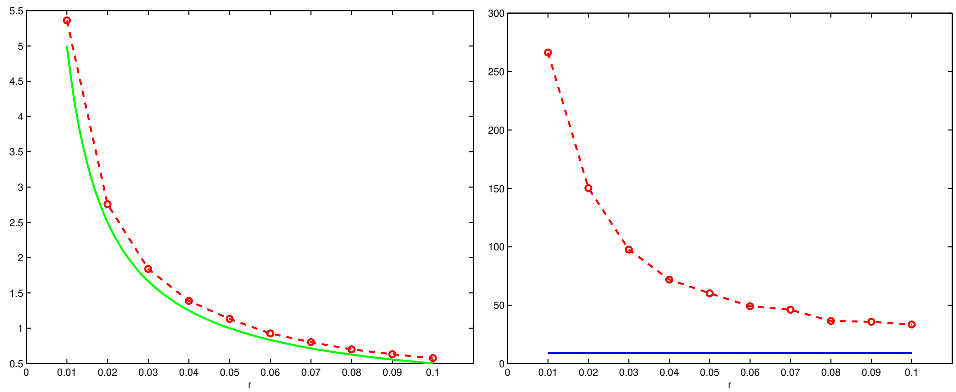

Figure 4: Stationary mean population size as a function of r for $\lambda = 0.1$ (left) and $\lambda = 0.9$ (right).

To study the accuracy of the above heuristics, Figure 4 shows numerically simulated values for the stationary mean population size for varying scan radius r. The solid line on the left is the light-traffic approximation (5.1); the solid line on the right is the heavy-traffic approximation (5.2). The plots suggest that the light-traffic approximation is relatively accurate. On the other hand, the heavy-traffic approximation, being insensitive to the value of the scan radius r, appears quite rough.

6 Conclusions

This paper studied a spatial queueing system on a circle, polled at random locations by a myopic server that can only observe customers in a bounded neighborhood. Using a novel quadratic Lyapunov functional of the measure-valued population process, we showed that the system is positive recurrent under a natural stability condition, and proved the geometric ergodicity of the system for light-tailed interpolling times. The behavior of the stationary system was discussed in terms of light-traffic and heavy-traffic heuristics, complemented with numerical simulations. The quadratic Lyapunov functional studied in this paper appears a promising tool for the analysis of more general spatial birth-and-death processes.

Acknowledgments

The results in this paper have been presented at the 15th INFORMS Applied Probability Conference, Cornell University, 12–15 July 2009. We thank Sergey Foss for introducing this problem and for insightful discussions during the 30th Finnish Summer School on Probability Theory. This work was initiated when L. Leskelä was employed by the Eindhoven University of Technology and F. Unger by the Centrum Wiskunde & Informatica (CWI), the Netherlands. L. Leskelä has been financially supported by the Academy of Finland.

References

[1] Altman, E. and Levy, H. (1994). Queueing in space. *Adv. Appl. Probab.*, 26(4):1095–1116.

[2] Bertsimas, D. J. and van Ryzin, G. (1993). A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Oper. Res.*, 39(4):601–615.

[3] Borovkov, K. A. and Odell, D. A. (2007). On spatial thinning-replacement processes based on Voronoi cells. *Adv. Appl. Probab.*, 39(2):293–306.

[4] Coffman, Jr., E. G. and Gilbert, E. N. (1987). Polling and greedy servers on a line. *Queueing Syst.*, 2(2):115–145.

[5] Daley, D. J. and Vere-Jones, D. (2003). *An Introduction to the Theory of Point Processes. Vol. I.* Probability and its Applications (New York). Springer-Verlag, New York, second edition. Elementary theory and methods.

[6] Ferrari, P. A., Fernández, R., and Garcia, N. L. (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stoch. Proc. Appl.*, 102(1):63–88.

[7] Foss, S. (2009). Some open problems related to stability. 100 Years of Queueing – The Erlang Centennial, 1-3 Apr 2009, Copenhagen, http://arXiv.org/abs/0909.0462.

[8] Foss, S. and Last, G. (1996). Stability of polling systems with exhaustive service policies and state-dependent routing. *Ann. Appl. Probab.*, 6(1):116–137.

[9] Garcia, N. L. and Kurtz, T. G. (2006). Spatial birth and death processes as solutions of stochastic equations. *ALEA Lat. Am. J. Probab. Math. Stat.*, 1:281–303 (electronic).

[10] Hou, Z. and Liu, Y. (2004). Explicit criteria for several types of ergodicity of the embedded $M/G/1$ and $GI/M/n$ queues. *J. Appl. Probab.*, 41(3):778–790.

[11] Kendall, D. G. (1951). Some problems in the theory of queues. *J. Roy. Statist. Soc. Ser. B.*, 13:151–173; discussion: 173–185.

[12] Kroese, D. P. and Schmidt, V. (1994). Single-server queues with spatially distributed arrivals. *Queueing Syst.*, 17(1-2):317–345.

[13] Kroese, D. P. and Schmidt, V. (1996). Light-traffic analysis for queues with spatially distributed arrivals. *Math. Oper. Res.*, 21(1):135–157.

[14] Leskelä, L. (2010). Stochastic relations of random variables and processes. *J. Theor. Probab.* http://arxiv.org/abs/0806.3562.

[15] Litvak, N. and Adan, I. (2001). The travel time in carousel systems under the nearest item heuristic. *J. Appl. Probab.*, 38(1):45–54.

[16] Meester, R. and Quant, C. (1999). Stability and weakly convergent approximations of queueing systems on a circle. Preprint 1093, Department of Mathematics, Utrecht University.

[17] Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer. Available online at http://probability.ca/MT/.

[18] Robert, P. (2009). Stochastic evolutions of point processes. http://arXiv.org/abs/0908.3256.

[19] Rudin, W. (1987). *Real and Complex Analysis*. McGraw–Hill, third edition.

(continued from the back cover)

HELSINKI UNIVERSITY OF TECHNOLOGY INSTITUTE OF MATHEMATICS
RESEARCH REPORTS

The reports are available at *http://math.tkk.fi/reports/* .
The list of reports is continued inside the back cover.

A578 Jarno Talponen
Special symmetries of Banach spaces isomorphic to Hilbert spaces
September 2009

A577 Fernando Rambla-Barreno, Jarno Talponen
Uniformly convex-transitive function spaces
September 2009

A576 S. Ponnusamy, Antti Rasila
On zeros and boundary behavior of bounded harmonic functions
August 2009

A575 Harri Hakula, Antti Rasila, Matti Vuorinen
On moduli of rings and quadrilaterals: algorithms and experiments
August 2009

A574 Lasse Leskelä, Philippe Robert, Florian Simatos
Stability properties of linear file-sharing networks
July 2009

ISBN 978-952-248-077-4 (print)
ISBN 978-952-248-081-1 (PDF)
ISSN 0784-3143 (print)
ISSN 1797-5867 (PDF)