Electronic Journal of Probability

Vol. 12 (2007), Paper no. 3, pages 75–99.

Journal URL
http://www.math.washington.edu/~ejpecp/

Asymptotic distributions and chaos for the supermarket model

Malwina J Luczak *
Department of Mathematics, London School of Economics,
Houghton Street, London WC2A 2AE, UK.
e-mail: m.j.luczak@lse.ac.uk

Colin McDiarmid
Department of Statistics, University of Oxford,
1 South Parks Road, Oxford OX1 3TG, UK.
e-mail: cmcd@stats.ox.ac.uk

Abstract

In the supermarket model there are n queues, each with a unit rate server. Customers arrive in a Poisson process at rate λn, where $0 < \lambda < 1$. Each customer chooses $d \geq 2$ queues uniformly at random, and joins a shortest one.

It is known that the equilibrium distribution of a typical queue length converges to a certain explicit limiting distribution as $n \to \infty$. We quantify the rate of convergence by showing that the total variation distance between the equilibrium distribution and the limiting distribution is essentially of order n^{-1}; and we give a corresponding result for systems starting from quite general initial conditions (not in equilibrium). Further, we quantify the result that the systems exhibit chaotic behaviour: we show that the total variation distance between the joint law of a fixed set of queue lengths and the corresponding product law is essentially of order at most n^{-1}.

Key words: Supermarket model, join the shortest queue, random choices, power of two

*supported in part by the Nuffield Foundation

choices, load balancing, equilibrium, concentration of measure, law of large numbers, chaos.

AMS 2000 Subject Classification: Primary 60C05; Secondary: 68R05, 90B22, 60K25, 60K30, 68M20.

Submitted to EJP on August 21 2005, final version accepted January 8 2007.

1 Introduction

We consider the following well-known scenario, often referred to as the 'supermarket model' (11; 12; 15; 16; 17). Let d be a fixed integer at least 2. Let n be a positive integer and suppose that there are n servers, each with a separate queue. Customers arrive in a Poisson process at rate λn, where $\lambda \in (0, 1)$ is a constant. Upon arrival each customer chooses d servers uniformly at random with replacement, and joins a shortest queue amongst those chosen. If there is more than one chosen server with a shortest queue, then the customer goes to the first such queue in her list of d. Service times are independent unit mean exponentials, and customers are served according to the first-come first-served discipline.

Recent work on the supermarket model includes (2; 3; 4; 9; 11; 12; 20). The survey (17) gives several applications and related results: an important application is to load balancing when computational tasks are dynamically assigned to servers in a large array. It is shown in (2; 3) that the system exhibits *propagation of chaos* given a suitable initial state, and in particular if it is in equilibrium. This means that the paths of members of any fixed finite subset of queues are asymptotically independent of one another, uniformly on bounded time intervals. This result implies a law of large numbers for the time evolution of the proportion of queues of different lengths, more precisely for the empirical measure on path space (2; 3). In particular for each fixed positive integer k_0, as n tends to infinity the proportion of queues with length at least k_0 at time t converges weakly to a function $v_t(k_0)$, where $v_t(0) = 1$ for all $t \geq 0$ and $(v_t(k) : k \in \mathbb{N})$ is the unique solution to the system of differential equations

$$\frac{dv_t(k)}{dt} = \lambda(v_t(k-1)^d - v_t(k)^d) - (v_t(k) - v_t(k+1)) \tag{1}$$

for $k \in \mathbb{N}$ (see (20)). Here one assumes appropriate initial conditions $v_0 = (v_0(k) : k \in \mathbb{N})$ such that $1 \geq v_0(1) \geq v_0(2) \geq \cdots \geq 0$, and $v_0 \in l_1$. Further, again for a fixed positive integer k_0, as n tends to infinity, in the equilibrium distribution the proportion of queues with length at least k_0 converges in probability to $\lambda^{(d^{k_0}-1)/(d-1)}$, and thus the probability that a given queue has length at least k_0 also converges to $\lambda^{(d^{k_0}-1)/(d-1)}$.

Recent results in (11) include rapid mixing and two-point concentration for the maximum queue length in equilibrium.

The main contribution of the present paper is to give quantitative versions of the convergence results for the supermarket model mentioned above, and to extend them to hold uniformly over all times. We rely in part on combinatorial techniques developed in (10; 11).

For each time $t \geq 0$ and each $j = 1, \ldots, n$ let $X_t^{(n)}(j)$ denote the number of customers in queue j, always including the customer currently being served if there is one. We shall keep the superscript 'n' in the notation in this section, but then usually drop it in later sections. We make the standard assumption of right-continuity of the sample paths. Let $X_t^{(n)}$ be the *queue-lengths vector* $(X_t^{(n)}(1), \ldots, X_t^{(n)}(n)) \in \mathbb{Z}_+^n$, where $\mathbb{Z}_+^n$ denotes the set of all n-vectors with components taking non-negative integer values. Note that the l_1-norm $\| X_t \|_1$ of X_t is the total number of customers present at time t, and the l_∞-norm $\| X_t \|_\infty$ is the maximum queue length.

For a given positive integer n, the n-queue process $(X_t^{(n)})$ is an ergodic continuous-time Markov chain. Thus there is a unique stationary distribution $\mathbf{\Pi}^{(n)}$ for the vector $X_t^{(n)}$; and, whatever the distribution of the starting state, the distribution of the vector $X_t^{(n)}$ at time t converges to $\mathbf{\Pi}^{(n)}$

as $t \to \infty$. As already noted in (2; 3) (and easily verified), the distribution $\mathbf{\Pi}^{(n)}$ is *exchangeable*, that is invariant under permuting the co-ordinates. We shall usually write $Y_t^{(n)}$ to denote the queue-lengths vector in equilibrium: we drop the subscript t when no explicit reference to a particular time is needed.

The probability law of a random element X will be denoted by $\mathcal{L}(X)$. The *total variation distance* between two probability distributions μ_1 and μ_2 is defined by $d_{\mathrm{TV}}(\mu_1, \mu_2) = \sup_A |\mu_1(A) - \mu_2(A)|$ where the supremum is over all measurable sets A in the underlying measurable space (see also the start of Section 2). Also, given a vector $v = (v(k) : k = 0, 1, \ldots)$ such that

$$1 = v(0) \geq v(1) \geq v(2) \geq \ldots \geq 0 \text{ and } v(k) \to 0 \text{ as } k \to \infty, \tag{2}$$

let $\mathcal{L}_v$ denote the law of a random variable V taking non-negative integer values, where $\mathbf{Pr}(V \geq k) = v(k)$ for each $k = 0, 1, 2, \ldots$. In fact, throughout this paper we shall work only with vectors $v \in l_1$.

Finally, throughout we use the asymptotic notations $O(), \Omega(), o(), \omega()$ in a standard way, to describe the behaviour of functions depending on the number of servers n as n tends to infinity; for instance $f(n) = \Omega(g(n))$ means that, for some constants $c > 0$ and n_0, we have $f(n) \geq cg(n)$ for all $n \geq n_0$.

We may now state four main results, two concerning approximating the distribution of a single typical queue length and two concerning collections of queues and chaos. These will be proved in the following sections, where we also present some further results.

1.1 Single queues

We first consider the n-queue system in equilibrium, and investigate how close the distribution of a typical queue length $Y^{(n)}(1)$ is to the limiting distribution. Let $\mathcal{L}_{\lambda,d}$ denote the law $\mathcal{L}_v$ where $v(k) = \lambda^{(d^k-1)/(d-1)}$ for each $k = 0, 1, \ldots$. Note that $\mathbf{Pr}(Y^{(n)}(1) \geq 1) = \lambda = v(1)$. It is known (and was mentioned earlier) that $\mathcal{L}(Y^{(n)}(1))$ tends to $\mathcal{L}_{\lambda,d}$ as $n \to \infty$: we now quantify this convergence.

Theorem 1.1. *For each positive integer n let $Y^{(n)}$ be a queue-lengths n-vector in equilibrium, and consider the length $Y^{(n)}(1)$ of queue 1. Then*

$$d_{TV}(\mathcal{L}(Y^{(n)}(1)), \mathcal{L}_{\lambda,d})$$

is of order n^{-1} up to logarithmic factors.

In fact, we shall see that the above total variation distance is $o(n^{-1} \ln^3 n)$ and is $\Omega(n^{-1})$. Also, we shall deduce directly from Theorem 1.1, together with a bound on the maximum queue length from (11) (given as (3) below), the following:

Corollary 1.2. *For each positive integer k, the difference between the kth moment $\mathbf{E}[Y^{(n)}(1)^k]$ and the kth moment of $\mathcal{L}_{\lambda,d}$ is of order n^{-1} up to logarithmic factors.*

Now we drop the assumption that the system is in equilibrium, and consider its behaviour from the beginning (that is, from time 0). We state one theorem here, Theorem 1.3. More general though less digestible results (Theorems 3.3 and 3.4) are stated and proved in Section 3, and

Theorem 1.3 will follow easily from Theorem 3.4. We assume in Theorem 1.3 that the initial queue lengths are iid and not too large; and see that the law of a typical queue length is close to $\mathcal{L}_{v_t}$, where $v_t = (v_t(k) : k \in \mathbb{N})$ is the solution to the system (1) subject to the natural initial conditions.

Given a queue-lengths vector x, that is a vector with non-negative integer co-ordinates, we let $u(k, x)$ be the proportion of queues with length at least k.

Theorem 1.3. *There is a constant $\varepsilon > 0$ such that the following holds. Let the random variable X take non-negative values, and suppose that $\mathbf{E}[e^{X/\varepsilon}]$ is finite. For each positive integer n, let $(X_t^{(n)})$ be an n-queue process where the random initial queue lengths $X_0^{(n)}(1), \ldots, X_0^{(n)}(n)$ are iid and stochastically at most X. Let $v_t^{(n)} = (v_t^{(n)}(k) : k = 0, 1, 2, \ldots)$ be the unique solution to the system (1) subject to the initial conditions $v_0^{(n)}(k) = \mathbf{E}[u(k, X_0^{(n)})]$ for $k = 0, 1, \ldots$.*

Then the law of $X_t^{(n)}$ is exchangeable for each $t \geq 0$, and

$$\sup_{t \geq 0} d_{TV}(\mathcal{L}(X_t^{(n)}(1)), \mathcal{L}_{v_t^{(n)}}) = O(n^{-\varepsilon}).$$

1.2 Collections of queues

The above results concern the distribution of a single queue length. We now consider collections of queues and propagation of chaos. The term "propagation of chaos" comes from statistical physics (7), and the original motivation was the evolution of particles in physical systems. The subject has since then received considerable attention culminating in the work of Sznitman (18).

Our results below establish chaoticity for the supermarket model. As before, we first discuss the equilibrium distribution. We see that for fixed r the total variation distance between the joint law of r queue lengths and the product law is at most $O(n^{-1})$, up to logarithmic factors. More precisely and more generally we have:

Theorem 1.4. *For each positive integer n, let $Y^{(n)}$ be a queue-lengths n-vector in equilibrium. Then, uniformly over all positive integers $r \leq n$, the total variation distance between the joint law of $Y^{(n)}(1), \ldots, Y^{(n)}(r)$ and the product law $\mathcal{L}(Y^{(n)}(1))^{\otimes r}$ is at most $O(n^{-1} \ln^2 n (2 \ln \ln n)^r)$; and the total variation distance between the joint law of $Y^{(n)}(1), \ldots, Y^{(n)}(r)$ and the limiting product law $\mathcal{L}_{\lambda,d}^{\otimes r}$ is at most $O(n^{-1} \ln^2 n (2 \ln \ln n)^{r+1})$.*

Note that since the distribution of $Y^{(n)}$ is exchangeable, for any distinct indices $j_1, \ldots, j_r$ the joint distribution of $Y^{(n)}(j_1), \ldots, Y^{(n)}(j_r)$ is the same as that of $Y^{(n)}(1), \ldots, Y^{(n)}(r)$. Also, if d were 1 we would have 'exact' independence of all queues in equilibrium. Further, note that the above result can yield a bound less than 1 only if $r < \ln n / \ln \ln \ln n$; but as long as $r = o(\ln n / \ln \ln \ln n)$ the bound is $O(n^{-1+o(1)})$. We shall mention in Section 4 how this result relates to Sznitman's treatment of chaos in (18) – see in particular the inequalities (19) and (20) below.

Now we drop the assumption that the system is in equilibrium, and consider its behaviour from time 0. We state one theorem here, Theorem 1.5. More general though less digestible results (Theorems 4.1 and 4.2) are stated and proved in Section 4, and Theorem 1.5 will follow easily from Theorem 4.2. As with Theorem 1.3 earlier, we assume in Theorem 1.5 that the initial queue lengths are iid and not too large: now we see that the joint law of r queue lengths is close to the product law, that is we have chaotic behaviour, uniformly for all times.

Given an n-queue process $(X_t^{(n)})$ and a positive integer $r \leq n$, let $\mathcal{L}_t^{(n,r)}$ denote the joint law of $X_t^{(n)}(1), \ldots, X_t^{(n)}(r)$, and let $\tilde{\mathcal{L}}_t^{(n,r)}$ denote the product law $\mathcal{L}(X_t^{(n)}(1))^{\otimes r}$ of r independent copies of $X_t^{(n)}(1)$.

Theorem 1.5. *Let the random variable X take non-negative values, and suppose that the moment generating function $\mathbf{E}[e^{s_0 X}]$ is finite for some $s_0 > 0$. For each integer $n \geq r$, let $(X_t^{(n)})$ be an n-queue process where the random initial queue lengths $X_0^{(n)}(1), \ldots, X_0^{(n)}(n)$ are iid and stochastically at most X. Then the law of $X_t^{(n)}$ is exchangeable for each $t \geq 0$, and uniformly over positive integers $r \leq n$*

$$\sup_{t \geq 0} d_{TV}\left(\mathcal{L}_t^{(n,r)}, \tilde{\mathcal{L}}_t^{(n,r)}\right) = O\left(n^{-1}(\ln n)^{r+2}\right).$$

We remark also that an $O(n^{-1})$ upper bound on the total variation distance between the law of a finite r-tuple of queues and the product law is known on fixed time intervals for iid initial queue lengths - see (3), Theorem 3.5. However, the bound in (3) grows exponentially in time, and does not extend to the equilibrium case. We are not aware of any earlier equilibrium bounds or time-uniform bounds like those given in Theorems 1.4, 1.5, 4.1 and 4.2.

Finally, let us mention at this point that connections between rapid mixing, concentration of measure and chaoticity have been observed earlier in various other contexts. The reader is referred to (6; 8; 18; 19) and references therein for more information.

2 Preliminaries

This section contains some definitions and some results from (11) needed in our proofs. We start with two lemmas which show that the supermarket model is rapidly mixing. Lemma 2.1 upper bounds the total variation distance and Lemma 2.2 upper bounds the Wasserstein distance between $\mathcal{L}(X_t)$ and the equilibrium distribution $\mathbf{\Pi}$.

We say that a real-valued function f defined on a subset A of $\mathbb{R}^n$ is 1-Lipschitz if $|f(x) - f(y)| \leq \| x - y \|_1$ for all $x, y \in A$. Let us recall two equivalent definitions of the total variation distance $d_{\mathrm{TV}}(\mu_1, \mu_2)$ between two probability distributions μ_1 and μ_2 on $\mathbb{Z}_+^n$, and two corresponding definitions of the Wasserstein distance $d_{\mathrm{W}}(\mu_1, \mu_2)$. We have $d_{\mathrm{TV}}(\mu_1, \mu_2) = \inf \mathbf{Pr}(X \neq Y)$, and $d_{\mathrm{W}}(\mu_1, \mu_2) = \inf \mathbf{E}[\| X - Y \|_1]$, where in each case the infimum is over all couplings of X and Y where $\mathcal{L}(X) = \mu_1$ and $\mathcal{L}(Y) = \mu_2$. Also,

$$d_{\mathrm{TV}}(\mu_1, \mu_2) = \frac{1}{2} \sup_{\|\phi\|_\infty \leq 1} \left| \int \phi \, d\mu_1 - \int \phi \, d\mu_2 \right|$$

where the supremum is over measurable functions ϕ with $\| \phi \|_\infty \leq 1$; and

$$d_{\mathrm{W}}(\mu_1, \mu_2) = \sup_{f \in \mathcal{F}_1} \left| \int f \, d\mu_1 - \int f \, d\mu_2 \right|$$

where the supremum is over the set $\mathcal{F}_1$ of 1-Lipschitz functions f on $\mathbb{Z}_+^n$. The total variation distance between the corresponding laws on $\mathbb{Z}_+^n$ is always at most the Wasserstein distance: to see this, note for example that $\mathbb{I}_{X\neq Y} \leq \| X - Y \|_1$.

Recall that $0 < \lambda < 1$ and $d \in \mathbb{N}$ are fixed.

Lemma 2.1. *For each constant $c > 0$ there exists a constant $\eta > 0$ such that the following holds for each positive integer n. For each time $t \geq 0$ let*

$$\delta_{n,t} = \mathbf{Pr}(\| X_0^{(n)} \|_1 > cn) + \mathbf{Pr}(\| X_0^{(n)} \|_\infty > \eta t).$$

Then

$$d_{TV}(\mathcal{L}(X_t^{(n)}), \mathbf{\Pi}) \leq ne^{-\eta t} + 2e^{-\eta n} + \delta_{n,t}.$$

Lemma 2.2. *For each constant $c > \frac{\lambda}{1-\lambda}$ there exists a constant $\eta > 0$ such that the following holds for each positive integer n. Let $M^{(n)}$ denote the stationary maximum queue length. Consider any distribution of the initial queue-lengths vector $X_0^{(n)}$ such that $\| X_0^{(n)} \|_1$ has finite mean. For each time $t \geq 0$ let*

$$\delta_{n,t} = 2\mathbf{E}[\| X_0^{(n)} \|_1 \mathbf{1}_{\|X_0^{(n)}\|_1 > cn}] + 2cn\ \mathbf{Pr}(\| X_0^{(n)} \|_\infty > \eta t).$$

Then

$$d_W(\mathcal{L}(X_t^{(n)}), \mathbf{\Pi}) \leq ne^{-\eta t} + 2cn\mathbf{Pr}(M^{(n)} > \eta t) + 2e^{-\eta n} + \delta_{n,t}.$$

We now introduce the natural coupling of n-queue processes (X_t) with different initial states. Arrival times form a Poisson process at rate λn, and there is a corresponding sequence of uniform choices of lists of d queues. Potential departure times form a Poisson process at rate n, and there is a corresponding sequence of uniform selections of a queue: potential departures from empty queues are ignored. These four processes are independent. Denote the arrival time process by $\mathbf{T}$, the choices process by $\mathbf{D}$, the potential departure time process by $\tilde{\mathbf{T}}$ and the selection process by $\tilde{\mathbf{D}}$.

Suppose that we are given a sequence of arrival times $\mathbf{t}$ with corresponding queue choices $\mathbf{d}$, and a sequence of potential departure times $\tilde{\mathbf{t}}$ with corresponding selections $\tilde{\mathbf{d}}$ of a queue (where all these times are distinct). For each possible initial queue-lengths vector $x \in \mathbb{Z}_+^n$ this yields a deterministic queue-lengths process (x_t) with $x_0 = x$: let us write $x_t = s_t(x; \mathbf{t}, \mathbf{d}, \tilde{\mathbf{t}}, \tilde{\mathbf{d}})$. Then for each $x \in \mathbb{Z}_+^n$, the process $(s_t(x; \mathbf{T}, \mathbf{D}, \tilde{\mathbf{T}}, \tilde{\mathbf{D}}))$ has the distribution of a queue-lengths process with initial state x. The following lemma from (11) shows that the coupling has certain desirable properties. Part (c) is not stated explicitly in (11), but it follows easily from part (b). Inequalities between vectors in the statement are understood to hold component by component. 'Adding a customer' means adding an arrival time and corresponding choice of queues.

Lemma 2.3. *Fix any 4-tuple $\mathbf{t}, \mathbf{d}, \tilde{\mathbf{t}}, \tilde{\mathbf{d}}$ as above, and for each $x \in \mathbb{Z}_+^n$ write $s_t(x)$ for $s_t(x; \mathbf{t}, \mathbf{d}, \tilde{\mathbf{t}}, \tilde{\mathbf{d}})$. Then*

(a) for each $x, y \in \mathbb{Z}_+^n$, both $\| s_t(x) - s_t(y) \|_1$ and $\| s_t(x) - s_t(y) \|_\infty$ are non-increasing.

(b) if $0 \leq t < t'$ and $s_t(x) \leq s_t(y)$ then $s_{t'}(x) \leq s_{t'}(y)$;

(c) if $\mathbf{t}'$ and $\mathbf{d}'$ are obtained from $\mathbf{t}$ and $\mathbf{d}$ by adding some extra customers then, for all $t \geq 0$, $s_t(x; \mathbf{t}', \mathbf{d}', \tilde{\mathbf{t}}, \tilde{\mathbf{d}}) \geq s_t(x; \mathbf{t}, \mathbf{d}, \tilde{\mathbf{t}}, \tilde{\mathbf{d}})$.

Next, let us consider the equilibrium distribution, and note some upper bounds on the total number of customers in the system and on the maximum queue length established in (11).

Lemma 2.4. *(a) For any constant $c > \frac{\lambda}{1-\lambda}$, there is a constant $\eta > 0$ such that for each positive integer n, in equilibrium the queue-lengths process $(Y_t^{(n)})$ satisfies*

$$\mathbf{Pr}\left(\| Y_t^{(n)} \|_1 > cn\right) \leq e^{-\eta n}$$

for each time $t \geq 0$.

(b) For each positive integer n, in equilibrium the maximum queue-length $M_t^{(n)}$ satisfies

$$\mathbf{Pr}\left(M_t^{(n)} \geq j\right) \leq n\lambda^j$$

for each positive integer j and each time $t \geq 0$.

We shall require an upper bound on the maximum length $M^{(n)}$ of a queue in equilibrium, from Section 7 of (11). Let $i^* = i^*(n)$ be the smallest integer i such that $\lambda^{\frac{d^i-1}{d-1}} < n^{-1/2} \ln^2 n$. Then $i^* = \ln\ln n/\ln d + O(1)$; and with probability tending to 1 as $n \to \infty$, if $d = 2$ then $M^{(n)}$ is i^* or $i^* + 1$, and if $d \geq 3$ then $M^{(n)}$ is $i^* - 1$ or i^*. The bound we need is that if $r = r(n) = O(\ln n)$ then

$$\mathbf{Pr}(M^{(n)} \geq i^* + 1 + r) = e^{-\Omega(r \ln n)}. \tag{3}$$

Now we state some concentration of measure results for the queue-lengths process $(X_t^{(n)})$. Let us begin with the equilibrium case, where we use the notation $Y^{(n)}$. Recall that $\mathcal{F}_1$ denotes the set of 1-Lipschitz functions on $\mathbb{Z}_+^n$. (We suppress the dependence on n.)

Lemma 2.5. *There is a constant $c > 0$ such that the following holds. Let $n \geq 2$ be an integer and consider the n-queue system in equilibrium. Then for each $f \in \mathcal{F}_1$ and each $u \geq 0$*

$$\mathbf{Pr}\left(|f(Y^{(n)}) - \mathbf{E}[f(Y^{(n)})]| \geq u\right) \leq ne^{-cu/n^{\frac{1}{2}}}.$$

Let $\ell(k,x)$ denote $|\{j : x(j) \geq k\}|$, the number of queues of length at least k. Thus $\ell(k,x) = nu(k,x)$. Tight concentration of measure estimates for the random variables $\ell(k,Y)$ may be obtained directly from the last lemma.

Lemma 2.6. *Consider the stationary queue-lengths vector $Y^{(n)}$, and for each non-negative integer k let $\ell(k) = \mathbf{E}[\ell(k, Y^{(n)})]$. Then for any constant $c > 0$,*

$$\mathbf{Pr}(\sup_k |\ell(k, Y^{(n)}) - \ell(k)| \geq cn^{\frac{1}{2}} \ln^2 n) = e^{-\Omega(\ln^2 n)}.$$

Also, there exists a constant $c > 0$ such that

$$\sup_k \mathbf{Pr}(|\ell(k, Y^{(n)}) - \ell(k)| \geq cn^{\frac{1}{2}} \ln n) = o(n^{-2}).$$

Furthermore, for each fixed integer $r \geq 2$

$$\sup_k |\mathbf{E}[\ell(k, Y^{(n)})^r] - \ell(k)^r| = O(n^{r-1} \ln^2 n).$$

The first and third parts of this lemma form Lemma 4.2 in (11): the second part follows directly from the preceding lemma. We now present a time-dependent result, which will be essential in the proof of Theorem 3.3.

Lemma 2.7. *There is a constant $c > 0$ such that the following holds. Let $n \geq 2$ be an integer, let $f \in \mathcal{F}_1$, let $x_0 \in \mathbb{Z}_+^n$, and let $X_0^{(n)} = x_0$ almost surely. Then for all times $t \geq 0$ and all $u > 0$,*

$$\mathbf{Pr}(|f(X_t^{(n)}) - \mathbf{E}[f(X_t^{(n)})]| \geq u) \leq n\ e^{-\frac{cu^2}{nt+u}}. \tag{4}$$

We shall also use the following extension of Bernstein's inequality, which will follow easily for example from Theorem 3.8 of (14).

Lemma 2.8. *Let f be a 1-Lipschitz function on $\mathbb{R}^n$. Let $\mathbf{Z} = (Z_1, \ldots, Z_n)$ where $Z_1, \ldots, Z_n$ are independent real-valued random variables, such that Z_j has variance at most σ_j^2 and range at most b, for each $j = 1, \ldots, n$. Let $v = \sum_{j=1}^n \sigma_j^2$. Then for each $u > 0$*

$$\mathbf{Pr}(|f(\mathbf{Z}) - \mathbf{E}[f(\mathbf{Z})]| \geq u) \leq 2 \exp\left(-\frac{u^2}{2v + \frac{2}{3}bu}\right).$$

Proof. Given fixed numbers $z_1, \ldots, z_{j-1}$ where $1 \leq j \leq n$, for each z let

$$g(z) \quad = \quad \mathbf{E}[f(z_1, \ldots, z_{j-1}, z, Z_{j+1}, \ldots, Z_n)].$$

Then the function g is 1-Lipschitz, so the random variable $g(Z_j)$ has variance at most σ_j^2 and range at most b. Thus, in the terms of Theorem 3.8 of (14), the 'sum of variances' is at most v, and the 'maximum deviation' is at most b; and so we may use that theorem to deduce the result here. $\square$

Let $u_t(k) = \mathbf{E}[u(k, X_t)]$, the expected proportion of queues of length at least k at time t. Also let $u(k)$ denote $\mathbf{E}[u(k, Y)]$, the expected proportion of queues with at least k customers when the process is in equilibrium. Then $u(0) = 1$ and $u(1) = \lambda$. Whatever the initial distribution of X_0, for each positive integer k,

$$\frac{du_t(k)}{dt} = \lambda\left(\mathbf{E}[u(k-1, X_t)^d] - \mathbf{E}[u(k, X_t)^d]\right) - (u_t(k) - u_t(k+1)); \tag{5}$$

and for Y in equilibrium

$$0 = \lambda\left(\mathbf{E}[u(k-1, Y)^d] - \mathbf{E}[u(k, Y)^d]\right) - (u(k) - u(k+1)). \tag{6}$$

In (11), this last fact is used to show that $u(k)$ is close to $\lambda^{(d^k-1)/(d-1)}$; more precisely, for some constant $c > 0$, for each positive integer n

$$\sup_k |u(k) - \lambda^{(d^k-1)/(d-1)}| \leq cn^{-1} \ln^2 n. \tag{7}$$

3 Distribution of a single queue length

3.1 Equilibrium case

In this subsection we prove Theorem 1.1. Let us note that the equilibrium distribution $\mathbf{\Pi}$ is exchangeable, and thus all queue lengths are identically distributed.

We begin by showing that (a) in equilibrium the total variation distance between the marginal distribution of a given queue length and the limiting distribution $\mathcal{L}_{\lambda,d}$ is small; and (b) without assuming that the system is in equilibrium, a similar result holds after a logarithmic time, given a suitable initial distribution. [In fact, the result (a) will follow easily from (b), since by Lemma 2.4 and (3), if X_0 is in equilibrium and we set $c_0 > \lambda/(1-\lambda)$ then the quantity δ_n in Proposition 3.1 below is $o(n^{-1})$.] Part (a) includes the upper bound part of Theorem 1.1 above.

Proposition 3.1.

(a) The equilibrium length $Y^{(n)}(1)$ of queue 1 satisfies

$$d_{TV}(\mathcal{L}(Y^{(n)}(1)), \mathcal{L}_{\lambda,d}) = O\left(n^{-1}\ln^2 n \ \ln\ln n\right).$$

(b) For any constant $c_0 > 0$ there is a constant c such that the following holds. Let $\delta_n = \mathbf{Pr}(\| X_0^{(n)} \|_1 > c_0 n) + \mathbf{Pr}(\| X_0^{(n)} \|_\infty > c_0 \ln n)$. Then uniformly over $t \geq c \ln n$

$$\max_{j=1,\dots,n} d_{TV}(\mathcal{L}(X_t^{(n)}(j)), \mathcal{L}_{\lambda,d}) \leq \delta_n + O(n^{-1}\ln^2 n \ \ln\ln n).$$

Let us note here that better bounds may be obtained in the simple case $d = 1$. Here the n queues behave independently, and $\mathcal{L}(Y(1)) = \mathcal{L}_{\lambda,d}$ for each n. The arrival rate at each queue is always λ, regardless of the state of all the other queues. Then it follows from the proof of Lemma 2.1 in (11) (which is stated as Lemma 2.1 here) that we can drop the term $\mathbf{Pr}(\| X_0 \|_1 > c_0 n)$ in the mixing bound of that lemma, and so part (b) of Proposition 3.1 holds with the bound $\delta_n + O(n^{-1}\ln^2 n \ln\ln n)$ replaced by $\mathbf{Pr}(\| X_0 \|_\infty > c_0 \ln n) + O(n^{-K})$ for any constant K.

Proof of Proposition 3.1. Since the distribution of Y is exchangeable, $\mathbf{Pr}(Y(1) \geq k) = u(k)$ for each non-negative integer k. Note that $u(0) - u(1) = 1 - \lambda$. Part (a) now follows easily from (7), since

$$\begin{aligned} d_{\mathrm{TV}}(\mathcal{L}(Y(1)), \mathcal{L}_{\lambda,d}) &= \frac{1}{2}\sum_{k=1}^{\infty} |u(k) - u(k+1) - \lambda^{1+d+\cdots+d^{k-1}} + \lambda^{1+d+\cdots+d^k}| \\ &\leq \sum_{k=1}^{k_0+1} |u(k) - \lambda^{1+d+\cdots+d^{k-1}}| + u(k_0+1) + \lambda^{1+d+\cdots+d^{k_0}}, \end{aligned}$$

for any positive integer k_0. But if $k_0 \geq \ln\ln n/\ln d + c$, where $c = -\ln\ln(1/\lambda)/\ln d$ then $\lambda^{1+d+\cdots+d^{k_0}} \leq n^{-1}$, and so also $u(k_0+1) = O(n^{-1}\ln^2 n)$.

For part (b) note that

$$\begin{aligned} d_{\mathrm{TV}}(\mathcal{L}(X_t(j)), \mathcal{L}_{\lambda,d}) &\leq d_{\mathrm{TV}}(\mathcal{L}(X_t(j)), \mathcal{L}(Y(j))) + d_{\mathrm{TV}}(\mathcal{L}(Y(j)), \mathcal{L}_{\lambda,d}) \\ &\leq d_{\mathrm{TV}}(\mathcal{L}(X_t), \mathcal{L}(Y)) + d_{\mathrm{TV}}(\mathcal{L}(Y(1)), \mathcal{L}_{\lambda,d}). \end{aligned}$$

Thus, by part (a) and by Lemma 2.1, there exists a constant $\eta = \eta(c_0)$ such that for each time $t \geq \eta^{-1} c_0 \ln n$

$$\max_j d_{\mathrm{TV}}(\mathcal{L}(X_t(j)), \mathcal{L}_{\lambda,d}) \leq \delta_n + ne^{-\eta t} + 2e^{-\eta n} + O(n^{-1} \ln^2 n \; \ln\ln n),$$

and the result follows with $c_1 = \max\{2, c_0\}\ \eta^{-1}$. ∎

We now show that the $O\left(n^{-1} \ln^2 n \; \ln\ln n\right)$ upper bound on the total variation distance between the equilibrium distribution of a given queue length and $\mathcal{L}_{\lambda,d}$ in Proposition 3.1 (a) is fairly close to optimal. The next lemma will complete the proof of Theorem 1.1.

Lemma 3.2. *For an n-queue system in equilibrium, the expected proportion $u(2)$ of queues of length at least 2 satisfies $u(2) \geq \lambda^{d+1} + \Omega(n^{-1})$. Hence*

$$d_{TV}(\mathcal{L}(Y^{(n)}(1)), \mathcal{L}_{\lambda,d}) = \Omega\left(n^{-1}\right).$$

Proof. Let $F_t = \ell(1, Y_t)$ be the number of non-empty queues at time t, and write F for F_1. We shall show that the variance of F is $\Omega(n)$, and from that we shall complete the proof quickly.

Recall that we model departures by a Poisson process at rate n (giving potential departure times) together with an independent selection process that picks a uniformly random queue at each event time of this process. If the queue selected is nonempty, then the customer currently in service departs; otherwise nothing happens.

Let Z be the number of arrivals in $[0, 1]$. By the last part of Lemma 2.3 we have the monotonicity result that for all non-negative integers x and z,

$$\mathbf{Pr}(F \leq x | Z = z) \geq \mathbf{Pr}(F \leq x | Z = z + 1). \tag{8}$$

Let the integer $x = x(n)$ be a conditional median of F given that $Z = \lfloor \lambda n \rfloor$. (It is not hard to see that $x = \lambda n + O((n \ln n)^{\frac{1}{2}})$ but we do not use this here.) Since $\mathbf{Pr}(F \leq x | Z = \lfloor \lambda n \rfloor) \geq \frac{1}{2}$ we have by (8) that

$$\begin{aligned} \mathbf{Pr}(F \leq x) &\geq \mathbf{Pr}(F \leq x | Z \leq \lambda n)\mathbf{Pr}(Z \leq \lambda n) \\ &\geq \mathbf{Pr}(F \leq x | Z = \lfloor \lambda n \rfloor)\mathbf{Pr}(Z \leq \lambda n) \\ &\geq \frac{1}{4} + o(1). \end{aligned}$$

We shall find a constant $\delta > 0$ such that $\mathbf{Pr}(F \geq x + \delta n^{\frac{1}{2}}) \geq \frac{1}{14} + o(1)$, which will show that the variance of F is $\Omega(n)$ as required.

Let A be the event that $Z \geq \lambda n + (\lambda n)^{\frac{1}{2}}$. Then, by the central limit theorem, $\mathbf{Pr}(A) \geq \frac{1}{7} + o(1)$. We shall need to condition on A.

If $Z \leq \lambda n$ then all the customers arriving during $[0, 1]$ will be called *basic*. Otherwise, we split the Z arriving customers at random into $\lfloor \lambda n \rfloor$ *basic* customers plus the $Z - \lfloor \lambda n \rfloor$ *extra* customers. Let $\tilde{F}$ be the corresponding number of non-empty queues at time 1 when just the basic customers are counted. (That is, we pretend that the extra customers never arrive.) Then

$$\mathbf{Pr}(\tilde{F} \geq x | A) = \mathbf{Pr}(F \geq x | Z = \lfloor \lambda n \rfloor) \geq \frac{1}{2}.$$

Call a queue *shunned* if it is empty at time 0, none of the basic customers has it on their list of d choices (so it stays empty throughout the interval $[0,1]$ if we ignore any extra customers), and further there are no potential departures from it in this period. We shall see that with high probability there are linearly many such queues. Let S be the number of shunned queues. Let $0 < \eta < (1-\lambda)e^{-(1+d\lambda)}$. We now prove that

$$\mathbf{Pr}(S < \eta n) = e^{-\Omega(n^{\frac{1}{2}})}. \tag{9}$$

Choose $0 < \alpha < 1-\lambda$ and $0 < \beta < \alpha e^{-\lambda d}$ such that $0 < \eta < \beta/e$. Let B be the event that the number $n - F_0$ of queues empty at time 0 is at least αn. Since $\mathbf{E}[F_0] = u(1)n = \lambda n$ (as we noted earlier), by Lemma 2.5 we have $\mathbf{Pr}(\overline{B}) = e^{-\Omega(n^{\frac{1}{2}})}$. Let Q consist of the first $\lceil \alpha n \rceil \wedge (n - F_0)$ empty queues at time 0, and let R be the set of queues $j \in Q$ such that no basic customer has j in their list. Then

$$\mathbf{E}[|R| \mid B] \geq \lceil \alpha n \rceil (1 - \frac{1}{n})^{d\lfloor \lambda n \rfloor} \sim \alpha e^{-\lambda d} n.$$

By the independent bounded differences inequality (see for example (13)), $\mathbf{Pr}(|R| \leq \beta n \mid B) = e^{-\Omega(n)}$. Let R' be the set of the first $\lceil \beta n \rceil \wedge |R|$ queues in R. For each queue in R', independently there is no attempted service completion during $[0,1]$ with probability e^{-1}. Then, since $\eta < \beta/e$, conditional on $|R| \geq \beta n$ with probability $1 - e^{-\Omega(n)}$ at least ηn of the queues in R' have no attempted service completion during $[0,1]$, and so are shunned. Putting the above together yields (9), since

$$\mathbf{Pr}(S < \eta n) \leq \mathbf{Pr}(S < \eta n \mid |R| \geq \beta n) + \mathbf{Pr}(|R| < \beta n \mid B) + \mathbf{Pr}(\overline{B}).$$

Let H (for 'hit') be the number of shunned queues which are the first choice of some extra customer. Our next aim is to show (10), which says that when the event A occurs usually H is large.

Let $\mathcal{C}$ be the set of extra customers whose first choice is a shunned queue. Let $z = \lfloor \frac{2}{3}\eta(\lambda n)^{\frac{1}{2}} \rfloor$. Let $\mathcal{C}'$ consist of the first $|\mathcal{C}| \wedge z$ customers in $\mathcal{C}$. Let $\tilde{Z}$ denote a binomial random variable with parameters $\lfloor (\lambda n)^{\frac{1}{2}} \rfloor$ and η. Then

$$\mathbf{Pr}(|\mathcal{C}'| < z \mid S \geq \eta n, A) = \mathbf{Pr}(|\mathcal{C}| < z \mid S \geq \eta n, A) \leq \mathbf{Pr}(\tilde{Z} < z) = e^{-\Omega(n^{\frac{1}{2}})}.$$

Now if no shunned queue is first choice for more than two customers in $\mathcal{C}'$, then $H \geq |\mathcal{C}'| - H'$, where H' is the number of shunned queues which are first choice for two customers in $\mathcal{C}'$. But, given A and $S = s$, $\mathbf{E}[H'] \leq \binom{z}{2}/s$, and the probability that some shunned queue is first choice for more than two customers in $\mathcal{C}'$ is at most $\binom{z}{3}/s^2$. So, setting $\delta = \frac{1}{3}\eta\lambda^{\frac{1}{2}}$, it follows from the above, using Markov's inequality, that

$$\mathbf{Pr}(H \geq \delta n^{\frac{1}{2}} \mid S \geq \eta n, A) = 1 - O(n^{-\frac{1}{2}})$$

and so

$$\mathbf{Pr}(H \geq \delta n^{\frac{1}{2}} \mid A) = 1 - O(n^{-\frac{1}{2}}). \tag{10}$$

Next, we claim that $F \geq \tilde{F} + H$. To see this, start with the basic customers, and for each of the H 'hit' shunned queues throw in the first (extra) customer to hit it. With these customers we have exactly $\tilde{F}+H$ non-empty queues at time 1. If we now throw in any remaining extra customers, then by Lemma 2.3 (c) we have $F \geq \tilde{F}+H$ as claimed. Now

$$\mathbf{Pr}(F < x + \delta n^{\frac{1}{2}} \mid A) \leq \mathbf{Pr}(\tilde{F} < x \mid A) + \mathbf{Pr}(H < \delta n^{\frac{1}{2}} \mid A) \leq \frac{1}{2} + o(1).$$

Hence,

$$\mathbf{Pr}(F \geq x + \delta n^{\frac{1}{2}}) \geq \mathbf{Pr}(F \geq x + \delta n^{\frac{1}{2}} | A)\mathbf{Pr}(A) \geq \frac{1}{14} + o(1),$$

which shows that the variance of F is $\Omega(n)$. Thus we have completed the first part of the proof.

By expanding $\ell(1,Y)^d = [\ell(1) + (\ell(1,Y) - \ell(1))]^d$ we find

$$0 \leq \mathbf{E}[\ell(1,Y)^d] - \ell(1)^d = \sum_{s=2}^{d} \binom{d}{s} \mathbf{E}[(\ell(1,Y) - \ell(1))^s] \ell(1)^{d-s}.$$

If $d \geq 3$, then by the second bound in Lemma 2.6 we can upper bound

$$|\sum_{s=3}^{d} \binom{d}{s} \mathbf{E}[(\ell(1,Y) - \ell(1))^s] \ell(1)^{d-s}| = O(n^{d-3/2} \ln^3 n).$$

Also, $\mathbf{E}[(\ell(1,Y) - \ell(1))^2] = \mathrm{Var}(F) = \Omega(n)$, and so it follows that $\mathbf{E}[\ell(1,Y)^d] - \ell(1)^d = \Omega(n^{d-1})$, that is $\mathbf{E}[u(1,Y)^d] - \lambda^d = \Omega(n^{-1})$. But by (6) with $k=1$ (since $u(1) = \lambda$, as we noted before),

$$\lambda \mathbf{E}[u(1,Y)^d] - u(2) = 0,$$

and so

$$d_{\mathrm{TV}}(\mathcal{L}(Y(1)), \mathcal{L}_{\lambda,d}) \geq u(2) - \lambda^{d+1} = \lambda(\mathbf{E}[u(1,Y)^d] - \lambda^d) = \Omega\left(n^{-1}\right),$$

as required. □

Proof of Corollary 1.2. Let k be a fixed positive integer. Let X denote the random variable $Y^{(n)}(1)$, and let the random variable Z have distribution $\mathcal{L}_{\lambda,d}$. Let $m_1 = \lfloor 2 \ln \ln n / \ln d \rfloor$. Then $\mathbf{E}[Z^k \mathbb{I}_{Z \geq m_1}] = e^{-\Omega(\ln^2 n)}$. We need a similar bound for $\mathbf{E}[X^k \mathbb{I}_{X \geq m_1}]$.
By the bound (3), $\mathbf{Pr}[X \geq m_1] = e^{-\Omega(\ln n \ \ln \ln n)}$. Also, by Lemma 2.4 (b), for any positive integer m

$$\mathbf{E}[X^k \mathbb{I}_{X \geq m}] \leq \mathbf{E}[(M^{(n)})^k \mathbb{I}_{M^{(n)} \geq m}] \leq n \sum_{j \geq m} j^k \lambda^j.$$

For m sufficiently large we have $j^k \lambda^{j/2} \leq 1$ for all $j \geq m$, and then the last bound is at most

$$n \sum_{j \geq m} \lambda^{j/2} = n\lambda^{m/2}/(1-\lambda^{\frac{1}{2}}).$$

Now let $m_2 = \lceil 4 \ln n / \ln \frac{1}{\lambda} \rceil$. Then for n sufficiently large $\mathbf{E}[X^k \mathbb{I}_{X \geq m_2}] \leq e^{-\ln n + O(1)}$. Putting the above together we have

$$\mathbf{E}[X^k \mathbb{I}_{X \geq m_1}] \leq m_2^k \mathbf{Pr}[X \geq m_1] + \mathbf{E}[X^k \mathbb{I}_{X \geq m_2}] \leq e^{-\ln n + O(1)}.$$

Hence

$$\begin{aligned} |\mathbf{E}[X^k] - \mathbf{E}[Z^k]| &\leq m_1^k d_{\mathrm{TV}}(X,Z) + \mathbf{E}[X^k \mathbb{I}_{X \geq m_1}] + \mathbf{E}[Z^k \mathbb{I}_{Z \geq m_1}] \\ &\leq e^{-\ln n + O(\ln\ln n)}. \end{aligned}$$

■

3.2 Non-equilibrium case

Here we aim to prove Theorem 1.3, where we still consider a single queue length but we no longer assume that the system is in equilibrium. We shall first prove a rather general result, namely Theorem 3.3; and then deduce Theorem 3.4, from which Theorem 1.3 will follow easily. We consider the behaviour of the system from time 0, starting from general exchangeable initial conditions.

Theorem 3.3 below shows that uniformly over all $t \geq 0$ the law of a typical queue length at time t is close to $\mathcal{L}_{v_t}$, where $v_t = (v_t(k) : k = 0, 1, 2, \ldots)$ is the unique solution to the system (1) of differential equations subject to the natural initial conditions. The upper bound on the total variation distance involves three quantities δ_n, γ_n and s_n defined in terms of the initial distribution for $X_0^{(n)}$. Here δ_n concerns the total number of customers and the maximum queue length, γ_n concerns concentration of measure, and s_n concerns how close the proportions of queues of length at least k are to the equilibrium proportions.

Theorem 3.3. *Let $\theta > 1$. There exists a constant $c_2 > 0$ such that for any sufficiently large constant c_1 and any constant c_0 there is a constant $\varepsilon > 0$ for which the following holds.*

Assume that the law of $X_0^{(n)}$ is exchangeable (and thus the law of $X_t^{(n)}$ is exchangeable for all $t \geq 0$). Define

$$\delta_n = n^{-1}\mathbf{E}[\| X_0^{(n)} \|_1 \, \mathbb{I}_{\|X_0^{(n)}\|_1 > c_0 n}] + \mathbf{Pr}(\| X_0^{(n)} \|_\infty > 2\varepsilon \ln n);$$

and

$$\gamma_n = \sup_f \mathbf{Pr}(|f(X_0^{(n)}) - \mathbf{E}[f(X_0^{(n)})]| \geq (c_1 n \ln n)^{1/2})$$

where the supremum is over all non-negative functions $f \in \mathcal{F}_1$ bounded above by n.

Let $v_t^{(n)} = (v_t^{(n)}(k) : k = 0, 1, 2, \ldots)$ be the unique solution to the system (1) subject to the initial conditions $v_0^{(n)}(k) = \mathbf{E}[u(k, X_0^{(n)})]$ for $k = 0, 1, \ldots$. Also let $s_n \geq 0$ be given by $s_n^2 = \sum_{k\geq 1}(v_0^{(n)}(k) - \lambda^{(d^k-1)/(d-1)})^2\theta^k$. Then

$$\sup_{t\geq 0} d_{TV}(\mathcal{L}(X_t^{(n)}(1)), \mathcal{L}_{v_t^{(n)}}) \leq \varepsilon^{-1}\left(n^{-\varepsilon} + n^{-c_2}s_n + \delta_n + n\gamma_n\right).$$

It is known (20) that the vector $v_t^{(n)}$ satisfies (2) for each $t \geq 0$, and so $\mathcal{L}_{v_t^{(n)}}$ is well-defined. Note that the above result is uniform over all positive integers n, all exchangeable initial distributions of $X_0^{(n)}$, and all times $t \geq 0$.

We shall prove Theorem 3.3 shortly, but first let us give a corresponding result for a particular form of initial conditions, which we describe in three steps. For each n, we start with an initial vector x which is not 'too extreme' ; then we allow small independent perturbations Z_j, where we quantify 'small' by bounding the moment generating function; and finally we perform an independent uniform random permutation of the n co-ordinates, in order to ensure exchangeability. If say $x = \mathbf{0}$ and the Z_j are identically distributed (and non-negative) then we may avoid the last 'permuting' step in forming $X_0^{(n)}$, as that last step is simply to ensure that the distribution of $X_0^{(n)}$ is exchangeable.

Theorem 3.4. *For any constant $c_0 \geq 0$ there is a constant $\varepsilon > 0$ such that the following holds. Let $\beta > 0$. Suppose the initial state $X_0^{(n)}$ is obtained as follows. Take an integer n-vector x satisfying $\| x \|_1 \leq c_0 n$ and $\| x \|_\infty \leq \varepsilon \ln n$. Let the 'perturbations' $Z_1, \ldots, Z_n$ be independent random variables each taking integer values, each with variance at most β and each satisfying $\mathbf{E}[e^{Z_j/\varepsilon}] \leq \beta$. Denote $\tilde{Z} = (\tilde{Z}_1, \ldots, \tilde{Z}_n)$ where $\tilde{Z}_j = (x(j) + Z_j)^+$. Finally let $X_0^{(n)}$ be obtained from $\tilde{Z}$ by performing an independent uniform random permutation of the n co-ordinates.*

Define $v_t^{(n)}$ as in Theorem 3.3 above. Then for each $t \geq 0$, the law of $X_t^{(n)}$ is exchangeable and

$$\sup_{t \geq 0} d_{TV}(\mathcal{L}(X_t^{(n)}(1)), \mathcal{L}_{v_t^{(n)}}) = O(n^{-\varepsilon}).$$

Note that if we set x to be the zero vector above then we obtain the simpler result we stated earlier as Theorem 1.3. It remains here to prove Theorems 3.3 and 3.4.

Proof of Theorem 3.3. We aim to prove that there is a constant $c_2 > 0$ such that for any constant c_1 large enough and any constant c_0, there is a constant $\epsilon > 0$ such that we have

$$\sup_{t \geq 0} \| u_t - v_t \|_\infty = O(n^{-\epsilon} + n^{-c_2} s_n + \delta_n + n\gamma_n), \tag{11}$$

uniformly over the possible distributions for $X_0^{(n)}$. Here u_t refers to the vector $u_t(k), k = 0, 1, \ldots,$ where as before $u_t(k) = \mathbf{E}[u(k, X_t)]$.

To see that the theorem will follow from (11), consider a state x with $\| x \|_\infty \leq m$. Couple (X_t) where $X_0 = x$ with $(\tilde{X}_t)$ where $\tilde{X}_0 = \mathbf{0}$ in the way described in the previous section. By Lemma 2.3, we always have $\| X_t - \tilde{X}_t \|_\infty \leq m$. Hence always $u(k + m, X_t) \leq u(k, \tilde{X}_t)$, and so $\mathbf{E}[u(k + m, X_t)] \leq \mathbf{E}[u(k, \tilde{X}_t)] \leq u(k)$ (where the last inequality again uses Lemma 2.3). Thus, dropping the condition that $X_0 = x$, for any non-negative integers m and k, we have $\mathbf{E}[u(k + m, X_t)| \, \| X_0 \|_\infty \leq m] \leq u(k)$; and so

$$u_t(k + m) = \mathbf{E}[u(k + m, X_t)] \leq u(k) + \mathbf{Pr}(\| X_0 \|_\infty > m).$$

We shall choose a (small) constant $\epsilon > 0$ later, and let $m = \lceil 2\epsilon \ln n \rceil$. Note that $u(k) = o(n^{-1})$ if $k = \Omega(\ln n)$ by (3), and that $\mathbf{Pr}(\| X_0 \|_\infty > m) \leq \delta_n$. But now we may complete the argument as in the proof of Proposition 3.1 (a).

It remains to prove (11). First we deal with small t. We begin by showing that there is a constant $\tilde{c} > 0$ such that for each positive integer k, each integer $n \geq 2$, each $t > 0$ and each $w > 0$,

$$\mathbf{Pr}(|\ell(k, X_t) - nu_t(k)| \geq w) \leq ne^{-\frac{\tilde{c}w^2}{nt+w}} + \mathbf{Pr}(|\mathbf{E}[\ell(k, X_t)|X_0] - nu_t(k)| > w/2). \tag{12}$$

To prove this result, note that the left side above can be written as

$$\sum_{x\in\mathbb{Z}_+^n} \mathbf{Pr}(|\ell(k,X_t)-nu_t(k)|\geq w \mid X_0=x)\Pr(X_0=x).$$

Let $S=\{x\in\mathbb{Z}_+^n : |\mathbf{E}[l(k,X_t)|X_0=x]-nu_t(k)|\leq w/2\}$. Then by Lemma 2.7, there is a constant $\tilde{c}$ such that for each $x\in S$

$$\begin{aligned}
& \mathbf{Pr}(|\ell(k,X_t)-nu_t(k)|\geq w \mid X_0=x) \\
\leq\; & \mathbf{Pr}(|\ell(k,X_t)-\mathbf{E}[l(k,X_t)|X_0=x]|\geq w/2|\ X_0=x) \\
\leq\; & ne^{-\frac{\tilde{c}w^2}{nt+w}}.
\end{aligned}$$

Also, the contribution from each $x\in\mathbb{Z}_+^n\setminus S$ is at most $\mathbf{Pr}(X_0=x)$, and summing shows that the total contribution is at most

$$\mathbf{Pr}(X_0\in\mathbb{Z}_+^n\setminus S)=\mathbf{Pr}(|\mathbf{E}[\ell(k,X_t)|X_0]-nu_t(k)|>w/2).$$

This completes the proof of (12).

Now the function $\mathbf{E}[\ell(k,X_t)|X_0=x]$ is a 1-Lipschitz function of x by Lemma 2.3, and is non-negative and bounded above by n, so

$$\mathbf{Pr}(|\ell(k,X_t)-nu_t(k)|\geq w)\leq ne^{-\frac{\tilde{c}w^2}{nt+w}}+\sup_f \mathbf{Pr}(|f(X_0)-\mathbf{E}[f(X_0)]|>w/2),$$

where the supremum is over all functions $f\in\mathcal{F}_1$ such that $0\leq f(x)\leq n$ for each $x\in\mathbb{Z}_+^n$. Hence, replacing w by nw,

$$\begin{aligned}
\sup_{k\in\mathbb{N}}\mathbf{Pr}(|u(k,X_t)-u_t(k)|\geq w) \;\leq\; & ne^{-\frac{\tilde{c}w^2n}{2t}}+ne^{-\frac{\tilde{c}wn}{2}} \\
+\; & \sup_f \mathbf{Pr}(|f(X_0)-\mathbf{E}[f(X_0)]|>nw/2),
\end{aligned}$$

since $e^{-\frac{a}{b+c}}\leq e^{-\frac{a}{2b}}+e^{-\frac{a}{2c}}$ for $a,b,c>0$. Now let $c_1\geq 2/\tilde{c}$ be a sufficiently large constant and set $w=2\left(c_1n^{-1}(1+t)\ln n\right)^{\frac{1}{2}}$. Then by the last inequality, we have, uniformly over $t\geq 0$

$$\begin{aligned}
& \sup_{k\in\mathbb{N}}\mathbf{Pr}(|u(k,X_t)-u_t(k)|\geq w) \\
\leq\; & n\ e^{-(4+o(1))\ln n}+\sup_f\mathbf{Pr}(|f(X_0)-\mathbf{E}[f(X_0)]|\geq(c_1n\ln n)^{\frac{1}{2}}) \\
\leq\; & o(n^{-2})+\gamma_n;
\end{aligned}$$

that is, uniformly over $t\geq 0$

$$\sup_{k\in\mathbb{N}}\mathbf{Pr}\left(|u(k,X_t)-u_t(k)|\geq\left(4c_1n^{-1}(1+t)\ln n\right)^{\frac{1}{2}}\right)\leq o(n^{-2})+\gamma_n. \tag{13}$$

Hence, arguing as in the proof of Lemma 4.2 in (11) (Lemma 2.6 above), uniformly over $t\geq 0$

$$\sup_{k\in\mathbb{N}}\{|\mathbf{E}[u(k,X_t)^d]-u_t(k)^d|\}=O(n^{-1}(1+t)\ln n+\gamma_n).$$

We shall use this result only for $0 \leq t \leq \ln n$. By (5) we have, uniformly over $0 \leq t \leq \ln n$ and over k,

$$\frac{du_t(k)}{dt} = \lambda(u_t(k-1)^d - u_t(k)^d) - (u_t(k) - u_t(k+1)) + O(n^{-1}\ln^2 n + \gamma_n).$$

Now recall that v_t is the unique solution to (1) subject to the initial conditions $v_0(k) = u_0(k)$ for each k. Also, $2\lambda d + 2$ is a Lipschitz constant of equation (1) under the infinity norm. More precisely, $2\lambda d + 2$ is a Lipschitz constant of the operator A defined on the space of all vectors $v = (v(k) : k \in \mathbb{N})$ with components in $[0, 1]$ by

$$(Av)(k) = \lambda(v(k-1)^d - v(k)^d) - (v(k) - v(k+1)).$$

Equation (1) may be expressed succinctly in terms of A as

$$\frac{dv_t}{dt} = Av_t.$$

Thus by Gronwall's Lemma (see for instance (1)) there exists a constant $c_3 > 0$ such that uniformly over $0 \leq t \leq \ln n$

$$\| u_t - v_t \|_\infty \leq c_3(n^{-1}\ln^2 n + \gamma_n)\ln n \; e^{(2\lambda d+2)t}.$$

Let $\varepsilon_1 = \frac{1}{2d+2}$ so $0 < \varepsilon_1 < \frac{1}{2\lambda d+2}$. Then there exists $\varepsilon_2 > 0$ such that, uniformly over t with $0 \leq t \leq \varepsilon_1 \ln n$, we have

$$\| u_t - v_t \|_\infty = O(n^{-\varepsilon_2} + n\gamma_n). \tag{14}$$

For larger t we introduce the equilibrium distribution into the argument. Note that $\sup_k |u_t(k) - u(k)| \leq n^{-1} d_{\mathrm{W}}(\mathcal{L}(X_t), \mathbf{\Pi})$. Consider Lemma 2.2 with $c = \max\{c_0, \frac{2\lambda}{1-\lambda}\}$, let $\eta > 0$ be as there, and suppose $2\varepsilon \leq \eta\varepsilon_1$. By (3), there is a constant $c_4 > 0$ such that $\mathbf{Pr}(M \geq \ln\ln n/\ln d + c_4) = o(n^{-2})$. Hence by Lemma 2.2, there is a $\varepsilon_3 > 0$ such that, uniformly over $t \geq \varepsilon_1 \ln n$,

$$\| u_t - u \|_\infty = o(n^{-\varepsilon_3} + \delta_n).$$

Also, by inequality (7) there is a constant c_5 such that for each n

$$\sup_k |u(k) - \lambda^{(d^k-1)/(d-1)}| \leq c_5 n^{-1}\ln^2 n.$$

Let $s_n(t) \geq 0$ be given by $s_n^2(t) = \sum_{k\geq 1}(v_t^{(n)}(k) - \lambda^{(d^k-1)/(d-1)})^2\theta^k$; thus $s_n(0) = s_n$. We now use Theorem 2.12 in (4), which says that for some constants $\gamma_\theta > 0$ and $C_\theta < \infty$,

$$s_n(t)^2 \leq e^{-\gamma_\theta t} C_\theta s_n^2.$$

Hence we deduce that there exists a constant $c_6 > 0$ such that for all $t \geq 0$

$$\sup_k |v_t(k) - \lambda^{(d^k-1)/(d-1)}| \leq c_6 e^{-t/c_6} s_n.$$

Combining the last three inequalities, we see that for some $\varepsilon_4 > 0$, uniformly over $t \geq \varepsilon_1 \ln n$, we have

$$\| u_t - v_t \|_\infty = O(n^{-\epsilon_4} + n^{-\varepsilon_1/c_6} s_n + \delta_n).$$

This result together with (14) completes the proof of (11) (with $c_2 = \varepsilon_1/c_6$), and thus of the whole result. ∎

Proof of Theorem 3.4. Clearly X_0 is exchangeable, and thus by symmetry so is X_t for each $t \geq 0$. Let us ignore the last step in constructing the initial state, which involves the random permutation. Note that this does not affect $u_t = (u_t(k) : k = 0, 1, 2, \ldots)$, and similarly $v_0(k) = \mathbf{E}[u(k, X_0)] = \mathbf{E}[u(k, \tilde{Z})]$.

We shall choose a small constant $\varepsilon > 0$ later, and assume that $\mathbf{E}[e^{Z_j/\varepsilon}] \leq \beta$. By Markov's inequality, for each n and each $b \geq 0$,

$$\mathbf{Pr}(\max_j Z_j \geq b) \leq \sum_j e^{-b/\varepsilon}\mathbf{E}[e^{Z_j/\varepsilon}] \leq n\beta e^{-b/\varepsilon}. \tag{15}$$

Let us take $b = b(n) = \ln^2 n$, and let $A = A(n)$ be the event that $\max_j Z_j \leq b$. Then by (15) we have $\mathbf{Pr}(\bar{A}) = O(e^{-\ln^2 n})$. Given a function $f \in \mathcal{F}_1$ such that $0 \leq f(x) \leq n$ for all $x \in \mathbb{Z}_+^n$, let $\mu = \mathbf{E}[f(X_0)]$ and let $\tilde{\mu} = \mathbf{E}[f(X_0)|A]$. Then uniformly over such f

$$|\tilde{\mu} - \mu| = |\tilde{\mu} - \mathbf{E}[f(X_0)|\bar{A}]|\mathbf{Pr}(\bar{A}) \leq n\mathbf{Pr}(\bar{A}) = O(e^{-\ln^2 n}).$$

Also, the variance of Z_j given $Z_j \leq b$ is at most $var(Z_j)/\mathbf{Pr}(Z_j \leq b)$, which is at most 2β for n sufficiently large that $\mathbf{Pr}(Z_j \leq b) \geq \frac{1}{2}$. Thus given $Z_j \leq b$, the variance of $\tilde{Z}_j$ is also at most 2β, and further $\tilde{Z}_j$ has range at most $\frac{3}{2}b$, assuming that $2\varepsilon \leq \ln n$. Hence, by Lemma 2.8, for each sufficiently large integer n and each $w > 0$,

$$\mathbf{Pr}(|f(X_0) - \tilde{\mu}| \geq w|A) \leq 2e^{-\frac{w^2}{4n\beta + bw}},$$

and so the quantity $n\gamma_n$ in Theorem 3.3 is $O(n^{-1})$ provided the constant c_1 is large enough.

Note that $\mathbf{E}[Z_j] \leq \varepsilon\beta$ for each j. We shall choose $0 < \varepsilon \leq 1$ below. Let $c = \max\{c_0 + \beta + 1, \frac{1}{1-\lambda}\}$. Then $\mathbf{E}[\| X_0 \|_1 \mathbf{1}_{\|X_0\|_1 > cn}] = o(1)$ since $c > c_0 + \beta$. Now assume that $0 < \varepsilon \leq 1$ is such that the bound in Theorem 3.3 holds with this value of ε and with $c_0 = c$. Then by (15), with $b = 2\varepsilon \ln n$,

$$\mathbf{Pr}(\| X_0 \|_\infty > 2\varepsilon \ln n) \leq n\beta e^{-2\ln n} = O(n^{-1}).$$

Thus the quantity δ_n in Theorem 3.3 is $O(n^{-1})$.

It remains to upper bound the term s_n. Note that

$$s_n \leq \sum_{k \geq 1} \theta^k \left((v_0^{(n)}(k)\right)^2 + \sum_{k \geq 1} \theta^k \lambda^{2(d^k - 1)/(d-1)}.$$

The second of these terms certainly takes a finite value independent of n, so it suffices to ensure that the first term is $O(n^{c_2/2})$ where the constant $c_2 > 0$ is from Theorem 3.3. Now $v_0^{(n)}(k) \leq u(\lfloor k/2 \rfloor, x) + \max_j \Pr(Z_j \geq k/2)$, and so $v_0^{(n)}(k)^2 \leq 2u(\lfloor k/2 \rfloor, x)^2 + 2\max_j \Pr(Z_j \geq k/2)^2$. Hence

$$\begin{aligned} \sum_{k \geq 1} \theta^k \left((v_0^{(n)}(k)\right)^2 &\leq 2\sum_{k \geq 1} \theta^k (u(\lfloor k/2 \rfloor, x))^2 \\ &+ 2\sum_{k \geq 1} \theta^k (\max_j \mathbf{Pr}(Z_j \geq k/2))^2. \end{aligned}$$

Using Markov's inequality again, $\mathbf{Pr}(Z_j \geq k/2) \leq \beta e^{-k/2\varepsilon}$, and so the second term is bounded by a constant uniformly in n, provided $\theta < e^{1/\varepsilon}$. This is certainly true if $\varepsilon > 0$ is small enough. As for the first term, if we let $k_0 = 2\varepsilon \ln n + 2$, we have

$$\begin{aligned} \sum_{k\geq 1} \theta^k (u(\lfloor k/2 \rfloor, x))^2 &\leq \sum_{1\leq k\leq k_0} \theta^k (u(\lfloor k/2 \rfloor, x))^2 \\ &\leq k_0 \, \theta^{k_0} \quad \leq n^{c_2/2} \end{aligned}$$

if ε is sufficiently small. Combining the above estimates, we obtain the required bound on s_n. ∎

4 Asymptotic chaos

4.1 Equilibrium case

In this subsection we prove Theorem 1.4 concerning chaoticity of the queue-lengths process, where the system is stationary. This result quantifies the chaoticity in terms of the total variation distance between the joint law of the queue lengths and the corresponding product laws. In the proof we first bound another natural measure of these distances, following the treatment in (18) – see (19) and (20) below.

Proof of Theorem 1.4. By Lemma 2.5, there exists a constant $c > 0$ such that

$$\sup_{f\in\mathcal{F}_1} \mathbf{Pr}\left(|f(Y) - \mathbf{E}f(Y)| \geq cn^{1/2}\ln n\right) = o(n^{-2}).$$

Let $\mathcal{F}_1'$ be the set of functions $f \in \mathcal{F}_1$ with $\| f \|_\infty \leq n$. By the last result, uniformly over positive integers a,

$$\sup_{f_1,\ldots,f_a\in\mathcal{F}_1'} \mathbf{E}\left[\prod_{s=1}^{a} |f_s(Y) - \mathbf{E}[f_s(Y)]|\right] \leq (cn^{1/2}\ln n)^a + o(an^{a-2}). \tag{16}$$

Let $\mathcal{G}_1$ denote the set of measurable real-valued functions ϕ on $\mathbb{Z}_+$ with $\| \phi \|_\infty \leq 1$. For any measurable real-valued function ϕ on $\mathbb{Z}_+$ let $\bar{\phi}$ be the function on $\mathbb{Z}_+^n$ defined by setting $\bar{\phi}(y) = \frac{1}{n}\sum_{i=1}^n \phi(y_i)$ for $y = (y_1,\ldots,y_n) \in \mathbb{Z}_+^n$. Observe that

$$\bar{\phi}(y) = \sum_{k=0}^{\infty} \phi(k)(u(k,y) - u(k+1,y)) = \phi(0) + \sum_{k=1}^{\infty}(\phi(k) - \phi(k-1))u(k,y).$$

Observe also that since the distribution of Y is exchangeable we have $\mathbf{E}[\bar{\phi}(Y)] = \mathbf{E}[\phi(Y(1))]$: let us call this mean value $\hat{\phi}$.

Let $f(y) = (n/2)\bar{\phi}(y) = \frac{1}{2}\sum_{i=1}^n \phi(y_i)$. It is easy to see that if $\phi \in \mathcal{G}_1$ then $f \in \mathcal{F}_1'$. Hence by (16), uniformly over positive integers a,

$$\sup_{\phi_1,\ldots,\phi_a\in\mathcal{G}_1} \mathbf{E}\left[\prod_{s=1}^{a} |\bar{\phi}_s(Y) - \hat{\phi}_s|\right] \leq (2cn^{-1/2}\ln n)^a + o(a2^a n^{-2}) \leq (3cn^{-1/2}\ln n)^a \tag{17}$$

for n sufficiently large. But, by writing $\bar{\phi}_s(Y)$ as $(\bar{\phi}_s(Y) - \hat{\phi}_s) + \hat{\phi}_s$, we see that

$$\mathbf{E}\left[\prod_{s=1}^{a} \bar{\phi}_s(Y)\right] - \prod_{s=1}^{a} \hat{\phi}_s$$

can be expanded as

$$\sum_{A\subseteq\{1,\ldots,a\},|A|\geq 2} \mathbf{E}\left[\prod_{s\in A}(\bar{\phi_s}(Y)-\hat{\phi}_s)\right]\prod_{s\in\{1,\ldots,a\}\setminus A}\hat{\phi}_s.$$

Hence by (17), uniformly over all positive integers $r \leq n^{\frac{1}{2}}/\ln n$

$$\begin{aligned}
&\sup_{\phi_1,\ldots,\phi_r\in\mathcal{G}_1} |\mathbf{E}\left[\prod_{s=1}^{r}\bar{\phi_s}(Y)\right]-\prod_{s=1}^{r}\hat{\phi}_s| \\
\leq\ & \sum_{A\subseteq\{1,\ldots,r\},|A|\geq 2}\mathbf{E}\left[\prod_{s\in A}|\bar{\phi_s}(Y)-\hat{\phi}_s|\right] \\
\leq\ & \sum_{a=2}^{r}\binom{r}{a}(3cn^{-\frac{1}{2}}\ln n)^a \leq \sum_{a=2}^{\infty}(\frac{3ec}{a}\ rn^{-\frac{1}{2}}\ln n)^a \\
=\ & O(r^2n^{-1}\ln^2 n).
\end{aligned} \tag{18}$$

Now

$$\mathbf{E}\left[\prod_{s=1}^{r}\bar{\phi_s}(Y)\right] = n^{-r}\mathbf{E}\left[\prod_{s=1}^{r}\sum_{j=1}^{n}\phi_s(Y(j))\right] = \mathbf{E}\left[\prod_{s=1}^{r}\phi_s(Y(s))\right]+O(r^2n^{-1})$$

uniformly over r and all r-tuples $\phi_1,\ldots,\phi_r\in\mathcal{G}_1$ since when we expand the middle expression there are at most r^2n^{r-1} terms for which the values of j are not all distinct. Hence from (18), uniformly over all positive integers $r \leq n$,

$$\sup_{\phi_1,\ldots,\phi_r\in\mathcal{G}_1}|\mathbf{E}[\prod_{s=1}^{r}\phi_s(Y(s))]-\prod_{s=1}^{r}\hat{\phi}_s| = O(r^2n^{-1}\ln^2 n). \tag{19}$$

(For $r > n^{\frac{1}{2}}/\ln n$ the result is trivial.) Further, by Proposition 3.1(a),

$$\sup_{\phi\in\mathcal{G}_1}|\hat{\phi}-\mathbf{E}_{\mathcal{L}_{\lambda,d}}(\phi)| = O(n^{-1}\ln^2 n\ \ \ln\ln n),$$

where $\mathbf{E}_{\mathcal{L}_{\lambda,d}}(\phi)$ denotes the expectation with respect to the limiting measure $\mathcal{L}_{\lambda,d}$. Thus

$$\sup_{\phi_1,\ldots,\phi_r\in\mathcal{G}_1}|\prod_{s=1}^{r}\hat{\phi}_s-\prod_{s=1}^{r}\mathbf{E}_{\mathcal{L}_{\lambda,d}}(\phi_s)| = O(rn^{-1}\ln^2 n\ \ \ln\ln n)$$

uniformly over all positive integers r (again noting that the result is trivial for large r). It follows that, uniformly over all positive integers $r \leq n$

$$\sup_{\phi_1,\ldots,\phi_r\in\mathcal{G}_1}|\mathbf{E}[\prod_{s=1}^{r}\phi_s(Y(s))]-\prod_{s=1}^{r}\mathbf{E}_{\mathcal{L}_{\lambda,d}}(\phi_s)| = O(n^{-1}\ln^2 n\ (r^2+r\ln\ln n)). \tag{20}$$

Following Sznitman (18), the vector $Y=(Y(1),\ldots,Y(n))$ is chaotic (in total variation) since the left hand side of (19) tends to 0 as $n\to\infty$ for each fixed positive integer r. Thus (19) quantifies the chaoticity of the equilibrium queue lengths Y in terms of this definition. Similarly, the inequality (20) quantifies Y being $\mathcal{L}_{\lambda,d}$-chaotic. Let us note that, up to factors logarithmic in n, the bound in (20) is of the same order as the time-dependent bound of Theorem 4.1 in (5) obtained for a related class of models.

The results (19) and (20) yield bounds on the total variation distance between the joint law of $Y(1),\ldots,Y(r)$ and the product law $\mathcal{L}(Y(1))^{\otimes r}$, and between the joint law of $Y(1),\ldots,Y(r)$ and the product law $\mathcal{L}_{\lambda,d}^{\otimes r}$ respectively, as follows. [In general, even for random variables $Y(s)$ taking values in $\mathbb{Z}_+$ and $r=2$, it does not follow that if the left hand side of (19) tends to 0 as $n\to\infty$ then the total variation distance between the joint law of $Y(1),\ldots,Y(s)$ and the product law $\mathcal{L}(Y(1))^{\otimes r}$ tends to 0.] Putting ϕ_s as the indicator of the set $\{k_s\}$, we obtain, uniformly over positive integers $r\le n$

$$\sup_{k_1,\ldots,k_r}\left|\mathbf{Pr}(\wedge_{s=1}^r(Y(s)=k_s))-\prod_{s=1}^r\mathbf{Pr}(Y(s)=k_s)\right|=O(r^2n^{-1}\ln^2 n) \tag{21}$$

and

$$\sup_{k_1,\ldots,k_r}\left|\mathbf{Pr}(\wedge_{s=1}^r(Y(s)=k_s))-\prod_{s=1}^r\mathcal{L}_{\lambda,d}(\{k_s\})\right|=O(n^{-1}\ln^2 n\ (r^2+r\ln\ln n)). \tag{22}$$

But by (3), there exists a constant $c>0$ such that

$$\mathbf{Pr}(\max_j Y(j)>\ln\ln n/\ln d+c)=O(n^{-1}).$$

Hence in total variation distance, for any fixed $\epsilon>0$, the joint law of $Y(1),\ldots,Y(r)$ differs from the product law $\mathcal{L}(Y(1))^{\otimes r}$ by at most $O(n^{-1}\ln^2 n\ ((1+\epsilon)\ln\ln n/\ln d)^r)$, and by at most $O(n^{-1}\ln^2 n\ ((1+\epsilon)\ln\ln n/\ln d)^{r+1})$ from the law $\mathcal{L}_{\lambda,d}^{\otimes r}$. ∎

4.2 Non-equilibrium case

We now no longer assume that the system is in equilibrium, and show that under quite general exchangeable initial conditions, chaotic behaviour occurs in the system, uniformly for all times. We need to prove Theorem 1.5. We first state two general results, Theorem 4.1 and 4.2, and note that Theorem 1.5 will follow easily from the latter. We then prove Theorem 4.1 and deduce Theorem 4.2.

Given an n-queue process $(X_t^{(n)})$ and a positive integer $r\le n$, let $\mathcal{L}_t^{(n,r)}$ denote the joint law of $X_t^{(n)}(1),\ldots,X_t^{(n)}(r)$, and let $\tilde{\mathcal{L}}_t^{(n,r)}$ denote the product law of r independent copies of $X_t^{(n)}(1)$. The following result shows that, as long as initially there are not too many customers in the system, the maximum queue is not too long and the system has sufficient concentration of measure, there will be chaotic behaviour uniformly for all times.

Theorem 4.1. *For each sufficiently large constant c_1 and each constant c_0, there is a constant c such that the following holds.*

Let n be a positive integer. Assume that the law of $X_0^{(n)}$ is exchangeable (which implies that the law of $X_t^{(n)}$ is exchangeable for all $t\ge 0$). Let $m\ge 0$, let

$$\delta_{n,m}=\mathbf{Pr}(\|X_0^{(n)}\|_1>c_0n)+\mathbf{Pr}(\|X_0^{(n)}\|_\infty>m),$$

and let

$$\gamma_n = \sup_f \mathbf{Pr}(|f(X_0^{(n)}) - \mathbf{E}[f(X_0^{(n)})]| \geq (c_1 n \ln n)^{1/2}),$$

where the supremum is over all non-negative functions $f \in \mathcal{F}_1$ bounded above by n. Then for each positive integer $r \leq n$ we have

$$\sup_{t\geq 0} d_{TV}\left(\mathcal{L}_t^{(n,r)}, \tilde{\mathcal{L}}_t^{(n,r)}\right) \leq c[(n^{-1}\ln n)(m+\ln n) + \gamma_n](m + 2\ln\ln n)^r + (r+1)\delta_{n,m}.$$

We shall see later that it is straightforward to deduce Theorem 1.4 from Theorem 4.1. Theorem 4.2 below is a straightforward consequence of Theorem 4.1 and shows that, in particular, there is chaotic behaviour uniformly for all times $t \geq 0$, when the initial state is obtained by perturbing a "nice" queue-lengths vector by a set of 'small' independent random variables.

Theorem 4.2. *Let $c_0 > 0$, $s_0 > 0$ and $\alpha > 0$ be any constants. Then there is a constant $c > 0$ such that the following holds.*

Suppose that the initial state $X_0^{(n)}$ is obtained as follows. Take $x \in \mathbb{Z}_+^n$ with $\| x \|_1 \leq c_0 n$ and let $m = \| x \|_\infty$. Let the 'perturbations' $Z_1, \ldots, Z_n$ be independent random variables each taking integer values, each with variance at most α and each satisfying $\mathbf{E}[e^{s_0 Z_j}] \leq \alpha$. Let $\tilde{Z} = (\tilde{Z}_1, \ldots, \tilde{Z}_n)$ where $\tilde{Z}_j = (x(j) + Z_j)^+$. Finally let $X_0^{(n)}$ be obtained from $\tilde{Z}$ by performing an independent uniform random permutation of the n co-ordinates.

Then for each $t \geq 0$, $X_t^{(n)}$ is exchangeable, and for each positive integer $r \leq n$ we have we have

$$\sup_{t\geq 0} d_{TV}\left(\mathcal{L}_t^{(n,r)}, \tilde{\mathcal{L}}_t^{(n,r)}\right) \quad \leq \quad cn^{-1}(m + 2\ln n)^{r+2}.$$

Note that Theorem 1.5 will follow from this last result on taking x as the zero vector. It remains to prove Theorems 4.1 and 4.2.

Proof of Theorem 4.1. First we deal with small t. We start by arguing in a similar way to the proofs of Theorems 3.3 and 1.4. As in the inequality (13), there exists a constant $c_1 > 0$ such that uniformly over $t \geq 0$, over n, and over all non-negative functions $f \in \mathcal{F}_1$ bounded above by n, we have

$$\mathbf{Pr}(|n^{-1}f(X_t) - \mathbf{E}[n^{-1}f(X_t)]| \geq (c_1(1+t)n^{-1}\ln n)^{1/2}) \leq o(n^{-2}) + \gamma_n.$$

With notation as in the proof of Theorem 1.4, since the distribution of $X_t^{(n)}$ is exchangeable, we have $\mathbf{E}[\bar{\phi}(X_t^{(n)})] = \mathbf{E}[\phi(X_t^{(n)}(1))]$: let us call this mean value $(\hat{\phi})_t$. Uniformly over integers a and r with $2 \leq a \leq r$,

$$\sup_{\phi_1,\ldots,\phi_a \in \mathcal{G}_1} \left[\mathbf{E}\prod_{s=1}^{a} |\bar{\phi}_s(X_t^{(n)}) - (\hat{\phi}_s)_t|\right] \leq O((1+t)n^{-1}\ln n) + r\gamma_n,$$

Using the above bound and following the steps in the proof of Theorem 1.4, we find that uniformly over r

$$\sup_{\phi_1,\ldots,\phi_r \in \mathcal{G}_1} \left|\mathbf{E}\left[\prod_{s=1}^{r} \phi_s(X_t^{(n)}(s))\right] - \prod_{s=1}^{r}(\hat{\phi}_s)_t\right|$$

$$= O\left(r^2[(1+t)n^{-1}\ln n + r\gamma_n]\right).$$

Let $M^{(n)}$ denote the maximum queue length in equilibrium. Then, arguing as in the proof of Theorem 3.3, given that $\| X_0^{(n)} \|_\infty \le m$, we have by Lemma 2.3 and (3) that

$$\mathbf{Pr}(\| X_t^{(n)} \|_\infty \ge m+k) \le \mathbf{Pr}(M^{(n)} \ge k) = O(n^{-1})$$

if $k = \lfloor \ln\ln n/\ln d \rfloor + c_2$ for some suitable constant c_2. Hence, uniformly over $t \ge 0$, over n, over positive integers $2 \le a \le r$, over $m \ge 0$ and over exchangeable distributions for $X_0^{(n)}$, we have

$$d_{\mathrm{TV}}\left(\mathcal{L}_t^{(n,r)}, \tilde{\mathcal{L}}_t^{(n,r)}\right) = O\left([(1+t)n^{-1}\ln n + \gamma_n](m + 2\ln\ln n)^r\right) + \delta_{n,m}. \tag{23}$$

The above takes care of all $t = O(m + \ln n)$. More precisely, we use the bound (23) for $0 \le t \le \eta^{-1}(m + 2\ln n)$, where $\eta > 0$ is as in Lemma 2.1 with $c = c_0$.

To handle larger t we introduce the equilibrium distribution. Let $Y^{(n)}$ denote a queue-lengths vector in equilibrium. Denote by $\mathcal{M}^{(n,r)}$ the joint law of $Y^{(n)}(1), \ldots, Y^{(n)}(r)$; and denote by $\tilde{\mathcal{M}}^{(n,r)}$ the product law of r independent copies of $Y^{(n)}(1)$. Then for $t \ge \eta^{-1}(m + 2\ln n)$ we have by Lemma 2.1 and Theorem 1.4 that

$$\begin{aligned}
& d_{\mathrm{TV}}(\mathcal{L}_t^{n,r}, \tilde{\mathcal{L}}_t^{n,r}) \\
\le\; & d_{\mathrm{TV}}(\mathcal{L}_t^{n,r}, \mathcal{M}_t^{n,r}) + d_{\mathrm{TV}}(\mathcal{M}_t^{n,r}, \tilde{\mathcal{M}}_t^{n,r}) + d_{\mathrm{TV}}(\tilde{\mathcal{M}}_t^{n,r}, \tilde{\mathcal{L}}_t^{n,r}) \\
\le\; & d_{\mathrm{TV}}(\mathcal{M}_t^{n,r}, \tilde{\mathcal{M}}_t^{n,r}) + (r+1) d_{\mathrm{TV}}(\mathcal{L}(X_t), \mathcal{L}(Y)) \\
\le\; & O(n^{-1}\ln^2 n (2\ln\ln n)^r) + (r+1)\delta_{n,m}.
\end{aligned}$$

This completes the proof of Theorem 4.1. ∎

Let us indicate briefly how to deduce Theorem 1.4 from Theorem 4.1. By Lemma 2.4 and equation (3), if we choose a constant $c_0 > \frac{\lambda}{1-\lambda}$ and $m = \ln\ln n/\ln d + \tilde{c}$ for a sufficiently large constant $\tilde{c}$ then the term $\delta_{n,m}$ in Theorem 4.1 is $O(n^{-1})$. Further, by Lemma 2.5, if the constant c_1 is sufficiently large then the term γ_n is also $O(n^{-1})$, and Theorem 1.4 follows.

Proof of Theorem 4.2. We fix $s_0 > 0$ and $\alpha > 0$ such that $\mathbf{E}[e^{s_0 Z_j}] \le \alpha$ for each j. Let $\beta = \ln(1+\alpha)$, so that $\beta > 0$. Note that

$$\| X_0^{(n)} \|_1 = \| \tilde{Z} \|_1 \le \| x \|_1 + \sum_{j=1}^n Z_j^+$$

and

$$\mathbf{E}[e^{s_0 Z_j^+}] \le 1 + \mathbf{E}[e^{s_0 Z_j}] \le (1+\alpha) = e^\beta.$$

Thus

$$\begin{aligned}
\mathbf{Pr}\left(\| X_0 \|_1 \ge (c_0 + \frac{2\beta}{s_0})n\right) \;&\le\; \mathbf{Pr}(\sum_{j=1}^n Z_j^+ \ge \frac{2\beta}{s_0} n) \\
&\le\; e^{-2\beta n}\mathbf{E}[e^{s_0 \sum_{j=1}^n Z_j^+}] \\
&=\; e^{-2\beta n}\prod_{j=1}^n \mathbf{E}[e^{s_0 Z_j^+}] \;\le\; e^{-\beta n}.
\end{aligned}$$

Also,

$$\begin{aligned}\mathbf{Pr}\left(\| X_0 \|_\infty > (m + \frac{2}{s_0}\ln n)\right) &\leq \sum_{j=1}^{n} \mathbf{Pr}(Z_j > \frac{2}{s_0}\ln n) \\ &\leq e^{-2\ln n} \sum_{j=1}^{n} \mathbf{E}[e^{s_0 Z_j}] = O(n^{-1}).\end{aligned}$$

Further, arguing as in the proof of Theorem 3.4, by Lemma 2.8 the constant c_1 can be chosen large enough so that the term γ_n in Theorem 4.1 is $O(n^{-1})$. Hence the result follows from Theorem 4.1 with this value of c_1, with c_0 replaced by $c_0 + \frac{2\beta}{s_0}$ and with m replaced by $m + \frac{2}{s_0}\ln n$. We have

$$\begin{aligned}\sup_{t\geq 0} d_{\mathrm{TV}}\left(\mathcal{L}_t^{(n,r)}, \tilde{\mathcal{L}}_t^{(n,r)}\right) &\leq cn^{-1}\ln n(m+\ln n)(m+\ln n+2\ln\ln n)^r \\ &\leq cn^{-1}(m+2\ln n)^{r+2}.\end{aligned}$$

∎

Acknowledgement We would like to thank the referee for detailed and helpful comments.

References

[1] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence.* Wiley, New York (1986). MR0838085

[2] C. Graham. Kinetic limits for large communication networks. In *Modelling in Applied Sciences* (2000), N. Bellomo and M. Pulvirenti (eds.), 317 – 370. Birkhäuser, Basel. MR1763158

[3] C. Graham. Chaoticity on path space for a queueing network with selection of the shortest queue among several. *J. Appl. Probab.* **37** (2000) 198 – 201. MR1761670

[4] C. Graham. Functional central limit theorems for a large network in which customers join the shortest of several queues, 2004, arXiv math.PR/0403538. MR2105045

[5] C. Graham and S. Méléard. Chaos hypothesis for a system interacting through shared resources. *Probab. Theor. Relat. Fields* **100** (1994) 157 – 173. MR1296426

[6] A. Guionnet and B. Zegarlinski. Lectures on logarithmic Sobolev inequalities. Séminaire de Probabilités XXXVI. *Lect. Notes Math.* **1801** (2002), Springer-Verlag, Berlin. MR1971582

[7] M. Kac. Foundation of kinetic theory. *Proc. 3rd Berkeley Symp. on Math. Stat. and Probab.* **3** (1956) 171–197. Univ. of Calif. Press. MR0084985

[8] M. Ledoux. *The Concentration of Measure Phenomenon.* AMS, Providence, Rhode Island (2001). MR1849347

[9] M.J. Luczak. A quantitative law of large numbers via exponential martingales. In *Stochastic Inequalities and Applications* (2003), E. Giné, C. Houdré and D. Nualart (eds.), Birkhäuser Series Progress in Probability **56** 93 – 111. Birkhäuser, Basel. MR2073429

[10] M.J. Luczak and C. McDiarmid. On the power of two choices: balls and bins in continuous time. *Ann. Appl. Probab.* **15** (2005) 1733 – 1764. MR2152243

[11] M.J. Luczak and C. McDiarmid. On the maximum queue length in the supermarket model. *Ann. Probab.* **34** (2006) 493 – 527. MR2223949

[12] M.J. Luczak and J.R. Norris. Strong approximation for the supermarket model. *Ann. Appl. Probab.* **15** (2005) 2038 – 2061. MR2152252

[13] C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics* (1989), J. Siemons (ed.), 148 – 188. London Mathematical Society Lecture Notes Series 141. CUP, Cambridge. MR1036755

[14] C. McDiarmid. Concentration. In *Probabilistic Methods for Algorithmic Discrete Mathematics* (1998), M. Habib, C. McDiarmid, J. Ramirez and B. Reed (eds.), 195–248. Springer-Verlag, Berlin. MR1678578

[15] M. Mitzenmacher. Load balancing and density dependent jump Markov processes. *Proc. 37th Ann. Symp. Found. Comp. Sci.* (1996) 213 – 222. MR1450619

[16] M. Mitzenmacher. The power of two choices in randomized load-balancing. Ph.D. thesis. Berkeley, 1996, http://www.eecs.harvard.edu/ michaelm/.

[17] M. Mitzenmacher, A.W. Richa and R. Sitaraman. The power of two random choices: a survey of techniques and results. In *Handbook of Randomized Computing* Vol. **1** (2001), S. Rajasekaran, P.M. Pardalos, J.H. Reif and J.D.P. Rolim (eds.) 255–312. Kluwer Academic Publishers, Dordrecht. MR1966907

[18] A. Sznitman. Topics in propagation of chaos. In *École d'Été de Probabilités de Saint-Flour XIX - 1989* (1989) 165 – 251. Springer-Verlag, Berlin. MR1108185

[19] C. Villani. Optimal transportation, dissipative PDE's and functional inequalities. *Lect. Notes Math.* **1813**, Springer-Verlag, Berlin, to appear.

[20] N.D. Vvedenskaya, R.L. Dobrushin and F.I. Karpelevich. Queueing system with selection of the shortest of two queues: an asymptotic approach. *Prob. Inform. Transm.* **32** (1996) 15 – 27. MR1384927