Adv. Appl. Prob. **46,** 59–75 (2014)
Printed in Northern Ireland

STABILITY OF MULTI-DIMENSIONAL BIRTH-AND-DEATH PROCESSES WITH STATE-DEPENDENT 0-HOMOGENEOUS JUMPS

MATTHIEU JONCKHEERE,* *University of Buenos Aires*
SEVA SHNEER,** *Heriot-Watt University*

Abstract

We study the conditions for positive recurrence and transience of multi-dimensional birth-and-death processes describing the evolution of a large class of stochastic systems, a typical example being the randomly varying number of flow-level transfers in a telecommunication wire-line or wireless network. First, using an associated deterministic dynamical system, we provide a generic method to construct a Lyapunov function when the drift is a smooth function on $\mathbb{R}^N$. This approach gives an elementary and direct proof of ergodicity. We also provide instability conditions. Our main contribution consists of showing how discontinuous drifts change the nature of the stability conditions and of providing generic sufficient stability conditions having a simple geometric interpretation. These conditions turn out to be necessary (outside a negligible set of the parameter space) for piecewise constant drifts in dimension two.

Keywords: Birth-and-death process; positive recurrence; transience; fluid limit

2010 Mathematics Subject Classification: Primary 93E15
Secondary 60K25

1. Introduction

We study the stochastic stability of multi-dimensional birth-and-death processes $X = (X_1, \dots, X_N)$ on $\mathbb{Z}_+^N$ (N being an integer greater than 1) with state-dependent birth (also called arrival) and death (also called service) rates (respectively $\lambda(x) = (\lambda_i(x))_{i=1,\dots,N}$ and $\phi(x) = (\phi_i(x))_{i=1,\dots,N}$, with $x = (x_1, \dots, x_N)$) being 0-homogeneous functions, i.e. $\lambda(\alpha x) = \lambda(x)$ and $\phi(\alpha x) = \phi(x)$ for any $\alpha > 0$ and for any $x \in \mathbb{R}_+^N$.

The main motivation for this research stems from the study of queueing networks. In the last two decades an enormous amount of literature has emerged, aiming at the finest possible description of the condition under which various queueing systems are stable, see e.g. [15], [18], [20], and [22]. For a very large class of cases, the dynamics of these queueing (or more generally bandwidth-sharing) networks can be described with state-dependent service rates depending on the proportion of customers of each class present in the network; hence, satisfying the 0-homogeneity assumption stated previously. For instance, data wire-line communication networks can be represented (at a sufficiently large time scale) as processor-sharing networks, with processing rates that depend on the proportion of users at each node of the network [1]. In wireless networks the service rates depend only asymptotically on the proportion of customers,

Received 8 February 2011; revision received 3 April 2013.
* Postal address: CONICET, Mathematics Department, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires (1428) Pabellón I, Ciudad Universitaria, Buenos Aires, Argentina.
** Postal address: Department of AMS, Heriot-Watt University, Edinburgh EH14 4AS, UK.
Email address: v.shneer@hw.ac.uk

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

but this requires the same analysis, as far as only stability is concerned. The 0-homogeneity assumption is also relevant for describing load-balancing schemes between a set of servers or computer systems. In contrast to the previous examples, very few results have been discussed in the literature concerning simple schemes such as joining the shortest queue when the death rates are not constant. This is usually considered to be a very difficult issue. (For the constant case, see [9].)

Good overviews of stochastic-stability methods and their applications may be found in, for example, [3], [11], [17], and [20]. Let us discuss briefly some of the main ideas in this area.

A general framework for analyzing stochastic stability consists of applying the Foster–Lyapunov criteria, which are based on finding a suitable test function, usually referred to as a Lyapunov function, having a negative (when stability is concerned) or positive (in case instability is proven) mean drift outside of a compact subset of the state space [8], [18], [25]. When further restrictive assumptions are made on death rates, an appropriate Lyapunov function has been found in many cases. This is, for example, possible for rates being the solution of specific optimization problems [1], or for small dimensions (two and three) when the rates of the process are constant on sub-faces of the orthant [8]. For more complex systems, however, finding a Lyapunov function can be a formidable task and there is no general method for constructing such a function.

An alternative tool for deriving stability conditions is to establish that the system of interest is stochastically comparable to another system that is easier to analyze. This approach was first used in the multi-class queueing context by Rao and Ephremides [19] and Szpankowski [23], and later refined by Szpankowski [24], to characterize the stability of buffered random access systems. It was later generalized to birth-and-death processes with state-dependent transitions with fixed birth rates and decreasing death rates with uniform limits in [2]. These specific assumptions are, however, not satisfied for various cases of interest.

Finally, many stability results are obtained using the so-called ODE (ordinary differential equation) methods. A powerful exposition of these ideas applied to controlled random walks can be found in [17, Chapter 10] and in [10]. The use of ODE methods is often coupled with the analysis of fluid limits or fluid models (fluid models being a set of deterministic trajectories including possible limits): first the convergence of (sub-sequences of) a scaled version of the process towards an element of the fluid model is proven; then (under restrictive conditions) the stability of the fluid model is proven to imply the stability (positive recurrence in our case) of the stochastic process. Stability conditions for a wide class of multi-class queueing networks with work-conserving service disciplines (see [5] and [6]) have been derived using these steps. Often, it is difficult to prove directly the convergence towards a fluid limit.

When the drift $\delta = \lambda - \phi$ of the system is Lipschitz-continuous, the state of the stochastic network (under an appropriate space–time scaling) converges to a deterministic system whose evolution is represented by the differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) = \delta(x(t)), \tag{1}$$

and we can prove that $x(t)$ is the state of the fluid limit at time t.

It turns out that such a convergence does not hold in general when the drift cannot be extended to a continuous function on $\mathbb{R}_+^N$. When the drift vector field is discontinuous, the trajectories of a fluid-equivalent system enter sliding modes and the differential equation (1) has to be replaced by a new dynamical system defined piecewise by differential equations $(\mathrm{d}/\mathrm{d}t)x(t) = \tilde{\delta}(x(t))$, where $\tilde{\delta}$ is a convex combination of drifts of points in the neighborhood of the discontinuity.

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

In terms of stability conditions, such a phenomenon was already emphasized in [8] where $\tilde{\delta}$ was called the 'second vector-field' (see also [10] and [17]). Unfortunately, $\tilde{\delta}$ is difficult to compute in many cases (it depends crucially on the statistical assumptions made) and has not been characterized in general.

In this paper we restrict our attention to the study of birth-and-death processes representing networks without internal routeing but having state-dependent arrival and service rates. Our contribution is three-fold.

First, we aim at unifying various results appearing recurrently in the literature concerning the case of processes with continuous drifts. To this end, we present a generic way to construct a Lyapunov function using the deterministic differential equations driving the fluid limits dynamics when the drift is continuous. The advantage of finding a Lyapunov function explicitly is that it potentially gives much more precise information on the nature of the convergence of the process towards its stationary regime [18]. It also provides a simple understanding of the meaning of such fluid limits for obtaining the stochastic stability. We also show that in the case of drifts that are conservative vector fields, the complexity of the problem is reduced considerably. We complement these results with instability counterparts.

Second, we aim at deriving general sufficient conditions for stability in the case of discontinuous drifts. To the best of the authors' knowledge, these conditions have not been obtained previously. The case of discontinuous drifts is shown to be (both in the existing literature and by a few simple examples in our paper) a very challenging theoretical problem and generally extremely demanding in terms of computation time. For this case we provide sufficient conditions that have a natural geometric interpretation and are shown to be useful in some important examples.

Third, we use these conditions to get a sharp geometric characterization of the stability set in the case of piecewise constant drifts in dimension two. We provide, in particular, an algorithm allowing us to determine whether the process is stable or not, for all fixed birth rates outside a set of dimension one.

The paper is organized as follows. In Section 2 we describe the model in detail and discuss the methodology used in the subsequent analysis. In Section 3 we examine the case when the drift vector field can be extended to a continuous function. Section 4 is devoted to deriving sufficient stability conditions in the case of discontinuous drifts. We start by showing that stability conditions and the fluid limits are in this case very cumbersome and then proceed to presenting our approach in a generic scenario. In Section 5 we show how this approach may be applied to the processes in dimension two with piecewise constant drifts in order to obtain a sharp geometric characterization of the stability region. Section 6 illustrates our various results and shows that our sufficient conditions are not necessary in dimension three. Section 7 concludes the paper.

2. The model

Let N be an integer greater than 1. We denote by $\mathbb{A}_+^N$ the positive orthant of $\mathbb{A}^N$ (where $\mathbb{A}$ in this paper will be $\mathbb{Z}$ or $\mathbb{R}$) while $\mathbb{A}_{+,*}^N$ stands for $\mathbb{A}_+^N \setminus \{0\}$. This means that, for example, the set $\mathbb{R}_+^N$ consists of real vectors $(x_1, \dots, x_N)$ such that $x_i \geq 0$ for all $i = 1, \dots, N$.

Let e_i be the vector in $\mathbb{Z}_+^N$ defined by $(e_i)_i = 1$ and $(e_i)_j = 0$, $j \neq i$. If not specified otherwise, $|\cdot|$ denotes the usual Euclidean norm. The notation $x \leq y$ is used for the coordinate-wise ordering: for all i, $x_i \leq y_i$, and we denote by $\langle x, y\rangle$ the usual scalar product of two vectors in $\mathbb{R}^N$. A process X or a trajectory u started in x at time 0 will be denoted by X^x and u^x, respectively.

Assume now that X is a continuous-time Markov process on $\mathbb{Z}_+^N$ with the following transition rates:

$$q(x, x + e_i) = \lambda_i(x),$$
$$q(x, x - e_i) = \phi_i(x), \tag{2}$$

where $\lambda = (\lambda_i)_{i=1\dots N}$ and $\phi = (\phi_i)_{i=1,\dots,N}$ are vectors consisting of positive 0-homogeneous functions from $\mathbb{R}_+^N$ to $\mathbb{R}$.

The drift function $\delta = (\delta_1, \dots, \delta_N) = \lambda - \phi$ is bounded, which guarantees that the process X is nonexplosive. Hence, we may assume that X and all other stochastic processes treated in the sequel have paths in the space $D = D(\mathbb{R}_+, \mathbb{Z}_+^N)$ of right-continuous functions from $\mathbb{R}_+$ to $\mathbb{Z}_+^N$ with finite left limits. Recall that a stochastic process with paths in D can be viewed as a random element on the measurable space $(D, \mathcal{D})$, where $\mathcal{D}$ denotes the Borel σ-algebra generated by the standard Skorokhod topology [13]. We also define $\mathcal{F}_t$ as the induced filtration $\sigma\{X_s,\ s \leq t\}$ and all the martingales constructed are adapted to that filtration.

We are interested in conditions on the drift vector field δ ensuring that the process is either stable (positive recurrent) or unstable (transient or null-recurrent). In subsequent sections we shall find such conditions with the use of the so-called Foster–Lyapunov criterion, that we recall here. Proofs under different conditions may be found in, for example, [8], [18], [20], and [25]. The following version of the criterion is proved in [20].

Theorem 1. *Assume that $\{X_t\}$ is a Markov process on $\mathcal{X}$. If there exist a (so-called Lyapunov) function $L\colon \mathcal{X} \to \mathbb{R}_+$, a constant $K > 0$, and an integrable stopping time τ such that*

(1) $\mathbb{E}^x[L(X_\tau)] - L(x) \leq -\gamma \mathbb{E}^x[\tau]$, *for* $L(x) > K$,

(2) *the set* $F = \{x\colon L(x) \leq K\}$ *is finite,*

(3) $\mathbb{E}^x[L(X(1))] < \infty$, *for all* $x \in \mathcal{X}$,

then the Markov process X is stable, i.e. positive recurrent.

3. Smooth drift

In this section we consider the case when the drift of the system may be extended to a smooth function. Our goal is to unify the existing results by using an ODE method without proving that (1) or a similar equation represents the behavior of a scaled version of the process under consideration. We also provide a way to construct an explicit Lyapunov function.

3.1. Stability conditions

Let $(\mathrm{ODE})_x$ be the following deterministic differential equation:

$$\frac{\mathrm{d}}{\mathrm{d}t} u(t) = \delta(u(t)), \qquad u(0) = x.$$

We denote its solution by u^x. Define $\mathcal{S} = \{x\colon |x| = 1\}$.

Theorem 2. *Assume that δ is a Lipschitz-continuous vector field from $\mathbb{R}_+^N$ to $\mathbb{R}_+^N$. Let*

$$T_x = \inf\{t\colon u^x(t) = 0\}.$$

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

Assume that for all $x \in \mathcal{S}$ there exists a continuous solution of $(\mathrm{ODE})_x$ *such that* $T_x < \infty$. *Assume in addition that*

$$\sup_{x\in\mathcal{S}} T_x < \infty. \tag{3}$$

Then X is positive recurrent and $x \mapsto T_x$ is a Lyapunov function.

Proof. Note first that the homogeneity of the drift implies that, if $u^x(t)$ is a solution of $(\mathrm{ODE})_x$, then we can define $u^{Kx}(t) = Ku^x(t/K)$ as a solution of $(\mathrm{ODE})_{Kx}$. Indeed,

$$u^{Kx}(0) = Ku^x(0) = Kx,$$

$$\frac{\mathrm{d}}{\mathrm{d}t}u^{Kx}(t) = K\frac{\mathrm{d}}{\mathrm{d}t}u^x\left(\frac{t}{K}\right) = \delta\left(u^x\left(\frac{t}{K}\right)\right) = \delta\left(Ku^x\left(\frac{t}{K}\right)\right) = \delta(u^{Kx}(t)).$$

As δ is Lipschitz-continuous, the flow $(t, x) \to u^x(t)$ is C^1 on $\mathbb{R} \times \mathbb{R}^N_*$, i.e. continuously differentiable in t and x (see, for instance, [12, Theorem 1]).

We are going to show that $F : x \to T_x$ is a suitable Lyapunov function for proving the positive recurrence of X. Owing to the assumptions of the theorem, F is a positive finite function. Since $u^{Kx}(t) = Ku^x(t/K)$, it follows that $F(Kx) = KF(x)$, i.e. F is a 1-homogeneous function. This, in particular, implies that $F(x) \to \infty$, when $|x| \to \infty$.

We then observe that

$$\frac{\mathrm{d}}{\mathrm{d}t}F(u^x(t)) = \lim_{h\to 0^+}\frac{F(u^x(t+h)) - F(u^x(t))}{h} = -1 \quad \text{for all } t < T_x. \tag{4}$$

Indeed, the difference between $F(u^x(t+h))$ and $F(u^x(t))$ (where h is positive) is negative and is equal in absolute value to the time needed to reach $u^x(t+h)$ from $u^x(t)$, which is exactly h. Hence, the latter equality follows.

The drift of F is given by

$$\Delta F(x) = \sum_y q(x, y)(F(y) - F(x)),$$

where $q(x, y)$ is defined in (2).

It will be convenient to approximate the variations of F using its gradient (when it is differentiable). For that purpose, define for each $\kappa > 0$ the functions

$$T_{x,\kappa} = \inf\{t : |u^x(t)| \le \kappa\}.$$

It is clear that $T_{x,\kappa} \to T_x$ for each fixed x as $\kappa \to 0$. It is also clear that, for each fixed x, functions $T_{x,\kappa}$ increase when κ decreases. Note that, due to the continuity of $u^x(t)$ in t,

$$|u^x(T_{x,\kappa})| = \kappa.$$

We can now prove that $T_{x,\kappa}$ is differentiable. Examine the above equality: $u^x(t)$ is a differentiable function, while the norm of a differentiable function is also differentiable. Hence, we conclude that $T_{x,\kappa}$ is differentiable for all values of $\kappa > 0$. In order to prove the differentiability of T_x (and $F(x)$), it remains to show that the convergence $T_{x,\kappa} \to T_x$ is uniform in x. It follows from the following sequence of equalities that

$$\sup_x |T_{x,\kappa} - T_x| = \sup_{x\,:\,|x|=\kappa} |T_x| = \kappa \sup_{x\in\mathcal{S}} |T_x| \to 0,$$

as $\kappa \to 0$, due to (3). The last equality in the sequence follows from 1-homogeneity of T_x.

Recall now that, due to (2),

$$\Delta F(x) = \sum_{i=1}^{N} \lambda_i(x)(F(x+e_i) - F(x)) + \phi_i(x)(F(x-e_i) - F(x)).$$

As F is 1-homogeneous and C^1, we have

$$F(x+e_i) - F(x) = |x|\left(F\left(\frac{x}{|x|} + \frac{e_i}{|x|}\right) - F\left(\frac{x}{|x|}\right)\right) \le \frac{\partial F}{\partial x_i}\left(\frac{x}{|x|}\right) + \varepsilon,$$

for a positive ε and for all large enough x, where we used the uniform continuity of a continuous function $\partial F/\partial x_i$ on a compact set $\mathcal{S}$.

We can now conclude that

$$\Delta F(x) \le \langle \nabla F(x), \delta(x) \rangle + \varepsilon,$$

for $|x|$ large enough, where $\varepsilon > 0$ may be chosen to be sufficiently small.

Note also that, with the use of (4), we can obtain

$$-1 = \frac{\mathrm{d}}{\mathrm{d}t} F(u^x(0)) = \langle \nabla F(u^x(0)), \delta(u^x(0)) \rangle = \langle \nabla F(x), \delta(x) \rangle.$$

To summarize, the function F is such that $F(x) \to \infty$ as $|x| \to \infty$ and that $\Delta F(x) \le -1 + \varepsilon < 0$, for $|x|$ large enough. After an examination of the Foster–Lyapunov criterion, it can be concluded that X is positive recurrent.

Remark 1. Note that our approach is very similar to that of the fluid-limits approximation as developed, for instance, in [10]. However, we do not need to prove that the differential equation we look at is indeed the one that represents the behavior of our birth-and-death process on the fluid scale.

Remark 2. It is worth mentioning that our analysis is valid due to the 0-homogeneity assumption. Without this assumption, the stability of the birth-and-death processes cannot be described using $(\mathrm{ODE})_x$. A counter-example (with continuous drifts) can be found in [2].

3.2. Instability conditions

In this section, we consider a reverse statement establishing instability relying on the previously considered dynamical system. It is much more challenging to state generic instability conditions based on $(\mathrm{ODE})_x$, without a direct use of fluid limits, i.e. without needing to prove the convergence of a scaled version of the process towards the trajectories of $(\mathrm{ODE})_x$ (see [16] for the construction of a Lyapunov function proving the transience of multi-class queueing networks with routeing).

We use hereafter explicitly the convergence to the fluid limits (see also [10] for a proof of convergence towards the fluid limit in the discrete-time setting). The next theorem is hence essentially a combination of proving the convergence of the scaled process and the extended version of the (instability part of the) Foster–Lyapunov criterion (see, for instance, [8, Theorem 2.2.7]).

Theorem 3. *Assume that δ is a Lipschitz function. Assume further that there exists a strictly positive time T and a number $a > 1$ such that, for all x with $|x| = 1$, a solution u^x of* $(\mathrm{ODE})_x$ *is defined on an interval* $[0, \tau_x]$ *with $\tau_x > T$ and satisfies $|u^x(T)| \ge a$. Then X is transient.*

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

Proof. Note first that the Lipschitz condition ensures that $(\text{ODE})_x$ has a unique solution u^x for each $x \in \mathbb{R}^N_{+,*}$ on an interval $[0, \tau_x]$. Furthermore, since 0 is necessarily a stable point if it is an equilibrium point, the conditions of the theorem imply that the trajectories $u^x(t)$ did not hit 0 before time T. Using standard arguments, we first prove the convergence of the process towards its fluid limit uniformly in its initial point.

Fix $\eta > 0$. Using the martingale decomposition, X^{Kx} can be decomposed as

$$\begin{aligned}\frac{X^{Kx}(Kt)}{K} &= \frac{Kx}{K} + \frac{1}{K}\int_0^{Kt} \delta(X^{Kx}(s))\,\mathrm{d}s + \frac{M_{Kt}}{K}\\ &= x + \frac{1}{K}\int_0^{Kt} \delta(X^{Kx}(s))\,\mathrm{d}s + \frac{M_{Kt}}{K}\\ &= x + \int_0^{t} \delta\Big(\frac{X^{Kx}(Ks)}{K}\Big)\,\mathrm{d}s + \frac{M_{Kt}}{K},\end{aligned}$$

where M is a martingale that satisfies

$$\mathbb{E}\Big(\sup_{0\le s\le t} \frac{M_{Kt}}{K}\Big) \le A\Big(\frac{t}{K}\Big)^{1/2} \le \eta,$$

for all K large enough and with A being a positive constant. Now define

$$f_K(t) = \mathbb{E}\bigg[\sup_{0\le s\le t}\bigg|\frac{X^{Kx}(Ks)}{K - u^x(s)}\bigg|\bigg]$$

and recall that, due to $(\text{ODE})_x$,

$$u^x(s) = x + \int_0^s \delta(u^x(v))\,\mathrm{d}v.$$

With the use of the Lipschitz condition and the homogeneity of the drift we now obtain

$$\begin{aligned}f_K(t) &\le \eta + \mathbb{E}\bigg[\sup_{0\le s\le t}\bigg|\int_0^s \delta\Big(\frac{X^{Kx}(Kv)}{K}\Big)\,\mathrm{d}v - \int_0^s \delta(u^x(v))\,\mathrm{d}v\bigg|\bigg]\\ &\le \eta + L\int_0^t f_K(s)\,\mathrm{d}s,\end{aligned}$$

where L is the Lipschitz constant. Gronwall's lemma allows us to conclude that

$$\sup_x \mathbb{E}\sup_{s\le t}\bigg|\frac{X^{Kx}(Ks)}{K} - u^x(s)\bigg| \le \eta\exp(Lt),$$

for each interval $[0, t]$ included in the interval $[0, \tau_x]$ where the ODE has a solution.

Take ε such that $a - 1 - \varepsilon > 0$. This further implies that there exists K_0 such that, for all x and $K \ge K_0$,

$$\mathbb{E}|X^{Kx}(KT)| - K|x| - (|u^x(T)| - |x|)K \ge -\varepsilon K.$$

Hence, since $|x| = 1$, $\mathbb{E}|X^{Kx}(KT)| - K|x| \ge (a - 1 - \varepsilon)K > 0$. This bound being uniform in x with $|x| = 1$, we can make use of the extended Foster–Lyapunov criterion (see, for instance, [8, Theorem 2.2.7]) to conclude the proof.

Remark 3. We believe that Theorem 3 can be extended to the case where the drift is locally Lipschitz outside a neighborhood of 0 but the proof details become more involved as we need to introduce stopping times to control that the scaled process is not entering this neighborhood. This falls outside the scope of this paper.

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

3.3. Conservative drifts

Finding an explicit form for the function F might be difficult in general. However, under slightly stronger assumptions on the drift function δ, we can construct an explicit Lyapunov function directly from the vector field δ. This is a well-known fact in the theory of deterministic dynamical systems.

Proposition 1. *Assume that δ is a conservative vector field, i.e. $\delta = -\nabla V$, and assume that*

$$V(x) \geq a > 0.$$

Assume further that there exists $\varepsilon > 0$ such that

$$|\delta(x)| \geq \varepsilon \quad \textit{for all } x.$$

Then $V(x)$ is a Lyapunov function and X is positive recurrent.

Proof. Note first that since δ is 0-homogeneous, V is 1-homogeneous. Using the 1-homogeneity, there exists $\kappa < \varepsilon^2$ such that, for $|x|$ large enough, we can estimate the drift of V (as in the proof of Theorem 2) by

$$\Delta V(x) \leq \langle \nabla V(x), \delta(x) \rangle + \kappa = \langle \nabla\ V(x), -\nabla V(x) \rangle + \kappa = -|\nabla\ V(x)|^2 + \kappa \leq -\varepsilon',$$

for a constant $\varepsilon' > 0$. Furthermore, $|V(x)| \to \infty$ for $|x| \to \infty$ since it is a 1-homogeneous and strictly positive function. We can therefore apply the Foster–Lyapunov criterion.

Remark 4. A vector field $\delta(x_1, x_2) = (\delta_1(x_1, x_2), \delta_2(x_1, x_2))$ (on a completely connected set) is conservative if and only if $(\mathrm{d}/\mathrm{d}x_1)\delta_2(x) = (\mathrm{d}/\mathrm{d}x_2)\delta_1(x)$.

4. Discontinuous drifts

4.1. Complexity of the fluid limits

So far we restricted ourselves to the case when the drift vector field is continuous. The situation changes dramatically when this condition is dropped. When the drift vector field is discontinuous, the trajectories of a fluid-equivalent system near a point of discontinuity may enter 'sliding modes' and the differential equation (1) has to be replaced by a new dynamical system defined piecewise by differential equations $(\mathrm{d}/\mathrm{d}t)x(t) = \tilde{\delta}(x(t))$, where $\tilde{\delta}$ is a convex combination of drifts around neighborhoods of the discontinuities.

Let us give a simple example of this phenomenon. Consider the following transitions with fixed birth rates λ_1, λ_2 and death rates given by the following bandwidth allocation:

$$\phi_1(x) = \mathbf{1}_{x_2=0} + a_1\, \mathbf{1}_{x_2>0},$$
$$\phi_2(x) = \mathbf{1}_{x_1=0} + a_2\, \mathbf{1}_{x_1>0},$$

where $\mathbf{1}.$ denotes the indicator function.

Suppose that $\lambda_1 < a_1$ and $\lambda_2 < (1 - \rho_1) + a_2\rho_1$, where $\rho_i = \lambda_i / a_i$. This condition is known to be sufficient for stability of such a model in dimension two, and it has been obtained through different methods (see, for instance, [2], [4], and [8]).

We could further prove that the process $X^K(Kt)/K$ converges in distribution when $K \to \infty$ towards a process $x(t)$ satisfying the differential equations

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) = \delta(x(t)) = (\lambda - \phi)(x(t)), \quad \text{for } x(t) > 0,$$
$$\frac{\mathrm{d}}{\mathrm{d}t}x_1(t) = 0, \quad \text{for } x_2(t) = 0,$$
$$\frac{\mathrm{d}}{\mathrm{d}t}x_2(t) = \lambda_2 - (1 - \rho_1) - a_2\rho_1, \quad \text{for } x_1(t) = 0.$$

The stability condition is then easily interpreted when considering the convergence to 0 of the obtained fluid limit. This example shows, however, that even in a very simple case the fluid limit satisfies an equation different from (1).

In the next subsection we develop an approach that allows us to find stability conditions in the case of discontinuous drift vector fields without the use of fluid approximation.

4.2. Sufficient stability conditions

This section is devoted to identifying rather general conditions on the drift vector field ensuring stability even in the presence of discontinuities for δ. These conditions lead to useful geometric stability conditions in dimension two, which are discussed in Section 5.

We start by considering a general vector field of 0-homogeneous drifts such that the number of discontinuities is finite. We construct a Lyapunov function by pasting together local Lyapunov functions and using a smoothing technique. This method was first used in [7].

Define a closed sphere with radius ε and center x by $\mathcal{B}_\varepsilon(x)$. For a point x, denote the set of drifts in a neighborhood of x by $\mathcal{D}_\varepsilon(x)$, i.e.

$$\mathcal{D}_\varepsilon(x) = \{\delta(y) \colon y \in B_\varepsilon(x)\}.$$

Also define the following set of vectors:

$$\mathcal{D}^*_{a,\varepsilon}(x) = \{\eta \in \mathbb{R}^N \colon \langle \eta, v\rangle < -a, \text{ for all } v \in \mathcal{D}_\varepsilon(x) \cup \{-x\}\}.$$

We now state an assumption on the vector field $\delta(x)$ that we shall prove to be sufficient to characterize the stability region of the process.

Assumption (A_1). For all $x \neq 0$, there exist $\varepsilon > 0$ and $a > 0$ such that

$$\mathcal{D}^*_{a,\varepsilon}(x) \neq \varnothing.$$

We now state the main result of this section.

Theorem 4. *Assumption (A_1) implies that X is positive recurrent.*

Before presenting a rigorous proof, we would like to explain the result intuitively. If the sets $\mathcal{D}_\varepsilon(x)$ are finite for all x, Assumption (A_1) may be better understood using a simple geometric interpretation. Using Farkas' lemma (see e.g. [21, p. 200]), we can state that either Assumption (A_1) is true or x is in the cone induced by the vectors of $\mathcal{D}_\varepsilon(x)$, i.e. there exist nonnegative weights α_i such that

$$\sum_{i \in \mathcal{I}} \alpha_i \delta^i = x,$$

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

with a nonempty index set $\mathcal{I}$ and vectors $\delta_i \in \mathcal{D}_\varepsilon(x)$. It is hence natural to expect that if x is never contained in the cone induced by the drifts $\delta(y)$ at points y close to x (which is exactly Assumption (A_1)), then the process is stable.

Proof of Theorem 4. Consider first the function H and the vector-field v_u constructed by

$$v_u(x) = \arg \max_{\eta \in \mathcal{D}^*_{a,u}(x)} \langle x, \eta \rangle, \qquad H(x,u) = \max_{\eta \in \mathcal{D}^*_{a,u}(x)} \langle x, \eta \rangle.$$

The function H is a natural candidate for a Lyapunov function but it is discontinuous which complicates drastically the drift calculations and precludes having a negative drift in all points. We overcome this difficulty by considering a smoothed version of H. Let κ_ε be a C^∞-probability density supported on the sphere $\mathcal{B}_\varepsilon(0)$ and introduce

$$F(x) = \int_{u \in \mathcal{B}_\varepsilon(0)} H(x,u) \kappa_\varepsilon(u)\, \mathrm{d}u.$$

Then F is clearly a C^∞-function. We shall now prove that F is a suitable Lyapunov function. First, note that by Assumption (A_1), $H(x,u) \geq a$; hence, $F(x) \geq a$ for all x. This, together with the observation that the function F is 1-homogeneous, implies that $F(x) \to \infty$ as $|x| \to \infty$.

Note further that the compactness of the sphere guarantees the existence of $u \leq \varepsilon$ (arbitrarily small) such that, for all x,

$$|\nabla F(x) - v_u(x)| \leq -\varepsilon. \tag{5}$$

Using Assumption (A_1) again, we obtain

$$\langle \delta(x), v_u(x) \rangle \leq -a. \tag{6}$$

Hence, using (5),

$$\langle \delta(x), \nabla F(x) \rangle \leq \langle \delta(x), v_u(x) \rangle + C\varepsilon. \tag{7}$$

Again using the fact that F is 1-homogeneous and combining (6) and (7), we obtain

$$\Delta F(x) = \sum_y q(x,y)(F(y) - F(x)) \leq \langle \delta(x), \nabla F(x) \rangle + \varepsilon',$$

for large $|x|$. Hence, F is a Lyapunov function.

5. Piecewise constant drift in dimension two

In this section we apply the general result of Theorem 4 to a particular case of a discontinuous drift function. We assume that the state space of the underlying process is $\mathbb{N}^2$ and that the rate functions are piecewise constant. Together with the assumption that the rate functions are 0-homogeneous, this means that there is a finite number of cones inside which the rate functions are constant.

We start by introducing some notation that will be used throughout this section. Assume that there are N vectors $v_1, \ldots, v_N$ such that $v_1 = e_1$, $v_N = e_2$ (where e_1 and e_2 are vectors co-directed with one of the axes), and such that $\delta(x) = \delta^k$ for any $x = Av_k + Bv_{k+1}$ with positive A and B. This means that the drift at any point of the cone defined by v_k and v_{k+1} is equal to δ^k.

Note that we do not require the vectors v_k to be different. (This means that a cone reduces to a line when two consecutive vectors v_k and v_{k+1} are equal.)

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

We also introduce certain sets that will be crucial for the definition of the stability region. For each $k = 1, \dots, N-1$, let

$$U_k^1 = \{\delta \colon \delta^k = Av_k + Bv_{k+1}, \text{ for some } A \geq 0, B \geq 0, A + B > 0\},$$

and for each $k = 2, \dots, N$, let

$$U_k^2 = \{\delta \colon \alpha\delta^k + (1-\alpha)\delta^{k-1} = Av_k, \text{ for some } A \geq 0, \alpha \in [0, 1]\}.$$

We are ready to state the main result of this section.

Theorem 5. *Assume that*

$$\delta \in \mathcal{S} = \bigcap_{k=1}^{N-1} (\overline{U_k^1} \cap \overline{U_k^2}).$$

Then the Markov process X is positive recurrent. Conversely, if δ belongs to the interior of the complement of $\mathcal{S}$, then X is transient or null-recurrent.

Proof. We start by proving the first part of the theorem. In order to do this, let us verify the conditions of Theorem 4. It is clear that only two situations are possible:

(i) the vector x belongs to the interior of a cone defined by vectors v_k and v_{k+1}; in this case $x = Av_k + Bv_{k+1}$ for some $A > 0$, $B > 0$,

(ii) the vector x is collinear to a vector v_k; in this case $x = Av_k$ for some $A > 0$.

We consider these two situations separately. In case (i), thanks to Theorem 4, we need to show the existence of a vector η such that

$$\langle \eta, Av_k + Bv_{k+1} \rangle > 0, \qquad \langle \eta, \delta^k \rangle < 0.$$

In geometric terms this means that there exists a vector η such that vectors $Av_k + Bv_{k+1}$ and δ^k belong to different half-planes separated by the line normal to vector η. It is easy to see that the existence of such a vector is guaranteed by the fact that $\delta \notin U_k^1$.

Consider now situation (ii). Again applying Theorem 4, we see that we need to show the existence of a vector η such that

$$\langle \eta, Av_k \rangle > 0, \qquad \langle \eta, \delta^k \rangle < 0, \qquad \langle \eta, \delta^{k-1} \rangle < 0.$$

If we interpret this again in geometric terms, it is equivalent to the existence of a vector η such that vectors $-Av_k$, δ^k, and δ^{k-1} belong to the same half-plane defined by the line normal to η. Direct computations show that it follows from the fact that $\delta \notin U_k^2$. The proof of the positive recurrence under the assumption that $\delta \in \mathcal{S}$ is now complete.

Let us now show that if $\delta \in \operatorname{int}(\bigcup_{k=1}^{N-1} U_k^1 \cup U_k^2)$, then the Markov process X is not positive recurrent. We are going to prove that, under the given conditions, the process is actually not rate stable which prevents stability. Assume that the process is started in a cone k and that $\delta \in U_k^1$. The strong law of large numbers (SLLN) then implies that, with a positive probability, the process stays in the cone $\{Av_k + Bv_{k+1},\ A \geq 0, B \geq 0\}$.

Assume that X is positive recurrent. This implies that $X_t^x/t \to 0$ almost surely. Using the martingale decomposition of the process, write

$$\frac{X_t^x}{t} = \frac{x}{t} + \frac{M_t}{t} + \int_0^t \delta(X_s^x)\,\mathrm{d}s.$$

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

Owing to the boundedness of the transitions, the martingale M_t is such that $\mathbb{E}[M_t^2] \leq Ct$, which implies the convergence in L^2 and in probability of M_t/t to 0, which in turn implies the almost sure convergence along a subsequence. Conditioning on the fact that the process stays in the cone, we obtain

$$(1/t_n)\int_0^{t_n} \delta(X^x(s))\,\mathrm{d}s \to 0, \qquad t_n \to \infty,$$

which combined with the ergodic theorem for the positive recurrent Markov process X implies that there exists $\alpha \geq 0$ such that

$$0 = \alpha\delta^{k-1} + (1-\alpha)\delta^k,$$

which contradicts $\delta \in U_k^1$.

Suppose now that $\delta \in U_k^2$. In this case the SLLN implies that, with a strictly positive probability, the process stays in the set

$$\{Av_{k-1} + Bv_k,\ A \geq 0, B \geq 0\} \cup \{Av_k + Bv_{k+1},\ A \geq 0, B \geq 0\}.$$

Proceeding in a similar way as in the previous case, there exists $\tilde{\alpha}$ such that

$$0 = \tilde{\alpha}\delta^k + (1-\tilde{\alpha})\delta^{k+1},$$

which contradicts $\delta \in U_k^2$.

5.1. Fluid limits

An interesting situation occurs within the framework of this section when there exist k and $\alpha \in (0, 1)$ such that $\alpha\delta^k + (1-\alpha)\delta^{k+1} = Av_k$ for some $A \neq 0$. In this case we know that the fluid limits with an initial state in the cone defined by v_{k-1} and v_k or in the cone defined by v_k and v_{k+1} enter a so-called 'sliding mode' with their trajectory reaching the ray defined by v_k after a finite time and not leaving it after this time. The new drift $\tilde{\delta}$ obtained during the sliding mode on v_k must be a convex combination of δ^k and δ^{k+1} and must also be collinear with v_k. Hence, we can explicitly calculate $\tilde{\delta}$ by solving the following system in α and A:

$$\alpha\delta^k + (1-\alpha)\delta^{k+1} = Av_k. \tag{8}$$

The existence of a solution with a strictly negative A is necessary to get a stable system while the existence of a solution with strictly positive A is sufficient to get instability of the process, which corresponds respectively to the case where the fluid limits converge to zero or infinity.

This is a very particular scenario, as in the case of a dimension higher than two, (8) is generally undetermined and stability conditions cannot be characterized directly.

5.2. Algorithm

Recall that we defined the drift vector $\delta(\cdot)$ to be equal to $\lambda(\cdot) - \phi(\cdot)$. It is often the case in queueing and telecommunications applications that the vector field λ (representing the arrival rate) is assumed to be constant and the question is for which values of λ the system under consideration is going to be stable. We present here an algorithm to construct the stability set when λ is fixed. Define (within the description of the algorithm and the examples later on) the different drift values as $\delta^k = \lambda - \psi_k$. The algorithm is given below.

Step 1. Draw the points representing various values of ψ_k.

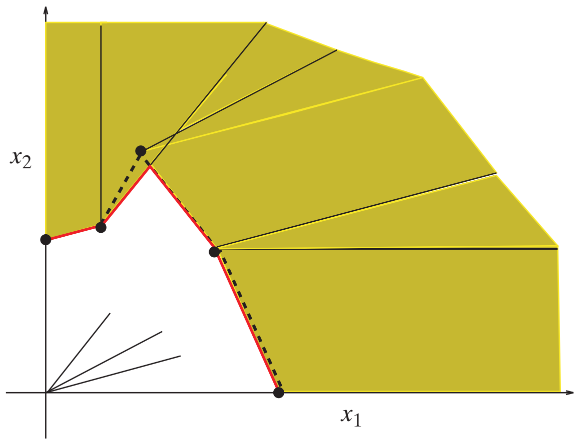

FIGURE 1: Generic construction of the stability set for piecewise constant drifts in dimension two.

Step 2. Connect ψ_1 to ψ_2, ψ_2 to ψ_3, and so on.

Step 3. For each k, draw the cone defined by vectors v_k and v_{k+1} based on point ψ_k. The compact set obtained is the stability region.

A generic illustration of this algorithm is given in Figure 1.

6. Examples

6.1. Continuous drifts

The results obtained in Section 3 allow us to study numerically the positive recurrence of processes even when the drifts are too complicated to get an explicit solution for the associated ODE. We now give an example of such a situation.

Example 1. An example of a wireless network with two types of users competing for the same bandwidth could lead to the following death rates (using Shannon's formula and a state-dependent allocation policy):

$$\phi_1(x) = \log\left(1 + \frac{x_1/|x|}{N + x_2/|x|}\right), \qquad \phi_2(x) = \log\left(1 + \frac{x_2/|x|}{N + x_1/|x|}\right),$$

where N is the thermal noise.

Let us consider two possible vectors of arrival (birth) rates: $(0.4, 0.8)$ and $(0.5, 0.8)$. The associated ODE can be solved numerically for any value λ_1, λ_2 allowing us to conclude for the positive recurrence of the process in the first case and the transience in the second case, as shown in Figure 2 using the following properties of the trajectories of the ODE.

- In the first case, all trajectories started from any state on the sphere hit 0 in a bounded time.
- In the second case, all trajectories started from the sphere do not reach a sufficiently small neighborhood of 0, from which we can conclude that all fluid-limit solutions stay outside of a ball of radius ε and center 0. Moreover, all trajectories do reach a state with a norm larger than 1 before a finite time T.

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

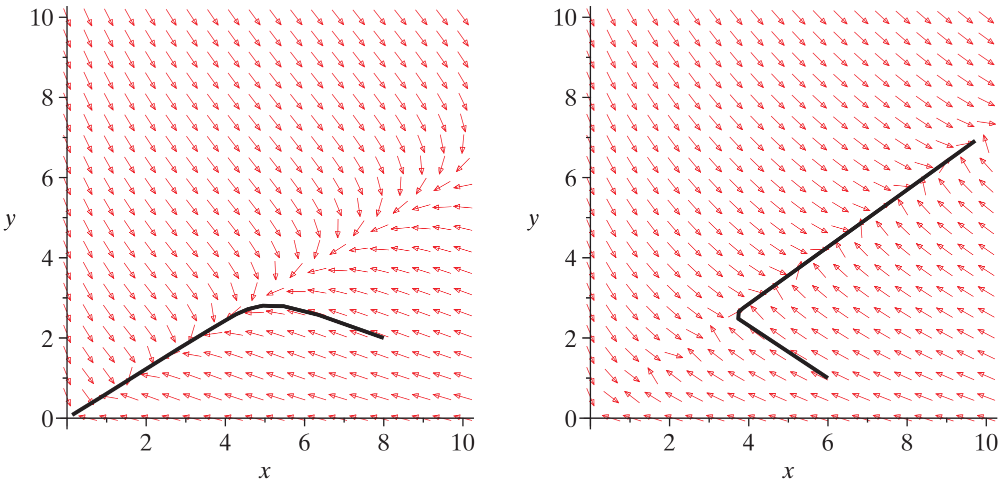

FIGURE 2: Example 1: drift vector fields for birth rates (0.4, 0.8) and (0.5, 0.8).

6.2. Discontinuous drifts

We will give here a few examples where Theorem 5 provides interesting results.

Example 2. We describe here how Theorem 5 can be used to obtain the well-known stability results for the so-called coupled-processors problem. Consider the allocation described as the most basic example with discontinuous drifts in Section 4:

$$\phi_1(x) = \mathbf{1}_{x_2=0} + a_1\,\mathbf{1}_{x_2>0}, \qquad \phi_2(x) = \mathbf{1}_{x_1=0} + a_2\,\mathbf{1}_{x_1>0}\,.$$

It is clear that in this case the algorithm of Section 5.2 allows us to recover the well-known stability region for this problem [4], [8] (see Figure 3): we should have that for all $\alpha > 0$ $\alpha(\lambda_1 - 1) + (1-\alpha)(\lambda_1 - a_1) < 0$ and symmetrically in the other coordinate. Solving in α, we obtain $\lambda_1 < a_1$ and $\lambda_2 < (1-\rho_1) + a_2\rho_1$ or $\lambda_2 < a_2$ and $\lambda_1 < (1-\rho_2) + a_1\rho_2$, with $\rho_i = \lambda_i / a_i$.

Remark 5. Note that, as illustrated in Figure 1, the stability set may not be convex.

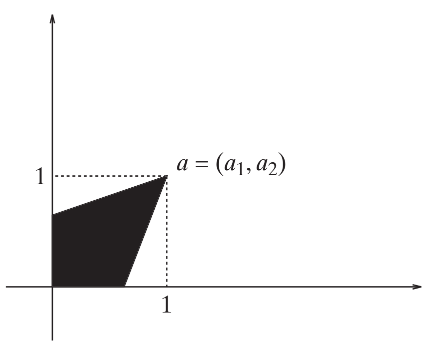

FIGURE 3: Stability region (set of birth rates λ) for the two coupled processors of Example 2.

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

Example 3. We now look at the same type of death rates as in Example 2 but with different birth rates. Consider for instance a queueing or manufacturing system with two processors and two types of traffic:

- some dedicated traffic arriving with intensity λ_i to processor i,
- some flexible traffic with intensity ν that can be routed to either processor 1 or 2 depending on the congestion level of both processors.

Assume that the flexible traffic is actually routed to the processor with the smallest number of jobs in process (and to processor 1, say, if the processors are equally loaded, this last assumption having no impact on the stability conditions). Assume further that the arrivals of jobs of each type of traffic are following a Poisson process (independent of each other and of everything else) and that the processing times are exponentially distributed. The presence of interference or switching costs between different type of tasks raise the allocation of service (or death rates) described in Example 2 with $a_i < 1$.

Using the notation of Section 5, the vectors v_k are

$$v_1 = (1, 0), \qquad v_2 = (1, 0), \qquad v_3 = (1, 1), \qquad v_4 = (0, 1), \qquad v_5 = (0, 1),$$

and the drifts are

$$\delta^1 = (\lambda_1 - 1, \lambda_2 + \nu), \qquad \delta^2 = (\lambda_1 - a_1, \lambda_2 + \nu - a_2),$$
$$\delta^3 = (\lambda_1 + \nu - a_1, \lambda_2 - a_2), \qquad \delta^4 = (\lambda_1 + \nu, \lambda_2 - 1).$$

Using the results of Section 5, the interior of the stability region can be described as follows:

- if $\lambda_2 + \nu > a_2$ then $\lambda_1 - a_1 < \nu + \lambda_2 - a_2$, and $\lambda_1 + \lambda_2 + \nu < a_1 + a_2$,
- if $\lambda_2 + \nu < a_2$, then $\lambda_1 < 1 + ((a_1 - 1)/a_2)(\lambda_2 + \nu)$.

Symmetric conditions with the indices 1 and 2 reversed should also hold.

6.3. Bounds on the stability region of three coupled processors

Consider a process of dimension three where the death rates of each dimension depend on whether the other coordinates are strictly positive or zero, so that for all $i \neq j \neq k$ and $x_i > 0$,

$$\phi_i(x) = \begin{cases} a_i, & x_j = 0,\ x_k = 0, \\ a_{ij}, & x_j > 0,\ x_k = 0, \\ 1, & x_j > 0,\ x_k > 0, \end{cases}$$

which leads to the following drifts:

$$\delta(x) = \begin{cases} \delta^i : \delta^i_i = \lambda_i - a_i,\ \delta^i_j = \lambda_j,\ \delta^i_k = \lambda_k, & \text{for } x_j = 0,\ x_k = 0, \\ \delta^{ij} : \delta^{ij}_i = \lambda_i - a_{ij},\ \delta^{ij}_j = \lambda_j - a_{ji},\ \delta^{ij}_k = \lambda_k, & \text{for } x_j > 0,\ x_k = 0, \\ \delta = (\lambda_i - 1)_{i=1,\dots,3} & \text{for } x_j > 0,\ x_k > 0. \end{cases}$$

Let us assume that $a_i \geq a_{ij} \geq 1$, so that $\phi = (\phi_1, \phi_2, \phi_3)$ is partially decreasing.

It was shown in [8, Theorem 4.4.4] and [2, Theorem 3] that the stability region is a union of six regions corresponding to the six possible permutations of the coordinates. The first of these regions corresponding to the identity permutation is the set of $(\lambda_1, \lambda_2, \lambda_3)$ such that

$$\delta_1 < 0, \qquad \delta^{23}_2 < \lambda_1(1 - a_{23}), \qquad \delta^3_3\pi_{00} + \delta^{13}_3\pi_{10} + \delta^{23}_3\pi_{01} + \delta_3\pi_{11} < 0,$$

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press

where

$$\pi_{00} = \mathbb{P}(Y_1 = 0, Y_2 = 0), \qquad \pi_{01} = \mathbb{P}(Y_1 = 0, Y_2 > 0),$$
$$\pi_{10} = \mathbb{P}(Y_1 > 0, Y_2 = 0), \qquad \pi_{11} = \mathbb{P}(Y_1 > 0, Y_2 > 0),$$

and $Y = (Y_1, Y_2)$ is a random vector distributed according to the stationary distribution of a process in which coordinate three would be always strictly positive. Quite heavy simulations are needed in order to compute the stability region. On the other hand, the sufficient conditions obtained in Section 4 do not need any computation and can be written as the complement of the following set:

$$\delta > 0,$$

or there exist i, j, and $\alpha_1, \alpha_2 \geq 0$ such that

$$\langle \alpha_1 \delta^{ij} + \alpha_2 \delta, e_i + e_j \rangle > 0,$$

or there exist i and $(\alpha_l)_{l=1,\dots,4} \geq 0$ such that

$$\delta_i^i \alpha_1 + \delta_i^{ij} \alpha_2 + \delta_i^{ik} \alpha_3 + \delta_i \alpha_4 > 0.$$

7. Conclusions

We derived various computable criteria of stability and instability for continuous drifts in any dimension and for piecewise constant drifts in dimension two, together with generic sufficient conditions for discontinuous drifts in any dimension. An important direction of future research is to systematically characterize the second vector field and the stability conditions in dimension three and more.

Acknowledgements

This work was carried out when both authors were employed by the Eindhoven University of Technology. The authors are very grateful to Mark Peletier for fruitful discussions and enlightening suggestions.

References

[1] Bonald, T., Massoulié, L., Proutière, A. and Virtamo, J. (2006). A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Systems* **53,** 65–84.
[2] Borst, S., Jonckheere, M. and Leskelä, L. (2008). Stability of parallel queueing systems with coupled service rates. *Discrete Event Dyn. Syst.* **18,** 447–472.
[3] Bramson, M. (2008). *Stability of Queueing Networks* (Lecture Notes Math. **1950**). Springer, Berlin.
[4] Cohen, J. W. and Boxma, O. J. (1983). *Boundary Value Problems in Queueing System Analysis* (North-Holland Math. Studies **79**). North-Holland, Amsterdam.
[5] Dai, J. G. (1995). On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Prob.* **5,** 49–77.
[6] Dai, J. G. and Meyn, S. P. (1995). Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Trans. Automatic Control* **40,** 1889–1904.
[7] Dupuis, P. and Williams, R. J. (1994). Lyapunov functions for semimartingale reflecting Brownian motions. *Ann. Prob.* **22,** 680–702.
[8] Fayolle, G., Malyshev, V. A. and Menshikov, M. V. (1995). *Topics in the Constructive Theory of Countable Markov Chains.* Cambridge University Press.
[9] Foley, R. D. and McDonald, D. R. (2001). Join the shortest queue: stability and exact asymptotics. *Ann. Appl. Prob.* **11,** 569–607.
[10] Fort, G., Meyn, S., Moulines, E. and Priouret, P. (2008). The ODE method for stability of skip-free Markov chains with applications to MCMC. *Ann. Appl. Prob.* **18,** 664–707.

[11] Foss, S. and Konstantopoulos, T. (2004). An overview of some stochastic stability methods. *J. Operat. Res. Soc. Japan* **47,** 275–303.
[12] Hirsch, M. W. and Smale, S. (1974). *Differential Equations, Dynamical Systems, and Linear Algebra* (Pure Appl. Math. **60**). Academic Press, New York.
[13] Kallenberg, O. (2002). *Foundations of Modern Probability*, 2nd edn. Springer, New York.
[14] Kelly, F. P., Maulloo, A. K. and Tan, D. K. H. (1998). Rate control for communication networks: shadow prices, proportional fairness and stability. *J. Operat. Res. Soc.* **49,** 237–252.
[15] Massoulié, L. (2007). Structural properties of proportional fairness: stability and insensitivity. *Ann. Appl. Prob.* **17,** 809–839.
[16] Meyn, S. P. (1995). Transience of multiclass queueing networks via fluid limit models. *Ann. Appl. Prob.* **5,** 946–957.
[17] Meyn, S. (2008). *Control Techniques for Complex Networks*. Cambridge University Press.
[18] Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer, London.
[19] Rao, R. R. and Ephremides, A. (1988). On the stability of interacting queues in a multiple-access system. *IEEE Trans. Inf. Theory* **34,** 918–930.
[20] Robert, P. (2003). *Stochastic Networks and Queues* (Appl. Math. (New York) **52**). Springer, Berlin.
[21] Rockafellar, R. T. (1970). *Convex Analysis* (Princeton Math. Ser. **28**). Princeton University Press.
[22] Stolyar, A. L. (1995). On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Process. Relat. Fields* **1,** 491–512.
[23] Szpankowski, W. (1988). Stability conditions for multidimensional queueing systems with computer applications. *Operat. Res.* **36,** 944–957.
[24] Szpankowski, W. (1994). Stability conditions for some distributed systems: buffered random access systems. *Adv. Appl. Prob.* **26,** 498–515.
[25] Tweedie, R. L. (1976). Criteria for classifying general Markov chains. *Adv. Appl. Prob.* **8,** 737–771.

https://doi.org/10.1239/aap/1396360103 Published online by Cambridge University Press