Stability of the two queue system

Iain M. MacPhee and Lisa J. Müller

University of Durham
Department of Mathematical Science
Durham, DH1 3LE, UK
(e-mail: i.m.macphee@durham.ac.uk, l.j.muller@durham.ac.uk)

Abstract. We describe ergodicity and transience conditions for a general two queue system with multiple service regimes, a dedicated traffic stream for each queue, a further stream which can be routed to either queue and where completed jobs can be fed back into the queues. There is only one class of jobs but the service times and feedback probabilities depend upon the configuration of the servers. Several different levels of control of the service regimes are considered. We use the semi-martingale methods described in [Fayolle *et al.*, 1995] and our results generalise those of [Kurkova, 2001].
Keywords: controlled queue systems.

1 Introduction

In this paper we consider a system which has two queues with servers that can be configured in several ways. Our main aim is to identify conditions under which we can give a queue length dependent policy for choosing the service configurations that guarantees the stability of the system.

The queues have independent Poisson arrival streams with rates λ_i, $i = 1$, 2 and there is an independent Poisson arrival stream with rate λ of jobs that can be sent to either queue (we will call this the *routeable* stream). We assume all jobs are of the same class and are served in the order they join their queues but their service times depend upon their queue and the service scheme in force while they are being served. Under server configuration k, at most one job is in service at each non-empty queue and all jobs in queue i have independent, exponentially distributed service times with mean μ_{ki}^{-1}, $i = 1$, 2 (so the server configuration k and the destination of a routeable job determines its service distribution). We label the server configurations by $k = 1, \ldots, K$, the queue to which the routable stream is directed by $j = 1$, 2 so that the finite set $\mathcal{R}$ of overall *management regimes* has members $\eta = (k, j)$. In addition the system has Jackson-type feedback with probabilities that depend upon the current management regime. Any job that completes service at queue i under regime η independently enters queue i' with probability $p_{ii'}^{\eta}$, $i' = 1$, 2 or leaves the system with probability $p_{i0}^{\eta} \equiv 1 - p_{i1}^{\eta} - p_{i2}^{\eta} \geq 0$. We will assume throughout this paper that we can instantaneously switch between different management regimes at the instants just after changes to queue lengths.

Example The model described above includes as a special case the model with two servers, where server i can be used to process jobs at either queue which it does at rate μ_i. This gives four service regimes $s_1s_2 = 1\,2,\ 2\,1,\ 1\,1$ and $2\,2$ (i.e. server 1 at queue 1, server 2 at queue 2; server 2 at queue 1, server 1 at queue 2; both servers at queue 1; both servers at queue 2). Given that the service rates are additive we get the pairs (μ_1, μ_2), (μ_2, μ_1), $(\mu_1 + \mu_2, 0)$ and $(0, \mu_1 + \mu_2)$ respectively.

The question we consider is whether for such a system with a given set of parameters, the management regime can be changed from time to time to ensure that the queue lengths remain stable or whether the queue lengths must grow indefinitely regardless of how the system is managed.

Similar systems but with fixed servers have been studied in the past using transform methods, often under strong symmetry assumptions on parameters, see [Feng *et al.*, 2002] and [Foley and McDonald, 2001] who give stability conditions for an n-dimensional JSQ model and carry out the large deviations analysis of system occupancy for the two dimensional system.

We define the model we consider in section 2 and state our results in section 3. We omit the proofs here to be able to describe the model in full length. The proof is done using the semi-martingale methods described in [Fayolle *et al.*, 1995] and can be found in [MacPhee and Müller,]. Our results generalise those of [Kurkova, 2001] as we consider multiple service regimes and do not require any symmetry.

2 Definitions

We now define the queueing system, its control, and the classes of control policies that we wish to investigate.

2.1 Events, blocks and control policies

As the Lyapunov function results we use are described in terms of discrete processes it is convenient to study a discrete time process which we now describe. To simplify comparison of the process dynamics under different management regimes we *uniformise* the continuous time jump process, following Serfozo [Serfozo, 1979], by choosing a constant $\rho \geq \max_k\{\lambda + \lambda_1 + \lambda_2 + \mu_{k1} + \mu_{k2}\}$ and introducing a fictitious *bell* event which has exponential inter-event times with rate $\rho - (\lambda + \lambda_1 + \lambda_2 + \mu_{k1} + \mu_{k2})$ at any given queue lengths when regime (k, j) is used (so the total event rate has the same value ρ in all states under all regimes). We now consider the uniformised discrete time process Ξ on state space $\mathbf{Z}_0^2 \equiv \{(x, y) \in \mathbf{Z}^2 : x \geq 0, y \geq 0\}$, obtained by considering the queue lengths at bell events, arrival times of new jobs and at service completions and consequent re-entry to queues. We will use $\alpha = (x, y) \in \mathbf{Z}_0^2$ to denote a typical state vector for Ξ.

It is also necessary to define the policies by which the management regimes at each state are selected. Our main interest will be in policies which choose the same regimes over large sets of states, specifically cone shaped blocks for which we need some notation. Let e_i denote the unit vector in the axis i direction and for non-zero $z \in \mathbf{R}^2$ let $|z|$ denote the length of z and $\arg_u(z)$ the argument relative to non-zero vector $u \in \mathbf{R}^2$ (the angle anticlockwise from u to z). For any non-zero u, $v \in \mathbf{R}^2$ let $\ell(u) = \{z \in \mathbf{R}^2 : z = tu, t > 0\}$ denote the half-line in the direction u and

$$\mathcal{C}(u, v) \equiv \{z \in \mathbf{R}^2 : |z| > 0,\ 0 < \arg_u(z) < \arg_u(v)\} \tag{1}$$

the cone swept anticlockwise from direction u to direction v. The closure of such a cone will be denoted $\bar{\mathcal{C}}(u, v)$. We give specific labels to the positive parts of the axes, $\mathcal{A}_i \equiv \ell(e_i)$ as we will consider them as blocks subsequently. It will also be convenient to define two special versions of the argument, one relative to each axis $\mathcal{A}_i$. Let $R : \mathbf{R}^2 \to \mathbf{R}^2$ be reflection in the line $z_1 = z_2$ i.e. $R(z_1, z_2) = (z_2, z_1)$ and define

$$\arg_1(z) = \arg_{e_1}(z)\ , \qquad \arg_2(z) = \arg_1\big(R(z)\big) \tag{2}$$

so $\arg_2(z)$ is the angle measured clockwise from e_2 to z.

A policy for controlling this discrete event system is a sequence $\Pi = \{\pi_n : n \geq 0\}$ of transition probabilities π_n from $\mathcal{H}_n$, the process history at time n, to $\mathcal{R}$, the set of regimes i.e. for any history α_0, η_0, $\ldots$, α_{n-1}, η_{n-1}, α_n the next action is selected according to the distribution $\pi_n(\alpha_0, \eta_0, \ldots, \alpha_n, \cdot)$. This definition includes non-stationary, non-Markov randomised policies though they offer no performance benefits when applied to stationary Markov processes, see e.g. Blackwell [Blackwell, 1965]. Let $\xi_i(n)$ denote the length of queue i at time n and $\xi(n) = (\xi_1(n), \xi_2(n))$. A policy Π along with an initial distribution for the queues determines a stochastic process $(\Xi, \Pi) = \{(\xi(n), \eta_n) : n \geq 0\}$ which will only be Markov when $\pi_n(\alpha_0, \eta_0, \ldots, \alpha_n, \cdot)$ is a distribution dependent only on α_n.

A policy Π which selects an action $a(\alpha)$ with probability 1 whenever the system state is α, where a is a map from $\mathbf{Z}_0^2$ to $\mathcal{R}$, is a deterministic stationary policy. Our main interest is in a class of these that we call *block pure policies*, denoted Π^b, where the state space $\mathbf{Z}_0^2$ is partitioned into a small number of disjoint blocks, always lines or cones, such that a is constant on each block $\mathcal{C}(u, v)$. We also investigate a generalisation of these, *block randomised* policies, denoted Π^r, where for each block the distribution $\pi_n^r(\alpha, \cdot)$ is the same at every state α in the block (so the Π^b are degenerate cases of the Π^r). With such policies the process (Ξ, Π^r) is Markov due to our assumptions about Poisson arrivals and exponential service times.

2.2 The queues and their mean drifts

The process (Ξ, Π) has bounded jumps, specifically $\pm e_i$ and $\pm(e_2 - e_1)$ and so all moments of its jump distributions exist under any policy but in this two dimensional case our results can be stated in terms of their first moments. For each regime η let

$$M^{\eta} = \mathbf{E}(\xi(n+1) - \xi(n) \mid \mathcal{H}_n, \pi_n = \eta) \tag{3}$$

denote the *mean drift* vector for any period when the policy selects regime η. We have, for $k = 1, \ldots, K$ at states $\alpha \in \mathbf{Z}_+^2 \equiv \{(x, y) \in \mathbf{Z}^2 : x > 0, y > 0\}$

$$\begin{aligned} M^{\eta} &= (M_1^{\eta}, M_2^{\eta}) \\ &= \begin{cases} \rho^{-1}\big(\lambda + \lambda_1 + \mu_{k2}p_{21}^{\eta} - \mu_{k1}p_{10}^{\eta},\ \lambda_2 + \mu_{k1}p_{12}^{\eta} - \mu_{k2}p_{20}^{\eta}\big), \ \eta = (k,1) \\ \rho^{-1}\big(\lambda_1 + \mu_{k2}p_{21}^{\eta} - \mu_{k1}p_{10}^{\eta},\ \lambda + \lambda_2 + \mu_{k1}p_{12}^{\eta} - \mu_{k2}p_{20}^{\eta}\big), \ \eta = (k,2) \end{cases} \end{aligned} \tag{4}$$

It is convenient to assume that when queue i is empty the policy selects a regime η chosen from among those with $\mu_{ki} = 0$ (this is equivalent to having non-idling servers). This ensures that equation (4) is also correct for histories ending in states $\alpha \in \mathcal{A}_1 \equiv \{(x, 0) : x > 0\}$ and $\alpha \in \mathcal{A}_2 \equiv \{(0, y) : y > 0\}$ for such service regimes. We will sometimes use the notation M' and M'' to denote mean drifts for the system under appropriate regimes for $\mathcal{A}_1$ and $\mathcal{A}_2$ respectively.

Now consider any policy Π allowing randomisation. The mean drift of our process Ξ under Π when the current state is $\alpha \in \mathbf{Z}_+^2$ is a 2-dimensional vector M^{Π} lying in the convex set

$$\mathcal{M} = \left\{ \sum_{\eta} p_{\eta} M^{\eta} \ : p_{\eta} \in [0, 1] \text{ and } \sum_{\eta} p_{\eta} = 1 \right\} \tag{5}$$

the convex hull of the regime mean drifts. The extreme points of $\mathcal{M}$ are a subset of the regime mean drifts M^{η}. When three or more of the M^{η} are distinct it may happen that the two-dimensional interior,

$$\mathrm{Int}_2(\mathcal{M}) \equiv \{z \in \mathcal{M} : B(z, \epsilon) \subset \mathcal{M} \text{ for some } \epsilon > 0\},$$

(where for $z \in \mathbf{R}_+^2$, $B(z, \epsilon) = \{z' \in \mathbf{R}^2 : |z - z'| < \epsilon\}$) is non-empty.

3 Classification of the system

The behaviour of the system depends on whether the convex set $\alpha + \mathcal{M}$ can be separated from the origin by a line through α. Any set of parameters for the process (Ξ, Π) falls into one the following four exclusive cases:

C1 $(0, 0) = \underline{0} \notin \mathcal{M}$ and there exists a state $\alpha \in \mathbf{Z}_+^2$ and a line

$$L_v(\alpha) \equiv \{\beta \in \mathbf{R}^2 : v^T(\beta - \alpha) = 0\} \tag{6}$$

with normal vector v through α separating $\alpha + \mathcal{M}$ from the origin $\underline{0}$. If there exists one such $\alpha \in \mathbf{Z}_+^2$ then there is an infinite cone of such α.

C2 $\underline{0} \notin \mathcal{M}$ and there exists no $\alpha \in \mathbf{Z}_+^2$ and line $L_v(\alpha)$ which separates $\alpha + \mathcal{M}$ from $\underline{0}$.

C3 $\mathrm{Int}_2(\mathcal{M})$ is non-empty, $\underline{0} \in \mathcal{M}$ and there exists no $\alpha \in \mathbf{Z}_+^2$, $v \in \mathbf{R}^2$ such that the line $L_v(\alpha)$ separates $\alpha + \mathrm{Int}_2(\mathcal{M})$ from the origin.

C4 $\underline{0}$ is a boundary point of $\mathcal{M}$ and either $\mathrm{Int}_2(\mathcal{M}) = \emptyset$ or the tangent line to $\alpha + \mathcal{M}$ through α separates the origin from $\alpha + \mathrm{Int}_2(\mathcal{M})$ for each α in a cone within $\mathbf{Z}_+^2$.

See Figure 1 for examples of C1-C4.

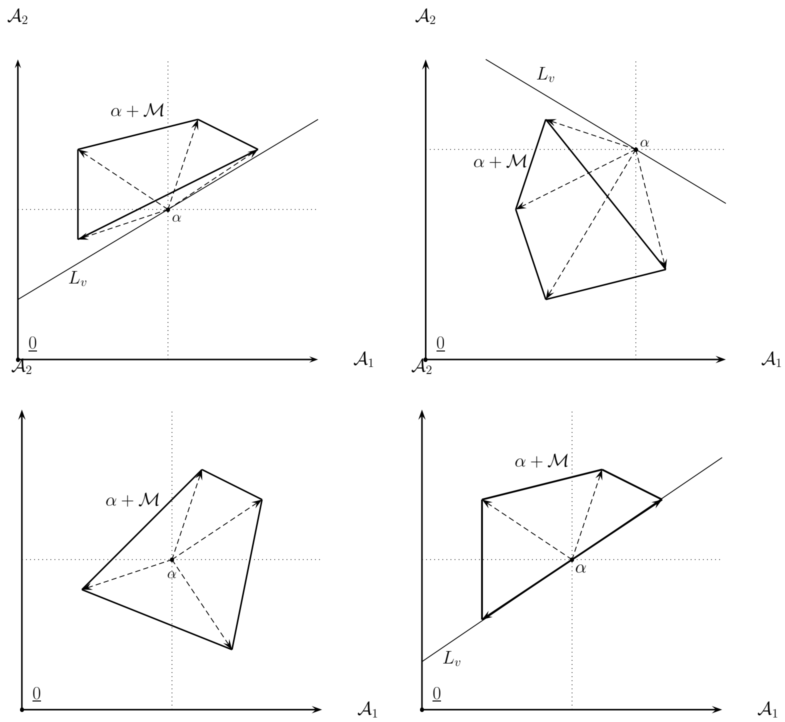

Fig. 1. From top left: C1, C2, and below C3, C4.

Note: the cases in C4 are critical but we will say very little about them in this paper.

We start stating our results by giving sufficient conditions for instability or stability respectively of the system under fully randomised controls in cases C1 and C2 respectively. Next we show that in case C3 there is always a block pure policy that makes (Ξ, Π^b) ergodic and we also show that randomisation allows the use of fewer blocks. Finally in this section we consider some situations with even lower levels of control.

3.1 Fully randomised controls

The following two results apply when even the most general policy Π is used to control the queueing system. They imply that in cases C1 and C2 the control policy used does not affect the stability or otherwise of the process.

Theorem 1 *If $\underline{0} \notin \mathcal{M}$ and there exists an $\alpha \in \mathbf{Z}_+^2$ and $v \in \mathbf{R}^2$ such that the line $L_v(\alpha)$, see (6), separates $\alpha + \mathcal{M}$ from the origin $\underline{0}$ then the process (Ξ, Π) is unstable, in the sense that the total number of queued jobs almost surely goes to ∞ linearly in time for any policy Π.*

The conditions of the theorem can be pictured in an alternative way. Specifically there exists a state $\alpha \in \mathbf{Z}_+^2$ such that the line segment from $\underline{0}$ to α does not intersect $\alpha + \mathcal{M}$ (it follows that if there is any such pair α, v then there is an infinite cone of points α' such that $L_v(\alpha')$ separates $\underline{0}$ and $\alpha' + \mathcal{M}$).

Theorem 2 *If $\underline{0} \notin \mathcal{M}$ and there is no $\alpha \in \mathbf{Z}_+^2$, $v \in \mathbf{R}^2$ such that $L_v(\alpha)$ separates $\alpha + \mathcal{M}$ from $\underline{0}$ then (Ξ, Π) is stable, in the sense that the total number of queued jobs remains bounded in mean, under every policy Π.*

The alternative description of the conditions here is that for every $\alpha \in \mathbf{Z}_+^2$ the line segment joining $\underline{0}$ to α intersects $\alpha + \mathcal{M}$. From this it follows there is some $v \in \mathbf{R}_+^2$ such that $\underline{0}$ and $\alpha + \mathcal{M}$ are in the same halfspace created by $L_v(\alpha)$.

3.2 Block controls

In case C3 it does make a difference which policy is used for running the system. In fact we can show that block pure policies Π^b with at most a handful of blocks are adequate to ensure stability of the process. Under policies of this type the process (Ξ, Π^b) is Markov so we can now talk about ergodicity and transience.

Theorem 3 *If $\underline{0} \in Int_2(\mathcal{M})$ then there is a block pure policy Π^b with at most five blocks such that the Markov chain (Ξ, Π^b) is ergodic.*

Theorems 2 and 3 imply the following result.

Corollary 1 *If $\underline{0}$ is a boundary point of $\mathcal{M}$, $Int_2(\mathcal{M})$ is non-empty and there exists no $\alpha \in \mathbf{Z}_+^2$, $v \in \mathbf{R}^2$ such that $L_v(\alpha)$ separates $\alpha + Int_2(\mathcal{M})$ from $\underline{0}$ then there is a policy Π^b with at most three blocks such that (Ξ, Π^b) is ergodic.*

In Theorem 3 the number of blocks required to achieve ergodicity can be reduced if block randomised policies Π^r are used.

Corollary 2 *If $\underline{0} \in Int_2(\mathcal{M})$ and a block randomised policy Π^r is used then at most four blocks are necessary to ensure that (Ξ, Π^r) is ergodic.*

Example [Foley and McDonald, 2001] consider a model which has fixed servers (their service rate drops to 0 when their queues are empty), no feedback and is strictly JSQ. Their stability criterion for $N = 2$ queues is that $\rho_{\max} \le 1$ where

$$\rho_{\max} = \max\{\lambda_1/\mu_1,\ \lambda_2/\mu_2,\ (\lambda + \lambda_1 + \lambda_2)/(\mu_1 + \mu_2)\ \}.$$

For the policy which sends the routable stream to the queue with minimum weighted work our model has two regimes depending upon where the routable traffic is sent and these have drift vectors

$$M^1 = \tfrac{1}{\rho}(\lambda + \lambda_1 - \mu_1, \lambda_2 - \mu_2) \text{ and } M^2 = \tfrac{1}{\rho}(\lambda_1 - \mu_1, \lambda + \lambda_2 - \mu_2).$$

As $\rho(M^2 - M^1) = (-\lambda, \lambda) \perp (1, 1)$ the line segment joining these two drift vectors has the form $z_1 + z_2 = (\lambda + \lambda_1 - \mu_1 + \lambda_2 - \mu_2)/\rho$ which can only intersect $\mathbf{R}^2_-$ when $\rho_{\max} < 1$. The case $\rho_{\max} = 1$ is critical and we see that our conditions are equivalent to those of Foley and McDonald in this case.

The simplicity of the classification based on the convex hull $\mathcal{M}$ confirms that this geometrical approach combined with the Lyapunov function method is a natural technique for studying stability of multi-queue systems though of course large deviations results like those [Foley and McDonald, 2001] are not achievable thisway.

3.3 Low levels of control

The results of [Fayolle *et al.*, 1995] can also be used to classify the process for any control policy that is block homogeneous for any small number of blocks. It soon becomes evident to anybody who attempts this that there are many ways for the process to remain stable and many more for it to be transient. To illustrate this we now spell out the possible behaviours of the queueing system with four blocks, specifically the axes $\mathcal{A}_1$, $\mathcal{A}_2$ and two cones, $\mathcal{C}_1 = \mathcal{C}(e_1, d) \cup \ell(d)$ and $\mathcal{C}_2 = \mathcal{C}(d, e_2)$ (see (1) for this notation), that partition $\mathbf{Z}^2_+$. The two cones are not assumed to be symmetric i.e. the vector $d \in \mathbf{R}^2_+$ need not be parallel to (1, 1).

We assume that in each of the $\mathcal{A}_i$ and $\mathcal{C}_i$, $i = 1, 2$ a single management regime is used (different blocks may have a common regime) with mean drift vectors M^1, M^2 in blocks $\mathcal{C}_1$, $\mathcal{C}_2$ respectively and M', M'' in blocks $\mathcal{A}_1$, $\mathcal{A}_2$ respectively. This assumption about and notation for the regimes on the $\mathcal{A}_i$ we will use in all further sections but the $\mathcal{C}_i$ are specific to this section.

We first label the M^i according to the angles φ_i they make relative to the axes $\mathcal{A}_i$, $i = 1, 2$. For each M^i angle $\varphi_i = 0$ is in the direction of $\mathcal{A}_i$ and φ_1 increases clockwise while φ_2 increases anticlockwise i.e. $\varphi_i = 2\pi - \arg_i(M^i)$. We label the directions of the M^i as **A** when $0 < \varphi_i < \pi$, **B** when $\pi \le \varphi_i \le \frac{3\pi}{2}$ and **D** when $\frac{3\pi}{2} < \varphi_i \le 2\pi$. The various cases of this model are labelled with *label of* M^1/*label of* M^2 so a label **B**/**A** means M^1 has a positive y and a negative x component and M^2 has x component negative with y of either sign. Figure 2 illustrates this labelling scheme for the directions of the M^i from origins α_i.

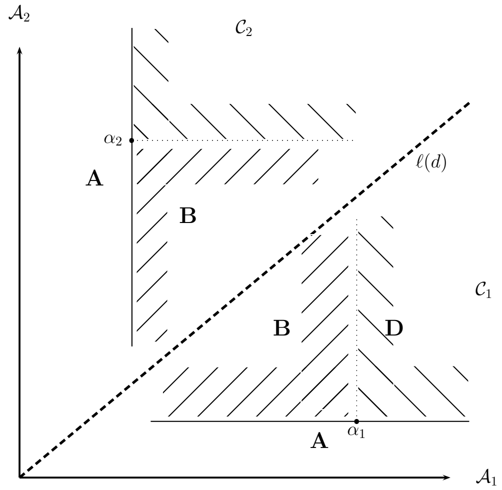

Fig. 2. Graphical explanation of the labels.

From the results in [Fayolle *et al.*, 1995] on the random walk in the positive quadrant, we have (in their terminology): (i) if a drift M^i has an $\mathbf{A}$ label then axis $\mathcal{A}_i$ is an *ergodic* face; (ii) that face will be *outgoing*, *ingoing* or *neutral* according to the sign of the *second vector field* (which is scalar in this case); (iii) if M^i has a $\mathbf{B}$ or $\mathbf{D}$ label then face $\mathcal{A}_i$ is transient and there is no second vector field. In this two dimensional case the sign of the second vector field depends only upon the angles of M' and M^1 for $\mathcal{A}_1$, M'' and M^2 for $\mathcal{A}_2$. As with the angles φ_i it is convenient to name the angles $\psi_1 = \arg_1(M')$, $\psi_2 = \arg_2(M'')$ that M', M'' make relative to axes $\mathcal{A}_1$, $\mathcal{A}_2$ respectively, so $\psi_i = 0$ is in the $\mathcal{A}_i$ direction and ψ_1 increases anticlockwise while ψ_2 increases clockwise. Now, following the sign of the second vector field, we modify the labels for M^i, $i = 1, 2$ to

$$\mathbf{A}^+ : \varphi_i + \psi_i < \pi, \quad \mathbf{A}^- : \varphi_i + \psi_i > \pi, \quad \mathbf{A}^0 : \varphi_i + \psi_i = \pi \tag{7}$$

Using this labelling system we can identify 25 different cases to deal with. It turns out that in many of the cases we get the same result for all choices of the two cones i.e. all slopes $d' \equiv d_2/d_1 \in (0, \infty)$ of the line $\ell(d)$ separating them. Theorem 4 classifies these invariant cases.

Theorem 4 *The system is*

(1) ergodic in cases $\mathbf{A}^-/\mathbf{A}^- \cup \mathbf{B}$, $\mathbf{B}/\mathbf{A}^-$, $\mathbf{B}/\mathbf{B}$ *with* $\left|\frac{M_x^1}{M_y^1}\right| > \left|\frac{M_x^2}{M_y^2}\right|$

(2) transient in cases $\mathbf{A}^+/\mathbf{A}\cup\mathbf{B}\cup\mathbf{D}$, $\mathbf{A}\cup\mathbf{B}\cup\mathbf{D}/\mathbf{A}^+$, $\mathbf{B}/\mathbf{B}$ *with* $\left|\frac{M_x^1}{M_y^1}\right| < \left|\frac{M_x^2}{M_y^2}\right|$, $\mathbf{D}/\mathbf{B}$, $\mathbf{B}/\mathbf{D}$, $\mathbf{D}/\mathbf{D}$;

(3) null recurrent in cases $\mathbf{A}^0/\mathbf{A}^0 \cup \mathbf{A}^+ \cup \mathbf{B}$, $\mathbf{A}^+ \cup \mathbf{B}/\mathbf{A}^0$, $\mathbf{B}/\mathbf{B}$ *with* $\left|\frac{M_x^1}{M_y^1}\right| = \left|\frac{M_x^2}{M_y^2}\right|$.

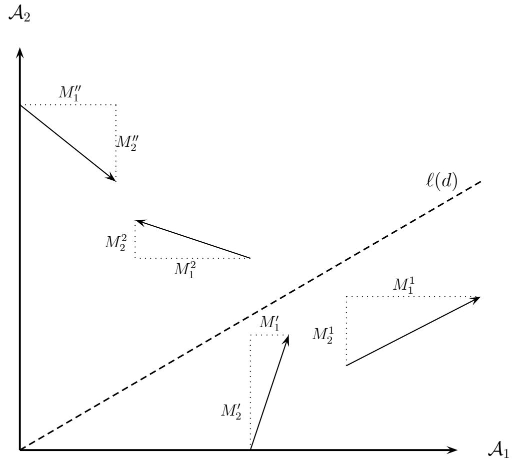

Fig. 3. Example of case $\mathbf{D}/\mathbf{A}^-$ where $\ell(d)$ is important.

For systems with no control over the service regimes there still may be some control over the routable traffic stream. The next theorem shows that there are sets of parameters such that a change to the slope of the switching line $\ell(d)$ can change Ξ from a transient to an ergodic process. We describe in detail only the case $\mathbf{D}/\mathbf{A}^0 \cup \mathbf{A}^-$, depicted in Fig. 3, as case $\mathbf{A}^0 \cup \mathbf{A}^-/\mathbf{D}$ is very similar. The relative slopes of M^1, $\ell(d)$ and M^2 are crucial so we label two key conditions:

E1: $M_2^1 < d' M_1^1$ (so $\ell(d)$ is steeper than M^1); E1$'$: $M_2^1 > d' M_1^1$;
E2: $-M_2^2 \leq d'(-M_1^2)$ (including cases with $M_2^2 \geq 0$ and implies $-M^2$ is not steeper than $\ell(d)$).

Theorem 5 *In case* $\mathbf{D}/\mathbf{A}^0 \cup \mathbf{A}^-$ *the ergodicity or non-ergodicity of the Markov chain* Ξ *also depends on the slope* $d' > 0$ *of the line* $\ell(d)$ *separating* $\mathcal{C}_1$ *and* $\mathcal{C}_2$ *as follows:*

(a) if E1 holds then Ξ *is transient,*
(b) if E1$'$ holds then Ξ*'s excursions into* $\mathcal{C}_1$ *have finite mean time and* Ξ *is*
  (i) ergodic if E2 holds and M^2 *is* $\mathbf{A}^-$ *or if E2 does not hold and* $(-M_2^2)M_1^1 < M_2^1(-M_1^2)$ *(so* M^1 *is steeper than* $-M^2$*);*
  (ii) null recurrent if E2 holds and M^2 *is* $\mathbf{A}^0$ *or if E2 does not hold and* $(-M_2^2)M_1^1 = M_2^1(-M_1^2)$*;*
  (iii) transient if E2 does not hold and $(-M_2^2)M_1^1 > M_2^1(-M_1^2)$.

The case $\mathbf{A}^0 \cup \mathbf{A}^-/\mathbf{D}$ *is simply the reflection of the above in the line* $\ell(1,1)$.

Note: this theorem says nothing about the cases where M^1 is parallel to $\ell(d)$ but in practice this will not be a major problem if the slope of the line $\ell(d)$ is under user control.

References

[Blackwell, 1965]D. Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, pages 226–235, 1965.

[Fayolle *et al.*, 1995]G. Fayolle, V. A. Malishev, and M. V. Menshikov. *Topics in the Constructive Theory of countable Markov Chains.* Cambridge University Press, Cambridge, 1995.

[Feng *et al.*, 2002]W. Feng, K. Adachi, and A. Kowada. A two-queue and two-server model with a threshold-based control service policy. *European Journal of Operational Research*, pages 593–611, 2002.

[Foley and McDonald, 2001]R. D. Foley and D. R. McDonald. Join the shortest queue: stability and exact asymptotics. *Annals of applied Probability*, pages 569–607, 2001.

[Kurkova, 2001]I. A. Kurkova. A load-balanced network with two servers. *Queueing Systems*, pages 379–389, 2001.

[MacPhee and Müller,]I.M. MacPhee and L.J. Müller. Stability classification of a controlled multi-queue system with many service and routing regimes. *submitted to Queueing Systems*.

[Serfozo, 1979]R.F. Serfozo. An equivalence between continuous and discrete time markov decision processes. *Operations Research*, 27:616–620, 1979.