



UNIVERSIDAD NACIONAL AUTÓNOMA DE MÉXICO

Doctorado en Ciencias Matemáticas

DYNAMICS OF STOCHASTIC NETWORKS AND FLOWS: CONVERGENCE,
COUPLING AND LARGE DEVIATIONS.

TESIS
QUE PARA OPTAR POR EL GRADO DE:
DOCTOR EN CIENCIAS

PRESENTA:
SERGIO IVÁN LÓPEZ ORTEGA

DIRECTOR Y CODIRECTOR DE LA TESIS
Dr. Pablo Ferrari, Universidad de Buenos Aires
Dr. Matthieu Jonckheere, Consejo Nacional de Investigaciones Científicas y
Técnicas, Argentina

MIEMBROS DEL COMITÉ TUTOR
Dra. Begoña Fernández Fernández, Facultad de Ciencias, UNAM
Dra. Ana Meda Guardiola, Facultad de Ciencias, UNAM

MÉXICO, D. F., Agosto de 2013.



Universidad Nacional
Autónoma de México

Dirección General de Bibliotecas de la UNAM

Biblioteca Central



UNAM – Dirección General de Bibliotecas
Tesis Digitales
Restricciones de uso

DERECHOS RESERVADOS ©
PROHIBIDA SU REPRODUCCIÓN TOTAL O PARCIAL

Todo el material contenido en esta tesis esta protegido por la Ley Federal del Derecho de Autor (LFDA) de los Estados Unidos Mexicanos (México).

El uso de imágenes, fragmentos de videos, y demás material que sea objeto de protección de los derechos de autor, será exclusivamente para fines educativos e informativos y deberá citar la fuente donde la obtuvo mencionando el autor o autores. Cualquier uso distinto como el lucro, reproducción, edición o modificación, será perseguido y sancionado por el respectivo titular de los Derechos de Autor.

Dynamics of stochastic networks and flows:
convergence, coupling and large deviations.

Sergio Iván López Ortega

Advisors: Inés Armendáriz, Pablo Ferrari & Matthieu Jonckheere

Version 5.0

Agradecimientos

Introducción

En esta tesis estudiamos la convergencia y grandes desviaciones de algunos modelos específicos de redes y flujos. La técnica de acoplamiento es usada a lo largo de toda la tesis.

En el Capítulo 1 estudiamos a la cola Browniana: un modelo de almacenamiento que toma valores continuos no negativos y que evoluciona a tiempo continuo. Para este modelo, Harrison y Williams [HW90] probaron que si el proceso de arribos tiene la distribución de un movimiento Browniano entonces el proceso de salidas efectivas tiene también la distribución de un movimiento Browniano. Es decir, obtuvieron una versión del Teorema de Burke [Bur56] en este contexto. En este trabajo probamos un resultado acerca de la convergencia bajo la operación de cola Browniana: si empezamos con un proceso continuo arbitrario que satisface algunas condiciones generales, y lo pasamos a través de una hilera de colas en tándem, ese proceso convergerá débilmente a un movimiento Browniano. Obtenemos este resultado bajo el supuesto de cargas iniciales independientes y con distribución exponencial para todas las colas. Este capítulo está basado en un trabajo conjunto con Inés Armendáriz y Pablo Ferrari.

En el Capítulo 2 estudiamos el modelo de crecimiento polinuclear (PNG, por sus siglas en inglés). Éste es un modelo físico para un cristal que crece capa por capa en una superficie unidimensional debido a posicionamientos aleatorios de partículas. Para el proceso definido en la línea aplicamos una técnica de acoplamiento similar a una utilizada en el contexto de colas con valores discretos en tiempo continuo por [MP95], con el fin de probar la convergencia del modelo. Además, obtenemos un resultado de convergencia local, para un proceso que contiene más información que el modelo PNG: se guarda la información generada por las trayectorias de las partículas, no sólo sus posiciones en un tiempo dado. Este capítulo está basado en un trabajo conjunto con Inés Armendáriz y Pablo Ferrari.

El Capítulo 3 es una introducción a las redes de ancho de banda compartido. Primero caracterizamos a los procesos de nacimiento y muerte reversibles y a las redes de Whittle,

después estudiamos el concepto de insensibilidad, y finalmente introducimos el modelo específico de Redes de ancho de banda.

En el capítulo final 4, probamos que la asignación de ancho de banda *Proportional fair* (PF) bajo un régimen estacionario y Markoviano comparte las mismas grandes desviaciones que la asignación *modified Proportional fair allocation (mPF)*, y por lo tanto las mismas que la asignación *Balanced Fairness*, debido al trabajo previo de Massoulié [Mas07]. Esto es probado en dos pasos: probamos ergodicidad geométrica para PF y mPF, y después controlamos la diferencia de los procesos asociados a PF y mPF processes empezando en el origen, utilizando argumentos de martingalas. Finalmente presentamos una prueba independiente para el caso de redes monótonas. Este resultado resuelve una conjetura propuesta por Massoulié [Mas07]. Este capítulo está basado en un trabajo conjunto con Matthieu Jonckheere aceptado para publicarse en la revista *Mathematics of Operations Research*. Una versión online está disponible en <http://arxiv.org/abs/1207.5908>

Bibliography

- [Bur56] P. Burke. The output of a queuing system. *Operations Research*, 4(6):pp. 699–704, 1956.
- [HW90] M. Harrison and R. Williams. On the quasireversibility of a multiclass brownian service station. *The Annals of Probability*, 18(3), 1990.
- [Mas07] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.
- [MP95] T. Mountford and B. Prabhakar. On the weak convergence of departures from an infinite series of $\cdot/M/1$ queues. *Ann. Appl. Probab.*, 5(1):121–127, 1995.

Structure

In this thesis we study the convergence and large deviations of some specific models of networks and flows. The coupling technique is used along the entire thesis.

In Chapter 1, our model is given by the Brownian queue: a storage model taking non negative real values and evolving in continuous time. For this model, Harrison and Williams [HW90] proved that if the arrival process has Brownian distribution, then the departure process has Brownian motion as well: that is, they obtained a Burke's theorem [Bur56] version in this context. We prove a result about the convergence under the queue operation: starting with a continuous arbitrary process satisfying some mild conditions, we show that when we pass it through a tandem network of queues, the resulting process converges weakly to a Brownian motion. This issue is proven under the assumption of independent and exponential initial workloads for all queues. This chapter is based on joint work with Inés Armendáriz and Pablo Ferrari.

In Chapter 2, we study the polynuclear growth model (PNG). This is a physical model for a crystal growing layer by layer, through the random deposition of particles. For the process defined on the line, we apply a coupling technique resembling the one used in the context of discrete-valued continuous time queues by Mountford and Prabhakar [MP95], in order to prove the convergence of the model. We also obtain a local convergence result for a process that keeps track of the paths of the particles up to a given time, instead of just their positions at this time. This chapter is based on joint work with Inés Armendáriz and Pablo Ferrari.

Chapter 3 is an introduction to Bandwidth-sharing networks. First, we characterize reversible birth and death processes and Whittle networks, then we study the concept of insensibility, and finally we introduce a specific model of Bandwidth-networks that we will study.

In Chapter 4, we prove that the proportional fair allocation (PF) under stationary and Markovian regime shares the same large deviations asymptotics as modified proportional

fair allocation (mPF), and hence the same as balanced fairness, due to the previous work of Massoulié [Mas07]. This is proved in two steps: we show geometric ergodicity for PF and mPF, and then we control the difference of PF and mPF processes starting from the origin, using martingales arguments. At the end, an independent proof is given for the case of monotone networks. This result solves a conjecture proposed by Massoulié [Mas07]. This chapter is based in a joint paper with Matthieu Jonckheere accepted for publication in the *Mathematics of Operations Research* journal. An online version is available at <http://arxiv.org/abs/1207.5908>

Bibliography

- [Bur56] P. Burke. The output of a queuing system. *Operations Research*, 4(6):pp. 699–704, 1956.
- [HW90] M. Harrison and R. Williams. On the quasireversibility of a multiclass brownian service station. *The Annals of Probability*, 18(3), 1990.
- [Mas07] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.
- [MP95] T. Mountford and B. Prabhakar. On the weak convergence of departures from an infinite series of $\cdot/M/1$ queues. *Ann. Appl. Probab.*, 5(1):121–127, 1995.

Contents

- Bibliography vi
- Bibliography viii
- 1 A non stationary Brownian analogue to Mountford-Prabhakar’s theorem 3**
- 1.1 Introduction 3
- 1.2 A Brownian analogue to Burke’s theorem 6
- 1.3 A non stationary Brownian analogue to Mountford-Prabhakar’s theorem . 11
- 1.4 Open problems 14
 - 1.4.1 A version of Mountford-Prabhakar’s theorem for general initial workloads 14
 - 1.4.2 The Totally Asymmetric Brownian Exclusion Process 15
- Bibliography 18
- 2 Convergence of the Polynuclear Growth Model dynamics on the line 21**
- 2.1 Introduction 21
 - 2.1.1 Hammersley process 21
 - 2.1.2 Polynuclear growth dynamics 25
- 2.2 Convergence of the PNG model on the line 30
- 2.3 Hammersley reflection 38
 - 2.3.1 Stationary distribution under the Hammersley reflection 38
- 2.4 Open problems 39
 - 2.4.1 Scaling 39
 - 2.4.2 Multiclass version 39
- Bibliography 40

3	Reversible birth and death processes, Whittle networks, Insensibility and Bandwidth-sharing networks	41
3.1	Reversible birth and death processes	41
3.1.1	Fixed birth rates and state dependent death rates processes	44
3.2	Whittle networks	46
3.2.1	Invariant measures for Whittle networks	48
3.3	Insensitivity	49
3.3.1	Residual times in renewal processes	50
3.3.2	Coupling construction for a processor sharing queue	51
3.3.3	Insensitivity of processor sharing networks	53
3.4	Bandwidth-sharing networks.	54
3.4.1	Some examples of bandwidth-sharing networks	56
3.4.2	The α -fairness allocations	57
3.4.3	Gradient allocations.	58
3.4.4	Monotone networks with general service time distributions.	60
	Bibliography	61
4	Large deviations for Proportional Fairness allocation	63
4.1	Why to calculate large deviations for Proportional fairness allocation?	63
4.1.1	Insensitivity in telecommunication systems	63
4.1.2	The Proportional fairness election.	65
4.2	The Freidlin-Wentzell approach	67
4.3	Large deviations for Markovian dynamics.	70
4.4	Large deviations for monotone networks with general service time distributions.	79
4.5	Open problems	82
4.5.1	Large deviations for discriminatory Processor Sharing	82
4.5.2	Large deviations approximations by reversible processes	83
	Bibliography	83

Chapter 1

A non stationary Brownian analogue to Mountford-Prabhakar's theorem

1.1 Introduction

In 1956, Burke [Bur56] obtained one of the most fundamental results for queueing theory. The first part of Burke's theorem states that given an arrival process Poisson with rate $\lambda < 1$, and an independent service Poisson process with rate 1, which together define a $M/M/1$ queue, then the departure process is again Poisson with the parameter λ . The second part of Burke's theorem states a *factorization property*: the length of the queue at time t is independent of future arrivals and past departures. After [Bur56] was published, several extensions of this result have followed, such as a version of Burke's theorem for multiclass queues with priority [MP10], another related to the Asymmetric Simple Exclusion Process [FF94], and a version applied to dual processes in inventory and queueing processes [DMO03].

A tandem queue is a system of queues where there is an arrival process A^1 , and a sequence $\{S^n\}_{n \geq 1}$ of service processes, all independent. The system is defined recursively. The initial queue is fed from the arrival process A^1 , and has departures determined by the service process S^1 . For $n \geq 2$, the arrival process for the n -th queue is defined as the departure process of the $(n - 1)$ -th queue and the departures are determined by the service process S^n .

When the initial arrival process is Poisson, Burke's theorem allows us to treat a tandem system of queues at any fixed time as if the queues acted independently. For example, take a two nodes system of tandem queues, Q_t^1 and Q_t^2 , with arrivals Poisson(λ), $\lambda < 1$,

service processes $\text{Poisson}(1)$, all independent, and sample the initial length of each queue from its stationary measure, namely a $\text{Geometric}(\lambda)$ random variable. Because of the first part of Burke's theorem, the departure process of Q_t^1 is a $\text{Poisson}(\lambda)$ process. Moreover, due to the second part of Burke's theorem, the departure process of the first queue prior to time t is independent of the length of the queue at time t . Then the Q_t^2 variable is independent from the variable Q_t^1 , and it follows that the invariant measure of the system is a product measure. The factorization property from Burke's theorem has thus enabled the analysis of more complex systems, for example Jackson's networks [Asm03]. Using similar ideas, a stationary version of a tandem queue system (a double-infinite line of queues) can be constructed and it has indeed the distribution of the holes between particles in the stationary version of the totally asymmetric simple exclusion process (TASEP) [FM06].

When the initial arrival process is not Poisson, a natural question is whether it is possible to prove convergence to the stationary distribution. Assuming an existence result, Anantharam [Ana93] proved the uniqueness of a stationary ergodic fixed point for the $\cdot/M/K$ queue. Next, Mountford and Prabhakar proved the attractiveness of the Poisson distribution in the class of ergodic stationary point processes on the line with rate $\lambda < 1$ [MP95]. In order to obtain this result, they used a coloring coupling technique based on an argument of Ekhaus and Gray (unpublished, cited by [MP95]).

In [OY01] some Brownian analogues of Burke's theorem and further generalizations for geometric functionals of the Brownian Motion were introduced. The proof of some of these versions goes back to [HW90] in the context of multiclass stations, and it relies on weak convergence arguments, or alternatively, on some properties of the Brownian Motion.

One way to construct a queue is by the use of the *reflective mapping*: Take two Poisson processes A, S with rates $\lambda < 1$ and 1 respectively. Construct a simple random walk X by increasing 1 each time there is a mark of A , and decreasing 1 each time there is a mark of S . This random walk with negative drift is called the *free process*. Now, take the free process X starting at 0 given by some arbitrary non-negative integer (called the *initial workload* of clients) and map it to another process $Q = \mathcal{R}(X)$ in the following way: Q has the same jumps as X , except in case where there is a negative jump of X that would have made Q cross below the 0 axis. In that case, the jump is ignored. See Figure 1.1.




Figure 1.1: Construction of a queue by the reflective mapping with empty initial workload.

The reflective mapping is the procedure of mapping a free function f to a non-negative valued function $\mathcal{R}(f)$ that follows the trajectory of f , such that a completely inelastic collision occurs at the barrier 0. This operator is very different of the absolute value operator that maps a function f to its absolute value $|f|$, as in this case where a completely elastic collision happens at 0. This mapping can be generalized to the setting of càdlàg functions, and we will give the formal definition and its properties in Section 1.2. The general definition is used to study properties of reflected Lévy processes, which is an entire area of study in probability, see [Asm03] for an introduction. In the case where the free function is continuous, the resulting $\mathcal{R}(f)$ is continuous too [Har90], and can be used to represent a storage flow system taking continuous values that increase and decrease continuously above 0. A well known example is the case when the free process is given by a Brownian motion with negative drift. In this case $\mathcal{R}(f)$ is the Brownian queue (also called reflected Brownian motion or regulated Brownian motion), see [Har90] for an introduction.

In this work we present analogues of both Burke's and Mountford-Prabhakar's theorems, for a continuous valued queue which has a Brownian motion as the service process. We first prove a version of Burke's theorem. This was first obtained by Harrison and Williams [HW90] in the context of Brownian networks using weak convergence. Another proof using Pitman's Representation theorem and path properties of the Brownian

motion was obtained by O’Connell and Yor [OY01] (that proof was originally suggested by Harrison, as O’Connell and Yor [OY01] remark). We present here an independent and direct proof using weak convergence. It turns out to be a particular case of [HW90]. We also show that the Brownian motion is an attractor, under the Brownian queue operator dynamics, in a wide class of continuous-valued processes with a particular set of initial conditions: all the tandem queues begin with independent workloads, identically distributed as the stationary distribution of the Brownian queue, namely an exponential random variable (see [Har90]). The coloring coupling used in the discrete valued case [MP95] is no longer suitable and we introduce an ad-hoc coupling technique that takes advantage of simple path properties of the Brownian motion.

This chapter is organized as follow: In Section 1.2 we introduce the reflective mapping to define the continuous queue, and prove the Brownian analogue to Burke’s theorem, in Section 1.3 we prove a version of Mountford-Prabhakar’s theorem in our context, and in the Section 1.4 we discuss some related open problems.

1.2 A Brownian analogue to Burke’s theorem

In this section we introduce the queueing operation when both arrival and service processes are defined by continuous functions.

Definition 1. *Denote the space of real càdlàg functions $f : [0, \infty) \rightarrow \mathbb{R}$ by $D[0, \infty)$. We define the reflective Skorokhod mapping (sometimes called regulator mapping) as the operator $\mathcal{R} : D[0, \infty) \rightarrow D[0, \infty)$ given by*

$$\mathcal{R}(f)(t) = f(t) - \inf_{0 \leq s \leq t} \{f(s) \wedge 0\}.$$

Let us state some fundamental properties of this mapping:

Lemma 1. *Let $\mathcal{R} : (D[0, \infty), \tau_D) \rightarrow (D[0, \infty), \tau_D)$ be the reflective mapping, where τ_D is the Skorokhod topology. Then*

(a) *For every function f and constant $c \geq 0$, $\mathcal{R}(cf) = c\mathcal{R}(f)$ (homothecy property).*

(b) *The reflective mapping is Lipschitz continuous, with Lipschitz constant equal to 2.*

Proof.

The proof of (a) follows by trivial verification. To prove (b), we prove first an observation about the continuity of the supremum mapping:

Observation 1. Let us define the supremum mapping $\mathcal{S} : (D[0, \infty), \tau_D) \rightarrow (D[0, \infty), \tau_D)$ by

$$\mathcal{S}(f)(t) := \sup_{0 \leq s \leq t} f(s).$$

Then, for fixed $T \geq 0$, the supremum mapping is Lipschitz continuous in the uniform topology on $[0, T]$ with Lipschitz constant equal to 1.

Proof.

To prove this observation, let $f, g \in D[0, \infty)$ be two functions, $t \in [0, T]$ fixed, and let us assume without loss of generality that $\mathcal{S}(f)(t) \geq \mathcal{S}(g)(t)$. Then

$$\mathcal{S}(f)(t) - \mathcal{S}(g)(t) \leq \mathcal{S}(f)(t) - g(s) \quad \forall 0 \leq s \leq t.$$

By definition, we know that there exist $u \in [0, t]$ such that $f(u) = \mathcal{S}(f)(t)$. Then, for any $t \in [0, T]$ we have found $u = u(t) \in [0, T]$ such that

$$|\mathcal{S}(f)(t) - \mathcal{S}(g)(t)| \leq |f(u) - g(u)|.$$

By taking the particular case where f and g are two different constant, we see that the bound is tight, so the observation follows. □

Now, note that

$$\begin{aligned} \|\mathcal{R}(f)(t) - \mathcal{R}(g)(t)\|_{[0, T]} &= \|f(t) + \sup_{0 \leq s \leq t} \{(-f(s)) \vee 0\} - [g(t) + \sup_{0 \leq s \leq t} \{(-g(s)) \vee 0\}]\|_{[0, T]} \\ &\leq \|f(t) - g(t)\|_{[0, T]} \\ &\quad + \|\sup_{0 \leq s \leq t} \{(-f(s)) \vee 0\} - \sup_{0 \leq s \leq t} \{(-g(s)) \vee 0\}\|_{[0, T]}, \\ &\leq \|f(t) - g(t)\|_{[0, T]} + \|\{(-f(t)) \vee 0\} - \{(-g(t)) \vee 0\}\|_{[0, T]}, \\ &\leq 2 \|f(t) - g(t)\|_{[0, T]}, \end{aligned}$$

where the second inequality follows from the observation above and the third one holds because $\|h^+ - i^+\|_{[0, T]} \leq \|h - i\|_{[0, T]}$ for every pair h, i of continuous real functions.

To see that the bound is tight, define $f \equiv 0$ and $g := -1_{[\frac{1}{3}, \frac{1}{2})} + 1_{[\frac{1}{2}, 1]}$, both on $[0, 1]$. Then $\mathcal{R}(f) \equiv 0$ and $\mathcal{R}(g) = (2) 1_{[\frac{1}{2}, 1]}$ so $\|f - g\|_{[0, 1]} = 1$ and $\|\mathcal{R}(f)(t) - \mathcal{R}(g)(t)\|_{[0, 1]} = 2$.

The technical details concerning the extension from the continuity of this mapping in the uniform topology to the continuity in the Skorokhod topology, can be found in Whitt [Whi02], pp 439. □

We use the reflective mapping to define the reflection of one function on another one.

Definition 2. Let $f, g \in D[0, \infty)$ be such that $f(0) > g(0)$. Define

$$U_g(f)(t) = f(t) - \inf_{0 \leq s \leq t} \left\{ (f(s) - g(s)) \wedge 0 \right\},$$

and

$$L_f(g)(t) = g(t) + \inf_{0 \leq s \leq t} \left\{ (f(s) - g(s)) \wedge 0 \right\}.$$

In this case, $U_g(f)$ is a function that has f as free process and g as a lower barrier, and $L_f(g)$ is the function that has g as free process and f as an upper barrier. Note that $U_g(f) \geq f$ and $L_f(g) \leq g$, see Figure 1.2. We immediately obtain that

$$U_g(f) = g + \mathcal{R}(f - g), \quad L_f(g) = f - \mathcal{R}(f - g).$$




Figure 1.2: Reflection of a function on another one.

We are ready to introduce the queueing operation in the context of càdlàg functions.

Definition 3. Let $f, g \in D[0, \infty)$ be such that $f(0) > g(0)$. We will say that f is the **arrival process**, g the **service process**, and $f(0) - g(0)$ the **initial workload** of the queue. We define the queueing operator $\mathcal{Q} : D[0, \infty)^2 \rightarrow D[0, \infty)$ by

$$\mathcal{Q}(f, g) = L_f(g) = f - \mathcal{R}(f - g),$$

and we will call $L_f(g)$ the **departure process** of the queue. Other important processes are the **queue length process** defined as $f - L_f(g)$, and the **free process** defined as $f - g$.

Note that due to the continuity of the reflective mapping (Lemma 1) the queueing operator is continuous. See Figure 1.3 for an illustration.




Figure 1.3: The càdlàg queue.

The above definition matches the definition given by O’Connell and Yor [OY01] of the stationary version of the Brownian queue for $t \geq 0$ when the initial workload $f(0) - g(0)$ is chosen with the stationary distribution of the queue, as Norros [SN01] pointed out. As we mentioned in Section 1.1, we can represent a standard queue as the result of applying the reflective mapping to the random walk obtained as free process when we have Poisson arrival and service processes. Therefore, the reflective mapping provides a unified way to characterize a queue in both the discrete and continuous valued settings.

Now we turn our attention to a similar result to Burke’s theorem, where we have Brownian motions as arrival and service processes. Let us denote by \Rightarrow the weak convergence of processes, see [Bil68] for a classical reference on this topic. We will need the following corollary of the Functional Central Limit theorem for Poisson processes.

Lemma 2. *Let $\{P^n\}_{n \in \mathbb{N}}$ be a sequence of Poisson processes with rate $r_n > 0$. Suppose that $r_n \rightarrow r \in (0, \infty)$. Then*

$$\frac{P^n(nt) - r_n n t}{\sqrt{n}} \Rightarrow \sqrt{r} B(t),$$

where $\{B(t)\}_{t \geq 0}$ is a standard Brownian motion.

Proof.

Let $P(t)$ be a Poisson process with intensity 1. Let us denote equality in law by $=^d$.

Note that $\{P^n(t) : t \geq 0\} =^d \{P(r_n t) : t \geq 0\}$. Then,

$$\frac{P^n(nt) - n r_n t}{\sqrt{n}} =^d \frac{P(n r_n t) - n r_n t}{\sqrt{n}} = \left(\frac{P((n r_n)t) - (n r_n)t}{\sqrt{n r_n}} \right) \sqrt{r_n}.$$

The result now follows by an application of the functional CLT.

□

We state now the main result of this section.

Theorem 1. *Let B_t^1, B_t^2 be standard Brownian motions, \mathcal{E} an exponential variable with mean $\frac{1}{c}$, and x a real number. Let us assume that the random elements defined are independent. Then*

1. $D_t = \mathcal{Q}(B_t^1 + x, B_t^2 + ct + x - \mathcal{E})$ has the law of a standard Brownian motion starting at $x - \mathcal{E}$.
2. $\{D_s : 0 \leq s < t\}$ and Q_t are independent.

Proof.

1) Define $\lambda_n := 1 - \frac{c}{\sqrt{n}}$, so that $\sqrt{n}(1 - \lambda_n) = c$. For each $n > [c^2] + 1$, let $A_n(t)$ be a Poisson process with parameter λ_n , $S(t)$ a Poisson process with parameter 1, and G_n a Geometric random variable with parameter λ_n , all independent from each other.

Next, let

$$A_n^*(t) := \frac{A_n(nt) - nt\lambda_n}{\sqrt{n}}.$$

Applying Lemma 2, $A_n^*(t)$ converges weakly to a Brownian motion. We also define

$$S_n^*(t) := \frac{S(nt) - nt\lambda_n}{\sqrt{n}}.$$

Then

$$S_n^*(t) = \frac{S(nt) - nt}{\sqrt{n}} + ct \Rightarrow B(t) + ct,$$

where $B(t)$ is a Brownian motion, because $\sqrt{n}(1 - \lambda_n) = c$. Finally, let $G_n^* := \frac{G_n}{\sqrt{n}}$, so that G_n^* converges to an exponential mean $\frac{1}{c}$ random variable.

Since A_n, S_n and G_n are independent, we conclude that $(A_n^*(t) + x, S_n^*(t) + x - G_n^*)$ converges weakly to $(B^1(t) + x, B^2(t) + ct + x - \mathcal{E})$, where B^1 and B^2 are independent standard Brownian motions and \mathcal{E} is an exponential random variable of mean $\frac{1}{c}$. By the Skorokhod continuity of the queueing operator, we can apply the continuous mapping theorem (see [Whi02]) to conclude that $\mathcal{Q}(A_n^*(t) + x, S_n^*(t) + x - G_n^*)$ converges weakly to $\mathcal{Q}(B^1(t) + x, B^2(t) + ct + x - \mathcal{E})$.

On the other hand, let $D_n(t) := \mathcal{Q}(A_n(t), S(t) - G_n)$. By Burke's theorem, D_n has the law of a Poisson process with rate λ_n . Define the process

$$D_n^*(t) := \frac{D_n(nt) - nt\lambda_n}{\sqrt{n}}.$$

By Lemma 2 again, D_n^* converges weakly to an standard Brownian motion. Since

$$\begin{aligned} D_n^*(t) &= \frac{\mathcal{Q}(A_n(nt), S(nt) - G_n) - nt\lambda_n}{\sqrt{n}} \\ &= \mathcal{Q}\left(\frac{A_n(nt) - nt\lambda_n}{\sqrt{n}}, \frac{S(nt) - nt\lambda_n - G_n}{\sqrt{n}}\right) = \mathcal{Q}(A_n^*(t), S_n^*(t) - G_n^*), \end{aligned}$$

by the uniqueness of the weak limit for stochastic processes, we conclude that $\mathcal{Q}(B^1(t) + x, B^2(t) + ct + x - \mathcal{E})$ is equal in distribution to $B(t) + ct + x - \mathcal{E}$.

2) Define Q_n as the length of the queue process constructed from the arrival A_n and S service processes. From Burke's theorem, we have that $\{D_n(s) : 0 \leq s < t\}$ is independent of $Q_n(t)$.

On the other hand, the queue length process Q constructed from the arrival $B^1(t) + x$ and the service $B^2(t) + ct + x - \mathcal{E}$ processes, has the explicit formulation

$$Q(t) = B^1(t) + x - L_{B^1(t)+x}(B^2(t) + ct + x - \mathcal{E}),$$

so by an argument similar to the one applied to prove 1), we conclude that Q_n converges weakly to Q and the independence of $\{D_n(s) : 0 \leq s < t\}$ and $Q_n(t)$ is inherited by D and $Q(t)$.

□

1.3 A non stationary Brownian analogue to Mountford-Prabhakar's theorem

We now turn our attention towards obtaining a Brownian analogue of Mountford-Prabhakar's theorem. We first present a monotonicity result for the reflection dynamics. Here $\|\cdot\|_{[0,T]}$ denotes the supremum norm on $[0, T]$.

Lemma 3. *Denote the space of continuous real functions by $C[0, \infty)$. Let $f^1, f^2, g \in C[0, \infty)$ be such that $f^1(0), f^2(0) > g(0)$, and let $L_{f^1}(g), L_{f^2}(g)$ be as in Definition 2. Then, for any $T > 0$,*

$$\|L_{f^1}(g) - L_{f^2}(g)\|_{[0,T]} \leq \|f^1 - f^2\|_{[0,T]}.$$

Proof.

We have

$$\begin{aligned} \|L_{f^1}(g) - L_{f^2}(g)\|_{[0,T]} &= \| [g_t + \inf_{0 \leq s \leq t} \{(f_s^1 - g_s) \wedge 0\}] - [g_t + \inf_{0 \leq s \leq t} \{(f_s^2 - g_s) \wedge 0\}] \|_{[0,T]} \\ &= \| \sup_{0 \leq s \leq t} \{(g_s - f_s^1) \vee 0\} - \sup_{0 \leq s \leq t} \{(g_s - f_s^2) \vee 0\} \|_{[0,T]}. \end{aligned}$$

Then

$$\begin{aligned} \|L_{f^1}(g) - L_{f^2}(g)\|_{[0,T]} &\leq \| (g_t - f_t^1) \vee 0 - (g_t - f_t^2) \vee 0 \|_{[0,T]} \\ &\leq \| f^1 - f^2 \|_{[0,T]}. \end{aligned}$$

The first inequality follows from the Observation 1 and the second one holds because $\|h^+ - i^+\|_{[0,T]} \leq \|h - i\|_{[0,T]}$ for every pair h, i of continuous real functions. □

Next theorem is the main result of this chapter.

Theorem 2. *Let A^0 be a continuous process with paths $A^0(\cdot, \omega) : [0, \infty) \rightarrow \mathbb{R}$ that do not explode in finite time almost surely, and $A^0(0, \omega) = 0$. Let $\{W^n\}_{n \in \mathbb{N}}$ be a family of standard Brownian motions, $\{\mathcal{E}^n\}_{n \in \mathbb{N}}$ a family of exponential random variables with common mean $\frac{1}{c}$, $c > 0$, all independent. We define recursively the sequence of processes*

$$A^n = \mathcal{Q}\left(A^{n-1}, W^n + ct - \sum_{i=1}^n \mathcal{E}^i\right), \quad n \geq 1,$$

where \mathcal{Q} is the queueing operator (Definition 3). Then $A^n + \sum_{i=1}^n \mathcal{E}^i$ converges weakly to a Brownian motion.

The proof relies on a coupling argument: we show that if different arrival processes are run through the same services, the resulting trajectories are eventually locally coupled. Since we know that there exists an stationary distribution under the reflecting dynamics of the tandem queues, given by the Theorem 1, we conclude the convergence result. The idea of the proof is the next: we prove that if we begin with two different arrival process, at some step of the tandem dynamics we will have positive workload during a fixed period. Then the departures will coincide with the services in such period, and since we are using the same services processes, the departures of both systems will coincide. This coupling persists in the following iterations of the reflective dynamics.

Fix $T > 0$ and let B^0 be an independent Brownian motion. Define

$$B^n = \mathcal{Q}\left(B^{n-1}, W^n + ct - \sum_{i=1}^n \mathcal{E}^i\right), \quad \forall n \geq 1.$$

That is, we apply the tandem queues dynamics to the process B^0 using the same service processes $\{W^n + ct\}_{n \in \mathbb{N}}$ and initial workloads $\{\mathcal{E}^n\}_{n \in \mathbb{N}}$. By Theorem 1, $B^n + \sum_{i=1}^n \mathcal{E}^i$ has the law of a Brownian motion for every n , so in order to get convergence it is enough to prove that the trajectories of A^n y B^n eventually couple on $[0, T]$.

We will use the following version of Borel-Cantelli's lemma.

Lemma 4. *Let $(\Omega, \mathcal{F}, \mathbb{P})$ a probability space, $\{\mathcal{F}_n\}_{n \in \mathbb{N}}$ a filtration such that $\mathcal{F} = \bigcup_{n \in \mathbb{N}} \mathcal{F}_n$ and $O_n \in \mathcal{F}_n$, $n \in \mathbb{N}$. Then*

$$\{O_n \text{ i.o.}\} = \left\{ \sum_{n \in \mathbb{N}} \mathbb{P}(O_{n+1} | \mathcal{F}_n) = \infty \right\} \quad \text{a.s..}$$

Proof. See [Kal02]. □

Proof of Theorem 2.

Let us first make the observation that if

$$W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i \leq A_t^n, B_t^n, \quad \forall t \in [0, T],$$

then $A_t^{n+1} = B_t^{n+1} = W_t^n + ct - \sum_{i=1}^{n+1} \mathcal{E}^i$, for $t \in [0, T]$, and the trajectories couple. If the two queues have positive workload during a fixed period, then the departures coincide with the services for both queues and this coupling persists in the following iterations of the tandem dynamics.

For $n \in \mathbb{N}$, define the event

$$O_n := \left\{ \omega \in \Omega : W_t^n + ct - \sum_{i=1}^n \mathcal{E}^i \leq A_t^{n-1}, B_t^{n-1}, \forall t \in [0, T] \right\},$$

and the σ -algebra $\mathcal{F}_n := \sigma(\{A^0, B^0, W^i, \mathcal{E}^i : i \leq n\})$. Define for $n \in \mathbb{Z}_+$, $\delta_n := \|A_t^n - B_t^n\|_{[0, T]}$. By the observation above, O_n is a coupling event. And by Lemma 4 it is enough to prove that $\sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) = \infty$.

Since the process A_0 does not explode in finite time a.s. we have that $\delta_0 < \infty$ a.s., hence:

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) &= \left(\sum_{k=1}^{\infty} 1_{\{k-1 \leq \delta_0 < k\}} \right) \left(\sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) \right) \\ &= \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} 1_{\{k-1 \leq \delta_0 < k\}} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) \\ &= \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} 1_{\{k-1 \leq \delta_0 < k\}} | \mathcal{F}_n). \end{aligned}$$

By Lemma 3, we know that $\delta_n(\omega) \leq \delta_0(\omega)$ for all $\omega \in \Omega$, and then:

$$\sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) = \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} 1_{\{k-1 \leq \delta_0 < k\}} 1_{\{\delta_n < k\}} | \mathcal{F}_n).$$

Note now that $\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - \delta_n, \forall t \in [0, T]\} \subseteq O_{n+1}$. Hence:

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{E}(1_{O_{n+1}} | \mathcal{F}_n) &\geq \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{E}(1_{\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - \delta_n, \forall t \in [0, T]\}} 1_{\{k-1 \leq \delta_0 < k\}} 1_{\{\delta_n < k\}} | \mathcal{F}_n) \\ &\geq \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{E}(1_{\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - k, \forall t \in [0, T]\}} 1_{\{k-1 \leq \delta_0 < k\}} 1_{\{\delta_n < k\}} | \mathcal{F}_n) \\ &= \sum_{k=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{E}(1_{\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - k, \forall t \in [0, T]\}} 1_{\{k-1 \leq \delta_0 < k\}} | \mathcal{F}_n) \\ &= \sum_{k=1}^{\infty} \left[1_{\{k-1 \leq \delta_0 < k\}} \sum_{n=1}^{\infty} \mathbb{E}(1_{\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - k, \forall t \in [0, T]\}} | \mathcal{F}_n) \right]. \end{aligned}$$

Define

$$X_n^k := \mathbb{E}(1_{\{W_t^{n+1} + ct - \sum_{i=1}^{n+1} \mathcal{E}^i < B^n - k, \forall t \in [0, T]\}} | \mathcal{F}_n).$$

By the Brownian version of Burke's Theorem 1, we have that B^n is a Brownian motion for all n and then the random variables $\{X_n^k\}_{n \in \mathbb{N}}$ are identically distributed. Moreover, since the reflective dynamics are Markovian with respect to the n -th step in the tandem queue, we have that $\{X_n^k\}_{n \in \mathbb{N}}$ are independent. By simple properties of the Brownian motion, they are strictly positive random variables. Then its sum $\sum_{n=1}^{\infty} X_n^k$ almost surely diverges for every k and we are done by the almost sure finiteness of δ_0 . \square

1.4 Open problems

1.4.1 A version of Mountford-Prabhakar's theorem for general initial workloads

We are interested in proving a general version of Theorem 2. Suppose we have an ergodic arrival process A^0 , which we run as before through of a system of Brownian services, and sample the initial workloads according to the stationary distribution of the n -th queue. We would then obtain a stationary version of the whole system of tandem queues, defined for all times. We would want to prove Theorem 2 in this setting.

1.4.2 The Totally Asymmetric Brownian Exclusion Process

A question related to the previous results is whether it is possible to construct a continuous analogue of the Totally Asymmetric Simple Exclusion Process (TASEP) in the real line, using the scaling limit for the discrete queues introduced in the last section. We first state the following classical result.

Theorem 3 (Kolmogorov's extension theorem). *Let A be a non empty set. Let us suppose that for each $k \in \mathbb{N}$ and subset $\{a_1, a_2, \dots, a_k\} \subseteq A$, we have a probability measure $\nu_{(a_1, a_2, \dots, a_k)}(\cdot)$ in $(\mathbb{R}^k, \mathcal{B}(\mathbb{R}^k))$. Let*

$$\mathcal{M}_A = \{\nu_{(a_1, a_2, \dots, a_k)} : a_1, \dots, a_k \in A, k \geq 1\}$$

be the family of these measures and assume that it satisfies the next conditions:

1. For each $C_1, \dots, C_k \in \mathcal{B}(\mathbb{R})$, $a_1, \dots, a_k, a_{k+1} \in A$, $1 \leq k < \infty$,

$$\nu_{(a_1, \dots, a_k, a_{k+1})}(C_1 \times \dots \times C_k \times \mathbb{R}) = \nu_{(a_1, \dots, a_k)}(C_1 \times \dots \times C_k).$$

2. For every permutation σ in the symmetric group \mathcal{S}_k ,

$$\nu_{(a_{\sigma(1)}, \dots, a_{\sigma(k)})}(C_{\sigma(1)} \times \dots \times C_{\sigma(k)}) = \nu_{(a_1, \dots, a_k)}(C_1 \times \dots \times C_k).$$

Then there exists a probability measure P in $(\Omega := \mathbb{R}^A, \mathcal{F} := (\mathcal{B}(\mathbb{R}))^A)$ and an \mathbb{R} -valued stochastic process $\{X_a : a \in A\}$, such that $\nu_{(a_1, a_2, \dots, a_k)}(\cdot) = P[(X_{a_1}, \dots, X_{a_k})^{-1}(\cdot)]$.

The proof can be found in Kallenberg [Kal02].

□

We now define an exclusion process $\{\eta^i(t) : i \in \mathbb{Z}, t \geq 0\}$ where $\eta^i(t)$ denotes the position of particle i at time t . Let U be a random variable with arbitrary distribution and $\{\mathcal{E}_j\}_{j \in \mathbb{Z} \setminus \{0\}}$ a collection of independent exponential random variables with media $\frac{1}{\lambda} > 0$. We define the sequence of random variables $\{X_i\}_{i \in \mathbb{Z}}$ by $X_0 = U$, $X_i = \mathcal{E}_i$ for $i > 0$ and $X_i = -\mathcal{E}_i$ for $i < 0$. Set $\eta^i(0) = \sum_{r=0}^i X^r$, so that the initial configuration of particles is given by a rate λ Poisson process on the line. Let $\{B^i\}_{i \in \mathbb{Z}}$ be a sequence of Brownian motions independent of U and $\{\mathcal{E}_j\}_{j \in \mathbb{Z} \setminus \{0\}}$.

Next, we define a probability measure for each finite vector $((i_1, t_1), \dots, (i_k, t_k))$ in $(\mathbb{Z} \times [0, \infty))^k$. Let (j_1, \dots, j_k) be the string consisting of the coordinates of (i_1, \dots, i_k) arranged in non-decreasing order, and τ the permutation in the symmetric group \mathcal{S}_k such

that $j_{\tau(l)} = i_l$ for $1 \leq l \leq k$. Let $\eta^{j_1}(t) := B^{j_1}(t) + \sum_{r=0}^{j_1} X_r$, and for $j_1 < l \leq j_k$, let $\eta^l(t) := \mathcal{Q}(\eta^{l-1}(t), \sum_{r=0}^l X_r + B^l(t) + ct)$. Given times $(t_1, \dots, t_k) \in [0, \infty)^k$, define

$$\nu(A_1 \times \dots \times A_k)_{(i_1, t_1), \dots, (i_k, t_k)} := \mathbb{P}(\eta^{j_1}(t_{\tau(1)}) \in A_1, \dots, \eta^{j_k}(t_{\tau(k)}) \in A_{j_k}),$$

with \mathbb{P} being the probability measure induced by the sequence of independent Brownian motions $\{B^i\}_{i \in \mathbb{Z}}$.

In order to prove that these measures determine a process $\{\eta^i(t) : i \in \mathbb{Z}, t \geq 0\}$ it is enough to check that the consistency conditions of Kolmogorov's extension theorem are satisfied. We do this below.

Let $\sigma \in \mathcal{S}_k$. Note that the non-decreasing re-arrangements of (i_1, \dots, i_k) and $(i_{\sigma(1)}, \dots, i_{\sigma(k)})$ are equal, and therefore, the probability measures $\nu_{(i_1, t_1), \dots, (i_k, t_k)}$ and $\nu_{(i_{\sigma(1)}, t_{\sigma(1)}), \dots, (i_{\sigma(k)}, t_{\sigma(k)})}$ coincide.

We need to verify that for $((i_1, t_1), \dots, (i_k, t_k)) \in \left(\mathbb{Z} \times [0, \infty)\right)^k$, $(i_{k+1}, t_{k+1}) \in \mathbb{Z} \times [0, \infty)$ and $A_1, \dots, A_k \in \mathcal{B}(\mathbb{R}^k)$ we have

$$\nu(A_1, \dots, A_k, \mathbb{R})_{(i_1, t_1), \dots, (i_k, t_k), (i_{k+1}, t_{k+1})} = \nu(A_1, \dots, A_k)_{(i_1, t_1), \dots, (i_k, t_k)}.$$

Let $j_1 \leq \dots \leq j_k \leq j_{k+1}$ be the non-increasing reordering of i_1, \dots, i_{k+1} . We need to consider two cases: when $i_{k+1} > j_1$ and when $i_{k+1} = j_1$.

Case $i_{k+1} > j_1$.

Suppose $i_{k+1} = j_l$, for $l > 1$. We set

$$\begin{aligned} \nu(A_1 \times \dots \times A_k \times \mathbb{R})_{(i_1, t_1), \dots, (i_k, t_k), (i_{k+1}, t_{k+1})} &= \mathbb{P}(\eta^{j_1}(t_{\tau(1)}) \in A_{j_1}, \dots, \eta^{j_{l-1}}(t_{\tau(l-1)}) \in A_{j_{l-1}}, \\ &\quad \eta^{i_{k+1}}(t_{\tau(k+1)}) \in \mathbb{R}, \eta^{j_{l+1}}(t_{\tau(l+1)}) \in A_{j_{l+1}}, \\ &\quad \dots, \eta^{j_k}(t_{\tau(k)}) \in A_{j_k}) \\ &= \mathbb{P}(\eta^{j_1}(t_{\tau(1)}) \in A_{j_1}, \dots, \eta^{j_{l-1}}(t_{\tau(l-1)}) \in A_{j_{l-1}}, \\ &\quad \eta^{j_{l+1}}(t_{\tau(l+1)}) \in A_{j_{l+1}}, \dots, \eta^{j_k}(t_{\tau(k)}) \in A_{j_k}) \\ &= \nu(A_1 \times \dots \times A_k)_{(i_1, t_1), \dots, (i_k, t_k)} \end{aligned}$$

where the second equality follows from the consistency of the finite-dimensional distributions of the independent Brownian motions $\{B^{j_1}, \dots, B^{j_k}\}$, and the third equality is due to the fact that the string (i_1, \dots, i_k) has $(j_1, \dots, j_{l-1}, j_{l+1}, \dots, j_k)$ as its non-decreasing reordering. Note that this argument does not work in the case $l = 1$ because the processes η^i are defined in terms of η^{j_1} .

Case $i_{k+1} = j_1$.

In this case, due to the consistency of the finite-dimensional distributions of the independent Brownian motions, we have

$$\begin{aligned} \nu(A_1, \dots, A_k, \mathbb{R})_{(i_1, t_1), \dots, (i_k, t_k), (i_{k+1}, t_{k+1})} &= \mathbb{P}(\eta^{i_{k+1}}(t_{i_{k+1}}) \in \mathbb{R}, \\ &\quad \eta^{j_2}(t_{\tau(2)}) \in A_{j_2}, \dots, \eta^{j_{k+1}}(t_{\tau(k+1)}) \in A_{j_{k+1}}) \\ &= \mathbb{P}(\eta^{j_2}(t_{\tau(2)}) \in A_{j_2}, \dots, \eta^{j_{k+1}}(t_{\tau(k+1)}) \in A_{j_{k+1}}). \end{aligned}$$

By our version of Burke's theorem 1 we know that

$$\eta^{j_1+1} = \mathcal{Q}\left(\eta^{j_1}, \sum_{r=0}^{j_1+1} X_r + B^{j_1+1} + ct\right)$$

is distributed as $\sum_{r=0}^{j_1+1} X_r + W^{j_1+1}$, where W^{j_1+1} is a Brownian motion independent of $\{B^i : i > j_1 + 1\}$.

Applying this argument $j_2 - j_1$ times, we set that η^{j_2} has the same distribution as $W^{j_2} + \sum_{r=0}^{j_2+1} X_r$, where W^{j_2} is a Brownian motion independent of $\{B^i : i > j_2\}$. The result follows. □

By Kolmogorov's extension theorem, we then have a process $\eta = \{\eta^i(t) : i \in \mathbb{Z}, t \in \mathbb{R}\}$ such that the finite-dimensional distributions can be characterized as drifted Brownian motions interacting by exclusion in the front. For each fixed i , P^i moves as an undrifted Brownian motion beginning at $\eta^i(0)$, where $\{\eta^i(0)\}_{i \in \mathbb{Z}}$ is a Poisson process on the line. We call this process the Totally Asymmetric Brownian Exclusion Process (TABEP).

Some properties of the TABEP are:

1. The TABEP is selfsimilar with exponent of selfsimilarity $\frac{1}{2}$.
2. The Poisson process is an invariant distribution for the TABEP.

The first property is inherited from the selfsimilarity of Brownian motion and the homothetic property of the reflective mapping. The second one can be obtained by weak convergence: because the product of Bernoulli random variables is invariant for the Totally Asymmetric Exclusion Process (TASEP) [Lig85], and hence the scaling of the TASEP used in the proof of Theorem 1 in different times leads to the same distribution of particles for the TABEP.

Finally, we define a multiclass TABEP. A Brownian motion beginning at $x > 0$ has probability $p(x, T) < 1$ of not having hit 0 by time $T > 0$, for all values x and T . So for the TABEP, given a fixed interval $[0, T]$, there will be an infinite number of particles that do not hit the particle in front during $[0, T]$, and then it is possible to construct a multiclass TABEP using finite windows.

Let $\mathcal{L} = \{1, \dots, k\}$ the set of classes. We put $\eta^i(t)$ for the position of the i -th particle and $C^i(t) \in \mathcal{L}$ for its class at time t . Given a finite number of particles $\{\eta^i(t) : t \in [0, T], 1 \leq i \leq n\}$ such that η^1 and η^n do not hit η^{-1} and η^{n+1} respectively during $[0, T]$, we define the multiclass TABEP by letting the position of particles evolve as in the TABEP, and at each hitting time $\eta^i(t) = \eta^{i+1}(t)$, we redefine the classes for the particles by $C^i(t) = \min(C^i(t^-), C^{i+1}(t^-))$ and $C^{i+1}(t) = \max(C^i(t^-), C^{i+1}(t^-))$. This dynamics defines a continuous model similar to the multiclass TASEP, see [FM07].

Some natural questions related to this model are: What is the asymptotic evolution of the multiclass TABEP? Which distributions are invariant for the multiclass TABEP? Are those distributions weak limits of the invariant distributions for the multiclass TASEP? We would be interested in taking the density of particles to infinite and studying the resulting distribution of classes, in the saturated limiting model.

Bibliography

- [Ana93] V. Anantharam. Uniqueness of stationary ergodic fixed point for a $\cdot/M/K$ node. *Ann. Appl. Probab.*, 3(1):154–172, 1993.
- [Asm03] S. Asmussen. *Applied Probability and Queues*. Springer, 2003.
- [Bil68] P. Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, 1968.
- [Bur56] P. Burke. The output of a queuing system. *Operations Research*, 4(6):pp. 699–704, 1956.
- [DMO03] M. Draief, J. Mairesse, and N. O’Connell. Joint Burke’s theorem and RSK representation for a queue and a store. pages 69–82 (electronic), 2003.
- [FF94] P. A. Ferrari and L. R. G. Fontes. The net output process of a system with infinitely many queues. *Ann. Appl. Probab.*, 4(4):1129–1144, 1994.
- [FM06] P. A. Ferrari and J. Martin. Multiclass processes, dual points and m/m/1 queues. *Markov Processes and Related Fields*, 2006.

- [FM07] P. A. Ferrari and J. Martin. Stationary distributions of multi-type totally asymmetric exclusion processes. *The Annals of Probability*, 35(3):pp. 807–832, 2007.
- [Har90] M. Harrison. *Brownian Motion and Stochastic Flow Systems*. Krieger Pub Co, 1990.
- [HW90] M. Harrison and R. Williams. On the quasireversibility of a multiclass brownian service station. *The Annals of Probability*, 18(3), 1990.
- [Kal02] O. Kallenberg. *Foundations of modern probability*. Springer, 2002.
- [Lig85] T. Liggett. *Interacting Particle Systems*. Springer, 1985.
- [MP95] T. Mountford and B. Prabhakar. On the weak convergence of departures from an infinite series of $\cdot/M/1$ queues. *Ann. Appl. Probab.*, 5(1):121–127, 1995.
- [MP10] J. Martin and B. Prabhakar. Fixed points for multiclass queues. *Arrive: 1003.3024v1*, 31, 2010.
- [OY01] N. O’Connell and M. Yor. Brownian analogues of Burke’s theorem. *Stochastic Process. Appl.*, 96(2):285–304, 2001.
- [SN01] P. Salminen and I. Norros. On busy periods of the unbounded brownian storage. *Queueing Systems*, 39:317–333, 2001.
- [Whi02] W. Whitt. *Stochastic Process Limits*. Springer, 2002.

Chapter 2

Convergence of the Polynuclear Growth Model dynamics on the line

In this chapter we prove the convergence of the Polynuclear growth model (PNG) dynamics defined on the line to its stationary distribution.

This chapter is organized as follows. In Section 2.1 we present an introduction to the PNG model and a related process: the Hammersley process. In Section 2.2 we prove the convergence to the stationary distribution of the PNG model. In Section 2.3 we introduce a reflection related to the PNG model, the *Hammersley reflection*. Finally, in Section 2.4, we present some related open problems.

2.1 Introduction

We present in this section an introduction to both the Hammersley process and the PNG model and we explain how they are related.

2.1.1 Hammersley process

The Hammersley process is an interacting particle system developed with the purpose of solving a problem concerning random permutations: the *Ulam's problem*. We next describe this motivation.

Consider the group \mathcal{S}_n of permutations of $\{1, \dots, n\}$, where $\pi \in \mathcal{S}_n$ is written in the notation $\pi = (x_1, \dots, x_n)$, meaning that

$$\pi(1) = x_1, \quad \pi(2) = x_2, \quad \dots, \quad \pi(n) = x_n.$$

Given $\pi \in \mathcal{S}_n$, let $L_n = L_n(\pi)$ the length of a longest increasing subsequence of π . For instance, for the permutation (6 2 4 3 7 1 5), a maximal increasing subsequence is (2 4 7) (another one is (2 3 5)). Ulam's problem consists in finding $\alpha(n)$ and c such that

$$\lim_{n \rightarrow \infty} \frac{L_n(\pi)}{\alpha(n)} = c,$$

when the permutation π is chosen uniformly in \mathcal{S}_n . One solution was proposed by Vershik and Kerov [VK77], and independently by Logan and Shepp [LS77], who proved that

$$\frac{\mathbb{E}(L_n)}{\sqrt{n}} \rightarrow 2, \tag{2.1}$$

by means of combinatorial techniques.

Another approach to the problem was proposed by Hammersley [Ham72]. Note that a random permutation of n elements can be constructed by the following procedure: Given fixed $x > 0, t > 0$, choose n uniform independent points in the rectangle $[0, x] \times [0, t]$. Let $x_1 \leq x_2 \leq \dots \leq x_n$ and $t_1 \leq t_2 \leq \dots \leq t_n$ be the coordinates of these points, and $\pi \in \mathcal{S}_n$ such that the points themselves are $(x_1, t_{\pi(1)}), \dots, (x_n, t_{\pi(n)})$. This construction can be extended to the positive halfplane by using a space-time Poisson process \mathcal{N} with intensity 1 in $\mathbb{R} \times [0, \infty)$. Let us call $M(x, t)$ the number of points in $[0, x] \times [0, t]$, so that $M(x, t)$ has distribution $\text{Poisson}(xt)$. Consider all possible up-right paths joining $(0, 0)$ and (x, t) going through marks of the process \mathcal{N} , and let $L(x, t)$ be the maximum number of marks employed by any such path. Let π be the permutation in $\mathcal{S}_{M(x,t)}$ determined by these points. See Figure 2.1 for an illustration.

It is easy to check that a path realizing the maximum $L(x, t)$ corresponds to a longest increasing subsequence of π , so $L(x, t) = L_{M(x,t)}(\pi)$.

Define $g(t) := \mathbb{E}[L(t, t)]$. By considering paths from $(0, 0)$ to $(t + s, t + s)$ passing by (t, t) , it can be shown that g is superadditive: that is, $g(t + s) \geq g(t) + g(s)$ for all $s, t \geq 0$. By the subadditive ergodic theorem it follows that

$$\frac{\mathbb{E}[L(t, t)]}{t} \rightarrow c.$$

This approach of studying the amount of marks attached to up-right paths is known as *Last passage percolation* (LPP), see [Mar06] for a review.

Aldous and Diaconis [AD95] remarked that Hammersley's solution involves an underlying interacting particle system. Let us show how it arises.

To fix ideas, let us work on $0 \leq x \leq X$ and $0 \leq t \leq T$. Let $t_1 \leq \dots \leq t_{M(x,t)}$ be the second coordinates of the marks of \mathcal{N} in $[0, X] \times [0, T]$. We consider a system of




Figure 2.1: Last passage percolation

particles with dynamics evolving as t grows. Here π is the permutation determined by the marks, defined in previous paragraph. At time 0, we begin with an empty configuration of particles, maintained till time t_1 . The first particle appears at time t_1 at the place $x_{\pi^{-1}(1)}$ where $(x_{\pi^{-1}(1)}, t_1) \in \mathcal{N}$. For $t_1 \leq t < t_2$ the configuration stays the same. At time t_2 , there are two possibilities: If the new mark $(x_{\pi^{-1}(2)}, t_2)$ lies to the left of $x_{\pi^{-1}(1)}$, then the existing particle moves to $x_{\pi^{-1}(2)}$. Otherwise, a new particle appears at $x_{\pi^{-1}(2)}$. Given a configuration of particles at time t_k^- , we update the system at this time by moving the nearest existing particle at the right of $x_{\pi^{-1}(k)}$ to this position, if such particle exists, and otherwise we create a new particle at $x_{\pi^{-1}(k)}$. See Figure 2.2. This dynamics define the *Hammersley process*.

This construction can be extended to $\mathbb{R} \times [0, \infty)$ for any initial configuration of particles, not only the empty one. Let us denote by $H(\cdot, t)$ the point process associated to configuration at time t . In particular, when at time 0 we have the empty configuration then $H(x, t) = L(x, t)$. On the other hand, if we start at time 0 with an arbitrary configuration $H(\cdot, 0)$, then we have

$$H(x, t) = \sup_{0 \leq z \leq x} [H(z, 0) + L((z, 0), (x, t))], \quad x, t \geq 0,$$

where, matching the previous definition of $L(x, t)$,

$$L((z, 0), (x, t)) := \max\{\text{number of } \mathcal{N} - \text{marks visited by } P\},$$

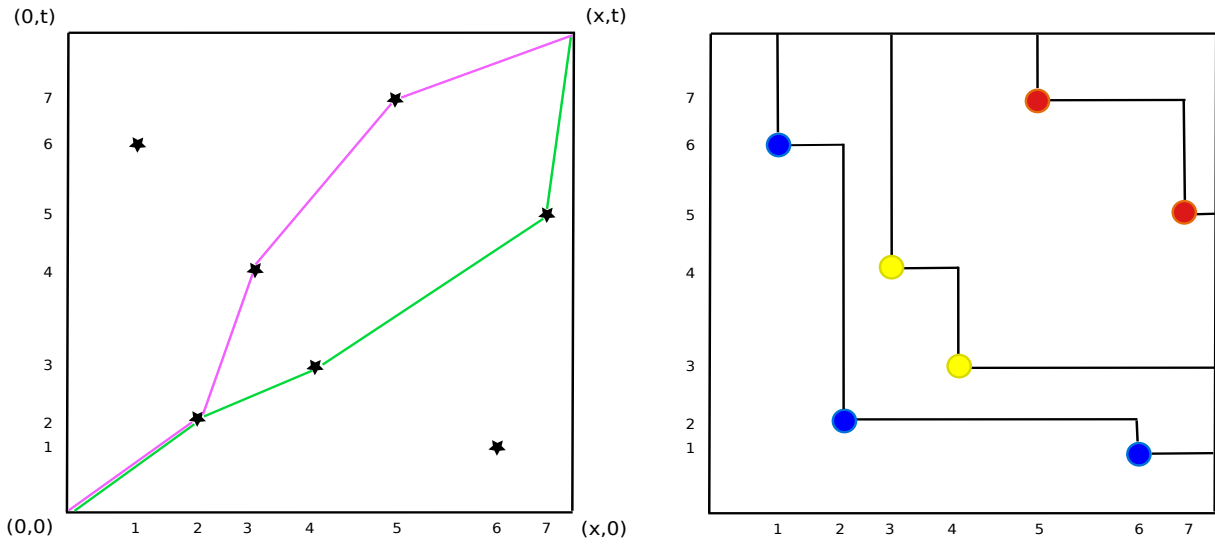


Figure 2.2: Bijection between LPP and Hammersley process

where the maximum is taken over all upright paths from $(z, 0)$ to (x, t) .

Aldous and Diaconis studied some hydrodynamic limit of the Hammersley process to prove the known result of Equation 2.1 by more direct arguments than the ones given in [VK77] and [LS77].

Without any reference to LPP, we can informally describe the Hammersley process as follows: each particle at position $r \in \mathbb{R}$ waits an exponentially distributed time with rate given by the distance to the nearest particle to its left located at $q < r$, and then chooses a site to jump to uniformly in $[q, r]$. See Figure 2.3, where the graphical representation of Hammersley process is oriented in the usual way for an interacting particle system, that is not the same as in previous pictures (that corresponds to the canonical orientation in LPP).




Figure 2.3: Graphical construction of Hammersley's process.

Aldous and Diaconis [AD95] showed that for any $\lambda > 0$, the rate λ Poisson process is an invariant measure for the Hammersley dynamics.

2.1.2 Polynuclear growth dynamics

In this subsection, we introduce the Polynuclear growth model (PNG). We first we define the model formally and then we give a detailed construction in two different spaces of states: the droplet and the line. This introduction will follow closely Prahöfer and Spohn [PS00].

The PNG model is a mathematical description of a crystal growing layer by layer in a flat substrate through random deposition of particles. These particles nucleate on existing plateaus of the crystal, forming new islands. These islands spread at linear speed to their sides and, when two islands of same level meet, they coalesce. On the top of the levels, new islands emerge.

The configuration of the model is codified by a *height function* $h(x, t)$, (x, t) , $x \in \mathbb{R}$, $t \geq 0$ the space-time coordinates. Given a nucleation event at (x_0, t_0) , the corresponding island nucleates to increase the height to

$$h_0 = h(x_0, t_0^-) + 1. \quad (2.2)$$

We are also assuming that the linear speed at which islands spread is 1, and that nucleation events are random and independent. A nice graphic simulation is available online at <http://www-wt.iam.uni-bonn.de/ferrari/animations/AnimationRSK.html>, in the webpage of Patrick Ferrari.

PNG droplet

Before introducing the model, we give a definition that we will need later on in the chapter.

Definition 4. For a point $(x, t) \in \mathbb{R} \times [0, \infty)$ we define

- its **light cone** as

$$\mathcal{V}_{(x,t)} = \{(y, s) \in \mathbb{R} \times [0, \infty) : |x - y| \leq s - t\},$$

- and its **backward light cone** as

$$\overleftarrow{\mathcal{V}}_{(x,t)} = \{(y, s) \in \mathbb{R} \times [0, \infty) : |x - y| \leq t - s\}.$$

Given two points $X = (x, t), Y = (y, t')$ in $\mathbb{R} \times [0, \infty)$, we will say that $X \leq_c Y$ in the light cone order if Y belong to \mathcal{V}_X . This defines a partial order.

The **PNG droplet** is the PNG model under the assumption that the space where nucleation events take place is the light cone of the origin $(0, 0)$:

$$\{(x, t) : -t \leq x \leq t\}.$$

This is a natural election for the PNG model: In order to determine the height $h(x, t)$ it is enough to know the nucleation events in its backward light cone. Label those events as (x_n, t_n) , with $n \in \mathbb{N}$. By 2.2, let

$$h_n := h(x_n, t_n^-) + 1 \quad n \in \mathbb{N}.$$

Then $h(x, t)$ is obtained as the maximum over all h_n for (x_n, t_n) occurring in the backward light cone:

$$h(x, t) = \max\{h_n : n \in \mathbb{N}, \quad |x - x_n| < t - t_n\},$$

with the rule that $h(x, t) = 0$ if the light cone is empty.

We now describe the droplet growth starting from the initially flat substrate: A single island starts spreading from the origin and then further nucleations take place above this ground layer. In this case, $h(x, t)$ is determined by the nucleation events inside the rectangle

$$R(x, t) = \{(x', t') \in \mathbb{R}^2 : |x'| \leq t' \text{ and } |x - x'| \leq t - t'\}.$$

In *light-like* coordinates, $u := \frac{(t'+x')}{2}$, $v := \frac{(t'-x')}{2}$, we can write rectangle $R(x, t)$ as $[0, U] \times [0, V]$ with $U = \frac{(t+x)}{2}$, $V = \frac{(t-x)}{2}$. The nucleation events in $R(x, t)$ are Poisson distributed with rate 1. We label them as (u_n, v_n) , $n = 1, \dots, N$, with N distributed Poisson $(\frac{t+x}{2} \cdot \frac{t-x}{2})$, such that $0 \leq u_1 < \dots < u_N \leq U$.

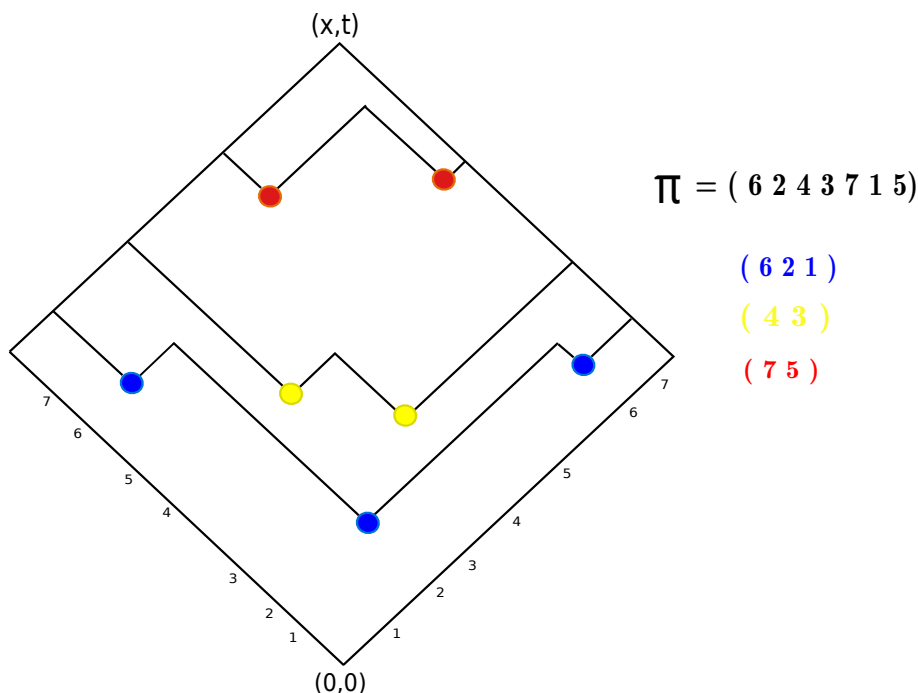


Figure 2.4: PNG Droplet.

Once again, if we label the v coordinates as $0 \leq v_1 \leq \dots \leq V$, the Poisson marks determine a permutation π by defining $\pi(i)$ such that $(u_i, v_{\pi(i)})$ is a mark.

In order to distinguish nucleation events corresponding to different height levels, we partition $(\pi(1), \dots, \pi(N))$ into decreasing subsequences according to the following procedure: for the first subsequence one starts with the first entry in the tuple and, going from left to right, one adds an entry to the subsequence if it is smaller than the last entry in this subsequence. Every used entry is marked as deleted. When the end of the tuple is reached, the first subsequence is completed. This procedure is repeated to obtain subsequences by using the undeleted entries of the original tuple.

As a result, the permutation π is partitioned into decreasing subsequences. The nucleation events of each such subsequence are on the same height level and distinct subsequences correspond to distinct heights. The space-time height lines, across which $h(x', t')$ increases by one unit, are then constructed by connecting the nucleation events belonging to the same subsequence by a zigzag line, as depicted in Figure 2.4.

In this example we have the permutation $(6, 2, 4, 3, 7, 1, 5)$. The decreasing subsequences are $(6, 2, 1), (4, 3), (7, 5)$. They define three height lines. For an arbitrary permutation π , $h(x, t)$ equals the number of decreasing subsequences, which is precisely the length, l , of the longest increasing subsequence of π . To see this, take any increasing

subsequence of π . By construction every element of this subsequence belongs to a different height line. Then $l \leq h(x, t)$. The reverse inequality is obtained by noticing that a nucleation event $(u_i, v_{\pi(i)})$ at level $k + 1$ is necessarily situated in the forward light cone of at least one nucleation event $(u_j, v_{\pi(j)})$ at level k , i.e. $i > j$ and $\pi(i) > \pi(j)$. We call $(u_j, v_{\pi(j)})$ a predecessor of $(u_i, v_{\pi(i)})$. Since there is at least one nucleation event in level $h(x, t)$, successive selections of predecessors in consecutive levels result in an increasing subsequence of length $h(x, t)$, therefore $l \geq h(x, t)$.

Notice that this mapping corresponds to the one defined in Subsection 2.1.1: the paths of particles in the Hammersley process are exactly the space-time step lines of the PNG model, after performing a 45° rotation. This remark, first made by Prähofer and Spohn [PS00], establishes a connection between the two models. Figure 2.5 illustrates the mapping.

The PNG droplet is a growth model in the KPZ class, for an introduction to the relationship of this PNG model with the KPZ equation and the Tracy-Widom distribution see [PS00].

PNG on the line

In this subsection we extend the definition of the PNG model in the droplet to a halfplane.

The initial configuration on \mathbb{R} is given by a signed point process $A = \{a_i = (\text{pos}(a_i), \text{sign}(a_i))\}_{i \in \mathbb{Z}}$, such that $(\text{pos}(a_i), \text{sign}(a_i)) \in \mathbb{R} \times \{-1, 1\}$. There is also a rate 2 Poisson process \mathcal{N} on $\mathbb{R} \times [0, \infty)$ (space-time). The *marks* of this process \mathcal{N} are the **nucleations events**.

Each particle a_i moves at constant speed $\text{sign}(a_i)$, so position $\text{pos}(a_i)$ is changing at linear rate. At a given nucleation event $(x_0, t_0) \in \mathcal{N}$, two new particles are generated simultaneously at the place x_0 : one travelling with speed $+1$ and the other with speed -1 . In the instant where a particle with speed 1 hits another particle (of speed -1) both of them are annihilated whether these particles were part of the initial configuration or not. See Figure 2.6.

By performing a 45° rotation, we may obtain a stationary version of the PNG process from a stationary version of Hammersley process in the plane. If we tag the particles in the Hammersley process, after the rotation, we obtain the lines of the PNG model (see Figure 2.5 in last subsection). Prähofer and Spohn proved in [PS04] that if $A^+ = \{a_i : \text{sign}(a_i) = +1\}$ and $A^- = \{a_i : \text{sign}(a_i) = -1\}$ are independent Poisson processes with intensity ρ^+ and ρ^- respectively, such that $\rho^+ \rho^- = 1$, and we apply the PNG dynamics then, at any given time t , the set of alive particles $A(t) = A^+(t) \cup A^-(t)$ is distributed

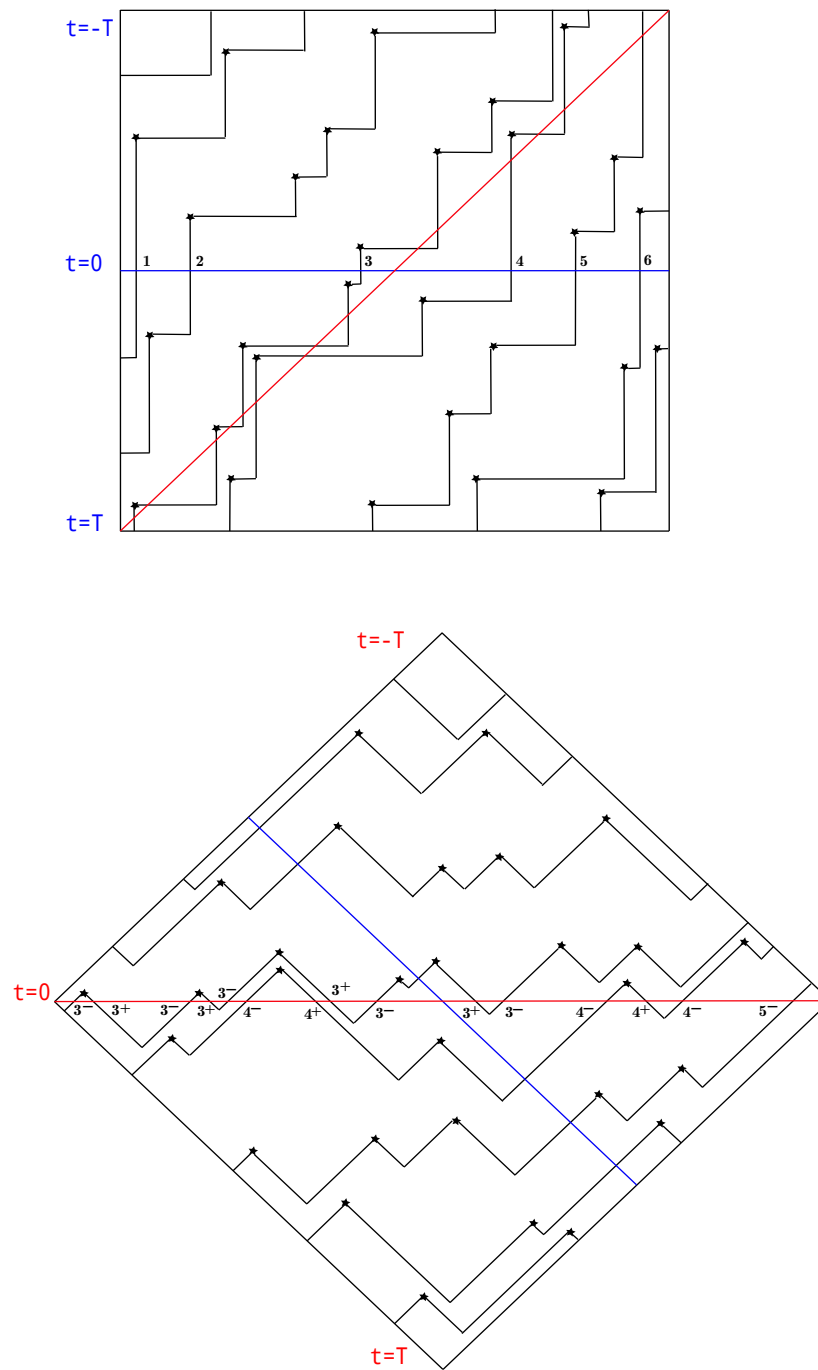


Figure 2.5: Relationship between Hammersley and PNG processes.




Figure 2.6: PNG model defined in the line.

as the initial equally to A (with the same parameters as A^+, A^-). That is, the sum of two independent Poisson processes is invariant under the PNG dynamics. Notice that the distribution of this sum of Poisson processes is the law of the jump times of a biased simple random walk on the line which jumps one unit upwards at rate ρ^+ and one unit downwards at rate ρ^- . We will call to this law of a signed point process as the **ring bells of a (simple) random walk** with parameters ρ^+ and ρ^- , for simplicity.

2.2 Convergence of the PNG model on the line

In this section we prove that any initial distribution converges under the PNG dynamics to the stationary distribution. We apply a coupling technique similar to the one used by Mountford and Prabhakar [MP95] to prove that the $\cdot \setminus M \setminus 1$ queueing operator dynamics is convergent.

Assume we start with two different configurations of a signed point process: a process A , distributed as the ring bells of a random walk with parameters ρ^+, ρ^- ($\rho^+ \rho^- = 1$), and an independent ergodic marked point process $E = E^+ \cup E^-$ with the same rates.

We will couple these processes by using a rate 2 Poisson process \mathcal{N} to define the PNG dynamics. We thus obtain two signed processes A_t and E_t , $t \geq 0$. We will prove that as t grows to infinity the density of particles in $A_t \Delta E_t$ goes to 0. In other words,

Theorem 4. *Let E be a signed ergodic process on the line with rates ρ^+, ρ^- for $+$ and $-$ marks, such that $\rho^+ \rho^- = 1$. Define E_t as the signed process with initial configuration E obtained by running the PNG dynamics. Then E_t converges weakly to the ring bells of a random walk with parameters ρ^+ and ρ^- .*

First, we explain briefly the coloured coupling that we will use. At time 0 we define red particles as $R = A^+ \cup E^-$, and blue particles as $B = A^- \cup E^+$. The sign of each particle is conserved, so at time 0 we have two signed processes: the red one $R = R^+ \cup R^-$ and the blue one $B = B^+ \cup B^-$. Particles generated by nucleation events are yellow particles $Y_t = Y_t^+ \cup Y_t^-$.

The colours will represent the discrepancies between configurations A_t and E_t , for each time t . We need to define some dynamics between coloured particles and describe the relationship between the coloured dynamics and the original processes A_t and E_t .

Each coloured particle moves with a speed equal to its sign. If a particle hits another particle of the same colour, nothing happens and both particles continue their paths (they cross each other). If a red particle hits a blue particle, both of them are annihilated. If a red or blue particle hits a yellow one, the yellow particle is annihilated and the red or blue changes its sign. If two yellow particles collide they annihilate each other. See Figure 2.7.

It is easy to check that for each $t \geq 0$, $A_t^+ = R_t^+$, $A_t^- = B_t^-$ and $E_t^+ = B_t^+$, $E_t^- = R_t^-$, so the technique of coloured particles is just a handy way of seeing the particles. Notice that there is no creation of red or blue particles, only annihilation. The coupling of two particles from A and E corresponds to the annihilation of particles in R and B . So we must prove that the red and blue particles disappear as t goes to infinity in the coloured dynamics.

Let us first show that space ergodicity is preserved under the dynamics.

Lemma 5. R_t, B_t and Y_t are ergodic for every $0 \leq t < \infty$.

Proof.

Let us denote by \mathcal{N}_t the restriction of \mathcal{N} to $\mathbb{R} \times [0, t]$. For a fixed time t , we have that

$$R_t = f_t(A^+, A^-, E^+, E^-, \mathcal{N}_t),$$

where f_t is a deterministic function.

Note that the space shift-operation commutes with f_t : If we denote by τ_c the space shift by c units,

$$\tau_c(x) := x + c,$$

we have that

$$\tau_c(R_t) = f_t(\tau_c(A^+), \tau_c(A^-), \tau_c(E^+), \tau_c(E^-), \tau_c(\mathcal{N}_t)).$$

Hence, the shift invariance of R_t is a consequence of the shift invariance of processes $A^+, A^-, E^+, E^-, \mathcal{N}_t$.

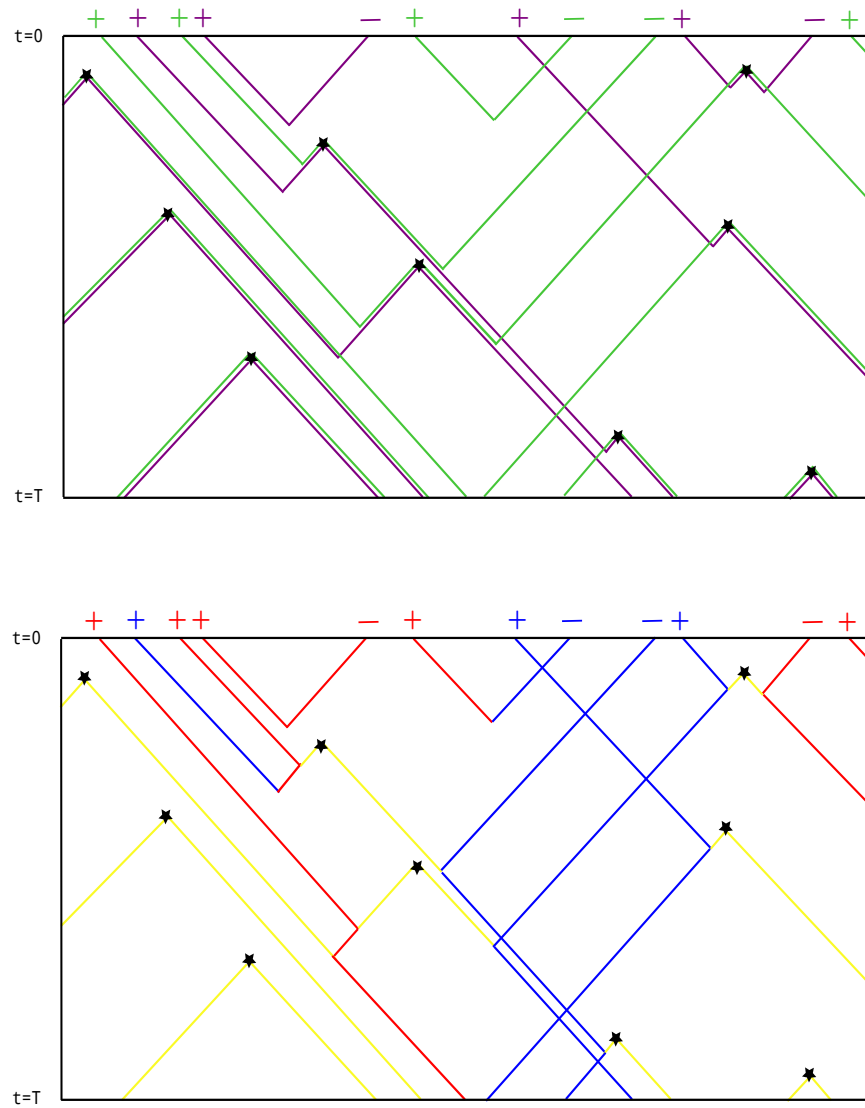


Figure 2.7: A simultaneous realization of two configurations using the same spatial marks and its colouring.

Since $\{(A^+, A^-), (E^+, E^-), \mathcal{N}_t\}$ is a family of independent ergodic processes, $(A^+, A^-, E^+, E^-, \mathcal{N}_t)$ is a jointly ergodic process. Since any shift-invariant function of R_t is a shift-invariant function of $(A^+, A^-, E^+, E^-, \mathcal{N}_t)$, it follows that the process R_t is ergodic. A similar argument works for the processes B_t and Y_t .

□

Proof of Theorem 4:

The proof will proceed by contradiction. Let us assume that R and B never annihilate each other, so that the density of red particles which are never annihilated is a positive random variable, denoted by $d_R(\infty, \omega)$. Such non-annihilating particles that never annihilate will be called **ever-red** and **ever-blue particles**, and denoted by ER and EB .

By Lemma 5, the density of red particles at a fixed time $t > 0$, $d_R(t, \omega)$, is deterministic:

$$d_R(t, \omega) = d_R(t) \quad a.s..$$

By definition, $d_R(t, \omega)$ is a non-increasing function of t , so we have that

$$d_R(\infty, \omega) = \lim_{t \rightarrow \infty} d_R(t, \omega) = \lim_{t \rightarrow \infty} d_R(t) =: d_{ER} \quad a.s.,$$

where the limit exists by the monotonicity and boundedness of $\{d_R(t)\}$.

Since red and blue particles cannot cross paths it is possible to define an order between ever-red and ever-blue particles.

Procedure 1. *We tag the particles to indicate an order. At time 0, the first red ever-red particle to the right of the origin will have tag 0, let us call x the position of that particle. Let l_x be the nearest ever-blue particle to the left of x , and r_x the nearest ever-blue particle to the right of x . All the ever-red particles in $(x - \text{pos}(l_x), x + \text{pos}(r_x))$ will have tag 0. Now, we tag the blue particle in $\text{pos}(r_x)$ as 1, and look for its nearest ever-red particle to the right which will have tag 2. The ever-blue particle l_x will have tag -1 and the nearest ever-red particle to its left will have tag -2 . All ever-blue particles between particles with tag 1 and 2 will have tag 1, and all ever-blue particles between particles with tags -1 and -2 will have tag -1 . We do this inductively, with the result that ever-red particles get even tags, ever-blue particles get odd tags and the order between particles with different tags is preserved through the dynamics. See Figure 2.8.*




Figure 2.8: Relative order of ever-red and ever-blue particles.

We will need the following lemmas that state some basic facts about ergodic process on the line.

Lemma 6. *Let A be an ergodic point process in the line with density $0 < d_A < \infty$.*

1. *For $c > 0$, let \bar{A} an ergodic subset of A with positive density $d_{\bar{A}}$. Let*

$$\bar{A}_c := \{a \in \bar{A} : (\text{pos}(a) - c, \text{pos}(a) + c) \cap A = \{a\}\}$$

Then $d_{\bar{A}_c} \nearrow d_{\bar{A}}$ as c goes to 0.

2. *Let B be another ergodic process on the line with the same density d_A . Tag the particles of A and B as we did in Procedure 1. Let \tilde{A} be an ergodic subset of A with positive density. For $c > 0$, let*

$$\tilde{A}_c := \{a \in \tilde{A} : \exists u \in \tilde{A} \text{ with } \text{pos}(u) \in (\text{pos}(a), \text{pos}(a) + c) \text{ such that } \text{tag}(u) > \text{tag}(a)\}.$$

Then $d_{\tilde{A}_c} \nearrow d_{\tilde{A}}$, as c goes to ∞ .

Proof.

1. Note that $0 < c_1 < c_2$ implies $\bar{A}_{c_2} \subseteq \bar{A}_{c_1}$. For $m \geq 1, i \in \mathbb{Z}$ define

$$X_i := \#[i, i + 1) \cap \bar{A}, \quad X_i^m := \#[i, i + 1) \cap \bar{A}_{\frac{1}{m}}.$$

Note that $X_i^m \nearrow X_i$ almost surely, because an ergodic point process with finite density does not have accumulation points. On the other hand, by the ergodic theorem, we have that $d_{\bar{A}_{\frac{1}{m}}} = \mathbb{E}(X_1^m)$, and then

$$\lim_{m \rightarrow \infty} d_{\bar{A}_{\frac{1}{m}}} = \lim_{m \rightarrow \infty} \mathbb{E}(X_1^m) = \mathbb{E}(\lim_{m \rightarrow \infty} X_1^m) = \mathbb{E}(X_1) = d_{\bar{A}}$$

where the interchanging of limits is justified by the monotone convergence theorem, as $X_i^m \nearrow X_i$ a.s. and $d_{\bar{A}} < \infty$.

2. Once again, we take advantage of a monotonicity property: the inequality $0 < c_1 < c_2$ implies $\bar{A}_{c_1} \subseteq \bar{A}_{c_2}$. For $m \geq 1, i \in \mathbb{Z}$ define

$$X_i := \#[i, i+1) \cap \bar{A} \quad X_i^m := \#[i, i+1) \cap \bar{A}_m.$$

The process B has positive density so $X_i^m \nearrow X_i$ almost surely. Since $d_{\bar{A}} < \infty$, by the ergodic theorem we get

$$\lim_{m \rightarrow \infty} d_{\bar{A}^m} = \lim_{m \rightarrow \infty} \mathbb{E}(X_1^m) = \mathbb{E}(\lim_{m \rightarrow \infty} X_1^m) = \mathbb{E}(X_1) = d_{\bar{A}}.$$

□

Lemma 7. *For all $t \geq 0$, there exists $\delta > 0$ such that following set has density larger than δ :*

$$\mathcal{C}_t = \left\{ r \in ER : \exists u \in ER \text{ with } \begin{aligned} & \text{pos}(u) \in (\text{pos}(r), \text{pos}(r) + c_1), \text{tag}(u) > \text{tag}(r), \\ & \#R_t \cap (\text{pos}(r) - c_2, \text{pos}(r) + c_2) = \#R_t \cap (\text{pos}(u) - c_2, \text{pos}(u) + c_2) = 1 \end{aligned} \right\}$$

for some $0 < c_2 < c_1$, where all these constants do not depend on t .

Proof.

Fix $t > 0$. By Lemma 6,1), we can choose $c_2 > 0$ small enough such that the set

$$\mathcal{B}_t = \left\{ r \in ER : \#R_t \cap (\text{pos}(r) - c_2, \text{pos}(r) + c_2) = 1 \right\}$$

has density bigger than $\frac{d_{ER}}{2}$. Then using Lemma 6,2), we can choose $c_1 > c_2$ big enough such that the set

$$\mathcal{C}_t = \left\{ r \in \mathcal{B}_t : \exists u \in \mathcal{B}_t, \text{ with } \begin{aligned} & \text{pos}(u) \in (\text{pos}(r), \text{pos}(r) + c_1) \text{ such that } \text{tag}(u) > \text{tag}(r) \end{aligned} \right\}$$

has density bigger than $\frac{\text{dens}(\mathcal{B}_t)}{2} > \frac{d_{ER}}{4}$. Let $\delta := \frac{d_{ER}}{4}$, and we are done.

□

We are now ready to finish the proof of Theorem 4. We will show that if $d_R(t) > \epsilon$ for all $t \geq 0$, then there exist $\Delta > 0$ and $K > 0$, independent of t , such that $d_R(t) - d_R(t+\Delta) > K$.

Fix $t > 0$. By Lemma 7 we know that

$$\mathcal{C}_t = \left\{ r \in ER : \exists u \in ER \text{ with } \begin{aligned} & \text{pos}(u) \in (\text{pos}(r), \text{pos}(r) + c_1), \text{tag}(u) > \text{tag}(r), \\ & \#R_t \cap (\text{pos}(r) - c_2, \text{pos}(r) + c_2) = \#R_t \cap (\text{pos}(u) - c_2, \text{pos}(u) + c_2) = 1 \end{aligned} \right\}$$

has density larger than $\delta > 0$, for some $0 < c_2 < c_1$.

Since R_t , B_t and Y_t are ergodic point processes with finite density, for every $r \in \mathcal{C}_t$ there is a finite number of yellow particles in the interval $(r, r + c_1)$. For $k \geq 0$, let

$$\mathcal{C}_t^k := \left\{ r \in ER : \exists u \in ER \text{ with } \begin{aligned} & \text{pos}(u) \in (\text{pos}(r), \text{pos}(r) + c_1), \text{tag}(u) > \text{tag}(r), \\ & \#R_t \cap (\text{pos}(r) - c_2, \text{pos}(r) + c_2) = \#R_t \cap (\text{pos}(u) - c_2, \text{pos}(u) + c_2) = 1, \\ & \#Y_t \cap (\text{pos}(r), \text{pos}(r) + c_1) \leq k \end{aligned} \right\}.$$


By the same argument used in the proof of Lemma 6, $\text{dens}(\mathcal{C}_t^k) \nearrow \text{dens}(\mathcal{C}_t)$ as $k \nearrow \infty$, and we can choose k large enough so that \mathcal{C}_t^k has density bigger than $\frac{\delta}{2}$.

Let $r_1 \in \mathcal{C}_t^k$. We know that there exists at least one red particle r_2 in $(\text{pos}(r_1), \text{pos}(r_1) + c_1)$, such that $R \cap (\text{pos}(r_2) - c_2, \text{pos}(r_2) + c_2) = \{r_2\}$ and $\text{pos}(r_2) = \text{pos}(r_1) + d$, with $d \leq c_1$, we fix such red particle r_2 . Hence, there are no other red particles in $(\text{pos}(r_1) - c_2, \text{pos}(r_1) + c_2)$, neither in $(\text{pos}(r_2) - c_2, \text{pos}(r_2) + c_2)$. Define the event E_{r_1} , by the simultaneous occurrence of the following events (see Figure 2.9):

- In the isosceles triangle Δ_L with base $(\text{pos}(r_1) - c_2, \text{pos}(r_1))$ (at time t) and height on the time interval $(t, t + \frac{c_2}{2})$ there are exactly $k + 1$ marks of \mathcal{N} , and these marks are completely ordered in the light cone order (recall Definition 4 in the Subsection 2.1.2).
- In the isosceles triangle Δ_R with base $(\text{pos}(r_2), \text{pos}(r_2) + c_2)$ (at time t) and height on the time interval $(t, t + \frac{c_2}{2})$ there are exactly $k + 1$ marks of \mathcal{N} , and these marks are completely ordered in the light cone order.
- Let denote by Δ_C the isosceles triangle with base $(\text{pos}(r_1) - c_2, \text{pos}(r_2) + c_2)$ (at time t) and height on the time interval $(t, t + \frac{d+2c_2}{2})$. Then, there are no marks of \mathcal{N} on $\Delta_C \setminus (\Delta_L \cup \Delta_R)$.

The third event has probability bounded by below by $e^{-\frac{(c_1+2c_2)^2}{4}}$ while the first and the second ones have some positive probability depending only on c_2 , and they are all independent events depending only on location of the marks of \mathcal{N} . Hence E_{r_1} has some positive probability $\mathbb{P}(E_{r_1}) \geq p(c_1, c_2) > 0$, provided that $r_1 \in \mathcal{C}_t^k$.

Let us show that if $r_1 \in \mathcal{C}_t^k$ and E_{r_1} occurs, there is a collision of a red and a blue particle. For instance, assume that there are no red particles in the configuration at time t in the interval $(\text{pos}(r_1), \text{pos}(r_2))$. Inside $\Delta_C \setminus (\Delta_L \cup \Delta_R)$ there are no marks, so the only

Figure 2.9: Event of coupling E_r .

particles that could interact with the red particles r_1 and r_2 and the blue particles between them, are the yellow particles that were already present at time t . Since there are at most k of these particles, we have at most k changes of direction for each red particle. Now, because triangles Δ_L and Δ_R contain $k + 1$ marks each, and the particles yellow thereby created cannot interact (because they are ordered), every time one of the red particles r_1 or r_2 goes in the wrong direction (defined as the direction that increases distance between them), it eventually meets a yellow mark that corrects the direction. Note that we choose $k + 1$ instead of k in the definition of E_{r_1} because the red particles r_1, r_2 could initially have the wrong directions. Then, in the absence of blue particles, particle r_1 and r_2 would have met before time $t + \frac{d+2c_2}{2}$. But we made sure there is at least one blue particle in between that annihilates one of the red particles before that time, see Figure 2.9.

If we have more red particles in $(pos(r_1), pos(r_2))$, the statement still holds. Since the yellow marks in Δ_L and Δ_R has bigger cardinality than all the particles in $[pos(r_1), pos(r_2)]$, by the same argument all red particles will be trapped in Δ_C (possibly changing relative positions). By the existence of a blue particle, one red particle will be annihilated before time $t + \frac{d+2c_2}{2}$.

We hence arrive at the contradiction

$$d_R(t) - d_R\left(t + \frac{d + 2c_2}{2}\right) \geq \frac{\delta}{2} p(c_1, c_2) = K > 0, \quad \forall t > 0$$

where the constant K does not depend on t . \square

2.3 Hammersley reflection

In this section, we prove some extension of Theorem 4. We first define a new reflection operator, closely related to the particle paths in the Hammersley process, find its invariant measure and finally show some local convergence result.

Definition 5. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous function. Let \mathcal{N} be a rate 2 Poisson process on $A = \{(x, t) \in \mathbb{R} \times \mathbb{R} : t \leq f(x)\}$. Define

$$R = \bigcup_{\omega \in \mathcal{N} : \mathcal{V}_\omega \subseteq A} \mathcal{V}_\omega,$$

where \mathcal{V}_ω is the (forward) light cone of ω , (Definition 4). We define the Hammersley reflection of f as the boundary set ∂R , that can be parametrized as a real continuous function, and we denote that function by $R(f)$, see Figure 2.10.




Figure 2.10: Hammersley reflection of a function f .

2.3.1 Stationary distribution under the Hammersley reflection

A natural question is whether there exists an invariant distribution under this operation. That is, we look for a process X such that the increments of the reflected process $R(X)$ are distributed as the increments of X . This is indeed the case, it follows from the following result proven by P. A. Ferrari [Fer96] (Theorem 10.2) stated for the continuous time, discrete state space Hammersley process, but the scaling used in the proof of Theorem 1 gives us the result for the continuous state space version (as already noted in [Fer96]).

Theorem 5. *Let H be a stationary Hammersley process, with rate 1 marks \mathcal{N} on $\mathbb{R} \times [0, t)$ and initial distribution of particles $\text{Poisson}(\lambda)$, $\lambda > 0$. Then the trajectory of a tagged particle $X^x(\cdot)$, $X^x(0) = x$, is compound Poisson, that is:*

$$X^x(t) \stackrel{d}{=} x + \sum_{i=1}^{N_t} E_i,$$

N_t is a rate 1 Poisson process in \mathbb{R} independent of the sequence of exponential(λ) $\{E_i\}_{i=1}^{\infty}$ i.i.d. random variables.

In the stationary regime, we can picture the paths of two neighbouring particles in the Hammersley process as the 45° rotation of a function f . This would be the trajectory of the leftmost particle, and the 45° rotation of its reflection $R(f)$, associated to the trajectory of the rightmost particle of the pair. We can thus obtain the invariant distribution of the Hammersley reflection from Theorem 5.

Corollary 1. *An invariant distribution for the increments of a process, under the Hammersley reflection, is given by the increments of a process Y with a path given by a particle travelling at speed $+1$ or -1 , and switching between these speeds at independent exponential($\sqrt{2}$) intervals.*

2.4 Open problems

2.4.1 Scaling

In the stationary version of the Hammersley process, each particle follows a compound Poisson path. By scaling time and space (Euler scaling) we will obtain a surface, presumably a Brownian sheet. How can the convergence result of the Hammersley reflection can be translated to the dynamics of the scaled version of the process?

2.4.2 Multiclass version

In [Fer04] P.L. Ferrari constructed a multi-layer version of the PNG model and analyzed asymptotics by considering an initial condition given by a flat substrate. This multi-layer PNG model is an analogue of the multi-line process for the Hammersley process and can be defined for arbitrary initial conditions. By the results of P.A. Ferrari and J. Martin [FM09], we know that the multi-line dynamics is isomorphic to the multiclass dynamics, so we would expect a similar relation for the multiclass PNG model. Is it possible to compute the stationary distribution for this multiclass version of the PNG model?

Bibliography

- [AD95] D. Aldous and P. Diaconis. Hammersley's interacting particle process and longest increasing subsequences. *Probab. Th. Rel. Fields*, 103:199–213, 1995.
- [Fer96] P. A. Ferrari. Limit theorems for tagged particles. *Markov Process. Related Fields*, 2(1):17–40, 1996.
- [Fer04] P. L. Ferrari. Polynuclear growth on a flat substrate and edge scaling of goe eigenvalues. *Communications in Mathematical Physics*, 252(1-3):77–109, 2004.
- [FM09] P. A. Ferrari and J. Martin. Multiclass Hammersley-Aldous-Diaconis process and multiclass-customer queues. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(1):250–265, 2009.
- [Ham72] J. M. Hammersley. A few seedlings of research. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. I: Theory of statistics*, pages 345–394. Univ. California Press, 1972.
- [LS77] B.F Logan and L.A Shepp. A variational problem for random young tableaux. *Advances in Mathematics*, 26(2):206 – 222, 1977.
- [Mar06] J. Martin. Last-passage percolation with general weight distribution. *Markov Process. Related Fields*, 12(2):273–299, 2006.
- [MP95] T. Mountford and B. Prabhakar. On the weak convergence of departures from an infinite series of $\cdot/M/1$ queues. *Ann. Appl. Probab.*, 5(1):121–127, 1995.
- [PS00] M. Prähofer and H. Spohn. Statistical self-similarity of one-dimensional growth processes. *Phys. A*, 279(1-4):342–352, 2000. Statistical mechanics: from rigorous results to applications.
- [PS04] M. Prähofer and H. Spohn. Exact scaling functions for one-dimensional stationary kpz growth. *Journal of Statistical Physics*, 115(1-2):255 – 279, 2004.
- [VK77] A.M. Vershik and S.V. Kerov. Asymptotics of the plancherel measure of the symmetric group and the limiting form of young tables. *Soviet Math. Dokl.*, 18:527–531, 1977.

Chapter 3

Reversible birth and death processes, Whittle networks, Insensibility and Bandwidth-sharing networks

This chapter gives some preliminaries on stochastic networks and then introduces the mathematical definition of a bandwidth-sharing network. In Section 3.1 we give classical results about reversible birth and death processes on \mathbb{R}^N . This allows us to introduce the notion of a Whittle network and characterize its stationary measure in Section 3.2. After that, in Section 3.3, we study the concept of *insensitivity* in processor sharing queues and Whittle networks relating the stationary distribution of Markovian and non-Markovian models. Finally, in Section 3.4, we introduce the bandwidth-sharing network mathematical model, and we give some toy examples and commonly used specific models. We use some content of Jonckheere [Jon10].

3.1 Reversible birth and death processes

In this section we give a brief review of some results about birth and death processes on \mathbb{Z}_+^N concerning reversibility and stationary distributions.

Definition 6. *Let X a state dependent birth and death process on \mathbb{Z}_+^N with jump rates $\{q(x, y)\}$. We say that X is reversible if there exists a distribution π such that*

$$\pi(x)q(x, y) = \pi(y)q(y, x), \quad \forall x, y \in \mathbb{Z}_+^N. \quad (3.1)$$

Equation (3.1) is usually known as **detailed balance equations**. Note that if π is reversible with respect to the process X , then π is stationary for X .

This definition of reversibility corresponds indeed to the idea of a process being reversible in time, that is, to have the same distribution for dynamics forwards or backwards in time:

Theorem 6. *Fix a time $\tau > 0$ and consider the reverse time process*

$$X_t^\tau := X_{\tau-t}.$$

Then X is reversible if and only if X_t^τ is ergodic and

$$(X_{\tau-t_1}, \dots, X_{\tau-t_n}) \stackrel{d}{=} (X_{t_1}, \dots, X_{t_n}),$$

for all time sequences t_1, \dots, t_n .

This is a very classic result whose proof can be consulted in [Asm03], for instance. \square

Now we derive a criterion for a process to be reversible depending only on its transition rates. This allows us to check if a process is reversible without calculating explicitly the stationary distribution.

Theorem 7 (Kolmogorov criterion). *The following statements are equivalent:*

1. *There exists a positive distribution π satisfying detailed balance equations (3.1) for the transition rates q .*
2. *For all $x_0, \dots, x_n \in \mathbb{Z}_+^N$ with $x_0 = x_n$, we have that*

$$\prod_{i=1}^n q(x_{i-1}, x_i) = \prod_{i=1}^n q(x_i, x_{i-1}). \quad (3.2)$$

3. *For each path $x_0, x_1, \dots, x_n \in \mathbb{Z}_+^N$, such that x_i communicates with x_{i+1} for all i , the quantity (defined for non-zero transitions)*

$$\frac{\prod_{i=1}^n q(x_{i-1}, x_i)}{\prod_{i=1}^n q(x_i, x_{i-1})}$$

depends only on x_0 and x_n .

Proof. [Kel79]

We prove first that (1) implies (3.2). Since there exists π which satisfy Equations (3.1), we have that $\prod_{i=1}^n q(x_{i-1}, x_i)\pi(x_{i-1}) = \prod_{i=1}^n \pi(x_i)q(x_i, x_{i-1})$. Since $x_0 = x_n$, we can

cancel the π factors of both sides and we are done.

For the rest of the proof, we do the following: Let $\gamma = x_0, \dots, x_n$ be a path with no zero transitions and define the quantity:

$$I_\gamma(q) = \frac{\prod_{i=1}^n q(x_{i-1}, x_i)}{\prod_{i=1}^n q(x_i, x_{i-1})}.$$

Note that this functional satisfies that

- For two paths $\gamma = x_0, \dots, x_n$, $\tau = x_n, x_{n+1}, \dots, x_{n+k}$, and the composition path $\tau \gamma = x_0, \dots, x_{n+k}$, we have that $I_{\tau \gamma}(q) = I_\gamma(q)I_\tau(q)$.
- Given a path $\gamma = x_0, \dots, x_n$ define the inverse path $\gamma^{-1} = x_n, \dots, x_0$. Then $I_{\gamma^{-1}}(q) = [I_\gamma(q)]^{-1}$.
- For two different transitions q_1, q_2 , it happens that $I_\gamma(q_1 q_2) = I_\gamma(q_1)I_\gamma(q_2)$.

So, Theorem 7 is indeed a version of Morera's vectorial calculus theorem: For a continuous function defined on a simply connected domain, it is equivalent to have integral equal to zero over any closed curve or that the value of the integral over any curve only depends on the integration limit points. We include the proof for completeness.

We prove now that (2) implies (3). Let $\gamma = x_0, x_1, \dots, x_{n-1}, x_n$, and $\tau = x_0, y_1, \dots, y_{n-1}, x_n$ be two paths. Then by (2), for the closed curve $\tau^{-1}\gamma$ we have

$$1 = I_{\tau^{-1}\gamma}(q) = I_{\tau^{-1}}(q)I_\gamma(q) = [I_\tau(q)]^{-1}I_\gamma(q),$$

so in this case $I_\gamma(q)$ depends only on x_0 and x_n .

Let us check that (3) implies (2). For a closed curve $C = x_0, x_1, \dots, x_n, x_n = x_0$, let l be an index such that $x_l \neq x_0$. Define $\gamma = x_0, \dots, x_l$ and $\tau = x_l, \dots, x_n$. Then

$$I_C(q) = I_{\tau \gamma}(q) = I_\tau(q)I_\gamma(q) = I_{\gamma^{-1}}(q)I_\gamma(q) = [I_\gamma(q)]^{-1}I_\gamma(q) = 1,$$

where the third equality is due to the assumption (3).

Finally we prove that (3) implies (1). Fix $x_0 \in \mathbb{Z}_+^N$ and define a measure π by

$$\pi(x_0) = 1, \tag{3.3}$$

$$\pi(x) = I_{x_0}^x(q) \quad x \in \mathbb{Z}_+^N \setminus x_0. \tag{3.4}$$

Note that π satisfies (3.2), since

$$I_{x_0}^x(q) \frac{q(x, y)}{q(y, x)} = I_{x_0}^y(q)$$

by choosing as representant to calculate $I_{x_0}^x$ some path from x_0 to x , and for $I_{x_0}^y$ the same path plus the jump from x to y .

□

Last part of the proof of Theorem 7, allow us to express easily the stationary measure of a reversible process in terms of its transitions.

Theorem 8. *If there exists a positive measure π satisfying the detailed balance equations (3.1) for q , an invariant measure is given by:*

$$\pi(x_0) = 1 \tag{3.5}$$

$$\pi(x) = \prod_{i=1}^n \frac{q(x_{i-1}, x_i)}{q(x_i, x_{i-1})}, \quad x \in \mathbb{Z}_+^N \setminus x_0, \quad x_n = x \tag{3.6}$$

□

3.1.1 Fixed birth rates and state dependent death rates processes

The steady analysis of many queuing networks boils down to find the stationary measure of some associated multidimensional birth and death process. In our case, we are interested in processes with fixed birth rates and state dependent death rates:

$$\begin{aligned} q(x, x + e_i) &= \lambda_i, & i = 1, \dots, N, \\ q(x, x - e_i) &= \phi_i(x), & i = 1, \dots, N. \end{aligned}$$

We can characterize the reversibility of these processes in terms of a function called the *balance function*. We use this function many times in the sequel.

Proposition 1. *The following statements are equivalent:*

1. *There exists π solution of the detailed balance equations (3.1).*
2. *The transitions $\phi_i(x)$ verify*

$$\phi_i(x)\phi_j(x - e_i) = \phi_j(x)\phi_i(x - e_j), \quad \forall i, j \in \{1, \dots, N\}, x \in \mathbb{Z}_+^N. \tag{3.7}$$

3. There exists a strictly positive function Φ , the **balance function**, such that the transitions $\phi_i(x)$ write

$$\phi_i(x) = \frac{\Phi(x - e_i)}{\Phi(x)}, \quad \forall i \in \{1, \dots, N\}, x \in \mathbb{Z}_+^N. \quad (3.8)$$

Proof.

First, note that (2) is equivalent to equations (3.2), and then equivalent to (1). If we choose the path $x, x - e_i, x - e_i - e_j$, equation (3.2) is exactly (2), after cancelling λ_i and λ_j from both sides. And, for any path, we can derive (3.2) by applying (2) recursively and multiplying by the corresponding λ 's.

We prove now that (2) and (3) are equivalent. Let $x, x - e_{i_n}, \dots, x - e_{i_n} - e_{i_{n-1}} \dots x - e_{i_2}, x - e_{i_n} - e_{i_{n-1}} \dots x - e_{i_1} = 0$ be a direct path from state x to state 0, i.e., a path of length n , where $n = \sum_{i=1}^N x_i$. Then, property (2) implies that the expression

$$\Phi(x) := \frac{1}{\phi_{i_1}(x)\phi_{i_2}(x - e_{i_1}) \dots \phi_{i_n}(x - e_{i_{n-1}} \dots - e_{i_1})}, \quad (3.9)$$

is independent of the considered direct path. In particular, the rates are uniquely characterized by the function Φ , and it satisfies (3.8).

Conversely, if there exists a function Φ such that the rates $\phi_i(x)$ satisfy (3.8), it can be easily verified that these rates are balanced.

□

We have an equivalent analytic characterization of such type of processes being reversible. This is an analogue of the vectorial calculus result that states that a continuous function defined on a simply connected domain has antiderivative, if and only if, its integral over any closed curve is zero.

Definition 7. We say that a fixed birth rates and state dependent death rates process is **discrete gradient** if there exists a function $\tilde{P} : \mathbb{R}^N \rightarrow \mathbb{R}$ (called the **discrete potential**) such that for all $x \in \mathbb{N}^N$:

$$\log(\phi(x)) = \mathbf{D}\tilde{P}(x) := (\tilde{P}(x) - \tilde{P}(x - e_i))_{i=1, \dots, N}. \quad (3.10)$$

Proposition 2. A fixed birth rates and state dependent death rates process is reversible if and only if it is discrete gradient.

Proof.

Define the discrete potential \tilde{P} as

$$\tilde{P}(x) := -\log \Phi(x), \quad \forall x \in \mathbb{Z}_+^N, \quad (3.11)$$

where $\Phi(x)$ is the balance function. Then the multiplicative property for the balance function (3.8) is transferred into the additive property for the discrete gradient (3.10). \square

As a corollary of Theorem 8 and Equation 3.11, we can express the stationary measure of the process in terms of its potential. As a consequence, we are able to characterize the large deviations behavior of the stationary measure, in the case where the potential scales adequately:

Corollary 2. *The process \tilde{X} associated with a reversible allocation with discrete potential \tilde{P} is reversible (in the usual sense) and its stationary measure (when it exists) is:*

$$\pi(x) = C \lambda^x \exp(-\tilde{P}(x)).$$

Assume further that $\frac{1}{n} \tilde{P}(nx) \rightarrow \gamma(x)$ as $n \rightarrow \infty$. Then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \pi(nx) = -\left(\gamma(x) - \sum_{i=1}^N x_i \log(\lambda_i)\right).$$

\square

3.2 Whittle networks

In this section we give the define processor sharing queues and processor sharing networks. After that, we study the important particular case of reversible processor sharing networks: the Whittle networks. We use results from Section 3.1 to characterize its stationary distribution.

A processor sharing queue is a queue where the service capacity is equally shared between all the clients present in a given time. If $Q(t)$ denotes the number of clients and $c(t)$ the capacity, at time t , hence any client is served at time t with a service rate $\frac{c(t)}{Q(t)}$. We extend this notion to a network.

Let us consider here an open queuing network of N nodes with routing of clients inside the network. At any node i , there are exogenous arrivals as a Poisson process of intensity ν_i . In all the following, we will focus on the process $X(t) = (x_1(t), \dots, x_N(t))$ representing the number of clients at each node of a network.

Definition 8 (Processor-sharing networks). *We call processor-sharing network to a network of processor-sharing nodes with state-dependent capacities. The dynamics of a processor sharing network with exponential service requirements can be characterized by the following principles:*

- *Customer arrive to node j at rate ν_j .*
- *Whenever the network is in state x , the time to the next movement of a single unit from node i to node j is exponentially distributed with rate $p_{ij}\phi_i(x)$ (p_{ij} represents the routing probability from node i to node j while $\phi_i(x)$ represents the capacity of the network at node i).*
- *At node i , customers require exponential services of rate $\phi_i(x)$. After service completion at node i , a costumer is routed to node j with probability p_{ij} and leaves the network with probability $p_i := 1 - \sum_{j \neq i} p_{ij}$.*
- *The transition rates of the number of customers or units at each node are then given by:*

$$q(x, y) = \begin{cases} p_{jk}\phi_j(x) & \text{if } y = x - e_j + e_k \\ p_j\phi_j(x) & \text{if } y = x - e_j \\ \nu_j & \text{if } y = x + e_j \\ 0 & \text{otherwise} \end{cases}$$

Definition 9 (Whittle network). *A Whittle network is a processor sharing network X (Definition 8) with the balance property:*

$$\phi_i(x)\phi_j(x - e_i) = \phi_j(x)\phi_i(x - e_j), \quad i, j = 1, \dots, N, \quad x_i > 0, \quad x_j > 0. \quad (3.12)$$

We assume moreover that $\phi_i(x) > 0$ if and only if $x_i > 0$.

The balance property is directly linked to the reversibility of the process representing the number of customers in the network (see Proposition 1).

These networks are somehow the most general instances of tractable multiclass networks as the capacity (or speed) ϕ_i of node i may depend on the **whole state** of the system $x = (x_1, \dots, x_N)$, where x_i is the number of customers at node i . The particular case of networks where the service rate at each node depends only of the number of clients at this node, are called *Jackson networks*.

We will relax the assumption of exponential services. An important property will arise for Whittle networks: if we have non-exponential services at each node but all other characteristics are preserved we will obtain the same stationary measure as in the exponential case.

3.2.1 Invariant measures for Whittle networks

For simplicity, we focus thereafter on the case of open networks, that is, networks where every customer eventually leaves the system. However, this assumption is not necessary and the same analysis can be extended for closed or mixed networks.

Theorem 9 (Invariant measure). *Let X be an open Whittle network. Then its invariant measure is given by*

$$\pi(x) = \Phi(x) \prod_{i=1}^N \rho_i^{x_i}, \quad (3.13)$$

where Φ is the balance function of the network.

The effective arrival rate λ_i at node i is uniquely defined by the **traffic equations**:

$$\lambda_i = \nu_i + \sum_{j \neq i} \lambda_j p_{ji}, \quad i = 1, \dots, N,$$

and $\rho_i = \lambda_i / \mu_i$, the **traffic intensity** at node i .

Theorem 10. *Let X be an open Whittle network. Define another Markov jump process \tilde{X} , called the adjoint process, by the following transitions:*

$$\begin{aligned} x &\mapsto x + e_i && \text{at rate } \lambda_i, \\ x &\mapsto x - e_i && \text{at rate } \phi_i(x)1(x_i > 0), \end{aligned}$$

The process \tilde{X} is reversible and has the same stationary measure as X .

Proof.

We prove both theorems simultaneously. It follows from the results on Section 3.1, that the adjoint process \tilde{X} is reversible. Hence its stationary measure, π , is given by Theorem 8. Then it is enough to verify the global balance equations for the original process X using the measure π . That is, to verify that for every $x \in \mathbb{Z}_+^N$ it happens

$$\begin{aligned} &\sum_j \pi(x - e_j) \nu_j + \sum_j \pi(x + e_j) p_j \phi_j(x + e_j) + \sum_{j \neq k} \sum_k \pi(x + e_j - e_k) p_{jk} \phi_j(x + e_j - e_k) \\ &= \sum_j \pi(x) \nu_j + \sum_j \pi(x) p_j \phi_j(x) + \sum_{j \neq k} \sum_k \pi(x) p_{jk} \phi_j(x). \end{aligned} \quad (3.14)$$

Let us call A_1, A_2, A_3 the ordered summands in the lefthand side of 3.14, and B_1, B_2, B_3 the ordered summands in the righthand side of 3.14.

By the balance equations for \tilde{X} , we have

$$\lambda_j \pi(x) = \pi(x + e_j) \phi_j(x + e_j), \quad (3.15)$$

which leads to

$$\begin{aligned}
\sum_j \phi_j(x)\pi(x) &= \sum_j \pi(x - e_j)\lambda_j, \\
&= \sum_j \pi(x - e_j)\left(\nu_j + \sum_{k \neq j} p_{kj}\lambda_k\right), \\
&= \sum_j \pi(x - e_j)\nu_j + \sum_j \sum_{k \neq j} p_{kj}\phi(x - e_j + e_k)\pi(x - e_j + e_k).
\end{aligned}$$

Since $p_j = 1 - \sum_k p_{jk}$, we have that $\sum_{i=1}^N \phi_i(x)\pi(x)$ is equal to $B_2 + B_3$, so we have proved that $B_2 + B_3 = A_1 + A_3$.

On the other hand, by summing the traffic equations we have that $\sum_j \nu_j = \sum_j \lambda_j p_j$. Hence:

$$\begin{aligned}
\sum_j \nu_j \pi(x) &= \sum_j \pi(x)\lambda_j p_j, \\
&= \sum_j p_j \phi_j(x + e_j)\pi(x + e_j),
\end{aligned}$$

where we used 3.15 in the second equality. Then $B_1 = A_2$ and the result follows. \square

A stochastic network is generally not Markovian, e.g. when the service requirements distributions are not exponentially distributed. We show in the next section how we can nevertheless use the Markovian analysis developed so far for Whittle networks to deal with more general cases.

3.3 Insensitivity

To deal with non-Markovian networks can be a difficult task. However, if we construct a non-Markovian network by changing the Poisson service processes by general service processes but keeping the mean of the service times and keeping all the other features of the network, there are some cases when this new system shares the stationary distribution with the original Markovian system. In this section, we characterize the networks when this property holds in terms of Whittle networks, defined in Section 3.2.

Definition 10. *A network with Poisson arrivals is said to be **insensitive** when its stationary distribution depends on the service time distribution only via its mean.*

A key question is to characterize the rates $(\phi_j)_{j=1,\dots,N}$ for which a PS network is a Whittle network for any distribution of services and hence to find a necessary and sufficient condition of insensitivity. To proceed with the analysis of this property, we first recall a basic result from renewal theory.

3.3.1 Residual times in renewal processes

Let $N_t = \{T_i\}_{i=0}^\infty$ a point process, i.e., $T_0 = 0$, the sequence $(T_{i+1} - T_i)$ is i.i.d., and let us assume that $T_1 - T_0$ has absolutely continuous distribution function θ and mean m .

We want to compute the distribution of the remaining time to the next point of the point process starting from an arbitrary time t , i.e. the distribution of $R_t = t_{N_t+1} - t$, that we will denote by $\bar{\theta} = \bar{\theta}(t)$. This can also be thought as the residual time until the next arrival seen by an observer arriving uniformly between two points of N . We have the following result about its distribution.

Proposition 3. *Given a stationary renewal point process $N_t = \{T_i\}_{i=0}^\infty$, with the distribution of $T_1 - T_0$ denoted by $\equiv \theta$, the distribution of the residual lifetime is*

$$\bar{\theta}(t) = \frac{1 - \int_t^\infty (1 - \theta(s)) ds}{m}.$$

Proof.

A formal proof can be given using renewal theory (see [Asm03]). We give here a sketch of the proof based on the renewal equation.

Given the stationarity of N_t , it is clear that the distribution of R_t does not depend on t . Using the stationarity we have that

$$E(N_{t+s}) = E(N_t) + E(N_s) \quad \forall s, t \geq 0$$

which implies that there exists a constant c such that $E(N_t) = ct$.

Since this expectancy (the *renewal function*) satisfy the renewal equation, we have that

$$ct = \bar{\theta}(t) + \int_0^t c(t-y)\theta(dy),$$

or

$$\bar{\theta}(t) = ct - \int_0^t c(t-y)\theta(dy).$$

Integrating by parts,

$$\int_0^t c(t-y)\theta(dy) = c(t-y)\theta(y)\Big|_0^t + c \int_0^t \theta(y)dy = c \int_0^t \theta(y)dy,$$

then we have that

$$\bar{\theta}(t) = c \int_0^t (1 - \theta(y))dy.$$

Since $\bar{\theta}(\infty) = 1$, we deduce that

$$1 = c \int_0^t (1 - \theta(y))dy = cm,$$

so we have that

$$\bar{\theta}(t) = \frac{\int_0^t (1 - \theta(y))dy}{m} = \frac{1 - \int_t^\infty (1 - \theta(s))ds}{m}.$$

□

3.3.2 Coupling construction for a processor sharing queue

In this subsection we prove the insensitive property for a single node with a processor sharing policy, using an explicit coupling based on the residual times of the customers.

Consider a processor sharing queue with arrivals as a Poisson process with state dependent rate $\lambda(n)$ and service rates $\phi(n)$, where $n \geq 0$ is the number of clients in the queue. Customers have a generally distributed (absolutely continuous) service requirement with mean 1 and distribution Θ . Processor sharing means that the residual service time of each customer decreases with speed $\frac{\beta(n)}{n}$, when there are n customers in the system.

To keep a Markovian description of the system, it is now necessary to record the residual service times of each customer. The state space S is the product of \mathbb{N} and the union of \mathbb{R}_+^n , for each $n \in \mathbb{N}$ (here the ordering of customers is irrelevant). We denote by $X_t = \{(n_t, R_t^i)\}_{i=1}^{n_t}$ the number of customers and their residual service times at time t . Note that X is not a Markov jump process but belongs to a wider class of Markov process, called as piece-wise deterministic Markov processes (see [DC99] for a definition).

Let $\bar{\theta}^n$ be the product distribution on \mathbb{R}^n , where $\bar{\theta}$ is the distribution of the residual life time for a renewal point process with inter-arrival distribution θ (see Proposition 3). To have a distribution π on S , define $\bar{\theta}_\pi = \sum_n \pi(n)\bar{\theta}^n$.

Theorem 11 ([Zac07]). *If the distribution π is solution of the detailed balanced equations with respect to the transitions $\lambda(\cdot)$ and $\phi(\cdot)$, then $\bar{\mu}_\pi$ is invariant for X_t .*

Proof.

Suppose without loss of generality that θ has mean $m = 1$. Define $\hat{\theta}_n = \bar{\theta}^{n-1} \times \theta$. We are going to couple X_t with a process \hat{X} which has the following dynamics:

- when the workload of a customer is completed, it is immediately replaced by a new customer, having an independent workload with distribution θ .
- there are no arrivals.

For any distribution p , $\bar{\mu}_p$ is stationary for \hat{X} . Let P and \hat{P} the transition kernels of both process and \mathcal{D} the set of functions such that both kernels are continuous at $t = 0$.

For $f \in \mathcal{D}$ we have:

$$\begin{aligned}
\bar{\theta}_\pi P_h f - \bar{\theta}_\pi \hat{P}_h f &= E^{\bar{\theta}_\pi} f(X_h) - E^{\bar{\theta}_\pi} f(\hat{X}_h) \\
&= h \sum_n \pi(n) [\lambda(n)(\hat{\theta}_{n+1} f - \bar{\theta}_n f) + \phi(n)(\bar{\theta}_{n-1} f - \hat{\theta}_n f)] + o(h) \quad (3.16) \\
&= h \sum_n [\pi(n)\lambda(n) - \phi(n+1)\pi(n+1)] (\hat{\theta}_{n+1} f - \bar{\theta}_n f) + o(h) \\
&= o(h).
\end{aligned}$$

In equation (3.16), we have used that:

- The service discipline is processor sharing, then all customers are symmetric (since they receive the same amount of the total service).
- Given the symmetry of service, a completion of service for the original process occurs before time h with probability

$$\begin{aligned}
P\left(\min_{i=1,\dots,n} \bar{\theta}_i \leq h \frac{\beta(n)}{n}\right) &= 1 - P\left(\bar{\theta}_i \geq h \frac{\beta(n)}{n}\right)^n \\
&= 1 - \left[1 - \int_0^{h \frac{\beta(n)}{n}} (1 - \theta(s)) ds\right]^n, \\
&= h \beta(n) + o(h).
\end{aligned}$$

where we know the distribution of $\bar{\theta}_i$ from Proposition 3.

- Since the arrival is Poisson, there is no need to keep track of the time since the last arrival. An arrival occurs before h with probability $\alpha(n)h + o(h)$. When an arrival occurs, the new law of the residual times is changed into $\hat{\theta}_{n+1}$ (keeping in mind that the order of the customers does not matter).

- The probability of more than one event (arrival or service completion) is of order $o(h)$.

The stationarity of $\bar{\theta}_\pi$ with respect to \hat{X} gives us that

$$|\bar{\theta}_\pi P_h f - \bar{\theta}_\pi f| = o(h),$$

uniformly on \mathcal{D} , and we conclude that $\bar{\theta}_\pi$ is stationary for X .

□

3.3.3 Insensitivity of processor sharing networks

Bonald and Proutiere [BP02] gave a necessary and sufficient condition for a processor sharing network to be insensitive:

Theorem 12. *A processor sharing network with fixed arrival rates is insensitive if and only if it is a Whittle network.*

In other words, partial reversibility or reversibility of the adjoint process is a sufficient and necessary condition of insensitivity.

Proof.

The sufficiency of the theorem can be proved following the same approach as for a single PS queue (Theorem 11), combined with the detailed balance equations for the adjoint process. The second part of the proof, the necessary conditions, uses ideas from [BP02].

We use an induction of the number of classes. If $N = 1$, note that the process and the adjoint process coincide, once one has get rid of the loops, and both are reversible.

Suppose the result is true for $N - 1$ classes. Assume for simplicity that there are no transitions for one node to itself. Consider a network of N classes and take the service requirements of node N to be:

- 0, with probability $1 - \alpha$,
- $1/\alpha$, with probability α .

Remark that such a distribution has mean 1. Such network with these service requirements at node N is equivalent to a new network with:

$$\begin{aligned}
\tilde{\nu}_i &= \nu_i + (1 - \alpha)p_{Ni}\nu_N, & i = 1 \dots N - 1, \\
\tilde{p}_{ij} &= p_{ij} + (1 - \alpha)p_{iN}p_{Ni}, & i, j = 1 \dots N - 1, \\
\tilde{\nu}_N &= \alpha\nu_N, \\
\tilde{p}_{iN} &= \alpha p_{iN}, & i = 1 \dots N - 1, \\
\tilde{\phi}_i(x) &= \phi(x), & i = 1 \dots N - 1, \\
\tilde{\phi}_N(x) &= \alpha\phi(x).
\end{aligned}$$

Insensitivity implies that the network has the same invariant measure for all $\alpha > 0$. Letting α goes to zero, we obtain:

$$\begin{aligned}
\tilde{\nu}_i &= \nu_i + p_{Ni}\nu_N, & i = 1 \dots N - 1, \\
\tilde{p}_{ij} &= p_{ij} + p_{iN}p_{Ni}, & i, j = 1 \dots N - 1, \\
\tilde{\nu}_N &= 0, \\
\tilde{p}_{iN} &= 0, & i = 1 \dots N - 1, \\
\tilde{\phi}_i(x) &= \phi(x), & i = 1 \dots N - 1, \\
\tilde{\phi}_N(x) &= 0.
\end{aligned}$$

The limiting balance equations are the balance equations of a $N - 1$ dimensional process. Using the induction assumption, the associated adjoint process is reversible.

Note that the transitions \bar{q} of the adjoint process associated with the $N - 1$ first queues are equal to the transitions of the adjoint process associated with the original N queues. Then for all states such that $x_N = y_N$, we have that

$$\pi(x)\bar{q}(x, y) = \pi(y)\bar{q}(y, x).$$

Since node N was chosen arbitrarily, the equality is valid for all states.

□

3.4 Bandwidth-sharing networks.

Bandwidth-sharing networks describe the evolution of the number of flows (or calls) in a communication network where different classes of traffic compete for the bandwidth. They have become a standard modeling tool over the past decades for modeling

communication networks [BP03, MR02] and have been used in particular to represent the flow level dynamics of a wide range of wireline and wireless networks [BMPV06], generalizing more traditional voice traffic models [Kel79]. In this section we give its formal definition, some toy examples and specific models widely used. Finally we define gradient allocations and monotone networks, and give some examples.

Next, we define mathematically a bandwidth-sharing network.

Definition 11. *A bandwidth-sharing network consists in*

- *some finite set of nodes \mathcal{A} , where each node $a \in \mathcal{A}$ has capacity $C_a \geq 0$,*
- *a set of N routes (or classes) through those nodes, where a route is a sequence of different nodes,*
- *N independent arrival processes,*
- *N independent service processes,*
- *and a function $\phi : \mathbb{R}_+^N \rightarrow \mathbb{R}_+^N$, $\phi(x) = \{\phi_i(x)\}_{i=1}^N$, that satisfies*

$$\sum_{i \text{ route} : a \in i} \phi_i(x) \leq C_a \quad \forall a \in \mathcal{A}.$$

*That function is called a **bandwidth-sharing allocation**.*

Informally, the mechanism of the network is the following: Let us associate each arrival and service process to a route i and call them A_i and S_i , respectively. Given the state of the network as the current number of clientes using each route of network $x = (x_1, \dots, x_n)$, the state will increase to $x + e_i$ when an arrival event occurs at A_i or decrease to $x - e_i$ when a service event occurs at S_i and $x_i > 0$. In intervals where the state x is constant, each route i is using some constant capacity $\phi_i(x)$, as if a perfect fluid flow passes through all the nodes of the route (an idealization of packets transference).

Indeed, these models boil down to a particular class of processor sharing networks with state-dependent service rates: they may depend on the number of flows within the same class, as well as on the numbers of flows in all the other classes.

Assuming that class- i flows arrive subject to a Poisson process of intensity λ_i and require exponentially distributed service times of mean μ_i (the arrival processes of all classes being mutually independent), and in the absence of internal routing, the stochastic process

$X = (X_1, \dots, X_N)$ describing the number of flows (or calls) in progress in the network is a state dependent multi-dimensional birth and death process with transition rates:

$$\begin{aligned} q(x, x - e_i) &= \mu_i \phi_i(x), \\ q(x, x + e_i) &= \lambda_i, \end{aligned}$$

where $x = (x_1, \dots, x_N)$ is the number of flows in each class of traffic.

The service rates of the N traffic classes $\phi = (\phi_i(\cdot))_{i=1}^N$ encodes the particularities of the network resulting from the specific topology, technology, radio conditions, interference and multi-diversity effects and the packet protocols and congestion control mechanisms in use.

3.4.1 Some examples of bandwidth-sharing networks

Consider the linear network represented in Figure 3.1 with four routes sharing three links. While the fourth route passes through all links, each of the routes 1, 2 and 3 only uses one of the links. This gives the following capacity constraints:

$$\begin{aligned} \phi_4(x) + \phi_1(x) &\leq c_1, \\ \phi_4(x) + \phi_2(x) &\leq c_2, \\ \phi_4(x) + \phi_3(x) &\leq c_3. \end{aligned}$$




Figure 3.1: Linear network.

Another example is the tree communication network represented in Figure 3.2. We have two traffic routes, each passing through a dedicated link, followed by a common link. If each dedicated link has a capacity $c_i \leq 1$, $i = 1, 2$, and the common link has capacity 1, the flow on each route gets a capacity $\phi_i(x)$ that lies in the polyhedron \mathcal{C} :

$$\begin{aligned} \phi_i(x) &\leq c_i, \quad i = 1, 2, \\ \sum_{i=1}^2 \phi_i(x) &\leq 1. \end{aligned}$$




Figure 3.2: Tree network.

For examples modelling wireline networks, the vector $\phi(x)$ is usually assumed to belong for all x to a polyhedron describing the capacities constraints of each link that are used by the different routes.

The wireless network is an example as well. In [LSD10], the rate region \mathcal{C} of a 2-station network functioning under the 802.11e protocol is studied. The rate region is the set of achievable throughput vectors at the fine time scale (packet level). Their findings show that the rate region, (which exact expression depends in a complicated manner of the probabilities of transmitting) is generally not convex but is however log-convex. In the sequel, we always assume that \mathcal{C} is log-convex and contains the set $\{\eta : \sum \eta_i \leq c\}$ for some $c > 0$. This is not a restriction for applications.

In general, like for these examples, the capacity constraints determine the space over which a network controller can choose a desired allocation function.

3.4.2 The α -fairness allocations

Some specific bandwidth allocations have received a lot of attention in recent years. Mo and Walrand [MW00] introduced the following family of utility functions:

$$U^\alpha(x, \eta) = \sum_{i=1}^N x_i \frac{\eta_i^{1-\alpha}}{1-\alpha}, \quad \alpha \in (0, \infty) \setminus \{1\}, \quad (3.17)$$

$$U^1(x, \eta) = \sum_{i=1}^N x_i \log(\eta_i), \quad \alpha = 1, \quad (3.18)$$

where the allocation associated is defined by the optimization problem associated to a capacity set \mathcal{C} as follows:

$$\phi^\alpha(x) = \arg \max_{\eta \in \mathcal{C}} U^\alpha(\eta, x). \quad (3.19)$$

The idea behind of this definition is that α is a parameter of “fairness”: the bigger α is, the more fair the allocation is. Consistently, α -fairness allocations include as special cases the **Max-min** fair allocation (Max-min) for α growing to infinity, the **Max** allocation when α goes to zero, the **Proportional fair** allocation (PF) when $\alpha = 1$ (introduced originally by Kelly et. al. in [KMT98]) and the **Potential delay minimization** allocation (MinD) when $\alpha = 2$, see [MR02].

These particular cases have different objectives: Max-min tries to maximize the minimum rate achieved by a flow, Max maximizes the total flow through the network, PF maximizes the total flow through the network but defining the value of a flow allocation $\phi_i(x)$ by $\log(\phi_i(x))$, and MinD minimizes the total potential delay, where the potential flow transfer time in a route is defined to be linearly proportional to the reciprocal of the rate allocation $\frac{1}{\phi_i(x)}$. A common feature for α positive and finite, is the search of a big current total flow while some degree of fairness is kept. An interesting characterization of the α allocations in terms of information divergence measures (some particular cases of Csiszár f -divergence) is given by [UK11].

This framework has been generalized to the weighted α -fair allocations, which provide flexibility to model different levels of fairness in the network. Specifically, the weighted proportional fair allocation $\eta(x)$ for state vector x maximizes

$$\sum_{i=1}^N w_i x_i \log(\eta_i), \quad \eta \in \mathcal{C},$$

where the weights w_i are class-dependent control parameters.

It has been argued in [KMT98] that a good approximation of current congestion control algorithms such as TCP (the Internet’s predominant Transfer Control Protocol) can be obtained by using the weighted proportional fair allocation, which solves an optimization problem for each vector x of instantaneous numbers of flows.

3.4.3 Gradient allocations.

In the characterization of the large deviations of the stationary regime in Chapter 4, we rely on two properties of the allocation function ϕ playing a crucial role and being

intrinsically related: being a gradient allocation or a discrete gradient allocation (defined in Definition 7).

Definition 12. *A bandwidth allocation is called*

- **gradient** if there exists a function $P : \mathbb{R}_+^N \rightarrow \mathbb{R}$ (that we call a continuous potential) such that

$$\log(\phi(x)) = \nabla P(x), \quad \forall x \in \mathbb{R}_+^N \setminus \{0\}.$$

- **discrete gradient** if there exists a function $\tilde{P} : \mathbb{R}^N \rightarrow \mathbb{R}$ (a discrete potential) such that

$$\log(\phi(x)) = \mathbf{D}\tilde{P}(x) := (\tilde{P}(x) - \tilde{P}(x - e_i))_{i=1,\dots,N} \quad \forall x \in \mathbb{N}^N,$$

We show now one example of a gradient allocation. Recall that the proportional fair allocation is defined by the optimization problem associated to a capacity set \mathcal{C} as follows:

$$\phi^{PF}(x) = \arg \max_{\eta \in \mathcal{C}} U^1(\eta, x) = \arg \max_{\eta \in \mathcal{C}} \langle x, \log(\eta) \rangle. \quad (3.20)$$

Following Massoulié [Mas07], observe that the proportional fair bandwidth allocation is gradient. Let $\delta_{\mathcal{A}}^*$ the support function of a bounded convex set \mathcal{A} i.e.:

$$\delta_{\mathcal{A}}^*(x) = \max_{\eta \in \mathcal{A}} \langle x, \eta \rangle.$$

Proposition 4. *Assume that the set \mathcal{C} is log-convex, i.e., the set $\log(\mathcal{C})$ is convex, then*

$$\log(\phi^{PF}(x)) = \nabla \delta_{\log(\mathcal{C})}^*(x), \quad \forall x \in \mathbb{R}_+^N \setminus \{0\}.$$

Proof.

The function δ^* is sub-differentiable because it is convex and finite (see [Roc70]) for all $x \in \mathbb{R}_+^N$. The unicity of the sub-gradient comes from the strict concavity of the log function and implies the differentiability. □

We denote the Proportional fairness potential by $P^{PF} := \delta_{\log(\mathcal{C})}^*$.

By Section 3.2, we know that every Whittle network is a discrete gradient allocation. We show an important example of a gradient allocation: the **Balanced fair** allocation (BF) defined in [BP02] as the allocation ensuring the reversibility of the Markov process X and maximizing the probability of the network to be empty (among the ‘reversible’ allocations). We define formally the BF allocation:

Definition 13 (Balanced fair allocation).

$$\log(\phi^{BF}(x)) = \mathbf{D}\tilde{P}(x),$$

where the potential \tilde{P} is recursively defined by:

$$\tilde{P}(x) = 0, \quad \tilde{P}(x) = \max\{a > 0 : \tilde{P}(x - e_i) - a \in \log(\mathcal{C}), \forall i = 1, \dots, N\}.$$

Since it is gradient, it allows a closed form expression for the stationary distribution of the numbers of flows in progress. In addition, the Balanced fair allocation gives a good approximation of the Proportional fair allocation while being easily evaluated, which is attractive for performance evaluation. See Bonald et. al. [BMPV06] for a comparison analysis of BF, Max-Min and PF allocations.

3.4.4 Monotone networks with general service time distributions.

In this section, we consider a processor sharing network (i.e. a set of processor sharing nodes) with generally distributed flow (or service) sizes. The dynamics of these networks are described for instance in [BP04, Zac07]. This means that a flow is served at node i with speed $\frac{\phi_i(X_i)}{X_i(t)}$. We need the notion of strong monotonicity introduced for instance in [BP04].

Definition 14. ϕ is strongly decreasing, if the function

$$\psi_i(x) = \begin{cases} \frac{\phi_i(x)}{x_i} & \text{if } x_i > 0 \\ 0 & \text{otherwise,} \end{cases}$$

is decreasing in x_j for all $j \neq i$. In that case, we say that the network is **monotone**.

The intuition of a network being monotone is that the allocation cannot increase the capacity of any route without decreasing the capacity per flow of another route. This is verified on all tree topologies and for many wireless networks instances.

We recall from [BP04] the following proposition:

Proposition 5. Let X_t and \tilde{X}_t two processes associated with the allocations ϕ and $\tilde{\phi}$ on the same set of routes. Suppose that ϕ is strongly decreasing, and that $\frac{\phi_i(x)}{x_i} \geq \frac{\tilde{\phi}_i(x)}{x_i}$ for all $x \in \mathbb{Z}_+^N$. Then, for all t and for all service time distribution we have that:

$$X_t \leq_{st} \tilde{X}_t.$$

Proof.

Remark that the processes X_t and \tilde{X}_t are not Markov in general. However, we have a sample-path comparison between the processes by recalling the coupling construction of Subsection 3.3.2.

For both allocations, we use the same arrival processes and record the residual service times of each flow through the network. Define the state space as

$$S := \left(\mathbb{N} \times \bigcup_{n=1}^{\infty} \mathbb{R}_+^n \right)^N,$$

and the processes that saves the residual times by

$$Y_t := \left((n_t^i, \{R_t^{i,j}\}_{j=1}^{n_t^i}) \right)_{i=1}^N, \quad \tilde{Y}_t := \left((\tilde{n}_t^i, \{\tilde{R}_t^{i,j}\}_{j=1}^{\tilde{n}_t^i}) \right)_{i=1}^N,$$

where n_t^i denotes the number of flows at route i and $R_t^{i,j}$ the residual time of j -th flow in the route i , at time t .

Let us suppose that processes X_t and \tilde{X}_t start at positions $x \leq y$. Then, the rates at which each flow is served in the network verify

$$\frac{\phi_i(x)}{x_i} \geq \frac{\phi_i(y)}{y_i} \geq \frac{\tilde{\phi}_i(y)}{y_i}, \quad \forall i \in \{1, \dots, N\}, \quad (3.21)$$

where the first inequality is due to the fact that ϕ is strongly decreasing. Then X_0 has a smaller position than \tilde{X}_0 with a bigger rate of services and the same coupled arrivals, so

$$Y_t(\omega) \leq \tilde{Y}_t(\omega) \quad a.s.$$

for $0 \leq t < \tau_1$, where τ_1 is the next jumping time of X or \tilde{X} . The inequality is preserved for every finite $t > 0$ by induction over over the jumping times of X and \tilde{X} . This is a coupling construction (X, \tilde{X}) where $X_t \leq \tilde{X}_t$, and hence $X_t \leq_{st} \tilde{X}_t$.

□

Bibliography

- [Asm03] S. Asmussen. *Applied Probability and Queues*. Springer, 2003.
- [BMPV06] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Syst. Theory Appl.*, 53(1-2):65–84, 2006.

- [BP02] T. Bonald and A. Proutière. Insensitivity in processor-sharing networks. *Performance Evaluation*, 49(14):193 – 209, 2002.
- [BP03] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Syst. Theory Appl.*, 44(1):69–100, 2003.
- [BP04] T. Bonald and A. Proutière. On stochastic bounds for monotonic processor sharing networks. *Queueing Syst. Theory Appl.*, 47(1/2):81–106, 2004.
- [DC99] F. Dufour and O. Costa. Stability of piecewise-deterministic Markov processes. *SIAM J. Control Optim.*, 37(5):1483–1502 (electronic), 1999.
- [Jon10] M. Jonckheere. *Introduction to Stochastic Networks*. Lecture notes, course given at the Universidad de Buenos Aires, 2010.
- [Kel79] F. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.
- [KMT98] F. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *The Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [LSD10] D. Leith, V. Subramanian, and K. Duffy. Log-convexity of rate region in 802.11e w lans. *IEEE Communications Letters*, 14(1):57–59, 2010.
- [Mas07] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.
- [MR02] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans. Netw.*, 10(3):320–328, 2002.
- [MW00] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. on Netw.*, 8(5):556–567, 2000.
- [Roc70] T. Rockafellar. *Convex analysis*, volume 28. Princeton landmarks in mathematics, 1970.
- [UK11] M. Uchida and J. Kurose. An information-theoretic characterization of weighted -proportional fairness in network resource allocation. *Information Sciences*, 181(18):4009 – 4023, 2011.
- [Zac07] S. Zachary. A note on insensitivity in stochastic networks. *Journal of Applied Probability*, 44(1):238–248, 2007.

Chapter 4

Large deviations for Proportional Fairness allocation

4.1 Why to calculate large deviations for Proportional fairness allocation?

4.1.1 Insensitivity in telecommunication systems

In 1917 an important problem concerning telephone exchanges was considered by Erlang on [Erl18] while he was working on the Copenhagen Telephone Company, [Sto07]. Let us suppose that we have K telephonic lines, telephonic calls that arrive at rate λ , and the duration time of each call is exponential(μ) distributed. Calls arriving while all lines are busy are lost. Let $\rho := \frac{\lambda}{\mu}$ and $E_K(\rho)$ the fraction of calls that are lost in steady state. What is the value of $E_K(\rho)$? We give the following the answer obtained by Erlang.

We can model the number of busy lines as a one-dimensional birth and death process with positive rates of jumps given by

$$\begin{aligned}q(x, x + 1) &= \lambda_i, & \text{for } x = 0, 1, \dots, K - 1, \\q(x, x - 1) &= x \mu & \text{for } x = 1, \dots, K.\end{aligned}$$

This system is usually known as the $M \setminus M \setminus K$ queue, that is, the queue with Markovian arrival and service processes and K servers.

On the other hand, let us consider the $M \setminus M \setminus \infty$ queue, that is, the one-dimensional birth and death process with positive rates of jumps given by

$$\begin{aligned}q(x, x + 1) &= \lambda_i, & \text{for } x = 0, 1, \dots \\q(x, x - 1) &= x \mu & \text{for } x = 1, 2, \dots\end{aligned}$$

It is straightforward to note that the $M \setminus M \setminus \infty$ queue is reversible with respect to the Poisson distribution:

$$\pi(x) = \frac{\rho^x e^{-\rho}}{x!}, \quad x = 0, 1, \dots$$

by checking the Balance equations (Equation 3.1).

The stationary distribution of the $M \setminus M \setminus K$ queue, denoted by π^K , can be obtained as the stationary distribution of an $M \setminus M \setminus \infty$ queue conditioned to take values on $\{0, 1, \dots, K\}$, see [Asm03] for the details. Then

$$\pi^K(x) = \frac{\frac{\rho^x}{x!}}{1 + \rho + \dots + \frac{\rho^K}{K!}}, \quad x = 0, 1, \dots, K,$$

and $E_K(\rho)$ correspond to the probability of having all lines occupied,

$$E_K(\rho) = \frac{\frac{\rho^K}{K!}}{1 + \rho + \dots + \frac{\rho^K}{K!}}. \quad (4.1)$$

Equation 4.1 is usually known as *Erlang's B loss formula*.

Erlang's B loss formula characterizes easily the probability of having a lost call, depending only on the parameter ρ , usually known as the *traffic intensity*. It is an easy to apply formula, solved in practise by some recurrence equations depending on K . It was indeed applied short time after its publication in telephonic exchange systems, see [Sto07]. A very surprising fact is that the formula is a sharp approximation to the performance of real systems. The approximation of the arrival process by a Poisson process when the population of clients is big enough, was already studied by Erlang in a previous work [Erl09]. However, assuming that the service time of a client has exponential distribution was not justified. The subyacent reason of the robustness of the formula is the insensitivity property of the system: the $M \setminus M \setminus K$ is a Whittle network and then it is insensitive, by Theorem 12, so the formula depends on the distribution of the service time only by its traffic intensity.

Since queueing systems with the insensitivity property have stationary distribution that can be estimated by simple statistics estimators, such property became a key property. For a wide class of modern and more complex telecommunication models, known as Kelly-Whittle networks, the insensitive property still holds (see [Bon07] for a survey of such models). However, there are models widely used in practise where the insensitivity property does not hold. Despite that, in this chapter we prove that a weak version of this property, called *weak insensitivity*, still holds in a general context. This includes the allocation of bandwidth-sharing capacity of the Internet network as a direct application.

4.1.2 The Proportional fairness election.

In the general context, the selection of a specific bandwidth allocation is motivated by several properties of the resulting process X . Among those properties, the following ones are of particular interest:

1. *maximal stability*: one expects the bandwidth sharing mechanism to stabilize the system whenever it can be stabilized,
2. *decentralized protocol*: the bandwidth allocation can be implemented in the network using decentralized schemes,
3. *insensitivity*: when changing the traffic conditions¹ and in particular the service time distribution (but keeping the mean flow size fixed), one could expect the stationary measure of X to remain the same, in which case the system and the bandwidth allocation are said to be insensitive,
4. *weak insensitivity*: the large deviations characteristics do not depend on the service time distribution, except for its mean and coincide with the large deviations of the most efficient insensitive allocation.

Table 4.1 show which properties are satisfied for the bandwidth sharing allocations introduced in Chapter 3: Max, Max-Min, PF, α -fair, and BF allocations.

	Maximum stability	Decentralized protocol	Insensitivity
Max	NO	YES	NO
Max-Min	YES	YES	NO
PF	YES	YES	NO
α -fair	YES	YES	NO
BF	YES	?	YES

Table 4.1: Properties of allocations

Max is an inestable allocation while α -fair [BM01] and Max-min [dVLK99] allocations reach maximum stability (that is, to be positive recurrent for all λ in the interior of the capacity set).

¹When the size distributions of the flows are not exponentially distributed, the process X is not Markov by itself anymore and the dynamics have to be defined using the residual service time.

By the characterization theorem of insensitive allocations as Whittle networks (Theorem 12), we know that except BF, all allocations fail to satisfy insensitivity (except on the case of very particular topologies). On the other hand, the BF allocation satisfies maximum stability and insensitive but it is not known whether it can be implemented with a decentralized protocol, [BMPV06].

This leaves us weak insensitivity as a central criterion for choosing a bandwidth sharing allocation. It was shown in [Mas07] that an appropriate modification of the proportional fair allocation, called modified proportional fair allocation (mPF),

$$\log(\phi^{mPF}(x)) = \mathbf{D}P^{PF}(x),$$

(with P^{PF} being the continuous potential associated with PF) coinciding asymptotically (point-wise) with PF, has the same large deviations characteristics as BF. On the other hand, it has been recently proven that an insensitive allocation being maximal stable, is asymptotically equivalent (point-wise) to PF [Wal11].

However, it remained as an open problem to prove that the large deviations characteristics of the stationary measure of the PF allocation itself coincide with those of mPF, and BF as it was conjectured in [Mas07]. The major difficulty of this problem is that the stationary measure of a network under the proportional fair allocation does not have a closed form in general, except in the very particular cases of symmetric hypergrids, see [Pro03].

Our main contribution consists in overcoming these difficulties for the specific processes we are studying. We first show that for Poisson arrivals and exponentially distributed flow sizes, the stationary distribution π^{PF} of the number of flows associated with PF and the stationary distribution π^{mPF} associated with the mPF allocation have the same large deviations characteristics. More precisely:

Theorem 13. *For all $x \in \mathbb{Z}_+^N$:*

$$\left| \frac{1}{n} \log \left(\frac{\pi^{PF}(\{nx\}^\uparrow)}{\pi^{mPF}(\{nx\}^\uparrow)} \right) \right| = O(n^{-\frac{1}{2}+\epsilon}), \quad \forall \epsilon > 0,$$

where $\pi(\{nx\}^\uparrow) = \sum_{k \in \mathbb{Z}_+^N: k \geq nx} \pi(k)$. In the particular case that $x \in \mathbb{N}^N$ the bound can be improved to $O(n^{-1})$.

By recalling the next result which established that BF and mPF share large deviations asymptotics,

Proposition 6. (*Massoulié, 2007*)

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \pi^{BF}(nx) = \sum_{i=1}^N x_i \log(\lambda_i) - P^{mPF}(x).$$

we conclude that PF, BF and mPF share the large deviations rate function, in the case of Markovian services.

As a consequence, we obtain the explicit tail asymptotics for PF allocation:

Corollary 3. *For all $x \in \mathbb{Z}_+^N$:*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log(\pi^{PF}(\{nx\}^\uparrow)) = - \sum_{i=1}^N \log\left(\frac{\phi_i(x)}{\lambda_i}\right).$$

The proof of Theorem 13 is achieved by first proving the geometric ergodicity of both the mPF allocation and the PF allocation. For that purpose, we exhibit appropriate Lyapunov functions, relying on some structural results of PF described in [Mas07]. This then allows us to use simple martingale arguments.

For monotone networks and generally distributed flow size, we give a more direct proof establishing that the large deviations characteristics are actually insensitive to the service time distribution. This shows that the proportional fair allocation indeed satisfy properties (1) and (2) and (4) at least on monotone topologies.

4.2 The Freidlin-Wentzell approach

To understand the links between our results and the classical theory of large deviations for Markov processes, it is enlightening to recall the principles of the Freidlin and Wentzell [FW84] theory for birth and death processes on \mathbb{Z}^N , $N \in \mathbb{N}$, with rates being Lipschitz and with bounded logarithms.

Let Y^n be a multi-dimensional birth and death process with transition rates:

$$q\left(\frac{x}{n}, \frac{x}{n} - \frac{e_i}{n}\right) = n\phi_i\left(\frac{x}{n}\right),$$

$$q\left(\frac{x}{n}, \frac{x}{n} + \frac{e_i}{n}\right) = n\lambda_i.$$

Suppose additionally that the death rates are 0-homogeneous, i.e. $\phi(az) = \phi(z)$, $\forall z \in \mathbb{R}^N/n, a > 0$. Define the logarithmic moment-generating of the increment of the process for $z \in \mathbb{R}^N/n, y \in \mathbb{R}^N$ by

$$H_n(z, y) := \frac{d}{dt} E^z[\exp\langle y, Y^n(t) - z \rangle].$$

Using the structure of the generator and the 0-homogeneity of the rates we have that

$$H_n(z, y) = n \left(\sum_i \lambda_i (e^{y_i/n} - 1) + \phi_i(z) (e^{-y_i/n} - 1) \right).$$

Then we obtain that $H_n(z, y/n) = nH(z, y/n)$ with

$$H(z, y) := \langle e^y - 1, \lambda \rangle + \langle e^{-y} - 1, \phi(z) \rangle.$$

Now, define the Fenchel-Legendre L transform of H by $L(x, y) := \sup_{\theta \in \mathbb{R}^N} \langle y, \theta \rangle - H(x, \theta)$, the action functional $S : \mathcal{C}[0, T] \rightarrow \mathbb{R}$ as

$$S_T(r) = \begin{cases} \int_0^T L(r_s, \dot{r}_s) ds & \text{if } r \text{ is absolutely continuous,} \\ \infty & \text{otherwise,} \end{cases}$$

and the quasi-potential R by

$$R(x) = \min_{r, T, r(0)=0, r(T)=x} S_T(r, \dot{r}).$$

The original results of Freidlin and Wentzell are stated for Lipschitz-continuous transitions (extended to \mathbb{R}^N):

Theorem 14 (Freidlin and Wentzell, [FW84]). *Assume that for each i , the function $\log \phi_i(\cdot)$ is bounded and Lipschitz continuous and that the function L is such that:*

$$\lim_{|x-\tilde{x}| \rightarrow 0} \frac{|L(x, y) - L(\tilde{x}, y)|}{1 + L(x, y)} \rightarrow 0,$$

uniformly in x and \tilde{x} . Let π^n the stationary measure of the birth and death processes Y^n . Then, a large deviations principle holds for the family of probabilities $\pi^n(\cdot)$ with rate function $R(\cdot)$.

In the case of general multi-dimensional birth and death processes, it is difficult to solve the variational problem from which the quasi-potential R is defined. However, assuming that the rates are gradient greatly simplifies the expression of the quasi-potential:

Proposition 7. *Assume that*

$$\log(\phi(x)) = \nabla P(x), \quad \forall x \in \mathbb{Z}^N,$$

for some differentiable function $P : \mathbb{R}^n \rightarrow \mathbb{R}$. Then the quasi-potential of the birth and death process is equal to

$$R(x) = -\langle \log(\lambda), x \rangle + P(x).$$

Proof.

We first show that:

$$R(x) = \max_{x_s: x_0=0, x_T=x} \int_0^T \langle \dot{x}_s, \log\left(\frac{\lambda}{\phi(x_s)}\right) \rangle ds.$$

The first simplification comes from the time homogeneity of the process. Using the Euler-Lagrange principle, combined with the time homogeneity (which implies that $\nabla_x L = 0$) the minimizing path satisfies the Beltrami identity:

$$L - \langle y, \nabla_y L \rangle = 0.$$

Now, observe that if $\theta_y = \arg \max \langle y, \theta \rangle - H(x, \theta)$, we obtain that $\nabla_y L = \theta_y$ and $L - \langle y, \nabla_y L \rangle = H(x, \theta_y) = 0$. Solving the last equation leads to

$$\theta_{\dot{x}} = \frac{1}{2} \log\left(\frac{\phi(x)}{\lambda}\right),$$

which allows us to conclude since

$$L(x, \dot{x}) = \langle \dot{x}, \nabla_y L(x, \dot{x}) \rangle = \langle \dot{x}, \theta_{\dot{x}} \rangle.$$

The expression for the quasi-potential then follows from integration.

□

We insist on the fact that this theory allows to prove the large deviations rate of the stationary measure of birth and death processes only for smooth allocations and does not provide a rate of convergence. As underlined for instance in [SW95], it is very demanding to extend these results to processes on state spaces with boundaries since the technical conditions of the classical theory are never fulfilled. (Remark that boundness of the logarithm is never verified for birth and death processes defined in the orthant. Moreover, the Lipschitz assumption is not verified in our case for most network topologies).

4.3 Large deviations for Markovian dynamics.

Before proving Theorem 13, we illustrate it on a toy example.

Example 1 (Single link). *Consider first a single link shared by N classes. This corresponds to choosing $\mathcal{C} = \{\sum \frac{\eta_i}{\mu_i} \leq 1\}$, where the $\{\mu_i\}$ correspond to the mean flow sizes. In this case, the proportional fair allocation coincides with the balanced fairness allocation. Its service rates are given by:*

$$\phi_i(x) = \mu_i \frac{x_i}{|x|}.$$

Note that the continuous and the discrete potential functions associated with the allocation ϕ do not coincide:

$$\delta_{\log(\mathcal{C})}^*(x) = \sum_{i=1}^N \frac{x_i}{|x|} \log \left(\frac{x_i}{|x|} \right),$$

while the discrete potential is given by:

$$P(x) = \log \left(\binom{|x|}{x_1, \dots, x_N} \right).$$

One can however prove using the Stirling formula that

$$\frac{1}{n} (P(nx) - \delta_{\log(\mathcal{C})}^*(nx)) \rightarrow 0, \text{ as } n \rightarrow \infty.$$

An extension of these formulae can be obtained for hypergrids topologies [BP03].

In this section, we call X the process associated with the PF allocation, denoted ϕ and corresponding to the continuous gradient of the potential P^{PF} . We further denote by \tilde{X} the process associated to the allocation $\tilde{\phi}$ itself corresponding to the discrete gradient of P^{PF} , i.e.:

$$\log(\phi(x)) = \nabla P^{PF}(x), \tag{4.2}$$

$$\log(\tilde{\phi}(x)) = \mathbf{D}P^{PF}(x) = (P^{PF}(x) - P^{PF}(x - e_i))_{i=1\dots N}. \tag{4.3}$$

Structural properties of the PF allocation.

From the structural representation of PF, we know that $\tilde{\phi}$ is a perturbation of the original rates ϕ in the sense that for all $x_i > 0$ (see Lemma 9 in [Mas07]):

$$|\log(\phi_i(x)) - \log(\tilde{\phi}_i(x))| = |\log(\phi_i(x)) - (P^{PF}(x) - P^{PF}(x - e_i))| \leq \frac{1}{x_i}, \quad \forall i. \tag{4.4}$$

Geometric ergodicity.

We first consider the dynamics corresponding to the Proportional fairness allocation. For a Markov jump process with countable state space E , rates of jump $\{q(x, y)\}$ and a function $F : E \rightarrow \mathbb{R}$, define the drift of F as

$$\Delta F(x) := \sum_{y \in E} q(x, y)(F(y) - F(x)).$$

Proposition 8 (Geometric ergodicity of X). *Suppose $\lambda \in \text{int}(\mathcal{C})$, then there exists a function $G : \mathbb{R}_+^N \rightarrow \mathbb{R}$ and constants K and $\gamma > 0$ such that for $|x| > K$:*

$$\Delta G(x) \leq -\gamma G(x).$$

Denote by π^{PF} the stationary distribution of X , and P_t^0 the distribution at time t of process X starting from zero. Hence:

$$\left| P_t^0(x) - \pi^{PF}(x) \right| \leq K_1 e^{-K_2 t},$$

for some constants $K_1, K_2 > 0$.

Proof.

The proof has two steps. We first construct a Lyapunov function with bounded drift. We then use this Lyapunov function to construct a new one verifying a geometric drift inequality.

First step: let $F(x) = (\sum_{i=1}^N \frac{x_i^2}{\lambda_i})^{1/2}$. Observe that F is a norm in \mathbb{R}^N (hence is positive, 1-homogeneous, and diverges to infinity when $|x| \rightarrow \infty$). Furthermore, it is C^2 for all $x \neq 0$ and:

$$\frac{\partial F(x)}{\partial x_i} = \frac{x_i}{\lambda_i F(x)}.$$

Then, there exists $K > 0$ such that for $|x| > K$

$$\frac{1}{|x|} \sup_{z=x, x \pm e_i} \left| \frac{\partial^2 F(z)}{\partial^2 x_i} \right| \leq \epsilon.$$

Using that F is 1-homogeneous, this leads for $|x| > K$ to:

$$\begin{aligned} \Delta F(x) &\equiv \sum_{i=1}^N \lambda_i (F(x + e_i) - F(x)) + \phi_i(x) (F(x - e_i) - F(x)), \\ &\leq \langle \lambda - \phi(x), \nabla F(x) \rangle + \epsilon, \\ &= \left\langle \lambda - \phi(x), \frac{x}{F(x)\lambda} \right\rangle + \epsilon, \\ &= \frac{|x|}{F(x)} \left\langle \lambda - \phi(x), \frac{x}{|x|\lambda} \right\rangle + \epsilon. \end{aligned}$$

Using the definition of Proportional fairness (as the allocation maximizing $U^1(\cdot, x)$ and the strict concavity of the log function), for all $\eta \in \mathcal{C}$,

$$\sum_{i=1}^N \frac{x_i}{\eta_i} (\eta_i - \phi_i(x)) = \left\langle \left(\frac{\partial U(x, \eta)}{\partial \eta_i} \right), \eta - \phi \right\rangle < 0,$$

which implies that there exists $\gamma_1 > 0$ such that:

$$\left\langle \lambda - \phi(x), \frac{x}{|x|} \frac{1}{\lambda} \right\rangle \leq -\gamma_1,$$

which in turn implies (together with the fact that F is norm-like) that there exists $\gamma_2 > 0$ such that $\Delta F(x) \leq -\gamma_2$ for all $|x| > K$. Remark also that for all x , $|\Delta F(x)| \leq C$.

Second step: we now calculate the drift of $G(x) := \exp(\delta F(x))$:

$$\begin{aligned} \Delta G(x) &= \sum_{i=1}^N \lambda_i (e^{\delta F(x+e_i)} - e^{\delta F(x)}) + \phi_i(x) (e^{\delta F(x-e_i)} - e^{\delta F(x)}), \\ &= G(x) \sum_{i=1}^N \lambda_i (e^{\delta(F(x+e_i)-F(x))} - 1) + \phi_i(x) (e^{\delta(F(x-e_i)-F(x))} - 1), \\ &\leq G(x) (\delta \Delta F(x) + C_1 \delta^2 \sum_{i=1}^N \exp(c|\Delta F(x)|) + O(\delta^3)), \\ &\leq G(x) (-C_2 \delta + C_3 \delta^2) \leq -\gamma G(x). \end{aligned}$$

By a continuous time version of the results of [MT93], this implies that :

$$|P^0(X(t) = x) - \pi^{PF}(x)| \leq K_1 \exp^{-K_2 t}.$$

□

We now deal with the dynamics of the modified proportional fair allocation. In the following proposition, we use the same functions F and G defined in the Proportional fairness case.

Proposition 9 (Geometric ergodicity of \tilde{X}). *Suppose $\lambda \in \text{int}(\mathcal{C})$, then there exist constants \tilde{K} and $\tilde{\gamma} > 0$ such that for $|x| > \tilde{K}$:*

$$\Delta G(x) \leq -\tilde{\gamma} G(x).$$

Denote by π^{mPF} the stationary distribution of X , and \tilde{P}_t^0 the distribution at time t of process \tilde{X} starting from zero. Hence:

$$\left| \tilde{P}_t^0(x) - \pi^{mPF}(x) \right| \leq \tilde{K}_1 e^{-\tilde{K}_2 t},$$

for some constants $\tilde{K}_1, \tilde{K}_2 > 0$.

Proof.

We proceed as previously. The only difference consists in proving that the perturbations of the rates are small enough to be negligible in the drift calculations.

For $|x| > \tilde{K}$, using the bounds on the difference of ϕ and $\tilde{\phi}$, we obtain that:

$$|\phi_i(x) - \tilde{\phi}_i(x)| \leq c_0 |\log(\phi_i(x)) - \log(\tilde{\phi}_i(x))| \leq \frac{c_0}{x_i}, \quad \text{for } x_i > 0.$$

Hence:

$$\begin{aligned} \Delta F(x) &= \sum_{i=1}^N \lambda_i (F(x + e_i) - F(x)) + \tilde{\phi}_i(x) (F(x - e_i) - F(x)), \\ &\leq \langle \lambda - \tilde{\phi}(x), \nabla F(x) \rangle + \epsilon, \\ &\leq \left\langle \lambda - \phi(x), \frac{x}{F(x)\lambda} \right\rangle + \left\langle \phi(x) - \tilde{\phi}(x), \frac{x}{F(x)\lambda} \right\rangle + \epsilon, \\ &\leq -\gamma_1 + \left\langle \phi(x) - \tilde{\phi}(x), \frac{x}{F(x)\lambda} \right\rangle, \\ &\leq -\gamma_1 + \frac{c}{|x|}. \end{aligned}$$

which gives (together with the fact that F is norm-like) that there exists $\gamma_2 > 0$ such that $\tilde{\Delta}F(x) \leq -\gamma_2$ for all $|x| > \tilde{K}$. Remark also that $|\tilde{\Delta}F(x)| \leq C$. We can conclude exactly as in the previous proof. □

Before proving the main result of this section, we establish a useful lemma using the geometric ergodicity of the processes.

Lemma 8. *There exists $C > 0$ such that for $t_n = Cn$, we have*

$$\log \left| \frac{\mathbb{P}^0(X_{t_n} \geq nx)}{\pi^{PF}(\{nx\}^\uparrow)} \right| \leq C_1 \exp\{-C_2 n\}$$

for some $C_1, C_2 > 0$, where $\pi^{PF}(\{k\}^\uparrow) = \sum_{j \in \mathbb{Z}_+^N: \geq k} \pi^{PF}(j)$. The analogue inequality is true for the process \tilde{X} .

Proof.

We need to bound from below the stationary probabilities of the process. This is easy since the rates are bounded. Let us define a process with the same arrival rates and the service rates equal to φ , the maximum of ϕ_i for all coordinates. It is clear that

$$\pi^{PF}(\{nx\}^\uparrow) \geq K_1 \left(\frac{\lambda}{\varphi} \right)^n,$$

for some constant $K_1 > 0$.

We have

$$\begin{aligned} \log \left| \frac{\mathbb{P}^0(X_{t_n} \geq nx)}{\pi^{PF}(\{nx\}^\uparrow)} \right| &\leq \left| 1 - \frac{\mathbb{P}^0(X_{t_n} \geq nx)}{\pi^{PF}(\{nx\}^\uparrow)} \right| = \left| \frac{\mathbb{P}^0(X_{t_n} \geq nx) - \pi^{PF}(\{nx\}^\uparrow)}{\pi^{PF}(\{nx\}^\uparrow)} \right| \\ &\leq \frac{K_2 e^{-K_3 t_n}}{\pi^{PF}(\{nx\}^\uparrow)} \\ &\leq \frac{K_2 e^{-K_3 t_n}}{K_1 \left(\frac{\lambda}{\varphi}\right)^n} = \frac{K_2}{K_1} \exp\{-K_3 t_n + \log\left(\frac{\varphi}{\lambda}\right)n\} \end{aligned}$$

where first inequality comes by $\log(x) \leq |1 - x|$ for all $x > 0$, and the second one by the geometric ergodicity of X . We take $C > 0$ such that $K_3 C > \log\left(\frac{\varphi}{\lambda}\right)$. Note that the same argument remains valid for process \tilde{X} taking its respective transition rates and stationary distribution. □

Change of measure and control of the martingale.

To relate the distribution of X and \tilde{X} , we recall the Proposition B.6 513 of Shwartz and Weiss [SW95]:

Proposition 10. *Let Y_t^i $i = 1, 2$ be two multidimensional birth-death process in \mathbb{Z}_+^N with step directions e_j , bounded step rates $\{q_j^i(x)\}_{x \in \mathbb{Z}_+^N}$ and law P^i . Assume that for all x and j , $q_j^1(x) = 0$ if and only if $q_j^2(x) = 0$. Then we can relate the distributions of the processes by $dP^2 = M_t dP^1$ where*

$$M_t = \exp \left(\int_0^t \sum_{j=1}^N \log \left(\frac{q_j^2(X_{s^-})}{q_j^1(X_{s^-})} \right) dN_s^j - \int_0^t \sum_{j=1}^N (q_j^2(X_s) - q_j^1(X_s)) ds \right).$$

Here $\{N^j\}$ denote a family of counting processes describing the jumps of the processes in the j -th direction and M_t is a càdlàg martingale.

In our case, $\phi_i(x)$, $\tilde{\phi}_i(x)$ are both positive whenever $x_i > 0$, and 0 otherwise. All the rates of X and \tilde{X} are bounded so we meet the conditions of the previous Proposition and we can write that $d\tilde{P}(\omega) = M_t dP(\omega)$, with

$$\begin{aligned} M_t &= \exp \left(\int_0^t \sum_{j=1}^N 1_{\{X_{s^-} > 0\}} \log \left(\frac{\tilde{\phi}_j(X_{s^-})}{\phi_j(X_{s^-})} \right) dN_s^j - \int_0^t \sum_{j=1}^N 1_{\{X_s > 0\}} (\tilde{\phi}_j(X_s) - \phi_j(X_s)) ds \right), \\ &= \prod_{j=1}^N M_t^j, \end{aligned}$$

and each of the factors M^j is itself a martingale (with the obvious definition of M^j).

Remark 1. *An important observation for the following is that the counting processes N^j can be seen as thinning (according to X_s) of some Poisson processes \hat{N}^j which are all independent.*

We denote by $M_{s,t}$ the last expression with the integrals running from s to t , so $M_t = M_{0,t}$. The change of measure formula is then $E^0[1_{\{\tilde{X}_t \geq nx\}}] = E^0[M_t 1_{\{X_t \geq nx\}}]$. We can now prove our main result.

Proof of Theorem 13 :

We consider first the case where $x_i > 0$ for all $i = 1 \dots, N$.

Define the sequence of stopping times:

$$\begin{aligned} \tau_1 &= 0, \\ \tau_i &= \inf\{t > \tau_{i-1}, \tilde{X}_t \geq nx\}, \text{ for } i \text{ even}, i \geq 2, \\ \tau_i &= \inf\{t > \tau_{i-1}, \tilde{X}_t < nx\}, \text{ for } i \text{ odd}, i \geq 3. \end{aligned}$$

Observe that if $\tilde{X}_t \geq nx$, then necessarily, there exists k even (a.s. finite) such that $\tau_k \leq t \leq \tau_{k+1}$. Using the Markov property and the martingale property, we then have

$$\begin{aligned} E^0(M_t 1_{\{\tilde{X}_t \geq nx\}}) &= E^0 \sum_{k \text{ even}} 1_{\tau_k \leq t < \tau_{k+1}} M_t 1_{\{\tilde{X}_t \geq nx\}} \\ &= \sum_{k \text{ even}} E^0 1_{\tau_k \leq t < \tau_{k+1}} M_{0,\tau_k} M_{\tau_k,t} 1_{\{\tilde{X}_t \geq nx\}}, \\ &= \sum_{k \text{ even}} E^0 (E(1_{\tau_k \leq t < \tau_{k+1}} M_{0,\tau_k} M_{\tau_k,t} 1_{\{\tilde{X}_t \geq nx\}} | \tilde{\mathcal{F}}_{\tau_k})) \\ &= \sum_{k \text{ even}} E^0 (M_{0,\tau_k} E(1_{\tau_k \leq t < \tau_{k+1}} M_{\tau_k,t} 1_{\{\tilde{X}_t \geq nx\}} | \tilde{X}_{\tau_k})), \end{aligned}$$

with $\{\tilde{\mathcal{F}}_k\}$ being the natural filtration of process \tilde{X} .

We define $g_k(t) = \sup_y E(1_{\tau_k \leq t < \tau_{k+1}} M_{\tau_k,t} 1_{\{\tilde{X}_t \geq nx\}} | \tilde{X}_{\tau_k} = y)$, so

$$E^0(M_t 1_{\{\tilde{X}_t \geq nx\}}) \leq \sum_{k \text{ even}} g_k(t) E^0(M_{0,\tau_k}) = \sum_{k \text{ even}} g_k(t).$$

Remark that on $\{\tilde{X}_t \geq nx\} \cap \{\tau_k \leq t < \tau_{k+1}\}$ (k even) we have $\{\tilde{X}_s \geq nx : s \in [\tau_k, t]\}$. Recall that x is such that $x_i > 0$ for all i . Using the assumption on ϕ and $\tilde{\phi}$:

$$\begin{aligned} E^0 \left(1_{\tau_k \leq t < \tau_{k+1}} 1_{\{\tilde{X}_t \geq nx\}} \exp \left(\int_{\tau_k}^t \sum_{j=1}^N 1_{\{X_{s-} > 0\}} \log \left(\frac{\tilde{\phi}_j(X_{s-})}{\phi_j(X_{s-})} \right) dN_s^j - \int_{\tau_k}^t \sum_{j=1}^N 1_{\{X_{s-} > 0\}} (\tilde{\phi}_j(X_s) - \phi_j(X_s)) ds \right) \right) \\ \leq E^0 \left(1_{\tau_k \leq t < \tau_{k+1}} \exp \left\{ \sum_{j=1}^N \frac{C_1}{n} (N_t^j + t) \right\} 1_{\{\tilde{X}_t \geq nx\}} \right), \\ \leq E^0 \left(1_{\tau_k \leq t < \tau_{k+1}} \exp \left\{ \frac{C_1 N}{n} (\bar{N}_t + t) \right\} 1_{\{\tilde{X}_t \geq nx\}} \right), \end{aligned}$$

where \bar{N}_t is a Poisson process with parameter $\bar{\lambda}$ equal as the maximum of the parameters of the Poisson processes $\{N_t^j\}_{j=1}^N$. Summing these inequalities and using Hölder's inequality for $p > 1$:

$$E^0(M_t 1_{\{\tilde{X}_t \geq nx\}}) \leq \exp \left\{ \frac{C_1 N t}{n} \right\} \mathbb{P}^0(\tilde{X}_t \geq nx)^{\frac{p-1}{p}} \exp \{ \bar{\lambda} t [e^{\frac{C_1 p N}{n}} - 1] \}^{\frac{1}{p}}.$$

We now choose $t_n = C_2 n$ such the result of Lemma 8 holds, and the sequence $p_n = n$ to obtain

$$\begin{aligned} \frac{1}{n} \log \mathbb{P}^0(X_{t_n} \geq nx) &\leq \frac{C_1 N C_2}{n} + \frac{\bar{\lambda} C_2 (\exp(C_1 N) - 1)}{n} + \frac{n-1}{n^2} \log \mathbb{P}^0(\tilde{X}_{t_n} \geq nx), \\ \frac{1}{n} \log(\pi^{PF}(\{nx\}^\uparrow)) + \frac{C_3}{n} \exp(-C_4 n) &\leq O(n^{-1}) + \left(\frac{n-1}{n} \right) \frac{1}{n} \log(\pi^{mPF}(\{nx\}^\uparrow)) + \frac{\tilde{C}_3(n-1)}{n^2} \exp(-\tilde{C}_4 n) \end{aligned}$$

So we have

$$\frac{1}{n} \log(\pi^{PF}(\{nx\}^\uparrow)) \leq \frac{1}{n} \log(\pi^{mPF}(\{nx\}^\uparrow)) + O(n^{-1}).$$

We also have the converse of this last inequality using the same arguments interchanging X and \tilde{X} .

Assume now that $x_i > 0$ for $i \in \mathcal{U}$, while $x_i = 0$ for $i \in \mathcal{S}$, where $\mathcal{U}, \mathcal{S} \subseteq \{1, \dots, N\}$. If classes of \mathcal{U} are independent of classes of \mathcal{S} , the result is obvious from the case where x is strictly positive. Hence suppose that there exists some dependence between classes of \mathcal{U} and \mathcal{S} .

We will use the following fact:

Lemma 9. *Given the definition of Proportional fairness, if class $i \in \mathcal{S}$ and class $j \in \mathcal{U}$ are not independent, then there exist $K > 0$ such that if $y_i \leq n\epsilon$ and $y_j \geq n$ for all $j \in \mathcal{U}$, then:*

$$|\phi_i(y) - \tilde{\phi}_i(y)| \leq K\epsilon.$$

Proof.

If at least two classes are not independent, there exists a Lagrange multiplier $\alpha > 0$ and some positive constants $c_i > 0$ and c_j (with at least one j with $c_j > 0$) such that

$$\frac{\partial U^1(y, \eta_i)}{\partial \eta_i} = \frac{x_i}{\eta_i} = \alpha c_i,$$

$$\frac{\partial U^1(y, \eta_j)}{\partial \eta_j} = \frac{x_j}{\eta_j} = \alpha c_j.$$

Combined with the fact that there exists $c > 0$ such that $0 < c \leq \sum_{j=1}^N \eta_j \leq C$, we obtain that

$$\tilde{c} \frac{y_i}{y_i + \sum_{j \in \mathcal{U}_i} y_j} \leq \phi_i(y) \leq \tilde{C} \frac{y_i}{y_i + \sum_{j \in \mathcal{U}_i} y_j},$$

where $\mathcal{U}_i \subset \mathcal{U}$. Hence if $y_i \leq n\epsilon$ and $y_j \geq n$ for all $j \in \mathcal{U}$

$$\phi_i(y) \leq \tilde{C} \frac{\epsilon n}{1+n} \leq \tilde{K}\epsilon.$$

Using the control on the modified proportional fair allocation for $x_i \geq 1$:

$$\tilde{\phi}_i(y) \leq \exp\left(\frac{1}{y_i}\right) \phi_i(y) \leq e\phi_i(y) \leq e\tilde{K}\epsilon.$$

□

Now we consider a finer classification of indexes in \mathcal{S} . Let $\{\epsilon_n\}$ be a sequence such that ϵ_n goes to 0 as n grows, and $\{t_n = C_2 n\}$ where $C_2 > 0$ is chosen to meet conditions of Lemma 8. For each n , define the sets

$$\bar{\mathcal{S}}_n = \{i \in \mathcal{S} : X_i(t_n) \geq n\epsilon_n\}, \quad \underline{\mathcal{S}}_n = \{i \in \mathcal{S} : X_i(t_n) < n\epsilon_n\}.$$

We look at the set of events:

$$\mathcal{A} = \{X_i(t_n) \geq nx_i, i \in \mathcal{U}_R, X_i(t_n) \geq n\epsilon_n, i \in \bar{\mathcal{S}}_n, X_i(t_n) \leq n\epsilon_n, i \in \underline{\mathcal{S}}_n\}.$$

Recall that

$$\left| \log \left(\frac{\phi_i(y)}{\tilde{\phi}_i(y)} \right) \right| \leq 1,$$

and remark that on $A \in \mathcal{A}$, using the previous lemma, the process counting the number of downwards jumps in direction $j \in \mathcal{S}$ which are not common for X and \tilde{X} is dominated by a Poisson process N^{ϵ_n} of intensity $C_5 \epsilon_n$ independent of N^j , $j \in \mathcal{S}$. Hence:

$$\begin{aligned} & E^0 1_{\tau_k \leq t < \tau_{k+1}} 1_A M_{\tau_k, t} \\ & \leq E^0 \left(\exp \left(\sum_{j \in \mathcal{U}} \frac{C_1}{n} (N_t^j + t) + \sum_{j \in \bar{\mathcal{S}}_n} \frac{C_1}{n\epsilon_n} (N_t^j + t) + \sum_{j \in \underline{\mathcal{S}}_n} [C_5 N_t^{\epsilon_n} + (C_6 \epsilon_n) t] \right) 1_{\tau_k \leq t < \tau_{k+1}} 1_A \right). \end{aligned}$$

and,

$$\begin{aligned}
E^0(M_t 1_A) &\leq E^0\left(\exp\left(\sum_{j \in \mathcal{U}} \frac{C_1}{n}(N_t^j + t) + \sum_{j \in \bar{\mathcal{S}}_n} \frac{C_1}{n\epsilon_n}(N_t^j + t) + \sum_{j \in \mathcal{S}_n} [C_5 N_t^{\epsilon_n} + (C_6 \epsilon_n)t]\right) 1_A\right) \\
&\leq E^0\left(\exp\left(\frac{|\mathcal{U}|C_1}{n}(\bar{N}_t + t) + \frac{|\bar{\mathcal{S}}_n|C_1}{n\epsilon_n}(\bar{N}_t + t) + |\mathcal{S}_n|[C_5 N_t^{\epsilon_n} + C_6 \epsilon_n t]\right) 1_A\right) \\
&\leq E^0\left(\exp\left\{\frac{C_1 p N}{n}(\bar{N}_t + t)\right\}\right)^{\frac{|\mathcal{U}|}{pN}} E^0\left(\exp\left\{\frac{C_1 p N}{n\epsilon_n}(\bar{N}_t + t)\right\}\right)^{\frac{|\bar{\mathcal{S}}_n|}{pN}} \\
&\quad E^0\left(\exp\{pN[C_5 N_t^{\epsilon_n} + C_6 \epsilon_n t]\}\right)^{\frac{|\mathcal{S}_n|}{pN}} \mathbb{P}^0(\tilde{X}_t \geq nx)^{\frac{p-1}{p}}, \tag{4.5}
\end{aligned}$$

where we used Hölder's inequality for $1 = \frac{|\mathcal{U}|}{pN} + \frac{|\bar{\mathcal{S}}_n|}{pN} + \frac{|\mathcal{S}_n|}{pN} + \frac{p-1}{p}$.

We now bound the speed of convergence, in the large deviations scale, of each of the three first factors on the right-hand side of inequality (4.5). Choosing $t_n = C_2 n$, $p_n = C_7 \log n$, $\epsilon_n = n^{-\frac{1}{2}}$, we have:

$$\begin{aligned}
\frac{1}{n} \log E^0\left(\exp\left\{\frac{C_1 N p_n}{n}(\bar{N}_{t_n} + t_n)\right\}\right)^{\frac{|\mathcal{U}|}{N p_n}} &= \frac{1}{n} \log\left(\exp\left\{\frac{|\mathcal{U}|C_1 t_n}{n}\right\} \exp\{\bar{\lambda} t_n [e^{\frac{C_1 N p_n}{n}} - 1]\}\right)^{\frac{|\mathcal{U}|}{N p_n}} \\
&= \frac{|\mathcal{U}|C_1 C_2}{n} + \frac{|\mathcal{U}|\bar{\lambda} t_n}{N n p_n} [e^{\frac{C_1 N p_n}{n}} - 1] \\
&= \frac{|\mathcal{U}|C_1 C_2}{n} + \frac{|\mathcal{U}|\bar{\lambda} C_2}{N C_7 \log n} [e^{\frac{C_1 C_7 N \log n}{n}} - 1] \\
&= O(n^{-1}).
\end{aligned}$$

We do similarly for the second factor:

$$\begin{aligned}
\frac{1}{n} \log E^0\left(\exp\left\{\frac{C_1 N p_n}{\epsilon_n n}(\bar{N}_{t_n} + t_n)\right\}\right)^{\frac{|\bar{\mathcal{S}}_n|}{N p_n}} &= \frac{1}{n} \log\left(\exp\left\{\frac{|\bar{\mathcal{S}}_n|C_1 t_n}{\epsilon_n n}\right\} \exp\{\bar{\lambda} t_n [e^{\frac{C_1 N p_n}{\epsilon_n n}} - 1]\}\right)^{\frac{|\bar{\mathcal{S}}_n|}{N p_n}} \\
&= \frac{|\bar{\mathcal{S}}_n|C_1 C_2}{\epsilon_n n} + \frac{|\bar{\mathcal{S}}_n|\bar{\lambda} t_n}{N n p_n} [e^{\frac{C_1 N p_n}{\epsilon_n n}} - 1] \\
&= \frac{|\bar{\mathcal{S}}_n|C_1 C_2}{\epsilon_n n} + \frac{|\bar{\mathcal{S}}_n|\bar{\lambda} C_2}{N C_7 \log n} [e^{\frac{C_1 C_7 N \log n}{\epsilon_n n}} - 1] \\
&= O((n\epsilon_n)^{-1}) = O(n^{-\frac{1}{2}}).
\end{aligned}$$

Finally, for the third factor we have:

$$\begin{aligned}
\frac{1}{n} \log E^0 \left(\exp\{Np_n[C_5Nt_n^{\epsilon_n} + C_6\epsilon_n t_n]\} \right)^{\frac{|\underline{\mathcal{S}}_n|}{Np_n}} &= \frac{1}{n} \log \left(\exp\{C_6|\underline{\mathcal{S}}_n|\epsilon_n t_n\} \exp\{\epsilon_n t_n [e^{C_5Np_n} - 1]\} \right)^{\frac{|\underline{\mathcal{S}}_n|}{Np_n}} \\
&= \frac{C_6\epsilon_n |\underline{\mathcal{S}}_n| t_n}{n} + \frac{|\underline{\mathcal{S}}_n|\epsilon_n t_n}{Np_n n} [e^{C_5Np_n} - 1] \\
&= C_6C_2|\underline{\mathcal{S}}_n|\epsilon_n + \frac{C_2|\underline{\mathcal{S}}_n|\epsilon_n}{C_7N \log n} [e^{C_5C_7N \log n} - 1] \\
&= C_6C_2|\underline{\mathcal{S}}_n|\epsilon_n + \frac{C_2|\underline{\mathcal{S}}_n|\epsilon_n [n^{C_5C_7N} - 1]}{C_7N \log n} \\
&= O(\epsilon_n) + O\left(\frac{\epsilon_n n^\epsilon}{\log n}\right) = O(n^{-\frac{1}{2}+\epsilon}),
\end{aligned}$$

where we chose $C_7 > 0$ small enough such that $C_7C_5N < \epsilon$, for any given $\epsilon > 0$.

Then, along the same lines as in the case where x has all entries positive, it follows

$$\frac{1}{n} \log(\pi^{PF}(\{nx\}^\uparrow)) \leq \frac{1}{n} \log(\pi^{mPF}(\{nx\}^\uparrow)) + O(n^{-\frac{1}{2}+\epsilon}) \quad \forall \epsilon > 0.$$

We also have the converse of last inequality using the same arguments interchanging X and \tilde{X} . \square

4.4 Large deviations for monotone networks with general service time distributions.

In this section, we consider a processor sharing network (i.e. a set of processor sharing nodes) with a Proportionally fair bandwidth allocation and generally distributed flow service times. This means that a flow is served at node i with speed $\frac{\phi_i^{PF}(X_t)}{X_i(t)}$. The dynamics of these networks are described for instance in [BP04, Zac07]. We recall definitions from Section 3.4.4.

We assume in this section that the network is monotone, see Definition 14. This is verified on all tree topologies for instance and for many wireless networks instances. In the proof of the following theorem, we use the monotonicity of the network to obtain stochastic comparisons. Recall Proposition 5 for the stochastic comparison, from 3.20 that we define P^{PF} as $P^{PF}(x) = \max_{\eta \in \log(C)} \langle \eta, x \rangle$, and from Proposition 4 that $\log(\phi_i(x)) = \nabla P^{PF}(x)$.

Now define the discrete potential functions $\bar{\Psi}$ and $\underline{\Psi}$ by the recursive formula:

$$\bar{\Psi}(x) = - \max_{i=1\dots N} (\bar{\Psi}(x - e_i) - \log(\phi_i(x))),$$

$$\underline{\Psi}(x) = - \min_{i=1\dots N} \{ \underline{\Psi}(x - e_i) - \log(\phi_i(x)) \}.$$

We call supPF and infPF the reversible allocations associated with the discrete potentials $\bar{\Psi}$ and $\underline{\Psi}$ and define the rate function $R(x) = \langle \log(\lambda), x \rangle - P(x)$. Let \bar{X} and \underline{X} the process associated with the discrete potentials $\bar{\Psi}$ and $\underline{\Psi}$.

Define further the invariant measures (not necessarily stationary) $\bar{\pi}$ and $\underline{\pi}$ of the processes $\bar{X} \prec X \prec \underline{X}$. Since we have constructed these processes from a discrete potential, we have that (see Proposition 2):

$$\bar{\pi}(x) = \lambda^x \exp(-\bar{\Psi}(x)),$$

$$\underline{\pi}(x) = \lambda^x \exp(-\underline{\Psi}(x)).$$

We can now state the main result of this section.

Theorem 15. *Suppose the network monotone, i.e. ϕ is strongly monotone. If λ is in the interior of the capacity set \mathcal{C} , then the allocations Proportional fairness, supPF and infPF are stable and admit the same large deviation characteristics with rate function R . More precisely:*

$$\left| \frac{1}{n} \log(\underline{\pi}(nx)) - \frac{1}{n} \log(\pi^{PF}(nx)) \right| = O(\log(n)n^{-1}),$$

$$\left| \frac{1}{n} \log(\bar{\pi}(nx)) - \frac{1}{n} \log(\pi^{PF}(nx)) \right| = O(\log(n)n^{-1}).$$

Proof.

Given the definitions of the discrete potentials $\bar{\Psi}$ and $\underline{\Psi}$, observe that there exists two paths $\bar{\mathcal{P}}(x)$ and $\underline{\mathcal{P}}(x)$ from 0 to x such that:

$$\bar{\Psi}(x) = - \sum_{k:(i_k, z_k) \in \bar{\mathcal{P}}(x)} \log(\phi_{i_k}(z_k)),$$

$$\underline{\Psi}(x) = - \sum_{k:(i_k, z_k) \in \underline{\mathcal{P}}(x)} \log(\phi_{i_k}(z_k)),$$

where, with a slight abuse of notations, the indexes $i(z)$ correspond to the indexes defined by the specific paths $\bar{\mathcal{P}}(x)$ and $\underline{\mathcal{P}}(x)$.

Recall now that P is the potential associated with the proportional fair allocation, and since we can define this potential up to an additive constant, let us choose it such

that $P(0) = 0$. So, we can write that for any path $\mathcal{P}(x)$ going from 0 to x , $P(x) = P(x) - P(0) = \sum_{k:(i_k, z_k) \in \mathcal{P}(x)} \mathbf{D}_{i_k} P(z_k)$. Choosing $\mathcal{P}(x) = \bar{\mathcal{P}}(x)$:

$$\begin{aligned} |\bar{\Psi}(nx) - P(nx)| &= \left| \sum_{k:(i_k, z_k) \in \bar{\mathcal{P}}(nx)} \mathbf{D}_{i_k} P(z_k) - \log(\phi_{i_k}(z_k)) \right|, \\ &\leq \sum_{k:(i_k, z_k) \in \bar{\mathcal{P}}(nx)} |\mathbf{D}_{i_k} P(z_k) - \log(\phi_{i_k}(z_k))|. \end{aligned} \quad (4.6)$$

Recalling from (4.4) that for $x_i \geq 1$:

$$|P(x) - P(x - e_i) - \log(\phi_i(x))| \leq \frac{1}{x_i},$$

we can bound (4.6) by:

$$\begin{aligned} |\bar{\Psi}(nx) - P(nx)| &\leq \sum_{(i_k, z_k) \in \bar{\mathcal{P}}(nx)} \frac{1}{z_{i_k}}, \\ &\leq C \sum_{i=1}^n \frac{1}{i} \leq \bar{C} \log(n). \end{aligned} \quad (4.7)$$

Similarly,

$$|\bar{\Psi}(nx) - P(nx)| \leq \underline{C} \log(n). \quad (4.8)$$

Now observe that:

$$\begin{aligned} \frac{1}{n} \log(\bar{\pi}(nx)) &= \langle \log(\lambda), x \rangle - \frac{1}{n} \bar{\Psi}(nx), \\ \frac{1}{n} \log(\underline{\pi}(nx)) &= \langle \log(\lambda), x \rangle - \frac{1}{n} \underline{\Psi}(nx). \end{aligned}$$

Using (4.7) and (4.8) we obtain that

$$\begin{aligned} \frac{1}{n} \log(\bar{\pi}(nx)) &= \frac{1}{n} \langle \log(\lambda), nx \rangle - \frac{1}{n} P(nx) + O(\log(n)n^{-1}), \\ \frac{1}{n} \log(\underline{\pi}(nx)) &= \frac{1}{n} \langle \log(\lambda), nx \rangle - \frac{1}{n} P(nx) + O(\log(n)n^{-1}), \end{aligned}$$

and hence, by the 1-homogeneity of R ,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log(\bar{\pi}(nx)) = \lim_{n \rightarrow \infty} \frac{1}{n} \log(\underline{\pi}(nx)) = R(x).$$

Assume λ is in the interior of the capacity set \mathcal{C} . Using the definition of P , this implies that there exists $a > 0$ such that

$$R(x) = \langle \log(\lambda), x \rangle - P(x) = \langle \log(\lambda) - \log(\phi(x)), x \rangle \leq -a|x|,$$

which implies that the invariant measure $\bar{\pi}$ and $\underline{\pi}$ are summable. This hence proves that the stationary distributions of both processes \bar{X} and \underline{X} are well defined for Markovian dynamics while the reversibility condition (i.e. the balance property of the service rates) implies the insensitivity of the stationary distribution to the service time distribution [BP02, Zac07]. Therefore, for any service time distribution, the networks with allocation infPF and supPF admit the stationary distribution $C_1\bar{\pi}$ and $C_2\underline{\pi}$ where C_1 and C_2 are normalizing constants. By the assumption of monotonicity of the network, we obtain that:

$$P(|\bar{X}| \geq n) \geq P(|X| \geq n) \geq P(|\underline{X}| \geq n), \quad (4.9)$$

and by Cramér's theorem we conclude:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log P(|X| \geq n) = \lim_{n \rightarrow \infty} \frac{1}{n} \log P(|X| = n).$$

It has further been proven in [Mas07] that mPF and BF have $R(x)$ as rate functions. □

We proved that the stationary measure of the number of flows in progress in a bandwidth sharing network functioning under the Proportional fair allocation shares the same large deviations characteristics with the stationary measure of the number of flows in progress of the same network under the Balanced fair allocation. This formalizes the idea that in long excursions Proportional fair allocation behaves similarly to the most efficient insensitive allocation.

4.5 Open problems

4.5.1 Large deviations for discriminatory Processor Sharing

Let us examine the simple context of a node with N routes through it. In this case, Proportional fair allocation reduces to Processor sharing routine, defined in Section 3.2. Namely:

$$\phi_i^{PS}(x) = \frac{x_i}{\sum_{j=1}^N x_j}, \quad \forall i = 1, \dots, N.$$

We can generalize this allocation to the *discriminatory Processor Sharing* allocation (dPS), defined by

$$\phi_i^{dPS}(x) = \frac{x_i}{\sum_{j=1}^N \omega_j x_j}, \quad \forall i = 1, \dots, N.$$

for some fixed weights $0 < \omega_j < \infty$, $i = 1, \dots, N$. The main idea is to keep the sharing mechanism to allocate bandwidth while giving different preference to the routes. We are interested in to calculate the large deviations of the stationary distribution of a network under this sharing routine.

4.5.2 Large deviations approximations by reversible processes

To find the large deviations asymptotics of Proportional fairness allocation in Markovian and stationary regime, our technique consisted in approximating our allocation by a reversible allocation sharing the same behaviour asymptotically. Since reversible process are more manageable in terms of stationary distributions, it was possible to compute explicitly all the tail behaviour. Is that method possible for more general allocations?

Bibliography

- [Asm03] S. Asmussen. *Applied Probability and Queues*. Springer, 2003.
- [BM01] T. Bonald and T. Massoulié. Impact of fairness on internet performance. In *Proceedings of the ACM Sigmetrics*, pages 82–91, 2001.
- [BMPV06] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Syst. Theory Appl.*, 53(1-2):65–84, 2006.
- [Bon07] T. Bonald. Insensitive traffic models for communication networks. *Discrete Event Dynamic Systems*, 17(3):405–421, 2007.
- [BP02] T. Bonald and A. Proutière. Insensitivity in processor-sharing networks. *Performance Evaluation*, 49(14):193 – 209, 2002.
- [BP03] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Syst. Theory Appl.*, 44(1):69–100, 2003.
- [BP04] T. Bonald and A. Proutière. On stochastic bounds for monotonic processor sharing networks. *Queueing Syst. Theory Appl.*, 47(1/2):81–106, 2004.
- [dVLK99] G. de Veciana, T. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting services with rate control - could the internet be unstable? *INFOCOM*, pages 802–810, 1999.
- [Erl09] K. Erlang. Sandsynlighetsregningog telefonsamtaler (in danish). *Nytt tidsskrift for Matematik B*, (20), 1909.
- [Erl18] K. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *The Post Office Electrical Engineers Journal*, (10), 1918.

- [FW84] M. Freidlin and A. Wentzell. *Random perturbations of dynamical systems*. Springer, 1984.
- [Mas07] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.
- [MT93] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
- [Pro03] A. Proutière. *Insensibilité et bornes stochastiques dans les réseaux de file d’attente. Application à la modélisation des réseaux de télécommunications au niveau flot*. PhD Thesis, École Polytechnique, 2003.
- [Sto07] K. Stordahl. The history behind the probability theory and the queuing theory. *Teletronikk*, (2):123–140, 2007.
- [SW95] A. Shwartz and A. Weiss. *Large Deviations for performance analysis. Queues, communication and computing*. Chapman & Hall, 1995.
- [Wal11] N. Walton. Insensitive, maximum stable allocations converge to proportional fairness. *Queueing Systems*, 68(1):51–60, 2011.
- [Zac07] S. Zachary. A note on insensitivity in stochastic networks. *Journal of Applied Probability*, 44(1):238–248, 2007.

Bibliography

- [AD95] D. Aldous and P. Diaconis. Hammersley’s interacting particle process and longest increasing subsequences. *Probab. Th. Rel. Fields*, 103:199–213, 1995.
- [Ana93] V. Anantharam. Uniqueness of stationary ergodic fixed point for a $\cdot/M/K$ node. *Ann. Appl. Probab.*, 3(1):154–172, 1993.
- [Asm03] S. Asmussen. *Applied Probability and Queues*. Springer, 2003.
- [Bil68] P. Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, 1968.
- [BM01] T. Bonald and T. Massoulié. Impact of fairness on internet performance. In *Proceedings of the ACM Sigmetrics*, pages 82–91, 2001.
- [BMPV06] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Syst. Theory Appl.*, 53(1-2):65–84, 2006.
- [Bon07] T. Bonald. Insensitive traffic models for communication networks. *Discrete Event Dynamic Systems*, 17(3):405–421, 2007.
- [BP02] T. Bonald and A. Proutière. Insensitivity in processor-sharing networks. *Performance Evaluation*, 49(14):193 – 209, 2002.
- [BP03] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Syst. Theory Appl.*, 44(1):69–100, 2003.
- [BP04] T. Bonald and A. Proutière. On stochastic bounds for monotonic processor sharing networks. *Queueing Syst. Theory Appl.*, 47(1/2):81–106, 2004.
- [Bur56] P. Burke. The output of a queueing system. *Operations Research*, 4(6):pp. 699–704, 1956.
- [DC99] F. Dufour and O. Costa. Stability of piecewise-deterministic Markov processes. *SIAM J. Control Optim.*, 37(5):1483–1502 (electronic), 1999.

- [DMO03] M. Draief, J. Mairesse, and N. O’Connell. Joint Burke’s theorem and RSK representation for a queue and a store. pages 69–82 (electronic), 2003.
- [dVLK99] G. de Veciana, T. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting services with rate control - could the internet be unstable? *INFOCOM*, pages 802–810, 1999.
- [Erl09] K. Erlang. Sandsynlighetsregning og telefonsamtaler (in danish). *Nytt tidsskrift for Matematik B*, (20), 1909.
- [Erl18] K. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *The Post Office Electrical Engineers Journal*, (10), 1918.
- [Fer96] P. A. Ferrari. Limit theorems for tagged particles. *Markov Process. Related Fields*, 2(1):17–40, 1996.
- [Fer04] P. L. Ferrari. Polynuclear growth on a flat substrate and edge scaling of goe eigenvalues. *Communications in Mathematical Physics*, 252(1-3):77–109, 2004.
- [FF94] P. A. Ferrari and L. R. G. Fontes. The net output process of a system with infinitely many queues. *Ann. Appl. Probab.*, 4(4):1129–1144, 1994.
- [FM06] P. A. Ferrari and J. Martin. Multiclass processes, dual points and m/m/1 queues. *Markov Processes and Related Fields*, 2006.
- [FM07] P. A. Ferrari and J. Martin. Stationary distributions of multi-type totally asymmetric exclusion processes. *The Annals of Probability*, 35(3):pp. 807–832, 2007.
- [FM09] P. A. Ferrari and J. Martin. Multiclass Hammersley-Aldous-Diaconis process and multiclass-customer queues. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(1):250–265, 2009.
- [FW84] M. Freidlin and A. Wentzell. *Random perturbations of dynamical systems*. Springer, 1984.
- [Ham72] J. M. Hammersley. A few seedlings of research. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. I: Theory of statistics*, pages 345–394. Univ. California Press, 1972.

-
- [Har90] M. Harrison. *Brownian Motion and Stochastic Flow Systems*. Krieger Pub Co, 1990.
- [HW90] M. Harrison and R. Williams. On the quasireversibility of a multiclass brownian service station. *The Annals of Probability*, 18(3), 1990.
- [Jon10] M. Jonckheere. *Introduction to Stochastic Networks*. Lecture notes, course given at the Universidad de Buenos Aires, 2010.
- [Kal02] O. Kallenberg. *Foundations of modern probability*. Springer, 2002.
- [Kel79] F. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.
- [KMT98] F. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *The Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [Lig85] T. Liggett. *Interacting Particle Systems*. Springer, 1985.
- [LS77] B.F Logan and L.A Shepp. A variational problem for random young tableaux. *Advances in Mathematics*, 26(2):206 – 222, 1977.
- [LSD10] D. Leith, V. Subramanian, and K. Duffy. Log-convexity of rate region in 802.11e wlans. *IEEE Communications Letters*, 14(1):57–59, 2010.
- [Mar06] J. Martin. Last-passage percolation with general weight distribution. *Markov Process. Related Fields*, 12(2):273–299, 2006.
- [Mas07] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.
- [MP95] T. Mountford and B. Prabhakar. On the weak convergence of departures from an infinite series of $\cdot/M/1$ queues. *Ann. Appl. Probab.*, 5(1):121–127, 1995.
- [MP10] J. Martin and B. Prabhakar. Fixed points for multiclass queues. *ArXiv: 1003.3024v1*, 31, 2010.
- [MR02] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans. Netw.*, 10(3):320–328, 2002.
- [MT93] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
- [MW00] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. on Netw.*, 8(5):556–567, 2000.

- [OY01] N. O’Connell and M. Yor. Brownian analogues of Burke’s theorem. *Stochastic Process. Appl.*, 96(2):285–304, 2001.
- [Pro03] A. Proutière. *Insensibilité et bornes stochastiques dans les réseaux de file d’attente. Application à la modélisation des réseaux de télécommunications au niveau flot*. PhD Thesis, École Polytechnique, 2003.
- [PS00] M. Prähofer and H. Spohn. Statistical self-similarity of one-dimensional growth processes. *Phys. A*, 279(1-4):342–352, 2000. Statistical mechanics: from rigorous results to applications.
- [PS04] M. Prähofer and H. Spohn. Exact scaling functions for one-dimensional stationary kpz growth. *Journal of Statistical Physics*, 115(1-2):255 – 279, 2004.
- [Roc70] T. Rockafellar. *Convex analysis*, volume 28. Princeton landmarks in mathematics, 1970.
- [SN01] P. Salminen and I. Norros. On busy periods of the unbounded brownian storage. *Queueing Systems*, 39:317–333, 2001.
- [Sto07] K. Stordahl. The history behind the probability theory and the queuing theory. *Teletronikk*, (2):123–140, 2007.
- [SW95] A. Shwartz and A. Weiss. *Large Deviations for performance analysis. Queues, communication and computing*. Chapman & Hall, 1995.
- [UK11] M. Uchida and J. Kurose. An information-theoretic characterization of weighted -proportional fairness in network resource allocation. *Information Sciences*, 181(18):4009 – 4023, 2011.
- [VK77] A.M. Vershik and S.V. Kerov. Asymptotics of the plancherel measure of the symmetric group and the limiting form of young tables. *Soviet Math. Dokl.*, 18:527–531, 1977.
- [Wal11] N. Walton. Insensitive, maximum stable allocations converge to proportional fairness. *Queueing Systems*, 68(1):51–60, 2011.
- [Whi02] W. Whitt. *Stochastic Process Limits*. Springer, 2002.
- [Zac07] S. Zachary. A note on insensitivity in stochastic networks. *Journal of Applied Probability*, 44(1):238–248, 2007.