

A Nonequilibrium Dynamical Transition in the Asymmetric Exclusion Process

Arno Proeme



Doctor of Philosophy
The University of Edinburgh
2011

Abstract

Over the last few decades the interests of statistical physicists have broadened to include the detailed quantitative study of many systems – chemical, biological and even social – that were not traditionally part of the discipline. These systems can feature rich and complex spatiotemporal behaviour, often due to continued interaction with the environment and characterised by the dissipation of flows of energy and/or mass. This has led to vigorous research aimed at extending the established theoretical framework and adapting analytical methods that originate in the study of systems at thermodynamic equilibrium to deal with out-of-equilibrium situations, which are much more prevalent in nature.

This thesis focuses on a microscopic model known as the asymmetric exclusion process, or ASEP, which describes the stochastic motion of particles on a one-dimensional lattice. Though in the first instance a model of a lattice gas, it is sufficiently general to have served as the basis to model a wide variety of phenomena. That, as well as substantial progress made in analysing its stationary behaviour, including the locations and nature of phase transitions, have led to it becoming a paradigmatic model of an exactly solvable nonequilibrium system. Recently an exact solution for the dynamics found a somewhat enigmatic transition, which has not been well understood. This thesis is an attempt to verify and better understand the nature of that dynamical transition, including its relation, if any, to the static phase transitions.

I begin in Chapter 2 by reviewing known results for the ASEP, in particular the totally asymmetric variant (TASEP), driven at the boundaries. I present the exact dynamical transition as it was first derived, and a reduced description of the dynamics known as domain wall theory (DWT), which locates the transition at a different place. In Chapter 3, I investigate solutions of a nonlinear PDE that constitutes a mean-field, continuum approximation of the ASEP, namely the Burgers equation, and find that a similar dynamical transition occurs there at the same place as predicted by DWT but in disagreement with the exact result. Next, in Chapter 4 I report on efforts to observe and measure the dynamical transition through Monte Carlo simulation. No directly obvious physical manifestation of the transition was observed. The relaxation of three different observables was measured and found to agree well with each other but only slightly better with the exact transition than with DWT. In Chapter 5 I apply a numerical renormalisation scheme known as the Density Matrix Renormalisation Group (DMRG) method and find that it confirms the exact dynamical transition, ruling out the behaviour predicted by DWT. Finally in Chapter 6 I demonstrate that a perturbative calculation, involving the crossing of eigenvalues, allows us to rederive the location of the dynamical transition found exactly, thereby offering some insight into the nature of the transition.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified. The original work described in chapters 3-5 has been published in [1].

(Arno Proeme)

Acknowledgements

I first of all want to thank both my supervisors, Martin Evans and Richard Blythe, for their patience, encouragement, and support - in the end it made all the difference.

I am thankful to Jane Patterson for always looking out for us postgrads, from start to finish.

I would like to thank everyone in Rudsambee for helping make wednesday night the energising recharge night that it has become, it has helped me stay sane and prevented various kinds of atrophy. Thanks especially to John and Susan for all their hospitality, and to Rufus and Rosie for comments.

I thank the band Meshuggah for their phenomenal rhythmic precision to which I attribute increased productivity (and Tom for introducing me to them), and J.S. Bach for a reliable route to instant happiness.

I would like to thank some of those who have inspired me: Henrik Jensen for an infectious “... *state of wonder*” and a beautifully open mind, Chris Isham for conceptual clarity, for showing us that formalism is formalism, and for a rare and deepseated intellectual honesty, and Kim Christensen for his passionate rigour and for wiping away a little tear when the percolating cluster extends to infinity.

Further back I would like to thank Dick Hoekzema for planting some seeds, and Alexander Morozov and Wim van Saarloos for an early taste of nonequilibrium.

Anke & Ian, I will always be grateful for your kindness on so many occasions when I first came to these shores.

M'n lieve ouders, Rob en Olga, jullie wil ik bedanken voor de voortdurende steun, en voor alle kansen en ervaringen die ik dankzij jullie heb gehad, zonder jullie had ik dit natuurlijk nooit kunnen behalen...

And to my Quantum Transition π : glasses touching, you know how I feel...

Contents

Abstract	i
Declaration	iii
Acknowledgements	v
Contents	vii
1 Introduction	1
2 The Asymmetric Exclusion Process	7
2.1 Exclusion processes: history, genealogy, and applications	8
2.2 The open-boundary TASEP: model definition	10
2.3 Mathematical formulation	12
2.3.1 The master equation	13
2.3.2 Spectrum, relaxation times, and the gap	15
2.4 Mean-field analysis of the steady state	16
2.5 Mean-field dynamics: the Burgers equation	18
2.6 Exact solutions for the steady state	20
2.7 Preliminary evidence for dynamical scaling	22
2.8 Exact dynamics	27
2.8.1 Thermodynamic gap and dynamical phase diagram	28
2.9 Alternative dynamical description: domain wall theory	31
2.10 This thesis: the dynamical transition	33
3 Gap from Burgers equation	35
3.1 The Cole-Hopf transformation, diffusive modes, and the Burgers density	36
3.2 Thermodynamic diffusive spectrum and the Burgers gap	39
3.3 The stationary density profile	45
3.4 Relaxation of the longest-lived transient: soliton motion from nonlinear superposition of diffusive modes	48
3.5 Summary	50
4 Gap from Monte Carlo simulation	53
4.1 Monte Carlo simulation	54
4.1.1 Discrete-time Monte Carlo: random sequential updating	57
4.1.2 Continuous-time Monte Carlo: Gillespie's direct method	58
4.2 Measuring the gap	61
4.2.1 Relaxation of an observable	61

4.2.2	Fitting a sum of exponentials	63
4.2.3	Fitting a truncated sum of exponentials	64
4.2.4	Asymptotic fitting	67
4.3	Gap from relaxation of an initial overall density excess or deficit	72
4.4	Gap from occupancy fluctuations in the steady state	81
4.5	Gap from relaxation of a second class particle	82
4.6	Discussion	85
4.6.1	Approximate theoretical second class particle gap	86
5	Gap from DMRG	89
5.1	The density matrix renormalization group method	90
5.2	DMRG for the TASEP	94
5.2.1	Infinite system method DMRG for the TASEP	94
5.2.2	An infinite system method algorithm for the TASEP gap	97
5.3	Infinite system method gap results	100
5.4	Discussion	106
6	The dynamical transition as an eigenvalue crossing	109
6.1	An eigenvalue crossing scenario	110
6.2	A perturbative approach to the crossing	112
6.3	Identifying gap and second gap eigenvalues	115
6.3.1	Degeneracies of the unperturbed eigenvalues	117
6.4	Perturbative expressions for the gap and second gap	119
6.4.1	Perturbation theory for normal and non-normal matrices	119
6.4.2	Jordan normal form, degeneracy structure, and Lidskii perturbation	121
6.4.3	Eigenvalue perturbation by direct analysis of the characteristic polynomial	123
6.5	Rederiving the dynamical transition	127
6.6	Reinterpreting the dynamical transition	131
7	Conclusions	133
7.1	Summary of findings	133
7.2	The dynamical transition: interpretation and suggested investigations	135
	Bibliography	141

Chapter 1

Introduction

The last three or four decades have borne witness to a large number of important developments in our understanding of and control over fundamental physical phenomena. These include: the formation of novel states of matter such as Bose-Einstein condensates, superfluid, and superconducting states, the emergence of quantum computation, controlled quantum-mechanical entanglement, the dramatic slowing and trapping of light, the manipulation of matter on a nanoscopic scale, controlled plasma fusion, the controlled production of antimatter, the development of string theory, efforts at a unified description of quantum gravity, and the study of the cosmic microwave background, to name just a few.

A no less exciting feature of physics research over the same period has been the increasingly widespread study of less exotic systems which, though quantum mechanical effects play no significant role in their dynamics, nonetheless exhibit surprising and richly varied behaviour, even from very simple underlying interactions. This includes such traditional hydrodynamical systems as thermally-driven Rayleigh-Benard convection and mechanically-driven Taylor-Couette flow, which are studied as part of the ongoing investigation of turbulence [2]. However I wish to draw special attention to studies of systems that, fifty years ago, would not have been the subject of much attention. One useful way of categorising examples of these types of systems is according to whether they have an (approximately) continuous medium as substrate, or whether they consist of many distinct interacting constituents, which may be free to move in space or arranged in arbitrary spatial structures including networks, regular lattices, etc. In the first category comes the study of the dynamics of pattern formation in systems such as sand dunes, snow flakes, slime molds, persistent atmospheric phenomena like the Great Red Spot on Jupiter,

and coat patterning in animals [2–4]. In the second category we find, for example, the Ising model of ferromagnetism [5], related spin glass models including neural networks [6], models of automobile, insect and internet traffic, the swarm behaviour of flocks of birds [7, 8], and autocatalytic sets of chemical reactions [9].

Systems in both categories are characterised by some combination of the following: interactions between many individual constituents, leading to collective dynamics and emergent behaviour such as phase transitions; the continued influence of the environment, which may itself be homogeneous and unchanging or heterogeneous and time-dependent; and by causal feedback loops internal to the system or between the system and the environment, both of which lead to nonlinear effects such as amplification and spatiotemporal chaos. The systems may be deterministic or stochastic, and may exhibit qualitatively different features in either case.

Consider systems of the second kind, that is, systems consisting of many interacting constituents. This situation is of course familiar from the classical statistical mechanics theory of gases, wherein one foregoes determining the behaviour of the system by solving the laws of motion – an intractable problem for the average number of particles contained in a macroscopic volume of gas – and instead contents oneself with a probabilistic description of the stationary state of the system, subject to the condition of equilibrium, as follows. For any system allowed to exchange energy with an environment that is itself strictly homogeneous and time-independent, i.e. for any system allowed to equilibrate with such an environment, statistical mechanics [10] tells us that after a sufficiently long time the probability $P(\mathcal{C})$ of a state with energy E is given by the Boltzmann distribution:

$$P(\mathcal{C}) = \frac{e^{-E/k_{\text{B}}T}}{\mathcal{Z}} . \quad (1.1)$$

The partition function \mathcal{Z} not only provides a normalisation for the stationary weights:

$$\mathcal{Z} = \sum_{\mathcal{C}_i \in \mathcal{S}} P(\mathcal{C}_i) , \quad (1.2)$$

with \mathcal{S} the entire state space of the system, but once computed also allows all thermodynamic variables, such as the average energy, to be calculated. There is however no such general prescription currently in existence that gives the analogue of the Boltzmann distribution for an arbitrary *nonequilibrium* system. This is not surprising, since the statement that a system is out of equilibrium can mean any of the following (at least)¹:

¹For an excellent account of the development of equilibrium statistical mechanics and a thorough examination of its conceptual and historical foundations as well as of some of the goals and obstacles

-
- (i) The system is in some sense close to equilibrium, and is in the process of relaxing to the equilibrium state.
 - (ii) The system is in some sense far from equilibrium, and is in the process of relaxing to the equilibrium state.
 - (iii) The system is in some sense close to equilibrium, and is held there by a persistent driving force. This corresponds to a nonequilibrium steady state.
 - (iv) The system is in some sense close to the nonequilibrium steady state (iii), and is in the process of relaxing to this state.
 - (v) The system is in some sense far from the nonequilibrium steady state (iii), and is in the process of relaxing to this state.
 - (vi) The system is in some sense far from equilibrium and is held there by a persistent driving force, hence it is in a nonequilibrium steady state distinct from (iii)
 - (vii) The system is in some sense close to the nonequilibrium steady state (vi), and is in the process of relaxing to this state.
 - (viii) The system is in some sense far from the nonequilibrium steady state (vi), and is in the process of relaxing to this state.

Whereas (iii.) and (vi.) concern stationary behaviour, the other categories describe dynamical problems. Furthermore, for each above case of nonequilibrium a further distinction can be made according to whether there exists a spatiotemporal scale sufficiently small that it encompasses a region of the system that behaves locally as if it is nonetheless at equilibrium. The problem studied in this thesis and described below is that of the long-time relaxation to a far out-of-equilibrium steady state, and hence falls under (vii.).

Since there is not one single kind of nonequilibrium it stands to reason that there is not, nor could there ever be a single prescription, probabilistic or otherwise, for out-of-equilibrium stationary or dynamical behaviour without at least some further specification of the model under consideration. Models of nonequilibrium systems therefore have to be dealt with mostly on a case-by-case basis.

In this thesis I study a particular instance drawn from the class of stochastic models referred to collectively as the asymmetric exclusion process, or ASEP. These models describe the driven diffusion on a one-dimensional lattice of particles with short-range

involved in constructing a nonequilibrium formalism, I recommend above all [11].

interactions, held far out of equilibrium by persistent driving forces. The past decade and a half have seen the successful development of a method, known as the matrix product ansatz [12], for determining exactly the stationary distributions for models belonging to this class. Certain variants show rich stationary behaviour and exhibit first and second order phase transitions that manifest themselves as a loss in analyticity of a nonequilibrium analogue of the free energy, defined as the logarithm of the partition function, like one encounters in equilibrium statistical mechanics.

In addition the ASEP is described, in a continuum limit of the mean-field, or more formally in certain hydrodynamic scaling limits, by solutions to the Burgers equation, a nonlinear PDE. Finally, the models can be shown to be formally equivalent to a certain quantum-mechanical spin-chain problem, namely the XXZ spin- $\frac{1}{2}$ Heisenberg chain [13], with appropriate boundary conditions. This connection has been extremely fruitful, as it has allowed exact methods such as the Bethe ansatz to be used to obtain exact characterisations of the dynamical scaling of the various exclusion process variants.

Being an exactly solved nonequilibrium model for which the stationary behaviour is completely characterised has made the ASEP into a prototypical model of a genuinely far-out-of-equilibrium model. As such, any results on its exact properties may prove meaningful for a wider class of similarly nonequilibrium systems. In this thesis I study the variant known as the totally asymmetric exclusion process, or TASEP, with open boundaries. I focus not on the steady state itself, but on the process by which the system relaxes to it, in other words, on the approach to nonequilibrium stationarity. In particular, I investigate a recent exact finding of a kind of dynamical transition in the long-time relaxation behaviour.

I begin in chapter 2 by reviewing known results for the ASEP and in particular for the TASEP. I present the exact dynamical transition as it was first derived, and a reduced description of the dynamics known as domain wall theory (DWT), which locates the transition at a different place. In chapter 3 I investigate solutions of the Burgers equation and find that a similar dynamical transition occurs there at the same place as predicted by DWT but in disagreement with the exact result. Next, in chapter 4 I report on efforts to observe and measure the dynamical transition through Monte Carlo simulation. No directly obvious physical manifestation of the transition was observed. The relaxation of three different observables was measured and found to agree well with each other but only slightly better with the exact transition than with DWT. In chapter 5 I apply a numerical renormalisation scheme known as the Density Matrix Renormalisation Group (DMRG) method and find that it confirms the exact dynamical transition, ruling out the behaviour

predicted by DWT. Finally in chapter 6 I demonstrate that a perturbative calculation, involving the crossing of eigenvalues, allows me to rederive the location of the dynamical transition found exactly, thereby offering some insight into the nature of the transition. I conclude in chapter 7 by gathering together the available evidence and speculating on the nature of the transition.

Chapter 2

The Asymmetric Exclusion Process

In this chapter I review essential background material, which the work presented in rest of this thesis builds on and which I refer back to at various points. Extensive literature including review articles covering much of the material exist [7, 12, 14–17], so to avoid duplication of efforts I summarise the salient points and refer to the relevant literature for full details.

First I provide a general introduction to exclusion processes (section 2.1), and give a precise definition of the variant studied here, namely the totally asymmetric exclusion process (TASEP) with open boundaries evolving in continuous time (section 2.2). I present a master equation formulation of the problem (section 2.3) before discussing mean-field approximations for the steady state (section 2.4) and for the dynamics (section 2.5). I then highlight the main features of the exact solution for the steady state, including the location and nature of phase transitions (section 2.6). What follows in section 2.7 is a consideration of the numerical and analytical evidence for the dynamical scaling behaviour of the TASEP that was available *before* recent exact dynamical results [18]. As part of this and in order to ground the discussion, in this chapter and others, of spectral properties of the TASEP transition matrix, I include in section 2.7 a taster of the behaviour of finite-size spectra. In section 2.8 I summarise the recent exact dynamical results [18] by de Gier and Essler, which not only confirm the existence and location of boundary-driven crossovers in the dynamical scaling behaviour of the model, but also find a more subtle distinction in the long-time dynamics. This distinction constitutes a dynamical transition, and is the topic of this thesis. In section 2.9 I review domain

wall theory, which provides an alternative and remarkably successful characterisation of the dynamics. Finally, in section 2.10 I summarise the questions this thesis attempts to answer, thereby motivating the work presented in succeeding chapters.

2.1 Exclusion processes: history, genealogy, and applications




Figure 2.1: The asymmetric exclusion process on a one-dimensional lattice of indefinite size and with indefinite boundary behaviour. In discrete time the parameters p and q are hopping probabilities. In continuous time they are Poisson rates, i.e. expected inverse waiting times, for the same stochastic particle motion (see section 2.2 for more details). Hops onto occupied sites are not allowed.

From a physicist's point of view the asymmetric exclusion process, or ASEP, describes the stochastic biased diffusion, typically on a one-dimensional lattice, of particles obeying the simple exclusion interaction (figure 2.1). The sole effect of simple exclusion is to rule out the occupation of any lattice site by more than one particle. In contrast, relaxing this restriction to allow an arbitrary integer number of particles to occupy any site leads to the zero-range process, or ZRP [19] which, perhaps surprisingly, is isomorphic¹ to the ASEP, though only when both systems are subject to periodic boundary conditions. Asymmetric exclusion processes are also referred to as driven lattice gas models [20], for obvious reasons. The exclusion interaction impedes particle flow, since particles are unable to hop over or, in one dimension, otherwise bypass one another (e.g. by quantum mechanical tunneling), hence the diffusion is not only biased but also congested.

Though they have come to play an important role in theoretical statistical physics, the first occurrence of an exclusion process in the literature is thought to be a model for the dynamics of ribosomes [21]. Shortly thereafter and independently, they began to be studied much more formally as problems in applied probability, primarily by Spitzer and by Liggett [22, 23], resulting in [24] and [25]. Extensive results were obtained by Lavis for the special case of equilibrium lattice gases, i.e. for symmetrically driven lattice gases whose particles experience unbiased diffusion, and which are in contact with a

¹by equating each particle in the ZRP to an empty site in the ASEP, and each site in the ZRP (regardless of whether it is empty or occupied) to an occupied site in the ASEP.

homogeneous environment [26, 27]. However the last fifteen years of the last century, especially after [28], saw an explosion of interest amongst statistical physicists in solving the asymmetrically driven cases which, despite being held persistently out of equilibrium by an ever-present driving force on the particles, exhibit well-defined stationary behaviour, i.e. they possess nonequilibrium steady states. Hence currents of particles flow in the steady state, and if the system is in contact with an inhomogeneous environment there is persistent transport of mass through the system. The subsequent success achieved in characterising the behaviour of these models has led to the asymmetric exclusion process becoming paradigmatic in the development of exact, approximate, and numerical techniques that are beginning to constitute a statistical mechanical toolkit for analysing models of nonequilibrium systems.

The probability theory community on the one hand and the statistical physics community on the other have enthusiastically taken up the pursuit of exact results for the stationary and dynamical properties of these models. Their approaches can roughly be categorised as follows. Probability theorists have focused *a priori* on the setting of infinite lattices (processes on \mathbb{Z}) to derive invariant measures as well as inequalities and bounds for certain fluctuation statistics, sometimes in connection with the precise determination of formal hydrodynamical scaling limits, and they have proven statistical connections to distributions resulting from random matrix theory [29–38]. Statistical physicists on the other hand have usually considered the behaviour of observables such as the density profile and current in finite systems with periodic or open boundary conditions [7, 12, 14–17]. The work in this thesis belongs to the latter approach, wherein, notwithstanding the above comment on the distinction between the study of finite versus infinite lattices, the thermodynamic limit continues to play its usual role, familiar from equilibrium statistical mechanics, in analysis of the phase transitions and the dynamical scaling behaviour.

As well as attracting fundamental theoretical interest asymmetric exclusion processes have, due to their ready encapsulation of the essence of many naturally occurring systems, been used to model a wide variety of phenomena ranging from social behaviour such as traffic flow [8, 39, 40] to various kinds of biological systems such as molecular motors [41, 42] and the growth of a fungal filament [43], as well as the flow of hydrocarbons through a zeolite pore [44], and in queueing theory [45, 46].

To specify a particular exclusion process we begin by considering whether the diffusion has a directional bias. With the parameters p and q as shown in figure 2.1 the case of unbiased diffusion ($p = q$), is referred to as the *symmetric* simple exclusion process, or SSEP, general nonzero bias ($p \neq q$) as the *partially asymmetric* exclusion process, or




Figure 2.2: The ASEP on a finite lattice of indefinite size and with open boundaries. The open boundaries allow particle inflow with Poisson rates α and δ at the left and right boundaries respectively, and likewise particle outflow with rates β and γ .

PASEP, infinitesimal bias ($p = q \pm \delta/L$, $\delta \ll 1$) as the *weakly* asymmetric exclusion process, or WASEP, and diffusion strictly in the one direction, usually to the right, ($p = 1$, $q = 0$) the *totally* asymmetric exclusion process, TASEP.

We are free to choose continuous or discrete time dynamics, of which the latter is more general [47]. In either case the dynamics consist of a specification of the allowed stochastic motion in terms of local, i.e. sitewise update rules, and of the corresponding probabilities (discrete time) or rates (continuous time).

Finally, the boundary conditions play a key role in the definition of the model. The case of periodic boundary conditions is generally solved more easily, since there one has conservation of total particle number. The open boundary problem, where the total particle number fluctuates under the influence of variable particle inflow and outflow rates at the boundaries, is more complicated and exhibits richer behaviour including first and second order phase transitions.

2.2 The open-boundary TASEP: model definition

In this section I give a precise definition of the model on which all the work in this thesis is based, namely the continuous-time TASEP with open boundaries. Before focusing on the totally asymmetric case we consider the general ASEP of figure 2.1, but now on a finite lattice and with coupling to the open boundaries that allows particle inflow and outflow on both sides of the system (figure 2.2).

The ASEP is a stochastic process. Restricting the discussion to continuous time, this means allowed moves (particle hops) happen in random order, and at random times. In particular, the ASEP is described by the simultaneous and joint evolution of interacting Markov processes, in the sense set out by Spitzer in [22]. To be precise, assume that at

time t_0 the configuration of the system becomes such that it allows a particular unique move to take place, whereas before t_0 this was not possible. For example, the hopping of a particle from site i to site $i + 1$ (for a particular value of i) may not be possible until t_0 , if before that time there was either no particle on site i , or site $i + 1$ was occupied. The occurrence of each possible unique move that is allowed, given the state of the system, is then a Poisson process. Hence the waiting time τ for any such move to occur is exponentially distributed:

$$P(\tau) d\tau = r e^{-r\tau} d\tau, \quad (2.1)$$

with rate r equal to the inverse of the expected waiting time and dependent on the type of move, as detailed in figure 2.2. Equation (2.1) holds for all times $t_0 < t < t_1$, where the unique move in question becomes disallowed at time t_1 due to a change in the configuration of the system. The probability $P(\tau \leq \delta t)$ that the move occurs during an infinitesimal time interval δt located anywhere between t_0 and t_1 is then

$$P(\tau \leq \delta t) = r \delta t + \mathcal{O}(\delta t^2). \quad (2.2)$$

The rate corresponds to a temporal probability density (probability per unit time), since (2.2) implies

$$r = \lim_{\delta t \rightarrow 0} \frac{P(\tau \leq \delta t)}{\delta t}. \quad (2.3)$$

Hence the probability, for example for a particle in the bulk to hop to an unoccupied site on its right during δt , which as is shown in figure 2.1 and figure 2.2 occurs with Poisson rate $r = p$, is given in the limit $\delta t \rightarrow 0$ by:

$$P\left(\bullet \text{ --- } \xrightarrow{\delta t} \text{ --- } \bullet\right) = p \delta t. \quad (2.4)$$

The TASEP is obtained from the ASEP by choosing $p = 1$ and $q = 0$, which corresponds to an effective driving field with totally asymmetric rightward bias in the bulk. Additionally, for the TASEP the boundaries are also driven totally asymmetrically, so we set $\delta = \gamma = 0$. As a result at the left boundary only particle influx takes place, with rate α , and at the right boundary only particle outflux takes place, with rate β , as shown in figure 2.3. The open-boundary TASEP can therefore be thought of as connected at the boundaries to infinite particle reservoirs, with density α on the left, and density $1 - \beta$ on the right. When $\alpha \neq 1 - \beta$ there is a persistent gradient across the system in the chemical potential. The behaviour of the TASEP thus results from an interplay between driven diffusion in the bulk, which transports rightwards any particles in the system regardless of the behaviour at the boundaries, and coupling to variable particle supply at the left




Figure 2.3: The TASEP on a finite lattice of size L with open boundaries, which allow particle inflow only on the left, and outflow only on the right side of the system.

boundary and variable particle removal capacity at the right.

Before proceeding to a thorough formal characterisation of the model we can anticipate in intuitive terms some of the expected steady state behaviour, as follows. The left boundary attempts to impose density α on the bulk, and the right boundary attempts to impose density $1 - \beta$. When $\alpha > \beta$ there is ample supply of particles to the bulk of the system hence the limiting factor, i.e. bottleneck, that determines the steady state behaviour is the rate β at which those particles can be extracted, so that the steady state density is controlled by the right boundary and equal to $1 - \beta$. Conversely when $\beta > \alpha$ the limiting factor is particle supply, which depends on the injection rate α , so that the steady state density is regulated by the left boundary and equal to α .

Additionally, one expects that the transport capacity of the bulk also plays a role. The throughput of particles, i.e. the current, is limited and is highest if on average every other site is occupied, i.e. when the steady state density in the bulk is $\frac{1}{2}$, regardless of whether that density is being imposed by the left boundary or the right boundary. In fact when either boundary attempts to impose a steady state density higher than $\frac{1}{2}$, one expects the system to stop responding in the same way and to enter a new regime of behaviour.

As we shall see in what follows below, the two arguments of the preceding paragraphs offer an intuitive explanation for the existence of three distinct stationary regimes, and for the two types of phase transitions that separate them.

2.3 Mathematical formulation

At any given time the state, or configuration \mathcal{C} of the TASEP on a lattice of size L is uniquely and completely specified by the set $\{s_i\}_{i=1}^L$ of binary site occupation variables,

which obey

$$s_i = \begin{cases} 0 & \text{if site } i \text{ is empty,} \\ 1 & \text{if site } i \text{ is occupied by a particle.} \end{cases} \quad (2.5)$$

2.3.1 The master equation

Time evolution consists of stochastic transitions between states by means of the allowed particle moves described in section 2.2. The natural state descriptor for time evolution is therefore the discrete probability distribution $P_{\mathcal{C}}(t)$ over the 2^L -dimensional state space \mathcal{S} consisting of all possible lattice configurations. Because the dynamics are defined by construction to depend only on the state of the system at the current time, the process is Markovian. The temporal evolution of $P_{\mathcal{C}}(t)$ can therefore be written as a master equation:

$$\frac{dP_{\mathcal{C}}(t)}{dt} = \sum_{\mathcal{C}' \in \mathcal{S}} M(\mathcal{C}' \rightarrow \mathcal{C}) P_{\mathcal{C}'}(t), \quad (2.6)$$

where the 2^L -dimensional transition matrix M fully encodes the dynamics of the TASEP. Off-diagonal entries are simply the Poisson rates described in section 2.2 for transitions between distinct states $\mathcal{C}' \neq \mathcal{C}$. Diagonal entries $M(\mathcal{C} \rightarrow \mathcal{C})$ follow by conservation of probability and are given by the negative of the total rate of leaving \mathcal{C} :

$$M(\mathcal{C} \rightarrow \mathcal{C}) = - \sum_{\mathcal{C}' \neq \mathcal{C}} M(\mathcal{C} \rightarrow \mathcal{C}'). \quad (2.7)$$

A concrete example should clarify the structure of the transition matrix. For a system of size $L = 2$, the matrix on the right hand side of (2.8) catalogues the transition rates between the configurations depicted on the left side:

$$\begin{pmatrix} \text{--} \rightarrow \text{--} & \text{--} \bullet \rightarrow \text{--} & \bullet \text{--} \rightarrow \text{--} & \bullet \bullet \rightarrow \text{--} \\ \text{--} \rightarrow \text{--} \bullet & \text{--} \bullet \rightarrow \bullet & \bullet \text{--} \rightarrow \bullet & \bullet \bullet \rightarrow \bullet \\ \text{--} \rightarrow \bullet \text{--} & \text{--} \bullet \rightarrow \bullet \text{--} & \bullet \text{--} \rightarrow \bullet \text{--} & \bullet \bullet \rightarrow \bullet \text{--} \\ \text{--} \rightarrow \bullet \bullet & \text{--} \bullet \rightarrow \bullet \bullet & \bullet \text{--} \rightarrow \bullet \bullet & \bullet \bullet \rightarrow \bullet \bullet \end{pmatrix} = \begin{pmatrix} -\alpha & \beta & 0 & 0 \\ 0 & -(\alpha + \beta) & 1 & 0 \\ \alpha & 0 & -1 & \beta \\ 0 & \alpha & 0 & -\beta \end{pmatrix} \quad (2.8)$$

Operating in the same ordered physical basis as in (2.8), the transition matrix for a

system of general size L can conveniently be constructed according to:

$$M = h_l \otimes \mathcal{I}^{\otimes L-1} + \sum_{k=1}^{L-1} \mathcal{I}^{\otimes k-1} \otimes h_b \otimes \mathcal{I}^{\otimes L-k-1} + \mathcal{I}^{\otimes L-1} \otimes h_r , \quad (2.9)$$

where \mathcal{I} is the 2×2 identity matrix, and we make use of the elementary transition matrices

$$h_l = \begin{pmatrix} -\alpha & 0 \\ \alpha & 0 \end{pmatrix} \quad (2.10)$$

for particle inflow at the left boundary,

$$h_r = \begin{pmatrix} 0 & \beta \\ 0 & -\beta \end{pmatrix} \quad (2.11)$$

for particle outflow at the right boundary, and

$$h_b = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.12)$$

for transitions in the bulk. Introducing a bra-ket notation, we can represent $P_{\mathcal{C}}(t)$ as

$$|P(t)\rangle = \sum_{\mathcal{C} \in \mathcal{S}} P_{\mathcal{C}}(t) |\mathcal{C}\rangle \quad (2.13)$$

and write the master equation (2.6) as

$$\frac{d}{dt} |P(t)\rangle = M |P(t)\rangle . \quad (2.14)$$

Since M has a complete set of (right) eigenvectors $|\psi_i\rangle$ satisfying

$$M |\psi_i\rangle = \lambda_i |\psi_i\rangle , \quad (2.15)$$

$|P(t)\rangle$ can also be expressed as

$$|P(t)\rangle = \sum_i a_i(t) |\psi_i\rangle , \quad (2.16)$$

and there exists some similarity transformation D for the change of basis $\{P_{\mathcal{C}}(t)\} \rightarrow \{a_i(t)\}$. Due to the spatial asymmetry in the dynamics ($p \neq q$, $\delta \neq \alpha$ and $\gamma \neq \beta$), we have that

$$M(\mathcal{C} \rightarrow \mathcal{C}') \neq M(\mathcal{C}' \rightarrow \mathcal{C}) , \quad (2.17)$$

and hence that M is non-Hermitian. This implies the existence of a distinct set of left eigenvectors $\langle \phi_i | \neq | \psi_i \rangle^T$ obeying bi-orthogonality:

$$\langle \phi_i | \psi_j \rangle = \delta_{ij} . \quad (2.18)$$

Substituting (2.16) into the master equation (2.14) and making use of bi-orthogonality, one finds that the $a_i(t)$ must obey

$$a_i(t) = b_i e^{\lambda_i t} , \quad (2.19)$$

where the initial conditions are encoded in the b_i through

$$b_i = \langle \phi_i | P(0) \rangle . \quad (2.20)$$

The solution to the master equation (2.14) can therefore be written formally² as

$$|P(t)\rangle = \sum_i b_i e^{\lambda_i t} | \psi_i \rangle = e^{Mt} |P(0)\rangle . \quad (2.21)$$

2.3.2 Spectrum, relaxation times, and the gap

Because M is a stochastic matrix it is guaranteed by the Perron-Frobenius theorem (see e.g. [48]) to have eigenvalues satisfying

$$\lambda_0 = 0 > \text{Re}(\lambda_1) \geq \text{Re}(\lambda_2) \geq \dots . \quad (2.22)$$

This means that

$$\begin{aligned} |P_\infty\rangle &= \lim_{t \rightarrow \infty} |P(t)\rangle \\ &= \langle \phi_0 | P(0) \rangle | \psi_0 \rangle , \end{aligned} \quad (2.23)$$

i.e. the steady state depends only on the left and right eigenvectors $\langle \phi_0 |$ and $| \psi_0 \rangle$ with eigenvalue $\lambda_0 = 0$.

All remaining eigenvalues reflect the dynamical behaviour of the TASEP as it evolves towards the steady state. The spectrum (2.22) corresponds to a set of relaxation times $\{\tau_i\}_{i=0}^{2^L-1}$ given by

$$\tau_i = -\frac{1}{\text{Re}(\lambda_i)} , \quad (2.24)$$

²using $e^{Mt} \equiv 1 + Mt + \frac{M^2 t^2}{2} + \dots$

of which the longest (apart from the infinite relaxation time τ_0 associated with the stationary state) is τ_1 . The dominant transient as $t \rightarrow \infty$ in this relaxation spectrum is therefore the eigenmode associated with λ_1 .

The quantity known as the **spectral gap**, denoted by ε , plays a key role in this thesis and is defined as

$$\varepsilon \equiv \text{Re}(\lambda_1 - \lambda_0) . \quad (2.25)$$

Since $\lambda_0 = 0$ we have simply that

$$\varepsilon = \text{Re}(\lambda_1) , \quad (2.26)$$

so that the term ‘gap’ (equivalently ‘first gap’) in this thesis always refers to the real part of λ_1 , and is inversely related to the longest finite relaxation time:

$$\varepsilon = -\tau_1^{-1} . \quad (2.27)$$

Likewise the second gap is denoted by ε_2 and given by

$$\varepsilon_2 \equiv \text{Re}(\lambda_2 - \lambda_0) = \text{Re}(\lambda_2) = -\tau_2^{-1} . \quad (2.28)$$

Before discussing the exact solution of the master equation and what it reveals about both the stationary and dynamical behaviour of the TASEP, I show in the next section what insight we may gain from a mean-field analysis.

2.4 Mean-field analysis of the steady state

Mean-field analysis of the TASEP is straightforward, gives immediate insight into the behaviour of the model, and mostly yields qualitatively correct stationary phase behaviour. The analysis is well-documented [7, 12, 49], so I will only sketch the essentials here.

Given the master equation formalism of section 2.3.1, it can be shown that the current between any two sites $i, i + 1$, which in the steady state is the same everywhere, can be written as

$$J = \langle s_i(1 - s_{i+1}) \rangle , \quad (2.29)$$

where $\langle \dots \rangle$ denotes an average in the steady state. Since the model is ergodic [12], such an average is equivalent both to the average over the states of an infinite ensemble of TASEP




Figure 2.4: Mean-field phase diagram of the TASEP showing the low-density (LD), high-density (HD) and maximal-current (MC) phases, and the behaviour of the bulk density and current therein.

systems at a particular late time, and to the time average over an infinite observation time for a single TASEP. Making the mean-field approximation $\langle s_i s_{i+1} \rangle = \langle s_i \rangle \langle s_{i+1} \rangle$, this gives

$$J = \rho_i (1 - \rho_{i+1}), \quad (2.30)$$

where we have used the notation $\rho_i \equiv \langle s_i \rangle$. Equation 2.30 implies the recursive relation for the density profile

$$\rho_{i+1} = 1 - \frac{J}{\rho_i}, \quad (2.31)$$

and has fixed points satisfying

$$\rho_{\pm} = \frac{1}{2} \left(1 \pm \sqrt{1 - 4J} \right). \quad (2.32)$$

Solving (2.31) and (2.32) consistently with the boundary conditions $\rho_1 = \alpha$, $\rho_L = 1 - \beta$ for a system of size L reveals three different stationary phases, shown in figure 2.4. The first two, which confirm the intuitively expected behaviour mentioned in section 2.2, occur for $J < \frac{1}{4}$ and correspond to two distinct solutions of the fixed point equation (2.32):

$$\begin{array}{lll} \rho = \alpha & \text{and} & J = \alpha(1 - \alpha) & \text{in the Low-density (LD) phase,} \\ \rho = 1 - \beta & \text{and} & J = \beta(1 - \beta) & \text{in the High-density (HD) phase.} \end{array}$$

The third occurs for $J = \frac{1}{4}$, when (2.32) has only one solution, namely

$$\rho = \frac{1}{2} \quad \text{and} \quad J = \frac{1}{4} \quad \text{in the Maximal-current (MC) phase.}$$

The fixed point density values are reached only at one side of an infinite system, namely at the left boundary for $\rho = \alpha$, and at the right boundary for $\rho = 1 - \beta$. Nonetheless, because the decay of the density profile starting from the opposite boundary in each case is very rapid, they effectively give the bulk density in the thermodynamic limit. Hence the line separating the LD and HD phases i.e. when $\alpha = \beta < \frac{1}{2}$, is called the coexistence line since it corresponds to a first order phase transition, with the current exhibiting a discontinuity in the first derivative. It can be shown that the corresponding density profile interpolates between a low-density domain and a high-density domain, which are separated by a sharp interface, known as a shock, or a domain wall. Similarly the lines separating the LD and HD phases from the MC phase correspond to second order phase transitions, as for these parameter values the current is discontinuous in the second derivative.

2.5 Mean-field dynamics: the Burgers equation

A mean-field analysis can also tell us about the dynamical behaviour of the TASEP. In chapter 3 I pursue this route, therefore I explicitly show in this section how taking a continuum limit of the mean-field dynamics yields a description of the time-dependent density profile by the noiseless, i.e. deterministic Burgers equation with nonzero viscosity:

$$\partial_t \rho(x, t) = -(1 - 2\rho) \partial_x \rho + \frac{1}{2} \partial_x^2 \rho. \quad (2.33)$$

This description does not reflect the stochasticity of the ASEP because it results from a mean-field approximation, which in some sense averages over the underlying fluctuations.

As in section 2.4 we begin by writing down the current, which since we are dealing with the general time-dependent problem retains its spatial dependence. The average current between site i and site $i + 1$ is given by

$$\langle J_{i,i+1} \rangle = \langle s_i (1 - s_{i+1}) \rangle. \quad (2.34)$$

Making the mean-field approximation $\langle s_i s_{i+1} \rangle = \langle s_i \rangle \langle s_{i+1} \rangle$ gives

$$\langle J_{i,i+1} \rangle = \rho_i (1 - \rho_{i+1}), \quad (2.35)$$

where we have used the same notation $\rho_i \equiv \langle s_i \rangle$ as in section 2.4. To make the continuum approximation we begin by mapping the discrete lattice to real space. Given a constant lattice spacing δ , we have some freedom in how to choose this mapping. An obvious mapping is to associate site i with position x , site $i+1$ with position $x+\delta$, and the current $J_{i,i+1}$ with position $x+\delta/2$. Whilst this is completely valid it turns out a more natural choice, which leads to a simplification of the procedure, is to use the mappings:

$$\begin{aligned}\rho_i &\leftrightarrow \rho(x - \delta/2), \\ \rho_{i+1} &\leftrightarrow \rho(x + \delta/2), \\ \langle J_{i,i+1} \rangle &\leftrightarrow \langle J(x) \rangle.\end{aligned}\tag{2.36}$$

Together with the added notational convention $J(x) \equiv \langle J(x) \rangle$, this gives us the realspace analogue of (2.35), namely

$$J(x) = \rho(x - \delta/2)(1 - \rho(x + \delta/2)).\tag{2.37}$$

Next, we note that because of conservation of mass $J(x)$ must obey the continuity equation:

$$\partial_t \rho(x, t) = -\partial_x J(x).\tag{2.38}$$

Taylor expanding for $\rho(x \pm \delta/2)$ in (2.37) and substituting in the continuity equation we get

$$\partial_t \rho(x, t) = -(1 - 2\rho)\partial_x \rho + \frac{\delta}{2}\partial_x^2 \rho + O(\delta^2).\tag{2.39}$$

Taking $\lim_{\delta \rightarrow 0}$ on the right hand side of (2.39) reduces it to the inviscid Burgers equation $\partial_t \rho(x, t) = -(1 - 2\rho)\partial_x \rho$. To connect to the literature it is worth noting that this is equivalent to first computing $\lim_{\delta \rightarrow 0} J(x)$ to give the well known [28], [12] current-density relation for the ASEP,

$$J(x) = \rho(1 - \rho),\tag{2.40}$$

and then substituting this into the continuity equation. Hence both (2.40) and the inviscid Burgers equation itself should be understood as simultaneously a mean-field and continuum approximation of the ASEP.

The inviscid Burgers equation is an accurate description of the mean-field behaviour only over length scales exceedingly large compared to the lattice spacing. To provide a more accurate description of the mean-field behaviour over finite length scales we can refrain from taking the limit in our expression for $\partial_t \rho$ in (2.39) and retain the leading

correction due to nonzero δ to give

$$\partial_t \rho(x, t) \approx -(1 - 2\rho)\partial_x \rho + \frac{\delta}{2}\partial_x^2 \rho. \quad (2.41)$$

Setting $\delta = 1$ means we arrive at the viscous Burgers equation (2.33). Here, the discrete nature of the underlying process results in the appearance of an effective viscosity, arbitrarily rescaled in magnitude by our choice $\delta = 1$. Setting $\delta = 1$ does not violate the Taylor expansions leading to (2.39), as their validity rests not on δ being small *per se* but rather on $\rho(x)$ varying slowly over length scales δ . This is similar to the assumption of local equilibrium in formal hydrodynamic scaling limit approaches [50–52], namely that the ratio of microscopic and macroscopic lengthscales goes to zero.

2.6 Exact solutions for the steady state

The quest for exact solutions has resulted in a versatile method based on a matrix product representation of the stationary distribution, which was first used in [53] to calculate the exact current and exact density profiles for the open-boundary TASEP for both finite and infinite systems. The approach has been found to be remarkably versatile and has been adapted to solve a large variety of variants of the asymmetric exclusion process and related models. A comprehensive review that includes the calculation of exact quantities for the ASEP can be found in [12], and the method is also discussed in [7]. From the matrix product solution and from a complementary approach taken in [54] by solving a recursion relation introduced in [49], we know that the open-boundary TASEP possesses the exact stationary phase diagram shown in figure 2.5.

Exact density profiles agree qualitatively, though not precisely with the mean-field profiles. The behaviour of the bulk current and density are the same as in the mean field, and so are the location and order of the phase boundaries, with one exception. Namely, within the LD and HD phases a division into subphases I and II (indicated by dotted lines in figure 2.5) is made on the basis of a change in the functional form of the exact density profile, as follows.

In the I subphases the density profile, measured from one of the boundaries, approaches the bulk value exponentially with a lengthscale ξ that is a combination of a lengthscale ξ_α depending on α , and a lengthscale ξ_β depending on β as follows [15, 55]:

$$\xi^{-1} = |\xi_\alpha^{-1} - \xi_\beta^{-1}|, \quad (2.42)$$




Figure 2.5: Exact stationary phase diagram of the TASEP.

LD = LDI \cup LDII = low density phase, with current $J = \alpha(1 - \alpha)$.

HD = HDI \cup HDII = high density phase, with $J = \beta(1 - \beta)$.

MC = maximal current phase, with $J = 1/4$.

CL = coexistence line, with $J = \alpha(1 - \alpha) = \beta(1 - \beta)$.

where the functional form of both ξ_α and ξ_β is given by

$$\xi_\sigma^{-1} = -\ln 4\sigma(1 - \sigma). \quad (2.43)$$

On the line separating the LDI and LDII subphases ξ_β diverges (i.e. $\xi_\beta^{-1} \rightarrow 0$), while on the line separating the HDI and HDII subphases this happens for ξ_α . In the II subphases the density profile approaches the bulk value as an exponential with the remaining finite lengthscale (ξ_α for LDII and ξ_β for HDII), modulated by a power law [12, 54].

Consider for example the behaviour near the right boundary in the low-density phase. In the bulk, the density coincides with that of the left reservoir, $\rho = \alpha$. The mean occupancy of the lattice site positioned a distance j from the right boundary approaches in the thermodynamic limit $L \rightarrow \infty$ the form [12, 54, 56]

$$\rho_{L-j} \sim \begin{cases} \alpha + c_I(\beta) \left(\frac{\alpha(1-\alpha)}{\beta(1-\beta)} \right)^j & \alpha < \beta, \beta < \frac{1}{2} & \text{(LDI)} \\ \alpha + c_{II}(\alpha, \beta) \frac{[4\alpha(1-\alpha)]^j}{j^{3/2}} & \alpha < \beta, \beta > \frac{1}{2} & \text{(LDII)} \end{cases} \quad (2.44)$$

where c_I and c_{II} are functions of the boundary parameters whose precise details are not essential to the present discussion. The main thing to notice is that the lengthscale of the exponential decay from the right boundary increases with β until the point $\beta = \frac{1}{2}$ is reached, above which the decay length is constant, and the exponential is modulated by

a power law. The density profile at the left boundary exhibits the same kind of change in the high-density phase as α is increased through $\frac{1}{2}$ as a consequence of the particle-hole symmetry, $\rho_{i-1}(\alpha, \beta) = 1 - \rho_{L-i}(\beta, \alpha)$, exhibited by the model.

The line dividing regions I and II does not constitute a true phase transition, since the current is analytic within the LD and HD phases. Moreover, the distance from the boundary over which the density profile in I differs from that in II by a given amount vanishes as the system size $L \rightarrow \infty$. Hence in the thermodynamic limit the difference between the two subphases I and II effectively becomes irrelevant, whereas a true phase transition in contrast is only then realised.

Though I do not consider them here, renormalisation approaches have also proven successful in determining the locations of the static phase transitions [57].

2.7 Preliminary evidence for dynamical scaling

As we saw in section 2.3.2, the spectrum of the transition matrix directly corresponds to a hierarchy of relaxation times. The longest relaxation time τ_1 sets a characteristic time scale for the asymptotic relaxation to the steady state, as well as for the first-order or dominant fluctuations around the steady state. Relaxation times generally depend on system size, and usually grow as the system gets larger. Particular kinds of scaling of the longest relaxation time with system size have been found to coincide with similar underlying dynamical mechanisms independently of the model under consideration. This has led to the identification of certain dynamical universality classes [58]. We can characterise the dynamical scaling by a dynamical exponent, z , for the leading order dependence of τ_1 on L through

$$\tau_1 \propto L^z, \tag{2.45}$$

and hence, since as we saw in section 2.3.2 the gap ε is related to the longest relaxation time by $\tau_1 = -\varepsilon^{-1}$, through

$$\varepsilon \propto L^{-z}. \tag{2.46}$$

Thus universality classes are identified by the value of the dynamical exponent, for example diffusive behaviour is associated with $z = 2$. Of course this scheme assumes that the gap ultimately vanishes, i.e. that the longest relaxation time diverges in the thermodynamic limit. If on the other hand the gap approaches a constant value, we say $z = 0$. Even though the dynamical scaling is characterised only in the thermodynamic limit, we expect to gain some insight into the behaviour of the open-boundary TASEP

by examining the spectral structure of the transition matrix given in equation (2.9) for finite systems.

For small system sizes the exact spectra can readily be computed numerically using existing Linear algebra libraries such as *ARPACK* [59], which is implemented in *Mathematica* as well as in the statistical analysis software *R*. Whereas e.g. *Maple* currently (v.12) uses non-specialised *NAG* routines for diagonalisation of nonsymmetric matrices, *ARPACK* provides the Arnoldi algorithm which is particularly well suited to diagonalisation with the aim of obtaining only one or a few of the eigenvalues at the extremes of the spectrum. The amount of available instant access memory becomes a bottleneck to reaching system sizes larger than $L = 12$ (on a computer with 2GB of memory), unless the implementing software makes use of an *ARPACK* option to avoid explicit construction in memory of the 2^L dimensional matrix. *Mathematica* currently (v.7) suffers from this drawback, whereas *R* does not.

Figure 2.6 shows the complete spectrum of the TASEP transition matrix for $L = 3$, obtained by exact numerical diagonalisation. It consists of $2^L = 8$ eigenvalues, including the zero eigenvalue whose associated eigenvector is the stationary distribution. For $\beta \lesssim 0.125$ the spectrum contains a single complex conjugate pair, which is joined for $\beta \gtrsim 0.125$ by another. The appearance and formation of these conjugate pairs is a generic feature of the spectra for small system sizes.

It is worthwhile examining the low lying spectrum for several other small system sizes in order to ensure our observations are generic. Moreover we may then consider scaling with system size, which will lead us to discussion of the dynamical exponent. Figure 2.7 shows only the first few eigenvalues of the TASEP transition matrix with real parts closest to zero.

Several features are worthy of note. The first is the behaviour of the gap at $\alpha = \beta$, at which point there is a first order phase transition (the coexistence line in the stationary phase diagram, figure 2.5). We expect the gap to go to zero in the thermodynamic limit, corresponding to the emergence of an additional mode with infinite lifetime (possibly one of many) and an accordingly modified stationary state and probability distribution. Plotting the gap as a function of L at coexistence (see figure 2.8) suggests it would not be surprising that indeed

$$\lim_{L \rightarrow \infty} \varepsilon(\alpha = \beta) = 0, \quad (2.47)$$

however for convincing evidence one needs to use the finite-size values to perform a rigorous extrapolation prediction for the expected thermodynamic behaviour. This can




Figure 2.6: Real parts of the spectrum of the TASEP transition matrix as a function of β . Strictly real eigenvalues are in black, complex eigenvalues in blue. If viewing in Adobe Acrobat Reader (v.7 and upwards), the reader may manipulate the spectrum in 3D to observe eigenvalues in the complex plane.

be achieved as follows.

One wants to obtain the best possible estimate of the asymptotic value of ε , based on a small set of exact values $\varepsilon(L)$ for $2 \leq L \leq L'$, where L' is the largest system size for which the transition matrix can conveniently be numerically diagonalised. The true convergent behaviour of $\varepsilon(L)$, though unknown, can be described as

$$\varepsilon(L) = a_0 + a_1 L^{-\omega_1} + a_2 L^{-\omega_2} + \dots, \quad (2.48)$$

where our desired asymptotic value is $a_0 = \varepsilon$, and where the ω_i are real and positive but otherwise unconstrained. One approach to determine a_0 would be to perform least-squares fits of polynomials in $\frac{1}{L}$ of similar form to (2.48), however due to the inclusion of free parameters $\omega_1, \omega_2, \dots$ in the exponents and the small number of known values of $\varepsilon(L)$, the extrapolated value then depends strongly on the number of terms used in the fit. Hence this approach is ambiguous.

A more appropriate method for extrapolating this kind of sequence, that is, a convergent sequence for which one has a small number of finite-size values that are known exactly, is the Bulirsch-Stoer (BST) algorithm [60–62]. The BST algorithm uses rational




Figure 2.7: Real parts of the five TASEP eigenvalues closest to zero (including the ground state) as a function of β , for $\alpha = 0.3$ and for the system sizes indicated. Strictly real eigenvalues are in black, complex eigenvalues in blue. Note the diminishing vertical scale for increasing system size, which indicates that all relaxation times grow with system size.




Figure 2.8: The gap as a function of system size for $\alpha = \beta = 0.3$, i.e. on the first order transition boundary between the high and low density phases.

functions to iteratively generate a sequence of extrapolant estimates that rapidly converge to the true value of a_0 . It does this under the convenient assumption that instead of the more general (2.48), the convergent behaviour is adequately described by a member of the family of polynomials

$$\varepsilon(L) = a_0 + a_1 L^{-\omega} + a_2 L^{-2\omega} + \dots, \quad (2.49)$$

parametrised by a single real positive ω . The difference $\delta^{(n)}$ between successive extrapolant estimates resulting from n iterations of the algorithm,

$$\delta^{(n)} \equiv a_0^{(n)} - a_0^{(n-1)}, \quad (2.50)$$

depends on the choice of ω , and it can be shown [62] that as $n \rightarrow \infty$, the proximity of the n^{th} estimate to the true value a_0 obeys

$$\left| a_0 - a_0^{(n)} \right| \leq \left| \delta^{(n)} \right|. \quad (2.51)$$

Since n is limited by the number of exact finite-size values available, then in order to obtain as accurate an extrapolated value as possible ω is chosen so as to minimise $|\delta^{(n)}|$. This selects the member of the family (2.49) that most closely matches the true convergence behaviour (2.48).

In [63] the BST procedure was applied for a range of values of α and β , and it was deemed reasonable to conclude that

$$z \approx \begin{cases} 2 & \text{for } \alpha = \beta \\ 3/2 & \text{for } \alpha > 1/2 \text{ and } \beta > 1/2 \end{cases} \quad (2.52)$$

These dynamical scaling results indicate that the dynamics on the coexistence line are of Edwards-Wilkinson type [58], i.e. diffusive, whereas throughout the entire maximal current phase they are of Kardar-Parisi-Zhang type [64], i.e. subdiffusive. Moreover, the dynamical exponent for the latter case is identical to that of the TASEP with periodic boundary conditions [65–67]. In the low and high density phases the gap and therefore also the longest relaxation time were both found in [63] to be finite and independent of system size. There the relaxation is exponentially fast, with $z = 0$.

The second feature worthy of note becomes clear only after checking the spectra for a range of boundary parameters and system sizes, and it is that there are no generic degeneracies that persist for arbitrary combinations of α and β . As is evident from [68, 69],

this is in contrast to the TASEP with periodic boundary conditions, i.e. on a ring, for which there are many doublets as well as higher order degeneracies, all of which were found to be quantitatively predictable. There, doublets were found to result from obvious physical symmetries of the model such as translation and charge conjugation, combined with reflection. It was found possible also to account for higher order degeneracies of the periodic model, but only by considering certain underlying symmetries of solutions to Bethe ansatz equations. These latter symmetries so far lack a clear physical interpretation.

For the open-boundary TASEP on the other hand one observes apparent degeneracies, in particular doublets, only for special combinations of α and β , and they are of two types. The first type corresponds to points at which two real eigenvalues appear to merge to form a complex conjugate pair, such as for $\beta \approx 0.26$ in figure 2.7(d). The second type corresponds to apparent crossings of real eigenvalues, such as observed for $\beta \approx 0.16$ in figure 2.7(d). One expects that both types can be understood as manifestations of the behaviour of exceptional points of a non-Hermitian operator [70], namely the transition matrix.

Though not pursued in this thesis, it would be intriguing to investigate the origin of these apparent degeneracies and to discover whether they are reflected in certain obvious physical symmetries or symmetries of solutions to the Bethe ansatz equations, as was found to be the case for the periodic-boundary TASEP. One major concern in this undertaking would be to determine which apparent crossings of real eigenvalues, including the merging points associated with the emergence of complex conjugate pairs, constitute true crossings and therefore genuine degeneracies, as opposed to so-called avoided crossings (also known as level-repulsions) [70–72] which are false, i.e. merely *apparent* degeneracies. Theoretical background for this spectral behaviour may be found in [70].

2.8 Exact dynamics

It has been shown [7, 15] that variants of the asymmetric exclusion process can be mapped onto quantum spin chain models, so that in each case the transition matrix is represented by an appropriate transformation as a Hamiltonian. This allows analytical techniques from that field to be fruitfully applied to determine the spectrum. For example, it then becomes possible in many cases to exactly compute the gap by means of a Bethe ansatz calculation. For the periodic system this was achieved quite some time ago [65, 66]. Only more recently was the Bethe ansatz successfully applied to the open system, starting with

an exact determination of the behaviour of the gap as a function of boundary parameters for the TASEP [18, 73], the results for which prompted the work of this thesis and are presented in section 2.8.1. Thereafter the behaviour of the gap was computed for the PASEP [74]. The analysis for the open boundary TASEP, which can be mapped onto a spin- $\frac{1}{2}$ Heisenberg XXZ chain, builds on [75–82].

2.8.1 Thermodynamic gap and dynamical phase diagram

In this section I summarise the results of [18, 73]. Once the gap has been calculated, one can plot the dynamical phase diagram by investigating the points at which the gap vanishes or displays a nonanalyticity as $L \rightarrow \infty$. The former reveals coexistence of two phases with infinite lifetimes, and indeed it was found that the gap vanishes as L^{-2} along the HD-LD coexistence line [18, 73]. Furthermore, the gap vanishes as $L^{-3/2}$ in the entirety of the MC phase, reflecting the generic long-range (power-law) correlations seen in this phase [56]. Both results confirm the preliminary work discussed in section 2.7.

In the remainder of this thesis I focus on the high- and low-density phases³, in which the gap remains finite in the limit $L \rightarrow \infty$. There is a region, marked LDI' and HDI' on Fig. 2.9, within which the gap assumes the analytic expression

$$\varepsilon_1(L) = -\alpha - \beta + \frac{2}{(ab)^{\frac{1}{2}} + 1} - \frac{\pi^2}{(ab)^{\frac{1}{2}} - (ab)^{-\frac{1}{2}}} L^{-2} + \mathcal{O}(L^{-3}) \quad (2.53)$$

in which

$$a = \frac{1 - \alpha}{\alpha} \quad \text{and} \quad b = \frac{1 - \beta}{\beta} . \quad (2.54)$$

Within the low-density phase ($\alpha < \beta, \alpha < \frac{1}{2}$), this form of the gap applies for values of $\beta < \beta_c$ where

$$\beta_c(\alpha) = \left[1 + \left(\frac{\alpha}{1 - \alpha} \right)^{\frac{1}{3}} \right]^{-1} . \quad (2.55)$$

Likewise, in the high-density phase ($\beta < \alpha, \beta < \frac{1}{2}$), the region within which the gap is given by (2.53) is bounded by $\alpha < \alpha_c$ where

$$\alpha_c(\beta) = \left[1 + \left(\frac{\beta}{1 - \beta} \right)^{\frac{1}{3}} \right]^{-1} . \quad (2.56)$$

³de Gier and Essler [18, 73, 74] refer to these as “massive” phases by analogy with the quantum spin chains; I shall only use the terminology associated with the static phase diagram to avoid confusion.




Figure 2.9: Dynamical phase diagram according to [18]. The solid lines show the static phase boundaries, along which the gap vanishes in the thermodynamic limit. The de Gier-Essler line (dotted) indicates a dynamical transition, dividing the stationary LD phase into dynamical subphases LDI' and LDII', and likewise the HD phase into dynamical subphases HDI' and HDII'.

In the remainder of the low-density phase, $\alpha < \frac{1}{2}, \beta > \beta_c$ the gap takes the form

$$\varepsilon_1(L) = -\alpha - \beta_c + \frac{2}{(ab_c)^{\frac{1}{2}} + 1} - \frac{4\pi^2}{(ab_c)^{\frac{1}{2}} - (ab_c)^{-\frac{1}{2}}} L^{-2} + \mathcal{O}(L^{-3}). \quad (2.57)$$

Finally, we have by symmetry that when $\beta < \frac{1}{2}, \alpha > \alpha_c$,

$$\varepsilon_1(L) = -\alpha_c - \beta + \frac{2}{(acb)^{\frac{1}{2}} + 1} - \frac{4\pi^2}{(acb)^{\frac{1}{2}} - (acb)^{-\frac{1}{2}}} L^{-2} + \mathcal{O}(L^{-3}). \quad (2.58)$$

In these expressions,

$$a_c = \frac{1 - \alpha_c}{\alpha_c} \quad \text{and} \quad b_c = \frac{1 - \beta_c}{\beta_c}. \quad (2.59)$$

The boundaries between the dynamic subphases are shown dotted in Fig. 2.9.

There are a number of interesting points to notice about the distinction between the static and dynamic phase diagrams, Figs. 2.5 and 2.9 respectively. First, in the LDI' and HDI' regions, the gap (2.53) does not exhibit any nonanalytic behaviour along the LD-HD phase coexistence line $\alpha = \beta < \frac{1}{2}$. It is also analytic across the static subphase boundaries $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{2}$. On the other hand, the gap is nonanalytic along the lines $\beta = \beta_c > \frac{1}{2}$ when $\alpha < \frac{1}{2}$, and $\alpha = \alpha_c > \frac{1}{2}$ when $\beta < \frac{1}{2}$: I will refer to these dynamic transition lines




Figure 2.10: Exact $L \rightarrow \infty$ gap (black, dashed) as a function of β for $\alpha = 0.3$, described by the HDI' equation (2.53) (red) for $\beta < \beta_c$ ($\beta_c \approx 0.57$, indicated), and by equation (2.57) (green) for $\beta > \beta_c$.

collectively as the *de Gier-Essler line*. Specifically as β is increased from zero at some fixed α , the gap first approaches zero at the LD-HD phase boundary, then increases in magnitude until β reaches the critical value β_c , and then remains constant as β is further increased: see figure 2.10. I regard this constancy of the gap above a critical threshold as the signature of a dynamical transition in the TASEP, and note that similar transitions also occur in the PASEP [74].

Although there is no *a priori* reason for the static and dynamic subphases to coincide, the fact that they do not suggests that the mechanisms generating these subdivisions are physically distinct.

Furthermore, de Gier and Essler find that for $L \rightarrow \infty$ the second gap ε_2 collapses onto the first gap. However the leading order finite L corrections to the two asymptotic gaps are such that they obey

$$\lim_{\beta \uparrow \beta_c} \varepsilon_2(L) = \lim_{\beta \downarrow \beta_c} \varepsilon_1(L) \neq \lim_{\beta \uparrow \beta_c} \varepsilon_1(L). \quad (2.60)$$

In chapter 6 I propose that this finding can feasibly be interpreted as indicative of a crossing of the gap and the second gap for L very large, and I explore the repercussions of that proposal.

2.9 Alternative dynamical description: domain wall theory

In this section I summarise domain wall theory, which provides an alternative characterisation of ASEP dynamics, and compare it to the exact relaxation result.

The relaxation behaviour of the ASEP for any initial conditions and outside the maximal current phase can be seen in simulation to consist, on a coarse-grained scale, of the stochastic competition between a domain of density $1 - \beta$ enforced by the right boundary and a domain of density α enforced by the left boundary. For $\alpha = \beta$ the wall separating these domains performs a random walk through the bulk of the system, reflecting off the boundaries. For any $\beta > \alpha$ though with $\alpha < \frac{1}{2}$, i.e. in the low-density regime, the wall travels to the right boundary and the lower density domain dominates the system in the steady state. The long-time behaviour then consists of fluctuations of the position of the wall close to the right boundary, and vice versa at the left boundary for $\beta < \alpha$.

This picture is central in the domain wall theory (DWT) description [15, 55, 83], in which the collective relaxational dynamics are taken to be effectively reducible to a single coordinate describing the position of an interface, or wall. The wall separates a domain of density ρ^- and current $j^- = \rho^-(1 - \rho^-)$ to the left from a domain of density ρ^+ and current $j^+ = \rho^+(1 - \rho^+)$ to the right. Each domain is taken to possess the steady state characteristics imposed by the boundary reservoir on that particular side of the wall, with ρ^+ and ρ^- therefore as in Fig. 2.11(a). For example, for $\alpha < \frac{1}{2}$, $\beta > \frac{1}{2}$, the wall is taken to separate a low density domain on the left and a maximal current domain on the right. The effective theory is expected to be exact along the line $\alpha = \beta < 1/2$ where the exact properties of the system such as density profile [53, 84] and current fluctuations [85] are recovered.

Let us consider the case $\alpha < 1/2$ and $\beta < 1/2$. The motion of the wall is then described as a random walker coupled to and driven by the domains that it separates, with left and right hopping rates D^- and D^+ respectively given by imposing mass conservation on the fluxes into and out of the wall ([83]):

$$D^- = \frac{j^-}{\rho^+ - \rho^-} = \frac{\alpha(1 - \alpha)}{1 - \alpha - \beta} \quad D^+ = \frac{j^+}{\rho^+ - \rho^-} = \frac{\beta(1 - \beta)}{1 - \alpha - \beta}, \quad (2.61)$$

For $\alpha > \beta$ the random walk is biased to the left and in the stationary state the domain wall is localised at the left boundary and the bulk density is given by $\rho^+ = 1 - \beta$. For $\alpha < \beta$ the random walk is biased to the right and in the stationary state the domain wall

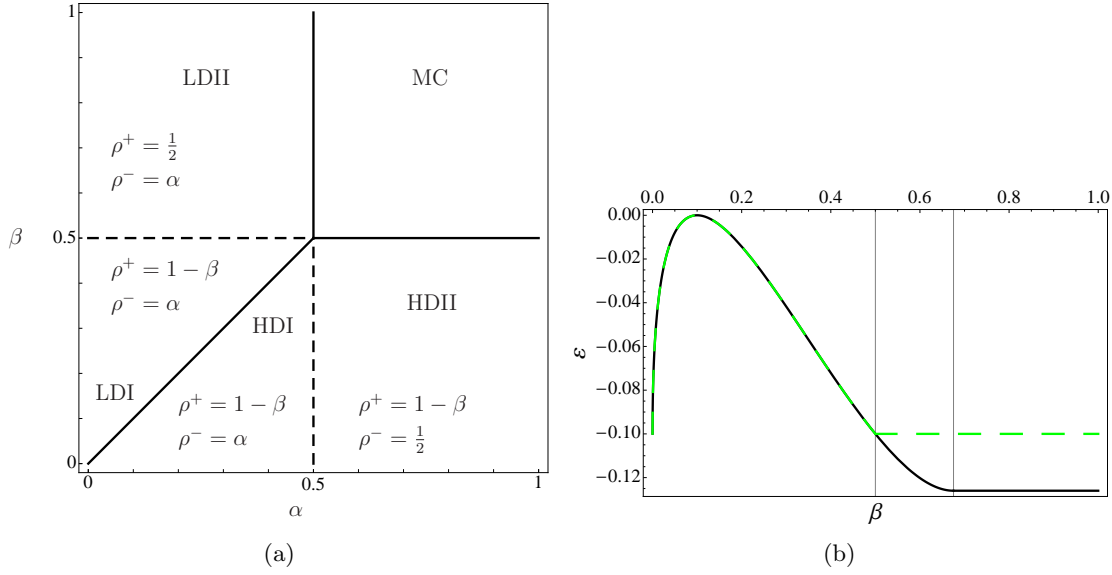


Figure 2.11: (a) Bulk densities of left and right domains used in domain wall theory (DWT) description of relaxational dynamics (b) Comparison of the gap for $L \rightarrow \infty$ from domain wall theory (green, dashed), given by (2.62), to the infinite L exact gap (black), for $\alpha = 0.1$ and $\beta_c \approx 0.675$

is localised at the right boundary and the bulk density is given by $\rho^- = \alpha$. Thus the first order transition at $\alpha = \beta$ is correctly predicted.

Moreover, the gap in the resulting spectrum for $L \rightarrow \infty$ is given [83] by

$$\varepsilon_{DWT}^{\infty} = -\left(D^+ + D^- - 2\sqrt{D^+ D^-}\right). \quad (2.62)$$

Remarkably, this expression is identical to (2.53) to order $1/L^2$. Thus in the region $\alpha < 1/2, \beta < 1/2$ DWT correctly predicts the thermodynamic gap.

When $\beta > 1/2$ so that $\rho^+ = 1/2$, the microscopic dynamics may be chosen as $j^+(\ell)$, where ℓ is the size of the right hand domain, and $j^+(\ell)$ is the stationary current in a TASEP of size ℓ in the maximal current phase. This implies the large ℓ behaviour

$$j^+ \simeq \frac{1}{4}\left(1 + \frac{3}{2\ell}\right) \quad (2.63)$$

The microscopic hopping rates D^{\pm} are now position dependent and the position dependence coming from (2.63) results in a modification of the density profile to an exponential spatial decay modulated by a power law with power $3/2$, similar to equation 2.44 for LDII.

In brief, DWT is remarkably successful, correctly predicting the static phase diagram

(including subphases), and the exact thermodynamic gap function found by de Gier and Essler in the region $\alpha < 1/2$ and $\beta < 1/2$. It differs in that the dynamic subphases are not given by the de Gier-Essler lines, but the static subphase transition lines. Thus, the region within which the gap is constant is different in the two theories, as indicated by Fig. 2.11(b).

2.10 This thesis: the dynamical transition

The exact stationary properties of the open-boundary TASEP are well-understood, and previous to [18, 73] it was thought that the only relevant dynamical behaviour consisted of boundary-driven crossovers between Edwards-Wilkinson diffusive scaling, and Kardar-Parisi-Zhang scaling associated with interface growth. The determination in [18] of a previously unsuspected boundary-induced dynamical feature, namely the dynamical transition discussed in section 2.8.1, introduced into the dynamical phase diagram a surprising structure that one suspects constitutes a more fine-grained distinction than the dynamical scaling crossovers. It is therefore desirable to understand the nature of the dynamical transition, i.e. to gain some insight into the underlying origin and perhaps the microscopic mechanism.

The work presented in this thesis probes the dynamical transition from various angles in an effort to improve our understanding. In chapter 3 I present a mean-field study of the gap obtained by analysing the Burgers equation, and compare the dynamical transition there observed with the exact result. In chapter 4 I report to what extent the dynamical transition manifests itself in Monte Carlo simulation. In chapter 5 I show how a numerical renormalisation scheme, the density matrix renormalisation group, may be applied to the open-boundary TASEP and whether its prediction of the gap confirms the dynamical transition. Then, in chapter 6, I consider the possible involvement of the second excited state, i.e. the second-longest-lived transient, in the dynamical transition. This is achieved by the perturbative investigation of a seemingly novel kind of dynamical crossover scenario, corresponding to a crossing of the gap and second gap eigenvalues. I conclude in chapter 7 by drawing together the available evidence so far, and speculate on the underlying mechanism for the dynamical transition. I also discuss a side issue that arises during the study of the dynamical transition, namely the differences in spectral manifestations of static and dynamical phase transitions in systems described by finite state-space Markov models.

Chapter 3

Gap from Burgers equation

In chapter 2 I set the challenge of understanding the transition encountered in the exact Bethe ansatz solution of the TASEP gap. One possible way to shed light on the dynamical transition is to see whether it also occurs in a mean-field approximation, and if it does, to ask whether it reveals any distinction between the nature of the long-time relaxation above versus below the transition.

As we saw in chapter 2, making a mean-field approximation for the TASEP and taking the continuum limit of the site occupancies means that the time evolution of the density profile $\rho(x, t)$ is described by the viscous Burgers equation. In this chapter I therefore present an asymptotic analysis ($t \rightarrow \infty$ and $L \rightarrow \infty$) of $\rho(x, t)$ obtained by solving the Burgers equation with boundary conditions corresponding to the open boundaries of the TASEP.

First, in section 3.1 I show how the Cole-Hopf transformation linearises the Burgers equation, turning it into the linear diffusion equation. The solution $\rho(x, t)$ to the Burgers equation may then be written down by applying the inverse of the Cole-Hopf transformation to a linear combination of the diffusive modes $\phi(x, t)$. This demonstrates that the mean-field gap corresponds to the difference between the two largest eigenvalues in the associated diffusive spectrum. In section 3.2 I show how applying the appropriate boundary conditions determines the diffusive spectrum, and that by identifying the behaviour for $L \rightarrow \infty$ of the two largest eigenvalues as a function of α and β , one finds that the mean-field gap exhibits a dynamical transition similar to that found in the exact gap, though at a different location. In section 3.3 I find boundary-matching eigenfunctions for finite systems, and construct the stationary state. In section 3.4 I focus on the relaxation behaviour of the long-time density profile, which corresponds to a moving domain wall and

is found to differ subtly above compared to below the transition. Finally, in section 3.5 I compare the dynamical transition findings from the mean-field analysis to the exact transition result, as well as to the domain wall theory characterisation of the dynamics and its prediction of a transition.

The Burgers equation has been studied extensively, both as a simplified model for hydrodynamical turbulence [86, 87] and, more recently and with added noise, to describe interface growth in nonequilibrium systems described by the Kardar-Parisi-Zhang equation, to which it is closely related [64, 88, 89]. The work by H.C. Fogedby on the deterministic [90, 91] and noisy [92, 93] Burgers equations in particular deals with this connection, and elucidates the dynamical interplay between diffusive and soliton modes. However most of the literature focuses on the time-asymptotic formation and structure of shocks on the interval $(-\infty, \infty)$, with constant-field boundary conditions at infinity. This is true of PDE-analysis literature on the Burgers equation, including by Burgers himself and [94, 95], and even of works set specifically in the context of the ASEP such as for example [96–98], including probability-theoretic treatments of the ASEP such as [99–101]. The limit $L \rightarrow \infty$ of the solution on the finite interval $[0, L]$ with *variable* boundary fields, which describes the behaviour of the open-boundary TASEP in the mean field, is rarely, if ever treated.

As far as I am aware only the recent [102] considers mean-field predictions of the dynamical exponents for the open-boundary TASEP, though without explicitly calculating the gap. According to an argument given in [102] in the mean field one obtains $z = 1$ in the LD regime, i.e. a vanishing gap. This disagrees with the exact result by de Gier and Essler as well as with the explicit analysis in this chapter of the mean-field gap, both of which find a gapped phase, corresponding to $z = 0$.

3.1 The Cole-Hopf transformation, diffusive modes, and the Burgers density

Our aim is to solve the viscous Burgers equation,

$$\partial_t \rho(x, t) = -(1 - 2\rho) \partial_x \rho + \frac{1}{2} \partial_x^2 \rho, \quad (3.1)$$

with boundary conditions corresponding to the open boundaries of the TASEP, namely

$$\rho(0, t) = \alpha \quad \text{and} \quad \rho(L, t) = 1 - \beta, \quad (3.2)$$

and with an arbitrary initial profile $\rho(x, 0)$. A standard approach to obtain the solution (see, for example [103, 104]) is by means of the Cole-Hopf transformation [105, 106], which amounts to the variable change $\rho(x, t) \rightarrow u(x, t)$ according to

$$\rho(x, t) = \frac{1}{2} [1 + \partial_x \ln u(x, t)] . \quad (3.3)$$

It is straightforward to show, by substituting (3.3) for $\rho(x, t)$ in (3.1), that (3.1) is then satisfied iff $u(x, t)$ obeys the diffusion equation,

$$\partial_t u(x, t) = D \partial_x^2 u(x, t) , \quad (3.4)$$

with diffusion constant $D = \frac{1}{2}$. Some finite-size nonstandard situations may lead to a spatially-dependent $D(x)$, for example when the TASEP has bond disorder in the form of site-dependent hopping rates, as investigated in [107].

The boundary conditions on $u(x, t)$ are determined by applying the Cole-Hopf transformation to the boundary conditions (3.2) on $\rho(x, t)$, giving:

$$u'(0, t) - a u(0, t) = 0 \quad \text{and} \quad u'(L, t) - b u(L, t) = 0 , \quad (3.5)$$

where

$$a = 2\alpha - 1 \quad \text{and} \quad b = 1 - 2\beta . \quad (3.6)$$

We proceed by standard separation of variables, writing $u(x, t) = \phi(x) T(t)$ so that we get

$$\frac{1}{2} \frac{\phi''}{\phi} = \lambda = \frac{\dot{T}}{T} . \quad (3.7)$$

The spatially-dependent part (after taking the factor $\frac{1}{2}$ to the time-dependent part) satisfies

$$d_x^2 \phi(x) = \lambda \phi(x) , \quad (3.8)$$

together with the boundary conditions

$$\phi'(0) - a \phi(0) = 0 \quad \text{and} \quad \phi'(L) - b \phi(L) = 0 \quad (3.9)$$

obtained from the boundary conditions (3.5) on $u(x, t)$ using $T(0) = 1$. Imposing the boundary conditions (3.9) determines the diffusive spectrum, which given that (3.8) is a Sturm-Liouville problem (see e.g. [108]) we expect to consist of a discrete set of eigenvalues λ_n that are all real and which can be ordered such that $\lambda_0 > \lambda_1 > \lambda_2 > \dots$, together with a complete set of mutually orthogonal eigenfunctions $\phi_n(x)$. The time-dependent part

for each mode n is required to satisfy

$$2\dot{T}_n(t) = \lambda_n T(t), \quad (3.10)$$

which is achieved by $T_n(t) = e^{\frac{\lambda_n t}{2}}$.

In section 3.2 I explicitly determine the diffusive spectrum, however here we first consider the general solution for $u(x, t)$, which can be written as a linear superposition of the diffusive modes, each weighted by its own time-dependence:

$$u(x, t) = \sum_{n=0}^{\infty} c_n u_n(x, t) \quad (3.11)$$

$$= \sum_{n=0}^{\infty} c_n e^{\frac{\lambda_n t}{2}} \phi_n(x), \quad (3.12)$$

where the coefficients c_n are determined by an initial profile $\rho(x, 0)$. Using the Cole-Hopf transformation (3.3) to map $u(x, t)$ back to $\rho(x, t)$ we obtain the general solution for the Burgers density:

$$\rho(x, t) = \frac{1}{2} \left[1 + \partial_x \ln \left(\sum_{n=0}^{\infty} c_n e^{\frac{\lambda_n t}{2}} \phi_n(x) \right) \right], \quad (3.13)$$

which can be written as

$$\rho(x, t) = \frac{1}{2} \left[1 + \frac{\phi'_0(x)}{\phi_0(x)} + \partial_x \ln \left(\sum_{n=0}^{\infty} \frac{c_n}{c_0} e^{\frac{(\lambda_n - \lambda_0)t}{2}} \frac{\phi_n(x)}{\phi_0(x)} \right) \right], \quad (3.14)$$

and, to show concretely the effect which we consider in section 3.4 of the Cole-Hopf transformation on the linear superposition of the diffusive modes, as

$$\rho(x, t) = \frac{1}{2} \left[1 + \frac{\phi'_0(x)}{\phi_0(x)} + \frac{\sum_{n=0}^{\infty} c_n e^{\frac{(\lambda_n - \lambda_0)t}{2}} \phi_0(x) \left(\frac{\phi_n(x)}{\phi_0(x)} \right)'}{\sum_{n=0}^{\infty} c_n e^{\frac{(\lambda_n - \lambda_0)t}{2}} \phi_n(x)} \right]. \quad (3.15)$$

It is worth noting that the relaxation times present in the problem are determined by the relative positions $\lambda_n - \lambda_0$ of the diffusive eigenvalues within the spectrum, rather than by their absolute values. The final term in (3.14) and hence in (3.15) vanishes as $t \rightarrow \infty$ since $\lambda_n < \lambda_0$ for all $n > 0$ so that all contributions to the sum decay to zero. Thus the stationary profile is

$$\rho(x) \equiv \lim_{t \rightarrow \infty} \rho(x, t) = \frac{1}{2} \left[1 + \frac{\phi'_0(x)}{\phi_0(x)} \right]. \quad (3.16)$$

In section 3.2 we determine the eigenfunctions $\phi_n(x)$ and obtain an explicit expression for the mean-field stationary profile according to (3.16). This is then compared with the exact stationary profile and the discrete-space mean-field stationary profile, both of which were reviewed in chapter 2.

For now, we return our attention to the form (3.14) of the full time-dependent solution to the viscous Burgers equation. The transient behaviour in the long-time regime can be identified by recognising that each successive term inside the sum is exponentially smaller than the previous term. To leading order, therefore, the approach to stationarity is characterised by keeping the first two diffusive modes, ϕ_0 and ϕ_1 :

$$\rho(x, t) \approx \rho(x) + \frac{1}{2} \partial_x \ln \left(1 + \frac{c_1}{c_0} e^{\frac{(\lambda_1 - \lambda_0)t}{2}} \frac{\phi_1(x)}{\phi_0(x)} \right) \quad (3.17)$$

$$\approx \rho(x) + \frac{1}{2} \frac{c_0}{c_1} e^{\frac{(\lambda_1 - \lambda_0)t}{2}} \left(\frac{\phi_1(x)}{\phi_0(x)} \right)'. \quad (3.18)$$

where we have expanded $\ln(1 + y)$ for $y \propto e^{\frac{(\lambda_1 - \lambda_0)t}{2}}$ small. It is clear from (3.17) that the timescale of the longest relaxation is set by the difference between the largest two eigenvalues in the diffusive spectrum. To be precise, (3.18) allows us to define the Burgers gap, ε_B , as the asymptotic rate of the exponential decay of the density profile to its stationary form:

$$\varepsilon_B \equiv \lim_{t \rightarrow \infty} \frac{d}{dt} \ln[\rho(x, t) - \rho(x)] = \frac{\lambda_1 - \lambda_0}{2}. \quad (3.19)$$

Hence to find ε_B and determine whether it exhibits a thermodynamic transition we need to determine λ_0 and λ_1 and their behaviour for $L \rightarrow \infty$. This we do in the succeeding section.

3.2 Thermodynamic diffusive spectrum and the Burgers gap

To determine the thermodynamic diffusive spectrum and the two largest eigenvalues therein we return to the 2nd order equation (3.8) for the spatial part of $u(x, t)$. From the separation of variables procedure (3.7) there are three distinct possibilities to consider for the diffusive eigenvalues λ , namely $\lambda > 0$, $\lambda < 0$ and $\lambda = 0$. We take each in turn, and summarise afterwards.

$\lambda = 0$

Solving equation 3.8 for $\lambda = 0$ gives the solution $\phi(x) = Ax + B$. Applying the boundary conditions (3.9) then gives $ab = \frac{a-b}{L}$, so that

$$\lim_{L \rightarrow \infty} ab = \lim_{L \rightarrow \infty} (2\alpha - 1)(1 - 2\beta) = 0. \quad (3.20)$$

Hence $\lambda = 0$ is part of the infinite-volume diffusive spectrum only for $\alpha = \frac{1}{2}$ or $\beta = \frac{1}{2}$ (or both).

$\lambda > 0$

To find positive eigenvalues we represent $\lambda > 0$ as $\lambda = \gamma^2$, where γ is real. The general solution to equation 3.8 for the spatial eigenfunctions $\phi(x)$ can then be written as

$$\phi_\gamma(x) = A \cosh(\gamma(x + C)) + B \sinh(\gamma(x + C)). \quad (3.21)$$

The third constant C , any real value of which means (3.21) is a solution to (3.8) for $\lambda > 0$, is included to facilitate obtaining a solution in section 3.4 for L large but finite. This will be achieved by two convenient choices, namely $C = 0$ and $C = -L$, each of which, as we will see in section 3.4, is an optimal choice for matching behaviour of the solution at one of the boundaries.

Now, still for general C , we see that imposing the boundary conditions (3.9) on (3.21) leads, for the left boundary, to

$$\tanh(\gamma C) = \frac{aA - \gamma B}{\gamma A - aB}, \quad (3.22)$$

and for the right boundary, to

$$\tanh(\gamma(L + C)) = \frac{bA - \gamma B}{\gamma A - bB}, \quad (3.23)$$

where, as before, we have

$$a = 2\alpha - 1 \quad \text{and} \quad b = 1 - 2\beta. \quad (3.24)$$

So far we have not made use of the freedom to choose C . For either choice $C = 0$ or

$C = -L$, (3.22) and (3.23) jointly simplify to

$$\tanh(\gamma L) = \frac{(a-b)\gamma}{ab - \gamma^2}. \quad (3.25)$$

In the limit $L \rightarrow \infty$ the left-hand side of (3.25) approaches ± 1 , depending on the sign of γ . Assuming $\gamma > 0$ reduces (3.25) to a quadratic with solutions

$$\gamma_+ = \{-a, b\}, \quad (3.26)$$

whereas for $\gamma < 0$ the resulting quadratic gives

$$\gamma_- = \{a, -b\}. \quad (3.27)$$

Of course $\gamma_+^2 = \gamma_-^2 = \lambda$, hence all that is left to determine is when these eigenvalues form part of the spectrum. Regardless of the choice γ_+ or γ_- , requiring that the resulting solutions are consistent with the assumed sign in each case means that the contributions of positive eigenvalues to the thermodynamic spectrum are

$$\lambda_\infty^+ = a^2, \quad \text{as long as } \alpha < \frac{1}{2} \text{ regardless of } \beta, \quad (3.28)$$

as well as

$$\lambda_\infty^+ = b^2, \quad \text{as long as } \beta < \frac{1}{2} \text{ regardless of } \alpha. \quad (3.29)$$

Clearly since the only other part of the spectrum left to consider consists of negative eigenvalues, (3.28) and (3.29) provide the two largest eigenvalues in the diffusive spectrum when $\alpha < \frac{1}{2}$ and $\beta < \frac{1}{2}$, and only one of the largest when only one of those conditions holds.

$\lambda < 0$

Again representing λ as $\lambda = \gamma^2$, solutions for the spatial eigenfunctions $\phi(x)$ with $\lambda < 0$ have purely imaginary γ , $\gamma = ik$ (k real), and can be written as

$$\phi(x) = A \cos(k(x+C)) + B \sin(k(x+C)). \quad (3.30)$$




Figure 3.1: LHS (solid, black) and rhs (dashed, various colours) of equation 3.31, intersections of which determine the solutions k that give negative diffusive eigenvalues $\lambda = (ik)^2$. Each colour corresponds to a different combination of the signs of $(a - b)$ and ab appearing in the rhs of (3.31) and hence to a different region of the phase diagram, as shown in figure 3.2. Here I have chosen $\alpha = 0.2$, $\beta = 0.3$ for the green region, $\alpha = 0.3$, $\beta = 0.6$ for the red region, and $\alpha = 0.4$, $\beta = 0.9$ for the blue region. For the blue and green regions the smallest positive k solution (and hence the largest negative λ) occurs for $\frac{\pi}{L} < k < \frac{2\pi}{L}$, whilst for the red region this takes place for $0 < k < \frac{\pi}{L}$.

Applying boundary conditions and following the same procedure regarding the choice of C as above, one finds that the allowed values of k satisfy

$$\tan(kL) = \frac{(a-b)k}{ab+k^2}. \quad (3.31)$$

To determine the behaviour of k for $L \rightarrow \infty$ and for all the possible choices of α and β (and hence a and b) we resort, in figures 3.1 and 3.2, to a graphical analysis of (3.31). These figures show that throughout the LD and HD phases the smallest positive solution k that solves (3.31) scales as $1/L$. Hence the largest negative diffusive eigenvalue $\lambda = (ik)^2$ scales as $1/L^2$, so that for $L \rightarrow \infty$ the largest contribution to the spectrum from the negative eigenvalues is therefore

$$\lambda_{\infty}^{-} = 0. \quad (3.32)$$




Figure 3.2: Showing the colours used in figure 3.1 to denote the regions of the phase diagram, which determine the behaviour of the rhs of equation 3.31, namely: $(a - b) > 0$, $ab > 0$ (blue), $(a - b) < 0$, $ab > 0$ (red) and $(a - b) < 0$, $ab < 0$ (green).

Burgers gap

Figure 3.3 shows an overview of the findings for the two largest eigenvalues in the thermodynamic diffusive spectrum, on the basis of which we can now state the $L \rightarrow \infty$ behaviour of the Burgers gap as α and β are varied. When both α and β are smaller than $\frac{1}{2}$, there are two isolated positive eigenvalues $\lambda = \{a^2, b^2\}$. In this region, then, in accordance with (3.19) the gap behaves as

$$\begin{aligned} \varepsilon_B &= -\frac{|a^2 - b^2|}{2} = -\frac{|(2\alpha - 1)^2 - (2\beta - 1)^2|}{2} \\ &= -2|\alpha(1 - \alpha) - \beta(1 - \beta)|. \end{aligned} \quad (3.33)$$

Like the exact gap (2.53), this vanishes along the HD-LD coexistence line $\alpha = \beta < \frac{1}{2}$. However, unlike the exact gap, the Burgers gap has a nonanalyticity along this line due to the modulus sign in 3.33.

If α or β is increased above $\frac{1}{2}$, one of the two solutions with positive λ ceases to exist, and the second-largest eigenvalue, originating from the negative part of the diffusive spectrum, is $\lambda_1 = 0$ in the thermodynamic limit. Hence then the size of the Burgers gap is simply equal to half the value of the largest eigenvalue. That is, within the low-density phase $\alpha < \frac{1}{2}$, $\beta > \frac{1}{2}$,

$$\varepsilon_B = -\frac{(1 - 2\alpha)^2}{2} \quad (3.34)$$




Figure 3.3: Showing the two largest eigenvalues in the diffusive spectrum associated with the viscous Burgers equation, the difference between which gives the mean-field Burgers gap ε_B . A saturation of the resulting form of ε_B to a constant value at $\alpha = \frac{1}{2}$ and at $\beta = \frac{1}{2}$ subdivides the high density (HD) and low density (LD) stationary phases, and corresponds to a mean-field analogue of the exact dynamical transition.

and within the high-density phase $\alpha > \frac{1}{2}$, $\beta < \frac{1}{2}$,

$$\varepsilon_B = -\frac{(1-2\beta)^2}{2}. \quad (3.35)$$

As a result the Burgers gap saturates to a constant value at $\alpha = \frac{1}{2}$ and at $\beta = \frac{1}{2}$, which is the same signature of a dynamical transition found in the exact solution at $\alpha_c(\beta)$ and at $\beta_c(\alpha)$.

The Burgers gap thus exhibits nonanalyticities along the coexistence line $\alpha = \beta < \frac{1}{2}$, and along the lines $\alpha < \frac{1}{2}$, $\beta = \frac{1}{2}$ and $\alpha = \frac{1}{2}$, $\beta < \frac{1}{2}$. As well as showing the largest eigenvalues in the diffusive spectrum, figure 3.3 therefore simultaneously constitutes the mean-field dynamical phase diagram resulting from our analysis of the Burgers equation. Note that here the dynamic subphases LDI and LDII (likewise HDI and HDII) correspond with the exact *static* subphases.

One way to visualise the nonanalyticities in the Burgers gap is to plot it as a function of β at some fixed α . We show this plot in Fig. 3.4, along with the exact gap for comparison. We note the qualitative similarities, namely the vanishing at the static transition and saturation in the LDII phase, even though quantitatively they are rather different.




Figure 3.4: Comparison of the gap for $L \rightarrow \infty$ from the spectrum of the viscous Burgers equation (3.1) (red, dashed) to the infinite L exact gap according to (2.53), (2.57) (black, solid), for $\alpha = 0.1$ for which $\beta_c \approx 0.675$ according to (2.55)

3.3 The stationary density profile

Now that we have determined the first two diffusive modes ϕ_0 , ϕ_1 , and the two largest eigenvalues λ_0 , λ_1 in the thermodynamic spectrum, we would first of all like to compute the stationary density profile for finite systems. Also, since a dynamical transition was found in the Burgers gap, we would like to construct the longest-time transient density profile from the first two diffusive modes, and study its relaxation above and below the Burgers transition.

We can compute the stationary density profile predicted by the viscous Burgers equation by returning to the expression (3.16), namely

$$\rho(x) \equiv \lim_{t \rightarrow \infty} \rho(x, t) = \frac{1}{2} \left[1 + \frac{\phi_0'(x)}{\phi_0(x)} \right]. \quad (3.16)$$

The largest eigenvalue always comes from the positive part from the spectrum, hence the general solution for $\phi_0(x)$ is given in (3.21). However applying the boundary condition (3.25) to the general solution we determined the eigenvalue only thermodynamically, hence we do not know the precise eigenvalues nor the precise eigenfunctions for finite L . This situation suggests we try to obtain some leading-order modification for L large but finite to the known asymptotic eigenvalue and eigenfunction.

This is where the freedom to choose C , hinted at in the discussion of the general solution (3.21) of $\phi(x)$ for $\lambda > 0$, plays a role. Though I will not prove so here, it turns out that for each of the eigenfunctions associated with $\lambda = a^2$ and with $\lambda = b^2$, one of the two choices of $C = 0$ and $C = -L$ is optimal, in that it allows us to determine the leading order behaviour of the eigenfunction for finite but large L . As we shall see, in each case the

optimal choice of C ensures that the eigenfunction exactly matches the conditions at *one* of the boundaries for any finite L , while the other boundary condition is satisfied exactly only for $L \rightarrow \infty$. For the inferior choice of C one does not directly obtain a meaningful expression for the resulting finite-size density profile. Additional work is then needed to obtain the leading-order correction for finite L , for example by means of an ansatz for the eigenvalue. Making use of the freedom afforded by the general solution with variable C spares us this effort.

Instead of labeling the eigenfunctions by the positions of their associated eigenvalues in the spectrum, which vary as a function of α and β , it is useful to denote by ϕ_a the eigenfunction always associated with the positive eigenvalue $\lambda = \gamma^2 = (-a)^2$, for which the optimal choice turns out to be $C = -L$, and likewise by ϕ_b the eigenfunction for $\lambda = \gamma^2 = b^2$, for which the optimal choice is $C = 0$:

$$\phi_a(x) = A \left(\cosh(a(L-x)) - \left(\frac{b}{a}\right) \sinh(a(L-x)) \right), \quad (3.36)$$

$$\phi_b(x) = A \left(\cosh(bx) + \left(\frac{a}{b}\right) \sinh(bx) \right). \quad (3.37)$$

One can check that $\phi_a(x)$ matches exactly the condition at the right boundary for any L , whilst $\phi_b(x)$ does so at the left boundary.

In the LD regime $\lambda_0 = a^2$ so we use $\phi_0 = \phi_a$ in (3.16), which gives

$$\rho(x) = \frac{1}{2} \left[1 - \frac{a \tanh(a(L-x)) - \frac{b}{a}}{1 - \frac{b}{a} \tanh(a(L-x))} \right]. \quad (3.38)$$

Using the identity

$$\tanh(A+B) = \frac{\tanh A + \tanh B}{1 + \tanh A \tanh B} \quad (3.39)$$

means (3.38) can be rewritten as

$$\rho_{\text{LD}}(x) = \frac{1}{2} \left[1 - a \tanh \left(a(L-x) + \tanh^{-1} \left(\frac{-b}{a} \right) \right) \right]. \quad (3.40)$$

Similarly, in the HD regime $\lambda_0 = b^2$, so we use $\phi_0 = \phi_b$ in (3.16), which gives

$$\rho_{\text{HD}}(x) = \frac{1}{2} \left[1 + b \tanh \left(bx + \tanh^{-1} \left(\frac{a}{b} \right) \right) \right]. \quad (3.41)$$

Again, a and b are directly related to the boundary parameters through

$$a = 2\alpha - 1 \quad \text{and} \quad b = 1 - 2\beta. \quad (3.42)$$




Figure 3.5: For $\alpha = 0.3$, $\beta = 0.4$, $L = 100$: Burgers equation prediction of the stationary density profile in the LD phase, given by (3.40) (blue curve). The exact density profile on the lattice is plotted in comparison (black \bullet).

As expected, the expressions for the continuum mean-field stationary profiles (3.40) and (3.41) differ from the discrete-space mean-field profiles on the lattice discussed in [7, 12, 107]. They do however agree with the continuum soliton form discussed in [7] (namely equation 16 therein) and in [92]. It seems useful that here the precise way the boundary parameters enter has been made explicit.

It is straightforward to verify that $\rho_{\text{LD}}(x)$ matches exactly the boundary density at $x = L$, and $\rho_{\text{HD}}(x)$ exactly the density at $x = 0$, for any system size. On the other hand each expression for the stationary density profile matches the other boundary ($x = 0$ and $x = L$ respectively) exactly only for $L \rightarrow \infty$. The choice of C that led to (3.40) and (3.41) appears natural, or optimal, not just because it easily leads to expressions satisfying the boundary conditions in a controlled way, but also because if it is chosen the other way around ($C = 0$ for $\lambda = a^2$ and $C = -L$ for $\lambda = b^2$) then the resulting expressions predict zero density.

Comparing ρ_{LD} to the exact density profile – figure 3.5 – highlights that in the LD regime the mean-field profile deviates most seriously from the exact solution by underestimating the occupancy close to the right boundary, as one expects from known discrete-space mean-field results. I have checked that there is no discrepancy compared to the exact results when $\alpha + \beta = 1$, where the profile is flat and we expect mean-field results to be exact.

3.4 Relaxation of the longest-lived transient: soliton motion from nonlinear superposition of diffusive modes

Having determined the continuum stationary density profile in the mean field, we also want to investigate the behaviour as a function of boundary parameters of the mean-field density profile associated with the longest-lived transient modes. The contributions of the individual diffusive modes $\phi_n(x)$ to $u(x, t)$ decay with a uniform rate at all points in space, in accordance with (3.12). The same holds therefore for the entire linear combination of all modes, so that the relaxation of u consists of the superposition of profiles each decaying at a uniform rate everywhere in the system.

However, as we will see the logarithmic nonlinearity combined with the derivative in the Cole-Hopf transform (3.3) means that the effect of the decay of the diffusive modes becomes spatially localised. In this section we will see that this produces a highly localised time-dependent rate of relaxation for the Burgers density, leading in the long-time regime to collective motion of a stable feature recognised in the literature [92] as a soliton solution. The transmutation of diffusive modes into soliton motion for the Burgers equation is also discussed in [93].

The longest-lived transient behaviour of the density profile is described by the first two diffusive modes $\phi_0(x)$ and $\phi_1(x)$ according to (3.17), namely

$$\rho(x, t) \approx \rho(x) + \frac{1}{2} \partial_x \ln \left(1 + \frac{c_1}{c_0} e^{\varepsilon_B t} \frac{\phi_1(x)}{\phi_0(x)} \right), \quad (3.43)$$

with ε_B the Burgers gap, defined as in section 3.2:

$$\varepsilon_B = \frac{\lambda_1 - \lambda_0}{2}. \quad (3.44)$$

In section 3.1 after keeping only the first two modes we made an additional approximation for the \ln in (3.43) to make explicit the asymptotic time dependence and hence define the Burgers gap according to (3.19). However in order to study the asymptotic behaviour of the density profile it becomes necessary to keep all contributions from the two modes, so that concretely the longest-lived soliton solution is given by writing (3.43), this time

without further approximation, as

$$\rho(x, t) \approx \rho(x) + \frac{1}{2} \left[\frac{\frac{c_1}{c_0} e^{\varepsilon_B t} \left(\frac{\phi_1(x)}{\phi_0(x)} \right)'}{1 + \frac{c_1}{c_0} e^{\varepsilon_B t} \left(\frac{\phi_1(x)}{\phi_0(x)} \right)} \right]. \quad (3.45)$$

The same expression is also given directly from the full solution written as in (3.15) by keeping all contributions from the first two modes. Unlike (3.45), the further approximated (3.18) used to define the Burgers gap does not have a time-decaying term in the denominator. Only (3.45) is properly normalised to give finite densities for finite t . It may arise that (3.45) exactly describes the evolution of $\rho(x, t)$ for *all* time, namely if the initial conditions are such that the coefficients c_n in the nonlinear superposition (3.13) are zero for all $n > 1$.

Let us consider the low density regime. As shown in figure 3.3, in the LDI phase we are below the transition in the Burgers gap found in section 3.2. The two largest eigenvalues are then $\lambda_0 = a^2$ and $\lambda_1 = b^2$, and the associated eigenfunctions for L large but finite $\phi_0 = \phi_a$ and $\phi_1 = \phi_b$ are given by (3.36) and (3.37). We can plot (3.45) in the LDI phase, and observe the soliton solution – see figure 3.6(a) – which consists of a sharp and stable domain wall separating a region of density α on the left from a domain of density $1 - \beta$ on the right, and travelling with constant velocity towards the right boundary.

Whereas in the LDI phase the two largest eigenvalues were positive, in the LDII phase, when $\beta > \frac{1}{2}$, there is only one positive eigenvalue, $\lambda = a^2$. The second largest eigenvalue comes from the negative part of the spectrum and is 0 for $L \rightarrow \infty$, but genuinely negative for finite L . As for ϕ_a and ϕ_b before, we want to obtain an expression for finite L for the eigenfunction $\phi_k(x)$ from (3.30) associated with the largest eigenvalue in the negative part of the spectrum. The largest negative eigenvalue $\lambda = (ik)^2$ is given by the smallest positive value of k that solves the condition (3.31) when $\alpha < \beta > \frac{1}{2}$ and $\alpha + \beta < 1$, i.e. in the ‘red’ region of the phase diagram as illustrated in figure 3.2. As we saw in figure 3.1 this value of k , which we denote by k^* , lies just below $\frac{\pi}{L}$. We can therefore make the following ansatz for k^* :

$$k^* = \frac{\pi}{L} - \frac{\Delta}{L}, \quad (3.46)$$

where Δ is taken to be positive and small on the scale of π , i.e. $\Delta \ll \pi$. We then want to determine Δ by substituting the ansatz (3.46) into the condition (3.31) that any k has to satisfy. This yields

$$\tan(\pi - \Delta) = \frac{(a - b) \left(\frac{\pi - \Delta}{L} \right)}{ab + \left(\frac{\pi - \Delta}{L} \right)^2}. \quad (3.47)$$

Using the periodicity of \tan and the fact that it is an asymmetric function, and taking

the dominant contribution for large L on the rhs gives

$$-\tan(\Delta) \approx \frac{(a-b)}{abL}(\pi - \Delta). \quad (3.48)$$

Finally, using the assumption that Δ is small and hence expanding the tan and rearranging gives

$$\Delta = \frac{\pi}{1 - \frac{ab}{a-b}L}, \quad (3.49)$$

and hence

$$k^* = \frac{\pi}{L - \frac{a-b}{ab}}. \quad (3.50)$$

In the red region therefore, the first excited state $\phi_1(x)$ for large L is $\phi(x)$ from equation 3.30 with $k = k^*$ given by (3.50):

$$\phi_1(x) = \cos(k^*x) + \frac{a}{k^*} \sin(k^*x). \quad (3.51)$$

For the k -mode either choice $C = 0$ or $C = -L$ yields equally good results, and in (3.51) C has been set to 0.

Plotting (3.45) for the LDI phase with $\phi_0 = \phi_a$ and ϕ_1 as in (3.51) we see in figure 3.6(b) that the resulting behaviour at late times of the density decay above the Burgers transition also consists of the motion, with constant velocity, of a domain wall. The main difference between this domain wall profile and that below the transition is the ‘hump’: whereas below the transition the profile at all points to the right of the wall are at the same density as the right boundary, above the transition the profile immediately to the right of the wall significantly exceeds the right boundary density, before decreasing only gradually to the boundary value.

3.5 Summary

The asymptotic analysis of the Burgers equation in this chapter shows that in the continuum limit of the mean-field approximation of the TASEP, the dominant relaxational mode has an interface between two regions, and that that interface travels through the system with constant velocity. Although the existence of this soliton solution is very well established as the elementary excitation ([92]) of the Burgers equation, what has not previously been studied to my knowledge is the precise effect on the relaxation time of this soliton due to the coupling to variable driving at the boundaries.




Figure 3.6: Steady state (black, solid) and long-time transient (red, dashed) density profiles $\rho(x,t)$, given by the mixture (3.45) of the first two modes in the solution to the viscous Burgers equation. Initial conditions are chosen so that domain walls are initially near the center of the system, after which they travel to the right. Successive dashed curves for linearly increasing times depict shape of solution as domain wall approaches the right boundary travelling at a constant velocity. For example, in (a) successive dashed curves, starting from the leftmost, describe the transient profile at $t = 0, 200, 400$, and $t = 600$. Similarly in (b) the dashed curves are for $t = 0, 100, 200$, and $t = 300$.

By studying the asymptotic relaxation time of the soliton solution, i.e. the thermodynamic Burgers gap ε_B , I have found that a dynamical transition similar to the one in the exact gap does indeed occur in the mean-field. Like the exact dynamical transition it is characterised by a nonanalyticity in the gap, in particular an independence with regards to one of the boundary parameters. This takes place at $\alpha = \frac{1}{2}$ and at $\beta = \frac{1}{2}$, as opposed to at the larger values $\alpha_c(\beta)$ and $\beta_c(\alpha)$.

The mean-field dynamical transition therefore occurs at the same point as the dynamical transition observed in domain wall theory. This is not particularly surprising since, as we saw in chapter 2, the basis of domain wall theory consists of a mean-field picture of the effect of driving at the boundaries. However in contrast to the Burgers results, domain wall theory correctly predicts the exact gap behaviour within the mean-field transition boundaries $\alpha = \frac{1}{2}, \beta = \frac{1}{2}$. The analysis presented in this chapter therefore elucidates the commonalities between domain wall theory and full mean-field theory, and strengthens the suggestion that the reason why domain wall theory fails to predict the exact dynamical transition is because it fails to accurately take into account the precise nature of the coupling of the domain wall to the boundaries.

The mean-field dynamical transition also occurs at the same point as the distinction between the exact static subphases, which is characterised by a change in the functional form of the lengthscale over which the exact stationary profile approaches its boundary

value. In the mean-field stationary solution there is no such distinction, because the single mode $\phi_a(x)$, which has an intrinsic lengthscale depending only on α , is associated with stationarity throughout the LD phase.

I also find that the Burgers gap for $L \rightarrow \infty$ remains finite throughout the LD phase, corresponding to a dynamical exponent $z = 0$ in agreement with the exact result. This contradicts with the claim in [102] that the thermodynamic mean-field gap vanishes with an effective dynamical exponent $z = 1$ resulting from the traversal of a kinematic wave of the full length of the system. Clearly both the exact solution and the mean-field result presented here indicate that the longest-lived transient behaviour does not scale with system size in this way. Instead, as can be confirmed from simulation, in the LD phase the long-time behaviour consists of fluctuations in the occupancy of sites near to the right boundary, on a timescale independent of system size.

Beyond determining the occurrence and location of the dynamical transition in the mean-field, I briefly investigated the long-time transient relaxation profiles. In section 3.4 we saw how the nonlinear superposition of diffusive modes according to the truncated expansion (3.17) (or, equivalently, equation 3.43) for the behaviour of the Burgers density $\rho(x, t)$ at late times gave rise to soliton solutions, consisting of moving domain walls. It should be possible to show, from the solutions obtained, that the velocity of these walls corresponds to the velocity of the shock as described by domain wall theory and the soliton solution analysed in [92]. Above the transition, the domain wall profiles showed a qualitative change in the shape of the profile (a ‘hump’), on the same side of the wall as the boundary that ceases to influence the relaxation time. Though it is not immediately clear what the underlying physics associated with this change is, it speculatively suggests, for example in the LD phase, that above the dynamical transition there is an accumulation, or congestion of particles in the bulk of the system with an effective density exceeding that at the right boundary. The long-time relaxation then consists of the clearing of this excess density.

Chapter 4

Gap from Monte Carlo simulation

Already in chapter 2 we saw that the spectral gap is inversely proportional to the longest relaxation time. This property allowed us in chapter 3 to identify the gap from the solution to the Burgers equation, used to approximate the TASEP. There we saw that the gap set the timescale for the dominant contribution to the nonlinear relaxation of the density profile in the long-time regime. So far I have not, however, said anything about how the longest relaxation time manifests itself directly in the observable and measurable behaviour of the TASEP itself, and hence discussion of the gap has remained rather abstract.

In this chapter I show how the role of the longest-lived mode in the TASEP dynamics may be made more concrete. This is done by identifying relaxation and correlation functions that yield direct measurements of the gap. In particular, I investigate whether the dynamical transition behaviour of the gap reported by de Gier and Essler can be observed in simulation.

Correlation functions for the TASEP have previously been investigated in [109], [110], and [111], though the latter deals with the TASEP subject to a parallel updating scheme, which corresponds to a discrete-time master equation [112]. Though [109] extracts estimates of the dynamical exponents from simulation, none of these works explicitly studies the dependence of the gap on boundary parameters.

The relaxation and correlation functions in question are computed on the basis of data generated by Monte Carlo simulation of the process. Hence I begin in section 4.1 by discussing two possible methods that may be used to simulate the process.

The first method, random sequential updating (see [112] and section 4.1.1), is straightforward to implement and commonplace in the ASEP literature. However, although exact for stationary properties it fails in general to reproduce exactly the time-dependent behaviour, and suffers from significant computational inefficiency.

The second method, varyingly referred to as Gillespie's direct method [113, 114], the n -fold way or the Bortz-Kalos-Lebowitz (BKL) algorithm [115] is detailed in section 4.1.2. It overcomes the inefficiencies inherent in random sequential updating and ensures the generation of stochastic trajectories whose ensemble properties are described by the continuous-time master equation (2.14) and the transition matrix (2.9). Further references to this and other Monte Carlo methods for simulating explicit dynamics are provided at key points in section 4.1.

In section 4.2 I explain how measuring the relaxation of an observable gives access to the gap. This is followed by results and discussion for three particular choices of observable, corresponding to the relaxation of an overall density excess (section 4.3), the correlations of density fluctuations in the steady state (section 4.4), and the relaxation of a second class 'tracer' particle trapped in the system (section 4.5). Finally in section 4.6 I discuss the obtained results, and in section 4.6.1 I compare the second class particle simulation results to the gap predicted by an approximate theoretical calculation of the second class particle motion.

4.1 Monte Carlo simulation

Faced with the analytical intractability of computing properties of systems that are composed of a 'large' number (which in many cases means more than two) of interacting constituents, be they gas particles, grains of sand, neurons, stockmarket traders, or stars in a globular cluster, Monte Carlo simulation is often an extremely useful tool.

A model of the system in question may, but need not, include a specific dynamical prescription, i.e. a set of rules governing time evolution. In the 'molecular' dynamics approach these rules reflect as closely as possible what are thought to be the true, exact equations of 'motion', which may themselves be classical or quantum mechanical, deterministic or stochastic. In other approaches, the dynamical rules can become rather coarse approximations of what are considered to be the true dynamics. The rules are chosen to be simple enough to be simulated using Monte Carlo methods or integrated numerically, whilst remaining faithful to the essential underlying symmetries

and conservation laws.

With increasing coarseness and simplification of the dynamical rules as well as restriction of the spatial degrees of freedom one enters the domain of the so called toy models, such as the Ising and Potts, and Hubbard and Heisenberg spin chain models, which have become central to theoretical condensed matter physics. Here the underlying laws may deliberately be distorted, disentangled, and deconstructed in order to perform, again using Monte Carlo methods, the kind of detailed quantitative thought experiments that are so essential to theorising.

If a model does *not* include an explicit dynamical prescription, this may be because it is deemed impossible to numerically integrate or to simulate with sufficient accuracy to calculate quantities of interest. Lacking the required computational resources, this was effectively the situation faced by Boltzmann, Gibbs, and other early contributors to statistical mechanics. Their solution to circumvent this challenge lies of course in the *a priori* probabilistic description that we are now so familiar with, and in the derivation of the stationary canonical distributions that are at the core of equilibrium statistical mechanics [10].

Once the physically relevant canonical distribution $P(X)$ giving the stationary probability of any state X is known, one may determine the stationary properties of an equilibrium model *without* explicitly solving or indeed simulating the dynamics. This is achieved by employing a Monte Carlo algorithm to effect importance sampling, that is, to generate a ‘representative’ sequence of states $X_1 \rightarrow X_2 \rightarrow \dots$ drawn according to $P(X)$, i.e. drawn in such a way that the frequency of occurrence $f_n(X)$ of X in a sequence of length n obeys

$$\lim_{n \rightarrow \infty} \frac{1}{n} f_n(X) = P(X). \quad (4.1)$$

The equilibrium average of some observable \mathcal{O} is then computed according to

$$\langle \mathcal{O} \rangle \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathcal{O}(X_t). \quad (4.2)$$

For comprehensive reviews of the nature and use of Monte Carlo methods in equilibrium statistical physics to compute stationary properties I refer to [116–118].

Depending on which of the canonical distributions one wants to sample, different rules are employed to generate the sequence of states $\{X_t\}_{t=1}^T$. If the model under investigation includes no explicit dynamical prescription, then the Monte Carlo rules may themselves be interpreted as an artificial dynamics, with t in (4.2) taking the role of time. This

furnishes the model with an artificial time evolution which, though in the first instance a tool to effect importance sampling, may be studied in its own right as a description of relaxation towards equilibrium [119].

For example, one may study the stationary properties at equilibrium of the Ising model (for which no dynamics are prescribed) by executing a Metropolis algorithm with spin-exchange rules that conserve total magnetisation, though not total energy, whenever a spin flips (so called Kawasaki dynamics [120]). Interpreted dynamically these rules describe the time evolution not only of a system already having reached a canonical equilibrium, where it is subject to fluctuations in the total energy but also, if the initial state X_1 lies outside the stationary ergodic set [10], of the approach to that equilibrium. Alternatively, one may implement single spin-flip rules that are nonconserving of both energy and magnetisation (so called Glauber dynamics [121]). This in turn corresponds to a system at or relaxing to a grand-canonical equilibrium.

The dynamics of the TASEP of course *are* explicitly defined, namely first and foremost as a set of $k(t)$ Poisson processes (see section 2.2), where at any given time $k(t)$ is equal to the number of pairs of sites (counting the left and right boundary reservoirs as ‘sites’) between which a particle is free to hop. This primary specification of the dynamics is encoded in the master equation formalism of section 2.3, and in particular in the transition matrix, equation 2.9.

For our purposes a simulation of the TASEP must be consistent with the primary specification, and crucially it must reproduce both the exact stationary distribution as well as the time evolution prescribed by the continuous-time master equation (2.14).

In [112] we find a comparison of a number of common update procedures for the asymmetric exclusion process. The use of a synchronous updating algorithm (also called parallel updating), in which all sites may change simultaneously during a single update step, simulates a discrete time process which is inconsistent with the TASEP model studied here. Ordered sequential updating, whereby moves are attempted in sequence from one side of the lattice to the other during successive update steps, does implement a continuous-time dynamics, however the ordering effects a different time evolution to that prescribed by the master equation (2.14) and leads to a different stationary distribution.

In contrast, the *random* sequential updating scheme (also described in [112], and detailed in section 4.1.1), in which sites are updated sequentially but in random order on the lattice, is commonly used to simulate the ASEP. With contemporary computational resources it easily yields accurate results for stationary properties. The scheme is

commonplace in the ASEP literature, to the extent that a newcomer to the field would not be blamed for mistaking it, e.g. on the basis of [49], as fundamental to the model and as strictly equivalent to the dynamics defined in section 2.2. In reality, it is just a particular algorithm for sampling the stationary distribution.

Though it leads to the correct stationary distribution, there are two objections to random sequential updating. First, that while it is exact for stationary properties it fails, for reasons related to the fact that it operates with a fixed time increment, to reproduce exactly the time-dependent behaviour specified by the master equation (for details see [122]). Second, it is computationally inefficient over a significant range of boundary parameters.

For dynamical studies requiring accurate measurement of the exact time-dependent behaviour¹ therefore, a method that does accurately reproduce the time-evolution specified by the master equation, such as the method described in [113, 114] and [115], is preferable. In section 4.1.2 I explain in detail how this method, known as Gillespie’s direct method, the n -fold way, the Bortz-Kalos-Lebowitz (BKL) algorithm [115], or simply as continuous-time Monte Carlo, may be used to simulate the open-boundary TASEP. Moreover, it has the additional advantage that it offers substantial computational efficiency gains over the random sequential update algorithm, by avoiding a weakness inherent to that approach.

Note that since the random sequential update scheme (which is an instance of a *random selection method*, see [122, 123]) as well as Gillespie’s direct method both implement a prescribed dynamics, they appear in the literature under the header either of dynamical Monte Carlo or, depending on the author, kinetic Monte Carlo methods. The former term commonly denotes an overarching category that includes the latter as a subset. Since for modern practitioners the distinction is uninformative at best I support taking this tendency to its logical conclusion and eliminating ‘kinetic’ whilst retaining ‘dynamical’ to denote any simulation of an explicitly prescribed dynamics. For enlightening overviews of dynamical (and kinetic) Monte Carlo methods, see [122–127].

4.1.1 Discrete-time Monte Carlo: random sequential updating

The simplest algorithm available to us to simulate the TASEP is a random sequential updating scheme, as follows. For each elementary iteration

¹For example [109], which studies time-dependent correlation functions for the open-boundary TASEP.

1. an integer i is chosen at random from $\{0, 1, 2, \dots, L\}$;
2. one of the following takes place:
 - If $1 \leq i < L$ we implement bulk dynamics for the pair of sites $i, i + 1$, namely if site i is occupied by a particle and $i + 1$ unoccupied, then with probability 1 the particle hops from i to $i + 1$.
 - If $i = 0$ (which does not correspond to a lattice site) we implement left boundary dynamics, namely if site $i + 1 = 1$ is unoccupied then with probability α a particle enters the system by hopping from the infinite particle reservoir on the left onto site 1.
 - If $i = L$ we implement right boundary dynamics, namely if site L is occupied then with probability β this particle exits the system by hopping from site L and disappearing into the infinite particle reservoir on the right.

For a system of size L these steps are iterated $L + 1$ times to give the evolution during one Monte Carlo sweep (MCS) — the basic unit of time in simulation. This choice mirrors the specification from section 2.2 that the Poisson processes for all ‘active’ pairs of sites are concurrent, and that hence the total number of attempted moves per unit time is proportional to the system size. Equivalently, this amounts to ensuring that the attempt probability per unit time for each site is independent of system size.

Since the random sequential updating scheme measures algorithmic time in fixed intervals, it is sometimes referred to as discrete-time Monte Carlo, even though it implements a continuous-time transition matrix. This is in contrast to the continuous-time Monte Carlo method discussed below in section 4.1.2, which operates with a continuously variable time interval. Though the random sequential updating scheme correctly samples the stationary distribution and can therefore be used to calculate steady-state properties to high precision, it is not suitable for high-accuracy dynamical studies.

4.1.2 Continuous-time Monte Carlo: Gillespie’s direct method

The random sequential update scheme of section 4.1.1 is easy to implement but does not give exact time-evolution, and is inefficient, especially when the state of the system allows only a small number of possible moves. This occurs at high and at low overall densities. In these situations the algorithm wastes time by attempting a large number of unsuccessful moves.

A more efficient approach is one that effects a successful move with every single attempt, after which the time variable is increased by a stochastic, exponentially distributed waiting time. In [113, 114] Gillespie provides theoretical grounding for a method that accomplishes this, and which produces stochastic trajectories in accordance with the desired master equation.

In fact [113, 114] outline two possible algorithmic realisations (see also [122, 126]): the *direct* method, and the *first-reaction* method. The procedure adopted here and outlined below corresponds to the former, which is also known as the *n-fold way* or the Bortz-Kalos-Lebowitz (BKL) algorithm [115].

The procedure involves maintaining a list of all possible moves, given the current state, and grouping the possible moves into classes according to their transition rates. For the TASEP there are three classes of moves:

1. A bulk particle hopping to the unoccupied neighbour site on its right, with rate $r_1 = 1$,
2. A particle hopping from the left reservoir onto site 1, with rate $r_2 = \alpha$,
3. A particle hopping from site L into the right reservoir, with rate $r_3 = \beta$.

The total rate R_i for each event class is then given by

$$R_i = \sum_{j=1}^{n_i} r_i = n_i r_i, \quad (4.3)$$

where n_i is the number of distinct opportunities that exist for a transition of type i to change the current state of the system. For example, if there are 10 particles in the bulk that are free to hop to their neighbouring sites on the right, then $n_1 = 10$. Hence we have

$$R_1 = n_1 \quad (4.4)$$

$$R_2 = \alpha \delta(1 - s_1) \quad (4.5)$$

$$R_3 = \beta \delta(s_L) \quad (4.6)$$

The next step is to compute the cumulative rates Q_k (cumulative over the event classes), according to

$$Q_k = \sum_{i=1}^k R_i, \quad (4.7)$$

with Q_3 giving the total rate for all allowed transitions away from the current state, and

with Q_0 set to 0. These are used to pick the class of the event that is to take place by drawing a random number u uniformly on the interval $[0, Q_3]$ and finding j such that

$$Q_{j-1} \leq u < Q_j . \quad (4.8)$$

This ensures that each event type is selected with probability

$$p_i = \frac{R_i}{Q_3} , \quad (4.9)$$

as desired.

The ordering of the event classes and hence of the values Q_k for $k < 3$ is arbitrary. However, for computational efficiency in finding j it is desirable to choose the order where Q_1 is largest, if indeed it is significantly larger than Q_2 and Q_3 . This is generally the case for the TASEP, hence the choice adopted above.

Once the event class i is chosen, the particular move to execute is chosen randomly with equal probability from the n_i instances. Finally time is increased by an interval

$$\Delta t = -\frac{1}{Q_3} \log(\nu) \quad (4.10)$$

where ν is drawn uniform random on $[0, 1]$. This ensures [113–115, 127] that algorithmic time is proportional to real time and that the simulation mimics events proceeding as independent Poisson processes, according to which one expects the event interval distribution

$$P(\Delta t) = e^{-Q_3 \Delta t} , \quad (4.11)$$

which is consistent with (4.10). Since algorithmic time in this approach varies continuously the method is sometimes referred to as continuous-time Monte Carlo, especially in contrast to the discrete-time random sequential update method of section 4.1.1.

Finally, it is worth mentioning that for Monte Carlo simulation of more complicated dynamical systems than the TASEP, including biochemical reaction networks with many different event classes and transition rates, optimisations [128–131] exist for ordering the Q_i and for deciding (4.8).

4.2 Measuring the gap

Having discussed simulation methodology in section 4.1, I show in this section what kind of simulation-based measurements provide access to the gap.

4.2.1 Relaxation of an observable

Consider an observable, \mathcal{O} , which takes the value \mathcal{O}_C when the system is in a configuration C belonging to the 2^L dimensional state space \mathcal{S} . By means of simulation one may measure $\langle \mathcal{O}(t) \rangle$, namely the average, over an ensemble of identically prepared systems evolving according to the master equation (2.14), of the value of \mathcal{O} at time t . This quantity is given, in the notation of chapter 2, by

$$\langle \mathcal{O}(t) \rangle \equiv \sum_{C \in \mathcal{S}} P_C(t) \mathcal{O}_C \quad (4.12)$$

$$= \langle 1 | \hat{\mathcal{O}} | P(t) \rangle, \quad (4.13)$$

where we have used

$$\hat{\mathcal{O}} |C\rangle \equiv \mathcal{O}_C |C\rangle, \quad (4.14)$$

and where $\langle 1 |$ is simply the 2^L -dimensional vector $(1, 1, 1, \dots)$ in the physical basis, i.e.

$$\langle 1 | \equiv \sum_{C \in \mathcal{S}} \langle C |, \quad (4.15)$$

with $\langle C |$ such that $\langle C | C' \rangle = \delta_{CC'}$. Then, bearing in mind the decomposition (2.21) of $|P(t)\rangle$ over the right eigenvectors $|\psi_i\rangle$ of the transition matrix M given in (2.9):

$$|P(t)\rangle = \sum_i b_i e^{\lambda_i t} |\psi_i\rangle, \quad (4.16)$$

and denoting by \mathcal{O}_i the projection of \mathcal{O} onto $|\psi_i\rangle$:

$$\hat{\mathcal{O}} |\psi_i\rangle = \hat{\mathcal{O}} \sum_{C_j \in \mathcal{S}} \langle \phi_i | c_j | C_j \rangle \quad (4.17)$$

$$= \sum_{C_j \in \mathcal{S}} c_j \mathcal{O}_{C_j} \langle \phi_i | C_j \rangle \quad (4.18)$$

$$\equiv \mathcal{O}_i |\psi_i\rangle \quad (4.19)$$

$$\Rightarrow \mathcal{O}_i = \langle 1 | \hat{\mathcal{O}} | \psi_i \rangle, \quad (4.20)$$

we have that $\langle \mathcal{O}(t) \rangle$ may be written as

$$\langle \mathcal{O}(t) \rangle = \sum_i b_i \mathcal{O}_i e^{\lambda_i t} . \quad (4.21)$$

Since $\lambda_0 = 0$ and all other eigenvalues have nonzero and negative real parts (at least when L is finite, see section 2.3), the average *stationary* value of \mathcal{O} can be identified as

$$\langle \mathcal{O}(\infty) \rangle \equiv \lim_{t \rightarrow \infty} \langle \mathcal{O}(t) \rangle = b_0 \mathcal{O}_0 . \quad (4.22)$$

To measure the relaxation of \mathcal{O} to its stationary value we can define the relaxation function

$$\Phi(t) \equiv \frac{\langle \mathcal{O}(t) \rangle - \langle \mathcal{O}(\infty) \rangle}{\langle \mathcal{O}(0) \rangle - \langle \mathcal{O}(\infty) \rangle} \quad (4.23)$$

$$= \sum_{i=1}^{2^L-1} c_i e^{\lambda_i t} , \quad (4.24)$$

where the coefficients

$$c_i \equiv \frac{b_i \mathcal{O}_i}{\sum_{i=1}^{2^L-1} b_i \mathcal{O}_i} \quad (4.25)$$

are normalised so that $\Phi(0) = 1$, and which decays to zero:

$$\lim_{t \rightarrow \infty} \Phi(t) = 0 . \quad (4.26)$$

For t large but finite the dominant contribution comes from the longest lived decay mode and $\Phi(t)$ obeys

$$\Phi(t) \underset{t \text{ large}}{\sim} e^{\lambda_1 t} . \quad (4.27)$$

Recalling that λ_1 is equal to the gap ε , it should be possible to determine the gap by measuring $\Phi(t)$ for suitably chosen observable and initial conditions.

Though it is not immediately obvious what might constitute suitable choices, clearly one prerequisite for (4.27) to hold is that the observable must have a non-zero projection onto the gap eigenmode ($\mathcal{O}_1 \neq 0$). This is equivalent to demanding that the observed quantity is not associated (exclusively) with some initial transient, and hence that it must continue to change in time until the steady state is reached. It seems safe to assume that choosing to observe for example the overall density, defined as the fraction of sites occupied, satisfies this criterion since it is a function of the time evolution of all the sites in the system.

Likewise and at the same time the system must be initialised so that $b_1 \neq 0$, so as not to eliminate that mode from the dynamics from the outset. In practice, it seems safe to assume that this prerequisite is satisfied by initialising the system with all sites occupied with the same probability, which is taken to be significantly higher than the stationary probability. Such a gross, system-wide perturbation to the stationary profile is likely to project onto a large collective of modes, including the gap mode.

Going beyond merely satisfying the minimum prerequisite criteria, one would like to specifically target the gap mode in order to improve the signal-to-noise ratio of $\Phi(t)$ when measured in simulation. This would be accomplished by choosing initial conditions that suppress the contributions to $\Phi(t)$ from higher modes, i.e. to get as close as possible to $b_{i>1} = 0$, and by choosing an observable that is more sensitive to the behaviour in the long-time regime than to initial transients, i.e. to get as close as possible to $\mathcal{O}_{i>1} = 0$.

If by some very particular choice of observable we actually had that $\mathcal{O}_{i>1} = 0$, or if by some equally particular choice of the initial conditions that $b_{i>1} = 0$, then extracting the gap value would be exceedingly simple since in this case

$$\Phi(t) \propto e^{\lambda_1 t} \quad (4.28)$$

would hold for all t . Unfortunately, since the eigenvectors of the transition matrix are not known (apart from numerical diagonalisation for small system sizes) it is not possible to construct such simplifying initial conditions or to choose such a simplifying observable. Hence instead of the ideal situation (4.28) of a single exponential decay one encounters in practice the mix of modes (4.24), which becomes purely exponential only in the long-time regime.

4.2.2 Fitting a sum of exponentials

Given the above, some further thought suggests three possible approaches to extract the gap value from a measurement of $\Phi(t)$. The first is to fit data for all measurement times (up to some latest time, determined by computational economy) from a simulation for system size L to the sum of exponentials

$$\sum_{i=1}^{2^L-1} d_i e^{\gamma_i t}, \quad (4.29)$$

with $\gamma_i < 0$. Fitting technicalities aside one expects this approach to yield the most accurate gap value *in principle*, since it fits the data to the true functional form (4.24) of

$\Phi(t)$.

Unfortunately the problem of finding the best fit of a high-dimensional nonlinear function of this form is highly ill-conditioned [132, 133] in the sense that satisfactory convergence is highly sensitive to choosing ‘good’ initial values for the parameters to be estimated. Ultimately this approach therefore involves considerable effort spent obtaining estimates $\gamma_{i>1}$ of all higher eigenvalues $\lambda_{i>1}$ and of all coefficients d_i , just to obtain an accurate estimate of λ_1 . Moreover, since the number of fitting terms scales as 2^L this approach quickly becomes prohibitive as the system size increases, not least due to the demand that the number of data points be sufficient to avoid overfitting.

4.2.3 Fitting a truncated sum of exponentials

The second approach is to fit data from intermediate times onward to a sum of only two or three (let us say n) exponential terms [134], i.e. to use the fitting function (4.29) truncated after n terms:

$$\sum_{i=1}^n d_i e^{\gamma_i t}. \quad (4.30)$$

This approach was used in [135] to determine the critical exponent of the three-dimensional Ising model from Monte Carlo simulation.

One expects the accuracy of the gap value obtained using this approach to depend on the extent to which the following two conditions are met, namely

1. The initial conditions and the observable project much more strongly onto a set of ℓ modes, which are a subset of the full spectrum. Formally: With Λ the full spectrum of the transition matrix, there exists a set Ω_ℓ of eigenvalues, $\Omega_\ell \subset \Lambda$, corresponding to ℓ modes for which the projection of the initial conditions, b_i , and that of the observable, \mathcal{O}_i , jointly satisfy

$$\mathcal{O}_i b_i \gg \mathcal{O}_j b_j, \quad (4.31)$$

for each $\lambda_i \in \Omega_\ell$ and for all the modes not in Ω_ℓ : $\lambda_j \in (\Lambda \setminus \Omega_\ell)$.

If this first condition holds then already for short times (4.30) is a good approximation to the true form of $\Phi(t)$, provided one fits $n = \ell$ exponential terms. However as before, without knowing the eigenvectors of the transition matrix it remains impossible to predict for a given choice of observable and initial conditions whether (4.31) is satisfied. Now consider the second condition, namely

2. The m largest relaxation times are each much longer than those of all higher modes. Formally: the first m eigenvalues λ_i form a set Λ_m by virtue of their satisfying

$$\operatorname{Re}(\lambda_i) \gg \operatorname{Re}(\lambda_j), \quad (4.32)$$

for each $\lambda_i \in \Lambda_m$ and for all $\lambda_j \in (\Lambda \setminus \Lambda_m)$, i.e. with $1 < i \leq m$ and $m < j \leq 2^L$.

If this second condition holds then after an initial timescale $t' \approx -\operatorname{Re}(\lambda_{m+1})^{-1}$ the modes with $\lambda_{k>m}$ will have decayed sufficiently such that for $t > t'$ the relaxation is effectively dominated by the m modes of Λ_m and (4.30) becomes a good approximation to $\Phi(t)$, provided one uses $n = m$. However, without knowing the full spectrum or the spectral statistics it is impossible to know *a priori* whether such a separation of timescales is present, let alone how many exponential terms are involved and hence whether the approach will succeed.

It is also possible for both conditions (4.31) and (4.32) to be satisfied at the same time. If this is the case, then with exception of the unlikely event that exactly the same modes are involved, i.e. with the exception that $\Omega_\ell = \Lambda_m$, then the expected behaviour of $\Phi(t)$ and the likely success of fitting (4.30) using a given number of exponential terms is more complicated to predict.

Clearly what is needed is some measure of the suitability of a fit using the truncated sum of exponentials (4.30), and hence a gauge of the likely accuracy of the gap value arrived at by this approach. This is provided by the normalisation check

$$\sum_{i=1}^n d_i = 1, \quad (4.33)$$

which can be used to empirically determine the optimal number of terms n to fit with. This is done by performing fits for increasing n , starting with $n = 1$, and calculating the normalisation sum in each instance.

To see how this works, imagine that there is indeed a well defined separation of timescales, with n' modes dominating after t' , or alternatively that just n' modes are selected at the outset by the initial conditions and choice of observable; or finally that both factors collude, with effectively the same result. Then this procedure should demonstrate a significant jump in the value of the normalisation (4.33) towards unity (and in the gap estimate towards the true gap value) as n is increased from $n' - 1$ to n' , followed by relatively little improvement as more than n' terms are used to fit.

In [135] it was found by checking $\sum_{i=1}^n d_i \approx 1$ that $n = 3$ is a sufficient number of terms

to accurately approximate the relaxation of the 3D Ising model for all but the shortest times, provided the observable \mathcal{O} was chosen to be the magnetisation. When observing the energy however, it was found that three terms were not (yet) sufficient to describe $\Phi(t)$ at intermediate times.

If on the other hand neither (4.31) nor (4.32) holds, then the distribution of relaxation times is relatively homogeneous so that there is no clear cut-off time t' separating a quick initial relaxation involving many modes from a longer-lived relaxation dominated by a small number of modes, nor will a smaller number of modes be selected from the outset. In this case the normalisation check (4.33) would yield a value incrementally closer to unity with every additional exponential term included in the fit. Likewise the accuracy of the estimated gap value would also improve gradually with increasing n , until soon one faces the ill-conditioned problem, familiar from the first approach, of trying to fit a prohibitive number of exponential terms.

An additional complication arises from the freedom to choose the shortest time t' , i.e. the lower bound determining what data to fit to. Consistent with the empirical approach outlined above one should vary t' independently of n' in order to determine the most accurate gap estimate, where that accuracy is judged by the normalisation condition or by a goodness-of-fit measure such as the coefficient of determination (the R^2 value). However, in practice there are instances where one encounters a normalisation result equally close to unity – or a similar R^2 value – for a fit of $n'' = n' + 1$ terms to data starting at $t'' < t'$, as for a fit of n' terms to data starting at t' . A significant difference in the two corresponding gap estimates γ'_1, γ''_1 would cast doubt on the robustness of the method.

Clearly then, and notwithstanding the success of [135], the accuracy of this approach is strongly dependent not just on the spectrum but also on the choice of observable, and on the initial conditions. Unless the model exhibits strong separation of timescales, or unless a very specific observable is chosen, or the system initialised in a very specific way, unambiguously obtaining a reliable gap estimate by fitting a truncated sum of exponentials may well be problematic. This situation is exacerbated by the additional free parameter t' .

We can investigate the feasibility of this approach for the TASEP by comparing the known gap for a small system size to the gap arrived at by fitting n exponential terms to $\Phi(t)$. Since for L small the spectrum can be calculated by numerical diagonalisation, $\Phi(t)$ can be computed explicitly, i.e. without performing a simulation, and by choosing a distribution for the coefficients c_i in (4.24). For sake of simplicity these are taken uniformly

as

$$c_i = \frac{1}{2L} \quad (4.34)$$

For $L = 2$ a sum of 3 exponential terms is able to recover the gap, no matter if only short-time data is used in the fit – not surprising since in this case $\Phi(t)$ contains only 3 modes. However already for $L = 8$ I found significant deviation from the exact gap with $n = 2$, $n = 3$, or $n = 4$ terms unless data for long times was included, which, if obtained from simulation, would be computationally costly.

As opposed to for the Ising model therefore, fitting to a truncated exponential appears not to be an efficient way to determine the gap for the TASEP. Moreover, we note that even the Ising results of [135], given by 2- and 3-term fits, have been criticised in [136, 137] as giving inaccurate estimates of the critical exponent. In [137] in particular the author discusses difficulties associated with extracting the longest relaxation time, and advocates fitting with a single exponential in the asymptotic regime, or using a spectral density method.

For these reasons we consider another approach which, though computationally more costly in terms of the generation of simulation data, should provide a less ambiguous, more systematic and hopefully accurate estimate of the TASEP gap.

4.2.4 Asymptotic fitting

The approach of section 4.2.3 consisted of attempts to fit to data an appropriate fitting function for the behaviour of $\Phi(t)$ including for intermediate and potentially for short times. However, even when the model exhibits effective separation of timescales, thereby allowing in principle for this type of analysis, the presence of free parameters raises ambiguities in determining the most appropriate fitting function, and in selecting the subset of data to fit to. Thus even in situations where the method could work, there is doubt about its robustness and about the accuracy of the estimated gap value.

In the approach described in this section we remove any ambiguity about the appropriate form of the fitting function by instead focusing explicitly on the long-time regime. Regardless of whether there is an effective separation of timescales or not, we know that asymptotically $\Phi(t)$ always approaches a purely exponential decay, with exponent equal to the gap. If one generates simulation data to measure $\Phi(t)$ for large enough times therefore, $\ln(\Phi(t))$ should approach a linear relationship, and the slope should approach the true value of the gap.

To demonstrate the advantage of this approach over those of previous sections we return to the exact behaviour of $\Phi(t)$ for $L = 8$, computed in section 4.2.3. This corresponds to a perfect (noiseless) measurement of $\Phi(t)$ with initial conditions distributed uniformly over the modes $-c_i$ according to (4.34). Figure 4.1(a) shows that the gap mode comes to dominate some time after $t \approx 10$. To reveal explicitly the convergence of the slope to the true gap value, it is informative to investigate the local slope $\ln(\Phi(t))'$ of $\ln(\Phi(t))$, that is, simply

$$\ln(\Phi(t))' \equiv \frac{\ln(\Phi(t + \delta t)) - \ln(\Phi(t))}{\Delta t}, \quad (4.35)$$

for points separated in time by Δt , where for all the results presented in this chapter I have taken $\Delta t = 1$. When plotted against $1/t$, the asymptotic local slope value is then given by the (extrapolated) intersect with the y -axis. This is shown in figure 4.2(b), which confirms that $\ln(\Phi(t))'$ converges to the exact gap value shown.

It should be understood that no explicit extrapolation of the finite-time data is implied. Instead, the aim of the approach is to obtain sufficiently clean data for large enough times such that the local slope obtained directly from the finite-time data agrees with the exact gap to within the desired accuracy, whatever that happens to be. Clearly this places strong demands on the quality of the data and hence on computational resources.




Figure 4.1: $\Phi(t)$ for a system of size $L = 8$ and with $\alpha = 0.3$, $\beta = 0.68$, computed as in section 4.2.3 from exact numerical diagonalisation and with mode coefficients c_i of (4.24) chosen uniformly according to (4.34). The local slope in (b) converges for $t \rightarrow \infty$ onto the y -axis to the exact gap, $\lambda_1 \approx -0.098$ (marked by horizontal line), hence the asymptotic analysis reveals what was not possible to obtain in section 4.2.2 by fitting a truncated sum of exponentials to short or intermediate-time data.

For noiseless data, as in figure 4.1, clearly the best estimate of the gap is given by the local slope value computed for the latest measurement time t_ℓ , and the accuracy of

that estimate improves the later the time. However for realistic and hence noisy data originating from Monte Carlo simulation the value at time t_ℓ is not necessarily the most accurate, and one must effect some kind of averaging over the long-time data to obtain the best possible gap estimate, given the data.

Including data for $t < t_\ell$ will improve the accuracy of the gap estimate, as long as the data still lies ‘sufficiently’ far in the asymptotic regime so as not to be affected significantly by the presence of higher modes. For example, in the case of figure 4.1 this would suggest including data for $t > 20$ (i.e. for $1/t < 0.05$).

Before considering how to determine the *shortest* time to include whilst still remaining ‘sufficiently’ far in the asymptotic regime, as well as how to effect the averaging, note that in practice the noise in $\ln(\Phi(t))$ grows rapidly for t very large. Hence a useful first step to improve accuracy of the final gap estimate is to exclude data for all times *larger* than some time t_ℓ .

For consistency it is desirable to choose the cut-off t_ℓ systematically, according to an objective criterion. I demonstrate the method chosen to accomplish this by first adding Gaussian noise with mean 0 and standard deviation 2×10^{-6} to the exact $\Phi(t)$ that was computed for figure 4.1. The resulting $\Phi(t)$, shown in figure 4.2, mimics the expected result of a Monte Carlo measurement.




Figure 4.2: $\Phi(t)$ for a system of size $L = 8$ and with $\alpha = 0.3$, $\beta = 0.68$, computed as in section 4.2.3 and figure 4.1 from exact numerical diagonalisation and with mode coefficients c_i of (4.24) chosen according to (4.34), but with added Gaussian noise (mean 0, standard deviation 2×10^{-6}). The exact gap is marked in (b) by a horizontal line.

We may quantify the quality of data by defining

$$\Delta\Phi(t) \equiv \frac{\ln(\Phi(t + \delta t))' - \ln(\Phi(t))'}{\ln(\Phi(t))' - \ln(\Phi(t - \delta t))'} \quad (4.36)$$

Namely, as $\ln \Phi$ approaches a straight line (and hence Φ a purely exponential relaxation) for large t , $\Delta\Phi$ approaches 1. For later times, noise causes significant deviations of $\Delta\Phi$ from 1, thus $\Delta\Phi$ provides a measure of the noise and may be used to discard data for $t > t_\ell$, where t_ℓ is defined as the first time for which

$$|\Delta\Phi(t)| > r, \quad (4.37)$$

for some constant r , which constitutes the cut-off criterion. Applied to the data of figure 4.2 this yields figure 4.3, showing that data for $t > t_\ell = 75$ is discarded if the criterion (4.37) is applied with $r = 5$.




Figure 4.3: $\Delta\Phi(t)$ versus t for a system of size $L = 8$, with $\alpha = 0.3$, $\beta = 0.68$, and with added noise. This quantifies the noise in $\Phi(t)$ of figure 4.2. With the cut-off criterion (4.37) set at $r = 5$, data for $t > t_\ell = 75$ is discarded.

One might be concerned that the asymmetric choice $|\Delta\Phi(t)| > r$ as a cut-off criterion rather than for example the symmetrised $|\Delta\Phi(t) - 1| > r$ represents a bias in the selection of the cut-off point and the discarding of noisy data. In practice, however, this is not significant as it produces very similar cut-off points.

Subsequently we must decide the shortest time t_s for which to include $\Phi(t)$ data to compute the gap estimate, keeping in mind that we wish to avoid times so short that $\Phi(t)$ is not well approximated by a pure exponential. In order to accomplish this in a systematic

way I make use of a goodness-of-fit statistic, namely the coefficient of determination, also known as the R^2 value, defined as

$$R^2 \equiv 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} = \frac{\sigma_y^2 - \text{ss}_r}{\sigma_y^2}, \quad (4.38)$$

where $\{y_i\}_{i=1}^n$ are the data, \hat{y}_i is the predicted value of y_i according to the fitted model, and $\bar{y} \equiv \frac{1}{n} \sum_{i=1}^n y_i$ the sample mean of y . Hence R^2 measures the size of the residual sum of squares, ss_r , relative to the sample variance of the data, and presents the result as a fraction of that same variance. The closer R^2 is to 1, the better the fit, with $R^2 = 0$ when the model is no better match to the data than a line with zero slope, constant value equal to the sample average at each point, and unity when the fit is perfect, i.e. when the sum of residuals is zero [138].

In fact I use R_{adj}^2 , a version of R^2 adjusted for varying number of model parameters, k and sample size, n , according to [138]:

$$R_{\text{adj}}^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2 / (n - k)}{\sum_i (y_i - \bar{y})^2 / (n - 1)}. \quad (4.39)$$

To determine the effectively purely exponential regime, i.e. to determine its lower bound t_s , I proceed as follows. I fit a single exponential $a e^{bt}$ to $\Phi(t)$ data for $t_s \leq t \leq t_\ell$ for a range of values of t_s , and choose the one whose corresponding fit gives the largest value of R^2 . This definition of the ‘shortest’ time for which the decay is effectively purely exponential coincides with the *best* fit to the data of a single exponential. The best estimate of the gap given the data is therefore taken to be the exponent parameter b from the fit for that value of t_s .

To clarify the procedure used to extract the best gap estimate, figure 4.4 shows this method applied to the data of figure 4.2 (trimmed to exclude data for $t > t_\ell$). The goodness-of-fit increases rapidly for $8 < t_s < 20$, reflecting the fact that for $t_s < 20$ many higher modes contribute significantly to the decay whilst for $t \approx 20$ a single exponential becomes a much better fit. For $t > 20$ R^2 decreases slowly, due to increased noise for later-time data. The exact gap is recovered with equivalent accuracy as in the noiseless case, despite the noise injected into $\Phi(t)$.

As far as I am aware this particular approach for extracting a gap value from Monte Carlo data has not appeared elsewhere in the literature. For additional discussion of some of the issues raised in this section and in sections 4.2.2 and 4.2.3 I direct the reader to [135, 139, 140].




Figure 4.4: For the data of figure 4.2, t_s is determined by fitting a single exponential to data starting at t_s and ending at t_ℓ , and observing which choice of t_s yields the largest R^2 value. From (a) we see that this corresponds to $t_s = 27$ (marked by vertical line). Figure (b) shows that the best gap estimate therefore lies at $t = 27$, and matches very closely the exact gap $\lambda \approx -0.098$, marked by a horizontal line.

Having described how the gap may be measured in general and having presented my particular method for analysing relaxation data to extract the gap value, I present in the succeeding sections 4.3, 4.4 and 4.5 three sets of results corresponding to particular choices of observables, the aim in each case being to compare the gap estimates with the exact gap computed by de Gier & Essler.

4.3 Gap from relaxation of an initial overall density excess or deficit

We may choose to observe the overall density, i.e. the fraction of occupied sites on the lattice. With N the total number of particles in the system, this amounts to choosing the observable \mathcal{O} as

$$\mathcal{O} = \rho(t) \equiv N(t)/L. \quad (4.40)$$

Denoting the relaxation function Φ defined in (4.23) under this choice of \mathcal{O} by Φ_ρ , we have

$$\Phi_\rho(t) \equiv \frac{\langle \rho(t) \rangle - \rho(\infty)}{\langle \rho(0) \rangle - \rho(\infty)}, \quad (4.41)$$

where $\rho(\infty)$ is simply the stationary overall density, given by summing up the exact expression for the density profile according to the matrix product ansatz [12, 53]:

$$\begin{aligned} \rho(\infty) &= \sum_{i=1}^L \rho_i(\infty) \\ &= \sum_{i=1}^{L-1} \left(\sum_{p=0}^{L-i-1} \frac{(2p)!}{p!(p+1)!} \frac{\langle W | C^{L-1-p} | V \rangle}{\langle W | C^L | V \rangle} + \frac{\langle W | C^{i-1} | V \rangle}{\langle W | C^L | V \rangle} \sum_{p=2}^{L-i+1} \beta^{-p} \frac{(p-1)(2(L-i)-p)!}{(L-i)!(L-i+1-p)!} \right) \\ &\quad + \beta^{-1} \frac{\langle W | C^{L-1} | V \rangle}{\langle W | C^L | V \rangle}. \end{aligned} \quad (4.42)$$

Here $\langle W | C^L | V \rangle$ normalises stationary weights in a system of size L , and hence plays the role of a nonequilibrium partition function. It is given^[12, 53] by

$$\langle W | C^L | V \rangle \equiv \sum_{p=1}^L \frac{p(2L-1-p)!}{L!(L-p)!} \frac{\beta^{-(p+1)} - \alpha^{-(p+1)}}{\beta^{-1} - \alpha^{-1}}. \quad (4.43)$$

In [110] it was observed that the power spectrum of $\rho(t)$ exhibits oscillations, which were related to the time taken for a fluctuation to traverse the entire lattice. In [111] the same quantity was studied, but for the open-boundary TASEP with parallel updating. Finally, [109] also measured similar correlation functions from simulation, and showed that they agree with a description of the dynamics that combines domain wall motion on long time scales with local relaxation of density fluctuations on smaller timescales.

The aim here is to obtain gap measurements below and above the de Gier-Essler transition for finite system sizes and then, if possible, to extrapolate to the thermodynamic gap behaviour to compare to the de Gier-Essler result. To this end, Monte Carlo simulations are initialised with uniform density, that is, with each site occupied with equal probability $\rho(0)$. For values of α and β corresponding to the low-density phase, where the left boundary dominates and imposes the stationary bulk density α , $\rho(0)$ is chosen as $1-\beta$ (i.e. equal to the right boundary reservoir density). Conversely in the high-density phase where the stationary bulk density imposed by the right boundary boundary is $1-\beta$, the choice is $\rho(0) = \alpha$. Depending on α and β these initialisations correspond to either an overall density excess, or an overall density deficit.

For each simulation, i.e. for each system size L and for each value of β with α kept fixed, $\Phi_\rho(t)$ is measured by performing an ensemble average, i.e. by averaging over many runs, as $\rho(t)$ relaxes to the value of $\rho(\infty)$ computed numerically from (4.42). Each run for given values of L and β is initialised with the same *average* density, as described above, though not with the same exact microscopic configuration, hence the ensemble average

is simultaneously an average over noise in the initial conditions as well as over stochastic trajectories.

Finite size gap estimates obtained from an initial set of simulations performed using the random sequential updating scheme of section 4.1.1 are shown in figure 4.5. These results show qualitative and, as a function of system size, increasing quantitative




Figure 4.5: Finite size gap estimates obtained from random sequential update Monte Carlo simulations by measuring $\Phi_\rho(t)$ according to equation 4.41, as a function of β for $\alpha = 0.3$. Successively higher lying curves are for $L = 10$ (blue, \bullet), $L = 20$ (purple, \blacksquare), $L = 30$ (gold, \blacklozenge) and $L = 50$ (green, \blacktriangle). The $L = \infty$ gaps predicted by de Gier & Essler (thick, black) and domain wall theory (thick, dashed, green) are shown for comparison. The dynamical transition is at $\beta_c \approx 0.57$ (marked), whereas domain wall theory predicts a transition at $\beta = 0.5$ (marked). Note, for $\beta < 0.5$ domain wall theory and de Gier-Essler coincide.

agreement with the expected thermodynamic behaviour over the entire range of β . Nonetheless for the largest system size reached ($L = 50$) there are still significant finite size effects, hence it is not possible to directly observe agreement with the de Gier-Essler results rather than with domain wall theory. Moreover, the small number of distinct system sizes investigated means that extrapolating from the data to the $L = \infty$ behaviour would yield unreliable results.

Hence a second set of simulations were performed, this time using the continuous-time Monte Carlo scheme of section 4.1.2. This allowed data for significantly larger system sizes and for a larger number of intermediate sizes to be obtained using the same

computational resources. For these simulations I chose to focus on a narrow region of β around the dynamical transition, in order to finely resolve the gap behaviour around β_c . The resulting finite size gap estimates are shown in figure 4.6 and figure 4.7.




Figure 4.6: Finite size gap estimates as a function of β for $\alpha = 0.3$, for $L = 10$ (blue, ●), $L = 20$ (purple, ■) and $L = 30$ (gold, ◆) obtained from continuous-time Monte Carlo simulations by measuring $\Phi_\rho(t)$ according to equation 4.41. The $L \rightarrow \infty$ de Gier-Essler exact gap (black, very thick) and domain wall theory gap (green, dashed) are shown for comparison. In addition, the exact gap from numerical diagonalisation for $L = 10$ is shown (black, thick). The dynamical transition is at $\beta_c \approx 0.57$ (marked), whereas domain wall theory predicts a transition at $\beta = 0.5$.

This second set of results demonstrates excellent agreement with the exact gap for $L = 10$ from numerical diagonalisation (figure 4.6), confirming the validity of the approach. Gap estimates for larger system sizes approach the expected thermodynamic behaviour (figure 4.7). However for the largest system size, $L = 150$, the gap estimate actually overshoots the thermodynamic gap, i.e. is closer to zero. This overestimation is systematic, i.e. it occurs over a wide range of β , hence the discrepancy is not simply the result of noise.

The reason for this discrepancy can be understood by taking a closer look at the asymptotic analysis used to arrive at the gap estimates. A typical measurement of $\Phi_\rho(t)$ yields data like that shown in figure 4.8 (c.f. the analogous figure 4.1 for the idealised case).




Figure 4.7: Finite-size gap estimates as a function of β for $\alpha = 0.3$ as in figure 4.6, but with successively higher lying curves showing finite size gap estimates for $L = 20, 30, 40, \dots, 100$, and 150.

In the analysis associated with figure 4.8, data for $t > 407$ is discarded following the procedure outlined in section 4.2.4. Following this procedure, the best estimate of the gap is then given by determining the choice of t_s (the shortest time for which to use data to fit a single decaying exponential to) that gives the largest value of R^2 . The outcome is shown in figure 4.9. As expected from the hypothetical example of section 4.2.4, the goodness-of-fit first increases rapidly as a function of t_s (namely, for $80 < t_s < 120$). This is followed (for $t_s > 140$) by a slow decrease in R^2 due to increased noise in long-time data.

We are now closer to explaining the overestimation observed in figure 4.7 (i.e. the fact that the gap is closer to zero than the exact result). Going back to figure 4.8(b) we see that the local slope of $\ln(\Phi_\rho(t))$ approaches its asymptotic value from above (in contrast to the idealised case of figure 4.1, where this occurs from below). This is true for $L = 150$ also, i.e. there the asymptotic local slope is also approached from above. This explains why the gap estimates are *overestimates*: since the true gap value is only reached asymptotically, and from above, any analysis based on fitting to finite-time data will estimate a gap value larger (i.e. more positive hence closer to zero and with smaller absolute value) than the true one.

Crucially, I found the behaviour observed in figure 4.8 to be generic for all system




Figure 4.8: $\Phi_\rho(t)$ data from simulation for $L = 50$, $\alpha = 0.3$, $\beta = 0.64$. Note that data for $t > 407$ is deemed to be overly noisy (see text) and hence excluded from (b) and from further analysis. Gridlines mark the same equal-time intervals in (b) as in (a).




Figure 4.9: For the data of figure 4.8, the best gap estimate is given by fitting a single exponential to data starting from $t_s = 140$ (marked in both (a) and (b)).

sizes and for all values of β examined in the narrow range around β_c . Hence there is good reason to suspect that all finite-size gaps shown in figure 4.7, not just the ones for $L = 150$, are in fact overestimates.

The degree of inaccuracy depends on how close the average local slope $\ln(\Phi_\rho(t))'$ over the regime $t_s < t < t_\ell$ is to the true gap value. This in turn depends on how well separated the higher modes are from the gap mode. It is thought [18] that as $\beta \rightarrow \alpha$, not just the gap itself but a large number of higher modes all vanish diffusively (i.e. with leading order L^{-2}) for $L \rightarrow \infty$. The closer to coexistence therefore, the less well separated the timescales are expected to be, and the larger the discrepancy (assuming equal computational investment) with the true gap. This can be observed in figure 4.7, since it is for β closer to α that the finite-size gap estimates exceed the exact thermodynamic gap. The gap estimates for $L = 10$ are very accurate, presumably, following the above argument, because there the timescales involved are short enough that the measured signal is nearly purely exponential,

and hence the results from fits are very close to the asymptotic gap results.

Despite the observed inaccuracies – whose origins I feel have been elucidated – one may still try to extrapolate from the finite-size gap estimates to the $L = \infty$ behaviour in order to compare the thermodynamic gap according to simulation with the de Gier-Essler result. From the latter we have the asymptotic series solution for the gap, according to which the gap scales with L to leading order as

$$a + bL^{-2}. \quad (4.44)$$

This gives us a function to fit to the finite-size gap estimates for each value of β and to extrapolate to the thermodynamic gap, a , predicted by simulation.

Of course for small system sizes the leading order expression (4.44) is a bad fit for the gap. Including higher order terms and hence fitting more parameters should allow a more accurate value of a to be extracted.

Using R_{adj}^2 as a goodness-of-fit measure we may legitimately judge the relative effectiveness of including additional terms. As expected, comparing fits of the quadratic (4.44), of a cubic, and a quartic in $1/L$ shows that the adjusted R^2 value generally improves as the number of terms increases. An example comparison is shown in figure 4.10.




Figure 4.10: As a function of β : Finite size gap estimates for $\alpha = 0.3$, $\beta = 0.54$ (blue points), best fit of $a + bL^{-2}$ (dashed, red curve) with adjusted $R^2 = 0.9934$, and best fit of $a + bL^{-2} + cL^{-3} + dL^{-4}$ (green curve) with adjusted $R^2 = 0.9998$.

The prospect of merely gradual incremental improvement in goodness-of-fit with the inclusion of additional terms appears to render arbitrary the decision of how many to use. However, it turns out the quartic $a + bL^{-2} + cL^{-3} + dL^{-4}$ is the optimal choice, since this is the lowest order fitting function for which one finds that the adjusted R^2 values are independent of β . For lower order fitting functions on the other hand there is a strong dependence and hence a non-systematic bias in the accuracy of the extrapolated gap value.

Extrapolated gap values (using the quartic) are shown in figure 4.11. As expected,




Figure 4.11: As a function of β and with $\alpha = 0.3$: extrapolated gap values for $L = \infty$ based on measurements of $\Phi(t)$ according to equation 4.41 (blue datapoints and curve), the exact thermodynamic gap predicted by de Gier & Essler (thick, black), and the gap predicted by domain wall theory (dashed, green). We have $\alpha = 0.3$ and hence $\beta_c \approx 0.57$ (marked). Note, for $\beta < 0.5$ domain wall theory and de Gier-Essler coincide.

the overestimation of the finite-size gaps is reflected in an overestimation of the thermodynamic gap. This is unambiguous for $\beta < 0.59$ yet given the discussion above it must also be true for $\beta \geq 0.59$, though for larger β the discrepancy is expected to be smaller. Notwithstanding these inaccuracies, it appears that the simulation results corroborate the de Gier-Essler result, even if only to the extent that they convincingly rule out the domain wall theory prediction of a transition at $\beta = 0.5$. The simulation results, though failing to show a sharp transition and marred by non-systematic, statistical errors, could at first sight be read as suggesting a transition point somewhat larger than β_c . If

one imagines the extrapolated gap result decreased in value by a magnitude that grows as $\alpha \rightarrow \beta$ however, that would sharpen the descent of the gap to the plateau value it undoubtedly reaches for $\beta > 0.64$ (see figure 4.5), and hence move the then more readily identifiable transition point closer to β_c .

It is an unfortunate feature, inherent to the chosen method of analysing relaxation measurements, that there are essentially uncontrolled inaccuracies that lead to an overestimation of the finite size gap estimates and hence to a distorted prediction of the thermodynamic gap behaviour. Nonetheless I do advocate the method's utility, for several reasons.

First, it is thoroughly justified on theoretical grounds to focus as it does on the asymptotic regime, where the gap dominates. As a consequence, the only factor limiting the accuracy that can be obtained is the computational resources one is willing and able to expend on obtaining good quality data for late times.

Second, it is simple in its application, provided one implements a systematic scheme (such as the one discussed in section 4.2.4) to balance quality of data with proximity to the purely exponential regime. In particular, it avoids the ambiguities and difficulties associated with more complicated fitting methods described in section 4.2.3 and section 4.2.2.

Third, and finally, although the results of figure 4.11 are somewhat unsatisfying and inconclusive it is doubtful, for the reasons given in section 4.2.3, that non-asymptotic methods such as employed in [141] would fare better.

Reasonable results were obtained for a model whose spectral structure and hence (lack of) separation of timescales is imposed by the driving at the boundaries. This suggests asymptotic methods like the one used here may well find more use in analysing Monte Carlo data from models of other nonequilibrium systems. Nonetheless the method would benefit from improvement. Most obviously, finite-time data could be used to extrapolate (in a functionally appropriate way) to the asymptotic value to yield more accurate finite-size values than currently result from effecting an average of the local slope over available late-time data by means of a single exponential fit.

4.4 Gap from occupancy fluctuations in the steady state

In section 4.3 I showed how the TASEP gap may be measured by observing how the fraction of occupied sites relaxed from an initial nonstationary situation. There, $\Phi(t)$ was computed as an ensemble average, i.e. as an average over runs with the same average initial conditions.

In this section however I show to what extent the dynamical transition is manifested in the steady state fluctuations by considering the *autocorrelation* of the total occupancy, which amounts to the choice of observable

$$\mathcal{O}(t) = N(t_0 + t)N(t_0) \quad (4.45)$$

for some t_0 . This quantity has the advantage that it can be sampled in a single long run once the simulation has reached the steady state (rather than requiring the simulation to be rerun from the same initial conditions every time). The corresponding relaxation function $\Phi(t)$, which we will call $\chi(t)$, is then given by

$$\chi(t) = \frac{\langle N(t_0 + t)N(t_0) \rangle - \langle N(t_0) \rangle^2}{\langle N^2(t_0) \rangle - \langle N(t_0) \rangle^2}, \quad (4.46)$$

where here all averages are taken in the steady state.

In practice, the initial condition used in simulation (all sites initialised with the bulk stationary density), though constituting a state relatively close to the stationary profile, still projects significantly onto higher modes. The contribution of these modes to the dynamics becomes small only after some relaxation time t' . Thus the value of t_0 should be taken larger than t' . As noted above, we can use multiple values of $t_0 > t'$ from the same simulation run to obtain a better sample of $\chi(t)$ without needing to restart the simulation and wait for the initial relaxation to take place.

This scenario was simulated, again using the continuous-time Monte Carlo method of section 4.1.2, for $\alpha = 0.3$ and for a range of system sizes. It became apparent that finite size gap estimates were indistinguishable from the results of section 4.3 for the same system size and value of β (see figure 4.12). If anything, the gap estimates obtained by measuring the autocorrelation of the total occupancy are somewhat noisier after an equivalent investment of computational time.

It seemed uninformative therefore to invest in obtaining a large dataset extending to numerous larger system sizes, as was done in section 4.3, since the extrapolation is




Figure 4.12: As a function of β and with $\alpha = 0.3$: the gap determined by measuring $\Phi_\rho(t)$ – the relaxation of an initial overall density excess or deficit (solid curves, same data as figure 4.6), and gap values obtained from measuring $\chi(t)$ – the fluctuations of the total occupancy in the steady state (dashed curves). Successively higher lying curves connecting finite L datapoints are for $L = 10, 20$, and 30 . The $L \rightarrow \infty$ de Gier-Essler exact gap is shown (thick, black) for comparison, as is the exact gap for $L = 10$ from numerical diagonalisation (black) and the gap according to domain wall theory (green, dashed). The dynamical transition is at $\beta_c \approx 0.57$ (marked).

expected to suffer from the same shortcomings (though with additional noise) and not to reveal anything new. It is reassuring however that identical results may be obtained by means of these different measures, and it supports the notion that lack of agreement with the de Gier-Essler gap is due to the discussed shortcomings in the data analysis, rather than due to the choice of observables.

4.5 Gap from relaxation of a second class particle

The success of domain wall theory in predicting the exact gap for $\beta < 0.5$, and the knowledge that extensive literature [15, 96, 142, 143] identifies a second class particle as tracking the position of a shock / domain wall, even in a finite system [144], formed the inspiration for a third measure of the relaxational dynamics. In the hope of gaining some insight into the nature of the dynamical transition, we consider out of curiosity the relaxation of the position of a second class particle.

A second class particle (scp) behaves like a hole with respect to ordinary (first class) particles, and like a first class particle with respect to holes. Representing first class particles by '1', second class particles by '2', and superimposing the relevant Poisson transition rate over the arrow indicating each move, the dynamics are then as follows:

$$12 \xrightarrow{1} 21 \tag{4.47}$$

$$21 \xrightarrow{0} 12 \tag{4.48}$$

$$20 \xrightarrow{1} 02 \tag{4.49}$$

$$02 \xrightarrow{0} 20 \tag{4.50}$$

On the infinite lattice such a second class particle has been proven [96] to track the position of a shock separating a density ρ_- to the left and ρ_+ to the right, with $\rho_- < \rho_+$. The suggestion in [144] to identify the position of a shock in a single realisation of the open system by injecting a single second class particle motivates our third choice of observable and we measure

$$\Phi_{\text{scp}}(t) \equiv \langle x_{\text{scp}}(t) \rangle - \langle x_{\text{scp}}^{\infty} \rangle, \tag{4.51}$$

with x_{scp}^{∞} the stationary position of the second class particle, taken to be its position at the longest time reached in simulation (I have not attempted to calculate the stationary position analytically). The averages indicated by angular brackets are calculated over an ensemble of runs.

The simulations are initialised with the second class particle occupying the leftmost site (site 1) on the lattice. If at any time it happens to be displaced and ejected from the system at the left boundary by a first class particle entering the system, this is recorded and it is reinserted at the next available opportunity when site 1 is empty. Eventually the scp travels to the right boundary. Once it reaches the right boundary and is occupying site L , it is not allowed to exit the system through the right boundary (i.e. reflecting boundaries), and stays on site L until a first class particle displaces it, forcing it to site $L - 1$. The motion of the scp following these dynamical rules and boundary conditions is illustrated in figure 4.13.

Measuring $\Phi_{\text{scp}}(t)$ according to (4.51) yields an exponential decay as the value at the longest time recorded is approached. As with the other measures, we can extract the asymptotic exponent by the procedure of section 4.2.4. This yields the finite-size gap estimates shown in figure 4.14

Surprisingly, the finite-size gap estimates given by the relaxation of the scp for large system sizes ($L = 50, 150$) agree extremely well with the same results obtained by




Figure 4.13: Position as a function of time of a second class particle (scp) initially injected at the left boundary of a system of size $L = 150$ with $\alpha = 0.3$ and $\beta = 0.6$. The smooth curve represents an ensemble average over 10^4 runs, the irregular curve the time series of one such run. On the y-axis, -1 corresponds to the particle being at the left boundary, 0 to it being at the right boundary.

measuring the relaxation of an initial density excess or deficit in section 4.3. The scp data overestimate the true finite size gap by the same amount as the initial relaxation data, which is perhaps not surprising given the same method was used to extract the gap estimates.

The close agreement of these sets of data does however suggest that the motion of the second class particle constitutes a remarkably effective dimensional reduction of the long-time dynamics. Moreover, tracking the scp in simulation shows (notwithstanding the systematic overestimation) that it follows the exact gap behaviour not just for $\beta < 0.5$, but also above.

This is in contrast to domain wall theory, where the scp is assumed to separate two particular types of domain for $\beta < 0.5$, above which one of the domain changes character and hence the dynamics of the scp. Instead, these simulations appear to capture the true motion of the scp including, crucially, fluctuations in its position as it approaches the right boundary.




Figure 4.14: As a function of β and with $\alpha = 0.3$: finite size gap estimates obtained from Monte Carlo measurements of Φ_{scp} – the relaxation of a second class particle (scp) to its stationary position (unconnected datapoints, large \bullet for $L = 50$ and large \blacksquare for $L = 150$). Finite size gap estimates obtained from Φ_{ρ} – relaxation of overall density (same data as figure 4.7) are shown for comparison. Colours of scp gaps coincide with colours of relaxation gaps for same system size.

Speculatively, these effects might form the much sought after basis of a justification for modifying domain wall theory in the manner first proposed in [18]. In other words, the fluctuations – due to its interaction with the boundary – in the position of the scp and arguably therefore in the position of an associated domain wall, could somehow be responsible for shifting the dynamical transition away from $\beta = 0.5$ (as predicted by domain wall theory) and towards $\beta_c(\alpha)$ calculated by de Gier & Essler. It is not apparent however how exactly this effect could be quantified to reproduce the relationship $\beta_c(\alpha)$ (equation 2.55), starting from the microscopic dynamics.

4.6 Discussion

In this chapter I have shown that Monte Carlo simulation gives access to the TASEP gap. First, in section 4.1, I discussed simulation methodology and detailed the use of an efficient dynamical Monte Carlo algorithm. Then, in section 4.2, I argued how one might best obtain finite-size gap estimates from simulation data. In sections 4.3, 4.4 and 4.5 I demonstrate that finite size gap estimates obtained from measurements of three different

quantities all agree, and that the resulting extrapolated thermodynamic gap provides evidence in favour of the exact gap computed by de Gier and Essler rather than that predicted by domain wall theory. I also present an argument for the observed discrepancy between simulation results and the exact gap.

As was shown in section 4.5, it appears possible to determine the gap behaviour by measuring the relaxation of a single second class particle. This apparent effective reduction of the long-time dynamics onto a single co-ordinate is rather remarkable. In the hope of shedding some light on the origin of this connection, I conclude this chapter by reporting in section 4.6.1 on an approximate calculation of the gap associated with the motion of a second class particle. The result bears qualitative resemblance to the exact gap and exhibits a dynamical transition. The exercise also informs us that the Markovian approximation made in the analysis leads to a transition at $\beta = \frac{1}{2}$, and furthermore that the gap, though constant, is grossly overestimated for $\beta > \frac{1}{2}$.

4.6.1 Approximate theoretical second class particle gap

The second class particle interacts with the rest of the system, thereby inducing temporal correlations in the local environment encountered by the particle. This means that the dynamics of the second class particle are history-dependent. However, one can try to neglect these interactions and assume that the local environment encountered by the particle from one point in time to the next is uncorrelated and equal to the steady state of the TASEP in the particle's absence. It becomes possible under this Markovian approximation to write down an exact transition matrix for the motion of the second class particle, and to find the associated gap.

The first step is to calculate the entries of the transition matrix, which involves using the matrix product ansatz to calculate exact stationary probabilities for all the possible local environments (neighbour configurations) that the second class particle may encounter. This matrix is then diagonalised to find the gap. Finite-size gaps for a few system sizes $L \leq 80$ are shown in figure 4.15.

To obtain a prediction of the thermodynamic gap behaviour based on the approximate motion of the second class particle, one wants to extrapolate from the finite-size results to $L = \infty$. It was computationally feasible to obtain numerous finite-size gaps from diagonalisation for systems up to $L = 300$ starting from $L = 20$ and increasing in increments of 10. In contrast to sections 4.3 and 4.4, extrapolation was performed using the Bulirsch-Stoer (BST) algorithm ([60] and see section 2.7) instead of assuming *a priori* asymptotic




Figure 4.15: As a function of β with $\alpha = 0.1$: exact de Gier-Essler gap (thick, black), domain wall theory gap (dashed, green), and gap from the motion of a second class particle in the steady state under a Markovian approximation (orange, thin) for $L = 20, 40, 60,$ and 80 (successively higher lying curves are for increasing L). The dynamical transition is at $\beta_c \approx 0.68$ (marked).

convergence according to the de Gier-Essler result and fitting $a+bL^{-2}+cL^{-3}+\dots$ to finite-size results. The predicted thermodynamic gap given by BST extrapolation is shown in figure 4.16.

The finite-size scp gaps grossly overestimate the exact gap. In particular, they do so by considerably more than the gap predicted by domain wall theory, and for all $\beta > \alpha$ rather than just for $\beta > 0.5$. This is no great surprise, given the gross nature of the Markovian approximation. It is in fact surprising that the scp gap qualitatively reproduces the exact gap behaviour to the extent that it does.

Interestingly, the thermodynamic scp gap (figure 4.16) unambiguously displays a transition at $\beta = 0.5$, just like domain wall theory and the mean-field result of chapter 3. Comparison with scp gap results from simulation from the section 4.5 suggests that the effect of neglecting correlations is to shift the transition to $\beta = 0.5$. This strengthens the argument, made at the end of that section, that somehow taking into account correlations, in particular at the boundary, may shift the transition away from $\beta = 0.5$ and towards $\beta_c(\alpha)$, and that the nature of these correlations would give insight into the dynamical transition itself.




Figure 4.16: As a function of β with $\alpha = 0.1$: exact de Gier-Essler gap (thick, black), domain wall theory gap (dashed, green), and thermodynamic gap from the motion of a second class particle in the steady state under a Markovian approximation (blue, \bullet), given by Bulirsch-Stoer extrapolation of finite-size gaps. The dynamical transition is at $\beta_c \approx 0.68$ (marked).

Chapter 5

Gap from numerical density-matrix-based renormalisation

In chapter 4 we saw that Monte Carlo measurements of relaxation functions, though informative, provided us with estimates of the thermodynamic gap behaviour for the TASEP that narrowly fell short of conclusively ruling out the domain wall theory prediction in favour of the dynamical transition predicted by de Gier and Essler.

In this chapter I present results for the gap that were arrived at using a numerical real-space renormalisation scheme known as the density matrix renormalisation group (DMRG) method [145–147]. The DMRG method has successfully been applied to calculate to high accuracy the critical points and critical exponents in reaction-diffusion models [148]. It has also previously been used to determine the largest relaxation time and hence the gap for the TASEP in [149], though there the emphasis was on determining the dynamical exponents rather than the functional dependence on boundary parameters.

Section 5.1 provides an introduction to the DMRG method, i.e. to density-matrix-based renormalisation, which those who are familiar with the method may wish to skip but which otherwise provides essential conceptual grounding for what follows.


Section 5.2 focuses on the TASEP. In section 5.2.1 I go into some detail about how the DMRG method may be used to compute the TASEP gap, introducing notation but not yet providing a concrete algorithm. In the technical exposition of section 5.2.2 I detail exactly how to compute the TASEP gap using the so-called *infinite system* algorithm, which forms

the core of the DMRG method. My aim in so doing is twofold: first, to document, for the sake of reproducibility, the manner in which the gap results presented in section 5.3 were obtained, and second, to provide the potential first-time DMRG practitioner with a concrete and detailed example that demonstrates unambiguously how to construct a working implementation.

Finally, in section 5.3 I present finite-size gap estimates from density-matrix-based renormalisation and show that extrapolations for $L \rightarrow \infty$ confirm the de Gier-Essler result. In section 5.4 I discuss issues arising in the implementation of the algorithm, and suggest possible modifications that would improve the results as well as give additional insight into the long-time dynamics.

5.1 The density matrix renormalization group method

The DMRG method began life as an investigation into why numerical real-space renormalization-group (RG) schemes yield poor results when applied to quantum lattice models [150] in a straightforward manner. The essence of such a scheme is as follows.



One begins by writing down the Hamiltonian H for a block of sites A that is sufficiently small so that H may be diagonalised numerically. During each renormalisation step that follows one constructs the Hamiltonian H' for a new block A' consisting of two of the old blocks ($A' = AA$). H' is then renormalised (i.e. its dimensionality is reduced) by changing its basis to m of its own eigenvectors. Finally, once H' is renormalised the new block is considered as an old block ($A = A'$), we set $H = H'$, and the renormalisation step begins anew. In this way the system continues to grow without the dimensionality of H ever exceeding m^2 , so that it can always be diagonalised.

The choice of which eigenvectors to use as a reducing basis depends on the state of interest, also known as the *target state*. If the aim is to characterise the thermodynamic behaviour of the ground state or the first few excited states then the m states with lowest


energies are used. However, it is also possible to target other single states or small clusters of states by including the m eigenstates closest in energy to that of the state(s) of interest.

The underlying idea therefore is that of an approximation of the true Hamiltonian by an effective Hamiltonian, operating in a reduced Hilbert space. This is achieved by a truncation, during each growth-renormalisation cycle, that aims to minimise the error made in the approximation by retaining the m basis states deemed most important to accurately represent the target state for larger system sizes. Clearly the error will be smaller for larger m , since more more states are then retained, and whereas m is often taken to be constant, in some schemes (e.g. [151]) it is allowed to vary between one renormalisation step and the next.

However, White and Noack found in [150] that this procedure failed to give correct ground state energies even for a very simple one-dimensional single-particle-in-a-box model. The solution was found to lie in varying the boundary conditions applied to the outer (i.e. non-adjoining) edges of the two A blocks, diagonalising H' for each choice, and in each case keeping as part of the reducing basis the $m/4$ lowest-energy eigenvectors. With four possible combinations of boundary conditions this leads once again to a reducing basis of size m . This modification resulted in excellent agreement of the ground state energy and of a few higher-lying energies with the exact spectrum.

The key insight of White therefore was that in order for the reducing basis to be able to accurately represent low-lying eigenstates of the further enlarged four-block $A'A'$ system, it must be chosen from amongst the low-energy eigenstates generated by imposing all possible choices of external boundary conditions on the two-block system.

Instead of manually choosing boundary conditions on a case-by-case basis and in an ad hoc fashion, thereby possibly risking an incomplete representation of the target state, it was shown already in [150] that there is a more elegant alternative. It consists of surrounding each of the two single A blocks with additional B blocks, before joining all the blocks together to form a so-called *superblock*, described by the Hamiltonian H^{sup} . Soon after [150] it became clear [152] that for many-particle models the most appropriate choice is in fact for each of the additional B blocks to be a single site, and that for open boundary systems each original A block requires such an additional site on one side only, namely the one facing inwards to the rest of the lattice:



We reassign H' to stand in this new approach for the Hamiltonian that describes each of the two block-plus-site combinations $A\bullet$ and $\bullet A$, which are therefore implicitly assumed to be identical (barring reflection). During the next renormalisation step each of these combinations will constitute a new block A' , which will in turn be combined with a single site. In order to then be able to write down a Hamiltonian that remains diagonalisable, we want to renormalise H' during the current cycle. To do so we diagonalise H^{sup} , and obtain the reducing basis for H' by projecting the m lowest-energy eigenstates of H^{sup} onto the parts of the basis of H^{sup} that are associated with the block-plus-site combination relevant to H' (either $A\bullet$ or $\bullet A$). The basis of H' is then changed to this reducing basis, truncating the Hilbert space.

It is worth giving some insight into the rationale behind the superblock approach. The additional sites are conceived of as mimicking the environment outside the original A blocks. The eigenstates of H^{sup} can therefore be thought of as incorporating all the possible states that the environment may be in, or equivalently, as having been generated by the application of the full range of possible boundary conditions on the original A blocks. Low-lying eigenstates of H^{sup} correspond to low-energy states of the entire lattice, i.e. of the superblock $A\bullet\bullet A$. In particular, they do so regardless of what the local states are within the A blocks, and may include states that might otherwise have been missed in the ad hoc application of an incomplete set of boundary conditions. That is why these states are excellent candidates to represent low-energy eigenstates of the further enlarged system that appears in the the next renormalisation step, and are chosen as the reducing basis for H' .

There is however an important complication that arises in models allowing for more than a single particle on the lattice. Whereas single-particle eigenstates of H^{sup} each project onto a single state of the original A blocks, so that choosing m eigenstates of H^{sup} immediately gives m states that can be used to reduce H' , this is not the case for eigenstates of H^{sup} for a many-particle model. Then a complete set of block states is in principle needed to represent any given superblock state, prompting the question of which block states should be chosen as part of the reducing basis.

The answer is provided in [152], where White showed that the argument regarding the coupling of the superblock to the environment may be reformulated in terms of density matrices and thereby be made simultaneously more rigorous as well as more flexible, lending the DMRG method its name.

To understand the role of the density matrices, imagine that we wish to target any particular state $|\psi\rangle$ of the system. This means we want during every growth-renormalisation cycle to renormalise H' , which describes each block-plus-site combination, by choosing as a reducing basis those block-plus-site states that are in some sense optimal given that the superblock as a whole is in the state $|\psi\rangle$. This state itself can be represented as

$$|\psi\rangle = \sum_{i,j} \psi_{ij} |i\rangle |j\rangle, \quad (5.1)$$

where the $|i\rangle$ are the states associated with one of the block-plus-site combinations (i.e. one half of the superblock), and the $|j\rangle$ the states of the rest of the superblock. If one computes the reduced density matrix [153], for example for the left $A\bullet$ combination, according to

$$\rho_L = \sum_j \psi_{ij} \psi_{i'j}, \quad (5.2)$$

then the m eigenvectors of ρ_L with the largest eigenvalues are an optimal choice for a reducing basis, since they are the *most likely* states of the left half of the superblock when the superblock as a whole is in the state $|\psi\rangle$ [152, 154]. This choice of reducing basis can additionally be justified by showing [152, 154] that it minimizes the error made in approximating the true $|\psi\rangle$, and maximises the preservation of entanglement between the block and the rest of the lattice [146].

For further conceptual clarification and the earliest detailed exposition of the DMRG algorithms in their mature form, presented in the context of Heisenberg spin chains, see White's [154]. For details on the DMRG method that go beyond the early papers already cited, see the review by White [145] and its contemporary collection [147], as well as the more recent excellent review [146], and the wildly informative [155].

Of particular interest for the work presented in this thesis are [149] for previous DMRG results for the TASEP gap, and [148] for methodological issues that arise in applying DMRG to the non-Hermitian transition matrices that arise in classical stochastic reaction-diffusion models.

5.2 DMRG for the TASEP

The DMRG method as described in section 5.1 refers to the so-called *infinite system method* (ISM), which is the primary DMRG method and used to generate states that approach their true thermodynamic behaviour after many growth-renormalisation cycles.

The *finite system method* (FSM), on the other hand, is a computationally intensive refinement that generates states approaching their true finite-size behaviour for a fixed system size. An FSM algorithm iteratively refines the state and operators for a fixed system size that were originally generated by an ISM algorithm. Instead of growing the system as a whole, one side of the system (the left, say) is grown while the other side is shrunk, keeping the system size fixed. This is repeated until the right side can shrink no longer, after which it grows at the expense of the left side. This continues until the left side too can shrink no longer, at which point the process reverses. One iteration of the algorithm is complete when the sides are of equal size, as they would be at each stage of an ISM algorithm. During each grow-and-shrink step the operators for the growing block are updated and renormalised using a similar density-matrix-based procedure as in the ISM. For details see [145, 147, 148, 154].

I discuss the relevance of the finite system method for the TASEP gap in section 5.3, but until then focus exclusively on the infinite system method. In section 5.2.1 I outline a procedure for computing the TASEP gap using the infinite system method, which seems appropriate given that we want to determine the thermodynamic gap behaviour.

5.2.1 Infinite system method DMRG for the TASEP

In this section and in section 5.2.2 one should in fact read ‘transition matrix’ at any mention of a Hamiltonian – the notational convenience allows for ease of discussion between sections and when dealing with DMRG reference works.

To begin, we diagonalise numerically the Hamiltonian H (i.e. the transition matrix H) given by equation 2.9 for a sufficiently small system. Each iteration of the infinite system algorithm that is subsequently applied incorporates

A growth step, where the system is enlarged by the insertion at the centre of two new lattice sites. As a result the underlying state space grows by a factor 2^2 , as does the dimensionality of H .

A renormalisation step, where the dimensionality of the state space is reduced by projecting H for the enlarged system onto a suitably chosen basis, $\{|\omega_i\rangle\}_{i=1}^{2m}$.

In this way H remains of fixed dimension $2m$ and hence diagonalisable as the effective system size grows. The gap is computed at each iteration, for as long as the algorithm remains stable.

As stated above this constitutes the so-called infinite system method, which aims to reach sufficiently large effective system size to exhibit the thermodynamic behaviour, removing the need for extrapolation from finite size data. In practice the algorithm may become unstable before agreement to the desired precision with the thermodynamic behaviour is reached, in which case extrapolation is required. As we shall see in the results section, section 5.3, this is what I pursued for the TASEP gap.

A ‘suitable’ choice for the reducing basis $\{|\omega_i\rangle\}_{i=1}^{2m}$ is one that ensures that the eigenstates of the repeatedly renormalized Hamiltonian remain adequate representations of both ground and gap states. This is dealt with in the following way.

The system is treated as consisting of a left block of sites A_L , governed by the Hamiltonian H_L , and a right block of equal size, A_R , governed by H_R . The Hamiltonian H for the entire system, can be composed by combining H_L , H_R , and an interaction term that couples the two blocks.

H is non-Hermitian, and is diagonalised using the Arnoldi method [156] as implemented in the ARPACK package [59] and accessed using *Mathematica* to find the gap as well as the ground state ($\langle\phi_0|$ and $|\psi_0\rangle$) and first excited state ($\langle\phi_1|$ and $|\psi_1\rangle$) respectively left and right eigenvectors. Before proceeding, we note that to gain accuracy in the gap as well as in the gap eigenvectors $\langle\phi_1|$ and $|\psi_1\rangle$, we follow [157] by first finding $\langle\phi_0|$ and $|\psi_0\rangle$ before shifting the ground state by an arbitrary but large amount $\Delta > 0$ down into the spectrum by performing

$$H \rightarrow H + \Delta|\psi_0\rangle\langle\phi_0|. \quad (5.3)$$

The gap as well as the gap eigenvectors are given by applying Arnoldi diagonalisation to find the ground state of the shifted Hamiltonian. It can be observed when diagonalising for small system sizes that this trick reproduces the accuracy of drastically less efficient routines, yet still allows the speed of the Arnoldi method to be used.

The system is enlarged by adding one site to each block, causing a twofold growth in the Hilbert space associated with each side. H_L and H_R grow accordingly, and in order to perform the necessary renormalisation we turn to the density matrices.

One forms reduced density matrices for the left and right sides based on the ground state and first excited state eigenvectors (the target states):

$$\rho_L = \frac{1}{4} \text{Tr}_R \{ |\psi_0\rangle\langle\psi_0| + |\phi_0\rangle\langle\phi_0| + |\psi_1\rangle\langle\psi_1| + |\phi_1\rangle\langle\phi_1| \}, \quad (5.4)$$

$$\rho_R = \frac{1}{4} \text{Tr}_L \{ |\psi_0\rangle\langle\psi_0| + |\phi_0\rangle\langle\phi_0| + |\psi_1\rangle\langle\psi_1| + |\phi_1\rangle\langle\phi_1| \}. \quad (5.5)$$

Here the trace Tr_L over the left side of the system for example sums over those parts of the eigenvector basis that are associated with that side. The symmetric combination of eigenvectors in (5.4) and (5.5), also used in [149], was found in [148] to be optimal for non-Hermitian matrices if one wants to find the gap. Even though we only care about the gap state, the ground state is included as a target state because it stabilises the algorithm, which might otherwise fail if the gap and ground state eigenvalues lie too close together.

As discussed in section 5.1, the reduced density matrices provide us with the basis to renormalise the Hamiltonian. To be specific, ρ_L and ρ_R are diagonalised using the appropriate dense matrix diagonalisation routine from the LAPACK package [158], accessed using *Mathematica*. For a given side of the system, the left, say, the renormalisation is accomplished by projecting H_L for the enlarged left block onto the eigenvectors of ρ_L with the m largest eigenvalues ω_i .

One argument supporting this choice of $\{|\omega_i^L\rangle\}_{i=1}^{2m}$ as basis set is, as explained in section 5.1, that for some general $\rho_L = \text{Tr}_R\{|\psi\rangle\langle\psi|\}$, the eigenvalues ω_i represent the probability that the left side of the system is in state $|\omega_i^L\rangle$ given that the system as a whole is in state $|\psi\rangle$. Hence our choice (5.4) amounts to projecting onto the m states the left side is most likely to be in when the system as a whole is equally likely to be in the ground as in the first excited state.

One subtlety has been left unspecified, namely that the rightmost site in the left block as well as the leftmost site in the right block are always expressed in a physical basis. This allows us to write the terms connecting H_L and H_R to the inserted sites as well as the coupling term between the blocks when constructing H prior to diagonalisation. It is also the reason why, though we keep only m eigenvectors of ρ_L and ρ_R , we project onto a $2m$ -dimensional basis set: the factor of two corresponds to the physical basis of the most recently inserted site in each block.

This provides sufficient background to the exposition in 5.2.2 of a concrete step-by-step recipe for computing the TASEP gap using the infinite system method.

5.2.2 An infinite system method algorithm for the TASEP gap

In what follows we make use of three elementary transition matrices. In the basis

$$\begin{pmatrix} - \rightarrow - & \bullet \rightarrow - \\ - \rightarrow \bullet & \bullet \rightarrow \bullet \end{pmatrix} \quad (5.6)$$

we have

$$h_l = \begin{pmatrix} -\alpha & 0 \\ \alpha & 0 \end{pmatrix} \quad \text{and} \quad h_r = \begin{pmatrix} 0 & \beta \\ 0 & -\beta \end{pmatrix} \quad (5.7)$$

for the left and right boundary sites respectively, and in the basis

$$\begin{pmatrix} - \rightarrow - & \bullet \rightarrow - & \bullet \rightarrow - & \bullet \bullet \rightarrow - \\ - \rightarrow \bullet & \bullet \rightarrow \bullet & \bullet \rightarrow \bullet & \bullet \bullet \rightarrow \bullet \\ - \rightarrow \bullet & \bullet \rightarrow \bullet & \bullet \rightarrow \bullet & \bullet \bullet \rightarrow \bullet \\ - \rightarrow \bullet \bullet & \bullet \rightarrow \bullet \bullet & \bullet \rightarrow \bullet \bullet & \bullet \bullet \rightarrow \bullet \bullet \end{pmatrix} \quad (5.8)$$

we have

$$h_b = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (5.9)$$

for pairs of bulk sites. The algorithm then proceeds as follows.

1. *Initialise block Hamiltonians* H_L, H_R

The algorithm begins by considering the partitioning of our system into two blocks, each of size ℓ :

$$A_L^\ell \ A_R^\ell \quad (5.10)$$

Using the 2×2 identity matrix \mathcal{I} to augment the dimensionality of the elementary matrices, the Hamiltonian for the left block can be written as

$$H_L = h_l \otimes \mathcal{I}^{\otimes \ell-1} + \sum_{k=1}^{\ell-1} \mathcal{I}^{\otimes k-1} \otimes h_b \otimes \mathcal{I}^{\otimes \ell-k-1}, \quad (5.11)$$

and for the right block, as

$$H_R = \sum_{k=1}^{\ell-1} \mathcal{I}^{\otimes k-1} \otimes h_b \otimes \mathcal{I}^{\otimes \ell-k-1} + \mathcal{I}^{\otimes \ell-1} \otimes h_r. \quad (5.12)$$

2. Construct and diagonalise the Hamiltonian for the entire system, i.e. the superblock, H , to find the gap and the target states $|\psi_0\rangle$, $\langle\phi_0|$, $|\psi_1\rangle$, $\langle\phi_1|$

The Hamiltonian for the entire superblock, H , is given by

$$H = H_L \otimes \mathcal{I}^{\otimes s_R} + \mathcal{I}^{\otimes s_L-1} \otimes h_b \otimes \mathcal{I}^{\otimes s_R-1} + \mathcal{I}^{\otimes s_L} \otimes H_R. \quad (5.13)$$

For the first iteration of the algorithm H is in a physical basis hence $\dim(H_L) = \dim(H_R) = 2^\ell$ and $s_L = s_R = \ell$, with ℓ chosen small enough so that H can be diagonalised exactly. In all later iterations the basis used to describe each side has been reduced to $\dim(H_L) = \dim(H_R) = 2m < 2^\ell$, which means that $s_L = s_R = \log_2 2m$.

Now we diagonalise H , first using the Arnoldi algorithm to find the left and right ground state eigenvectors $\langle\phi_0|$, $|\psi_0\rangle$, which are the first two of our four *target states*¹. As described in section 5.2.1, we do the same for the shifted Hamiltonian $H + \Delta|\psi_0\rangle\langle\phi_0|$ to accurately find the gap and associated eigenvectors $|\psi_1\rangle$ and $\langle\phi_1|$.

3. Form reduced density matrices ρ_L , ρ_R

We want to form the reduced density matrices

$$\rho_L = \frac{1}{4} \text{Tr}_R \{ |\psi_0\rangle\langle\psi_0| + |\phi_0\rangle\langle\phi_0| + |\psi_1\rangle\langle\psi_1| + |\phi_1\rangle\langle\phi_1| \}, \quad (5.14)$$

$$\rho_R = \frac{1}{4} \text{Tr}_L \{ |\psi_0\rangle\langle\psi_0| + |\phi_0\rangle\langle\phi_0| + |\psi_1\rangle\langle\psi_1| + |\phi_1\rangle\langle\phi_1| \}. \quad (5.15)$$

First consider a single target state $|\psi\rangle$, being an eigenvector of H . If $|i\rangle$ labels basis states of the left side of the superblock and $|j\rangle$ basis states of the right side, $|\psi\rangle$ can be represented, as in equation 5.1, as

$$|\psi\rangle = \sum_{i,j} \psi_{ij} |i\rangle|j\rangle, \quad (5.16)$$

where

$$\psi_{i'j'} = \langle j' | \langle i' | \psi \rangle. \quad (5.17)$$

¹Only for the first iteration of the algorithm is $\langle\phi_0|$ trivially known to be $(1, 1, 1, \dots, 1)$

The reduced density matrix for the left side, $\rho_L = \text{Tr}_R\{|\psi\rangle\langle\psi|\}$, is then defined by

$$\rho_{Lii'} = \sum_j \psi_{ij} \psi_{i'j}, \quad (5.18)$$

and for the right side,

$$\rho_{Rjj'} = \sum_i \psi_{ij} \psi_{ij'}, \quad (5.19)$$

We can replace the notation ψ_{ij} by a simple labelling ψ_k of the elements of $|\psi\rangle$ as a column vector, where the basis is implicitly given by the order of combination of bases in the composition of H above. Concretely this means that for the single target state $|\psi\rangle$,

$$\rho_L = \begin{pmatrix} \sum_k^{s_R} \psi_k \psi_k & \sum_k^{s_R} \psi_k \psi_{s_R+k} & \cdots & \sum_k^{s_R} \psi_k \psi_{(s_L-1)s_R+k} \\ \sum_k^{s_R} \psi_{s_R+k} \psi_k & \sum_k^{s_R} \psi_{s_R+k} \psi_{s_R+k} & \cdots & \sum_k^{s_R} \psi_{s_R+k} \psi_{(s_L-1)s_R+k} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_k^{s_R} \psi_{(s_L-1)s_R+k} \psi_k & \sum_k^{s_R} \psi_{(s_L-1)s_R+k} \psi_{s_R+k} & \cdots & \sum_k^{s_R} \psi_{(s_L-1)s_R+k} \psi_{(s_L-1)s_R+k} \end{pmatrix} \quad (5.20)$$

and

$$\rho_R = \begin{pmatrix} \sum_k^{s_L} \psi_{(k-1)s_R+1} \psi_{(k-1)s_R+1} & \sum_k^{s_L} \psi_{(k-1)s_R+1} \psi_{(k-1)s_R+2} & \cdots & \sum_k^{s_L} \psi_{(k-1)s_R+1} \psi_{ks_R} \\ \sum_k^{s_L} \psi_{(k-1)s_R+2} \psi_{(k-1)s_R+1} & \sum_k^{s_L} \psi_{(k-1)s_R+2} \psi_{(k-1)s_R+2} & \cdots & \sum_k^{s_L} \psi_{(k-1)s_R+2} \psi_{ks_R} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_k^{s_L} \psi_{ks_R} \psi_{(k-1)s_R+1} & \sum_k^{s_L} \psi_{ks_R} \psi_{(k-1)s_R+2} & \cdots & \sum_k^{s_L} \psi_{ks_R} \psi_{ks_R} \end{pmatrix}. \quad (5.21)$$

We now form ρ_L and ρ_R using all four target states according to (5.14), (5.15).

4. *Diagonalise ρ_L , ρ_R and form reducing projectors P_L , P_R*

Diagonalise ρ_L and ρ_R using a dense matrix diagonalisation routine. Form the matrices O_L and O_R that have as their columns the m eigenvectors with largest eigenvalues of ρ_L and ρ_R respectively. Then the operators that project onto the reducing basis, P_L and P_R are given by

$$P_L = O_L \otimes \mathcal{I} \quad (5.22)$$

$$P_R = \mathcal{I} \otimes O_R \quad (5.23)$$

5. *Enlarge system and construct new Hamiltonians \tilde{H}_L , \tilde{H}_R*

We now insert two central lattice sites to give the enlarged superblock of size $L' = 2\ell + 2$:

$$A_L \begin{matrix} \ell \\ \bar{1} \end{matrix} \begin{matrix} \ell \\ \bar{2} \end{matrix} A_R \quad (5.24)$$

The left side of the enlarged system is governed by the effective² Hamiltonian

$$\tilde{H}_L = H_L \otimes \mathcal{I}^{\otimes 2} + \mathcal{I}^{\otimes s_L-1} \otimes h_b, \quad (5.25)$$

and the right side by

$$\tilde{H}_R = h_b \otimes \mathcal{I}^{\otimes s_R-1} + \mathcal{I}^{\otimes 2} \otimes H_R. \quad (5.26)$$

6. Renormalise \tilde{H}_L, \tilde{H}_R

Renormalise the partial Hamiltonians for the enlarged system by transforming them to the appropriate reduced density matrix eigenbasis using the reduced projection operators:

$$H_L = P_L^T \tilde{H}_L P_L \quad (5.27)$$

$$H_R = P_R^T \tilde{H}_R P_R \quad (5.28)$$

7. Return to step 2

5.3 Infinite system method gap results

The infinite system method (ISM) DMRG algorithm detailed in section 5.2.2 was implemented in *Mathematica* and used to generate gap values for fixed α and for a range of values of β , in each case starting with system size $L_{\text{init}} = 8$ and stopping only when the algorithm became so unstable that complex gap values were generated. The Arnoldi algorithm as provided by the ARPACK library [59], which was found in [148] to be the preferred diagonalisation routine for the non-Hermitian matrices arising in stochastic reaction-diffusion models, was accessed from within *Mathematica* and used to diagonalize the superblock Hamiltonian during each iteration of the ISM algorithm. Similarly, the LAPACK library [158] provided the appropriate routine for the less numerically demanding task of diagonalising the density matrices, which are symmetric.

Because the DMRG results are ultimately to be compared to both the domain wall theory prediction and the de Gier-Essler gap, the ideal choice is for α to be as small as possible, since then the difference between the two gap predictions is largest so that the DMRG results may most clearly be seen as adhering to either the one or the other. Unfortunately the stability of the ISM algorithm was found to follow the opposing trend,

²The Hamiltonian is exact for the first iteration of the algorithm

that is, to be worse for small α . An empirically optimal compromise of $\alpha = 0.4$ was used in all results presented in this section.

The size of the reducing basis was chosen to be $m = 16$. Since the state space associated with each side of the system for $L_{\text{init}} = 8$ is $2^4 = 16$, this means the full complement of eigenvectors of the reduced density matrices for this system size were used as a reducing basis, and hence there was not yet any error due to truncation introduced during the first renormalisation step, which is desirable. Similarly matched choices (L_{init}, m) are $(4, 4)$, $(6, 8)$, $(8, 16)$ and $(10, 32)$.

Assuming one uses the optimal choice of m for each of the choices of L_{init} , it is not made explicit in the primary literature cited in section 5.1 whether the most accurate DMRG results are expected if one initialises a DMRG algorithm by diagonalising the smallest possible system size, or if one begins with the largest system size that can still be diagonalised and grow-renormalises from there onward. Intuitively one expects the best results to come from the latter, since there the initial superblock states $|\psi_0\rangle$, $|\phi_0\rangle$, $|\psi_1\rangle$ and $|\phi_1\rangle$ are closer to their thermodynamic behaviour. It is worth investigating to what extent this assumption is justified for the TASEP.

To that end, figure 5.1 shows gap estimates as a function of system size, generated by applying the ISM algorithm for three of the different choices of L_{init} and m listed two paragraphs previously. Although the data for $L_{\text{init}} = 6$ lies closest to the de Gier-Essler gap at each iteration (i.e. for each value of L in figure 5.1), which suggests superior convergence, the algorithm becomes unstable earlier than is the case for the other two choices of L_{init} , namely around $L = 200$.

The ISM algorithm is expected to become unstable when at any point it attempts to resolve two eigenvalues in the spectrum of the renormalised Hamiltonian that are ‘too close’ together [146]. The point at which this occurs depends, as we saw, not on the true spectrum for given α , β , but on the particular parameters chosen for the run, such as how many states are kept.

Since the infinite system algorithm becomes most accurate for large L , it is desirable to obtain data for as large L as possible by choosing parameters that maximise stability. The preferable choice out of the three sets of initialisation parameters shown in figure 5.1 is therefore whichever one out of the two remaining that stays stable for the greatest number of iterations.

One can also evaluate the performance of the algorithm and hence the expected gap accuracy resulting from each parameter choice by studying the truncation error ϵ , which




Figure 5.1: DMRG results for the gap for $\alpha = 0.4$, $\beta = 0.75$, using the ISM method of section 5.2.2 and with the following initialisation choices: $L_{\text{init}} = 6$, $m = 8$ (blue \bullet), $L_{\text{init}} = 8$, $m = 16$ (purple \blacksquare) and $L_{\text{init}} = 10$, $m = 32$ (green \blacklozenge). The thermodynamic de Gier-Essler gap is shown as a constant (thick, black).

quantifies in some sense the error made in the renormalisation during each iteration of the algorithm. The truncation error is defined (see e.g. [147]) as the total discarded weight of the states of the density matrix that are excluded from the reducing basis:

$$\epsilon \equiv 1 - \sum_{k=1}^m \omega_k, \quad (5.29)$$

where the ω_k are the eigenvalues of one of the reduced density matrices, ordered in decreasing magnitude.

Figure 5.2 shows the truncation error for both sides of the system and for each iteration of the algorithm, i.e. for each system size L . As expected, the truncation error is zero for the first renormalisation step. The greatest loss of accuracy results from renormalisation steps early on, peaking around $L = 25$. In an ISM study of the one-dimensional Ising model in a transverse field, similar behaviour was observed for the difference between both ground and first excited states and their exact finite-size values, which grew quickly for small system sizes before saturating for larger L [159]. Interestingly, for the open-boundary TASEP with $\alpha = 0.4$, $\beta = 0.75$, the right side of the system suffers more from the approximation than does the left. One could speculate that this asymmetry is due to




Figure 5.2: Truncation errors associated with renormalisations of the left side Hamiltonian (blue \bullet) and of the right side Hamiltonian (purple \blacksquare) during a run of the ISM algorithm to calculate the gap for $\alpha = 0.4$, $\beta = 0.75$, initialised with $L_{\text{init}} = 6$, $m = 8$.

the fact that for these values of the boundary parameters the density profile is relatively flat everywhere except close to the right boundary. Operators for the right side of the system could therefore be expected to suffer more from a truncation of the basis than those for the left side. Both these features were found to be generic for all the choices of L_{init} and m discussed earlier.

Crucially, the typical value of the truncation error was found to decrease by roughly three orders of magnitude each time the initial system size was increased and more states were kept, following the sequence $L_{\text{init}} = 4 \rightarrow 6 \rightarrow 8 \rightarrow 10$. This confirms that the most accurate gap results are expected from ISM algorithms initialised with as large as possible an initial system size for which the exact Hamiltonian can still be diagonalised directly. A similar increase was observed in the Ising model study [159].

The *Mathematica* implementation of the algorithm allowed for rapid and flexible prototyping and data analysis, however it suffers from a shortcoming in its own internal implementation of the Arnoldi algorithm. The ARPACK library allows for diagonalisation of a non-Hermitian matrix without requiring that that matrix be explicitly constructed. Instead, the user may provide the result of acting upon a given vector with the matrix in question, which can be achieved without explicit construction thereof, thereby reducing memory requirements and speeding up the algorithm. Unfortunately *Mathematica* does




Figure 5.3: Stable DMRG data (red crosses) and exact thermodynamic gap (black line), all for $\alpha = 0.4$ (for which $\beta_c \approx 0.57$), and $\beta = 0.75$.

not implement this functionality, hence the results of this section were generated without this performance optimisation, which is deemed crucial by many practitioners (see section I in [146], and [155]).

It is partly for this reason (i.e. time and memory constraints) that the largest feasible initial system size was $L_{\text{init}} = 8$, with $m = 16$ states kept at each iteration, rather than the 32 used in [148].

The quality of convergence of the gap values that remain after discarding from the first complex gap value onwards was computed in the same way as described for Monte Carlo simulation data in section 4.2.4, i.e. according to equation 4.36 but with $\Delta\epsilon$ in place of $\Delta\phi$. Gap values were discarded starting from the first system size for which $\Delta\epsilon > 1.1$, which constitutes a much stricter cut-off than for the Monte Carlo data ($r = 1.1$ versus $r = 5$). For each choice of α and β this yields data of the type shown in Fig. 5.3.

The DMRG gap at the largest system size reached in the DMRG runs ($L = 250$ for the $\beta = 0.75$ example of Fig. 5.3) is not yet representative of the thermodynamic behaviour. As was done with the finite-size gap estimates from Monte Carlo simulations of chapter 4, we may extrapolate to the thermodynamic gap by using an appropriate fitting function and fitting to a subset of the data, both determined as in section 4.3 by varying the lower bound L_{min} of the datawindow and using the value that gives the largest adjusted coefficient of determination, R_{adj}^2 , defined in equation 4.39.




Figure 5.4: (a) Adjusted coefficient of determination R_{adj}^2 and (b) extrapolated gap, as a function of L_{min} for DMRG data corresponding to $\alpha = 0.4$, $\beta = 0.75$. The optimal choice of L_{min} lies just below 200, for which the extrapolated gap (red points) very closely matches the de Gier-Essler result (black line) and clearly excludes the DWT gap (green, dashed line).

Whereas Monte Carlo data in chapter 4 was extrapolated using a fourth-order function, system sizes reached in the DMRG study were sufficiently large that the most appropriate fitting function was found to be the leading order exact asymptotic form $a + \frac{b}{L^2}$. The example data of Fig. 5.3 is subjected to this analysis in Fig. 5.4(a), which shows a clear optimal value for L_{min} lying in the higher end of the full range, as expected.

Fig. 5.4(b) shows the extrapolated gap as a function of L_{min} , and it becomes clear by comparison with Fig. 5.4(a) that for the optimal choice of L_{min} this corresponds very closely to the exact gap (rather than to the domain wall theory). In practice, for a given value of β we take the average of the extrapolated gaps resulting from the five fits with largest R_{adj}^2 value.

Finally, carrying through this procedure for a range of values of β gives us the thermodynamic gap behaviour as predicted by DMRG, in Fig. 5.5. Though for $\beta > 0.5$ there are clearly remaining inaccuracies, I feel confident in stating that the DMRG results confirm the de Gier-Essler result and exclude the DWT gap and the possibility of a dynamical transition at $\beta = \frac{1}{2}$. Considering all DMRG gap values for $\beta > \beta_c$ as independent samples $\varepsilon_{\text{DMRG}}$ of the same value $\lim_{L \rightarrow \infty} \varepsilon(L, \alpha, \beta_c(\alpha)) = \varepsilon(0.4, \beta_c(0.4)) = -0.00121\dots$, we can state the agreement between the two by giving the average and standard deviation $\langle \varepsilon_{\text{DMRG}} \rangle = -0.00125 \pm 0.00009$.




Figure 5.5: Stable DMRG data (red crosses), exact thermodynamic gap (black curve) and DWT gap (green curve), all for $\alpha = 0.3$ (for which $\beta_c \approx 0.57$).

5.4 Discussion

Though DMRG results agree well with the de Gier-Essler gap, in fact very well for $\beta < 0.5$, one might ask how they might be improved for $\beta > 0.5$, especially given the extremely high accuracy achieved in many studies that use DMRG techniques, including in one that deals as we do here with the numerically less robust case of non-Hermitian transition matrices [148].

This is first of all due to the computational restrictions, for reasons discussed earlier in this section, on the largest initial system size and on the number of states kept at each iteration. Not only does using $L_{\text{init}} = 8$ and $m = 16$ lead to the ISM algorithm becoming unstable sooner than for $L_{\text{init}} = 10$ with $m = 32$, as figure 5.1 shows, it also yields different finite-size gap estimates for the same L . This means in turn that the extrapolations yield different gap values, showing how the estimated thermodynamic gap depends on particulars of the DMRG run, despite the fact that ISM runs for different parameter values should all converge to the same thermodynamic behaviour.

This highlights a second, more fundamental reason for the discrepancies, namely the finite-size states themselves. The infinite system method is not optimized to generate accurate finite-size gaps. For example, using $L_{\text{init}} = 8$ and $m = 16$, the gap from the renormalised Hamiltonian for $L = 10$ deviates from the gap given by exact numerical diagonalisation by an amount that is just larger than the limit of machine precision,

whilst by the second iteration this has grown to a difference in the fifth significant digit. Returning to the Ising model study [159], it was also found there that the verifiable deviation of the finite-size gaps generated by the ISM algorithm from their exact values increased as a function of system size, though it converged to a constant. All this warns us to be wary of using the finite-size gaps generated in an infinite system algorithm *at all* as a basis for extrapolation.

Instead, if for whatever computational reason it is not feasible, using an infinite system method, to reach system sizes large enough so that the states generated are sufficiently close to the thermodynamic behaviour for the intended purpose, then one should use the finite system method described in section 5.2 to generate accurate finite-size gaps, and extrapolate from those. In fact this is exactly how the remarkably accurate agreement with exact results was achieved in [148], as well as in the early numerically simpler Hermitian problems of Heisenberg spin chains [154].

The finite system method could therefore be used to improve the accuracy of the finite-size data and hence the quality of the extrapolants and agreement with the de Gier-Essler gap.

Chapter 6

The dynamical transition as an eigenvalue crossing

In chapters 3, 4 and 5 I investigated the dynamical transition by focusing exclusively on the behaviour of the gap, in isolation from the rest of the spectrum. This amounts to considering only how the transition is manifested in the asymptotic, longest-time relaxation behaviour and ignoring the potential role played by any earlier-time dynamics.

In this chapter I focus on another aspect of the dynamical transition, namely the interplay between the relaxational mechanisms, if identifiable, that are associated with the gap mode and the second gap mode, and the role this plays in defining the transition. In particular, in section 6.1 I suggest, on the basis of the results of de Gier and Essler, the possibility that the dynamical transition coincides with values of the boundary parameters at which eigenvalues corresponding to the gap and second gap cross.

To substantiate this hypothesis I propose in section 6.2 a perturbative approach, wherein one concentrates on the regime where α is very small and β close to one, for which the steady state corresponds to the almost always empty lattice. Then, in section 6.3, I identify candidate eigenvalues for the crossing by considering finite-size spectra. In section 6.4 I obtain perturbative expressions to leading order in α for the candidate gap and second gap eigenvalues within this regime. Finally, in section 6.5 I present an analysis showing that the eigenvalue crossing hypothesis reproduces with reasonable degree of accuracy the leading-order behaviour of the de Gier-Essler transition line.

The chapter closes with a discussion in section 6.6 of the implications of the apparent consistency of the crossing hypothesis and the perturbative approach for an understanding

of the dynamical transition.

6.1 An eigenvalue crossing scenario

In chapter 2 we saw that the dynamical transition is characterised by a nonanalyticity in the gap, namely a discontinuity in its first derivative, at $\alpha = \alpha_c(\beta)$ and at $\beta = \beta_c(\alpha)$, as well as by the fact that the gap is constant above the transition, i.e. for $\alpha > \alpha_c$ and for $\beta > \beta_c$.

Additionally, de Gier and Essler report in [18] on the asymptotic behaviour of the *second* gap, defined as $\varepsilon_2 \equiv \lambda_2 - \lambda_0 = \lambda_2$, which is found below the transition to obey

$$\varepsilon_2(L) = -\alpha - \beta + \frac{2}{(ab)^{\frac{1}{2}} + 1} - \frac{4\pi}{(ab)^{\frac{1}{2}} - (ab)^{-\frac{1}{2}}} L^{-2} + \mathcal{O}(L^{-3}), \quad (6.1)$$

with a, b defined as in (2.54) by

$$a = \frac{1 - \alpha}{\alpha} \quad \text{and} \quad b = \frac{1 - \beta}{\beta}. \quad (6.2)$$

Comparing (6.1) to the equivalent expression (2.53) for ε_1 , reproduced here for convenience:

$$\varepsilon_1(L) = -\alpha - \beta + \frac{2}{(ab)^{\frac{1}{2}} + 1} - \frac{\pi^2}{(ab)^{\frac{1}{2}} - (ab)^{-\frac{1}{2}}} L^{-2} + \mathcal{O}(L^{-3}), \quad (2.53)$$

it becomes clear that for $L \rightarrow \infty$ the second gap collapses onto the first gap, i.e. that

$$\lim_{L \rightarrow \infty} (\varepsilon_2 - \varepsilon_1) = 0. \quad (6.3)$$

Since the asymptotic forms of ε_2 for $\alpha > \alpha_c$ and for $\beta > \beta_c$ are not reported in [18], the equivalence (6.3) is proven only below the transition. It is thought¹ however to hold over the entire parameter range, i.e. throughout the entire LD and HD regimes.

For finite L however the leading-order corrections to the first and second gaps are different, as can be observed by again comparing (6.1) with (2.53). Taking the LD regime as an example this means, as noted in [18], that up to and including the L^{-2} terms it is the case that

$$\lim_{\beta \uparrow \beta_c} \varepsilon_2(L) = \lim_{\beta \downarrow \beta_c} \varepsilon_1(L) \neq \lim_{\beta \uparrow \beta_c} \varepsilon_1(L), \quad (6.4)$$

as illustrated in figure 6.1. The discontinuity in the gap at β_c originates in the leading order finite L correction term. Higher order terms are likely to ameliorate the

¹F. Essler, private communication




Figure 6.1: The leading-order behaviour for finite L of the gap ε_1 (solid, black) and second gap ε_2 (dashed, blue), as a function of β for $L = 100$ and with $\alpha = 0.3$, i.e. in the LD regime. The behaviour at β_c illustrates the relation between the gap and second gap expressed in equation 6.4, however the expression (6.5) used for ε_2 for $\beta > \beta_c$ is a conjecture and for illustrative purposes only.

discontinuity, which should therefore not be taken to be a feature of the exact finite L spectrum. The same applies to the behaviour of the second gap shown in figure 6.1. Even though no evidence is given in [18] for the existence of a discontinuity in the leading-order behaviour of the second gap, this scenario is thought² to be likely. The expression used in figure 6.1 to plot ε_2 for $\beta > \beta_c$ is simply the equivalent function for ε_1 (equation 2.57), modified to have a jump at the transition that is twice as large as that for ε_1 :

$$\varepsilon_2(L) = -\alpha - \beta_c + \frac{2}{(ab_c)^{\frac{1}{2}} - 1} - \frac{8\pi}{(ab_c)^{\frac{1}{2}} - (ab_c)^{-\frac{1}{2}}} L^{-2}. \quad (6.5)$$

I stress that (6.5) has not been derived in any way; it is an ad hoc conjecture purely to help illustrate the scenario described above.

The equivalence, to leading order, of the second gap below the transition and the gap above it suggests the possibility of a scenario wherein the eigenvalues associated with each of the two modes cross, i.e. intersect, at β_c . For example, in [74] de Gier and Essler speculate that an eigenvalue crossing takes place at an analogous dynamical transition line in the *partially* asymmetric exclusion process, and that it corresponds to a change in the detailed nature of the relaxation. No such crossing is observed in the spectra for small system sizes, where we see instead that the second gap corresponds to the real part of a complex conjugate pair of eigenvalues – see figure 6.2. In determining the behaviour of solutions to the Bethe ansatz equations for the TASEP in [18], de Gier and Essler observe that the second gap generically consists of a complex conjugate pair for all system sizes $L \leq L'(\alpha, \beta)$, where $20 \leq L'(\alpha, \beta) \leq 100$, and that for $L > L'$ the second gap is real and distinct.

6.2 A perturbative approach to the crossing

If the eigenvalue crossing hypothesis outlined in section 6.1 is true then it should in principle be possible to obtain the location of the dynamical transition, i.e. to arrive at $\alpha_c(\beta)$ and $\beta_c(\alpha)$, by determining the functional relationship that holds between α and β for all the points at which the gap and second gap cross. Making this determination requires, at first thought, that the gap and second gap are known exactly and in closed form for finite L , which unfortunately they are not. Although it is not obvious how one might obtain such expressions for general α and β , it may be that a perturbative approach provides a viable alternative, as follows.

²F. Essler, private communication.




Figure 6.2: Real part of the TASEP spectrum for $L = 3$ and $\beta = 0.3$, including strictly real eigenvalues (black \bullet) and complex eigenvalues (blue \times). Starting from $\alpha = 0$ the second gap is real until $\alpha \approx 0.1$ (indicated), where it splits into a complex conjugate pair of eigenvalues. In [18] it was observed that this behaviour is generic for small system sizes $L \leq L'(\alpha, \beta)$, where $20 \leq L'(\alpha, \beta) \leq 100$.

We may fix β and consider how the dynamical transition in the LD regime is approached by varying α (as opposed to fixing α and varying β as in figure 6.1). When β is close to 1 then the dynamical transition is reached already for small α – see figure 2.9, which is reproduced here as figure 6.3 for convenience. Hence β_c allows an expansion in α and is given to leading order by

$$\beta_c \approx 1 - \alpha^{\frac{1}{3}}. \quad (6.6)$$

Now consider the TASEP transition matrix, M . Without any loss of generality (i.e. for general α) we can write M for any system size L explicitly as a linear relation in α :

$$M = M^0 + \alpha M^1, \quad (6.7)$$

and identify, from equation 2.9,

$$M^0 = \sum_k \mathcal{I}^{\otimes k-1} \otimes h_b \otimes \mathcal{I}^{\otimes L-k-1} + \mathcal{I}^{\otimes L-1} \otimes h_r, \quad (6.8)$$

and

$$M^1 = h_1^1 \otimes \mathcal{I}^{\otimes L-1}, \quad (6.9)$$

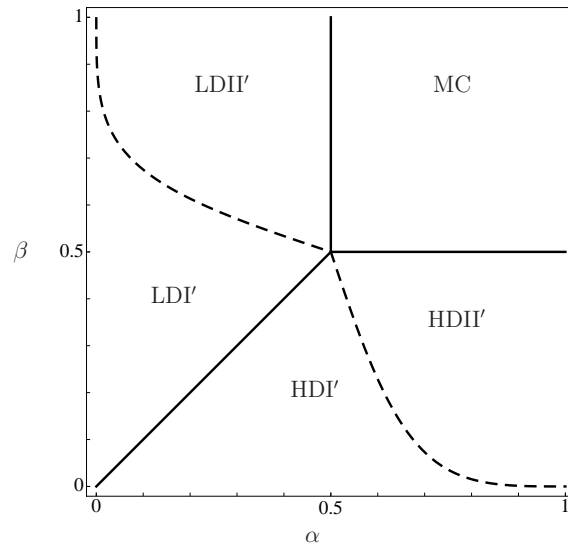


Figure 6.3: Dynamical phase diagram of the TASEP

where as a slight modification to equation 2.10 for h_1 we have

$$h_1^1 = \begin{pmatrix} -1 & 0 \\ 1 & 0 \end{pmatrix} \quad (6.10)$$

For example, for $L = 2$

$$M = M^0 + \alpha M^1$$

$$= \begin{pmatrix} 0 & \beta & 0 & 0 \\ 0 & -\beta & 1 & 0 \\ 0 & 0 & -1 & \beta \\ 0 & 0 & 0 & -\beta \end{pmatrix} + \alpha \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

Next, note that for α small equation 6.7 describes an analytic perturbation of the transition matrix M^0 . Since this amounts to an analytic perturbation of the characteristic polynomial from which the eigenvalues are derived, one expects it to be possible to find perturbative expressions for all eigenvalues, including for the gap and second gap, when α is small. Although it can be shown that these exist and are given by power series expansions in α , as we shall see in section 6.4 the fact that M^0 is non-normal means that finding them can be more difficult than the equivalent problem for perturbed Hermitian operators in quantum mechanics.

If the dynamical transition does in fact correspond to a crossing of the two gaps

then it should be possible, by equating the perturbative expressions for the two gaps and rearranging for one of the boundary parameters, β , say, to determine to leading order the functional relationship between α and β that determines the location of the transition line. If the result matches the expansion of β_c to leading order (equation 6.6), i.e. if by this approach one is able to rederive the location of the dynamical transition, that fact can be taken as evidence in support of the crossing hypothesis.

Our strategy in what follows will be to obtain, in section 6.4, perturbative expressions for the gap and second gap for a number of small system sizes, and to conjecture their forms for general L . First however, in order to ensure that the correct perturbative expressions are obtained, it is necessary to identify which eigenvalues in the spectrum correspond to the gap and second gap when the perturbation $\alpha > 0$ is turned on and when β is close to one; this is dealt with in section 6.3. The final step, in section 6.5, is to equate the forms for general L and thereby to determine whether the crossing scenario yields the transition line.

6.3 Identifying gap and second gap eigenvalues

In this section we want to determine the correct eigenvalues for which to obtain perturbative expressions in section 6.4. We are concerned therefore with identifying those eigenvalues in the unperturbed spectrum that, when perturbed by $\alpha > 0$, yield the gap and second gap when β is close to one. We begin by considering the gap and second gap for the unperturbed problem, i.e. the two eigenvalues closest to zero for $\alpha = 0$.

Without needing to compute the entire spectrum, a straightforward consideration reveals that for general L the two eigenvalues closest to zero in the spectrum of M^0 are always

$$\lambda_\beta \equiv -\beta \quad \text{and} \quad \lambda_1 \equiv -1. \quad (6.11)$$

To see why this is so, first note that M^0 is upper triangular for any L , and that therefore its eigenvalues are given by the diagonal entries. As discussed in section 2.3, for a continuous-time transition matrix a diagonal entry M_{ii}^0 is given by negative the sum of the rates for all transitions away from the configuration associated with column i , i.e. negative the total of all other entries in that column. Hence the corresponding eigenvalue λ_i is given by

$$\lambda_i = -M_{ii}^0 = - \sum_{j \neq i}^{2^L - 1} M_{ij}^0. \quad (6.12)$$




Figure 6.4: Real part of the TASEP spectrum for $L = 5$ as a function of small α and for $\beta = -0.99$. Each of the degenerate unperturbed eigenvalues $\lambda_\beta = -\beta$ and $\lambda_1 = -1$ sprouts a cluster of eigenvalues for $\alpha > 0$, some of which are entirely real (black \bullet), whilst other branches consist of complex conjugate pairs (blue \times).

Clearly $|\lambda_i|$ is smallest if only one transition is possible, namely either a particle exiting the system, which happens at rate β , or a particle moving within the bulk at rate 1, thus giving us the unperturbed eigenvalues of interest.

It is promising for our crossing hypothesis that these unperturbed eigenvalues are the same (i.e. they cross in a trivial sense) at the same point where de Gier and Essler find the transition for $\alpha = 0$, i.e. that

$$\lim_{\beta \rightarrow \beta_c(0)} \lambda_\beta = \lim_{\beta \rightarrow 1} \lambda_\beta = -1 = \lambda_1 ,$$

since this makes it feasible that for nonzero but small α , and β close to one, they will cross once more.

Having determined the gap and second gap for $\alpha = 0$, the question then becomes which eigenvalues play the same role for $\alpha > 0$. In figure 6.4 we see that for $\alpha > 0$ both λ_1 and λ_β give rise to a cluster of eigenvalues, which highlights the fact that for $\alpha = 0$ they are degenerate. This situation is analogous to the familiar spectral splitting observed in the perturbative treatment of the interactions of weak magnetic and electric fields with a magnetic moment, namely the Stark and Zeeman effects respectively. There,

as here, the perturbation lifts the degeneracy of the noninteracting regime (here, the empty lattice) so that an n -fold degenerate unperturbed eigenvalue splits into n branches. However, the analogy breaks down when one considers the form of the expansions for perturbed eigenvalues, which for non-Hermitian matrices contain fractional rather than integer powers and can be less straightforward to obtain, as we shall see in section 6.4. Moreover, whereas the Hermitian operators of quantum mechanics allow only real-valued shifts or splittings in the spectrum, in the non-Hermitian case the perturbative solutions generally include complex conjugate pairs, as is the case for the eigenvalue clusters shown in figure 6.4.

For the TASEP transition matrix we see in figure 6.4 that the -1 cluster and the $-\beta$ cluster partially overlap, with the top branch of the -1 cluster intersecting the two lowest branches of the $-\beta$ cluster. As β approaches 1 the clusters overlap more and more. For α small enough or for β sufficiently far away from 1, the gap is still given by the top branch of the $-\beta$ cluster, and the second gap by the second-largest branch in that same cluster. However as β approaches 1 this will change. First the top branch of the -1 cluster intersects with the second-largest branch of the $-\beta$ cluster, taking over the role of the second gap. At some point for β even closer to 1 it will intersect the top branch of the $-\beta$ cluster, i.e. with what was until then the gap itself. Hence the relevant eigenvalues to obtain perturbative expressions for as β approaches 1 and α approaches zero are the top branches of the -1 and $-\beta$ clusters.

Before proceeding in section 6.4 to obtain the perturbative expressions for these branches, it is useful to understand, from an intuitive calculation shown in section 6.3.1, that one can compute the exact degeneracies of λ_1 and λ_β , and thereby gain insight into their physical origins. This will prove relevant in interpreting the dynamical transition, as I discuss in section 6.6.

6.3.1 Degeneracies of the unperturbed eigenvalues


The unperturbed eigenvalues λ_β and λ_1 are both degenerate. That this is the case is easily understood by first of all realising that the number of appearances μ_i on the diagonal of $M^0 - \lambda$ of the factor $(\lambda_i - \lambda)$ is equal to the number of configurations for which the total transition rate away is equal to $-\lambda_i$. Since the characteristic polynomial $p_0(\lambda)$ of M^0 is

given by

$$\begin{aligned}
 p_0(\lambda) &= \text{Tr}(M^0 - \lambda) \\
 &= \prod_i^{2L} (\lambda - \lambda_i),
 \end{aligned}
 \tag{6.13}$$


then $(\lambda_i - \lambda)$ appears in $p_0(\lambda)$ raised to the power of μ_i , hence μ_i is the algebraic multiplicity of λ_i . If there is more than one state for which the total exit rate equals $-\lambda_i$, then $\mu_i > 1$ and λ_i is degenerate. A combinatorial argument readily gives us the multiplicities μ_β and μ_1 of respectively λ_β and λ_1 for general L , as follows.

The only configurations from which no other transition is possible apart from a particle hopping out at the right boundary at rate β are those consisting of a domain of occupied sites only on the right side of the system, and a domain of unoccupied sites only on the left:



For configurations fitting this description no particles can hop within the bulk, and since $\alpha = 0$ no particles can enter the system. The size N of the occupied domain can range from 1 to L , and with each occupied domain size implicitly determining the unoccupied domain size there are therefore L possible configurations, so $\mu_\beta = L$.

Likewise, the configurations that allow no other transitions apart from one particle hopping at rate 1 within the bulk of the system are those consisting of a completely occupied domain of size l satisfying $1 \leq l \leq L - 1$, positioned anywhere in the system as long as the L^{th} site is left free:



Note this includes single-particle configurations such as



The only enumerative freedom in constructing these configurations lies in choosing the positions of the left and right boundary sites of the l -domain. Once this is done there is only one possibility for the remainder of the configuration. Hence the number of consistent

configurations and therefore the multiplicity of λ_1 is given by

$$\mu_1 = \binom{L}{2} = \frac{L!}{(L-2)!2!} \quad (6.14)$$

where the ‘choose two’ in the binomial coefficient refers of course to the boundary sites. At first sight we appear to be overcounting our configurations by $L - 1$ by choosing out of L , rather than out of $L - 1$ as one would expect given that the L^{th} site remains free. However in following the above counting logic we leave out the single-particle cases, where the boundary sites are one and the same. Including these contributes an additional $L - 1$ configurations, an equal number as was already compensated for in our expression. Hence (6.14) is correct. Alternatively, the argument can be reformulated by considering all possible choices for the positions of the bonds to the rest of the system made by the leftmost and rightmost particles in the contiguous particle domain.

It is easy to verify from the spectra for small system sizes that μ_β and μ_1 correctly predict the degeneracies for $\alpha = 0$ as well as the sizes of the clusters that emerge for $\alpha > 0$.

6.4 Perturbative expressions for the gap and second gap

In this section I demonstrate how to arrive at perturbative expressions for the top branches of the $-\beta$ and -1 eigenvalue clusters of M , which we recall is given by

$$M = M^0 + \alpha M^1.$$

6.4.1 Perturbation theory for normal and non-normal matrices

In this subsection I refer to key perturbation-theoretic results, proofs and summaries of which may be found in [160–163], while a comprehensive treatment of formal issues such as well-posedness, radii of convergence of the power series, etc. can be found in [160, 164, 165].

Let $F(\alpha, \beta, \lambda)$ denote the characteristic polynomial of M , which is of degree $d = 2^L$:

$$F(\alpha, \beta, \lambda) = \lambda^d + k_{d-1}(\alpha, \beta) \lambda^{d-1} + \dots + k_1(\alpha, \beta) \lambda + k_0(\alpha, \beta). \quad (6.15)$$

The eigenvalues λ_i^0 of the unperturbed matrix M^0 are the roots of $F(0, \beta, \lambda)$ and are known, whereas the eigenvalues λ_i of M are the 2^L roots of (6.15) amongst which are the top branches of the -1 and $-\beta$ clusters, whose dependence on α we want to determine for

α small.

If the coefficients $k_i(\alpha, \beta)$ are analytic functions of α and β then the map $(\alpha, \beta) \mapsto \{\lambda_i\}_1^{2^L}$ constitutes a multi-valued algebraic function, $\lambda(\alpha, \beta)$, which is the solution to $F(\alpha, \beta, \lambda) = 0$. It has been known for some time that if the k_i satisfy the analyticity condition, then by the theory of algebraic functions the roots of $F(\alpha, \beta, \lambda)$, i.e. the eigenvalues λ_i of M , are themselves analytic functions of α and hence amenable to a perturbative approach. Hence for α small each of the λ_i can be expanded as a power series around some λ_i^0 .

Moreover, it can be shown that for an unperturbed eigenvalue with algebraic multiplicity m there exist m distinct single-valued analytic functions in a neighbourhood surrounding $\alpha = 0$. Hence the perturbation $\alpha > 0$ completely lifts all degeneracies, as we saw in figure 6.4, section 6.3.

For Hermitian operators, which includes quantum Hamiltonians, it can be shown that self-adjointness of the unperturbed operator implies that only integer powers appear in the expansions for perturbed eigenvalues. The derivation of these Taylor series, i.e. the computation of the so-called Rayleigh-Schrödinger coefficients, is well-established and at the core of finite-dimensional perturbation-theoretic methods in quantum mechanics and more widely in mathematical physics.

Perturbation theory for non-normal matrices such as the TASEP transition matrix M^0 , on the other hand, is less well-established. Here the expansions for the eigenvalues in general have fractional powers: they are Puiseux series given, for example for a perturbation around $\alpha = 0$ and with integer n , by

$$\lambda_i(\alpha, \beta) = \lambda_i^0 + \sum_{i=1}^{\infty} a_i \alpha^{\frac{i}{n}}. \quad (6.16)$$

This result is at least as early as [160], while the more recent [166] not only summarises results since then but also recovers and extends a calculational approach [167] that was translated from [168]. Since M^0 is non-normal it is not reducible to diagonal form by a unitary transformation, nor does it possess a complete set of 2^L eigenvectors. The algebraic multiplicity (i.e. the eigenvalue degeneracy) exceeds the geometric multiplicity (the number of independent eigenvectors), i.e. M^0 is defective. This is directly related to the appearance of fractional powers in expansions for eigenvalues of M and has implications for the leading-order behaviour of those expansions, as the following shows.

6.4.2 Jordan normal form, degeneracy structure, and Lidskii perturbation

Although M^0 cannot be diagonalised it can always be reduced to the more general, nearly-diagonal Jordan canonical form (Jordan normal form) matrix J [160] by means of a similarity transformation:

$$J = P^{-1}M^0P. \quad (6.17)$$

Like a diagonal matrix, the Jordan form of a matrix directly exposes the underlying spectral structure, including algebraic and geometric multiplicities. To highlight this structure we represent J partitioned into blocks \widehat{J}_i , arranged diagonally and with all offdiagonal elements equal to zero.

$$\begin{pmatrix} \widehat{J}_1 & & & \\ & \widehat{J}_2 & & \\ & & \ddots & \\ & & & \widehat{J}_l \end{pmatrix}. \quad (6.18)$$

Each \widehat{J}_i is associated with one out of a total of l distinct eigenvalues in the spectrum. Focusing on a particular eigenvalue λ_i and hence a particular \widehat{J}_i , we now drop the i subscript. The \widehat{J} block can in turn be decomposed into a diagonal arrangement of different blocks Γ_j :

$$\widehat{J} = \begin{pmatrix} \Gamma_1 & & & \\ & \Gamma_2 & & \\ & & \ddots & \\ & & & \Gamma_q \end{pmatrix}. \quad (6.19)$$

The blocks $\{\Gamma_j\}_{j=1}^q$ are the so-called Jordan blocks associated with λ . Each Jordan block itself is a simple near-diagonal listing of λ :

$$\Gamma_j = \begin{pmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{pmatrix}. \quad (6.20)$$

The Jordan blocks are identical in form, differing only in their dimensionality, n_j , equal to the number of appearances of λ on the diagonal. Clearly the overall algebraic multiplicity m of λ is therefore given by

$$m = \sum_{j=1}^q n_j. \quad (6.21)$$

Furthermore, each Jordan block is associated with a distinct genuine eigenvector of M^0 , or rather with a distinct pair consisting of a genuine left eigenvector and a genuine right eigenvector, so that q is in fact the geometric multiplicity of λ . The genuine eigenvector pair associated with each block form the basis of a generalized eigenspace, which it does not suit our purpose to discuss here.

In this brief exposition of Jordan normal form I have mostly followed the notation of [166], though I have ignored the possibility that identical Jordan blocks Γ_j with the same dimensionality n_j appear more than once in $\hat{\mathcal{J}}$.

We are of course interested in the relevance of the Jordan structure of M^0 for the effect of the perturbation of its eigenvalues. This is dealt with in principle by the contribution to formal matrix perturbation theory by Lidskii, described in [166]. There it is shown that when perturbed, a Jordan block associated with λ and of dimension n_j generically (but not always!) generates eigenvalues λ' satisfying the Puiseux series

$$\lambda' = \lambda + \sum_{s=1}^{\infty} a_s \alpha^{\frac{s}{n_j}}. \quad (6.22)$$

Furthermore, if (6.22) holds then a_1 is multi-valued and given by the n_j distinct n_j^{th} roots of an intermediate quantity computed, as described in [166], using the genuine left and right eigenvectors for the Jordan block under consideration.

These two theorems, which I shall refer to as the Lidskii result, directly relate the expansion behaviour of perturbed eigenvalues to the Jordan structure of the unperturbed matrix. They do this by revealing the common origin of the fractional powers in the series expansions and of the number of eigenvalues in the degeneracy-lifted clusters that sprout from degenerate eigenvalues in the unperturbed spectrum. Unfortunately the Lidskii result is not universal and though it is easy to test if it holds, obtaining the leading order perturbative behaviour when it does not cannot be achieved by a simple modification of the argument. In that case all that can be said without any further investigation is that the leading-order behaviour of the perturbation is $o(\alpha^{1/n_j})$, and still possesses a fractional power given the result (6.16), which is stronger than the Lidskii result and holds regardless.

For the sake of clarity and concreteness consider as an example the Jordan form of

and

$$\lambda'_\beta = \lambda_\beta + b \alpha^y, \quad (6.27)$$

with a , b , x , and y to be determined by requiring that λ'_1 and λ'_β individually solve, to leading order in α ,

$$F(\alpha, \beta, \lambda) = 0. \quad (6.28)$$

Naively one might be tempted to truncate $F(\alpha, \beta, \lambda)$ after the leading order term in α and hence require that the ansätze solve

$$p_0(\beta, \lambda) = -p_1(\beta, \lambda) \alpha. \quad (6.29)$$

However, one quickly realises by substituting either ansatz for λ into the full form of $F(\alpha, \beta, \lambda)$ that each $p_i \alpha^i$ term itself contributes, for example for λ'_1 :

$$p_{i,0}(\beta, a) \alpha^i + p_{i,1}(\beta, a) \alpha^{i+\nu_{i,1}(x)} + p_{i,2}(\beta, a) \alpha^{i+\nu_{i,2}(x)} + \dots, \quad (6.30)$$

where the $\nu_{i,j}$ are rational functions of x .

Let us denote by $\tilde{p}_i(\beta, a) \alpha^{\eta_i(x)}$ the dominant term appearing in the expression (6.30) for $p_i \alpha^i$, i.e. let the dominant power therein be

$$\eta_i(x) \equiv \min_j \{i + \nu_{i,j}(x)\} \quad (\text{with } \nu_{i,0}(x) = 0) \quad (6.31)$$

and the dominant coefficient,

$$\tilde{p}_i(\beta, a) \equiv \sum_j \delta_{\eta_i, i+\nu_{i,j}} p_{i,j}(\beta, a), \quad (6.32)$$

and note that the simplest case $\eta_i = i$ occurs only when $p_{i,0}$ happens to be nonzero. The same considerations apply for λ'_β , giving us $\eta_i(y)$ and $\tilde{p}_i(\beta, b)$, and allow us to observe that merely solving (6.29) to determine the perturbative expressions is justified only if

$$\eta_{i>2} > \eta_{i\leq 2}, \quad (6.33)$$

i.e. only if the leading order contributions to $F(\alpha, \beta, \lambda)$ for λ in the neighbourhood of $-\beta$ and $-\beta$ come solely from p_0 and from $p_1 \alpha$.

There is no *a priori* reason to assume this is true, so that contributions to the leading-order behaviour may come from any term in $F(\alpha, \beta, \lambda)$. Indeed upon substitution of the ansätze and using the additional information, guaranteed by (6.16), that x and y are both fractions, it turns out that for $L > 2$ the dominant contributions $\tilde{p}_i \alpha^{\eta_i}$ from $p_i \alpha^i$

terms obeying $0 \leq i \leq i_{\max}$ are all of the same order and need to be taken into account simultaneously, i.e. one finds

$$\eta_1 = \eta_2 = \dots \eta_{i_{\max}} \equiv \eta, \quad (6.34)$$

where i_{\max} is an increasing function of L . The leading-order solution to $F(\alpha, \beta, \lambda) = 0$ is therefore found by solving (6.34) for η , which establishes the leading-order powers x and y for the perturbative expressions. Finally, the leading-order coefficients are found by solving

$$\sum_{i=0}^{i_{\max}} \tilde{p}_i = 0. \quad (6.35)$$

for a and b , using $\tilde{p}_i(\beta, a)$ and $\tilde{p}_i(\beta, b)$ respectively.

This yields μ_1 distinct values of a and μ_β distinct values of b , as expected given that μ_1 and μ_β are the degeneracies of λ_1 and λ_β , which are lifted by the perturbation. Out of these solutions, which together describe the leading order behaviour of all eigenvalues in the $-\beta$ and -1 clusters, we focus on the largest positive corrections to the real parts, which yield the top branches of the clusters. This direct approach is equivalent to the more elaborate graphical analysis of the characteristic polynomial which is proposed in [166] as an extension of the basic Lidskii method.

The above analysis was performed for $2 \leq L \leq 7$ using *Mathematica* and, for the top branches of the -1 and $-\beta$ clusters, led to the perturbative expressions listed in Table 6.1. The perturbations for the other branches in a given cluster grow with the same power as the top branch; it is only the coefficients that differ and which bring parts of the cluster into the complex plane. The top branches however are always real. Figures 6.5 and 6.6 show how well the perturbative expressions for $L = 5$ match the exact spectrum. For a given value of α the λ_1 cluster shows a larger deviation from the leading-order perturbative expression. We can now see that the leading order of the perturbation to λ_β , which we were not able to find in section 6.4.2, is of order $\alpha^{1/2}$, independent of L , rather than $\alpha^{1/\mu_\beta} = \alpha^{1/L}$ as naively expected given the Lidskii result (6.22).

Though the results in Table 6.1 do not suggest any conjecture for the form of a and b for general L , a pattern is apparent for $3 \leq L \leq 7$ in the relationship between these coefficients for a given system size, namely that a and b obey

$$b(L) = (a(L) + 1)^{\frac{1}{2}} \left(\frac{\beta}{1 - \beta} \right)^{\frac{1}{2}}. \quad (6.36)$$

It can be verified numerically that this equality also holds, up to machine precision, for

L	x	y	a	b
2	1	$\frac{1}{2}$	$-\frac{\beta(\beta-2)}{(1-\beta)^2}$	$\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$
3	$\frac{1}{3}$	$\frac{1}{2}$	1	$2^{\frac{1}{2}}\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$
4	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{2}(1+\sqrt{5})$	$\left(\frac{1}{2}(3+\sqrt{5})\right)^{\frac{1}{2}}\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$
5	$\frac{1}{3}$	$\frac{1}{2}$	2	$3^{\frac{1}{2}}\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$
6	$\frac{1}{3}$	$\frac{1}{2}$	$\left(\frac{11+2\sqrt{7\times 19}\cos\left(\frac{1}{3}\tan^{-1}\left(\frac{39\sqrt{3}}{433}\right)\right)}{3}\right)^{\frac{1}{3}}$	$\left(\frac{5+2\times 7^{\frac{1}{2}}\cos\left(\frac{1}{3}\tan^{-1}(3\sqrt{3})\right)}{3}\right)^{\frac{1}{2}}\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$
7	$\frac{1}{3}$	$\frac{1}{2}$	$1+\sqrt{2}$	$(2+\sqrt{2})^{\frac{1}{2}}\left(\frac{\beta}{1-\beta}\right)^{\frac{1}{2}}$

Table 6.1: Perturbative expressions for top branches of $-\beta$ and -1 eigenvalue clusters, obtained by direct analysis of the characteristic polynomial.




Figure 6.5: For $L = 5$ and $\beta = 0.3$: perturbative expressions for the eigenvalue cluster originating in the unperturbed eigenvalue λ_β (red curves), the top branch of which is listed in Table 6.1, and the exact spectrum of M , real parts only (black \bullet).




Figure 6.6: For $L = 5$ and $\beta = 0.3$: perturbative expressions for the eigenvalue cluster originating in the unperturbed eigenvalue λ_1 (red curves), the top branch of which is listed in Table 6.1, and the exact spectrum of M , real parts only. Strictly real eigenvalues are shown as black \bullet , complex conjugate pairs as blue \times .

$L = 6$, despite the fact that perturbative expressions for this system size are the odd ones out in not having an obvious algebraic closed form.

Given that a and b describe the leading order behaviour within different eigenvalue clusters, it is rather remarkable that this relationship holds, and it seems indicative of some underlying connection between the dynamical behaviour associated with the two clusters, including for $L \rightarrow \infty$. Indeed in section 6.5 we see how this relationship (6.36), which I hereby conjecture holds for general L , links the scaling behaviour of $b(L)$ to that of $a(L)$ and thus helps determine the transition point predicted by the eigenvalue crossing hypothesis.

6.5 Rederiving the dynamical transition

Our hypothesis is that the location of the dynamical transition is defined by the crossing point of the top branches of the -1 and $-\beta$ eigenvalue clusters, i.e. that it occurs when

$$\lambda'_\beta = \lambda'_1. \quad (6.37)$$

Using the results of section 6.4, including the observed pattern (6.36) between the perturbation coefficients a and b , the crossing criterion (6.37) can be rewritten, for $L \geq 3$, as

$$-\beta + (a(L) + 1)^{\frac{1}{2}} \left(\frac{\beta}{1 - \beta} \right)^{\frac{1}{2}} \alpha^{\frac{1}{2}} = -1 + a(L) \alpha^{\frac{1}{3}}. \quad (6.38)$$

To substantiate the crossing hypothesis we want to determine $\beta(\alpha)$ to leading order as dictated by (6.38) in the limits $\alpha \rightarrow 0$ and $L \rightarrow \infty$, and compare with the location $\beta_c(\alpha)$ found by de Gier and Essler under the same conditions, namely

$$\beta_c \approx 1 - \alpha^{\frac{1}{3}}. \quad (6.39)$$

This suggests we make an ansatz similar to (6.39) for β in (6.38), namely

$$\beta = 1 - c \alpha^\gamma, \quad \text{with } \gamma > 0, \quad (6.40)$$

which gives

$$c \alpha^\gamma + (a(L) + 1)^{\frac{1}{2}} \left(\frac{1}{c} \alpha^{1-\gamma} - \alpha \right)^{\frac{1}{2}} = a(L) \alpha^{\frac{1}{3}}. \quad (6.41)$$

In order for (6.41) to remain finite as $\alpha \rightarrow 0$ we require $\gamma \leq 1$. The dominant behaviour then becomes

$$c \alpha^\gamma + \left(\frac{a(L) + 1}{c} \right)^{\frac{1}{2}} \alpha^{\frac{1-\gamma}{2}} = a(L) \alpha^{\frac{1}{3}}. \quad (6.42)$$


Noting that

$$\frac{1-\gamma}{2} < \gamma \text{ for } \gamma > \frac{1}{3} \quad (6.43)$$

$$\frac{1-\gamma}{2} > \gamma \text{ for } \gamma < \frac{1}{3}, \quad (6.44)$$

we find γ by a process of elimination, as follows. Assume that $\gamma > \frac{1}{3}$. This means, given the inequality (6.43), that for $\alpha \rightarrow 0$ the second term on the lhs in (6.42) dominates over the first term. Requiring that the second term matches the scaling of the rhs, i.e. that $\frac{1-\gamma}{2} = \gamma$, gives $\gamma = \frac{1}{3}$, in contradiction to the original assumption. Similarly, assuming the other inequality (6.44) holds implies the first term dominates the lhs, which leads again to $\gamma = \frac{1}{3}$, in contradiction to the second assumption. The two contradictions lead us to conclude that strict equality holds, i.e. that

$$\gamma = \frac{1}{3}. \quad (6.45)$$

Figure 6.7: Coefficients a from Table 6.1 for $3 \leq L \leq 7$.

This agrees with the power of the de Gier-Essler result and converts equation 6.42 to

$$c + \left(\frac{a(L) + 1}{c} \right)^{\frac{1}{2}} - a(L) = 0, \quad (6.46)$$

which has solutions

$$c(L) = \frac{1}{2} \left(a(L) - 1 \pm \sqrt{a^2(L) - 2a(L) - 3} \right). \quad (6.47)$$


We are interested in the behaviour of $c(L)$ as $L \rightarrow \infty$, however since we have no conjecture for the form of $a(L)$ for general L we must resort to a numerical investigation. Figure 6.7 shows the numerical values of the coefficients $a(L)$ from Table 6.1 for $2 \leq L \leq 7$, which suggests that $a(L)$ is convergent for $L \rightarrow \infty$,

$$a_0 \equiv \lim_{L \rightarrow \infty} a(L). \quad (6.48)$$

This means that $c(L)$ in turn is also convergent:

$$c \equiv \lim_{L \rightarrow \infty} c(L) = \frac{1}{2} \left(a_0 - 1 \pm \sqrt{a_0^2 - 2a_0 - 3} \right). \quad (6.49)$$

Before we attempt to determine a_0 it is intriguing to note that if $a_0 = 3$ then the two solutions of (6.49) are equal, and we have $c = 1$, in exact agreement with the de Gier-Essler result.



(a) The deviation $|\delta^{(n)}|$ at the last iteration is minimal at $\omega \approx 1.82$ (indicated). (b) BST extrapolant estimates for a_0 as a function of ω . The most accurate extrapolant is for $\omega \approx 1.82$, as determined in (a), and is equal to $a_0 \approx 3.01$.

Figure 6.8: Bulirsch-Stoer extrapolation results for a_0

Now we want to obtain the best possible estimate of the asymptotic value a_0 , based on the exact values $a(L)$ for $3 \leq L \leq 7$. This is done using the BST algorithm, as described in section 2.7. The best extrapolated value according to the BST algorithm is $a_0 \approx 3.01$, given by the choice $\omega \approx 1.82$ which minimises the error $|\delta^{(n)}|$, as shown in figure 6.8.

Despite the fact that we find a_0 very close to three, which corresponds to the de Gier-Essler result, the fact that it is not exactly equal means that (6.49) gives two solutions rather than one for c . This is anomalous, since the dynamical transition occurs only at one point for a given value of α or β . The values of c are distributed asymmetrically around the de Gier-Essler result $c = 1$:

$$c \approx \{0.90, 1.12\}. \quad (6.50)$$

To put this result in perspective: one finds for arbitrary subsets containing only four rather than all five of the finite-size values $a(L)$ for $3 \leq L \leq 7$ that the BST algorithm, now capable of three rather than four iterations, yields extrapolants all roughly equal to $a_0 \approx 2.8$. The observed change in going from three to four iterations suggests it is very reasonable to suspect that if exact finite-size values were calculated for a few more system sizes $L > 7$, thereby allowing more than four iterations of the algorithm and a closer approach to the true value, the BST estimate of the extrapolant might yet more closely approach the de Gier-Essler result $a_0 = 3$ and give the single solution $c = 1$.

6.6 Reinterpreting the dynamical transition

The results of section 6.5 make a strong case for the hypothesis that the dynamical transition coincides with a crossing, for large L , of the gap and the second gap. In what follows I consider repercussions of the crossing scenario and attempt to place it in the context of other spectral behaviour, including that at the static transitions. In particular I argue for an interpretation that is consistent with other information about the dynamical transition, including the physical picture given by looking purely at the thermodynamic behaviour of the gap in isolation from the rest of the spectrum.

As we saw in section 6.1, in the limit $L \rightarrow \infty$ the second gap is degenerate with the gap for *all* values of the boundary parameters outside the maximal current regime. The eigenvalue crossing result indicates that for L large but finite the same degeneracy still holds, but only for a restricted curve in the α - β plane corresponding, to leading order, to the location of the dynamical transition.

It is worth comparing the behaviour of the spectrum at the dynamical transition to that at the first-order static transition at $\alpha = \beta$. There the first L excited states correspond to a band of diffusive modes that are nondegenerate for finite L , but which vanish ([18, 83]) for $L \rightarrow \infty$ to give an $(L + 1)$ -fold degenerate ground state. For the static transition we therefore have a thermodynamic degeneracy, which we know for the TASEP [83] corresponds to the coexistence of two stationary phases, spatially separated by a domain wall.

If we interpret the finite-size degeneracy at the dynamical transition in the same vein, rather than as a merely formal, physically insignificant precursor of the thermodynamic equivalence of gap and second gap, then this suggests the genuine coexistence of two distinct relaxational mechanisms. Assuming that the gap and second gap modes do each correspond to a different kind of contribution to the relaxation process, for example with different spatial localisation, then the crossing suggests a change in which of the two mechanisms constitutes the dynamical ‘bottleneck’ and hence determines the longest relaxation time. Since the dependence of the second gap on the boundary parameters closely mimics that of the gap, and in fact matches it exactly for $L \rightarrow \infty$, the difference in the two mechanisms may be subtle.

Nonetheless, I feel the results of this chapter warrant some degree of speculation as to the possible nature of the two distinct relaxational mechanisms, as follows. First, remember that in section 6.3.1 we saw that of the two unperturbed eigenvalues closest to

zero, λ_β is associated with particles leaving the system at the right boundary, and λ_1 with particles hopping within the bulk. Furthermore, that the seemingly successful rederivation of the dynamical transition in section 6.5 is based on the assumption that the transition corresponds to a crossing of the top branch of the λ_β cluster, which constitutes the gap below the transition, with the top branch of the λ_1 cluster, which gives the gap above. Now, although the clear-cut argument that gave the unperturbed eigenvalues directly as transition rates associated with particular configurations no longer holds for nonzero α , provided α is small the eigenvalues change gradually as small modifications to λ_β and λ_1 , lifting their degeneracies in the process and giving two clusters. I speculate that although eigenvalues in these clusters do depend on both α and β , in the perturbative regime those in the λ_β cluster continue to be associated with relaxation at the right boundary, and those in the λ_1 cluster with relaxation in the bulk.

Based on the eigenvalue crossing result I therefore postulate that below the transition the relaxational mechanism possessing the longest relaxation time (the gap mechanism) takes place at or close to the right boundary, whereas above the transition the mechanism associated with the gap corresponds to relaxation in the bulk. This claim, which states that relaxation at the right boundary ceases to provide the dynamical ‘bottleneck’ mechanism above the transition, is consistent with the primary signature of the dynamical transition, namely the independence, above β_c , of the thermodynamic gap from β . In chapter 7 I discuss this interpretation in the context of other findings in the thesis, and refer back to it in order to propose a possible microscopic explanation for the dynamical transition.

Chapter 7

Conclusions

This thesis has been concerned entirely with the study of the dynamical transition encountered in [18] in an exact Bethe ansatz solution for the gap, i.e. the longest relaxation time, in the totally asymmetric exclusion process with open boundaries. In this final chapter I summarise in section 7.1 the main findings that have resulted from studying the transition using four different approaches, namely: mean-field analysis, Monte Carlo simulation, numerical density matrix renormalisation group (DMRG) calculation, and a perturbative approach. In section 7.2 I speculate on the nature of the transition based on the findings, and suggest further investigations that may shed light on it.

7.1 Summary of findings

In chapter 3 I determined the thermodynamic gap behaviour in the mean field by analysing the longest relaxation time obtained from solutions to the Burgers equation, matched to variable densities at the boundaries. The mean-field gap was found to exhibit the same signature of a dynamical transition as the exact gap, though at $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{2}$ just like the domain wall theory (DWT) gap, rather than at the exact locations $\alpha_c(\beta)$, $\beta_c(\alpha)$. As I discuss in section 7.2, this provides a basis of comparison for DWT, so that we may better understand its disagreement with exact results. In addition, the thermodynamic mean-field gap was found to be finite, in agreement with the exact result and correcting a mean-field argument to the contrary in [102]. The soliton solution for the density profile, consisting of a stable, sharp domain wall traversing the system at constant velocity, was found to take a subtly different form above the dynamical transition. I discuss in section 7.2 to what extent this may help us interpret the exact transition. It would be

worthwhile to show from the solutions that, as expected, the velocity of the Burgers domain wall is the same as that predicted by DWT.

Preliminary Monte Carlo simulations showed no immediately obvious qualitative difference in the long-time relaxation above compared to below the dynamical transition, which in both cases was observed to consist, on a coarse-grained scale, of fluctuations in the position of a domain wall localised at one of the boundaries of the system. Seeking concrete confirmation of the existence and location of the transition, I undertook in chapter 4 a detailed quantitative study of the relaxation behaviour of these fluctuations by examining three different quantities. Extensive Monte Carlo simulations showed that in the long-time limit, the relaxation of an initial system-wide density excess (or deficit), the autocorrelation of fluctuations in the total occupancy, and the relaxation of a single second class particle to its stationary position all agree reasonably well with the exact gap behaviour. Unfortunately, as explained in that chapter, the method used to extract gap values from any measurement, though seemingly the best alternative available, suffers from an intrinsic limitation that leads to a seemingly incorrigible systematic overestimation of the gap. This, as well as statistical errors, precluded obtaining a sufficiently accurate estimate of the gap to completely convincingly discriminate between the location of the exact gap and that of the gap predicted by DWT. At the end of the chapter 4 I show that the thermodynamic gap computed by making a Markovian approximation for the motion of the second class particle exhibits a dynamical transition at $\alpha = \frac{1}{2}$ and at $\beta = \frac{1}{2}$. As I discuss in section 7.2, this may also help interpret the dynamical transition and provide further evidence for the reason behind the disagreement of DWT with exact dynamics.

The density matrix renormalisation group (DMRG), a numerical method for computing low-lying eigenvalues for large systems, proved in chapter 5 to offer a more fruitful approach. Using the infinite-system method I obtained estimates of the gap that agreed convincingly with the exact gap. As noted there, there is further room for improvement of accuracy by use of the finite-system method.

Finally, in chapter 6 I focused on the potential involvement of higher excitations in the dynamical transition. Guided by the possibility, on the basis of [18], of an asymptotic crossing scenario for the exact gap and second gap, I proposed a perturbative rederivation of the location of the dynamical transition. By obtaining perturbative expressions for the candidate crossing eigenvalues for a number of small system sizes and asking where they predict an asymptotic crossing, I showed that the crossing hypothesis is consistent with the location of the dynamical transition. This suggests a possible interpretation of the

dynamical transition, which I discuss in section 7.2.

7.2 The dynamical transition: interpretation and suggested investigations

In [169] Schütz considered general features of a family of hard-core reaction-diffusion models, which includes the asymmetric exclusion process as well as related processes that additionally include creation and annihilation of particles. Subsequently the authors of [170] presented an analysis that showed that this family of models generically exhibits dynamical phase transition behaviour characterised by the longest relaxation time becoming independent of the boundary injection and extraction rates in a certain regime. The same authors also prove similar boundary-driven dynamical transitions in special cases, such as for asymmetric Glauber dynamics [171], the q -state Potts model [172], and the kinetic Ising model [173]. Surprisingly, they conclude in [174] that the ASEP does not feature a dynamical transition of this type, in direct contradiction to the Bethe ansatz results by de Gier and Essler and in disagreement with the findings of this thesis. Their analysis, particular that in [174], is therefore very much worth of scrutiny. It is conceivable that the approach taken there fails to capture the true underlying dynamics because it is based on a representation of state space in terms of linear combinations of shock-measures, which may not be an appropriate decomposition for the full dynamical problem. Nonetheless, given that they do find dynamical transitions in a wide class of related models, the existence of a dynamical transition in the TASEP and indeed in the PASEP as found in [74] comes as no surprise. However in all the above cases the physics associated with the dynamical transition is not understood. The results of this thesis may constitute the first case for which one hopes one can say something more about the nature of such a dynamical transition, which appears to be a general feature of boundary-driven systems in one dimension.

Chapters 4 and 5 confirmed the existence and location of the dynamical transition in the open-boundary TASEP. They provided evidence – DMRG results more so than Monte Carlo simulations – that the primary signature of the transition, namely the constancy of the thermodynamic gap, takes place at $\alpha_c(\beta)$ and $\beta_c(\alpha)$ as predicted by de Gier and Essler, rather than at $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{2}$, which are the domain wall theory (DWT) predictions and which I confirmed in chapter 3 also correspond to the dynamical transition in the mean-field, continuum approximation.

Monte Carlo simulations are often the most straightforward way to directly observe the

behaviour of a model and hence gain intuitive insight. In chapter 4, the only immediate observation from Monte Carlo simulation relevant to the interpretation of the dynamical transition was that the long-time relaxation is associated with density fluctuations at one of the boundaries which, on a coarse-grained scale, can be interpreted as fluctuations in the position of a domain wall localised at that boundary. Of the three detailed sets of Monte Carlo measurements, only the agreement with the exact gap of the third – the ensemble-averaged relaxation of a single second class particle (SCP) to its average stationary position – offers any further insight into the precise nature of the transition. Indeed the choice of this observable to track the fluctuations was inspired by the identification, in [144], of the position of a single SCP with the location of a shock or domain wall in a single non-coarse-grained realisation of the stochastic dynamics of the ASEP on a finite lattice. At the end of chapter 4 we find that a Markovian approximation for the motion of the SCP shifts the location of the associated dynamical transition, which we know from simulation agrees with the exact result, to the location predicted by mean-field and domain wall theory. This is not particularly surprising given that the Markovian approximation also disregards correlations. However by virtue of the identification of the SCP motion with that of the domain wall, and the fact that measurements that agreed with the exact gap came from when the SCP was close to the boundary, the above does suggest that at the basis of any successful modification of DWT to fully reproduce the exact dynamics there must be an accurate account of density fluctuations at the boundaries at late times.

This conclusion is supported by two other findings. First, that although DWT is exact for $\alpha < \frac{1}{2}$ and $\beta < \frac{1}{2}$ whereas, as I showed in chapter 3, mean-field dynamics are only qualitatively correct, the two have in common a dynamical transition at $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{2}$. Second, as was already pointed out by de Gier and Essler in [18], that DWT can be made to agree with the exact dynamics simply by extending the critical value at which the low-density|high-density domain wall picture crosses over to the low-density|maximal-current or maximal-current|high-density picture, from $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{2}$ to $\alpha_c(\beta)$ and $\beta_c(\alpha)$. In other words, all that DWT appears to get wrong is the location of the transition.

To summarise, DWT is almost exact, and I argue that the only reason it is not exact is because it relies on a purely mean-field argument [55, 84] to describe the coupling between the boundaries and the bulk. According to this argument the inability of kinematic waves to penetrate the system from the boundary whose driving parameter α or β exceeds $\frac{1}{2}$ (the “overdriven” side) is responsible for the effective decoupling of that boundary from the dynamics and hence, in retrospect, for the inexact location of the dynamical transition. Appropriately modifying DWT to take into account the true

nature of the coupling to the boundaries is bound to be subtle, more subtle than the above suggestion to simply postpone the crossover to the low-density|maximal-current or maximal-current|high-density picture until the exact critical value, since apart from being completely unjustified, this modification destroys the characterisation of the static subphase distinctions within the HD and LD phases, which DWT obtains in agreement with the exact stationary state [1, 55].

In chapter 6 I found from a perturbative approach that the dynamical transition feasibly corresponds to a crossing of the gap and second gap. I proposed that this amounts to a crossover between two different relaxational mechanisms, each associated with one of the modes. In this interpretation the transition point is where one mechanism takes over from the other as the ‘bottleneck’ dynamics that determines the longest relaxation time in the system. I also tentatively identify one mechanism as associated with relaxation at the boundary, and the other, which dominates above the transition, with relaxation in the bulk. At first sight the finding from Monte Carlo simulation that the long-time relaxation corresponds to density fluctuations near the boundary both above and below the transition somewhat weakens the argument that one of the mechanisms is associated with relaxation in the bulk. However, this may still be consistent if all the crossing result implies is that the relaxation does not take place exactly *at* the boundary, but at a finite distance away as $L \rightarrow \infty$.

On the basis of the above I now speculatively propose a possible microscopic interpretation of the dynamical transition, which attempts both to explain the constancy of the gap with respect to one of the boundary parameters, and which, inspired by the crossing result, involves a crossover between two different relaxational mechanisms. It is an intuitive argument, like the one originally given in [84] for the formation of a domain wall. Consider, as for the perturbative scenario in chapter 6, that α is very small and β close to 1, so that effectively only a single particle is on the lattice at any given time. Any particle that is injected travels uninterrupted to the right boundary, where it exits with rate β . The longest relaxation time is $1/\beta$. Now consider α larger, though still smaller than $\frac{1}{2}$. There is now a finite probability of having more than one particle on the lattice at any given time. For β ‘too’ small, i.e. $\beta < \beta_c(\alpha)$, the particles tend to accumulate at the right boundary, where the long-time relaxation therefore takes place, so that the longest relaxation time is still a decreasing function of β . However for β sufficiently large, the escape of particles at the boundary is no longer the dynamical bottleneck: congested transport through the bulk of the system is. As a result, increasing β further whilst keeping α constant no longer decreases the longest relaxation time, which is now determined by the congestion a finite distance away from the boundary and which in an

infinite system is therefore completely independent of β .

It is not clear how compatible these proposed microscopics are with the soliton solution of the Burgers equation above the mean-field transition line, which was found, in the LD regime, to consist of a significantly higher density than the right boundary density (a ‘hump’) appearing immediately next to the shock and extending most of the distance to the right boundary. In the long time regime the shock travels to the boundary, where in the mean field it disappears, but where in reality it persists to stochastically reenter the system and make excursions of varying length away from the boundary and into the system. Since the hump is of higher density than the boundary this suggests congestion immediately next to the boundary, in contradiction with the proposed microscopic mechanism.

However it may be the case that the analogue of the mean-field soliton but then in the exact dynamics, observable in simulation after coarse-graining, provides more insight. It should be possible to access this feature using the DMRG method, which here confirmed the transition but offered no insight into the underlying physics, by computing renormalised approximations of the thermodynamic left and right gap eigenvectors and of an operator that uses these to compute a density profile. In [175] this scheme was realised for the stationary density profile of the ASEP, and shows excellent agreement with the exact solution given by the matrix product ansatz. I see no *a priori* reason why it should not be possible to do the same for the first excited state, and thereby to calculate the longest-lived purely transient density profile. This would constitute the exact, discrete-space analogue of the elementary soliton excitation found in the Burgers equation and discussed in chapter 3.

As a basis for better understanding the true nature of the coupling of the bulk to the boundaries one would want to investigate the dependence on the boundary parameters of the density fluctuations at the boundaries. This could be done by an empirical characterisation, using Monte Carlo simulation, of certain statistical properties of the fluctuations, or could proceed analytically starting from the model dynamics. This may also allow it to be confirmed to what extent the microscopic mechanism proposed above plays a role in determining the transition. Furthermore, one could investigate in what way the dynamical transition is modified when the TASEP model is varied, such as by varying the hardness of the exclusion, i.e. allowing more than a single particle to occupy a given site, or by extending the range of the exclusion interaction to next nearest neighbour rather than nearest neighbour sites. This might give insight not only into the origin of the transition, but also into its robustness and in what sense, if any, it is a universal feature of boundary-driven systems in one dimension.

One intriguing side issue that was raised in the study of the dynamical transition regards differences in the spectral manifestation of static and dynamical phase transitions and the connection, if any, between them. Especially relevant to this discussion is the connection, which holds for all finite state-space Markov processes, between the transition matrix spectrum and the normalisation Z of the steady state weights, which for an equilibrium system defines the partition function [12, 176]:

$$Z_L = \prod_{\lambda_j \neq 0} (-\lambda_j) \quad (7.1)$$

where $(-\lambda_j)$ are the eigenvalues of the transition matrix. In [177, 178] the static phase transitions of the open-boundary TASEP were found to occur when the analogue of the free energy, given by the usual equilibrium thermodynamic relation to the partition function applied to (7.1),

$$f_L = \frac{\ln Z_L}{L} = \frac{1}{L} \sum_{i \neq 0} \ln(-\lambda_i), \quad (7.2)$$

exhibits nonanalyticities. The orders of the transition could be determined by considering the Lee-Yang zeroes [177, 178]. However no such behaviour was observed at the dynamical transition. Hence one wonders about the spectral distinction between the static and dynamical transitions in the TASEP, and whether this is in any sense a general feature.

As a starting point to understanding this phenomenon, we may consider an arbitrary eigenvalue λ_i in (7.2), possessing a nonanalyticity at a critical value of some parameter. Unless a significant number of other eigenvalues in the spectrum converge onto the same nonanalyticity in the thermodynamic limit, the effect of the λ_i nonanalyticity is lost as $L \rightarrow \infty$ and (7.2) and hence f_L remain analytic. Given the detectability of the static transitions but the absence of thermodynamic nonanalyticity for the dynamical transition, I suggest therefore that one expects a static transition to correspond to a significant spectral reordering that gives rise to nonanalyticities not just in the gap but in a finite fraction of the spectrum as $L \rightarrow \infty$, and that any spectral nonanalyticities at the dynamical transition, must for $L \rightarrow \infty$ be restricted to a set of eigenvalues with measure zero within the spectrum. Within a crossing scenario like the one that seems feasible for the first two TASEP gaps, both eigenvalues remain analytic. However this does not explain the analyticity of the free energy at the dynamical transition, since we know that for $L \rightarrow \infty$ both the gap and the second gap are nonanalytic not only at the static phase transitions but also at the dynamical transition. The question of whether there is any spectral signature of the dynamical transition, such as eigenvalue crossings, to be found in higher states remains open and may be answered only by studying properties of the entire thermodynamic spectrum. The apparent feasibility of the crossing of gap

and second gap at the dynamical transition, even though it occurs for large but finite L rather than in the thermodynamic limit, suggests a possible classification of dynamical transitions based on eigenvalue degeneracy, along the lines of [179] and, more recently, [180]. It would be interesting to investigate more broadly the connection between spectral degeneracies, crossings, and static as well as dynamical phase transitions.

Bibliography

- [1] A. Proeme, R.A. Blythe, and M.R. Evans. Dynamical transition in the open-boundary totally asymmetric exclusion process. *Journal of Physics A: Mathematical and Theoretical*, 44(3):035003, 2011.
- [2] M. Cross and H. Greenside. *Pattern Formation and Dynamics in Nonequilibrium Systems*. Cambridge University Press, 2009.
- [3] J.D. Murray. *Mathematical Biology II: Spatial Models and Biomedical Applications*, volume 2. Springer, 3rd edition, 2004.
- [4] H.M. Jaeger, S. Nagel, and R.P. Behringer. Granular solids, liquids, and gases. *Reviews of Modern Physics*, 68, 1996.
- [5] K. Christensen and N. Moloney. *Complexity and Criticality*. Imperial College Press, 2005.
- [6] A.C.C. Coolen, R. Kühn, and P. Sollich. *Theory of Neural Information Processing Systems*. Oxford University Press, 2005.
- [7] R. Stinchcombe. Stochastic non-equilibrium systems. *Advances in Physics*, 50(5):431–496, 2001.
- [8] D. Helbing. Traffic and related self-driven many-particle systems. *Reviews of Modern Physics*, 73(4):1067–1141, December 2001.
- [9] S. Kauffman. *At Home in the Universe*. Oxford University Press, 1995.
- [10] O. Penrose. *Foundations of Statistical Mechanics*. Dover, 2005.
- [11] L. Sklar. *Physics and Chance*. Cambridge University Press, 1993.
- [12] R.A. Blythe and M.R. Evans. Nonequilibrium steady states of matrix-product form: a solver’s guide. *Journal of Physics A: Mathematical and Theoretical*, 40(46):R333–R441, 2007.
- [13] S.L. Starr. Some properties for the low-lying spectrum of the ferromagnetic, quantum XXZ spin system. *arXiv:math-ph/0106024v1*, 2001.
- [14] B. Derrida. An exactly soluble non-equilibrium system: the asymmetric simple exclusion process. *Physics Reports*, 301(1):65–84, 1998.
- [15] G.M. Schütz. *Exactly Solvable Models for Many-Body Systems Far from Equilibrium*, volume 19 of *Phase Transitions and Critical Phenomena*. 2001.
- [16] O. Golinelli and K. Mallick. The asymmetric simple exclusion process: an integrable model for non-equilibrium statistical mechanics. *Journal of Physics A: Mathematical and General*, 39(41):12679–12705, 2006.
- [17] R. Stinchcombe. Primitive cooperative particle flow models. *Physica A: Statistical Mechanics and its Applications*, 372(1):1–11, 2006.

- [18] J. de Gier and F.H.L. Essler. Exact spectral gaps of the asymmetric exclusion process with open boundaries. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(12):P12011–P12011, 2006.
- [19] M.R. Evans and T. Hanney. Nonequilibrium statistical mechanics of the zero-range process and related models. *Journal of Physics A: Mathematical and General*, 38(19):R195–R240, 2005.
- [20] B. Schmittman and R.K.P. Zia. *Statistical Mechanics of Driven Diffusive Systems*, volume 17 of *Phase Transitions and Critical Phenomena*. Academic Press, 1995.
- [21] C.T. MacDonald, J.H. Gibbs, and A.C. Pipkin. Kinetics of biopolymerization on nucleic acid templates. *Biopolymers*, 6(1), 1968.
- [22] F. Spitzer. Interaction of Markov Processes. *Advances in Mathematics*, Vol. 5, 1970.
- [23] T.M. Liggett. Ergodic theorems for the asymmetric simple exclusion process. *Transactions of the American Mathematical Society*, 213, 1975.
- [24] T.M. Liggett. *Interacting Particle Systems*. Springer, 1985.
- [25] T.M. Liggett. *Stochastic Interacting Particle Systems: Contact, Voter and Exclusion Processes*. Springer, 1999.
- [26] D.A. Lavis and G.M. Bell. *Statistical Mechanics of Lattice Systems: Closed form and exact theories of cooperative phenomena*, volume 1. Springer, 1989.
- [27] D.A. Lavis and G.M. Bell. *Statistical Mechanics of Lattice Systems: Exact, Series and Renormalization Group Methods*, volume 2. Springer, 1999.
- [28] J. Krug. Boundary-induced phase transitions in driven diffusive systems. *Physical Review Letters*, 67(14):1882–1885, 1991.
- [29] F. Rezakhanlou. Hydrodynamic Limit for Attractive Particle Systems on \mathbb{Z}^d . *Communications in Mathematical Physics*, 140:417–448, 1991.
- [30] F. Rezakhanlou. Evolution of tagged particles in non-reversible particle systems. *Communications in Mathematical Physics*, 165(1):1–32, 1994.
- [31] F. Rezakhanlou. A central limit theorem for the asymmetric simple exclusion process. In *Annales de l'Institut Henri Poincaré/Probabilités et statistiques*, volume 38, pages 437–464. Elsevier, 2002.
- [32] T. Seppäläinen. Diffusive Fluctuations for One-Dimensional Totally Asymmetric Interacting Random Dynamics. *Communications in Mathematical Physics*, 229(1):141–182, 2002.
- [33] T. Nagao and T. Sasamoto. Asymmetric simple exclusion process and modified random matrix ensembles. *Nuclear Physics B*, 699(3):487–502, 2004.
- [34] T. Sasamoto. Fluctuations of the one-dimensional asymmetric exclusion process using random matrix techniques. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(07):P07007–P07007, 2007.

- [35] T. Imamura and T. Sasamoto. Dynamics of a Tagged Particle in the Asymmetric Exclusion Process with the Step Initial Condition. *Journal of Statistical Physics*, 128(4):799–846, 2007.
- [36] T. Sasamoto. Exact results for the 1D asymmetric exclusion process and KPZ fluctuations. *The European Physical Journal B*, 64(3-4):373–377, 2008.
- [37] Craig A. Tracy and Harold Widom. Formulas for ASEP with Two-Sided Bernoulli Initial Condition. *Journal of Statistical Physics*, 140(4):619–634, July 2010.
- [38] P.L. Ferrari and R. Frings. On the partial connection between random matrices and interacting particle systems. *arXiv:math-ph/1006.3946*, page 27, 2010.
- [39] M. Schreckenberg, A. Schadschneider, K. Nagel, and N. Ito. Discrete stochastic models for traffic flow. *Physical Review E*, 51(4):2939–2949, 1995.
- [40] T. Antal and G.M. Schütz. Asymmetric exclusion process with next-nearest-neighbor interaction: some comments on traffic flow and a nonequilibrium reentrance transition. *Physical Review E*, 62(1 Pt A):83–93, July 2000.
- [41] Y. Aghababaie, G. Menon, and M. Plischke. Universal properties of interacting Brownian motors. *Physical Review E*, 59(3):2578–2586, 1999.
- [42] S. Klumpp and R. Lipowsky. Traffic of Molecular Motors Through Tube-Like Compartments. *Journal of Statistical Physics*, 113(1/2), 2003.
- [43] K. Sugden and M.R. Evans. An exclusion process for modelling fungal hyphal growth. *Physica A: Statistical Mechanics and its Applications*, 384(1):53–58, 2007.
- [44] S. Chatterjee and G.M. Schütz. The importance of boundary effects in diffusion of hydrocarbon molecules in a one-dimensional zeolite channel. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(01):P01017, 2010.
- [45] P.A. Ferrari and J.B. Martin. Stationary distributions of multi-type totally asymmetric exclusion processes. *The Annals of Probability*, 35(3):807–832, May 2007.
- [46] M.R. Evans, P.A. Ferrari, and K. Mallick. Matrix Representation of the Stationary Measure for the Multispecies TASEP. *Journal of Statistical Physics*, 135(2):217–239, 2009.
- [47] N.G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 3rd edition, 2007.
- [48] E. Seneta. *Non-negative Matrices*. George Allen & Unwin, 1973.
- [49] B. Derrida, E. Domany, and D. Mukamel. An exact solution of a one-dimensional asymmetric exclusion model with open boundaries. *Journal of Statistical Physics*, 69(3):667–687, 1992.
- [50] J.L. Lebowitz, E. Presutti, and H. Spohn. Microscopic models of hydrodynamic behavior. *Journal of Statistical Physics*, 51(5):841–862, 1988.

- [51] G. Giacomin, J.L. Lebowitz, and E. Presutti. *Deterministic and Stochastic Hydrodynamic Equations Arising From Simple Microscopic Model Systems*, chapter 3, pages 1–46. *Stochastic partial differential equations: six perspectives*. American Mathematical Society, 1991.
- [52] G.L. Eyink, J.L. Lebowitz, and H. Spohn. Hydrodynamics and fluctuations outside of local equilibrium: driven diffusive systems. *Journal of Statistical Physics*, 83(3):385–472, 1996.
- [53] B. Derrida, M.R. Evans, V. Hakim, and V. Pasquier. Exact solution of a 1D asymmetric exclusion model using a matrix formulation. *Journal of Physics A: Mathematical and General*, 26:1493–1517, 1993.
- [54] G.M. Schütz. Generalized Bethe ansatz solution of a one-dimensional asymmetric exclusion process on a ring with blockage. *Journal of Statistical Physics*, 71(3-4):471–505, 1993.
- [55] A.B. Kolomeisky, G.M. Schütz, E.B. Kolomeisky, and J.P. Straley. Phase diagram of one-dimensional driven lattice gases with open boundaries. *Journal of Physics A: Mathematical and General*, 31:6911–6919, 1998.
- [56] B. Derrida and M.R. Evans. Exact correlation functions in an asymmetric exclusion model with open boundaries. *J. Phys. I France*, 3:311–322, 1993.
- [57] T. Hanney and R. Stinchcombe. Real-space renormalization group approach to driven diffusive systems. *Journal of Physics A: Mathematical and General*, 39(47):14535–14544, 2006.
- [58] G. Ódor. Universality classes in nonequilibrium lattice systems. *Reviews of modern physics*, 76(3):663–724, 2004.
- [59] R. Lehoucq, K. Maschhoff, D. Sorensen, and C. Yang. The ARPACK library.
- [60] M. Henkel and G.M. Schütz. Finite-lattice extrapolation algorithms. *Journal of Physics A: Mathematical and General*, 21:2617–2633, 1988.
- [61] J. Monroe. Extrapolation and the Bulirsch-Stoer algorithm. *Physical Review E*, 65(6):1–8, 2002.
- [62] R. Bulirsch and J. Stoer. *Introduction to numerical analysis*. Springer, 3 edition, 2002.
- [63] U. Bilstein and B. Wehefritz. Spectra of non-Hermitian quantum spin chains describing boundary induced phase transitions. *Journal of Physics A: Mathematical and General*, 30:4925–4938, 1997.
- [64] M. Kardar, G. Parisi, and Y.C. Zhang. Dynamic scaling of growing interfaces. *Physical Review Letters*, 56(9):889–892, 1986.
- [65] L.H. Gwa and H. Spohn. Bethe solution for the dynamical-scaling exponent of the noisy Burgers equation. *Physical Review A*, 46(2):844–854, 1992.
- [66] D. Kim. Bethe ansatz solution for crossover scaling functions of the asymmetric XXZ chain and the Kardar-Parisi-Zhang-type growth model. *Physical Review E*, 52(4):3512–3524, 1995.

- [67] O. Golinelli and K. Mallick. Bethe ansatz calculation of the spectral gap of the asymmetric exclusion process. *Journal of Physics A: Mathematical and General*, 37(10):3321–3331, 2004.
- [68] O. Golinelli and K. Mallick. Hidden symmetries in the asymmetric exclusion process. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(12):P12001, 2004.
- [69] O. Golinelli and K. Mallick. Spectral Degeneracies in the Totally Asymmetric Exclusion Process. *Journal of Statistical Physics*, 120(5-6):779–798, 2005.
- [70] W.D. Heiss. Exceptional points of non-Hermitian operators. *Journal of Physics A: Mathematical and General*, 37(6):2455–2464, 2004.
- [71] M. Bhattacharya and C. Raman. Detecting Level Crossings without Looking at the Spectrum. *Physical Review Letters*, 97(14):6–9, 2006.
- [72] M. Bhattacharya and C. Raman. Detecting level crossings without solving the Hamiltonian. I. Mathematical background. *Physical Review A*, 75(3):1–15, 2007.
- [73] J. de Gier and F.H.L. Essler. Bethe Ansatz Solution of the Asymmetric Exclusion Process with Open Boundaries. *Physical Review Letters*, 95(24):1–4, 2005.
- [74] J. de Gier and F.H.L. Essler. Slowest relaxation mode of the partially asymmetric exclusion process with open boundaries. *Journal of Physics A: Mathematical and Theoretical*, 41(48):485002, 2008.
- [75] J. Cao. Exact solution of XXZ spin chain with unparallel boundary fields. *Nuclear Physics B*, 663(3):487–519, 2003.
- [76] R.I. Nepomechie. Functional relations and Bethe Ansatz for the XXZ chain. *Journal of Statistical Physics*, 111(5):1363–1376, 2003.
- [77] R.I. Nepomechie and F. Ravanini. Completeness of the Bethe Ansatz solution of the open XXZ chain with nondiagonal boundary terms. *Journal of Physics A: Mathematical and General*, 36:11391–11401, 2003.
- [78] R.I. Nepomechie and F. Ravanini. Addendum to Completeness of the Bethe Ansatz solution of the open XXZ chain with nondiagonal boundary terms. *Journal of Physics A: Mathematical and General*, 37(5):1945–1946, 2004.
- [79] A. Nichols, V. Rittenberg, and J. de Gier. One-boundary TemperleyLieb algebras in the XXZ and loop models. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(03):P03003, 2005.
- [80] J. de Gier and P. Pyatov. Bethe ansatz for the TemperleyLieb loop model with open boundaries. *Journal of Statistical Mechanics: Theory and Experiment*, 2004():P002, 2004.
- [81] R. Murgan and R.I. Nepomechie. Bethe ansatz derived from the functional relations of the open XXZ chain for new special cases. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(05):P05007, 2005.
- [82] R. Murgan and R.I. Nepomechie. Generalized T Q relations and the open XXZ chain. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(08):P08002–P08002, 2005.

- [83] G.M. Schütz and M. Dudzinski. Relaxation spectrum of the asymmetric exclusion process with open boundaries. *Journal of Physics A: Mathematical and General*, 33:8351, 2000.
- [84] G.M. Schütz and E. Domany. Phase transitions in an exactly soluble one-dimensional exclusion process. *Journal of Statistical Physics*, 72(1-2):277–296, 1993.
- [85] B. Derrida, M.R. Evans, and K. Mallick. Exact diffusion constant of a one-dimensional asymmetric exclusion model with open boundaries. *Journal of Statistical Physics*, 79(5):833–874, 1995.
- [86] J.M. Burgers. *The nonlinear diffusive equation*. D. Reidel Publishing Company, Dordrecht, 1974.
- [87] S. Kida. Asymptotic properties of Burgers turbulence. *Journal of Fluid Mechanics*, 93(02):337–377, 1979.
- [88] T.J. Newman. Dynamical scaling in dissipative Burgers turbulence. *Physical Review E*, 55(6):6989–6999, 1997.
- [89] T. Kriecherbauer and J. Krug. Interacting particle systems out of equilibrium. *arXiv:0803.2796*, page 35, 2008.
- [90] H.C. Fogedby, A. Eriksson, and L. Mikheev. Continuum Limit, Galilean Invariance, and Solitons in the Quantum Equivalent of the Noisy Burgers Equation. *Physical Review Letters*, 75(10):1883–1886, 1995.
- [91] H.C. Fogedby. Nonequilibrium dynamics of a growing interface. *Journal of Physics: Condensed Matter*, 14:1557–1569, 2002.
- [92] H.C. Fogedby. Solitons and diffusive modes in the noiseless Burgers equation: Stability analysis. *Physical Review E*, 57(2):2331–2337, 1998.
- [93] H.C. Fogedby. Dynamical pinning and non-Hermitian mode transmutation in the Burgers equation. *The European Physical Journal B*, 20(2):153–156, 2001.
- [94] P.L. Sachdev. *Nonlinear Diffusive Waves*. Cambridge University Press, 1987.
- [95] P.L. Sachdev and C.S. Rao. *Large Time Asymptotics for Solutions of Nonlinear Partial Differential Equations*. Springer, 2009.
- [96] B. Derrida, S.A. Janowsky, J.L. Lebowitz, and E. R. Speer. Exact solution of the totally asymmetric simple exclusion process: shock profiles. *Journal of Statistical Physics*, 73(5):813–842, 1993.
- [97] G.M. Schütz. Critical phenomena and universal dynamics in one-dimensional driven diffusive systems with two species of particles. *Journal of Physics A: Mathematical and General*, 36(36):R339–R379, 2003.
- [98] G.M. Schütz. Exact Shock Measures and Steady-State Selection in a Driven Diffusive System with Two Conserved Densities. *Journal of Statistical Physics*, 117:55–76, 2004.
- [99] P.A. Ferrari. Shock fluctuations in asymmetric simple exclusion. *Probability Theory and Related Fields*, 91(1):81–101, March 1992.

- [100] P.A. Ferrari and L.R. Fontes. Shock fluctuations in the asymmetric simple exclusion process. *Probability Theory and Related Fields*, 319, 1994.
- [101] I. Corwin, P.L. Ferrari, and S. P ech e. Limit Processes for TASEP with Shocks and Rarefaction Fans. *Journal of Statistical Physics*, June 2010.
- [102] S. de Queiroz and R. Stinchcombe. Nonequilibrium processes: Driven lattice gases, interface dynamics, and quenched-disorder effects on density profiles and currents. *Physical Review E*, 78(3):1–11, 2008.
- [103] L. Debnath. *Nonlinear Diffusion-Reaction Phenomena, Burgers’ and Fisher’s Equations*. Nonlinear Partial Differential Equations for Scientists and Engineers. Birkh user, 1997.
- [104] W.A. Woyczynski. *Hopf-Cole Formula and its Asymptotic Analysis*, chapter 3. Number 3 in Burgers-KPZ Turbulence, Lecture Notes in Mathematics. Springer, 1998.
- [105] J.D. Cole. On a quasilinear parabolic equation occurring in aerodynamics. *Quart. Appl. Math.*, 9:225–236, 1951.
- [106] E. Hopf. The partial differential equation $u_t + uu_x = u_{xx}$. *Communications in Pure and Applied Mathematics*, 3:201–230, 1950.
- [107] R.J. Harris and R. Stinchcombe. Disordered asymmetric simple exclusion process: Mean-field treatment. *Physical Review E*, 70(1):1–15, 2004.
- [108] C. Harper. *Analytic methods in physics*. Wiley, 1999.
- [109] P. Pierobon, A. Parmeggiani, F. von Oppen, and E. Frey. Dynamic correlation functions and Boltzmann-Langevin approach for driven one-dimensional lattice gas. *Physical Review E*, 72(3):1–10, 2005.
- [110] D. Adams, R.K.P. Zia, and B. Schmittmann. Power Spectra of the Total Occupancy in the Totally Asymmetric Simple Exclusion Process. *Physical Review Letters*, 99(2):1–4, 2007.
- [111] J. de Gier, T. Garoni, and Z. Zhou. Autocorrelations in the totally asymmetric simple exclusion process and Nagel-Schreckenberg model. *Physical Review E*, 82(2):9, 2010.
- [112] N. Rajewsky, L. Santen, A. Schadschneider, and M. Schreckenberg. The asymmetric exclusion process: Comparison of update procedures. *Journal of Statistical Physics*, 92(1):151–194, 1998.
- [113] D.T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, December 1976.
- [114] D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [115] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems 1. *Journal of Computational Physics*, 17(1):10–18, 1975.

- [116] K. Binder. Applications of Monte Carlo methods to statistical physics. *Reports on Progress in Physics*, 60:487, 1997.
- [117] D.P. Landau and K. Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge University Press, 2000.
- [118] M. Newman and G. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 1999.
- [119] H. Muller-Krumbhaar and K. Binder. Dynamic properties of the Monte Carlo method in statistical mechanics. *Journal of Statistical Physics*, 8(1):1–24, May 1973.
- [120] K. Kawasaki. *Kinetics of Ising Models*, volume 2 of *Phase Transitions and Critical Phenomena*. 1972.
- [121] R.J. Glauber. Time-Dependent Statistics of the Ising Model. *Journal of Mathematical Physics*, 4(2):294, 1963.
- [122] A.P.J. Jansen. An Introduction To Monte Carlo Simulations Of Surface Reactions. *arXiv:cond-mat/0303028v1*, 2008.
- [123] J. Lukkien, J. Segers, P. Hilbers, R. Gelten, and A. Jansen. Efficient Monte Carlo methods for the simulation of catalytic surface reactions. *Physical Review E*, 58(2):2598–2610, 1998.
- [124] P. Kratzer. Monte Carlo and kinetic Monte Carlo methods - a tutorial. *arXiv:0904.2556*, 2009.
- [125] A.F. Voter. *Radiation Effects in Solids*, chapter 1, pages 1–23. Springer, 2007.
- [126] M.A. Gibson and J. Bruck. Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels. *Journal of Physical Chemistry A*, 2:1876–1889, 2000.
- [127] K.A. Fichthorn. Theoretical foundations of dynamical Monte Carlo simulations. *Journal of Chemical Physics*, 95(2), 1991.
- [128] J. Blue, I. Beichl, and F. Sullivan. Faster Monte Carlo simulations. *Physical Review E*, 51(2):R867–R868, 1995.
- [129] T. Schulze. Kinetic Monte Carlo simulations with minimal searching. *Physical Review E*, 65(3):1–3, 2002.
- [130] T. Opperstrup, V.V. Bulatov, G.H. Gilmer, M.H. Kalos, and B. Sadigh. First-Passage Monte Carlo Algorithm: Diffusion without All the Hops. *Physical Review Letters*, 97(23):10–13, December 2006.
- [131] A. Slepoy, A.P. Thompson, and S.J. Plimpton. A constant-time kinetic Monte Carlo algorithm for simulation of large biochemical reaction networks. *The Journal of Chemical Physics*, 128(20):205101, 2008.
- [132] K. Holmström. A review of the parameter estimation problem of fitting positive exponential sums to empirical data. *Applied Mathematics and Computation*, 126(1):31–61, February 2002.

- [133] W.J. Wiscombe and J.W. Evans. Exponential-sum fitting of radiative transmission functions. *Journal of Computational Physics*, 24:416–444, 1977.
- [134] D.P. Landau and K. Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*, page 101. Cambridge University Press, 2000.
- [135] S. Wansleben and D.P. Landau. Monte Carlo investigation of critical dynamics in the three-dimensional Ising model. *Physical Review B*, 43(7):6006–6014, 1991.
- [136] H. Heuer. Critical slowing down in local dynamics simulations. *Journal of Physics A: Mathematical and General*, 567, 1992.
- [137] H. Heuer. Critical dynamics of the three-dimensional Ising model: A Monte Carlo study. *Journal of Statistical Physics*, 72(3-4):789–827, 1993.
- [138] T.P. Ryan. *Modern Regression Methods*. Wiley series in probability and statistics, 1997.
- [139] J.K. Williams. Monte Carlo estimate of the dynamical critical exponent of the 2d kinetic Ising. *Journal of Physics A: Mathematical and General*, 18:49–60, 1985.
- [140] S. Tang and D.P. Landau. Monte Carlo study of dynamic universality in two-dimensional Potts models. *Physical Review B*, 36(1):567–573, 1987.
- [141] S. Wansleben and D.P. Landau. Dynamical critical exponent of the 3D Ising model. *Journal of Applied Physics*, 61(8):3968, 1987.
- [142] C. Boldrighini, G. Cosimi, S. Frigio, and M. Grasso Nunes. Computer simulation of shock waves in the completely asymmetric simple exclusion process. *Journal of Statistical Physics*, 55(3):611–623, 1989.
- [143] B. Derrida, J.L. Lebowitz, and E. R. Speer. Shock profiles for the asymmetric simple exclusion process in one dimension. *Journal of Statistical Physics*, 89(1-2):135–167, 1997.
- [144] K. Krebs, F.H. Jafarpour, and G.M. Schütz. Microscopic structure of travelling wave solutions in a class of stochastic interacting particle systems. *New Journal of Physics*, 5:145, 2003.
- [145] S.R. White. Strongly correlated electron systems and the density matrix renormalization group. *Physics Reports*, 301(1-3):187–204, 1998.
- [146] U. Schollwöck. The density-matrix renormalization group. *Reviews of Modern Physics*, 77(1):259–315, 2005.
- [147] I. Peschel, M. Kaulke, X. Wang, and K. Hallberg, editors. *Density-Matrix Renormalization. A New Numerical Method in Physics*. Lecture Notes in Physics. Springer, 1998.
- [148] E. Carlon, M. Henkel, and U. Schollwöck. Density matrix renormalization group and reaction-diffusion processes. *The European Physical Journal B*, 12(1):99–114, 1999.
- [149] Z. Nagy, C. Appert, and L. Santen. Relaxation times in the ASEP model using a DMRG method. *Journal of Statistical Physics*, 109(3):623–639, 2002.

- [150] S.R. White and R.M. Noack. Real-space quantum renormalization groups. *Physical Review Letters*, 68(24):3487–3490, 1992.
- [151] S. Chui and J. Bray. Computer renormalization-group technique applied to the Hubbard model. *Physical Review B*, 18(6):2426–2430, 1978.
- [152] S.R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69(19):2863–2866, 1992.
- [153] R.P. Feynman. *Statistical Mechanics: a set of lectures*. Westview Press, 1998.
- [154] S.R. White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48(14):10345–10356, 1993.
- [155] R.M. Noack. Diagonalization- and Numerical Renormalization-Group-Based Methods for Interacting Quantum Systems. In *AIP Conference Proceedings*, volume 1, pages 93–163, arXiv:cond-mat/510321v1, 2005.
- [156] G.H. Golub and C.F. van Loan. *Matrix Computations*. Johns Hopkins University Press, third edition, 1996.
- [157] E. Carlon, M. Henkel, and U. Schollwöck. Critical properties of the reaction-diffusion model $2A \rightarrow 3A$, $2A \rightarrow 0$. *Physical Review E*, 63(3):1–10, 2001.
- [158] The LAPACK library. <http://www.netlib.org/lapack/>.
- [159] Ö. Legeza and G. Fáth. Accuracy of the density-matrix renormalization-group method. *Physical Review B*, 53(21):14349–14358, 1996.
- [160] J.H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon, Oxford, 1965.
- [161] M. Reed and B. Simon. *Methods of Modern Mathematical Physics IV: Analysis of operators*. Academic Press, 1978.
- [162] P. Lancaster and M. Tismenetsky. *The theory of matrices*. Academic Press, San Diego, 2 edition, 1985.
- [163] F. Rellich. *Perturbation Theory of Eigenvalue Problems*. Gordon & Breach, 1969.
- [164] T. Kato. *Perturbation Theory for Linear Operators*. Springer-Verlag, Berlin, 1980.
- [165] H. Baumgärtel. *Analytic Perturbation Theory for Matrices and Operators*. Birkhäuser-Verlag, Basel, 1985.
- [166] J. Moro, J.V. Burke, and M.L. Overton. On the Lidskii-Vishik-Lyusternik perturbation theory for eigenvalues of matrices with arbitrary Jordan structure. *SIAM Journal on Matrix Analysis and Applications*, 18(4):793–817, 1997.
- [167] V.B. Lidskii. Perturbation theory of non-conjugate operators. *U.S.S.R. Comput. Math. and Math. Phys.*, 1:73–85, 1965.
- [168] V.B. Lidskii. Perturbation theory of non-conjugate operators. *Zh. vychisl. Mat. mat. Fiz.*, 6:52–60, 1965.
- [169] G.M. Schütz. Reaction-diffusion processes of hard-core particles. *Journal of Statistical Physics*, 79(1-2):243–264, 1995.

- [170] A. Aghamohammadi and M. Khorrami. Phase transitions in autonomous reaction-diffusion systems on a one-dimensional lattice with boundaries. *Journal of Physics A: Mathematical and General*, 34:7431, 2001.
- [171] M. Khorrami and A. Aghamohammadi. Phase transition in an asymmetric generalization of the zero-temperature Glauber model. *Physical Review E*, 63(4):1–4, 2001.
- [172] N. Majd, A. Aghamohammadi, and M. Khorrami. Phase transition in an asymmetric generalization of the zero-temperature q-state Potts model. *Physical Review E*, 64(4):1–5, 2001.
- [173] M. Khorrami and A. Aghamohammadi. Dynamical phase transition in one-dimensional kinetic Ising model with nonuniform coupling constants. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(10):P10019, October 2010.
- [174] M. Arabsalmani and A. Aghamohammadi. Phase transitions in systems possessing shock solutions. *Physical Review E*, 74(1), 2006.
- [175] Y. Hieida. Application of the Density Matrix Renormalization Group Method to a Non-Equilibrium Problem. *arXiv:cond-mat/9711072*, page 8, 1997.
- [176] M.R. Evans and R.A. Blythe. Nonequilibrium dynamics in low-dimensional systems. *Physica A: Statistical Mechanics and its Applications*, 313(1-2):110–152, 2002.
- [177] R.A. Blythe and M.R. Evans. Lee-Yang Zeros and Phase Transitions in Nonequilibrium Steady States. *Physical Review Letters*, 89(8):2–5, 2002.
- [178] R.A. Blythe and M.R. Evans. The Lee-Yang theory of equilibrium and nonequilibrium phase transitions. *Brazilian Journal of Physics*, 33(3):464–475, 2003.
- [179] M. Kac and C.J. Thompson. On the mathematical mechanism of phase transition. *Proceedings of the National Academy of Sciences*, 55(4):676, 1966.
- [180] B. Gaveau and L. S. Schulman. Theory of nonequilibrium first-order phase transitions for stochastic dynamics. *Journal of Mathematical Physics*, 39(3):1517, 1998.