Sorption in disordered porous media

Numerical study of the phenomena within a lattice gas model

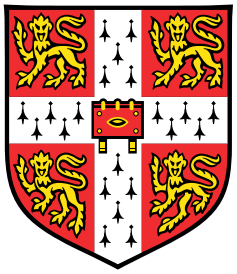

Zilvinas Rimas

Department of Chemistry
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Sidney Sussex College July 2017

Declaration

This thesis is the result of my own work carried out between May 2013 and April 2017 at the Department of Chemistry, Cambridge, under the supervision of Dr. S. N. Taraskin. It includes nothing which is the outcome of work done in collaboration except as declared in the Declaration and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Declaration and specified in the text.

This thesis does not exceed the specified limit of 60,000 words, including summary/abstract, tables, and footnotes, as defined by the Physical Sciences Degree Committee. It has been typeset in 12pt font using LaTeX2ε according to the specifications defined by the Board of Graduate Studies and the Physical Sciences Degree Committee.

Zilvinas Rimas
July 2017

Acknowledgements

First and foremost, I would like to thank my supervisor Dr. Sergei Taraskin for guidance and motivation throughout the years. In addition to his constant support and dedication, I have always felt trusted and free to explore the research paths of my own interest and choosing. I have greatly appreciated such privileges over the course of my work, but more importantly, I believe the resulted experience and the lessons I took from it will be of immense value to me in the long term.

I am also grateful to Dr. Tom Handford for his help during the starting months of my PhD as well as Sam Niblett, Alex Dear and Luke Stevens for their assistance and fruitful discussions. This thesis would also not have been possible without my friends and family. I thank Issabella, Liana and Petras for sharing with me the good times and sticking by my side through the difficult ones. I am especially grateful to Heidi for her support during the most challenging final months of this work.

Lastly, I thank the EPSRC for the financial support, Sidney Sussex College and the Department of Chemistry for the provided accommodation and office space.

Abstract

The lattice-gas model of sorption in disordered porous media is studied for a variety of settings, using existing, updated and newly developed numerical techniques. Firstly, we construct an efficient algorithm to calculate the exact partition function for small lattice-gas systems. The exact partition function is used for detailed analysis of the core features exhibited by such systems. We proceed to develop an interactive Monte Carlo (MC) simulation engine, that simulates sorption in a porous media sample and provides real-time visual data of the state space projection and the 3d view of the sample among other parameters of interest, as the external fields are manipulated. The use of such tool provides a more intuitive understanding of the system behaviour. The MC simulations are employed to study sorption in several porous solids: silica aerogel, Vycor glass and soil. We investigate how the phenomena depend on the microstructure of the original samples, how the behaviour varies with the external conditions, and how it is reflected in the paths that the system takes across its state space. Secondly, we develop two methods for estimation of the relative degeneracy (the number of microstates that have the same value of some macroscopic variables) in the systems that are too large to be handled exactly. The methods, based on a restricted infinite temperature sampling, obtain equidegenerate surfaces and the degeneracy gradient across the state space. Combined with the knowledge of an internal energy of a microstate, it enables us to construct the free energy map and thus the equilibrium probability distribution for the studied projection of the state space. Thirdly, the jump-walking Monte-Carlo algorithm is revisited and updated to study the equilibrium properties of systems exhibiting quasi-ergodicity. It is designed for a single processing thread as opposed to currently predominant algorithms for large parallel processing systems. The updated algorithm is tested on the Ising model and applied to the lattice-gas model for sorption in aerogel and Vycor glass at low temperatures, when dynamics of the system is significantly slowed down. It is demonstrated that the updated jump-walking simulations are able to produce equilibrium isotherms which are typically hidden by the hysteresis effect characteristic of the standard single-flip simulations. As a result, we answer the long standing question about the existence of the first-order phase transitions in Vycor. Finally, we investigate sorption in several distinct topology network representations of soil and aerogel samples and demonstrate that the recently developed analytical techniques for random networks can be used to achieve a qualitative understanding of the phenomena in real materials.

List of Publications

Chapter 3

(1) “A single-walker approach for studying quasi-ergodicsystems”
Z. Rimas and S. N. Taraskin. *arxiv:*1609.05683 (2016) (accepted for publication in *Sci. Rep.* (2017))

Chapter 6

(2) “Condensation of liquid in porous media”
Z. Rimas, T. P. Handford, and S. N. Taraskin. *III GEFENOL Summer School on Statistical Physics of Complex Systems*, IFISC, (2013)

Table of contents

List of figures

Chapter 1

Introduction

The free-energy landscapes of complex interacting systems [1] such as spin glasses [2–4], Lennard-Jones clusters [5], supercooled liquids [6, 7], random-field Ising model [8–10], lattice gas in porous media [11, 12], biological evolutionary dynamics of proteins [13–17], etc. are often characterised by many local minima separated by free-energy barriers. Understanding system's behaviour in such landscapes is a great challenge for experimental, analytical and numerical approaches [18].

The existence of several local minima in a free-energy landscape of a system gives rise to the entire class of phenomena associated with the fact that multiple phases, such system can be in, can coexist under some set of external conditions. An extremely rich field of research of phase transitions and critical phenomena [19] has been fascinating scientists since the early works of van der Walls and Maxwell [20, 21]. However, a set of more peculiar behaviours arise from the existence of not just multiple, but in fact a very large number of local free-energy minima as in the mentioned systems. The diffusion across the state space in this case becomes increasingly slow as the temperature is lowered. Such regime in the past has often been confused as corresponding to the thermodynamic metastable phases associated with the conventional first-order phase transition, which is expected in pure systems with only two free-energy minima. Instead, in this case such slowing down is known to be due to system inability to cross local potential barriers within the experimentally (or computationally) realizable time scale [11, 12, 22].

As far as computer simulations are concerned, a further challenge which arises from the extremely slow equilibration times across large regions of the state space is so called quasi-nonergodicity [18]. Although the systems in consideration are fundamentally ergodic, in the discussed regime of slow diffusion, the explored volume of the state space, within the simulation time scale, is often insufficient to reliably estimate statistical averages of properties of interest. And even if the local minima in the free-energy landscape are linked

by low-energy passages, as is likely in the multidimensional state space of such systems, the passages, so called 'entropic bottlenecks', are comparatively narrow and hence rarely sampled, thus not alleviating the problem of quasi-nonergodicity.

Within a real-world experimental environment, our ability to speed up the equilibration process while maintaining the desired external conditions (e.g. temperature, chemical potential, strength of magnetic field, etc.) is limited. However, in the world of computer simulations, capability to tweak the most fundamental properties of 'nature', such as the visit probabilities of the states or even the connectivity of the state space itself, provides us with the additional set of tools to investigate the systems, and to tackle problems such as quasi-nonergodicity.

Due to its discrete nature, the lattice-gas model is well suited for computer simulations. The model is widely used in physics, biophysics and physical chemistry and known by different names in various fields, for example, as the order–disorder model for alloys [23], as the regular solution model for liquid solutions [24], or as the Ising model for magnetic problems [25]. For pure liquids, the model was introduced by Frenkel [26] in 1932 and, in recent decades, it has been widely employed in analytical and computational approaches to study sorption in porous media [27–31]. It is one of the simplest, if not the simplest model, exhibiting a phase transition. However, although the formulation of the model is deceptively straightforward, its solutions are complex. The model has been exactly analytically solved in only a few 1d [32] and 2d cases [33, 34], hence further progress demands the development of numerical methods or some approximations in analytical approaches.

The original and most intuitive way of performing a computer simulation of some system is to exactly follow the kinetics that occur in the real-world experiments. Molecular Dynamics (MD) simulations [35] which numerically integrate the equations of motion do precisely that, aiming to authentically reproduce the behaviour of the studied system. There is a long and rich history of such studies, providing insights and answers to many challenges associated with phase behaviour. However the 'authenticity' comes at a cost: like the real world experiments in this field, MD simulations reproduce the phenomena associated with the long time-scale equilibration of the system and slow diffusion across its state space. Instead, the natural choice of computational tool, that provides us the promised access to tweak the underlying mechanics of the phenomena in question, is the Monte Carlo (MC) method [4]. The advantages of MC techniques are in the freedom with which the algorithms can be engineered to move around the system configuration space. Some of the options are, but not limited to: (i) altering the dynamics of the system by e.g. cluster-flip methods [36–38], (ii) deforming the free-energy landscape itself by e.g. multicanonical algorithms [39–44], (iii) extending the dimensionality of the state space by e.g. expanded ensemble, simulated

tempering algorithms [45, 46, 3], (iv) systematically introducing new connections within the state space by e.g. replica exchange methods [36, 47–50].

Despite the decades of steady improvement in computing power, and hundreds of algorithms developed to study disordered systems with complex free-energy landscapes, the challenges seem to be at least as tough as the tools built to tackle them are powerful. Such ubiquitous and fundamental problems as finding the lowest energy state, in a general case, cannot be solved without a brute-force method of enumerating all possible states. However within a lattice-gas model, the size of the state space grows exponentially with the system size, thus the improvement in computational power alone will not enable us to handle large systems. The development of new, case-specific algorithms is going to be required, and they can only come from physical insights into the configurational space, its shape, its local connectivity and global topology, of the system at hand.

In this work, we set out to numerically study sorption (adsorption and desorption) phenomena in disordered porous media, one of the systems exhibiting complex free-energy landscapes with numerous metastable states. The analysis is undertaken within the lattice-gas model. We aim to develop a deeper understanding of the system configurational space. In particular, how the sorption processes are reflected in the paths that the system takes across its state space, how the form, shape and orientation of the state space itself depends on the microscopic parameters of the original sample and the physical fields in play.

This Chapter briefly introduces the models of porous media of interest in Sec. 1.1. In particular, the characteristics and generation procedures of silica aerogel and Vycor glass model samples are described in Sec. 1.1.2 and Sec. 1.1.3, respectively. The current state of the research field of sorption phenomena in porous media is discussed in Sec. 1.2 and the lattice-gas model presented in Sec. 1.3. Finally, the relevant numerical methods are considered in Sec. 1.4 including the overview of the algorithms developed to study systems with complex free-energy landscapes in Sec. 1.4.3. The structure of the thesis is summarized in Sec. 1.5.

1.1 Porous media. Generation of model samples

Porous materials are solids with empty (pore) space inside the material and significant internal surface area (e.g. $10^2 - 10^3$ m^2/g in mesoporous silica [51, 52]). This property results in their wide applications as adsorbents, drug delivery systems, molecular sieves or nanoreactors in catalysis [53, 30]. Porous materials can be classified according to typical pore size as (i) microporous (pore size $\lesssim 2$ nm), (ii) mesoporous (pore size in the range from 2 to 50 nm), and (iii) macroporous (with pore size $\gtrsim 50$ nm) [53].

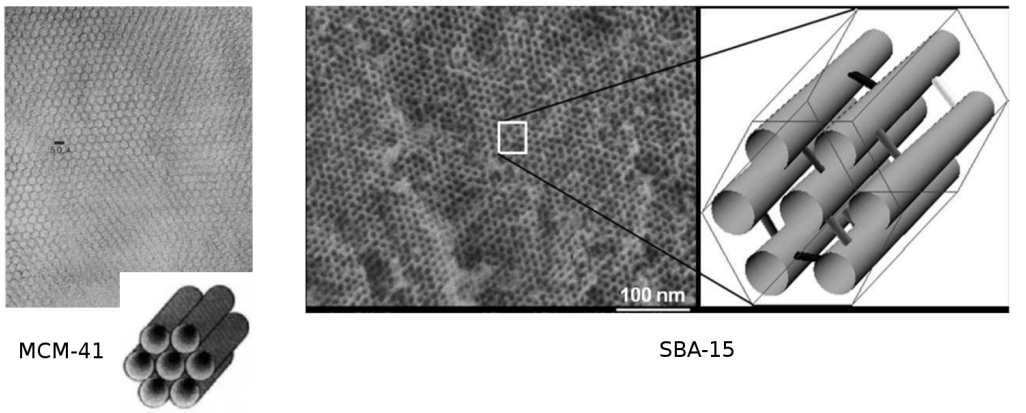

Fig. 1.1 Representative transmission electron micrograph and schematic drawings of MCM-41 [51] (left panel) and SBA-15 [52] (right panel).

In the last two decades, ordered porous silicas have been instrumental for the development of numerous theoretical and experimental techniques designed to study adsorption and desorption in porous media, hence we briefly describe them in the following Section. However, one of the main aims of this work is to study sorption phenomena in the disordered porous media, in particular silica aerogel and Vycor glass. In Sec. 1.1.2 and Sec. 1.1.3, we discuss those materials, their properties and construction of their structural model samples, which we use throughout the computational study of sorption phenomena.

1.1.1 Ordered porous media

The study of adsorption, also known as capillary condensation, in mesopores was drastically changed by the recent discovery of a class of highly ordered mesoporous materials known as the M41S family [51, 54], such as MCM-41 [51] and SBA-15 [52]. Since these discoveries, many similar materials with the highly ordered pore structures, e.g. cylindrical, interconnected spheres etc., have been developed. One of the main ways to understand sorption behaviour is to relate the properties of macroscopic phenomena, such as hysteresis between adsorption and desorption branches of isotherms, with microscopic properties of the material, i.e. distributions of pore sizes, shapes and the long-range structure of the media. Therefore, the development of materials with well defined structure and homogeneous pore properties is crucial in testing any theoretical model in this area.

The MCM-41 (Mobil Composition of Matter No. 41) [51] was synthesised from the calcination of aluminosilicate gels in the presence of surfactants. This mesoporous material consists of a regular hexagonal array of one-dimensional channels with typical channel diameter in the range 16 – 100 Å (see the left panel in Fig. 1.1). The distribution of the

channel diameters is quite narrow and the mean value of the diameter can be tailored through the choice of surfactants, auxiliary chemicals and reaction conditions.

Another example of synthetic porous material is SBA-15 (Santa Barbara Amorphous No.15) [52]. The structure of SBA-15 is similar to that of MCM-41, but the walls between pores are thicker (3–7 nm) and the pore sizes are slightly larger (adjustable between 6 and about 15 nm). One of the interesting features of SBA-15 is the presence of micropores in the pore walls. These micropores of $\simeq 2$ nm in diameter connect neighbouring cylindrical mesopores (see the right panel in Fig. 1.1).

We have constructed several basic structural models for SBA-15, with varying concentration of micropores connecting the mesopores. The structural samples were built on a custom lattice, specifically designed to reflect the topological properties of the material. A brief investigation of sorption dependence on the topology of the material, controlled by varying the concentration of the micropores, raised multiple interesting questions, and we hope to continue this direction of research in future work. However, at the present state, the investigation of sorption in ordered porous media is outside of the scope of this thesis.

Besides the examples mentioned above, there are many other synthetic porous materials, which could be relatively straightforwardly represented by lattice based structural models. Of particular interest would be KLE silica [55], which contains mesopores of almost ideal spherical shape (diameter $\simeq 140$ Å), that are arranged on a distorted FCC lattice. KLE/IL silica [56] is similar to KLE silica, but in addition, the spheres are connected by cylindrical mesopores (diameter $\simeq 26$ Å). Finally, SE3030 silica [57] with worm-like cylindrical channels (mode diameter $\simeq 95$ Å) is well suited to test how the models developed for the highly structured materials behave as a limited amount of disorder is introduced. Such studies could help to bridge the gap between our understanding of the sorption phenomena in ordered and disordered porous media, e.g. silica aerogel or Vycor glass.

1.1.2 Silica aerogel

Silica (silicon dioxide) aerogels consist of bonded silicon and oxygen atoms joined into long strands and then into beads randomly linked together with pockets of air between them. Silica aerogel is an amorphous form of "common sand" - nonflammable, nontoxic, and environmentally safe. It has remarkable physical properties such as extremely low thermal conductivity (down to 0.004 W/m/K) [60], high stress resistance, and very low density ($\rho < 2\text{g/m}^3$ [61]). The physical structure of silica aerogel has fractal properties across a wide range of scales [62–64], which is of particular importance when studying sorption phenomena in such materials [65, 66].

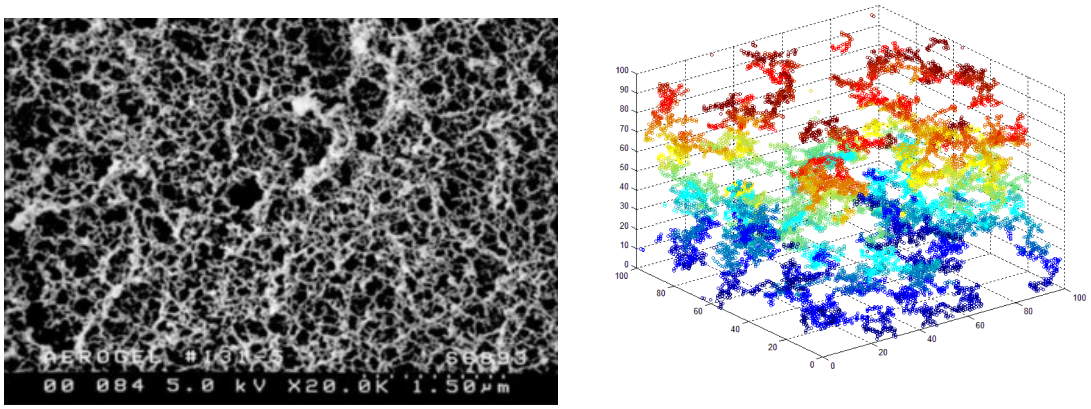

(a) (b)

Fig. 1.2 (a) Scanning electron micrograph of silica aerogel network. Note 1.5 micron scale shown at the bottom of the right panel [58]. (b) DLCA -generated model of silica aerogel [59]

To produce model samples of silica aerogel, we have used Diffusion-Limited Cluster Aggregation (DLCA) algorithm [67]. The algorithm is based on the process resembling the physical synthesis of the base-catalysed aerogels [68]. At first, a lattice with desired concentration of randomly distributed particles is set up. The particles then perform a random walk across the lattice sites, i.e. a randomly selected particle or cluster of particles is moved by one lattice spacing in a random direction. A cluster i is selected to move according to the probabilities,

$$p_i = \frac{n_i^\alpha}{\sum_i n_i^\alpha} \,, \tag{1.1}$$

where n_i is the number of particles(called further as matrix cells) belonging to that cluster and α is the aggregation parameter. Unless specified otherwise, we use $\alpha = -0.55 \simeq -1/D_f$ where $D_f = 1.78$ is the expected fractal dimension of the DLCA structures [69]. This choice ensures that the aggregates have an effective diffusion coefficient inversely proportional to their radius. When two particles or clusters of particles encounter one another, they irreversibly aggregate and proceed further as a single object. Such process continues until only a single cluster remains. The resulting structures well reproduce experimental neutron scattering data of the real aerogels [69].

It turns out that there are two distinct rules in use for the two clusters to be considered as colliding. The clusters may aggregate once some member particles of each cluster occupy adjacent lattice sites [67, 70]. Alternatively, the clusters can aggregate only when a proposed move would make some member particles from each cluster to occupy the same lattice site [71–73]. Despite the detailed analysis [59] of the two aggregation rules, a definitive answer could not be found as to which of the rules results in the model samples better

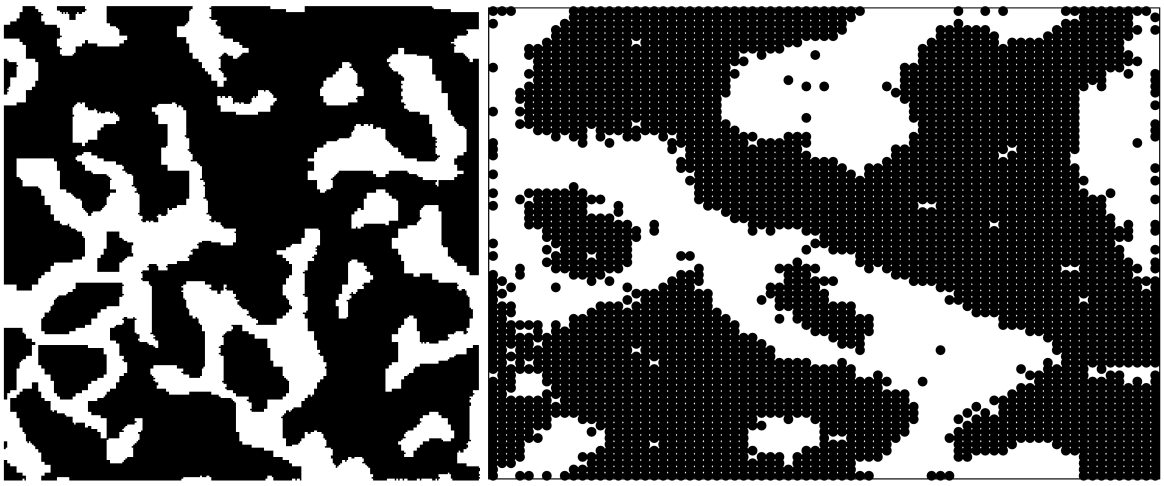

(a) (b)

Fig. 1.3 (a) Digitised transmission electron microscope image of a 2d-section through Vycor glass. [74]. (b) 2d-section through a reconstructed Vycor sample on a simple cubic lattice with lattice size $L = 75$ and lattice constant 8 Å [59]. Dark circles represent matrix cells in both panels.

reflecting the experimental data. Due to the more common use in literature, we have chosen to follow the latter convention in this work. An image of real silica aerogel and a structural model of this material obtained using the DLCA algorithm are shown in Fig. 1.2.

All our model samples of aerogel and other materials are constructed on a Body Centred Cubic (bcc) lattice, since it has been shown that phase transitions simulated on a simple cubic lattice under certain conditions exhibit non-physical artefacts such as faceted growth [71].

1.1.3 Vycor glass

Vycor is a controlled pore glass produced by Corning Inc. It is approximately 96% silica and 4% boron trioxide, with $25-30\%$ porosity, i.e. $25-30\%$ of its volume consists of empty pore space. Vycor glass is made by a multi-step process. First, a relatively soft alkali-borosilicate glass is melted and formed into a desired shape. The following heat treatment causes the material to separate into two intermingled chemically distinct phases. One of the phases is made up mostly of alkali and boric oxide and can be easily dissolved in acid. The other phase is mostly silica, which is insoluble. The sample is then submerged in a hot acid solution, which leaches away the first phase producing a porous glass material. The resulting solid (matrix) structure as well as the pore space percolate through the system [12]. Unlike in the case of silica aerogel, Vycor structure does not exhibit fractal characteristics,

and instead has a well defined scale of the pore radius $\simeq 40-60$ Å, with 96% of the pores in the glass being ± 3 Å away from the mean value [75].

There are two main approaches to produce structural models of Vycor glass for numerical studies. The process-based methods, e.g. MD simulation of Vycor synthesis [76], and reconstruction-based methods [74, 77], making use of experimentally obtained two-point autocorrelation functions of the real Vycor samples. We have decided to use one of the latter approaches, in particular, simulated annealing MC method [78–80]. The simulation box is coarse grained into bcc lattice consisting of matrix and pore cells. Starting with a random distribution of matrix cells with the desired concentration, the algorithm attempts trial moves by swapping a randomly chosen matrix cell with a randomly chosen void cell. The trial structure is then compared with the target structure using some spatial correlation function, available from the experimental scattering data. The trial move is accepted with the standard Metropolis probability [78]. A virtual temperature is used to determine the acceptance probability and thus to control the annealing process. Throughout the simulation, the temperature is gradually decreased, preventing the unfavourable configurations to be chosen, and eventually forcing the system to settle down in a local minimum corresponding to a model sample sufficiently similar to the target structure.

In general, it is known [80] that the two-point autocorrelation function does not uniquely describe the porous material and that unrelated structures may have very similar autocorrelation functions [81, 82]. Thus ideally, additional experimentally obtained correlation functions should be used when comparing the current and target structure, in order to capture complete set of structural characteristics. However, it has been shown [59] that the Vycor glass reconstructed using only the two-point autocorrelation function reproduces higher order properties as well. Therefore, in this work we have relied on the algorithm developed by S. Niblett [59] based solely on the two-point autocorrelation function. An example of the experimentally obtained and reconstructed structures are shown in Fig. 1.3. The Vycor models produced and used in this work have the lattice spacing equal to 10 Å.

1.2 Sorption in porous media

The quantification of vapour condensation and evaporation in porous media is of interest to many industrial and natural processes. Various forms of capillary-assisted evaporation and condensation processes occur across a wide range of scales varying from the wetting of nanometric textured surfaces [83] to the vapour transport enhancement in soil, rock pores, aerosol and dust [84]. Synthetic applications such as sintering [85] of materials are highly dependent on bridging effects resulting from capillary condensation. In addition to the

advantages of capillary condensation, it can also cause many problems in materials science applications such as atomic force microscopy [86] and microelectromechanical systems [87]. Rapid condensation processes play a crucial role in the onset of capillary-induced friction affecting mechanical behaviour of physical systems and industrial applications [88]. Therefore, qualitative and quantitative understanding of these phenomena is highly desirable.

The confined geometry of porous media significantly alters the condensation and evaporation behaviour when compared with what is expected in free space, e.g. it takes place at a lower pressure (or chemical potential) compared to the bulk saturation value [28, 29]. A common feature of sorption in mesoporous materials is the occurrence of a hysteresis loop between the adsorption and desorption isotherms below some critical hysteresis temperature [89, 90]. The characteristics of the hysteresis loop depend on the material examined, the adsorbate and the temperature [30, 29]. As far as the material itself is concerned, the shape and location of the isotherms are associated with (i) properties of individual pores (the distributions of pore sizes and shapes), (ii) connectivity between the pores, i.e. various topological properties of the pore network, and (iii) interactions between the surface of the material and the adsorbate. While the studies of fluids confined in single pores of simple geometry have clarified the mechanism for such transitions [27, 29, 91–93], the situation in real materials, such as mesoporous glasses and silica gels which consist of an interconnected network of pores of various shapes and sizes, is still under active investigation [94, 71, 95, 96, 30, 31, 97, 98].

Please note that some of the content and the structure of the following Sec. 1.2.1 and Sec. 1.2.2 are partially based on an excellent review by Coasne et. al. [31] on the current state of the field.

1.2.1 Ordered porous media

In this work, we mainly concentrate on disordered porous silicas, however, it is valuable to briefly summarize the progress achieved in simpler systems, because it provides great general insights about the sorption phenomena in any porous media.

Simple geometry systems, such as cylindrical or 2d-pores can often be represented by a small enough structural model, for the molecular simulations to handle. It has been shown that the hysteresis loop appears for the pores above certain diameter ($D > 30$ Å in simulations [99] and $D > 40$ Å in experimental data [100]). Below such diameter, the transition seems to be reversible and no hysteresis is observed. This is explained by the fact that for small pore diameters the liquid layer adsorbed near the pore surface occupies almost the entire width of the pore, thus eliminating clear interface between the adsorbed film and the gas phase inside the pore (see the left panel in Fig. 1.4). Moreover, for a given diameter

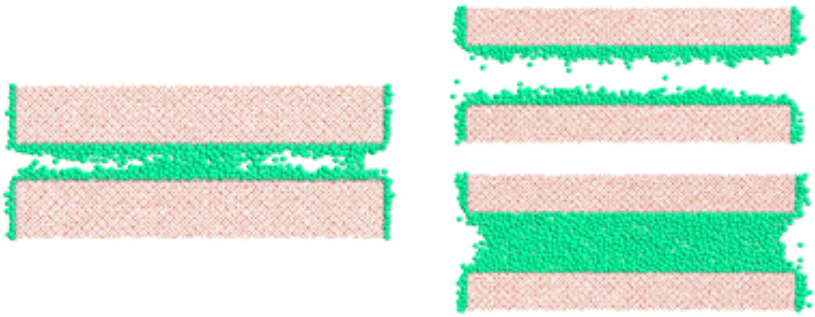

Fig. 1.4 Characteristic molecular configurations obtained using grand-canonical MC simulations for nitrogen adsorption in a MCM-41 pore with $D = 24$ Å (left panel) and $D = 48$ Å (right panel) [31]. The right panel shows two distinct fluid configurations at different pressure, before (upper panel) and after the adsorption (lower panel).

D of the pores, there exist the so-called capillary condensation temperature $T_{cc}(D)$, above which the condensation becomes reversible due to the high disorder induced by the finite temperature [101–103]. Below $T_{cc}(D)$, the irreversible transition is observed between the partially filled (liquid film on the surface of the pores) and completely filled configurations. Since the adsorption transition happens at a higher pressure than the desorption one, the hysteresis phenomena is observed. Such difference is explained [104] by the different shape of the gas–liquid interfaces encountered upon filling (cylindrical meniscus) and emptying (hemispherical meniscus) of the pores.

Classical models explaining capillary condensation in pores are mainly based on Kelvin equation [105, 104, 106]. However, such approach fails to account for the formation of a film adsorbed at the pore surface prior to capillary condensation. While it is reasonable in the case of large pores, the approximation becomes inaccurate for the pore diameters of the same order of magnitude as the thickness of the adsorbed film. Multiple corrections for the Kelvin equation based approaches have been proposed [107–111]. The resulting models significantly increase the prediction accuracy for certain systems, and even provide accurate quantitative description of the desorption process and the height of the hysteresis loop for argon sorption in MCM-41 pores. Nevertheless, such models fail to describe pore filling in small nanopores, since they ultimately rely on the macroscopic concept of surface tension, which breaks down in the absence of well-defined interface separating the two phases.

In addition to thermodynamic models, density functional theory (DFT) has been used to study sorption phenomena in pores [112]. In contrast to the models described above, DFT approach was able to predict the capillary critical point above which capillary condensation is replaced by continuous and reversible pore filling. The theory was developed to predict

adsorption isotherms and estimate pore size distributions for pores of simple geometries [113–115].

1.2.2 Disordered porous media

It is convenient to start the study of sorption in complex disordered porous media, such as silica aerogel or Vycor glass, from the analysis of simpler case such as defective pores [116–118], e.g. spherical cavities connected by cylindrical pores of smaller diameter (i.e. by constrictions). The filling of such systems starts by the usual liquid film formation on the surface of the pores, and continues by filling the narrower cylindrical regions at the pressure close to that obtained for the regular cylindrical pores having the same diameter. The remaining spherical nano-bubbles are filled by a metastable transition at a pressure higher than the equilibrium transition pressure. The desorption mechanism in such systems depends on the diameter of the cylindrical pores. If the spherical cavities empty together with the cylindrical constrictions, the process is called "pore blocking" [119, 120]. Otherwise, if the cavities empty by nucleation, before the cylindrical parts of the pore evaporate and without having a direct connection to the external reservoir, such process is called "cavitation" [118, 121, 117, 122, 123]. By which of the two ways a particular system undergoes the desorption depends on the relative stability of the liquid confined in the constriction and the liquid confined in the cavity.

Sorption phenomena in highly disordered porous materials significantly differ from what is observed for regular pores. The asymmetrical hysteresis loop, with the desorption branch much steeper than that of adsorption, has been attributed to the pore network effects [119, 120]. A more precise picture was obtained using DFT calculations [124], showing that the asymmetric hysteresis loops arise due to pore blocking. It was also demonstrated [125–127, 76] that interconnected porous materials such as Vycor and Controlled Pore Glasses cannot be described as an assembly of ideal cylindrical pores. Investigations by the mean-field DFT within the lattice-gas model [11, 128, 12] showed that most of the features of sorption in disordered porous materials can be explained by appearance of a free-energy landscape with a large number of metastable states. Such investigation revealed that classical van der Waals picture of condensation in pores breaks down for disordered porous materials. In particular, the hysteresis loops observed for disordered porous solids are not necessarily related to a true equilibrium phase transition, but can be a signature of out-of-equilibrium phenomena. By using the mean-field approach for coarse-grained models of aerogel [128], it was shown that the existence of the first-order phase transition depends on the porosity of the material (the transition is observed for $\phi = 0.98$ but not for $\phi = 0.87$). Since the collective phenomena are suppressed by the spatial disorder in porous media, the pore critical temperature T_{cc} is

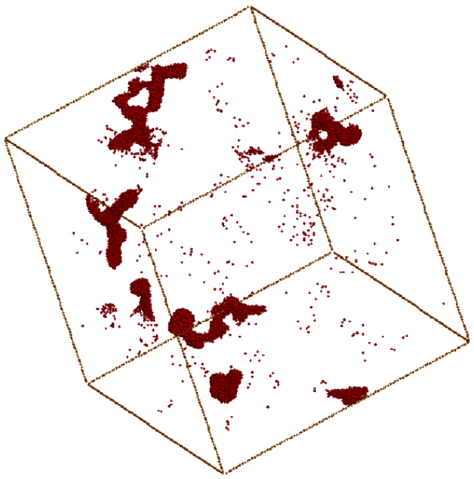

Fig. 1.5 Collective condensation events (avalanches) in a 87%-aerogel sample [128].

therefore lower than that for the bulk. It has also been shown that the hysteresis phenomenon is closely related to the irreversible chain reaction-like events, called avalanches [128] (see Fig 1.5). Their statistical properties as a function of the sample porosity, temperature and chemical potential have been investigated [96]. Despite the enormous progress that has been achieved in the area during the recent years, quantitative description of the phase transitions in disordered porous media is not yet complete. Neither the necessary conditions for the existence of the first-order phase transition, nor the dependence of T_{cc} on the morphological and topological pore properties have been fully explained.

In this work, we aim to continue developing the understanding of sorption phenomena in highly disordered porous media, such as silica aerogel and Vycor glass. We chose to proceed by developing and utilizing numerical sampling techniques designed to explore the complex free-energy landscapes of such systems within the lattice-gas model.

1.3 Lattice-gas model

Sorption phenomena in porous media can be studied using a lattice-gas model [129, 29, 27], which splits the sample space into cells typically arranged in a regular lattice. Each cell can be occupied either by matrix, liquid or vapour. It is assumed that vapour does not interact with liquid or matrix and thus a cell occupied by vapour is equivalent to an empty cell. If a cell is occupied by liquid it is called a liquid or, equivalently, a fluid cell. The state of the i-th cell is characterized by two indicator state variables η_i and τ_i. If the cell is occupied by a matrix ($\eta_i = 0$), it cannot become fluid and its state does not change during the simulation. Otherwise, if the cell is not occupied by a matrix ($\eta_i = 1$), it can be filled either by fluid

($\tau_i = 1$) or vapour ($\tau_i = 0$), and variables τ_i can change during the simulation. Such system has nearest neighbour interactions between two fluid cells and between a fluid and a matrix cell. The total number of fluid cells is $N_\mathrm{f} = \sum_{i=0}^{N} \tau_i$, where N represents the volume of the system, i.e. the total number of pore cells occupied either by fluid or vapour. The Hamiltonian of such system can be expressed as follows [129, 11],

$$\mathscr{H} = -w_\mathrm{ff} \sum_{\langle ij \rangle} \tau_i \tau_j \eta_i \eta_j - w_\mathrm{mf} \sum_{\langle ij \rangle} [\tau_i \eta_i (1 - \eta_j) + \tau_j \eta_j (1 - \eta_i)] \,, \tag{1.2}$$

where w_ff and w_mf are the parameters representing fluid-fluid and matrix-fluid interaction strength, respectively, and the sums are taken over all distinct pairs of the nearest-neighbour cells. The energy scale is defined by setting up $w_\mathrm{ff} = 1$. The only parameter of the model, wettability $y = w_\mathrm{mf}/w_\mathrm{ff}$, depends on the material of the matrix, the type of the fluid considered and the lattice-spacing of the structural model.

For given chemical potential μ and inverse temperature $\beta = \frac{1}{k_B T}$ (where k_B is the Boltzmann constant), the system can be in a set of states which form a grand canonical ensemble described by the following grand partition function,

$$Z_\mu = \sum_{\{\tau\}} e^{-\beta \tilde{\mathscr{H}}} \,, \tag{1.3}$$

where the reduced Hamiltonian $\tilde{\mathscr{H}} = \mathscr{H} - \mu N_\mathrm{f}$, and the sum runs over all possible fluid-occupation configurations $\{\tau\}$ of non-matrix cells, with the matrix variables η_i being quenched.

Following Ref. [71], throughout this work, unless specified otherwise, the value of the wettability is set to $y = 2$, which results in a relative agreement between experimental and simulation results for aerogel. In the case of $y = 1/2$, the lattice-gas model described by Hamiltonian 1.2 can be mapped onto spatially correlated site-diluted Ising model [129, 12], and in a general case, onto random-field Ising model [130] with spatial correlations in random fields imposed by matrix sites. If there are no matrix cells present, i.e. porosity $\phi = 1$, the model describes condensation in bulk, which can be mapped onto the classical Ising model.

1.4 Numerical methods

The theoretical approaches discussed above have proven to be highly valuable in the study of sorption phenomena in porous media. However, all of them rely on a variety of approximations, such as mean-field, absence of thermal noise effects in the case of DFT, or the existence of well defined interface between the phases for Kelvin-equation-based approaches.

Thus the numerical methods are used not only for independent studies of the phenomena but often also serve as a validation tool for some newly developed theoretical approaches.

Within the lattice-gas model of a system with N cells that can be either fluid or vapour, the state space is a N-dimensional cube, with 2^N edges marking the possible states. At any one point in time the system is in one of those states, and can jump across some side of such cube, as one of the cells switches from fluid to empty or *vice versa*. Beyond $N = 10^2$-size systems, enumerating all possible states and thus obtaining the exact partition function is computationally infeasible. The growing interest in understanding sorption inside the materials with a complex interconnected pore structure, and thus the need to study large systems (as opposed to small, fixed geometry, single-pore ones), points to the sampling-based numerical approaches, such as MC method.

1.4.1 Monte Carlo method

In order to model the trajectory of the system in the grand canonical potential landscape given by the lattice-gas model, the MC algorithm defines a set of rules by which the system evolves. In thermal equilibrium, the transition rates $R(\{\tau_i\} \to \{\tau_j\})$ from the state $\{\tau_i\}$ to the state $\{\tau_j\}$ should obey the detailed balance [131, 132],

$$R\left(\{\tau_i\} \to \{\tau_j\}\right) P\left(\{\tau_i\}\right) = R\left(\{\tau_j\} \to \{\tau_i\}\right) P\left(\{\tau_j\}\right) \ , \tag{1.4}$$

which follows from the master equation describing the time evolution of the probability $P(\{\tau_i\},t)$ for the system to be in a state $\{\tau_i\}$ at time t,

$$\frac{\mathrm{d}P(\{\tau_i\},t)}{\mathrm{d}t} = \sum_{\{\tau_j\}} \left[-P(\{\tau_i\},t)\,R\left(\{\tau_i\} \to \{\tau_j\}\right) + P\left(\{\tau_j\},t\right) R\left(\{\tau_j\} \to \{\tau_i\}\right)\right] \ , \tag{1.5}$$

because, in thermal equilibrium, the probability $P(\{\tau_i\},t)$ does not depend on time and thus the left-hand side of the Eq. (1.5) equals zero.

The detailed balance condition does not uniquely define the possible transitions and their rates. One of the simplest, most popular and most convenient to implement choices for systems with binary-state dynamics [133], such as spin and lattice-gas models, is the so-called Metropolis [134] single spin-flip local dynamics. It was originally developed for the Ising model [131, 132]. The locality here refers to the fact that only the nearest-neighbour interactions are considered in the models. The single spin-flip dynamics means that each consecutive configuration $\{\tau_j\}$ differs from the previous one $\{\tau_i\}$ only at a single lattice cell.

The transition rates are thus defined as follows,

$$R\left(\{\tau_i\} \to \{\tau_j\}\right) = \min\left[\exp\left(-\beta\left(\mathscr{H}(\{\tau_j\}) - \mathscr{H}(\{\tau_i\})\right)\right), 1\right] \ , \tag{1.6}$$

and if $\{\tau_i\}$ and $\{\tau_j\}$ differ at multiple lattice cells $R\left(\{\tau_i\} \to \{\tau_j\}\right) = 0$. In Eq. (1.6), the expression for $\mathscr{H}(\{\tau_i\})$ is given by Eq. (1.2). The form of the rate equation is derived from the detailed-balance condition (see Eq. (1.4)), where the equilibrium probabilities, $P(\{\tau_i\})$, are given by Boltzmann distribution,

$$P(\{\tau_i\}) \propto \exp\left(-\beta \mathscr{H}(\{\tau_i\})\right) \ . \tag{1.7}$$

For the system of size N, there are N possible transitions to occur, one for each site that can change its state. From such set of possible events, one is picked at random, and is accepted with a probability equal to the corresponding transition rate R. According to this protocol, the system traverses the state space and eventually reaches the equilibrium defined by the Boltzmann probability distribution (see Eq. (1.7)).

The classical MC method described here is most often used to obtain the equilibrium state of the system at some set of external parameter values. However, the method is not well suited for studying the temporal dynamics of the system, and thus out-of-equilibrium phenomena, since it is difficult to relate the real-world time scale with the MC time steps. In the next subsection, we present a different version of the MC method, the so-called kinetic Monte Carlo algorithm which resolves this problem.

1.4.2 Kinetic Monte Carlo method

The kinetic Monte Carlo (kMC) method, also known as the N-fold Monte Carlo [132, 135], at its core, is similar to the classical MC described in the previous Section, however it provides a physical time scale of the events in the simulation [136]. Moreover, under certain conditions, e.g. low temperatures, when the system does not change its state very often, kMC is vastly more efficient than the classical MC, and thus is well suited for the study of low-temperature behaviour that we are interested in throughout this thesis.

Within kMC, each possible transition is modelled as a homogeneous Poisson process with the rates defined as before (see Eq. (1.6)). Thus the next event after time t occurs stochastically at time $t + \Delta t$, where Δt has the cumulative distribution,

$$\text{Prob}\left(\Delta t < \Delta t'\right) = 1 - \exp\left(-\sum_{\{\tau_j\}} R\left(\{\tau_i\} \to \{\tau_j\}\right) \Delta t'\right) \ . \tag{1.8}$$

Here, Prob $(\Delta t < \Delta t')$ is the probability that the random variable Δt is less than some value $\Delta t'$. At time $t + \Delta t$, the occupancy of a single cell is changed, and the particular cell is chosen with probability proportional to the respective event transition rate and normalized by the sum of all N transition rates, in order to obtain a valid probability distribution [4].

Constructing, maintaining and choosing an event from a set of possibilities with an arbitrary probability distribution as described above, is a relatively computationally expensive task. In a general case, such procedure (a single evaporation or condensation event) would be an $\mathrm{O}(\log(N))$ complexity task, however throughout this project we have developed algorithms that, under certain conditions, decrease it to $\mathrm{O}(1)$, which allow us to study large systems relatively efficiently. The detailed description of the algorithm and other numerical developments are presented in Ch. 5.

1.4.3 Exploring complex free-energy landscapes

As discussed earlier, systems of disordered porous media such as Vycor glass and silica aerogel are characterised by a complex free-energy landscape with multiple metastable states [11, 128, 12]. At sufficiently low temperatures, the use of conventional single-flip MC methods to study such systems is prohibited by exponentially long transition times over the free-energy barriers. Under such conditions, the system can be trapped in one of the local free-energy minima, thus failing to ergodically explore the entire state space and achieve equilibrium. This behaviour is know as quasi-nonergodic and it has been, as well as a related broken-ergodicity problem, widely addressed in the literature (see e.g. Refs. [137, 2, 138, 18, 22, 139, 140]).

Since the dawn of ever cheaper and more accessible computing power in 1980s, numerous improvements for MC methods have been developed in order to alleviate the quasi-nonergodicity problem [141, 142, 139, 143]. The most popular currently used algorithms can be broadly classified into several groups. *Cluster-flip* methods improve single-flip ones by introducing non-local jumps in the state space by updating groups of interacting particles at once [36–38], thus accelerating the exploration of the state space. *Multicanonical algorithm* [39, 40] (MUCA) (also known as entropic sampling [41, 42] and adaptive umbrella sampling [43, 44]) assigns the non-Boltzmann weight factors to each state so that a uniform potential energy distribution is achieved, thus ensuring that the simulated system exhibits a free random walk in the energy dimension. This helps the system to escape from local energy minima. The MUCA requires either some prior knowledge of the weight factors for each state or running multiple iterations of the simulation with an increasing accuracy. A similar increasingly popular method known as *Wang-Landau* [144–146] (WL) algorithm recursively updates the weight factors within a single simulation in order to obtain a uniform

potential energy distribution and thus eventually leads to a random walk in energy dimension. Since the algorithm alters the potential surface while the simulation is running (as opposed to changes of potential in between the distinct iterations of the simulation in MUCA), the detailed-balance condition is not strictly obeyed. However, by gradually decreasing the incremental change in the potential surface the condition is achieved asymptotically. Another related method is *simulated tempering* [46, 3] (ST) (also known as method of expanded ensemble [45]). The ST method performs a random walk in temperature dimension, using a similar technique of reassigning weight factors to each state, as in the MUCA. This, in turn, induces a random walk in energy space allowing the system to leave the free-energy minima easier. *Replica exchange method* [36, 47, 48] (REM) (also known as parallel tempering [3], multiple Markov chain method [50], parallel annealing [49]) introduces multiple non-interacting copies (or replicas) of the original system at different temperatures, which are concurrently simulated by a conventional MC method. In REM, every few MC steps, pairs of replicas are allowed to exchange their states with specified transition probabilities set up to maintain the detailed balance. Given that the highest chosen temperature is sufficient to overcome all free-energy barriers, in theory, each replica is supposed to explore the entire state space ergodically. Finally, *infinite swapping* approach [147] aims to achieve an equivalent behaviour as a theoretical limit of infinitely frequent replica exchange method. By constructing a symmetric sampling distribution, the method entangles temperature and coordinate information, thereby increasing the flow of temperature-related information within the computational ensemble. All of the mentioned methods and any of their possible combinations have their strengths and weaknesses which have been thoroughly discussed in Refs. [139, 148, 18, 141, 149, 150].

The most intense recent developments were focused on REM and its modifications such as REM-WL [151] as well as generalised multidimensional REM [152] (MREM), which enables the exchange of replicas between the simulations not only at a range of different temperatures, but also while varying any other parameter of the Hamiltonian, e.g. chemical potential μ. The main observation which led to MREM is that as long as there is a number of non-interacting replicas of the original system, the Hamiltonian of the system does not have to be identical among the replicas and it can depend on any number of parameters with distinct parameter values for different replicas. Given correct exchange probabilities for any two replicas, the algorithm can maintain the detailed balance condition while naturally sampling such multidimensionally extended ensemble, and therefore overcome the free-energy barriers. However, while in theory MREM is a very powerful method to tackle quasi-nonergodicity problem, two main issues arise from the computational constraints. It has been shown that in a system of N particles, the number of required replicas increases as $N^{1/2}$ [47], which

prevents practical implementation of the MREM beyond 1 or 2 dimensions due to the large amount of processing power required. The second issue is caused by the slow cycling of each replica through the parameter space. In the simplest case of 1d REM, where multiple replicas are simulated at different temperatures, in order to ensure ergodicity, each replica should have enough time to cycle through all the temperatures. If the number of replicas is too small, the exchange rates will be negligible due to the small overlap of the sampled state space, but if the number of replicas is too large, then it will take too long to cross the entire parameter space while performing a random walk [153, 154]. Application of MREM drastically exacerbates this issue, by significantly expanding the parameter space through which each replica should traverse. In fact, these limitations and high processing power requirements of REM and its multidimensional counterpart led to development of combined methods, such as REM-MUCA [155] and REM-ST [156]. In these combined methods, REM is used only for obtaining the initial estimates of the weight factors used in MUCA (or ST), and the final accurate simulations are then performed by MUCA (or ST), which requires only a single copy of the system to be simulated. However, the combined methods heavily depend upon parameter fine-tuning and also face further problems [149, 141, 156, 155].

Throughout this project, the quasi-nonergodicity problem has been a constant challenge, and a large proportion of the work described in Ch. 2-4 has been dedicated to escape or overcome it. In Chs. 2 and 3, we explore several methods that help to alleviate the problem and test it on silica aerogel systems, and in Ch. 4 we apply one of the developed techniques to study sorption in Vycor glass.

1.5 Structure of the Thesis

In this thesis, we set out to study sorption phenomena in disordered porous media. We develop and apply a variety of numerical methods to examine the phenomena within the lattice-gas model.

We start Ch. 2 by introducing a visualisation technique to display the state space of the lattice-gas systems. The first half of the Chapter is dedicated to the study of relatively small systems by a method which relies on the exact enumeration of all possible states that the system can be in. We continue by developing kMC-based techniques to investigate larger systems. Finally, we propose several methods to obtain the degeneracy of the states in such systems and test them on the models of silica aerogel.

In Ch. 3, we update and explore the jump-walking MC algorithm designed to study systems with complex free-energy landscapes. By taking advantage of the current state of computing technology (as oppose to that available at the time when the method was originally

developed), we upgrade the algorithm and apply it to lattice-gas systems, which previously could not have been done. Using the updated algorithm, we examine the equilibrium isotherms for silica aerogel and Vycor glass systems in Chs. 3-4. The sorption mechanism in Vycor glass is investigated in Ch. 4, addressing the long-standing question about the existence of the first-order phase transition in such systems.

We develop a real-time interactive simulation engine in Ch. 5. The engine is comprised of a set of computational tools designed to provide a broad view of a simulated system in real-time, while simultaneously allowing the user to alter the parameters of the system. We discuss the structure, user interface and the implementation of the application. Lastly, we present several usage examples to illustrate the capabilities of the engine.

Sorption phenomena in lattice and network representations of soil and aerogel systems is studied in Ch. 6. We start by constructing a different type of disordered porous media model, namely soil networks. Sorption behaviour is examined by kMC simulations developing a qualitative description of the phenomena in such systems. Recently developed analytical techniques [98] have been successful in modelling sorption phenomena in the systems of random networks. Motivated by such results, we explore a possibility to represent the models of real disordered porous media by the particularly constructed random graphs. We analyse the properties of the system that determine the discrepancies between the sorption phenomena in the original models and their random network representations.

The conclusions and implications of the findings are given in Ch. 7. The main results are summarised and the possible directions for further development of this work are suggested.

Chapter 2

The state space of lattice-gas model

2.1 Introduction

All information that can be known about the system is encoded in its state space - a complete set of all possible configurations that the system can be in. In a lattice-gas model with N cells that can be either occupied by fluid or empty, the state space is a N-dimensional cube, with 2^N edges marking the states. At any one point in time the system is in one of those states, and can jump across some side of such cube, as one of the cells switches from fluid to empty or *vice versa*. Understanding how exactly the state space depends on a particular system, what are the subspaces that are most frequently occupied in equilibrium, how precisely the drift from one equilibrium position to another takes place, etc., could provide a deeper and more fundamental understanding of such systems. However, the state space, when considered as a multi-dimensional cube, is neither easy to visualize, nor does it provide valuable physical insights. In this Chapter, we explore different ways to probe, study and visualize the state space. We aim to investigate how its properties depend on the microstructure of the original sample. In addition, we study how system behaviour varies with the external conditions, and how it is reflected in the paths that the system takes across its state space.

In Sec. 2.2.1, we start from the analysis of a very small toy system, where the general methodology and visualisation techniques are introduced. We continue with larger systems that can still be handled exactly by enumerating all possible states. Throughout Sec. 2.2, we develop the basic understanding about the behaviour of the system and how it varies with the external conditions. It is also analysed how various features of the system and its behaviour are reflected in its state space. The limitations of the exact methods are discussed in Sec. 2.2.5. Based on the knowledge developed in Sec. 2.2, we continue the study of larger systems using MC simulations. We start Sec. 2.3 by validating MC techniques on the small samples that have previously been analysed exactly. In further sections, we explore the behaviour

of different porosity aerogel systems while varying the chemical potential and temperature. Finally, in Secs. 2.3.9-2.3.10, we develop several ways to analyse the degeneracy of the state space, and discuss the main challenges encountered in such study.

2.2 Exact methods

In this section, we introduce the core principles and develop the basic techniques that will guide the analysis of any system, its properties and behaviour throughout the rest of this work. Here, we are looking at relatively small systems, for which all possible microstates can be enumerated. We seek to find the relationships between the microstructure of the sample itself, the possible states that the system can visit and the behaviour of the system as the external variables are varied. In final sections, we also discuss the limitations of the discussed methodology.

2.2.1 Visualizing the state space

Everything that needs to be known about the system can be gathered by analysing the set of all possible microstates that this system can be in. Let us consider a simple example shown in Fig. 2.1. The system has $N = 4$ pore cells, which can be occupied by a fluid or remain empty. The cell on the left has two matrix neighbours and one pore-cell neighbour, similarly the two cells on the right have one matrix neighbour each as well as two pore-cell neighbours. Such system has $2^4 = 16$ possible microstates. Throughout this work, we are mainly interested in the occupation number N_f, i.e. the number of pore cells occupied by fluid, and the total interaction energy $\mathscr{H}$ of a microstate.

The occupation number N_f is the number of cells occupied by fluid in a particular microstate. As described in Sec. 1.3, the interaction energy $\mathscr{H}$ is the sum of all fluid-fluid and matrix-fluid interactions with w_ff and w_mf representing each interaction strength, respectively. Thus, we can simplify the lattice-gas model Hamiltonian (see Eq. (1.2)) to be applied to the off-lattice sample system:

$$\mathscr{H} = -w_\mathrm{ff} n_\mathrm{ff} - w_\mathrm{mf} n_\mathrm{mf} \,, \tag{2.1}$$

where n_ff and n_mf is the number of fluid-fluid and matrix-fluids interactions (links), respectively. As before (Sec. 1.3), the energy scale is defined by setting up $w_\mathrm{ff} = 1.0$ and wettability $y = w_\mathrm{mf}/w_\mathrm{ff}$ is set to $y = 2.0$. The resulting energies of each microstate are shown in Fig. 2.1. The number of the fluid cells N_f in the system and thus the energy $\mathscr{H}$ of the system can vary, but the chemical potential μ, temperature T and volume V (or equivalently N) of the system

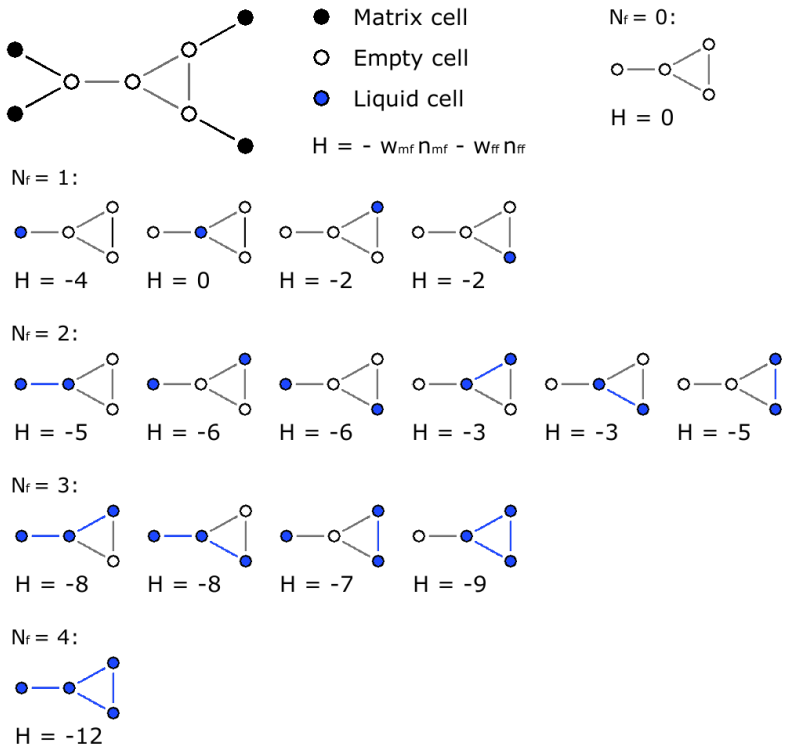

Fig. 2.1 A toy sample of a lattice-gas system and all the possible microstates that such system can be in. Open, solid blue and solid black circles represent empty, fluid and matrix cells, respectively.

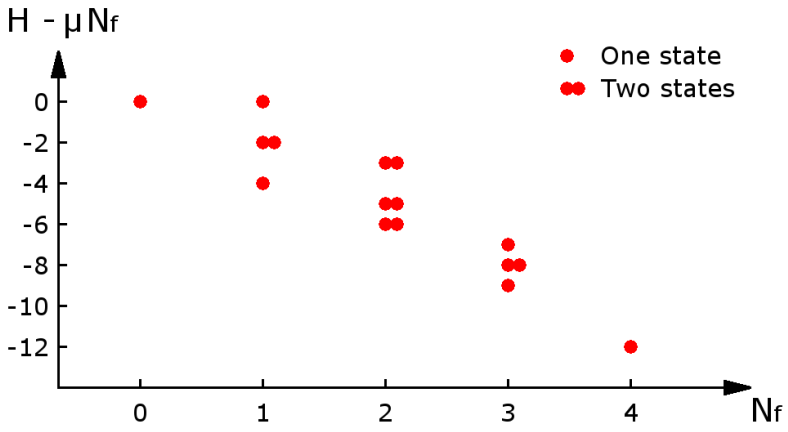

Fig. 2.2 State-space diagram of the toy system shown in Fig. 2.1 at $\mu = 0$. The degeneracy of each macrostate in this particular system, can only take two values: 1 (single circles) or 2 (double circles).

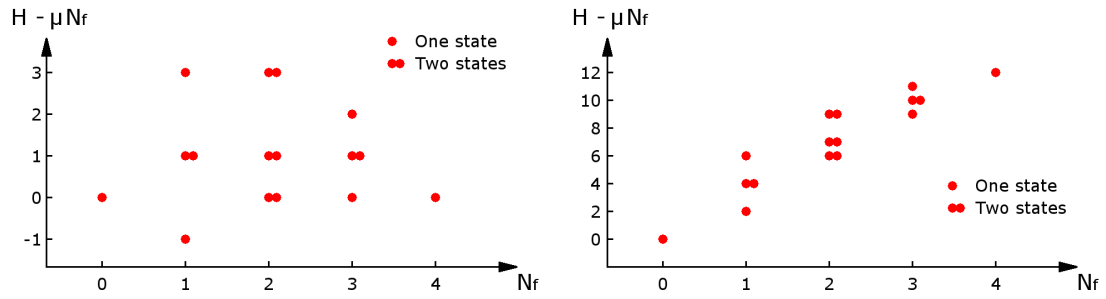

(a) $\mu = -3.0$ (b) $\mu = -6.0$

Fig. 2.3 State-space diagrams of the toy system shown in Fig. 2.1 at different values of μ as labelled in panels (a) and (b).

are fixed. For a given set of μ, T and N, the system can be in a set of states which form a grand canonical ensemble described by the grand partition function given by Eq. (1.3) in Sec. 1.3.

In order to visualise the multi-dimensional state space of the system, it is convenient to study its projection onto two dimensions representing the quantities of interest N_f and $\tilde{\mathscr{H}} = \mathscr{H} - \mu N_\mathrm{f}$. A similar visualization technique has been previously used to study random-field Ising model systems [10]. In general, microstates can have the same values of N_f and $\tilde{\mathscr{H}}$ and it is convenient to define a macrostate $\mathscr{Q}(N_\mathrm{f}, \tilde{\mathscr{H}})$, with the degeneracy $g(\mathscr{Q}(N_\mathrm{f}, \tilde{\mathscr{H}}))$, as a set of g microstates with the same values of N_f and $\tilde{\mathscr{H}}$. The 11 macrostates obtained from 16 microstates in our sample system at $\mu = 0$ are show in Fig. 2.2, each with degeneracy either 1 or 2 (represented by single or double circles). One can observe that the lowest-energy macrostate represents fully condensed system (with all pore cells occupied by fluid), and since the occupation probabilities of each state are governed by Boltzmann distribution (Eq. (1.7)), at low temperatures, this state dominates the system. Decrease in μ by $\Delta\mu$ shifts each state with N_f fluid cells towards the higher energies by $\Delta\mu N_\mathrm{f}$. Eventually, this causes states with lower fluid occupancy to become more favourable, as shown in Fig. 2.3a and Fig. 2.3b.

The probability, $P(\mathscr{Q}(N_\mathrm{f}, \tilde{\mathscr{H}}))$, to find the system in a particular macrostate $\mathscr{Q}(N_\mathrm{f}, \tilde{\mathscr{H}})$ is discussed in detail in the next Chapter. Here, we are primarily interested in the probability of the system to be in a state with a particular value of N_f, $P(N_\mathrm{f}) = \sum_{\tilde{\mathscr{H}}} P(\mathscr{Q}(N_\mathrm{f}, \tilde{\mathscr{H}}))$. At high temperatures, all the microstates are equally likely (see Eq. (1.7)) and the probability distribution is dependent only on the degeneracy of each macrostate, therefore in the infinite-temperature limit, $P(N_\mathrm{f})$ is simply proportional to the binomial coefficients, $\binom{N_\mathrm{f}}{N}$, which are independent of μ. Under such conditions, each cell is equally likely to be empty or filled, and thus $P(N_\mathrm{f})$ has a peak at $N_\mathrm{f} = N/2$. In large systems, this peak is increasingly more

2-bit	3-bit	4-bit
00	000	0000
01	001	0001
11	011	0011
10	010	0010
	110	0110
	111	0111
	101	0101
	100	0100
		1100
		1101
		1111
		1110
		1010
		1011
		1001
		1000

Fig. 2.4 Gray code for 2-, 3- and 4-bit binary numbers. A microstate of an N-cell system can be represented by N-digit binary number, and thus Gray code can be used to generate all possible configurations of the system in such order that each consecutive configuration differs from the previous one only at a single cell, as marked by digits in red.

pronounced, since under such conditions, the binomial distribution is well approximated by the Gaussian, and the width of the peak $\propto N^{-1/2}$. At very low temperatures, on the other hand, the occupation probability of a macrostate is determined by its exponential dependence on $\tilde{\mathscr{H}}$. Hence, $P(N_\mathrm{f})$ has peak(s) around the state(s) with the lowest values of $\tilde{\mathscr{H}}$, which tend to δ-function(s) as $T \to 0$. Visualizing the projection of the state space onto two dimensions as above allows us to better understand the behaviour of the system at various values of temperature and μ. More detailed examples of the state-space analysis are presented in the following sections.

2.2.2 Enumerating the microstates. Gray code

In this Chapter, we rely upon our ability to enumerate all possible microstates of the system. In this section, we will propose an efficient way to achieve this and discuss the limitations of such methodology.

In order to calculate any observable of interest in equilibrium at some particular values of μ and β, one needs to sum over all possible microstates of the system. This task is often very computationally expensive, since the observable can be a function of the variables characterising all the cells and the interactions between them. However, it is possible to

simplify the procedure significantly by choosing a particular path over the state space, which ensures that each consecutive configuration along such a path differs from the previous one only at a single site. This way, one can store the value of the observable at the previous configuration and alter it according to the change at a single cell that has been either filled with or emptied of the fluid. This leads to a constant-time operation for calculation of the observables for each microstate, as opposed to being dependent on some degree of N.

Let us consider a system of size N, where each cell can be either filled with fluid or empty. The state of such system can be written down by N-digit binary number, that we denominate as microstate signature. Gray code [157] (also known as reflected binary code) is a method to alter the digits of a binary number in such a way that all possible numbers from 0 to 2^{N-1} are obtained only by switching one digit at a time. Fig. 2.4 displays the examples of such sequences in the case of $N = 2,3,4$, with the most recently altered digit marked in red.

By employing efficient algorithm utilising bit-shift operations, it is possible to generate Gray code "on-the-fly", without storing any of the state variables for each microstate or even the microstate signatures themselves. The algorithm ”knows” only the current state and the partial sums of the observables that we aim to obtain. This resolves memory problems which would be faced even for very small systems, if any microstate level information was stored. However, the exact enumeration methods fundamentally cannot escape the second computational limitation, that is the processing power. Just to enumerate all the states in the system of size $N = 50$, taking a single operation per state and with a processing unit of 1GHz, the computation would take approximately 13 days. And even though it is possible to employ more powerful computers, since each added cell in the system doubles the computational power required, even the most powerful currently existing supercomputer would not be able to tackle the system of size $N = 100$, when employing such methodology. In some cases, it is also possible to reduce the number of states by exploiting the symmetries in the system. However, in this work, we are interested in analysis of random disordered natural systems such as soil, aerogel or Vycor glass, which exhibit no symmetries and thus no significant progress can be achieved this way.

2.2.3 Variation of projected state space with μ and β

Let us apply the same analysis as in Sec. 2.2.1 to a larger system shown in Fig. 2.5a. The system consists of 5×5 pore-space cells arranged in a simple square lattice with matrix boundary cells as shown in the Fig. 2.5a. In this case, the state space has $2^{25} \simeq 3.4 \times 10^6$ microstates. The microstates are grouped into macrostates and projected onto a plane of N_f and $\beta\tilde{\mathscr{H}}$ dimensions (see Fig. 2.5b). Here, we use $\beta\tilde{\mathscr{H}}$ instead of simply $\tilde{\mathscr{H}}$ as before, since it is the exact exponent entering the Boltzmann equation (see Eq. (1.7)), thus making it

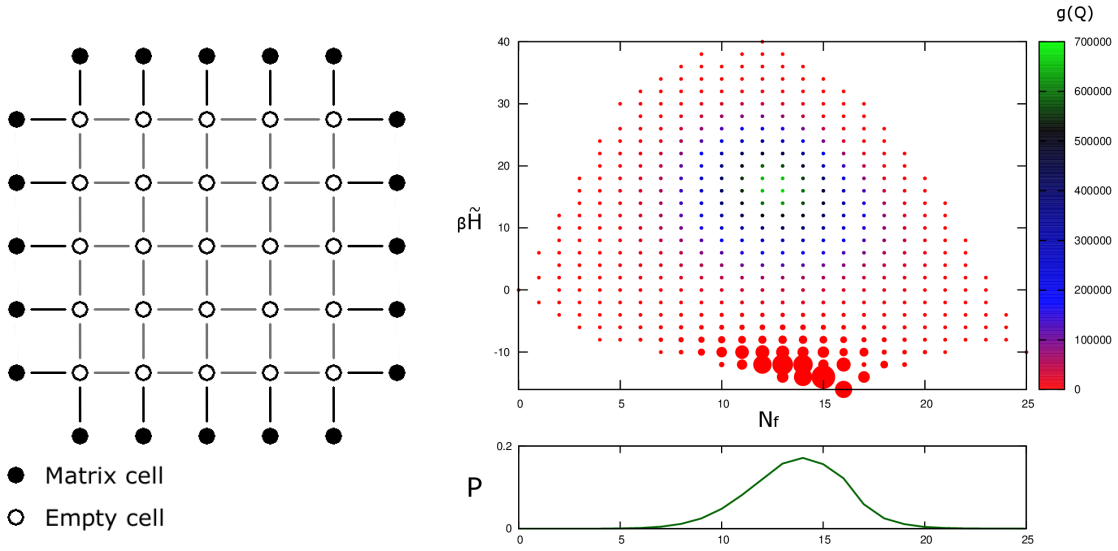

(a) (b)

Fig. 2.5 (a) Schematic representation of the 25-cell system on a square lattice. (b) State-space diagram of such system at $\mu = -3.0$ and $\beta = 2.0$. The size and colour of the circles represent the equilibrium probability and degeneracy of the corresponding macrostates, respectively. The colour scheme is given by vertical bar on the right of panel (b).

easier and more informative to compare between the state spaces produced from different systems at differing values of β. The colour and the size of each circle in Fig. 2.5b (the same colour and size schemes are used for all similar state-space diagrams in this Chapter) represent the degeneracy $g(\mathscr{Q})$ and the total occupation probability $P(\mathscr{Q})$ of each macrostate, respectively. However, the function determining the size of each circle is not linear. It is constructed in such a way that even very improbable states could be visible with a minimum constant-size circle, and the most probable states would be represented by logarithmically increasing size of the circle depending on the probability of the macrostate. The lower panel in Fig. 2.5b shows the probability distribution $P(N_\mathrm{f})$, obtained by simply summing up the probabilities of all macrostates at a particular value of N_f.

Such visualisation reveals several interesting features of the system. First of all, we observe that the lowest energy state at $\mu = -3.0$ is achieved when $N_\mathrm{f} = 16$, in particular, it is the microstate when all 16 of the boundary cells that touch the matrix are filled with fluid, and the inner 9 cells are left empty. Let us denominate the macrostate that contains this microstate as S_{16}. The degeneracy of S_{16} is $g(\mathscr{Q}(N_\mathrm{f} = 16, \beta\tilde{\mathscr{H}} = -16)) = 1$, since there is only one such configuration. The occupation probability of this individual microstate is the highest in the system and is approximately $\simeq 0.052$. However, even though the lowest-energy state with $N_\mathrm{f} = 15$ (S_{15}) has a higher energy than S_{16}, the total probability of macrostate S_{15}

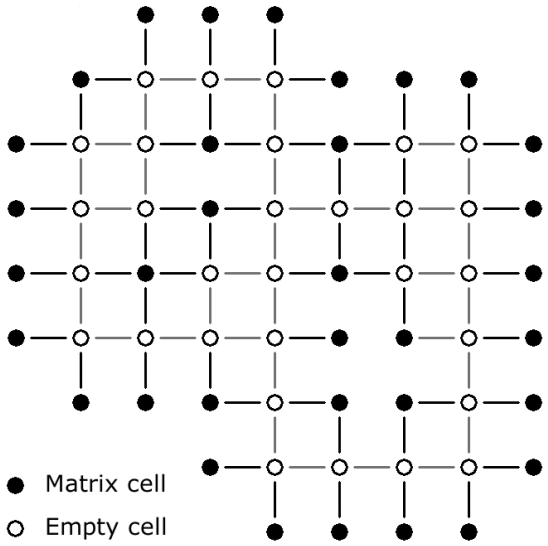

Fig. 2.6 Schematic representation of a 30-cell system. The particular system is discussed in more detail in Sec. 2.2.4.

is larger. The macrostate S_{15} differs from S_{16} simply by the presence of one empty cell out of 12 non-corner boundary cells. This results in degeneracy $g(\mathscr{Q}(N_\mathrm{f}=15,\beta\tilde{\mathscr{H}}=-14))=12$. Therefore, although the individual microstate at $(N_\mathrm{f}=15,\beta\tilde{\mathscr{H}}=-14)$ is $e^{-14-(-16)}=e^2\simeq 7.4$ times less likely than the microstate at $(N_\mathrm{f}=16,\beta\tilde{\mathscr{H}}=-16)$, the macrostate S_{15} is $12\times e^{-2}\simeq 1.6$ times more probable than S_{16}, and we can see this feature being reflected in the size of the corresponding circles representing the discussed macrostates.

Similarly, it follows from Fig. 2.5b that the least probable state also having degeneracy $g=1$ is realised for $N_\mathrm{f}=12$. This is the check-board configuration which avoids creating interactions between cells as much as possible. However, the upper half of the projected state space is almost never relevant when calculating $P(N_\mathrm{f})$ or other quantities of interest. This is due to the fact that at low temperatures only the lowest energy states are occupied while at high temperatures only the highest degeneracy macrostates are occupied. As it can be seen from Fig. 2.5b, the highest degeneracy macrostates are grouped around the middle of the projected state space, thus as the temperature is increased the system drifts towards those macrostates independently of μ.

Let us consider a more complex system shown in Fig. 2.6. It consists of 30 empty cells with 8 matrix cells in the interior region of the pore space. This particular system will be discussed in more detail in Sec. 2.2.4, while here, we focus on general impacts of variation in μ and β on any state space. Fig. 2.7 shows in more detail how the significance of each region of the projected state space changes as we vary μ and β. It is produced by compiling

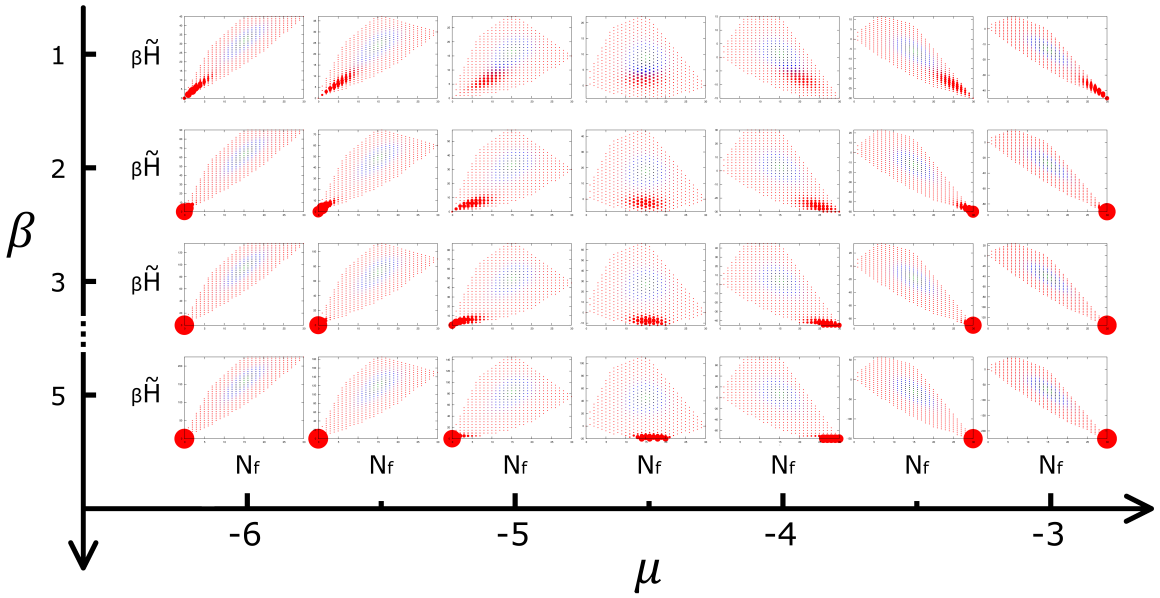

Fig. 2.7 A collection of the state-space diagrams of the 30-cell system for various values of β and μ.

the state-space diagrams such as that shown in Fig. 2.5b for several values of μ and β. The change in μ results in a vertical shear transformation of the state space, i.e. each state is moved along the $\beta\tilde{\mathscr{H}}$-axis proportionally to N_f coordinate of that state as discussed in the previous Section (cf. the panels in the same rows in Fig. 2.7). Such transformation maintains the state at $(N_\mathrm{f}=0,\beta\tilde{\mathscr{H}}=0)$ point fixed, while shifting the states that are to the right of this point. As μ increases, the states with larger values of N_f go further down in energy. In turn, at equilibrium, the system visits those states more frequently, thus shifting the equilibrium N_f toward the higher values. The change in β, on the other hand, preserves the same relative position of the states (cf. the panels in the same columns in Fig. 2.7), but stretches the $\beta\tilde{\mathscr{H}}$ axis, thus resulting in a wider (narrower) distribution of probabilities across the state space as the temperature is increased (decreased). This is reflected in change of the circle sizes for the panels in the same column but different rows in Fig. 2.7. At lower temperatures (higher β values), the most probable states are highly clustered around the bottom of the projected state space. This can be explained by analysing two consecutive macrostates at some particular value of μ. We know that the occupation probability of a macrostate depends on the balance between the increasing degeneracy as we move towards the centre of the state space, and the decreasing probability of each microstate as we move up the $\beta\tilde{\mathscr{H}}$-axis. As β is increased, the degeneracy of the macrostates does not change, but the difference in $\beta\tilde{\mathscr{H}}$ values of the considered macrostates is increasing. Hence, the most probable states shift down in the state space. Finally, once the macrostate with the highest probability is the one at the lowest

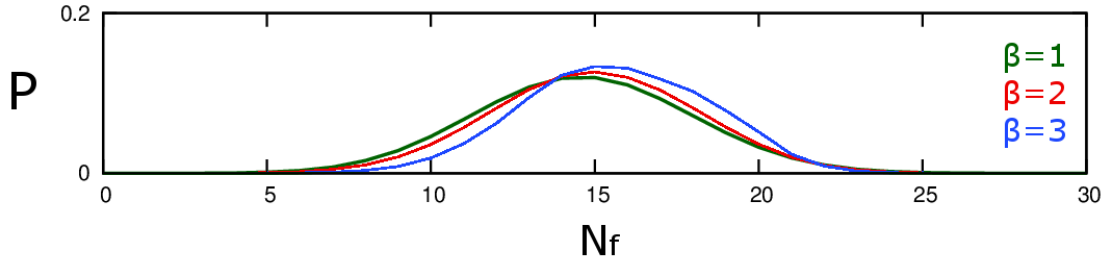

Fig. 2.8 The probability distribution, $P(N_\mathrm{f})$, at $\mu = -4.5$ for three different values of β as indicated in the figure.

energy, further lowering the temperature simply increases the probability difference between that state and the rest.

Another important observation can be made about the path that the system takes to traverse the state space as μ is varied, and how it results in the familiar $N_\mathrm{f} - \mu$ diagrams (sorption isotherms). Here, the path refers to the subset of the macrostates that are most probable at equilibrium, and how this subset drifts across the state space as external conditions are varied. At high temperatures, the transition from low to high values of N_f happens over a wide range of μ (see the $\beta = 1$ row in Fig. 2.7), i.e. the system is not completely empty even at $\mu = -6$ and not yet fully filled with fluid at $\mu = -3$, as indicated by wide regions of relatively large red circles for the most left and most right panels. While at low temperatures (the $\beta = 5$ row in Fig. 2.7), the transition happens almost entirely between $\mu = -5$ and $\mu = -3.5$, as reflected by presence of a single large red circle in the corresponding panels. This behaviour results in a known phenomenon of steepening of the sorption isotherms with decreasing temperature. However, in this particular system, even though the transition from low to high values of N_f is sharper at lower temperatures, it is still continuous and for any μ there is single-peaked distribution in N_f, thus no discontinuity in the order parameter, N_f, is observed. More detailed look at the phase transitions and the conditions on the shape of the state space required for them to occur will be analysed in the following Sections.

It is apparent that the two-dimensional visualisation of the state space allows to discriminate between the different sets of states that the system in equilibrium resides in, whereas the more frequently used $P(N_\mathrm{f})$ distribution would not. This can be seen from comparative analysis of the three top panels in the column corresponding to $\mu = 4.5$ in Fig. 2.7. The $P(N_\mathrm{f})$ distributions in $\beta = 1,2,3$ cases (see Fig. 2.8) are relatively similar, even though the subsets of the occupied states at equilibrium are substantially different, possibly resulting in very different properties of the system in those conditions. Thus the additional projection of the state space onto the $\beta\tilde{\mathscr{H}}$-axis allows us to distinguish between a wider range of different states of the system and to gain a better understanding of its behaviour. Such ability to

have a more detailed look at the system becomes especially useful when considering phase transitions, discussed in the following Section.

2.2.4 Phase transitions

The concept of phase transitions in porous media have been introduced and discussed in detail in Sec. 1.2. Usually, the notion of a phase transition is defined for an infinite system, however, it is very informative to consider it in small systems that can be numerically handled exactly, since this provides great insights into the phenomenon in general.

Fig. 2.9 displays the state space of 5×5 empty cells arranged in a square lattice with periodic boundary conditions. Such system inherently is a two-dimensional empty space. The shape of the projected state space indicates the presence of low-energy states at extremal values of N_f, and absence of low-energy states in between. At a low enough temperature, as shown in Fig. 2.9 for $\beta = 1.6$, such shape of the projected state space leads to a very sharp transition from completely empty system to fully occupied by fluid. In fact, at $\mu = -2.0$, we observe coexistence of both phases, since $P(N_\mathrm{f})$ has two peaks (see the lower panel in the middle column in Fig. 2.9), denoting two preferred configurations that the system can be in. Such phenomena exists in any system, whose projected state space has a concave lower boundary for some range of N_f values, or the low-energy states within some range of N_f cannot be accessed. At high temperatures, the path that the equilibrium states take across the state space is continuous as in Fig. 2.7 at $\beta = 1$, since the most probable states are free to shift freely in the entire range of N_f values, as μ is varied. However, at low enough temperature only the lowest energy states are important. In the absence of sufficiently low-energy states for some range of N_f, the system gets stuck in the "valleys" of the state space, and is not free to take any value of N_f. Thus concave lower boundary of the state space for any range of μ values leads to multi-peak $P(N_\mathrm{f})$ distribution and therefore to what we will denominate as a finite-system phase transition.

Let us consider how the shape of the state space depends on the microscopic structure of the sample. A completely empty sample as that shown in Fig. 2.9 results in characteristically concave lower boundary of the projected state space. At $\mu = -2.0$ the states representing a completely empty system and fully occupied system are at the same value of $\tilde{\mathscr{H}} = 0$. This is because the mean interaction energy per cell is -2, and thus the change in $\tilde{\mathscr{H}} = \mathscr{H} - \mu N_\mathrm{f}$ as the entire system is filled with fluid is zero. However, for low values of N_f, as we fill cells with fluid one by one, on average, we create less than 2 fluid-fluid connections, thus the system goes up in $\tilde{\mathscr{H}}$ even for the best possible packing of fluid cells (the lowest energy configurations). As the system gets close to being fully filled with fluid, any additional fluid

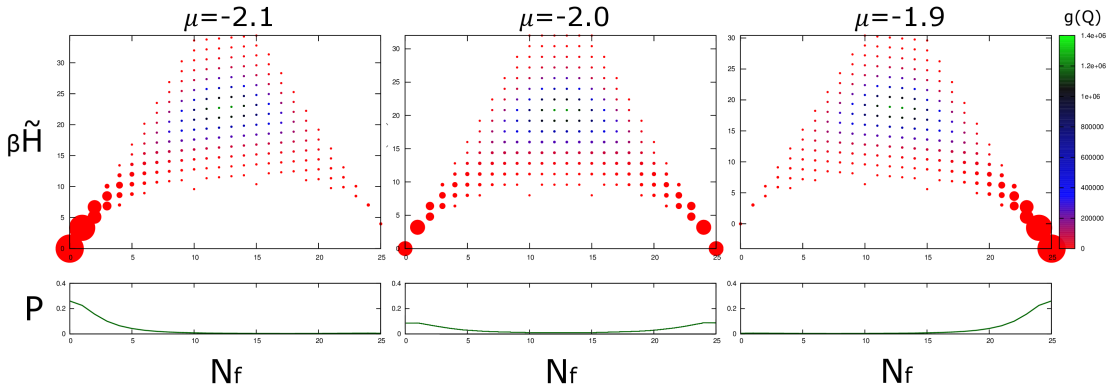

Fig. 2.9 State-space diagram of the 25-cell system similar to that shown in Fig. 2.5, but instead of the matrix cells at the edges, this system has periodic boundary conditions. The diagrams are shown for various values of μ (as indicated in the figure) at $\beta = 1.6$.

cell creates more than 2 fluid connections on average, thus creating the second "valley" in the state space at high values of N_f.

In a similar system with a single matrix cell added (see Fig. 2.10), the projected state space maintains a concave lower boundary, however, the curvature is significantly lower. In order to understand the changes in the structure of the projected state space it is again convenient to imagine the process of filling the system with fluid one cell at a time. Empty system always starts at $\tilde{\mathscr{H}} = 0$. However, the first fluid cell, instead of adding no interaction energy, can be near the matrix cell, and therefore the change in $\tilde{\mathscr{H}} = \mathscr{H} - \mu N_\mathrm{f}$ is zero for the lowest energy state at $\mu = -2.0$. Similarly, it is possible to find a way of adding an extra fluid cell, with at least 2 units of interaction strength for every following addition. Once the system is almost entirely full, average interaction strength per added fluid cell exceeds 2 as in the matrix-free system, thus creating a downwards slope in the state-space diagram. As expected (see Sec. 1.1), the phase transition in such system shifts toward the lower values of μ.

Further increase in concentration of matrix cells supports the same trends (see. Fig. 2.11). The negatively curved lower state-space boundary flattens, as the region of the state space corresponding to high values of N_f shifts downwards. This results in a phase transition shifted towards lower values of μ. Fig. 2.12 shows that, in this case, the finite-system phase transition occurs approximately at $\mu = -2.5$. However, the jump in fluid density does not correspond to transition from completely empty to fully occupied system as before. Instead, the low-density peak for distribution of N_f is located about $N_\mathrm{f} = 7$, while the high-density peak corresponds to the fully occupied system with $N_\mathrm{f} = 22$ (see the lower panels in Fig. 2.12). This is because the first cells filled by fluid are located near the matrix cells thus forming a fluid ”layer”

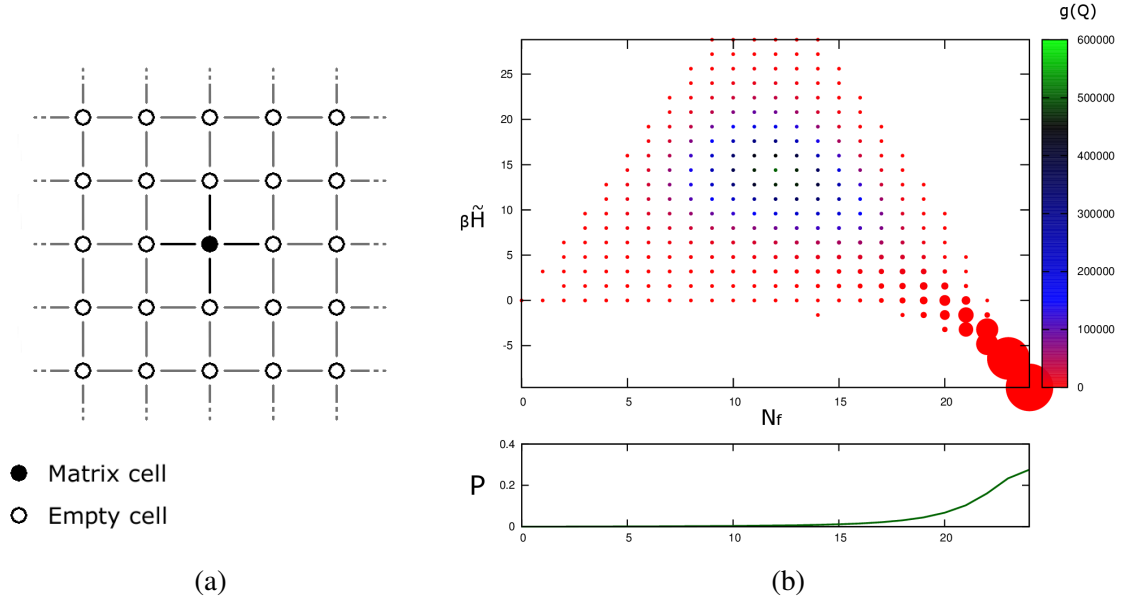

(a) (b)

Fig. 2.10 (a) Schematic representation of a 24-cell system with periodic boundary conditions. (b) State-space diagram of such system at $\mu = -2.0$ and $\beta = 1.6$.

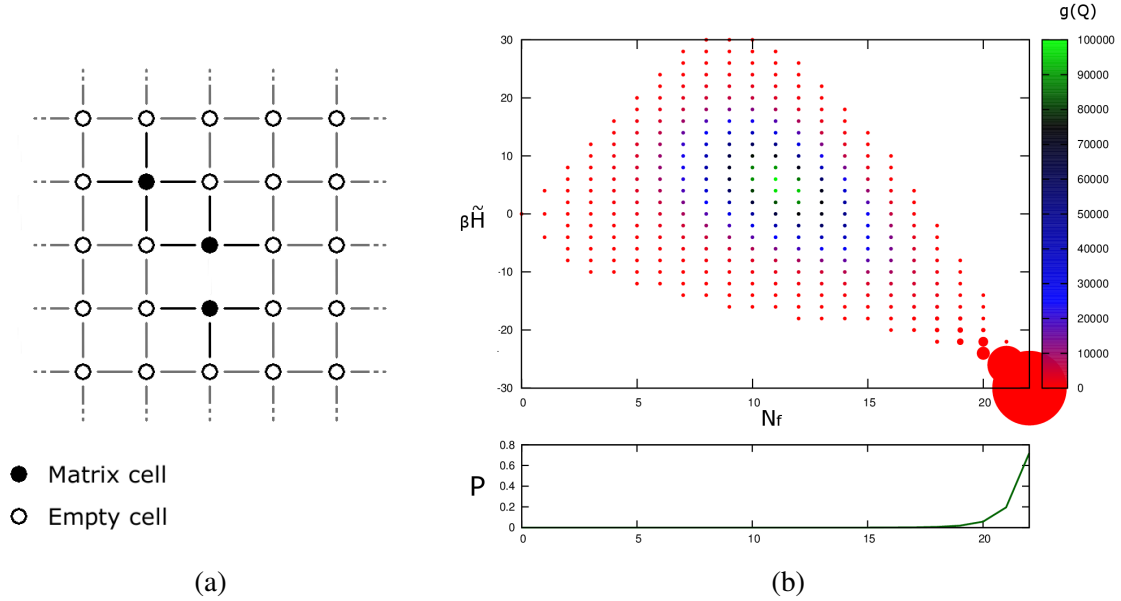

(a) (b)

Fig. 2.11 a) Schematic representation of a 22-cell system with periodic boundary conditions. b) State-space diagram of such system at $\mu = -2.0$ and $\beta = 2.0$.

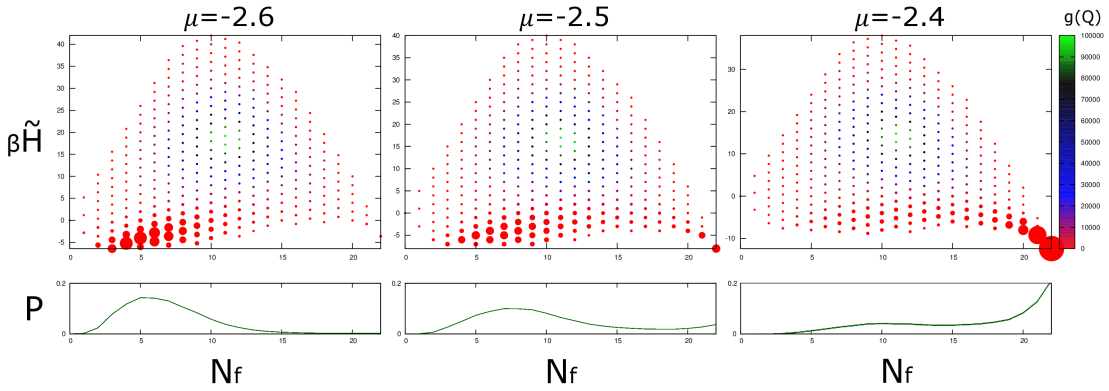

Fig. 2.12 A collection of the state-space diagrams of the 22-cell system shown in Fig. 2.11 for various values of μ (as indicated in the figure) at $\beta = 2.0$.

around these three matrix cells, and the phase transition happens between the configurations corresponding to the mostly empty system with such fluid layer and the completely filled system. Similar behaviour is observed in experiments examining adsorption in porous media as discussed in Sec. 1.1.

Finally, let us examine the system analysed in Sec. 2.2.3 for various β and μ values (Fig. 2.7). Fig. 2.13 shows the real-space lattice representation as well as the state-space diagram of the system. At the same value of $\mu = -2.0$, as in the case of previously analysed systems, we observe that the state space is heavily skewed downwards in $\tilde{\mathscr{H}}$ at the high values of N_f. This shows that, on average, the increase in the magnitude of interaction energy with every added fluid cell is significantly higher than 2. In fact, the middle column of Fig. 2.7 demonstrates that the system fully occupied by fluid has the same value of $\tilde{\mathscr{H}}$ as an empty one at $\mu = -4.5$, meaning that the average interaction energy of a fluid cell in a completely filled system is approximately equal to -4.5. The lower boundary of the projected state space is convex (see Fig. 2.7), and thus the system has a single-peaked distribution, $P(N_\mathrm{f})$, for all values of μ. Therefore, a finite-system phase transition does not occur in this system even at the lowest temperatures.

2.2.5 Flexibility, use and limitations of the exact methods

In the previous Sections of this Chapter, we have analysed a variety of small systems within the framework of the lattice-gas model. It is always possible to gain any desired information about a particular system simply by enumerating all of its microstates and summing over the variables of interest. However, as we have seen in Sec. 2.2.2, the exact enumeration method is highly limited by the size of the systems that are computationally tractable. Due to the exponential nature of the number of the microstates in the system, the exact methods in their

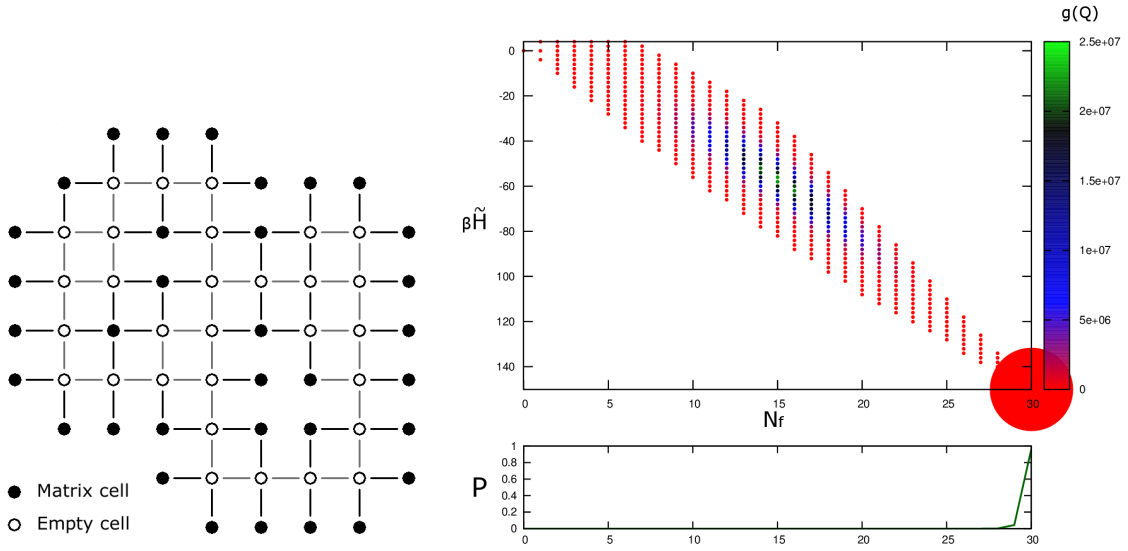

(a) (b)

Fig. 2.13 (a) Schematic representation of a 30-cell system. (b) State-space diagram of such system at $\mu = -2.0$ and $\beta = 2.0$. The projected state space of this system at different values of μ and β is shown in Fig. 2.7.

basic form can only be used to simulate small systems, independently of the computational resources that are available. Despite that, the information that we have obtained by such analysis greatly improves our understanding of any system and it's behaviour. Once the microstates are counted and assigned to the respective macrostates, any further examination of the system is reduced to summing over the chosen subset of macrostates. Having defined them according to the observables of interest (e.g. $\tilde{\mathscr{H}}$ and N_f), the result is equivalent to having the exact partition function of the system. Thus with a single procedure, we obtain the equilibrium properties of the system at any values of μ and β. The method is especially useful for gaining fundamental understanding and developing intuition about the interplay between the variables of interest, the microstructure of the system, its state space, and the resulting sorption phenomena. As we shall see in the following sections, most of the conclusions that have been obtained for small systems hold for larger ones and the main trends can be straightforwardly extrapolated.

Nevertheless, being restricted to small 2-dimensional systems prevents us from performing quantitative study of the phenomena observed in the experiments. Such constraints encourage us to look for methodologies that avoid direct enumeration of the microstates and, instead, to explore sampling techniques such as Monte-Carlo. In the previous sections of this Chapter, we have seen that it is possible to obtain information about the behaviour of the system simply from the shape of its state space. We have concluded that at sufficiently low

temperatures the system resides mostly in the lowest energy states, which also tend to have comparatively low degeneracy. Such observation supports additional motivation to explore sampling algorithms. In the following Sections of this Chapter and the rest of this work, we aim to develop the methodologies that bypass the limitations encountered by the exact enumeration algorithm, while simultaneously use the knowledge gained by this approach.

2.3 Monte Carlo methods

For large systems, that we are interested in, enumerating all of the microstates in the state space is a computationally impossible task as discussed above. In this Section, we therefore aim to investigate some of the techniques based on MC sampling. At the core of various sampling methods is the idea that it is possible to obtain desired information about the state space without visiting every single microstate. As we have seen in the previous Sections, often only a small fraction of the microstates significantly contribute to distribution $P(N_\mathrm{f})$, indicating that with an appropriate sampling strategy such methodology can be sufficient to determine the desired properties of the system. Moreover, it was shown that even just the knowledge about the general shape of the state space, provides us with useful information about system behaviour. Thus in the following sections, we concentrate on the large systems that are impossible to handle with the exact state-enumeration techniques, and aim to develop tools that help us to understand the relationships between the microstructure of a sample, the state space of such system and its behaviour within the lattice-gas model.

2.3.1 Comparison between the exact and MC methods

To check the accuracy and validity of the kinetic Monte Carlo (kMC) method introduced in Sec. 1.4.2, we first apply it to one of the small systems that we have analysed exactly in the previous section.

Fig. 2.14 compares between the exact and kMC methods applied to the previously described system, displayed in Fig. 2.11. The left panels in Fig. 2.14 show the results obtained by the exact state-enumeration method, the macrostates are colour-coded by degeneracy as previously in Fig. 2.12. In the right panels, the projected state space is displayed as it was sampled by the kMC method. The size of the circle that represents a particular macrostate is determined by the time fraction that the system has spent in that macrostate. If the system is ergodic, this time fraction is proportional to the equilibrium occupation probability of the state as discussed in Sec. 1.4.2, thus we expect the sizes of the circles to be identical between the left and right panels. In order to be able to make a more accurate comparison, the colour-

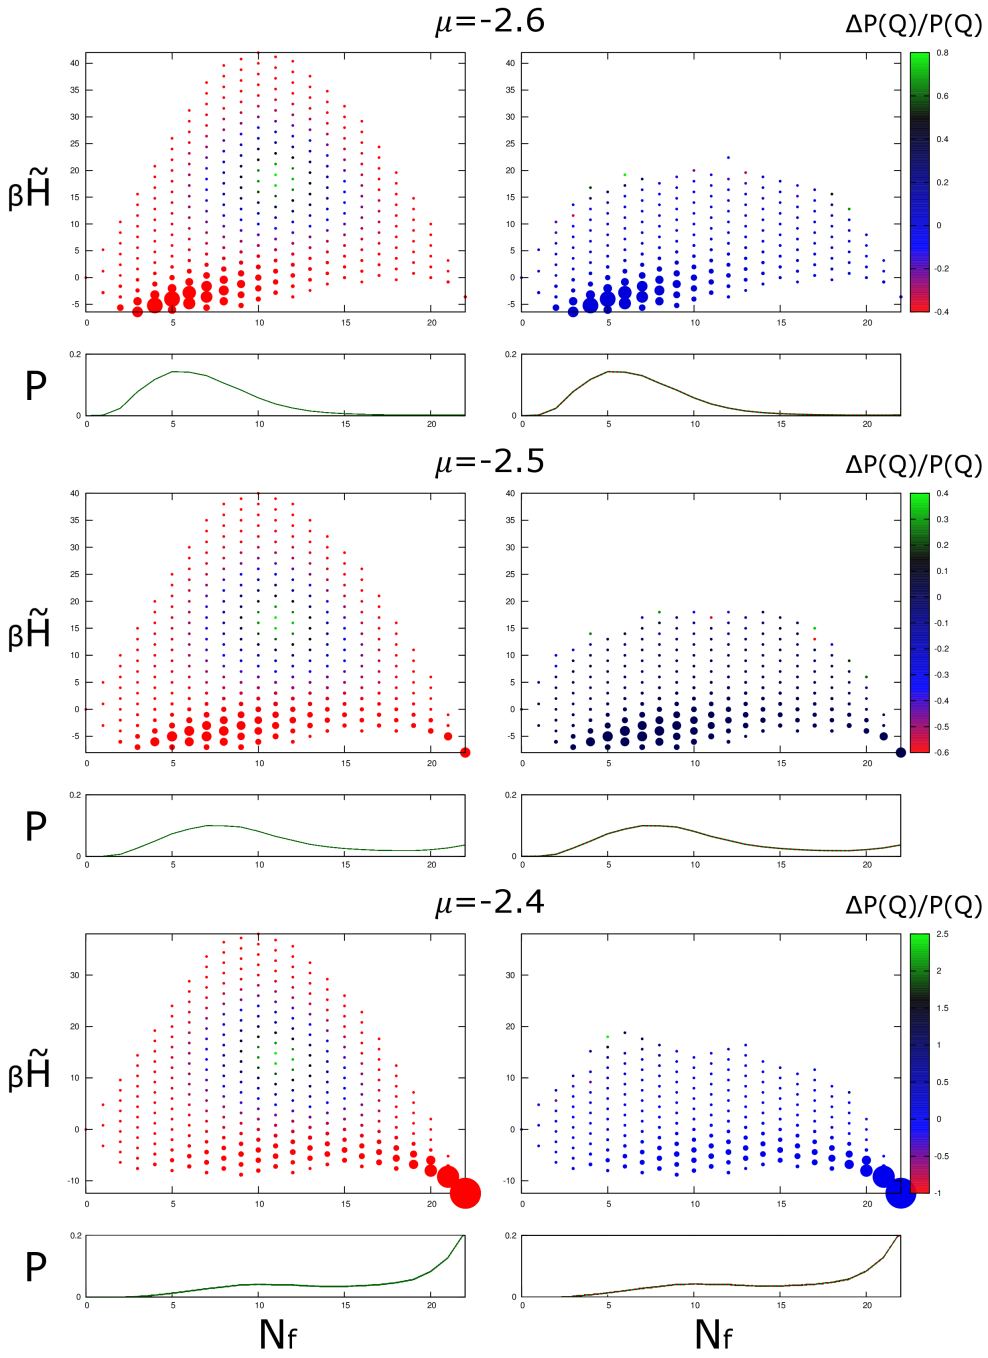

Fig. 2.14 Left panels: state-space diagrams of the 22-cell system shown in Fig. 2.11 obtained by exact state enumeration for various values of μ at $\beta = 2.0$. Right panels: the state-space diagrams as explored by the kMC simulations. The same circle-size code as in all similar diagrams is used for both panels. The circle-colour code in the left panels is also the same as in all similar diagrams but, in the right panels, the colour-code reflects the relative difference (see Eq. (2.2)) in occupation probability of the corresponding macrostates achieved by kMC and exact methods. The fluid density distribution, $P(N_f)$, is shown by solid green (left and right panels) and dashed-red (right panels) line for exact and kMC methods, respectively. Note, that the dashed-red lines overlap with green lines and thus are hardly visible.

coding of the circles in the right panels is based on the relative difference, $\Delta P(\mathscr{Q})/P(\mathscr{Q})$, between the equilibrium probabilities obtained by the exact and kMC methods, which is given by the following expression,

$$\frac{\Delta P(\mathscr{Q})}{P(\mathscr{Q})} = \frac{P_{\mathrm{kMC}}(\mathscr{Q}) - P_{\mathrm{exact}}(\mathscr{Q})}{P_{\mathrm{exact}}(\mathscr{Q})} \,. \tag{2.2}$$

To compare the resulting $P(N_{\mathrm{f}})$ distributions, below the state-space diagrams, the right panels display $P(N_{\mathrm{f}})$ graph, which contains two lines. The green line shows the distribution obtained by the exact-enumeration method, while the dashed red line represents the results obtained by the kMC method. Both lines identically overlap over the entire range of N_{f}.

Firstly, we observe that the kMC method, having a limited running time and thus accuracy, does not explore the entire state space of the system. The least favourable states at the top of the state space are never sampled. However, the colouring of the circles and overlapping green and dashed red lines in the right panels of Fig. 2.14 show that kMC method correctly obtains the probability of any macrostate that carries a non-negligible weight in equilibrium distribution. Detailed analysis of the results shows no systematic deviations in the macrostate probabilities obtained by the two methods. The large relative deviations appear only for the weakly sampled states at the top of the state space. Therefore, given the necessary assumptions (such as ergodicity) apply, the kMC method is a reliable tool that can be used to explore the properties of such systems. The explored fraction of the state space reduces with the size of the system. However, as long as the ergodicity is maintained, the weakly explored parts of the state space correspond to the states which only negligibly contribute to the occupation probability distribution. Detailed discussion of algorithms performance is presented in Ch. 5.

2.3.2 MC sampling in aerogel systems

Aerogel systems have been introduced in Sec. 1.1.2. Let us consider a cubic bcc-lattice sample of size $L = 18$ and porosity $\phi = 0.95$. Such sample contains $2 \times 18 \times 18 \times 18 \times 0.95 \simeq 11081$ pore-space cells that can be either occupied by fluid (fluid cells) or empty, and $2 \times 18 \times 18 \times 18 \times 0.05 \simeq 583$ matrix cells forming a rigid matrix skeleton. Even though it is a relatively small sample, the state space of such system is enormous containing $2^{11081} \simeq 5 \times 10^{3335}$ states. Exact-enumeration methods could never cope with this amount of states. Hence, in this section, we employ the kMC sampling to examine the behaviour of the system.

The kMC methods sample the state space by moving the system from one neighbouring state to another, thus requiring a starting state from which the sampling originates. Two

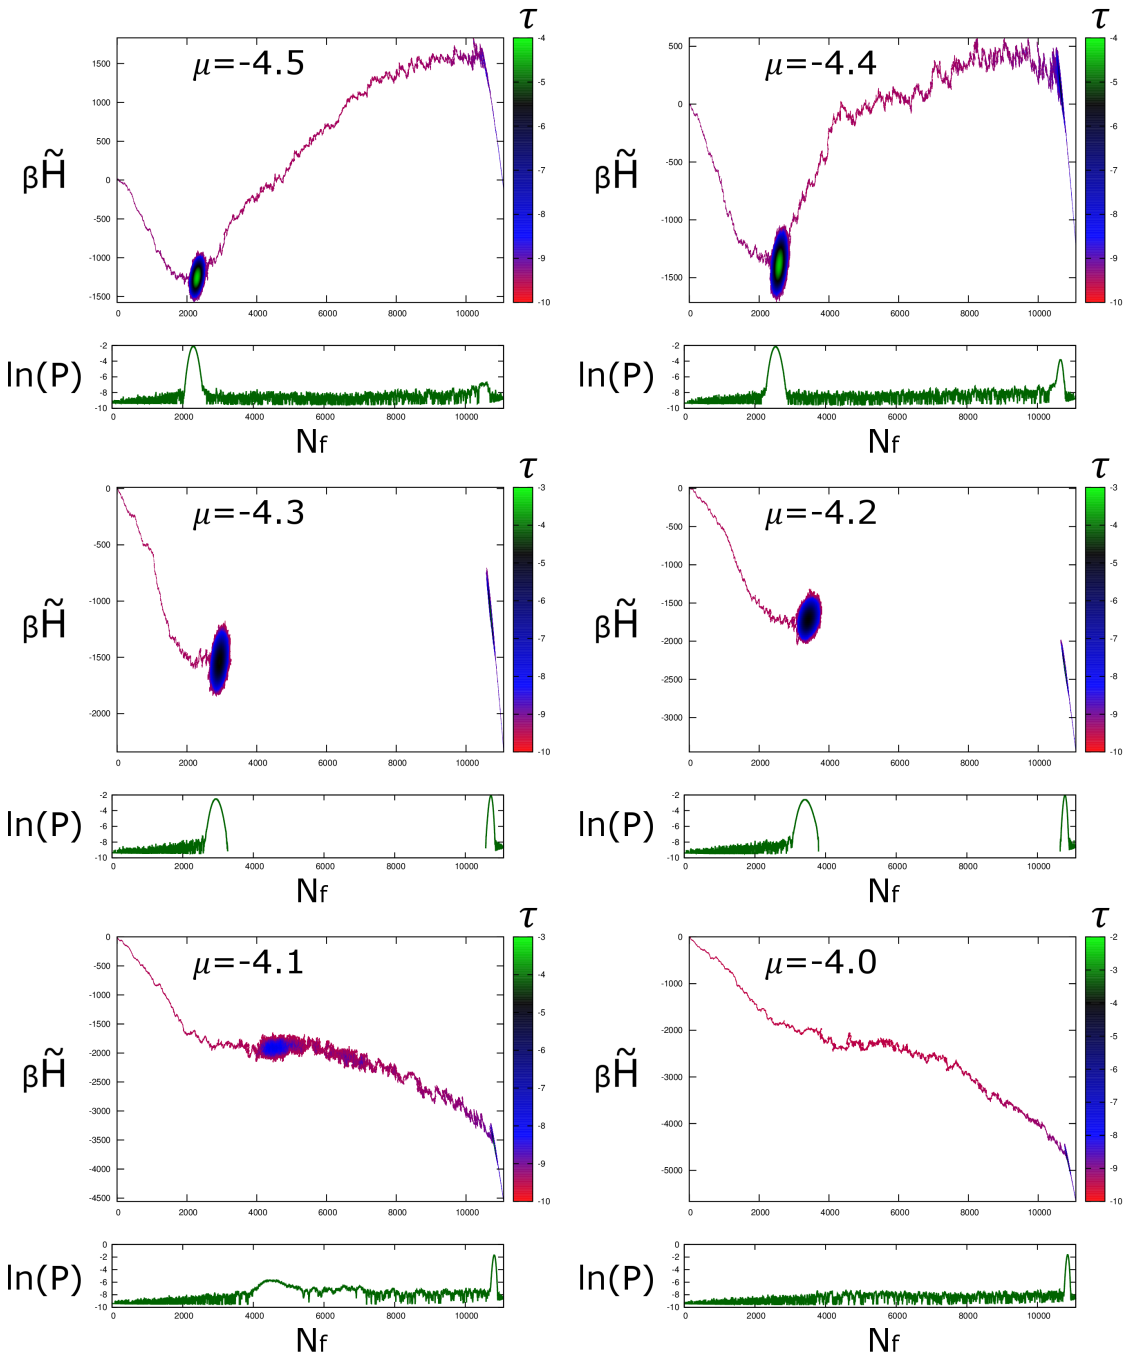

Fig. 2.15 State-space diagrams as explored by kMC method for the cubic bcc lattice system of size $L = 18$ and porosity $\phi = 0.95$ at various values of μ (as indicated in the figure) and $\beta = 1.0$. The lower panels show the $ln(P(N_f))$ distributions obtained by adding up the simulated occupation probabilities of the macrostates at some value of N_f. Each diagram displays the results of two independent simulation runs starting at completely empty and fully fluid initial states.

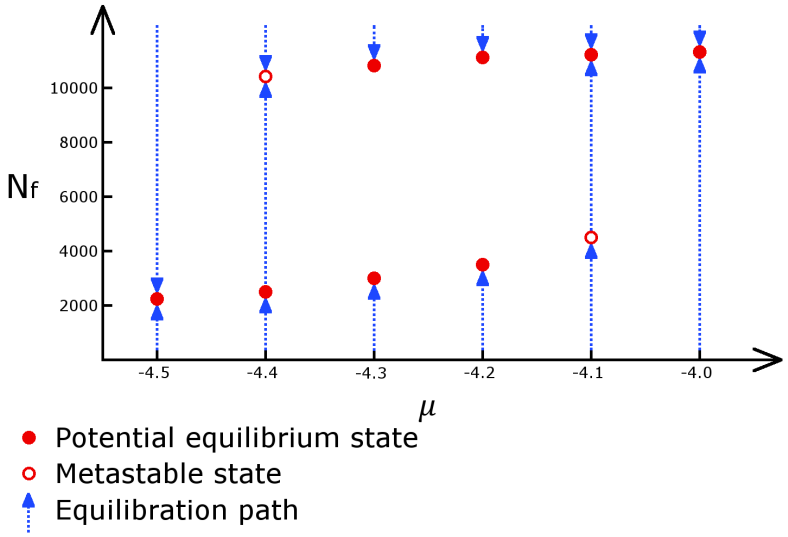

Fig. 2.16 Schematic representation of the results shown in Fig. 2.15. Dashed blue lines represent the equilibration paths through the parameter space that the system has taken to reach either a metastable or a potentially equilibrium state.

of the most natural options are a completely empty system and the one that is fully filled with fluid. We choose to perform a simulation starting from both of those states for each set of relevant values of μ and β. Fig. 2.15 shows the results of such simulations for a set of μ values at $\beta = 1.0$. The x- and y- axes are the same as previously, while the vertical axis in the lower panels is changed to $\ln P$. The colour scheme represents the logarithm of the time fraction, that the system has spent at a particular macrostate, $\tau = \ln(\Delta t(\mathscr{Q})/\mathscr{T})$. Here, $\Delta t(\mathscr{Q})$ is the kMC time that the system has spent at some macrostate $\mathscr{Q}$, while $\mathscr{T}$ is the total kMC simulation time.

Consider the top left panel of Fig. 2.15, with simulation results at $\mu = -4.5$. Both simulations that started at either completely empty or fully fluid systems have reached the same equilibrium state at $N_\mathrm{f} \simeq 2250$. The top right panel shows the same process at $\mu = -4.4$. The end result in this case is similar, i.e. both simulations end up at the same equilibrium state. However, the equilibration of the simulation starting at the fully fluid system has been slowed down by the metastable state at $N_\mathrm{f} \simeq 10500$. As it can be seen from $ln(P)$ plot underneath, the system has spent a substantial amount of time stuck in the metastable state, before escaping it and reaching the equilibrium state. As we have learned in the previous sections, increasing μ value "tilts" the system clockwise, thus making it more difficult for the system to climb from right to left over the potential barrier in between the two favourable states. At $\mu = -4.3$ (the left central panel in Fig. 2.15) the state space "tilts" further to the right and the two simulations do not end up at the same equilibrium state. Since the

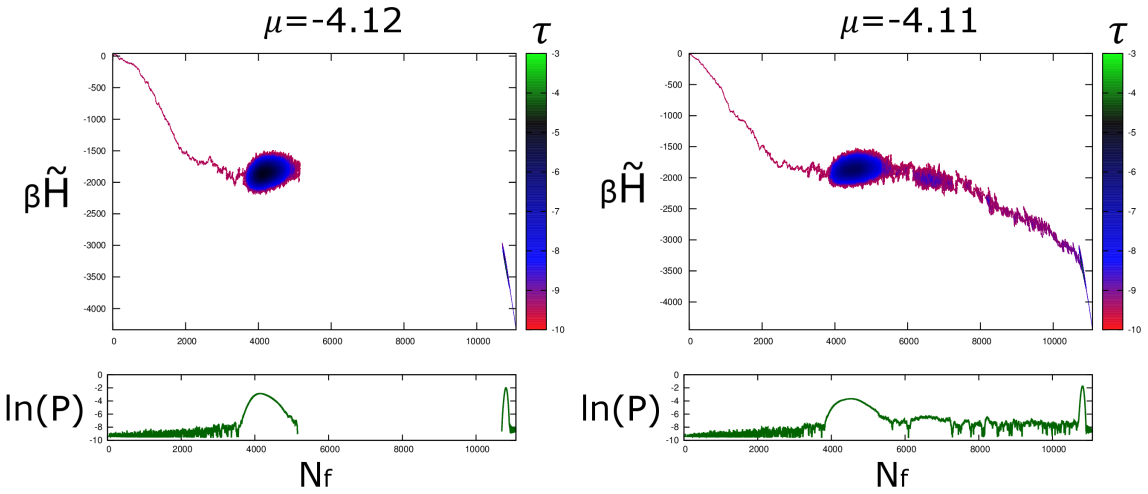

Fig. 2.17 State-space diagrams as explored by kMC method for the cubic bcc-lattice system of size $L = 18$ and porosity $\phi = 0.95$ at $\mu = -4.12, -4.11$ and $\beta = 1.0$. Each diagram consists of two independent simulation runs starting at a completely empty and fully fluid initial states.

system does not reach the same final state when started from distinct initial configurations the simulation is evidently non-ergodic. There is also no trivial way to determine the true equilibrium state, or even to choose which of the metastable states is more stable. From the diagram, we can see that the $\tilde{\mathscr{H}}$ of the microstates at the leftmost metastable state for $\mu = -4.3$ is lower, but there is no direct way to compare the degeneracy in two different subspaces of the state space. A similar behaviour can be observed at $\mu = -4.2$. Continuing to increase μ results in both simulation branches equilibrating to the state at around $N_\mathrm{f} \simeq 10750$, i.e. fully occupied system (see the bottom row in Fig. 2.15).

The more familiar $\mu - N_\mathrm{f}$ diagram summing up the results discussed above (and presented in Fig. 2.15) is shown in Fig. 2.16. It is apparent that connecting the equilibrium states going from left to right and from right to left results in the expected shape of the adsorption and desorption isotherms, respectively. The hysteresis loop at the centre of the diagram shows that the system under this simulation cannot overcome the grand-potential barriers, and the final state for some values of μ and β is path dependent. The problem with the non-ergodic behaviour of the simulation that we have encountered in this example is at the core of most studies performed by MC simulations. We will be faced with the challenges of similar nature all throughout this work, and tackling them is one of our main undertakings.

In order to take a closer look at the behaviour of the system at metastable states let us consider the same example at two very near values of μ, one at which equilibration does not occur, and the one in which it does (see Fig. 2.17). In both cases, the simulation ran for

the same amount of kMC time. At $\mu = -4.11$, the simulations that started at a completely empty system spent a substantial amount of time at a metastable state around $N_\mathrm{f} \simeq 4500$, before eventually equilibrating. The relatively quick transition between the metastable states can be seen as an avalanche, i.e. once the potential barrier is crossed, the system rapidly drifts through the state space towards a more favourable set of states. It is apparent that if we had interrupted the simulation just before it escaped the metastable state, or simply set a shorter simulation time $\mathscr{T}$, we would observe a diagram a lot like the one at $\mu = -4.12$. The same argument applies the other way around in the case of $\mu = -4.12$. In principle, it is possible that, if we set large enough $\mathscr{T}$, the system would equilibrate to the state at $N_\mathrm{f} \simeq 10750$. However, due to the exponential nature of the systems sensitivity to the height of the barriers between the metastable states, even a small difference in barrier height results in a large change in probability that the simulation overcomes it. In principle, the nature of the process is stochastic, and technically there exists some $\mathscr{T}$ for which the system will equilibrate from any starting configuration. However, once the potential barrier reaches a certain height, the time scales in question become not only computationally infeasible, but also too long for natural process, as confirmed by the experimentally observed hysteresis [29]. We seek to examine these and similar questions about the robustness of the results presented here in more detail in the following section.

2.3.3 Equilibration path and robustness of the simulation results

In the previous section, we have introduced the basic behaviour of the kMC simulation in a system representing a cubic sample of aerogel. We have observed that a system starting in either completely empty or fully fluid initial microstate tends to traverse the state space towards some metastable state. At a metastable state, the simulation can either get stuck indefinitely or eventually overcome the grand-potential barrier and experience a quick transition to another metastable state. Here, we would like to further investigate the path that the simulation chooses to traverse through the state space and the robustness of the terminal state that the system resides at the end of the simulation.

Let us further investigate the same system (bcc-lattice representation of a cubic aerogel sample of size $L = 18$ and porosity $\phi = 0.95$), that has been examined in Sec. 2.3.2. Fig. 2.18 shows the simulation results at $\beta = 1.0$ as before and the value of μ is varied. However, in this case six simulation runs are executed and recorded at each set of conditions and, therefore, we can see three equilibration paths traversing the state space instead of only one. We observe that the simulation paths from the initial microstate to a metastable state or between the two metastable states (in the case of two central panels) do not necessarily overlap on the projected state-space diagram, while the locations of the metastable states themselves are

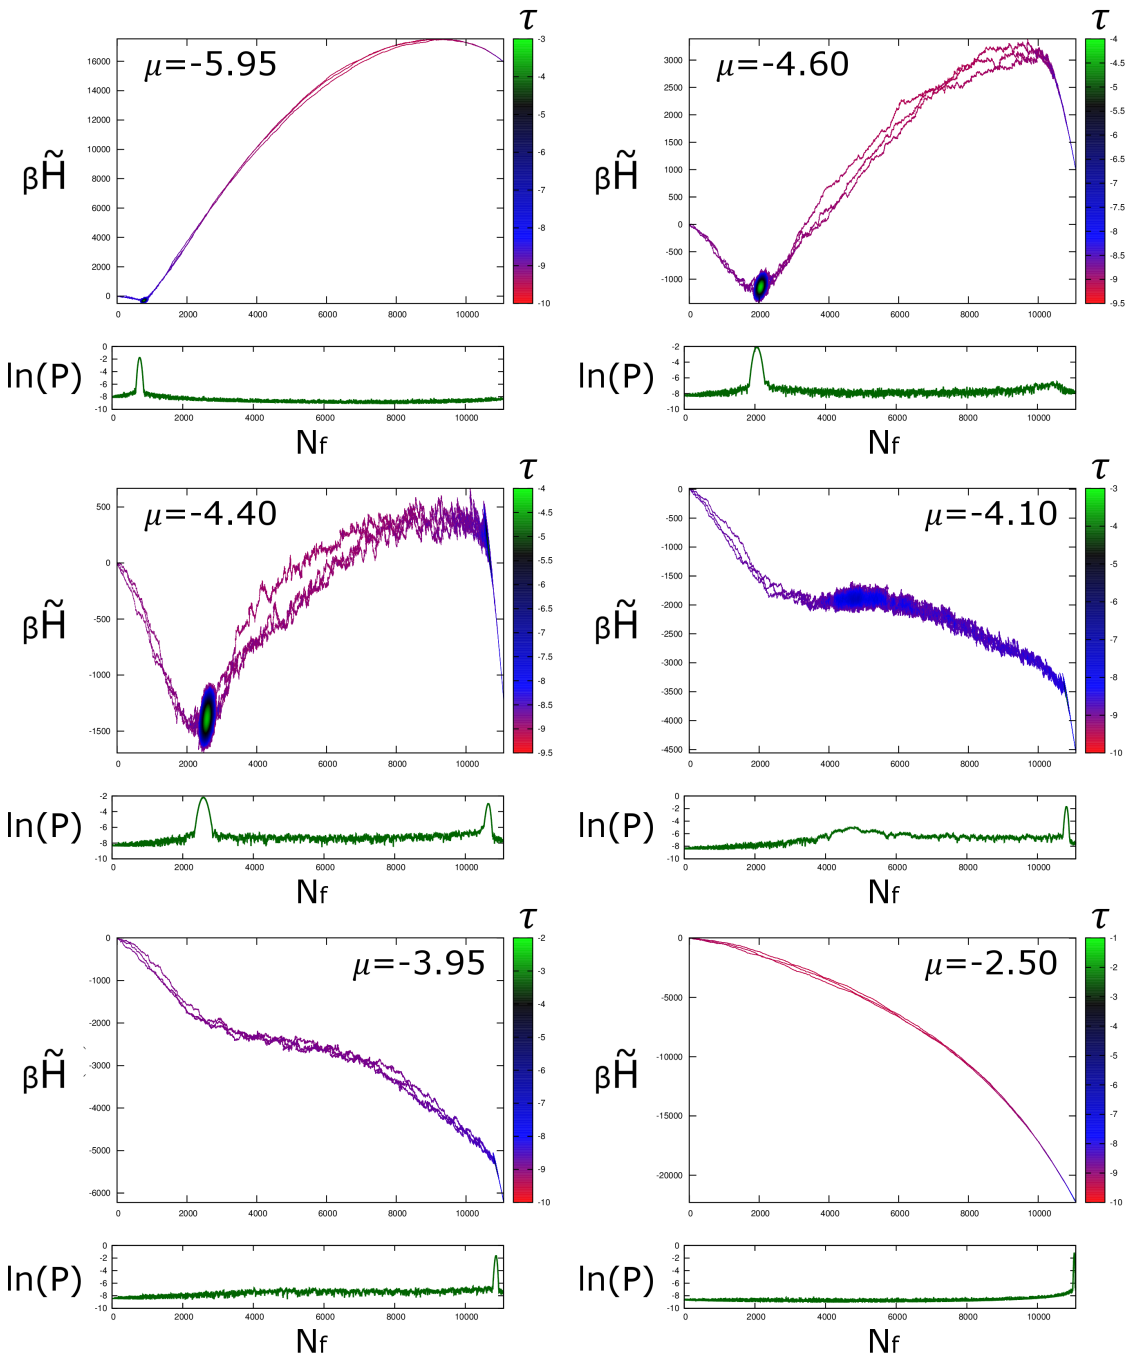

Fig. 2.18 State-space diagrams as explored by kMC method for the cubic bcc-lattice system of size $L = 18$ and porosity $\phi = 0.95$ at various values of μ and $\beta = 1.0$. Each diagram represents the results from six independent simulation runs with three runs starting at completely empty and the other three at fully fluid initial states.

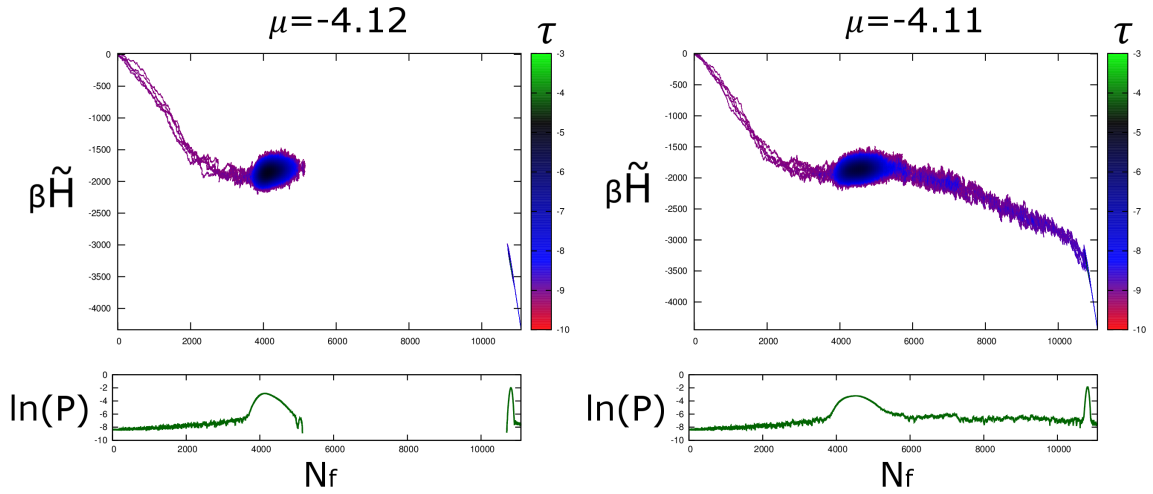

Fig. 2.19 State-space diagrams as explored by kMC method for the cubic bcc lattice system of size $L = 18$ and porosity $\phi = 0.95$ at $\mu = -4.12, -4.11$ and $\beta = 1.0$. Each diagram shows the results from ten independent simulation runs with a half of them starting at completely empty and another half at fully fluid initial states.

robust. This demonstrates the stochastic nature of such transitions, i.e. even though the start and end states are similar, there is no strict order by which the cells get emptied or filled with fluid. Furthermore, the spread of the paths indicates that the grand-potential landscape in the orthogonal direction to the path of an avalanche is relatively flat, allowing for trajectory variation between the different executions of the transition.

As expected, the time that system spends traversing the state space between the metastable states depends on the gradient of the grand-potential landscape that the avalanche is passing through. The $\ln(P)$ graphs in the top left and bottom right panel ($\mu = -5.95$ and $\mu = -2.50$ respectively) of Fig. 2.18 show very fast traversal through the state space as opposed to a relatively slow drift in the two central panels ($\mu = -4.40$ and $\mu = -4.10$, respectively). In the case of $\mu = -4.10$, it can be difficult to distinguish whether the system has been stuck at a shallow metastable state or simply traversed over a relatively flat region of the state space.

To further examine the robustness of the metastable states, we perform a set of repeated simulations for the two values of μ as discussed at the end of Sec. 2.3.2, and were shown in Fig. 2.17. As we can see from Fig. 2.19, at $\mu = -4.12$, none of the simulations managed to cross the potential barrier, while at $\mu = -4.11$ all of them did (it is easier to deduce this directly from the detailed simulation output than from a picture). We have performed multiple sets of such simulations for the variety of conditions and found the same terminal states. Even though it is, in principle, possible to obtain varying results at some very particular value of μ, in general, such behaviour is negligibly rare. Thus, throughout the rest of the project,

we assume that the terminal states of the simulations are robust. In the cases where such assumption is not valid (e.g. high temperatures), we discuss it in the context of the particular conditions.

2.3.4 The effects of temperature variation

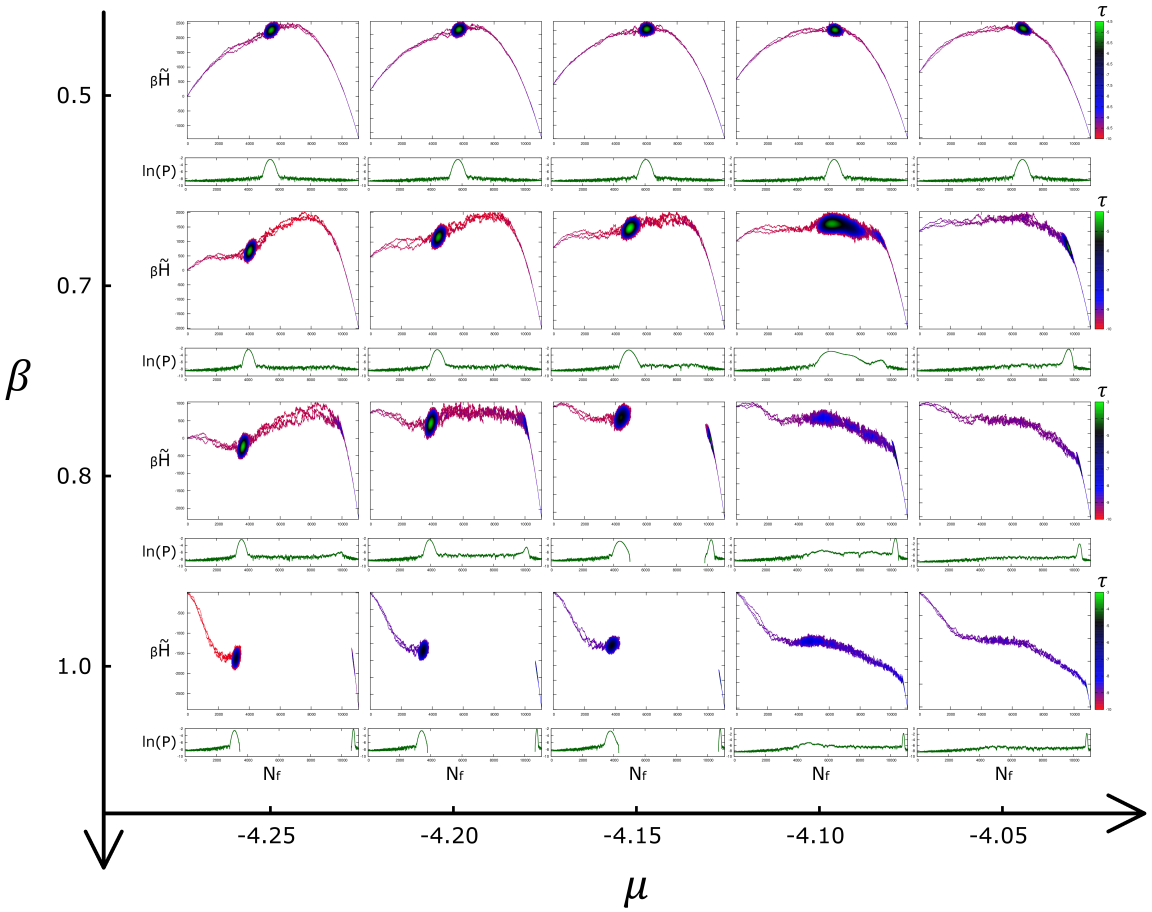

Fig. 2.20 A set of state-space diagrams as explored by kMC method for the cubic bcc-lattice aerogel system of size $L = 18$ and porosity $\phi = 0.95$ at various values of μ and β. Each diagram consists of six independent simulation runs with a half of them starting at completely empty and another half at fully fluid initial states.

In Sec. 2.3.2, we have introduced the basic principles which govern simulations of the aerogel systems. We have studied equilibration of the system at various values of μ starting from either a completely empty or a fully fluid system states. In this section, we explore the behaviour of the system as the temperature is varied.

Fig. 2.20 shows the state-space diagrams for a variety of μ and β values. We observe that at the highest temperature ($\beta = 0.5$), the equilibrium state is reached from both starting

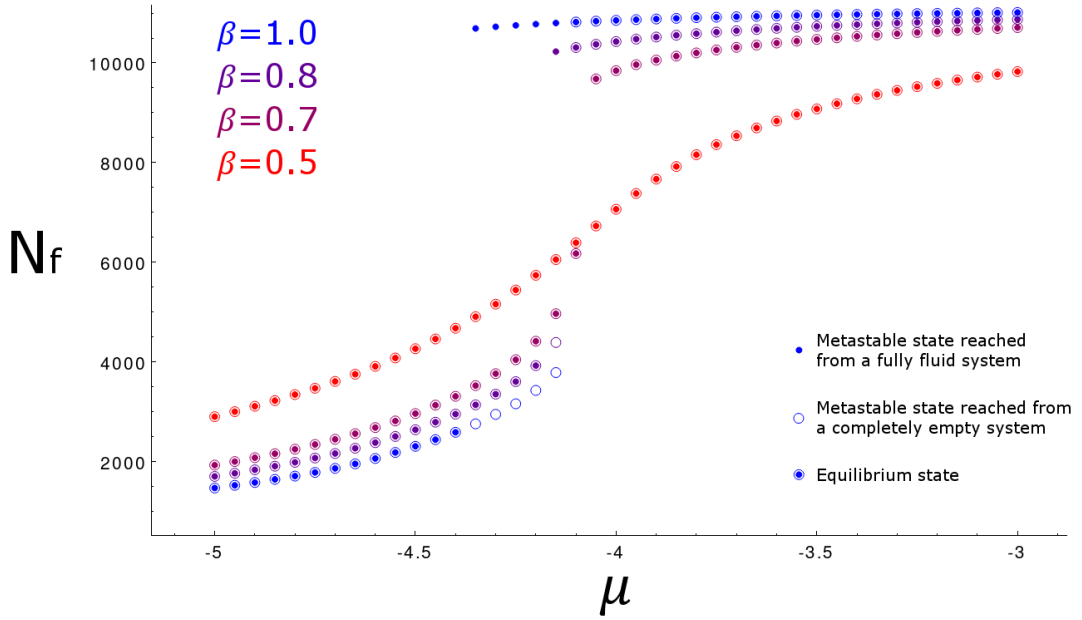

Fig. 2.21 Schematic representation of the results shown in Fig. 2.20 at wider range of μ values. The marked metastable states (see the figure legend) represent the subspace of the state space which were reached from only one of the two starting configurations. At each value of μ where two of such metastable states are shown, one of them is the equilibrium state, however, with the given information, it is impossible to deduce which, since that requires the knowledge of degeneracies as well as the energies of the states.

states. As expected from the findings in Sec. 2.2.3, the equilibrium state at this temperature is in high-degeneracy region of the state space (see the green spot in the top row of panels in Fig. 2.20), and shifts continuously as μ is varied. At $\beta = 0.7$, the system still undergoes a relatively continuous transition from low to high values of N_f, however, the widely explored region of the state space at $\mu = -4.10$ implies the proximity to the critical parameter values. At such values of β and μ, the system undergoes large fluctuations in the order parameter (i.e. in fluid density), which can be directly observed in the state-space diagram (see Fig. 2.20 at $\beta = 0.7$ and $\mu = -4.10$). With further decrease in temperature, for a certain range of μ values (e.g. at $\mu = -4.15$ for $\beta = 0.8$ in the third row of panels and at $\mu = -4.25,\ -4.20$ and -4.15 for $\beta = 1.0$ in the bottom row of panels) the system cannot fully equilibrate from either completely empty or fully fluid starting states. As discussed previously, failure to equilibrate creates sharp phase transitions as the value of μ is varied.

Fig. 2.21 displays the summary of the results in a more common form of $N_f - \mu$ diagram. We can recognise the familiar shape of the adsorption-desorption isotherms formed by the metastable states. The range of μ values at which the system does not fully equilibrate

represents the hysteresis loop discussed in Sec. 1.2.2. The hysteresis loop appears below some characteristic temperature and expands as the temperature is lowered further. Such findings qualitatively match the previous theoretical and experimental results discussed in detail in Sec. 1.2.

2.3.5 The effects of aerogel porosity

In order to study the impact of varying porosity on the sorption phenomena, we have created $\phi = 0.99$ and $\phi = 0.87$ porosity samples of aerogel. As in the case of the previously analysed $\phi = 0.95$ sample, the linear size of the system is $L = 18$, resulting in systems with $N = 11548$ and $N = 10148$ pore cells for $\phi = 0.99$ and $\phi = 0.87$, respectively.

As in the previous Section discussing the results for $\phi = 0.95$ porosity sample, we have performed similar analysis on $\phi = 0.99$ and $\phi = 0.87$ samples. Fig. 2.22 shows the results for the three samples at several temperatures. The trend of opening up and growing hysteresis loop as the temperature is decreased remains for all porosity samples. However, we observe that the formation of the hysteresis loop starts at higher temperatures in $\phi = 0.99$ sample and lower temperatures in $\phi = 0.87$ sample. This phenomenon is explained [29] by the fact that the increased disorder (i.e. a higher fraction of matrix cells) in the material impairs the interaction between the pore cells, thus weakening the cooperative behaviour. This way, the disorder hinders the formation of sharp phase transitions. It is also apparent that the shape of the hysteresis loop changes from a roughly symmetrical rectangle in $\phi = 0.99$ sample, to a triangular in $\phi = 0.87$ sample. Finally, the centre of the hysteresis loop moves towards the lower values of μ as porosity decreases. Since the simulated aerogel has a matrix-fluid interaction strength higher than that of the fluid-fluid interaction (the wettability $y = 2.0$), the increasing proportion of the matrix cells accelerates fluid adsorption into the system, thus shifting the transition towards the lower values of μ.

The qualitative observations made above are consistent with the previous results in the field. Fig. 2.23 taken from the paper by Detcheverry et al. [128] based on local mean-field (i.e. density functional) theory studying sorption phenomenon in silica aerogels. The left (right) panel shows adsorption and desorption isotherms for $\phi = 0.95$ ($\phi = 0.87$) porosity sample, for a varying temperature from lower (in blue) to higher (in black). As in our kMC based analysis, with the decreasing porosity, the hysteresis loop moves towards lower μ values and its shape becomes more triangular. Fig. 2.23 also confirms the trends established in the previous section. At sufficiently low temperatures, the hysteresis loop appears and grows while simultaneously shifting towards lower μ values.

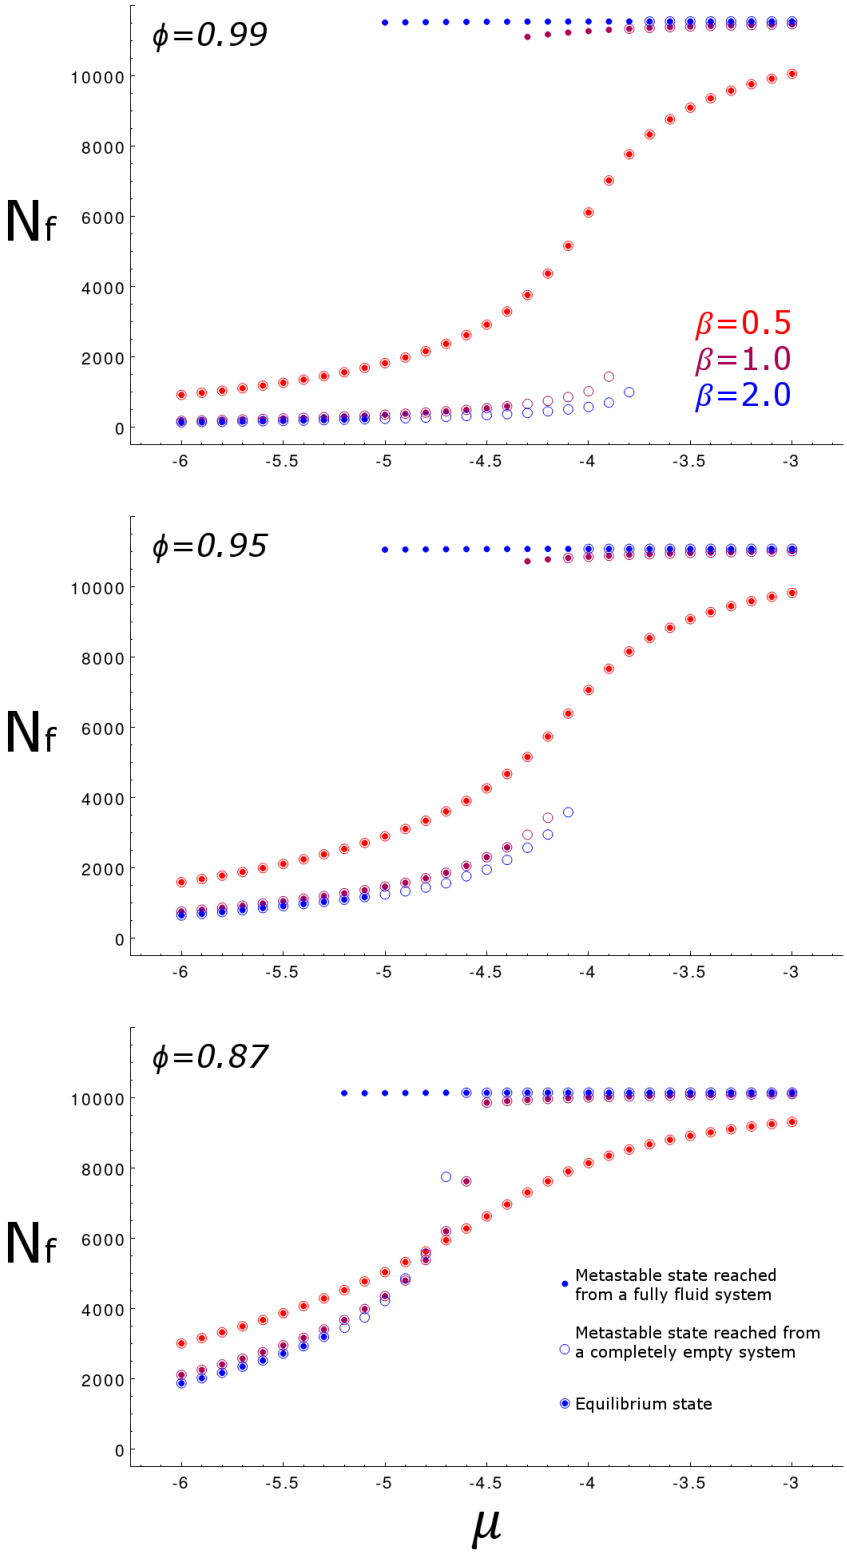

Fig. 2.22 The $N_f - \mu$ diagrams obtained in the same way and the same symbols used as those shown in Fig. 2.21 for three different values of aerogel porosity, that are indicated at the top left corner of each panel.

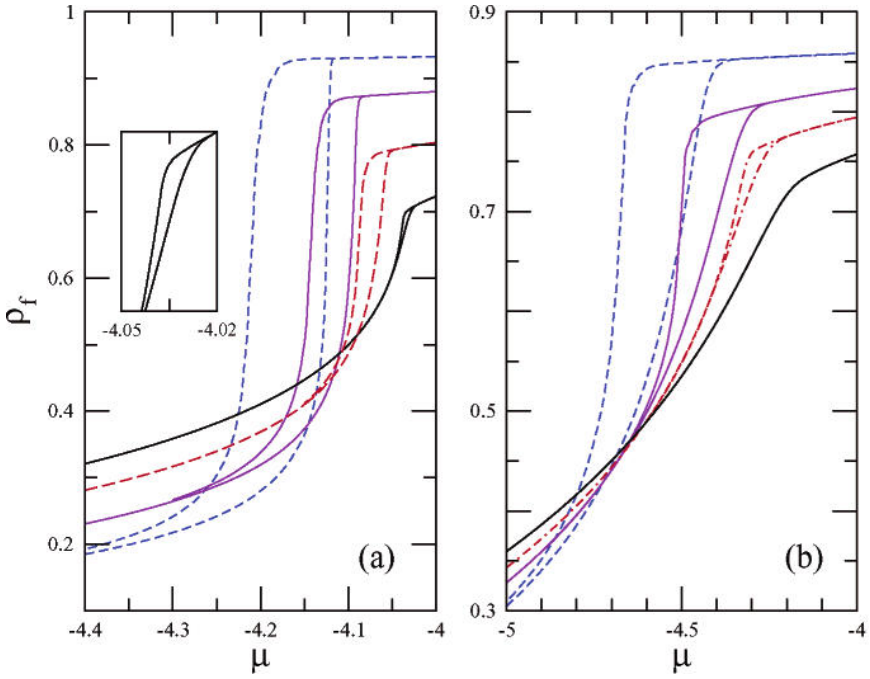

Fig. 2.23 The hysteresis loops for different porosities and temperatures obtained by Detcheverry et.al. [128] by using the local mean-field (i.e., density functional) theory in DLCA generated aerogel samples. The porosity of the samples is $\phi = 0.95$ and $\phi = 0.87$ for (a) and (b) panels respectively. The corresponding temperature is in ascending order from blue ($\beta^{-1} = 1$) to purple ($\beta^{-1} = 1.4$) to red ($\beta^{-1} = 1.7$ in (a) and $\beta^{-1} = 1.6$ in (b)) to black ($\beta^{-1} = 1.9$ in (a) and $\beta^{-1} = 1.8$ in (b)) lines. The inset in the (a) panel enlarges the highest temperature isotherm (in black) near the hysteresis loop.

2.3.6 State space of the aerogel systems

Observations made in the previous two sections encourage us to draw parallels between the impact of varying porosity and that of varying temperature on the sorption phenomenon. As the porosity is increased or the temperature decreased, the disorder caused by spatial inhomogeneities in the sample and temporal thermal fluctuations, respectively, are reduced, thus strengthening the collective behaviour of the system. Heightened collective behaviour leads to large correlated fluctuations that one observes close to criticality, and therefore, to the abrupt phase transitions. The question that naturally arises from the above analysis, is whether the lowering of the temperature can always counteract the decrease in porosity, i.e. is it always possible to observe phase transitions at low enough temperatures, or is there a limiting porosity below which, a system cannot support them. This, and related questions have been continuously studied for over a decade [29], as it was described in detail in Sec. 1.2. While a lot of progress has been achieved for certain materials and regular-geometry pore structures, it proved to be difficult to find definite answers about the sorption behaviour in

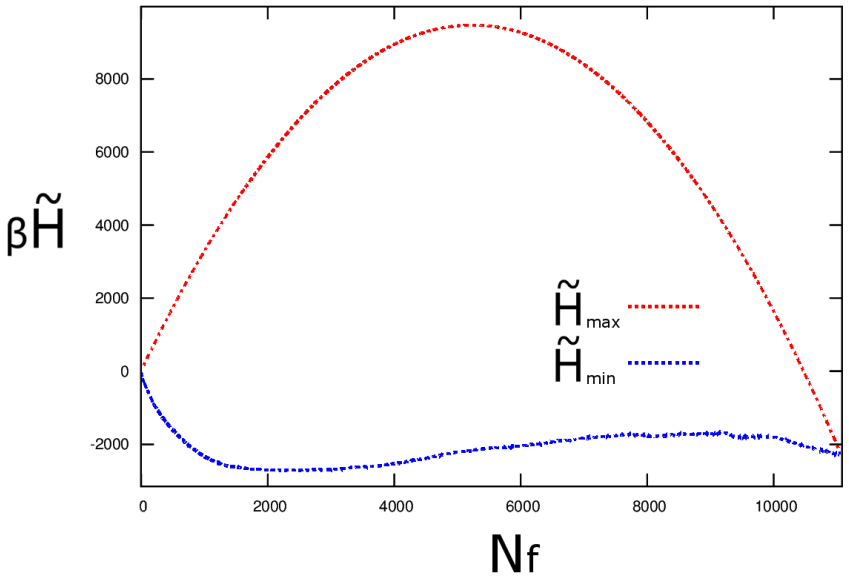

Fig. 2.24 Upper (the red dotted curve) and lower (the blue dotted curve) projected-state-space boundaries for the $\phi = 0.95$ aerogel sample at $\mu = -4.3$.

random natural materials such as silica aerogels and Vycor glass. However, the analysis based on the 2-dimensional projection of the state space, that we have introduced in the beginning of this Chapter, could provide us with very useful insights into the behaviour of the systems.

In Sec. 2.2.4, we have shown that the shape of the state space determines whether the system can undergo a phase transition at some low enough temperature. However, we have seen that in large systems, kMC sampling for some fixed values of μ and β only explores a small part of the state space. In order to gain a better understanding about the entire state space of a particular system, we first seek to find some rough estimates for the upper and lower limit of $\tilde{\mathscr{H}}$ at each value of N_f, and secondly, compile the explored subspaces of the multiple simulation runs onto a single diagram, representing all the states that the system has visited.

At sufficiently high temperatures, $\beta^{-1} \gg 1$, each lattice site of the pore space is equally likely occupied by fluid, i.e. averaged fluid density over time at any cell is simply $\bar{\rho} = N_\mathrm{f}/N$. The Hamiltonian of the system in this regime can be approximated by the mean-field expression, $\mathscr{H}_\mathrm{max}(\bar{\rho})$, thus giving an estimate $\tilde{\mathscr{H}}_\mathrm{max}(N_\mathrm{f})$ for the high-temperature boundary at a fixed value of μ. The obtained function $\tilde{\mathscr{H}}_\mathrm{max}(N_\mathrm{f})$ goes roughly through the mid-points of the projected state space in $\mathscr{H}$ dimension, and corresponds to the location of the states with the highest degeneracy, where the system would reside at an infinitely high temperature. Despite the fact that $\tilde{\mathscr{H}}_\mathrm{max}(N_\mathrm{f})$ marks the vertical mid-point of the projected state space, we

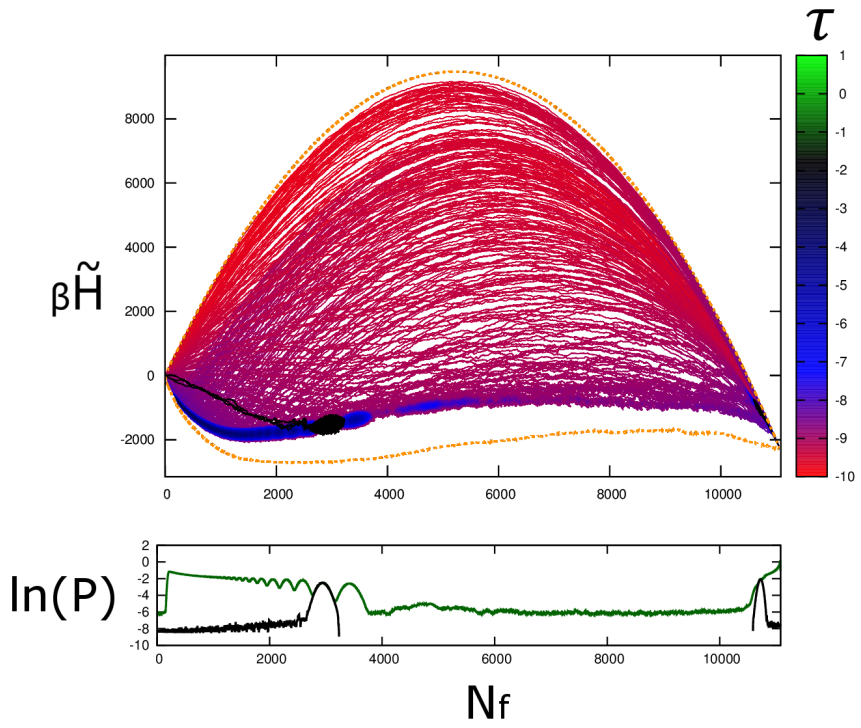

Fig. 2.25 Compiled visited state-space diagram for the $\phi = 0.95$ sample, constructed from the simulations at $\mu = -8.0, -7.9, .. -0.1, 0.0$ and fixed value of $\beta = 1.0$. The diagram is displayed for $\mu = -4.3$, and the results of the simulations at this particular value of μ are drawn in black. Upper and lower state-space boundaries are shown in orange. The lower panel shows $\ln(P)$ for the compiled diagram in green and for a single execution at $\mu = -4.3$ in black.

will call it the upper projected-state-space boundary, because the equilibrium state of the system has to reside below it at any finite temperature. The approximate low-temperature boundary of the projected state space $\tilde{\mathscr{H}}_{\min}(N_f)$ is obtained by a systematic exploration of the low-energy state-space region as described in detail in Sec. 2.3.7. Fig. 2.24 shows the upper and lower projected-state-space boundaries obtained for the $\phi = 0.95$ porosity sample at $\mu = -4.30$.

The projected-state-space boundaries add an important reference point when interpreting the visited macrostate diagram of an individual kMC simulation run at fixed values of β and μ. In order to see the results of a particular simulation in a context of the entire state space, we compiled the visited subspaces of multiple simulation runs for different values of μ and fixed temperature onto a single diagram. Fig. 2.25 shows the compiled visited state-space diagram for the $\phi = 0.95$ porosity sample, when the simulations were ran for a range of values $\mu = -8.0, -7.9, .. -0.1, 0.0$. The diagram is drawn for $\mu = -4.30$, and the states that were visited for this particular value of μ are marked in black, as an example of

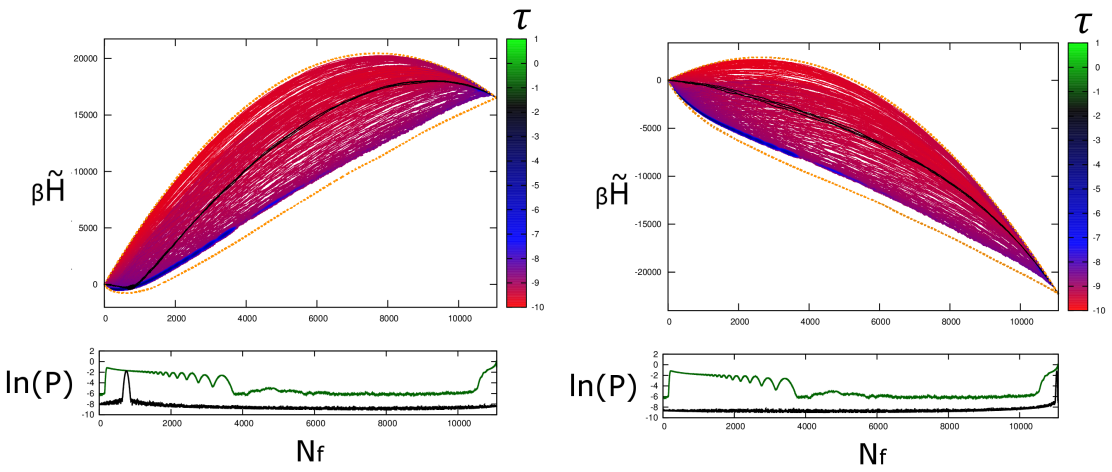

(a) $\mu = -6.0$ (b) $\mu = -2.5$

Fig. 2.26 Compiled state-space diagrams obtained in the same way as those shown in the Fig. 2.25, but drawn for different values of μ. The results of the particular single simulation run at (a) $\mu = -6.0$ and (b) $\mu = -2.5$ (b) are drawn in black.

a single simulation run. The panel below the diagram shows the sum of all the probability distributions in green, and for the particular run at $\mu = -4.30$ in black. The upper and lower energy state boundaries are marked in orange.

The dark blue area in the lower left of Fig. 2.25 (and the probability distribution peaks at $N_f \leq 4000$) denotes the locations of the equilibrium states of the simulation runs with the values of $\mu \leq -4.20$. Similarly, the dark region in the lower right corner of the diagram (and the probability distribution peak at $N_f > 10000$) represents the equilibrium states for the simulation runs at $\mu > -4.20$. The red and purple lines across the centre of the state space are the equilibration paths for the very high or very low values of μ (see $\mu = -5.95$ and $\mu = -2.50$ diagrams in Fig. 2.18). As predicted, the system never crosses upper projected-state-space boundary, while for very low and very high values of μ, the equilibration paths get asymptotically close to it. For a sufficiently tilted state space (see Fig. 2.26), the $\mathscr{H}$ gradient is almost parallel to the equilibration path. In turn, the path is determined by the degeneracy gradient alone, and thus the system drifts towards the highest degeneracy, i.e. the upper state-space boundary (see Fig. 2.26).

Fig. 2.25 demonstrates the path that the equilibrium state takes as μ is varied, as well as the explored subspace of the state space during the equilibration process. The simulations that compose the results presented in Fig. 2.25 are performed at a fixed value of $\beta = 1.0$. In order to see the relative location of the path that the equilibrium state takes at different temperatures, it is convenient to overlap the compiled diagrams constructed from the results of simulations made at several different temperatures. Fig. 2.27a shows two overlapping

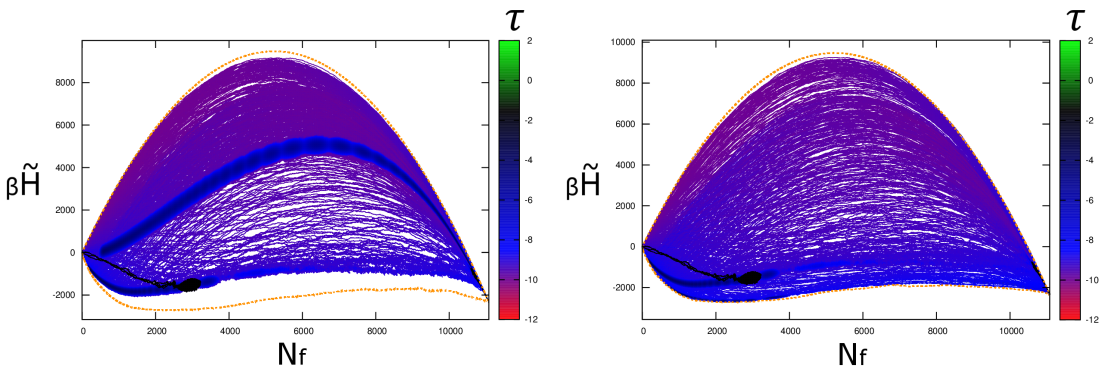

(a) $\beta = 0.5$ and $\beta = 1.0$ (b) $\beta = 1.0$ and $\beta = 2.0$

Fig. 2.27 Compiled state-space diagrams obtained in the same way as those shown in the Fig. 2.25 for two distinct values of β (as indicated under the panels) overlapped on each other. The diagrams are displayed for $\mu = -4.3$.

compiled state-space diagrams at $\beta = 1.0$ and $\beta = 0.5$. The thick blue band across the centre of the state space is the path that the equilibrium state takes at $\beta = 0.5$ for varying values of μ. We can see that it is located in the region of the state space that is significantly higher in $\tilde{\mathscr{H}}$ and contains higher degeneracy macrostates. It is also apparent that the band is continuous across the entire range of N_f values, confirming that the system at this temperature does not experience a sharp phase transition. Fig. 2.27a is a similar diagram constructed from the results of the simulations at $\beta = 1.0$ and $\beta = 2.0$. First of all, we can see that the lower projected-state-space boundary is very closely approached by the simulation results at $\beta = 2.0$. The dark blue regions on the left-hand side of the diagram just above the $\tilde{\mathscr{H}}_\mathrm{min}(N_\mathrm{f})$ boundary signify the equilibrium states at lower μ values for the simulations at $\beta = 2.0$. As expected at a lower temperature, the middle region of the state space, $4000 < N_\mathrm{f} < 11000$, does not contain any equilibrium states (there are no dark blue spots there), indicating that there is a sharp phase transition at this temperature. Further lowering the temperature, however, does not allow us to consistently find new states below the $\tilde{\mathscr{H}}_\mathrm{min}(N_\mathrm{f})$ boundary, confirming that the estimation procedure obtains a sufficiently accurate lower projected-state-space boundary for such samples.

2.3.7 Algorithm to estimate the low-temperature state-space boundary

The approximate lower projected-state-space boundary of the state space $\tilde{\mathscr{H}}_\mathrm{min}(\rho)$ is obtained through a systematic exploration of the state space by random variation of β and μ values. The algorithm records the lowest found value of $\mathscr{H}$ for each $\rho \in (0,1)$ according to the following algorithm.

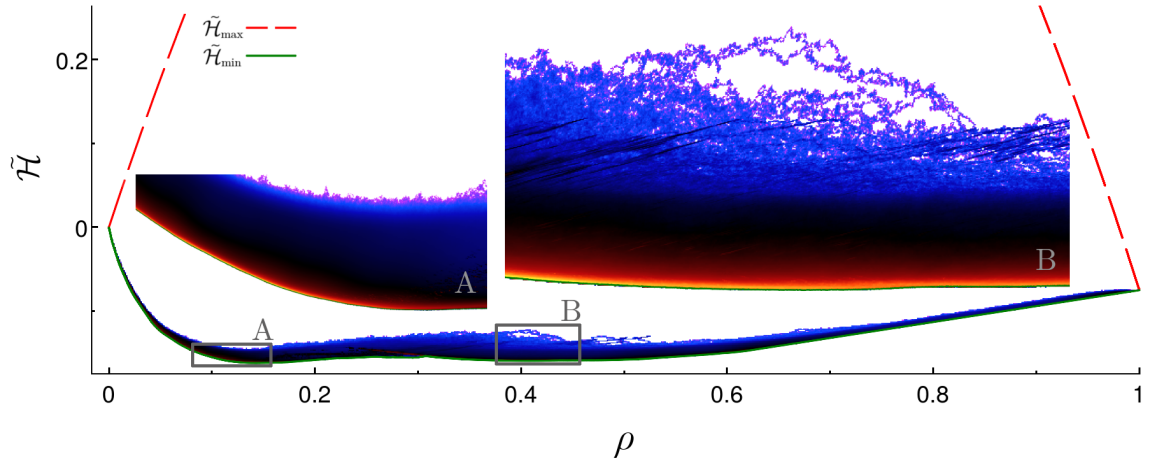

Fig. 2.28 The macrostate visit histogram for a typical execution of the procedure that obtains the lower projected-state-space boundary of the state space. Simulation parameters as follows: $\mu_{\min} = -10$, $\mu_{\max} = 0$, $\Delta\mu_{\text{step}} = 0.001$, $\beta_{\min} = 2$, $\beta_{\max} = 20$, n in the range $[10^4, 10^5]$ depending on temperature, and $m = 20$. The system studied in this figure is a $\phi = 0.95$ porosity aerogel sample on a bcc lattice of size $18 \times 18 \times 26$. The insets show magnification of two regions (A and B) near the low-temperature boundary.

(i) Run the standard single-flip MC simulation and record the energy $\mathscr{H}(\rho)$ of lowest energy state for each visited $\rho \in (0,1)$.

(ii) For every n-th MC step, randomly choose a value of β_{new} from the uniform distribution, $\beta_{\text{new}} \sim U(\beta_{\min}, \beta_{\max})$, and increment in chemical potential $\Delta\mu$ as either $\Delta\mu_{\text{step}}$ or $-\Delta\mu_{\text{step}}$, each with probability $p = 1/2$.

(iii) Set $\beta = \beta_{\text{new}}$ and $\mu = \mu_{\text{new}}$. Here, $\mu_{\text{new}} = \mu_{\max}$ if $\mu + \Delta\mu > \mu_{\max}$, $\mu_{\text{new}} = \mu_{\min}$ if $\mu + \Delta\mu < \mu_{\min}$ or $\mu_{\text{new}} = \mu + \Delta\mu$, otherwise.

(iv) Repeat (i)-(iii) until no new lowest value of $\mathscr{H}(\rho)$ at any ρ is found for m system (performing random walk in μ) traversals from $\mu_{\min}$ to $\mu_{\max}$ and back.

(v) Compute $\tilde{\mathscr{H}}_{\min}(\rho) = \mathscr{H}_{\min}(\rho) - \mu\rho$ for any μ of interest.

The system under this algorithm performs a bounded random walk in μ dimension. The bounds $\mu_{\min}$ and $\mu_{\max}$ are chosen such that the pore-space lattice sites of the system are either all empty or occupied by fluid at $\mu_{\min}$ and $\mu_{\max}$, respectively, ensuring that the entire range of $\rho \in [0,1]$ is explored. The step size, $\Delta\mu_{\text{step}}$, is chosen as small as possible, provided that the system still has enough time to diffuse from $\mu_{\min}$ to $\mu_{\max}$ and back multiple times throughout the simulation. The random variations in temperature provide an effective way to explore (search for and escape from) metastable states. The value of $\beta_{\min}^{-1}$ is chosen high

enough, so that the system is insensitive to the local energy landscape ruggedness at such high temperature and explored states for any $\mu \in [\mu_{\min}, \mu_{\max}]$ span the wide range of ρ and $\tilde{\mathscr{H}}$ values. The low-temperature limit, $\beta_{\max}^{-1}$, the frequency, n^{-1}, of μ and β updates and the number of traversals m, are selected after the empirical exploration of algorithm's performance, ensuring that the further increase in $\beta_{\max}$, n and m does not improve the resultant values of $\mathscr{H}_{\min}(\rho)$. This procedure for finding $\mathscr{H}_{\min}(\rho)$ needs to be performed only once for a given aerogel model structure, since, for any particular value of μ, the estimate for $\tilde{\mathscr{H}}_{\min}(\rho)$ can be found by means of linear transformation given in step (v) of the algorithm.

The algorithm was tested for a wide range of all parameter values in order to demonstrate that the lower projected-state-space boundary, or equivalently, the low-temperature boundary, remains relatively unchanged if sufficiently high values of m are used. Fig. 2.28 illustrates macrostate visit histogram after a typical execution of the algorithm for 0.95-porosity aerogel sample containing $18 \times 18 \times 26$ cubic unit cells. The most to least visited macrostates are coloured in yellow to blue, on the white background of the unexplored state space. It should be emphasised that $\tilde{\mathscr{H}}_{\min}(\rho)$ is not the strict state-space boundary, since in order to guarantee that no states exist with $\tilde{\mathscr{H}}(\rho) < \tilde{\mathscr{H}}_{\min}(\rho)$, an exhaustive search of the state space would be required, which is computationally infeasible for large systems.

Computation of the lowest-energy state, also known as ground state, is a widely discussed problem in computational physics for a variety of systems e.g. atomic clusters [158], proteins [13, 14], spin glasses [159], etc. In many of such systems, e.g. 3d spin glasses, the problem is NP-complete [159] and therefore MC-based algorithms such as simulated annealing are used to tackle the problem. However, in the case of the random-field Ising model a polynomial-time algorithm has been developed, that obtains the exact ground state of the system. The algorithm is based on the mapping to the maximum flow problem [160, 161]. Further schemes based on this result have been developed to obtain the ground state configuration for the entire range of external field values [162], thus producing the zero-temperature isotherm for the given system. Due to the mapping of the lattice-gas model to the random-field Ising model with spatial correlations in random fields [130], it might be possible to apply such techniques to obtain the low-temperature state-space boundary discussed in this section.

2.3.8 Aerogel-porosity impact on the state space

Let us now come back to consider different porosity samples, employing the knowledge and the tools that we have developed in the previous section.

Fig. 2.29 shows the compiled state-space diagrams for the three different porosity samples that we have seen before ($\phi = 0.99,\ 0.95,\ 0.87$). The left-column diagrams are constructed

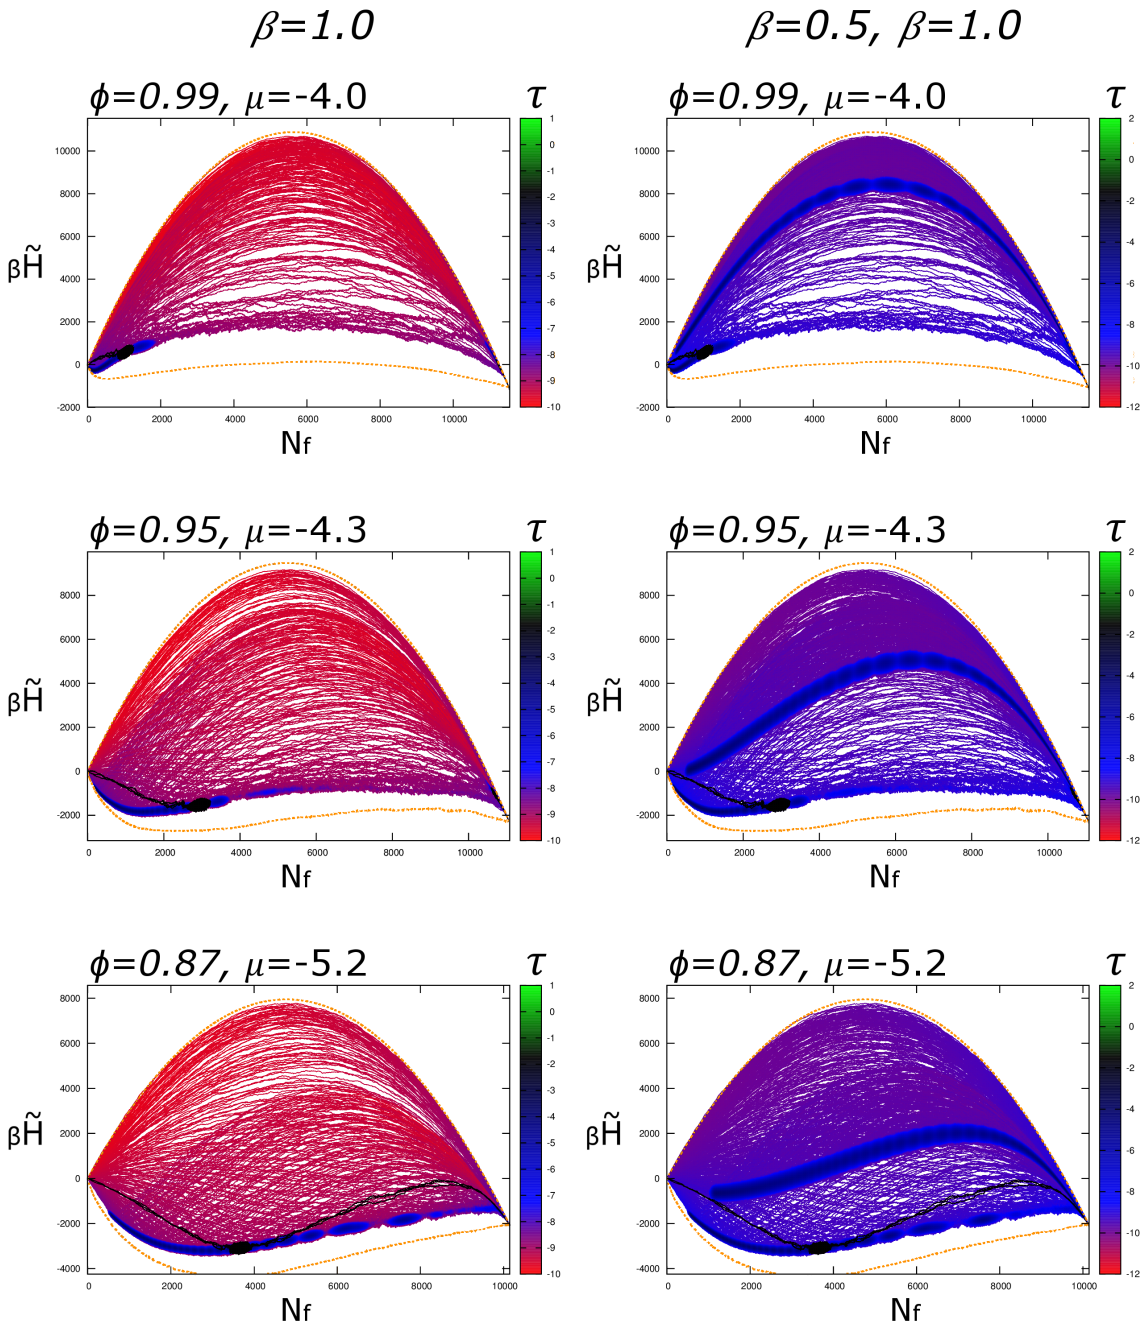

Fig. 2.29 The three panels on the left display the compiled state-space diagrams obtained in the same way as those shown in the Fig. 2.25. The panels on the right display the compiled state-space diagrams for $\beta = 0.5$ and $\beta = 1.0$ overlapped on each other (as seen previously in Fig. 2.27). The corresponding porosity as well as the value of μ, for which each diagram is displayed, are indicated at the top of each panel.

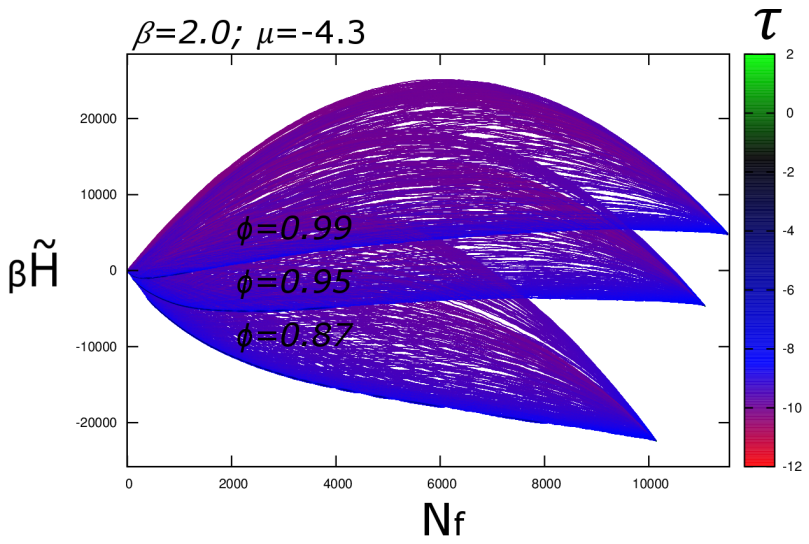

Fig. 2.30 Compiled state-space diagrams obtained in the same way as those shown in the Fig. 2.25, for three different porosity samples. The state-space projection is drawn for $\mu = -4.3$.

using the results of simulations at $\beta = 1.0$. We can see that the shape of the state space, as the porosity is varied, mimics the behaviour of the small systems (see Fig. 2.9 for high porosity, Fig. 2.12 for medium porosity and Fig. 2.7 for low porosity). The higher-porosity system ($\phi = 0.99$) exhibits two deep minima, and the absence of the states at the lower energies for the range of intermediate values of N_f. As discussed, this causes the occurrence of the abrupt phase transitions at sufficiently low temperature. Meanwhile, the lower porosity ($\phi = 0.87$) sample has positively curved lower state-space boundary, which ensures that there always exist some states for each value of N_f, allowing for a continuous equilibrium state path across the state space as μ is varied. Such behaviour is also confirmed by the blue band in the lower part of the $\phi = 0.87$ sample diagram, indicating the location of the equilibrium state at various values of μ (there is a continuous set of dark blue spots along the low boundary of explored state space). The right-column panels in Fig. 2.29 display the compiled state-space diagrams of simulation runs at $\beta = 0.5$ and $\beta = 1.0$ overlapped on top of each other. The thick blue band in the middle of the state space indicate that for all three considered porosities the equilibrium density continuously changes with μ at $\beta = 0.5$.

In addition to affecting the shape of the state space, porosity also influences its location and orientation. As we have seen, in order for the state-space diagrams to be "horizontal", a particular value of μ has to be chosen. We say that the state space is horizontal, if the relatively empty (the state with a thin fluid layer surrounding the matrix cells) and fully fluid states are at a similar value of $\tilde{\mathscr{H}}$. Horizontal position roughly indicates the region of μ at

which most of the equilibrium transition occurs, and when the two peaks in the $ln(P(N_f))$ distribution can be observed at low temperatures. Thus, it is important to understand how the orientation of the state space, and thus μ at which it is horizontal, depends on the porosity of the sample. Fig. 2.30 displays all compiled state-space diagrams of the three considered porosity samples. As we have previously seen, at $\beta = 2.0$, the compiled state-space diagram explores almost identical region of the state space as that enclosed by the upper and lower projected-state-space boundaries. Therefore, in order to keep the figure simple, the boundaries themselves here are not displayed. The figure is produced at $\mu = -4.3$, which maintains the state-space diagram for the $\phi = 0.95$ porosity sample horizontal. In turn, the state spaces for the $\phi = 0.99$ and $\phi = 0.87$ porosity samples appear tilted anticlockwise and clockwise respectively. The relative orientation of the state spaces also explains why at this particular value of μ the $\phi = 0.99$ porosity sample is close to empty, while the $\phi = 0.87$ porosity sample is almost completely filled with fluid. As it was discussed in Sec. 2.2.3, the orientation of the lower state-space boundary at some value of N_f depends on the maximum interaction energy that can be added with one additional fluid cell. The fact that the $\phi = 0.95$ porosity sample has a roughly horizontal lower state-space boundary for $N_f > 2000$ at $\mu = -4.3$ indicates that after forming the initial fluid layer around the matrix cells ($N_f < 2000$), on average a single empty cell turning to fluid adds -4.3 units of interaction energy. In comparison, an empty system on average always adds $-q/2$ units of interaction energy, where q is the coordination number of the underlying lattice, i.e. $q = 8$ for bcc lattice.

So far we have established that the curvature of the lower state-space boundary determines the existence of a sharp phase transition in the system, that the chemical potential of the transition depends on the orientation of the state space itself, and how all of it is related to the microstructure of the sample. However, in order to predict the behaviour of the system at intermediate temperatures we ought to develop a better understanding about the degeneracy of the macrostates and how it varies across the state space.

2.3.9 Equidegenerate lines

In this Chapter, so far we have studied how the shape of the state space determines the set of all possible ways that the system can behave, the paths that the equilibrium state takes across the state space as μ is varied, and how it all depends on the microstructure of the sample. We have obtained a detailed description of the behaviour in the high- and low-temperature limits. However, the subset of the state space that is explored at equilibrium for the intermediate temperatures depends on the states degeneracy as much as it does on their energy ($\tilde{\mathscr{H}}$). Given a particular microstate, finding its $\tilde{\mathscr{H}}$ is a matter of algebra, while obtaining the degeneracy is a significantly more difficult task. Without being able to count all the microstates of the

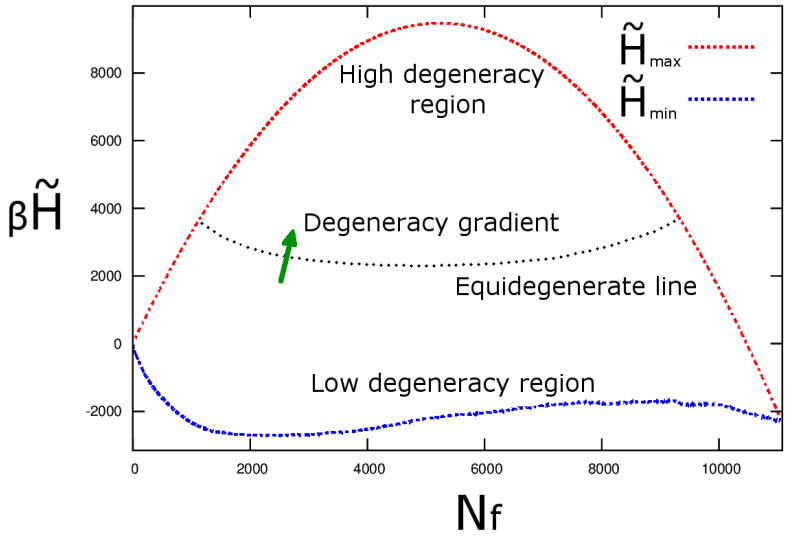

Fig. 2.31 Schematic projected-state-space diagram with the upper and lower projected-state-space boundaries for the $\phi = 0.95$ sample of aerogel, shown for $\mu = -4.3$. The dashed black line represents an equidegenerate line which is orthogonal to the degeneracy gradient across the projected state space.

system and thus having to rely on the sampling methods, there is no trivial way to obtain absolute value of the degeneracy at some point in the state space. However, it is possible to compare the relative degeneracies of two nearby regions of the state space, since the sampling frequency is a function of the density of states at each region (subspace). In this section, we develop a methodology which is aimed to obtain the equidegenerate lines in the state space, i.e. the lines across the projected state space along which the degeneracy of the macrostates is constant.

Fig. 2.31 is the schematic representation of the projected state space for the already analysed $\phi = 0.95$ sample of aerogel. We know that the highest state density region is in the middle of the projected state space (i.e. at the top of the displayed lower half of the projected state space), while the lowest state density is at around the perimeter. To explore the variation of the degeneracy in between we use the concept of equidegenerate lines. In order to obtain them we alter the standard kMC sampling algorithm in the following ways:

- the standard Boltzmann sampling is replaced by an infinite-temperature or simply degeneracy-based sampling. As before, the next state that the system will go to is chosen out of the set of neighbouring states, however, instead of the usual Boltzmann probability distribution, each state is assigned an equal weight, irrespective of $\tilde{\mathscr{H}}$.

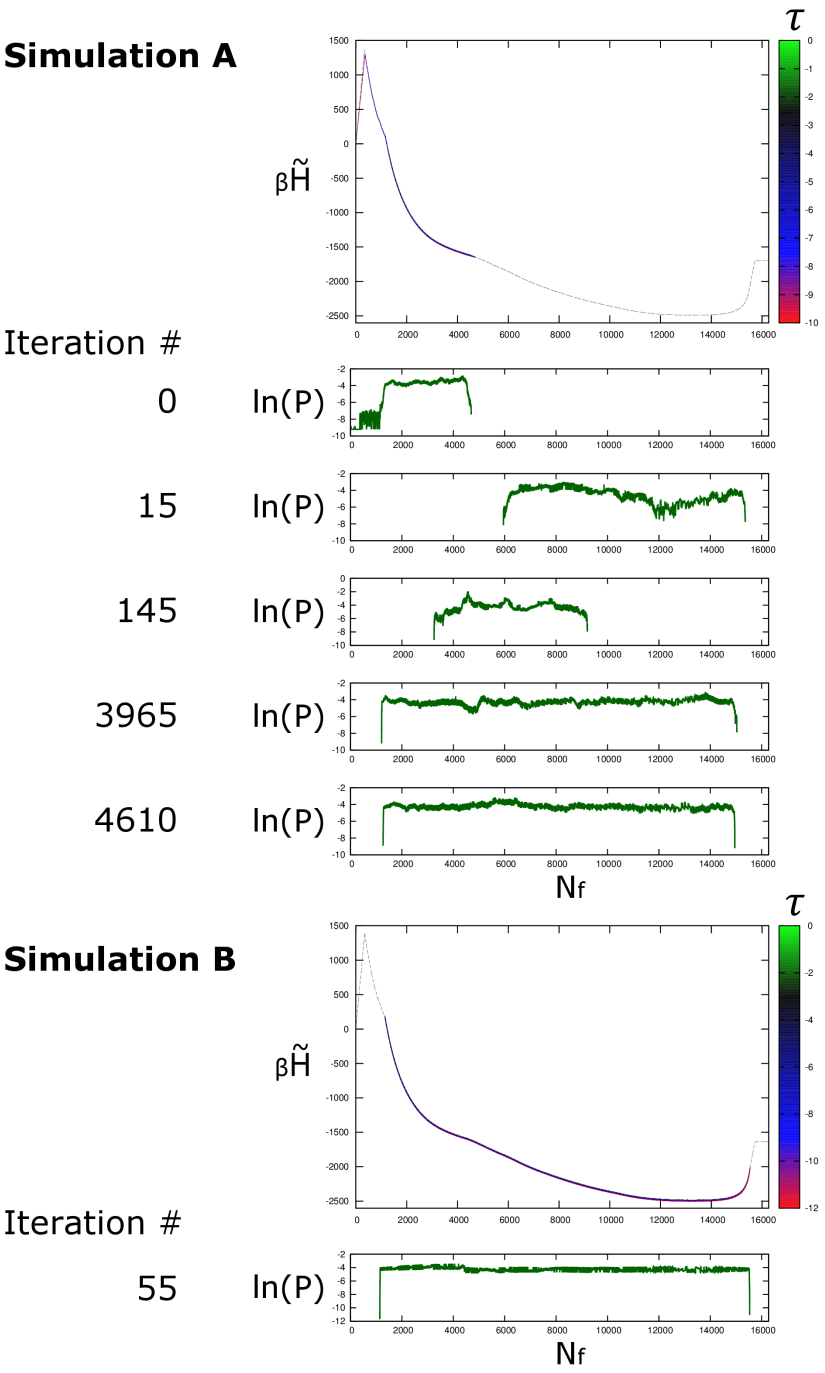

Fig. 2.32 The intermediate results for the iterative procedure that generates the equidegenerate lines. Simulation A is the initial part of the procedure with a relatively short equilibration time designed to converge quickly to the approximate results, while Simulation B is the final part of the process with a longer equilibration time designed to fine-tune the equidegenerate lines. The entire projected-state-space diagram is shown for the *Iteration 0* of the Simulation A, and *Iteration 55* of the Simulation B. The rest of the iterations are only shown in $ln(P)$ diagrams. The system studied in this figure is a $\phi = 0.95$ porosity aerogel sample on a bcc lattice of size $18 \times 18 \times 26$. The projected-state-space diagrams are displayed for $\mu = -4.0$.

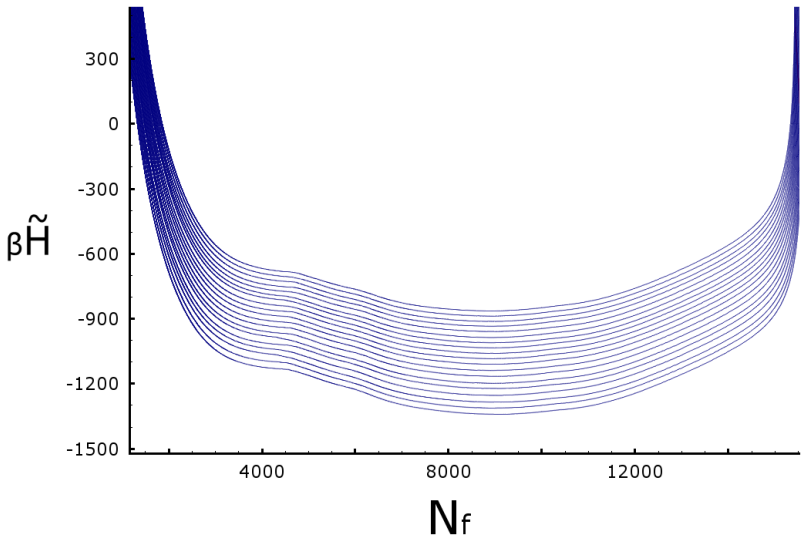

Fig. 2.33 Equidegenerate lines obtained by the iterative procedure described in Fig. 2.32. The projected state space for $\phi = 0.95$ porosity sample is displayed at $\mu = -4.1$.

- If the sampling protocol described above is set up to sample the state space unrestricted, it would simply drift towards the maximum degeneracy region and remain there indefinitely. Therefore, we implement a *simulation limit*, i.e. an arbitrary boundary across the state space which acts as an upper simulation limit, preventing the sampling algorithm to see any states that are above it. Otherwise, the simulation behaves in exactly the same fashion as before.

By construction, the gradient of the degeneracy across the state space is orthogonal to the equidegenerate lines. Thus, if the simulation limit happens to coincide with the equidegenerate line, the algorithm would sample a thin layer of the states just below the sampling limit with a uniform probability distribution. However, if the simulation limit allows the sampling of higher degeneracy region at some part of the state space than at others, i.e. simulation limit is not parallel to some equidegenerate line, the sampled subspace drifts toward the highest degeneracy region that is permitted by the simulation limit and resides there. Thus, we have developed an iterative procedure, which is aimed to obtain equidegenerate lines by incrementally altering the simulation limit until the space just below it is sampled approximately homogeneously. The procedure can be summarised as follows:

- an arbitrary simulation limit is chosen roughly parallel to the lower projected-state-space boundary.
- Infinite-temperature sampling is performed for a certain kMC time period.
- $P(N_\mathrm{f})$ distribution is obtained.

- The simulation limit is altered at every N_f by either moving it up or down if $P(N_\mathrm{f})$ is either lower or higher than the mean value of $P(N_\mathrm{f})$, respectively.
- The procedure is repeated until the resulting $P(N_\mathrm{f})$ is uniform within the expected statistical error.

Fig. 2.32 shows the results of multiple iterations. It is evident that during the initial iterations of Simulation A the simulation limit is not parallel to an equidegenerate line. Meanwhile, at later iterations almost the entire range of possible N_f values is explored, and the algorithm simply performs a random walk just below the entire simulation limit. Once a satisfactory simulation limit is obtained, it is possible to increase the accuracy of the results by performing several additional iterations using longer kMC waiting times. Simulation B panel in Fig. 2.32 thus demonstrates the resulting state-space diagram and $P(N_\mathrm{f})$ distribution. It is evident that $P(N_\mathrm{f})$ is roughly uniform, granting that the current simulation limit is parallel to some equidegenerate line. We repeat the procedure after slightly lifting the simulation limit across the entire range of N_f. However, the number of iterations that is required to obtain further equidegenerate lines is significantly reduced, since the algorithm starts with a fairly accurate initial guess. An example of a resulting set of equidegenerate lines is shown in Fig. 2.33.

It is worth mentioning that the procedure, that the algorithm for obtaining the equidegenerate lines is based upon, is highly sensitive. As we see from Fig. 2.33, the relevant range of $\tilde{\mathscr{H}}$ is of the order 10^3. Meanwhile, if the simulation limit differs from the equidegenerate line by $\Delta\tilde{\mathscr{H}} \simeq 2$ for some value of $N_\mathrm{f} = N_\mathrm{f}^*$, the resulting $P(N_\mathrm{f})$ distribution would have over an order of magnitude sized peak (or dip) at N_f^*. Thus, in turn ensuring that in the next iteration such deviation is fixed, and that the resultant equidegenerate lines are highly accurate.

2.3.10 Degeneracy map

To complete our study of the state space we would like to obtain a full map of the degeneracy of the macrostates, i.e. the state density at any point in the projected state space. In this section, we demonstrate another method aimed to obtain a global picture of the degeneracy from the local observations of its gradient, that the kMC simulations can provide.

The procedure that we have designed to find a local gradient of the degeneracy at some point $(N_\mathrm{f}^\dagger, \tilde{\mathscr{H}}^\dagger)$ in the projected state space can be summarized as follows:

- record a microstate along some equidegenerate line at $N_\mathrm{f} = N_\mathrm{f}^\dagger$.

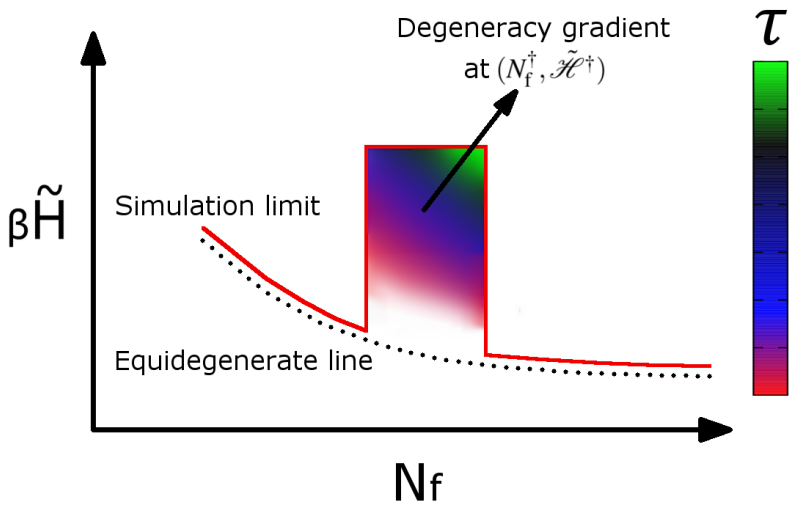

Fig. 2.34 Schematic representation of the procedure designed to obtain the degeneracy gradient at a particular point of the projected state space. The precise description of the procedure is given in the text.

- Alter the simulation limit in the way that is shown in Fig. 2.34, thus creating a semi-rectangular region surrounded from three sides ("a column"). Choose the height of the altered column to be just above $\tilde{\mathscr{H}} = \tilde{\mathscr{H}}^\dagger$, and the width of the column depending on the desired precision. Note that at low energy a too narrow width of the column could prevent the system to explore the entire "depth" of the state space inside the semi-rectangular region, thus breaking the ergodicity even within the allowed subspace of the state space, and spoiling the results. However, such simulation regime is likely to be accompanied by locally irregular and discontinuous results, and therefore can be easily detected.

- Run the infinite-temperature simulation starting from the recorded microstate. Since, on average, the explored states drift towards higher degeneracy region, the simulation will be trapped in the column.

- Once equilibrated, find the mean horizontal and vertical gradients of the kMC time spent at the macrostates surrounded by the simulation limit. If the process was ergodic, the result is also equal to the degeneracy gradient around $(N_f^\dagger, \tilde{\mathscr{H}}^\dagger)$.

Such procedure can be repeated for all desired coordinates across the projected state space. Since the state-space region of the simulation is highly limited in size, the equilibration times are relatively low, thus allowing us to obtain a dense map of degeneracy gradients for the entire projected state space. Fig. 2.35 shows an example of such a map for the $\phi = 0.95$

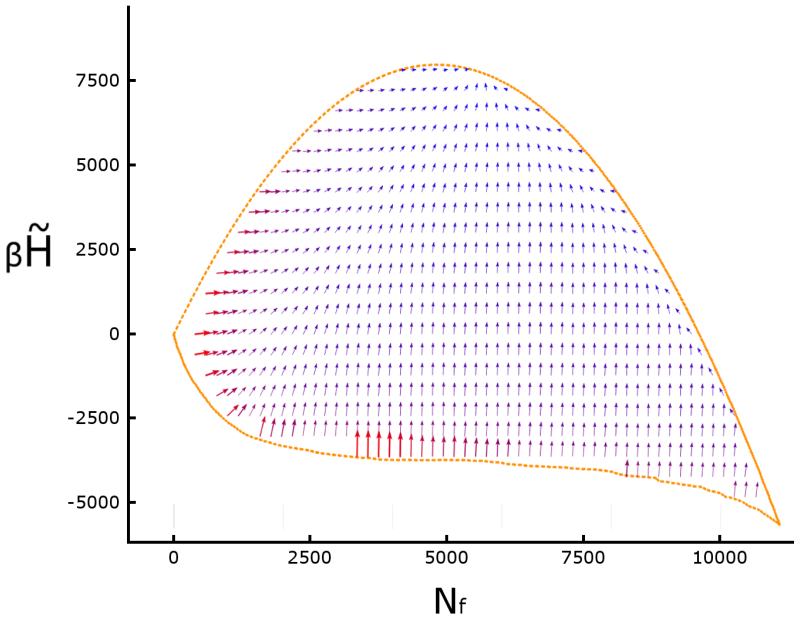

Fig. 2.35 The map of the state density (degeneracy of the macrostates) gradient for the $\phi = 0.95$ sample of aerogel. The projected state space is displayed for $\mu = -4.0$. The upper and lower projected-state-space boundaries are shown in orange.

aerogel sample. The size and colour of the arrows qualitatively represent the magnitude of the degeneracy gradient. If we had started such procedure from an equidegenerate line, it is straightforward to integrate the gradient along $\tilde{\mathscr{H}}$ dimension to obtain the relative values of the degeneracy across the entire projected state space. Fig. 2.36 shows the resultant degeneracy map. As we can see the scale of degeneracies is enormous, measured in thousands of orders of magnitude. Therefore, it is clear that the procedure would obtain pretty much identical results if started from a zero-temperature isotherm. Even though the isotherm does not necessarily coincide with any equidegenerate line, it can be expected that the density of states along its path does not vary more than a few orders of magnitude, which would be negligible in the context of the degeneracy map as a whole. The advantage of using an isotherm is that it is much easier to obtain than an equidegenerate line.

Having obtained an estimate of the density of states across the state space and knowing $\tilde{\mathscr{H}}$ of each macrostate, it is possible to predict the subspace of the state space that the system in equilibrium would reside for any given values of μ and β. To obtain the equilibrium state is simply a matter of applying the Boltzmann factor to the density of states and finding the maximum of the resultant function. In turn, the metastable states are represented by the local maxima of such function. Fig. 2.37 shows such probability maps for the $\phi = 0.95$ porosity

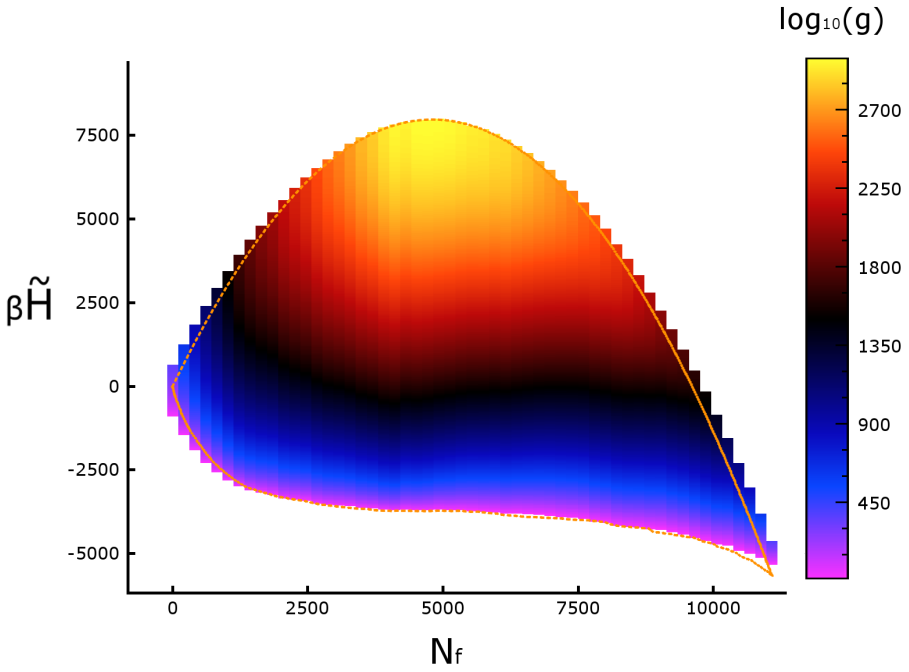

Fig. 2.36 The map of the state density (degeneracy of the macrostates) for the $\phi = 0.95$ sample of aerogel. The projected state space is displayed for $\mu = -4.0$. The upper and lower projected-state-space boundaries are shown in orange.

aerogel model. The location of the maxima indicated by the bright regions of the probability maps is consistent with the results obtained in the rest of this Chapter. Detailed comparison between the direct simulation results and such probability maps is yet to be investigated.

Despite the findings described above, the developed methodology has several limitations. At the very low temperatures, which is often of particular interest, the state space is relatively sparse, and therefore the relative errors in the degeneracy map that we have obtained increase. On top of that, if the simulation limit is moved sufficiently low in the state space, the ergodicity of the simulation cannot be guaranteed, since certain configurations start to be cut off from the rest of the state space, because the only path to reach them might need to go through the higher energy states. As we can see from Fig. 2.35 and Fig. 2.36, the degeneracy grows very rapidly at the lower part of the state space, and thus in most cases the mentioned problems vanish relatively close to the lower state-space boundary, however, in general some extra care need to be taken when studying systems at the very low temperatures.

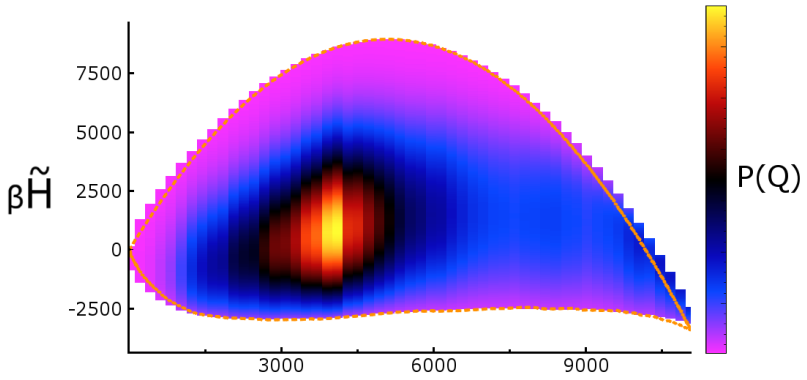

(a) $\beta = 0.7$

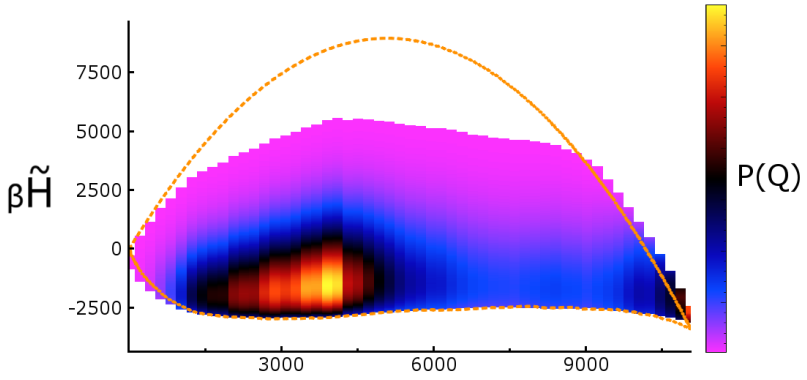

(b) $\beta = 1.0$

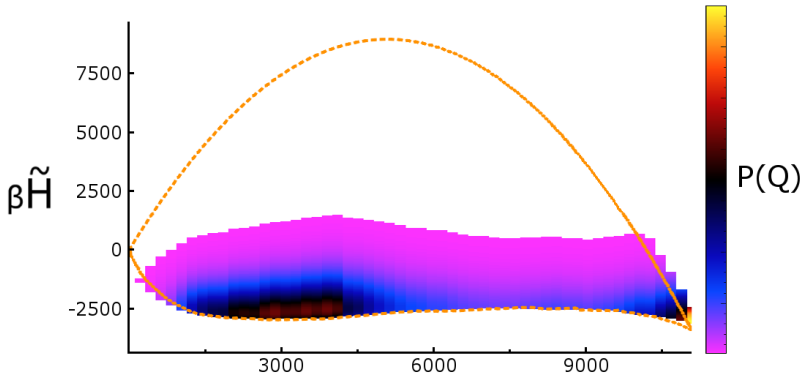

(c) $\beta = 1.5$

Fig. 2.37 The probability map for the $\phi = 0.95$ sample of aerogel. The projected state space is displayed for $\mu = -4.2$. The upper and lower projected-state-space boundaries are shown in orange.

2.4 Conclusions

In this Chapter, we have studied sorption processes in the lattice-gas systems, their state space, its properties and how it all relates to the microstructure of the sample itself. We have developed a way to visualize the state space of the system, that provides a more complete view of the sorption processes. Using the visualization techniques, we have established the grounds for our study by analysing relatively small systems, while employing the exact microstate enumeration methods. The confidence in exact methods allowed us to develop the basic intuition about systems behaviour, that we could rely upon as we continued the study using MC simulations of larger systems. We have found how the DLCA generated aerogel samples behave as μ and β are varied, investigated the potential phase transitions and how it all depends on the state space of the system. The state spaces of different porosity samples have been studied and the basic criteria to determine the existence of the phase transition have been established. We concluded the study by obtaining the density of states map, in theory granting us a complete knowledge of the system and its behaviour under any external conditions.

To summarise, the main results of this Chapter are the following:

(i) introduction of the concept of projected state space and its use for analysis of sorption in lattice-gas models;

(ii) development of efficient exact state-enumeration method based on Gray code and its application to small models;

(iii) conjecture about connection between the shape of the lower projected-state-space boundary and presence/absence of discontinuous phase transition in lattice-gas models;

(iv) development of the efficient algorithm for estimating the lower projected-state-space boundary;

(v) development of the efficient algorithms for calculation of equidegenerate lines, degeneracy-gradient maps and density of macrostates for projected state space;

(vi) checking all the algorithms and calculation of the properties and characteristics of the state space for models of aerogel of different porosity.

The knowledge of the relevant systems, their sorption behaviour and the corresponding state spaces, that we have gained in this Chapter explains the basic principles as well as many of the subtleties in this area of research. However, as mentioned in Sec. 2.3.10 there exist limitations to the applicability of the developed methodologies. The errors in quantitative

results are likely to increase at the very low temperatures or as the equilibrium state is nearing other sparse regions of the state space. Ultimately most of the MC simulation based methods employed in this Chapter in one way or another suffer from the quasi-nonergodicity and related problems. The complications arise once a subset of states is cut off from the rest of the state space by either a natural energy barrier (see Sec. 2.3.2) of an artificial simulation limit (see Sec. 2.3.10). In the following Chapters, we aim to overcome such problems as we continue to study lattice-gas systems using MC methods. Despite the complications mentioned above that we have encountered in this Chapter, the understanding and intuition developed here support strong grounds for the rest of this work and guide further research by providing the next set of questions to investigate.

Chapter 3

A single-walker approach for studying quasi-ergodic systems

3.1 Introduction

As we have seen in the previous Chapter, the exponentially large number of states, that the system can be in, restricts exact analysis to only small systems with the number of particles $N \sim 10^2$. The state space of the larger systems can be explored approximately by means of numerical techniques such as Monte Carlo (MC) sampling. However, at sufficiently low temperatures the use of conventional MC methods is prohibited by exponentially long transition times over the free-energy barriers (for literature review on the subject see Sec. 1.2). Such system can be trapped in one of the local free-energy minima and thus fail to ergodically explore the entire state space and achieve equilibrium. This behaviour is know as quasi-nonergodic and it has been discussed in Sec. 1.2 and Sec. 1.4.3, as well as widely addressed in the literature [137, 2, 138, 18, 22, 139, 140].

Over last several decades, numerous MC methods have been developed in order to alleviate the quasi-nonergodicity problem [141, 142, 139, 143] (for a brief review see Sec. 1.4.3). In this Chapter, we revisit one of the possible approaches developed in Refs. [163, 164], i.e. the jump-walking (JW) algorithm. As suggested by its name, the JW algorithm allows the system evolving in quasi-nonergodic regime (e.g. at low temperature) to jump into the states sampled from a different distribution (j-distribution) constructed for the system in ergodic regime (high-temperature) prior the beginning of the main simulation. This enables the algorithm to address the quasi-nonergodicity problem, i.e. to overcome free-energy barriers by means of such jumps and ergodically explore the state space even at low temperatures. However, the original JW algorithm used for studying dynamics of small atomic

clusters [165–167] had several drawbacks. First, it did not strictly satisfy the detailed balance [168, 169, 22, 170] and thus did not guarantee to achieve correct (Boltzmann) limiting distribution for finite-length simulations. Second, the acceptance rate for jumps can be very low due to small overlap between high- and low-temperature distributions [171]. Third, high amount of slowly accessible computer memory on hard drive (rather than random-access memory (RAM) at that time) required for construction of the j-distribution highly limited the applicability of the method only to relatively small systems [165–167]. The main rout in further development of the JW algorithm and removing these constraints was through the use of multiple walkers. In particular, the constraints of the JW method were addressed by developing several algorithms such as smart-walking [172], smart-darting [173], cool-walking [171], replica-exchange method [36, 169, 48] (REM) (also known as parallel tempering [3], multiple Markov chain method [50], parallel annealing [49]) and combination of REM with other algorithms[155, 156]. The key feature of the majority of such approaches is in parallel running and exchange of multiple non-interacting replicas of the original system [3, 22, 143] in contrast to a single-replica simulation employed by the JW.

Since the early 2000's the JW and its direct successors have been fully replaced by REM and similar algorithms. This was caused mainly by three factors: (i) parallel computing becoming cheap and accessible, (ii) growing requirements for ever larger systems to be simulated (and therefore stored in memory if one uses j-walking), and (iii) given the necessity to use the hard-drive storage, the implementation of the JW algorithm which maintains the detailed balance was significantly less straightforward than in the case of REM (see a direct comparison of the two methods in Ref. [22]). Currently, the multiple-replica methods and their combinations with multicanonical algorithm dominate studies of quasi-nonergodic behaviour (see Sec. 1.4.3).

In this Chapter, we revisit the original JW algorithm, develop it further for the current state of computing technology and explore its new capabilities for studying quasi-nonergodic behaviour of spin and lattice-gas models, which it has never been applied to. We demonstrate, that all the three major drawbacks of the original JW algorithm can be eliminated for a single-replica (single-walker) version of the algorithm. This gives an opportunity to study the low-temperature dynamics of lattice models of relatively large size ($\sim 10^5$ particles) using just a single-core simulations as contrasted to massively parallel ($\gtrsim 10^3$ cores [174, 175]) simulations by REM or similar methods.

The description of the original JW algorithm and its development are presented in Secs. 3.2-3.3. The performance of the improved JW algorithm is discussed for the Ising and lattice-gas models in Sec. 3.4. The conclusions are given in Sec. 3.5.

3.2 The original JW algorithm

Consider a system described by a Hamiltonian $H = H(Q)$ which is in equilibrium with the thermal bath at temperature T (measured in energy units), where Q denotes a state from the set $\{Q\}$ of all possible states the system can be in (canonical ensemble). Assume that the target temperature T is low enough so that the system is quasi-nonergodic at T. The JW algorithm [163, 164] explores the state space by means of the standard MC sampling at T augmented by jumps to the states which have been sampled at higher temperature $T' > T$. The temperature T' is chosen high enough to ensure that the system is ergodic at T'. The states, to which the jumps are attempted to, are sampled according to the equilibrium Boltzmann distribution (thus constructing the j-distribution),

$$P_{\mathrm{B}}(Q,\beta') = \frac{e^{-\beta' H(Q)}}{\sum_{\{Q\}} e^{-\beta' H(Q)}} = \frac{1}{Z(\beta')} e^{-\beta' H(Q)} \,, \tag{3.1}$$

where $Z(\beta')$ stands for the partition function at the inverse temperature $\beta' = (T')^{-1}$. These additional jumps from states Q_i (conventionally sampled at T) to states Q_j (sampled at T') allow for the ergodic sampling of the state space at temperature T if they preserve the detailed balance,

$$\mathbb{P}(Q_i \to Q_j) P_{\mathrm{B}}(Q_i,\beta) = \mathbb{P}(Q_j \to Q_i) P_{\mathrm{B}}(Q_j,\beta) \,, \tag{3.2}$$

with $\mathbb{P}(Q_i \to Q_j)$ being the transition probability between states Q_i and Q_j.

The transition probability, $\mathbb{P}(Q_i \to Q_j)$, can be split into sampling and acceptance probabilities,

$$\mathbb{P}(Q_i \to Q_j) = P_{\mathrm{s}}(Q_j,\beta') P_{\mathrm{acc}}(\beta,\beta',Q_i,Q_j) \,, \tag{3.3}$$

where $P_{\mathrm{s}}(Q_j,\beta')$ is proportional to the Boltzmann probability at T'. The detailed balance condition given by Eq. (3.2) imposes a constrain on the acceptance probability P_{acc},

$$\frac{P_{\mathrm{acc}}(\beta,\beta',Q_i,Q_j)}{P_{\mathrm{acc}}(\beta,\beta',Q_j,Q_i)} = e^{(\beta'-\beta)(H(Q_j)-H(Q_i))} \,, \tag{3.4}$$

which is satisfied by the following choice [163],

$$P_{\mathrm{acc}}(\beta,\beta',Q_i,Q_j) = \min(1, e^{(\beta'-\beta)(H(Q_j)-H(Q_i))}) \,. \tag{3.5}$$

The standard Metropolis expression for P_{acc} is recovered from Eq. (3.5) as T' goes to infinity ($\beta' \to 0$), since the j-distribution is then uniformly random. On the other hand, if the state space is sampled at the same temperature as that of the j-distribution ($\beta = \beta'$), the detailed

balance is maintained with the acceptance probability independent of the state parameters, since the j-distribution itself would already encompass the required transition probability.

3.3 Development of the algorithm

Although the original JW algorithm successfully tackled the quasi-nonergodicity problem, its implementation faced technical constraints at time of its creation. As mentioned in Sec. 3.1, they were related to (i) inability to maintain the detailed balance, (ii) inefficiency of jumps between low- and high-temperature states with increasing temperature difference and (iii) low access rate of states in the j-distribution stored on hard drive. In this Section, we discuss solutions to these problems and suggest an updated protocol for the JW algorithm, which addresses the mentioned issues. Namely, we show that (i) the detailed balance can be maintained exactly within the new protocol presented below, (ii) the standard temperature-sequence approach can be employed for the JW to increase the jump-acceptance probability and (iii) RAM can be used for accessing the j-distribution.

In the past, the detailed-balance problem for the JW algorithm was addressed by two approaches. The first one [176] requires multiple parallel simulations at high temperature T', as well as one run at the target temperature, $T < T'$. The j-distribution is sampled at runtime by randomly choosing one of the systems simulated at T' and attempting a transition to its current state. Use of several systems for sampling at high temperature prevents possible correlations, which may appear if sampling from the same simulation is done before it loses memory of the previously sampled state. While this method does not use large amount of memory since the states are sampled at runtime, it is based, similarly to REM, on multiple simulations running at the same time and it was replaced by well established REM [22].

The second approach [163] is to perform the initial sampling of the j-distribution at T' prior to the main simulation and save the recorded states. During the main simulation at T, a random state is periodically chosen from the saved j-distribution. The state could be accepted or rejected according to the Metropolis rule (Eq. (3.5)). Due to insufficient amount of RAM, the set of states, $\{Q^{\mathrm{rec}}(T')\}$, recorded at temperature T', was kept on the hard-drive storage, thus severely limiting access speed and manipulation capabilities. In order to maintain the detailed balance, one had to ensure that the same state from $\{Q^{\mathrm{rec}}(T')\}$ is not chosen twice [177, 22]. Within this strategy, the probability of picking the same element in $\{Q^{\mathrm{rec}}(T')\}$ more than once, decreases as the ratio of the initially recorded and eventually required states increases. However, even if significantly more states were recorded than the simulation at T ultimately needed, this probability would remain finite and the detailed balance would be only approximately obeyed.

Below, we suggest a new algorithm which resolves the detailed-balance problem for a single-walker JW method. The key point of our approach is in using RAM for j-distribution which permits easy manipulations with recorded states. In particular, the recorded states stored in RAM can be easily removed once chosen for a possible jump-transition. Moreover, only necessary number of states from j-distribution for its unbiased sampling can be stored thus avoiding any storage overhead.

A further challenge faced by JW as well as REM and related methods arises due to the transition acceptance probability vanishing as the difference between temperatures T and T' increases [22]. This hinders access to the states sampled at T' and thus prevents efficient equilibration of the system. The standard solution to this problem for the JW algorithm [177] is to run the simulation sequentially at several temperatures, $\{T_\alpha\}$ $(\alpha = 1,\ldots,n)$, with $T_1 > T_2 > \ldots > T_n$. The initial temperature T_1, is set high enough so that the system is fully ergodic at T_1 (see Sec. 3.4.4). Any two consecutive temperatures T_α and $T_{\alpha+1}$ are chosen in such a way that the Boltzmann distributions at these temperatures overlap, i.e. there exist states which the system can visit with non-vanishing probabilities at both temperatures. At each temperature T_α, the target states Q_j for the jump transitions are sampled from $\{Q^{\mathrm{rec}}(T_{\alpha-1})\}$, and simultaneously the new j-distribution is constructed by recording the states to $\{Q^{\mathrm{rec}}(T_\alpha)\}$. This way the system remains ergodic at each successive temperature down to T_n. In case of REM, a typical solution to the equivalent problem is to employ multiple different temperature replicas running concurrently [37, 47].

There has been a lot of work done to optimise the set of selected temperatures for REM [153, 154] as well as the JW [177] and other methods [178]. Typically, the temperatures are selected to follow geometric sequence [179]. If the free-energy landscapes are sufficiently similar across the temperature range (which is not necessary the case e.g. for lattice polymers [180] and Lennard-Jones clusters [22]), the resultant overlaps between the subsets of the state space explored at different consecutive temperatures are comparable. This, in turn, leads to approximately the same acceptance probabilities for replica exchanges or JW jumps. However, in general, the free-energy landscape can be significantly different for different problems (e.g. for proteins [14] and structural glasses [181, 7]). Therefore, the temperature-sequence optimisation strategy can depend on a particular system and preliminary simulations are often necessary to examine the subsets of available states at each temperature and optimise $\{T_\alpha\}$ [169]. The question of such optimisation is outside the scope of this work. Instead, we provide a general-purpose multi-stage protocol for the JW algorithm given an optimised set of temperatures $\{T_\alpha\}$. By using this protocol the JW MC simulation running on a single thread ergodically explores the state space of the system exhibiting quasi-nonergodic behaviour below a certain characteristic temperature T_c.

The protocol is as follows:

(i) Run the standard single-flip MC simulation at temperature $T_1 > T_c$. Record the current state of the system at equal intervals between the regular MC steps, so that R states are recorded into array $\{Q^{\mathrm{rec}}(T_1)\}$ until the total required number, S, of MC steps is executed.

(ii) Decrease the temperature to T_α, where T_α is the next element in the optimised set of temperatures $\{T_\alpha\}$, and run the standard single-flip MC simulation again. Choose R step indexes at random from 1 to S and, once each of them is reached in the simulation, attempt a transition to a state randomly picked from $\{Q^{\mathrm{rec}}(T_{\alpha-1})\}$ with the acceptance probability P_{acc} defined by Eq. (3.5). Whether the transition is accepted or not, remove the suggested state from $\{Q^{\mathrm{rec}}(T_{\alpha-1})\}$, so that it cannot be chosen again. In the meantime, continue to record the states as in (i) into the new array $\{Q^{\mathrm{rec}}(T_\alpha)\}$, which now contains an ergodic sample of the state space at T_α. Proceed until S MC steps are executed.

(iii) Set $\{Q^{\mathrm{rec}}(T_\alpha)\}$ as the new $\{Q^{\mathrm{rec}}(T_{\alpha-1})\}$ from which to sample the j-distribution. Repeat (ii) and (iii) until the final temperature T_n in the list $\{T_\alpha\}$ is reached.

The algorithm provides us with two free parameters: R and S. The constraint on precision naturally sets the required number of MC steps, S, to be completed. The interval between the attempted jumps, S/R, is determined by the desired frequency of the jumps, $(R/S) \times P_{\mathrm{acc}}$, and limited by how many states R it is feasible to store in memory. For REM, it has been shown that given there were no computational constraints, equilibration accelerates as the replica-exchange frequency increases [182, 183]. However, since the JW algorithm simulates only one copy of the system at any time, the jump transition, unlike the replica-exchange process in REM, requires to modify larger amounts of data associated with the state of the system. Computational cost of the transition is thus system dependent and could impose an upper bound on the optimal jump rate. Since the attempted states are removed from the storage independently of their acceptance, it is efficient to keep P_{acc} as high as possible, minimising the number of stored states. Therefore, we are only free to adjust R in order to set the required jump rate.

For the purpose of clarity, the above analysis assumes that the average acceptance probability P_{acc} is the same for each $\{T_\alpha\}$. However, in general, P_{acc} depends on temperature and the constant jump-rate can be achieved by adjusting the number of stored states $\{R_\alpha\}$ for each temperature T_α. All simulation parameters used for testing the algorithm are given in Sec. 3.4.4.

Finally, it is straightforward to generalise the JW method for parameter-dependent Hamiltonians in a similar way to that used in the case of multidimensional REM [152]. For a system described by a Hamiltonian $H = H(Q, \{\lambda_m\})$, where $\{\lambda_m\}$ is a set of parameters (e.g. chemical potential in the case of grand-canonical ensemble), an equivalent formalism leads to a generalised version of the expression for P_{acc},

$$P_{\text{acc}}(\beta, \beta', Q_i, \{\lambda_m\}, Q_j, \{\lambda'_m\}) = \min\left(1, e^{\beta'(H(Q_j,\{\lambda'_m\})-H(Q_i,\{\lambda'_m\}))-\beta(H(Q_j,\{\lambda_m\})-H(Q_i,\{\lambda_m\}))}\right). \tag{3.6}$$

This allows us to sample the non-local transitions from a distribution constructed not only at a different temperature, but also based on Hamiltonian characterised by different values of parameters λ'_m. An example of such analysis is presented in Sec. 3.4.2.

To summarise, the updated JW algorithm described above maintains the detailed balance by storing the j-distribution in RAM and, given an optimised set of temperatures, is able to explore ergodic behaviour of the system for relatively low temperatures in lattice models as demonstrated in Sec. 3.4.

3.4 Application of the algorithm

In order to demonstrate the updated JW MC method and further clarify its performance, the algorithm has been implemented for two systems: a standard 3d Ising model and lattice-gas model for sorption of fluid in porous media. Both models, their implementation details and simulation results are described in the following subsections.

3.4.1 Ising model

The Ising model is one of the simplest models exhibiting phase transitions, making it well suited to test the performance of the JW algorithm. Since the memory required to record the states of such system grows linearly with the system size N, the total memory needed to store the multiple recorded states has been prohibitively high and the original JW method [163, 164] has never been implemented for the Ising or other spin models. Here, we demonstrate that the updated JW algorithm is capable of handling relatively large systems ($N \sim 10^5$) using only the memory of a regular contemporary personal computer and a single CPU core.

As a test example we study a 3d Ising model on a body centred cubic (bcc) lattice of $N = 2 \times L^3$ ferromagnetically interacting spins (with periodic boundary conditions), where L

is the number of unit cells along the edge of a cubic sample. The dimensionality and lattice type are convenient for comparison with the lattice-gas models discussed in Sec. 3.4.2.

The Hamiltonian $\mathscr{H}_{\mathrm{I}}$ of the Ising model with the interaction strength between two neighbouring spins $J = 1$ in external field H (measured in units of J) is given by,

$$\mathscr{H}_{\mathrm{I}} = -J \sum_{\langle ij \rangle} \tau_i \tau_j - H \sum_{i}^{N} \tau_i \ , \tag{3.7}$$

where the first sum is taken over all the nearest-neighbour spin pairs $\langle ij \rangle$ and $\tau_i \in \{+1, -1\}$ is the spin variable for a spin on site i.

In order to visualise the multi-dimensional state space of the spin system, it is convenient to study its projection onto two dimensions representing magnetisation $M(\{\tau_i\}) = N^{-1} \sum_i^N \tau_i$ (the order parameter) and energy $\mathscr{H}_{\mathrm{I}}(\{\tau_i\})$ of a microstate $\{\tau_i\}$ (see the upper panels in Fig. 3.1). In general, several spin configurations can have the same values of M and $\mathscr{H}_{\mathrm{I}}$ and it is convenient to define a macrostate $Q(M, \mathscr{H}_{\mathrm{I}})$, with the degeneracy $g(Q(M, \mathscr{H}_{\mathrm{I}}))$, as a set of g microstates with the same values of $M(\{\tau_i\})$ and $\mathscr{H}_{\mathrm{I}}(\{\tau_i\})$. As MC sampling is running, the number of MC steps (or total time for kMC [135]) the system spends in a macrostate $Q(M, \mathscr{H}_{\mathrm{I}})$ is recorded. In the ergodic regime, this quantity is proportional to the probability distribution for system to be in a certain macrostate and below we refer to it as distribution of visits. The probability to visit some macrostates can be very small and domains of the state space containing such states (far tails of the probability distribution) cannot be explored by the JW algorithm for finite value of parameter S [22] (see Sec. 3.3). The results for the distribution of visits are displayed on the M-$\mathscr{H}_{\mathrm{I}}$ plane using the following colour-scheme. On the white background of unexplored state space, the most to the least visited macrostates are coloured in yellow to dark blue, respectively.

Fig. 3.1a illustrates the results of the updated JW MC simulations on a $18 \times 18 \times 18$ bcc lattice ($N = 11664$) for zero external field $H = 0$ and several temperatures. At sufficiently high temperatures, all transitions between microstates are approximately equally likely, and therefore mainly the degeneracy of the macrostates determines which part of the state space is visited within the simulation time (see the single-peak distribution of visits around zero magnetisation for $\beta = 0.13$ in the upper panel of Fig. 3.1a). As temperature decreases and occupation of the low-energy states becomes exponentially more probable (according to Eq. (3.1)), the subspace of explored states shifts to the lower energies (cf. distributions at $\beta = 0.13$ and $\beta = 0.16$). Moreover, with decreasing temperature, the distribution of visits becomes bi-modal (see e.g. the distribution for $\beta = 0.16$) reflecting the symmetry of the spin system with respect to the change of the spin orientation. This is due to the fact that the degeneracy of macrostates with approximately zero magnetisation decreases (with

decreasing energy) more rapidly than of those with finite magnetisation. The two peaks in the distribution of visits become more pronounced with further decrease in temperature (cf. distributions at $\beta = 0.17$ and $\beta = 0.19$) and at $T = 0$ only two lowest energy states retain non-zero probability, each corresponding to all spins aligned either up or down.

This picture can be described by using the language of the free-energy landscapes, referring to the dependence of the relative free energy $f(M) \equiv \beta(F(M) - F)$ solely on M, where

$$F = -\beta^{-1} \ln Z = -\beta^{-1} \ln \left[\sum_{\{\tau\}} e^{-\beta \mathscr{H}_{\mathrm{I}}} \right] , \tag{3.8}$$

$$F(M) = -\beta^{-1} \ln Z(M) = -\beta^{-1} \ln \left[\sum_{\mathscr{H}_{\mathrm{I}}} g(Q(M, \mathscr{H}_{\mathrm{I}})) e^{-\beta \mathscr{H}_{\mathrm{I}}} \right] . \tag{3.9}$$

The summation for free energy is taken over all microstates $\{\tau\}$ in Eq. (3.8) and over all macrostates $\{Q(M, \mathscr{H}_{\mathrm{I}})\}$ with fixed magnetisation M in Eq. (3.9). It follows from Eqs. 3.8-3.9 that the value of $f(M) = -\ln(P(M))$ is related to the probability $P(M)$ for the system to be in a state with fixed magnetisation, $P(M) = Z(M)/Z$. This probability can be estimated from the simulations by integrating the distribution of visits over the energies for fixed value of M. Thus obtained free-energy landscapes, $f(M)$, are shown in the bottom panels of Fig. 3.1.

For temperatures above some critical value, T_{c}, the relative free energy $f(M)$ has a single minimum at zero magnetisation (see yellow line marked by solid squares in the bottom panel of Fig. 3.1a). At $T \approx T_{\mathrm{c}}$, two symmetric minima at finite values of M separated by a barrier at $M \simeq 0$ appear (see red line marked by open squares). With decreasing temperature, the barrier between minima increases, the minima become deeper (see lines marked by the circles) and their positions tend to $M = \pm 1$ at zero temperature.

The picture described above cannot be seen by means of MC simulations governed only by the standard single-flip mechanism, such as kMC, because the transitions through the free-energy barrier become exponentially rare with decreasing temperature. Eventually, within the standard single-flip MC simulations, the system settles down in one of the free-energy minima shown in Fig. 3.1a, thus violating the equilibrium and breaking ergodicity. This does not happen for the JW MC simulations for which the multiple-flip jump transitions to the states recorded at higher temperatures are possible. These transitions enable exploration of all the states proportionally to the canonical distribution at the current temperature thus maintaining an ergodic sampling of the state space and all free-energy minima present at sufficiently low temperatures become achievable within a single JW simulation.

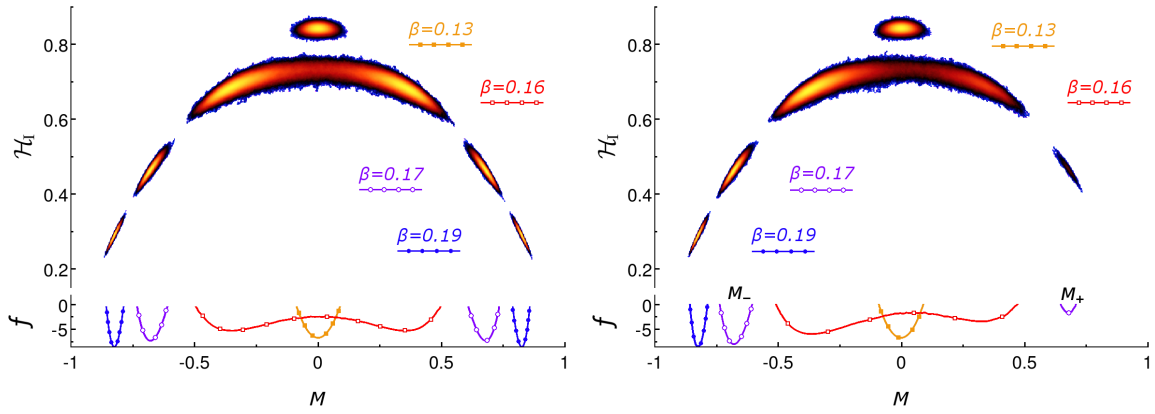

(a) $H = 0$ (b) $H = -5 \times 10^{-4}$

Fig. 3.1 (a) Upper panel: the probability distribution of visits to macrostates characterised by energy (per spin) $\mathscr{H}$ (vertical axis) and magnetisation M (horizontal axis) for the zero-field Ising model described in the text. The distributions of visits at four labelled temperatures are shown. Two distributions for $\beta = 0.17$ and $\beta = 0.19$ consist of two disconnected regions each located at approximately the same horizontal level. The most visited macrostates are shown in yellow while the less visited ones are in the dark-blue. The white background refers to unvisited macrostates. Lower panel: Free-energy landscapes, i.e. $f(M)$, obtained for four distributions shown in the upper panel. The curves marked by different symbols refer to temperatures labelled accordingly in the upper panel. (b) The same as in (a) but for non-zero external field $H = -5 \times 10^{-4}$. The values M_+ and M_- in the lower panel indicate the positions of minima of $f(M)$ for $\beta = 0.17$. The values of other parameters used in simulations are given in Sec. 3.4.4.

In the presence of a finite external field, e.g. $H < 0$, (see Fig. 3.1b), similarly to the zero-field regime, the free-energy minima develop at positive, M_+, and negative, M_-, values of magnetisation (see e.g. the solid line marked by open circles for $\beta = 0.16$ in the bottom panel of Fig. 3.1b). The negative external field breaks the symmetry and makes the minimum at M_- deeper. The JW algorithm samples the states at both free-energy minima according to the canonical probability distribution. This means that only the energy difference at these minima is significant for their relative occupation and the barrier between the minima is irrelevant for the JW sampling. As the macrostates around M_+ become ever less probable with decreasing temperature, the jump-transitions from this minimum to the minimum at M_- are favoured over the reverse ones (see the asymmetric bi-modal distributions of visits at $\beta = 0.16$ and $\beta = 0.17$ in the upper panel of Fig. 3.1b). Therefore, the amount of time that the system spends exploring the M_+ minimum gradually decreases with decreasing temperature (because the energy difference between minima increases) until it eventually

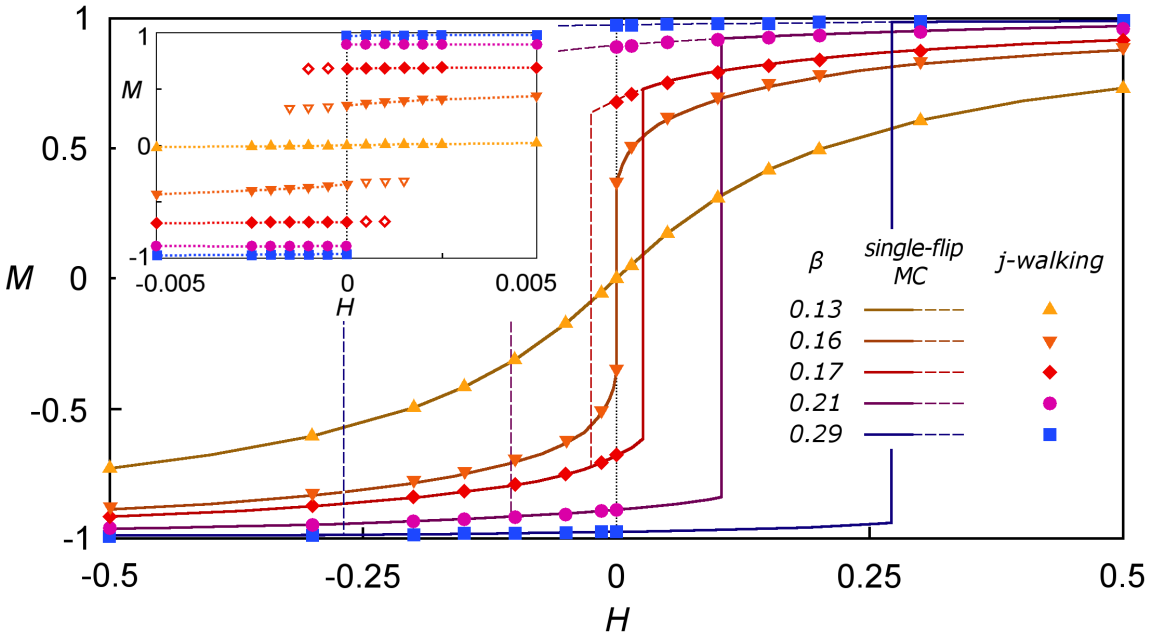

Fig. 3.2 Magnetisation reversal, i.e. the mean magnetic moment M *vs* external field H, for the Ising model defined on bcc lattice with the same parameters as in Fig. 3.1, for several temperatures as marked in the legend. The data represented by solid (dashed) lines refer to increasing (decreasing) magnetic field and were obtained by the single-flip conventional kMC simulations. For equilibrium isotherm at $\beta = 0.13$ (▲), the solid and dashed lines coincide. For $\beta = 0.16$ (▼), a small hysteresis loop exists but is not distinguishable on the scale of the graph. The symbols represent results of the JW MC simulations. The inset shows magnified region around $H = 0$. The solid (open) symbols in the inset refer to the mean magnetisation at energetically favourable (unfavourable) minimum of the free energy. The dotted horizontal lines in the inset serve as a guide for eye representing the value of magnetisation at the deepest minimum in the free-energy landscape. The dotted vertical line marks the transition at $H = 0$.

vanishes completely for the precision set up within the JW algorithm (see the single-peaked distribution at $\beta = 0.19$ in Fig. 3.1b).

An example of equilibrium behaviour of the Ising model obtained by means of the JW MC simulations contrasted with the results of the single-flip kMC is presented in Fig. 3.2. The standard way to investigate $M(H)$ is to run a conventional single-flip MC simulations at fixed temperature gradually incrementing the external field from sufficiently negative (so that $M \simeq -1$) to high positive values ($M \simeq 1$) and then back again. Under this protocol, a hysteresis loop is observed [184] (the areas between solid and dashed lines in Fig. 3.2 for $\beta = 0.17, 0.19, 0.21$). In this hysteresis loop (in some systems found experimentally as well), the magnetisation of the sample starts to align with the external field and then reverse in response to the change in H. Such a hysteresis loop is still observed even if the

change of the field takes place adiabatically slowly. This implies that the magnetic material is not in equilibrium and is stuck in a state which it cannot leave on the single-flip MC (or experimental) time-scales with the thermal energy available, i.e. the system gets stuck in a metastable state.

The JW MC algorithm provides a possibility to avoid trapping in metastable states, i.e. to remove the hysteresis loop, by sampling the state space according to the canonical equilibrium distribution. Indeed, Fig. 3.2 (see the solid symbols corresponding to the JW data) shows a clear and expected phase transition at $H = 0$ below the critical temperature T_c (for bcc lattice,$T_c^{-1} \simeq 0.157371(1)$ [185]). Therefore, the JW MC simulations produce equilibrium isotherms, i.e. $M(H)$, which are obfuscated by hysteresis phenomenon in the single-flip MC simulations or experiment. The inset in Fig. 3.2 magnifies the transition region. In this region, for fixed temperature (below critical) and H close to zero, the JW MC simulations detect two minima of the free-energy, energetically favourable (solid symbols) and unfavourable (open symbols), which are also shown in lower panel of Fig. 3.1b. The probability to visit these minima follows the canonical equilibrium distribution and thus the accessibility of the energetically unfavourable minimum becomes exponentially small with increasing energy difference between the minima. Eventually, the energetically unfavourable minimum cannot be detected by the JW sampling with fixed precision. This occurs for relatively small deviations of the external field from zero, e.g. $|H| \sim 10^{-3}$ for $\beta = 0.17$. For the single-flip MC simulations, the height of the barrier between the free-energy minima matters and energetically unfavourable minima, in this case, act as traps, i.e. they become metastable states from which the system can escape only due to action of relatively large external fields, e.g. $H \sim 10^{-1}$ for $\beta = 0.17$.

The behaviour of the Ising model described above is well established and served as a test for the JW MC simulations. As we can see from the results presented in Figs. 3.1-3.2, the JW MC simulations reproduce all expected phenomena and, in addition, allow the magnetisation equilibrium isotherms, i.e. $M(H)$, to be obtained for sufficiently low temperatures.

3.4.2 Lattice-gas model

In order to test the applicability of the updated JW MC algorithm to relatively large and complex systems, we have implemented the method for a lattice-gas model to study fluid sorption in porous media.

In this Section, porous media is represented by two models: (i) a small lattice toy model for which the exact numerical solution is available (see Sec. 3.4.3) and (ii) structural models of silica aerogel.

The structural models of silica aerogel [186, 62, 68] were generated using the diffusion-limited cluster-cluster aggregation (DLCA) algorithm [187] run on bcc lattice consisting of $N_{\rm tot} = 2 \times L^3$ sites. Each realization of the modelled aerogel sample can be characterized by porosity, $\phi = {N_{\rm tot}}^{-1} \sum_i^{N_{\rm tot}} \eta_i$, which for real silica aerogels varies from 80% to 99.8% [68], and fractal dimension D. In this Section, we use the models for which the DLCA algorithm was tuned to create cubic model samples with porosity $\phi = 95\%$ and $D \simeq 1.78$ that accurately reproduce the structural properties of the aerogels when compared to experimental neutron-scattering data [187, 188]. In order to avoid fictitious "pinning" effects and simulate realistic desorption mechanism [189, 71], an empty space slab ($\eta = 1$) of width equal to 4 lattice spacings, followed by a layer of enforced vapour ($\eta = 1$, $\tau = 0$) was attached to both sides of aerogel sample in z-direction. In the x- and y-directions periodic boundary conditions were imposed. The use of the bcc lattice allows to avoid lattice artefacts such as the faceted growth regime observed for the simple cubic lattices [71].

In the lattice-gas model [190, 191, 129, 4], each out of $N_{\rm tot}$ lattice sites can be occupied by a fluid or a matrix particle, as described by the occupancy variables τ_i and $1 - \eta_i$, respectively, which are equal to unity (zero) for occupied (unoccupied) sites. The matrix sites do not change their state, i.e. cannot be occupied by the fluid, and their concentration is quantified by porosity, $\phi = {N_{\rm tot}}^{-1} \sum_i^{N_{\rm tot}} \eta_i$. In the simulation setup, the porous material is assumed to be an open system of $N = \phi N_{\rm tot}$ pore sites connected to a reservoir of fluid particles. The number of the fluid particles $N_{\rm f} = \sum_i^{N_{\rm tot}} \eta_i \tau_i$ in the system and thus the energy $\mathscr{H}$ (with the lattice-gas Hamiltonian given by Eq. (1.2)) of the system can vary, but the volume V (or equivalently N) of a particular system is fixed. The chemical potential μ and temperature T can be altered depending on the protocol of the simulation, but each such change is followed by the equilibration period, when no measurements are made. For a given set of μ, T and N, the system can be in a set of states which form a grand canonical ensemble described by the grand partition function given by Eq. (1.3).

For the JW MC algorithm, the jump-transition acceptance probability (see Eq. (3.6)) for such system is given by the following expression,

$$P_{\rm acc}(\beta, \beta', Q_i, \mu, Q_j, \mu') = \min \left[1, e^{\beta'(\tilde{\mathscr{H}}(Q_j,\mu') - \tilde{\mathscr{H}}(Q_i,\mu')) - \beta(\tilde{\mathscr{H}}(Q_j,\mu) - \tilde{\mathscr{H}}(Q_i,\mu))} \right] , \quad (3.10)$$

where the non-local multiple-flip transitions from a current microstate of the system simulated at β and μ to a state in the j-distribution previously sampled at β' and μ' are allowed. However, to assure clarity and simplicity of the illustrative examples, the results presented below refer to a value of μ fixed within a single JW simulation.

The quantity of interest in our analysis is the mean equilibrium fluid density (order parameter), $\langle \rho \rangle = \langle N_{\rm f} \rangle / N$ as a function of chemical potential and temperature. The angular

brackets mean thermodynamic averaging over the states $\{\tau\}$ in the grand canonical ensemble, i.e.

$$\langle\rho\rangle = Z_\mu^{-1} \sum_{\{\tau\}} \rho(\{\tau\}) e^{-\beta\tilde{\mathscr{H}}} \ . \tag{3.11}$$

The lattice-gas model is a discrete model and thus the density ρ takes $N+1$ discrete values, $\rho(N_\mathrm{f}) = N_\mathrm{f}/N = 0, 1/N, \ldots, (N-1)/N, 1$. Therefore, similarly to Eq. (3.9) it is convenient to rearrange the summation in Eq. (3.11) by first summing over all macrostates $Q(\rho, \tilde{\mathscr{H}})$ (constructed in the same way as $Q(M, \mathscr{H}_1)$ in Sec. 3.4.1) with fixed number of fluid sites (i.e. fixed ρ), but variable energy $\tilde{\mathscr{H}}$, and then over fluid densities, ρ,

$$\langle\rho\rangle = Z_\mu^{-1} \sum_\rho \rho \sum_{\{\tau(\rho)\}} e^{-\beta\tilde{\mathscr{H}}(\rho)} = \sum_\rho \rho \sum_{\tilde{\mathscr{H}}} P(Q(\rho, \tilde{\mathscr{H}})) \ . \tag{3.12}$$

The probability, $P(Q(\rho, \tilde{\mathscr{H}}))$, introduced in Eq. (3.12), to find the system in a certain macrostate $Q(\rho, \tilde{\mathscr{H}})$ is given by

$$P(Q(\rho, \tilde{\mathscr{H}})) = Z_\mu^{-1} g(Q(\rho, \tilde{\mathscr{H}})) e^{-\beta\tilde{\mathscr{H}}(\rho)} \ , \tag{3.13}$$

where $g(Q(\rho, \tilde{\mathscr{H}}))$ stands for the degeneracy of macrostate $Q(\rho, \tilde{\mathscr{H}})$, and thus the probability of the system to be in a state with a particular ρ is $P(\rho) = \sum_{\tilde{\mathscr{H}}} P(Q(\rho, \tilde{\mathscr{H}}))$.

Similarly to $F(M)$ in the case of the Ising model (see Sec. 3.4.1), the grand-potential landscape can be defined as the ρ-dependence of the grand potential $\Omega(\rho) = -\beta^{-1} \ln Z_\mu(\rho)$, where $Z_\mu(\rho) = Z_\mu P(\rho)$. The shape of $\Omega(\rho)$ can describe possible phases of the system. For example, if $\Omega(\rho)$ has one minimum then the system is characterised by a single phase. However, if two minima appear in $\Omega(\rho)$ this might be an indication of coexistence of two phases with different densities. In order to visualize the grand-potential landscape, similarly to $f(M)$ in Sec. 3.4.1, we plot $\omega(\rho) \equiv \beta(\Omega(\rho) - \Omega) = -\ln(P(\rho))$ [12] (calling this quantity as grand-potential landscape), where $\Omega = -\beta^{-1} \ln Z_\mu$ is the total grand-potential of the system (see the bottom panels in Fig. 3.3).

The state space of the lattice-gas model can be sampled by means of the JW MC algorithm in a similar way as that of the Ising model. The distribution of visits of different macrostates is displayed in the ρ-$\tilde{\mathscr{H}}$ plane (see the upper panels in Fig. 3.3), where the same colour-scheme is used as in Fig. 3.1. Approximate high- and low-temperature boundaries of the projected state space are also computed (see Sec. 2.3.7 for details) and displayed for reference in the upper panels of Fig. 3.3 (the red (dashed) and green (solid) lines, respectively).

Fig. 3.3 presents the results of the JW MC simulations for sorption of fluid in a model sample of aerogel for two different values of μ and several temperatures. At low temperatures

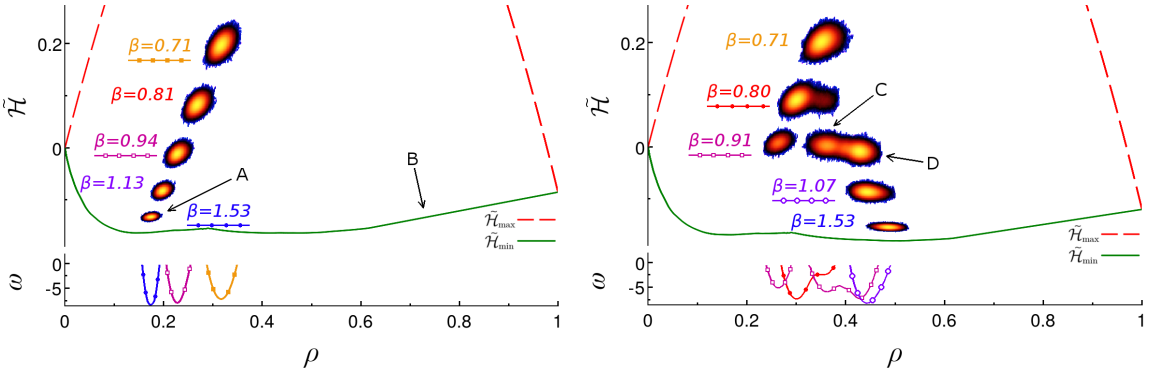

(a) $\mu = -4.18$ (b) $\mu = -4.145$

Fig. 3.3 Upper panel: The probability distribution of visits of macrostates characterised by relative energy per pore site, $\tilde{\mathscr{H}}$ (vertical axis), and density ρ (horizontal axis) for the lattice-gas model applied to aerogel sample as described in the text. The distributions at five labelled temperatures for (a) $\mu = -4.18$ and (b) $\mu = -4.145$ are shown using the same colour scheme for density of visits as in Fig. 3.1. The estimates for the upper and lower boundaries of the projected state space are shown by the dashed (red) and solid (green) lines. The arrows A, B, C and D point to the states characterised by different fluid densities (see Fig. 3.5 for graphical illustrations of fluid configurations). Bottom panels: Grand-potential landscapes, i.e. $\omega(\rho)$, obtained for three temperatures: (a) $\beta = 0.71$ (▪), $\beta = 0.94$ (□) and $\beta = 1.53$ (•); (b) $\beta = 0.80$ (•), $\beta = 0.91$ (□) and $\beta = 1.07$ (○) (see Sec. 3.4.4 for the values of other parameters).

for both values of μ shown in Fig. 3.3, the grand-potential landscape has two minima, as can be inferred from the shape of the low-temperature boundary, $\tilde{\mathscr{H}}_{\min}(\rho)$ (the solid green line in the upper panels). The low-density minimum (e.g. at $\rho \simeq 0.15$ for $\mu = -4.18$) represents the fluid distribution in the system for which the sites occupied by fluid are concentrated only around and nearby the quenched matrix sites (see Fig. 3.5A). The minimum at the higher values of ρ corresponds to the aerogel sample filled with fluid, and only the added surface layers left unoccupied (see Fig. 3.5B). As temperature is decreased, the explored part of the state space shifts down until the system eventually settles in a minimum (see the sequences of distribution of visits corresponding to gradually decreasing temperature in the top panels of Fig. 3.3). For sufficiently low or high values of μ, the system straightforwardly descents to the low- or high-density minimum, respectively (the descent to the low-density minimum is shown in Fig. 3.3a). However, at the intermediate values of μ (see Fig. 3.3b) the state space sampled by the JW MC simulation splits into several regions and thus the distribution of visits has several peaks (see e.g. the distribution at $\beta = 0.91$ in the top panel of Fig. 3.3b). In this case, the right peak in the distribution (see colour-map for $\beta = 0.80$ in Fig. 3.3b) becomes dominant with decreasing temperature and eventually the system settles down in the

high-density minimum of the grand-potential. The single-flip MC simulations would not be able to follow such equilibration of the system for $\mu = -4.145$. This is due to development of the large entropic barrier between two minima leading to significant slowing down of single-flip dynamics. Without the jump-transitions provided by the JW MC simulations, the region of the state space corresponding to the high-density peak could not be sufficiently visited and system would remain in the low-density state, i.e. in the metastable state.

Fig. 3.3b also shows the existence of the secondary split of the high-density peak in the distribution of visits into two peaks (see the two neighbouring bright spots on the right in the colour map for $\beta = 0.91$) corresponding to two minima in the grand-potential landscape at $\rho \simeq 0.35$ and $\rho \simeq 0.45$ (see the solid line marked by the open squares in the bottom panel). This feature (two high-density peaks in the distribution of visits) is an aerogel-sample specific (see Fig. 3.5). In addition to these two high-density configurations at $\beta = 0.91$, there is a low-density one with $\rho \simeq 0.25$ (see the corresponding minimum of $\omega(\rho \simeq 0.25)$ in the bottom panel). For this configuration, mainly the pore sites surrounding the matrix sites are occupied by fluid.

The JW MC simulations performed for fixed temperature and different values of chemical potential provide equilibrium isotherms, i.e. $\rho(\mu)$, for sorption in aerogel (see the data points labelled by circles, squares and diamonds in Fig. 3.4a). The solid symbols refer to the densities corresponding to the deepest minimum of $\omega(\rho)$ while the open symbols (see the inset in Fig. 3.4a) are related to other minima of $\omega(\rho)$ available for given values of μ and β. A typical evolution of the grand-potential landscape with chemical potential for relatively low temperature ($\beta = 0.89$) is shown in Fig. 3.4b. In two bottom panels ($\mu = -4.147$ and -4.146), the low-density minimum (marked by solid square) is the deepest one and system mainly has the density corresponding to this minimum. One (for $\mu = -4.147$) and two (for $\mu = -4.146$) higher-density minima marked by open squares are weakly occupied in equilibrium. They are separated by a large barrier (not fully shown) from the low-density minimum and cannot be detected by the single-flip MC simulations due to large height of this barrier as compared to temperature. For $\mu = -4.145$, the low- and intermediate-density minima become of approximately the same depth (they both marked by solid squares) and the system, in equilibrium, can be in one of them with approximately equal probability but, again, this cannot be detected by the single-flip MC simulations due to high barrier between the minima. In the relatively narrow range of chemical potential, $-4.145 \leqslant \mu \leqslant -4.144$, the most likely fluid density is related to the middle minimum of $\omega(\rho)$ (see the inset in Fig. 3.4a) which becomes approximately of the same depth as the high-density minimum at $\mu \simeq -4.144$ (both minima are marked by solid squares in Fig. 3.4b). For greater values of chemical potential, $\mu \geqslant -4.144$, the system in thermal equilibrium is in the high-density

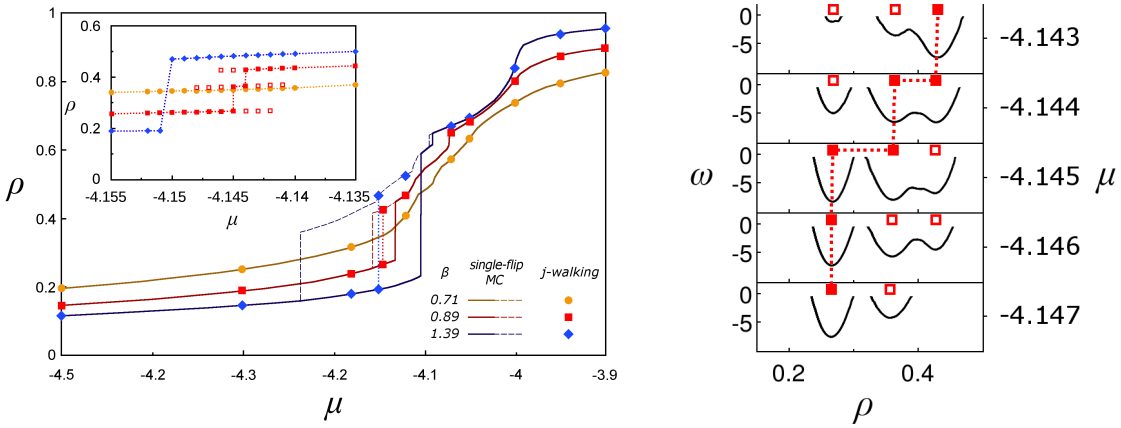

Fig. 3.4 (a) The dependence of the mean fluid density ρ in structural model of aerogel *versus* chemical potential μ for several temperatures as marked in the figure. The results of the single-flip conventional kMC simulations for adsorption (increasing μ) and desorption (decreasing μ) are shown by solid and dashed lines, respectively. The symbols refer to results of the JW MC simulations. For equilibrium isotherm at $\beta = 0.71$, the solid and dashed lines coincide with each other and with the solid circles. The inset shows magnification of the isotherms obtained by the JW MC simulations for the values of μ where the discontinuous transitions occur at low temperatures. The solid (open) symbols correspond to the deepest minimum (other minima) in the grand-potential landscape, $\omega(\rho)$. (b) Grand-potential landscapes, $\omega(\rho)$, for $\beta = 0.89$ and several values of μ specified on the right from each panel. The solid and open squares mark the deepest and other minima of $\omega(\rho)$, respectively, in the same way as in (a).

state (see the top panel in Fig. 3.4b where the high-density minimum is marked by solid square and two other states with lower density by open squares).

In contrast to the JW MC simulations, both experimental studies [192, 97] and single-flip MC simulations [193, 194] are not able to access the equilibrium for sufficiently low temperatures due to the presence of the hysteresis effect. Indeed, the results of the single-flip kMC simulations for adsorption (solid lines in Fig. 3.4a) and desorption (dashed lines) for the same aerogel model exhibit hysteresis (area between solid and dashed lines) for $\beta = 0.89$ and $\beta = 1.39$, although both adsorption and desorption curves coincide producing the equilibrium isotherm (solid line marked by circles representing the JW MC data) for $\beta = 0.71$. Therefore, obtaining the equilibrium isotherms for sorption in aerogel and other porous materials is often problematic, and requires techniques such as mean-field scanning curves [71]. However, despite those efforts conclusive answers have not yet been obtained to even the key questions such as about the existence of a discontinuous phase transition in aerogel samples of different

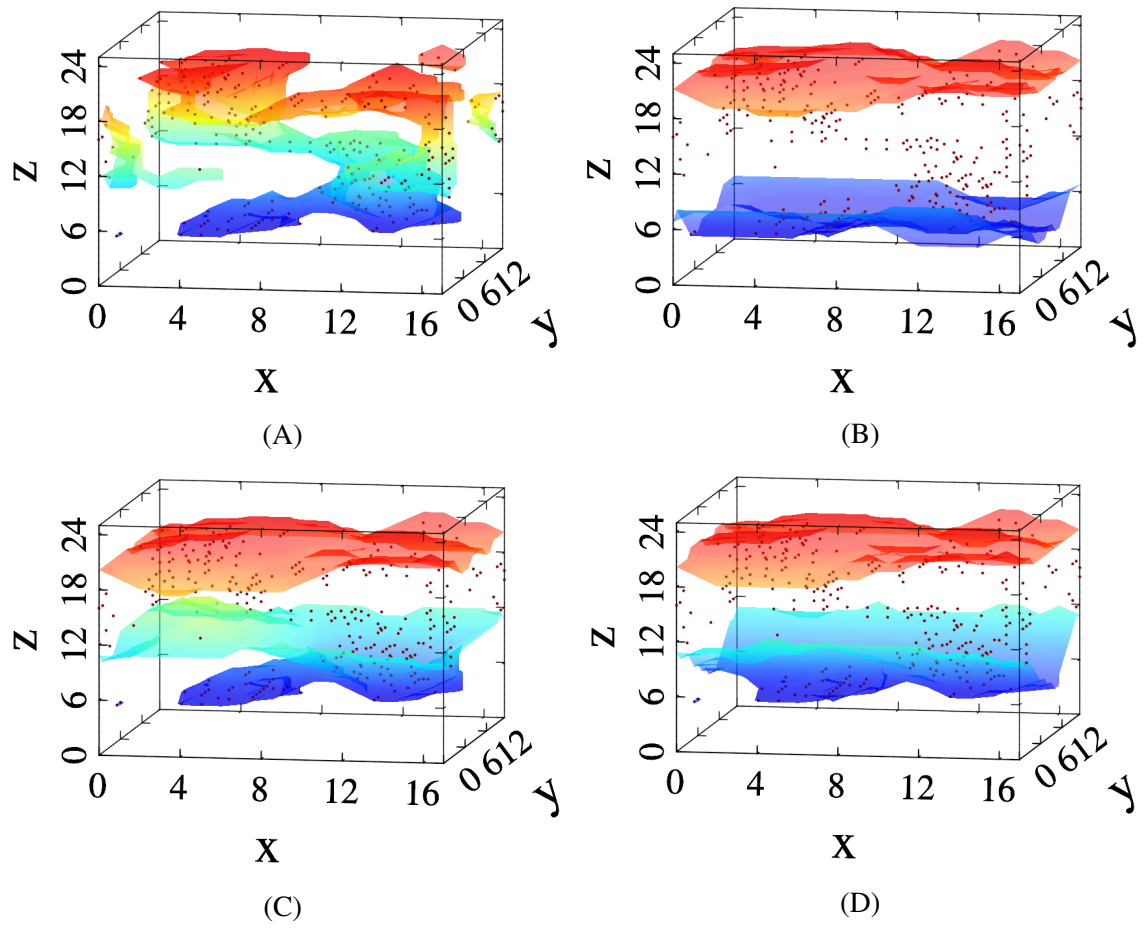

Fig. 3.5 Fluid configurations for $\mu = -4.18$ (upper panel) and $\mu = -4.145$ (lower panel) labelled by A, B and C, D in Fig. 3.3, respectively. The brown dots mark the matrix particles and the surfaces, coloured according to their z coordinate, represent the interfaces between the fluid and gas phases in the system.

porosity [97]. The JW algorithm is designed to avoid hysteresis and thus can be an invaluable tool in detecting phase transitions and resolving such problems. However a detailed analysis of equilibrium sorption isotherms in aerogel samples of different porosity is outside the scope of this work.

Fluid configurations

Figs. 3.5A-3.5D illustrate distributions of fluid in aerogel sample for the states characterised by different fluid densities marked by arrows A, B, C and D, respectively, in Fig. 3.3. The coloured surfaces show the boundaries between pore sites occupied and not occupied by fluid. The boundary layers are coloured according to their position in z-direction from blue (small z) to red (large z). The brown dots represent the matrix sites. Fig. 3.5A corresponds

to the low-density minimum of $\tilde{\mathscr{H}}_{\min}(\rho)$, e.g. at $\rho \simeq 0.15$ for $\mu = -4.18$, when the fluid occupies the sites neighbouring the matrix sites. Fig. 3.5B describes the state corresponding to the high-density minimum of $\tilde{\mathscr{H}}_{\min}(\rho)$. For this state, the aerogel sample (for $5 \leq z \leq 20$) is filled by fluid, and only the added bulk layers (with $1 \leq z \leq 4$ and $21 \leq z \leq 25$) are left unoccupied, i.e fluid occupies the space between two approximately horizontal surfaces.

Figs. 3.5C-3.5D illustrate the split, for intermediate values of μ around $\mu \simeq -4.145$, of high-density peak into two peaks which can be explained as follows. For this particular aerogel sample, the region of space with $z \simeq 5$ mainly populated by matrix sites is surrounded by a relatively empty space. This structural feature leads to existence of two fluid configurations of different density shown in Figs. 3.5C-3.5D and corresponding to two most right minima in the grand-potential landscape (solid line with open squares in Fig. 3.3). The fluid configuration of lower density with $\rho \simeq 0.35$ consists of pocket of pore sites occupied by fluid surrounding matrix sites near $z \simeq 5$ and separated by unoccupied pore space from the rest of aerogel sample occupied by fluid (see Fig. 3.5C). The fluid configuration of higher density with $\rho \simeq 0.45$ is similar to that with $\rho \simeq 0.35$ except that the empty pore sites separating the fluid pocket at $z \simeq 5$ from the rest of the fluid, in this case, are also occupied by the fluid (see Fig. 3.5D).

3.4.3 Ergodicity measures

In order to test the updated JW algorithm we studied sorption of fluid in a bulk system (without matrix sites) modelled by small bcc lattice consisting of $2 \times 2 \times 3$ (total number of sites, $N = 2 \times 12 = 24$) primitive unit cells with periodic boundary conditions. The bulk system is known to exhibit the first-order phase transition at sufficiently low temperatures. Therefore, significant slowing down of dynamics around the transition point is expected and thus all advantages of the JW algorithm as compared to the single-flip methods can be observed. Moreover, for a system of such size, it is possible to numerically calculate the grand partition function and other thermodynamic quantities exactly which can serve for testing and validation of MC algorithms.

By applying the exact state-enumeration procedure developed in Sec. 2.2, we were able to obtain the exact equilibrium occupation probabilities for all macrostates in the toy model. The results are presented in Fig. 3.6. Similarly to the zero-field Ising model (see Fig. 3.1), the most degenerate states are concentrated in the middle of the projected state space (i.e. $\rho \simeq 1/2$) but if the temperature is low enough then the occupation probabilities are higher for the states in the "corners" characterised by low energies. Exactly this situation is shown in Fig. 3.6 when for $\beta = 1$ the most probable states are either fully empty or fully occupied system.

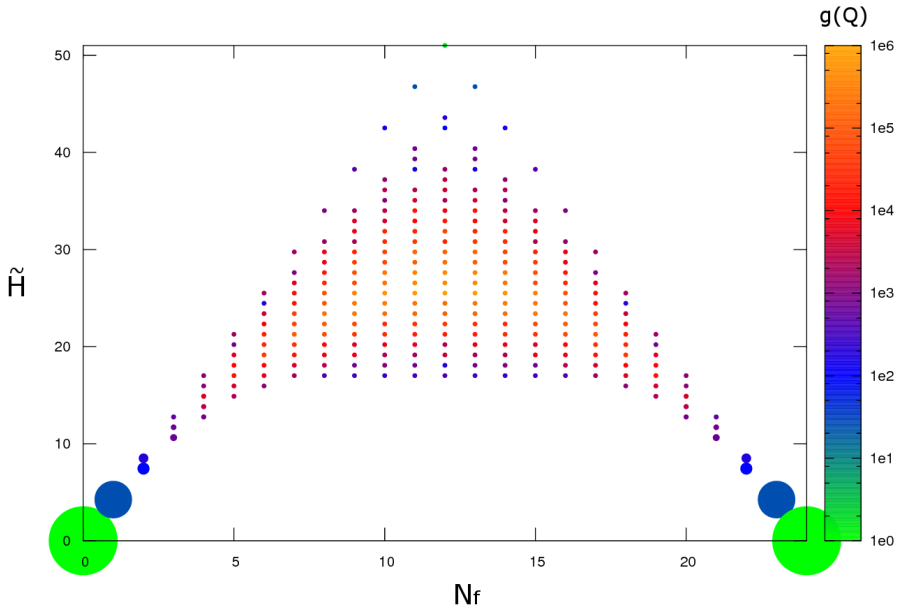

Fig. 3.6 The state space, i.e. $\tilde{\mathscr{H}}$ *vs* number of sites occupied by fluid, $N_{\rm f}$, of the $2\times2\times3$ bcc lattice containing 24 sites with periodic boundary conditions for $\mu=-4$. The colour of the data points reflects the values of the macrostate degeneracies, $g(Q)$, according to the colour scheme shown by the colour bar. The size of the data points proportional to the occupation probabilities of the macrostates for $\beta=1$.

The exact results for occupation probabilities in small models can serve for testing the MC algorithms [195–197]. We tested the updated JW algorithm by calculating the ergodicity measure, $\chi^2(\beta,S)$, given by the following expression [171],

$$\chi^2(\beta,S)=\sum_Q\left[P(Q(\rho,\tilde{\mathscr{H}}),S)-P_{\rm exact}(Q(\rho,\tilde{\mathscr{H}}))\right]^2\,, \tag{3.14}$$

where $\rho=N_{\rm f}/N$, $P(Q(\rho,\tilde{\mathscr{H}}),S)$ is the probability to visit macrostate Q by JW for finite MC process of S kMC steps and $P_{\rm exact}(Q(\rho,\tilde{\mathscr{H}}))$ is the exact thermodynamic value for such probability. The ergodicity measure calculated at different temperatures is shown in Fig. 3.7 both for updated JW and single-flip processes. The small values of $\chi^2(\beta,S)$ being characteristic of ergodic behaviour are seen for both single-flip and JW processes at high temperatures. However, the single-flip dynamics ceases to be ergodic for $\beta\gtrsim2$ while the JW process still remains to be ergodic. The ergodicity measure depends on the simulation time, i.e on S, and in the ergodic regime it decays $\propto S^{-1}$ [22] as we checked for the JW process (not shown).

The inverse-time decay of ergodicity measures is a general feature following from "diffusive" nature of all contributions in correlation functions [22] when a system is in ergodic

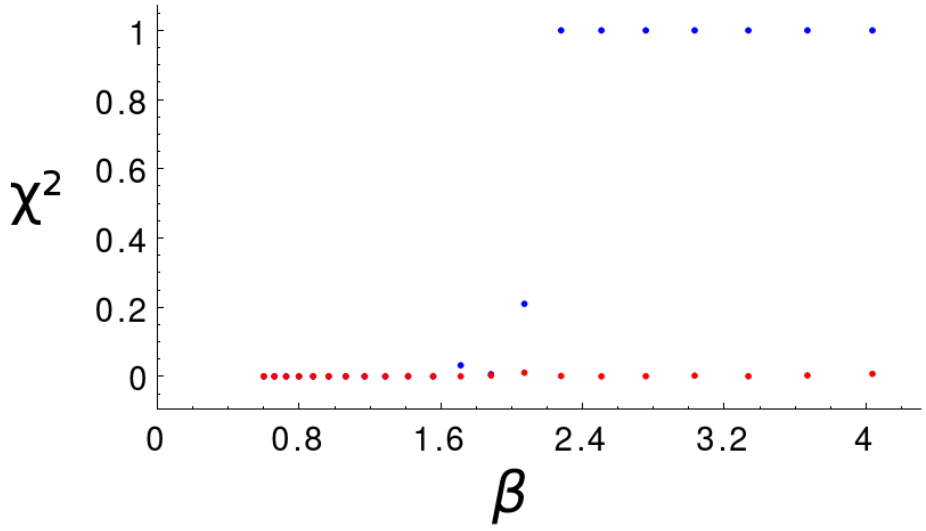

Fig. 3.7 Dependence of χ^2 *vs* β for the toy model system described in the text. The results of JW simulation and single-flip kMC are shown in red and blue, respectively.

regime. In order to reveal such a dependence for our toy model, we used an alternative and more general (not requiring exact solution) ergodicity metric d_k introduced in Ref. [22]. Let M to be a number of independent executions of simulation that produce N-dimensional density sequences, $\{\rho_i^m(n)\}$, with m numbering the sequence, $m = 1,\ldots,M$, and n counting the lattice sites, $n = 1,\ldots,N$. The ergodicity metric d_k at the k-th MC step is defined as,

$$d_k = \frac{2}{M(M-1)} \sum_{i=2}^{M} \sum_{j}^{i-1} \sum_{n=1}^{N} \left[\overline{\rho_k^i}(n) - \overline{\rho_k^j}(n)\right]^2 , \tag{3.15}$$

where $\overline{\rho_k^i}(n)$ is the time average fluid density at the n-th lattice site up until k-th MC step during the i-th simulation, i.e.

$$\overline{\rho_k^m}(n) = \sum_{i=0}^{k} \rho_i^m(n) \, . \tag{3.16}$$

We computed d_k metric for the toy model at different temperatures both for single-flip and JW processes (see Fig. 3.8). As follows from Fig. 3.8, the JW process is ergodic for all studied temperatures (cf. the slope of the thick blue line match with the red dotted line representing the k^{-1} dependence). In contrast, the single-flip processes (see the thin solid lines in Fig. 3.8) loose ergodic behaviour at low temperatures.

It should be mentioned that some systems exhibit weak-ergodicity breaking [140] which is different from quasi-nonergodic behaviour studied here and for which a single-trajectory time averaging does not necessarily coincide with the ensemble averaging even in the limit of

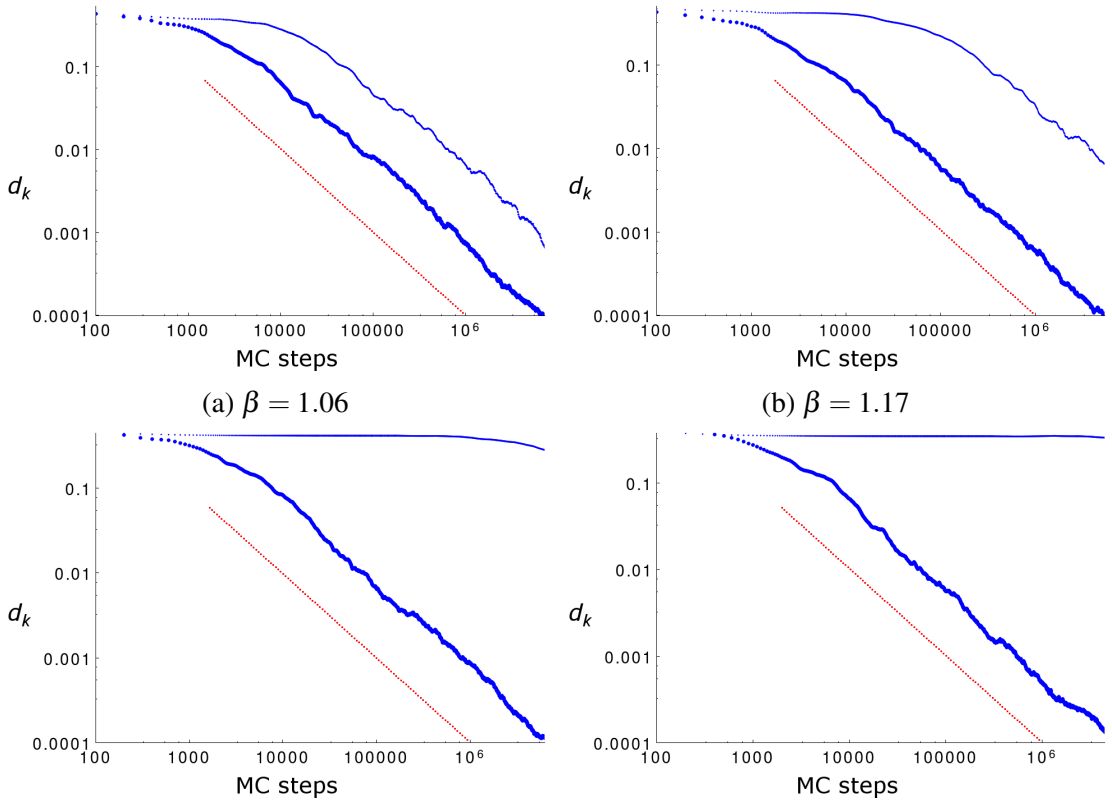

(a) $\beta = 1.06$ (b) $\beta = 1.17$

(c) $\beta = 1.56$ (d) $\beta = 1.71$

Fig. 3.8 Ergodic measure d_k *vs* the number of kMC steps k (the length of the simulation run), computed for the toy model at different temperatures β. The results for the JW and single-flip kMC simulations are shown in thick and thin blue lines respectively. As a guide for an eye $d_k \propto k^{-1}$ line is displayed in red. $M = 100$.

infinite observation time. This is a typical feature of diffusion models (thermal or athermal) with scale-free distribution of waiting times (with divergent mean sojourn time). In this situation, the limiting (infinite time) distribution for state occupation probabilities can differ from the Boltzmann distribution [198] and MC methods (including JW) might be useful for its sampling and studying non-ergodic dynamics.

3.4.4 Simulation parameters and correlation time

The updated JW simulations for Ising and lattice-gas models were performed with parameter S being in the range $S \in [2.2, 2.6] \times 10^9$ (the value of S depends on temperature and is measured in number of kMC steps), and $R = 7 \times 10^4$. The geometric sequence of simulation temperatures, $T_{i+1}^{-1} = kT_i^{-1}$ (where $k = 1.007$), with $T_1^{-1} = 0.5$ and $T_1^{-1} = 0.6$ was used for the Ising and lattice-gas models, respectively.

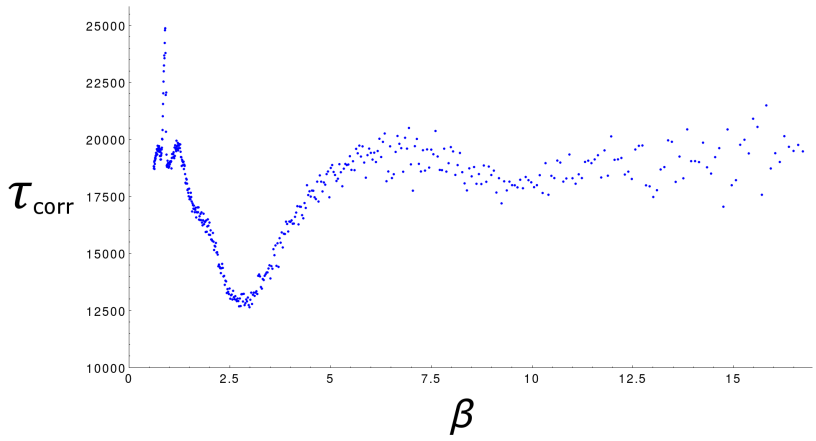

Fig. 3.9 Integrated correlation time τ_{corr} measured in kMC steps (vertical axis) for a range of inverse temperatures, β, obtained during the JW simulations of aerogel sample at $\mu = -4.145$.

In order to justify the choice of simulation parameters, we checked that the length of the JW runs was enough to avoid correlations between the states in the time series of MC chain. This was done by calculating the integrated correlation time τ_{corr} and demonstrating that $\tau_{\text{corr}} \ll S$. This quantity describes the time scale, for which the MC simulated system still holds some memory of its previous states [4]. For a time sequence of fluid densities ρ_i, that the system had at the i-th time step, the value of τ_{corr} is defined as,

$$\tau_{\text{corr}} = \sum_{k=0}^{S} \frac{c_\rho(k)}{c_\rho(0)} \,, \tag{3.17}$$

where $c_\rho(k) = \langle \rho_i \rho_{i+k} \rangle - \langle \rho_i \rangle^2$ is the time autocorrelation function of fluid density with separation interval k. The integrated correlation time for fluid density ρ, measured in the number of kMC steps, was obtained for the JW simulations of the aerogel sample at $\mu = -4.145$ and is shown in Fig. 3.9. The peak in τ_{corr} at $\beta \approx 0.9$ marks the critical region, when the multiple metastable states were being explored (See Fig. 3.3b) causing large density fluctuations and slow relaxation. The simulation time scale S exceeds the integrated correlation times τ_{corr} by several orders of magnitude for any range of temperatures.

All simulations have been performed using one core of Intel Xeon X5680 CPU unit. The maximum amount of RAM required by the JW simulation was below 20GB. Each stage corresponding to a temperature in $\{T_\alpha\}$ ran for approximately 20 minutes.

3.5 Conclusions

To conclude, in this Chapter, we have revisited the JW algorithm [163] which has been originally designed to tackle quasi-nonergodicity problem, a known problem in dynamics of interacting particles. However, due to demanding (at the time of its creation) memory requirements the algorithm lost in competition with other approaches such as REM. Bearing in mind significant improvements in RAM available in contemporary computers we have updated the JW algorithm, tested it for the Ising model and demonstrated its performance for sorption in lattice-gas models (a toy model and aerogel model). The main results of this Chapter can be summarised as follows,

(i) the development of the single-walker JW MC algorithm aiming to alleviate the quasi-nonergodicity problem characteristic of complex systems at low temperatures;

(ii) testing the algorithm for the 3d Ising model;

(iii) testing the JW MC algorithm on a small lattice-gas system and application of the algorithm to the lattice-gas model of silica aerogel for studying sorption and equilibrium behaviour at low temperatures.

The efficiency of the JW algorithm stems from the fact that only a single non-local (multiple-flip) jump is needed to cross the free-energy barriers as opposed to multiple replica exchanges in the case of REM. Even though each particular replica within REM is allowed to cross energy barriers by configuration exchange with a replica at a slightly different temperature, REM overall simply swaps the two replicas, and computer resources continue to be used for exploring the same free-energy minima. Therefore, if the minimum is deep enough the algorithm based on REM can spend vast amount of time randomly diffusing in the temperature dimension while exploring the same local free-energy minimum. In contrast, as temperature decreases the JW algorithm naturally drives the system to escape from the energetically unfavourable local minima and therefore, the system is maintained in the equilibrium throughout the entire duration of the simulation.

The principal differences between the two methods, the JW and REM, suggest the types of problems that can be tackled by each approach. Since the REM replicas tend to explore the local minima for extended periods of time throughout the simulation, the method is well suited for tasks such as mapping of the entire free-energy landscape, weight factor estimation for multicanonical algorithms and, more generally, for investigating the properties of the whole state space of the system. The JW algorithm is more appropriate for studying equilibrium behaviour of the system. As temperature decreases, the system simulated by the JW algorithm naturally approaches the most favourable explored minimum, thus this

method is well suited to search for the global free-energy minimum and investigate system's behaviour at the lowest temperatures. Since significantly less CPU resources is required by the JW algorithm as discussed in Sec. 3.3, it can either enable to study systems that have previously been prohibited by the lack of computing power, or to allow a more efficient investigation of the parameter space, provided that the analysis is mainly focused on the equilibrium behaviour of the system. In particular, the JW MC algorithm can be used on a single computing thread to investigate equilibrium properties of a system, especially near discontinuous phase transitions where the standard single-flip MC simulations are inefficient due to critical slowing down of dynamics and REM is too expensive in terms of computational resources.

Successful application of the updated JW algorithm to lattice-spin and lattice-gas models (see Sec. 3.4) might promise its relevance to other models facing quasi-nonergodic behaviour, such as spin-facilitated models for glassy dynamics [7] and lattice models for proteins [180]. Another possible area for application of the updated JW method could be in studying the systems exhibiting weak ergodicity breaking [140] with non-Boltzmann limiting distributions [198].

The description of the improved JW MC method was intentionally given as simple and general as possible. Therefore, despite certain advantages in its current form, the algorithm leaves much room for further optimisation and improvements. For example, it is possible to employ kMC-like (as oppose to MC-like) protocol for picking the state from the j-distribution, by adjusting the jump probabilities accordingly and removing the possibility of jump rejection entirely. Further potential computational improvements are discussed in Ch. 5. Moreover, due to key similarities with REM, numerous developments that studies of REM provided throughout the last two decades are readily available to implement for the updated JW algorithm (e.g. the temperature set $\{T_\alpha\}$ and jump frequency optimisation as well as combinations with Wang-Landau, multicanonical or simulated tempering algorithms).

Chapter 4

Sorption in Vycor glass systems

4.1 Introduction

Due to its well known structure and pore characteristics, Vycor glass has been extensively studied [199–202] as one of the prototype materials for the disordered porous media (see Secs. 1.1.3 and 1.2.2 for more details). The hysteresis loops in Vycor glass have been shown [200] to be of H2-type according to the IUPAC classification, which indicates that sorption phenomena in such systems cannot be explained by the independent domain theory [203]. Sorption behaviour in Vycor has been further examined by measuring adsorption and desorption scanning curves, which confirmed [120, 124, 201] that such systems display characteristics of interconnected pore structure. That said, the nature of the phase transition has not been fully understood and is still under active investigation [12, 204–206].

Previous research in the area has been primarily concerned with out-of-equilibrium isotherms, e.g. the debate regarding role of pore blocking and cavitation mechanisms (see Sec. 1.2.2 for details) during the desorption transition in Vycor [201, 204]. The nature of avalanches during the adsorption transition has also been examined [199]. The equilibrium isotherms, however, are usually obfuscated by the hysteresis phenomena in both experimental and theoretical studies. In this Chapter, we use JW MC algorithm to study the equilibrium behaviour in Vycor-glass systems. We investigate the nature of condensation and evaporation events that comprise the equilibrium isotherms. Such events may be localised within a single pore, involve regions that are larger than the material correlation length or even span the entire system. By examining the structure of the isotherms and the fluid configurations throughout the transition, we aim to address the questions about the existence of the first-order phase transition in Vycor-glass systems.

In Sec. 4.2 the equilibrium isotherms for Vycor-glass systems are obtained using JW MC algorithm. The behaviour is analysed by inspecting the isotherms and the corresponding

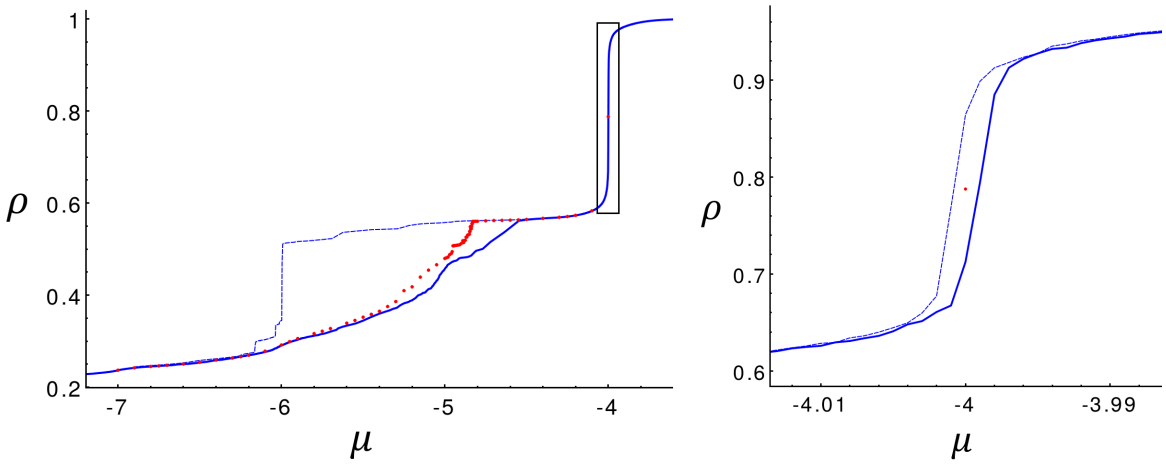

(a) (b)

Fig. 4.1 Sorption isotherms in V_{30} model at $\beta = 10$ and $y = 2$. (a) Adsorption (thick solid line) and desorption (thin dashed line) isotherms obtained by single-flip kMC method shown in blue. Equilibrium isotherm produced by JW MC algorithm are displayed in red. (b) Magnified region of the main diagram (as indicated by the black rectangle in (a)) at $\mu \simeq -4$, when the transition occurs in the bulk slabs as discussed in the text.

fluid configurations. The lower projected-state-space boundaries for two different values of wettability parameter are found in Sec. 4.3. Conclusions are given in Sec. 4.4.

4.2 Equilibrium isotherms

In this Section, we use JW MC method to study sorption phenomena in Vycor-glass systems. In particular, we obtain and analyse equilibrium isotherms in such systems aiming to develop a detailed understanding about the nature of the phenomena.

The cubic models of Vycor glass with porosity $\phi = 0.28$ and sizes $L = 30$ (sample V_{30}) and $L = 40$ (sample V_{40}) were produced as described in Sec. 1.1.3. Additional empty space slabs were added in z-direction, while the periodic boundary conditions were imposed in x- and y-directions as in the case of aerogel discussed in Sec. 3.4.2. As previously, the models are constructed on bcc lattice, thus resulting in the systems containing N pore cells with $N = 2 \times 30 \times 30 \times 38 \times \phi \simeq 2 \times 10^4$ and $N = 2 \times 40 \times 40 \times 48 \times \phi \simeq 4 \times 10^4$ for V_{30} and V_{40} systems, respectively.

As discussed in Sec. 1.3, the wettability y is parameter of the model. To study a particular system the suitable value of y depends on the adsorbate and on the distance corresponding to the lattice spacing of the structural model. It has been shown [59] that sorption isotherms produced by the models used in this work provide the best fit to the experimental results [206]

for adsorption of CH_2Br_2 in Vycor glass with wettability parameter $y \simeq 1.3$. In order to cover a wider range of y values the following analysis is performed for $y = 1$ and $y = 2$.

4.2.1 Sorption in Vycor glass with $y = 2.0$

Fig. 4.1a displays the results of single-flip kMC and JW MC simulations of V_{30} model system with $y = 2$ at $\beta = 10$. The equilibrium isotherm obtained using JW MC method is observed inside the hysteresis loop, between the adsorption and desorption isotherms produced by the single-flip kMC simulation. The steep jump in fluid density from $\rho \simeq 0.6$ to $\rho \simeq 1.0$ at $\mu \simeq -4$ corresponds solely to the fluid-vapour transition in the empty slabs that surround the system in z-coordinate. For the values of $\mu < -4$, the part of the system occupied by Vycor ($z \in [4,34]$) is completely filled with fluid, while the empty slabs ($z \in [0,4)$ and $z \in (34,38]$) are in vapour phase (such fluid configuration is displayed in panel (H) of Fig. 4.3, with empty space being above red and below blue surfaces). In contrast, for $\mu > -4$ the entire simulation box is occupied by fluid (see fluid configuration in panel (I) of Fig. 4.3, with fluid being between red and blue surfaces). Note, that the hysteresis loop for the single-flip kMC isotherms displayed in Fig. 4.1b is remarkably narrow. This is a consequence of two factors. Firstly, the developed kMC implementation that is described in detail in Sec. 5.2 allows for very efficient equilibration. Secondly, filling or emptying of the slabs is equivalent to the first-order phase transition in the bulk, hence, the grand-potential landscape between the two states is not rugged (i.e. does not contain multiple metastable states). Therefore, the equilibration in this region of the state space is significantly easier than in the case of the transitions between the different fluid occupations within the disordered porous media. In the following text, we are mainly interested in the behaviour of the system below $\mu = -4.0$, thus, we present only partial $\mu - \rho$ diagrams such as shown in Fig. 4.2a.

As discussed in Ch. 2, for any system with some matrix cells and $y > 0.5$, the first cells that turn into fluid at low values of μ are those neighbouring the matrix. The developed fluid layer thickens as μ increases, eventually filling small voids in the Vycor structure starting from the narrow crevices associated with the high curvature of the fluid-vapour interface. This initial process ($\mu \lesssim -5.5$) is governed by the moving menisci of the fluid in relatively narrow channels toward the more "open" pore spaces. Such transitions are continuous and reversible in the pores where the remaining vapour has physical access to the rest of the pore space. The other "phase" of the behaviour ($-5.5 \lesssim \mu \lesssim -4.5$) is examined further in Fig. 4.2. Detailed analysis of the structure of the equilibrium sorption isotherm reveals a ladder-like behaviour displayed in Fig. 4.2b. The states labelled (A)–(H) correspond to the fluid configurations in the respective panels of Fig. 4.3. In this regime, the remaining vapour is concentrated in large isolated pockets of pore space with relatively low curvature

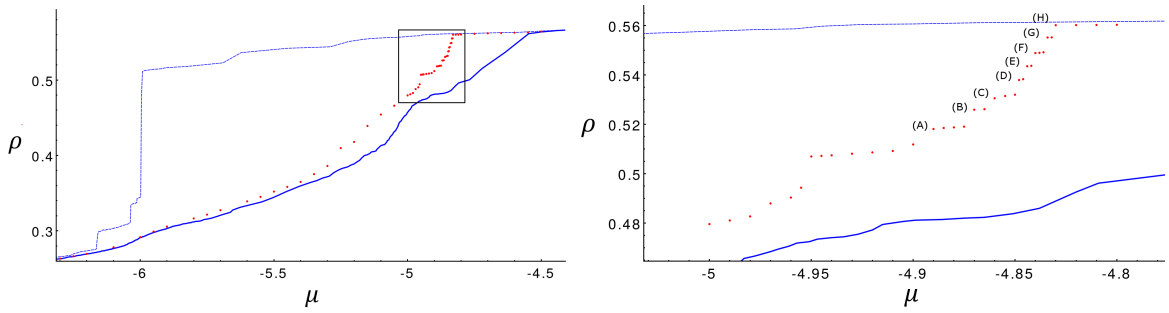

(a) (b)

Fig. 4.2 (a) The main hysteresis loop in the V_{30} system with $y = 2$. (b) Further magnification of the diagram indicated by the black rectangle in (a). The states labelled by the black letters (A)–(H) correspond to the fluid configurations displayed in Fig. 4.3. The same line styles are used as in Fig. 4.1.

as displayed in panel (A) of Fig. 4.3. Such "bubbles" are filled with fluid independently from each other producing the sudden jumps in density (avalanches between the equilibrium states) and forming the mentioned ladder-like structure of the isotherm as displayed in Fig. 4.2b. Each avalanche is associated with a separate spatially isolated transition, and it is possible to track them one by one as shown in Fig. 4.3. The avalanches are microscopic, since their size depends solely on the microstructure of the Vycor model, and not on the size of the system. In between the jumps in density, the system continues to follow the current state, as can be seen by the roughly horizontally oriented groups of data points in Fig. 4.2b.

An equivalent analysis performed on the larger V_{40} model system (not shown), confirms that the volumes affected by a single avalanche retain the same size, despite that the V_{40} system is over twice the volume of V_{30}. Detailed statistical analysis of the jumps in the equilibrium isotherm is outside of the scope of this work. However, the developed understanding of the sorption mechanism strongly indicates that the size of the region affected by a single avalanche is governed by the structural correlation length of the material. Thus, we conclude that such Vycor-glass systems do not exhibit first-order phase transition for $y = 2$.

4.2.2 Sorption in Vycor glass with $y = 1.0$

Two main factors influence the sorption behaviour (within the lattice-gas model) in porous materials, i.e. the structure of the material and value of wettability, y. It is known [12] that for $y = 0.5$ the bulk-sorption-like behaviour is recovered (if the pore space percolates through the system), since under such conditions the interaction energy per fluid volume of two neighbouring fluid cells is the same as that of a matrix cell and a fluid cell. Thus, the first-order phase transition definitely exists for $y = 0.5$ in Vycor glass with porosity $\phi = 0.28$.

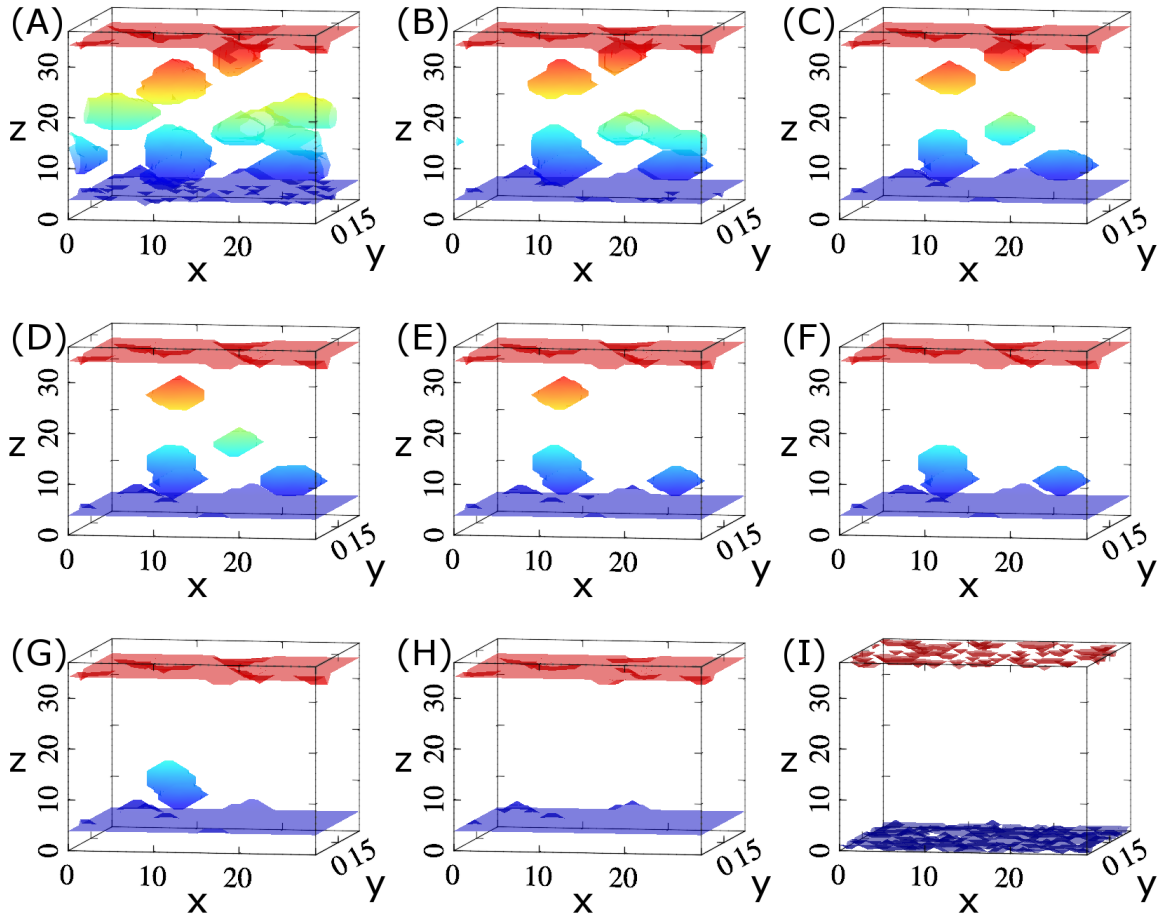

Fig. 4.3 Fluid configurations of V_{30} model system with $y = 2$ corresponding to the states labelled by (A)–(H) in Fig. 4.2. The configuration (I) represents the system fully occupied by the fluid at the values of $\mu > 4$, as discussed in the text. The "bubbles" (shown by different colour depending on their z-coordinate) in the interior region of the simulation box in panels (A)–(G) correspond to the pore space not occupied by fluid. The matrix cells are not displayed for clarity. The same colour-scheme for z-coordinate is used as that in Fig. 3.5.

Since in the previous Section we have established that Vycor systems with $y = 2$ do not exhibit the first-order phase transition, here, we aim to test the case of $y = 1$.

Fig. 4.4 presents a similar analysis performed on V_{30} model with $y = 1$ as that in the previous Section for $y = 2$. As before, the initial fluid cells appear for sufficiently low values of μ and they are located next to the matrix cells, i.e. on the matrix walls. However, further filling transitions are governed solely by the minimisation of the interface area between the regions of matrix/fluid and vapour, without the necessity for the fluid cells to cover the matrix surface entirely. This results in larger independent pockets of vapour remaining at the intermediate values of μ. Despite the differences, the ladder-like behaviour of the equilibrium

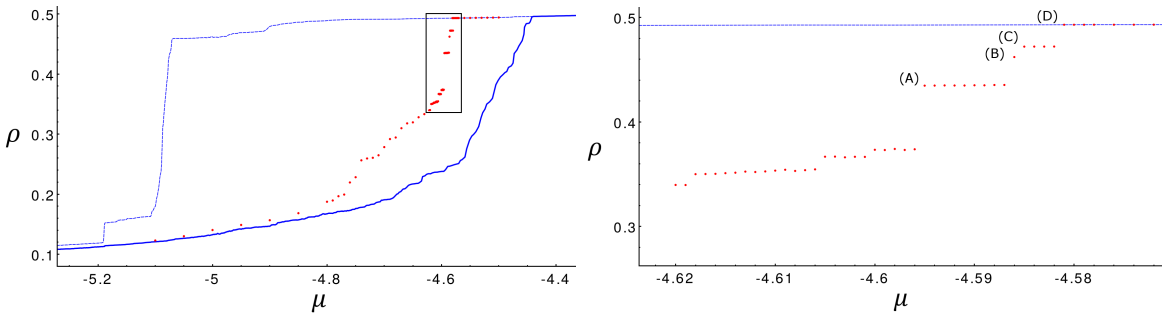

(a) (b)

Fig. 4.4 (a) The main hysteresis loop in the V_{30} system with $y = 1$. (b) Further magnification of the diagram indicated by the black rectangle in (a). The same line styles are used as in Fig. 4.1.

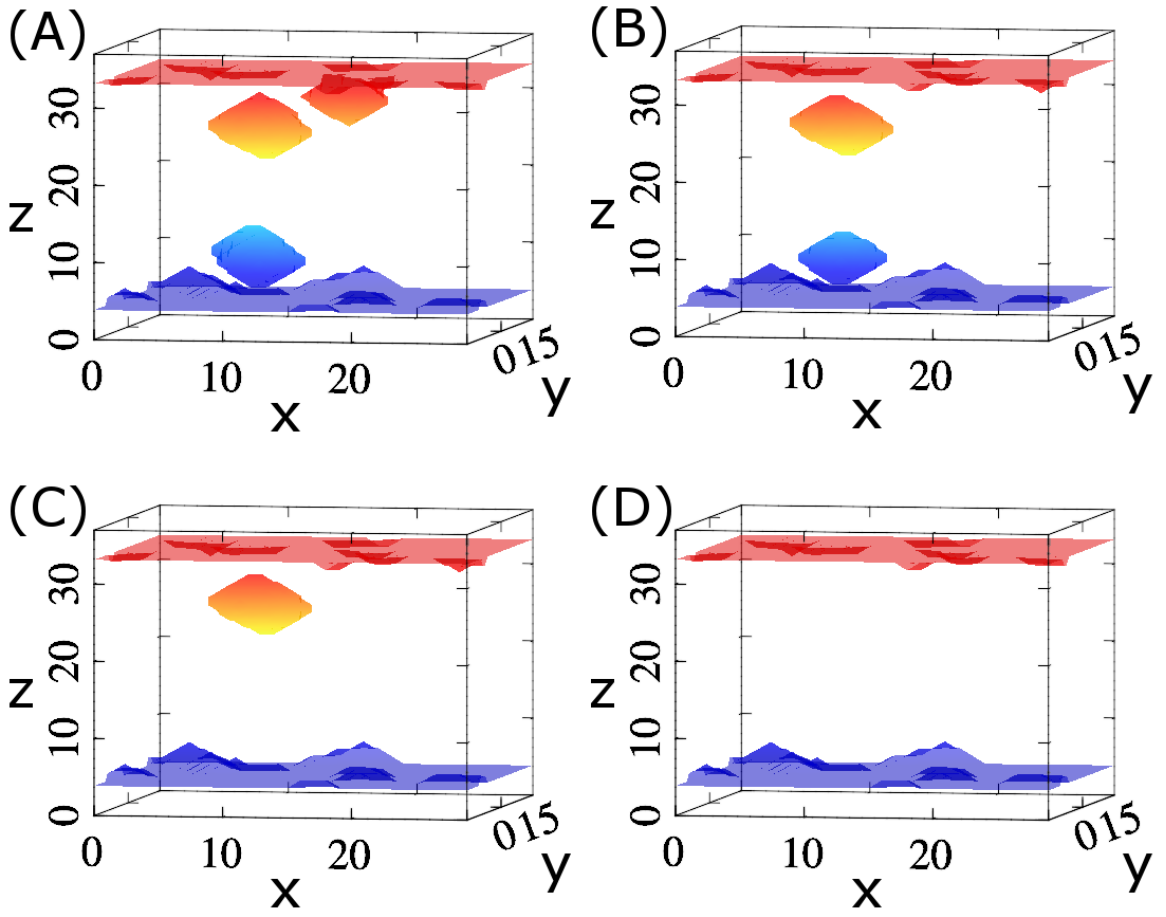

Fig. 4.5 Fluid configurations of V_{30} model system with $y = 1$ corresponding to the states labelled (A)–(D) in Fig. 4.4.

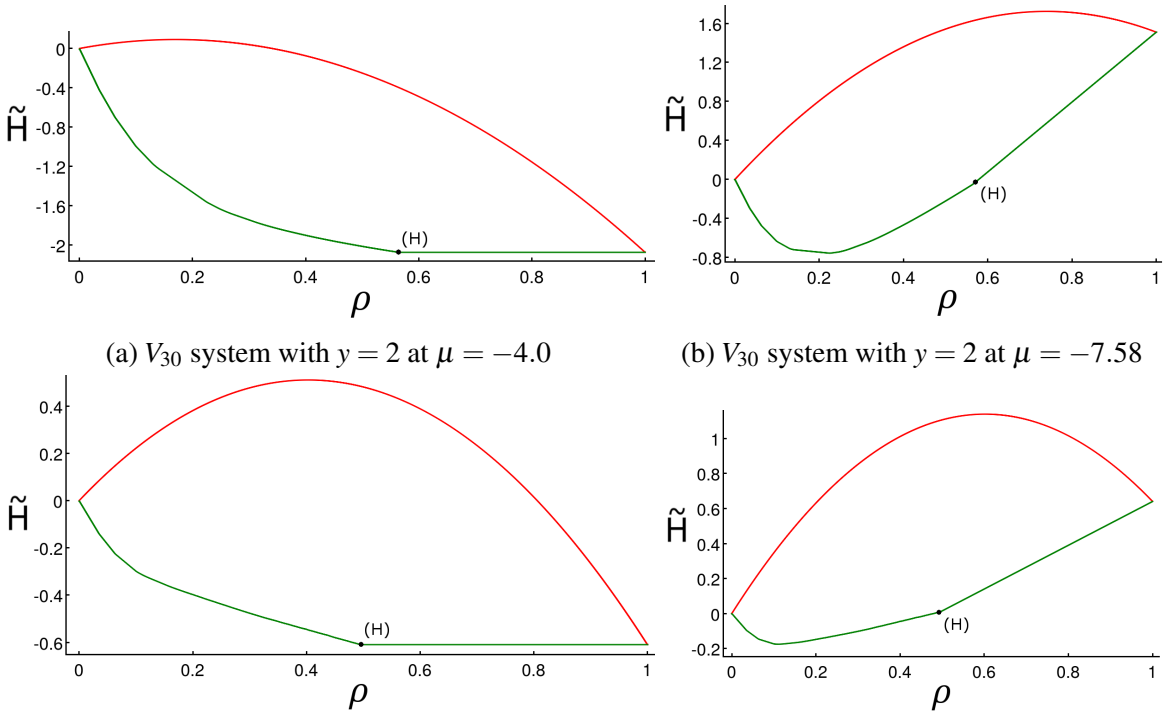

(a) V_{30} system with $y = 2$ at $\mu = -4.0$

(b) V_{30} system with $y = 2$ at $\mu = -7.58$

(c) V_{30} system with $y = 1$ at $\mu = -4.0$

(d) V_{30} system with $y = 1$ at $\mu = -5.25$

Fig. 4.6 The upper (in red) and lower (in green) projected-state-space boundaries for V_{30} model system obtained by the method described in Secs. 2.3.6-2.3.7. The point (H) marked on the projected-state-space diagram corresponds to the state of the system and the fluid configuration labelled by (H) in Figs. 4.2 and 4.3.

isotherm is retained. The fluid configurations displayed in Fig. 4.5 corresponding to the states labelled (A)–(D) in Fig. 4.4 confirm that the sorption is governed by the same processes as in the case of $y = 2$. Thus we conclude that Vycor systems do not exhibit the first-order phase transition neither with $y = 2$ nor with $y = 1$.

4.3 Lower projected-state-space boundary

In order to test the conclusions reached in the previous Section regarding the absence of the first-order phase transition in Vycor systems, here, we apply the projected-state-space boundary analysis developed in Ch. 2.

The upper and lower projected-state-space boundaries displayed in Figs. 4.6a-4.6b for the system V_{30} with wettability parameter $y = 2$ are found using the algorithms described in Secs. 2.3.6-2.3.7. The upper boundary is displayed for a visual reference to the entire projected state space of the system. The point marked by (H) on the lower projected-state-

space boundary corresponds to the state of the system in which the volume of the simulation box occupied by the Vycor glass structure is completely filled with fluid, but the empty slabs added in z-direction are occupied by vapour (such fluid configuration is depicted in panel (H) of Fig. 4.3). To the right of the point (H), the segment of the projected-state-space boundary is a straight line. At $\mu = -4.0$ this segment is exactly horizontal because, on average, each additional fluid cell in the bulk establishes 4 fluid-fluid interactions with its nearest neighbours (note, that in z-direction, instead of the periodic boundary conditions, an enforced vapour layer covers the surface of the simulation box). This also provides a qualitative explanation why the phase transition in this region of the space happens at $\mu = -4.0$. The rest of the lower projected-state-space boundary has a positive curvature along its entire length, as can be seen in Fig. 4.6b. The convex boundary indicates the absence of the first-order phase transition in the system as discussed at length in Ch. 2, thus supporting the findings of the previous section.

The projected-state-space boundaries for V_{30}-system with $y = 1$ are shown in Figs. 4.6c-4.6d. It is apparent that even though the curvature of the lower projected-state-space boundary is significantly reduced, it is still positive, thus confirming the results obtained in the previous Section.

4.4 Conclusions

As discussed in Ch. 2, the existence of a macroscopic phase transition depends on the level of disorder in the system. The disorder in lattice-gas systems is determined by the temperature, microstructure of the sample and wettability parameter. Vycor glass has a well defined porosity and structure, that has been reproduced in our models (see Sec. 1.1.3). We also know that the first-order phase transition does exist in bulk systems and thus can occur in systems with $y = 0.5$. Hence, at sufficiently low temperature, Vycor systems exhibit the first-order phase transition if the wettability parameter y is set below some upper critical value y_c [12]. In this Chapter, we investigated whether the values of y used in our models of Vycor are within such range.

JW MC algorithm was used to obtain equilibrium sorption isotherms in two models of Vycor glass, each for two different values of wettability parameter $y = 1$ and $y = 2$. We have examined the isotherms in detail by inspecting the fluid configurations that correspond to the equilibrium states of the system at different values of μ. We have also applied the analysis of the lower projected-state-space boundary developed in Ch. 2. The main findings can be summarised as follows,

(i) qualitative understanding of the equilibrium sorption behaviour in Vycor systems with $y = 1$ and $y = 2$ has been developed. At low values of μ, the fluid cells are mainly concentrated in the neighbourhood of the matrix surface. As μ is increased the fluid continuously fills increasingly larger voids of pore space in the Vycor structure. Such process continues until only the isolated, relatively smooth pockets of vapour remain in the system. Finally, such bubbles condensate one by one leading to a ladder-like shape of the equilibrium sorption isotherm;

(ii) the vapour pockets associated with the discontinuous segments of the isotherm are limited in size by the pore structure of Vycor and do not span volumes beyond a characteristic space scale of the pores in the system. In fact, at the studied values of y such pores effectively act as isolated pockets of bulk, thus implying that Vycor systems in given conditions do not exhibit first-order phase transition;

(iii) the lower projected state-space boundary for both values of wettability parameter is found to be convex, confirming the absence of the macroscopic phase transition in the studied systems.

The Vycor models used in this Chapter have been shown [59] to best fit the experimental sorption isotherms [206] for CH_2Br_2 with the wettability parameter set to $y = 1.3$. Since the absence of the first-order phase transition has been established in systems with $y = 1$ and $y = 2$, thus implying that $y_c < 1$ for our models of Vycor, we conclude that Vycor glass systems with the same adsorbate do not undergo macroscopic phase transition at any temperature.

Chapter 5

Interactive MC simulation engine

5.1 Introduction

In this Chapter, we discuss a unified set of computational tools, that comprise the interactive MC simulation engine developed to aid in our study of sorption phenomena in porous media. The original aims of the application were to provide a user with a more holistic view of the system in question. Most of the standard measurements used in the field, be it sorption isotherms, correlation time, mean interaction energy, ergodicity measures, etc., are designed to accurately answer a very particular question about the system. However, except when specifically searching for correlations, one rarely looks at multiple aspects of system behaviour at once. The simulation engine was designed and developed to serve this function, simultaneously allowing the user to interact with the model system in real-time.

The primary data streamed to the user were chosen to be (i) state-space visit histogram, (ii) 3d view of the fluid configuration in the simulation box and (iii) the standard $\mu - \rho$ diagram. User interaction with the system, which is continuously simulated by MC method, is achieved through manipulating the values of β and μ.

We have chosen to built the application within Qt software development framework because it provides sophisticated user interface tools. The macrostate visit histogram as well as the $\mu - \rho$ diagram were produced using QCustomPlot, which is a Qt C++ widget for plotting and data visualization. Finally, the 3d view of the fluid configurations was developed using MathGL, which is a cross-platform C++ library for high-quality scientific graphics. All mentioned tools are used under the GPL license.

In Sec. 5.2, several optimisations for the standard kMC algorithm are developed, that significantly improve the efficiency of the simulation engine. The main controls and visual outputs of the application are described in Sec. 5.4, and its implementation is discussed in

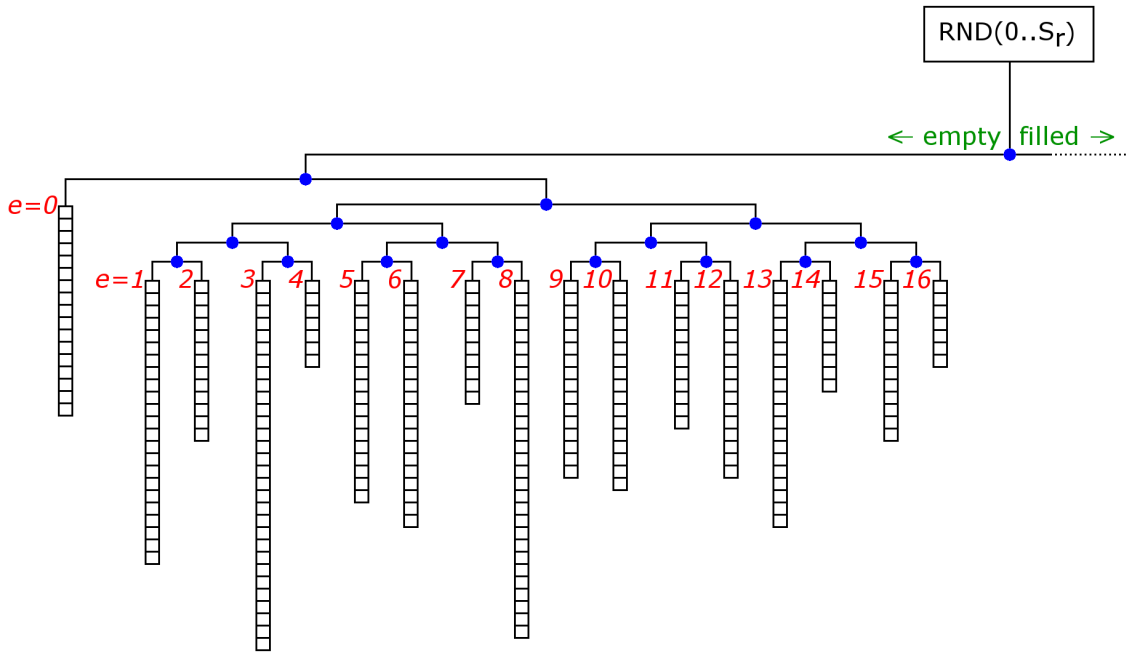

Fig. 5.1 The kMC implementation scheme for bcc lattice with $y = 2.0$. The next event (evaporation or condensation of some cell) is picked by, first, choosing what type of transition occurs (from filled to empty or *vice versa*), then choosing what the local environment e of the transforming cell is, and, finally, picking one of the cells in that environment. Blue dots represent the conditional statements, that need to be processed while choosing the next event. The entire operation requires only a single random number.

Sec. 5.5. Finally, several usage examples are provided in Sec. 5.6 and the concluding remarks given in Sec. 5.7.

5.2 Optimised Kinetic Monte Carlo implementation

At the core of the kinetic Monte Carlo (kMC) method is a task of choosing the next event out of a set of all possibilities, each with some known rate R, as it was described in Sec. 1.4.2. In our case, the possible set of events are the N transitions corresponding to the N cells in the system, that can turn from empty to filled or *vice versa*. The expressions for the rates $R\left(\{\tau_i\} \to \{\tau_j\}\right)$ of each transition were derived in Sec. 1.4.2 and depend on the current values of temperature, chemical potential, and how many fluid or matrix neighbours the particular cell has. The general problem of choosing an event out of N possibilities is of $\mathrm{O}(\ln(N))$ complexity. Indeed, the usual implementation is based on heap, i.e. tree-like data structure, which has all possible events with the corresponding rates represented as the

"child"-nodes at the very top of the tree. Every "parent" node of two children is represented by the sum of the two corresponding rates. The algorithm picks a random number in the interval from 0 to the sum of all rates S_r, and goes up the tree picking which event happens next. Once the probability of some event changes, the algorithm goes back down the tree structure adjusting each affected parent node value. Such operations are of $\mathrm{O}(\ln(N))$ complexity.

If all possible events had the same rate, i.e. $R = 1/N$, the procedure of choosing the next one would be reduced to $\mathrm{O}(1)$ complexity, simply by picking a random number from a discrete uniform distribution in the interval $(1,N)$. Although in our case, the probabilities of the events are not equal, under certain assumptions, the number of distinct rates can be finite. Therefore it is possible to construct a combined algorithm, which maintains $\mathrm{O}(1)$ complexity of the choosing procedure.

The probability of a cell to switch from empty to fluid or *vice versa* depends only on the external thermodynamic parameters and cells interactions with its nearest neighbours. Each site in bcc lattice has 8 neighbours which can be either empty, fluid or matrix. Hence, for the given values of β, μ and y (wettability), there is a finite number of distinct local environments (set of nearest neighbours and their kind) that a cell can be in, and thus there exists a finite number of corresponding probabilities (rates) for the transitions to happen. For example, if $y = 2.0$, any cell can have an absolute value of interaction energy with its local environment spanning from 0 to 16, corresponding to zero neighbouring matrix or fluid cells and 8 neighbouring matrix cells, respectively. It is also apparent that each value in the range from 0 to 16 can be attained by some combination of neighbouring matrix and fluid cells. Thus for $y = 2.0$, cells in such system have 17 distinct (in terms of interaction energy) local environments, and hence $2 \times 17 = 34$ different possible transition rates to fill or empty any cell. Similarly, for $y = 1.0$ there are only 18 distinct transition rates.

In order to take an advantage of such discrete and finite probability distribution, we construct a hybrid data structure shown in Fig. 5.1. All possible transitions are grouped by the current state (empty or filled) and the local environment e of the respective cell. Thus the rates of the transitions within the same environment are equal, and the sum of them is equal to the environment rate $R_e = R(e) \times n(e)$, where $n(e)$ is the number of the cells in an environment e, and $R(e)$ is the rate corresponding to one of such cells changing its state (see Eq. (1.6) derived in Sec. 1.4.1). Blue dots in Fig. 5.1 mark the parent nodes in the data structure. Each parent node i contains the value V_i equal to the sum of its two child node values. Thus the 8 lowest-level parent nodes, seen in the Fig. 5.1, have values equal to the sums $R_1 + R_2$, $R_3 + R_4, \ldots R_{15} + R_{16}$. The uppermost parent node S_r is therefore the sum of all N transition rates.

Each environment contains an array with the id's of the cells, that belong to the corresponding environment. In addition to retrieving any element, the arrays support operations of adding and removing only the last element. All three operations are of $\mathrm{O}(1)$ complexity. Within such data structure, a single transition is performed as follows.

(i) Generate a random number r from a uniform real distribution in the interval $(0; S_\mathrm{r})$.

(ii) Go down the data structure shown in Fig. 5.1 by comparing r with the value V_i of the respective parent node in the heap structure (parent nodes marked by blue dots in Fig. 5.1). If $r < V_i$, then proceed towards the left child branch of the current node, otherwise set $r = r - V_i$ and proceed towards the right child branch.

(iii) Once some environment e is reached, choose the transition from an array of cells in this environment. The array index of the chosen cell is given by $\lfloor \frac{r}{R(e)} \rfloor$.

(iv) Change the state of the chosen cell, and move it to the corresponding environment on the other side of the data structure (from empty to filled, or *vice versa*). The move is executed by removing the respective element from the array and replacing it by the current last element of that array, reducing it by one element in the process. Such operation is allowed since the order of the elements denoting possible transitions within an environment is not important. Finally, add the chosen cell to the end of the array corresponding to its new environment.

(v) Similarly, alter the environments of the affected nearest neighbours of the chosen cell. If the emptying transition has been executed, all nearest neighbours shift their environment towards the lower value of e, and *vice versa*.

(vi) Adjust the values V_i of all affected parent nodes.

In the case of $y = 2.0$ and the resulting data structure shown in Fig. 5.1, a single random number and at most 6 conditional statements are required to pick the next transition independently of the size of the system. Similarly, a finite and fixed number of computations is required to execute the entire process, thus ensuring $\mathrm{O}(1)$ complexity of the operation. It is important to note, that while this result gives the rate of MC steps independent of the system size, the observed performance of a simulation nevertheless decreases with the size of the system. The perceived slowdown is due to the fact, that in order to traverse the state space (e.g. to equilibrate the system after some change in external conditions), the number of MC steps required is proportional to the size of the system.

The data structure shown in Fig. 5.1 has been further optimized for the particular processes that we study. Namely, the $e = 0$ environment has been chosen to be especially efficient to

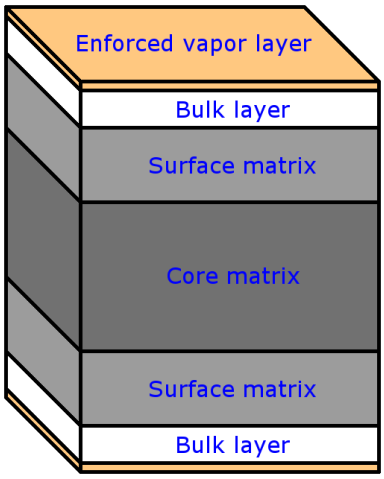

Fig. 5.2 Schematic representation of a simulation box and its partitions. The model sample is shown in light and dark grey depending on the proximity to the bulk layer. Bulk and the enforced vapour layers surround the sample material in z- (vertical) coordinate. Periodic boundary conditions are imposed horizontally in x- and y-coordinates.

access and alter. The studied materials, such as Vycor and silica aerogel, contain large voids of bulk space, which tend to stay empty throughout most of the simulation at low values of μ. Thus even if the external conditions are such, that any particular cell in the bulk is not likely to be chosen for a transition, due to the high number of them, the $e = 0$ environment is used relatively frequently. Therefore, its efficient access significantly speeds up an overall performance of the algorithm, especially at intermediate and high temperatures.

By exploiting the specifics of the perspective transition probability distribution, the developed kMC method implementation significantly improved the performance of the algorithm. Although some sacrifices in flexibility of wettability parameter y had to be made, within the scope of this work, a finite number of the available values of y has been sufficient in all relevant projects.

5.3 Sample preparation

The cubic samples of disordered porous materials, i.e. Vycor glass and silica aerogel, are generated as described in Sec. 1.1. However, in the case of most simulations additional preparation of the simulation box is required. As discussed in Sec. 3.4.2, an empty space slab (bulk layer) of 4 bcc-lattice spacings followed by a layer of enforced vapour is added to the sample in z-coordinate. For a more detailed analysis of the adsorption and desorption transitions, the porous media sample cube was further partitioned into three volume elements:

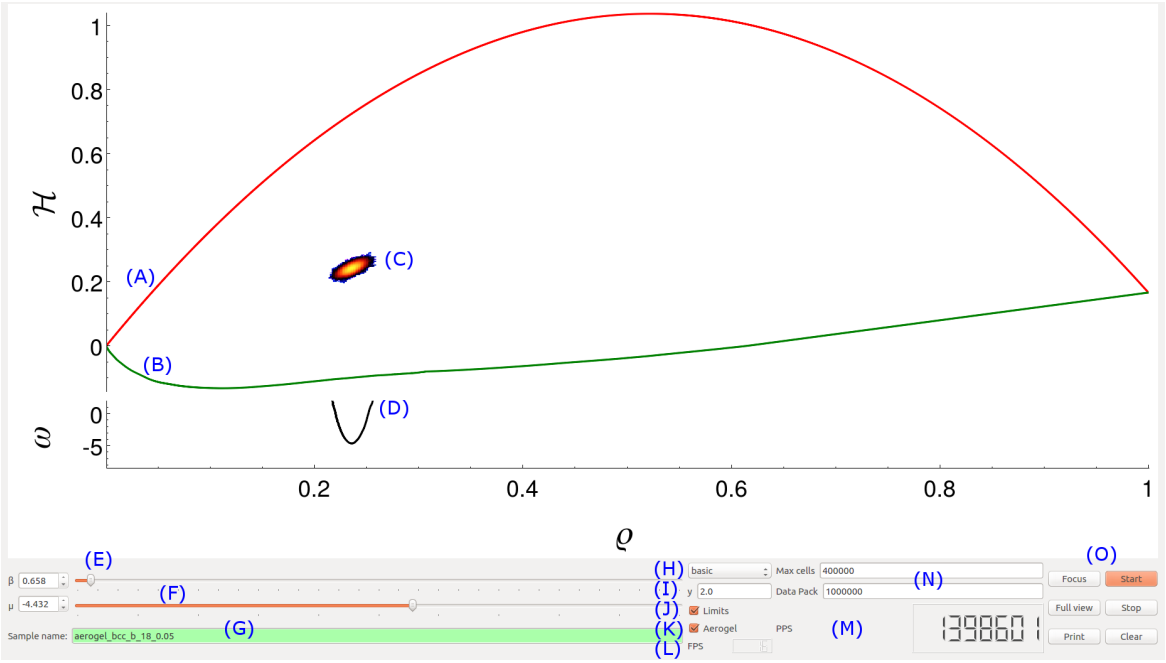

Fig. 5.3 The main screen view of the simulation engine displaying the macrostate-visit histogram. The functionality of the components marked by blue letters is explained in the text.

two equal surface matrix layers and core matrix between them (see Fig. 5.2). The volume of the core matrix is equal to the sum of volumes of surface matrix. The simulation engine is implemented to differentiate between the cells that belong to the bulk layer, core and surface matrix partitions as it is discussed in the following sections. Such capability allows for a more precise analysis of the system behaviour during the simulation.

5.4 Visual output

The simulation engine provides three visual output screens in real time. In this section, we briefly describe the output and the user interface of the engine.

Fig. 5.3 shows an example view of the engines main screen depicting the projected state space of the simulated system and the core part of the user interface. While the simulation is running the histogram of visited macrostates is being compiled and updated on the screen in real time. Below, we describe each component seen on the screen marked by the blue letters in Fig. 5.3.

(A) High-temperature boundary of the projected state space calculated as described in Sec. 2.3.6.

(B) Low-temperature boundary of the projected state space calculated as described in Sec. 2.3.7.

(C) The histogram of macrostate visits is coloured from the most visited (yellow) to the least visited (blue) macrostates on the background of the unexplored region of the projected state space in white.

(D) The grand-potential landscape, $\omega(\rho) \equiv \beta(\Omega(\rho) - \Omega) = -\ln(P(\rho))$, obtained as described in Sec. 3.4.2.

(E) β controller, that can be adjusted during the simulation by entering a precise number, using a slider or by keyboard short-cuts "l" and ",".

(F) μ controller, that can be adjusted during the simulation by entering a precise number, using a slider or by keyboard short-cuts "." and "/".

(G) Sample name to be entered before the simulation starts. The green (red) colouring of the field indicates that the data file describing a sample with such name exists (does not exist) in the input directory.

(H) The simulation engine can run under several different regimes including basic kMC, JW (jump-walking), low-temperature boundary scan, automatic μ-scan to obtain adsorption and desorption isotherms, infinite-temperature limit simulation, etc. This drop-down list controller allows the user to choose one of the regimes before the start of the simulation.

(I) Wettability controller. Currently available values are $y = 1.0$ and $y = 2.0$.

(J) Enabling/disabling this tick-box turns on/off the high- and low-temperature limits.

(K) Enabling/disabling this tick-box switches the display between the lattice-gas and Ising model regimes.

(L) Frames per second display shows how fast the engine is able to update the main screen view.

(M) kMC steps per second display shows how many MC steps the engine is generating at the moment.

(N) Some internal efficiency parameters of the engine controlling the display resolution and the size of data blocks that are used in communication between the different

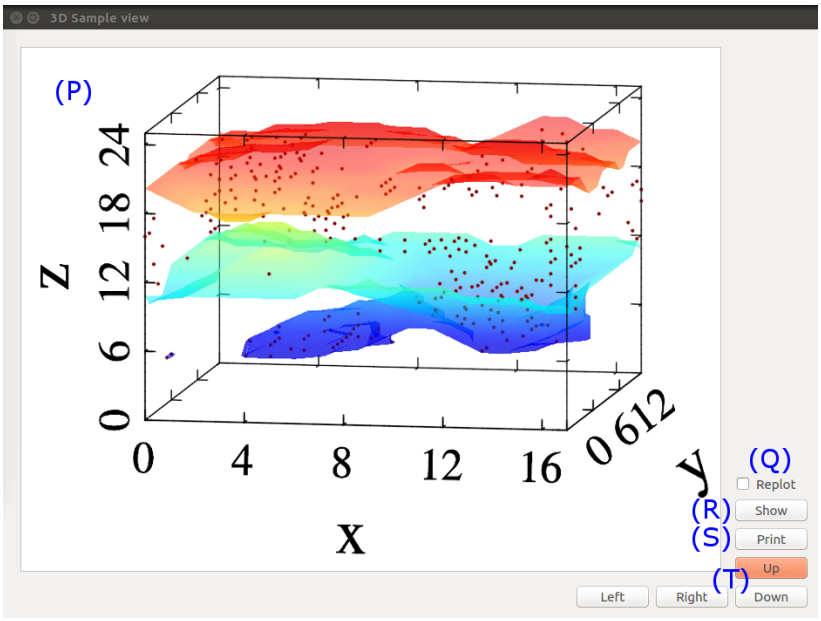

Fig. 5.4 The 3d sample-view screen of the simulation engine. The functionality of the components marked by blue letters is explained in the text.

components of the engine. These parameters are primarily designed to balance between the quality of the visual output and the simulation speed.

(O) Control buttons that start and stop the simulation, adjust the focus (also by key "f") of the projected-state-space view (e.g. if the user has zoomed in onto some particular area of the state space), clear all the data from the visit histogram and $\omega(\rho)$ graph (also by key "c"), and save the current view into a *.png* file.

The second visual output screen displays the 3d view of the simulation box as shown in Fig. 5.4. The ability to visualise the actual 3d structure of the fluid configuration in the metastable state, that the system has just settled down to, is of great value when developing an intuitive understanding of system behaviour. However, in general, displaying a complex 3d object on a 2d screen is a very challenging task. If the object does not have a clearly defined, convex, or at least familiar geometric shape, the observer attempting to make sense of such 2d presentation relies on a variety of assumptions. The 3d porous media samples that we are investigating, and the resulting fluid configurations inside them are notoriously difficult to display, since they consist of numerous interconnecting irregular surfaces with the empty voids in between. In our experience the highest clarity was achieved when fluid configurations were represented by semi-transparent surfaces that indicate the interface between the fluid and empty parts of the simulation box. The surfaces are coloured according to their z-coordinate for clarity. The matrix particles are marked by brown dots. In Fig. 5.4,

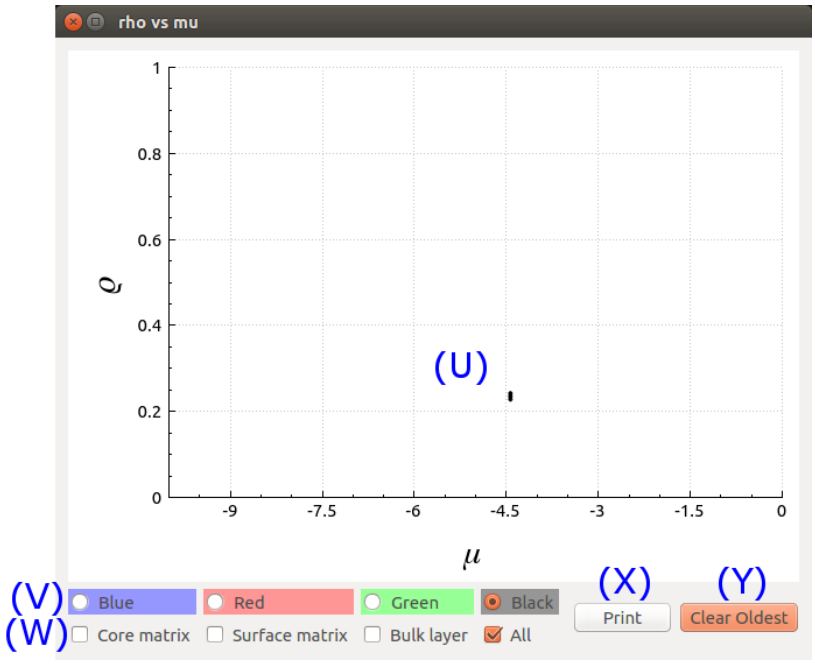

Fig. 5.5 The $\rho - \mu$ screen view of the simulation engine. The functionality of the components marked by blue letters is explained in the text.

we can see a roughly flat, horizontal surface in the upper part of the simulation box displayed in red. It separates the vapour phase above it and the fluid phase below. A more complex surface with a pocket of fluid sticking out below itself is shown in blue. Thus the fluid phase is between the two surfaces covering the matrix cells and the vapour phase is just above and below them, mostly in the bulk layer regions of the simulation box. Below, we describe the components of the 3d-view screen marked by the blue letters in Fig. 5.4.

(P) The 3d space rendered by MathGL library as described above.

(Q) Enabling this tick-box turns on the automatic re-rendering and redrawing the 3d view as frequently as possible.

(R) If the above tick-box is unchecked, a request to re-render and display an updated 3d view can be requested by pressing the "Show" button.

(S) "Print" button saves the current 3d view to the *.png* file.

(T) The simulation box can be rotated around two axes to obtain a 3d view from a different angle.

Finally, the third visual output screen is the $\mu - \rho$ diagram shown in Fig. 5.5. In addition to the usual density ρ recorded on the diagram as μ and/or β is varied, the simulation engine

is also able to output the separate densities for the bulk, surface and core matrix partitions of the simulation box, that were discussed in Sec. 5.3 and Fig. 5.2. Below, we describe the components of the $\mu - \rho$ diagram screen marked by the blue letters in Fig. 5.5.

(U) $\mu - \rho$ diagram updated in real time as μ and β are varied or ρ is updated by the simulation engine.

(V) The colour of the graph can be changed at any time, e.g. to distinguish between the adsorption and desorption isotherms.

(W) The density ρ can be split into its constituent parts corresponding to the bulk layer, surface and core matrix partitions, as described in Sec. 5.3 and Fig. 5.2. The respective densities are shown on the graph in different symbols.

(X) "Print" button saves the current $\mu - \rho$ diagram to the *.png* file.

(Y) "Clear" button removes the previous data. In particular, it removes the data that belongs to a graph generated before the earliest change of colour.

The three visual output screens provide a broad set of data about the current state of the system. The user is not only able to see the current thermodynamic variables, but also their distributions in the projected state space and the real-space fluid configuration that corresponds to the current state. Moreover, the ability to interact with the system in real-time, allows the user to observe the transitions from one equilibrium state to another in a remarkably direct way, and thus develop a more thorough understanding about the behaviour of the system.

5.5 Implementation structure

In general, building any real-time interactive graphical application requires to resolve several common problems that arise due to the computational constraints. In this section, we briefly describe the architecture of the simulation engine, how it addresses the relevant problems, and the possible modifications that could improve the current performance of the application.

One of the main issues that arise in graphical output applications which display the results of some demanding computational task (such as MC simulation), is the sharing of the computing resources between the graphical presentation and the core task. Moreover, the desired interactivity of the application further complicates the situation. In our case, if the application were built in the straightforward way, the MC simulation would have to pause every time the user decided to refocus the main screen on a different part of the state space,

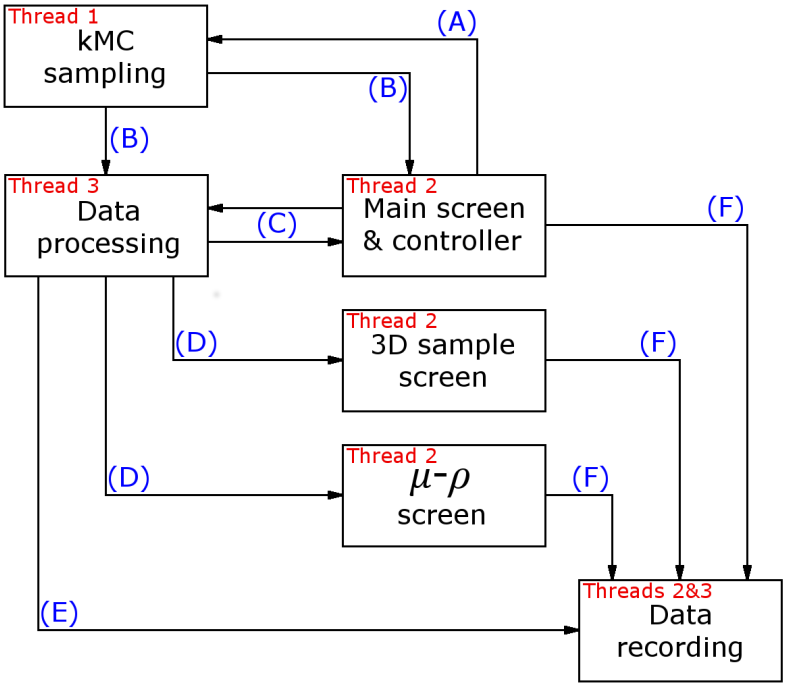

Fig. 5.6 Schematic representation of the simulation engine architecture and the data flows within it. The arrows marked by blue letters indicate communication between the connected components of the engine as described in the text.

render a 3d view, or alter the value of μ, since that requires redrawing of the entire projected state space displayed on the main screen. In order to avoid such problems and to run the MC simulation uninterrupted, the application is built to rely on separate CPU threads for different tasks. One thread is dedicated entirely for the MC simulation, another for most of the data processing, preparing it to store and display, and the third one for rendering the graphical output.

The architecture of the simulation engine is shown Fig. 5.6. The arrows marked by blue letters depict the corresponding data flows between the main components of the engine. The simultaneous data exchange between the different threads was implemented through the "signals and slots" mechanism provided by Qt framework. The implementation structure of the application is best explained by taking a closer look at such interactions between the different parts of the system, as described below.

(A) The user enters the initial information to the controller, such as simulation regime, sample name and type, initial β and μ values, etc., and starts the process. The controller in turn sends this data to the core simulation component of the engine, operated by *Thread 1*, which starts the MC simulation of the system.

(B) The output in the form of macrostate visit durations and the current state variables of each cell, is grouped into larger data bundles and sent back directly to the main screen display as well as the data processing unit operated by *Thread 3*.

(C) When the user changes the value of μ, clears the screen, or zooms into some part of the state space, the visit histogram data has to be altered and prepared for display, which is done by the data processing component. Therefore, such operations are followed by the controller sending the request for the updated data, and the data processing unit responding once ready. This way, the main screen display and the controller do not freeze while the data is being prepared, and stay responsive in the meantime. An additional advantage of such architecture is that only the latest request of the same kind is being worked on, e.g. if the value of μ is changed too frequently to be updated on the display, the engine does not necessarily execute all operations one after another, but instead, prioritizes the latest order and forgets the rest.

(D) The data processing unit converts the raw output received from the simulation component into a form that is ready for display on 3d sample- view and $\mu - \rho$ diagram screens. This way, the work that is dedicated to the *Thread 2*, responsible for the graphics rendering, is minimised, and thus a smoother user experience is achieved, which results from sufficiently high rate of frames per second.

(E) In addition to the real-time visual output, the simulation engine also records several other properties of the system, such as correlation times, internal energy statistics and data for ergodicity measures. Such information is saved in the data files on the hard-drive throughout the simulation.

(F) Finally, each graphical output screen has a button "Print", which saves the current view snapshot on the hard-drive in the form of *.png* file, as described in the previous section.

The details of the processes outlined above depend on the particular regime under which the application is operating, but the architecture and the main features of the system remain as described. The primary goals when designing the application was to minimize the work done by the core simulation and the graphical rendering components. As discussed previously, the aim was to run the MC simulation uninterrupted in the background, independently on the user interface and display issues that may arise. However, one of the unfortunate constraints of the Qt framework is that only a single thread can be used to support the graphical output itself, which includes all display panels in the application. In order to maintain the desired rate of frames per second, such restriction demanded to minimize the work dedicated to the graphical output components (*Thread 2*). Thus the intermediate data processing component

(*Thread 3*) was implemented to prepare the data in the form that is ready to display, and only then send it to the output screens. The resulting system functions seamlessly, without any interruptions to the MC simulation part, and for most systems maintains steady 15-30 frames per second update rate for the main and the $\mu - \rho$ graphical output screens. However, producing the real-time 3d view turned out to be a more difficult task. The rendering of the overlapping transparent surfaces in a 3d space, is highly computationally demanding. Usually, it is performed by the dedicated graphics card operated by OpenGL environment or by one of the wrappers written on top of OpenGL, to simplify its use. However, at the time of development, Qt framework did not provide a free version of its 3d-graphics package "Qt Data Visualization". One of several tested alternatives was a freely available library called MathGL [207], a project built by Dr. A. A. Balakin, designed for high-quality scientific graphics. Combined with the Qt framework we were able to adapt MathGL to our real-time visualisation purposes. However, originally MathGL was not designed for real-time streaming. It renders its visual output using CPU instead of the graphics card, which significantly slows down the performance. As a result of the single CPU core restriction mentioned above, it is impossible to achieve a seamless user experience if the 3d sample view is rendered continuously, since that prevents the simultaneous user interaction through the controller. For this purpose, when simulating large systems, it is recommended to disable continuous rendering of the 3d view and rely on output upon request, e.g. once some metastable state of interest is achieved. Further improvements of the application are therefore mostly related to the 3d view and its features. Decoupling the 3d rendering from the rest of the graphics would grant an immediate improvement in user experience. However, a major increase in visual quality and performance of the 3d-view screen, can only come from some application of OpenGL and the use of the graphics card. The new releases of Qt framework, including Qt3d component updates, which came out with the Qt 5.8 version in January, 2017, might provide the valuable tools for such development.

5.6 Examples of use

In this section, we present several examples of possible use of the engine, which demonstrate some of the available functionality. For the purpose of clarity, we disabled displaying the matrix distribution in the 3d view of the simulation box (as previously seen in Fig. 5.4), and retained only the interfaces between the fluid and vapour phases. Unless specified otherwise, the system used in the following examples is the cubic aerogel sample of size $18 \times 18 \times 26$ and porosity $\phi = 95\%$, generated as described in Sec. 1.1.2 and Sec. 5.3.

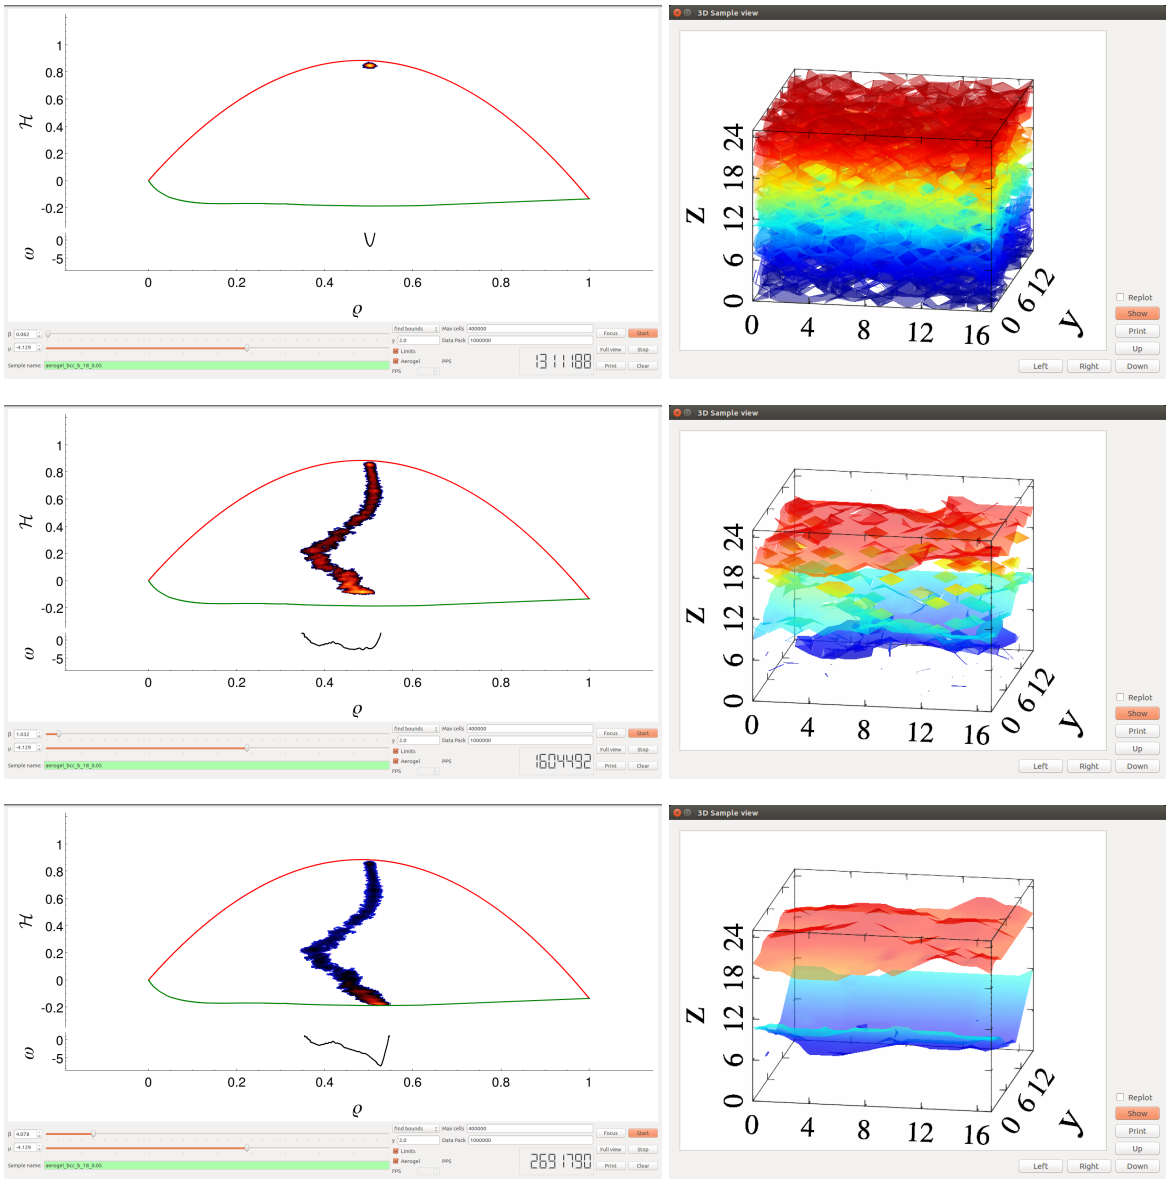

Fig. 5.7 Screen-shots of the main and 3d-view screens as the system is driven from high-temperature state ($\beta = 0.062$) shown in the top panel, through the intermediate-temperature ($\beta = 1.032$) state (middle panel), to the low-temperature state ($\beta = 4.078$) in the bottom panel.

5.6.1 Cooling down

We start with an example in which the system is initiated at a high temperature ($\beta = 0.062$) as shown in the upper panel of Fig. 5.7. We can see that the sampled states are at the top of the state space, just below the infinite-temperature limit. The 3d-view screen displays many intersecting surfaces within the simulation box, indicating that empty and fluid cells do

not agglomerate in any ordered fashion, but instead are randomly distributed, as expected at high temperatures. Simply by dragging the β-control slider, we reduce the temperature to $\beta = 1.032$ (central panel in Fig. 5.7). We can observe the path that the system has taken across its projected state space during the transition. Such paths depend on the speed with which the external parameters, e.g. temperature, is altered. This way, it is possible to explore equilibration behaviour dynamically, by changing the external parameters at some steady rate in a cyclical manner. However, for this example, the change was slow enough to allow the system to equilibrate at most of the intermediate positions, thus obtaining an equilibrium path across the state space for $\mu = -4.129$ as the temperature is being reduced. It is apparent that at $\beta = 1.032$ the fluid configuration has formed well-defined surfaces with only a few minor fluctuations. Further lowering the temperature to $\beta = 4.078$ (bottom panel in Fig. 5.7) moves the explored part of the projected state space close to the low-temperature boundary and removes the fluctuations entirely. The resultant fluid configuration consists of a layer of fluid that is entirely localized between two roughly horizontal surfaces in blue and red.

5.6.2 Zooming in

The second example (see Fig. 5.8) continues from a similar state to the one, which the system in the first example has left off. After clearing the accumulated data, the engine builds a new state space visit histogram at $\beta = 2.0$ as shown in the top panel of Fig. 5.8. The explored region of the projected state space can be seen as a narrow line along the low temperature boundary. Simply by scrolling the mouse wheel, it is possible to zoom in closer, as shown in the second panel of Fig. 5.8, however such operation retains the original resolution of the screen. By pressing "Focus" button on the control panel, or a key "f", the engine optimizes resolution for the current field of view, which results in the visual output shown in the third panel of Fig. 5.8. Such view helps to appreciate the density of states even at the lowest energy regions of the state space. For example, there appears to be no fine structure in the macrostate visit histogram, that would indicate large fluctuations in the density of states at this region of the state space. Just to further demonstrate the level of detail which is available, the bottom panel of Fig. 5.8 shows a close-up view of the macrostate visit histogram at different value of β. At the edges of the explored state space it is even possible to discern the individual paths that the system has taken during the simulation. The patterns and the structure that can be observed in such diagrams provide us with the valuable intuition about this complex drifting process, which is conventionally reflected by just a single data point in a standard $\mu - \rho$ diagram.

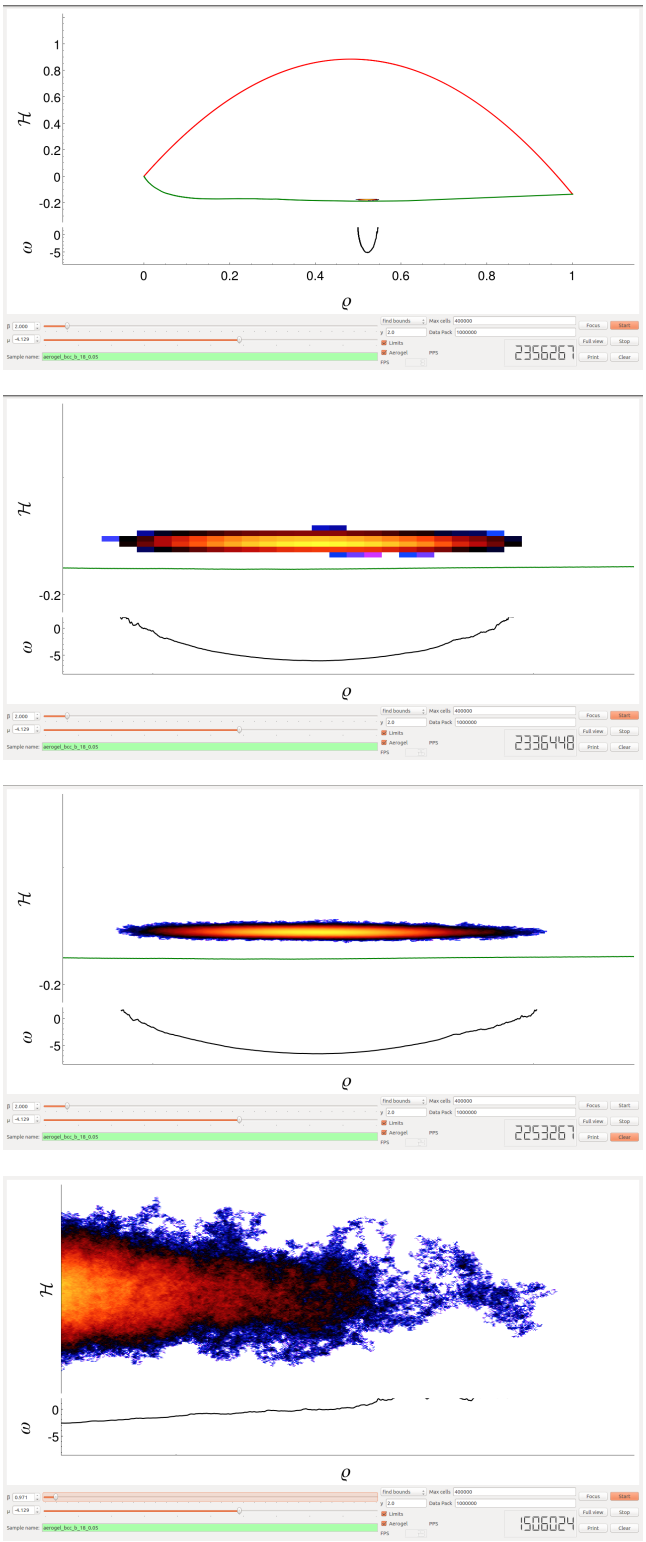

Fig. 5.8 Screen-shots of the main screen of the simulation engine, that demonstrate the zoom-in and refocus procedures. The upper panel shows the entire projected state space with a low-temperature metastable state at the bottom. The second panel from the top demonstrates the same metastable state zoomed-in. The third screen-shot is taken after refocussing the view shown in the second panel. The bottom panel shows an even further close-up view of the macrostate visit histogram.

5.6.3 Generating isotherms

The final example demonstrates how the standard adsorption and desorption isotherms can be generated by the simulation engine, and how such process can be inspected in more detail through the additional visual outputs that the application provides. The change of μ can be set to automatic regime, in which case the full hysteresis loop is obtained at a constant rate of change of μ. Otherwise, the user can simply move the μ-control slider as in the first example regarding change of β. Fig. 5.9 shows the state-space diagram and the 3d-view screens while μ is altered from $\mu = -5.0$ to $\mu = -3.8$ at $\beta = 1.2$. We can observe the state-space boundaries shifting dynamically while μ is being adjusted as discussed in Ch. 2. The fluid configuration changes from the low-density state, where only a thin layer of fluid covers the matrix sites, to a high-density one, where almost the entire simulation box is filled with fluid. Due to the finite temperature, we can observe small fluctuations within the volume of the simulation box as discussed in the previous examples. The isotherm obtained during such operation can be seen in the $\mu - \rho$ screen of the application as shown in Fig. 5.10. At this temperature the hysteresis loop cannot be clearly observed, since the adsorption and desorption isotherms largely overlap (not shown). However, at $\beta = 10$, similar operations as described above, allow us to obtain distinct adsorption and desorption isotherms. If the colour is changed to red in between the phases of increasing and decreasing μ, the resulting $\mu - \rho$ diagram displays the hysteresis loop as shown in Fig. 5.11.

It is also straightforward to obtain the scanning curves for any system, by simply reversing the change of μ in the middle of adsorption or desorption processes. Moreover, at any point of this process, we are able to see the exact location in the projected state space that the system is in, as well as the spacial configuration of the fluid within the simulation box. Such capabilities provide us with the understanding about the nature of the metastable states within the hysteresis loop, since we are able to see the differences between the fluid configurations that correspond to the distinct metastable states, detected on the $\mu - \rho$ diagram.

5.7 Conclusions

In this Chapter, we have presented a simulation engine, designed and developed to investigate sorption processes in porous media. Given certain assumptions that apply to the systems of our interest, a kMC implementation has been optimized as discussed in Sec. 5.2. The improved kMC algorithm proved to be significantly more efficient than its standard counterparts. The capabilities of the tools provided by the simulation engine as well as the architecture of the application itself have been discussed in detail. Finally, several usage examples have been

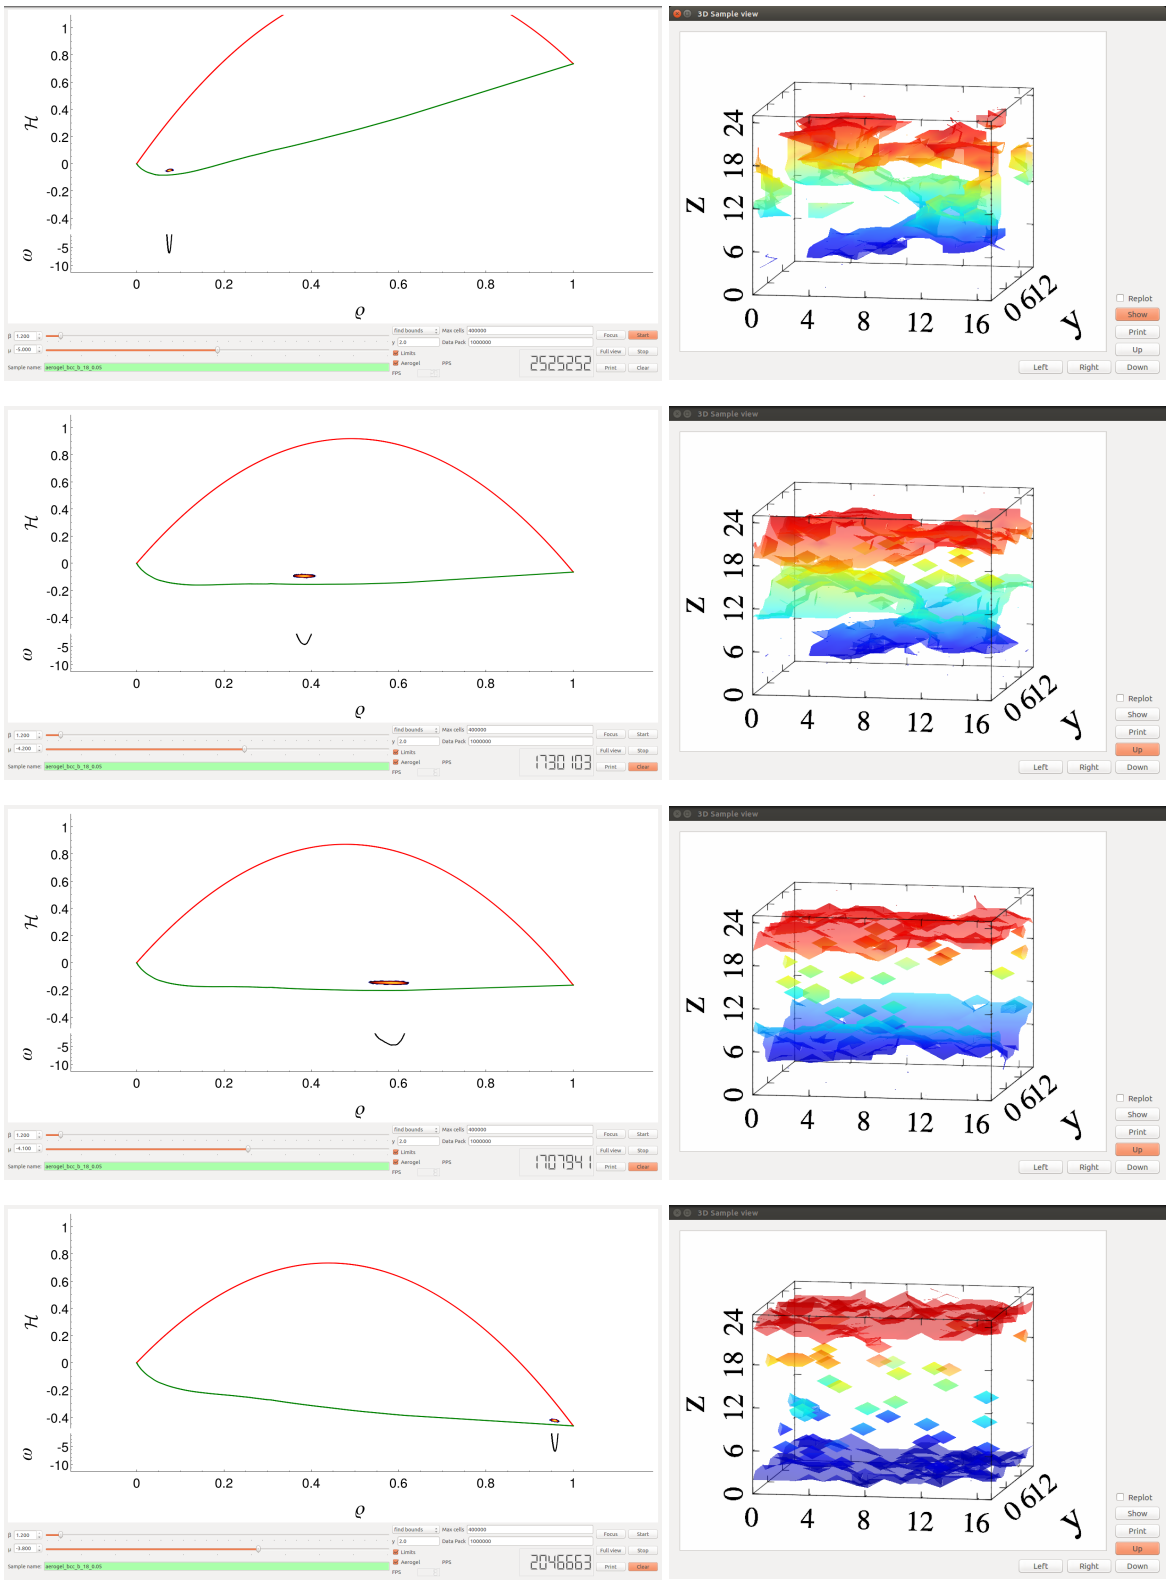

Fig. 5.9 Screen-shots of the applications main and 3d-view screens while the system is driven from $\mu = -5$ shown in the top panel, to $\mu = -3.8$ in the bottom panel. $\beta = 1.2$.

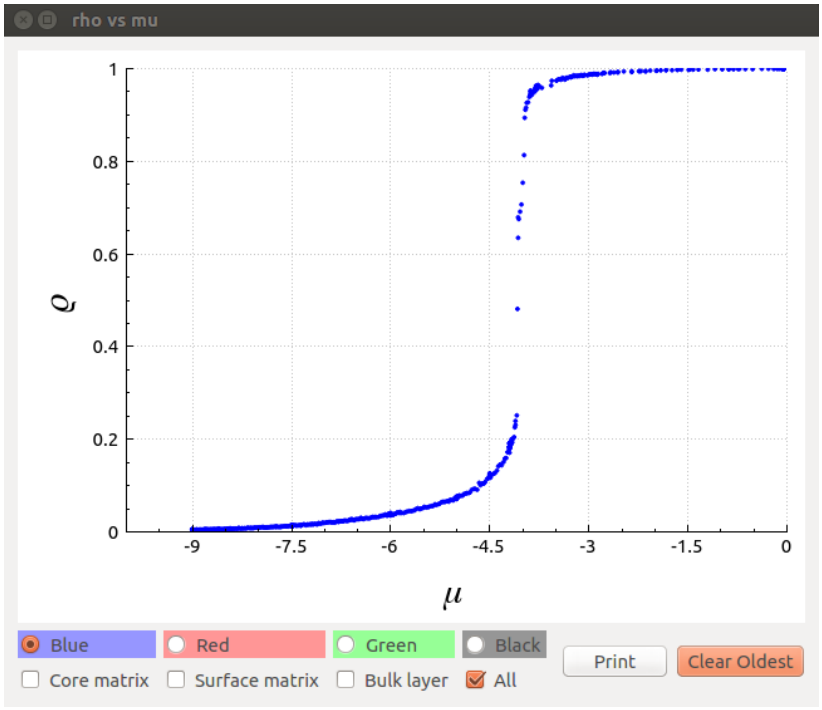

Fig. 5.10 The isotherm displayed on $\rho - \mu$ screen of the application, generated at $\beta = 1.2$ as described in the text and Fig. 5.9.

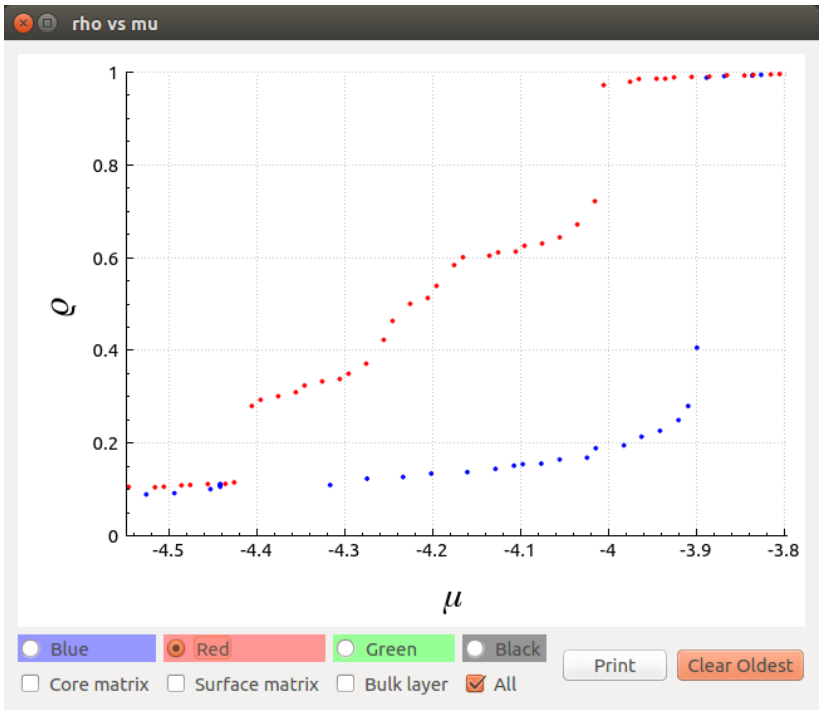

Fig. 5.11 Hysteresis loop for the adsorption (blue points) and desorption (red points) isotherms at $\beta = 10$ displayed on the $\rho - \mu$ screen of the application, obtained as described in the text.

described that demonstrate the potential of such engine. The main results of this Chapter can be summarised as follows,

(i) development of the $O(1)$-complexity kMC algorithm by means of constructing the system-dependent hybrid data structure;

(ii) development of the interactive simulation engine based on the optimised kMC and allowing to study and visualise in real time the dynamics of certain complex systems;

(iii) testing the engine for systems exhibiting binary dynamics.

While new features, efficiency updates and user interface were being improved throughout the entire length of the PhD project, the core components of the application were developed simultaneously with the contents of Ch. 2. Therefore, the later work, the results of which were discussed in parts of Ch. 2 and most of Chs. 3-4, significantly benefited from the use of the simulation engine.

The breadth and depth of the information provided by the application has helped us to address countless number of issues quickly and efficiently. Simply by manipulating a few buttons and sliders, we were able to gain insights into questions such as:

- how the scanning curves behave in this sample;
- how dense the metastable states within the hysteresis loop are;
- which fluid configurations the metastable states correspond to;
- what approximately the low temperature boundary for a particular system is;
- what the impact of the cooling rate is on a metastable state that the system settles down to;
- what approximately the critical parameters of the system are;
- how large the fluctuations of the order parameter at the critical point are;
- how robust the paths across the projected state space during the phase transitions are;
- in what order the fluid fills the sample and empties from the sample during the adsorption and desorption processes.

However, the most important contributions by this tool came not from answering the questions, but rather by guiding us to ask the right questions in the first place. If we were given an object in a box, and asked to describe it, we would not start by measuring its

diameter or elasticity along some particular axis, testing its radiation spectra or vibrational frequencies. Instead, we would open the box. The set of tools that comprise the simulation engine were designed for precisely such purpose in mind. The simulation engine enables the user to interact with and explore the system from multiple aspects at once, which helps to raise further questions and often provide the design ideas for the accurate measurements to answer them.

Despite the significant progress that has been achieved in this sub-project, the current state of the application leaves a lot of room for improvement. Beyond the already mentioned possible developments in graphics and efficiency (see Sec. 5.5), the user interface would benefit from certain amount of polish, to ease the learning process and make the user experience more intuitive. From the start, the application was designed to enable and encourage expanding functionality. It is straightforward to automatize manipulation of the already existent controls, e.g. to write a subroutine that automatically produces the scanning curves. Furthermore, the architecture of the application makes it also convenient to implement new features within the engines environment, e.g. computation of the correlation times, ergodicity measures or even complex algorithms such as jump-walking.

Chapter 6

Sorption in random-network models of porous media

6.1 Introduction

There has been a lot of progress achieved recently in developing analytical techniques to analyse sorption dynamics within the lattice-gas model defined on certain complex-network topologies [208, 93, 98]. In particular, the exact solutions for equilibrium isotherms were obtained for bilayered Bethe lattice [208] and a random graph [98]. Unfortunately, most real materials either natural (such as soil and rocks) or synthesized (e.g. silica aerogel and Vycor glass) usually have complex structure and irregular, however not entirely random, pore-network topology. In this Chapter, we numerically examine the possibility to apply recently developed analytical techniques to such materials.

Since the exact solutions [98] can only be applied to the random networks of a given degree distribution, in order to study the real systems represented by a complex pore networks, we alter their representation to a particularly constructed random graph. Then, using kMC simulations, we analyse sorption dynamics in both, the original and reconstructed networks, and analyse how the change in topology impacts sorption behaviour. The central aim of the study is to extract the properties of the real systems which determine whether their representation by a particular random network can be used to accurately simulate sorption behaviour and, hence, analyse sorption in real systems analytically. The analysis is done for two types of porous media, i.e. soil and aerogel.

In the following section, we introduce soil, a type of porous media that has not yet been investigated in this work. We continue by describing the methodology to generate model systems from experimental soil imaging data, that will be used in the simulations.

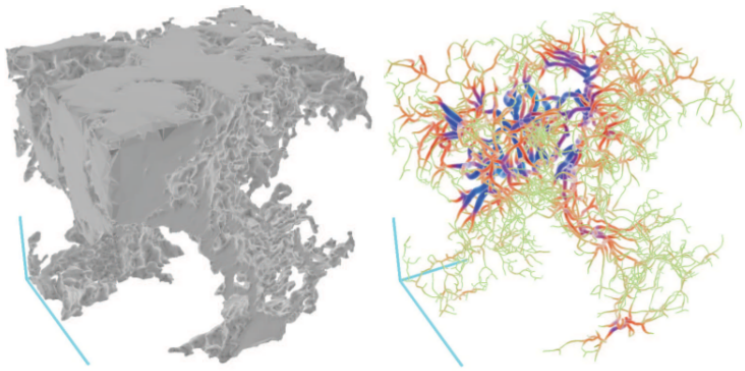

Fig. 6.1 Left: an X-ray micro-tomography image of a cubic soil sample with pore space shown in grey. Right: the network representation of the skeletonized pore space for the soil sample shown in the left panel [209].

Theoretical analysis of sorption behaviour in such systems is presented in Sec. 6.3. The results are discussed in Sec. 6.4 and the conclusions are given in Sec. 6.5.

6.2 Soil

In nature, soils are heterogeneous assemblies of matter forming porous media. The soil models analysed in this Chapter were obtained by using images of soil samples obtained and digitally processed in the James Hutton Institute, Invergowrie, Dundee [209]. Several representative soil samples of cubic shape of size $3.5 \times 3.5 \times 3.5$ cm were scanned by an X-ray micro-tomography device producing soil-density maps which can be used for construction of the 3d networks (Fig. 6.1). The soil pore space can be modelled as a network by applying a procedure called skeletonization (see Ref. [209] for more detail). This methodology leads to a simple representation of the soil pore structure by means of the medial lines which pass through the pore space as illustrated by red lines in Fig. 6.2a. The intersection points between such lines are labelled as network nodes (node cells) with coordination number $q > 2$. Each channel between two nodes is represented by a set of two-coordinated nodes ($q = 2$) called channel cells. The number of nodes in the channel is proportional to the length of the channel. In addition to the node and channel cells in such network, there are terminal cells, i.e. the end-points of channels (node with coordination number $q = 1$). The resulting network, below referred to as *skeletonized soil*, is characterised by the distribution in coordination number shown in Fig. 6.2b.

It is known [209] that the pore space in soil is highly heterogeneous and geometrical characteristics of cells in skeletonized network are broadly distributed. The effects of such

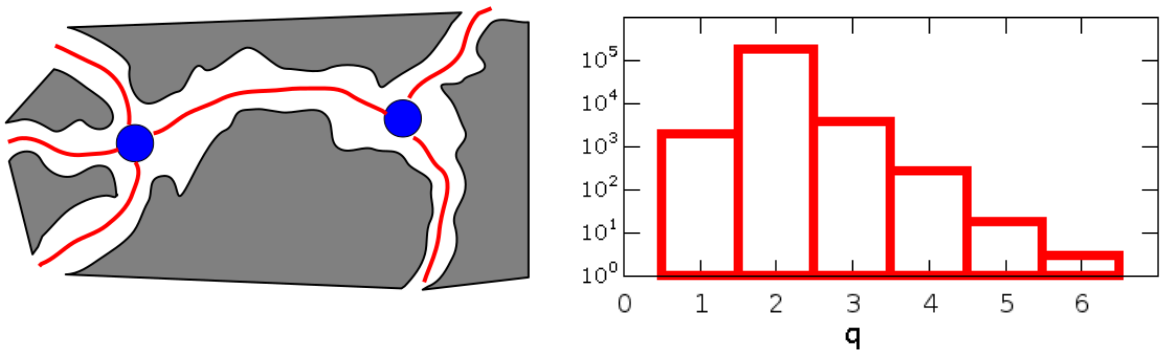

(a) (b)

Fig. 6.2 (a) Schematic illustration of skeletonization resulting in replacement of pore-space channels (shown in white) by the medial red lines. The blue nodes mark intersection of several medial lines [209]. (b) The distribution of node coordination number in a typical skeletonized soil network

heterogeneity on sorption have been studied in Ref. [210]. In this Chapter, we mainly concentrate on the effect of the network topology and thus use a simplified model for description of sorption in soil. According to this model, we assume that all the cells in the network independent of their coordination numbers, when filled with fluid, equally interact with matrix, i.e. the matrix-fluid interaction strength, w_{mf}, is identical for all pore cells in the network. The fluid-fluid interaction strength, w_{ff}, is also assumed to be the same for all interacting pairs of fluid cells. Under these assumptions, the effective Hamiltonian describing sorption can be written as (cf. Eq. (1.2)),

$$\tilde{\mathscr{H}} = \mathscr{H} - \mu \sum_i \tau_i = w_{\mathrm{ff}} \sum_{\langle i,j \rangle} \tau_i \tau_j - (w_{\mathrm{mf}} + \mu) \sum_i \tau_i \,. \tag{6.1}$$

It follows from Eq. (6.1) that the matrix-fluid interaction does not have a significant effect and just leads to the shift in the value of chemical potential. In contrast, the fluid-fluid interactions are determined by the network topology and can result in peculiar behaviour of sorption.

6.3 Avalanches in skeletonized soil samples

In this Section, we investigate the avalanche-like processes that are most prominent in skeletonized soil and random network systems. We categorise two conceptually distinct processes by which the avalanches can proceed and discuss how they manifest themselves in the studied systems.

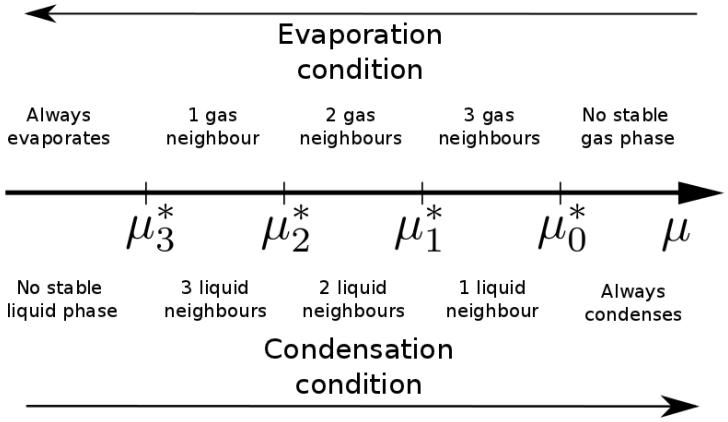

Fig. 6.3 Schematic representation of the sorption dynamics in the system governed by type-A transitions.

6.3.1 Type-A transition

In order to understand the underlying micro-scale behaviour of the system, let us consider the conditions required for a single cell to switch state. From Eq. 6.1, we find the reduced energy difference between two configurations $\mathscr{C}_0$ (current state) and $\mathscr{C}_k$ (perspective future state), which differ only at the cell k:

$$\tilde{\mathscr{H}}_k - \tilde{\mathscr{H}}_0 = -w_{\mathrm{mf}}\sigma_k - w_{\mathrm{ff}}\sigma_k \sum_{\langle i|k\rangle} \tau_i - \mu\,\sigma_k \ , \tag{6.2}$$

where $\sum_{\langle i|k\rangle}$ is a sum over the nearest neighbours of k^{th} cell, and the spin variable,

$$\sigma_k = 1 - 2\tau_k = \begin{cases} +1 & \text{if } \tau_k = 0 \text{ in } \mathscr{C}_0 \\ -1 & \text{if } \tau_k = 1 \text{ in } \mathscr{C}_0 \end{cases} \ .$$

From Eq. (6.2), it follows that the configuration $\mathscr{C}_k$ becomes as likely as configuration $\mathscr{C}_0$ once $\tilde{\mathscr{H}}_k - \tilde{\mathscr{H}}_0 = 0$, which can be rewritten in terms of characteristic value, $\mu^*_{0\leftrightarrow k}$, of the chemical potential,

$$\mu^*_{0\leftrightarrow k} = -\left(w_{\mathrm{mf}} + w_{\mathrm{ff}} \sum_{\langle i|k\rangle} \tau_i \right) \ . \tag{6.3}$$

Therefore, we expect the equilibrium $\mathscr{C}_0 \leftrightarrow \mathscr{C}_k$ transitions to occur at $\mu = \mu^*_{0\leftrightarrow k}$, where $\mu^*_{0\leftrightarrow k}$ is defined by Eq. (6.3). We can also explicitly see how $\mu^*_{0\leftrightarrow k}$ dependence on the nearest neighbour states results in hysteresis phenomena, since it implicitly depends on the past states of the system.

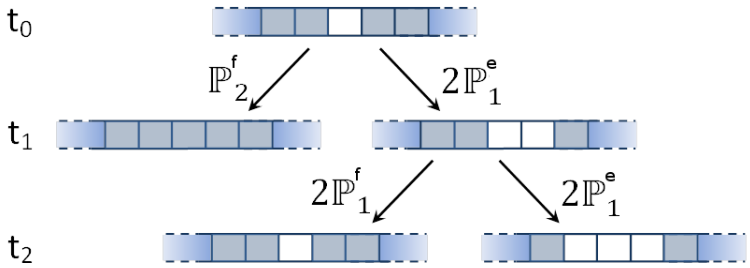

Fig. 6.4 The scheme of possible sequence of events at the beginning of a type-B transition. At the beginning of the process, t_0, a cell in a channel spontaneously evaporates, thus opening a "hole" in an otherwise all-fluid channel. At time $t_1 > t_0$, one of the two possible events could happen with certain probabilities: either the "hole" has closed (with the probability $\mathbb{P}^f_2$), or expanded into two cells (with the probability $2\mathbb{P}^e_1$)

Let us consider a long channel filled with fluid and terminating at one end by cell with $q = 1$, while μ is continuously decreased. If the geometrical properties of the cells are the same, the last cell in the channel with $q = 1$ (the end-point) is going to be the first to empty (Eq. 6.3), due to the fact that it has only one fluid neighbour, while the other cells in the channel have two. As soon as the upcoming phase switch is favourable for the cells with one fluid neighbour interaction, i.e. $\mu < \mu^*_{0\leftrightarrow k} = \mu^*_1 = w_{\mathrm{mf}} + w_{\mathrm{ff}}$, we expect to observe a chain reaction (an avalanche), because once the end-point empties, the next one will have only a single fluid neighbour and therefore the channel will continue to evaporate until a node with at least two fluid neighbours is encountered. Note, that if the variation of μ is reversed, the condensation/evaporation process also reverses and the system goes back to the previous state. We refer to this process as *type-A transition*. This type of chain reaction is an essential feature in skeletonized soil systems and 1d pores [93], and helps to explain most of the transitions between phases observed in Sec. 6.4. Schematic representation of the sorption behaviour in a system governed by type-A transitions is shown in Fig. 6.3.

The formalism can be straightforwardly generalised from the analysed 2-regular network with $q = 1$ end-points to any q-regular network containing ”defect” nodes characterised by coordination numbers less than q. In such networks, the defect nodes act as the "weak" points during the evaporation process. We analyse such system computationally in Sec. 6.4.

In the following text, notation $\mu^*_{0\leftrightarrow k}$ is replaced by μ^*_i for a cell which has i fluid neighbours.

6.3.2 Type-B transition

Let us consider a system with a fluid channel, having both of its ends connected to the fluid environment, e.g. to the nodes with $q > 2$. When μ decreases, clearly, the chain reaction

does not start before $\mu > \mu_1^* = -(w_{\mathrm{mf}} + w_{\mathrm{ff}})$ and the transition has to be over by the time when $\mu < \mu_2^* = -(w_{\mathrm{mf}} + 2w_{\mathrm{ff}})$, because at this point even the cells with two fluid neighbours would spontaneously evaporate. However, in the middle of this range, the system is stable in both empty and fluid state.

Due to the stochastic nature of the system, among a large number of cells or during long periods of time, there exist cells which randomly switch their state, thus creating a "hole" in otherwise uniformly fluid channel. The following sequence of possible events is shown in Fig. 6.4. Let us define the value of the chemical potential, $\mu^\dagger$, as the value at which the probabilities for the empty "hole" to close or to start expanding are the same, i.e. $\mathbb{P}_2^{\mathrm{f}}(\mu^\dagger) = 2\mathbb{P}_1^e(\mu^\dagger)$, where the upper index denotes the state that the system goes into (fluid, f, or empty, e), and lower index refers to the number of fluid-fluid interactions the considered cell currently has. Bearing in mind, that

$$\mathbb{P}_2^{\mathrm{f}} \propto e^{\beta(w_{\mathrm{mf}}+2w_{\mathrm{ff}}+\mu)},$$
$$\mathbb{P}_1^e \propto e^{-\beta(w_{\mathrm{mf}}+w_{\mathrm{ff}}+\mu)},$$

we obtain the following expression for the characteristic value of the chemical potential, $\mu_{f\to e}^\dagger$,

$$\mu_{f\to e}^\dagger = -\frac{1}{2}\left(2w_{\mathrm{mf}} + 3w_{\mathrm{ff}} - \frac{\ln(2)}{\beta}\right). \tag{6.4}$$

The above calculation neglects the possibility for the "hole" to close if it initially expands in 2 empty cells. This approximation is valid due to the fact that $\mathbb{P}_1^{\mathrm{f}} \propto e^{\beta(w_{\mathrm{mf}}+w_{\mathrm{ff}}+\mu)}$ is approximately 12 orders of magnitude smaller than $\mathbb{P}_1^e$ in the range of investigated conditions. Therefore, it is evident that once the "hole" of size 2 is formed, the chain reaction begins and the fluid in the channel proceeds to empty until the fluid-vapour boundary reaches the ends of the channel. In the following text, we refer to this process as *type-B transition* (it can be also called cavitation).

Similarly, we can calculate the characteristic value of $\mu_{e\to f}^\dagger$ for the transition when the system is initially empty and μ is increasing, i.e.

$$\mu_{e\to f}^\dagger = -\frac{1}{2}\left(2w_{\mathrm{mf}} + w_{\mathrm{ff}} + \frac{\ln(2)}{\beta}\right). \tag{6.5}$$

It is straightforward to generalize the results shown above, by considering a network of randomly connected nodes with fixed coordination number q (q-regular graph). The analysis, identical to the $q = 2$ case, predicts the type-B emptying and filling transitions in a q-regular

graph at:

$$\mu^{\dagger}_{f\to e}(q) = -\frac{1}{2}\left(2w_{\mathrm{mf}} + (2q-1)w_{\mathrm{ff}} - \frac{\ln(q)}{\beta}\right), \tag{6.6}$$

$$\mu^{\dagger}_{e\to f}(q) = -\frac{1}{2}\left(2w_{\mathrm{mf}} + w_{\mathrm{ff}} + \frac{\ln(q)}{\beta}\right). \tag{6.7}$$

From the above equations, we can deduce that variation in w_{mf} simply shifts the whole hysteresis loop (as expected), while variation in w_{ff} affects the width (indicated by the difference between $\mu^{\dagger}_{f\to e}(q)$ and $\mu^{\dagger}_{e\to f}(q)$) and the location of the loop.

It is important to note that, by definition, the value of $\mu^{\dagger}$ does not represent the exact value of the chemical potential at which the phase transition takes place, but rather serves as an estimate of the upper bound if μ is increasing, and lower bound if μ is decreasing. The estimate is more accurate for the regime of short MC waiting times (between the increments of μ) and small systems, because under such conditions the "holes" are created very rarely, and thus the phase change occurs only if $\mu \simeq \mu^{\dagger}$. In contrast, at large MC waiting times the "holes" are formed more often at the same value of μ and thus the system has more "attempts" to start the chain reaction. Therefore, the transition occurs at lower values of μ than predicted by $\mu^{\dagger}$, when μ is increasing (i.e. for adsorption), and at higher vales of μ for desorption.

Unlike the type-A transition described in the previous Section, the type-B transition is irreversible, because once the "hole" expands to two empty cells, it is extremely unlikely to close. In this sense, type-A transition could be considered as deterministic, because it always happens at the same exact value of μ and proceeds relatively slowly at $\mu \simeq \mu^{*}$, since the energies of the empty and fluid states are similar at such μ. In contrast, the type-B transition has a more stochastic nature because, at microscopic level (considering one particular channel), it does not have the exact μ of transition, and once the chain reaction starts, it proceeds instantaneously and is irreversible.

6.4 Results and discussion

The results section is divided into two parts. In Sec. 6.4.1, we start with the analysis of skeletonized soil samples and attempt to explain all the discrepancies between the sorption behaviour in the original samples and their random-network representations. The random networks are constructed using the configuration method [211], which maintains the same distribution of node coordination numbers, but connects them in a random order. In section 6.4.2, we present a similar analysis for the bcc lattice models of soil and aerogel. We

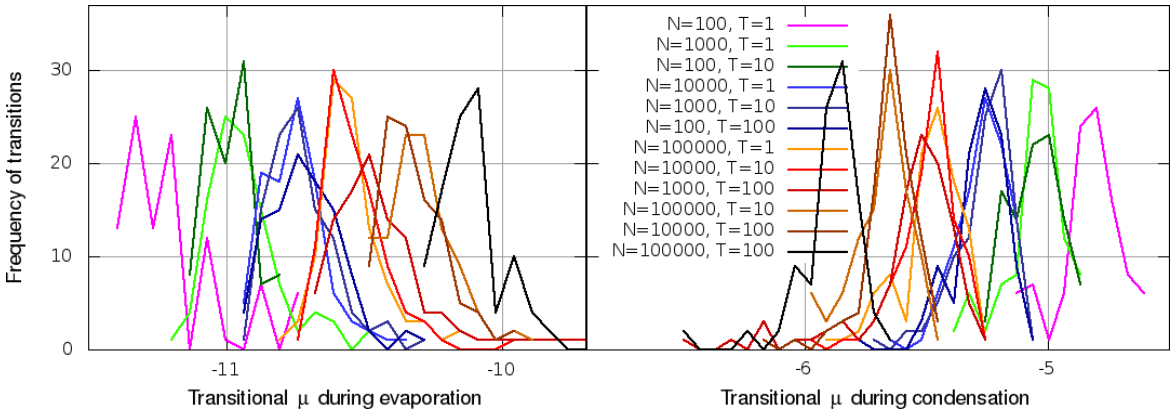

Fig. 6.5 The effect of the system size N and the waiting time T on the sorption behaviour in 3-regular network. Due to the fact that the location in μ of the transitions is stochastic, we present the distributions of μ values at which the transition occurred. 100 independent simulations were completed for each set of parameters. Note that various shades of the same colour approximately overlap since they represent the same value of $T_i \times N_i$.

proceed to discuss the differences between the two representations of porous media and main sources of the discrepancies between the sorption behaviour in the original samples and their random-network representations. Unless specified otherwise the model parameters used in the rest of this Section were as follows: $w_{\mathrm{ff}} = 2$, $w_{\mathrm{mf}} = 3$ for skeletonized soil models and $w_{\mathrm{ff}} = 1$, $w_{\mathrm{mf}} = 2$ for bcc lattice model systems.

6.4.1 Sorption in skeletonized soil samples

The sorption behaviour in skeletonized soil samples is a highly complex phenomenon due to the heterogeneities in topology, the geometric properties of cells and correlations between the two. In order to understand the basic principles that govern sorption in soil, we start the analysis from a simple q-regular network.

Effect of the system size and waiting time

As discussed in Sec. 6.3, the q-regular network system is governed by chain reactions (avalanches), in particular, the irreversible type-B transitions analysed in Sec. 6.3.2. Therefore, the system is stochastic and cannot be described by some single value of μ at which the transition takes place, instead we examine their distribution.

As the kMC waiting time, T, (kMC time interval between the two consecutive changes in the value of μ), increases, the transitions takes place at lower values of μ during adsorption and higher values of μ during desorption (see the changing position of transition μ distributions in Fig. 6.5). Such result is expected, since with longer T, there exist more

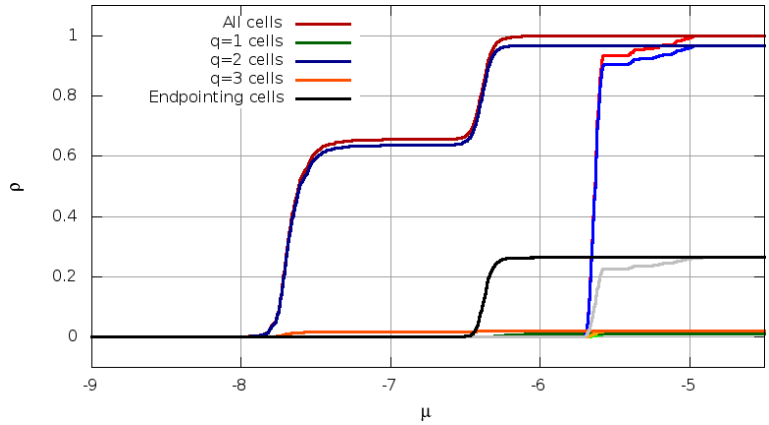

Fig. 6.6 Sorption isotherms in skeletonized soil network. Isotherms for cells characterised by different coordination numbers are shown by curves of different colour as indicated in the figure legend. The lighter and darker shades of the same colour represents adsorption and desorption isotherms respectively. The simulated system is described in the text. $\beta = 0.5$

opportunities for the "hole" in the lattice to appear, and thus for the chain reaction to be initiated, as described in Sec. 6.3.2.

Similar results follow from the same arguments regarding the size of the system. An increase in total number, N, of pore cells in the system leads to increase of the probability for the "hole" to appear. Moreover, the increase in the waiting time and in the size of the system are equivalent not only qualitatively but also quantitatively. Fig. 6.5 shows that the distributions of transition μ for two systems approximately overlap if the product $T \times N$ is the same for both of them (cf. the curves in various shades of the same colour).

Impact of disorder in coordination number

The distribution of coordination numbers for skeletonized soil network is shown in figure 6.2b. The majority of the cells are of coordination number $q = 2$, a small fraction of $q = 1$ and $q = 3$, and very few of higher coordination number. Such composition of cells creates a structure where small fraction of nodes ($q = 3$) or end-points ($q = 1$) are connected by long homogeneous channels of two-coordinated ($q = 2$) cells.

Fig. 6.6 shows the adsorption and desorption isotherms for such a system. The isotherms for the entire network are shown in red, while the isotherms for cells characterised by a particular coordination number are shown in different colours. The two-step behaviour in the desorption isotherm can be explained by evaporation of the end-points at a higher value of μ, which causes evaporation of fluid from all cells in the channels that are leading to an end-point (*end-point channel*). It is a manifestation of the type-A transition described in Sec. 6.3.1. Given the geometrical parameters of the system, the developed theory predicts

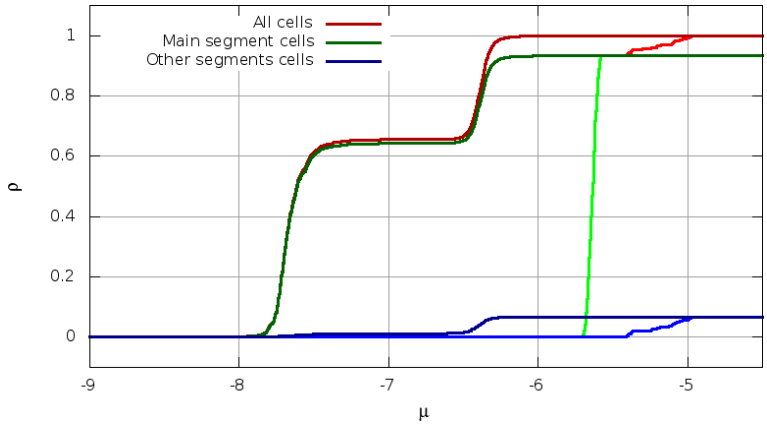

Fig. 6.7 Sorption in skeletonized soil network. Illustration of the ladder-like effect for adsorption caused by disconnected network components. Isotherms for the main network component and the rest of the system are shown in different colours as described in the legend. The lighter and darker shades of the same colour represents adsorption and desorption isotherms, respectively.

the transition to occur around $\mu_1^* = -6.33$ (see Eq. (6.3)), which is exactly what we observe in simulations. In order to confirm that the evaporation of fluid from the cells belonging to the end-point channels is solely responsible for this effect, the sorption isotherms for such cells are plotted in Fig. 6.6 separately in black. Indeed, the location of the step for the black curve is the same as for the desorption isotherm representing the entire system (red line). The remaining fluid part of the system is a network of higher coordination number nodes connected by a series of $q = 2$ cells. Once the end-point cells and end-pointing channels evaporate, the system becomes a 3-regular graph with a small fraction of higher coordination number nodes. Thus the second desorption step is governed by the type-B transition and will be discussed in more detail in the following sections.

Unconnected components of the network

In Fig. 6.6, in addition to the two-step desorption, we observe another ladder-like feature in the adsorption isotherm. Fig. 6.7 analyses exactly the same sorption isotherms as those shown in Fig. 6.6 but from the perspective of different unconnected network components. The real soil network consists of one large cluster of nodes connected by channels both containing approximately 95% of all cells and many small fragments of the network separate from the main cluster. Fig. 6.7 demonstrates that the ladder-like shape of the part of the adsorption isotherm is caused by the unconnected parts of the network filled with the fluid later at larger values of μ than the main cluster (the largest connected component). The unconnected components can be considered as separate systems, because they are not affected by the

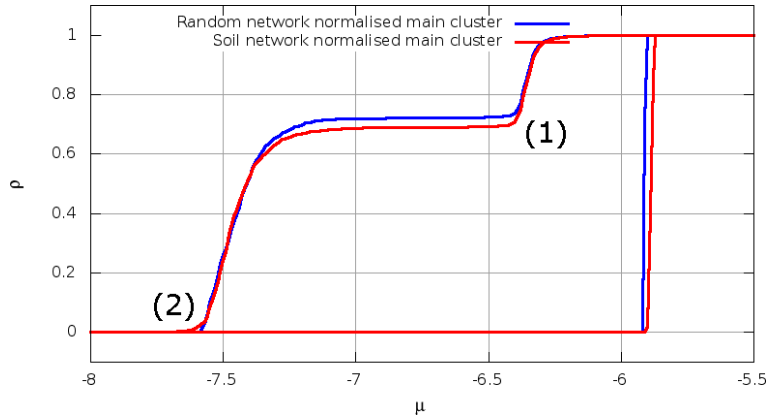

Fig. 6.8 The sorption isotherms for the largest connected component (red curves) and its random-network representation (blue curves). Numbers in parenthesis mark the main differences between the real and random network representation.

avalanches in the main cluster. As we have seen in Sec. 6.4.1 (see Fig. 6.5), the stochastic transition during condensation phase occurs later for smaller systems. This explains the delayed transition in the unconnected smaller parts of the system.

During the desorption phase the unconnected parts of the network mainly contribute to the first step of the process (at $\mu \simeq -6.33$), because such components mainly consist of the end-point channels which empty early during the transition. Therefore, the unconnected components increase the relative size of the first step in the desorption isotherm for the entire system.

We note that a large fraction of the unconnected components are located near the surface of the sample. Therefore, it is expected that in the bulk of the material and thus in large samples the effect of the unconnected components would be significantly diminished.

Comparison of sorption in skeletonized soil models and in their random-network representations

In this Section, we compare sorption in skeletonized soil network and its random-network representation. The random network was created by means of configuration model [212] according to which the nodes, with the same distribution in coordination number as in original skeletonized soil network, are randomly connected to each other. The kMC simulations of sorption were performed for both networks at the same values of parameters and results are displayed in Fig. 6.8.

Apart from the discrepancy during the adsorption, which is expected, because the type-B transition is inherently stochastic, there are only two other differences seen in the diagram:

(1) the size of the first step for desorption isotherm is larger, i.e the plateau value of the density is smaller, in skeletonized soil network,

(2) the very end of the desorption isotherm in soil network has a tail, while in a random network, it finishes sharply.

Both of these features are caused by the spatial correlations of the coordination number in the original skeletonized soil network.

Even though the number of cells with $q = 1$ and thus the number of end-point channels is the same in both systems, there exist regions of their higher density in the original sample. This results in increased probability that two end-point channels emanate from the same $q = 3$ node. Therefore, once the evaporation finishes at these two channels it does not stop at the $q = 3$ node, but continues to penetrate deeper into the system. In this way, not only the end-point channels change the phase, but also some part of the system deeper inside the network does. This can be seen in Fig. 6.6, as the initial drop in overall ρ is larger than the drop in ρ solely due to the end-point channels, which is not the case in the random network (not shown).

The second feature is caused by a similar clustering of nodes with higher coordination numbers. Since a smaller system during the desorption phase is likely to experience the type-B transition later than a larger system (see Sec. 6.3.2,6.4.1), a cluster of nodes with higher coordination numbers connected by short channels remains in fluid phase longer than an equivalent cluster connected by long channels. Therefore, such spatial heterogeneity within a sample leads to this discrepancy.

While the spatial correlations of the end-point channels can be partially explained by the surface effects of finite-size samples and are expected to be reduced in the bulk and in larger samples, the spatial clustering of the nodes with higher coordination numbers is inherent for the real-topology systems and cannot be trivially accounted for by a random-network representation. Having said that, the discrepancy between the sorption isotherms which is caused by such spatial heterogeneity is limited as we have seen in this Section.

Multiple random graphs can be constructed for any particular node degree distribution, however for large enough systems that were investigated in this Chapter the sorption properties do not vary significantly.

6.4.2 Sorption in lattice models

In this Section, we use an alternative structural model for soil sample. This is a coarse-grained bcc-lattice-gas model used for analysis of sorption in aerogel and Vycor (see Chs. 2 and 4). According to this model, a bcc lattice is superimposed onto the soil sample and each node of

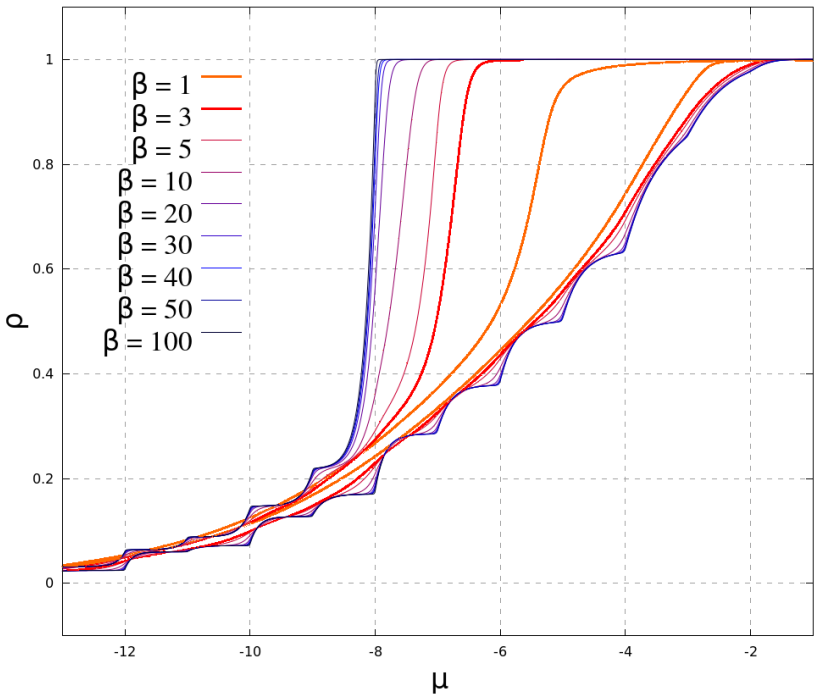

Fig. 6.9 Sorption isotherms simulated by kMC at different temperatures for bcc-lattice model of a soil sample. The lines of different colour refer to isotherms calculated at different temperatures as indicated in the figure legend.

the lattice is associated either with the matrix or pore cell depending on the mean grey-colour intensity at corresponding location in original soil images. Different threshold values could potentially alter the resultant pore structure and therefore the sorption behaviour, however, this aspect has not been investigated in this work. The results of kMC simulations (using the lattice-gas Hamiltonian given by Eq. (1.2) for bcc-lattice model of a soil sample are presented in Fig. 6.9. As follows from even qualitative comparison of isotherms shown in Fig. 6.9 with those found for skeletonized soil networks (see e.g. Fig. 6.6), two models of soil samples exhibit very different sorption behaviour. This is not surprising, because the skeletonized soil networks were designed to explore the effect of the network topology on sorption. In contrast, the lattice models are expected to provide a more complete picture of the sorption behaviour. The isotherms in such samples reflect the effects of the matrix-fluid interaction on the surface area inside the material as well as the impact of large regions of empty (pore) space in addition to the already mentioned topological properties of the material.

Figure 6.9 demonstrates sorption behaviour in a bcc-lattice model of a soil sample. The ladder-like shape of the low-temperature isotherms is the artefact of the lattice model. The adsorption is governed by the type-A transition and thus the system behaves exactly as described in Fig. 6.3 with the correction for varying matrix-fluid interaction. The first cells

to be filled with fluid are those which have the maximum number of matrix-fluid interactions, i.e. the empty pore cells completely surrounded by the matrix. For $w_{\mathrm{mf}} = 2$, this happens at $\mu = -16$ as empty cell in bcc lattice has maximum $q = 8$ nearest matrix neighbours, so that $qw_{\mathrm{mf}} = \mu$. The next increase in amount of adsorbed fluid is observed at $\mu = -14$ when the cells with 7 matrix neighbours are filled with fluid. At $\mu = -13$, the cells with 6 matrix and 1 fluid or 5 matrix and 3 fluid neighbours become filled with fluid, so that $6w_{\mathrm{mf}} + w_{\mathrm{ff}} = \mu$ or $5w_{\mathrm{mf}} + 3w_{\mathrm{ff}} = \mu$, respectively, for $w_{\mathrm{ff}} = 1$. Note that there is no step at $\mu = -15$ (not shown), because all the cells which are already filled with fluid do not have any non-matrix neighbours. Such pattern in behaviour (see the low-temperature adsorption isotherms in Fig. 6.9 exhibiting clear ladder-like pattern) continues on until most of the matrix surface cells are covered by a thin layer of fluid. Adsorption continues and the layer gets thicker, as the empty cells with the most fluid neighbours become filled with fluid. The rest of the system is filled with fluid at $\mu = -4$, i.e. when even empty cells with $q/2 = 4$ (which is a sufficient condition for condensation in the whole system) fluid neighbouring cells become filled with fluid.

The desorption is also governed by similar processes. Initially, the fluid is pinned across the entire empty space. The first to evaporate are the regions of the pore space distant from the matrix cells. At $\mu = -8$, it becomes insufficient to have 8 fluid neighbours to maintain the fluid phase, and thus the avalanche follows evaporating all the cells except for those which have 4 matrix neighbours or 3 matrix and 2 fluid ones, etc. This avalanche is represented in Fig. 6.9 by a steep drop in fluid density which occurs at $\mu \simeq -8$ for sufficiently low temperatures.

The ladder-like shape of the low-temperature isotherms seen in Fig. 6.9 is caused by δ-functional form of distribution of both fluid-fluid and matrix-fluid interaction strength. In real materials, these parameters are continuously distributed and this results in smooth shapes of sorption isotherms. Similar smoothing takes place with increasing temperature as shown in Fig. 6.9. The hysteresis loop has a "triangular" shape and can be classified as of H2 type. Similar shapes of the hysteresis were observed experimentally for sorption in Vycor glass [199–202]. Overall, the lattice-gas model provides qualitatively reasonable picture of sorption in soil and gives a physical insight into the mechanisms that determine the shape of the hysteresis loop and how it varies as we alter the microscopic parameters of porous media.

Together with the new aspects in sorption behaviour, the lattice models for soil samples also bring new challenges regarding our aim to represent soil as a random network of cells with a given distribution of coordination numbers. In the following Sections, we examine the main discrepancies between the original bcc embedded samples and their random-network representations.

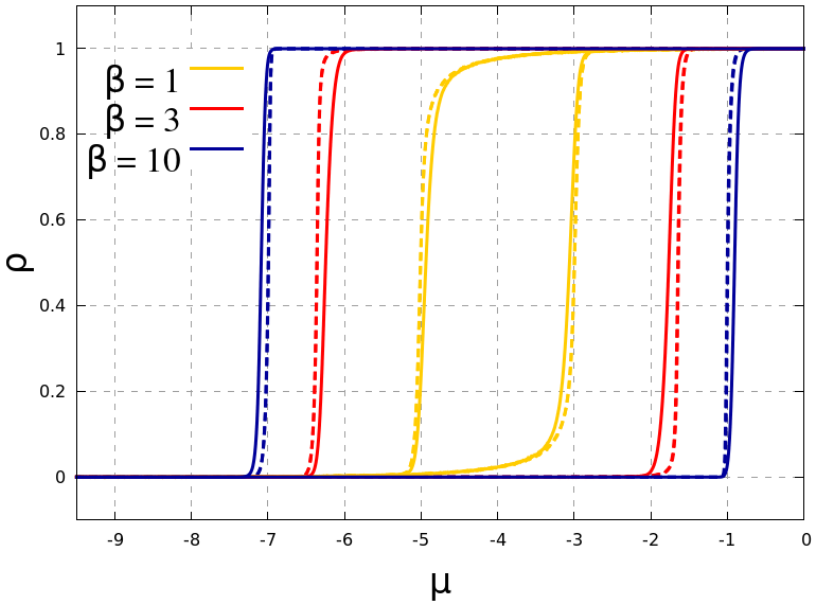

Fig. 6.10 Sorption isotherms in an empty bcc lattice (solid line) and random-network representation of it (dashed line) at different temperatures as indicated in the figure legend.

Sorption in empty bcc lattice

We start by studying the sorption behaviour in an empty bcc lattice, i.e. containing no matrix cells, and in a q-regular random graph ($q = 8$) constructed out of the same number of cells with the same degree distribution.

Fig. 6.10 displays sorption isotherms in an empty bcc lattice (solid lines), and in its random-network representation (dashed lines). Detailed analysis of multiple different systems over a large range of parameter space consistently confirmed the existence of the two main discrepancies between sorption in the original bcc samples and their random network representations.

- The first difference between the sorption behaviour in the two representations originates from the fact that a network based on a regular lattice (in our case bcc lattice) have high clustering while a random network, if it is large enough, has almost none. This feature is particularly important during the initial stages of the phase transition while the process is governed by nucleation. Let us consider the adsorption isotherm. As chemical potential increases, some cells start to be spontaneously filled and then emptied by fluid. Meanwhile, the nearest neighbours of such temporarily fluid cells are more likely to be filled with fluid, as they already have one fluid neighbour. Such process goes on until enough of the neighbouring cells are filled with fluid in such a way that they can remain relatively stable in fluid phase due to the multiple fluid-

fluid interactions. The avalanche of condensation then follows through the rest of the system. In a network based on a regular lattice, a cell is likely to have two nearest neighbours which are also linked through another short pathway. In fact, in a bcc lattice there are multiple 4-member rings going through each cell. At the beginning of the condensation process, this makes it very likely for a relatively stable fluid quadruplet to form, where each cell has at least two fluid-fluid interactions. This is, in turn, followed by formation of fluid sextuplets having 4 fluid connections each, etc. In contrast, a large enough random network is very unlikely to contain small rings, and thus locally, the condensation has to proceed in a tree like structure, with no additional fluid interactions reinforcing it. Moreover, such structure of fluid cells always has multiple cells with only one fluid neighbour that are highly susceptible to evaporation. This is the reason why both condensation and evaporation start later in random networks when compared to the same processes in the lattice-based networks (best seen in Fig. 6.10 for $\beta = 1$ and $\beta = 3$). Clustering in random and lattice-based networks is discussed in more detail in the next subsection.

- The second source of discrepancy is the difference in the average chemical distance between the cells in lattice-based and random networks. While in a lattice-based network an average chemical distance grows as a power law of the number of cells in the network, in a random graph, it grows logarithmically. Therefore, the average chemical distance in a random network is orders of magnitude lower than that in the lattice-based network of the same size. This factor becomes particularly important at the later stage of the phase transition, once a stable nucleation centres have already been established and the chain reaction is propagating through the system. In lattice based networks, the avalanche has to travel through the real space, layer after layer of cells, while in a random network, after the threshold has been reached, the filling avalanche spans the entire network almost instantaneously, since all the cells are chemically close to each another. This phenomenon leads to sharper changes in fluid densities on both isotherms (see Fig. 6.10) for random-network model as compared to the lattice network. Detailed analysis of chemical distances of the relevant networks is presented in the next section 6.4.2.

The described phenomena qualitatively explain all the discrepancies between sorption behaviour in empty lattice-based and random networks. However, obtaining a quantitative description of the differences in the behaviour appears to be complicated and requires further investigation.

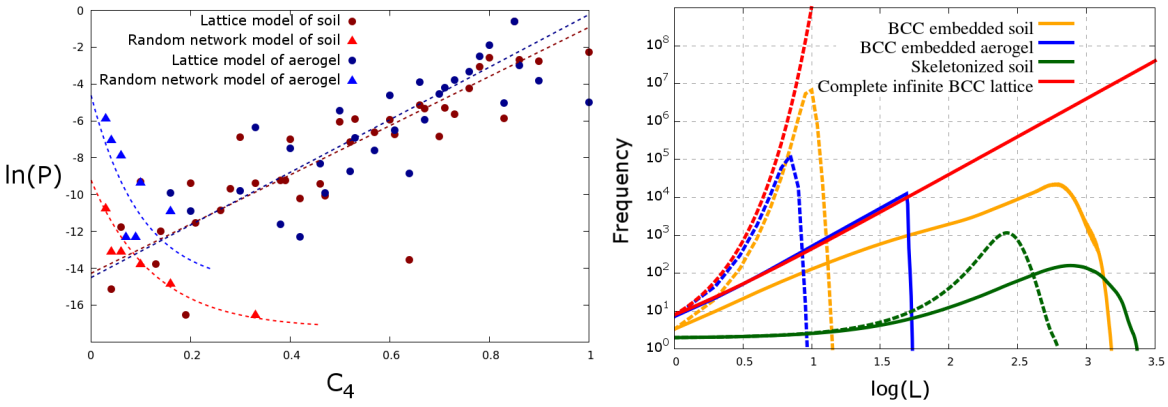

(a) (b)

Fig. 6.11 (a) Second-order clustering coefficient for different type model samples. Regression results are displayed in dashed lines of the corresponding colour to illustrate the difference in clustering behaviour of the studied systems. (b) Distribution of chemical distance between the cells in original (solid lines) and random (dashed lines) network representations of porous media as indicated in the figure legend.

Clustering coefficient and chemical distance

In this Section, we present analysis of network properties which can account for discrepancies between sorption in original networks and their random-network representations.

The standard way [213] to analyse clustering of nodes in a network is by counting how many nearest neighbours of a particular cell, are the nearest neighbours to each other. However, in bcc lattice, by construction, there are no triplets of cells connected all together. Thus, we use the second-order clustering coefficient based on the 4-cell loops instead of triplets. A cell is said to belong to a 4-cell loop if two of its nearest neighbours have another nearest neighbour in common. We define the coefficient, $C_4(i)$, for a cell i as follows,

$$C_4(i) = \frac{2n_4}{q(q-1)}\,, \tag{6.8}$$

where n_4 is the number of successfully formed 4-cell loops that the i-th cell belongs to. For each two distinct neighbours of a cell i, we check whether they have at least one more common neighbour. If they do, we increment n_i. The maximum number of such 4-member loops is the number of distinct pairs of neighbours that the i-th cell has, i.e., for a cell with coordination number q, it is $q(q-1)/2$. Fig. 6.11a presents the probability density function of C_4 for lattice models of two types of porous media (soil and aerogel) systems as well as for their random-network representations. As expected, clustering in random networks is substantially lower than in lattice-based networks. This result reinforces our hypothesis (see

the previous Subsection) about the reasons for differences between sorption in lattices and random networks.

It is important to note that clustering is practically zero in the skeletonized soil networks as well as in their random-network representations. Therefore, the discrepancies between the sorption isotherms in those systems are significantly lower than between bcc embedded samples and their respective random network representations.

The distributions of chemical distance, L, (the shortest path along network edges between two cells) in the network representations of porous media are shown in figure 6.11b. As mentioned before, the chemical distances between the cells in random networks are orders of magnitude shorter than in their lattice counterparts. In fact, the averages are: the mean chemical distance in lattice model of soil, $L_{\text{bcc soil}} = 600 \pm 200$ cell sizes, the mean chemical distance in random-network model of soil, $L_{\text{RND soil}} = 9.4 \pm 0.8$ cells, the mean chemical distance in lattice model of aerogel, $L_{\text{bcc aerogel}} = 40 \pm 10$ cells and the mean chemical distance in random-network model of aerogel, $L_{\text{RND aerogel}} = 6.7 \pm 0.7$ cells. The difference between $L_{\text{bcc aerogel}}$ and $L_{\text{RND aerogel}}$ in the aerogel samples is lower than that for soil purely because the samples that we are studying are smaller. As shown in the figure 6.11b, the distribution of chemical distances in the aerogel is practically identical to the theoretical one for an infinite bcc lattice. It is seen in Fig. 6.11b that the skeletonized soil samples exhibit much more similar chemical distance distributions for their lattice-based and random-network representations. The lack of significant discrepancy is caused by the high proportion of $q = 2$ nodes, which in turn results in a more similar sorption behaviour in the original network and in its random-network representation.

Comparison of sorption in lattice models and in their random-network representations

In a heterogeneous environment such as soil or aerogel, the discussed differences in sorption behaviour between lattice and random-network representations manifest themselves stronger than in an empty space. This is caused by randomising the distribution of the matrix, as the random-network representations of the samples are constructed.

Figure 6.12 presents the comparison between the sorption isotherms for lattice models of soil and aerogel samples, and their random-network models. As seen in Fig. 6.10, both condensation and evaporation in lattice models start ahead of the respective processes in the random networks due to higher clustering and thus easier nucleation, but proceeds slower due to the longer chemical distances between the cells. However, in the systems with the matrix-fluid interaction there exists a significantly stronger cause for such discrepancies. Sharper transitions in the random network representations of the samples are caused by the extra randomisation making the new system relatively more homogeneous than the original

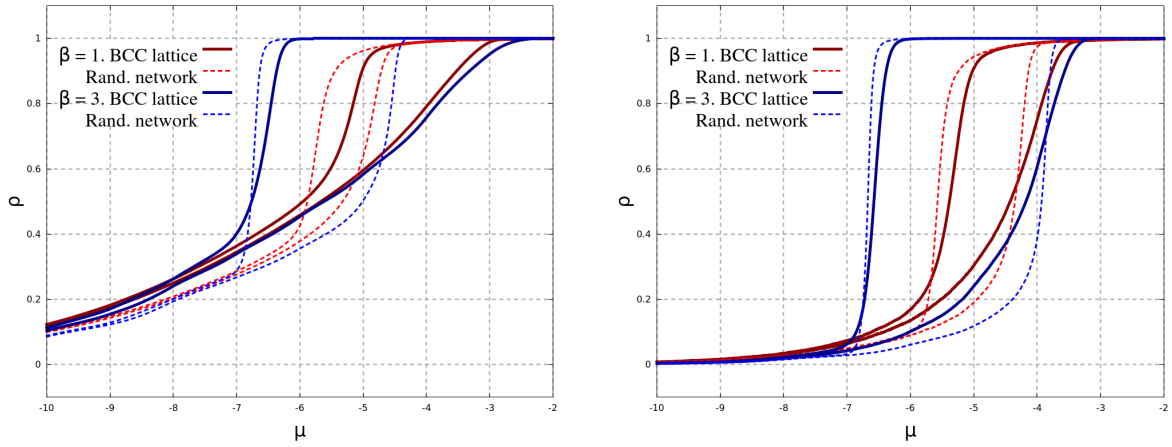

(a) (b)

Fig. 6.12 Comparison between the sorption isotherms of in lattice (solid curves) and random-network (dashed curves) models of (a) soil and (b) aerogel for $\beta = 1$ (red curves) and $\beta = 3$ (blue curves).

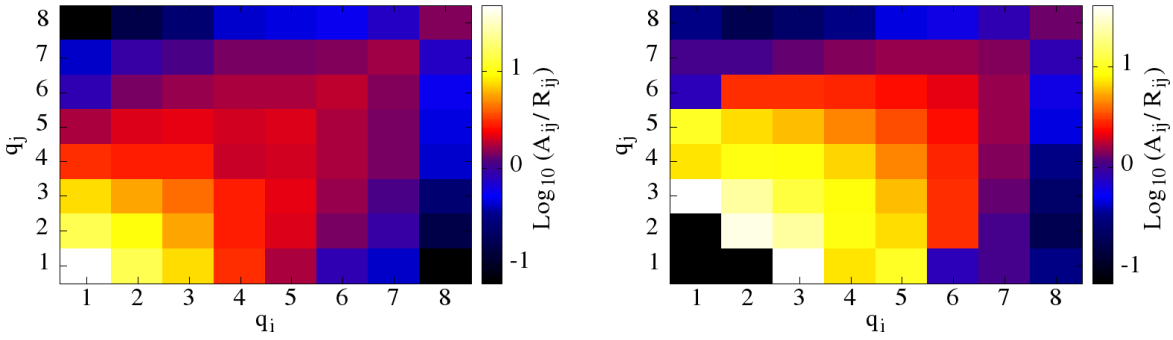

(a) Lattice model of soil (b) Lattice model of aerogel

Fig. 6.13 The ratio of the number of connections, A_{ij}, between the cells in lattice models with coordination numbers, q_i and q_j and number of similar connections R_{ij} in a random-network counterparts of lattice models is shown by different colours in log-scale. The colour scheme is represented by the vertical bar.

sample. In the original sample, various regions of the same system becomes susceptible to propagation of avalanches at different values of μ, thus smearing out the transition over a wider range in μ. In contrast, the reconstructed random-network systems are significantly more self-similar and thus all parts of them become susceptible to the avalanche propagation at similar values of μ, hence sharpening the transition.

The hypothesis about the presence of correlations in the connectivity between nodes of lattice models characterised by different coordination numbers was explicitly tested. Fig. 6.13 presents the ratio of the number, A_{ij}, of connections between the cells in lattice models with coordination number q_i and cells with coordination number q_j and the theoretically expected number, R_{ij}, of such connections in a random network characterised by the same distribution of coordination numbers. It follows from this figure that the cells with low coordination number tend to connect with other cells also having low coordination number and vice versa. This is expected as the bulk of the void space inside the pores contains many cells interconnected together, and the cells neighbouring the matrix are near the other cells that are neighbouring the matrix as well, since the structure of the matrix is continuous in space. Therefore, it can be possible to overcome the discrepancy in sorption behaviour which was caused by this phenomenon if the random network is constructed according to the known correlations in the original system. Such future improvement could potentially be very valuable since it is possible to adjust the technique described in Ref. [98] to account for such correlations when calculating sorption isotherms for a random network analytically.

6.5 Conclusions

In this Chapter, sorption of fluid in several models of real materials such as soil and aerogel was studied. Several types of structural models were analysed. In particular, experimentally obtained images of soil samples were processed and converted to

(i) skeletonized soil networks in which pore space is mimicked by chains of two-coordinated cells connected to each other by higher coordination number cells. Skeletonized systems have been taken from the previous work [209] and have not been done as part of this thesis;

(ii) lattice models in which the original soil images represented by 3d simple cubic lattice of voxels (characterised by grey colour of different intensity) were coarse grained and mapped onto bcc lattice of cells (with each cell, depending on the dominant colour of voxels in the cell, representing either matrix or pore space of original soil sample).

The structural models of aerogel were constructed using the DLCA algorithm on bcc lattice (see Sec. 1.1.2) and thus, by construction, are lattice models. In addition to skeletonized and lattice networks, we also created random-network counterparts for each structural model by randomly connecting nodes characterised by the same distribution of coordination number as in the original network. The random-network models served for testing the effects of topological disorder and applicability of the theoretical concepts [98] developed for description of sorption in random networks.

The sorption of fluid in all these models was studied numerically by means of kMC. We found that continuous isotherms observed at relatively high temperatures exhibit the hysteresis loop with decrease in temperature. The shape of the loop was found to depend on type of porous media. The abrupt changes in density found in the hysteresis regime were analysed theoretically and described in terms of A- and B-type transitions. Theoretical estimates for the values of chemical potential at which such changes occur were established and confirmed numerically. The qualitative understanding of the hysteresis phenomenon was achieved and used further in examining the impact of the microscopic parameters of the system on the macroscopic sorption behaviour.

The comparative analysis of sorption in the original systems and in their random-network counterparts was undertaken. In all cases, certain differences were observed. They can be summarised as follows:

- Sorption in skeletonized soil samples and their random-network representations differs due to spatial heterogeneity in the distribution of nodes with a certain coordination number in the skeletonized network. The discrepancy is partially caused by the surface effects of finite-size samples, and hence it would be mitigated in the bulk of the material. The scale of the remaining discrepancies is limited.

- The original lattice models and their random-network representations have very different clustering coefficient and chemical distance distributions of the networks. This causes the delay in significant change of fluid density with changing μ in the random-network representation of the system, and once begun, it proceeds quicker than in the original systems.

- Finally, randomizing the network with cells that have some matrix-fluid interaction makes the newly constructed system more homogeneous than the original one, in which the matrix cells are distributed not entirely randomly across the space. As a result, the transitions appear to be sharper in the random-network representations of the soil and aerogel systems than in the original networks.

We conclude that the random-network representation of the discussed porous media can be used to obtain qualitative understanding of the sorption behaviour in the original system, and thus the analytical techniques developed in Ref. [98] may be of great value. The sorption behaviour is more similar in skeletonized soil networks and their random-network representations. However, by construction, sorption behaviour in both such systems lacks certain features, such as appearance of an initial fluid layer around the matrix cells. Thus the skeletonized soil systems are suitable for examining some particular aspects of the sorption behaviour, but the results cannot be straightforwardly compared with the experimentally found isotherms. To achieve quantitative understanding of the differences between the sorption processes in the lattice models and their random-network representations, further investigation is required. One of the promising areas for future work is to examine sorption in the systems constructed accounting for correlations in connectivity between the nodes of different degree.

Finally, this study has strongly indicated that a more detailed investigation of the impact of the global topology on sorption behaviour in the system could provide a deeper understanding and more accurate quantitative description of the phenomena.

Chapter 7

Conclusions

In this thesis, we have studied sorption phenomena in disordered porous media within the lattice-gas model. We developed and applied multiple numerical methods to examine such systems and their behaviour.

In Ch. 2 we introduced a technique designed to visualize the state space of a lattice-gas system. The technique is based on a projection of an N-dimensional state space onto a 2-dimensional plane. It has greatly helped to gain intuition about the systems of interest and study their behaviour throughout the entire duration of the project. We started Ch. 2 by the study of relatively small systems using an exact microstate-enumeration method. With the complete knowledge of the system that such method provides, we have investigated the relationships between the structure of a system, external conditions and the resulting equilibrium state. A detail description of how the low-energy state-space boundary depends on the microstructure of the system has been provided. Furthermore, we have discussed the relationship between the shape of the state space and the morphology of the possible sorption transitions within the system. We concluded that the concave low-energy state-space boundary over some range of values of absorbed fluid density, ρ, leads to a multi-peak equilibrium state probability distribution in ρ. Hence, at sufficiently low temperatures, such systems can undergo abrupt phase transitions manifested by macroscopic avalanche-like events. In the second part of Ch. 2, we have introduced kinetic-Monte-Carlo-based techniques to investigate systems that are too large to be handled by the exact methods. We studied several silica aerogel models of varying porosity. The basic trends established in the first part of this chapter have been confirmed for the larger systems. The methodology to obtain upper and lower state-space boundaries have been introduced. Such boundaries provide an excellent reference point for the visualized paths that the system takes across its state space during the simulation. Finally, several methods to obtain the degeneracy of the states in such systems have been proposed, developed and tested on silica aerogel models.

The jump-walking Monte-Carlo (JW MC) algorithm has been revisited and updated in Ch. 3. The original algorithm was designed to tackle quasi-nonergodicity problem and was used to study dynamics of small atomic clusters. However, in its previous form it had several drawbacks. In Ch. 3, we have addressed the known issues, developed its implementation considering the current state of computing technology, and proved its applicability to large systems with binary-state dynamics, such as Ising and lattice-gas models. We have shown that the method is capable of finding equilibrium isotherms in systems with complex free-energy landscape, where the standard single-flip MC based approaches fail. Furthermore, unlike the majority of the contemporary methods used to study the quasi-nonergodic systems, the JW algorithm does not rely on simultaneous simulation of the system at multiple different temperatures. Hence, the computing resources needed are substantially lower. In Chs. 3-4, we have applied the JW MC algorithm to the models of silica aerogel and Vycor-glass. The equilibrium isotherms in several Vycor systems have been examined in detail in Ch. 4. We have concluded that system behaviour can be conceptually categorised into two stages. The system undergoes the first stage at relatively small values of chemical potential, μ, or vapour pressure, when the fluid continuously accumulates in (or recedes from) the proximity of the matrix surface. During the second stage, at higher values of μ, the equilibrium isotherms exhibit ladder-like structure associated with multiple localised avalanches. The size of the regions that change their phase during a single avalanche depend on the characteristic space scale of the pore structure and is independent of the system size. The macroscopic, system spanning avalanches are not observed, hence such systems do not exhibit first-order phase transition. The lower projected-state-space boundary of Vycor systems retains positive curvature in the range of relevant values of wettability, thus supporting the conclusions.

In Ch. 5, we have described the real-time interactive MC simulation engine, that was developed to help us study sorption phenomena. The engine was designed to provide a wide range of data about the simulated systems. In addition, it allows the user to manipulate the values of temperature and chemical potential, while observing system response in real-time. Such direct feedback loop allowed us to develop a deeper intuition of the sorption phenomena, and provided valuable insights into system behaviour, in general. The application was developed to be widely applicable and easily customisable. It is relatively straightforward to expand engines functionality, add to the set of available systems and even alter the core of MC simulation process itself. Finally, the application can serve as a platform to develop and test prototypes of complex MC algorithms, as it has in the case of JW MC.

An off-lattice model for a disordered porous media, i.e. skeletonized soil, is introduced in Ch. 6. The model is based on a 3d-network of interconnected channels, which are mainly composed of 2-coordinated cells (with higher coordination number cells at the intersections

between the channels). Sorption behaviour in such systems is examined by kinetic MC simulations, developing a qualitative description of the phenomena. Motivated by the recently developed analytical techniques [98], which model sorption phenomena in the systems of random graphs, we have explored the possibility to represent the models of real disordered porous media by particularly constructed random networks. The sorption phenomena in different representations of soil and silica aerogel models is examined in detail. We conclude that random-network representations of skeletonized soil can be used to obtain qualitative understanding of system behaviour in the original model. The recommendations are given for further directions of research to address the remaining discrepancies between the sorption behaviour in the bcc lattice models of disordered media and their random-network representations.

Looking ahead, we suggest the following potentially fruitful directions of future work.

(i) Further development of the techniques designed to obtain the degeneracy of the macrostates within the projected state space of a system. In particular, testing and improving the algorithm performance in the relatively low degeneracy region near the lower projected-state-space boundary.

(ii) It is fundamentally challenging to compare the free-energies of two metastable states that are not naturally sampled simultaneously. The standard approaches are based on either multiple replicas of the system or reshaping the free-energy landscape itself. JW MC method uses the ”jump” process between the macrostates (i.e. particularly defined subsets of microstates that partition the state space) of the system to link the states of interest. Broader application of such technique based on adding new connections within the configurational space of the system could lead to the development of further novel sampling algorithms.

(iii) Statistical analysis of lower projected-state-space boundary, and how it depends on the properties of the model structure, such as porosity and wettability in aerogel and Vycor systems.

(iv) Statistical analysis of the avalanches that comprise the equilibrium isotherms obtained by JW MC method could provide important insights into system behaviour.

(v) The interactive simulation engine would benefit from decoupling of the workload dedicated to the 3d view rendering from the maintenance of the graphical user interface. Furthermore, the performance of the 3d-view screen could be significantly improved by some application of OpenGL and the use of the graphics card.

In Ch. 1, we claimed that novel solutions to the contemporary challenges in the studied field of research are likely to arise from the physical insights into the configurational space of the system of interest. Thus, one of the primary goals of this project was to deepen our understanding about the state space for the sorption phenomena within the lattice-gas model. We approached this task by developing multiple numerical tools to visualize, analyse and measure various characteristics of the state space of the studied disordered porous media. As predicted, the gained understanding led to the development of new and improvement of existing algorithms as well as several physical results about the studied phenomena. We highly recommend to continue further investigation of the systems with complex free-energy landscapes using the approach based on an in-depth study of the configurational space of the relevant system.

References

[1] W. Janke, editor. *Rugged Free Energy Landscapes: Common Computational Approaches to Spin Glasses, Structural Glasses and Biological Macromolecules*. Lect. Notes Phys. Springer, Berlin, 2008.

[2] K. Binder and A. P. Young. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.*, 58:801, 1986.

[3] E. Marinari, G. Parisi, and J. J. Ruiz-Lorenzo. Numerical simulations of spin glass systems. In A.P. Young, editor, *Spin Glasses and Random Fields*, page 59. World Scientific, Singapore, 1998.

[4] D. P. Landau and K. Binder. *A guide to monte carlo simulations in statistical physics*. Cambridge University Press, 3-rd edition edition, 2005.

[5] D. Wales. Surveying a complex potential energy landscape: Overcoming broken ergodicity using basin-sampling. *Chem. Phys. Lett.*, 584:1, 2013.

[6] Ryoichi Yamamoto and Walter Kob. Replica-exchange molecular dynamics simulation for supercooled liquids. 61:5473–5476, 2000.

[7] F. Ritort and P. Sollich. Glassy dynamics of kinetically constrained models. *Adv. Phys.*, 52:219, 2003.

[8] F. Detcheverry, M. L. Rosinberg, and G. Tarjus. Metastable states and t=0 hysteresis in the random-fieldising model on random graphs. *The European Physical Journal B - Condensed Matter and Complex Systems*, 44:327–343, 2005.

[9] M L Rosinberg, G Tarjus, and F J Pérez-Reche. Stable, metastable and unstable states in the mean-field random-field ising model at t = 0. *Journal of Statistical Mechanics: Theory and Experiment*, 10:P10004, 2008.

[10] F. J. Pérez-Reche, M. L. Rosinberg, and G. Tarjus. Numerical approach to metastable states in the zero-temperature random-field ising model. *Phys. Rev. B*, 77:064422, 2008.

[11] E. Kierlik, P. A. Monson, M. L. Rosinberg, L. Sarkisov, and G. Tarjus. Capillary condensation in disordered porous materials: Hysteresis versus equilibrium behavior. *Phys. Rev. Lett.*, 87:055701, 2001.

[12] Hyung-June Woo and P. A. Monson. Phase behavior and dynamics of fluids in mesoporous glasses. *Phys. Rev. E*, 67:041207, 2003.

[13] Angel E. Garcia and Jose N. Onuchic. Folding a protein in a computer: An atomic description of the folding/unfolding of protein a. 100:13898–13903, 2003.

[14] Jinzen Ikebe, Koji Umezawa, and Junichi Higo. Enhanced sampling simulations to construct free-energy landscape of protein–partner substrate interaction. *Biophys. Rev.*, 8:45, 2016.

[15] Philip A. Romero and Frances H. Arnold. Exploring protein fitness landscapes by directed evolution. *Nat Rev Mol Cell Biol*, 10:866–876, 2009.

[16] José Nelson Onuchic and Peter G Wolynes. Theory of protein folding. *Current Opinion in Structural Biology*, 14:70–75, 2004.

[17] Adrian W.R. Serohijos, Zilvinas Rimas, and Eugene I. Shakhnovich. Protein biophysics explains why highly abundant proteins evolve slowly. *Cell Reports*, 2(2):249–256, 2012.

[18] C. Chipot and A. Pohorille. *Free Energy Calculations. Theory and Applications in Chemistry and Biology*. Springer, 2007.

[19] Didier Sornette. *Critical phenomena in natural sciences: chaos, fractals, selforganization and disorder: concepts and tools*. Springer Science & Business Media, 2006.

[20] J. D. van der Waals. Thesis, 1873.

[21] J Clerk Maxwell. Statique expérimentale et théorique des liquides soumis aux seules forces moléculaires, 1874.

[22] J. P. Neirotti, David L. Freeman, and J. D. Doll. Approach to ergodicity in monte carlo simulations. *Phys. Rev. E*, 62:7445–7461, 2000.

[23] Kerson Huang. *Statistical Mechanics, 2nd Edition*. John Wiley & Sons, New York, 1987.

[24] Bo Sundman and John Ågren. A regular solution model for phases with several components and sublattices, suitable for computer applications. *Journal of Physics and Chemistry of Solids*, 42(4):297 – 301, 1981.

[25] Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, 1925.

[26] J. Frenkel. *Kinetic Theory of Liquids*. Oxford University Press , Oxford, 1946.

[27] R Evans. Fluids adsorbed in narrow pores: phase equilibria and structure. *J. of Phys.: Cond. Matt.*, 2:8989–9007, 1990.

[28] S. J. Gregg and K. S. W. Sing. *Adsorption, Surface Area and Porosity*. Academic Press, second edition, 1982.

[29] L. D. Gelb, K. E. Gubbins, R. Radhakrishnan, and M. Sliwinska-Bartkowiak. Phase separation in confined systems. *Rep. Prog. Phys.*, 62:1573–1659, 1999.

[30] T. Horikawa, D.D. Do, and D. Nicholson. Capillary condensation of adsorbates in porous materials. *Advances in Colloid and Interface Science*, 169:40–58, 2011.

[31] B. Coasne, A. Galarneau, R. J. M. Pellenq, and F. Di Renzo. Adsorption, intrusion and freezing in porous silica: the view from the nanoscale. *Chem. Soc. Rev.*, 42:4141, 2013.

[32] R. J. Baxter. *Exactly Solved Models in Statistical Mechanics*. Academic Press, 1982.

[33] H. E. Stanley. *Introduction to phase transitions and critical phenomena*. International series of monographs on physics. Oxford University Press, 1971.

[34] Lars Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Physical Review*, 65(3-4):117, 1944.

[35] Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Academic press, 2001.

[36] Robert H. Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607–2609, 1986.

[37] Robert H. Swendsen and Jian-Sheng Wang. Nonuniversal critical dynamics in monte carlo simulations. 58:86–88, 1987.

[38] Ulli Wolff. Collective monte carlo updating for spin systems. 62, 1989.

[39] B. A. Berg and T. Neuhaus. Multicanonical algorithms for 1st order phase transitions. 267:249–253, 1991.

[40] B. A. Berg and T. Neuhaus. Multicanonical ensemble: A new approach to simulate first-order phase transitions. *Phys. Rev. Lett.*, 68:9–12, 1992.

[41] J. Lee. New monte carlo algorithm: Entropic sampling. 71:211–214, 1993.

[42] W.H. Hao and H.A. Scheraga. Monte carlo simulation of a first-order transition for protein folding. *J. Phys. Chem.*, 98:4940–4948, 1994.

[43] M. Mezei. Adaptive umbrella sampling - self-consistent determination of the non-boltzmann bias. *J. Comput. Phys*, 68:237–248, 1987.

[44] C. Bartels and M. Karplus. Probability distributions for complex systems: Adaptive umbrella sampling of the potential energy. *J. Phys. Chem. B*, 102:865–880, 1998.

[45] A.P. Lyubartsev, A.A. Martinovski, S.V. Shevkunov, and P.N. Vorontsov-Velyaminov. New approach to monte carlo calculation of the free energy - method of expanded ensemble. *J. Chem. Phys.*, 96:1776–1783, 1992.

[46] E. Marinari and G. Parisi. Simulated tempering - a new monte carlo scheme. *Europhys. Lett.*, 19:451–458, 1992.

[47] K. Hukushima and K. Nemoto. *J. Phys. Soc. Jpn.*, 65:1604, 1996.

[48] David J. Earl and Michael W. Deem. Parallel tempering: Theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.*, 7:3910–3916, 2005.

[49] K. Kimura and K. Taki. Time-homogeneous parallel annealing algorithm. *Proceedings of the 13th IMACS World Congress on Computation and Applied Mathematics*, 2:827–828, 1991.

[50] M. C. Tesi, E. J. Janse van Rensburg, E. Orlandini, and S. G. Whittington. Monte carlo study of the interacting self-avoiding walk model in three dimensions. *J. Stat. Phys.*, 82:155–181, 1996.

[51] C. T. Kresge, M. E. Leonowicz, W. J. Roth, J. C. Vartuli, and J. S. Beck. Ordered mesoporous molecular sieves synthesized by a liquid-crystal template mechanism. *Nature*, 359:710–712, 1992.

[52] Dongyuan Zhao, Qisheng Huo, Jianglin Feng, Bradley F. Chmelka, and Galen D. Stucky. Nonionic triblock and star diblock copolymer and oligomeric surfactant syntheses of highly ordered, hydrothermally stable, mesoporous silica structures. *Journal of the American Chemical Society*, 120:6024–6036, 1998.

[53] Akira Taguchi and Ferdi Schüth. Ordered mesoporous materials in catalysis. *Microporous and Mesoporous Materials*, 77:1–45, 2005.

[54] J. S. Beck, J. C. Vartuli, W. J. Roth, M. E. Leonowicz, C. T. Kresge, K. D. Schmitt, C. T. W. Chu, D. H. Olson, E. W. Sheppard, S. B. McCullen, J. B. Higgins, and J. L. Schlenker. A new family of mesoporous molecular sieves prepared with liquid crystal templates. *Journal of the American Chemical Society*, 114:10834–10843, 1992.

[55] Arne Thomas, Helmut Schlaad, Bernd Smarsly, and Markus Antonietti. Replication of lyotropic block copolymer mesophases into porous silica by nanocasting: learning about finer details of polymer self-assembly. *Langmuir*, 19:4455–4459, 2003.

[56] Daibin Kuang, Torsten Brezesinski, and Bernd Smarsly. Hierarchical porous silica materials with a trimodal pore system using surfactant templates. *Journal of the American Chemical Society*, 126:10534–10535, 2004.

[57] Bernd Smarsly, Christine Göltner, Markus Antonietti, Wilhelm Ruland, and Ernst Hoinkis. Sans investigation of nitrogen sorption in porous silica. *The Journal of Physical Chemistry B*, 105:831–840, 2001.

[58] Steven M. Jones. Aerogel: Space exploration applications. *Journal of Sol-Gel Science and Technology*, 40:351–357, 2006.

[59] Sam Niblett. A computational study of capillary condensation in disordered porous media. PartIII project, Cambridge.

[60] David R. Lide. *CRC Handbook of Chemistry and Physics*. CRC Press, 88 edition, 2007.

[61] LLNL. Lab's aerogel sets world record. *Science & Technology Review, Lawrence Livermore National Laboratory*, 2003.

[62] Dale W. Schaefer and Keith D. Keefer. Structure of random porous materials: Silica aerogel. *Phys. Rev. Lett.*, 56:2199–2202, 1986.

[63] René Vacher, Thierry Woignier, Jacques Pelous, and Eric Courtens. Structure and self-similarity of silica aerogels. *Phys. Rev. B*, 37:6500–6503, 1988.

[64] J L Rousset, A Boukenter, B Champagnon, J Dumas, E Duval, J F Quinson, and J Serughetti. Granular structure and fractal domains of silica aerogels. *Journal of Physics: Condensed Matter*, 2:8445, 1990.

[65] Andrzej B. Jarzbski, Jarosław Lorenc, and Lucjan Pajak. Surface fractal characteristics of silica aerogels. *Langmuir*, 13:1280–1285, 1997.

[66] D. A. Donatti, D. R. Vollet, A. Ibañez Ruiz, A. Mesquita, and T. F. P. Silva. Mass fractal characteristics of silica sonogels as determined by small-angle x-ray scattering and nitrogen adsorption. *Phys. Rev. B*, 71:014203, 2005.

[67] T. A. Witten and L. M. Sander. Diffusion-limited aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.*, 47:1400–1403, 1981.

[68] A. Soleimani Dorcheh and M.H. Abbasi. Silica aerogel; synthesis, properties and characterization. *Journal of Materials Processing Technology*, 199:10–26, 2008.

[69] Anwar Hasmy, Marie Foret, Eric Anglaret, Jacques Pelous, René Vacher, and Rémi Jullien. Small-angle neutron scattering of aerogels: simulations and experiments. *Journal of Non-Crystalline Solids*, 186:118–130, 1995.

[70] Paul Meakin. Formation of fractal clusters and networks by irreversible diffusion-limited aggregation. *Phys. Rev. Lett.*, 51:1119–1122, 1983.

[71] F. Detcheverry, E. Kierlik, M. L. Rosinberg, and G. Tarjus. Local mean-field study of capillary condensation in silica aerogels. *Phys. Rev. E*, 68:061504, 2003.

[72] Rafael Salazar and Lev D. Gelb. An investigation of the effects of the structure of gel materials on their adsorptive properties using a simple lattice–gas model. *Molecular Physics*, 102:1015–1030, 2004.

[73] Anwar Hasmy and Rémi Jullien. Percolation in cluster-cluster aggregation processes. *Phys. Rev. E*, 53:1789–1794, 1996.

[74] M.E. Kainourgiakis, E.S. Kikkinides, and A.K. Stubos. Diffusion and flow in porous domains constructed using process-based and stochastic techniques. *Journal of Porous Materials*, 9:141–154, 2002.

[75] Thomas H Elmer. Porous and reconstructed glasses. *ASM International, Engineered Materials Handbook.*, 4:427–432, 1991.

[76] Lev D. Gelb and K. E. Gubbins. Pore size distributions in porous glasses: A computer simulation study. *Langmuir*, 15:305–308, 1999.

[77] P. Levitz. Off-lattice reconstruction of porous media: critical evaluation, geometrical confinement and molecular transport. *Advances in Colloid and Interface Science*, 76:71–106, 1998.

[78] C. L. Y. Yeong and S. Torquato. Reconstructing random media. *Phys. Rev. E*, 57:495–506, 1998.

[79] C. L. Y. Yeong and S. Torquato. Reconstructing random media. ii. three-dimensional media from two-dimensional cuts. *Phys. Rev. E*, 58:224–233, 1998.

[80] Y. Jiao, F. H. Stillinger, and S. Torquato. A superior descriptor of random textures and its predictive capacity. *Proceedings of the National Academy of Sciences*, 106:17634–17639, 2009.

[81] C. J. Gommes, Y. Jiao, and S. Torquato. Microstructural degeneracy associated with a two-point correlation function and its information content. *Phys. Rev. E*, 85:051140, 2012.

[82] Anthony P. Roberts. Statistical reconstruction of three-dimensional porous media from two-dimensional images. *Phys. Rev. E*, 56:3203–3212, 1997.

[83] Kenneth K. S. Lau, José Bico, Kenneth B. K. Teo, Manish Chhowalla, Gehan A. J. Amaratunga, William I. Milne, Gareth H. McKinley, and Karen K. Gleason. Superhydrophobic carbon nanotube forests. *Nano Letters*, 3:1701–1705, 2003.

[84] John S. Gierke, Neil J. Hutzler, and David B. McKenzie. Vapor transport in unsaturated soil columns: Implications for vapor extraction. *Water Resources Research*, 28:323–335, 1992.

[85] Masato Kumagai and Gary L. Messing. Controlled transformation and sintering of a boehmite sol-gel by -alumina seeding. *Journal of the American Ceramic Society*, 68:500–505, 1985.

[86] Brandon L. Weeks, Mark W. Vaughn, and James J. DeYoreo. Direct imaging of meniscus formation in atomic force microscopy using environmental scanning electron microscopy. *Langmuir*, 21:8096–8098, 2005.

[87] U. Srinivasan, M. R. Houston, R. T. Howe, and R. Maboudian. Alkyltrichlorosilane-based self-assembled monolayer films for stiction reduction in silicon micromachines. *Journal of Microelectromechanical Systems*, 7:252–260, 1998.

[88] Ebrahim Shahraeeni and Dani Or. Pore-scale evaporation-condensation dynamics resolved by synchrotron x-ray tomography. *Phys. Rev. E*, 85:016317, 2012.

[89] K. Morishige and M. Shikimi. Adsorption hysteresis and pore critical temperature in a single cylindrical pore. *The Journal of Chemical Physics*, 108:7821–7824, 1998.

[90] Chunyan Fan, Yonghong Zeng, D. D. Do, and D. Nicholson. An undulation theory for condensation in open end slit pores: critical hysteresis temperature & critical hysteresis pore size. *Phys. Chem. Chem. Phys.*, 16:12362–12373, 2014.

[91] D. Wallacher, N. Künzner, D. Kovalev, N. Knorr, and K. Knorr. Capillary condensation in linear mesopores of different shape. *Phys. Rev. Lett.*, 92:195704, 2004.

[92] J. R. Edison and P. A. Monson. Dynamics of capillary condensation in lattice gas models of confined fluids: A comparison of dynamic mean field theory with dynamic monte carlo simulations. *J. Chem. Phys.*, 138:234709, 2013.

[93] T. P. Handford, F. J. Perez-Reche, and S. N. Taraskin. Capillary condensation in one-dimensional irregular confinement. *Phys. Rev. E*, 88:012139, 2013.

[94] Apollo P. Y. Wong and M. H. W. Chan. Liquid-vapor critical point of ^{4}He in aerogel. *Phys. Rev. Lett.*, 65:2567–2570, 1990.

[95] T. Lambert, C. Gabay, L. Puech, and P. E. Wolf. Optical study of 4he condensation into a silica aerogel. *J. Low Temp. Phys.*, 134:293–302, 2004.

[96] F. Detcheverry, E. Kierlik, M. L. Rosinberg, and G. Tarjus. Helium condensation in aerogel: Avalanches and disorder-induced phase transition. *Phys. Rev. E*, 72:051506, 2005.

[97] Geoffroy J. Aubry, Fabien Bonnet, Mathieu Melich, Laurent Guyon, Panayotis Spathis, Florence Despetis, and Pierre-Etienne Wolf. Condensation of helium in aerogel and athermal dynamics of the random-field ising model. *Phys. Rev. Lett.*, 113:085301, 2014.

[98] T. P. Handford, A. Dear, F. J. Perez-Reche, and S. N. Taraskin. Effect of disorder on condensation in the lattice gas model on a random graph. *Phys. Rev. E*, 90:012144, 2014.

[99] Michal Kruk, Mietek Jaroniec, Ji Man Kim, and Ryong Ryoo. Characterization of highly ordered mcm-41 silicas using x-ray diffraction and nitrogen adsorption. *Langmuir*, 15:5279–5284, 1999.

[100] B. Coasne, A. Galarneau, F. Di Renzo, and R. J. M. Pellenq. Molecular simulation of nitrogen adsorption in nanoporous silica. *Langmuir*, 26:10872–10881, 2010.

[101] Kunimitsu Morishige and Masataka Ito. Capillary condensation of nitrogen in mcm-41 and sba-15. *The Journal of Chemical Physics*, 117:8036–8041, 2002.

[102] Philippe Trens, Nathalie Tanchoux, Anne Galarneau, Daniel Brunel, Bice Fubini, Edoardo Garrone, François Fajula, and Francesco Di Renzo. A macrothermodynamic approach to the limit of reversible capillary condensation. *Langmuir*, 21:8560–8564, 2005.

[103] Michal Kruk and Mietek Jaroniec. Determination of mesopore size distributions from argon adsorption data at 77 k. *The Journal of Physical Chemistry B*, 106:4732–4739, 2002.

[104] Leonard H. Cohan. Sorption hysteresis and the vapor pressure of concave surfaces. *Journal of the American Chemical Society*, 60:433–435, 1938.

[105] Richard Zsigmondy. Über die struktur des gels der kieselsäure. theorie der entwässerung. *Zeitschrift für anorganische Chemie*, 71:356–377, 1911.

[106] A. Graham Foster. The sorption of condensible vapours by porous solids. part i. the applicability of the capillary theory. *Trans. Faraday Soc.*, 28:645–657, 1932.

[107] Milton W. Cole and W. F. Saam. Excitation spectrum and thermodynamic properties of liquid films in cylindrical pores. *Phys. Rev. Lett.*, 32:985–988, 1974.

[108] W. F. Saam and Milton W. Cole. Excitations and thermodynamics for liquid-helium films. *Phys. Rev. B*, 11:1086–1105, 1975.

[109] B.V Derjaguin and N.V Churaev. Polymolecular adsorption and capillary condensation in narrow slit pores. *Journal of Colloid and Interface Science*, 54:157–175, 1976.

[110] J. Crassous, E. Charlaix, and J.-L. Loubet. Capillary condensation between high-energy surfaces. an experimental study with a surface force apparatus. *EPL (Europhysics Letters)*, 28:37, 1994.

[111] Konstantin G. Kornev, Inna K. Shingareva, and Alexander V. Neimark. Capillary condensation as a morphological transition. *Advances in Colloid and Interface Science*, 96(1–3):143–167, 2002.

[112] R. Evans, U. Marini Bettolo Marconi, and P. Tarazona. Fluids in narrow pores: Adsorption, capillary condensation, and critical points. *The Journal of Chemical Physics*, 84:2376–2399, 1986.

[113] Peter I. Ravikovitch, Gary L. Haller, and Alexander V. Neimark. Density functional theory model for calculating pore size distributions: pore structure of nanoporous catalysts. *Advances in Colloid and Interface Science*, 76:203–226, 1998.

[114] P. I. Ravikovitch, D. Wei, W. T. Chueh, G. L. Haller, and A. V. Neimark. Evaluation of pore structure parameters of mcm-41 catalyst supports and catalysts by means of nitrogen and argon adsorption. *The Journal of Physical Chemistry B*, 101:3671–3679, 1997.

[115] Alexander V Neimark, Peter I Ravikovitch, and Aleksey Vishnyakov. Bridging scales from molecular simulations to classical thermodynamics: density functional theory of capillary condensation in nanopores. *Journal of Physics: Condensed Matter*, 15:347, 2003.

[116] L. Sarkisov and P. A. Monson. Modeling of adsorption and desorption in pores of simple geometry using molecular dynamics. *Langmuir*, 17(24):7600–7604, 2001.

[117] Peter I. Ravikovitch and Alexander V. Neimark. Density functional theory of adsorption in spherical cavities and pore size characterization of templated nanoporous silicas with cubic and three-dimensional hexagonal structures. *Langmuir*, 18:1550–1560, 2002.

[118] Aleksey Vishnyakov and Alexander V. Neimark. Monte carlo simulation test of pore blocking effects. *Langmuir*, 19:3240–3247, 2003.

[119] Geoffrey Mason. The effect of pore space connectivity on the hysteresis of capillary condensation in adsorption—desorption isotherms. *Journal of Colloid and Interface Science*, 88:36–46, 1982.

[120] G. Mason. Determination of the pore-size distributions and pore-space interconnectivity of vycor porous glass from adsorption-desorption hysteresis capillary condensation isotherms. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 415:453–486, 1988.

[121] Matthias Thommes, Bernd Smarsly, Matthijs Groenewolt, Peter I. Ravikovitch, and Alexander V. Neimark. Adsorption hysteresis of nitrogen and argon in pore networks and characterization of novel micro- and mesoporous silicas. *Langmuir*, 22:756–764, 2006.

[122] Chunyan Fan, D. D. Do, and D. Nicholson. On the cavitation and pore blocking in slit-shaped ink-bottle pores. *Langmuir*, 27:3511–3526, 2011.

[123] Alexander V. Neimark and Aleksey Vishnyakov. The birth of a bubble: A molecular simulation study. *The Journal of Chemical Physics*, 122(5):054707, 2005.

[124] P. C. Ball and R. Evans. Temperature dependence of gas adsorption on a mesoporous solid: capillary criticality and hysteresis. *Langmuir*, 5:714–723, 1989.

[125] R.J.M. Pellenq and P. E. Levitz. Capillary condensation in a disordered mesoporous medium: a grand canonical monte carlo study. *Molecular Physics*, 100:2059–2077, 2002.

[126] Roland J.-M. Pellenq, Bernard Rousseau, and Pierre E. Levitz. A grand canonical monte carlo study of argon adsorption/condensation in mesoporous silica glasses. *Phys. Chem. Chem. Phys.*, 3:1207–1212, 2001.

[127] Lev D. Gelb and K. E. Gubbins. Characterization of porous glasses: simulation models, adsorption isotherms, and the brunauerемmettteller analysis method. *Langmuir*, 14:2097–2111, 1998.

[128] F. Detcheverry, E. Kierlik, M. L. Rosinberg, and G. Tarjus. Mechanisms for gas adsorption and desorption in silica aerogels: the effect of temperature. *Langmuir*, 20:8006–8014, 2004.

[129] E. Kierlik, M. L. Rosinberg, G. Tarjus, and E. Pitard. Mean-spherical approximation for a lattice model of a fluid in a disordered matrix. *Mol. Phys.*, 95:341–351, 1998.

[130] J. P. Sethna, K. A. Dahmen, and O. Perković. Random-Field Ising Models of Hysteresis. In G. Bertotti and I. D. Mayergoyz, editors, *The science of hysteresis II*. Academic Press, New York, 2006.

[131] Kurt Binder and Walter Kob. *Glassy Materials and Disordered Solids: An Introduction to Their Statistical Mechanics*. World Scientific, 2005.

[132] M. E. J. Newman and G. T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 2006.

[133] James P. Gleeson. Binary-state dynamics on complex networks: Pair approximation and beyond. *Phys. Rev. X*, 3:021004, 2013.

[134] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.

[135] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz. A new algorithm for monte carlo simulation of ising spin systems. *J. Comp. Phys.*, 17:10–18, 1975.

[136] Santiago A. Serebrinsky. Physical time scale in kinetic monte carlo simulations of continuous-time markov chains. *Phys. Rev. E*, 83:037701, 2011.

[137] R.G. Palmer. Broken ergodicity. *Adv. Phys.*, 31:669–735, 1982.

[138] Vladimir A. Sharapov, Dario Meluzzi, and Vladimir A. Mandelshtam. Low-temperature structural transitions: Circumventing the broken-ergodicity problem. *Phys. Rev. Lett.*, 98:105701, 2007.

[139] A.E. Mark C. D. Christ and W.F. van Gunsteren. Basic ingredients of free energy calculations: A review. *J. Comput. Chem.*, 31:1569–1582, 2010.

[140] Ralf Metzler, Jae-Hyung Jeon, Andrey G. Cherstvy, and Eli Barkai. Anomalous diffusion models and their properties: non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking. *Phys. Chem. Chem. Phys.*, 16:24128, 2014.

[141] Yukito Iba. Extended ensemble monte carlo. *Int. J. Mod. Phys. C*, 12:623, 2001.

[142] Alastair D. Bruce and Nigel B. Wilding. *Computational Strategies for Mapping Equilibrium Phase Diagrams*, volume 127, pages 1–64. John Wiley & Sons, Inc., 2004.

[143] Yukito Iba, Nen Saito, and Akimasa Kitajima. Multicanonical mcmc for sampling rare events: an illustrative review. *Ann. Inst. Stat. Math.*, 66:611–645, 2014.

[144] Fugao Wang and D. P. Landau. Efficient, multiple-range random walk algorithm to calculate the density of states. 86:2050–2053, 2001.

[145] Chenggang Zhou and R. N. Bhatt. Understanding and improving the wang-landau algorithm. 72:025701, 2005.

[146] Chenggang Zhou, T. C. Schulthess, Stefan Torbru, and D. P. Landau. Wang-landau algorithm for continuous models and joint density of states. 96:120201, 2006.

[147] N. Plattner, J. D. Doll, P. Dupuis, H. Wang, Y. Liu, and J. E. Gubernatis. An infinite swapping approach to the rare-event sampling problem. *J Chem Phys*, 135:134111, 2011.

[148] A. I. Zeifman. Quasi-ergodicity for non-homogeneous continuous-time markov chains. *J. Appl. Prob.*, 26:643–648, 1989.

[149] Yuko Okamoto. Generalized-ensemble algorithms: enhanced sampling techniques for monte carlo and molecular dynamics simulations. *J. Mol. Graph. Model.*, 22:425–439, 2004.

[150] S. Reynal and H. T. Diep. Fast flat-histogram method for generalized spin models. *Phys. Rev. E*, 72:056710, 2005.

[151] Alexandra Valentim, Julio C S Rocha, Shan-Ho Tsai, Ying Wai Li, Markus Eisenbach, Carlos E Fiore, and David P Landau. Exploring replica-exchange wang-landau sampling in higher-dimensional p arameter space. *J. Phys.: Conference Series*, 640(1):012006, 2015.

[152] Yuji Sugita, Akio Kitao, and Yuko Okamoto. Multidimensional replica-exchange method for free-energy calculations. *J. Chem. Phys.*, 113:6042–6051, 2000.

[153] Helmut G Katzgraber, Simon Trebst, David A Huse, and Matthias Troyer. Feedback-optimized parallel tempering monte carlo. *J. Stat. Mech.*, page P03018, 2006.

[154] Elmar Bittner, Andreas Nußbaumer, and Wolfhard Janke. Make life simple: Unleash the full power of the parallel tempering algorithm. *Phys. Rev. Lett.*, 101:130603, 2008.

[155] Yuji Sugita and Yuko Okamoto. Replica-exchange multicanonical algorithm and multicanonical replica-exchange method for simulating systems with rough energy landscape. *Chem. Phys. Lett.*, 329:261–270, 2000.

[156] Ayori Mitsutake and Yuko Okamoto. Replica-exchange simulated tempering method for simulations of frustrated systems. *Chem. Phys. Lett.*, 332:131–138, 2000.

[157] F. Gray. Pulse code communication, u.s. patent no. 2,632,058, 1953. U.S. patent no. 2,632,058.

[158] L T Wille and J Vennik. Computational complexity of the ground-state determination of atomic clusters. *Journal of Physics A: Mathematical and General*, 18:L419, 1985.

[159] F Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15:3241, 1982.

[160] Alexander H. Hartmann and Heiko Rieger. *Maximum-Flow Methods*. Wiley-VCH Verlag GmbH Co. KGaA, 2003.

[161] Andrew T. Ogielski. Integer optimization and zero-temperature fixed point in ising random-field systems. *Phys. Rev. Lett.*, 57:1251–1254, 1986.

[162] Carlos Frontera and Eduard Vives. Studying avalanches in the ground state of the two-dimensional random-field ising model driven by an external field. *Phys. Rev. E*, 62:7470–7473, 2000.

[163] D.D. Frantz, D.L. Freeman, and J.D. Doll. Reducing quasi-ergodic behavior in monte carlo simulations by j-walking: Applications to atomic clusters. *J. Chem. Phys.*, 93:2769, 1990.

[164] D.D. Frantz, D.L. Freeman, and J.D. Doll. Extending j walking to quantum systems: Applications to atomic clusters. *J. Chem. Phys.*, 97:5713–5731, 1992.

[165] C. J. Tsai and K. D. Jordan. Monte carlo simulation of (h2o)8: Evidence for a low energy s4 structure and characterization of the solid-liquid transition. *J. Chem. Phys.*, 95:3850–3853, 1991.

[166] Gustavo E. López and David L. Freeman. A study of low temperature heat capacity anomalies in bimetallic alloy clusters using j-walking monte carlo methods. *J. Chem. Phys.*, 98:1428–1435, 1993.

[167] Gustavo E. López. Study of the solid–liquid transition for ar55 using the j-walking monte carlo method. *J. Chem. Phys.*, 104:6650–6653, 1996.

[168] C. J. Geyer and E. A. Thompson. Annealing markov chain monte carlo with applications to ancestral inference. *J. Am. Stat. Assoc.*, 90:909, 1995.

[169] Koji Hukushima and Koji Nemoto. Exchange monte carlo method and application to spin glass simulations. *J. Phys. Soc. Jpn.*, 65:1604–1608, 1996.

[170] Sheldon B. Opps and Jeremy Schofield. Extended state-space monte carlo methods. *Phys. Rev. E*, 63:056701, 2001.

[171] Scott Brown and Teresa Head-Gordon. Cool walking: A new markov chain monte carlo sampling method. *J. Comp. Chem.*, 24:68–76, 2003.

[172] Ruhong Zhou and B. J. Berne. Smart walking: A new method for boltzmann sampling of protein conformations. *J. Chem. Phys.*, 107:9185–9196, 1997.

[173] Ioan Andricioaei, John E. Straub, and Arthur F. Voter. Smart darting monte carlo. *J. Chem. Phys.*, 114:6994–7000, 2001.

[174] Thomas Vogel, Ying Wai Li, Thomas Wüst, and David P. Landau. Generic, hierarchical framework for massively parallel wang-landau sampling. *Phys. Rev. Lett.*, 110:210603, 2013.

[175] Thomas Vogel, Ying Wai Li, Thomas Wüst, and David P. Landau. Scalable replica-exchange framework for wang-landau sampling. *Phys. Rev. E*, 90:023302, 2014.

[176] A. Matro, D. L. Freeman, and R. Q. Topper. Computational study of the structures and thermodynamic properties of ammonium chloride clusters using a parallel jump-walking approach. *J. Chem. Phys.*, 104:8690, 1996.

[177] D. D. Frantz. Magic numbers for classical lennard-jones cluster heat capacities. *The Journal of Chemical Physics*, 102:3747–3768, 1995.

[178] A. Valentim, M.G.E. da Luz, and Carlos E. Fiore. Determining efficient temperature sets for the simulated tempering meth od. *Computer Physics Communications*, 185:2046–2055, 2014.

[179] David A. Kofke. On the acceptance probability of replica-exchange monte carlo trials. *J. Chem. Phys.*, 117:6911–6914, 2002.

[180] Thomas Wüst and David P. Landau. Optimized wang-landau sampling of lattice polymers: Ground state search and folding thermodynamics of hp model proteins. *J. Chem. Phys.*, 137:064903, 2012.

[181] A. Cavagna. Supercooled liquids for pedestrians. *Phys. Rep.*, 476:51, 2009.

[182] Daniel Sindhikara, Yilin Meng, and Adrian E. Roitberg. Exchange frequency in replica exchange molecular dynamics. *J. Chem. Phys.*, 128:024103, 2008.

[183] Daniel Sindhikara, Daniel J. Emerson, and Adrian E. Roitberg. Exchange often and properly in replica exchange molecular dynamics. *J. Chem. Theor. Comp.*, 6:2804–2808, 2010.

[184] J. P. Sethna, K. A. Dahmen, and O. Perković. Random-Field Ising Models of Hysteresis. In G. Bertotti and I. D. Mayergoyz, editors, *The science of hysteresis II*. Academic Press, New York, 2006.

[185] P.H. Lundow, K. Markström, and A. Rosengren. The ising model for the bcc, fcc and diamond lattices: A comparison. *Phil. Mag.*, 89:2009–2042, 2009.

[186] Boris M Smirnov. Aerogels. *Sov. Phys. Uspekhi*, 30:420–432, 1987.

[187] Abdelali Rahmani, Claude Benoit, Rémi Jullien, Gérard Poussigue, and Abdellah Sakout. Dynamical properties of a diffusion-limited cluster - cluster aggregation model. *J. Phys.: Cond. Matt.*, 8:5555–5567, 1996.

[188] Anwar Hasmy, Eric Anglaret, Marie Foret, Jacques Pelous, and Rémi Jullien. Small-angle neutron-scattering investigation of long-range correlations in silica aerogels: Simulations and experiments. *Phys. Rev. B*, 50:6006–6016, 1994.

[189] M. L. Rosinberg, E. Kierlik, and G. Tarjus. Percolation, depinning, and avalanches in capillary condensation of gases in disordered porous solids. *Europhys. Lett.*, 62:377–383, 2003.

[190] T. D. Lee and C. N. Yang. Statistical theory of equations of state and phase transitions. ii. lattice gas and ising model. *Phys. Rev. Lett.*, 87:410–419, 1952.

[191] Daniel H. Rothman and Stéphane Zaleski. Lattice-gas models of phase separation: interfaces, phase transitions, and multiphase flow. *Rev. Mod. Phys.*, 66:1417–1479, 1994.

[192] G. J. Aubry, F. Bonnet, M. Melich, L. Guyon, F. Despetis, and P. E. Wolf. Critical behavior of the liquid gas transition of ^{4}He confined in a silica aerogel. *J. Low Temp. Phys.*, 171:670, 2013.

[193] L. Sarkisov and P.A. Monson. Hysteresis in monte carlo and molecular dynamics simulations of adsorption in porous materials. *Langmuir*, 16:9857–9860, 2000.

[194] L. Sarkisov and P. A. Monson. Lattice model of adsorption in disordered porous materials: Mean-field density functional theory and monte carlo simulations. *Phys. Rev. E*, 65:011202, 2001.

[195] Paul D. Beale. Exact distribution of energies in the two-dimensional ising model. *Phys. Rev. Lett.*, 76:78–81, 1996.

[196] T. Papakonstantinou and A. Malakis. Parallel tempering and 3d spin glass models. *J. Phys.: Conference Series*, 487:012010, 2014.

[197] Yihui Ren, Stephen Eubank, and Madhurima Nath. From network reliability to the ising model: A parallel scheme for estimating the joint density of states. *Phys. Rev. E*, 94:042125, 2016.

[198] G. Bel and E. Barkai. Weak ergodicity breaking in the continuous-time random walk. *Phys. Rev. Lett.*, 94:240602, 2005.

[199] J. H. Page, J. Liu, B. Abeles, E. Herbolzheimer, H. W. Deckman, and D. A. Weitz. Adsorption and desorption of a wetting fluid in vycor studied by acoustic and optical techniques. *Phys. Rev. E*, 52:2763–2777, 1995.

[200] K. Morishige. Hysteresis critical point of nitrogen in porous glass: Occurrence of sample spanning transition in capillary condensation. *Langmuir*, 25:6221–6226, 2009.

[201] Richard Cimino, Katie A. Cychosz, Matthias Thommes, and Alexander V. Neimark. Experimental and theoretical studies of scanning adsorption–desorption isotherms. *Colloids and Surfaces A: Physicochem. Eng. Aspects*, 437:76–89, 2013.

[202] F. Bonnet, M. Melich, L. Puech, and P. E. Wolf. Light scattering study of collective effects during evaporation and condensation in a disordered porous material. *EPL*, 101:16010, 2013.

[203] D. H Everett. *The Solid-Gas Interfaces*. Marcel Dekker, New York, 1967.

[204] Hyung-June Woo, F. Porcheron, and P. A. Monson. Modeling desorption of fluids from disordered mesoporous materials. *Langmuir*, 20:4743–4747, 2004.

[205] A. C. Mitropoulos, K. L. Stefanopoulos, E. P. Favvas, E. Vansant, and N. P. Hankin. On the formation of nanobubbles in vycor porous glass during the desorption of halogenated hydrocarbons. *Sci. Rep.*, 5:10943, 2015.

[206] E. S. Kikkinides, M. E. Kainourgiakis, and A. K. Stubos. Origin of hysteresis of gas adsorption in disordered porous media: lattice gas model versus percolation theory. *Langmuir*, 19:3338–3344, 2003.

[207] A. A. Balakin. Mathgl, url: http://mathgl.sourceforge.net.

[208] T. P. Handford, F. J. Perez-Reche, and S. N. Taraskin. Zero-temperature random-field ising model on a bilayered bethe lattice. *Phys. Rev. E*, 88:022117, 2013.

[209] F.J. Perez-Reche, S.N. Taraskin, W. Otten, M.P. Viana, L. da F. Costa, and C.A. Gilligan. Prominent effect of soil network heterogeneity on microbial invasion. *Phys. Rev. Lett.*, 109:098102, 2012.

[210] Zilvinas Rimas. Condensation of liquid in porous media. PartIII project, Cambridge.

[211] E. A. Bender and E. R. Canfield. The asymptotic number of labelled graphs with given degree sequences. *Journal of Combinatorial Theory, Series A*, 24:296–307, 1978.

[212] M. Boguñá, R. Pastor-Satorras, and A. Vespignani. Cut-offs and finite size effects in scale-free networks. *The European Physical Journal B*, 38:205–209, 2004.

[213] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.