Susanna Röblitz

Statistical Error Estimation and Grid-free Hierarchical Refinement in Conformation Dynamics

Dissertation
am Fachbereich Mathematik und Informatik
der Freien Universität Berlin
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften

vorgelegt von

Diplom-Mathematikerin Susanna Röblitz

Berlin, Juli 2008

Betreuer/Gutachter:	Prof. Dr. Dr. h.c. Peter Deuflhard Fachbereich für Mathematik und Informatik, FU Berlin, Arnimallee 2-6, 14195 Berlin
Zweitgutachter:	Dr. Wilhelm Huisinga Senior Lecturer, Hamilton Institute, National University of Ireland, Maynooth, Co. Kildare, Ireland
Tag der Disputation:	17. Dezember 2008

In memory of my father Detlef (1954-2007)

Contents

Introduction

Among dynamical systems that possess a unique stationary distribution, there exist *metastable* dynamical systems characterized by the fact that their invariant set can be decomposed into *almost invariant* subsets. Once such a system is within one of these subsets, it stays there for "a long time" before it rapidly switches to another almost invariant subset.

In *molecular* dynamical systems, which are the focus of our applications here, the almost invariant subsets are named *metastable conformations*, and the stationary density is the Boltzmann density. Since such dynamical systems spend most of the time in one of the conformations, their geometric structures and dynamical properties are essential for the molecule's long-term behavior. This is of interest in computational drug design, where, for example, the binding capacities of different ligands to certain target molecules are examined in order to decide which new ligands are worth to be tested in laboratory experiments. To make it computationally feasible, a reduced description of the molecular system in terms of conformations is needed.

For this purpose, a multi-scale method called *conformation dynamics* has been developed. Its main objective is the identification of metastable conformations together with their life times and transition patterns. In this mixed deterministic/stochastic approach, the dynamics is modeled as a Markov process in a discretized finite state space, which results in a nearly completely decomposable transition probability matrix. The application of a cluster algorithm reveals the metastable conformations [21, 23].

The idea of conformation dynamics goes back to the work of Dellnitz and coworkers on almost invariant sets of dynamical systems [14, 15]. The method has first been realized for Hamiltonian systems as they appear in molecular dynamics by Deuflhard et al. [20], but the approach suffered some conceptual problems. A reformulation in terms of statistical mechanics and the definition of a spatial transition operator by Schütte et al. [89, 88, 90] were the key steps until the algorithmic concept became applicable to the identification of biomolecular conformations [18, 22]. The transfer operator approach led to the insight that metastable conformations can be identified with eigenfunctions of the transfer operator corresponding to eigenvalues near the *Perron root* $\lambda = 1$. These eigenvalues are denoted as Perron cluster eigenvalues and are usually well-separated from the rest of the spectrum. A discretization of the configuration space yields a discretized transfer

operator, the transition probability matrix P, whose entries must be approximated by Monte-Carlo quadrature. Since dynamical systems arising in molecular simulations are mostly high-dimensional, the number of discretized states tends to increase exponentially with the dimension. To avoid this *curse of dimensionality*, Weber introduced *mesh-free* methods [109]. The eigenvectors corresponding to the Perron cluster eigenvalues form an invariant subspace, the *Perron subspace*. The Robust Perron Cluster Analysis (PCCA+) [23] constructs a linear transformation between a basis of the Perron subspace and the discretized *almost* characteristic functions of metastable conformations. Once these conformations have been identified, transition probabilities between them can be computed, which requires the knowledge of the discrete stationary density. However, whenever a dynamical system exhibits metastabilities, the computation of its discrete stationary density as eigenvector of the transition probability matrix is highly ill-conditioned, whereas the computation of the partial densities restricted to the conformations is well-conditioned. This led to the development of a new aggregation/disaggregation method, where the aggregation step is not based on the matrix but on an enhanced sampling algorithm that covers only small patches of the sampling space [112].

Upon application of conformation dynamics to larger molecules, new difficulties came into sight. They are the focus of the present thesis.

(i) The main deficiency in former analyses of conformation dynamics was the lack of *error estimators* for the variables of interest, for example the Perron subspace. Weber's *normE*-criterion [109] gives only error bounds for the stationary distribution, which, due to ill-conditioning, yields impractically large bounds. However, since the entries of the transition probability matrix are determined by Monte-Carlo quadrature, it seems appropriate to assume random perturbations. As will be shown in this thesis, the errors in the matrix elements can efficiently be described by a few parameters of an appropriate probability distribution. By the use of matrix perturbation theory [17, 105], Markov chain theory [70, 9, 53], and statistics [77, 78, 91], these errors can be related to the errors in the Perron subspace and other variables of interest.

(ii) The approximation of matrix elements by Monte-Carlo quadrature is only one possible source of errors. The second main source is the quality of the *discretization*. These two aspects are closely related. Weber already pointed out that the sampling error is large, whenever there exist energy barriers within the sampling domain. For this reason, he suggested the idea of a *hierarchical refinement* of the current decomposition. Its practical realization, however, had been left as an open problem. Questions that will be answered here include error estimators, refinement criteria, the generation of new basis functions, and possible stopping criteria. They build the basis for a new sampling scheme that involves adaptive sampling and hierarchical refinement of basis functions.

(iii) The introduction of domain decomposition concepts and approximation errors

caused by truncated samplings led to a loss of reversibility in the resulting transition matrix. Thus, the perturbation analysis, originally developed for reversible metastable Markov chains, no longer applies. Surprisingly, PCCA+ turned out to be robust, a fact that will be proved here by a modified perturbation analysis based on the sensitivity of invariant subspaces.

(iv) Weber realized mesh-free discretizations in terms of radial basis functions [109]. Such basis functions increase the acceptance ratio in Monte-Carlo sampling methods. Furthermore, they additionally introduce a mass matrix S in the resulting eigenvalue problem. In other words, the standard eigenvalue problem $PX = X\Lambda$ turns into a generalized eigenvalue problem $PX = SX\Lambda$. While the matrix P can be interpreted as transition probability matrix of a "standard" Markov chain, the matrix pair (P, S) does not allow for such an interpretation. Nevertheless, we will speak of "generalized" Markov chains. The appearance of a mass matrix and the characterization of metastabilities in terms of almost characteristic functions require a modification in the derivation of a coarse-grained transition probability matrix that describes transition probabilities between metastable conformations. This approach will lead to a *new objective function* for PCCA+.

(v) Moreover, the identification of metastable conformations from matrix pairs often fails. The reason for this fact lies in the dependence of the condition number of the Perron subspace on the shape parameter of the radial basis functions. This relation will be examined in the thesis.

The solution of the above listed problems is crucial for the applicability of conformation dynamics to larger molecular systems. Our methods include molecular dynamics simulations, which are the most time-consuming part of the algorithms. Alternatively, one could restrict the analysis to given data, for example to long term trajectories, but this method prevents the possibility of adding missing information by extended sampling. If no access to supercomputers such as Blue Gene [33] or distributed computing platforms such as Folding@Home [84] is available, efficient sampling methods are especially important. In other words, we have to generate as few trajectories as possible, but as many trajectories as necessary. This is the key to running conformation dynamics algorithms for larger molecules on single workstations or computer clusters.

Outline The thesis is organized as follows. **Chapter 1** contains the mathematical background of conformation dynamics. It explains the steps that lead from a continuous dynamical system to a discrete transition probability matrix, from which conformations are recovered by PCCA+. **Chapter 2** summarizes perturbation bounds on relevant matrix functions. **Chapter 3** shows that transition matrices stemming from molecular simulations can be considered as perturbations of a certain class of matrices with special properties. This explains the stability of PCCA+ even for non-reversible Markov chains. Moreover, **Chapter 3** refers to

the limitations of the perturbation theory in case of generalization to matrix pairs as they arise from the discretization by radial basis functions. **Chapter 4** illustrates how the matrix errors can be quantified from simulation data such that the perturbation theory becomes of practical use. Since the error estimation is based on statistics, not only error bounds but also *error distributions* can be computed. As it turns out, the error bounds allow conclusions about the error contributions from different basis functions. This is the starting point for **Chapter 5**, where it is demonstrated how adaptive sampling and hierarchical refinement systematically reduce the error in the output. Finally, in **Chapter 6**, we apply the methods to the analysis of a small biomolecule, thus demonstrating the advantages of the newly developed techniques over previous ones.

Chapter 1

Conformation Dynamics

The main challenge in molecular modeling is to bridge the different scales in space and time between computational methods and real life phenomena due to limited computer power. Computational methods can cover the dynamics of all atoms in detail but they are restricted to short simulation times in the range of femtoseconds (fs) and to medium-sized molecules. A possible way to circumvent these problems is to perform a model reduction like conformation dynamics. This method is a mixed deterministic/stochastic approach that combines data-based coarse graining with efficient sampling techniques. In the presence of metastabilities, it is virtually impossible to explore all physically relevant parts of the molecular state space with one single trajectory because it would get trapped in basins of attraction of the potential energy surface. Therefore, the state space is decomposed such that different regions can be sampled independently. The dynamical process is then modeled as a Markov chain in this finite state space. Since the states are located in different regions of the sampling space, including transition states, the trapping problem does not occur. Hence the resulting Markov chain does not suffer from the problem of broken ergodicity. The Markov chain is defined by a transition probability matrix, which is the basis for the identification of conformations. The first section of this chapter explains the mathematical background of the steps that lead from the continuous dynamical system to the transition probability matrix. The recovering of metastable conformations from this matrix via the Robust Perron Cluster Analysis will be the topic of the second section. While the first section mainly contains well-known facts about conformation dynamics, which is necessary for the understanding of subsequent chapters, the second section includes a new contribution to the coarse graining of discrete transition operators in the case of radial basis functions.

1.1 From continuous to discrete systems

The starting point for our mathematical algorithms is the constitutional formula of a certain molecule. The parameterization by a force field like Merck Molecular Force Field (MMFF) [39, 40] makes it feasible for potential energy calculations and

thus for molecular dynamics simulations. This is the basis for the algorithm which leads to a reduced description of the dynamics in terms of a finite dimensional state space.

1.1.1 Molecular dynamics

Consider the canonical ensemble (constant number of particles, constant volume, constant temperature) in equilibrium, i. e., the positions $q \in \Omega$ and momenta $p \in \mathbb{R}^d$ of all d atoms are distributed according to the Boltzmann (or canonical) density,

$$\mu(q,p) = \frac{1}{Z} \exp(-\beta H(q,p)).$$

H denotes the Hamiltonian function, which is the sum of potential energy $V(q)$ and kinetic energy $K(p)$, $\beta = 1/(k_B T)$ denotes the inverse temperature with Boltzmann constant k_B, and Z is the corresponding partition function. The canonical density can be split into a distribution of momenta and positions, $\mu(q,p) = \eta(p)\pi(q)$, with

$$\pi(q) = \frac{1}{Z_q} \exp(-\beta V(q)) \quad \text{and} \quad \eta(p) = \frac{1}{Z_p} \exp(-\beta K(p)).$$

The positional density $\pi(q) : \Omega \to [0,1]$ induces an inner product in $L^2(\pi)$,

$$\langle f, g \rangle_\pi \equiv \int_\Omega f(q) g(q) \pi(q)\, dq.$$

Equivalently, we write $\langle f \rangle_\pi$ if we refer to $\int_\Omega f(q)\pi(q)\, dq$.

Molecular properties are determined by ensemble averages. However, in order to derive these averages, one must examine single molecules. For this purpose, one has to specify some microscopic dynamics. The corresponding equations of motion depend on the model assumption and requirements. The dynamics is supposed to model internal fluctuations within the unique stationary distribution.

The dynamics of choice is the Hamiltonian dynamics with randomized momenta. It simulates the relaxation of the molecule after an initial interaction with the heat bath (surrounding).

Hamiltonian dynamics with randomized momenta Consider the deterministic Hamiltonian system

$$\dot{q} = p, \qquad \dot{p} = -\nabla V(q)$$

with corresponding flow Φ^τ, i.e. $(q(t), p(t)) = \Phi^\tau(q(0), p(0))$. The process

$$q_{n+1} = \Pi_q \Phi^\tau(q_n, p_n), \qquad n \in \mathbb{N},$$

with p_n chosen randomly from the canonical distribution of momenta, represents a discrete time Markov process on Ω. Here, Π_q is the projection of a state (q,p) onto its position coordinate q. It has been shown that $\pi(q)$ is invariant w.r.t. that process [88].

The corresponding one-step transition function is given by

$$p^1(\tau, f, g) = \int_\Omega T^\tau f(q) g(q) \pi(q)\, dq,$$

where T^τ is the transfer operator introduced by Schütte [88],

$$T^\tau f(q) = \int_{\mathbb{R}^d} f(\Pi_q \Phi^{-\tau}(q,p)) \eta(p) dp.$$

It gives rise to a semigroup of propagators $T^{n\tau} : L^1(\pi) \to L^1(\pi)$ by $T^{n\tau} = (T^\tau)^n$. Thus the knowledge of T^τ allows the prediction of long term dynamic behavior. Moreover, T^τ is self-adjoint in $L^2(\pi)$.

Remark 1.1.1. In classical molecular dynamics (MD), the discrete time process is given by $(q_{n+1}, p_{n+1}) = \Phi^\tau(q_n, p_n)$. Since the energy $H(q,p)$ is conserved over time, the process generates a *micro-canonical* molecular dynamics trajectory (constant energy instead of constant temperature).

1.1.2 Metastable conformations

A metastable conformation is characterized by an average geometric structure that is preserved over a long period of time before the molecule switches to another metastable conformation. Thus, metastable conformations are defined from a dynamical point of view. They are not dynamically stable, but still *metastable*. Although we will mostly speak of conformations and skip the term "metastable", the concept should not be confused with geometric conformations, a term that is often used in the chemical literature to denote single molecular configurations.

From a mathematical point of view, conformations are *almost characteristic* functions $\chi : \Omega \to [0,1]$ defined in position space which are *almost invariant* under the transfer operator T^τ [88, 109]. Assume there are n_C such conformations. Then

$$T^\tau \chi_k(q) \approx \chi_k(q), \qquad k = 1, \ldots, n_C. \tag{1.1}$$

The use of $\approx$ instead of equality should emphasize the attribute *almost* in connection with invariance. Metastability analysis aims at a decomposition of the position space Ω into n_C metastable conformations $\chi_1, \ldots, \chi_{n_C}$ such that these functions are non-negative ($\chi_k(q) \geq 0$) and form a *partition of unity*,

$$\sum_{k=1}^{n_C} \chi_k(q) = 1 \quad \forall\, q \in \Omega,\ k = 1, \ldots, n_C.$$

Conformation dynamics comprises the description of the dynamic behavior of molecules in terms of transition probabilities between metastable conformations. In some cases, average holding times and transition rates can be derived as well.

1.1.3 Discretization

The first step in modeling complex dynamical systems is a discretization of Ω in terms of a finite number of states, denoted by indices $i = \{1, \ldots, N\}$. It is a well known fact that the application of mesh based methods to high dimensional approximation problems, as they appear in conformation dynamics, leads to the "curse of dimensionality". For this reason, Weber [109] introduced the concept of mesh-free methods. The notation "mesh-free" refers to a particle based discretization method. Given a number of *nodes* $\{q_i\}_{i=1}^N \in \Omega$, they define a set of functions $\phi_i(q) \equiv \phi_i(q_i, q) : \Omega \to [0, 1]$ which form a partition of unity. These functions can be either characteristic functions or almost characteristic functions. The special kind of discretization has different algorithmic consequences, which will play a role in later chapters. In either case, the conformations are expressed as linear combinations of basis functions,

$$\chi_j(q) = \sum_{i=1}^{N} \chi_{\text{discr}}(i, j)\phi_i(q), \qquad j = 1, \ldots, n_C. \tag{1.2}$$

Since we will nearly always work with the discrete membership vectors $\chi_{\text{discr}}(:, j)$, we will from now on skip the subscript "discr" and identify $\chi_j \equiv \chi_{\text{discr}}(:, j)$. Whenever a distinction becomes necessary, we will write $\chi_j(q)$ for the continuous function.

Voronoi tessellation A special kind of mesh-less discretization is the decomposition of Ω into N subsets. These subsets are described by characteristic basis functions $\{\phi_i(q)\}_{i=1}^N$ with $\phi_i : \Omega \to \{0, 1\}$. The basis functions are defined as Voronoi tessellation of Ω,

$$\phi_i(q) = \begin{cases} 1, & \text{if } d(q, q_i) = \min_{j=1,\ldots,n_C} d(q, q_j) \\ 0, & \text{else} \end{cases},$$

where $d(x, y)$ denotes some appropriate distance function, for example the Euclidean metric.

A Galerkin discretization of condition (1.1) in terms of (1.2) leads to the standard eigenvalue problem

$$P^\tau \chi_j \approx \chi_j, \qquad j = 1, \ldots, n_C, \tag{1.3}$$

where

$$P^\tau(i, j) = \frac{\langle \phi_i, T^\tau \phi_j \rangle_\pi}{\langle \phi_i \rangle_\pi}, \quad i, j = 1, \ldots, N.$$

The denominator is the statistical weight of basis function $\phi_i(q)$,

$$w_i \equiv \int_\Omega \phi_i(q)\pi(q)\, dq. \tag{1.4}$$

The vector $\mathbf{w} = [w_1, \ldots, w_N]^\top$ is the left eigenvector of P^τ corresponding to the eigenvalue $\lambda = 1$,

$$\mathbf{w}^\top P^\tau = \mathbf{w}^\top.$$

It is denoted as *stationary density* or *stationary distribution* of P^τ.

If the positional density $\pi(q)$ is restricted to a certain basis function and normalized appropriately, we obtain the partial density

$$\pi_i(q) \equiv \frac{\phi_i(q)\pi(q)}{\int_\Omega \phi_i(q)\pi(q)\,dq}. \tag{1.5}$$

The partial densities $\pi_i(q)$ are Boltzmann densities corresponding to modified potentials $V_i(q)$ [109],

$$\pi_i(q) = \frac{\exp(-\beta\, V_i(q))}{\int_\Omega \exp(-\beta\, V_i(q))\,dq}$$

with

$$V_i(q) = \begin{cases} V(q), & \text{if} \quad \phi_i(q) = 1 \\ \infty, & \text{if} \quad \phi_i(q) = 0 \end{cases}. \tag{1.6}$$

Using these notations results in

$$P^\tau(i,j) = \int_\Omega T^\tau \phi_j(q)\pi_i(q)\,dq.$$

The matrix $P^\tau \in \mathbb{R}^{N\times N}$ is a stochastic transition probability matrix that describes the transition behavior between the states. Metastable conformations will be identified from this matrix via the Robust Perron Cluster Analysis, which will be described in Section 1.2.

Radial basis functions Alternatively to characteristic basis functions, Ω can be discretized by *almost characteristic* basis functions, for example radial basis functions

$$\phi_i(q) = \frac{\exp(-\alpha\, d^2(q, q_i))}{\sum_{j=1}^N \exp(-\alpha\, d^2(q, q_j))}. \tag{1.7}$$

The shape parameter α influences the overlap between the basis functions. The larger α, the smaller the overlap. As α goes to infinity, the Voronoi tessellation is recovered.

A Galerkin discretization of condition (1.1) in terms of (1.2) leads to the generalized eigenvalue problem

$$P^\tau \chi_j \approx S\chi_j, \qquad j = 1, \ldots, n_C, \tag{1.8}$$

where

$$P^\tau(i,j) \equiv \frac{\langle \phi_i, T^\tau \phi_j \rangle_\pi}{\langle \phi_i \rangle_\pi} \tag{1.9}$$

and

$$S(i,j) \equiv \frac{\langle \phi_i, \phi_j \rangle_\pi}{\langle \phi_i \rangle_\pi}. \tag{1.10}$$

Note that S is also a row-stochastic matrix with stationary distribution $\mathbf{w}$. In terms of partial densities (1.5), the matrix entries can be rewritten as

$$P^\tau(i,j) = \int_\Omega T^\tau \phi_j(q)\pi_i(q)\, dq, \tag{1.11}$$

$$S(i,j) = \int_\Omega \phi_j(q)\pi_i(q)\, dq. \tag{1.12}$$

The corresponding modified potential reads

$$V_i(q) = V(q) - \frac{1}{\beta}\log(\phi_i(q)). \tag{1.13}$$

Thus, the sampling of the partial density $\pi_i(q)$ can be considered as umbrella sampling [106] with modified potential $V_i(q)$. Note that the expression for P^τ is the same as in the Voronoi case. The matrix pair (P^τ, S) is then examined by the Robust Perron Cluster Analysis to identify metastable conformations. In the following, we will skip the index τ at the matrix P. However, if distinction of matrices according to different simulation times τ becomes necessary, it will be added.

1.1.4 Simulation

The first step of the sampling routine comprises the generation of sampling points $\{q_k^{(i)}\}_{k=1}^{n_i}$ via hybrid Monte-Carlo (HMC) [26, 32] according to the partial density $\pi_i(q)$. In HMC sampling, the proposal step is realized via a short term molecular dynamics simulation. If the cells of the Voronoi tessellation become too small, the MD trajectories might often lead to end points outside the cell, which results in a small acceptance ratio of the sampling routine. To diminish this problem, one could apply umbrella sampling [106] which drives the dynamics towards the center of the cell. Umbrella sampling increases the acceptance ratio but requires the additional storage of weights for all sampling points. This disadvantage can be avoided by directly incorporating the umbrella sampling functions into the discretization, which gave rise to the idea of radial basis functions. High acceptance ratios can be considered as the main advantage of that approach.

For radial basis functions, the horizontal sampling is used for an approximation of (1.12) by Monte-Carlo quadrature,

$$S(i,j) \approx \frac{1}{n_i}\sum_{k=1}^{n_i} \phi_j(q_k^{(i)}). \tag{1.14}$$

Monte-Carlo integration of (1.11) includes the generation of random momenta $\{p_l^{(ik)}\}_{l=1}^{m_{ik}}$ and the application of the flow Φ^τ to the sampling points $(q_k^{(i)}, p_l^{(ik)})$,

$$P(i,j) \approx \frac{1}{n_i}\frac{1}{m_{ik}}\sum_{k=1}^{n_i}\sum_{l=1}^{m_{ik}} \phi_j(\Pi_q \Phi^\tau(q_k^{(i)}, p_l^{(ik)})). \tag{1.15}$$

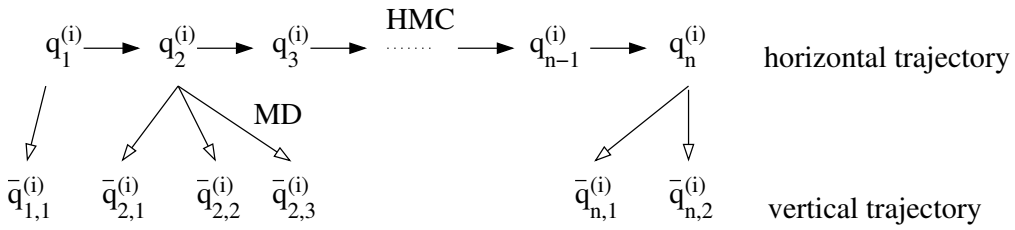

Figure 1.1: For every basis function $\phi_i(q)$, the partial density $\pi_i(q)$ is approximated by a horizontal trajectory which is generated by hybrid Monte-Carlo. The ensemble of points is then propagated by molecular dynamics in order to estimate transition probabilities. Here $\overline{q}_{jk}^{(i)} = \Pi_q \Phi^\tau(q_j^{(i)}, p_k^{(ij)})$.

In case of a Voronoi tessellation, the transition probability $P(i,j)$ from a state i to another state j is simply the portion of points $\{q_k^{(i)}\}_{k=1}^{n_i}$ located in basis function j after propagation. For radial basis functions, one has to compute the values of all basis functions in the end points $\{\Pi_q \Phi^\tau(q_k^{(i)}, p_l^{(ik)})\}$ and to sum up these vectors over all points. Both discretizations require the computation of distances between sampling points and the centers of the basis functions. Moreover, in case of a Voronoi tessellation, one has to compute the argmin-value of these distances, whereas in case of radial basis functions one has to apply the exponential function and to compute the sum of these values, compare (1.7). Thus, radial basis functions require slightly more computation time than Voronoi cells. Furthermore, much less storage space is needed if one wants to store the values of the basis functions for every sampling point in case of Voronoi cells compared to radial basis functions. This can simplify, for example, a-posteriori analysis based on resampling techniques.

The overall sampling scheme is illustrated in Figure 1.1. The points $\{q_k^{(i)}\}_{k=1}^{n_i}$ will be referred to as *horizontal* points or horizontal trajectory, whereas the points $\{\overline{q}_{kl}^{(i)} = \Pi_q \Phi^\tau(q_k^{(i)}, p_l^{(ik)})\}$ will be denoted as *vertical* points or vertical trajectory. Similarly, the process of generating the horizontal trajectory will be denoted as *horizontal sampling*, whereas the process of generating vertical sampling points will be denoted as *vertical sampling*. Depending on the sampling method, some horizontal points could be propagated several times with different initial momenta, whereas other points might not be propagated at all.

In the described sampling scheme, all rows of P and S are computed separately. The discretization in terms of basis functions $\phi_i(q)$ represents a partitioning of the sampling space such that different regions are sampled independently. Hence, we may as well speak of *domain decomposition* in Ω.

1.1.5 Error analysis

Since the elements of the matrices P and S are computed by Monte-Carlo quadrature, they suffer from truncation errors. Instead of computing exact matrices P

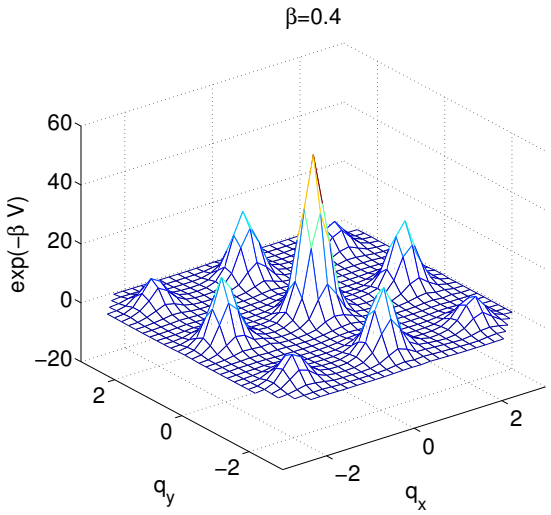

Figure 1.2: Toy example: Stationary density for $\beta = 0.4$.

and S, we end up with perturbed matrices $\widetilde{P} = P + E$ and $\widetilde{S} = S + F$. Weber [109] proposed a method to estimate $\|E\|_\infty$ or $\|F\|_\infty$, respectively, the so-called *normE*-criterion. Since $\|E\|_\infty = \max_{i=1,\ldots,N} \|E(i,:)\|_1$, the norm of the error matrix can be related to the row wise error that results from the sampling error of a single basis function ϕ_i. To quantify this error, one starts a number c of well dispersed Markov chains in basis function ϕ_i. All chains result in c different rows $P_k(i,:)$, $k = 1, \ldots, c$, from which the estimate

$$\|E(i,:)\|_1 \leq \max_{l,k=1,\ldots,c} \|P_k(i,:) - P_l(i,:)\|_1, \quad i = 1, \ldots, N. \tag{1.16}$$

is obtained. In practice, this kind of worst-case error estimator turns out to be too pessimistic because it is sensitive w.r.t. outliers. The sampling effort to make $\|E\|_\infty$ smaller than a certain bound is in most cases extremely large. Moreover, the error does not necessarily decrease with increasing number c of sampling chains. This motivated the construction of more appropriate estimators that will be presented in Chapter 4.

1.1.6 A toy example

Throughout the chapters, we will often return to the following "toy system". We consider a two dimensional periodic domain $\Omega = [-\pi, \pi]^2$ and an artificial potential energy

$$V(q) = -5\cos(3q_x) + 0.5q_x^2 - 5\cos(3q_y) + 0.5q_y^2, \tag{1.17}$$

which defines a Hamiltonian dynamical system w.r.t. positions q and momenta p,

$$\dot{q} = p, \qquad \dot{p} = -\nabla V(q).$$

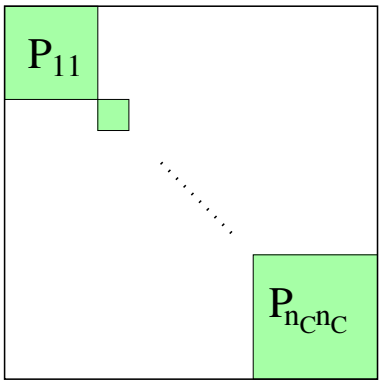

Figure 1.3: Structure of an unperturbed transition matrix with n_c stable clusters.

The potential V is very similar to the free energy landscape of a pentane molecule restricted to the space spanned by two dihedral angles. This low dimensional example allows fast calculations and visualization. Unless stated differently, we will always consider the system at inverse temperature $\beta = 0.4$. The change of the dynamical behavior w.r.t. β has been considered by several authors [45, 30, 29] but will not be in the scope of the present thesis. For $\beta = 0.4$, the corresponding stationary density $\pi(q) \propto \exp(-\beta V(q))$ decomposes into nine well separated peaks that represent the metastable conformations, see Figure 1.2. To apply molecular dynamics for this system means that we use the Verlet algorithm [32] with time step $\tau = 0.01$. Other quantities will be given explicitly in the concrete examples.

1.2 Robust Perron Cluster Analysis

The transition probability matrix P represents a Markov chain on the state space spanned by the basis functions. If discretized appropriately, the metastabilities inherent in the continuous system can be recovered from the structure of P. If two basis functions belong to the same metastable region, the transition probability between them will be large, whereas the transition probability between basis functions from different metastable regions will be small. Thus, after an appropriate permutation of states, the matrix P will become nearly block-diagonal. An assignment of states to the blocks corresponds to an assignment of basis functions to metastable regions. This is the task of the cluster algorithm.

1.2.1 Perturbation theory for reversible Markov chains

To explain the basic idea of the clustering algorithm, assume that Ω can be decomposed into stable parts, i.e. the corresponding Markov chain is decomposable. The transition probability matrix of a decomposable Markov chain can be reordered such that it becomes block diagonal, see Figure 1.3. Such matrices have an n_C-fold eigenvalue $\lambda = 1$, the so-called *Perron root*, where n_C corresponds to the number

of decoupled Markov chains. The corresponding eigenvectors $\{X_j\}_{j=1}^{n_C}$, the *Perron eigenvectors*, are piecewise constant on the blocks and can thus be used to identify the clusters [21]. They can be written as linear combinations of the *characteristic vectors* of the metastable conformations,

$$X_j = \sum_{k=1}^{n_C} a_{kj}\chi_k,$$

where

$$\chi_k(i) = \begin{cases} 1, & \text{if state } i \text{ belongs to cluster } k \\ 0, & \text{else} \end{cases}.$$

The perturbation theory developed in [21] shows that if, under some additional conditions, the matrix $\widetilde{P}$ is an $\mathcal{O}(\varepsilon)$-perturbation of a block-diagonal matrix P,

$$\widetilde{P} = P + \epsilon P^{(1)} + O(\epsilon^2),$$

then the transition matrix $\widetilde{P}$ will have a Perron cluster of eigenvalues

$$\widetilde{\lambda}_1 = 1, \quad \widetilde{\lambda}_2 = 1 - O(\epsilon), \quad \ldots \quad \widetilde{\lambda}_{n_C} = 1 - O(\epsilon)\ ,$$

where $\epsilon > 0$ denotes some perturbation parameter, which is scaled as

$$\epsilon = 1 - \widetilde{\lambda}_2\ .$$

Moreover, there will be a set of corresponding eigenvectors $\widetilde{X} = [\widetilde{X}_1, \ldots, \widetilde{X}_{n_C}] \in \mathbb{R}^{N\times n_C}$ of the form

$$\widetilde{X}_i = X_i + \epsilon X_i^{(1)} + O(\epsilon^2)\ .$$

In [21], the result

$$X_i = \sum_{j=1}^{n_C} \widetilde{\alpha}_{ij}\,\chi_j + \epsilon \underbrace{\sum_{j=k+1}^{N} \frac{1}{1-\lambda_j}\Pi_j P^{(1)} X_i}_{R_i} + \mathcal{O}(\epsilon^2) \tag{1.18}$$

has been obtained, where Π_j denotes the π-orthogonal projection on the eigenspace of $\widetilde{P}$ corresponding to the eigenvalue $\widetilde{\lambda}_i$ [52]. Unfortunately, the statement that R_i equals zero as given in [23] turned out to be wrong (see [55]). Thus, perturbations of order ϵ in the matrix elements cause errors of order ϵ in the characteristic functions of the conformations.

The proof of (1.18) exploits that $\widetilde{P}$ is generalized symmetric,

$$\widetilde{D}^2\widetilde{P} = \widetilde{P}^\top\widetilde{D}^2, \tag{1.19}$$

where $\widetilde{D}^2 = \mathrm{diag}(\widetilde{w}_i)$ and $\widetilde{\mathbf{w}}$ is the unique stationary distribution of $\widetilde{P}$,

$$\widetilde{\mathbf{w}}^\top\widetilde{P} = \widetilde{\mathbf{w}}^\top.$$

Equation (1.19) is just another formulation of the *detailed balance condition*,

$$\widetilde{w}_i \widetilde{P}(i,j) = \widetilde{w}_j \widetilde{P}(j,i), \quad \forall\, i,j = 1,\ldots,N. \tag{1.20}$$

However, this assumption is often violated in practical calculations due to the application of domain decomposition methods and truncated sampling. Therefore, a modified perturbation theory will be presented in Chapters 2 and 3.

Example 1.2.1. This example was part of a personal discussion with Ruedi Seiler from TU Berlin and should illustrate the perturbation theory described above.

Consider $P(\epsilon) = P(0) + \epsilon P^{(1)}$ with

$$P(0) = \frac{1}{2}\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix} \qquad \text{and} \qquad P^{(1)} = \begin{pmatrix} 0 & -1 & 0.5 & 0.5 \\ -1 & 0 & 0.5 & 0.5 \\ 0.5 & 0.5 & -1 & 0 \\ 0.5 & 0.5 & 0 & -1 \end{pmatrix}$$

Let us first consider the unperturbed matrix $P(0)$. It has a twofold eigenvalue $\lambda = 1$ with corresponding eigenvectors

$$X_1(0) = [1,\, 1,\, 0,\, 0]^\top \quad \text{and} \quad X_2(0) = [0,\, 0,\, 1,\, 1]^\top,$$

and a twofold eigenvalue $\lambda = 0$ with eigenvectors

$$X_3(0) = [-1,\, 1,\, 0,\, 0]^\top \quad \text{and} \quad X_4(0) = [0,\, 0,\, -1,\, 1]^\top,$$

The orthogonal projections $\Pi_3(0)$ and $\Pi_4(0)$ on the corresponding eigenspaces are

$$\Pi_3(0) = \frac{1}{2}\begin{pmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \Pi_4(0) = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix}.$$

According to the perturbation theory, the perturbed Perron eigenvectors can be written as a linear combination of the unperturbed vectors plus error terms $\epsilon R_i + O(\epsilon^2)$. In this example, R_1 and R_2 are indeed zero, but this is not satisfied in general as will be shown below. There are some interesting facts in the current example. First note that the eigenvalues of the perturbed matrix $P(\epsilon)$ are given by

$$\{\lambda_i(\epsilon)\}_{i=1}^4 = \{1,\, 1-2\epsilon,\, \epsilon,\, -\epsilon\}.$$

For $\epsilon = 1/3$, the second eigenvalue degenerates, i. e. the corresponding eigenvectors are no longer analytic in ϵ. As long as $\epsilon < 1/3$, the first two eigenvectors of the perturbed matrix are independent of ϵ,

$$X_1(\epsilon) = [1,\, 1,\, 1,\, 1]^\top \quad \text{and} \quad X_2(\epsilon) = [1,\, 1,\, -1,\, -1]^\top.$$

Even more interesting is the fact that they span the same subspace as the unperturbed eigenvectors (that is why $R_{1,2}$ equals zero). Thus, PCCA+ will produce

the same membership vectors χ in the perturbed case as in the unperturbed case, namely

$$\chi_1(\epsilon) = [1,\, 1,\, 0,\, 0]^\top \quad \text{and} \quad \chi_2(\epsilon) = [0,\, 0,\, 1,\, 1]^\top,$$

despite the fact that ϵ can be "large".

Now consider a different perturbation in the form $P(\hat{\epsilon}) = P(0) + \hat{\epsilon}\hat{P}^{(1)}$ with

$$\hat{P}^{(1)} = \begin{pmatrix} 0 & -1 & 0.5 & 0.5 \\ -1 & 0 & 0.75 & 0.25 \\ 0.5 & 0.75 & -0.5 & -0.75 \\ 0.5 & 0.25 & -0.75 & 0 \end{pmatrix}.$$

If $\hat{\epsilon} > 0.5$, $P(\hat{\epsilon})$ will not be positive. As long as $\hat{\epsilon} \leq 0.5$, $P(\hat{\epsilon})$ will be stochastic and reversible. In this case, the term R_2 does not vanish, i.e.

$$\hat{X}_2^{(1)} = [0,\, 0,\, -1,\, 1]^\top$$

is not in the span of the first two eigenvectors.

1.2.2 Basic algorithm

The first cluster algorithm, which was based on the sign structure of the Perron-eigenvectors, was PCCA [21]. Deuflhard and Weber [23] developed the more robust variant PCCA+ by exploiting the simplex structure of the eigenvectors for clustering. In case of a decomposable Markov chain, the rows of the Perron eigenvectors can be considered as vertices of an $(n_C - 1)$-dimensional simplex. Each state can be assigned to one of the n_C vertices and thus to one of the n_C blocks. Under perturbation, the simplex will be disturbed. There will be rows which cannot be assigned uniquely to one of the blocks. Instead, PCCA+ aims at an assignment of states $i \in \{1, \ldots, N\}$ to clusters $j \in \{1, \ldots, n_C\}$ with certain grades of *membership* $\chi(i, j) \in [0, 1]$.

Definition 1.2.2. [109] A set of vectors $\{\chi_j\}_{j=1}^{n_C}$ with $\chi_j \in \mathbb{R}^N$ are called **membership vectors** if they meet the following properties.

1. $\chi_j(i) \geq 0 \quad \forall i \in \{1, \ldots, N\},\, j \in \{1, \ldots, n_C\}$ (positivity),
2. $\sum_{j=1}^{n_C} \chi_j(i) = 1 \quad \forall i \in \{1, \ldots, N\}$ (partition of unity),
3. $\forall j \in \{1, \ldots, n_C\}\, \exists i \in \{1, \ldots, N\} : \chi_j(i) = 1$ (vertex condition).

The membership vectors can be interpreted in the sense of assigning a state i to a cluster j with a certain probability $\chi_j(i)$. Note that the characteristic vectors $\chi_j : \Omega \to \{0, 1\}$ introduced in Section 1.2.1 are a special instance of membership vectors. They are also called *crisp* or *hard* membership vectors, whereas general membership vectors are also denoted as *soft* membership vectors or *soft* characteristic vectors.

The cluster algorithm PCCA+ aims at the generation of a regular transformation matrix $\mathcal{A} \in \mathbb{R}^{n_C \times n_C}$ such that

$$\chi = X\mathcal{A} \tag{1.21}$$

becomes a matrix of membership vectors. Here, X are the right eigenvectors of the matrix P corresponding to the Perron cluster of eigenvalues near $\lambda_1 = 1$. The existence of such a transformation cannot be ensured for general transition matrices. One of the four constraints on χ (positivity, partition of unity, vertex condition, invariant subspace) must be dropped to obtain an approximate solution [109]. Therefore, PCCA+ searches for a matrix $\mathcal{A}$ that optimizes some objective function depending on the missing constraint [109, 23].

Originally, the objective was to maximize *metastability*, which is defined in [21] as

$$\text{trace}(\underbrace{D_c^{-2}\chi^\top D^2 P\chi}_{\mathcal{P}}), \tag{1.22}$$

where $D_c^2 = \text{diag}(\chi^\top \mathbf{w})$. The matrix $\mathcal{P}$ is called *coupling matrix* and can be considered as a coarse grained transition probability matrix. In case of crisp membership vectors $\chi : \Omega \to \{0, 1\}$, it describes transition probabilities between set-based metastable conformations. Thus, the objective is to maximize the holding probabilities within the conformations.

1.2.3 A new objective function

In case of non-overlapping basis function, i.e. $S = I$, the transition probability matrix P can be used to propagate discrete densities $\mathbf{x}(t)$ ($\mathbf{x}_i(t) \geq 0$, $\sum_i \mathbf{x}_i(t) = 1$),

$$\mathbf{x}(t + \tau) = P^\top(\tau)\mathbf{x}(t).$$

Such densities can be restricted to the space of metastable conformations via the membership functions χ,

$$\mathbf{x}_c(t) = \chi^\top \mathbf{x}(t).$$

This leads immediately to the question whether there exists a coarse grained propagator matrix P_c such that propagation and restriction commutate in the sense that

$$\chi^\top P^\top(\tau)\mathbf{x}(t) = P_c^\top(\tau)\chi^\top \mathbf{x}(t). \tag{1.23}$$

This task has been examined in detail in [58, 60]. Indeed, under the condition that χ is an invariant subspace of P (which is satisfied for PCCA+ because χ is a linear transformation of eigenvectors), the coarse propagator matrix

$$P_c = (D_c^{-2}\chi^\top D^2\chi)^{-1} D_c^{-2}\chi^\top D^2 P\chi = (\chi^\top D^2\chi)^{-1}\chi^\top D^2 P\chi$$

satisfies (1.23). The proof is based on the assumption that the Perron eigenvectors of P are orthogonal. Then they can be normalized such that $X^\top D^2 X = I$ and

consequently $P_c = \mathcal{A}^{-1}\Lambda\mathcal{A}$. Thus, the Perron cluster eigenvalues are maintained, which preserves the timescales of conformation dynamics. This can be considered as the main advantage of PCCA+. Note that for crisp membership vectors $\chi_i : \Omega \to \{0, 1\}$ the original definition (1.22) is recovered ($P_c = \mathcal{P}$).

Remark 1.2.3. The matrix

$$\chi_D^\dagger = (\chi^\top D^2 \chi)^{-1} \chi^\top D^2$$

can be considered as weighted pseudo-inverse and has already been examined by G. W. Stewart [101]. The solution of the weighted least squares problem

$$\|D(y - \chi b)\|^2 \to \min_b,$$

where D^2 is a diagonal matrix with positive diagonal elements, is given by $b = \chi_D^\dagger y$.

The coarse graining procedure generalizes to the case of overlapping basis functions by taking into account the mass matrix S. The appropriate coupling matrix is given by

$$P_c = (\underbrace{D_c^{-2}\chi^\top D^2 S\chi}_{\mathcal{S}})^{-1} D_c^{-2}\chi^\top D^2 P\chi = (\chi^\top D^2 S\chi)^{-1}\chi^\top D^2 P\chi. \tag{1.24}$$

This is the coarse analog to the matrix $S^{-1}P$, which can be considered as matrix representation of the transfer operator T^τ in $\mathcal{V}_N = \text{span}\{\phi_1, \ldots, \phi_N\}$, compare Eq. (1.8). Theoretically, the matrix S is generalized symmetric w.r.t. the stationary density $\mathbf{w}$, i.e. $D^2 S = S^\top D^2$ [1]. In this case, the subspace basis X with $PX = SX\Lambda$ can be normalized such that $X^\top D^2 SX = I$. Then

$$P_c = \mathcal{S}^{-1}\mathcal{P} = \mathcal{A}^{-1}\Lambda\mathcal{A}.$$

Optimization of $\text{trace}(P_c)$ or $\det(P_c)$ makes no sense because the expressions only depend on Λ but not on $\mathcal{A}$ or χ, respectively. On the other hand, the matrices $\mathcal{P}$ and $\mathcal{S}$ are stochastic whereas P_c is not, since it can have negative entries. Thus the interpretation of P_c as a Markov chain transition matrix will fail. However, if $\mathcal{S}$ were the identity, then $P_c = \mathcal{P}$ would be a row-stochastic matrix. This motivated the new objective to find a transformation matrix $\mathcal{A}$ such that $\mathcal{S}$ gets close to the identity matrix,

$$\|\mathcal{S}(\mathcal{A}) - I_{n_C}\| \to \min.$$

This corresponds to a minimization of the off-diagonal entries in $\mathcal{S}$. Since $\text{trace}(\mathcal{S}) \leq n_C = \text{trace}(I_{n_C})$, an appropriate objective function is given by

$$F(\mathcal{A}) = \text{trace}(\mathcal{S}(\mathcal{A})) \to \max.$$

[1] The perturbation of symmetry by the sampling scheme will be considered in detail in Section 3.1.3.3.

This equals the original objective function (1.22) with P replaced by S. Thus, the implementation of the algorithm PCCA+ need not to be modified, but requires only a different input argument. The new objective function can be interpreted as making the conformational membership vectors as crisp as possible.

Remark 1.2.4. Different objective functions that force $\mathcal{S}$ towards the identity are possible, for example the objective function used by Shalloway et al. [54]

$$F(\mathcal{A}) = \sum_i \log(\mathcal{S}(i,i)),$$

for which a gradient-based optimization routine has been developed [114]. Alternatively, the objective function

$$F(\mathcal{A}) = \det(\mathcal{S})$$

can be related to holding times in metastable macro-states [109].

1.2.4 Characteristic life times

From [18], the characteristic life times or holding times of conformations are known to behave (in first order) like

$$\tau_i \doteq \frac{\tau}{1 - P_c(i,i)}, \tag{1.25}$$

where τ is the time step used for vertical propagation to approximate P^τ. Clearly, these times are extremely sensitive to perturbations of $P_c(i,i) < 1$ when $P_c(i,i)$ is close to 1.

Therefore, the numbers τ_i should only be interpreted qualitatively, but chemical knowledge must be added to find the true quantitative values for characteristic life times of molecular conformations. Alternatively, one could use a sampling-based approach to obtain statistics of holding times, from which the characteristic life times can be approximated directly by parameter estimation [57].

Chapter 2

Matrix Perturbation Theory

The central question of this chapter is the following: How does a matrix function change when the matrix entries are subject to perturbations? The matrix functions we are concerned with are the Perron cluster eigenvalues and the corresponding invariant or deflating subspace, respectively. These are the quantities that influence the decision for the number of metastable conformations and the assignment of discrete states to the conformations. In other words, the accuracy of these quantities influences the accuracy of the clustering. Emphasis is put on perturbation bounds and the derivation of condition numbers that give insight into the sensitivity of the eigenvalue problems. In some cases, there exist sharper, maybe even optimal bounds, which are mostly more complicated. Instead, we decided to present bounds that are easier to calculate and thus of more practical use.

The sensitivity analysis of eigenvectors and invariant subspaces goes back to Stewart [97, 98, 99, 105], Golub and Wilkinson [37]. Later on, the theory was refined and became part of numerical software [16, 17, 4, 49]. The material in this chapter summarizes these contributions. It prepares the ground for subsequent chapters and does not contain major new aspects in the perturbation analysis of matrix eigenvalue problems. The main focus lies on first order perturbation bounds, which will allow for efficient error estimation in subsequent chapters.

2.1 The standard eigenvalue problem

In this section we will deal with the perturbation of eigenvalues and invariant subspaces of a matrix $A \in \mathbb{R}^{N \times N}$. We are concerned with the question whether the eigenvectors of a perturbed matrix $\widetilde{A} = A + E$ inherit the structure from the unperturbed eigenvectors. At this point an essential difference in comparison to the previous perturbation theory comes into play. We are not able to derive perturbation results for single eigenvectors because they are not unique in the unperturbed case. Instead, we can apply existing perturbation results for invariant subspaces. Since the goal of PCCA+ is to generate a linear transformation in the form of (1.21), the clustering is independent of the choice of basis vectors. In other words,

single columns of X do not play a role, only the subspace they span is of interest. The sensitivity of the Perron root and the corresponding invariant subspace are the topic of the present section.

2.1.1 Preliminaries

In the following, we will introduce the notion of distance between subspaces and the separation of matrices, which are necessary do derive condition numbers for cluster eigenvalues and the corresponding invariant subspace.

Angles between subspaces Denote by $\sigma(A)$ the singular values of a matrix A. Given two subspaces $\mathcal{X}$ and $\mathcal{Y}$ of same dimension with orthonormal bases X and Y, respectively, the largest principal angle is given by [36]

$$\theta_{\max}(\mathcal{X}, \mathcal{Y}) = \arccos \sigma_{\min}(Y^H X) = \arcsin \sigma_{\max}(X_\perp^H Y).$$

The largest principal angle is related to the notion of *distance* between equidimensional subspaces,

$$\operatorname{dist}(\mathcal{X}, \mathcal{Y}) = \sin(\theta_{\max}).$$

In the following, we identify $\sin \Theta(\mathcal{X}, \mathcal{Y}) = \operatorname{diag}(\sin \theta_i(\mathcal{X}, \mathcal{Y}))$. Sometimes, we use the symbol $\mathcal{R}$ if we refer to the column space of a matrix, for example $\mathcal{X} = \mathcal{R}(X)$.

Real valued subspaces The above definition of subspace angles is also valid for complex valued subspaces. Let X and Y be unitary bases where complex vectors appear as conjugate pairs. Then

$$\theta_{\max}(X, Y) = \theta_{\max}([\operatorname{real}(X),\, \operatorname{imag}(X)], [\operatorname{real}(Y),\, \operatorname{imag}(Y)]).$$

A problem arising in the nearly reversible case is that the eigenvectors might become complex valued, which is undesirable for the application of PCCA+. A possible way to circumvent this problem is to work with the real Schur decomposition of A and apply PCCA+ to the real Schur vectors. This is verified by the fact that the real Schur vectors span the same subspace as the corresponding complex Schur vectors which again span the same subspace as the corresponding eigenvectors.

Theorem 2.1.1. *(***Real Schur Decomposition***)*([105], Exc. I.3.24) *If $A \in \mathbb{R}^{N \times N}$, then there exists an orthogonal matrix $U \in \mathbb{R}^{N \times N}$ such that*

$$U^\top A U = T, \tag{2.1}$$

where T is block triangular with 1×1 and 2×2 blocks on its diagonal. The 1×1 blocks contain the real eigenvalues of A, and the eigenvalues of the 2×2 blocks are the complex eigenvalues of A.

Since the eigenvalues of T are easy to compute, the Schur decomposition is a target for iterative eigenvalue algorithms such as the Jacobi-Davidson method [96, 59] or the implicitly re-started Arnoldi method [63]. Since the following theory is valid for general matrices, it is based on the complex Schur decomposition. However, keep in mind that the basis of the invariant subspace can also be written in terms of real vectors.

Spectral projectors and the separation of two matrices Assume that the matrix A has a cluster of n_C eigenvalues with corresponding *block Schur decomposition*

$$[X_1, X_2]^H A [X_1, X_2] = \begin{pmatrix} L_1 & H \\ 0 & L_2 \end{pmatrix}, \tag{2.2}$$

where $X_1 \in \mathbb{R}^{N \times n_C}$. The columns of X_1 and X_2 are orthonormal and form bases for the invariant subspace $\mathcal{X}_1$ and its complement $\mathcal{X}_2$, respectively. The eigenvalues of L_1 are exactly the eigenvalues of the cluster we are interested in. The spectral projector belonging to L_1 is defined as

$$P = (X_1, X_2) \begin{pmatrix} I_{n_C} & R \\ 0 & 0 \end{pmatrix} (X_1, X_2)^H, \tag{2.3}$$

where R satisfies the *Sylvester equation*

$$L_1 R - R L_2 = H. \tag{2.4}$$

This equation has a unique solution if and only if L_1 and L_2 have no eigenvalues in common [105]. The norm

$$\|P\|_2 = (1 + \|R\|_2^2)^{1/2}$$

plays an important role in the bounds. However, in practice, one uses the cheaper overestimate

$$\|P\|' \equiv (1 + \|R\|_F^2)^{1/2}.$$

To develop the perturbation theory, we need two more definitions. The first one is the definition of the *Sylvester operator* $\mathbf{T}$,

$$\mathbf{T}_{F,G} : Q \mapsto QF - GQ.$$

The operator $\mathbf{T}_{L_1,L_2}(Q) = QL_1 - L_2 Q$ will be referred to as $\mathbf{T}_A(Q)$. Second, the *separation* of two matrices F and G is defined as [99]

$$\mathrm{sep}(F, G) := \begin{cases} \|\mathbf{T}_{F,G}^{-1}\|^{-1}, & 0 \notin \lambda(\mathbf{T}) \\ 0, & 0 \in \lambda(\mathbf{T}) \end{cases}, \tag{2.5}$$

where $\lambda(\mathbf{T})$ denotes the spectrum of $\mathbf{T}$. In practice, the norm is taken to be the operator norm subordinate to a consistent family of matrix norms $\|\cdot\|_p$ such that

$$\mathrm{sep}_p(F, G) = \inf_{\|Q\|_p = 1} \|\mathbf{T}_{F,G}(Q)\|_p.$$

Note that $\mathrm{sep}_F(F,G)$ is the smallest singular value of $\mathbf{T}$ [105].

A formulation for $\mathbf{T}_{F,G}(Q)$ can be obtained by using the Kronecker product and the operator *vecr* that stacks the rows of a matrix into a long column vector. Then

$$\mathrm{vecr}(\mathbf{T}_{F,G}(Q)) = K_{\mathbf{T}_{F,G}} \cdot \mathrm{vecr}(Q), \tag{2.6}$$

where $K_{\mathbf{T}_{F,G}}$ is given by

$$K_{\mathbf{T}_{F,G}} = I_{N-n_C} \otimes F^\top - G \otimes I_{n_C}.$$

Thus the action of the inverse $\mathbf{T}_{F,G}^{-1}$ is given by the inverse matrix $K_{\mathbf{T}_{F,G}}^{-1}$. This formulation will play an important role in the stochastic perturbation theory.

2.1.2 Perturbation of right invariant subspaces

We are now concerned with the question of how the eigenvector matrix X_1 of A corresponding to the Perron cluster eigenvalues changes under perturbations of A.

Perturbation bound Let us first summarize the classical perturbation results as presented in, e.g., [105]. Let $\mathcal{X}_1$ be an invariant subspace of A, and let $\tilde{A} = A + E$ be a perturbed matrix. The goal is to show that for sufficiently small E there is an invariant subspace $\tilde{\mathcal{X}}_1$ of $\tilde{A}$, that approaches $\mathcal{X}_1$ as E approaches zero, and to bound their difference in terms of E.

Let X_1 and X_2 form orthonormal bases for $\mathcal{X}_1$ and $\mathcal{X}_2$, respectively. Then $\mathcal{X}_1$ is a *simple* invariant subspace of A ([105], Def. V.1.2) if

$$\lambda(L_1) \cap \lambda(L_2) = \emptyset,$$

where L_1 and L_2 are the matrices of the block Schur decomposition (2.2).

Let $\widetilde{A} = A + E$ be a perturbation of A with

$$[X_1, X_2]^H E [X_1, X_2] = \begin{pmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{pmatrix}. \tag{2.7}$$

The existence of an invariant subspace $\widetilde{\mathcal{X}}_1$ for $\widetilde{A}$ and a bound on $\sin\Theta(\mathcal{X}_1, \widetilde{\mathcal{X}}_1)$ are specified by the following theorem.

Theorem 2.1.2. ([99], Thm.4.11.) *Let A be a matrix with a simple invariant subspace $\mathcal{X}_1$ and corresponding block Schur representation (2.2). Consider the perturbation $\widetilde{A} = A + E$ where E satisfies (2.7). Moreover, set*

$$\gamma = \|E_{21}\|, \qquad \eta = \|H\| + \|E_{12}\|, \qquad \delta = \mathrm{sep}(L_1, L_2) - \|E_{11}\| - \|E_{22}\|.$$

The conditions

$$\delta > 0 \qquad \textit{and} \qquad \frac{\gamma\eta}{\delta^2} < \frac{1}{4} \tag{2.8}$$

ensure that the eigenvalues of L_1 and L_2 do not coalesce under perturbation such that there exists a simple invariant subspace $\widetilde{\mathcal{X}}_1$ for $\widetilde{A}$ with an orthonormal basis $\widetilde{X}_1$ given by

$$\widetilde{X}_1 = (X_1 + X_2 Q)(I + Q^H Q)^{-\frac{1}{2}}.$$

Q is the solution of the Sylvester equation

$$\mathbf{T}_{\widetilde{A}}(Q) = Q(L_1 + E_{11}) - (L_2 + E_{22})Q = E_{21} - Q(H + E_{12})Q \tag{2.9}$$

and satisfies

$$\|Q\| < 2\frac{\gamma}{\delta}. \tag{2.10}$$

Remark 2.1.3. The conditions (2.8) can be reformulated as a condition on $\|E\|$ by introducing the corresponding left invariant subspace Y_1. The bound reads ([99], Thm.4.14.)

$$\|E\| \leq \frac{1}{4}\frac{\operatorname{sep}(L_1, L_2)}{\|\sec\Theta(\mathcal{X}_1, \mathcal{Y}_1)\|}.$$

Thus, the theorem can only be applied if E is small enough. Otherwise it may happen that $\delta < 0$.

Since

$$\sin\theta_i = \sigma_i(X_2^H \widetilde{X}_1) = \sigma_i(Q(I + Q^H Q)^{-\frac{1}{2}}) = \frac{\sigma_i(Q)}{\sqrt{1 + \sigma_i(Q)^2}},$$

one obtains $\sigma_i(Q) = \tan\theta_i$ and thus

$$\|\sin\Theta(\mathcal{X}_1, \widetilde{\mathcal{X}}_1)\| \leq \|\tan\Theta(\mathcal{X}_1, \widetilde{\mathcal{X}}_1)\| < 2\frac{\gamma}{\delta}. \tag{2.11}$$

The parameter δ indicates the separation of eigenvalues between L_1 and L_2. If L_1 and L_2 are diagonalizable ($L_1 X = X\Lambda_1$ and $Y^H L_2 = \Lambda_2 Y^H$), one obtains [105]

$$\operatorname{sep}_F(L_2, L_1) \geq \frac{\min_{\lambda_2\in\Lambda_2,\,\lambda_1\in\Lambda_1}|\lambda_2 - \lambda_1|}{\kappa_2(X)\kappa_2(Y)}.$$

If some of the eigenvalues are close, then the invariant subspace $\widetilde{\mathcal{X}}_1$ may be distant from $\mathcal{X}_1$. The converse need not be true, i.e. $\operatorname{sep}(L_1, L_2)$ can be small even though the eigenvalues of L_1 and L_2 are well separated [105]. However, if additionally η is small, the coupling between L_1 and L_2 is small and the subspace according to L_1 is well-conditioned.

Remark 2.1.4. The above perturbation results on eigenvalues and eigenvectors suggest the possibility of matrices with ill-conditioned eigenvalues and well conditioned eigenvectors. An example is given in [97]. Therefore, the occurrence of an eigenvalue gap cannot be used as the only criterion for the determination of the size of the perturbed invariant subspace. An additional criterion based on the condition of the invariant subspace should be applied.

First order approximation If E is small compared to A and the spectrum of L_1 is well separated from the spectrum of L_2, then $\|Q\|_F$ is of the order of $\|E\|_F$ [99] and a first-order approximation of (2.9) is given by

$$\mathbf{T}_A(Q) \doteq E_{21}. \tag{2.12}$$

Consequently

$$Q \doteq \mathbf{T}_A^{-1}(E_{21}).$$

With the definitions from Theorem 2.1.2 and (2.5) it follows

$$\|Q\| \leq \frac{\gamma}{\operatorname{sep}(L_1, L_2)} + \mathcal{O}(\|E\|_F^2). \tag{2.13}$$

This local bound is tighter than the global bound (2.10).

Using the vectorized expression (2.6) and the equality [107]

$$\operatorname{vecr}(K_1 K_2 K_3) = (K_1 \otimes K_3^\top)\operatorname{vecr}(K_2),$$

we obtain

$$\boxed{\operatorname{vecr}(Q) \doteq \underbrace{K_{\mathbf{T}_A}^{-1}(X_2^H \otimes X_1^\top)}_{\equiv T} \operatorname{vecr}(E).} \tag{2.14}$$

This illustrates the approximated linear relationship between the error in the entries of A and the error in the subspace represented by Q.

2.1.3 Perturbation of eigenvalues

Now we turn to the perturbation of the cluster eigenvalues of A, which coincide with the eigenvalues of L_1.

Perturbation bound Explicit expressions for $\widetilde{L}_1$ and $\widetilde{L}_2$ in terms of L_1 and L_2 can be derived [99, 105], but it is difficult to interpret them, in particular the eigenvalues of $\widetilde{L}_1$. Different representations can be obtained by working with the spectral resolution of A ([99], Thm.4.13),

$$[Y_1, X_2]^H A [X_1, Y_2] = \begin{pmatrix} L_1 & 0 \\ 0 & L_2 \end{pmatrix}. \tag{2.15}$$

where

$$[X_1, Y_2]^{-1} = [Y_1, X_2]^H.$$

The space $\mathcal{Y}_1 = \mathcal{R}(Y_1)$ is the unique left invariant subspace corresponding to $\mathcal{X}_1$, and L_1 is called the *generalized Rayleigh quotient*. The relation between Y_1 and X_1 is given by

$$Y_1 = X_1 + X_2 R^H,$$

where R solves $L_1R - RL_2 = H$. Moreover,

$$\|Y_1\|_2 = \|Y_2\|_2 = \|\sec\Theta(\mathcal{X}_1, \mathcal{Y}_1)\|_2 \leq \\ \|\sec\Theta(\mathcal{X}_1, \mathcal{Y}_1)\|_F = \|Y_1\|_F = \|Y_2\|_F.$$

Let E be a perturbation matrix with

$$[Y_1, X_2]^H E[X_1, Y_2] = \begin{pmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{pmatrix}. \tag{2.16}$$

The existence of a generalized Rayleigh quotient for the matrix $\widetilde{A} = A+E$ is verified by the following theorem, in which the star $\star$ is used to distinguish the quantities from Theorem 2.1.2.

Theorem 2.1.5. ([105], Thm.V.2.8.) *Assume that the matrix A has a simple invariant subspace $\mathcal{X}_1$ with corresponding spectral resolution (2.15). Consider the perturbation $\widetilde{A} = A + E$ where E satisfies (2.16). Set*

$$\gamma_\star = \|F_{21}\|, \quad \eta_\star = \|F_{12}\|, \quad \delta_\star = \mathrm{sep}(L_1, L_2) - \|F_{11}\| - \|F_{22}\|.$$

Under the condition

$$\delta_\star > 0 \quad \textit{and} \quad \frac{\gamma_\star\eta_\star}{\delta_\star^2} < \frac{1}{4}, \tag{2.17}$$

there exists a unique matrix $Q_\star$ with

$$\|Q_\star\| < 2\frac{\gamma_\star}{\delta_\star} \leq 2\frac{\|E\|_p}{\delta_\star}$$

such that

$$\widetilde{L}_1 = L_1 + F_{11} + F_{12}Q_\star$$

is a generalized Rayleigh quotient of $A + E$. The matrix $Q_\star$ is the solution of the equation

$$Q_\star(L_1 + F_{11}) - (L_2 + F_{22})Q_\star = F_{21} - Q_\star F_{12}Q_\star.$$

Assume that L_1 contains the eigenvalues of interest. In our applications, i.e. for metastable dynamical systems, they are closely clustered and thus highly sensitive to perturbations. However, the average eigenvalue [38]

$$\mu(L_1) \equiv \mathrm{trace}(L_1)/n_C$$

does not suffer from ill-conditioning of single eigenvalues as long as the spectra of L_1 and L_2 are well separated [50, 117]. Let $\mu(\widetilde{L}_1)$ be the mean of the perturbed eigenvalues. Since $|\mu(\widetilde{L}_1) - \mu(L_1)| \leq \|\widetilde{L}_1 - L_1\|$, we obtain

$$|\mu(\widetilde{L}_1) - \mu(L_1)| \leq \|F_{11}\| + 2\frac{\gamma_\star\eta_\star}{\delta_\star}. \tag{2.18}$$

First order approximation If E is small enough, one can disregard the term $F_{12}Q_\star = \mathcal{O}(\|E\|_p^2)$ to obtain

$$\begin{aligned}\mu(\widetilde{L}_1) &\doteq \mu(L_1) + \mu(F_{11}) \\ &= \mu(L_1) + \mu(Y_1^H E X_1) \\ &= \mu(L_1) + \mu((X_1^H + R X_2^H) E X_1)\end{aligned}$$

With $\Delta_\mu = \mu(\widetilde{L}_1) - \mu(L_1)$ the vectorized equation reads

$$\boxed{\Delta_\mu \doteq \underbrace{\left[((\bar{X}_1 + \bar{X}_2 R^\top) \otimes X_1) \frac{\operatorname{vecr}(I_{n_C})}{n_C}\right]^\top}_{\equiv \mathbf{t}^\top} \operatorname{vecr}(E),} \tag{2.19}$$

where $\bar{X}$ denotes the complex conjugate of X. This equation will be the starting point for statistical condition estimates.

2.2 The generalized eigenvalue problem

In this section we will deal with the perturbation of eigenvalues and deflating subspaces of matrix pairs $(A, B) \in \mathbb{R}^{N\times N} \times \mathbb{R}^{N\times N}$. If B was non-singular, the generalized eigenvalue problem $A\mathbf{x} = \lambda B\mathbf{x}$ could be reduced to the standard form $B^{-1}A\mathbf{x} = \lambda\mathbf{x}$. However, an ill-conditioned B might rule out this approach and necessitates a perturbation theory for matrix pairs.

2.2.1 Preliminaries

Throughout this section we assume that the matrix pair (A, B) is regular, i.e. the polynomial $\det(\beta A - \alpha B)$ is not identically zero. If

$$\beta A\mathbf{x} = \alpha B\mathbf{x} \tag{2.20}$$

for $(\alpha, \beta) \neq (0, 0)$ and $x \neq 0$, then (α, β) is a *generalized eigenvalue* of the matrix pair (A, B) with *(right) eigenvector* $\mathbf{x}$. Since all nonzero multiples of such a pair (α, β) also satisfy the eigenvalue equation (2.20), the eigenvalues are assigned to equivalence classes,

$$\langle \alpha, \beta \rangle = \{\tau(\alpha, \beta) : \tau \in \mathbb{C} \setminus \{0\}\}.$$

The set of all eigenvalues of a pair (A, B) is denoted by $\mathcal{L}[(A, B)]$.

An n_C-dimensional subspace $\mathcal{X}$ is called *(right) deflating* subspace of the pair (A, B) if $A\mathcal{X}$ and $B\mathcal{X}$ are contained in a n_C-dimensional subspace $\mathcal{Y}$, the *(left) deflating* subspace. Deflating subspaces allow to reduce the generalized eigenvalue problem to smaller subproblems. One possibility is the reduction to the *generalized Schur decomposition*,

$$V^H A U = R, \quad V^H B U = T$$

with unitary matrices U and V and triangular matrices R and T. For brevity, we will use the notation $(A,B)U = (AU, BU)$ and $V^H(A,B) = (V^HA, V^HB)$. The Schur decomposition will in general be complex, but realness can be preserved if R is allowed to have two-by-two blocks along the diagonal.

Theorem 2.2.1. *(***Generalized Real Schur Decomposition***)*([36], Thm.7.7.2) *If $(A,B) \in \mathbb{R}^{N\times N} \times \mathbb{R}^{N\times N}$, then there exist orthogonal matrices $U \in \mathbb{R}^{N\times N}$ and $V \in \mathbb{R}^{N\times N}$ such that*

$$V^\top(A,B)U = (R,T),$$

where $R \in \mathbb{R}^{N\times N}$ is in real Schur form and $T \in \mathbb{R}^{N\times N}$ is upper triangular.

The Schur decomposition represents an equivalence transformation in that $\mathcal{L}[(R,T)] = \mathcal{L}[(A,B)]$. Moreover, the n_C leading columns of U and V span a pair of deflating subspaces if the element $T(n_C+1, n_C) = 0$. Since the eigenvalues of (R,T) are easy to compute, the generalized Schur decomposition is a target for iterative generalized eigenvalue algorithms [95, 31, 82]. In the following, we will deal with the complex Schur form because the theory is valid for general complex-valued matrices.

U and V can be chosen such that the eigenvalues appear in any order along the diagonals of R and T. Let $V = [Y_1, Y_2]$ and $U = [X_1, X_2]$ be a conformal partitioning w.r.t. a cluster of n_C eigenvalues. Then the corresponding *generalized block Schur decomposition* is given by

$$[Y_1,Y_2]^H(A,B)[X_1,X_2] = \left(\begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix}, \begin{pmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{pmatrix}\right) \tag{2.21}$$

The subspaces $\mathcal{X}_1 = \mathcal{R}(X_1)$ and $\mathcal{Y}_1 = \mathcal{R}(Y_1)$ form a pair of deflating subspaces associated with the cluster of eigenvalues of (A_{11}, B_{11}). For a proof of existence of such a decomposition, see for example [99], Thm.2.1.

In the following, we will work with the norm [98]

$$\|(W,Q)\|_{\mathcal{F}} \equiv \max\{\|W\|_F, \|Q\|_F\}.$$

With $\|(W,Q)\|_F = \sqrt{\|W\|_F^2 + \|Q\|_F^2}$ we have

$$\|(W,Q)\|_{\mathcal{F}} \leq \|(W,Q)\|_F \leq \sqrt{2}\|(W,Q)\|_{\mathcal{F}}. \tag{2.22}$$

One can measure the *separation of two matrix pairs*, $\mathrm{dif}_{A,B}$, as

$$\mathrm{dif}[(A_{11},B_{11}),(A_{22},B_{22})] \equiv \inf_{\|(W,Q)\|_{\mathcal{F}}=1} \|(WA_{11} - A_{22}Q, WB_{11} - B_{22}Q)\|_{\mathcal{F}}. \tag{2.23}$$

Note that in general $\mathrm{dif}[(A_{11},B_{11}),(A_{22},B_{22})] \neq \mathrm{dif}[(A_{22},B_{22}),(A_{11},B_{11})]$. Therefore,the following notation has been introduced [17, 49]:

$$\begin{aligned} \mathrm{dif}_l[(A_{11},B_{11}),(A_{22},B_{22})] &= \mathrm{dif}[(A_{11},B_{11}),(A_{22},B_{22})] \\ &= \mathrm{dif}_u[(A_{22},B_{22}),(A_{11},B_{11})]. \end{aligned} \tag{2.24}$$

The operator

$$\mathbf{T}_{(A_{11},B_{11}),(A_{22},B_{22})} : (W,Q) \mapsto (WA_{11} - A_{22}Q, WB_{11} - B_{22}Q)$$

is called the *generalized Sylvester operator* and will be referred to as $\mathbf{T}_{A,B}$. If $\mathbf{T}_{A,B}$ is nonsingular and $\mathbf{T}_{A,B}(W,Q) = (E,F)$, then [98]

$$\|(W,Q)\|_{\mathcal{F}} \leq \frac{\|(E,F)\|_{\mathcal{F}}}{\mathrm{dif}_{A,B}}. \tag{2.25}$$

There also exists a vectorized form of $\mathbf{T}_{A,B}(W,Q)$:

$$\mathrm{vecr}(\mathbf{T}_{A,B}(W,Q)) = \underbrace{\begin{pmatrix} I_{N-n_C} \otimes A_{11}^{\top} & -A_{22} \otimes I_{n_C} \\ I_{N-n_C} \otimes B_{11}^{\top} & -B_{22} \otimes I_{n_C} \end{pmatrix}}_{K_{\mathbf{T}_{A,B}}} \begin{pmatrix} \mathrm{vecr}(W) \\ \mathrm{vecr}(Q) \end{pmatrix}. \tag{2.26}$$

This will later be useful for the derivation of a first order approximation.

Let (R_r, R_l) be the solution of the *generalized Sylvester equation*

$$A_{11}R_r - R_l A_{22} = A_{12}, \quad B_{11}R_r - R_l B_{22} = B_{12}. \tag{2.27}$$

Then

$$p = (1 + \|R_l\|_F^2)^{1/2}, \qquad q = (1 + \|R_r\|_F^2)^{1/2} \tag{2.28}$$

are the norms of the left and right spectral projectors belonging to $\mathcal{L}[(A_{11}, B_{11})]$.

In the following, we will deal with the influence of a perturbation

$$(\tilde{A}, \tilde{B}) = (A + E, B + F)$$

on the eigenvalues and deflating subspaces.

2.2.2 Perturbation of deflating subspaces

Let us first consider the perturbation of a pair of right and left complementary eigenspaces (X_1, Y_2).

Perturbation bound In the following we present the perturbation result for the right deflating subspace as it has been presented by Stewart and Sun [105].

Theorem 2.2.2. ([105], Thm.2.14.) *Let $\mathcal{X}_1$ be an eigenspace of the regular pair (A,B) and let the pair have the generalized block Schur decomposition (2.21). Consider a perturbation $(\tilde{A}, \tilde{B}) = (A + E, B + F)$ with*

$$[Y_1, Y_2]^H (E,F)[X_1, X_2] = \left(\begin{pmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{pmatrix}, \begin{pmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{pmatrix} \right). \tag{2.29}$$

Set

$$\begin{aligned}\gamma &= \|(E_{21}, F_{21})\|_{\mathcal{F}},\\ \eta &= \|(A_{12}+E_{12}, B_{12}+F_{12})\|_{\mathcal{F}},\\ \delta &= \operatorname{dif}[(A_{11},B_{11}),(A_{22},B_{22})] - \max(\|E_{11}\|_F+\|E_{22}\|_F, \|F_{11}\|_F+\|F_{22}\|_F).\end{aligned}$$

If $\delta > 0$ *and*

$$\frac{\eta\gamma}{\delta^2} < \frac{1}{4},$$

then there exist matrices W *and* Q *as solutions of the generalized Sylvester equation*

$$\mathbf{T}_{\widetilde{A},\widetilde{B}}(W,Q) = (E_{21} - W(A_{12}+E_{12})Q, F_{21} - W(B_{12}+F_{12})Q) \tag{2.30}$$

which satisfy

$$\|(W,Q)\|_{\mathcal{F}} < \frac{2\gamma}{\delta}. \tag{2.31}$$

The columns of

$$\widetilde{X}_1 = (X_1 + X_2Q)(I+Q^HQ)^{-1/2}) \quad \textit{and} \quad \widetilde{Y}_2 = (Y_2 - Y_1W^H)(I+WW^H)^{-1/2}$$

span right and left complementary eigenspaces of $(A+E, B+F)$ *corresponding to the regular pairs*

$$(\tilde{A}_{11},\tilde{B}_{11}) = (A_{11}+E_{11}+(A_{12}+E_{12})Q, B_{11}+F_{11}+(B_{12}+F_{12})Q),$$

$$(\tilde{A}_{22},\tilde{B}_{22}) = (A_{22}+E_{22}-W(A_{12}+E_{12}), B_{22}+F_{22}-W(B_{12}+F_{12})).$$

The spectrum of $(\tilde{A}_{11},\tilde{B}_{11})$ *is disjoint to the spectrum of* $(\tilde{A}_{22},\tilde{B}_{22})$.

Stewart and Sun [105] point out that the norm $\|\cdot\|_{\mathcal{F}}$ can be replaced by any norm that allows the conditions of the theorem to be verified.

The bound on Q bounds the sine of the largest canonical angle between $\mathcal{X}_1$ and $\tilde{\mathcal{X}}_1$,

$$\sin\theta_i(\mathcal{X}_1,\tilde{\mathcal{X}}_1) \le \tan\theta_i(\mathcal{X}_1,\tilde{\mathcal{X}}_1) \le \|Q\|_2 < 2\frac{\gamma}{\delta}.$$

First order approximation If E and F are small compared to A and B and if $\operatorname{dif}_{A,B}$ is large enough, then $\|Q\|_F$ and $\|W\|_F$ are on the order of $\|E\|_F$ and $\|F\|_F$ and a first-order approximation of (2.30) is given by

$$\mathbf{T}_{A,B}(W,Q) \doteq (E_{21},F_{21}). \tag{2.32}$$

Consequently

$$(W,Q) \doteq \mathbf{T}_{A,B}^{-1}(E_{21},F_{21}).$$

Together with (2.25) this yields the asymptotic bound

$$\|(W,Q)\|_{\mathcal{F}} \le \frac{\|(E_{21},F_{21})\|_{\mathcal{F}}}{\operatorname{dif}_{A,B}} + \mathcal{O}(\|(E,F)\|_F^2),$$

which is similar to the bound (2.13) for the standard eigenvalue problem. The vectorized form of (2.32) reads

$$\boxed{\begin{pmatrix}\operatorname{vecr}(W)\\ \operatorname{vecr}(Q)\end{pmatrix} \doteq \underbrace{K_{\mathbf{T}_{A,B}}^{-1}\begin{pmatrix}Y_2^H \otimes X_1^\top & 0\\ 0 & Y_2^H \otimes X_1^\top\end{pmatrix}}_{\equiv Z}\begin{pmatrix}\operatorname{vecr}(E)\\ \operatorname{vecr}(F)\end{pmatrix}.} \tag{2.33}$$

This will later be useful for the statistical error analysis in Section 4.5.

2.2.3 Perturbation of eigenvalues

In the following, we deal with the perturbation of the cluster eigenvalues, which correspond to the eigenvalues of the matrix pair (A_{11}, B_{11}).

Perturbation bound Similarly to the standard eigenvalue problem, the components A_{11} and B_{11} are called *Rayleigh components* of the generalized eigenvalue problem $A - \lambda B$. To provide bounds on these Rayleigh components, we return to the spectral resolution [99]. If $\operatorname{dif}[(A_{11}, B_{11}), (A_{22}, B_{22})] \neq 0$, there are unique matrices V_1 and U_2 such that (X_1, U_2) and (V_1, Y_2) are nonsingular and

$$[V_1, Y_2]^H (A, B)[X_1, U_2] = \left(\begin{pmatrix}A_{11} & 0\\ 0 & A_{22}\end{pmatrix}, \begin{pmatrix}B_{11} & 0\\ 0 & B_{22}\end{pmatrix}\right). \tag{2.34}$$

U_2 and V_1 are computed by

$$U_2 = X_2 - X_1 R_r, \qquad V_1 = Y_1 + Y_2 R_l^H,$$

where (R_r, R_l) solves (2.27). Let E and F be perturbations such that

$$[V_1, Y_2]^H (E, F)[X_1, U_2] = \left(\begin{pmatrix}E_{11} & E_{12}\\ E_{21} & E_{22}\end{pmatrix}, \begin{pmatrix}F_{11} & F_{12}\\ F_{21} & F_{22}\end{pmatrix}\right). \tag{2.35}$$

Note that the matrices on the right hand side are different from the matrices in (2.29), except E_{21} and F_{21}. The following theorem provides the existence of Rayleigh components $\widetilde{A}_{11}$ and $\widetilde{B}_{11}$ for $(A + E) - \lambda(B + F)$.

Theorem 2.2.3. ([105], Thm.VI.2.15.) *Let $\mathcal{X}_1$ be an eigenspace of the regular pair (A, B), and let the pair have the spectral resolution (2.34) which is perturbed according to (2.35). Set*

$$\gamma_\star = \|(E_{21}, F_{21})\|_{\mathcal{F}}, \quad \eta_\star = \|E_{12}, F_{12})\|_{\mathcal{F}},$$

$$\delta_\star = \operatorname{dif}[(A_{11}, B_{11}), (A_{22}, B_{22})] - \max(\|E_{11}\|_F + \|E_{22}\|_F, \|F_{11}\|_F + \|F_{22}\|_F).$$

If

$$\delta_\star > 0 \quad \textit{and} \quad \frac{\eta_\star \gamma_\star}{\delta_\star^2} < \frac{1}{4},$$

then there exist matrices $W_\star$ and $Q_\star$ satisfying

$$\|(W_\star, Q_\star)\|_{\mathcal{F}} \leq 2\frac{\gamma_\star}{\delta_\star}$$

such that

$$(\widetilde{A}_{11}, \widetilde{B}_{11}) = (A_{11} + E_{11} + E_{12}Q_\star, B_{11} + F_{11} + F_{12}Q_\star) \quad (2.36)$$
$$(\widetilde{A}_{22}, \widetilde{B}_{22}) = (A_{22} + E_{22} - W_\star E_{12}, B_{22} + F_{22} - W_\star F_{12}) \quad (2.37)$$

are the Rayleigh components of $(A + E) - \lambda(B + F)$.

The theorem states that the Rayleigh components $\widetilde{A}_{11}$ and $\widetilde{B}_{11}$ are accurate to terms $E_{12}Q_\star$ and $F_{12}Q_\star$ which are of order $\|E\|_F^2$ and $\|F\|_F^2$, respectively.

Assume that B_{11} is non-singular, i.e. (A_{11}, B_{11}) has only finite eigenvalues. As for the standard eigenvalue problem, we are interested in the perturbation of the mean eigenvalue

$$\mu(B_{11}^{-1}A_{11}) \equiv \text{trace}(B_{11}^{-1}A_{11})/n_C.$$

Let $\widetilde{\mu} = \mu(\widetilde{B}_{11}^{-1}\widetilde{A}_{11})$. Then [49]

$$\begin{aligned} |\mu - \widetilde{\mu}| &\leq \|B_{11}^{-1}A_{11} - \widetilde{B}_{11}^{-1}\widetilde{A}_{11}\|_F \\ &\leq \|B_{11}^{-1}\|_2\|A_{11} - \widetilde{A}_{11}\|_F + \|B_{11}^{-1} - \widetilde{B}_{11}^{-1}\|_F\|\widetilde{A}_{11}\|_2 \\ &\leq \|B_{11}^{-1}\|_2(\|A_{11} - \widetilde{A}_{11}\|_F + \|B_{11} - \widetilde{B}_{11}\|_F\|\widetilde{B}_{11}^{-1}\|_2\|\widetilde{A}_{11}\|_2) \end{aligned} \quad (2.38)$$

Theorem 2.2.3 provides bounds on $\|A_{11} - \widetilde{A}_{11}\|_F$ and $\|B_{11} - \widetilde{B}_{11}\|_F$. Equivalently to (2.18) we will use the bounds

$$\|A_{11} - \widetilde{A}_{11}\|_F \leq \|E_{11}\| + 2\frac{\eta_\star\gamma_\star}{\delta_\star}, \quad \|B_{11} - \widetilde{B}_{11}\|_F \leq \|F_{11}\| + 2\frac{\eta_\star\gamma_\star}{\delta_\star}.$$

Different bounds in terms of $\|(E, F)\|$ [99] or in terms of the spectral projectors p and q from (2.28) [49] can be derived as well.

In practice, the terms $\|\widetilde{B}_{11}^{-1}\|_2$ and $\|\widetilde{A}_{11}\|_2$ are difficult to evaluate because the exact size of the perturbation is unknown. Since

$$\widetilde{B}_{11}^{-1} = B_{11}^{-1} - B_{11}^{-1}F_{11}B_{11}^{-1} + \mathcal{O}(\|F\|^2),$$

we can estimate

$$\|\widetilde{B}_{11}^{-1}\|_2 \leq \|B_{11}^{-1}\|_2(1 + \|B_{11}^{-1}\|_2\|F_{11}\|_s),$$

which is actually a first order approximation. However, the error contained in this bound is compensated by using the sub-multiplication rule in the derivation of (2.38). Moreover,

$$\|\widetilde{A}_{11}\|_2 \leq \|A_{11}\|_2 + \|E_{11}\|_s + 2\frac{\eta_\star\gamma_\star}{\delta_\star}.$$

First order approximation By ignoring the terms $E_{12}Q_\star$ and $F_{12}Q_\star$ in (2.36) we obtain

$$\begin{aligned}\widetilde{B}_{11}^{-1}\widetilde{A}_{11} &\doteq (B_{11}+F_{11})^{-1}(A_{11}+E_{11})\\ &= (B_{11}^{-1}-B_{11}^{-1}F_{11}B_{11}^{-1}+\mathcal{O}(\|F\|^2))(A_{11}+E_{11})\\ &= B_{11}^{-1}A_{11}+B_{11}^{-1}E_{11}-B_{11}^{-1}F_{11}B_{11}^{-1}A_{11}+\mathcal{O}(\|(E,F\|^2).\end{aligned}$$

Consequently, the mean cluster eigenvalue difference $\Delta_\mu = \mu(\widetilde{B}_{11}^{-1}\widetilde{A}_{11}) - \mu(B_{11}^{-1}A_{11})$ can be written in vectorized form as

$$\Delta_\mu \doteq \frac{1}{n_C}\mathrm{vecr}(I_{n_C})^\top\left[((B_{11}^{-1}V_1^H)\otimes X_1^\top)\mathrm{vecr}(E) - ((B_{11}^{-1}A_{11}B_{11}^{-1}V_1^H)\otimes X_1^\top)\mathrm{vecr}(F)\right]. \tag{2.39}$$

This is often a very good approximation. With

$$\begin{aligned}\mathbf{z}_1^\top &\equiv \frac{1}{n_C}\mathrm{vecr}(I_{n_C})^\top((B_{11}^{-1}V_1^H)\otimes X_1^\top),\\ \mathbf{z}_2^\top &\equiv \frac{1}{n_C}\mathrm{vecr}(I_{n_C})^\top((B_{11}^{-1}A_{11}B_{11}^{-1}V_1^H)\otimes X_1^\top),\end{aligned} \tag{2.40}$$

the vectorized form can be written briefly as

$$\boxed{\Delta_\mu \doteq \mathbf{z}_1^\top\mathrm{vecr}(E) - \mathbf{z}_2^\top\mathrm{vecr}(F).} \tag{2.41}$$

We will return to this formulation in Section 4.5.

Chapter 3

Metastable, Nearly Reversible Markov Chains

The discretization of a reversible, metastable dynamical system leads to a symmetric, nearly decomposable transition probability matrix that describes the essential dynamic behavior between the discretized states. The assignment of states to different metastable regions comes as the result of the cluster algorithm PCCA+, which reveals the hidden block structure of the matrix. PCCA+ is based on the fact that the constant level patterns of the dominant eigenvectors of a symmetric, block-diagonal transition probability matrix are well-conditioned w.r.t. small perturbations of the matrix. However, the matrices resulting from our applications are only nearly reversible. The reason is that the matrices are computed row-wise and that trajectories initiated in different basis functions have different statistical weights. Thus, a simple symmetrization by evaluating trajectories in both directions is impossible. Nevertheless, as long as the number of Monte-Carlo sampling points in the computation of matrix entries is large enough, the truncation errors and thus the deviation from reversibility will be small. Moreover, the existence of transition states increases the departure from the completely decomposable Markov chain. Despite these deficiencies, PCCA+ turned out to be robust. To explain this behavior, we will show that the matrices generated in our simulations are close to some model matrices that have the desired properties. In the first section, we will consider Markov chains as they result from a discretization by Voronoi cells. The second section deals with the matrix pair (P, S) as it results from a discretization by radial basis functions. Such a matrix pair cannot be interpreted as a transition matrix of a discrete Markov chain. Nevertheless, it can be considered as a perturbation thereof. That means, under certain conditions the metastable structure can be recovered from a matrix pair. However, it is the main result of this chapter that the shape parameter α of radial basis functions must be large enough to maintain metastability.

3.1 Standard Markov chains

If the molecular state space is discretized by Voronoi cells, the continuous dynamics can be modeled as a Markov chain on the finite state space. The perturbation theory in Section 1.2.1 was based on the assumption that the states of the resulting transition probability matrix can be reordered such that the matrix becomes nearly block diagonal, whereas each block corresponds to one metastable conformation. Such Markov chains are known as *nearly completely decomposable* or *nearly reducible* Markov chains [10, 100, 70]. Moreover, reversibility in the continuous system is reflected by the fact that the resulting matrix meets the detailed balance condition.

However, there are two deficiencies in the former theory. First, the detailed balance assumption is often violated for matrices stemming from numerical simulations. This is due to truncation errors in Monte-Carlo sampling methods and the lack of a trajectory-based symmetrization method. Second, in many of our applications the departure from the completely decomposable block-structured transition matrix is too large due to the existence of *transition states*. Transition states are not only an artefact of the discretization, but they have an analog in the continuous systems in form of entropic or enthalpic energy barriers on the potential energy surface.

3.1.1 Properties of transition matrices

In our applications the states of the Markov chain stem from a discretization of some molecular potential energy surface. Regions in configuration space that separate different metastable conformations can be energetic or entropic barriers. Energy barriers can often be identified with saddle points of the potential energy surface, whereas entropic barriers are characterized by narrow valleys connecting different metastable regions. Roughly described, the discretization gives rise to two different kinds of states: cluster states $\mathcal{C}$, which are located near the center of a metastable conformation, and transition states $\mathcal{T}$, which are located near energetic or entropic barriers. The non-unique assignment of transition states to the conformations gave rise to the idea of soft clustering. Particles located in transition regions will rapidly move to one of the nearby conformations, whereas particles in cluster states will tend to stay in the states belonging to that specific conformation and rarely switch to a different metastable region.

This behavior is reflected by the corresponding transition probability matrix $\widetilde{P}$. Assume that the states have been ordered into **c**luster states and **t**ransition states and consider the matrix

$$\begin{pmatrix} \widetilde{P}_{cc} & \widetilde{P}_{ct} \\ \widetilde{P}_{tc} & \widetilde{P}_{tt} \end{pmatrix}.$$

If, additionally, the cluster states are ordered according to the conformations, the sub-matrix $\widetilde{P}_{cc}$ will take the well-known nearly-block-diagonal form [100]. Moreover, the entries in $\widetilde{P}_{ct}$ will be small because transitions from the conformations to

transition states are rare events. From now on, we assume that the Markov chains we deal with satisfy the following properties.

1. $\widetilde{P}$ is a primitive stochastic matrix.
2. The Markov chain is nearly reversible, i.e. there exists a small number $\varepsilon_r > 0$ such that
$$|\widetilde{w}(i)\widetilde{P}(i,j) - \widetilde{w}(j)\widetilde{P}(j,i)| \leq \varepsilon_r \quad \forall\, i,j. \tag{3.1}$$
This corresponds to a slight violation of the detailed balance condition.
3. The states in $\mathcal{C}$ can be reordered such that the sub-matrix $\widetilde{P}_{cc}$ is of block-diagonally dominant form
$$\widetilde{P}_{cc} = \begin{pmatrix} \widetilde{P}_{11} & E_{12} & \cdots & E_{1n_C} \\ E_{21} & P_{22} & \cdots & E_{2n_C} \\ \vdots & \vdots & & \vdots \\ E_{n_C1} & E_{n_C2} & \cdots & \widetilde{P}_{n_Cn_C} \end{pmatrix}$$
with
$$\widetilde{P}_{kk}\mathbf{e} > 1 - \varepsilon_t, \quad k = 1, \ldots, n_C, \quad 0 \leq \varepsilon_t \ll 1, \tag{3.2}$$
and
$$\widetilde{P}_{kk}(i,j) > \kappa_c, \quad i,j \in \mathcal{C}_k, \quad \varepsilon_t \ll \kappa_c < 1. \tag{3.3}$$

Inequality (3.2) represents the condition for weak coupling between metastable conformations, whereas inequality (3.3) ensures that the dynamical system is rapidly mixing within a metastable conformation on a timescale much smaller than the transition time τ. However, this last inequality is critical for states which are not classified as transition regions but located at the boundary of metastable conformations. In other words, the larger the overlap between metastable conformations, the smaller κ_c, which complicates the calculation of partial densities. The existence of a unique stationary distribution $\mathbf{w}$ is a direct consequence of item 1 [9, 53]. Matrices with the above properties can be considered as perturbations of a simplified model problem that will be explained in the next section.

3.1.2 Decomposable Markov chains with transient states

The stochastic model matrix P, which will be the basis for the upcoming perturbation theory, has a form as illustrated in Figure 3.1. The cluster states $\mathcal{C}$ belong to one of the metastable conformations and can thus be assigned to one of the blocks in P_{cc}. The $n_t = |\mathcal{T}|$ states represented by the last rows and columns are transient states. Starting in a transient state, the Markov chain can switch to another state, but once it reaches a metastable cluster, it stays there. Note that the discrete stationary density $\mathbf{w}$ is zero for transient states.

Now we want to examine the eigenvalues and eigenvectors of such a transition probability matrix with transient states. The number of stable clusters is denoted

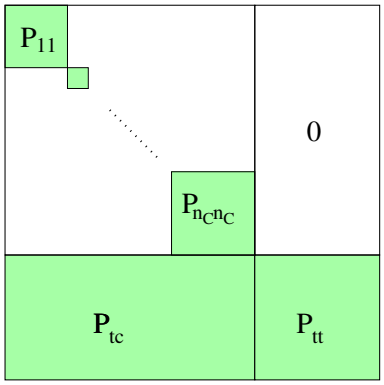

Figure 3.1: Structure of an unperturbed transition matrix with n_C stable clusters and a number of transition states, represented by the last columns and rows.

by n_C. Assume that the blocks P_{jj}, $j = 1, \ldots, n_C$, are primitive and that the corresponding Markov chains are rapidly mixing. That means, each block P_{jj} gives rise to an eigenvalue $\lambda_{j1} = 1$ and to eigenvalues λ_{jk} well-separated away from 1. Thus, the overall matrix P has an n_C-fold eigenvalue $\lambda = 1$ and a set of eigenvalues $\{\lambda_k\}_{k=n_C+1}^{N}$ with $|\lambda_k| < 1$. Of course, the complete matrix P is not primitive because any exponential P^m possesses the same structure as P and hence there exists no $m \in \mathbb{N}$ such that $P^m > 0$ element-wise.

Let us first consider the stationary distribution $\mathbf{w}$, which is defined as left eigenvector of P corresponding to the eigenvalue 1,

$$\mathbf{w}^\top P = \mathbf{w}^\top.$$

In fact, there is no unique stationary distribution, but only a unique left invariant subspace $\mathcal{X}$ spanned by the following vectors,

$$\mathcal{X} = \text{span}\{ \begin{bmatrix} \mathbf{w}_{C_1} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \ldots, \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \mathbf{w}_{C_{n_C}} \\ 0 \end{bmatrix} \}, \tag{3.4}$$

where the vectors $\{\mathbf{w}_{C_1}, \ldots, \mathbf{w}_{C_{n_C}}\}$ are the unique stationary distributions of the sub-chains. In other words, the Markov chain has several equilibrium states. Note that the last n_t components of all vectors are zero, which means that in every equilibrium state the probability of being in a transient state is zero.

Similarly, the right eigenvectors according to the Perron root are not unique, in contrast to the corresponding subspace. In the following, we call the subspace spanned by these eigenvectors the *Perron subspace*. We want to show that there exists a basis of membership vectors for this subspace. Denote by I_j, $j = 1, \ldots, n_C$, the set of indices which form the block P_{jj}. Moreover, the set $(t_1, \ldots t_{n_t})$ contains

the indices of the transient states. Define a matrix $\chi = [\chi_1, \ldots, \chi_{n_C}] \in \mathbb{R}^{N \times n_C}$ and decompose it into two parts,

$$\chi = \begin{pmatrix} \chi_c \\ \chi_t \end{pmatrix}, \quad \chi_c \in \mathbb{R}^{(N-n_t)\times n_C}, \quad \chi_t \in \mathbb{R}^{n_t \times n_C}.$$

For χ to be a basis of the Perron subspace, it must satisfy

$$P\chi = \lambda\chi, \quad \lambda = 1.$$

The equation can be rewritten in terms of χ_c and χ_t as

$$\begin{pmatrix} P_{cc} & 0 \\ P_{tc} & P_{tt} \end{pmatrix} \begin{pmatrix} \chi_c \\ \chi_t \end{pmatrix} = \begin{pmatrix} \chi_c \\ \chi_t \end{pmatrix},$$

which equals the following system of equations:

$$P_{cc}\chi_c = \chi_c \quad \text{and} \quad P_{tc}\chi_c + P_{tt}\chi_t = \chi_t.$$

Define the entries $\chi_j(i), i = 1, \ldots, N - n_t$, $j = 1, \ldots, n_C$, of the upper matrix χ_c as

$$\chi_j(i) = \begin{cases} 1, & \text{if } i \in I_j, \\ 0, & \text{if } i \notin I_j \text{ and } i \text{ is not a transition state.} \end{cases}$$

Thus, χ_c satisfies $P_{cc}\chi_c = \chi_c$. The matrix χ_t must satisfy the system of equations

$$(I - P_{tt})\chi_t = P_{tc}\chi_c,$$

which is equal to the systems

$$(I - P_{tt})\chi_t(:, j) = [\sum_{k\in I_j} P(t_1, k), \ldots, \sum_{k\in I_j} P(t_{n_t}, k)]^\top, \quad j = 1, \ldots, n_C. \tag{3.5}$$

Now assume that P_{tt} is irreducible. Otherwise, the states from $\mathcal{T}$ could be reordered such that P_{tt} becomes block-diagonal with irrreducible blocks. Then one could decompose χ_t accordingly and split the systems (3.5) into further subsystems.

To show that $I - P_{tt}$ is non-singular, we use the following Lemma.

Lemma 3.1.1. *Let $A \geq 0$ be an irreducible matrix with row sums strictly smaller than one, i.e. $\sum_j A(i,j) < 1 \, \forall i$. Then any eigenvalue λ of A is located within the unit disc ($|\lambda| < 1$).*

Proof. The Perron-Frobenius Theorem ([6], Thm.1.4.4) ensures the existence of a positive eigenvector $y > 0$ corresponding to an eigenvalue $\lambda_0 > 0$ that is maximal in modulus among all the eigenvalues of A. It remains to show that $\lambda_0 < 1$:

$$\lambda_0 y_i = \sum_j A(i,j)y_j \leq \max_k y_k \sum_j A(i,j) < \max_k y_k, \quad \forall i$$

The inequality is especially satisfied for $i = \text{argmax}_k(y_k)$, i.e. $\lambda_0 y_{\max} < y_{\max}$. Thus $\lambda_0 < 1$. □

The matrix P_{tt} is in fact non-negative. The assumption that the row sums of P_{tt} are strictly less than one is reasonable because there are transitions to the clusters, i.e. the row sums of P_{tc} are greater than zero. Consequently, the eigenvalues of P_{tt} are within the unit disc and thus $I - P_{tt}$ is regular ($I - P_{tt}$ is actually a nonsingular M-matrix). Therefore, equation (3.5) has a unique solution.

Lemma 3.1.2. *The solution vectors $\chi_t(:, j)$ of (3.5) are non-negative and satisfy $\sum_{j=1}^{n_C} \chi_t(:, j) = 1$.*

Proof. Since $(I - P_{tt})$ is a non-singular M-matrix, the inverse $(I - P_{tt})^{-1}$ exists and is non-negative ([6], Thm.1.5.2). Moreover, the right hand sides of (3.5) are non-negative, such that χ_t will also be non-negative. Furthermore, set $A \equiv (I - P_{tt})$, $\mathbf{x}_j \equiv \chi_t(:, j)$, $\mathbf{s} \equiv \sum_j \chi_t(:, j)$, and $\mathbf{b}_j \equiv [\sum_{k \in I_j} P(t_1, k), \ldots, \sum_{k \in I_j} P(t_{n_t}, k)]^\top$. We have to show that $\mathbf{s}$ is the vector of ones, i.e. $\mathbf{s} = \mathbf{e}$. Since P is stochastic, it holds $\sum_j \mathbf{b}_j = \sum_j A(:, j)$. Hence

$$\sum_j A(i, j) = \sum_j \mathbf{b}_j(i) = \sum_j \sum_k A(i, k)\mathbf{x}_j(k) = \sum_k A(i, k)\mathbf{s}(k).$$

Consequently, $\mathbf{s}$ satisfies $A(\mathbf{s} - \mathbf{e}) = 0$. Since A is regular, it follows $\mathbf{s} = \mathbf{e}$. □

To summarize, the vectors χ_j meet all properties of membership vectors; compare Definition 1.2.2. They can be interpreted in the sense of assigning a state i to a cluster j with a certain probability $\chi_j(i)$. Since the subspace is unique, any eigenvector basis X corresponding to the Perron root $\lambda = 1$ can be transformed linearly into such membership vectors, i.e. there exists a non-singular transformation matrix $\mathcal{A} \in \mathbb{R}^{n_C \times n_C}$ such that $\chi = X\mathcal{A}$.

3.1.3 Perturbation theory

In Markov chains stemming from our simulations, the clusters are not completely decoupled but interact either via the transition states (on small time scales) or directly (on larger time scales). They can be considered as a perturbation of the model matrix P from the previous section,

$$\widetilde{P} = P + \begin{pmatrix} E_{cc} & E_{ct} \\ 0 & 0 \end{pmatrix}. \tag{3.6}$$

From (3.2), it follows that $\|E\|_\infty < \varepsilon_t$. In the following, we will specify some perturbation results on the stationary density, the invariant subspace, and the cluster eigenvalue of this model problem.

3.1.3.1 Perturbation of the stationary distribution

It is well known that the stationary distribution of a metastable Markov chain is ill-conditioned because the second largest eigenvalue is not well-separated from the

Perron root $\lambda = 1$ [71]. Perturbation bounds have been derived for nearly uncoupled Markov chains by several authors, e.g. [72, 70, 103, 7], but only a few dealt with block-wise perturbations [104, 47, 46]. Therefore, we think it is worth to summarize and specify the results for our specific problem.

Stationary density of nearly transient states If the Markov chain has the properties described in Section 3.1.1, then the stationary density of transition states will be small. We will present a theorem for this fact which follows the ideas of Stewart [104] who examined a slightly different problem. To circumvent singularity of $I - P_{cc}$, the problem is transformed via a unitary similarity transformation.

Let $\{\lambda_i\}_{i=1}^{n_C} = 1$ be the (maybe multiple) Perron eigenvalue of P_{cc} and let U_1 be an orthonormal basis for the corresponding left invariant subspace such that every column of U_1 is non-negative. Such a basis can be obtained by normalizing the column vectors in (3.4). Let U_2 be a basis for the complementary subspace such that $U = (U_1, U_2)$ is orthonormal. Now apply the similarity transformation with the matrix $\mathrm{diag}(U, I)$ to P and let

$$B \equiv \begin{pmatrix} U^\top & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} P_{cc} & 0 \\ P_{tc} & P_{tt} \end{pmatrix} \begin{pmatrix} U & 0 \\ 0 & I \end{pmatrix} = \begin{pmatrix} B_{11} & 0 & 0 \\ B_{21} & B_{22} & 0 \\ B_{31} & B_{32} & B_{33} \end{pmatrix},$$

where $B_{11} = I$ and $B_{33} = P_{tt}$. Equivalently, let

$$F \equiv \begin{pmatrix} U^\top & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} E_{cc} & E_{ct} \\ 0 & 0 \end{pmatrix} \begin{pmatrix} U & 0 \\ 0 & I \end{pmatrix} = \begin{pmatrix} F_{11} & F_{12} & F_{13} \\ F_{21} & F_{22} & F_{23} \\ 0 & 0 & 0 \end{pmatrix}.$$

Let

$$\widetilde{\mathbf{w}}^\top = (\widetilde{\mathbf{w}}_c^\top, \widetilde{\mathbf{w}}_t^\top)$$

be the Perron vector of $\widetilde{P} = P + E$. Set

$$\beta = \|P\|, \quad \eta = \|E\|, \quad \gamma_i = \|(I - B_{ii})^{-1}\|, \quad \widetilde{\gamma}_i = \|(I - \widetilde{B}_{ii})^{-1}\|, \; i = 2, 3. \quad (3.7)$$

The following theorem bounds $\widetilde{\mathbf{w}}_t^\top$.

Theorem 3.1.3. *Let P be a stochastic matrix with decomposition*

$$P = \begin{pmatrix} P_{cc} & 0 \\ P_{tc} & P_{tt} \end{pmatrix},$$

where P_{tt} is non-singular, and E a perturbation matrix with decomposition

$$E = \begin{pmatrix} E_{cc} & E_{ct} \\ 0 & 0 \end{pmatrix}.$$

Assume that $\widetilde{P} = P + E$ is stochastic and has a unique stationary distribution $\widetilde{\mathbf{w}}^\top = (\widetilde{\mathbf{w}}_c^\top, \widetilde{\mathbf{w}}_t^\top)$. Given the notations from (3.7), assume that

$$\gamma_3 \eta \beta \widetilde{\gamma}_2 < 1.$$

Then $\widetilde{\mathbf{w}}_t$ will be small in that

$$\|\widetilde{\mathbf{w}}_t^\top\| \leq \frac{\gamma_3(\eta + \eta^2\widetilde{\gamma}_2)}{1 - \gamma_3\eta\beta\widetilde{\gamma}_2}.$$

The proof is given in the appendix A.1. If we take the ∞-norm, then $\beta = 1$ and $\eta \leq \varepsilon_t$. Since P_{tt} is small, γ_3 will be of moderate size. Moreover, as long as the coupling between the metastable states in P_{cc} is small, $\widetilde{\gamma}_2$ will be small. Consequently, the states in $\mathcal{T}$ remain nearly transient.

Condition of partial densities Assume that the stationary density $\widetilde{\mathbf{w}}$ is partitioned as

$$\widetilde{\mathbf{w}}^\top = (\widetilde{\mathbf{s}}_1^\top, \ldots, \widetilde{\mathbf{s}}_{n_C}^\top, \widetilde{\mathbf{s}}_t^\top),$$

where

$$\widetilde{\mathbf{s}}_i^\top = \frac{\widetilde{\mathbf{w}}_i^\top}{\|\widetilde{\mathbf{w}}_i^\top\|}$$

If $\widetilde{P}$ was reversible, the partial densities $\widetilde{\mathbf{s}}_i$ could be computed exactly by the following aggregation step [112]. Assume that the states have been reordered such that $\widetilde{P}_{cc}$ is nearly block diagonal. The row sums of $\widetilde{P}_{ct}$ plus the off-block-diagonal entries in $\widetilde{P}_{cc}$ are added to the diagonal of $\widetilde{P}_{cc}$. Then the entries of $\widetilde{P}_{ct}$ and the off-block-diagonal entries in $\widetilde{P}_{cc}$ are set to zero. The rows of $\widetilde{P}$ corresponding to transition states remain unchanged. The thus constructed matrix has the structure of the model matrix P from Section 3.1.2, and its partial densities $\mathbf{s}_i$, $i = 1, \ldots, n_C$, are equal to the partial densities $\widetilde{\mathbf{s}}_i$.

However, if $\widetilde{P}$ is not reversible, the constructed partial densities $\mathbf{s}_i$ are only approximations to $\widetilde{\mathbf{s}}_i$, but the relative error will be small as shown by the following theorem.

Theorem 3.1.4. *Let $\widetilde{P}$ be an irreducible stochastic matrix with the properties described in Section 3.1.1. Set $\varepsilon = \varepsilon_t/\kappa_c$. Moreover, let P be the matrix that is obtained from $\widetilde{P}$ by setting the off-blockdiagonal entries of $[\widetilde{P}_{cc}, \widetilde{P}_{ct}]$ to zero and adding the deleted row-sums to the diagonal. Then the partial steady state vectors $\mathbf{s}_k$ and $\widetilde{\mathbf{s}}_k$ of the matrices P and $\widetilde{P}$ satisfy*

$$\frac{\|\widetilde{\mathbf{s}}_k - \mathbf{s}_k\|_\infty}{\|\mathbf{s}_k\|_\infty} \leq (1+\varepsilon)^{N_k} - 1 = N_k\varepsilon + \mathcal{O}(\varepsilon^2), \quad k = 1, \ldots, n_C.$$

For ease of readability, the proof is shifted to the Appendix A.2.

Once the partial stationary density vectors $\widetilde{\mathbf{s}}_k$ have been computed, their aggregation into the overall steady state vector [70]

$$\widetilde{\mathbf{w}}^\top = (\widetilde{\xi}_1\widetilde{\mathbf{s}}_1^\top, \ldots, \widetilde{\xi}_t\widetilde{\mathbf{s}}_t^\top)$$

requires the computation of coupling factors $\widetilde{\xi}_k$. These coupling factors are the entries of the unique stationary density of the stochastic, irreducible aggregation matrix $\widetilde{C}$ defined by [70]

$$\widetilde{C}(i,j) = \widetilde{\mathbf{s}}_i^\top \widetilde{P}_{ij}\mathbf{e}.$$

However, the coupling factors are sensitive to perturbations in P [103]. Therefore, a more stable algorithm as proposed in [112] should be used.

3.1.3.2 Perturbation of cluster eigenvalues and invariant subspace

Theorem 2.1.5 in Section 2.1.3 has been stated in terms of the projected error matrices. This formulation will be of practical use in the statistical error analysis, but the expressions in the theorem are difficult to interpret for our model problem. Therefore, we repeat the theorem here in a slightly different form.

Theorem 3.1.5. [116, 99, 16, 4] *Let X_1 be an orthonormal basis of a simple invariant subspace $\mathcal{X}_1$ of a matrix $A \in \mathbb{R}^{n\times n}$ and let $\mathcal{Y}_1$ be its corresponding left invariant subspace. Let P and $\mathrm{sep}(L_1, L_2)$ be defined as in (2.3) and (2.5). Assume there is a perturbation matrix $E \in \mathbb{R}^{n\times n}$ with $\|E\|_F = \varepsilon_F$. Then there exists a perturbed right-invariant subspace $\widetilde{\mathcal{X}}_1$ and a generalized Rayleigh quotient $\widetilde{L}_1$ of $A + E$ such that*

$$\theta_{\max}(\mathcal{X}_1, \widetilde{\mathcal{X}}_1) \leq \frac{\varepsilon_F}{\mathrm{sep}(L_1, L_2)} + \mathcal{O}(\varepsilon_F^2)$$

and

$$\|L_1 - \widetilde{L}_1\|_F \leq \varepsilon_F \|P\|_2 + \mathcal{O}(\varepsilon_F^2).$$

Moreover,if

$$\|E\|_F < \frac{\mathrm{sep}(L_1, L_2)}{4\|P\|_2} \tag{3.8}$$

then

$$\theta_{\max}(\mathcal{X}_1, \widetilde{\mathcal{X}}_1) \leq \arctan\left(\frac{2\varepsilon_F}{\mathrm{sep}(L_1, L_2) - 4\varepsilon_F\|P\|_2}\right),$$

and

$$\|L_1 - \widetilde{L}_1\|_F < 2\varepsilon_F \|P\|_2.$$

Thus, errors of order ε_F in the transition probability matrix cause errors of order ε_F in the mean cluster eigenvalue and the corresponding invariant subspace. This is the same result as in case of reversible Markov chains, compare Section 1.2.1. Moreover, this form of the theorem verifies the notion of $1/\mathrm{sep}(L_1, L_2)$ as *condition number* of the invariant subspace. If we identify the matrix A with the model matrix from Section 3.1.2 and define E as in (3.6), then $\|E\|_\infty \leq \varepsilon_t$ and thus $\varepsilon_F \leq \varepsilon_t\sqrt{N \cdot |\mathcal{C}|}$. Thus, the size of the perturbation depends on the coupling between the clusters.

3.1.3.3 Reversibility and symmetrization

Although the continuous dynamical system is reversible, this property is lost for the matrix due to truncated sampling. That means, the detailed balance condition is just satisfied approximately. Therefore, the transition matrix can have complex eigenvalues. In the following we will show under which conditions the imaginary parts will be small.

Assume that $\widetilde{A}$ is an unsymmetric perturbation of a symmetric matrix $A \in \mathbb{R}^{N\times N}$ with eigenvalues λ_k, i.e. $E = \widetilde{A} - A$ is not symmetric. It can be written as the sum of the symmetric matrix $E_s = (E + E^\top)/2$ and the skew-symmetric part $E_u = (E - E^\top)/2$. The eigenvalues $\mu_k + i\nu_k$ of $\widetilde{A}$ may be complex, but they are all included in the union of regions

$$\mathcal{D}_k = \{\mu + i\nu : \, |\mu + i\nu - \lambda_k| \leq \|E\|_2 \text{ and } |\nu| \leq \|E_u\|_2\}.$$

This result can be found in Wilkinson [116]. The first inequality can be considered as a generalization of Gershgorin's Theorem, whereas the second inequality trims the disks at the top and bottom by horizontal lines at $\pm\|E_u\|_2$. As the perturbation becomes increasingly symmetric, these lines approach one another, restricting the imaginary parts of the eigenvalues of $\widetilde{A}$. Another version of the result was formulated by Kahan [51]. He proved that

$$\sqrt{\sum_{k=1}^{N} \nu_k^2} \leq \|E_u\|_F \quad \text{and} \quad \sqrt{\sum_{k=1}^{N} (\mu_k - \lambda_k)^2} \leq \|E_s\|_F + \sqrt{\|E_u\|_F^2 - \sum_{k=1}^{N} \nu_k^2}.$$

A common symmetrization for stochastic, nearly-reversible matrices [28] is

$$P_{\text{sym}} = \frac{1}{2}(P + D^{-2}P^\top D^2), \tag{3.9}$$

where $D^2 = \text{diag}(\mathbf{w})$ contains the stationary distribution $\mathbf{w}$ of P on its diagonal. This operation preserves the stationary distribution, i.e. $\mathbf{w}^\top P_{\text{sym}} = \mathbf{w}^\top$, and P_{sym} is generalized symmetric w.r.t. the matrix D, $D^2 P_{\text{sym}} = P_{\text{sym}}^\top D^2$. In contrast to P, P_{sym} has only real eigenvalues.

The eigenvalue problem for DPD^{-1} can be considered as an unsymmetric perturbation of the symmetric eigenvalue problem for $DP_{\text{sym}}D^{-1}$ with error matrix

$$E = D(P - P_{\text{sym}})D^{-1} = \frac{1}{2}D(P - D^{-2}P^\top D^2)D^{-1} = \frac{1}{2}(DPD^{-1} - D^{-1}P^\top D).$$

Note that in this case $E_s = 0$ and $E = E_u$. Thus, $\|E\|_F$ directly bounds the imaginary parts of the eigenvalues. The entries of E are given by

$$\frac{1}{2}\frac{1}{\sqrt{w_i\, w_j}}\left(w_i P(i,j) - w_j P(j,i)\right), \quad i,j = 1,\ldots,N.$$

If P is nearly reversible as in (3.1), one obtains $|E(i,j)| \leq \frac{1}{2}\varepsilon_r/\sqrt{w_i w_j}$ and thus $\|E\|_F \leq \frac{1}{2}\varepsilon_r\sqrt{\sum_{i,j} 1/(w_i w_j)}$. This term can be large if there exist many basis functions with small weight. Thus, the symmetrization should only be applied if E is small. Furthermore, a good estimate of the stationary density $\mathbf{w}$ is required for the symmetrization step. Since its computation as eigenvector of the transition matrix P is ill-conditioned, we will not apply the symmetrization in practice until a reliable estimate for $\mathbf{w}$ is available.

3.2 Generalized Markov chains

The term "generalized" in this section refers to a discrete dynamical process described by a matrix pair (P, S) as it arises in a discretization with radial basis functions. The membership vectors χ_i of metastable conformations are the solutions of the "generalized" eigenvalue problem

$$P\chi_i = \lambda_i S\chi_i, \quad , \lambda_i \approx 1, \qquad i = 1, \ldots, n_C.$$

Since P and S are both row-stochastic, $\lambda = 1$ is an eigenvalue of (P, S). Moreover, if both matrices are block-diagonal with n_C blocks and the same block structure, then $\lambda = 1$ is an n_C-fold eigenvalue. In the following, we want to examine how this multiple eigenvalue and the corresponding deflating subspace behave under perturbations.

3.2.1 Perron-Frobenius results for matrix pencils

Several generalizations of the classical Perron-Frobenius theory to the generalized eigenvalue problem $Ax = \lambda Bx$ have been presented in the literature. A first result has been given in [66]. There, the condition $(B^\top y \geq 0 \implies A^\top y \geq 0)$ has been proved to be sufficient for the existence of a positive eigenvalue and a corresponding nonnegative eigenvector. If furthermore either A or B has full column rank, then this eigenvalue is equal to the spectral radius $\rho(A, B)$. If A and B are square matrices and B^{-1} exists, then the condition is that of the classical Perron-Frobenius Theorem, i.e. $Z = B^{-1}A \geq 0$. Moreover, the irreducibility of Z ensures that the eigenvalue is simple and the corresponding eigenvalue unique up to a scalar multiple. However, both nonnegativity conditions are difficult to verify. Another generalization [5] derives the sufficient condition that $(B-A)^{-1}A$ is nonnegative and irreducible for the existence of an eigenvalue in $(0, 1)$ and a corresponding positive eigenvector. However, this condition is very restrictive since $B - A$ is not necessarily invertible. In [69] the relationship between these two different generalizations of the Perron-Frobenius theory has been discussed. It turned out that, under certain conditions, the main assumptions of both approaches are equivalent. A generalization of the Perron-Frobenius Theorem to regular matrix pairs of arbitrary index has been presented in [68]. The authors constructed projector chains in order to derive a new sufficient condition which guarantees that the finite spectral radius of (A, B) is an eigenvalue with a corresponding nonnegative eigenvector. Again, an additional irreducibility assumption ensures uniqueness. This approach can be considered as a generalization of [5], and several examples are presented where the new condition holds wheras the previous conditions are not satisfied. While the last two approches reduce to the classical Perron-Frobenius Theorem when $B = I$, the first approch does not.

In our applications, none of the above conditions can be ensured. Since, by construction, both matrices P and S are stochastic, $\lambda = 1$ is an eigenvalue corresponding to the (positive) constant eigenvector $\mathbf{e} = (1, \ldots, 1)^\top$. However, by the

same argument $S - P$ is singular which rules out the approach in [5]. As long as the parameter α in the construction of the radial basis functions is large enough, the matrix S will be diagonally dominant. Thus $S^{-1}P$ exists, but is seldom nonnegative which precludes the approach in [66]. The condition in [68] might be satisfied in some cases, but not in general. Thus, the main problem is that we cannot assume simplicity of $\lambda = 1$. But even worse, no conclusion about the spectral radius can be drawn.

3.2.2 Model problem

In case of radial basis functions, the resulting matrix pair $(\widetilde{P}, \widetilde{S})$ can often be considered as a perturbation of the model matrix pair (P, I) with identity matrix I and P as described in Section 3.1.2. Such model matrix pair has a subspace basis of membership vectors. As long as the perturbation is small enough and the eigenvalue cluster of (P, I) is well conditioned, the corresponding subspace of the matrix pair $(\widetilde{P}, \widetilde{S})$ will only differ slightly from the invariant subspace of P. This is verified by the following theorem, which is similar to Theorem 2.2.2 but expresses the errors in terms of $\|E\|_F$ and $\|F\|_F$, respectively.

Theorem 3.2.1. [99, 17, 49] *Let (A, B) be in generalized Schur form (2.21) such that (A_{11}, B_{11}) contains the cluster of n_C eigenvalues with left and right deflating subspaces $\mathcal{Y}_1$ and $\mathcal{X}_1$, respectively. Let* dif_l *and* dif_u*, p and q be defined as in (2.24) and (2.28), respectively. Furthermore, let $\widetilde{\mathcal{Y}}_1$ and $\widetilde{\mathcal{X}}_1$ be left and right deflating subspaces of the pertubed matrix pair $(A + E, B + F)$ with $\|(E, F)\|_F = \epsilon_F$. Then*

$$\sin\theta_{\max}(\mathcal{Y}_1, \widetilde{\mathcal{Y}}_1) \le \|Y_1 - \widetilde{Y}_1\|_F \le \frac{\epsilon_F}{\mathrm{dif}[(A_{11}, B_{11}), (A_{22}, B_{22})]} + \mathcal{O}(\epsilon_F^2),$$

$$\sin\theta_{\max}(\mathcal{X}_1, \widetilde{\mathcal{X}}_1) \le \|X_1 - \widetilde{X}_1\|_F \le \frac{\epsilon_F}{\mathrm{dif}[(A_{11}, B_{11}), (A_{22}, B_{22})]} + \mathcal{O}(\epsilon_F^2).$$

Moreover, if

$$\delta \equiv \epsilon_F \frac{4\max(p, q)}{\min(\mathrm{dif}_l, \mathrm{dif}_u)} < 1$$

then

$$\theta_{\max}(\mathcal{Y}_1, \widetilde{\mathcal{Y}}_1) \le \arctan\left(\frac{\delta}{p - \delta(p^2 - 1)^{1/2}}\right),$$

$$\theta_{\max}(\mathcal{X}_1, \widetilde{\mathcal{X}}_1) \le \arctan\left(\frac{\delta}{q - \delta(q^2 - 1)^{1/2}}\right),$$

and

$$|\mu - \widetilde{\mu}| \le \frac{1}{\sigma_{\min}(B_{11})}\left(1 + \frac{\sigma_{\max}(\widetilde{A}_{11})}{\sigma_{\min}(\widetilde{B}_{11})}\right) 3p\epsilon_F.$$

Note that the condition on δ is equivalent to an upper bound on $\|(E, F)\|_F$,

$$\|(E, F)\|_F \le \frac{\min(\mathrm{dif}_l, \mathrm{dif}_u)}{4\max(p, q)}. \tag{3.10}$$

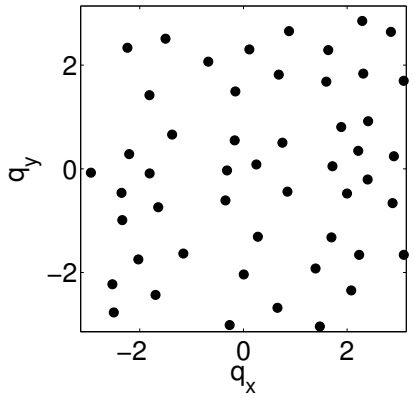

Figure 3.2: (Example 3.2.2) Discretization of the "toy system" with $N = 50$ nodes.

This is the same kind of result as for standard Markov chains. Errors of order ε_F in the transition matrices are carried forward to errors of the same order in the cluster eigenvalue and the deflating subspace. Moreover, the bounds in the theorem imply that $\mathrm{dif}[(A_{11}, B_{11}), (A_{22}, B_{22})]$ is the reciprocal of the condition number for deflating subspaces of a regular pair (A, B).

However, S can only be considered as a perturbation of the identy matrix if the shape parameter α in the construction of the radial basis functions is large enough, which guarantees that S is diagonally dominant. This completely coincides with the observations we made in our numerical computations. While we often observe a clustered eigenvalue in P even for a small value of α, no such cluster can be found in the generalized eigenvalue problem. One can say that the deflating subspace of the matrix pair (P, S) is more sensitive than the invariant subspace for the standard eigenvalue problem. Note that ([98], Thm.4.3)

$$\mathrm{sep}(A_{11}, A_{22}) \geq \mathrm{dif}[(A_{11}, I), (A_{22}, I)].$$

That means, even if S was the identity matrix, the condition number for the generalized eigenvalue problem would be larger than the condition number for the standard eigenvalue problem.

Alternatively, one could consider the standard eigenvalue problem

$$(P + I - S)X = X\Lambda.$$

Under the condition that $P + I - S \geq 0$ the theory for the standard stochastic eigenvalue problem is recovered. However, this will only be the case if S is again close to the identity matrix.

Example 3.2.2. We return to the "toy system" from Section 1.1.6 and generate a discretization. The goal is to place the nodes in regions where the stationary

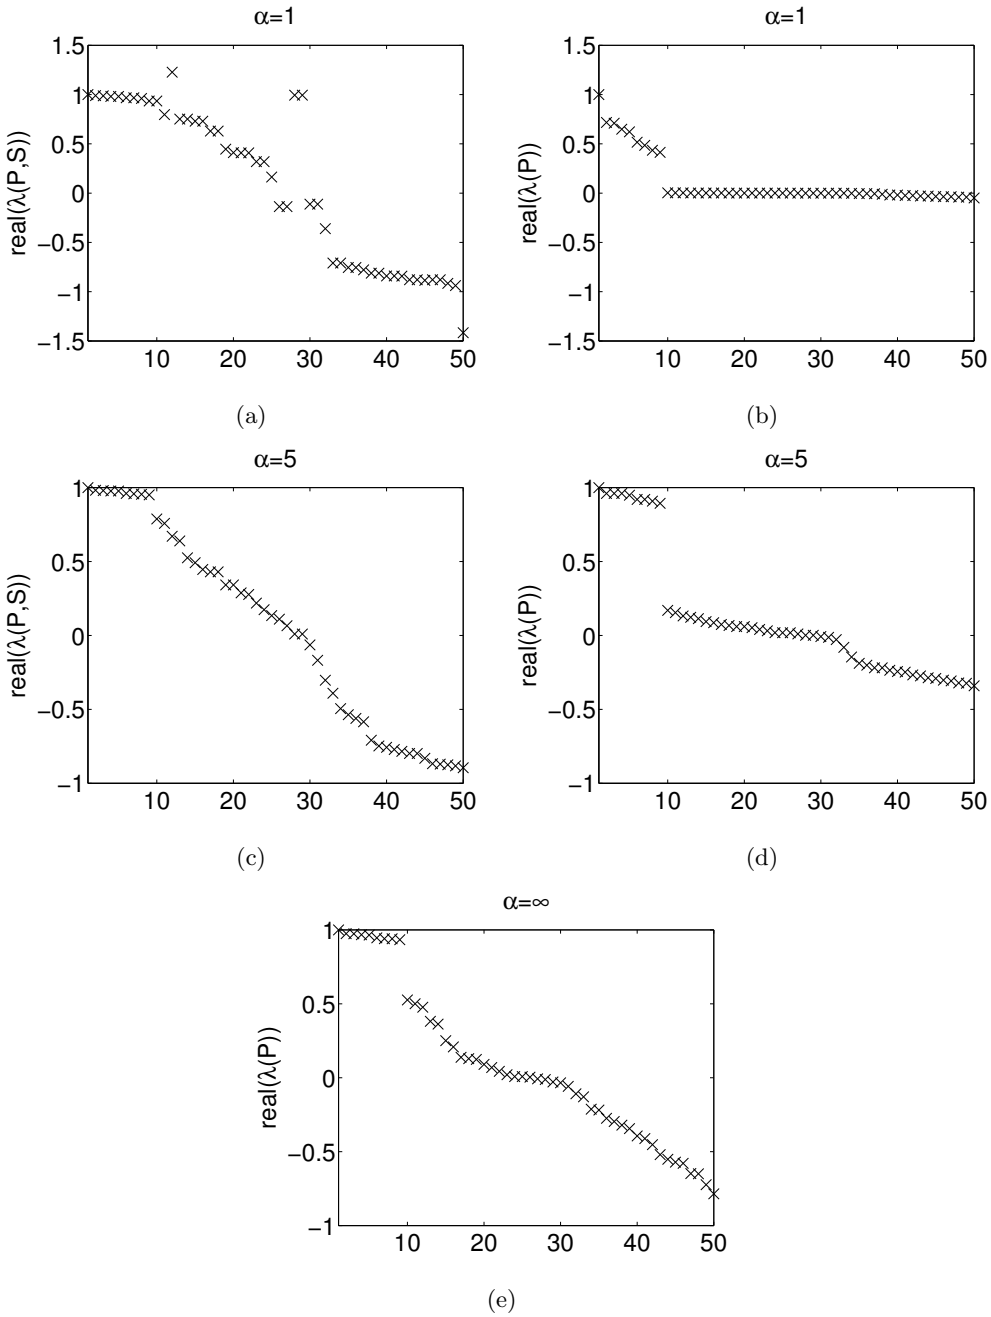

Figure 3.3: (Example 3.2.2) Real part of the spectrum of the matrix pair (P, S) (left hand side) and of P alone (right hand side) for different values of the shape parameter α. The Perron cluster degenerates if α becomes too small.

Table 3.1: (Example 3.2.2) Condition numbers for the eigenspace of the generalized and the standard eigenvalue problem for different values of the shape parameter α.

α	1.0	5.0	∞
$1/\mathrm{dif}_{P,S}$	$1.1e+17$	134.76	4.82
$1/\mathrm{sep}_P$	2.51	1.425	2.71

density is non-zero. For this purpose we first generated an HMC trajectory with 2000 sampling points at inverse temperature $\beta_{\mathrm{pre}} = 0.1$. This allows the trajectory to explore the complete state space. Then we applied k-means clustering with $k = 50$. The identified centers became the $N = 50$ nodes of the discretization, see Figure 3.2. We tested three different discretizations: radial basis functions with shape parameter $\alpha = 1$ and $\alpha = 5$, respectively, and a Voronoi discretization, which can be considered as the limit case as $\alpha \to \infty$. For all discretizations, we computed the transition matrices P and S by Monte-Carlo sampling. In every basis function, 2000 sampling points were generated by HMC (10 MD-steps with $\tau = 0.01$ as proposal step) and then propagated vertically once (50 MD-steps with $\tau = 0.01$). Afterwards, we considered the spectrum of the resulting matrices. The eigenvalues were ordered according to the difference $|1 - \lambda_i|$.

Although the spectrum of P has a cluster of 9 eigenvalues near the Perron root $\lambda = 1$ for all values of α, the corresponding eigenvalue gap disappears for the generalized eigenvalue problem if $\alpha = 1$, see Figure 3.3. For $\alpha = 1$, the eigenvalues λ_9 and λ_{10} appear as conjugate pair such that the gap cannot be recovered. As expected, the condition number of the corresponding invariant subspace of the matrix pair (P, S) becomes huge for small α, see Table 3.1, whereas it is roughly of same size for the standard eigenvalue problem for all values of α. Thus we close this chapter with the following advise.

> If the state space is discretized by radial basis functions with shape parameter α, the value of α must be chosen large enough to assure that the cluster problem is well-conditioned.

Chapter 4

Statistical Error Estimation

In order to apply the perturbation theory presented in the previous chapters, we need an estimate for the errors in the matrices P and S. Since the entries of the matrices are calculated by Monte-Carlo quadrature, it seems appropriate to assume random perturbations. This allows not only for the computation of bounds, but also for statistical analysis.

An appropriate norm for random matrices is the stochastic norm introduced by G. W. Stewart [102]. To obtain interpretable formulas one has to impose restrictions on the perturbation matrices E and F. For this purpose, we introduce the class of row-wise correlated matrices. This is in perfect line with our algorithm because the rows are generated independently by sampling single basis functions. In addition, the resulting formula for the stochastic norm gives insight into the contribution of different matrix rows to the error. This approach will be the basis for the adaptive sampling algorithm in the next chapter.

Since the stochastic norm only depends on first and second moments, it is independent of the form of the distribution. Thus, it would suffice to assume a multivariate normal (MVN) distribution for every matrix row and to estimate mean and covariance matrices. However, the MVN distribution is not the best choice because it allows for negative entries in the matrix elements. Therefore, we assume more appropriate distributions: the *Dirichlet* distribution for probability vectors or the *Pólya* distribution for count vectors. Since they are defined by only a few parameters, they allow for an efficient description of the error distribution. Moreover, as the number of sampling points increases the MVN distribution is recovered, which facilitates the sampling of matrix rows.

In this chapter we will demonstrate how the parameters of the distribution can be estimated from the sampling chains. Numerical examples verify the first order perturbation expansions from Chapter 2 and the MVN approximation for matrix rows.

4.1 Row-wise correlated random matrices

In the following we will analyze the random character of matrix elements and compute the norm and quadratic form of the corresponding error matrix. For brevity, the following considerations always refer to the transition probability matrix P but the results are also valid for the mass matrix S, which is a stochastic matrix as well.

4.1.1 Random matrices

The rows of the stochastic matrix P can be considered as random vectors which are computed from N different independent sampling runs. From the algorithmic point of view, they are approximated by the mean values

$$\widetilde{P}(i,:) = \frac{1}{n_i} \sum_{k=1}^{n_i} \Phi(\overline{q}_k^{(i)}), \tag{4.1}$$

where $\Phi : \Omega \to \mathbb{R}^N$ is the vector-valued function $(\phi_1, \ldots, \phi_N)$, and $\overline{q}_k^{(i)} \in \Omega$ is the kth vertical sampling point from sampling run $i \in \{1, \ldots, N\}$. The convergence towards

$$P(i,:) = \int_\Omega \Phi(q)\pi_i(q)dq$$

is verified by the central limit theorem.

Theorem 4.1.1. ([91], Sec. 1.9.1, Thm. B) *Let $\{\mathbf{X}_i\}$ be i.i.d. random vectors with mean $\boldsymbol{\mu}$ and covariance matrix C. Then*

$$\sqrt{n}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbf{X}_i - \boldsymbol{\mu}\right) \xrightarrow{d} N(\mathbf{0}, C),$$

where $\xrightarrow{d}$ denotes convergence in distribution.

In other words, the sequence of random vectors $S_n = \frac{1}{n}\sum \mathbf{X}_i$ is asymptotically normally distributed with mean $\boldsymbol{\mu}$ and covariance C/n for all n sufficiently large. Since the components of a multivariate normal vector are univariate normal distributions, the components of $\mathbf{X}_i = (x_{i1}, \ldots, x_{iN})$ are asymptotically normal with mean μ_j and variance $\sigma_{jj}^2 = C(j,j)$, $j = 1, \ldots, N$. Besides the existence of the first two moments, the central limit theorem does not use any assumption on the distribution of the random variables $\mathbf{X}_i$.

Based on the algorithmic realization and on the central limit theorem, we can make the assumption that the matrix $\widetilde{P}$ we obtain from the sampling routine equals the sum of a deterministic mean $\overline{P}$ and a random matrix E,

$$\widetilde{P} = \overline{P} + E,$$

where the entries of the error matrix E are distributed with mean zero and row wise covariance matrices $\{C_i\}_{i=1}^N$.

Definition 4.1.2. A matrix $E \in \mathbb{R}^{N\times N}$ is a **row-wise correlated random matrix**, if its entries are random variables with the following properties:

$$\mathbb{E}[E(i,j)] = 0 \quad \text{and} \quad \mathbb{E}[E(i,j)E(k,l)] = \delta_{ik}C_i(j,l), \qquad i,j,k,l = 1,\ldots,N. \tag{4.2}$$

The operator $\mathbb{E}$ is the expectation operator.

4.1.2 Stochastic norms

An appropriate norm for such random matrices is the stochastic norm introduced by G. W. Stewart [102],

$$\|E\|_s^2 \equiv \mathbb{E}(\|E\|_F^2).$$

First note that for random matrices with zero mean, the norm $\|\cdot\|_s$ simplifies a lot:

$$\begin{aligned}
\|E\|_s^2 &= \mathbb{E}\left[\sum_i \sum_j E(i,j)^2\right] \\
&= \sum_i \sum_j \mathbb{E}[E(i,j)^2] \\
&= \sum_i \sum_j \left(\mathrm{var}[E(i,j)] + \mathbb{E}[E(i,j)]^2\right) \\
&= \sum_i \sum_j \mathrm{var}[E(i,j)].
\end{aligned} \tag{4.3}$$

Similar to Thm. 2.3. in [102], we obtain the following expression for the quadratic form w.r.t. an arbitrary matrix $B \in \mathbb{R}^{N\times N}$,

$$\begin{aligned}
\mathbb{E}[(E^\top BE)(i,j)] &= \mathbb{E}\left[\sum_k \sum_l E(k,i)E(l,j)B(k,l)\right] \\
&= \sum_k \sum_l B(k,l)\mathbb{E}[E(k,i)E(l,j)] \\
&= \sum_k B(k,k)C_k(i,j).
\end{aligned}$$

The covariance matrices $\{C_i\}_{i=1}^N$ are symmetric positive semidefinite. Hence, there exist factorizations $C_i = \Sigma_i\Sigma_i^\top$. Σ_i can be computed from an eigenvalue decomposition of C_i. It is not necessarily triangular or square. The number of columns ν_i equals the number of positive eigenvalues of C_i. If C_i was positive definite, Σ_i would be the square, upper triangular Cholesky factor.

Remark 4.1.3. Since the Frobenius norm scales with the size of the matrix, one of the referees suggested to replace $\|A\|_F$ by $\|A\|_{F^\star} \equiv \frac{1}{\sqrt{p}}\|A\|_F$, where $p = \min\{m,n\}$ for an arbitrary matrix $A \in \mathbb{R}^{m\times n}$. Similarly, we would define $\|E\|_s^2 \equiv \mathbb{E}(\|E\|_{F^\star}^2)$. The theoretical results remain unaffected, but we expect some of the perturbation bounds to become less pessimistic.

4.2 Probability distributions

Given a set $D_i = \{\mathbf{p}_{i1}, \ldots, \mathbf{p}_{iK}\} \subset \mathbb{R}^N$ of possible transition probability vectors for row $P(i,:)$ resulting from different samplings, we could simply compute the maximum-likelihood estimate for the mean,

$$\hat{\boldsymbol{\mu}}_i = \frac{1}{K} \sum_{k=1}^{K} \mathbf{p}_{ik} \equiv \overline{\mathbf{p}}_i.$$

The estimate of the mean could be used, for example, to calculate an estimate of the mean cluster eigenvalue or an estimate of the Perron subspace. However, the maximum-likelihood estimates give no indication of the uncertainties in the transition probabilities. To quantify the size of the error matrix E, we additionally have to compute an estimate for the covariance matrix. To choose between the many different possible estimators, we use the fact that the distribution of probability vectors $\mathbf{p}_i$ follows the *Dirichlet distribution* $\mathrm{Dir}(\cdot)$. Similarly, the distribution of count vectors $\mathbf{z}_i$ follows the *Pólya distribution* $\mathrm{Poly}(\cdot)$, which is a compound of Dirichlet distribution and *multinomial distribution* $\mathrm{Mult}(\cdot)$. The properties of these distributions will be explained in the following.

4.2.1 The Dirichlet distribution

The Dirichlet distribution, $\mathrm{Dir}(\boldsymbol{\alpha})$, is a family of continuous multivariate probability distributions parameterized by a vector $\boldsymbol{\alpha}$ of positive real numbers. Let $\mathbf{p} = (p_1, \ldots, p_N)$ denote a random vector where

$$\sum_{i=1}^{N} p_i = 1, \qquad p_i > 0 \quad \forall\, i \in (1, \ldots, N).$$

Table 4.1: Moments of different distributions. The parameters n, w, α_i and p_i are explained in the text.

	$\mathbb{E}[X_i]$	$\mathrm{var}[X_i]$	$\mathrm{cov}[X_i, X_j]$
$\mathrm{Dir}(\boldsymbol{\alpha})$	α_i/w	$\dfrac{\alpha_i(w-\alpha_i)}{w^2(w+1)}$	$\dfrac{-\alpha_i\alpha_j}{w^2(w+1)}$
$\mathrm{Mult}(n, \mathbf{p})$	np_i	$np_i(1-p_i)$	$-np_ip_j$
$\mathrm{Poly}(n, \boldsymbol{\alpha})$	$n\alpha_i/w$	$(n^2+nw)\dfrac{\alpha_i(w-\alpha_i)}{w^2(w+1)}$	$(n^2+nw)\dfrac{-\alpha_i\alpha_j}{w^2(w+1)}$

Under the Dirichlet model of order $N \geq 2$ with parameters $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_N) > 0$, the probability density function at $\mathbf{p}$ is given by

$$f(\mathbf{p}; \boldsymbol{\alpha}) = \frac{1}{Z(\boldsymbol{\alpha})} \prod_{i=1}^{N} p_i^{\alpha_i - 1}.$$

Expected value, variance, and covariance are given in Table 4.1. The normalizing constant,

$$Z(\boldsymbol{\alpha}) = \frac{\prod_{i=1}^{N} \Gamma(\alpha_i)}{\Gamma(\sum_{i=1}^{N} \alpha_i)},$$

is the multinomial beta function, which is expressed in terms of the gamma function. In fact, the marginals p_i are beta distributed, $p_i \sim \beta(\alpha_i, w - \alpha_i)$, where

$$w \equiv \sum_{i=1}^{N} \alpha_i.$$

The value w can be understood as the *precision* of $\mathbf{p}$ [74]. If w is large, $\mathbf{p}$ is likely to be near the mean $\boldsymbol{\mu} \equiv \mathbb{E}[\mathbf{p}] = \boldsymbol{\alpha}/w$. When w is small, $\mathbf{p}$ is distributed more diffusely.

The Dirichlet distribution has a wide range of applications. In biology, Dirichlet distributions can represent proportions of amino acids when modeling sequences with hidden Markov models [93] or with allelic frequencies [61]; in document classification, they can model topic probabilities [65].

Given a training set of proportions $D = \{\mathbf{p}_1, \ldots, \mathbf{p}_K\}$, where $\mathbf{p}_k = (p_{k1}, \ldots, p_{kN})$, the classical way of estimating Dirichlet distribution parameters is to maximize the log-likelihood given by

$$\log \mathbb{P}(D|\boldsymbol{\alpha}) = K \log \Gamma(w) - K \sum_{i=1}^{N} \log \Gamma(\alpha_i) + \sum_{i=1}^{N} \left((\alpha_i - 1) \sum_{k=1}^{K} \log(p_{ki}) \right).$$

To maximize this function, one classically uses the Newton-Raphson algorithm [87, 78, 79, 80]. Since the Dirichlet distribution belongs to the exponential family, the function is globally concave [87] and the Newton-Raphson algorithm converges to the global optimum, assuming that it does not go outside the admissible region, which restricts the choice of initial parameter estimates [25, 87, 115]. An alternative convergent fixed-point iteration has been derived by Minka [74]. Since mean $\boldsymbol{\mu}$ and precision w are roughly decoupled in the maximum-likelihood objective, one can get simplifications and speedups by optimizing alternately the likelihood for w alone,

$$\mathbb{P}(D|w) \propto \left(\frac{\Gamma(w) \exp(w \sum_{i=1}^{N} \mu_i \log \overline{p}_i)}{\prod_{i=1}^{N} \Gamma(w\, \mu_i)} \right)^K,$$

and for $\boldsymbol{\mu}$ alone,

$$\mathbb{P}(D|\boldsymbol{\mu}) \propto \left(\prod_{i=1}^{N} \frac{\exp(w\, \mu_i \log \overline{p}_i)}{\Gamma(w\, \mu_i)} \right)^K.$$

Here,

$$\log \overline{p}_i = \frac{1}{K} \sum_{k=1}^{K} \log(p_{ki}).$$

The symbol $\propto$ means that terms not involving the variable to be optimized (w or $\boldsymbol{\mu}$) were omitted. This method, which we will refer to as Minka's ML algorithm, has been implemented in the MATLAB toolboxes FASTFIT and LIGHTSPEED [75, 76]. For further details, the reader is referred to [73, 74].

Given the estimated parameters $\hat{\boldsymbol{\alpha}}$ and $\hat{w}$, the estimator for the mean is

$$\hat{\boldsymbol{\mu}} = \frac{\hat{\boldsymbol{\alpha}}}{\hat{w}}, \tag{4.4}$$

and the estimator for the covariance matrix becomes

$$\hat{C} = \frac{1}{\hat{w}^2(\hat{w}+1)}[\hat{w}\mathrm{diag}(\hat{\boldsymbol{\alpha}}) - \hat{\boldsymbol{\alpha}}\hat{\boldsymbol{\alpha}}^\top] = \frac{1}{\hat{w}+1}[\mathrm{diag}(\hat{\boldsymbol{\mu}}) - \hat{\boldsymbol{\mu}}\hat{\boldsymbol{\mu}}^\top]. \tag{4.5}$$

The covariance matrices are positive semi-definite and represent rank-one updates of diagonal matrices. Therefore they can be factorized efficiently [92],

$$C_i = \Sigma_i \Sigma_i^\top.$$

For details about this factorization, the reader is referred to the Appendix C.2. If K is large enough, then, by the central limit theorem, the distribution of a row $\mathbf{p}_i = P(i,:)$ converges to a multivariate normal distribution with mean $\boldsymbol{\mu}_i$ and covariance matrix C_i.

4.2.2 The multinomial distribution

The multinomial distribution, $\mathrm{Mult}(n, \mathbf{p})$, is the probability distribution of the numbers of N outcomes in n independent trials. In other words, each trial results in exactly one of some fixed finite number N of possible outcomes, with probabilities $p_1, \ldots, p_N$, $p_i > 0$ for $i = 1, \ldots, N$, $\sum_{i=1}^{N} p_i = 1$. Let $\mathbf{z} = (z_1, \ldots, z_N)$ denote a random vector where z_i is the number of times outcome i was observed over n trials, i.e.

$$\sum_{i=1}^{N} z_i = n, \quad z_i \in \mathbb{N}.$$

Then $\mathbf{z}$ follows a multinomial distribution with parameters n and $\mathbf{p} = (p_1, \ldots, p_N)$, $\mathbf{z} \sim \mathrm{Mult}(n, \mathbf{p})$. The probability density function is given by

$$f(\mathbf{z}; n, \mathbf{p}) = \frac{\Gamma(n+1)}{\prod_{i=1}^{N} \Gamma(z_i+1)} \prod_{i=1}^{N} p_i^{z_i}.$$

Expected value, variance, and covariance are listed in Table 4.1. The conjugate prior of the multinomial distribution is the Dirichlet distribution, i.e., if $\mathbf{z}|\mathbf{p} \sim \mathrm{Mult}(n, \mathbf{p})$ and if the prior is a Dirichlet distribution with parameters $\boldsymbol{\alpha}$, $\mathbf{p} \sim \mathrm{Dir}(\boldsymbol{\alpha})$, then the posterior is also a Dirichlet distribution, $\mathbf{p}|\mathbf{z} \sim \mathrm{Dir}(\mathbf{z} + \boldsymbol{\alpha})$.

4.2.3 The Pólya distribution

The Pólya distribution, Poly$(n, \boldsymbol{\alpha})$, also denoted as Dirichlet-multinomial distribution, is a compound distribution where a probability vector $\mathbf{p}$ is drawn from a Dirichlet distribution with parameter vector $\boldsymbol{\alpha}$ and then a sample of discrete outcomes $\mathbf{z}$ is drawn from a multinomial distribution with probability vector $\mathbf{p}$. The probability density function of a vector of counts $\mathbf{z}$, $\sum_{i=1}^{N} z_i = n$, given the parameter vector $\boldsymbol{\alpha}$ is

$$\begin{aligned} f(\mathbf{z}; n, \boldsymbol{\alpha}) &= \int_{\mathbf{p}} f(\mathbf{z}; n, \mathbf{p}) f(\mathbf{p}; \boldsymbol{\alpha})\, d\mathbf{p} \\ &= \frac{n!}{\sum_{i=1}^{N} z_i!} \frac{Z(\boldsymbol{\alpha} + \mathbf{z})}{Z(\boldsymbol{\alpha})} \prod_{i=1}^{N} \frac{\Gamma(z_i + \alpha_i)}{\Gamma(\alpha_i)}. \end{aligned}$$

Properties of this distribution have for example been analyzed in [77]. Nowadays, the Pólya distribution is widely used in automated document classification and clustering [65, 27].

One can obtain maximum-likelihood estimates of the parameters from a training set of counts $D = \{\mathbf{z}_1, \ldots, \mathbf{z}_K\}$, where $\mathbf{z}_k = (z_{k1}, \ldots, z_{kN})$. Similar to the Dirichlet distribution, there exists an efficient iterative scheme that alternately optimizes for w and for $\boldsymbol{\mu}$. If n_k is the number of counts in sample $\mathbf{z}_k$ ($n_k = \sum_{i=1}^{N} z_{ki}$), then the likelihood for the precision w alone is given by

$$\mathbb{P}(D|w) \propto \prod_{k=1}^{K} \left(\frac{\Gamma(w)}{\Gamma(n_k + w)} \prod_{i=1}^{N} \frac{\Gamma(z_{ki} + w\,\mu_i)}{\Gamma(w\,\mu_i)} \right),$$

and for $\boldsymbol{\mu}$ alone by

$$\mathbb{P}(D|\boldsymbol{\mu}) \propto \prod_{k=1}^{K} \prod_{i=1}^{N} \frac{\Gamma(z_{ki} + w\,\mu_i)}{\Gamma(w\,\mu_i)}.$$

For further details, the reader is referred to [73, 74].

Under the Pólya model, one can also compute the moments of the normalized count vectors $\mathbf{z}_k/n_k$. They are given by $\mathbb{E}[\mathbf{z}_k/n_k] = \mathbb{E}[\mathbf{z}_k]/n_k$, $\mathrm{var}[\mathbf{z}_k/n_k] = \mathrm{var}[\mathbf{z}_k]/n_k^2$, and $\mathrm{cov}[\mathbf{z}_k/n_k, \mathbf{z}_j/n_j] \overset{n_k=n_j=n}{=} \mathrm{cov}[\mathbf{z}_k, \mathbf{z}_j]/n^2$; compare Table 4.1.

4.3 Parameter estimation

In Section 4.2, we have explained how the parameter vector $\boldsymbol{\alpha}$ of a Dirichlet or Pólya distribution can be estimated from a training set of outcome vectors. Now we want to demonstrate how these training sets are generated in our numerical simulations.

4.3.1 Horizontal sampling

In case of radial basis functions, the horizontal sampling points are used to compute an estimate for S. We are concerned with the question of how we can generate

independent samples for the rows of S from which we can then estimate the Dirichlet or Pólya parameters.

Single chain sampling Given a number of sampling points $\{q_k^{(i)}\}_{k=1}^{n_i}$ distributed according to a partial density $\pi_i(q)$, we can compute $S(i,j)$ by (1.14). The numbers $z_{ij} \equiv n_i S(i,j)$ can be interpreted as transition counts and thus as outcome of a multinomial distribution. If we assume a Dirichlet prior distribution with parameters α_{ij}, the posterior of the transition probabilities is a *Dirichlet* distribution with parameters $u_{ij} = \alpha_{ij} + z_{ij}$. This is possible because the Dirichlet parameters u_{ij} are not required to be integers. The values $n_i S(i,j)$ will sometimes be called *generalized* counts. If $\boldsymbol{\alpha}$ is small enough, the distribution of transition probabilities will become independent on the choice of the prior distribution with increasing number of transition counts [92].

Multiple chain sampling If a basis function inherits metastabilities, it could happen that a single horizontal chain is rapidly mixing in a partial domain but not exploring the complete partial density. This leads to a small estimated error in S although the approximated entries of the corresponding row are inexact. Such a *pseudo convergence* can be excluded by running multiple horizontal chains, say c chains. Each chain gives rise to a different estimate $\mathbf{s}_i^{(k)}$, $k = 1, \ldots, c$, for row i in S. These rows can be considered as a training set of count vectors, $D = \{n_i^{(1)}\mathbf{s}_i^{(1)}, \ldots, n_i^{(c)}\mathbf{s}_i^{(c)}\}$, and can be used to compute an estimate of the parameter vector $\boldsymbol{\alpha}_i$ in the *Pólya* distribution as explained in Section 4.2.3. As the chain lengths increase, the precision w will increase and the different estimates for the row will converge to the exact mean.

Remark 4.3.1. Alternatively, one could directly consider the training set of proportions $D = \{\mathbf{s}_i^{(1)}, \ldots, \mathbf{s}_i^{(c)}\}$ as outcome of a Dirichlet distribution. This approach implicitly assumes that the different chains sample from a unified Dirichlet distribution, which might not be true in case of metastabilities inside a basis function. Moreover, the information on the chain lengths $n_i^{(k)}$ is not modeled by this approach.

4.3.2 Vertical sampling

The vertical sampling points are used to compute an estimate for P. We are concerned with the question of how we can generate independent samples for the rows of P from which we can then estimate the parameters. There are two different kinds of sampling methods: The horizontal and vertical sampling can be performed consecutively, or they can be performed concurrently.

Consecutive sampling We assume that we are given a single horizontal sampling chain $\{q_k^{(i)}\}_{k=1}^{n_i}$ that has converged towards the correct partial density $\pi_i(q)$. (In case of multiple horizontal sampling chains, these trajectories are concatenated.)

Then the vertical sampling can be performed by drawing a fixed number of points randomly with replacement from the horizontal trajectory and propagating them with random initial momenta. Thus, we are in the case of single chain sampling. For Voronoi cells, we obtain a matrix $Z = (z_{ij})$ of transition counts where z_{ij} is the number of transition samples started in cell i and ended in cell j. Assuming a Dirichlet prior distribution with parameters α_{ij}, the posterior of the transition probabilities is a *Dirichlet* distribution with parameters $u_{ij} = \alpha_{ij} + z_{ij}$. For radial basis functions, we do not count transitions but we compute directly the values p_{ij} according to (1.15). However, the numbers $n_i p_{ij}$, where n_i is the number of propagated points from basis function i, can be interpreted as counts z_{ij} such that we can proceed in the same way as for Voronoi cells.

Concurrent sampling Concurrent sampling means that the points generated for a horizontal trajectory are immediately propagated vertically, which corresponds to multiple chain sampling. Thus, every sampling chain k gives rise to two proportion vectors $\mathbf{s}_i^{(k)}$ and $\mathbf{p}_i^{(k)}$, which again can be interpreted as count vectors $n_i\mathbf{s}_i^{(k)}$ and $n_i\mathbf{p}_i^{(k)}$. The two training sets $D_S = \{n_i^{(1)}\mathbf{s}_i^{(1)}, \ldots, n_i^{(c)}\mathbf{s}_i^{(c)}\}$ and $D_P = \{n_i^{(1)}\mathbf{p}_i^{(1)}, \ldots, n_i^{(c)}\mathbf{p}_i^{(c)}\}$ are then used to estimate the *Pólya* parameters $\boldsymbol{\alpha}_{S(i,:)}$ and $\boldsymbol{\alpha}_{P(i,:)}$ independently.

The following statement summarizes the previous explanations.

> Estimation of the parameter vector $\boldsymbol{\alpha}$ from single chain sampling requires an assumption on the Dirichlet prior distribution and treats the outcome as multinomial distributed. Multiple chain sampling does not need an assumption on the prior distribution but gives rise to a training set of outcomes (counts or proportions) from which the parameter vector $\boldsymbol{\alpha}$ is estimated directly.

4.4 Perturbation bounds

Once the parameters $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_N$ of the matrix rows have been estimated, the expected values $\overline{P}$ and $\overline{S}$ as well as the row-wise covariance matrices $\{C_{P,k}, C_{S,k}\}_{k=1}^N$ can be computed. Moreover, we are able to compute expected values of the variables of interest, for example a basis for the Perron subspace, $\overline{X}_1 = X_1|_{\overline{P},\overline{S}}$. With these tools at hand, the error matrices can be quantified.

4.4.1 Stochastic norms of projected error matrices

First consider the standard eigenvalue problem. Given the values for $\overline{X}_1$, $\overline{X}_2$, and $\{C_k\}$, we are able to estimate the stochastic norms of the matrices E_{ij}, $i, j \in \{1, 2\}$

(see Equation (2.7)):

$$\begin{aligned}
\|E_{ij}\|_s^2 = \mathbb{E}[\mathrm{trace}(E_{ij}^H E_{ij})] &= \mathbb{E}[\mathrm{trace}(X_j^H E^\top X_i X_i^H E X_j)] \\
&= \mathrm{trace}(\overline{X}_j^H \mathbb{E}[E^\top X_i X_i^H E]\overline{X}_j) \\
&= \mathrm{trace}(\overline{X}_j^H \sum_{k=1}^N \|\overline{X}_i(k,:)\|_2^2 C_k \overline{X}_j) \\
&= \sum_{k=1}^N \|\overline{X}_i(k,:)\|_2^2 \mathrm{trace}(\overline{X}_j^H C_k \overline{X}_j) \\
&= \sum_{k=1}^N \|\overline{X}_i(k,:)\|_2^2 \|\Sigma_k^\top \overline{X}_j\|_F^2. \qquad (4.6)
\end{aligned}$$

Note that these norms do not pose any assumptions on the distribution of elements $E(i,j)$ besides the fact that the first two moments exist.

For the generalized eigenvalue problem, similar formulas can be derived with a slightly different notation; compare Equations (2.7) and (2.29):

$$\|E_{ij}\|_s^2 = \sum_{k=1}^N \|\overline{Y}_i(k,:)\|_2^2 \|\Sigma_{P,k}^\top \overline{X}_j\|_F^2, \quad \|F_{ij}\|_s^2 = \sum_{k=1}^N \|\overline{Y}_i(k,:)\|_2^2 \|\Sigma_{S,k}^\top \overline{X}_j\|_F^2.$$

4.4.2 Toy example

To illustrate the error bounds in terms of stochastic norms, we analyze the sampling data stemming from the toy example in Section 1.1.6.

In the following, we want to examine error bounds for the mean cluster eigenvalue and for the invariant or deflating subspace, respectively. We are especially interested in how the bounds change, if the discretization changes between Voronoi decomposition and radial basis functions. To keep computational costs low, we restrict ourselves to $N = 16$ basis functions $\{\phi_i(q)\}_{i=1}^N$. However, the nodes must be located appropriately to resolve the metastabilities inherent in the system. This will be demonstrated in the following.

We tested the following three types of discretization:

Grid based Voronoi discretization: We decompose Ω into $N = s^2$ Voronoi sets by a uniform $s \times s$-grid,

$$\{[-\pi;-\pi],[-\pi;-\pi+2\pi/s],\ldots,[\pi-2\pi/s;\pi-2\pi/s]\},$$

where $s = 4$; see Figure 4.1(a).

Meshfree Voronoi discretization: We select $N = 16$ nodes via k-means from a pre-sampling trajectory, see Figure 4.2(a). The pre-sampling trajectory with 2000 sampling points was generated via hybrid Monte-Carlo at inverse temperature $\beta_{\mathrm{pre}} = 0.1$.

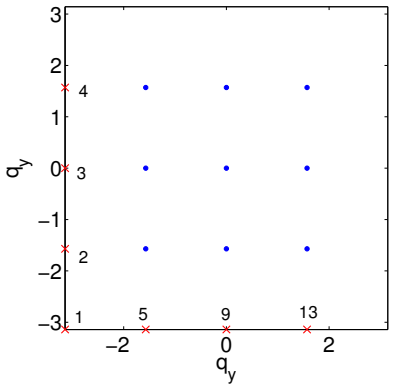

(a) Grid-based Voronoi discretization.

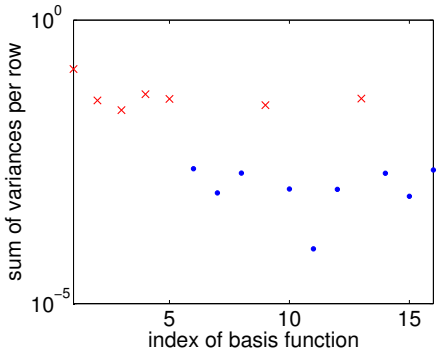

(b) Sum of variances in each row of P.

Figure 4.1: The grid-based Voronoi discretization comprises nodes in transition regions (marked by a cross) between metastable domains which results in large variances in the corresponding matrix rows.

Radial basis functions: We use the same nodes as in case of a meshfree Voronoi discretization, but this time we discretize by means of radial basis functions with $\alpha = 5$.

For each basis function, 2000 sampling points were generated via hybrid Monte-Carlo (generation of proposal point via 10 MD-steps with $\tau = 0.01$), distributed according to the partial densities $\pi_i(q)$ at inverse temperature $\beta = 0.4$. These points were then propagated by the Hamiltonian dynamics with randomized initial momenta drawn from the distribution $\eta(p)$ (50 MD-steps with $\tau = 0.01$). The rows of the transition matrices P and S were computed according to (1.15) and (1.14), respectively.

For every type of discretization, we repeated the experiment $c = 100$ times, resulting in training sets of 100 (generalized) count vectors for every matrix row, from which the parameter vectors $\boldsymbol{\alpha}_i$ were estimated. The matrices $\overline{P}$ with $\overline{P}(i,:) = \hat{\boldsymbol{\alpha}}_{P,i}/\hat{w}_{P,i}$ and $\overline{S}$ with $\overline{S}(i,:) = \hat{\boldsymbol{\alpha}}_{S,i}/\hat{w}_{S,i}$ were used as estimates for the exact matrices P and S. Then, the Perron cluster eigenvalues and the corresponding invariant or deflating subspace were calculated, together with the variables that are involved in the perturbation theory.

Numerical results The values for the characteristic variables and bounds for all problems are listed in Table 4.2. The notation is inherited from Chapter 2. The bar on some quantities indicates that these values were computed at the expected values $\overline{P}$ and $\overline{S}$. The row "global bound" contains the bound from (3.8) or (3.10), respectively. The bounds on $\|E\|_\infty$ and $\|F\|_\infty$ result from the former error estimator (1.16). The bound on $\sin(\theta)$ corresponds to (2.11), and the bound on Δ_μ was

Table 4.2: Values of variables and bounds involved in the perturbation analysis of eigenvalue problems stemming from different discretizations. The notation is adopted from Theorems 2.1.2, 2.1.5, 2.2.2, and 2.2.3.

	Voronoi cells		radial basis functions	
	grid based	meshfree	gen. eig.	stand. eig.
$\overline{\lambda}_9$	0.7223	0.9074	0.9197	0.9232
$\overline{\lambda}_{10}$	0.0106	0.0531	0.0213	0.5364
$\mathrm{sep}(\overline{L}_1, \overline{L}_2)$	0.7009	0.8395	–	0.3750
$\mathrm{dif}[(\overline{A}_{11}, \overline{B}_{11}), (\overline{A}_{22}, \overline{B}_{22})]$	–	–	0.2952	–
$(\|E\|_s, \|F\|_s)$	(0.6143,–)	(0.1543,–)	(0.1447.0.1545)	0.1581*
global bound	0.1418	0.1786	0.0647	0.0825
$(\|E\|_\infty, \|F\|_\infty)$	(1.9510,–)	(1.257,–)	(1.2476,0.6074)	0.7543*
$(\|E_{11}\|_s, \|F_{11}\|_s)$	(0.2418,–)	(0.0588,–)	(0.0583,0.0399)	0.0449*
$(\|E_{12}\|_s, \|F_{12}\|_s)$	(0.1845,–)	(0.0721,–)	(0.0642,0.0826)	0.0856*
$(\|E_{21}\|_s, \|F_{21}\|_s)$	(0.4274,–)	(0.0731,–)	(0.0756,0.0529)	0.0533*
$(\|E_{22}\|_s, \|F_{22}\|_s)$	(0.3197,–)	(0.0990,–)	(0.0877,0.1126)	0.1132*
δ	0.1394	0.6817	0.1427	0.2169
η	0.9608	0.6974	1.0424	0.3960
$\gamma\eta/\delta^2$	21.0	0.1098	3.8685	0.4485
bound on $\sin(\theta)$	–	0.2146	–	–
$\gamma_\star\eta_\star/{\delta_\star}^2$	0.3718	0.0089	0.1397	0.0739
bound on Δ_μ	–	0.0491	0.1751	0.0692

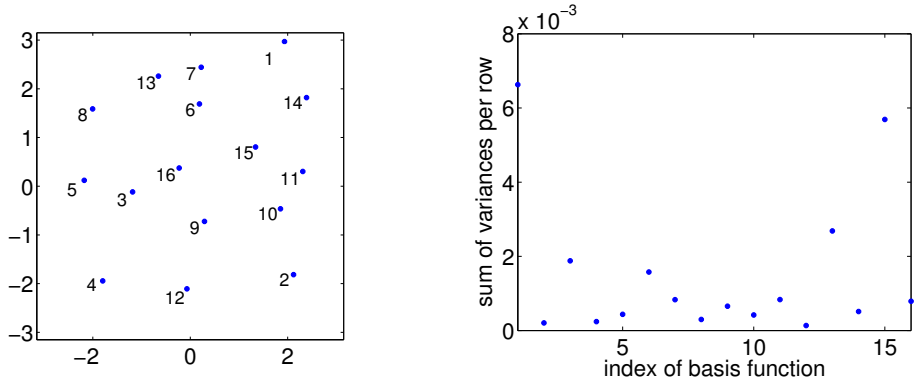

(a) Meshfree Voronoi discretization.

(b) Sum of variances in each row of P.

Figure 4.2: The meshfree Voronoi discretization suits better to the metastable subsets. The basis functions $\{\phi_i\}_{i=1}^N$ resolve the boundaries between metastable regions such that the row-wise variances are small.

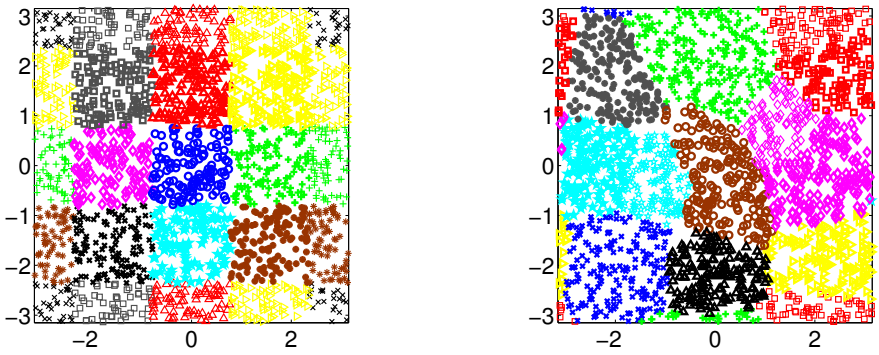

(a) Grid-based discretization. (b) Meshfree discretization.

Figure 4.3: Final clustering for grid-based and meshfree Voronoi discretization, respectively. The plots were generated by assigning 2000 uniformly distributed points to the cluster they belong to with maximal probability. If this probability is smaller than 0.8, the marker is drawn with a small linewidth, otherwise it is drawn in bold face.

computed according to (2.18) or (2.38), respectively. The $*$ in the last column signalizes that these are not the values for E or F, but for $G = E - F$ in the standard eigenvalue problem $(P + I - S)X = X\Lambda$. The $\star$ in the first column indicates values from the spectral resolution.

For the grid-based Voronoi discretization, the stochastic norm of the error matrix exceeds the global bound. Note that the former error estimator for $\|E\|_\infty$ is even larger. Unfortunately, $\gamma\eta/\delta^2 > 0.25$. Any bound on $\sin(\theta)$ thus becomes meaningless. The reason is that the row-wise variances are very large for basis functions whose nodes are located near a boundary between different clusters, see Figure 4.1. The assignment of such states to clusters is very sensitive, see Figure 4.3(a).

For the meshfree Voronoi discretization, the eigenvalue problem is better conditioned. The eigenvalue gap as well as the separation of the subspace are larger than in case of a grid-based discretization. Consequently, we have $\gamma\eta/\delta^2 < 0.25$, which implies that the global bounds are valid. Indeed, the basis functions resolve the boundaries between the metastable regions much better than in case of the grid based discretization, see Figure 4.3(b). Moreover, the mean cluster eigenvalue is approximated very well.

With radial basis functions, the eigenvalue gap is quite large and the estimated matrix errors are of similar size as for the meshfree Voronoi discretization, compare Table 4.2. However, the inverse condition number dif and thus the required upper bound (3.10) on $\max(\|E\|_s, \|F\|_s)$ is quite small. For the same reason, $\gamma\eta/\delta^2$ is larger than 0.25 such that we cannot apply the perturbation theorem 2.2.2. In case of the spectral resolution, the condition $\gamma_\star\eta_\star/{\delta_\star}^2 < 0.25$ is satisfied, but the bound on Δ_μ estimated via (2.38) yields the value 0.1751, which is quite large.

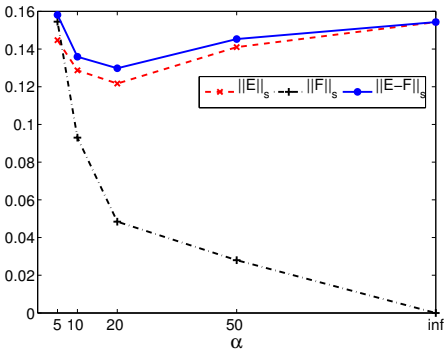

(a) Stochastic norm of the error matrices.

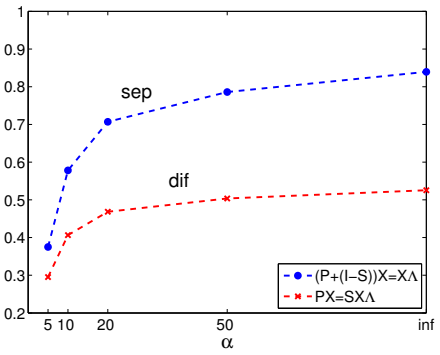

(b) Inverse condition numbers sep and dif for the standard and the generalized eigenproblem, respectively.

Figure 4.4: Error norms and condition numbers in dependence on α, the shape parameter of radial basis function. There seems to be an optimal value for α around 20 where the norms of the error matrices attain a minimum. Small values of α should be avoided because they imply large condition numbers.

Alternatively, we considered the standard eigenvalue problem $(\overline{P} + I - \overline{S})\overline{X}_1 = \overline{X}_1\overline{\Lambda}$ with error matrix $G \equiv E - F$. The eigenvalue gap in the corresponding spectrum is much smaller than for the generalized eigenvalue problem. However, the inverse condition number sep is larger. Nevertheless, $\|G\|_s = 0.1581$ is still larger than the global bound 0.0825, and $\gamma\eta/\delta^2 > 0.25$. Although the preconditions of the perturbation theorem are still not satisfied, the values are much closer to the required bounds than in case of the generalized eigenvalue problem. Moreover, the spectral resolution indicates that the mean cluster eigenvalue is approximated well $(\Delta_\mu < 0.0692)$.

In all cases, an enhanced sampling might reduce the error such that the perturbation theorems can be applied.

Dependence on the shape parameter In the following, we want to examine how our observations depend on the shape parameter α of the radial basis functions. The condition number of the invariant or deflating subspace, measured by $1/\text{sep}$ for the standard eigenvalue problem with $P + (I - S)$ and $1/\text{dif}$ for the generalized eigenvalue problem (P, S) decreases with increasing α, see Figure 4.4(b). In other words, the closer S to the identity, the better conditioned the eigenvalue problem. On the other hand, the stochastic norms of the error matrices E and $E - F$ seem to have a local minimum for $\alpha \in [10, 50]$, see Figure 4.4(a).

The reason is that there exists a value of α for which the basis functions approximate the conformations best. In our example, the conformations have nearly no overlap but large gradients. Thus α must be relatively large for a good approxi-

mation. This is also verified by the fact that for $\alpha = 1$ no clusters can be detected, i.e. there is no integer number n_C which leads to a well-conditioned invariant or deflating subspace. Since the results for $\alpha > 10$ do not deviate too much, one could as well decide for $\alpha = \infty$ (Voronoi tessellation) in our example.

A detailed statistical analysis to find the optimal value for α is out of range in practice because it would require a complete horizontal and vertical sampling for every α. However, the results are quite insensitive w.r.t. α, if α is near the optimum [109]. If no well-conditioned invariant or deflating subspace can be identified, it is recommended to repeat the simulation with a different value for α.

In summary, we may state that the perturbation bounds in case of radial basis functions are larger than in case of a Voronoi discretization because the condition number for the deflating subspace is larger than the condition number for the invariant subspace. Moreover, radial basis functions introduce an additional error in the mass matrix S, which is zero otherwise. Nevertheless, since perturbation bounds in terms of norms are often quite pessimistic, the analysis of distributions might give much better results. This will be examined in the following section.

4.5 Error distribution

Often one is not only interested in upper bounds on errors, but also in their distributions. Given a training set of data for each row of P and S, one can estimate the parameters $\boldsymbol{\alpha}_i$ and use them to sample more transition matrices. From these matrices, one can compute the variable of interest and analyze its distribution, which will be demonstrated in the following.

4.5.1 Direct method

The direct method starts with sampling the rows of E and F from the Dirichlet or the MVN distribution, see App. C.2. For each sampled pair of matrices, the variable of interest is computed. This method is very expensive because it requires solving the eigenvalue problem in every iteration.

4.5.2 1st order approximation

Whenever a linear relationship between the perturbation (E, F) and the variable of interest is available, samples of the variable can easily be obtained by applying the linear operator to the sample of (E, F). If all other variables are replaced by their expected values calculated from $(\overline{P}, \overline{S})$, then the linear relationship can be considered as the first order Taylor series expansion for the variable of interest about the values calculated at the expected values of the transition matrices.

Invariant subspace From Section 2.1.2 we know that, under certain conditions, there exists an invariant subspace $\widetilde{X}_1$ of the perturbed matrix $\widetilde{P} = P + E$ in the

form $\widetilde{X}_1 = X_1 + X_2 Q$, where a 1st order approximation of Q is given by

$$\operatorname{vecr}(Q) \doteq T \operatorname{vecr}(E) \tag{4.7}$$

with T defined as in (2.14). Remember that "vecr" stacks the rows of a matrix into one column (see Section 2.1.1). The matrix T needs to be computed only once. Then one can directly sample the rows of E from the Dirichlet or MVN distribution and compute Q by one matrix-vector multiplication. This is much faster than the direct method. The distribution of the largest singular value of Q equals the distribution of $\tan(\theta_{\max}(\overline{\mathcal{X}}_1, \widetilde{\mathcal{X}}_1))$.

Mean cluster eigenvalue of P According to Section 2.1.3, the deviation in the average eigenvalue of L_1 can be measured by the first-order approximation

$$\Delta_\mu \doteq \mathbf{t}^\top \operatorname{vecr}(E) \tag{4.8}$$

with $\mathbf{t}$ defined as in (2.19). Thus, a sample of Δ_μ can be calculated by one vector-vector multiplication with vectors of length N^2.

Stationary distribution Following the approach by Singhal and Pande [44], the first order Taylor series expansion for the stationary distribution $\mathbf{w}$ as a function of the transition probabilities p_{ij} is given by

$$\mathbf{w} \doteq \overline{\mathbf{w}} + \sum_{i=1}^{N} G_i^{\mathbf{w}} \Delta \mathbf{p}_i, \tag{4.9}$$

where $\Delta \mathbf{p}_i = E(i,:)^\top$ and

$$G_i^{\mathbf{w}} = \begin{bmatrix} \frac{\partial w_1}{\partial p_{i1}}\big|_{\overline{A}} & \cdots & \frac{\partial w_1}{\partial p_{iN}}\big|_{\overline{A}} \\ \vdots & \ddots & \\ \frac{\partial w_N}{\partial p_{i1}}\big|_{\overline{A}} & & \frac{\partial w_N}{\partial p_{iN}}\big|_{\overline{A}} \end{bmatrix}$$

with $\overline{A} = \overline{P}^\top - I$. The entries of the sensitivity matrix $G_i^{\mathbf{w}}$ are obtained from the solution of a system of linear equations (compare Section B.2),

$$\begin{bmatrix} \overline{A} \\ \mathbf{e}^\top \end{bmatrix} G_i^{\mathbf{w}} = -\overline{\mathbf{w}}_i \begin{bmatrix} I_N \\ 0 \end{bmatrix}.$$

Once the sensitivity matrices have been computed, every sample of $\mathbf{w}$ requires N matrix-vector multiplications with matrices of size $N \times N$.

Deflating subspaces Instead of solving the nonlinear Sylvester equation (2.30), it mostly suffices to solve the linearized system

$$\begin{pmatrix} \mathrm{vecr}(W) \\ \mathrm{vecr}(Q) \end{pmatrix} \doteq Z \begin{pmatrix} \mathrm{vecr}(E) \\ \mathrm{vecr}(F) \end{pmatrix} \tag{4.10}$$

with Z defined in (2.33). The matrix $Z \in \mathbb{R}^{2n_C(N-n_C)\times N^2}$ needs to be computed only once. Then samples of Q can be obtained by multiplying the lower $n_C(N-n_C)$ rows of Z with the vector $(\mathrm{vecr}(E), \mathrm{vecr}(F))^\top \in \mathbb{R}^{N^2}$.

Mean cluster eigenvalue of (P, S) There are two possibilities to sample from the distribution of Δ_μ. The first method does not rely on the linear relationship between $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$ and (E, F). According to (2.36), we can use the linearized and vectorized equations

$$\begin{aligned} \mathrm{vecr}(\widetilde{A}_{11}) &\doteq \mathrm{vecr}(A_{11}) + (V_1^H \otimes X_1^\top)\mathrm{vecr}(E), \\ \mathrm{vecr}(\widetilde{B}_{11}) &\doteq \mathrm{vecr}(B_{11}) + (V_1^H \otimes X_1^\top)\mathrm{vecr}(F) \end{aligned}$$

to sample the matrices $\widetilde{A}_{11}$ and $\widetilde{B}_{11}$ efficiently. Then we have to solve the reduced eigenvalue problem to obtain

$$\mu(\widetilde{A}_{11}, \widetilde{B}_{11}) \doteq \mathrm{mean}(\mathcal{L}[(\widetilde{A}_{11}, \widetilde{B}_{11})]).$$

This is in contrast to the standard eigenvalue problem, where $\mu(\widetilde{L}_1)$ was sampled directly without solving an eigenvalue problem. However, the computational effort to solve the system of size $n_C \times n_C$ is small compared to the direct approach which additionally requires the computation of the Schur decomposition in each step.

The solution of the small eigenvalue problem can be avoided by using the linear relationship (2.41),

$$\Delta_\mu \doteq \mathbf{z}_1^\top \mathrm{vecr}(E) - \mathbf{z}_2^\top \mathrm{vecr}(F). \tag{4.11}$$

This requires the multiplication of vectors $\mathrm{vecr}(E)$ and $\mathrm{vecr}(F)$ with vectors $\mathbf{z}_1$ and $\mathbf{z}_2$ of length N^2 in every sampling step, which can be done efficiently. These vectors need to be computed only once at the beginning of the simulation. For this purpose, n_C rows of the Kronecker-product matrices $(B_{11}^{-1}V_1^H)\otimes X_1^\top$ and $(B_{11}^{-1}A_{11}B_{11}^{-1}V_1^H)\otimes X_1^\top$ must be calculated (see Equation (2.40)), which is much cheaper than forming the complete Kronecker products.

4.5.3 Multivariate normal approximation

If the sampling size is large enough, the rows of E and F will be distributed according to the multivariate normal distribution with means $\mathbf{m}_i$ and covariance matrices C_i as given by (4.4) and (4.5). Once the covariance matrices have been factorized, sampling from the MVN distribution is much easier than from the Dirichlet distribution; see Section C.2. If additionally a linear relationship is available, simple closed form solutions for the distributions can be derived.

It is a well known fact from the theory of normal distributions that if one has a vector $\mathbf{y} \sim \mathrm{MVN}(\boldsymbol{\mu}, C)$, then the vector $\mathbf{y}' = F\mathbf{y} + \mathbf{b}$ is distributed as $\mathbf{y} \sim \mathrm{MVN}(F\boldsymbol{\mu} + \mathbf{b}, FCF^\top)$.

We know about E and F that their rows are pairwise uncorrelated, which can be expressed as

$$\mathrm{vecr}(E) \sim \mathrm{MVN}(0, \mathrm{diag}(C_{P,i})), \quad \mathrm{vecr}(F) \sim \mathrm{MVN}(0, \mathrm{diag}(C_{S,i})).$$

Given the factorizations of the covariance matrices $C_{P,i}$ and $C_{S,i}$,

$$C_{P,i} = \Sigma_{P,i}\Sigma_{P,i}^\top, \quad C_{S,i} = \Sigma_{S,i}\Sigma_{S,i}^\top,$$

samples of the variables of interest can just be generated by sampling vectors from the standard normal distribution and performing vector-vector or matrix-vector multiplications. The size of the matrices depends on the rank of the covariance matrices, which is related to the number of non-zero entries in P and S. We will denote by ν_P and ν_S the sum of the number of columns in the matrices $\Sigma_{P,i}$ and $\Sigma_{S,i}$, respectively.

Invariant subspace From samples of

$$\mathbf{y} \sim \mathrm{MVN}(0, I_{\nu_P})$$

we obtain samples for Q from (4.7) by setting

$$\mathrm{vecr}(Q) = T\mathrm{diag}(\Sigma_{P,i})\mathbf{y} \sim \mathrm{MVN}(0, T\mathrm{diag}(C_{P,i})T^\top).$$

Thus, every sampling step requires just one matrix-vector multiplication with a matrix $T\mathrm{diag}(\Sigma_{P,i})$ of size $n_C(N - n_C) \times \nu_P$.

Mean cluster eigenvalue of P Given samples $\mathbf{y}$ from the standard normal distribution, samples of Δ_μ can be obtained from (4.8) via

$$\Delta_\mu = \mathbf{t}^\top \mathrm{diag}(\Sigma_{P,i})\mathbf{y} \sim \mathrm{N}(0, \mathbf{t}^\top \mathrm{diag}(C_{P,i})\mathbf{t}).$$

Thus, Δ_μ follows a univariate normal distribution with variance $\sigma^2 = \mathbf{t}^\top \mathrm{diag}(C_{P,i})\mathbf{t}$ and can be sampled efficiently.

Stationary distribution Under the assumption that $\Delta\mathbf{p}_i = E(i,:)^\top \sim \mathrm{MVN}(0, C_{P,i})$, the linear relationship (4.9) yields

$$\mathbf{w} \sim \mathrm{MVN}(\overline{\mathbf{w}}, \sum_{i=1}^{N} G_i^{\mathbf{w}} C_{P,i} (G_i^{\mathbf{w}})^\top). \tag{4.12}$$

In order to sample from this distribution, we sample N standard MVN variables $\mathbf{y}_i \sim \mathrm{MVN}(0, I_{\nu_{P,i}})$, from which we get

$$\Delta\mathbf{w}_i = G_i^{\mathbf{w}} \Sigma_{P,i} \mathbf{y}_i \sim \mathrm{MVN}(0, G_i^{\mathbf{w}} C_{P,i} (G_i^{\mathbf{w}})^\top), \tag{4.13}$$

and finally

$$\mathbf{w} = \overline{\mathbf{w}} + \sum_{i=1}^{N} \Delta \mathbf{w}_i.$$

Deflating subspace According to (4.10), samples of Q can be obtained via

$$\operatorname{vecr}(Q) = Z(n_C(N-n_C) : 2n_C(N-n_C), :)\operatorname{diag}(\Sigma_{P,i}, \Sigma_{S,i})\mathbf{y},$$

where $\mathbf{y} \sim \mathrm{N}(0, I_{\nu_P+\nu_S})$.

Mean cluster eigenvalue of (P, S) Similar to the first order approximation, the distribution of $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$ can be obtained by solving a small generalized eigenvalue problem in every sampling step. This time, however, E and F are sampled directly from the MVN distribution such that

$$\begin{aligned}\operatorname{vecr}(\widetilde{A}_{11}) &\sim \mathrm{MVN}(\operatorname{vecr}(A_{11}), H\operatorname{diag}(C_{P,i})H^\top),\\ \operatorname{vecr}(\widetilde{B}_{11}^\top) &\sim \mathrm{MVN}(\operatorname{vecr}(B_{11}^\top), H\operatorname{diag}(C_{S,i})H^\top),\end{aligned}$$

where

$$H \equiv V_1^H \otimes X_1^\top.$$

From samples $\mathbf{y}_P \sim \mathrm{MVN}(0, I_{\nu_P})$ and $\mathbf{y}_S \sim \mathrm{MVN}(0, I_{\nu_S})$ we obtain samples of $\widetilde{A}_{11}$ and $\widetilde{B}_{11}$ by setting

$$\operatorname{vecr}(\widetilde{A}_{11}) = \operatorname{vecr}(A_{11}) + H\operatorname{diag}(\Sigma_{P,i})\mathbf{y}_P, \quad \operatorname{vecr}(\widetilde{B}_{11}) = \operatorname{vecr}(B_{11}) + H\operatorname{diag}(\Sigma_{S,i})\mathbf{y}_S.$$

By using the linear relationship (4.11), samples of Δ_μ could also be obtained directly by

$$\Delta_\mu = \mathbf{z}_1^\top \operatorname{diag}(\Sigma_{P,i})\mathbf{y}_P - \mathbf{z}_2^\top \operatorname{diag}(\Sigma_{S,i})\mathbf{y}_S.$$

4.5.4 Computational complexity

Table 4.3 summarizes the approximate running times for the different ways of computing error distributions. As before, N denotes the number of basis functions or the matrix size, respectively, and n_C the cluster size. The number of samples is denoted by L. Q is the time required to sample a single standard normal random variable. The constants ν_P and ν_S correspond to the overall numbers of independent normal random variables required to factorize the row covariance matrices of P and S, respectively.

The direct method requires the solution of the eigenvalue problem in all L sampling steps, which takes time $\mathcal{O}(Ln_CN^2)$ for the standard eigenvalue problem and time $\mathcal{O}(N^3)$ for the generalized eigenvalue problem. Details can be found in Section C.1. For a single eigenvector as the stationary distribution, the time reduces to $\mathcal{O}(N^2)$. Additionally, there are computational costs for the sampling of the error matrices E and F which amount to $\mathcal{O}(LN^2Q)$ for the Dirichlet distribution and

Table 4.3: Approximate running times for calculating the error distribution of different quantities from sampling either by the direct method or by 1st order approximations. The columns differ by the distribution from which the matrices are sampled. N denotes the number of basis functions or the matrix size, respectively, and n_C the cluster size. The number of samples is denoted by L. Q is the time required to sample a single standard normal random variable. The constants ν_P and ν_S correspond to the overall numbers of independent normal random variables required to factorize the row covariance matrices of P and S, respectively.

Direct method		
	Dirichlet	MVN
inv. subspace	$\mathcal{O}(Ln_CN^2 + LN^2Q)$	$\mathcal{O}(N^4 + Ln_CN^2 + L\nu_P(Q+N))$
$\mu(L_1)$	$\mathcal{O}(Ln_CN^2 + LN^2Q)$	$\mathcal{O}(N^4 + Ln_CN^2 + L\nu_P(Q+N))$
$\mathbf{w}$	$\mathcal{O}(LN^2 + LN^2Q)$	$\mathcal{O}(N^4 + LN^2 + L\nu_P(Q+N))$
defl. subspace	$\mathcal{O}(LN^3 + LN^2Q)$	$\mathcal{O}(N^4 + LN^3 + L(\nu_P+\nu_S)(Q+N))$
$\mu(A_{11}, B_{11})$	$\mathcal{O}(LN^3 + LN^2Q)$	$\mathcal{O}(N^4 + LN^3 + L(\nu_P+\nu_S)(Q+N))$
1st order approximation		
	Dirichlet	MVN
inv. subspace	$\mathcal{O}(n_C^2N^4 + LN^2(Q+n_CN))$	$\mathcal{O}(n_C^2N^4 + L\nu_P(Q+n_CN))$
$\mu(L_1)$	$\mathcal{O}(N^3 + LN^2(Q+1))$	$\mathcal{O}(N^3 + LQ)$
$\mathbf{w}$	$\mathcal{O}(N^3 + LN^2(Q+N))$	$\mathcal{O}(N^4 + L\nu_P(Q+N))$
defl. subspace	$\mathcal{O}(n_C^2N^4 + LN^2(Q+n_CN))$	$\mathcal{O}(n_C^2N^4 + L(\nu_P+\nu_S)(Q+n_CN))$
$\mu(A_{11}, B_{11})$	$\mathcal{O}(N^3 + LN^2(Q+1))$	$\mathcal{O}(N^3 + LQ)$

$\mathcal{O}(N^4 + LN\nu + LQ\nu)$ for the MVN distribution; see Section C.2. For the MVN distribution, there is an initial cost of $\mathcal{O}(N^4)$ for the factorization of the covariance matrices. This cost could be reduced if the structure of these matrices was exploited; see Section C.2.

Although the constants in the running time approximations are missing, we want to point out that the computational costs for the sampling in case of the generalized eigenvalue problem are approximately twice as large as the costs for the standard eigenvalue problem.

For the 1st order approximation of the invariant subspace, the calculation of the matrix T in Equation (4.7) takes an initial time $\mathcal{O}(n_C^2N^4)$. Thus, the time for the initial computation of the eigenvalue problem can be neglected. In case of a Dirichlet distributed matrix E, the sampling of E in all L sampling steps takes time $\mathcal{O}(LN^2Q)$. Additionally, the matrix-vector product in Equation (4.7) takes time $\mathcal{O}(Ln_CN^3)$. For the MVN distribution, the multiplication of the standard normal random vectors with the covariance factor matrices Σ_i does not need to be computed in every sampling step but must be computed only once at the beginning.

The sampling of a standard normal random vector $\mathbf{y} \in \mathbb{R}^{\nu_P}$ takes time $\mathcal{O}(L\nu_P Q)$, and multiplication with $T\mathrm{diag}(\Sigma_{P,i})$ takes time $\mathcal{O}(L\nu_P n_C N)$.

For the mean cluster eigenvalue, the construction of the vector $\mathbf{t}$ in (4.8) requires the solution of the Sylvester equation (2.4), which takes time $\mathcal{O}(N^3)$. Thus, the formation of the Kronecker product, which takes time $\mathcal{O}(n_C N^2)$, can be neglected. While the sampling of $\Delta\mu$ from a Dirichlet distributed E requires a vector inner product of size N^2 with a standard normal random vector $\mathbf{y} \in \mathbb{R}^{N^2}$ in every sampling step ($\mathcal{O}(LN^2Q+LN^2)$), the MVN distribution only requires one computation of the variance $\mathbf{t}^\top \mathrm{diag}(C_{P,i})\mathbf{t}$ ($\mathcal{O}(N^3)$) and the sampling of one standard normal random variable in every step ($\mathcal{O}(LQ)$).

The computation of the sensitivity matrices $G_i^{\mathbf{w}}$ for the stationary density in (4.9) and (4.12) takes time $\mathcal{O}(N^3)$. For the 1st order MVN approximation, however, this time is dominated by the factorization of the covariance matrices ($\mathcal{O}(N^4)$). For the Dirichlet distribution, the generation of $\Delta\mathbf{w}_i$ requires N matrix-vector products with a matrix of size $N \times N$ and standard normal random vectors $\mathbf{y} \in \mathbb{R}^N$ ($\mathcal{O}(LN^2Q + LN^3)$), whereas the samplings of $\mathbf{y}_i \in \mathbb{R}^{\nu_{P,i}}$ and the matrix-vector products in (4.13) can be performed in time $\mathcal{O}(L\nu_P Q + L\nu_P N)$.

To summarize, one can say that the direct method is only applicable for moderate numbers of the sampling size L. One might get the impression that the computational costs for the solution of the eigenvalue problems in the direct method and the costs for the generation of the error matrices and the following matrix-vector products in the 1st order approximation are comparable. However, the approximated computational complexities for the solution of eigenvalue problems include large constants that account for the number of iterations in the eigenvalue solver. Therefore, the required computing time for the eigenvalue problems is much larger than for the sampling in the 1st order approximation. Moreover, the use of the Dirichlet distribution in the 1st order approximation of subspaces is not advisable because the sampling effort scales with N^3. The MVN distribution, however, allows a reduced description in terms of fewer standard normal random variables (often $\nu_P+\nu_S \ll N^2$) and is therefore preferable. It especially proves its superiority in case of scalar variables as the mean cluster eigenvalue, where it reduces to a standard normal distribution.

4.5.5 Toy example

Consider again the "toy system" from Section 1.1.6. In Section 4.4.2, 100 transition probability matrices have been generated according to a discretization with 16 nodes. For both types of discretization, Voronoi cells and radial basis functions, we now generate 10000 samples of E from the estimated Dirichlet parameters and compute the desired statistics.

The standard eigenvalue problem First consider the distribution of the mean cluster eigenvalue Δ_μ. We compare the result from the exact calculation (solid line in Figure 4.5(a)) with the 1st order approximation (4.8) (dashed line) and the

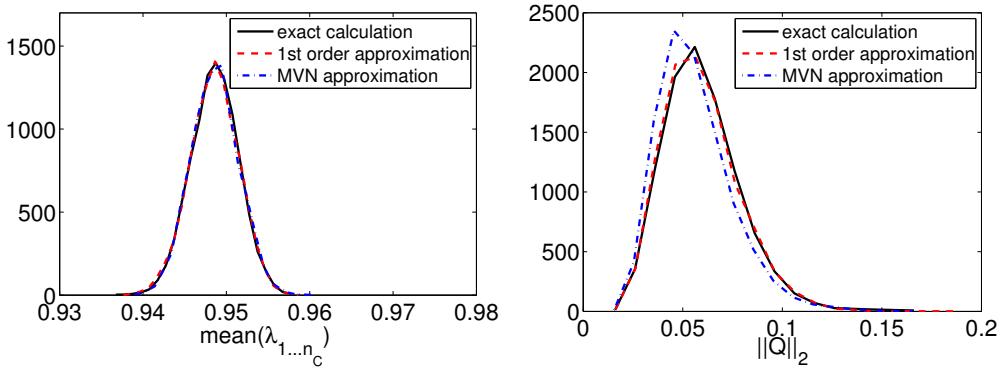

(a) Histogram of $\mu(\widetilde{L}_1)$.

(b) Histogram of $\|Q\|_2$.

Figure 4.5: Validation of different approximation schemes for $\|Q\|_2$ and $\mu(\widetilde{L}_1)$ from 10000 samplings. We observe that the 1st order approximation as well as the MVN approximation agree quite well with the exact calculation.

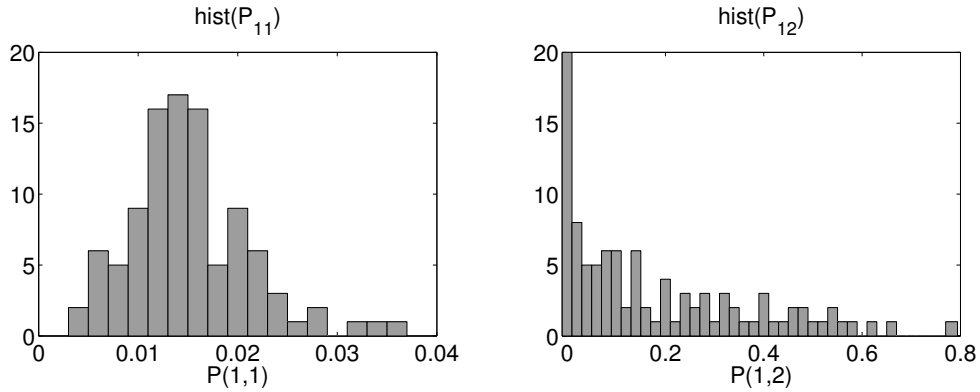

(a) Histogram of $P(1,1)$. The elements are distributed normally.

(b) Histogram of $P(1,2)$. The distribution is more similar to a log-normal than to a normal distribution.

Figure 4.6: Representative distributions of elements $P(i,j)$ over 100 different matrices. Although the distribution for larger matrix elements does not resemble a normal distribution, the row-wise MVN approximation gives good results.

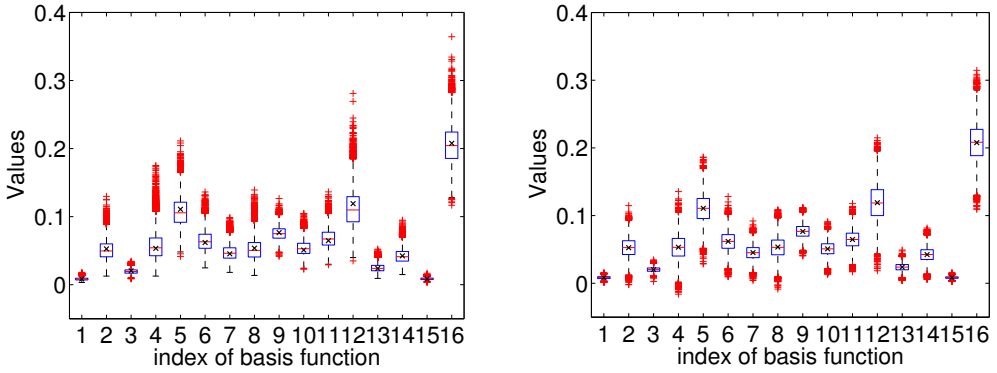

(a) Exact calculation. (b) MVN approximation.

Figure 4.7: Distribution of elements of the stationary density $\mathbf{w}$ resulting from 10000 Dirichlet samples of the error matrix E (left) and their MVN approximations resulting from 10000 direct samplings of $\mathbf{w}$ (right). The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The mean values are marked with a cross. We observe that the MVN approximation agrees rather well with the exact calculation, although it allows for negative entries in $\mathbf{w}$.

MVN approximation (dashed-dotted line). As it can be seen in Figure 4.5(a), the three histograms are nearly indistinguishable, which verifies the approximations. Moreover, the bound $|\mu(\widetilde{L}_1) - \mu(L_1)| \leq 0.0492$ from Section 4.4.2 is verified. The correctness of the MVN approximation is a bit surprising, because the distribution of some matrix elements $P(i,j)$ is far from being normal due to the occurrence of rare events, see Figure 4.6.

The histogram of $\|Q\|_2$ for samples of E from the Dirichlet distribution is represented by the solid line in Figure 4.5(b). It coincides very well with the 1st order approximation of Q (dashed line). As it can be seen in the figure, the MVN approximation (dashed-dotted line) slightly overestimates the 2-norm of Q. However, the approximation is good enough to give a quantitative estimate of the subspace error. Furthermore, the illustrated distributions verify the upper bound $2\gamma/\delta = 0.2146$ of $\|Q\|_2$ from Section 4.4.2.

Next, we analyze the stationary distribution. The results from the direct calculation are illustrated as box-plots in Figure 4.7(a). The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The MVN approximation gives

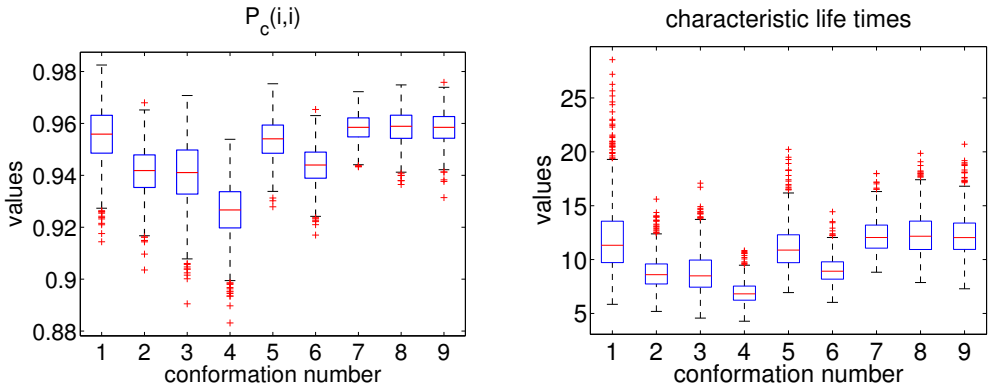

(a) Diagonal elements of the coupling matrix. (b) Characteristic life times of conformations.

Figure 4.8: Distribution of diagonal elements of the coupling matrix P_c (left) and corresponding characteristic life times (right) resulting from 1000 samples of the error matrix E. The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The standard deviation in the diagonal elements of P_c amounts to about 0.01 ($\approx 1\%$), whereas the standard deviation in the characteristic life times is about 1.8 ($\approx 17\%$).

very similar results, see Figure 4.7(b). Now the median values coincide with the mean values $\mathbf{w}_i$ because single elements are distributed according to normal distributions symmetrically around the mean values. However, the MVN approximation allows for negative entries in the stationary distribution, which is not the case in the exact calculations.

Furthermore, we compute the elements of the coupling matrix P_c according to (1.24). The membership vectors χ are obtained from PCCA+ via maximizing trace($\mathcal{S}$). The diagonal elements, which represent the probability of the system to stay within a certain conformation, have 95% confidence intervals of approximately ± 0.005, see Figure 4.8(a). This amounts to relative errors in the mean holding times of approximately 10% (Figure 4.8(b)).

Parameter sensitivity First of all, the quality of the MVN approximation mainly depends on the lengths of the sampling chains and thus remains valid also for a smaller number m of sampled matrices.

However, the number m of sampled rows influences the precision of the Dirichlet parameter vectors $\boldsymbol{\alpha}_i$. The larger m, the larger the precisions w_i. The results become more insensitive w.r.t. outliers if m increases. This can be illustrated by resampling sets of m matrices from the original set $\{P_i\}_{i=1}^{100}$ with replacement. The

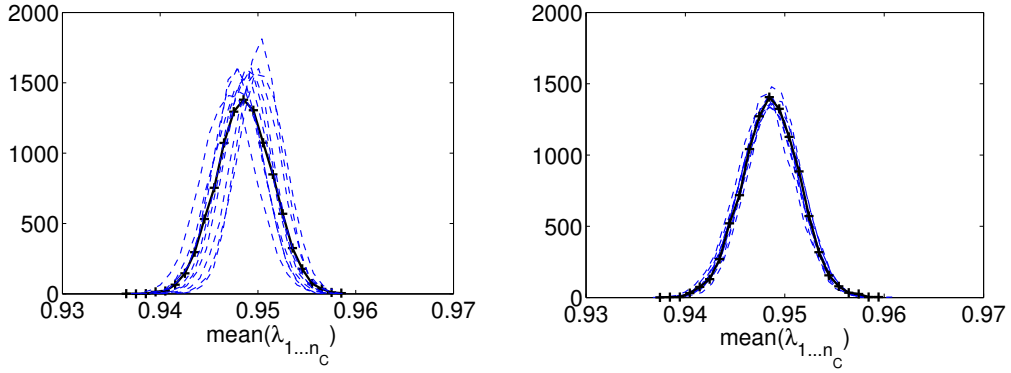

(a) $m = 10$ (b) $m = 100$

Figure 4.9: Histograms of the mean cluster eigenvalue μ for different collections of m matrices P_i from the original set of matrices $\{P_i\}_{i=1}^{100}$. The bold lines with +-markers represent the histograms for the original matrix set. The histograms have been obtained by sampling 10000 different values of μ via the MVN approach. The larger m, the smaller the deviation from the expected value.

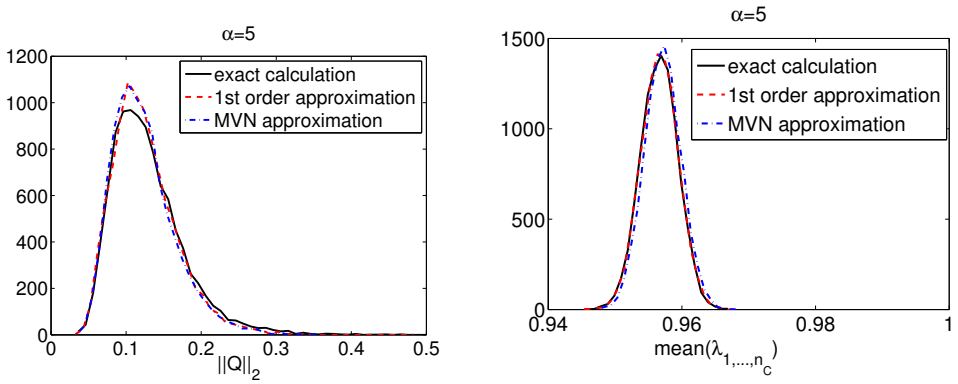

(a) Histograms of $\|Q\|_2$. (b) Histograms of $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$.

Figure 4.10: Histograms of $\|Q\|_2$ and $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$ from 10000 samplings. The 1st order approximation as well as the MVN approximation agree quite well with the exact calculation.

histogram plots of the mean cluster eigenvalue will look different for every such sampling because the estimated Dirichlet parameters will vary. However, the larger m, the smaller the variance and the more similar these histograms will be, see Figure 4.9.

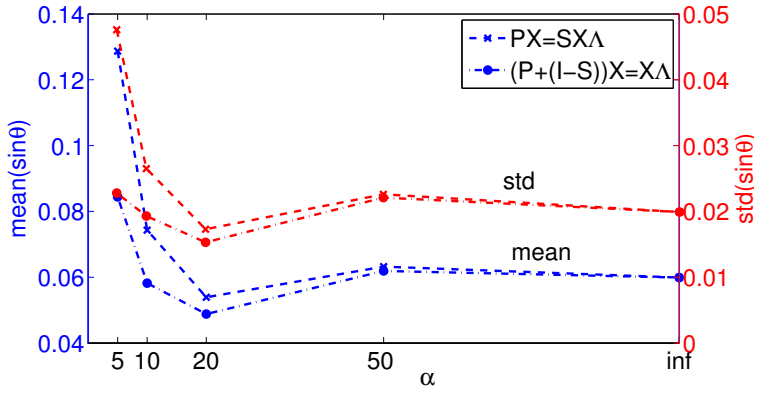

(a) $\sin(\theta_{\max})$

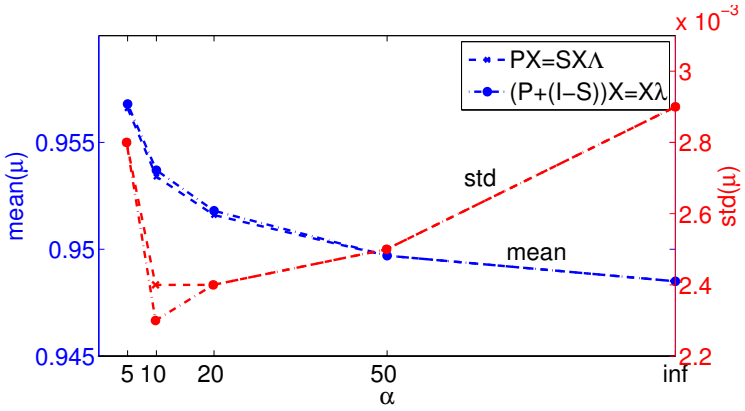

(b) $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$

Figure 4.11: Mean and standard deviation for $\sin(\theta_{\max})$ and $\mu(\widetilde{A}_{11}, \widetilde{B}_{11})$ with varying value of α, the shape parameter of the basis functions.

The generalized eigenvalue problem Figure 4.10 illustrates the validity of 1st order and MVN approximation for the generalized eigenvalue problem. We are especially interested in the question of how the sensitivity of the (generalized) eigenvalue problem changes with the shape parameter α. As already demonstrated in Section 4.4.2, the stochastic norms of the error matrices E and $E - F$ seem to have a local minimum for $\alpha \in [10, 50]$, see Figure 4.4(a). The same yields for the mean of the subspace deviation $\sin(\theta_{\max})$ (Figure 4.11(a)) and the mean of the standard deviation of the mean eigenvalue μ (Figure 4.11(b)). The reason is that the conformations are approximated best by basis functions with large value of α, compare Section 4.4.2. Thus, α should be chosen carefully in practical applications.

Chapter 5

Grid-free Hierarchical Refinement

Previous chapters have made clear that there are two main error sources in our approach to conformation dynamics: the quality of the discretization and the approximation of matrix elements by Monte-Carlo quadrature. This raises the following two questions:

- Given a discretization of the state space, how can one efficiently compute the entries of the transition probability matrix to sufficient accuracy?

- How must the discretization look like such that the sampling will be efficient and the clusters are represented correctly?

These two problems are coupled in the following way. Assume we are given a discretization such that one basis function covers two different conformations. Then the sampling of the stationary density converges slowly because trajectories tend to stay in one of the two conformations. Consequently, the entries in the corresponding row of the transition probability matrix will reveal large variances which implies large error bounds on the Perron subspace and other variables of interest. Given a fixed discretization, the goal is to distribute the sampling points among the basis functions in such a way that all basis functions contribute equally to the overall error. If a basis function requires a large amount of sampling points, this hints to the existence of metastabilities within that function. A decomposition of the basis function along the boundary between different conformations will resolve this problem on the next hierarchy level. In the present chapter, we want to demonstrate how an iteration of adaptive sampling and hierarchical refinement as illustrated in Figure 5.1 systematically reduces the stochastic norms of the error matrices and improves the accuracy of calculated variables.

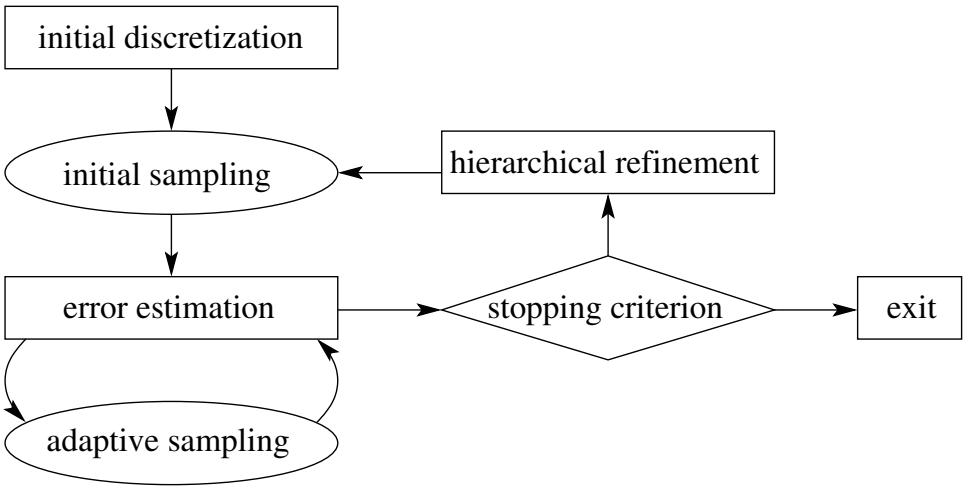

Figure 5.1: A sampling scheme that involves adaptive sampling and hierarchical refinement.

5.1 Adaptive sampling

In this section we are concerned with the first of the two questions cited above: Given a discretization of the state space, how can one efficiently compute the entries of the transition probability matrix to sufficient accuracy? So far we have considered a fixed number of horizontal trajectories with fixed chain length. However, since the matrix entries of P and S are evaluated by Monte-Carlo quadrature, the quadrature error will decrease with increasing number of sampling points at the worst case convergence rate of $\mathcal{O}(n^{-1/2})$ [86]. Thus, the estimated rows of the transition matrices P and S will get closer to the exact values, which will be reflected by an increased precision of the estimated parameters. In other words, the row wise variances will be reduced by an enhanced sampling, which improves the accuracy of calculated observables.

Here, adaptive sampling means that we add sampling points to the basis function that contributes most to the error. Within the basis functions, we always use standard Monte-Carlo methods. This is different to other concepts of adaptive sampling that would aim at choosing the sampling points within a basis function as elements of a low-discrepancy sequence such that the best possible order of convergence can be achieved with a fixed number of sampling points [81, 43]. These methods are also known as quasi-Monte-Carlo methods. They combine the low-discrepancy property of deterministic quadrature points with the advantage of Monte-Carlo methods that previous results must not be recomputed if more sampling points are added. However, quasi-Monte-Carlo methods are not in the scope of the present thesis. Moreover, it is unclear whether these methods can really outperform the standard Monte-Carlo approach in case of irregularly shaped domains

and non-smooth functions as they appear in our applications.

5.1.1 Horizontal sampling

The goal of adaptive horizontal sampling is to sample the partial Boltzmann distribution $\pi_i(q)$ within each basis function ϕ_i with the smallest possible number of sampling points. In case of a Voronoi discretization, we need a criterion that measures the distance to the desired distribution. This will be the Gelman-Rubin convergence factor r [34]. Gelman-Rubin convergence is a necessary condition that the distribution within a basis function is sampled correctly. In case of radial basis functions, we use a much stronger criterion. We want the errors in entries of the mass matrix S to be small. This requires an extensive sampling in the overlap region between different basis functions, which is difficult, if the overlap has small statistical weight. If the distribution in these overlap regions is approximated well, then it will also be accurate in the remaining parts of the basis functions.

Convergence towards the stationary distribution When we sample within a basis function, the horizontal sampling is said to have converged if the trajectory is rapidly mixing w.r.t. the partial density $\pi_i(q)$. That means, trajectories starting in different points will, after some relaxation time, sample the same distribution. Such a behavior can be analyzed via the Gelman-Rubin convergence factor. The Gelman-Rubin convergence indicator compares the within-chain-variances with the between-chain-variances and returns a potential scale reduction factor $r > 1$ that declines to 1 as the sampling size goes to infinity. When r is close to 1, each chain is close to the target distribution and the sampling can be stopped. When r is large, then further simulations may improve the inference about the target distribution. The decision to continue the sampling does not depend on other basis functions. The adaptive sampling procedure will be stopped as soon as r falls below a certain threshold TOL_{GR}.

Example 5.1.1. Let us revisit the example from Section 4.4.2 with the mesh-free discretization by 16 nodes. We start 5 trajectories within each basis function and compute the Gelman-Rubin convergence number r every 500 steps, ignoring the first 500 points. The horizontal sampling is stopped when $r < \mathrm{TOL}_{GR} = 1.01$. We perform this sampling strategy for the Voronoi discretization as well as for radial basis functions with $\alpha = 5$. Figure 5.2 shows the final distribution of sampling points among the basis functions for both discretizations. We observe that the points are almost equally distributed for the Voronoi discretization. In case of radial basis functions, some functions require much more sampling points than others. The reason is that the radial basis functions have larger support than Voronoi cells. In particular, if a radial basis function covers different metastable regions, the inherited metastability is much stronger than for a Voronoi cell. Figure 5.3 illustrates this behavior for the first basis function. Thus, a large Gelman-Rubin convergence number indicates metastabilities inside basis functions.

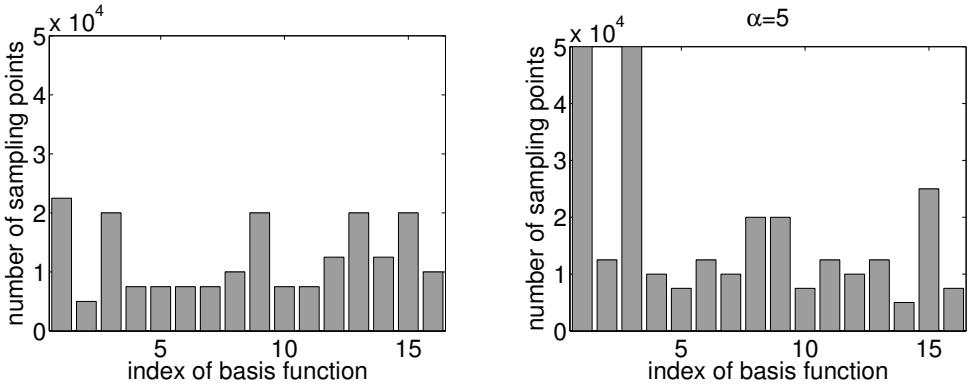

(a) Voronoi discretization. (b) Radial basis functions.

Figure 5.2: (Example 5.1.1) Distribution of sampling points among the basis functions to reach Gelman-Rubin convergence. In case of radial basis functions, metastabilities are more pronounced such that some functions require more sampling points than others.

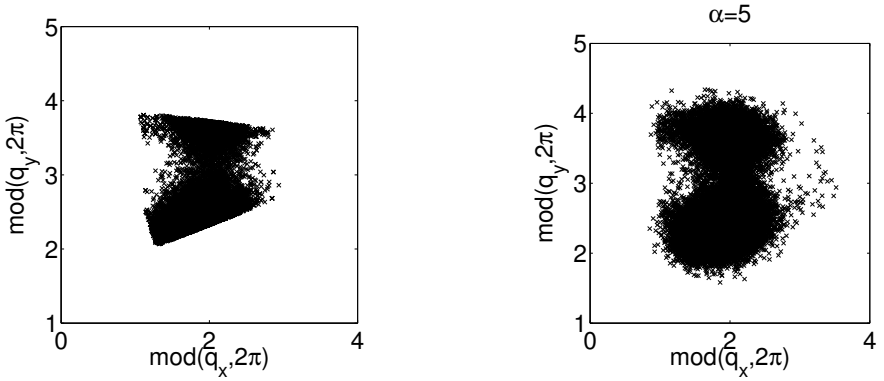

(a) Voronoi discretization. (b) Radial basis functions.

Figure 5.3: (Example 5.1.1) Representative horizontal trajectory for basis function no.1 with node $(1.9356, 2.9694)$. The radial basis function has larger support than the Voronoi cell such that the metastability is more pronounced.

Reducing the error in the mass matrix In case of radial basis functions, the horizontal sampling is used directly to compute the entries of the mass matrix S. Different horizontal trajectories started in different points will give different results for the corresponding row in S. From these candidate rows, one can estimate the parameters of the Dirichlet distribution and calculate the corresponding covariance matrix. Among other things, these covariance matrices enter the computation of F_{21} in the perturbation analysis described in Section 2.2.2. Since we do not yet know the matrix P, we have no information about the invariant subspace. Thus, we can only control the error matrix F in $S = \overline{S} + F$. The goal is to make $\|F\|_s$ as small as possible. According to (4.3), we have

$$\|F\|_s^2 = \mathbb{E}[\sum_i \sum_j F(i,j)^2] = \sum_i \sum_j \mathrm{var}[F(i,j)] = \sum_i \sum_j \mathrm{var}[S(i,j)]. \tag{5.1}$$

Single trajectory: First consider the case where a basis functions $\phi_i(q)$ is sampled by a single trajectory. The variance of $S(i,j)$ under the Dirichlet distribution with parameters u_{ij} is given by

$$\mathrm{var}[S(i,j)] = \frac{u_{ij}(w_i - u_{ij})}{w_i^2(w_i+1)} = \frac{\overline{S}(i,j)(1-\overline{S}(i,j))}{w_i+1}.$$

If m more samples are added and we assume that the expected transition probabilities and thus the parameters u_{ij} remain constant, then the variance sum $\sum_j \mathrm{var}[S(i,j)]$ will be reduced by

$$\Delta_i = \sum_{j=1}^{N} \left[\frac{\overline{S}(i,j)(1-\overline{S}(i,j))}{w_i+m+1} - \frac{\overline{S}(i,j)(1-\overline{S}(i,j))}{w_i+1} \right].$$

Thus, m more sampling points are added to the basis function k that reduces the variance sum the most,

$$k = \arg\max_i \Delta_i. \tag{5.2}$$

The sampling could be stopped as soon as $\|F\|_s$ becomes smaller than a given threshold TOL_F, a maximum number of sampling points is reached, or the decrease in $\|F\|_s$ is smaller than some value. A similar idea, applied to another variable, has been presented in [44].

Multiple trajectories: In case of multiple trajectories, each with length n_i, the variance of $S(i,j)$ under the Pólya distribution with parameters α_{ij} is given by

$$\mathrm{var}[S(i,j)] = \frac{(1+w_i/n_i)}{(w_i+1)} \overline{S}(i,j)(1-\overline{S}(i,j)).$$

An increasing of the chain lengths n_i will not decrease the variance, if n_i is already large and the chains locally mixing such that α_{ij} remains nearly constant. On the other hand, basis functions that inherit metastabilities are characterized by a small

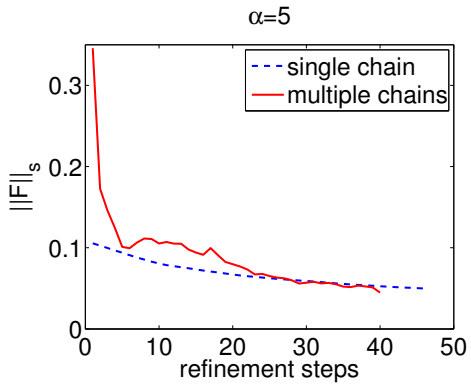

Figure 5.4: (Example 5.1.2) Reduction of $\|F\|_s$ by the adaptive horizontal sampling procedure for different numbers of sampling chains. In case of multiple chains, the existence of metastabilities within basis functions results in larger error estimates at the beginning.

precision w_i of the Dirichlet parameters $\boldsymbol{\alpha}_i$ and thus contribute much to the variance sum (5.1). Therefore, more sampling points will be added to the basis function that contributes most to the sum in (5.1) with index

$$k = \arg\max_i \sum_j \mathrm{var}[S(i,j)]. \tag{5.3}$$

Both sampling strategies (5.2) and (5.3) fall into the class of greedy algorithms, i.e. they make a locally optimal choice at each stage with the hope of finding the global optimum. The objective is to find the optimal distribution of sampling points such that $\|F\|_s$ becomes smaller than a certain threshold. However, we are not really interested in a global optimum because almost all solutions represent an improvement compared to the standard sampling method where the points are distributed evenly among the basis functions.

Example 5.1.2. Let us revisit the example from Section 4.4.2 with the discretization by $N = 16$ radial basis functions ($\alpha = 5$). We first start single trajectories of length 500 in every basis function. By assuming a Dirichlet prior with $\alpha_{ij} = 1/16$, the distribution of the probabilities $\mathbf{s}_i \equiv S(i,:)$ given the observed data is $\mathrm{Dir}(\mathbf{s}_i; \mathbf{u}_i)$ with $u_{ij} = \alpha_{ij} + 500S(i,j)$. According to (5.2) we determine the basis function that is expected to reduce the error most and add 500 more sampling points to the corresponding horizontal trajectory. Then the Dirichlet parameter vector $\boldsymbol{\alpha}$ is reestimated and the error analysis is repeated. This iteration is continued until $\|F\|_s \leq 0.05$, which happens after 45 iterations; see Figure 5.4. As shown in the histogram in Figure 5.5(a), the sampling points are distributed quite evenly among the basis functions. However, we cannot be sure that the horizontal chains really represent the local stationary distributions. Therefore, we repeat the experiment with

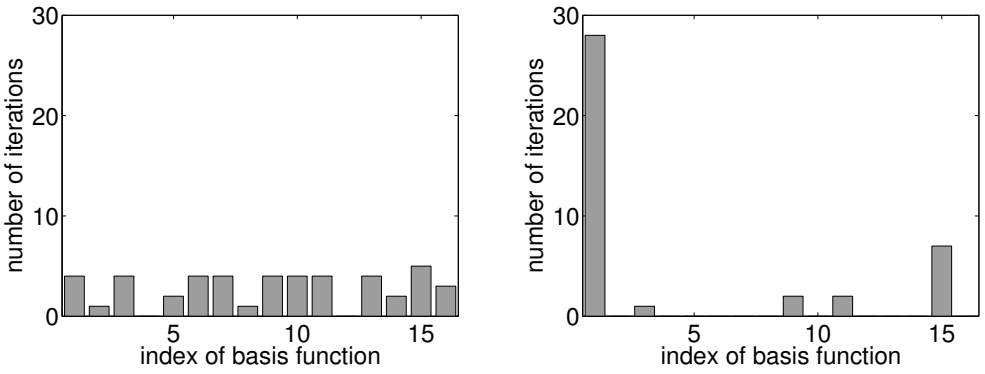

(a) Single chain. (b) 5 chains.

Figure 5.5: (Example 5.1.2) Distribution of horizontal sampling points among the basis functions for different numbers of chains. In case of multiple chains, basis functions with metastabilities require more sampling points and indicate insufficiencies in the discretization.

5 chains of length 500 each within every basis function. We consider the different rows $\{500\mathbf{s}_i^{(k)}\}_{k=1}^5$ as discrete outcomes from a Pólya distribution and estimate the parameters $\boldsymbol{\alpha}$ of the underlying Dirichlet distribution as described in Section 4.2.3. Then, according to (5.3), we determine the basis function that contributes most to the error and add 500 more sampling points to all 5 horizontal trajectories. After 40 iterations we end up with $\|F\|_s \leq 0.05$. The convergence behavior is illustrated in Figure 5.4. As one can see, the initial error is larger in case of multiple chains, which indicates the existence of basis functions with metastabilities. However, as the chain lengths increase the convergence rate becomes roughly the same as for a single chain. This indicates that the different chains are now mixing and give similar results for the corresponding row in S. Basis functions that inherit metastabilities require much more sampling points than others because the chains are not mixing at the beginning; see Figure 5.5(b). Thus, multiple chains are a useful tool to detect insufficiencies in the discretization.

In sum, one can say that a single chain can give misleading convergence results if it is too short. Moreover, a single chain is inadequate to decide whether a certain chain length is sufficient. Therefore, it is recommended to run multiple chains although the sampling effort is larger. This coincides with the numerous results about convergence towards the stationary distribution as, for example, the Gelman-Rubin convergence criterion. The reduction of $\|F\|_s$ can be used as alternative convergence criterion to Gelman-Rubin. It is of similar complexity, but more suitable to the problem in case of radial basis functions.

5.1.2 Vertical sampling

Both types of discretization, Voronoi cells or radial basis functions, give rise to a transition probability matrix P. In the same way as the horizontal sampling was enhanced to reduce the error in S , the vertical sampling can be enhanced to reduce the error in P. The advantage of vertical sampling is that the horizontal sampling is already given such that a perturbation analysis in terms of S **and** P can be performed. Thus, we can compute estimates for the invariant or deflating subspace and work with the projected error matrices. The goal is to reduce the error in P such that the preconditions of Theorem 2.1.2 or 2.2.2 become valid.

Consecutive vertical sampling Assume we are given a number of horizontal sampling points $\{q_k^{(i)}\}_{k=1}^{n_i}$ in basis function i which sample from the correct partial density $\pi_i(q)$. The vertical sampling can then be performed by drawing randomly points from the horizontal trajectory and propagating them with random initial momenta. This corresponds to single chain sampling. From the transition counts z_{ij} in case of Voronoi cells or $n_i p_{ij}$ for radial basis functions, respectively, we obtain an estimate for the Dirichlet parameter vector $\mathbf{u}_i$ by assuming a Dirichlet prior. In the following numerical examples, we will assume a Dirichlet prior with $\alpha_{ij} = 1/N$.

Moreover, in case of radial basis functions, we assume that F is fixed such that $\|F_{21}\|_s$ is much smaller than $\|E_{21}\|_s$. For Voronoi cells, $\|F_{21}\|_s$ is zero anyway. Since the upper bound on the subspace angle $\sin\theta(\mathcal{R}(X_1), \mathcal{R}(\overline{X}_1))$ decreases with decreasing $\|E_{21}\|_s$, the goal is to make this norm as small as possible. Since $\|E_{21}\|_s$ decomposes into a sum over all basis functions (compare Equation (4.6)), we can easily determine the one that contributes most to the norm. To be more precise, consider the row-wise covariance matrices (compare Equation (4.5)), which can be written as

$$C_i = \frac{1}{w_i + 1}[\mathrm{diag}(\overline{\mathbf{p}}_i) - \overline{\mathbf{p}}_i \overline{\mathbf{p}}_i^\top],$$

where $\overline{\mathbf{p}}_i = (\overline{P}(i,:))^\top$. If m more samples are added in basis function i and we assume that the expected transition probabilities remain constant, then the covariance matrix becomes

$$\hat{C}_i = \frac{1}{w_i + m + 1}[\mathrm{diag}(\overline{\mathbf{p}}_i) - \overline{\mathbf{p}}_i \overline{\mathbf{p}}_i^\top].$$

Thus, the ith summand in Equation (4.6) (standard eigenvalue problem) is expected to reduce from

$$p_i \equiv \|X_2(i,:)\|_2^2 \mathrm{trace}(X_1^H C_i X_1)$$

to

$$\hat{p}_i \equiv \|X_2(i,:)\|_2^2 \mathrm{trace}(X_1^H \hat{C}_i X_1).$$

For the generalized eigenvalue problem, X_2 is just replaced by Y_2. Consequently, we choose the state k with

$$k = \arg\max_i (p_i - \hat{p}_i). \tag{5.4}$$

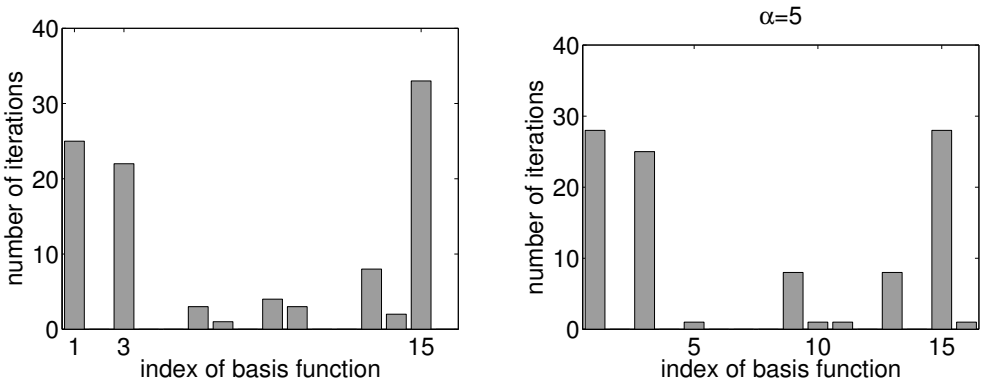

(a) Voronoi cells. (b) Radial basis functions.

Figure 5.6: (Example 5.1.3) Distribution of vertical sampling points among the basis functions. We do not observe any significant difference between Voronoi cells and radial basis functions.

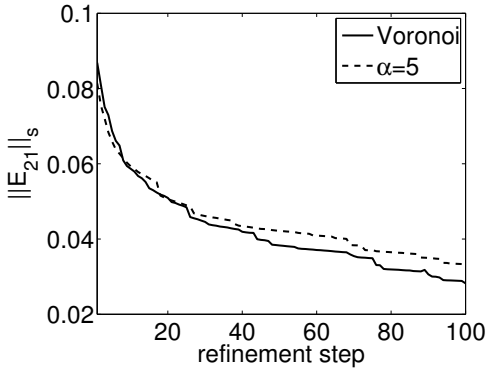

Figure 5.7: (Example 5.1.3) Reduction of $\|E_{21}\|_s$ for Voronoi cells and radial basis functions ($\alpha = 5$) by adaptive vertical sampling. We do not observe any significant difference between the two discretizations.

In other words, we will add sampling points to the basis function that is expected to reduce $\|E_{21}\|_s$ most.

Example 5.1.3. We start with the horizontal samplings that have been generated in Example 5.1.2. Since the Gelman-Rubin convergence factor r is small in all basis functions, the horizontal trajectories within one basis function can be concatenated into single trajectories. We want to demonstrate how $\|E_{21}\|_s$ increases by adaptively adding vertical sampling points to the basis function that is expected to reduce the error most. We first select randomly $m_0 = 100$ sampling points in every basis function from the horizontal trajectory and propagate them vertically with

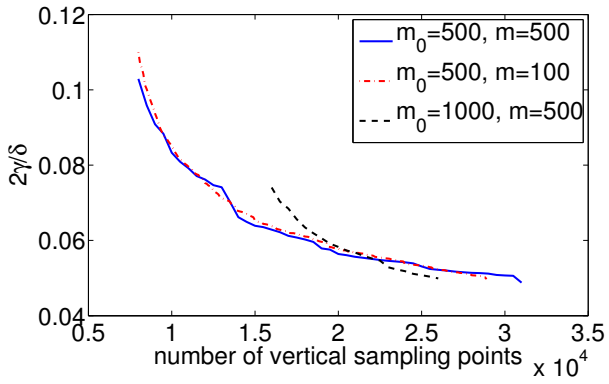

Figure 5.8: (Example 5.1.3) Convergence history of the adaptive vertical sampling for different values of m and m_0 in case of Voronoi cells. Larger values of m_0 and m improve the convergence, but could at the same time overshoot the mark.

random initial momenta. Then, we estimate the Dirichlet parameter vectors $\boldsymbol{\alpha}_i$ from the transition counts and the assumed prior. According to (5.4) we determine the index of the basis function that is expected to reduce $\|E_{21}\|_s$ most and select randomly $m = 100$ more sampling points from the horizontal trajectory for vertical propagation. Parameter estimation and index determination are continued over 100 iterations. As one can see from Figures 5.6 and 5.7, the convergence behavior does not depend on the type of discretization. Thus, the treatment of numbers $n_i p_{ij}$ as transition counts in case of radial basis functions is reasonable.

A comparison with Figure 4.2(a) shows that basis functions in transition regions (number 1, 3, and 15) require most sampling points.

Parameter sensitivity The adaptive vertical sampling consists of two main phases. In the first phase, a fixed number m_0 of points is chosen randomly from the horizontal trajectory in every basis function. Second, in each refinement step a fixed number m of points from a selected basis function is added for vertical sampling. The number m_0 determines the quality of the initial subspace basis. The larger m_0, the smaller the bound on $\|\sin\Theta\|$. If the initial bound is already close to the desired bound, fewer additional sampling steps might be required in the following adaptive sampling procedure; see Figure 5.8. This example is the same as the previous one based on a Voronoi disctretization, but this time the number of vertical sampling points is varied. Since the initial sampling will be improved by the following adaptive sampling, m_0 can in principle be very small. However, if it was too small, it could happen that some transitions between metastable regions do not occur and the perturbation analysis would fail. Hence, m_0 should be large enough to deliver an adequate initial solution but small enough to keep the sampling effort moderate. The smaller the parameter m, the more exactly the sampling points

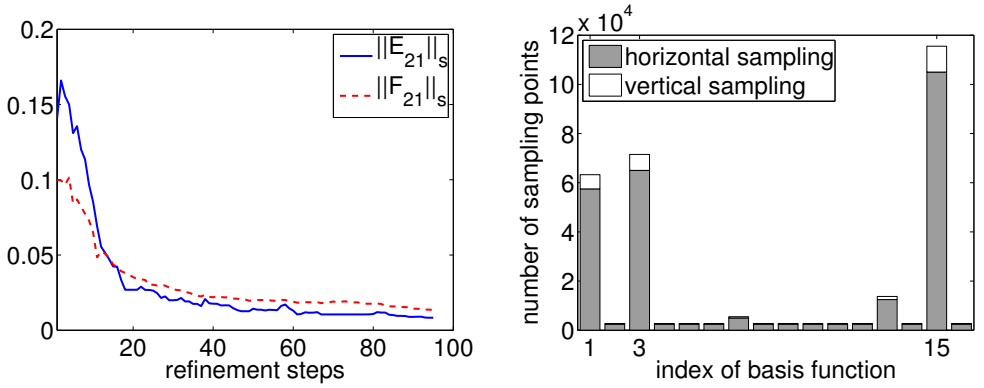

(a) Reduction of $\|E_{21}\|_s$ and $\|F_{21}\|_s$.

(b) Distribution of sampling points among the basis functions.

Figure 5.9: (Example 5.1.4) Results of the simultaneous horizontal and vertical sampling procedure. Although initially $\|E_{21}\|_s$ is larger than $\|F_{21}\|_s$, it decreases much faster such that less vertical sampling points are needed than horizontal points.

can be distributed among the basis functions. However, there is a computational overhead resulting from the calculation of the error bound after every refinement step. Therefore, m should not be too small. Furthermore, if m_0 and m were chosen too large, one could overshoot the mark, i.e. too many sampling points would be generated to reach a certain bound. In general, different values of m lead to similar distributions of sampling points.

Simultaneous horizontal and vertical sampling For radial basis functions, the horizontal sampling routine is stopped as soon as $\|F\|_s$ decreases to a certain threshold. The problem is that we do not know in advance which threshold on $\|F\|_s$ is necessary before having performed the vertical sampling. To circumvent this disadvantage, one could reduce the error in S and P simultaneously to meet a global convergence criterion.

By starting multiple horizontal chains and propagating points thereof vertically, we obtain a number of candidate rows for S and P. These rows are considered as outcomes of Pólya distributions and are used to estimate the underlying parameter vectors $\boldsymbol{\alpha}_i$. The basis function that contributes most to $\max(\|E_{21}\|_s, \|F_{21}\|_s)$ is selected for enhanced sampling,

$$k = \begin{cases} \arg\max_i \|Y_2(i,:)\|_2^2 \mathrm{trace}(X_1^H C_{S,i} X_1) & , \text{ if } \|E_{21}\|_s < \|F_{21}\|_s \\ \arg\max_i \|Y_2(i,:)\|_2^2 \mathrm{trace}(X_1^H C_{P,i} X_1) & , \text{ if } \|E_{21}\|_s > \|F_{21}\|_s \end{cases}. \tag{5.5}$$

Example 5.1.4. Previous experiments have shown that for radial basis functions $\|E_{21}\|_s$ decreases faster than $\|F_{21}\|_s$, i.e. to reach the same error tolerance for $\|E_{21}\|_s$

as for $\|F_{21}\|_s$, one needs fewer vertical sampling points than horizontal sampling points.

We started with 5 trajectories of length 500 in every basis function and propagated 50 randomly selected points in every trajectory vertically. Then we computed the 5 candidate rows for P and S and estimated the corresponding parameters of the Pólya distribution. Afterwards we computed the variables in Theorem 2.2.2 and determined the basis function that contributes most to $\max(\|E_{21}\|_s, \|F_{21}\|_s)$. The 5 horizontal trajectories in this basis function were then extended by further 500 sampling points, whereof 50 points were propagated vertically. Thus, the total number of horizontal sampling points is always ten times higher than that of vertical sampling points. The 5 candidate rows for S and P and the parameter vectors $\boldsymbol{\alpha}_i$ were then recalculated and the next basis function was determined. The iterations were stopped as soon as $\gamma\eta/\delta^2 < 0.25$, which happened after 93 iterations. The final bound on $\sin(\theta_{\max})$ amounts to $2\gamma/\delta = 0.1140$. At the beginning, $\|E_{21}\|_s$ is larger but it decreases faster than $\|F_{21}\|_s$; see Figure 5.9(a). Thus, the portion of vertical sampling points could have been chosen even smaller then 10 percent. The distribution of horizontal and vertical sampling points among the basis functions is illustrated in Figure 5.9(b). The final number of horizontal sampling points amounts to $2.725 \cdot 10^5$. Again, this example confirms the main insight: More sampling points are needed in transition regions (basis functions 1, 3, and 15) whereas only a few sampling points are required within metastable conformations.

5.1.3 Comparison of different sampling schemes

The simultaneous sampling as it has been performed in Section 5.1.2 has an important practical drawback. It requires a priori knowledge about the number of metastabilities, n_C. There are two possibilities to circumvent this problem. First, one could interrupt the sampling procedure after the first iteration and analyze the intermediate matrix pair (P, S) to select n_C according to the smallest condition number. However, the sampling chains must be rather long to ensure that all relevant parts of the sampling space have been captured. Otherwise, one might leave out some metastable regions, which leads to a wrong value for n_C. Since the following distribution of sampling points is optimized according to this value, one can obtain misleading results.

A second method that does not require the value n_C is to reduce $\|E\|_s$ and $\|F\|_s$ instead of the projected error matrices. However, since n_C and hence the subspace condition number are unknown in this case, it is difficult to define upper bounds on these norms and to select a stopping criterion for the sampling procedure. The same problem occurs in case of adaptive horizontal sampling, if the goal is to reduce $\|F\|_s$ alone.

Fore these reasons, we propose to perform the horizontal sampling first, and to start the vertical sampling afterwards. The basic algorithm for the adaptive vertical sampling is illustrated in Figure 5.10. In most cases, Gelman-Rubin convergence

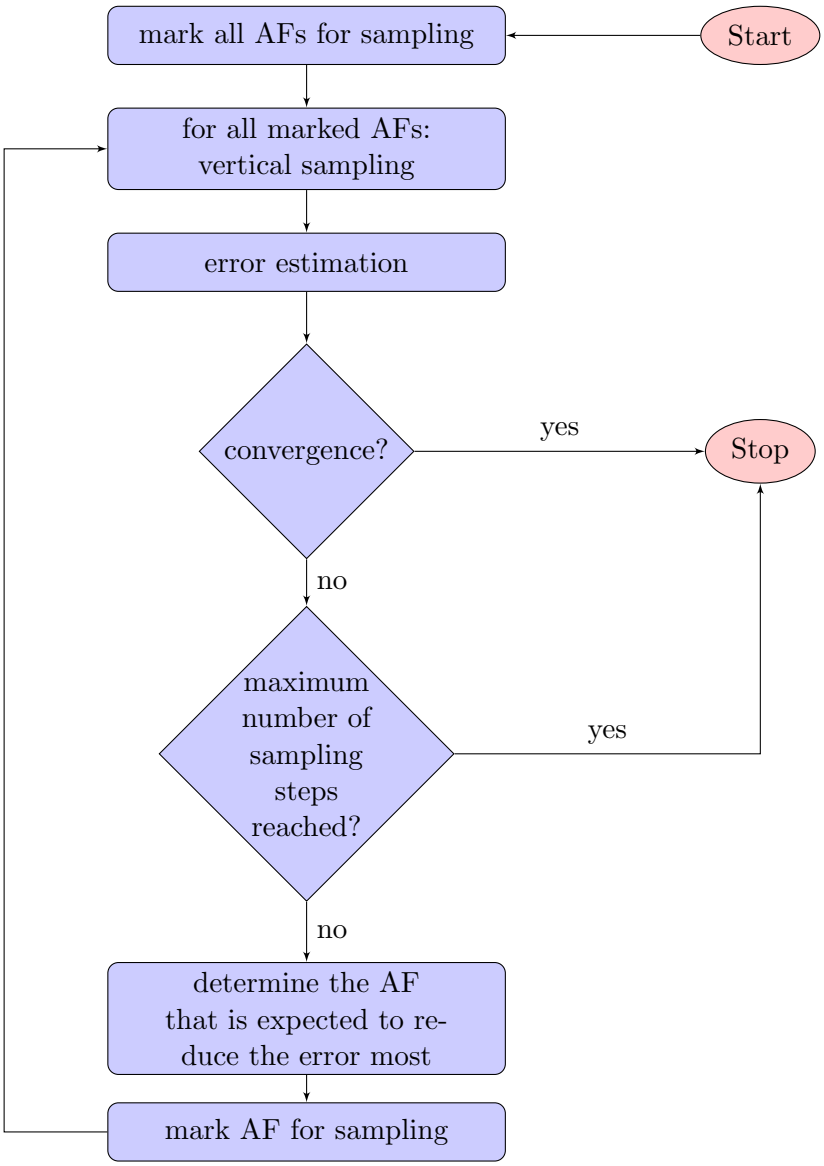

Figure 5.10: Algorithm for adaptive vertical sampling. The shortcut "AF" stands for ansatz or basis function.

implies a small error in S. Alternatively, one could first sample until Gelman-Rubin convergence is reached and then continue the horizontal sampling adaptively to further reduce $\|F\|_s$ as described in Section 5.1.1. Afterwards, the adaptive vertical sampling can be performed as described in Section 5.1.2. However, if Gelman-Rubin convergence cannot be achieved in some basis function within a maximum number of sampling steps, this hints to metastabilities within this function. Thus, it makes no sense to start a vertical sampling because the dynamical process is not rapidly mixing within this function. Instead, the basis function should be refined as it will be demonstrated in the following section.

5.2 Hierarchical refinement

In the present section we deal with the following question: How must the discretization look like such that the sampling will be efficient and the clusters are represented correctly? This question is closely related to the adaptive sampling procedure. Whenever the sampling effort for a certain basis function becomes too large, the discretization should be adapted by decomposing the function. The results from the previous section give rise to an algorithm in which alternating vertical sampling and hierarchical refinement lead to a final discretization that resolves the boundaries between different clusters. This algorithm is referred to as *vertical refinement* and explained in detail in Section 5.2.2. However, it turns out that this kind of algorithm is very expensive because it requires horizontal and vertical resampling of the refined functions. For this reason, a second algorithm, denoted as *horizontal refinement*, will be presented in Section 5.2.3. This algorithm arrives at a final discretization by alternating horizontal sampling and hierarchical refinement. Given the final discretization, an adaptive vertical sampling systematically reduces the error in the transition probability matrix. Thus, horizontal refinement decouples vertical sampling and hierarchical refinement, which makes it more efficient than vertical refinement and therefore applicable to larger systems.

The numerical examples are restricted to Voronoi based discretizations, but similar results could be derived for radial basis functions. However, since the condition number of the deflating subspace for the matrix pair (P, S) is larger than the condition number of the invariant subspace of P alone, the sampling effort for radial basis functions would be larger to achieve the same accuracy as for Voronoi cells. Moreover, for radial basis functions the evaluation of hierarchical basis functions in a single point q requires a recursive function call for *all* basis function. This is computationally more costly than in the Voronoi case, where only *one* basis function needs to be evaluated recursively.

5.2.1 Partitioning of basis functions

Whenever a basis function is selected for hierarchical refinement, its partitioning into new basis functions should satisfy two main properties. First, the decomposition must not influence the other basis functions, and second, the new set of basis

functions must again form a partition of unity and meet the positivity constraint. Weber [109] proposed a general way of constructing such a new basis. Assume we are given a set of basis functions $\{\phi_1, \ldots, \phi_N\} : \Omega \to [0, 1]$ whereof basis function k is selected for hierarchical refinement into s new basis functions. Then, a temporal set of basis functions $\{\widetilde{\phi}_{k1}, \ldots, \widetilde{\phi}_{ks}\} : \Omega \to [0, 1]$ is constructed such that this basis forms a partition of unity and meets the positivity constraint. Finally, the new basis functions are computed by

$$\phi_{ki}(q) := \phi_k(q)\widetilde{\phi}_{ki}(q), \quad i = 1, \ldots, s.$$

In fact, the new basis functions $\{\phi_1, \ldots, \phi_{k-1}, \phi_{k+1}, \ldots, \phi_N, \phi_{k1}, \ldots, \phi_{ks}\}$ form a partition of unity and meet the positivity constraint. Concerning the matrices P and S, hierarchical refinement means that row and column corresponding to the old basis function ϕ_k will be deleted and new rows and columns will be added. Unrefined basis functions must not be resampled, but their trajectories must be analyzed again to compute the interaction with the new basis functions.

The question remains of how to construct the temporal basis $\{\widetilde{\phi}_1, \ldots, \widetilde{\phi}_s\}$. Given the number of sub-states a basis function should be decomposed into, the task is to determine the center points for the functions on the next hierarchy level. Ideally, the boundary between the new functions coincides with the entropic or energetic barrier within the original basis function. However, in dimension 3 and higher it is impossible to detect these barriers with moderate effort. The original idea was to start a new pre-sampling at high temperature within the basis function and to select equidistant nodes thereof as new center points, which is computationally expensive. Therefore, an alternative node selection method became necessary. Finally, we decided to apply *k-means* clustering [41] to the horizontal trajectory that results from concatenating the single sub-trajectories. When the single chains are not mixing, i.e. if there are metastabilities within the basis function, k-means will probably deliver center points in different metastable regions. If there are at least two center points in different metastable regions, the boundary between these regions will be better resolved by the new basis functions than by the old ones. It might happen that the points selected by the k-means routine represent molecular configurations with low statistical weight, which complicates subsequent sampling methods. Therefore, the points from the horizontal trajectory of the decomposed basis function that are closest to the selected k-means center points are selected as nodes for the new basis functions.

5.2.2 Vertical refinement

In his doctoral thesis, Weber [109] pointed out that the convergence of Monte-Carlo sampling is slow whenever there exists an energy barrier within the modified potential $V_i(q)$. Horizontal trajectories starting on different sides of this barrier will seldomly or never cross it. When propagated vertically, these trajectories will show a different dynamic behavior resulting in different rows of the transition probability matrix. If the adaptive sampling method is applied to such a discretization, these

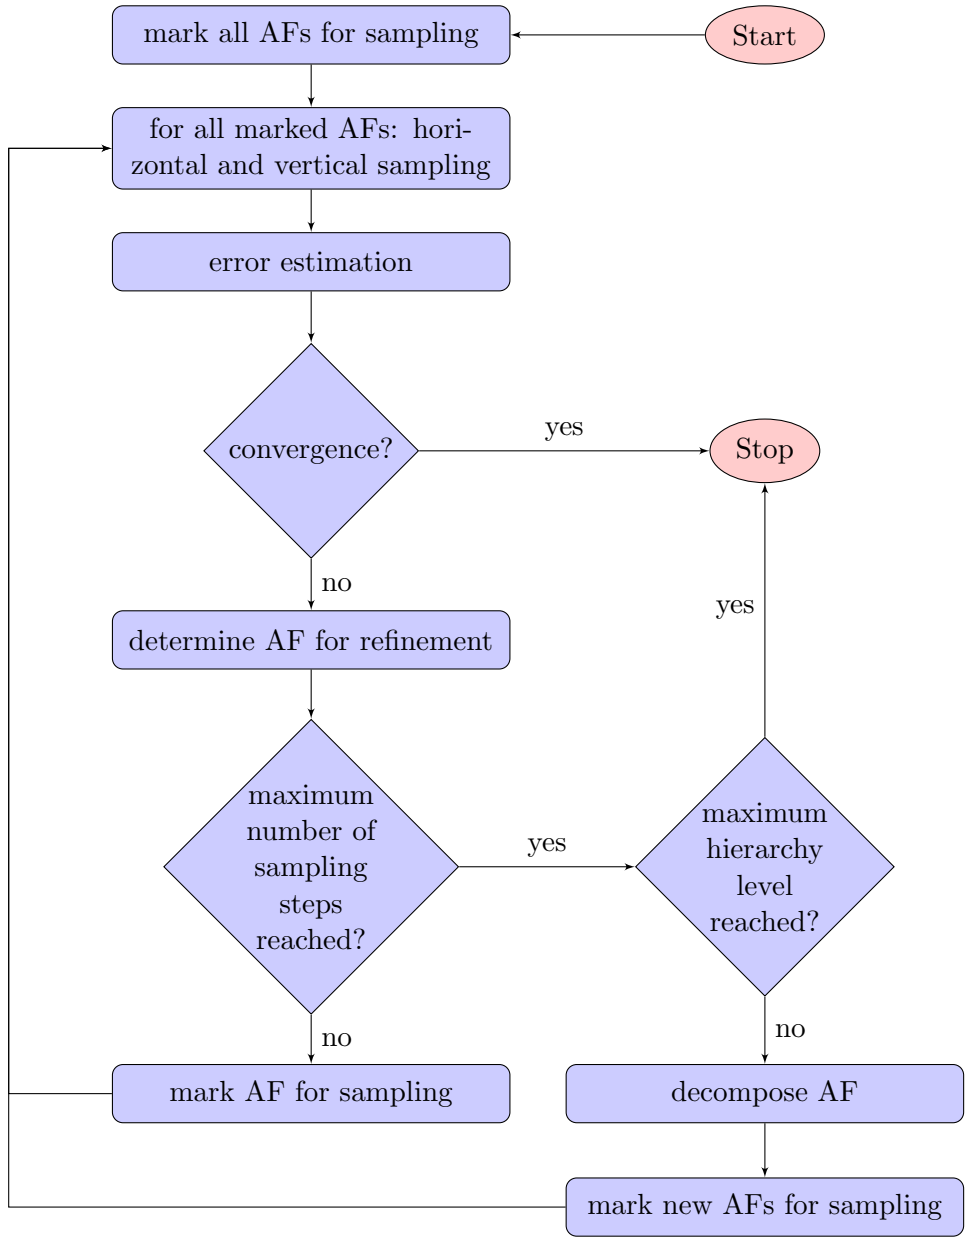

Figure 5.11: Algorithm for hierarchical refinement based on horizontal and vertical sampling. The shortcut "AF" stands for ansatz or basis function.

basis functions will require a large amount of sampling points. This will be used as indicator for a hierarchical refinement.

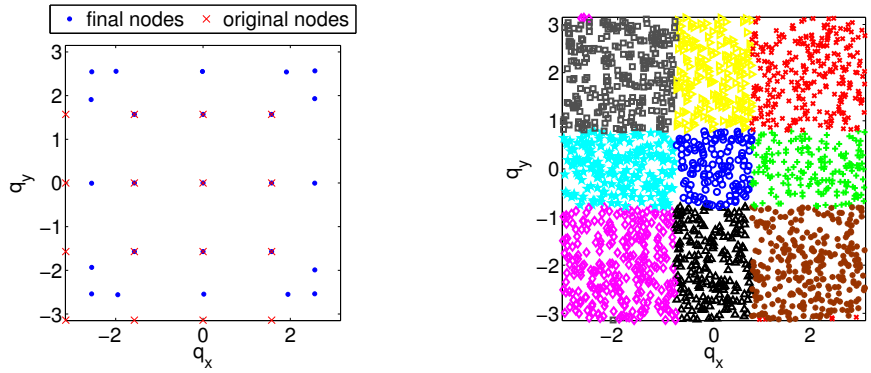

(a) Nodes before and after hierarchical refinement.

(b) Final clustering after vertical refinement.

Figure 5.12: (Example 5.2.1) Hierarchical refinement based on vertical sampling of a grid based starting discretization. The refinement steps shift the nodes away from the boundaries between metastable conformations such that these boundaries can be resolved by the final discretization.

Algorithmic details The basic scheme is presented in Figure 5.11. Given the matrix rows resulting from different sampling chains, one can estimate the parameter vectors $\{\boldsymbol{\alpha}_i\}_{i=1}^N$ and calculate the covariance matrices. The information is used to determine the basis function that contributes most to the error. Then, one can either extend the horizontal and vertical trajectories or refine the basis function. The hierarchical refinement will be stopped as soon as the preconditions of Theorem 2.1.2 or 2.2.2 are satisfied and the subspace angle $\sin\theta(\mathcal{X}_1, \overline{\mathcal{X}}_1)$ becomes smaller than a certain threshold $\mathrm{TOL}_{\sin(\Theta)}$.

Example 5.2.1. We return to the example from Section 4.4.2 with a Voronoi discretization, but now we aim at an automatic state space decomposition. We started with the same initial discretization as illustrated in Figure 4.1(a) and generated 5 horizontal sampling chains in every basis function, each one consisting of 1000 sampling points. Then the points were propagated over a time span $\Delta t = 0.5$ (50 MD-steps with $\tau = 0.01$), which results in 5 candidates for each row of P. We estimated the parameters $\boldsymbol{\alpha}_i$ of the Pólya distributions and computed the corresponding covariance matrices. Then we applied the perturbation theory and selected the state that contributes most to $\|E_{21}\|_s$ for refinement. This procedure was continued iteratively until the stopping criterion $2\gamma/\delta < \mathrm{TOL}_{\sin(\Theta)} = 0.1$ was satisfied. The initial and final nodes are shown in Figure 5.12(a). Each refinement step reduces the subspace error, see Figure 5.13. After nine refinement steps, the value of $\gamma\eta/\delta^2$ becomes smaller than 0.25 such that the error bound becomes valid; see Figure 5.13(b). After 10 refinement steps, we end up with $\gamma\eta/\delta^2 = 0.0839$, $\|E\|_s = 0.1077$, and $\|\sin\Theta\| \leq 2\gamma/\delta = 0.0868$. Now the boundaries between the

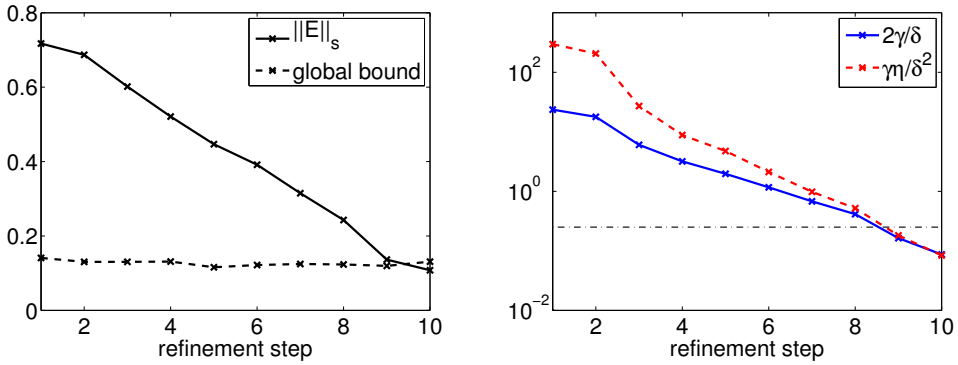

(a) Decreasing of $\|E\|_s$ toward the global bound (3.8).

(b) Decreasing of the subspace error bound $2\gamma/\delta$ and the precondition $\gamma\eta/\delta^2 < 0.25$ (dashed-dotted horizontal line).

Figure 5.13: (Example 5.2.1) Hierarchical refinement based on vertical sampling of a grid based starting discretization. Convergence history.

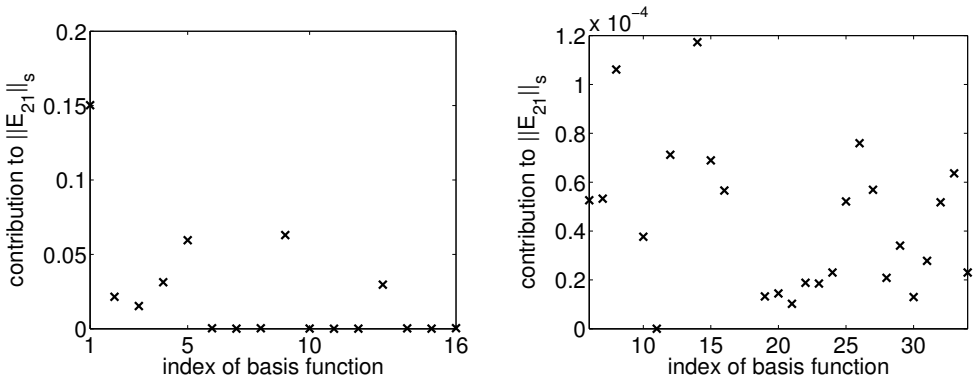

(a) Start discretization.

(b) Final discretization.

Figure 5.14: (Example 5.2.1) Hierarchical refinement based on vertical sampling of a grid based starting discretization. Contribution of basis functions to $\|E_{21}\|_s$ in the start discretization and after hierarchical refinement. Initially, the errors are distributed non-uniformly among the basis functions, whereas the error contributions are equilibrated in the final discretization.

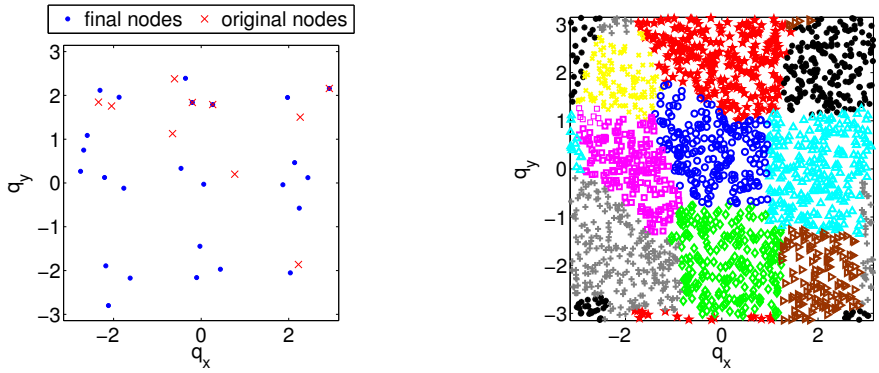

(a) Nodes before and after hierarchical refinement.

(b) Final clustering after vertical refinement.

Figure 5.15: (Example 5.2.2) Hierarchical refinement based on vertical sampling of a few initial basis functions. Distribution of nodes and clustering. Hierarchical refinement generates new nodes in metastable conformations which had not been resolved by the initial discretization.

clusters can be resolved; see Figure 5.12(b). As expected, the refinement strategy led to an equilibration of errors among the basis functions; see Figure 5.14.

The algorithm also works for irregularly distributed initial sampling points. If one uses a priori knowledge about the target number of clusters, and if the initial number of basis functions is larger than this number, then consecutive refinement steps will lead to a distribution of nodes in all metastable regions as will be illustrated in the following example.

Example 5.2.2. We selected randomly $N = 10$ nodes from a pre-sampling trajectory of length 2000 at $\beta_{\text{pre}} = 0.1$; see Figure 5.15(a). The horizontal and vertical sampling were performed with the same parameters as before. The basis function that contributed most to the stochastic norm of E_{21} was refined hierarchically until convergence; see Figure 5.16. Indeed, the final nodes cover all metastable regions, as illustrated in Figure 5.15(a). While the original discretization cannot resolve all metastable regions, the final nodes can do, see Figure 5.15(b).

Since the target number of clusters is seldomly known a priori, we will now propose a second refinement strategy which is solely based on the convergence of the horizontal sampling towards the stationary distribution.

5.2.3 Horizontal refinement

The goal is to decompose basis functions that inherit metastabilities. Since horizontal trajectories starting on different sides of an energy barrier will seldomly or

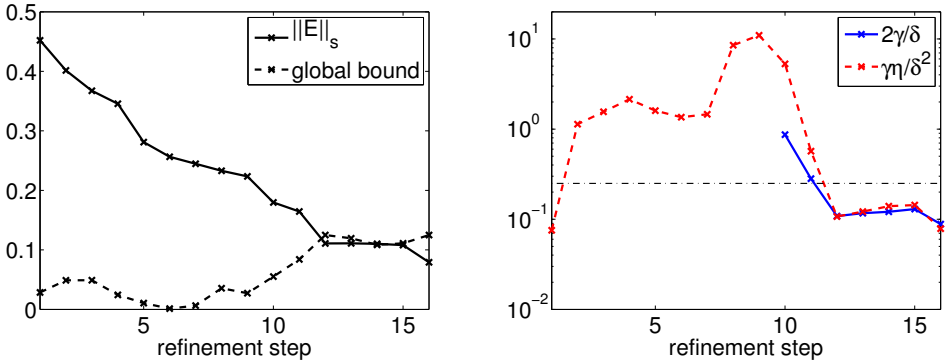

(a) Decreasing of $\|E\|_s$ towards the global bound (3.8).

(b) Decreasing of the subspace error bound $2\gamma/\delta$ (plotted only for $\delta > 0$) and the precondition $\gamma\eta/\delta^2 < 0.25$ (dashed-dotted horizontal line).

Figure 5.16: (Example 5.2.2) Hierarchical refinement based on vertical sampling of a few initial basis functions. Convergence history. It takes a few hierarchical refinement steps until all metastable conformations are detected and the preconditions of Theorem 2.1.2 become valid.

never cross it, metastabilities can also be detected by analyzing the variance of the horizontal sampling.

Algorithmic details The algorithm decomposes into two parts. First, the horizontal sampling is performed for each basis function and the discretization is refined such that these samplings converge; see Figure 5.17. Second, the adaptive vertical sampling is performed as in Figure 5.10.

When we sample within Voronoi cells, the horizontal sampling is said to have converged if the Gelman-Rubin convergence factor r is close to 1. For every basis function, the horizontal trajectories will either produce a small value for r or reach the maximal chain length. In the second case, the basis function will be marked for refinement. For radial basis functions, the function that contributes most to the norm $\|F\|_s$ will be marked for sampling. Then either the horizontal trajectories will be continued or the basis function will be marked for refinement. As explained in Section 5.1.2, we do not know in advance which bound on $\|F\|_s$ is necessary before having performed the vertical sampling. There is a possibility to circumvent this problem. Whenever the metastable conformations are geometrically well separated, a cluster analysis on S alone would reveal the conformations. In other words, we could compute the invariant subspace $\mathcal{R}(\overline{X}_1)$ of $\overline{S}$ in every iteration and stop as soon as the conditions of Theorem 2.1.2 are satisfied. Since the condition number sep is nearly twice as large as dif, one should nevertheless continue the iterations until $\|F\|_s$ is a factor smaller than the desired threshold.

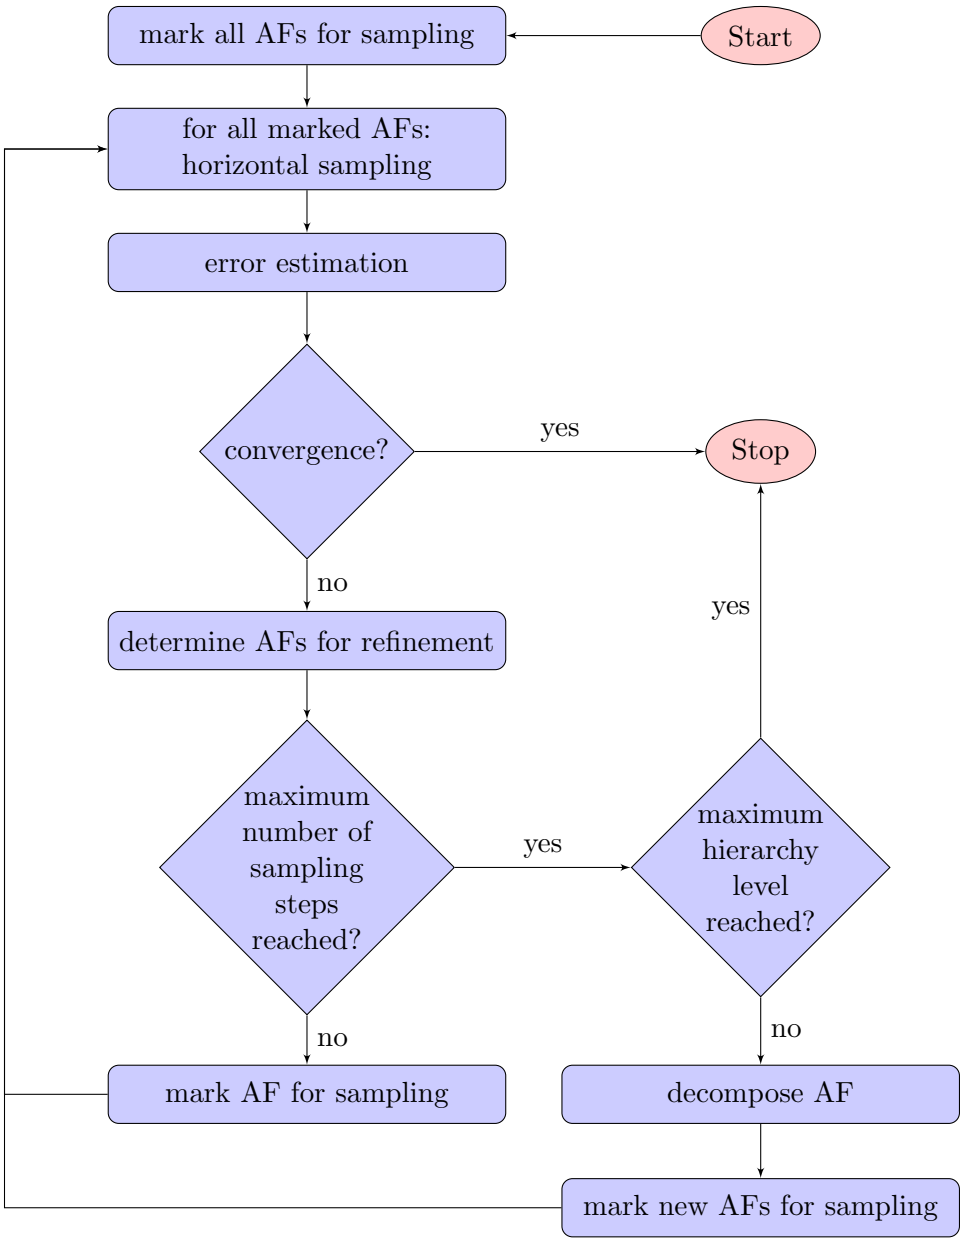

Figure 5.17: Algorithm for hierarchical refinement based on horizontal sampling. The shortcut "AF" stands for ansatz or basis function.

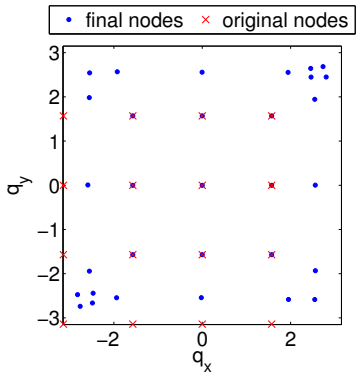

(a) Nodes before and after hierarchical refinement.

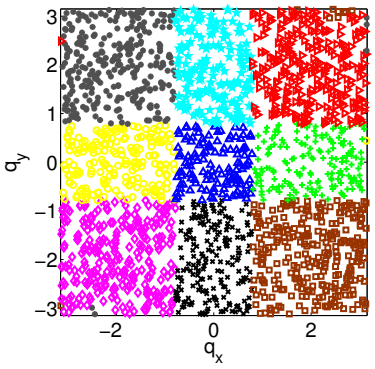

(b) Final clustering.

Figure 5.18: (Example 5.2.3) Hierarchical refinement based on horizontal sampling of a grid-based starting discretization. The final discretization is very similar to the discretization obtained by vertical refinement (compare Figure 5.12(a)).

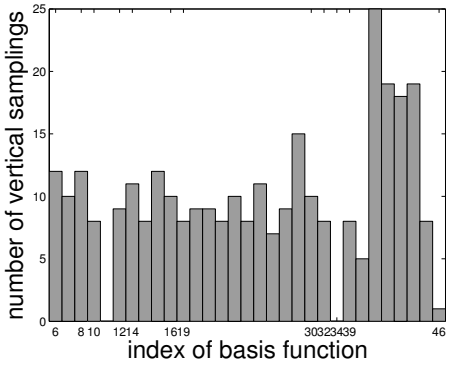

(a) Distribution of sampling points.

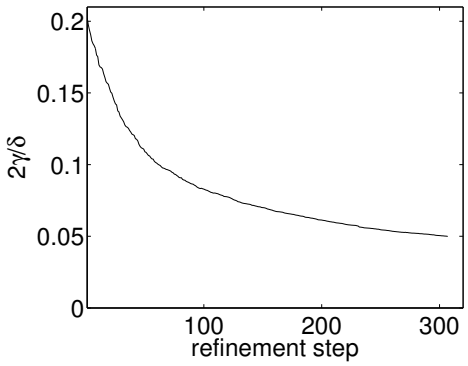

(b) Convergence history.

Figure 5.19: (Example 5.2.3) Adaptive vertical sampling w.r.t. the final discretization in Figure 5.18(a). The sampling points are distributed rather evenly among the basis functions which is a measure for the quality of the final discretization.

Example 5.2.3. First we consider the grid-based initial discretization from Section 4.4.2. In every basis function we generated 5 horizontal trajectories with a maximum of 2000 sampling steps each. Every 500 steps, we computed the Gelman-Rubin convergence factor r corresponding to the 5 different chains. If the number r did not become smaller than $\mathrm{TOL}_{GR} = 1.5$, the basis function was marked for refinement. The final discretization is very similar to the one achieved by vertical refinement, compare Figures 5.18(a) and 5.12(a). Finally, the horizontal trajectories within a basis function were concatenated to one single trajectory. Then, in every basis function 200 points from the horizontal trajectory were selected randomly with replacement and propagated vertically (50 MD steps with $\tau = 0.01$) with random initial momenta. These samples were used to estimate the Dirichlet parameters and to compute the subspace error. The basis function that was expected to reduce $\|E_{21}\|_s$ most was marked for further sampling. From the horizontal trajectory of this basis function, we selected 200 further points for vertical propagation and recomputed the corresponding row in P. This sampling procedure was continued iteratively until the bound $2\gamma/\delta$ on $\sin(\Theta)$ became smaller than 0.05. This happened after 307 refinement steps. Within these 307 steps, the bound on $\sin\Theta$ could be reduced from 0.2023 to 0.0499; see Figure 5.19(b). The histogram in Figure 5.19(a) indicates that the fewest sampling points are necessary for basis function number 11, which corresponds to the node $[0, 0]$, and basis function number 34, which corresponds to the node $(-2.5476, 2.5434)$. Indeed, these basis functions are located in centers of metastable regions. On the other hand, basis function number 41 requires many sampling points. It corresponds to the node $(-2.7623, -2.7388)$, which is located at the outer bound of the metastable cluster in the lower left corner (see Figure 5.18(a)) and thus close to the cluster in the upper right corner. This is exactly what one expects. Basis functions in transition regions require more sampling points than basis functions in metastable regions. However, by and large the sampling points are distributed quite regularly among the basis functions, which is a desired property of the discretization.

Example 5.2.4. Even if we start with only two basis function, the hierarchical refinement with $\mathrm{TOL}_{GR} = 1.3$ leads to a distribution of nodes which covers all metastable regions; see Figure 5.20(a). While it is obviously impossible to cluster Ω into metastable regions with these two basis functions, the final discretization can do so; see Figure 5.20(b). Subsequent adaptive vertical sampling with 200 sampling points in each iteration reduces the subspace error from 0.2384 to 0.0499 within 221 iterations, see Figure 5.21(b). The histogram in Figure 5.21(a) indicates that basis function number 29 is the critical one. It corresponds to the node $(-1.0121, -0.4105)$, which is located between two clusters. This is also verified by its membership vector

$$\chi = [0.0027,\ 0.0020,\ 0.0052,\ 0.0082,\ 0.0001,\ 0.4468,\ 0.0534,\ 0.0526,\ 0.4290],$$

which shows that the basis function belongs to cluster 6 and 9 with similar weight.

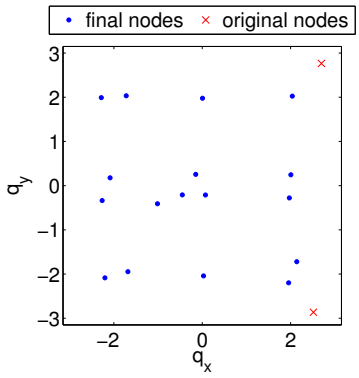

(a) Nodes before and after hierarchical refinement.

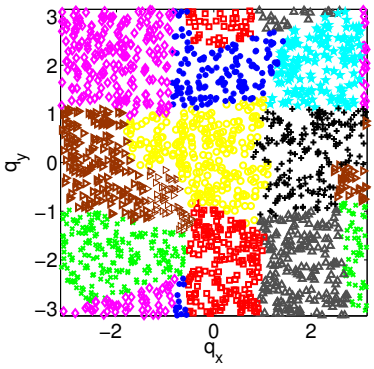

(b) Final clustering.

Figure 5.20: (Example 5.2.4) Hierarchical refinement based on horizontal sampling of a start discretization with two random nodes. Similar to vertical refinement, new nodes are generated in metastable conformations that are not resolved by the initial discretization (compare Figure 5.15(a)).

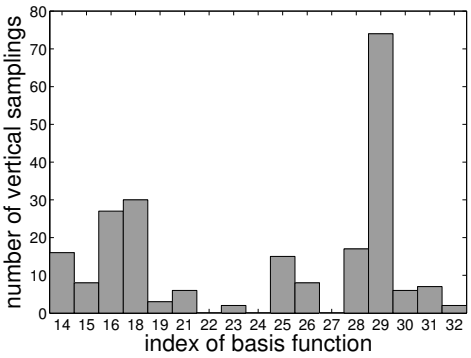

(a) Sampling histogram.

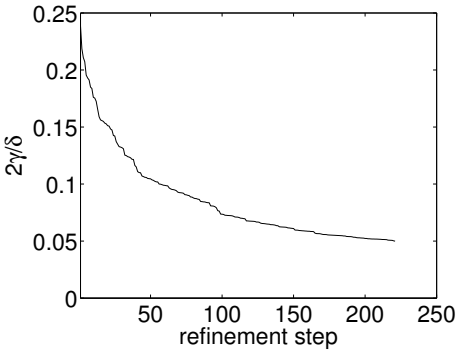

(b) Convergence history.

Figure 5.21: (Example 5.2.4) Adaptive vertical sampling w.r.t. the final discretization in Figure 5.20(a). The distribution of sampling points suggests a further refinement of basis function 29.

5.3 Discussion

As demonstrated in this chapter, there are several possibilities for adaptive sampling and hierarchical refinement. Which method shall we choose for the biomolecular application in the following chapter?

The first choice concerns the type of basis functions. Radial basis functions have some practical drawbacks. They can lead to an ill-conditioned eigenproblem and do not allow for an interpretation of the dynamical process as Markov chain. Therefore, we prefer Voronoi cells. Their only disadvantage compared to radial basis functions is the reduced acceptance ratio in HMC sampling. We did not observe this problem in our simulations, but whenever it occured, it could be diminished by the use of umbrella sampling.

Second, one has to decide between vertical and horizontal refinement. As already mentioned, horizontal refinement is computationally cheaper than vertical refinement. In contrast to vertical refinement, horizontal refinement of an insufficient start discretization does not use a priori knowledge about the number of metastabilities. The goal of horizontal refinement is to detect the correct number of metastabilities, whereas vertical refinement aims at a good approximation of a *fixed* number of metastabilities. Consequently, we will apply horizontal refinement for the biomolecular example. In case of Voronoi cells, the refinement will be based on the Gelman-Rubin convergence factor, which requires multiple chain sampling.

Given the final discretization, the natural way to proceed is to apply adaptive vertical sampling. For this purpose, the appropriate number of metastable conformations must be identified. We suggest to run an initial vertical sampling with a small number of sampling points in every basis function. From the initial guess for the matrix $\overline{P}$ it should be possible to select an appropriate value for n_C according to the spectral gap or the condition number. Then the vertical sampling can be continued adaptively in order to reduce the error in the corresponding invariant subspace, i.e. to reduce $\|E_{21}\|_s$. If it turns out that some basis function requires much more sampling points than others, it can further be refined hierarchically.

If there does not exist a number n_C of metastable conformations for which the invariant subspace is well-conditioned, the vertical sampling should be performed adaptively in order to reduce $\|E\|_s$ instead of $\|E_{21}\|_s$. Although the preconditions for the validity of perturbation bounds might not be satisfied in this case, error distributions for the invariant subspace could be derived.

The proposed sampling and refinement strategy will be illustrated in the next chapter, where the application of conformation dynamics to a small biomolecule is presented.

Chapter 6

A Biomolecular Application

The main challenge in molecular simulations is the high dimensionality of the state space. Besides rotation and translation of the complete molecule, positions and momenta of all atoms are changed during the molecular dynamics simulations. The positional degrees of freedom are comprised in form of internal coordinates: bond lengths, bond angles, and dihedral angles. Since bond lengths and angles mostly oscillate around mean values, molecular conformations can often be described by a small number of dihedral angles [1] whose selection requires some chemical intuition. The methods presented in this thesis are based on the assumption that the underlying dynamical system exhibits metastabilities w.r.t. these dihedral angles. In general, the sampling effort depends on the separation between the fast timescales that describe the molecule's motion within a metastable conformation, and the slow timescales on which the molecule switches between different conformations. If these time scales are clearly separated, there will be a cluster of eigenvalues close to the Perron root which is well separated from the rest of the spectrum. Then, the corresponding invariant or deflating subspace will be well conditioned such that the perturbation in the matrix entries can be of moderate size. In practice, however, many biomolecules do no exhibit such a behavior, at least not in such a pronounced way. This is mainly due to a low flexibility inside metastable conformations, very high energy barriers between them, and the existence of "sub-metastabilities". By sub-metastabilities we mean parts of the conformational space where the holding time of a trajectory is much shorter than in the identified metastable conformations, but at the same time much larger than in typical transition regions. In the discrete setting, this behavior is reflected in a spectrum that decays slowly from the Perron root. Sometimes, chemical knowledge is required to identify the conformations correctly. A typical example will be presented in the following, the molecule hexabromocyclododecane.

The simulations in this chapter have been performed with the software package ZIBgridfree [113], which has been developed at Zuse Institute for four years now. It includes the Merck Molecular Force Field (MMFF) [39, 40], which has been designed for small molecules.

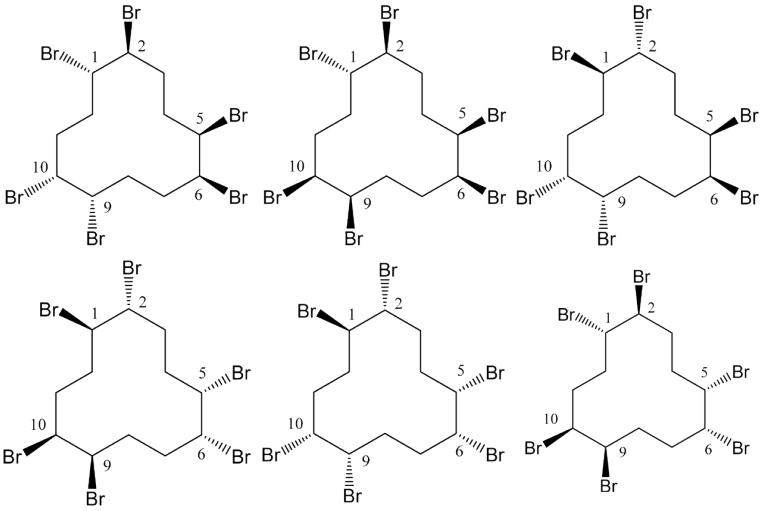

Figure 6.1: Different stereoisomers of technical HBCD. *Top row from left to right*: (+)-α-HBCD, (+)-β-HBCD, (+)-γ-HBCD. *Bottom row from left to right*: (-)-α-HBCD, (-)-β-HBCD, (-)-γ-HBCD.

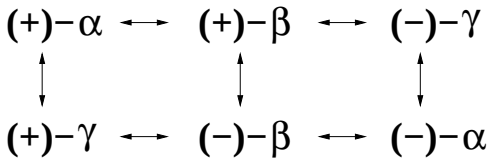

Figure 6.2: Possible interconversion reactions of HBCD isomers connected to the three sets of vicinal bromine atoms.

6.1 Hexabromocyclododecane

The molecule 1,2,5,6,9,10-hexabromocyclododecane (HBCD) is one of the major flame retardant additives to plastics and textiles. Besides trace amounts, which can be found in all parts of the environment, higher concentrations of the α-isomer have recently been found in sediments, clearing sludge, fish, and even humans [12, 13, 11]. Therefore, HBCD is currently one of the emerging environmental analytes of interest, and an EU risk assessment is under way.

Technical HBCD consists mainly of three diastereometric pairs of enantiomers; see Figure 6.1. Since, in the absence of a chiral environment, the (-)-enantiomers behave in the same way as the (+)-enantiomers, only one group must be investigated, here the (+)-enantiomers. It was shown experimentally [85, 42, 48] that a mixture containing α-, β-, and γ-HBCD in any composition interconverts towards an equilibrium that is dominated by α-HBCD. The reaction pathways are depicted in Figure 6.2. A configuration of (+)-α-HBCD, visualized by amira$^{\mathrm{TM}}$[2], is illustrated

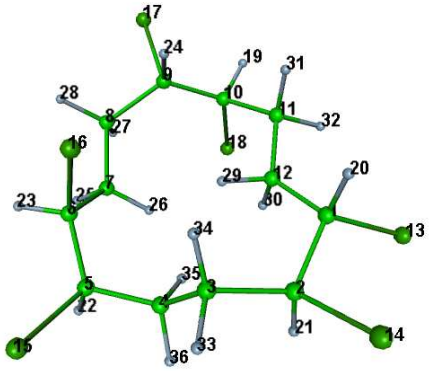

Figure 6.3: Structure of (+)-α-HBCD.

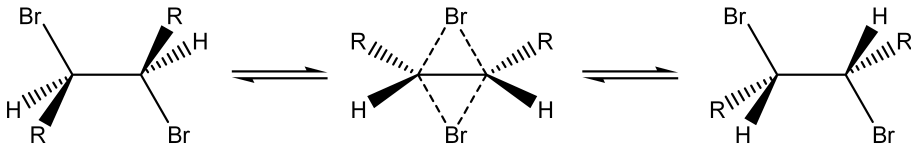

(1R,2R) (1S,2S)

Figure 6.4: Interconversion mechanism for HBCD. The interconversion is only possible, if the two bromine atoms are in anti-position. This process is largely independent from the structure of the rest R.

in Figure 6.3. The chemical mechanism of the transition between these structures is shown in Figure 6.4. Two vicinal bromine atoms change their positions with two hydrogen atoms under inversion of absolute configurations. This quantum mechanical process is largely independent from the structure of the rest R and can only take place, if the two bromine atoms are in anti-position. Anti-position means that the dihedral angle θ spanned by the bromine atoms and the connecting carbon atoms lies in the interval $[-\pi, -2\pi/3] \cup [2\pi/3, \pi]$, in contrast to the gauche-position where $\theta \in [-2\pi/3, 2\pi/3]$. Classical simulations can be used to characterize the interconversion process qualitatively by determining the part of the configurational space for which the dihedral angle is in anti-position compared to the part of the configurational space where the angle is in gauche-position. The more the anti-position is preferred, the faster the conversion at the corresponding angle will occur. The distribution of conformers is different for each diastereomer, but for the purpose of demonstration we will concentrate on the (+)-α-HBCD. In case of HBCD, there are three sets of vicinal bromine atoms and thus three dihedral angles (C1C2, C5C6, C9C10); compare Figures 6.3 and 6.1. Thus, we search for metastable conformations in the space spanned by these three dihedral angles. Since the interconversion of neat HBCD takes place above its melting point at about 433 K, this will be the temperature of interest. As an approximation, the simulations in this chapter are performed for the vacuum.

The question of dividing the conformational space into anti- and gauche-positions has already been addressed in [110]. Due to the existence of dynamical metastabilities inside the ring structure, the simulations were performed at high temperature (1500K) in order to avoid the trapping problem. Then, an interpolation approach was used to approximate the Boltzmann distribution at 433K. Since cooperation partners from industry were only interested in qualitative results, this approach was sufficient but an evaluation of the results is missing. With the help of adaptive sampling and hierarchical refinement we are now able to perform the simulations at 433K directly. Thus, we do not only obtain mean values, but also histograms at the desired temperature.

6.2 Simulation results

To verify the advantage of adaptive sampling and hierarchical refinement compared to the non-adaptive sampling method, we compare the results of both approaches. In the non-adaptive method, an initial discretization is generated and the horizontal sampling within every basis function is performed until the Gelman-Rubin convergence factor decreases to a certain tolerance or until a maximum number of sampling points is reached. All horizontal sampling points are then propagated once with random initial momenta to build the transition probability matrix P. In a second approach, the initial discretization is refined hierarchically until all horizontal samplings reach the tolerance for the Gelman-Rubin convergence factor. Then, the vertical sampling is performed adaptively until the number of sampling points from

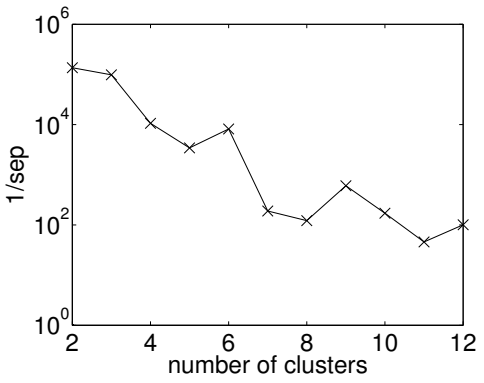

(a) Non-Adaptive approach: Condition number.

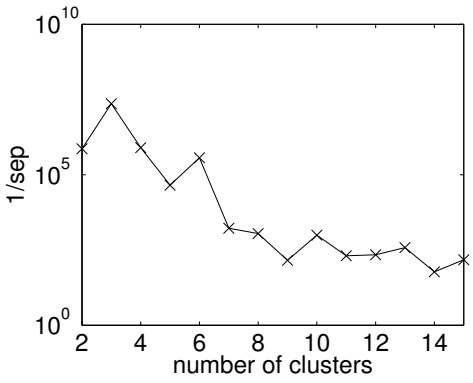

(b) Adaptive approach: Condition number.

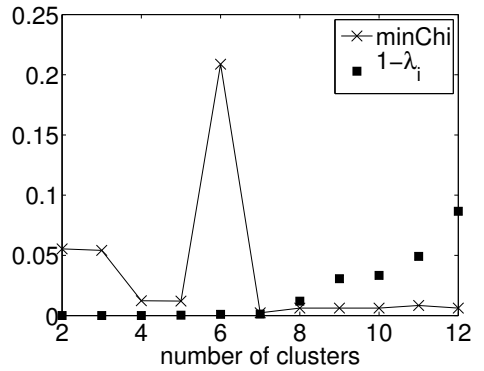

(c) Non-Adaptive approach: *minChi*-values and deviation from the Perron root.

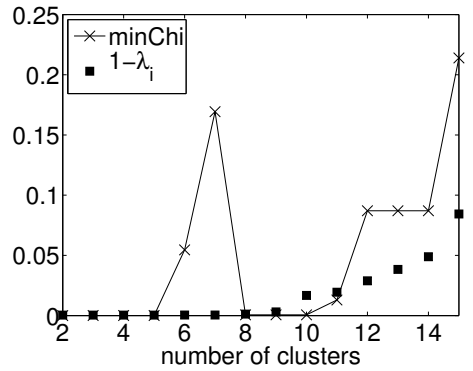

(d) Adaptive approach: *minChi*-values and deviation from the Perron root.

Figure 6.5: Condition numbers, *minChi*-values, and eigenvalues for the matrix stemming from the non-adaptive approach (figures on the left) and for the matrix resulting from initial vertical propagation after hierarchical refinement (figures on the right). Different criteria suggest different numbers of conformations.

the non-adaptive approach is reached. Finally, we compare the error distributions in the invariant subspace, the mean cluster eigenvalue, the cluster weights, and the characteristic life times of conformations. Moreover, the refined sampling results are used to examine the distribution of gauche- and anti-conformations. All HMC trajectories were generated with 60 MD steps and $\tau = 1.3\,\text{fs}$ as proposal step. The vertical sampling also consists of 60 MD steps with $\tau = 1.3\,\text{fs}$. This amounts to 78 fs, a time interval which is short enough such that the corresponding initial value problem remains well-conditioned [19].

Table 6.1: Identified conformations of HBCD by the non-adaptive approach.

conformation	C1C2	C5C6	C9C10
1	(-)-gauche	(-)-gauche	anti
2	anti	(+)-gauche	(-)-gauche
3	(+)-gauche	(+)-gauche	(+)-gauche
4	anti	(+)-gauche	(+)-gauche
5	(+)-gauche	(+)-gauche	(-)-gauche
6	(-)-gauche	(-)-gauche	(+)-gauche
7	anti	(-)-gauche	(+)-gauche
8	(-)-gauche	(+)-gauche	(+)-gauche
9	(+)-gauche	(-)-gauche	(+)-gauche
10	(-)-gauche	(-)-gauche	(-)-gauche
11	(+)-gauche	(-)-gauche	(-)-gauche

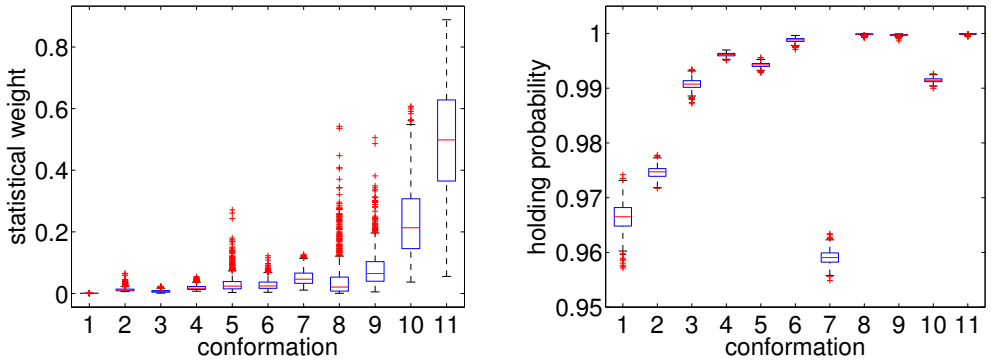

(a) Statistical weights.

(b) Holding probabilities.

Figure 6.6: Boxplots of statistical weights and holding probabilities for the 11 conformations identified by the non-adaptive approach. The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The statistics results from 1000 samples of the error matrix E.

Table 6.2: Mean and standard deviation (in parentheses) of variables computed by the non-adaptive approach.

conformation	statistical weight	holding probability	life time (ps)
1	0.0003 (0.0002)	0.9665 (0.0025)	2.342 (0.177)
2	0.0125 (0.0064)	0.9746 (0.0010)	3.079 (0.127)
3	0.0072 (0.0040)	0.9907 (0.0009)	8.493 (0.834)
4	0.0189 (0.0078)	0.9961 (0.0003)	20.282 (1.595)
5	0.0320 (0.0284)	0.9942 (0.0004)	13.621 (0.918)
6	0.0281 (0.0166)	0.9988 (0.0004)	71.337 (24.712)
7	0.0503 (0.0233)	0.9591 (0.0012)	1.909 (0.058)
8	0.0454 (0.0660)	0.9999 (0.0001)	930.5 (893.19)
9	0.0804 (0.0603)	0.9997 (0.0004)	343.9 (154.99)
10	0.2324 (0.1145)	0.9914 (0.0004)	9.137 (0.404)
11	0.4925 (0.1770)	0.9999 (0.0001)	1125.9 (627.12)

Table 6.3: Values of variables and bounds involved in the perturbation analysis of the eigenvalue problem for HBCD in the non-adaptive and the adaptive approach. Errors and condition numbers in these two approaches are of similar size.

	n_C	sep	$\|E\|_s$	global bound	$\|E_{21}\|_s$	δ
non-adapt.	8	0.0082	0.0768	0.0006	0.0050	-0.0671
non-adapt.	11	0.0219	0.0768	0.0016	0.0076	-0.0534
adapt.	9	0.0065	0.1058	0.0005	0.0059	-0.0988
adapt.	14	0.0183	0.1058	0.0010	0.0096	-0.0875

6.2.1 The non-adaptive approach

To find an initial discretization, we first generated an initial HMC pre-sampling trajectory at $T = 1500K$. From this trajectory, we selected 100 nodes located evenly in the space covered by the pre-sampling trajectory. The Voronoi cells defined by these nodes form the set of basis functions. In each basis function, 5 horizontal sampling chains were started at $T = 433K$ and continued until the Gelman-Rubin convergence factor reached the tolerance $\mathrm{TOL}_{GR} = 1.1$ or until the lengths of the chains extended to 5000 points each. In sum, 1302500 horizontal sampling points were generated. All points were propagated once to build the matrix P. The analysis of P suggests 8 or 11 metastable conformations (Figure 6.5(a)). The number of clusters should be selected such that the condition number is as small as possible. While the gap in the spectrum correlates with the condition number, the former *minChi*-criterion[1] does not give insight to the sensitivity of the subspace; see Figure 6.5(c). A larger subspace corresponds in general to a better separation, i.e. a smaller condition number, and at the same time to a larger value of $\|E_{21}\|_s$; see Table 6.3. However, these differences are mostly very small so that the choice of the number of clusters mainly depends on chemical intuition. If too few metastable conformations are selected, one might miss important structural information. Therefore, we decided for the selection of 11 metastable conformations.

As it turned out from the visualization of the corresponding configuration densities, all combinations of (+)-gauche and (-)-gauche appear as conformations, besides the combination ((-)-(+)-(-)), which is missing (Table 6.1). Thus, the state space might not have been sampled completely. In addition, three conformations contain the angle C1C2 in anti-position, and in one conformation, C9C10 is in anti-position. Although these conformations are quite unlikely, they are metastable in that the probability to stay in one of the conformations within the time interval $\tau = 78\,\mathrm{fs}$ is very large. While the accuracy of the computed holding probabilities (the diagonal entries of the matrix P_c from (1.24)) is quite high (Figure 6.6(b)), the corresponding characteristic life times are rather inexact (Table 6.2). The reason is that their computation from the holding probabilities via (1.25) is ill-conditioned. The same yields for the statistical weights of the conformations (Figure 6.6(a)). More sophisticated sampling methods are required to determine them correctly [112].

6.2.2 The adaptive approach

In this second numerical experiment, the basis functions that did not exhibit horizontal convergence in the first experiment were refined hierarchically. After 6 refinement steps, all basis functions converged. The final discretization now comprises 163 basis functions; see Table 6.4. Then the adaptive vertical sampling was performed.

Initially, 2500 horizontal sampling points in every basis function were selected for vertical propagation. An analysis of the subspace condition numbers after this initial

[1] The value $minChi = |\min_i \min_j \chi_j(i)|$ measures the deviation of the eigenvectors from the simplex structure as it occurs in the unperturbed case.

Table 6.4: Hierarchy levels in the simulation of HBCD in the adaptive approach. "AF" stands for ansatz function. The pool size refers to the total number of nodes in the hierarchy tree, whereas the leaves comprise only the nodes on the deepest hierarchy level.

level	sampled AFs	converged AFs	pool size	leaves
1	100	64	100	100
2	72	57	172	136
3	30	22	202	151
4	16	13	218	159
5	6	5	224	162
6	2	2	226	163

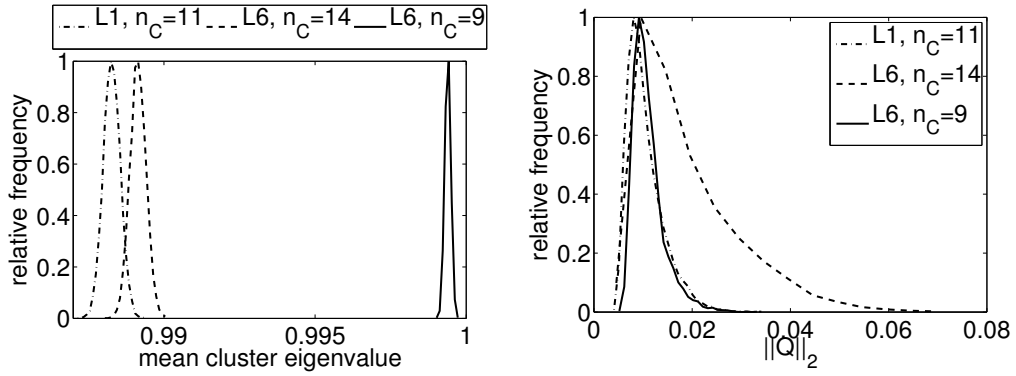

(a) Mean cluster eigenvalue.

(b) Invariant subspace.

Figure 6.7: Distribution of the mean cluster eigenvalue and the deviation of the subspace from the expected value for different discretizations (hierarchy level 1 or 6) and different numbers of conformations. The statistics have been obtained from 10000 samplings of the error matrix E.

Table 6.5: Identified conformations after hierarchical refinement in the adaptive approach.

conformation	C1C2	C5C6	C9C10
1	(-)-gauche	anti	(+)-gauche
2	(-)-gauche	(-)-gauche	anti
3	(+)-gauche	anti	(-)-gauche
4	(-)-gauche	(-)-gauche	(+)-gauche
5	(-)-gauche	(+)-gauche	(-)-gauche
6	anti	(+)-gauche	(-)-gauche
7	anti	(-)-gauche	(+)-gauche
8	anti	(+)-gauche	(+)-gauche
9	(-)-gauche	(+)-gauche	(+)-gauche
10	(-)-gauche	(-)-gauche	(-)-gauche
11	(+)-gauche	(+)-gauche	(-)-gauche
12	(+)-gauche	(+)-gauche	(+)-gauche
13	(+)-gauche	(-)-gauche	(+)-gauche
14	(+)-gauche	(-)-gauche	(-)-gauche

Table 6.6: Mean cluster weights for different numbers of identified conformations in the adaptive approach.

conformation	mean statistical weight	
	$n_C = 14$	$n_C = 9$
1	0.0001	–
2	0.0003	–
3	0.0041	–
4	0.0056	0.0021
5	0.0195	0.0044
6	0.0156	–
7	0.0174	–
8	0.0390	0.0111
9	0.0317	0.0374
10	0.0931	0.1019
11	0.2169	0.1757
12	0.1576	0.1727
13	0.1805	0.1922
14	0.2178	0.3026

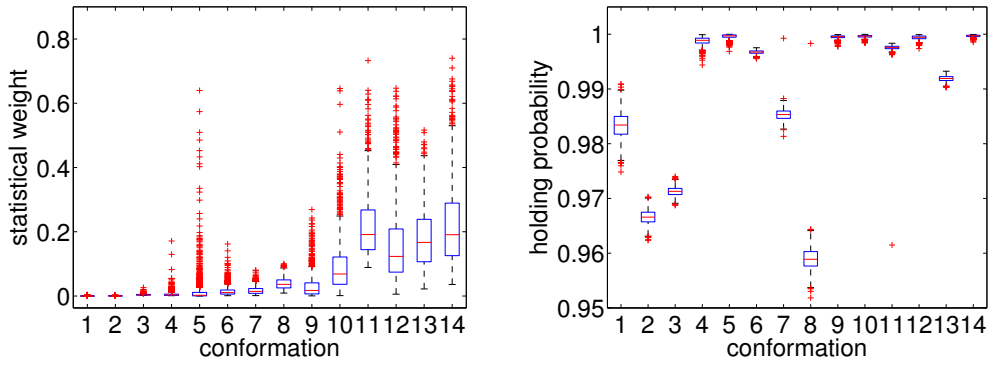

(a) Statistical weights.

(b) Holding probabilities.

Figure 6.8: Boxplots of statistical weights and holding probabilities for the 14 conformations identified in the adaptive approach. The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The statistics result from 1000 samples of the error matrix E.

Table 6.7: Mean and standard deviation (in parentheses) of variables computed in the adaptive approach.

conformation	statistical weight	holding probability	life time (ps)
1	0.0001 (0.0003)	0.9833 (0.0025)	4.76 (0.9)
2	0.0003 (0.0003)	0.9666 (0.0013)	2.34 (0.1)
3	0.0041 (0.0022)	0.9713 (0.0008)	2.72 (0.1)
4	0.0056 (0.0095)	0.9987 (0.0026)	95.96 (90.6)
5	0.0195 (0.0558)	0.9996 (0.0004)	645.41 (1340.7)
6	0.0156 (0.0161)	0.9967 (0.0003)	24.03 (2.3)
7	0.0174 (0.0122)	0.9853 (0.0010)	5.43 (3.2)
8	0.0390 (0.0176)	0.9589 (0.0024)	1.95 (1.4)
9	0.0317 (0.0388)	0.9995 (0.0003)	219.41 (145.9)
10	0.0931 (0.0840)	0.9996 (0.0002)	284.43 (217.1)
11	0.2169 (0.0989)	0.9975 (0.0012)	32.52 (4.4)
12	0.1576 (0.1154)	0.9994 (0.0004)	190.93 (154.3)
13	0.1805 (0.0928)	0.9919 (0.0013)	9.66 (0.7)
14	0.2187 (0.1235)	0.9997 (0.0002)	360.26 (249.5)

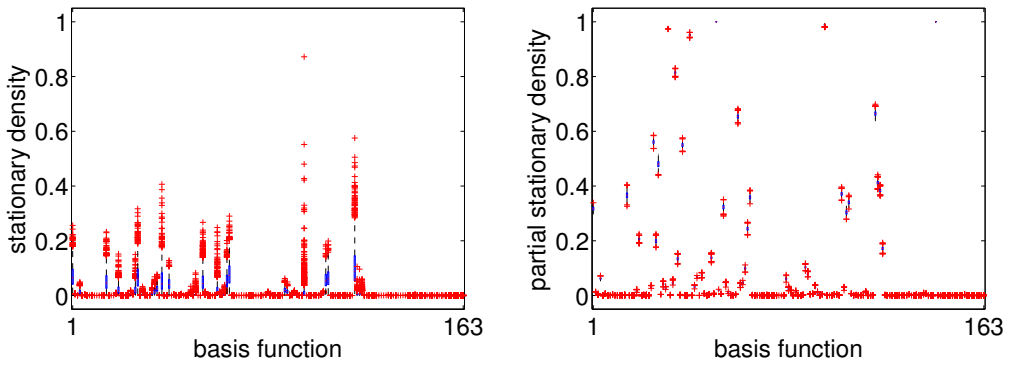

(a) Complete stationary density. (b) Partial stationary densities.

Figure 6.9: Boxplots of the stationary density computed as eigenvector of the matrix P (left) and of the partial stationary densities computed from the stochastic complements (right). The boxes have lines at the lower quartile, median, and upper quartile values. The dashed lines extending from each end of a box show the extent of the rest of the data. They extend out to the most extreme data value within 1.5 times the interquartile range of the sample. Data values beyond these lines are marked as outliers. The statistics have been obtained from 1000 samplings of the error matrix E. The overall density is ill-conditioned, resulting in a large variance, whereas the partial densities are well-conditioned.

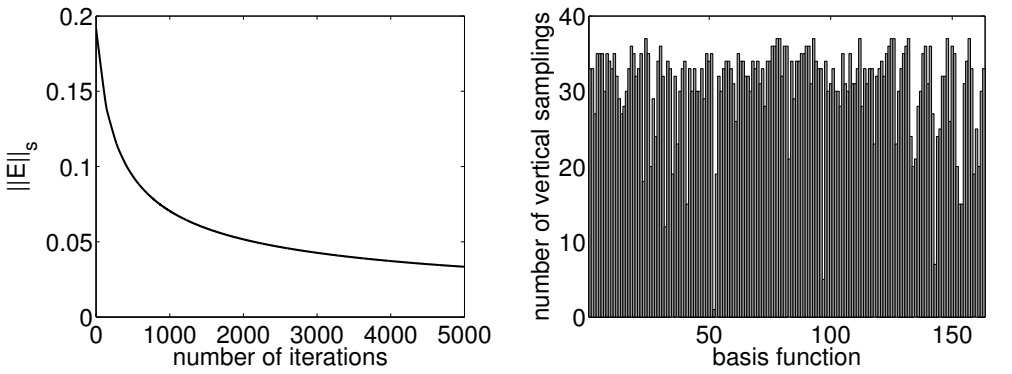

(a) Reduction of $\|E\|_s$. (b) Distribution of sampling points.

Figure 6.10: Convergence of $\|E\|_s$ during adaptive vertical sampling. The sampling points are distributed evenly among the basis functions, which hints to a sufficient discretization.

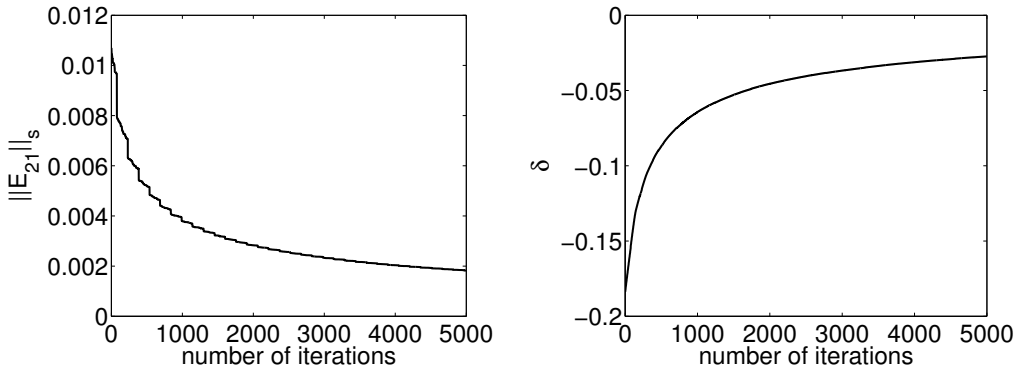

(a) Reduction of $\|E_{21}\|_s$.

(b) Convergence of δ.

Figure 6.11: Convergence of $\|E_{21}\|_s$ during adaptive vertical sampling ($n_C = 14$). Since the invariant subspace is ill-conditioned, the conditions for the validity of the perturbation bound are not satisfied ($\delta < 0$, compare Theorem 2.1.2).

sampling suggests the selection of 9 or 14 metastable conformations (Figure 6.5(b)). However, in both cases the condition number is quite large and, in particular, $\delta < 0$. Therefore, we decided to continue the vertical sampling adaptively in order to reduce $\|E\|_s$. From the basis function that was expected to reduce $\|E\|_s$ most, 2500 more sampling points were selected (with replacement) for vertical propagation. This procedure was repeated over 358 iterations such that finally the total number of vertical sampling points, $163 \cdot 2500 + 358 \cdot 2500 = 1302500$, equals the number of sampling points from the non-adaptive approach. Thus, $\|E\|_s$ could be reduced from 0.1916 to 0.1058; see Figure 6.10.

Based on the larger amount of sampling data, we now looked at the error distributions to decide for a number of conformations. If 9 conformations are selected, the mean cluster eigenvalue $\mu = 0.9994 \pm 1.9 \cdot 10^{-6}$ is nearly invariant under perturbations and the error in the corresponding invariant subspace is quite small; see Figure 6.7. In case of $n_C = 14$ conformations, the subspace error is slightly larger. However, among the 14 conformations, there are 6 conformations with one of the angles in anti-position, whereas 5 of these conformations are assigned to one of the other conformations in case of $n_C = 9$; compare Tables 6.5 and 6.6. Consequently, if one is interested in the computation of transition rates between gauche- and anti-positions, the analysis must be based on $n_C = 14$ metastable conformations.

The accuracy of statistical weights, holding probabilities and characteristic life times of conformations is given in Table 6.7. The holding probabilities are quite accurate whereas the statistical weights of the conformations as well as their characteristic life times are rather inexact; see also Figure 6.8. The reason is that these quantities are ill-conditioned. Furthermore, the errors in the overall stationary density are rather large (Figure 6.9(a)). This is due to the weak coupling between the

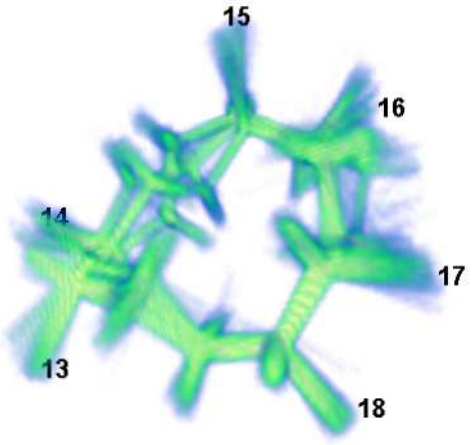

Figure 6.12: Density plot for conformation 14 from the adaptive approach.

conformations, which makes the computation of the left eigenvector of the transition probability matrix ill-conditioned. For the same reason, the partial densities are well-conditioned and can be computed with higher accuracy by stochastic complementation [70]; see Figure 6.9(b).

Extensive vertical sampling with 5000 iterations, thus increasing the number of vertical sampling points to $12.9 \cdot 10^6$, could not improve the results significantly; see Figures 6.10 and 6.11. Since the sampling points are distributed evenly among the basis functions, we do not get suggestions for further refinement. In both cases, $n_C = 9$ and $n_C = 14$, the condition number for the invariant subspace is quite large such that perturbation bounds are not valid.

Interpretation The ill-conditioning of the subspaces hints to the fact that HBCD does not have pronounced metastabilities w.r.t. the selected dihedral angles. Although it is possible to identify regions in sampling space with large holding probabilities, transitions between them are not fast processes. The typical flip-flop behavior of time series as it occurs for example for the dihedral angles of alkane molecules does not appear in such a pronounced way for the selected dihedral angles of HBCD.

Moreover, the identified conformations of HBCD exhibit low flexibility. To illustrate this behavior, Figure 6.12 shows the representative probability density of conformation 14. Besides some dislocation in the ring structure, the density has

Table 6.8: Free energy differences $\Delta_{ga}A$ between gauche- and anti-conformations of (+)-α-HBCD at $T = 433K$.

level	n_C	C1C2	C5C6	C9C10
1	11	-9.61 kJ/mol	-20.32 kJ/mol	-31.64 kJ/mol
6	14	-6.07 kJ/mol	-20.32 kJ/mol	-23.53 kJ/mol

only a small variance and the bromine atoms stay close to their mean positions over time.

Nevertheless, the adaptive as well as the non-adaptive approach lead to a common insight. The anti-configuration is less favorable then the gauche-configuration. If one of the three dihedral angles is in anti-position, it is probably the angle C1C2. Then C5C6 and C9C10 are most likely in (+)-gauche-position. On one hand, if one of the angles is in anti-position the holding probability is rather high (the conformation is metastable). On the other hand, conformations involving anti-positions are unlikely which hints to high energy barriers.

6.2.3 Evaluation of free energy differences

The mean membership basis $\overline{\chi}$ and the mean stationary density $\overline{\mathbf{w}}$ were used to compute the histograms of the three dihedral angles of interest. A comparison of the histograms in Figure 6.13 reveals that in the non-adaptive approach the (+)-gauche configurations were not explored by the horizontal trajectories, which hints to an insufficient discretization. The histograms from the adaptive approach were then used to approximate the free energy differences between gauche- and anti-configurations by

$$\Delta_{\mathrm{ga}}A(\beta) = -\frac{1}{\beta}\ln\left(\frac{N_{\mathrm{gauche}}}{N_{\mathrm{anti}}}\right),$$

where

$$\beta = \frac{1}{k_B\,T} = \frac{1}{T\cdot 0.008314403\frac{\mathrm{kJ}}{\mathrm{mol}\cdot K}} \approx 0.2778.$$

Here, N_{gauche} and N_{anti} denote the sum of histogram bin counts (bin width= $2\pi/500$) in gauche- or anti-position, respectively. The results are given in Table 6.8. From a chemical point of view, the angles C5C6 and C9C10 are symmetric such that the free energy differences should be equal. This symmetry is much better resolved by the adaptive approach that involves hierarchical refinement and adaptive sampling.

6.3 Discussion

In addition to the eigenvalue gap and the *minChi*-indicator, the condition number of the corresponding invariant subspace provides a further tool to detect the

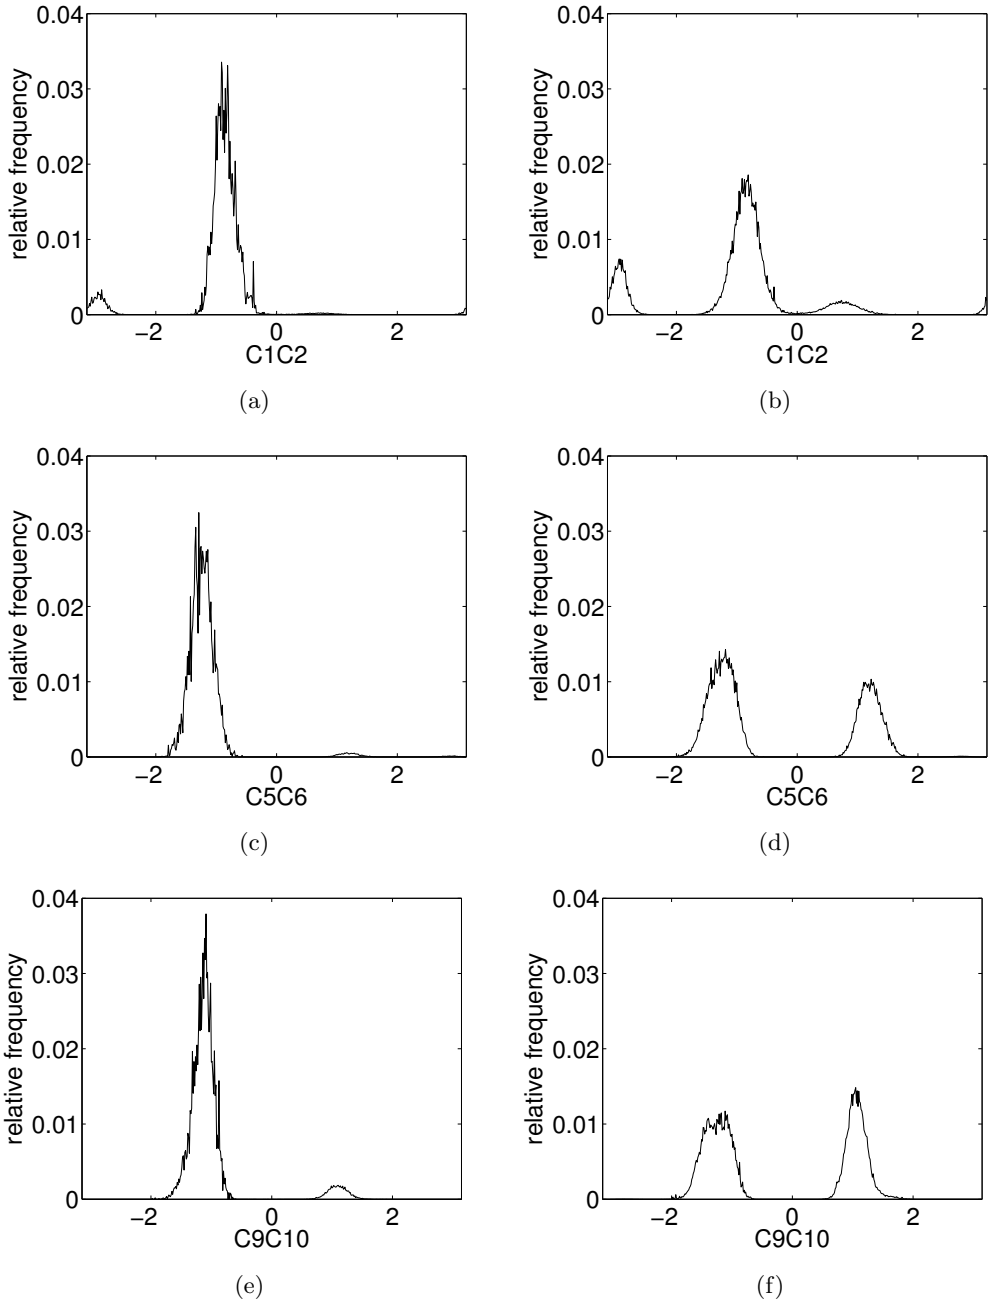

Figure 6.13: Histograms of dihedral angles C1C2, C5C6, and C9C10 (bin width= $2\pi/500$) as they result from the non-adaptive approach (left hand side) and from the adaptive approach (right hand side).

appropriate number of metastable conformations. This condition number is often very large such that an enormous sampling effort would be required to compute perturbation bounds. However, as already indicated by the low-dimensional example in previous chapters, error distributions are more appropriate and give insight into the accuracy of computed variables even for a moderate amount of sampling data. Moreover, by applying hierarchical refinement to an initial discretization, the probability of missing an important part of the conformational space decreases significantly.

The algorithm used in this thesis is optimized to find a discretization that can resolve all metastable conformations and to assign correct membership vectors to the basis functions. While the computation of observables for selected conformations requires a correct assignment of basis functions to clusters, the computation of histograms w.r.t. the complete molecular state space makes use of the stationary distribution. In the algorithm applied here, the stationary distribution has just been computed as left eigenvector of the transition probability matrix, which is ill-conditioned. Nevertheless, the resulting histograms contain all relevant conformations and reflect the symmetry between the angles C5C6 and C9C10 quite well. In order to achieve a higher accuracy, more sophisticated methods like ConfJump [108, 111] should be applied as post-processing.

As demonstrated for HBCD, it is often difficult to detect the correct number of metastable conformations. The reason is that the timescales of switching between metastable conformations and relaxation within them are not well-separated. This behavior is typical for many small biomolecules where transitions between different conformations involve the movement of massive atom groups such as aromatic ring structures. Further examples are fentanyl or epigallocatechin. However, the mathematical methods we have available now allow the conclusion whether a missing eigenvalue gap is due to insufficient discretization and missing sampling data or simply a property of the continuous system.

Conclusion

This thesis represents a further step in the development of conformation dynamics as a reliable and widely applicable method for the computation of molecular conformations. Since conformational membership vectors, as they are computed by the Robust Perron Cluster Analysis, form an invariant subspace of the discretized transfer operator, subspace-based error estimators are of particular interest. The extension of PCCA+ to non-reversible Markov chains, verified by the corresponding perturbation theory, and the modification of the objective function for the case of soft membership vectors represent a further generalization of the clustering method, thus continuing the development from PCCA over PCCA+ to "PCCA++".

The decomposition of the state space into basis functions and the approximation of integrals by Monte-Carlo quadrature give rise to row-wise correlated random matrices, for which stochastic norms are computed. Together with an appropriate statistical model for the distribution of matrix rows, this allows for the calculation of error bounds and error distributions of the invariant subspace and other variables of interest. Equilibration of errors among the basis functions can be achieved by enhanced sampling in regions where the trajectories are mixing slowly. Hierarchical refinement of such basis functions systematically improves the clustering into metastable conformations by reducing the error in the corresponding invariant subspace. In biomolecular applications, hierarchical refinement is essential for a complete exploration of the molecular state space.

The methods developed in this thesis are useful for but not limited to conformation dynamics. In fact, they are applicable to a broader class of problems which combine domain decomposition with Monte-Carlo quadrature. Possible application areas may include the chemical master equation or quantum dynamical systems [56].

Appendix A

Proofs

A.1 Proof of Theorem 3.1.3

Proof. Note that

$$\widetilde{\mathbf{p}}^\top = \widetilde{\mathbf{w}}^\top \operatorname{diag}(U, I)$$

is the Perron eigenvector of the matrix $\widetilde{B} = B + F$. Consequently,

$$\widetilde{\mathbf{p}}^\top (I - \widetilde{B}) = 0.$$

We decompose

$$\widetilde{\mathbf{p}}^\top = (\widetilde{\mathbf{p}}_1^\top, \widetilde{\mathbf{p}}_2^\top, \widetilde{\mathbf{p}}_3^\top)$$

where $\widetilde{\mathbf{p}}_1^\top = \widetilde{\mathbf{w}}_c^\top U_1 \geq 0$, $\widetilde{\mathbf{p}}_2^\top = \widetilde{\mathbf{w}}_c^\top U_2$, and $\widetilde{\mathbf{p}}_3^\top = \widetilde{\mathbf{w}}_t^\top$. Since both $\widetilde{\mathbf{w}}$ and U_1 are positive, we may renormalize $\widetilde{\mathbf{w}}$ such that $\|\widetilde{\mathbf{p}}_1\|_2 = 1$. In detail, the system reads

$$(\widetilde{\mathbf{p}}_1^\top, \widetilde{\mathbf{p}}_2^\top, \widetilde{\mathbf{p}}_3^\top) \begin{pmatrix} -F_{11} & -F_{12} & -F_{13} \\ -(B_{21}+F_{21}) & I-(B_{22}+F_{22}) & -F_{23} \\ -B_{31} & -B_{32} & I-B_{33} \end{pmatrix} = 0.$$

It follows

$$(\widetilde{\mathbf{p}}_2^\top, \widetilde{\mathbf{p}}_3^\top) \begin{pmatrix} I-\widetilde{B}_{22} & -F_{23} \\ -B_{32} & I-B_{33} \end{pmatrix} = (\mathbf{f}_{12}^\top, \mathbf{f}_{13}^\top),$$

where $\widetilde{B}_{22} = B_{22} + F_{22}$, $\mathbf{f}_{12}^\top = \widetilde{\mathbf{p}}_1^\top F_{12}$, and $\mathbf{f}_{13}^\top = \widetilde{\mathbf{p}}_1^\top F_{13}$. An equivalence transformation results in

$$(\widetilde{\mathbf{p}}_2^\top, \widetilde{\mathbf{p}}_3^\top) \begin{pmatrix} 0 & -F_{23} \\ I-B_{33}-B_{32}(I-\widetilde{B}_{22})^{-1}F_{23} & I-B_{33} \end{pmatrix} = (\mathbf{f}_{13}^\top + \mathbf{f}_{12}^\top (I-\widetilde{B}_{22})^{-1} F_{23}, \mathbf{f}_{13}^\top).$$

Consequently,

$$\widetilde{\mathbf{p}}_3^\top (I - B_{33} - B_{32}(I-\widetilde{B}_{22})^{-1}F_{23}) = \mathbf{f}_{13}^\top + \mathbf{f}_{12}^\top (I-\widetilde{B}_{22})^{-1} F_{23}.$$

By taking norms we obtain

$$\|\mathbf{f}_{13}^\top + \mathbf{f}_{12}^\top (I-\widetilde{B}_{22})^{-1} F_{23}\| \leq \eta + \eta^2 \widetilde{\gamma}_2,$$

$$\|(I - B_{33} - B_{32}(I - \widetilde{B}_{22})^{-1}F_{23})^{-1}\| \leq \frac{\gamma_3}{1 - \gamma_3\beta\widetilde{\gamma}_2\eta},$$

and thus

$$\|\widetilde{\mathbf{w}}_t^\top\| = \|\widetilde{\mathbf{p}}_3^\top\| \leq \frac{\gamma_3(\eta + \eta^2\widetilde{\gamma}_2)}{1 - \gamma_3\eta\beta\widetilde{\gamma}_2}.$$

This completes the proof. □

A.2 Proof of Theorem 3.1.4

Proof. The proof is based on the entry-wise perturbation bound given by O'Cinneide ([83], Thm.1). He shows that if the off-diagonal entries of an irreducible generator G of order n and its perturbation $\widetilde{G}$ satisfy

$$K_L G(i,j) \leq \widetilde{G}(i,j) \leq K_u G(i,j), \quad i \neq j,$$

where $0 < K_L \leq 1 \leq K_U$, then the steady-state vector $\mathbf{s}$ satisfies

$$\left(\frac{K_L}{K_U}\right)^n \mathbf{s}(j) \leq \widetilde{\mathbf{s}}(j) \leq \left(\frac{K_U}{K_L}\right)^n \mathbf{s}(j), \quad j = 1, \ldots, n.$$

We adopt his theorem in the following way. The partial density $\mathbf{s}_k$ satisfies

$$\mathbf{s}_k^\top(P_{kk} - I) = 0, \quad k = 1, \ldots, n_C,$$

where $P_{kk} = \widetilde{P}_{kk} + \mathrm{diag}(\widetilde{P}_{k\star}\mathbf{e})$ is the kth block in P_{cc}, and $\widetilde{P}_{k\star}$ is the kth row of blocks with $\widetilde{P}_{kk}$ removed. The matrix $(P_{kk} - I)$ takes the role of G in the theorem of O'Cinneide. The partial density $\widetilde{\mathbf{s}}_i$ can be obtained as steady state vector of the stochastic complement [70]

$$\widetilde{S}_{kk} = \widetilde{P}_{kk} + \widetilde{P}_{k\star}(I - \widetilde{P}_k)^{-1}\widetilde{P}_{\star k}.$$

Here again, $\widetilde{P}_{kk}$ is the kth block in the nearly block-diagonal matrix $\widetilde{P}_{cc}$. Moreover, $\widetilde{P}_k$ denotes the principal submatrix of $\widetilde{P}$ obtained by deleting the kth row and kth column of blocks from $\widetilde{P}$, and $\widetilde{P}_{\star k}$ is the kth column of blocks with $\widetilde{P}_{kk}$ removed. We have

$$\widetilde{\mathbf{s}}_k^\top(\widetilde{S}_{kk} - I) = 0,$$

and thus $\widetilde{S}_{kk} - I$ takes the role of $\widetilde{G}$. Now let us bound the entry-wise error.

$$\begin{aligned}\frac{\widetilde{G}(i,j)}{G(i,j)} &= \frac{\widetilde{S}_{kk}(i,j)}{P_{kk}(i,j)}, \quad i,j = 1, \ldots, N_k, \quad i \neq j, \quad k = 1, \ldots, n_C \\ &= \frac{\widetilde{P}_{kk}(i,j) + (\widetilde{P}_{k\star}(I - \widetilde{P}_k)^{-1}\widetilde{P}_{\star k})(i,j)}{\widetilde{P}_{kk}(i,j)}\end{aligned}$$

Since $\widetilde{P}_{k\star}(I - \widetilde{P}_k)^{-1}\widetilde{P}_{\star k}$ is non-negative ([70], Thm.2.1), we have

$$\frac{\widetilde{G}(i,j)}{G(i,j)} \geq 1 \equiv K_L.$$

Moreover [70]

$$\widetilde{P}_{k\star}(I-\widetilde{P}_k)^{-1}\widetilde{P}_{\star k}\mathbf{e}=\widetilde{P}_{k\star}\mathbf{e}=1-\widetilde{P}_{kk}\mathbf{e}\leq\varepsilon_t.$$

Consequently, $(\widetilde{P}_{k\star}(I-\widetilde{P}_k)^{-1}\widetilde{P}_{\star k})(i,j)\leq\varepsilon_t$ and

$$\frac{\widetilde{G}(i,j)}{G(i,j)}\leq\frac{\widetilde{P}_{kk}(i,j)+\varepsilon_t}{\widetilde{P}_{kk}(i,j)}\leq 1+\frac{\varepsilon_t}{\kappa_c}\equiv K_U.$$

Thus we obtain for $j=1,\ldots,N_k$, $k=1,\ldots,n_C$,

$$\frac{|\widetilde{\mathbf{s}}_k(j)-\mathbf{s}_k(j)|}{\mathbf{s}_k(j)}\leq\left(\frac{K_U}{K_L}\right)^{N_k}-1=\left(1+\frac{\varepsilon_t}{\kappa_c}\right)^{N_k}-1.$$

which yields the theorem. □

Appendix B

Sensitivity Analysis

The sensitivity analysis for right eigenvectors and eigenvalues in case of the standard eigenvalue problem has been derived in [44]. However, we believe that the equations for the generalized eigenvalue problem and the simplifications in case of the stationary density are worth to be presented.

B.1 Right eigenvectors and eigenvalues

The objective is to find $\partial \mathbf{v}_\lambda / \partial p_{ij}$ and $\partial \mathbf{v}_\lambda / \partial s_{ij}$, where $\mathbf{v}_\lambda$ is the solution of the eigenvector equation

$$\underbrace{(P - \lambda S)}_{A} \mathbf{v}_\lambda = 0.$$

We will first calculate $\partial \mathbf{v}_\lambda / \partial p_{ij}$. Differentiation of this equation w.r.t. a parameter p_{ij} gives

$$\left(\frac{\partial A}{\partial p_{ij}} + \frac{\partial A}{\partial \lambda} \frac{\partial \lambda}{\partial p_{ij}} \right) \mathbf{v}_\lambda + A \frac{\partial \mathbf{v}_\lambda}{\partial p_{ij}} = 0.$$

Using the fact that $\partial A / \partial \lambda = -S$, an evaluation of the partial derivatives at the expected values $\overline{p}_{ij}$ and $\overline{s}_{ij}$ leads to

$$\overline{A} \frac{\partial \mathbf{v}_\lambda}{\partial p_{ij}} \bigg|_{\overline{A}} = - \left(\frac{\partial A}{\partial p_{ij}} - \overline{S} \frac{\partial \lambda}{\partial p_{ij}} \right) \overline{\mathbf{v}}_\lambda. \tag{B.1}$$

The additional constraint

$$\mathbf{w}^\top \overline{\mathbf{v}}_\lambda = 1$$

uniquely determines $\overline{\mathbf{v}}$. Different normalizations are possible, see for example [44]. Differentiation leads to

$$\mathbf{w}^\top \frac{\partial \mathbf{v}}{\partial p_{ij}} = 0,$$

which, in combination with (B.1), gives

$$\begin{bmatrix} \overline{A} \\ \mathbf{w}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}_\lambda}{\partial p_{ij}} = \begin{bmatrix} -((\partial A / \partial p_{ij}) I - (\partial \lambda / \partial p_{ij}) \overline{S}) \overline{\mathbf{v}}_\lambda \\ 0 \end{bmatrix}.$$

The same system of equations results for $\partial \mathbf{v}_\lambda / \partial s_{ij}$ if p_{ij} is replaced with s_{ij}.

We have $\partial A / \partial p_{ij} = \mathbf{e}_i \mathbf{e}_j^\top$ and $\partial A / \partial s_{ij} = -\lambda \mathbf{e}_i \mathbf{e}_j^\top$. Thus it remains to calculate $\partial \lambda / \partial p_{ij}$ and $\partial \lambda / \partial s_{ij}$. As demonstrated in [44], this can be done efficiently by using the LU-factorization of A. In contrast to the calculations given there, the following lines account for the output of the corresponding MATLAB command.

The MATLAB command lu(A) delivers an LU-factorization,

$$A = LU, \tag{B.2}$$

where L is lower triangular with unit values along its diagonal, and U is upper triangular with a zero element in its last row: $u_{kk} = 0$. This zero element is due to the fact that the matrix A is singular. Its derivative w.r.t. p_{ij} and s_{ij} must therefore equal zero,

$$\frac{dl_{kk}}{dp_{ij}} = \frac{\partial l_{kk}}{\partial \lambda}\frac{\partial \lambda}{\partial p_{ij}} + \frac{\partial l_{kk}}{\partial p_{ij}} = 0, \quad \frac{dl_{kk}}{ds_{ij}} = \frac{\partial l_{kk}}{\partial \lambda}\frac{\partial \lambda}{\partial s_{ij}} + \frac{\partial l_{kk}}{\partial s_{ij}} = 0.$$

Define vectors $\mathbf{x}$ and $\mathbf{x}^a$ by

$$L^\top \mathbf{x} = \mathbf{e}_k, \quad U\mathbf{x}^a = 0, \quad \mathbf{x}^a_N = 1. \tag{B.3}$$

Differentiating (B.2) by a general parameter h gives

$$\frac{\partial A}{\partial h} = \frac{\partial L}{\partial h}U + L\frac{\partial U}{\partial h}.$$

Pre- and post-multiplying this equation by $\mathbf{x}$ and $\mathbf{x}^a$ and substituting their definitions (B.3) leads to

$$\mathbf{x}^\top \frac{\partial A}{\partial h}\mathbf{x}^a = \frac{\partial u_{kk}}{\partial h}.$$

Setting h equal to either λ or p_{ij} results in

$$\left(\mathbf{x}^\top \frac{\partial A}{\partial \lambda}\mathbf{x}^a\right)\frac{\partial \lambda}{\partial p_{ij}} + \left(\mathbf{x}^\top \frac{\partial A}{\partial p_{ij}}\mathbf{x}^a\right) = 0.$$

Since $\partial A / \partial \lambda = -S$ and $\partial A / \partial p_{ij} = \mathbf{e}_i \mathbf{e}_j^\top$, one obtains

$$\frac{\partial \lambda}{\partial p_{ij}} = \frac{\mathbf{x}_i \mathbf{x}^a_j}{\mathbf{x}^\top S \mathbf{x}^a}.$$

Similarly,

$$\frac{\partial \lambda}{\partial s_{ij}} = -\lambda \frac{\mathbf{x}_i \mathbf{x}^a_j}{\mathbf{x}^\top S \mathbf{x}^a}.$$

These expressions will be evaluated at the expected values of the parameters by replacing S by $\overline{S}$, λ by $\overline{\lambda}$, and $\mathbf{x}$ and $\mathbf{x}^a$ by $\mathbf{x}|_{\overline{A}}$ and $\mathbf{x}^a|_{\overline{A}}$, respectively. Finally, one obtains

$$\begin{bmatrix} \overline{A} \\ \mathbf{w}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}_\lambda}{\partial p_{ij}} = \begin{bmatrix} -\mathbf{e}_i (\overline{\mathbf{v}}_\lambda)_j - \dfrac{\mathbf{x}_i \mathbf{x}^a_j}{\mathbf{x}^\top \overline{S} \mathbf{x}^a} \overline{S} \overline{\mathbf{v}}_\lambda \\ 0 \end{bmatrix},$$

$$\begin{bmatrix} \overline{A} \\ \mathbf{w}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}_\lambda}{\partial s_{ij}} = \begin{bmatrix} \overline{\lambda} \mathbf{e}_i (\overline{\mathbf{v}}_\lambda)_j + \overline{\lambda} \dfrac{\mathbf{x}_i \mathbf{x}_j^a}{\mathbf{x}^\top \overline{S} \mathbf{x}^a} \overline{S} \overline{\mathbf{v}}_\lambda \\ 0 \end{bmatrix}.$$

The system of equations for $\partial \mathbf{v}_\lambda / \partial p_{ij}$ in case of the standard eigenvalue problem results by setting $S = I$, see [44].

B.2 Left eigenvectors and eigenvalues

The objective is to find $\partial \mathbf{v} / \partial p_{ij}$ and $\partial \mathbf{v} / \partial s_{ij}$, where $\mathbf{v}$ is the solution of the eigenvector equation

$$\underbrace{(P^\top - \lambda S^\top)}_{A} \mathbf{v}_\lambda = 0, \quad \mathbf{w}^\top \mathbf{v} = 1.$$

We have

$$\partial A / \partial p_{ij} = \mathbf{e}_j \mathbf{e}_i^\top, \quad \partial A / \partial s_{ij} = -\lambda \mathbf{e}_j \mathbf{e}_i^\top, \quad \frac{\partial \lambda}{\partial p_{ij}} = \frac{\mathbf{x}_j \mathbf{x}_i^a}{\mathbf{x}^\top S \mathbf{x}^a}, \quad \frac{\partial \lambda}{\partial s_{ij}} = -\lambda \frac{\mathbf{x}_j \mathbf{x}_i^a}{\mathbf{x}^\top S \mathbf{x}^a}.$$

Applying the same calculations as above results in

$$\begin{bmatrix} \overline{A} \\ \mathbf{w}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}_\lambda}{\partial p_{ij}} = \begin{bmatrix} -\mathbf{e}_j (\overline{\mathbf{v}}_\lambda)_i - \dfrac{\mathbf{x}_j \mathbf{x}_i^a}{\mathbf{x}^\top \overline{S} \mathbf{x}^a} \overline{S} \overline{\mathbf{v}}_\lambda \\ 0 \end{bmatrix},$$

$$\begin{bmatrix} \overline{A} \\ \mathbf{w}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}_\lambda}{\partial s_{ij}} = \begin{bmatrix} \overline{\lambda} \mathbf{e}_j (\overline{\mathbf{v}}_\lambda)_i + \overline{\lambda} \dfrac{\mathbf{x}_j \mathbf{x}_i^a}{\mathbf{x}^\top \overline{S} \mathbf{x}^a} \overline{S} \overline{\mathbf{v}}_\lambda \\ 0 \end{bmatrix}.$$

Sensitivity analysis of the stationary density The objective is to find $\partial \mathbf{v} / \partial p_{ij}$ and $\partial \mathbf{v} / \partial s_{ij}$, where $\mathbf{v}$ is the solution of the eigenvector equation

$$\underbrace{(P^\top - S^\top)}_{A} \mathbf{v} = 0, \quad \mathbf{e}^\top \mathbf{v} = 1.$$

Since P and S are alway stochastic matrices, $\lambda = 1$ is fixed and we have $\partial \lambda / \partial p_{ij} = \partial \lambda / \partial s_{ij} = 0$. Consequently,

$$\begin{bmatrix} \overline{A} \\ \mathbf{e}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}}{\partial p_{ij}} = \begin{bmatrix} -\mathbf{e}_j \overline{\mathbf{v}}_i \\ 0 \end{bmatrix},$$

$$\begin{bmatrix} \overline{A} \\ \mathbf{e}^\top \end{bmatrix} \cdot \frac{\partial \mathbf{v}}{\partial s_{ij}} = \begin{bmatrix} \mathbf{e}_j \overline{\mathbf{v}}_i \\ 0 \end{bmatrix}.$$

Appendix C

Computational complexity

C.1 Solving eigenvalue problems

Over the years, efficient solvers for eigenvalue problems have been developed and included into numerical software like ARPACK [62, 64]. We will briefly explain their main ideas and give rough estimates for the computational complexities.

C.1.1 Solving the standard eigenvalue problem

The MATLAB command $[V, D]$ =eigs(A,k,σ) returns a diagonal matrix D with the k eigenvalues of A closest to the scalar σ and a matrix V whose columns are the corresponding eigenvectors. The function eigs provides the reverse communication required by the Fortran library ARPACK, namely the routines DSAUPD, DSEUPD, DNAUPD, DNEUPD, ZNAUPD, and ZNEUPD [62, 64]. The underlying algorithm is the Implicitly Restarted Arnoldi Method (IRAM) [63]. This approach combines the implicitly shifted QR scheme with a k-step Arnoldi factorization to compute a few (k) eigenvalues with user-specified features. An Arnoldi factorization of length $m = k + p$ is compressed to a factorization of length k that retains the eigeninformation of interest,

$$AV_k = V_k H_k + f_k e_k^\top.$$

$V_k \in \mathbb{C}^{N\times k}$ has orthonormal columns, $V_k^H f_k = 0$, and $H_k \in \mathbb{C}^{k\times k}$ is upper Hessenberg with non-negative sub-diagonal elements. It is impossible to predict an optimal value for the number p of basis vectors. By default, the algorithm uses $p = 2k$. The computational costs (in floating point operations) for one iteration of IRAM amount to $\mathcal{O}(2N^2p + (k + p)^3)$ [3]. The extraction of eigenvalues and eigenvectors from the Arnoldi factorizations takes further time $\mathcal{O}(k^2)$ [36]. Under the assumption that the number of required iterations is much smaller than the problem size N and $n_C \ll N$, the computation of an n_C-dimensional invariant subspace takes time

$$\mathcal{O}(n_C N^2).$$

C.1.2 Solving the generalized eigenvalue problem

The MATLAB function eigs can also be applied to solve the generalized eigenvalue problem $AX = BX\Lambda$, but the matrix B is required to be symmetric. Since this is not the case in our applications, we have to use an alternative eigenvalue solver for the generalized non-hermitian eigenvalue problem. The Jacobi-Davidson style QZ algorithm [95, 31] is a generalization of the Jacobi-Davidson algorithm, in which a partial (generalized) Schur form of size k is computed. It can be interpreted as a subspace iteration variant of the QZ algorithm [36]. A MATLAB implementation has been provided by Gerard Sleijpen [94]. Each step in the iterative process requires the reduction of a projected systems to an ordered generalized Schur form, and an approximate solution of the Jacobi correction equation. The solution of the Jacobi correction equation requires the application of $(A - \tau B)^{-1}$ with shift τ. However, once an appropriate pre-conditioner has been derived, for example the LU factorization, which can be obtained in time $\mathcal{O}(2N^3/3)$ [36], the solution takes only time $\mathcal{O}(N^2)$ in every iteration. The number of times the correction equation must be solved is at least as large as the number k of desired eigenvectors. The ordered generalized Schur form can be obtained in time $\mathcal{O}(m^3)$, where m is the dimension of test and search spaces. In general, $m \leq k + 10$. Under the assumption that the number of required iterations is much smaller than the problem size N, the computation of an n_c-dimensional invariant subspace takes time

$$\mathcal{O}(2N^3/3 + n_C N^2).$$

For comparison only, the complete Schur decomposition by the overall QZ process takes time $\mathcal{O}(66N^3)$ [36].

C.1.3 Solving the Sylvester equation

The standard method to solve the Sylvester equation

$$L_1 R - R L_2 = H \tag{C.1}$$

with $L_1 \in \mathbb{R}^{n_C \times n_C}$, $L_2 \in \mathbb{R}^{(N-n_C)\times(N-n_C)}$, $H \in \mathbb{R}^{n_C \times (N-n_C)}$ is the Bartels-Stewart algorithm [8]. The main idea of this algorithm is to apply the Schur decomposition to transform (C.1) into a triangular system which can be solved efficiently by forward or backward substitution. The method takes time $\mathcal{O}(N^3)$. The effort to solve the generalized Sylvester equation (2.27) is about twice as large as for the standard equation.

C.2 Sampling from probability distributions

Algorithms for the fast and simple generation of random variables mostly assume that one has a fast way to generate normal variables. In the following, the time to sample from a normal distribution will be denoted by Q.

C.2.1 Sampling from the multivariate normal distribution

It is a well known fact from the theory of normal distributions that if one has a vector $\mathbf{y} \sim \mathrm{MVN}(\boldsymbol{\mu}, C) \in \mathbb{R}^N$, then the vector $\mathbf{y}' = F\mathbf{y} + \mathbf{b}$ is distributed as $\mathbf{y} \sim \mathrm{MVN}(F\boldsymbol{\mu} + \mathbf{b}, FCF^\top)$. Given the decomposition

$$C = \Sigma\Sigma^\top \tag{C.2}$$

and a sample $\mathbf{y} \sim \mathrm{MVN}(0, I)$, a sample $p \sim \mathrm{MVN}(\boldsymbol{\mu}, C)$ is given by

$$p = \Sigma\mathbf{y} + \boldsymbol{\mu}. \tag{C.3}$$

If C was positive definite, Σ would be the square, upper triangular Cholesky factor. If C is only positive semi-definite, Σ can be computed from an eigenvalue decomposition of C. Σ is not necessarily triangular or square in this case. Any eigenvectors whose corresponding eigenvalue is close to zero (within a small tolerance) are omitted. Thus, we obtain $\Sigma \in \mathbb{R}^{N\times\nu}$, where ν is the number of eigenvalues of C larger than the small tolerance. This method is implemented in the MATLAB function cholcov which has been used in our examples. The decomposition (C.2) is obtained with a cost of $\mathcal{O}(N^3)$ [36], but needs to be computed only once as long as the covariance matrix is fixed. Then, for every sample from the MVN distribution one has to draw ν independent samples from the standard normal distribution, which takes time $\mathcal{O}(Q\nu)$. The matrix-vector multiplication in Equation (C.3) takes an additional time $\mathcal{O}(N\nu)$. Altogether, the generation of L samples from the MVN distribution takes time $\mathcal{O}(N^3/3 + LN\nu + LQ\nu)$. If we want to sample a complete matrix $E \in \mathbb{R}^{N\times N}$ with MVN distributed rows, where the numbers of columns of the matrices $\Sigma_i \in \mathbb{R}^{N\times\nu_i}$ sum up to $\nu \equiv \sum_{i=1}^N \nu_i$, the overall sampling will take time

$$\mathcal{O}(N^4 + LN\nu + LQ\nu).$$

Remark C.2.1. Singhal and Pande [92] pointed out that the covariance matrices we are interested in (compare Table 4.1), are rank-one updates of diagonal matrices. Therefore, it is possible to implicitly calculate the matrix Σ in time $\mathcal{O}(N)$ and to perform the multiplication in Equation (C.3) in time $\mathcal{O}(N)$ by using slight modifications of the methods described by Gill *et al* [35]. We did not made use of this approach in our examples, but will apply the method for larger problem sizes in future work.

C.2.2 Sampling from the Dirichlet distribution

A method to sample a random vector $\mathbf{p} = (p_1, \ldots, p_N)$ from the N-dimensional Dirichlet distribution with parameters $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_N)$ follows from the following connection. If $Y_1, \ldots, Y_N$ are independent gamma random variables with parameters $\alpha_i > 0$, and $Y_0 = \sum_{i=1}^N Y_i$, then $p_i = Y_i/Y_0\,(i = 1, \ldots, N)$ are distributed by the Dirichlet distribution with parameters $\boldsymbol{\alpha}$ and independent of Y_0 ([24], Thm. XI.4.1). We use the gamma sampling algorithm implemented in the MATLAB function randg(). It is based on a rejection method by Marsaglia and Tsang [67]. Each

gamma sample requires in average one sample from the standard normal distribution and has a rejection constant of less than 1.052. Thus, by ignoring the time taken for mathematical functions such as exponentiation, square roots, and logarithms, the expected time to generate L Dirichlet samples is $\mathcal{O}(LNQ)$. The complete sampling of a matrix $E \in \mathbb{R}^{N \times N}$ with Dirichlet distributed rows thus takes time

$$\mathcal{O}(LN^2Q).$$

Bibliography

[1] A. Amadei, A. B. M. Linssen, and H. J. C. Berendsen. Essential dynamics of proteins. *Proteins*, 17:412–425, 1993.

[2] Amira—advanced visualization, data analysis and geometry reconstruction, user's guide and reference manual. Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Indeed–Visual Concepts GmbH and TGS Template Graphics Software Inc., 2000.

[3] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia, 2000. `http://www.cs.utk.edu/~dongarra/etemplates/`.

[4] Z. Bai, J. Demmel, and A. McKenney. On computing condition numbers for the nonsymmetric eigenproblem. *ACM Trans. Math. Softw.*, 19(2):202–223, 1993.

[5] R. B. Bapat, D. D. Olesky, and P. van den Driessche. Perron-Frobenius theory for a generalized eigenproblem. *Linear Multilinear Algebra*, 40:141–152, 1995.

[6] R. B. Bapat and T. E. S. Rhagavan. *Nonnegative Matrices and Applications*. Cambridge University Press, 1997.

[7] J. L. Barlow. Perturbation results for nearly uncoupled Markov chains with applications to iterative methods. *Numer. Math.*, 65:51–62, 1993.

[8] R. H. Bartels and G. W. Stewart. Solution of the matrix equation $AX + XB = C$. *Commun. Assoc. Comput. Machin.*, 15:820–826, 1972.

[9] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Number 31 in Texts in Applied Mathematics. Springer-Verlag New York, 1999.

[10] P. J. Courtois. *Decomposability; Queueing and Computer System Applications*. Academic Press, Orlando, Florida, 1977.

[11] A. Covaci, A. S. C. Gerecke, R. J. Law, S. Voorspoels, M. Kohler, N. V. Heeb, H. Leslie, C. R. Allchin, and J. DeBoer. Hexabromocyclododecanes (HBCD) in the environment and humans: a review. *Environ. Sci. Technol.*, 40:3680–3688, 2006.

[12] C. A. de Wit. An overview of brominated flame retardants in the environment. *Chemosphere*, 40:583–624, 2002.

[13] C. A. de Wit, M. Alaee, and D. C. G. Muir. Levels and trends of brominated flame retardants in the arctic. *Chemosphere*, 64:209–233, 2006.

[14] M. Dellnitz and A. Hohmann. A subdivision algorithm for the computation of unstable manifolds and global attractors. *Num. Math.*, 75:239–317, 1997.

[15] M. Dellnitz and O. Junge. On the approximation of complicated dynamical behavior. *SIAM J. Numer. Anal.*, 36(2):491–515, 1999. Also available as Preprint SC 96-35, Zuse Institutte Berlin, 1996.

[16] J. Demmel. Computing stable eigendecompositions of matrices. *Linear Algebra Appl.*, 79:163–193, 1986.

[17] J. Demmel and B. Kågström. Computing stable eigendecompositions of matrix pencils. *Linear Algebra Appl.*, 88/89:139–186, 1987.

[18] P. Deuflhard. From molecular dynamics to conformational dynamics in drug design. In M. Kirkilionis, S. Krömker, R. Rannacher, and F. Tomi, editors, *Trends in Nonlinear Analysis*, pages 269–287. Springer, 2003.

[19] P. Deuflhard and F. Bornemann. *Scientific Computing with Ordinary Differential Equations*. Springer, 2002.

[20] P. Deuflhard, M. Dellnitz, O. Junge, and Ch. Schütte. Computation of essential molecular dynamics by subdivision techniques. In P. Deuflhard, J. Hermans, B. Leimkuhler, A. E. Mark, S. Reich, and R. D. Skeel, editors, *Computational Molecular Dynamics: Challenges, Methods, Ideas*, volume 4 of *Lecture Notes in Computational Science and Engineering*, pages 98–115. Springer-Verlag, 1999.

[21] P. Deuflhard, W. Huisinga, A. Fischer, and Ch. Schütte. Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra Appl.*, 315:39–59, 2000.

[22] P. Deuflhard and Ch. Schütte. Molecular conformation dynamics and computational drug design. In J. M. Hill and R. Moore, editors, *Applied Mathemetics Entering the 21st Century. Proceedings ICIAM 2003*, pages 91–119, 2004. invited paper.

[23] P. Deuflhard and M. Weber. Robust Perron cluster analysis in conformation dynamics. *Linear Algebra Appl.*, 398:161–184, 2005.

[24] L. Devroye. *Non-Uniform Random Variate Generation*. Springer, New York, 1986.

[25] M. Dishon and G. Weiss. Small sample comparison of estimation methods for the beta distribution. *J. Stat. Comput. Simulation*, 11:1–11, 1980.

[26] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216–222, 1987.

[27] C. Elkan. Clustering documents with an exponential-family approximation for the Dirichlet compound multinomial distribution. In *Proceedings of the 23rd International Conference on Machine learning*, Pittsburgh, PA, 2006.

[28] S. P. Elmer, S. Park, and V. S. Pande. Foldamer dynamics expressed via Markov state models. I. Explicit solvent molecular-dynamics simulations in acetonitrile, chloroform, methanol, and water. *J. Chem. Phys.*, 123, 2005.

[29] A. Fischer. *An Uncoupling–Coupling Method for Markov Chain Monte Carlo Simulations with an Application to Biomolecules*. Doctoral thesis, Freie Universität Berlin, 2003.

[30] A. Fischer, Ch. Schütte, P. Deuflhard, and F. Cordes. Hierarchical uncoupling–coupling of metastable conformations. In T. Schlick and H. H. Gan, editors, *Computational Methods for Macromolecules: Challenges and Applications – Proc. of the 3rd Intern. Workshop on Algorithms for Macromolecular Modelling, New York, October 12-14, 2000.*, pages 235–259, New York, 2002. Springer.

[31] D. R. Fokkema, G. L. G. Sleijpen, and H. A. van der Vorst. Jacobi–Davidson style QR and QZ algorithms for the reduction of matrix pencils. *SIAM J. Sci. Comput.*, 20(1):94–125, 1998.

[32] D. Frenkel and B. Smit. *Understanding Molecular Simulation – From Algorithms to Applications*, volume 1 of *Computational Science Series*. Academic Press, 2002.

[33] A. Gara, M. A. Blumrich, D. Chen, G. L.-T. Chiu, P. Coteus, M. E. Giampapa, R. A. Haring, P. Heidelberger, D. Hoenicke, G. V. Kopcsay, T. A. Liebsch, M. Ohmacht, B. D. Steinmacher-Burow, T. Takken, and P. Vranas. Overview of the blue gene/l system architecture. *IBM J. Res. and Dev.*, 2005.

[34] A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Stat. Sci.*, 7:457–511, 1992.

[35] P. E. Gill, G. H. Golub, W. Murray, and M. A. Saunders. Methods for modifying matrix factorizations. *Math. Comput.*, 28(126):505–535, 1974.

[36] G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins University Press, 3rd edition, 1996.

[37] G. H. Golub and J. H. Wilkinson. Ill-conditioned eigensystems and the computation of the Jordan canonical form. *SIAM Rev.*, 18(4):578–619, 1976.

[38] T. Gudmundsen, C. Kenney, and A. J. Laub. Small-sample statistical estimates for the sensitivity of eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 18(4):868–886, 1997.

[39] T. A. Halgren. The representation of van der Waals (vdW) interactions in molecular mechanics force fields: potential form, combination rules, and vdW parameters. *J. Am. Chem. Soc.*, 114:7827–7843, 1992.

[40] T. A. Halgren. Merck molecular force field. *J. Comp. Chem.*, 17(I-V):490–641, 1996.

[41] J. A. Hartigan and M. A. Wong. A k-means clustering algorithm. *Adv. Appl. Stat.*, 28(1):100–108, 1979.

[42] N. V. Heeb, W. B. Schweizer, M. Kohler, and A. C. Gerecke. Structure elucidation of hexabromocyclododecanes – a class of compounds with a complex stereochemistry. *Chemosphere*, 61:65–73, 2005.

[43] F. J. Hickernell. Koksma-Hlawka inequality. *Encycl. Stat. Sci.*, 2006.

[44] N. Singhal Hinrichs and V. S. Pande. Calculation of the distribution of eigenvalues and eigenvectors in Markovian state models for molecular dynamics. *J. Chem. Phys.*, 126(244101), 2007.

[45] W. Huisinga. *Metastability of Markovian systems: A transfer operator based approach in application to molecular dynamics.* Doctoral thesis, Freie Universität Berlin, 2001.

[46] Xue J. Blockwise perturbation theory for nearly uncoupled Markov chains and its application. *Linear Algebra Appl.*, 326:173–191, 2001.

[47] Xue J. and Gao W. Blockwise perturbation theory for Markov chains. *J. Matrix Anal. Appl.*, 20(1):270–278, 1998.

[48] K. Janak, A. Covaci, S. Voorspoels, and G. Becher. Hexabromocyclododecane in marine species from the western scheldt estuary: Diastereoisomer- and enantiomer-specific accumulation. *Environ. Sci. Technol.*, 39:1987–1994, 2005.

[49] B. Kågström and P. Poromaa. Computing eigenspaces with specified eigenvalues of a regular matrix pair (A,B) and condition estimation: Theory, algorithms, and software. *Numer. Algorithms*, 12:369–407, 1996.

[50] W. Kahan. Conseving confluence curbs ill-condition. Computer Science Technical Report 6, University of California, Berkeley, 1972.

[51] W. Kahan. Spectra of nearly Hermitian matrices. *Proc. Am. Math. Soc.*, 48:11–17, 1975.

[52] T. Kato. *Perturbation Theory for Linear Operators*. Springer-Verlag, Berlin Heidelberg, 1984.

[53] M. Kijima. *Markov Processes for Stochastic Modeling*. Stochastic Modeling Series. Chapman and Hall, 1997.

[54] D. Korenblum and D. Shalloway. Macrostate data clustering. *Phys. Rev. E*, 67:056704, 2003.

[55] S. Kube and P. Deuflhard. Errata on "Robust Perron Cluster Analysis in Conformation Dynamics", December 2006. `http://www.zib.de/kube/korrektur`.

[56] S. Kube, C. Lasser, and M. Weber. Monte-Carlo sampling of Wigner functions and surface hopping quantum dynamics. ZIB-Report 07-17, Zuse Institute Berlin, 2007.

[57] S. Kube and M. Weber. Conformation kinetics as a reduced model for transition pathways. ZIB-Report 05-43, Zuse Institute Berlin, 2005.

[58] S. Kube and M. Weber. Coarse grained molecular kinetics. ZIB-Report 06-35, Zuse Institute Berlin, 2006.

[59] S. Kube and M. Weber. Identification of metastabilities in monomolecular conformation kinetics. ZIB-Report 06-01, Zuse Institute Berlin, 2006.

[60] S. Kube and M. Weber. A coarse graining method for the identification of transition rates between molecular conformations. *J. Chem. Phys.*, 126(2), January 2007.

[61] G. Laval, M. SanCristobal, and C. Chevalet. Maximum-likelihood and Markov chain Monte Carlo approaches to estimate inbreeding and effective size from allele frequency changes. *Genetics*, 164(3):1189–1204, 2003.

[62] R. B. Lehoucq, K. Maschhoff, D. C. Sorensen, and C. Yang. ARPACK software, 1996. `http://www.caam.rice.edu/software/ARPACK/`.

[63] R. B. Lehoucq and D. C. Sorensen. Deflation techniques for an implicitly restarted Arnoldi iteration. *SIAM J. Matrix Anal. Appl.*, 17(4):789–821, 1996.

[64] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM Publications, Philadelphia, 1989. `http://www.caam.rice.edu/software/ARPACK/UG/ug.html`.

[65] R. E. Madsen, D. Kauchak, and C. Elkan. Modeling word burstiness using the Dirichlet distribution. In *Proceedings of the 22nd International Conference on Machine learning*, Bonn, Germany, 2005.

[66] O. L. Mangasarian. Perron-Frobenius properties of $Ax - \lambda Bx$. *J. Math. Anal. Appl.*, 36:86–102, 1971.

[67] G. Marsaglia and W. W. Tsang. A simple method for generating gamma variables. *ACM Trans. Math. Softw.*, 26(3):363–372, 2000.

[68] V. Mehrmann, R. Nabben, and E. Virnik. Generalization of the Perron-Frobenius theory to matrix pencils. Preprint 369, DFG Research Center MATHEON, TU Berlin, 2007. Submitted to Lin. Alg. Appl.

[69] V. Mehrmann, D. D. Olesky, T. X. T. Phan, and P. van den Driesche. Relations between Perron-Frobenius results for matrix pencils. *Lin. Alg. Appl.*, 287:257–269, 1999.

[70] C. D. Meyer. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Rev.*, 31(2):240–272, 1989.

[71] C. D. Meyer. Sensitivity of the stationary distribution of a Markov chain. *SIAM J. Matrix Anal. Appl.*, 15:715–728, 1994.

[72] C. D. Meyer and G. W. Stewart. Derivatives and perturbations of eigenvectors. *J. Numer. Anal.*, 25(3):679–691, 1988.

[73] T. P. Minka. Beyound Newton's method. Technical report, Microsoft Research, 2000. `http://research.microsoft.com/~minka/papers/newton`.

[74] T. P. Minka. Estimating a Dirichlet distribution. Technical report, M.I.T., 2000. `http://research.microsoft.com/~minka/papers/dirichlet`.

[75] T. P. Minka. The Fastfit Matlab Toolbox. Efficient maximum-likelihood estimation using generalized Newton iterations, May 2004. `http://research.microsoft.com/~minka/software/fastfit/`, Version 1.2.

[76] T. P. Minka. The Lightspeed Matlab Toolbox. Efficient operations for Matlab programming, October 2006. `http://research.microsoft.com/~minka/software/lightspeed/`, Version 2.1.

[77] J. E. Mosimann. On the compound multinomial distribution, the multivariate β-distribution, and correlations among properties. *Biometrika*, 49(1):65–82, 1962.

[78] A. Narayanan. Algorithm AS 266: maximum likelihood estimation of the parameters of the Dirichlet distribution. *Appl. Statist.*, 40(2):365–374, 1991.

[79] A. Narayanan. Small sample properties of parameter estimation in the Dirichlet distribution. *Commun. Statist. Simula.*, 20(2-3):647–666, 1991.

[80] A. Narayanan. A note on parameter estimation in the multivariate beta distribution. *Comput. Math. Appl.*, 24(10):11–17, 1992.

[81] H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods.* SIAM, 1992.

[82] T. L. Noorden and J. Romme. Computing a partial generalized real Schur form using the Jacobi-Davidson method. *Numer. Linear Algebra Appl.*, 14(3):197–215, 2007.

[83] C. A. O'Cinneide. Entrywise perturbation theory and error analysis for Markov chains. *Numer. Math.*, 65:109–120, 1993.

[84] V. S. Pande, I. Baker, J. Chapman, S. P. Elmer, S. Khaliq, S. M. Larson, Y. M. Rhee, M. R. Shirts, C. D. Snow, E. J. Sorin, and B. Zagrovic. Atomistic protein folding simulations on the submillisecond time scale using worldwide distributed computing. *Biopolymers*, 68:91–109, 2003.

[85] M. Peled, R. Scharia, and D. Sondack. Thermal rearrangement of hexabromocyclododecane (HBCD). In J. R. Desmurs, B. Gerard, and M.J. Goldstein, editors, *Advances in Organobromine Chemistry*, volume 2, pages 92–99. Elsevier, Amsterdam, Netherlands, 1995.

[86] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods.* Springer Texts in Statistics. Springer-Verlag, New York, 1999.

[87] G. Ronning. Maximum likelihood estimation of Dirichlet distributions. *J. Stat. Comput. Simulation*, 32:215–221, 1989.

[88] Ch. Schütte. *Conformational Dynamics: Modelling, Theory, Algorithm, and Application to Biomolecules.* Habilitation thesis, Department of Mathematics and Computer Science, Freie Universität Berlin, 1999.

[89] Ch. Schütte, A. Fischer, W. Huisinga, and P. Deuflhard. A direct approach to conformational dynamics based on hybrid Monte Carlo. *J. Comput. Phys., Special Issue on Computational Biophysics*, 151:146–168, 1999.

[90] Ch. Schütte, W. Huisinga, and P. Deuflhard. Transfer operator approach to conformational dynamics in biomolecular systems. In B. Fiedler, editor, *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems.*, pages 191–223. Springer, 2001.

[91] R. J. Serfling. *Approximation Theorems of Mathematical Statistics.* Wiley series in probability and mathematical statistics. Wiley, 1980.

[92] N. Singhal and V. S. Pande. Error analysis and efficient sampling in Markovian state models for molecular dynamics. *J. Chem. Phys.*, 123(204909), 2005.

[93] K. Sjölander, K. Karplus, M. Brown, R. Hughey, A. Krogh, I. Mian, and D. Haussler. Dirichlet mixtures: a method for improved detection of weak but significant protein sequence homology. *Appl. Biosci.*, 12:327–345, 1996.

[94] G. L. G. Sleijpen. Jacobi-Davidson style QZ algorithm, 2002. `http://www.math.uu.nl/people/sleijpen/index.html`.

[95] G. L. G. Sleijpen, A. G. L. Booten, D. R. Fokkema, and H. A. Van der Vorst. Jacobi-Davidson type methods for generalized eigenproblems and polynomial eigenproblems. *BIT*, 36:3:595–633, 1996.

[96] G. L. G. Sleijpen and H. A. Van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.

[97] G. W. Stewart. Error bounds for approximate invariant subspaces, of closed linear operators. *SIAM J. Numer. Anal.*, 8(4):796–808, 1971.

[98] G. W. Stewart. On the sensitivity of the eigenvalue problem $Ax = \lambda Bx$. *SIAM J. Numer. Anal.*, 9(4):669–686, 1972.

[99] G. W. Stewart. Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM Rev.*, 15(4):727–764, 1973.

[100] G. W. Stewart. On the structure of nearly uncoupled Markov chains. In G. Iazeolla, P. J. Courtois, and A. Hordijk, editors, *Mathematical Computer Performance and Reliability*, pages 287–302, New York, 1984. Elsevier.

[101] G. W. Stewart. On scaled projections and pseudo-inverses. *Linear Algebra Appl.*, 112:189–194, 1989.

[102] G. W. Stewart. Stochastic perturbation theory. *SIAM Rev.*, 32(4):579–610, 1990.

[103] G. W. Stewart. On the sensitivity of nearly uncoupled Markov chains. In W. J. Stewart, editor, *Numerical Solution of Markov Chains*, pages 105–120, New York, 1991. Marcel Dekker, Inc.

[104] G. W. Stewart. On the perturbation of Markov chains with nearly transient states. *Numer. Math.*, 65:135–141, 1993.

[105] G. W. Stewart and Ji-guang Sun. *Matrix Perturbation Theory*. Computer Science and Scientific Computing. Academic Press, Boston, 1990.

[106] G. M. Torrie and J. P. Valleau. Monte Carlo study of a phase-separating liquid mixture by umbrella sampling. *J. Chem. Phys.*, 66(4):1402–1408, February 1977.

[107] C. F. van Loan. The ubiquitous Kronecker product. *J. Comput. Appl. Math.*, 123(1-2):85–100, 2000.

[108] L. Walter and M. Weber. ConfJump: a fast biomolecular sampling method which drills tunnels through high mountains. ZIB-Report 06-26, Zuse Institute Berlin, 2006.

[109] M. Weber. *Meshless Methods in Conformation Dynamics.* Doctoral thesis, Department of Mathematics and Computer Science, Freie Universität Berlin, 2006. Verlag Dr. Hut, München.

[110] M. Weber, R. Becker, R. Köppen, and V. Durmaz. Classical hybrid Monte-Carlo simulations of the interconversion of hexabromocyclododecane. ZIB-Report 07-31, Zuse Institute Berlin, 2007.

[111] M. Weber, S. Kube, A. Riemer, and A. Bujotzek. Efficient sampling of the stationary distribution of metastable dynamical systems. ZIB-Report 07-03, Zuse Institute Berlin, 2007.

[112] M. Weber, S. Kube, L. Walter, and P. Deuflhard. Stable computation of probability densities for metastable dynamical systems. *Multiscale Model. Simul.*, 6(2), 2007.

[113] M. Weber and H. Meyer. ZIBgridfree - adaptive conformation analysis with qualified support of transition states and thermodynamic weights. ZIB-Report 05-17, Zuse Institute Berlin, 2005.

[114] B. White and D. Shalloway. Practical macrostate data clustering. *Comp. Phys.*, 2007.

[115] N. Wicker, J. Muller, R. K. R. Kalathur, and O. Poch. A maximum likelihood approximation method for Dirichlet's parameter estimation. *Comput. Stat. Data Anal.*, 52:1315–1322, 2008.

[116] J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Clarendon Press, Oxford, 1965.

[117] J. H. Wilkinson. Sensitivity of eigenvalues. *Utilitas Mathematica*, 25:5–76, 1984.

Zusammenfassung

Das Verständnis von geometrischen Strukturen und dynamischen Eigenschaften molekularer Konformationen ist essentiell für die Vorhersage des Langzeitverhaltens von Molekülen. Die Identifikation metastabiler Konformationen sowie die Bestimmung von Übergangswahrscheinlichkeiten und Haltezeiten sind Bestandteil der Konformationdynamik. Dabei handelt es sich um eine Mehrskalenmethode, die auf eine reduzierte Beschreibung des Systems mittels einer stochastischen Übergangsmatrix führt. In der vorliegenden Dissertation wurde untersucht, wie man die Genauigkeit der Matrizen sowie der daraus berechneten Größen quantifizieren kann.

Im Mittelpunkt stehen dabei Fehlerschätzer für den invarianten Unterraum, da die rechten Eigenvektoren als Grundlage der Robusten Perron Cluster Analyse (PCCA+) zur Identifizierung der metastabilen Konformationen dienen. Die Zerlegung des Zustandsraumes in Basisfunktionen sowie die Approximation der Matrixeinträge mittels Monte-Carlo-Quadratur führen zu zeilenweise korrelierten Zufallsmatrizen. Mit Hilfe einer stochastischen Norm sowie einem geeigneten statistischen Modell für die Verteilung der Matrixzeilen können u.a. Fehlerschranken und -verteilungen für den invarianten Unterraum berechnet werden. Eine Equilibrierung des Fehlers zwischen den Basisfunktionen kann durch erweitertes Sampling in solchen Regionen erreicht werden, in denen die Trajektorien nur langsam mischen. Eine hierarchische Zerlegung dieser Basisfunktionen verbessert systematisch die Zerlegung in metastabile Konformationen, indem sie den Fehler im invarianten Unterraum reduziert. Diese Techniken gestatten eine Evaluierung der Simulationsergebnisse und ebnen den Weg zur Behandlung komplexerer Moleküle.

Desweiteren wurden Verallgemeinerungen der PCCA+ untersucht. Die Erweiterung der PCCA+ auf nicht-reversible Markov-Ketten sowie die Modifizierung der Zielfunktion für den Fall der weichen Clusterung setzen die Entwicklung von der PCCA über PCCA+ zu PCCA++ fort. Somit können neue Anwendungsfelder für dieses Cluster-Verfahren erschlossen werden.

Die Methoden wurden zwar in Rahmen der Konformationsdynamik entwickelt, jedoch lassen sie sich auf eine weite Problemklasse anwenden, in der Gebietszerlegungsverfahren mit Monte-Carlo-Quadratur kombiniert werden. Mögliche Anwendungsgebiete umfassen die chemische Master-Gleichung oder quantenchemische Systeme.

Danksagung

Ich möchte mich an dieser Stelle ganz herzlich bei Herrn Prof. Deuflhard bedanken, der mir in allen Phasen meiner Arbeit hilfreich zur Seite stand. Ich bin froh, dass es solche Institute wie das ZIB gibt, die eine angewandte Forschung unter besten Arbeitsbedingungen in einer inspirierenden Atmosphäre ermöglichen. Mein besonderer Dank gilt außerdem Marcus Weber, der als Kollege und Arbeitsgruppenleiter stets Zeit für mathematische Fachgespräche fand. Aufgrund der Erfahrungen aus seiner eigenen Promotion konnte er mir viele gute Hinweise geben. Weiterhin danke ich Alexander Riemer, der wesentlich zur Software-Entwicklung von ZIBgridfree beigetragen hat und dabei ein offenes Ohr für neue mathematische Ideen hatte. Zum Schluß möchte ich mich noch ganz besonders bei meinem Mann bedanken, der mich aufmunternd durch alle Höhen und Tiefen begleitet hat.