

Optimal control of molecular dynamics using Markov state models

Christof Schütte · Stefanie Winkelmann ·
Carsten Hartmann

the date of receipt and acceptance should be inserted later

Abstract A numerical scheme for solving high-dimensional stochastic control problems on an infinite time horizon that appear relevant in the context of molecular dynamics is outlined. The scheme rests on the interpretation of the corresponding Hamilton-Jacobi-Bellman equation as a nonlinear eigenvalue problem that, using a logarithmic transformation, can be recast as a linear eigenvalue problem, for which the principal eigenvalue and its eigenfunction are sought. The latter can be computed efficiently by approximating the underlying stochastic process with a coarse-grained Markov state model for the dominant metastable sets. We illustrate our method with two numerical examples, one of which involves the task of maximizing the population of α -helices in an ensemble of small biomolecules (Alanine dipeptide), and discuss the relation to the large deviation principle of Donsker and Varadhan.

1 Introduction

Optimal control of stochastic processes is an area of active research. From a theoretical perspective the theory of viscosity solutions and the Hamilton-Jacobi-Bellman (HJB) equation is well-established [1]. The applications are predominantly in Financial Mathematics where it is used to determine optimal investment strategies for the market. On the other hand, application in materials science and molecular dynamics are rare. Although the idea of numerically solving the HJB or *dynamic programming* equation using backward iterations is rather simple, the *curse of dimensionality* is often prohibitive and restricts the fields of possible application, excluding molecular dynamics and the alike.

Available methods for the numerical solution of high dimensional HJB equations include Markov chain approximations [2], monotone schemes [3, 4],

or methods designed for relatively specific problems [5,6]. Nonetheless we are not aware of a single article on optimal control of molecular dynamics (MD) using dynamic programming principles, although the interest in controlling molecular dynamics simulations has already started more than a decade ago with the development of Targeted MD or Steered MD [7,8], and laser control of (open) quantum systems [9]; see [10] for a recent survey of approaches from molecular physics, chemistry, and optical control of quantum molecular dynamics. While most experimental techniques are used in order to learn about a certain system from its response to an external forcing, like in atomic force microscopy [11], experimental works on controlling conformational changes or molecular switches has recently attracted a lot of attention (also among theorists).

In this article we propose an approach to solving the HJB for optimal control of stochastic molecular dynamics (MD) in high dimensions. The main idea is to first approximate the dominant modes of the molecular dynamics transfer operator by a low-dimensional, so-called Markov state model (MSM), and then solve the HJB for the MSM rather than the full MD. The type of optimal control problems that we consider here, and which appear relevant in molecular dynamics applications, belong to the class of ergodic stochastic control problems. Following ideas of Fleming and co-workers [12,13], we show that the optimal control of MD on an infinite time horizon can be reformulated as a linear eigenvalue problem that has deep relations to a Donsker-Varadhan large deviation principle (LDP) for the uncontrolled MD. The general strategy that is pursued in this article, namely, using low dimensional MSM for solving high dimensional optimal control problems is illustrated in Figure 1.

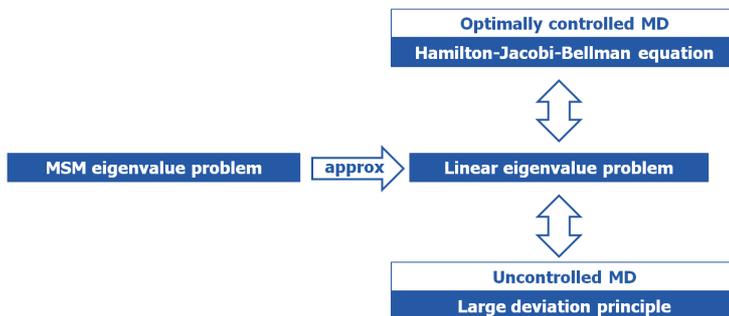


Fig. 1 Correspondence between HJB equation for controlled MD, LDP for uncontrolled MD, and the linear eigenvalue problem that is approximated by a Markov state model.

This article cannot give all the technical details. It is rather meant as a collection of the material that is required to introduce the main idea, namely, using MSMs to find approximate optimal control strategies for MD problems. Section 2 gives a short overview of stochastic molecular dynamics, Markov state models, and the role of optimal control in MD. Section 3 is devoted to ergodic control and its relation to the dominant eigenvalue of a certain elliptic

operator. How this can be turned into an optimal control algorithm using MSMs is outlined in Section 4. Section 5 briefly analyzes the connection to large deviation principles for the uncontrolled MD problem. Our findings are summarized in Section 6, including a short outlook for possible future research. The appendix contains the derivation of the HJB equation of ergodic control.

2 Molecular Dynamics and Markov State Models

Molecular dynamics models the dynamical behaviour of molecular systems, mostly in terms of the (Euclidean) coordinates of all atoms in the system. At each instance in time, the state x of the system is a vector in $\mathcal{S} \subseteq \mathbb{R}^n$ where n is three times the number N of atoms in the system. Most systems of interest contain hundreds to thousands of atoms which makes the state space \mathcal{S} huge. The key object of molecular dynamics is the gradient $-\nabla V(x)$ of a potential $V : \mathcal{S} \rightarrow \mathbb{R}$ that models the interactions between all atoms in the system. The potential V contains an enormous number of minima and saddle points (the number of minima is growing exponentially with the number of atoms). The deepest minima are located in wells that are separated by high energy barriers or extended diffusive transition regions. For the observation time scale of the conformation dynamics only the rare transitions between the minima are of interest, while on the time scale of the internal molecular vibrations these transitions are unlikely to occur, for the dynamics are confined to the energy minima. In this sense we speak of the potential V as a multiscale object.

In the following the energy landscape of potential V is assumed to be bounded below and to satisfy suitable growth conditions at infinity. Then, by the laws of statistical mechanics, the system visits all possible states according to the equilibrium distribution μ (also called the *invariant measure*) which is (via $d\mu/dx = \rho$) given by the density function $\rho(x) = \exp(-\beta V(x))/Z$ with Z being the normalization constant $Z = \int_{\mathcal{S}} \exp(-\beta V(x)) dx$ and the inverse temperature $\beta = 1/(k_b T)$ given by the temperature T of the environment of the system and the Boltzmann constant k_b . Let us assume that the main energy barriers in the system are of order 1, whereas the temperature $\varepsilon = 1/\beta$ is small compared to 1 or, more precisely, compared to the most important energy barriers. Then the wells of the energy landscape are the regions in which the dynamics spend most of their time and which are separated by transition regions of vanishing probability.

The simplest model of molecular dynamics is given by

$$dx_t = -\nabla V(x_t)dt + \sqrt{2\varepsilon} dw_t, \quad x_0 = x, \quad (1)$$

where $w_t \in \mathbb{R}^n$ denotes standard Brownian motion, $0 < \varepsilon \ll 1$ the temperature, and x_t the state of the molecular system at time $t > 0$. It is known that (under appropriate conditions on V) x_t converges to the equilibrium distribution μ in the sense that the strong law of large numbers holds, i.e.,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(x_s) ds = \int_{\mathcal{S}} f(x) d\mu(x) \quad \text{a.s.}, \quad (2)$$

for any $f \in L^1(\mu)$ and μ -almost all initial conditions $x_0 = x$.

Many molecular systems exhibit dominant metastable states, in that a typical MD trajectory will remain in the vicinity of the main wells for a long time before making a transition to another well. The wells are metastable or almost invariant subsets under the dynamics, and, as a consequence, the dynamical behaviour of the system on the longest timescales is essentially characterized by the rare transitions between these metastable sets. Metastable dynamics thus means that the system remains in metastable sets for *long* periods of time before it exits *quickly* to another metastable set; here the words "long" and "quickly" mainly state that the typical residence time has to be much longer than the typical transition time.

The most prominent toy system that displays metastability is the double-well potential shown in Figure 2 and that will serve as a paradigm throughout the text. If the energy barrier between the right (lower) well and the left one is ΔV then the rare transitions from the right to the left well happen on timescale $\exp(\Delta V/\varepsilon)$; more precisely, it follows from the Freidlin-Wentzell theory of large deviations [14] that

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log \mathbf{E}(\tau_x^\varepsilon) = \Delta V,$$

where τ_x^ε is the first exit time from the right well, given that it started at $x_0 = x$, and $\mathbf{E}(\cdot)$ is the expectation over the Brownian trajectories.¹

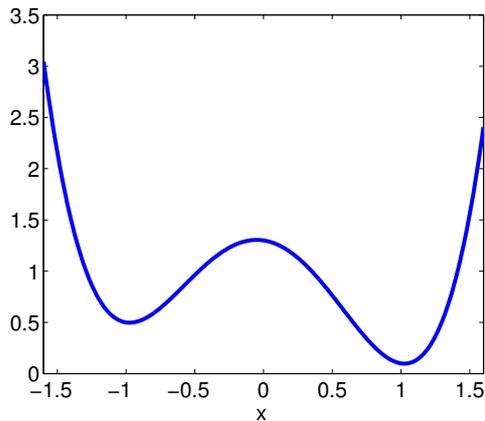


Fig. 2 The double well potential $V = V(x)$.

In the language of biophysics and biochemistry the metastable sets are called molecular conformations. Conformational transitions are essential to the function of proteins, nucleic acids and other macromolecules. These transitions

¹ It is a fundamental corollary of the Freidlin-Wentzell theory that, in the limit $\varepsilon \rightarrow 0$, the expected exit time $\mathbf{E}(\tau_x^\varepsilon)$ becomes independent of the initial value $x_0 = x$.

span large ranges of length scales, time scales and complexity, and include processes as important as folding [15, 16], complex conformational rearrangements between native protein substates [17, 18], and ligand binding [19]. MD simulations are becoming increasingly accepted as a tool to investigate both the structural and the dynamical features of these transitions at a level of detail that is beyond that accessible in laboratory experiments [20–22].

2.1 Markov State Models

Recent years have seen the advance of Markov state models (MSM) as low-dimensional models for metastable dynamical behaviour of molecular systems [23–27]. Recently the interest in MSMs has increased a lot, for it had been demonstrated that MSMs can be constructed even for very high dimensional systems [25]. They have been especially useful for modelling the interesting slow dynamics of biomolecules [21, 28–32] and materials [33] (there under the name "kinetic Monte Carlo"). If the system exhibits metastability and the jump process between the metastable sets are approximately Markovian, the corresponding MSM simply describes the Markov process that jumps between the sets with the aggregated statistics of the original process.

The key object of MSM theory is the transfer operator T_τ of the system. The transfer operator is the propagator of the system: $T_\tau f$ models the transport of the function f from time $t = 0$ to $t = \tau$ by the underlying dynamics. Mostly, the transfer operator is considered as a linear operator $T_\tau : \mathcal{H}_\mu \rightarrow \mathcal{H}_\mu$ on the weighted Hilbert space

$$\mathcal{H}_\mu = \left\{ \psi : \mathbb{R}^n \rightarrow \mathbb{R} \mid \int_{\mathbb{R}^n} |\psi|^2 d\mu < \infty \right\}, \quad (3)$$

with the weighted scalar product $\langle \psi_1, \psi_2 \rangle_\mu = \int_{\mathbb{R}^n} \psi_1 \psi_2 d\mu$ and μ denoting the invariant measure of the underlying dynamics. If the dynamics are given by (1) with $\mu \propto \exp(\beta V)$ then the transfer operator satisfies $T_\tau = \exp(\tau L)$ where L denotes the infinitesimal generator

$$L = \varepsilon \Delta - \nabla V(x) \cdot \nabla. \quad (4)$$

In the space \mathcal{H}_μ , the operators L and T_τ are known to be essentially self-adjoint which implies that their spectrum is real-valued. The spectrum of L is contained in $(-\infty, 0]$ with $\lambda = 0$ being a simple eigenvalue with eigenvector $\mathbf{1}$, the constant function. The location of the dominant metastable sets and the transition rates between these sets are encoded in the leading eigenvalues of T_τ , or L . In particular, the number of metastable sets is determined by the number of eigenvalues close to the maximum eigenvalue $\lambda = 0$ [27, 34].

MSMs are low-dimensional representations of the dynamics that approximate its longest timescales. If there are $m \ll n$ dominant metastable sets, an MSM is based on m disjoint core sets $C_1, \dots, C_m \subset \mathcal{S}$ that form the cores of the dominant metastable sets, i.e., the most attractive parts of the main wells

in the energy landscape. Given these sets, one then constructs the so-called committor functions q_1, \dots, q_m by

$$q_j(x) = \text{Prob}(\tau_x(C_j) < \tau_x(D_j)), \quad D_j = \cup_{k \neq j} C_k,$$

where $\tau_x(A)$ denote the first hitting time for the set A , assuming that x_t was initialized at $x_0 = x$. That is, the committor q_j gives the probability that the dynamics if started in x at time $t = 0$ hits the core set C_j next and before all other core sets.

By being probabilities, the committors form a partition of unity, $\sum_{j=1}^m q_j = \mathbf{1}$, and a basis of an m -dimensional subspace $D = \text{span}(q_1, \dots, q_m)$ of \mathcal{H}_μ . For instance, the Galerkin projection of the eigenvalue problem $Lu = \lambda u$ onto D reads

$$\hat{L}v = \hat{\lambda}Mv, \quad \hat{L}_{ij} = \langle q_i, Lq_j \rangle_\mu, \quad M_{ij} = \langle q_i, q_j \rangle_\mu. \quad (5)$$

The pair (\hat{L}, M) represents the MSM. For appropriately chosen core sets, the entries of $M^{-1}\hat{L}$ are known to approximate the transition rates of the underlying dynamics between the main metastable sets while its eigenvalues $0 = \hat{\lambda}_1 > \dots \geq \hat{\lambda}_m$ are close approximations of the leading eigenvalues $0 = \lambda_1 > \dots \geq \lambda_m$ of the infinitesimal generator L [32]. In this case, the MSM can be taken as a low-dimensional representation of the effective dynamics of the molecular system: instead of the complicated dynamical behaviour in some high dimensional state space we now can consider a Markov jump process on a finite state space. Since the metastable sets can be interpreted as biomolecular conformations, the entries of the MSM matrix $M^{-1}\hat{L}$ have the direct interpretation of conformational transition rates or residence times.

Remark 1 In real life molecular dynamics applications the state space is very high dimensional. In such cases the accurate computation of the committors as functions on state space is infeasible even on the most powerful computers. Yet the matrix pair (\hat{L}, M) can be accurately estimated from one long realization of x_t or from many independent short realizations [35].

Example 1 (Double well potential) For the illustration, the committor functions q_1 and q_2 for the dynamics (1) for the double well potential from Figure 2 and core sets $C_1 = [-1.1, -0.9]$ and $C_2 = [0.9, 1.1]$ are shown in Figure 3 ($\varepsilon = 1/4$). Here the infinitesimal generator has leading spectrum

$$0 = \lambda_1, \lambda_2 = -0.0186, \lambda_3 = -1.6647, \dots,$$

while the spectrum of the 2×2 matrix $M^{-1}\hat{L}$ is found to be

$$0 = \hat{\lambda}_1, \hat{\lambda}_2 = -0.0186.$$

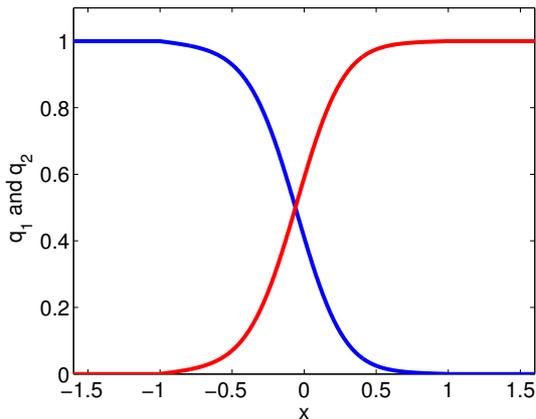


Fig. 3 Committor functions q_1 and q_2 for the dynamics (1) for the double well potential ($\beta = 4$) and core sets $C_1 = [-1.1, -0.9]$ and $C_2 = [0.9, 1.1]$.

2.2 Controlled MD

Biophysical and biochemical research aims at controlling the function of molecular systems, or even at designing molecular systems to perform certain functions. One direction of research focuses on steering the molecular system of interest from one conformation into another conformation. Mostly there is one conformation $C \subset \mathcal{S}$ with considerably higher equilibrium population² than all other conformations, and one wishes to drive the system into another conformation $B \subset \mathcal{S}$, e.g., by applying an external force as in force microscopy, or by photo or vibrational excitations as in pump-probe laser experiments. Instead of the dynamics (1) one considers the *controlled* process

$$dx_t = (u_t - \nabla V(x_t)) dt + \sqrt{2\varepsilon} dw_t, \quad x_0 = x_0, \quad (6)$$

with u_t denoting a suitable \mathbb{R}^n valued vector of control variables. The control is chosen such that the *cost function*

$$\liminf_{T \rightarrow \infty} \mathbf{E} \left(\frac{1}{T} \int_0^T f(x_s) ds \mid x_0 = x \right), \quad (7)$$

is maximized where f is any bounded function, measuring, e.g., the population of conformation B (this can be achieved by choosing $f = \mathbf{1}_B$, the indicator function of the conformation B , or f being the committor of conformation B). The expectation goes over all realizations of the process x_s starting in $x_0 = x \in C$ where x is any suitable reference point in C , e.g., the deepest minimum. To make sure that the maximum in (7) exists, the control must be

² Here the term *equilibrium population* refers to the probability $\mu(C) = \int_C \rho(x) dx$ to find the system in the conformation $C \subset \mathbb{R}^n$ when it is in thermal equilibrium, given by the stationary probability measure μ .

penalized, typically in form of a cost term $-\int_0^T |u_s|^2 ds$ that is proportional to the *work* done to the system between $t = 0$ and $t = T$, for too much external work may either alter the system irreversibly, e.g., by bond breaking, or it may be just too costly for practical purposes.

In order to consider both aspects, high output in the sense of (7) and low cost in terms of work done to the system, one can consider to solve

$$\sup \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left(\int_0^T (\sigma f(x_s) - \eta |u_s|^2) ds \mid x_0 = x \right),$$

where the maximization goes over all admissible controls u_t , and $\sigma, \eta > 0$ are (so far) arbitrary coefficients that weight the desired maximization of the output relative to the cost of the work done to the system.

3 From Ergodic Control to Eigenvalue Problems

When $u = 0$ in (6) the strong law of large numbers (2) is valid. A similar statement holds (under suitable assumptions) when $u \neq 0$ is a stationary feedback law of the form $u_t = c(x_t)$. We thus seek a control law $u = (u_t)_{t \geq 0}$ that maximizes the cost functional

$$J_x[u] = \liminf_{T \rightarrow \infty} \mathbf{E} \left(\frac{1}{T} \int_0^T g(x_s, u_s) ds \mid x_0 = x \right), \quad (8)$$

where we assume the cost function to be of the form

$$g(x, u) = \sigma f(x) - \eta |u|^2, \quad \sigma, \eta > 0. \quad (9)$$

3.1 A dynamic programming equation

The following result relates the solution of the above optimal control problem to the solution of a Hamilton-Jacobi-Bellman (HJB) equation.³

Theorem *Let $W \in C^2$ and $\lambda \in \mathbb{R}$ be the solution of*

$$\lambda = \sup_{c \in \mathbb{R}^n} \{ \varepsilon \Delta W + (c - \nabla V(x)) \cdot \nabla W + g(x, c) \}. \quad (10)$$

Further assume that

$$c^*(x) \in \operatorname{argmax}_{c \in \mathbb{R}^n} \{ \varepsilon \Delta W + (c - \nabla V(x)) \cdot \nabla W + g(x, c) \} \quad (11)$$

exists and defines an admissible Markov control by $u_t^ = c^*(x_t)$, such that*

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \mathbf{E} (W(x_t) \mid x_0 = x) = 0.$$

Then $J_x[u] \leq J_x[u^]$ among all admissible controls, where the optimal cost $J_x[u^*] = J[u^*]$ is independent of the initial value $x_0 = x$ and is equal to λ .*

³ For a brief derivations of the HJB equation and the interpretation of the value function W , we refer to Appendix A below.

Proof The proof is standard and consists basically in the application of Itô's formula to the value function W . A brief version is provided for the reader's convenience; for the technical details we refer to, e.g., [36]. Let

$$L^u = \varepsilon \Delta + (u - \nabla V(x)) \cdot \nabla$$

denote the infinitesimal generator of (6). Here the superscript indicates the explicit dependence on the control variable. We have to show that the solutions to (10) yield optimal controls that maximize (8)–(9) subject to (6).

Now choose a

$$c^*(x) \in \operatorname{argmax}_{c \in \mathbb{R}^n} \{L^c W + g(x, c)\}.$$

Obviously $u_t^* = c^*(x_t)$ is Markovian. Applying Itô's formula to W and taking expectations conditional on $x_0 = x$ yields

$$\begin{aligned} \lambda + \frac{W(x) - \mathbf{E}(W(x_t) | x_0 = x)}{t} &= \mathbf{E} \left(\frac{1}{t} \int_0^t (\lambda - L^{u_s} W(x_s)) ds \mid x_0 = x \right) \\ &\geq \mathbf{E} \left(\frac{1}{t} \int_0^t g(x_s, u_s) ds \mid x_0 = x \right), \end{aligned}$$

where the second line follows from dropping the $\sup\{\dots\}$ in (10). Taking the limit $t \rightarrow \infty$ gives

$$\lambda \geq \liminf_{t \rightarrow \infty} \mathbf{E} \left(\frac{1}{t} \int_0^t g(x_s, u_s) ds \mid x_0 = x \right).$$

But the right hand side of the inequality equals $J_x[u]$ where equality is attained for $u = u^*$. Hence the assertion is proved. \square

3.2 Formulation as an eigenvalue problem

The HJB equation (10) admits an interpretation in terms of an eigenvalue problem [37, 12]. To reveal it we first of all note that

$$\sup_{c \in \mathbb{R}^n} \{c \cdot \nabla W - \eta |c|^2\} = \frac{1}{4\eta} |\nabla W|^2$$

is unambiguously defined, so that (10) can be recast as

$$\lambda = LW + \frac{1}{4\eta} |\nabla W|^2 + \sigma f(x), \quad (12)$$

where L again denotes the infinitesimal generator of the control-free process (for $u = 0$) as given by (4). We can transform (12) into an eigenvalue equation for the function $\psi = \exp(W)$: Using chain rule, it follows that ψ solves

$$L\psi + \sigma f(x)\psi + \left(\frac{1}{4\eta} - \varepsilon \right) \frac{|\nabla \psi|^2}{\psi} = \lambda \psi,$$

where the nonlinearity can be removed by choosing $\eta = (4\varepsilon)^{-1}$. With this choice, solving the HJB equation (10) turns out to be equivalent to solving the linear elliptic eigenvalue problem

$$(L + \sigma f(x)) \psi = \lambda \psi. \quad (13)$$

By the Perron-Frobenius Theorem, the eigenfunction ψ to the maximum eigenvalue λ is positive. Hence the transformation from the eigenfunction ψ to the value function $W = \log \psi$ is well defined. Furthermore the operator $L + \sigma f$ is symmetric in the weighted Hilbert space \mathcal{H}_μ defined by (3), which implies that its spectrum in \mathcal{H}_μ is real-valued [38].

The representation of the value function $W = \log \psi_\sigma$ and its relation to the optimal control $u_t^* = c^*(x_t)$ by $c^*(x) = 2\varepsilon \nabla W(x)$ implies that the optimally controlled MD (6) takes the form

$$dx_t = -\nabla [V(x_t) - 2\varepsilon W(x_t)] dt + \sqrt{2\varepsilon} dw_t, \quad x_0 = x_0,$$

from which we see that the optimally controlled MD is again a diffusion in a transformed energy landscape

$$dx_t = -\nabla U(x_t) dt + \sqrt{2\varepsilon} dw_t, \quad x_0 = x_0, \quad (14)$$

with

$$U(x) = V(x) - 2\varepsilon W(x) = V(x) - 2\varepsilon \log \psi_\sigma(x). \quad (15)$$

This allows to compute the optimal target population

$$\pi_\sigma = \liminf_{T \rightarrow \infty} \mathbf{E} \left(\frac{1}{T} \int_0^T f(x_s^\sigma) ds \mid x_0 = x \right), \quad (16)$$

where x_t^σ denotes the optimally controlled MD, that is, the solution of (14), for a given σ . Assuming that U grows sufficiently at infinity, ergodicity implies that x_t^σ is sampling state space according to $\exp(-\beta U)$. Hence we have

$$\pi_\sigma = \frac{\langle \psi_\sigma, f \psi_\sigma \rangle_\mu}{\langle \psi_\sigma, \psi_\sigma \rangle_\mu}.$$

Remark 2 The logarithmic transformation that leads to the eigenvalue equation (13) is not unique, and for different parameter regimes it may be advantageous to use a scaled transformation. For instance, for studying the low-temperature regime $\varepsilon \ll 1$, a transformation of the form $\tilde{\psi} = \exp(W/\varepsilon)$ gives rise to a risk-sensitive control problem with the cost function [39,13]

$$\tilde{g}(x, u) = \frac{\sigma}{\varepsilon} f(x) - \frac{1}{4} |u|^2.$$

The result $\eta = (4\varepsilon)^{-1}$ of the specific logarithmic transformation considered above ought to be explained: The dominant energy barriers in the system are of order 1. The external work exerted by the control force, $\mathcal{W} = 1/2 \cdot \int_0^T |u_s|^2 ds$, should not be larger than the average kinetic energy of the system that scales

with ε ; otherwise one risks thermal dissociation of the molecular system. In order to allow the external work to overcome the dominant barriers efficiently and thus achieve order 1 improvements of the output, one needs to scale the work with ε^{-1} .

Remark 3 The operator $L + \sigma f$ is formally equivalent to a Schrödinger operator [40]. Define

$$S\varphi = \exp(-V/(2\varepsilon)) (L + \sigma f) (\varphi \exp(V/(2\varepsilon))) .$$

A straightforward computation then shows that S is a Schrödinger operator, namely,

$$S = \varepsilon\Delta + \mathcal{U}_\sigma, \quad \mathcal{U}_\sigma = \sigma f + \frac{1}{2}|\nabla V|^2 - \frac{1}{4\varepsilon}\Delta V .$$

Setting $\psi = \varphi \exp(V/(2\varepsilon))$, it can be readily seen that (13) is equivalent to

$$S\psi = \lambda\psi .$$

Example 2 (Double well potential, cont'd) For the illustration, let us return to the diffusion (1) in the double well potential shown in Figure 2. We choose the target function $f = \mathbf{1}_B$ with $B = (-\infty, -0.5)$ corresponding to the left well of the double well potential. The eigenvalue problem (13) can be solved by means of a finite element discretization to almost arbitrary accuracy. The results for the principal eigenvalue λ_σ , especially its dependence on σ are shown in Figure 4. We observe that the eigenvalue λ_σ grows almost linearly with σ for large σ (the graph displays that $\sigma^{-1}\lambda_\sigma$ tends to a constant line asymptotically) while for small σ we observe a linear regime of $\sigma^{-1}\lambda_\sigma$ that starts at a constant value $a > 0$. This behavior can be easily explained by noting that $\lambda_0 = 0$ and λ'_0 equals the expectation of the equilibrium cost (7), while the behaviour for large σ is due to the fact that the optimal cost in (8) saturates when σ is increased while ε is kept fixed. (Here the prime denotes the derivative with respect to the parameter σ ; cf. also Section 5 below)

Figure 5 shows the respective eigenvector ψ_σ as a function of x for different values of σ . The arbitrary prefactor of the eigenvector has been chosen such that $\max_x \psi_\sigma(x) = 1$.

Based on ψ_σ we can compute the transformed potential U from (15) that governs the optimally controlled MD. The potential U is shown in Figure 6. We observe that the well on the right hand side (which for $\sigma = 0$ and $U = V$ has been the deeper well of the double well potential V) is shifted upwards with increasing σ until $\sigma = 0.5$, for which the two wells become equal; for $\sigma > 0.5$, i.e., for low penalization of the control u relative to the cost function f , the left well becomes dominant and strongly attracts the dynamics.

4 Approximate Solution using MSMs

In real life molecular dynamics applications the state space is very high dimensional. Then finite element discretizations of (13) are infeasible even on

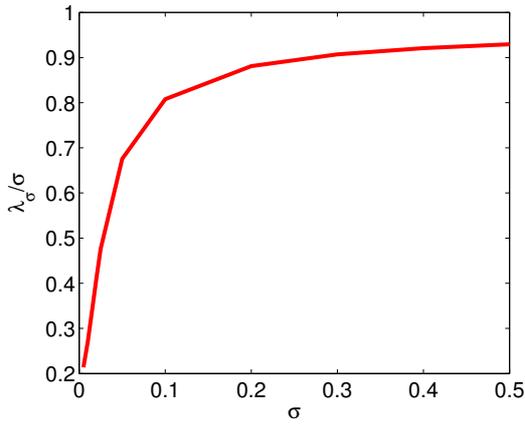


Fig. 4 The principal eigenvalue λ_σ/σ of the eigenvalue problem (13) versus σ .

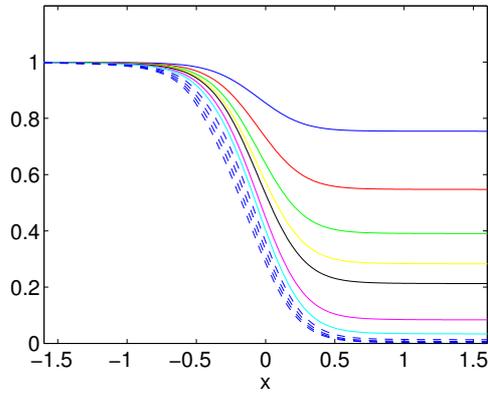


Fig. 5 Eigenvector ψ_σ of principal eigenvalue for different σ . The eigenvector is shown for the following values of σ : 0.005/0.01/0.015/0.02/0.025/0.05/0.1/0.2/0.3/0.4/0.5. The smallest value of σ belongs to the highest line on the right hand side. Increasing values of σ belong to lower and lower values of ψ on the right hand side. The highest values of σ are displayed as dashed lines.

the most powerful computers. Thus for such applications the step from the optimization problem (10) to the eigenvalue problem (13) does not seem to be a big step forward. We are in need of a method that allows to transform (13) to a feasible problem. We will now see that MSMs can do this job.

Let us assume that we have access to an MSM (M, \hat{L}) for the system with the dynamics (1) and that the MSM subspace is $D = \text{span}\{q_1, \dots, q_m\}$ with committor functions q_1, \dots, q_m . Then the eigenvalue problem $(L + \sigma f)\psi = \lambda\psi$ can be projected onto D . To this end, we represent the eigenvector ψ in the

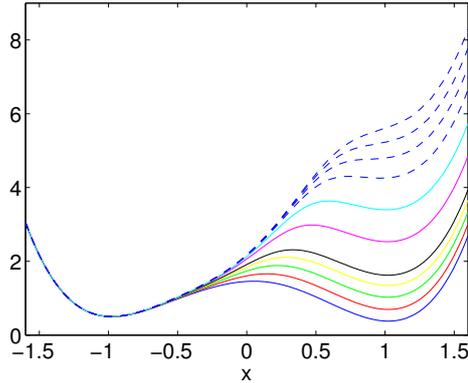


Fig. 6 The transformed potential $U(x)$ depending on σ . Same values of σ as shown in Figure 5. Again the highest values of σ are displayed as dashed lines.

basis of D , i.e., $\psi = \sum_j \alpha_j q_j$ with unknown coefficients α_i . This results in

$$(L + \sigma f)\psi = \sum_j (\alpha_j L q_j + \alpha_j f q_j) = \lambda \sum_j \alpha_j q_j.$$

Taking the scalar product with q_i $i = 1, \dots, m$ from the left yields the following generalized eigenvalue problem in m dimensions:

$$(\hat{L} + \sigma F)\alpha = \lambda M\alpha, \quad (17)$$

where $\alpha = (\alpha_1, \dots, \alpha_m)^T$ is the vector of the unknown coefficients, \hat{L} and M are the MSM matrices given in (5) and F is the $m \times m$ matrix representation of the multiplication operator f with entries

$$F_{ij} = \langle q_i, f q_j \rangle_\mu, \quad (18)$$

where $(f q_j)(x) = f(x) q_j(x)$ denotes pointwise multiplication.

As outlined in Remark 1 we typically do not have access to the committor functions in molecular dynamics applications, but to the matrix pair (M, \hat{L}) instead. Fortunately, the matrix F can also be accurately estimated from one long realization of the process (x_t) or from many independent shorter realizations without knowing the committor functions [35].

The optimal target population given in (16) can be computed based on the MSM approximation by

$$\pi_\sigma = \frac{\alpha_\sigma^T F \alpha_\sigma}{\alpha_\sigma^T M \alpha_\sigma},$$

where $\alpha_\sigma \in \mathbb{R}^m$ is the eigenvector of the principal eigenvalue of (17).

Example 3 (Double well potential, cont'd) Returning to the diffusive dynamics in the double well potential already discussed above with $f = \mathbf{1}_B$, $B = (-\infty, -0.5)$, corresponding to the left well of the double well potential, we

can compute the 2×2 matrix F from (18) based on the committor function shown in Figure 3, and use our knowledge of the MSM (M, \hat{L}) to solve the 2-dimensional eigenvalue problem (17). Figure 7 shows the ratio $\hat{\lambda}_\sigma/\sigma$ as a function of σ in comparison with the reference ratio λ_σ/σ that has been computed using a very accurate FEM discretization of the original eigenvalue problem (13). We observe that the MSM yields a very good approximation; in particular the asymptotic behaviour for $\sigma \rightarrow 0$ and the cross-over to the saturation regime for large σ are almost perfectly reproduced.

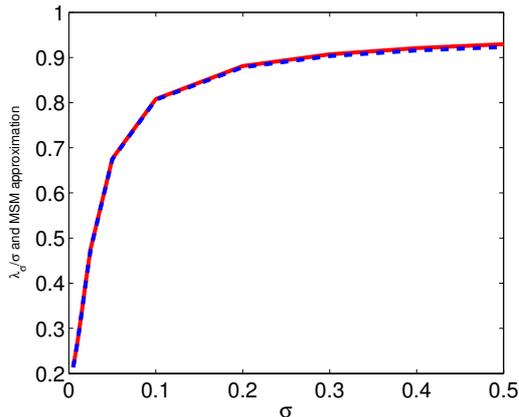


Fig. 7 The principal eigenvalue λ_σ depending on σ and its MSM approximation $\hat{\lambda}_\sigma$ as described in the text (dashed line).

Let the eigenvector corresponding to the principal eigenvalue $\hat{\lambda}_\sigma$ be denoted by $(\alpha_{1,\sigma}, \alpha_{2,\sigma})^T$. Then $\hat{\psi}_\sigma = \sum_{j=1,2} \alpha_{j,\sigma} q_j$ is the MSM approximation of the dominant eigenvector ψ_σ of (13). Figure 8 displays $\hat{\psi}_\sigma$ for different values of σ . Comparison with Figure 5 shows very good agreement with ψ_σ . The corresponding MSM approximation $\hat{U} = V - 2\varepsilon \log \hat{\psi}_\sigma$ of the transformed potential that shown in Figure 9 is in very good agreement with the reference potential in Figure 6 that is obtained from the accurate FEM solution.

In the above example the MSM approximation leads to a 2×2 eigenvalue problem. Its principal eigenvalue can be computed analytically. From this analytical form we find that $\hat{\lambda}_\sigma$ scales as follows: For small σ we find

$$\sigma^{-1} \hat{\lambda}_\sigma = a + b\sigma + \mathcal{O}(\sigma^2), \quad a = \mu((-\infty, -0.5)).$$

where $a = \mu((-\infty, -0.5)) \approx 0.171$ is the equilibrium population of the left hand well regarding the uncontrolled dynamics, and $b = \beta^2(1 + \gamma^2/4)/|\lambda_1|$ with $\beta = \langle f, u_1 \rangle_\mu$, and $\gamma = \langle u_1^2, u_1 \rangle_\mu$ with u_1 denoting the normalized first non-trivial eigenvector of L with eigenvalue λ_1 . This shows that the increase

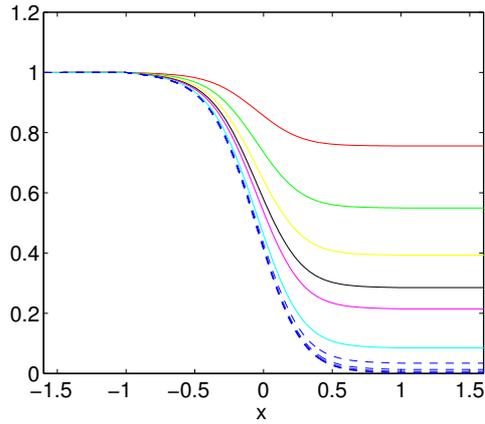


Fig. 8 MSM eigenvector $\hat{\psi}_\sigma$ of principal eigenvalue for different σ . The eigenvector is displayed for the same values of σ and in the same way as in Figure 5.

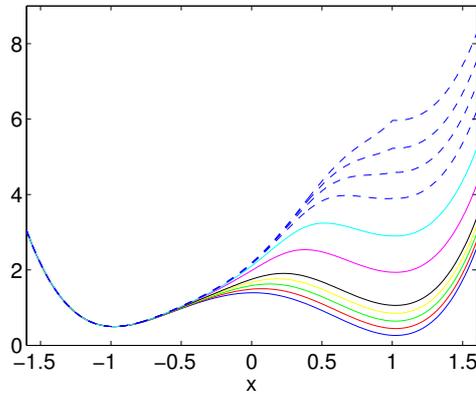


Fig. 9 The transformed potential $\hat{U} = \hat{U}(x)$ computed from the MSM described in the text.

of $\sigma^{-1}\hat{\lambda}_\sigma$ for small σ is the stronger the stronger the metastability of the system (i.e. the smaller $|\lambda_1|$). For large σ we find

$$\sigma^{-1}\hat{\lambda}_\sigma = a + \beta\left(\frac{1}{2}\gamma + (1 + \gamma^2/4)^{1/2}\right) + \mathcal{O}(\sigma^{-1}).$$

Here $\beta \approx 0.366$ and $\gamma \approx 1.640$ which yields $\sigma^{-1}\hat{\lambda}_\sigma = \kappa + \mathcal{O}(\sigma^{-1})$ with $\kappa \approx 0.945$ in good agreement with the reference FEM eigenvalue (see Figure 7).

In the calculation we have employed the explicit form of the committor functions in order to compute the transformed potential \hat{U} . Since computation of the committor functions is infeasible in real MD applications this renders the computation of \hat{U} infeasible in real life MD applications. However, this is no limitation, for in experiments only free energy differences between

conformations matter which can be computed from the MSM approximations directly without computing \hat{U} , cf. [18].

5 Deviations from ergodicity

The principal eigenvalue of (13) bears a deep relation with the cumulant generating function of the (time-integrated) cost function f and entails a large deviations principle for deviations from the ergodic limit of the cost function.

5.1 Relation to cumulant-generating functions

Let us for the last time return to the uncontrolled process x_t given by (1). We want to study the time-averaged exponential cost (cf. [41])

$$c(\sigma) = \liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{E}(\phi(y, t)), \quad (19)$$

with the function

$$\phi(y, t) = \mathbf{E} \left(\exp \left(\sigma \int_0^t f(x_s) ds \right) \middle| x_0 = y \right). \quad (20)$$

If we keep $t > 0$ fixed and notice that $\log \mathbf{E}(\exp(\sigma F))$ is the cumulant-generating function (CGF) of the random variable

$$F_t = \int_0^t f(x_s) ds,$$

it readily follows that

$$c(\sigma) = \liminf_{t \rightarrow \infty} \frac{\sigma}{t} \mathbf{E}(F_t) + \liminf_{t \rightarrow \infty} \frac{\sigma^2}{2t} \mathbf{E}((F_t - \mathbf{E}(F_t))^2) + \dots,$$

which, by ergodicity, implies that

$$c'(0) = \int_S f(x) d\mu(x).$$

Assuming that the limit in (19) exists, we can conclude that $\phi(y, t)$ grows exponentially at rate c as $t \rightarrow \infty$. That is, we have $\phi(y, t) \asymp \psi(y) \exp(ct)$ where, by the Feynman-Kac theorem, ϕ solves the linear evolution equation

$$\begin{aligned} \left(\frac{\partial}{\partial t} - L \right) \phi &= \sigma f(y) \phi \\ \phi(y, 0) &= 1. \end{aligned}$$

If we insert the asymptotic ansatz for ϕ in the last equation, we find that ψ is the eigenvector corresponding to the maximum eigenvalue of $L + \sigma f(x)$, i.e., the one that dominates the growth of $\phi(\cdot, t)$ for large t . More precisely,

$$(L + \sigma f(y)) \psi = c\psi,$$

which is readily seen to be equivalent to (13) with $\lambda = c$.

Remark 4 The specific structure of the eigenvalue problem (13) implies a large deviations principle that is due to Donsker and Varadhan [42, 43]; cf. also [44, 45]. It follows from Hölder's inequality that the CGF $\lambda_\sigma = c(\sigma)$ is convex, so that its Legendre transform is well-defined. Now let

$$c_*(\omega) = \sup_{\sigma} \{\omega\sigma - c(\sigma)\}$$

denote the Legendre transform of $c(\sigma)$ and let $I \subseteq \mathbb{R}$ be a any closed subset. Then F_t satisfies the large deviations principle (LDP)

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbf{P} \left(\frac{F_t}{t} \in I \right) = - \inf_{\omega \in I} c_*(\omega).$$

with the rate function c_* , which expresses the fact that the probability distribution p_t of $\omega = F_t/t$ decays like $p_t(\omega) \asymp \exp(-tc_*(\omega))$ as $t \rightarrow \infty$. In other words, the Donsker-Varadhan LDP determines the exponentially small likelihood of observing deviations from the law of large numbers as $t \rightarrow \infty$.

Example 4 (Double well potential, closing) Since we have computed the CGF $\lambda_\sigma = c(\sigma)$ for the diffusive dynamics in the double well potential already discussed above with $f = \mathbf{1}_B$, $B = (-\infty, -0.5)$, we can also compute its Legendre transform. The result is shown in Figure 10. The convex Legendre transform has its minimum at $\omega = \mu(B) \approx 0.171$ which determines the most probable value for F_t/t when $t \rightarrow \infty$. Another useful information that we obtain from the knowledge of c_* is that if we want to get $\mu(B)$ with a probability of 95% up to 10% correct based on long term simulation of x_t , then we have to use simulations of length $T \approx 30.000$.

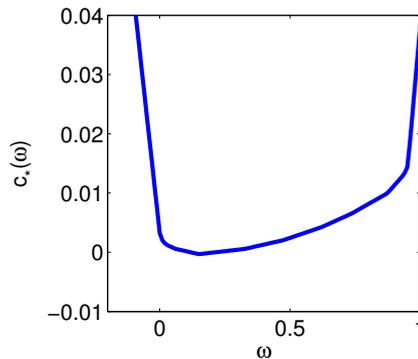


Fig. 10 Legendre transform $c_*(\omega)$ of the CGF c as introduced in the text.

5.2 Application to Alanine dipeptide

We consider the small biomolecule Alanine dipeptide (ADP), see Figure 11 (left panel). ADP with the CHARMM 27 force field was simulated in a box of 256 TIP3P water molecules. We performed a 1 μ s production run with Langevin dynamics at 300 K, using a friction constant of 5 ps^{-1} and options *rigibonds all* (i.e., all bond lengths were fixed). The conformation dynamics of the system can be monitored via the ϕ and ψ backbone dihedral angles. The corresponding invariant measure $\exp(-\beta V)/Z_V$ is shown in the right panel of Figure 11; specifically we have plotted the marginal density

$$\rho_V(\phi, \psi) = \frac{1}{Z_V} \int_{\mathcal{S}} \exp(-\beta V(x)) \delta[\Phi(x) - \phi] \delta[\Psi(x) - \psi] dx.$$

The plot shows that the density is highest at the α -helical and the β -sheet-like structures, where the fact that the conformations appear as sharp peaks in the contour plot indicates that the chosen temperature T is small compared to the most important energy barriers ΔV_{\max} of the system. More precisely, the average kinetic energy E satisfies $9k_B T = E \ll \Delta V_{\max}$. In other words, we are in a scenario with $\varepsilon \ll 1$, since ΔV_{\max} is a $\mathcal{O}(1)$ quantity, and $\varepsilon = k_B T$ in our MD model (1).

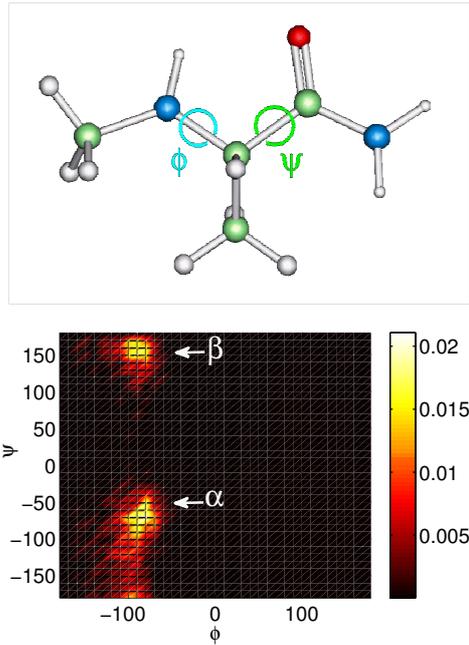


Fig. 11 Alanine dipeptide (ADP) and its invariant measure projected onto the ϕ - ψ plane, ρ_V . The metastable α and β conformations are the two indicated regions, carrying most of the invariant measure μ of the uncontrolled ADP system.

Our further considerations are based on the MSM (standard MSM with 250 states and a lag time of 500 time units) constructed in [35]. Our target function is $f = \mathbf{1}_B$ with $B = [-130, -50] \times [-100, -50]$, meaning that we want to maximize the population of the α -helix conformation while suppressing the β -sheet conformation without putting too much energy into the system. Figure 12 shows the resulting marginal invariant distribution of $\exp(-\beta U)/Z_U$,

$$\rho_U(\phi, \psi) = \frac{1}{Z_U} \int_{\mathcal{S}} \exp(-\beta U(x)) \delta[\Phi(x) - \phi] \delta[\Psi(x) - \psi] dx,$$

for optimally controlled MD computed from the eigenvalue problem (17) and the resulting target population π_σ for different σ .

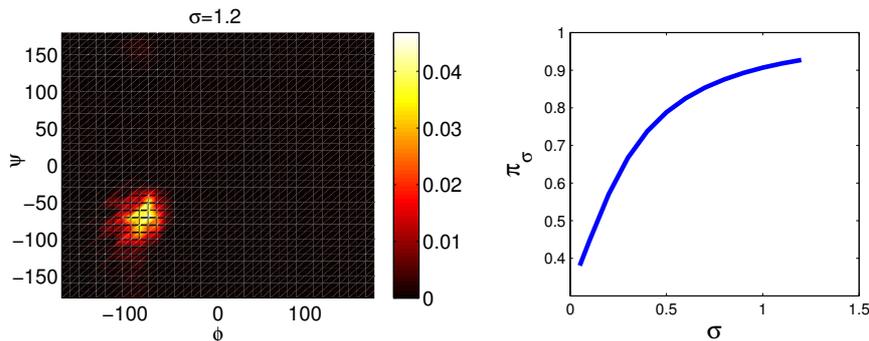


Fig. 12 Marginal invariant density ρ_U for optimally controlled MD ($\sigma = 1.2$; left) and optimal target population π_σ as a function of σ (right).

6 Summary and outlook

We have outlined a numerical scheme for solving high-dimensional stochastic control problems on an infinite time horizon that appear relevant in the context of MD. The scheme rests on the interpretation of the corresponding Hamilton-Jacobi-Bellman equation as a nonlinear eigenvalue problem that, using a logarithmic transformation, can be recast as a linear eigenvalue problem, for which principal eigenvalue and its eigenfunction are sought. We have given a proof of concept that (under certain conditions that are typically met in MD) the linear eigenvalue problem can be computed efficiently by approximating the underlying stochastic process with a coarse-grained Markov state model for the dominant metastable sets. A nice feature of the proposed scheme is that the MSM can be sampled from MD simulation data using the uncontrolled dynamics.

Our proposal is sketchy and a lot of work remains to be done. First and foremost, error bounds for the MSM approximation, along with an analysis of the backward stability of the approximation in terms of the optimal control are

still lacking. Furthermore the exact relation between the optimal cost (more precisely: its Legendre transform) and free energy of nonequilibrium steady states needs to be explored. Other open issues involve optimal control on finite time horizons or optimal stopping, both for non-degenerate (i.e., uniformly elliptic) and degenerate ergodic (i.e., hypoelliptic) diffusion processes.

Acknowledgements

The work of CS and CH is partially supported by the DFG Research Center MATHEON "Mathematics for Key Technologies" (FZT86) in Berlin.

A Hamilton-Jacobi-Bellman equation of ergodic control

The ergodic cost functional (8)–(9) is somewhat special in that the corresponding necessary optimality condition (i.e., the HJB equation) does not follow directly from a dynamic programming principle, but rather turns out to be the zero-discount limit of the following *discounted* control problem:

$$u_\alpha^* = \operatorname{argmax}_{u \in \mathcal{A}} \mathbf{E} \left(\int_0^\infty e^{-\alpha s} g(x_s, u_s) ds \right). \quad (21)$$

Here \mathcal{A} is the set of admissible Markov controls and $\alpha > 0$ denotes the discount factor that enforces convergence of the indefinite integral. The HJB equation for the discounted control problem is obtained as follows: let

$$V_\alpha(x) = \mathbf{E} \left(\int_0^\infty e^{-\alpha s} g(x_s^*, u_s^*) ds \mid x_0 = x \right)$$

be the optimal cost-to-go or value function of the discounted problem (21), for a given optimal control $u_t^* = u_{\alpha,t}^*$. Using the linearity and the tower property of the conditional expectation, it can be readily seen that

$$\begin{aligned} V_\alpha(x) &= \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} g(x_s^*, u_s^*) ds \right) + \mathbf{E}_{x_0} \left(\int_\tau^\infty e^{-\alpha s} g(x_s^*, u_s^*) ds \right) \\ &= \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} g(x_s^*, u_s^*) ds \right) + \mathbf{E}_{x_0} \left(e^{-\alpha \tau} \int_\tau^\infty e^{-\alpha(s-\tau)} g(x_s^*, u_s^*) ds \right) \\ &= \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} g(x_s^*, u_s^*) ds \right) + \mathbf{E}_{x_0} \left(e^{-\alpha \tau} \mathbf{E}_{x_\tau} \left(\int_\tau^\infty e^{-\alpha(s-\tau)} g(x_s^*, u_s^*) ds \right) \right) \\ &= \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} g(x_s^*, u_s^*) ds \right) + \mathbf{E}_{x_0} (e^{-\alpha \tau} V_\alpha(x_\tau)), \end{aligned}$$

where we have used the shorthand $\mathbf{E}_{x_0}(\cdot) = \mathbf{E}(\cdot | x_0 = x)$. Rearranging the last equation and dividing by $\tau > 0$, we obtain

$$0 = \frac{1}{\tau} \mathbf{E}_{x_0} \left(e^{-\alpha \tau} V_\alpha(x_\tau) - V_\alpha(x_0) + \int_0^\tau e^{-\alpha s} g(x_s^*, u_s^*) ds \right),$$

which upon using Itô's formula for $e^{-\alpha \tau} V_\alpha(x_\tau)$,

$$\mathbf{E}_{x_0} (e^{-\alpha \tau} V_\alpha(x_\tau) - V_\alpha(x_0)) = \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} (L^{u^*} - \alpha) V_\alpha(x_s) ds \right),$$

implies

$$0 = \frac{1}{\tau} \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} \left((L^{u^*} - \alpha)V_\alpha(x_s) + g(x_s^*, u_s^*) \right) ds \right).$$

If we relax the assumption on the control, we have

$$0 \geq \frac{1}{\tau} \mathbf{E}_{x_0} \left(\int_0^\tau e^{-\alpha s} \left((L^u - \alpha)V_\alpha(x_s) + g(x_s, u_s) \right) ds \right),$$

where equality is attained for $u = u_\alpha^*$. Letting $\tau \rightarrow 0$ formally yields

$$0 = \sup_{c \in \mathbb{R}^n} \{ (L^c - \alpha)V_\alpha + g(x, c) \}, \quad (22)$$

which is the HJB equation for the discounted control problem (21) [46].

The ergodic control functional can be treated as the limit of (22) when $\alpha \rightarrow 0$. To this end, we define $\lambda_\alpha = \alpha V_\alpha$ and $W_\alpha = V_\alpha - V_\alpha(0)$. Formally W_α satisfies the equation

$$0 = \sup_{c \in \mathbb{R}^n} \{ (L^c - \alpha)W_\alpha - \alpha V_\alpha(0) + g(x, c) \},$$

and under suitable assumptions, it can be shown that [47]

$$\lim_{\alpha \rightarrow 0} \lambda_\alpha = \lambda, \quad \lim_{\alpha \rightarrow 0} W_\alpha(x) = W(x),$$

where the pair (λ, W) solves the HJB equation

$$\lambda = \sup_{c \in \mathbb{R}^n} \{ L^c W + g(x, c) \}.$$

This is (10). Note that the reference point $\hat{x} = 0$ for which $W_\alpha(\hat{x}) = 0$ is arbitrary, since the limiting value function W is defined only up to an additive constant.

References

1. Fleming, W.H., Soner, H.M.: Controlled Markov Processes and Viscosity Solutions. Springer (2006)
2. Kushner, H.J., Dupuis, P.G.: Numerical Methods for Stochastic Control Problems in Continuous Time. Springer, New York (2001)
3. Barles, G., Souganidis, P.: Convergence of approximation schemes for fully nonlinear equations. *Asymptotic Anal.* **4**, 271–283 (1991)
4. Barles, G., Jakobsen, E.: Error bounds for monotone approximation schemes for parabolic Hamilton-Jacobi-Bellman equations. *Math. Comput.* **76**, 1861–1893 (2007)
5. Carlini, E., Falcone, M., Ferretti, R.: An efficient algorithm for Hamilton Jacobi equations in high dimension. *Comput. Visual. Sci.* **7**, 15–29 (2004)
6. Kappen, H.J.: Path integrals and symmetry breaking for optimal control theory. *J. Stat. Mech. Theor. Exp.* **2005**(11), P11,011 (2005)
7. Isralewitz, B., Gao, M., Schulten, K.: Steered molecular dynamics and mechanical functions of proteins. *Curr Opin Struct Biol* **11**, 224–230 (2001)
8. Schlitter, J., Engels, M., Krüger, P., Jacoby, E., Wollmer, A.: Targeted molecular dynamics. *Mol. Simul.* **10**, 291–309 (1993)
9. Rabitz, H., de Vivie-Riedle, R., Motzkus, M., Kompa, K.: Whither the future of controlling quantum phenomena? *Science* **288**(5467), 824–828 (2000)
10. Kühn, O., Wöste, L.: Analysis and Control of Ultrafast Photoinduced Reactions, *Series in Chemical Physics*, vol. 87. Springer (2007)
11. Engel, A., Müller, D.J.: Observing single biomolecules at work with the atomic force microscope. *Nature Structural and Molecular Biology* **7**, 715 – 718 (2000)
12. Fleming, W.H., McEneaney, W.M.: Risk-sensitive control on an infinite time horizon. *SIAM J. Control Optim.* **33**, 1881–1915 (1995)

13. Fleming, W.H., Sheu, S.J.: Asymptotics for the principal eigenvalue and eigenfunction of a nearly first-order operator with large potential. *Ann. Probab.* **25**, 1953–1994 (1997)
14. Freidlin, M., Wentzell, A.D.: *Random perturbations of dynamical systems*. Springer, New York (1998)
15. Jäger, M., Zhang, Y., Bieschke, J., Nguyen, H., Dendle, M., Bowman, M.E., Noel, J.P., Gruebele, M., Kelly, J.W.: Structure-function-folding relationship in a WW domain. *Proc. Natl. Acad. Sci. USA* **103**, 10,648–10,653 (2006)
16. Kobitski, A.Y., Nierth, A., Helm, M., Jäschke, A., Nienhaus, G.U.: Mg²⁺ dependent folding of a Diels-Alderase ribozyme probed by single-molecule FRET analysis. *Nucleic Acids Res.* **35**, 2047–2059 (2007)
17. Fischer, S., Windshuegel, B., Horak, D., Holmes, K.C., Smith, J.C.: Structural mechanism of the recovery stroke in the Myosin molecular motor. *Proc. Natl. Acad. Sci. USA* **102**, 6873–6878 (2005)
18. Noé, F., Krachtus, D., Smith, J.C., Fischer, S.: Transition Networks for the Comprehensive Characterization of Complex Conformational Change in Proteins. *J. Chem. Theo. Comp.* **2**, 840–857 (2006)
19. Ostermann, A., Waschipky, R., Parak, F.G., Nienhaus, U.G.: Ligand binding and conformational motions in myoglobin. *Nature* **404**, 205–208 (2000)
20. Schaeffer, D.D., Fersht, A., Daggett, V.: Combining experiment and simulation in protein folding: closing the gap for small model systems. *Curr. Opin. Struct. Biol.* **18**(1), 4–9 (2008). DOI 10.1016/j.sbi.2007.11.007. URL <http://dx.doi.org/10.1016/j.sbi.2007.11.007>
21. Noé, F., Schütte, C., Vanden-Eijnden, E., Reich, L., Weikl, T.R.: Constructing the full ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. USA* **106**, 19,011–19,016 (2009)
22. van Gunsteren, W., Dolenc, J., Mark, A.: Molecular simulation as an aid to experimentalists. *Curr. Opin. Struct. Biol.* **18**(2), 149–153 (2008). DOI 10.1016/j.sbi.2007.12.007. URL <http://dx.doi.org/10.1016/j.sbi.2007.12.007>
23. Schütte, C.: *Conformational dynamics: Modelling, theory, algorithm, and applications to biomolecules*. Habilitation thesis, Fachbereich Mathematik und Informatik, FU Berlin (1999)
24. Deuffhard, P., Huisinga, W., Fischer, A., Schütte, C.: Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra and its Applications* **315**, 39–59 (2000)
25. Schütte, C., Fischer, A., Huisinga, W., Deuffhard, P.: A direct approach to conformational dynamics based on hybrid Monte Carlo. *J. Comp. Physics Special Issue on Computational Biophysics* **151**, 146–168 (1999)
26. Schütte, C., Huisinga, W.: Biomolecular conformations can be identified as metastable sets of molecular dynamics. In: *Handbook of Numerical Analysis*, pp. 699–744. Elsevier (2003)
27. Bovier, A., Eckhoff, M., Gayraud, V., Klein, M.: Metastability and low lying spectra in reversible markov chains. *Comm. Math. Phys.* **228**, 219–255 (2002)
28. Noé, F., Horenko, I., Schütte, C., Smith, J.: Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states. *J. Chem. Phys.* **126**, 155,102 (2007)
29. Chodera, J., Singhal, N., Pande, V.S., Dill, K., Swope, W.: Automatic discovery of metastable states for the construction of Markov models of macromolecular conformational dynamics. *Journal of Chemical Physics* **126**, 155101 (2007)
30. Buchete, N.V., Hummer, G.: Coarse master equations for peptide folding dynamics. *Journal of Physical Chemistry B* **112**, 6057–6069 (2008)
31. Pan, A.C., Roux, B.: Building Markov state models along pathways to determine free energies and rates of transitions. *Journal of Chemical Physics* **129**, 064107 (2008)
32. Sarich, M., Noé, F., Schütte, C.: On the approximation quality of Markov state models. *Multiscale Modeling and Simulation* **8**(4), 1154–1177 (2010)
33. Voter, A.: Introduction to the kinetic Monte Carlo method. In: *Radiation Effects in Solids*. Springer, NATO Publishing Unit, Dordrecht, The Netherlands (2005)
34. Huisinga, W.: *Metastability of markovian systems a transfer operator based approach in application to molecular dynamics*. Phd thesis, Fachbereich Mathematik und Informatik, FU Berlin (2001)

35. Schütte, C., Noe, F., Lu, J., Sarich, M., Vanden-Eijnden, E.: Markov state models based on milestoneing. *J. Chem. Phys.* **134** (2011)
36. Borkar, V.S., Ghosh, M.K.: Ergodic control of multidimensional diffusions, II: Adaptive control. *Appl. Math. Optim.* **21**, 191–220 (1990)
37. Bensoussan, A., Nagai, H.: An ergodic control problem arising from the principal eigenvalue of an elliptic operator. *J. Math. Soc. Japan* **43**, 49–65 (1990)
38. Huisinga, W., Meyn, S.P., Schütte, C.: Phase transitions & metastability in Markovian and molecular systems. *Ann. Appl. Probab.* **14**, 419–458 (2004)
39. James, M.: Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Math. Control Signals Systems* **5**, 401–417 (1992)
40. Holland, C.J.: A minimum principle for the principal eigenvalue for second-order linear elliptic equations with natural boundary conditions. *Comm. Pure Appl. Math.* **31**, 509–519 (1978)
41. Whittle, P.: Risk-sensitivity, large deviations and stochastic control. *Eur. J. Oper. Res.* **73**, 295–303 (1994)
42. Donsker, M.D., Varadhan, S.S.: Asymptotic evaluation of certain Markov process expectations for large time, i. *Comm. Pure Appl. Math.* **28**, 1–47 (1975)
43. Ellis, R.S.: *Entropy, Large Deviations and Statistical Mechanics*. Springer, Berlin (1985)
44. Kurchan, J.: Fluctuation theorem for stochastic dynamics. *J. Phys. A: Math. Gen.* **31**, 3719–3729 (1998)
45. Lebowitz, J.L., Spohn, H.: A Gallavotti-Cohen-type symmetry in the large deviation functional for stochastic dynamics. *J. Stat. Phys.* **95**, 333–365 (1999)
46. Pham, H.: *Continuous-time Stochastic Control and Optimization with Financial Applications*. Springer, Berlin (2009)
47. Robin, M.: Long-term average cost control problems for continuous time Markov processes: a survey. *Acta Appl. Math.* **1**, 281–299 (1983)