

Meshless discretization of LQ-type stochastic control problems

Ralf Banisch Carsten Hartmann

October 1, 2013

We propose a novel Galerkin discretization scheme for stochastic optimal control problems on an indefinite time horizon. The control problems are linear-quadratic in the controls, but possibly nonlinear in the state variables, and the discretization is based on the fact that problems of this kind can be transformed into linear boundary value problems by a logarithmic transformation. We show that the discretized linear problem is dual to a Markov decision problem, the precise form of which depends on the chosen Galerkin basis. We prove a strong error bound in L^2 for the general scheme and discuss two special cases: a variant of the known Markov chain approximation obtained from a basis of characteristic functions of a box discretization, and a sparse approximation that uses the basis of committor functions of metastable sets of the dynamics; the latter is particularly suited for high-dimensional systems, e.g., control problems in molecular dynamics. We illustrate the method with several numerical examples, one being the optimal control of Alanine dipeptide to its helical conformation.

Contents

1. Introduction	2
1.1. Elementary notation and assumptions	5
2. Optimal control and logarithmic transformation	7
2.1. Duality between control and path sampling for diffusions	7
2.2. Duality between control and path sampling for jump processes	9
3. Discretization: Galerkin projection point of view	12
3.1. Galerkin projection of the Dirichlet problem	13
3.2. Interpretation in terms of a Markov decision problem	17

4. Numerical Results	22
4.1. 1D triple well potential	22
4.2. Alanine dipeptide	24
5. Conclusions	27
A. Proof of Theorem 2	28
B. Best-approximation error bound	30
C. Finite-volume approximation	32
D. Markov chain approximations	33
E. Sampling of the discretized running cost	34

1. Introduction

A large body of research is concerned with the question: How well can a continuous diffusion in an energy landscape be approximated by a Markov jump process (MJP) in the regime of low temperatures? Qualitatively, the approximation should be good if the system under consideration is metastable, for metastability precisely means that the diffusion process stays in the neighbourhood of the potential energy minima for a long time, and occasionally makes rapid transitions (jumps) between the wells. These metastable regions then become the states of the MJP, and the jump rates are determined by the frequency of the transitions. A rigorous mathematical proof of this fact, based on Gamma convergence, has recently been published for the special case of a symmetric double-well potential and in the limit of zero temperature [PSV12, AMP⁺12]. From a more practical point of view, Markov state models (e.g. see [PBB10, PWS⁺11, SNL⁺11, DSS10]) that are popular in the molecular dynamics community are approximations of metastable systems by MJPs that, in certain cases, can capture the correct transition rates between the metastable sets [SNS10, DSS12].

The situation is more complicated for controlled diffusions, when the processes are subject to external controls that are chosen so as to minimize a given cost criterion, in which case one has to approximate the stochastic processes (in an appropriate sense) as well as the corresponding cost functional and the resulting optimal control forces. Available numerical methods for solving stochastic optimal control problems include methods that solve the dynamic programming PDE or Hamilton-Jacobi-Bellman equation, such as Markov chain approximations [KD92], monotone schemes [BS91, BJ07] or finite elements [BH01]. A common problem for these PDE-based methods is that they become inefficient if the problems are high-dimensional. An alternative are direct methods that iteratively minimize the cost functional using, e.g. entropy minimization [Tod09], path integrals [Kap05] or policy iteration [BMZ09]. This class of methods works in high dimensions, but has difficulties if the solvers get stuck in local minima or if the search space is too large.

In this article we propose a Galerkin scheme for discretizing the dynamic programming equations that is meshless and hence can be applied to large-scale problems, assuming that the Galerkin basis is chosen in a clever way. Galerkin or, more generally, reduced basis methods are well established in terms of theory and numerical algorithms for linear elliptic equations, but to our knowledge very few results (e.g. [DD70, SWH12]) are available for the nonlinear dynamic programming equations of stochastic control; cf. also [KV01]. To close this gap we confine ourselves to a class of optimal control problems that are linear-quadratic (LQ) in the control variables, but possibly nonlinear in the states; they have the feature that they can be transformed to a linear problem by a suitable (logarithmic) transformation and are hence amenable to a discretization using Galerkin methods. LQ-type optimal control problems of this kind appear relevant in many applications, including molecular dynamics [SWH12, Sta04], material science [Ste10, AK11], quantum computing [PK02, RdVRMK00], or queuing networks [SR95, Hei95] to mention just a few; see also [DW04, DSW12] for an application in statistics.

A simple paradigm As an introductory example consider the one-dimensional diffusion process $(X_t)_{t \geq 0}$ satisfying the Itô stochastic differential equation

$$dX_t = b(X_t, t)dt + \sqrt{2\epsilon}dB_t, t \geq 0$$

where B_t is standard one-dimensional Brownian motion, $\epsilon > 0$ is noise intensity, called *temperature* in the following, and $b(\cdot, \cdot)$ is a smooth and Lipschitz continuous vector field. Specifically, we assume that b is of the form

$$b(x, t) = u_t - \nabla V(x),$$

with $V: \mathbb{R} \rightarrow \mathbb{R}$ being a smooth potential energy function that is bounded from below, and measurable control u , that will be chosen so as to minimize a certain cost criterion. As an example consider the situation depicted in Figure 1 and suppose that the control task is to force the particle in the left well to reach the right well in minimum time. For $u = 0$ and in the limit of small noise, the average of the transition time from the left to the right well is given by the Kramers-Arrhenius law [FW84, Ber11]

$$\mathbf{E}[\tau] \asymp \exp(\Delta V/\epsilon), \quad \epsilon \rightarrow 0,$$

where τ is the first exit time from the left well, ΔV is the energy barrier between the wells and $\mathbf{E}[\cdot]$ is the expectation over all realizations of the process. We can speed up the transitions by tilting the potential according to $V(x) \mapsto V(x) - ux$, making the barrier smaller. If we were allowed to apply arbitrarily large forces, we could make τ arbitrarily small, yet a real controller will seek to minimize the transition time without controlling too much. A natural choice in many cases is a penalization of the energy consumed by the controller, which leads to quadratic cost of the form [SWH12]

$$J(u) = \mathbf{E} \left[\tau + \frac{1}{4} \int_0^\tau |u_t|^2 dt \right].$$

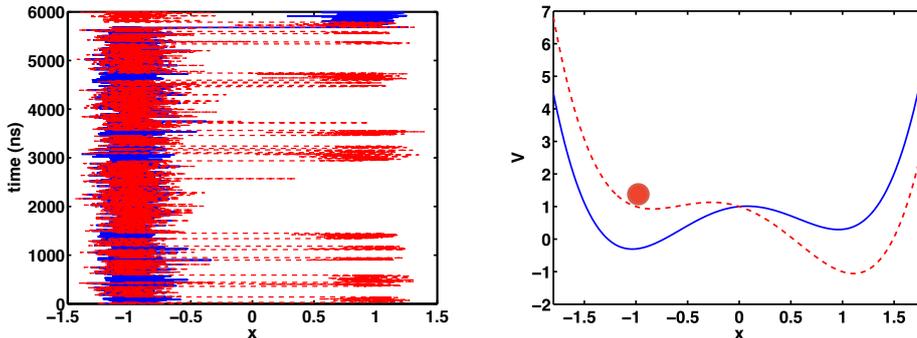


Figure 1: Two typical realization of the bistable system, with and without tilting (left panel). The corresponding potential energies are shown in the right panel.

(The factor $1/4$ is for convenience.) The showcase optimal control problem now reads:

$$\min_{u \in \mathcal{A}} J(u) \quad (1)$$

over a set \mathcal{A} of admissible (e.g. adapted) strategies and subject to

$$dX_t^u = (u_t - \nabla V(X_t^u))dt + \sqrt{2\epsilon}dB_t. \quad (2)$$

In this paper we deal with the question how to solve optimal control problems of the form (1)–(2), beyond simple one-dimensional examples. The typical application that we have in mind is molecular dynamics that is high-dimensional and displays vastly different time scales. This defines the basic requirements of the numerical method used to solve optimal control problems: it must handle problems with large state space dimension and it must be able to capture the relevant processes of the dynamics, typically the slowest degrees of freedom in the system.

For moderate controls, and if the temperature is small compared to the energy barrier, the dynamics in the above example basically consists of rare jumps between the potential wells, with the jump rate being controlled by u . Therefore an efficient discretization would be one that resolves only the jumps between the minima by a 2-state Markov jump process with adjustable jump rates, according to the value of the control variable. If the control u_t is given, the approximation of (2) by a 2-state MJP essentially boils down to an approximation of the dynamics by a time-inhomogeneous MSM, which requires only minor generalizations of the homogeneous MSM framework (see, e.g., [DSS12, SNS10]). When solving optimal control problems, however, the control becomes a function of the dynamics, for (2) enters as a constraint in the minimization problem (1), which makes the corresponding dynamic programming equations non-linear and renders the discretization less straightforward. The discretization scheme that we propose is based on the fact that the above control problem can be transformed into linear boundary value problems by a logarithmic transformation. The linear problem can then be discretized by standard means, including the discretization

by an MSM if the dynamics is metastable. As we will show below the discretized linear problem is dual to a Markov decision problem (i.e. a stochastic control problem for a continuous-time finite-state Markov process), which thus represents the natural Markovian discretization of the original stochastic control problem. The discretization is *meshless*, in that the number of states of the Markov model does not scale exponentially with the dimension of the continuous state space, hence our method avoids the curse of dimensionality of typical grid based methods.

Organization of the article The rest of the paper is organised as follows: In section 2 we introduce the class of optimal control problems studied and state the duality between optimal control and sampling for both continuous SDEs and MJPs. In section 3, the Galerkin projection method is introduced, and some results about the approximation error are discussed. We also give a stochastic interpretation of the discretized linear equation in terms of Elber’s milestoning process [FE04] and discuss special cases of the discretization, one of which being the known Markov chain approximation. Finally, we construct sampling estimators. Section 3 is the core parts of the paper and contains new results, including a strong L^2 error estimate for the Galerkin discretization. In section 4, we discuss numerical examples.

1.1. Elementary notation and assumptions

We implement the following notation and standing assumptions that will be used throughout the paper and that generalize the above example. Our optimal control problem has the following ingredients:

Dynamics Let $\mathcal{S} \subset \mathbb{R}^d$ bounded with smooth (e.g. Lipschitz) boundary and consider the potential energy function $V: \mathcal{S} \rightarrow \mathbb{R}$, that we assume to be two times continuously differentiable and bounded from below. We consider $X_t \in \mathcal{S}$ solving

$$dX_t^u = (u_t - \nabla V(X_t^u))dt + \sqrt{2\epsilon}dB_t, \quad t \geq 0, \quad (3)$$

where $B_t \in \mathbb{R}^d$ is d -dimensional Brownian motion under a probability measure P , and $u: [0, \infty) \rightarrow U \subset \mathbb{R}^d$ is a time dependent measurable and bounded function. We further impose reflecting boundary conditions at the set boundary $\partial\mathcal{S}$, so that the process cannot leave the set $\mathcal{S} \subset \mathbb{R}^d$; see [Gar96] for algorithmic issues.

Reversibility and invariant measure For test functions $\varphi: \mathcal{S} \rightarrow \mathbb{R}$ that are two times continuously differentiable, the infinitesimal generator of the uncontrolled process $X_t = X_t^0$ is defined as the second-order differential operator

$$L\varphi = \epsilon\Delta\varphi - \nabla V \cdot \nabla\varphi.$$

Define

$$d\mu(x) = \exp(-\epsilon^{-1}V(x))dx$$

to be the Boltzmann measure at temperature $\epsilon > 0$. Without loss of generality, we assume that μ is normalized, so that $\mu(\mathcal{S}) = 1$. For the subsequent analysis it will be convenient to think of L as an operator acting on a suitable subspace of

$$L^2(\mathcal{S}, \mu) = \left\{ \phi: \mathcal{S} \rightarrow \mathbb{R}: \int_{\mathcal{S}} |\phi(x)|^2 d\mu(x) < \infty \right\},$$

that is a weighted Hilbert space equipped with the scalar product

$$\langle v, w \rangle_{\mu} = \int_{\mathcal{S}} v(x)w(x)d\mu(x).$$

It can be readily seen that L is symmetric with respect to the weighted scalar product,

$$\langle Lv, w \rangle_{\mu} = \langle v, Lw \rangle_{\mu},$$

which implies that X_t is reversible with respect to the Boltzmann measure μ . Moreover μ is the unique invariant measure of the process X_t and satisfies

$$\int_{\mathcal{S}} (L\psi)d\mu = \int_{\mathcal{S}} \psi(L\mathbf{1})d\mu = 0$$

for all test functions $\psi \in L^2(\mathcal{S}, \mu)$.

Quadratic cost criterion We now introduce the cost criterion that the controller choosing u in (3) seeks to minimize. To this end let $A \subset \mathcal{S}$ be a bounded subset that is fully contained in the interior of \mathcal{S} and call $\tau_A < \infty$ the random stopping time

$$\tau_A = \inf\{t > 0: X_t \in A\}.$$

We define the cost functional

$$J(u) = \mathbf{E} \left[\int_0^{\tau_A} \left\{ f(X_t) + \frac{1}{4}|u_t|^2 \right\} dt \right], \quad (4)$$

where $f: \mathcal{S} \rightarrow \mathbb{R}$, called *running cost*, is any nonnegative function with bounded first derivative; the factor $1/4$ in the penalization term is merely conventional. Cost functionals of this form are called *indefinite time horizon cost*, because the terminal time τ_A is random. We will sometimes need the conditioned variant of the cost function,

$$J(u; x) = \mathbf{E}_x \left[\int_0^{\tau_A} \left\{ f(X_t) + \frac{1}{4}|u_t|^2 \right\} dt \right], \quad (5)$$

where $\mathbf{E}_x[\cdot] = \mathbf{E}[\cdot | X_0 = x]$ is a shorthand for the expectation over all realizations of X_t starting at $X_0 = x$, i.e. the expectation with respect to P conditional on $X_0 = x$.

Admissible control strategies We call a control strategy $u = (u_t)_{t \geq 0}$ *admissible* if it is adapted to the filtration generated by B_t , i.e., if u_t depends only on the history of the Brownian motion up to time t , and if the equation for X_t^u has a unique strong solution. The set of admissible strategies is denoted by \mathcal{A} .

Even though u_t may depend on the entire past history of the process up to time t , it turns out that optimal strategies are Markovian, i.e., they depend only on the current state of the system at time t . In our case, in which the costs are accumulated up to a random stopping time τ_A , the optimal strategies are of the form

$$u_t = \alpha(X_t^u)$$

for some function $\alpha: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Hence the optimal controls are time-homogeneous feedback policies, depending only on the current state X_t^u , but not on t .

2. Optimal control and logarithmic transformation

In this section we establish a connection between optimal control of continuous-time Markov processes and certain path sampling problems, the latter are associated with a linear boundary value partial differential equation (PDE) that can be discretized by standard numerical techniques for PDEs or Monte-Carlo. The duality between optimal control and path sampling goes back to Wendel Fleming and co-workers (e.g. [Fle77, FM95, Jam92]) and is based on a logarithmic transformation of the function

$$W(x) = \min_{u \in \mathcal{A}} J(u; x). \quad (6)$$

2.1. Duality between control and path sampling for diffusions

Our simple derivation of the duality between path sampling optimal control will be based on the Hamilton-Jacobi-Bellman equations of optimal control. To this end, we recall the dynamic programming principle for optimal control problems of the form (3)–(4) that we state without proof; for details we refer to, e.g., [FS06, Secs. VI.2].

Theorem 1. *Let $W \in C^2(\mathcal{S}) \cap C^1(\bar{\mathcal{S}}) \cap C(\bar{A})$ be the solution of*

$$\begin{aligned} \min_{c \in \mathbb{R}^d} \left\{ (L + c \cdot \nabla)W(x) + f + \frac{1}{4}|c|^2 \right\} &= 0, & x \in \mathcal{S} \setminus A \\ W(x) &= 0, & x \in A \\ \nu \cdot \nabla W(x) &= 0, & x \in \partial\mathcal{S}, \end{aligned} \quad (7)$$

where ν is the outward-pointing unit normal to $\partial\mathcal{S}$ at x . Then

$$W(x) = \min_{u \in \mathcal{A}} J(u; x)$$

where the minimizer $u^* = \operatorname{argmin} J(u)$ is unique and given by the feedback law

$$u_t = -2\nabla W(X_t^u). \quad (8)$$

Before we proceed with the derivation of the dual sampling problem, we shall briefly discuss some of the consequences of the dynamic programming approach. Equation (7) is the Hamilton-Jacobi-Bellman (HJB) equation, also called *dynamic programming equation* associated with the following optimal control task:

$$\min_{u \in \mathcal{A}} J(u) \quad \text{s.t.} \quad dX_t^u = (u_t - \nabla V(X_t^u))dt + \sqrt{2\epsilon}dB_t.$$

The function $W(x)$ is called *value function* or *optimal cost-to-go*. Existence and uniqueness of classical (i.e. smooth) solutions follow from our assumptions on the potential and the properties of \mathcal{S} , using the results in [FS06, Secs. VI.3–5].

Given the value function, and using the fact that optimal control is the gradient of two times the value function, the optimally controlled process X_t^* solves the SDE

$$dX_t^* = -\nabla U(X_t^*)dt + \sqrt{2\epsilon}dB_t. \quad (9)$$

with the new potential

$$U(x) = V(x) + 2W(x).$$

Note that X_t^* is reversible with respect to a tilted Boltzmann distribution having the density $\rho^* = \exp(-U/\epsilon)$. The reversibility follows from the fact that the value function does not depend on t , which would not be the case if the terminal time τ_A were a deterministic stopping time rather than a first exit time.¹

Logarithmic transformation and Feynman-Kac formula (part I) The approach that is pursued in this article is to discretize the HJB equation by first removing the non-linearity by a logarithmic transformation of the value function. Let

$$\phi(x) = \exp(-\epsilon^{-1}W(x)). \quad (10)$$

It follows by chain rule that

$$\epsilon \frac{L\phi}{\phi} = -LW + |\nabla W|^2, \quad \phi \neq 0, \quad (11)$$

which, together with the relation

$$-|\nabla W|^2 = \min_{c \in \mathbb{R}^m} \left\{ c \cdot \nabla W + \frac{1}{4}|c|^2 \right\},$$

implies that (7) is equivalent to the linear boundary value problem

$$\begin{aligned} (L - \epsilon^{-1}f)\phi(x) &= 0, & x \in \mathcal{S} \setminus A \\ \phi(x) &= 1, & x \in A \\ \nu \cdot \nabla \phi(x) &= 0, & x \in \partial \mathcal{S}. \end{aligned} \quad (12)$$

¹For finite time-horizon control problems the value function depends on the time $\tau_A - t$ remaining until the terminal time τ_A .

By the above assumptions and the strong maximum principle for elliptic PDEs it follows that (12) has a classical solution $\phi^\epsilon \in C^2(\mathcal{S}) \cap C(\bar{\mathcal{S}}) \cap C(\bar{A})$ that is uniformly bounded. Further note that the value function is uniformly bounded on \mathcal{S} , hence the log transformation (10)–(11) is well defined.

Now, by the Feynman-Kac theorem [Øks03, Thm. 8.2.1], the linear boundary value problem has an interpretation in terms of a sampling problem. The solution (12) can be expressed as the conditional expectation

$$\phi(x) = \mathbf{E}_x \left[\exp \left(-\frac{1}{\epsilon} \int_0^{\tau_A} f(X_t) ds \right) \right] \quad (13)$$

over all realizations of the following SDE on \mathcal{S} :

$$dX_t = -\nabla V(X_t)dt + \sqrt{2\epsilon}dB_t, \quad X_0 = x. \quad (14)$$

Remark 1. *The Neumann boundary condition in (7) and (12) amounts to the reflecting boundary conditions for the processes X_t^u and X_t at $\partial\mathcal{S}$.*

Remark 2. *In probabilistic terms, the logarithmic transformation amounts to a suitable change of measure of the underlying Markov process, by which the control variable is eliminated [PMR96]; see also [HS12] and the references given there.*

2.2. Duality between control and path sampling for jump processes

In the last section, we have established a connection between an optimal control problem and sampling of a continuous path observables $\phi(x)$. In this section, we will repeat the same construction for Markov jump processes, however, in reverse order: starting from a path observable for a Markov jump process, we derive the dual optimal control using a logarithmic transformation.

Let $(\hat{X}_t)_{t \geq 0}$ be a MJP on the discrete state space $\hat{\mathcal{S}} = \{1, \dots, n\}$ with infinitesimal generator $G \in \mathbb{R}^{n \times n}$. The entries of the generator matrix G satisfy

$$G_{ij} \geq 0 \text{ for } i \neq j \text{ and } G_{ii} = -\sum_{j \neq i} G_{ij},$$

where the off-diagonal entries of G are the jump rates between the states i and j .

Logarithmic transformation and Feynman-Kac formula (part II) In accordance with the previous subsection let $\hat{f} : \hat{\mathcal{S}} \rightarrow \mathbb{R}$ be nonnegative and define the stopping time

$$\tau_A = \inf\{t > 0 : \hat{X}_t \in A\}.$$

to be the first hitting time of a subset $A \subset \hat{\mathcal{S}}$. As before we introduce a function

$$\hat{\phi}(i) = \mathbf{E}_i \left[\exp \left(-\frac{1}{\epsilon} \int_0^{\tau_A} \hat{f}(\hat{X}_s) ds \right) \right],$$

with $\mathbf{E}_i[\cdot] = \mathbf{E}[\cdot | \hat{X}_0 = i]$ being the conditional expectation over the realizations of \hat{X}_t starting at $\hat{X}_0 = i$. We have the following Lemma that is the exact analogue of the Feynman-Kac formula for diffusions for the case of an MJP (see [GS75]).

Lemma 1. *The function $\hat{\phi}(i)$ solves the linear boundary value problem*

$$\begin{aligned} \sum_{j \in \hat{\mathcal{S}}} G_{ij} \hat{\phi}(j) - \epsilon^{-1} \hat{f}(i) \hat{\phi}(i) &= 0, \quad i \in \hat{\mathcal{S}} \setminus A \\ \hat{\phi}(i) &= 1, \quad i \in A. \end{aligned} \tag{15}$$

Now, in one-to-one correspondence with the log transformation procedure in the diffusion case, the function

$$\hat{W} = -\epsilon \log \hat{\phi}$$

can be interpreted as the value function of an optimal control problem for the MJP $(\hat{X}_t)_{t \geq 0}$. The derivation of the dual optimal control problem goes back to [She82, She85], and we repeat it here in condensed form for the reader's convenience (see also [FS06, Sec. VI.9]): First of all note that \hat{W} satisfies the equation

$$\begin{aligned} \exp(\hat{W}/\epsilon) G \exp(-\hat{W}/\epsilon) - \epsilon^{-1} \hat{f} &= 0, \quad i \in A^c \\ \hat{W}(i) &= 0, \quad i \in A. \end{aligned}$$

and define a new generator matrix by

$$G^v = (G_{ij}^v)_{i,j \in \hat{\mathcal{S}}}, \quad G_{ij}^v = \frac{G_{ij} v(j)}{v(i)},$$

with $v(i) > 0$ for all $i \in \hat{\mathcal{S}}$. The exponential term in above equation for \hat{W} can be recast as

$$\frac{(G\hat{\phi})(i)}{\hat{\phi}(i)} = \min_{v>0} \{-(G^v \hat{W})(i) + k^v(i)\}$$

where we have introduced the shorthand

$$k^v(i) = \epsilon(G^v(\log v))(i) - \epsilon \frac{(Gv)(i)}{v(i)},$$

and used the identity

$$\min_{y \in \mathbb{R}} \{e^{-y} + ay\} = a - a \log a, \quad a > 0.$$

As a consequence, (15) is equivalent (i.e. dual) to

$$\begin{aligned} \min_{v>0} \left\{ (G^v \hat{W})(i) + k^v(i) + \hat{f}(i) \right\} &= 0, \quad i \in A^c \\ \hat{W}(i) &= 0, \quad i \in A. \end{aligned} \tag{16}$$

which is the dynamic programming equation of a Markov decision problem, i.e. an optimal control problem for an MJP (e.g. see [FS06, Sec. VI.9]): Minimize

$$\hat{J}(v) = \mathbf{E} \left[\int_0^{\tau_A} \left\{ \hat{f}(\hat{X}_s^v) + k^v(\hat{X}_s^v) \right\} ds \right] \tag{17}$$

over all component-wise strictly positive controls v and subject to the constraint that the process $(\hat{X}_t^v)_{t \geq 0}$ is generated by G^v . It readily follows from the derivation of (16) that the minimizer exists and is given by

$$v^*(i) = \hat{\phi}(i).$$

The next lemma records some important properties of the controlled Markov jump process with generator G^v and the corresponding cost functional (17).

Lemma 2. *Let G^v and k^v be defined as above.*

- (i) *Let $(\hat{X}_t)_{t \geq 0}$ with generator G have a unique stationary distribution π and let G be reversible with respect to π . Then the tilted distribution $\pi^v(i) = Z_v^{-1} v^2(i) \pi(i)$, with Z_v an appropriate normalization constant, is the unique probability distribution such that G^v is reversible and stationary with respect to π^v .*
- (ii) *Let \hat{P} denote the probability measure on the space of trajectories generated by \hat{X}_t with initial condition $\hat{X}_0 = i$, and let \hat{Q} be the corresponding probability measure generated by \hat{X}_t^v with the same initial condition $\hat{X}_0^v = i$. Then \hat{Q} is absolutely continuous with respect to \hat{P} and the expected value of the running cost k^v is the Kullback-Leibler (KL) divergence between \hat{Q} and \hat{P} , i.e.,*

$$\mathbf{E}_{\hat{Q}} \left[\int_0^{\tau_A} k^v(\hat{X}_s^v) ds \right] = \int \log \frac{d\hat{Q}}{d\hat{P}} d\hat{Q}$$

where $\mathbf{E}_{\hat{Q}}[\dots]$ denotes expectation over all realizations of \hat{X}_t^v starting at $\hat{X}_0^v = i$.

Proof. We first show (i). By assumption we have $\pi(i)G_{ij} = \pi(j)G_{ji}$. Now, let π^v be such that $\pi^v(i)G_{ij}^v = \pi^v(j)G_{ji}^v$. We will show that π^v has the proposed form:

$$\pi^v(i)G_{ij}^v = \frac{v(j)}{v(i)} \frac{\pi^v(i)}{\pi(i)} \pi(i)G_{ij} = \frac{v(j)}{v(i)} \frac{\pi^v(i)}{\pi(i)} \pi(j)G_{ji} = \frac{v^2(j)}{v^2(i)} \frac{\pi(j)}{\pi(i)} \frac{\pi^v(i)}{\pi^v(j)} \pi^v(j)G_{ji}^v$$

But since $\pi^v(i)G_{ij}^v = \pi^v(j)G_{ji}^v$, we must have

$$\frac{\pi^v(j)}{\pi(j)v^2(j)} = \frac{\pi^v(i)}{\pi(i)v^2(i)} \quad \forall i \neq j.$$

This can only be true if the quantity $Z_v^{-1} = \frac{\pi^v(i)}{\pi(i)v^2(i)}$ is independent of i . This gives $\pi^v(j) = Z_v^{-1} v^2(j) \pi(j)$ as desired. The constant Z_v is uniquely determined by the requirement that π^v be normalized. Finally, from reversibility it follows directly that π^v is also a stationary distribution of G^v .

To show (ii), note that the running cost $k^v(i)$ can be written as

$$k^v(i) = \epsilon \sum_{j \neq i} G_{ij} \left\{ \frac{v(j)}{v(i)} \left[\log \frac{v(j)}{v(i)} - 1 \right] + 1 \right\}, \quad (18)$$

which is the KL divergence between \hat{Q} and \hat{P} (see [PMR96, Sec. 3.1.4]). The absolute continuity between \hat{Q} and \hat{P} simply follows from the fact that v in the definition of G^v was required to be component-wise strictly positive. \square

Remark 3. To reveal further similarities between the stochastic control problem (3)–(4) and the corresponding Markov decision problem, note that the quadratic penalization term in (4) equals the KL divergence between the reference measure P of the uncontrolled diffusion (14) and the corresponding probability measure Q that is induced by replacing B_t in (14) by

$$B_t^u = B_t + \sqrt{\frac{1}{2\epsilon}} \int_0^t u_s dt,$$

as can be shown using Girsanov’s theorem [Øks03, Thm. 8.6.8]. In other words, it holds that (cf. [HS12, HBS⁺13])

$$\mathbf{E}_Q \left[\frac{1}{4} \int_0^{\tau_A} |u_s|^2 ds \right] = \epsilon \int \log \frac{dQ}{dP} dQ$$

3. Discretization: Galerkin projection point of view

In this section we will develop a discretization for optimal control problems of the type discussed in Section 2. The discretization will approximate the continuous control problem with a control problem for a Markov jump process on finite state space.

Because of the nonlinearity of the problem, a general theory for discretizing continuous optimal control problems is unavailable. However, we saw in Section 2 that for the control problems we are interested in, a logarithmic transform to a linear PDE is available. For linear PDEs, discretization theory in terms of Galerkin projections onto finite-dimensional subspaces of the PDE solution space exists. Our strategy will therefore be the one indicated in Figure 2.

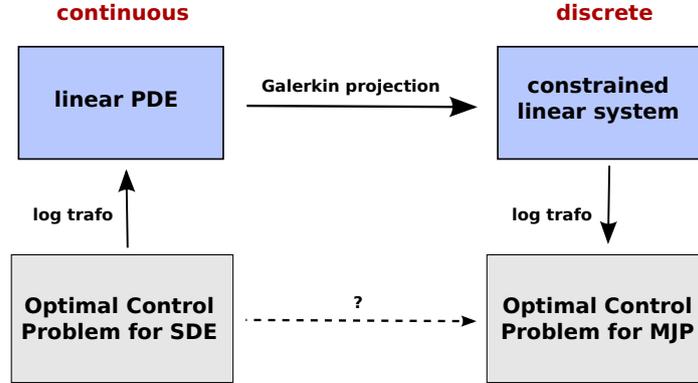


Figure 2: Discretization of continuous control problems via a logarithmic transform.

In the first part of this section, we will develop the Galerkin projection for general subspaces and obtain some control of the discretization error. To refine this control, we specify the subspace D we project onto. Specifically, we develop two possible choices

for D inspired by MSMs. The first choice specifies D as the space of step functions on a full partition of state space, and if that full partition is chosen as a lattice with spacing h , then, as expected, the discretization error vanishes for $h \rightarrow 0$. The second choice uses a core partition of state space, where the cores are the metastable regions of the uncontrolled dynamics. We will prove a novel error bound which gives us detailed control over the discretization error even if very few basis functions are used.

In the second part of this section, we will develop the stochastic interpretation of the resulting matrix equation as the backward Kolmogorov equation of a MJP, which enables us to identify the discrete control problem for the MJP, as it was developed in Section 2. We will study the resulting discrete control problem for the two choices of D specified earlier. In the full partition case, we will establish a connection to the finite volume approximation discussed in [LMHS11], and we will show that to first order in h , our discrete control problem coincides with the Markov chain approximation constructed by Kushner [KD92], confirming that the control problem itself converges to the continuous problem for $h \rightarrow 0$. In the incomplete partition case, we will make a connection to Transition Path Theory [VE06] and core set MSMs [SNL⁺11].

3.1. Galerkin projection of the Dirichlet problem

As discussed above, we consider the boundary value problem

$$\begin{aligned} (L - \epsilon^{-1}f)\phi(x) &= 0, & x \in A^c \\ \phi(x) &= 1, & x \in A \\ \nu \cdot \nabla \phi(x) &= 0, & x \in \partial \mathcal{S}, \end{aligned} \tag{19}$$

with L and f as given above and $A^c = \mathcal{S} \setminus A$. Following standard references we construct a Galerkin projection of (19), see e.g. [Bra12]. For this purpose, we introduce the H^1 -norm $\|\phi\|_{H^1} = \|\nabla u\|_\mu^2 + \|u\|_\mu^2$, the Hilbert space $V = \{\psi \in L^2(\mathcal{S}, \mu), \|\psi\|_{H^1} < \infty\}$ and the symmetric and positive bilinear form

$$\mathcal{B} : V \times V \rightarrow \mathbb{R}, \quad \mathcal{B}(\phi, \psi) = \epsilon^{-1} \langle f\phi, \psi \rangle_\mu + \epsilon \langle \nabla \phi, \nabla \psi \rangle_\mu.$$

Now if ϕ is a solution of (19), then it also solves the weak problem

$$\mathcal{B}(\phi, \psi) = 0 \quad \forall \psi \in V. \tag{20}$$

A *Galerkin solution* $\hat{\phi}$ is any function satisfying

$$\mathcal{B}(\hat{\phi}, \hat{\psi}) = 0 \quad \forall \hat{\psi} \in D, \tag{21}$$

with a predefined finite dimensional subspace $D \subset V$ that is adapted to the boundary conditions. In particular, we may choose basis functions $\chi_1, \dots, \chi_{n+1}$ with the following properties:

(S1) The χ_i form a partition of unity, that is $\sum_{i=1}^{n+1} \chi_i = \mathbf{1}$.

(S2) The χ_i are adapted to the boundary conditions in (19), that is $\nu \cdot \nabla \chi_i|_{\partial \mathcal{S}} = 0$, $\chi_{n+1}|_A = 1$ and $\chi_i|_A = 0$ for $i \in \{1, \dots, n\}$.

Then all elements of $D := \chi_{n+1} \oplus D_0$ with $D_0 = \text{lin}\{\chi_1, \dots, \chi_n\}$ will satisfy the boundary conditions in 19. Now define the matrices

$$F_{ij} = \frac{\langle \chi_i, f \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}, \quad K_{ij} = \frac{\langle \chi_i, L \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}.$$

Then (21) takes the form of a matrix equation for the coefficients $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i$:

$$\sum_{j=1}^{n+1} (K_{ij} - \epsilon^{-1} F_{ij}) \hat{\phi}_j = 0, \quad i \in \{1, \dots, n\} \quad (22)$$

$$\hat{\phi}_{n+1} = 1,$$

which is the discretization of (19).

Discretization error In order control the discretization error of the Galerkin method, we choose a norm $\|\cdot\|$ on V and introduce the two error measures:

1. The *Galerkin error* $\varepsilon = \|\phi - \hat{\phi}\|$, i.e. the difference between original and Galerkin solution measured in $\|\cdot\|$.
2. The *best approximation error* $\varepsilon_0 = \inf_{\hat{\psi} \in D} \|\phi - \hat{\psi}\|$, i.e. the minimal difference between the solution ϕ and any element $\hat{\psi} \in D$.

In order to obtain full control over the discretization error, we need to obtain bounds on ε , and we will do so by first obtaining a bound on the performance $\mathbf{p} := \varepsilon/\varepsilon_0$ and then a bound on ε_0 . The latter will depend on the choice of subspace D . For the former, standard estimates assume the following $\|\cdot\|$ -dependent properties of A :

- (i) Boundedness: $\mathcal{B}(\phi, \psi) \leq \alpha_1 \|\phi\| \|\psi\|$ for some $\alpha_1 > 0$
- (ii) Ellipticity: for all $\phi \in V$ holds $\mathcal{B}(\phi, \phi) \geq \alpha_2 \|\phi\|^2$ for some $\alpha_2 > 0$.

If both (i) and (ii) hold, Céa's lemma states that $\mathbf{p} \leq \frac{\alpha_1}{\alpha_2}$, see e.g. [Bra12]. For the *energy norm* $\|\phi\|_B^2 := \mathcal{B}(\phi, \phi)$ we have $\alpha_1 = \alpha_2 = 1$ and therefore $\mathbf{p} = 1$, thus the Galerkin solution $\hat{\phi}$ is the best-approximation to ϕ in the energy norm.

The next two Lemmas give a bound on \mathbf{p} if errors are measured in the L^2 -norm $\|\cdot\|_\mu$. In this case, $\mathcal{B}(\cdot, \cdot)$ is still elliptic but possibly unbounded. Later in this section, we will give examples for the choice of D and obtain bounds on ε_0 .

Theorem 2. *Let \mathcal{B} be elliptic. If Q is the orthogonal projection (with respect to $\|\cdot\|_\mu$) onto D_0 , we have*

$$\mathbf{p}^2 = \left(\frac{\varepsilon}{\varepsilon_0} \right)^2 \leq 1 + \frac{1}{\alpha_2^2} \|QBQ^\perp\|^2,$$

where $Q^\perp = 1 - Q$, $B : V \rightarrow V$ is the linear operator $\phi \mapsto \mathcal{B}(\cdot, \phi)$, and the operator norm is defined as $\|B\| = \sup_{\|x\|_\mu=1} \|Bx\|_\mu$.

Proof. In Appendix A.

Note that $\|QBQ^\perp\| \leq \|QB\|$ is always finite even though B is possibly unbounded since Q is the projection onto a finite-dimensional subspace. The bottom line of Theorem 2 is that if B leaves the subspace D almost invariant, then $\hat{\phi}$ is almost the best-approximation of ϕ in $\|\cdot\|_\mu$. The following Lemma gives a more detailed description. In the following, we will write $\|\cdot\| = \|\cdot\|_\mu$ for convenience.

Lemma 3. *Let*

$$\delta_L := \max_k \|Q^\perp L \chi_k\|, \quad \delta_f := \max_k \|Q^\perp \epsilon^{-1} f \chi_k\|$$

be the maximal projection error of the images of the χ_k 's under L and f respectively. Then

$$\|QBQ^\perp\| = \|Q^\perp BQ\| \leq (\delta_L + \delta_f) \sqrt{\frac{n}{m}},$$

where m is the smallest eigenvalue of \hat{M} .

Proof. The first statement is true since A is essentially self-adjoint. For the second statement, first of all

$$\|Q^\perp BQ\| = \|Q^\perp (\epsilon^{-1} f - L)Q\| \leq \|Q^\perp \epsilon^{-1} fQ\| + \|Q^\perp LQ\|$$

holds from the triangle inequality. We now bound the term involving L . Notice that for $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i \in D$:

$$\|Q^\perp L\hat{\phi}\| = \left\| \sum_i \hat{\phi}_i Q^\perp L \chi_i \right\| \leq \delta_L \sum_i |\hat{\phi}_i| = \delta_L \|\hat{\phi}\|_1.$$

Then, with $\hat{M}_{ij} := \langle \chi_i, \chi_j \rangle_\mu$:

$$\|Q^\perp LQ\| = \sup_{\phi = \phi_{\parallel} + \phi_{\perp} \in V} \frac{\|Q^\perp L\phi_{\parallel}\|}{\|\phi\|} \leq \sup_{\phi_{\parallel} \in D} \frac{\|Q^\perp L\phi_{\parallel}\|}{\|\phi_{\parallel}\|} \leq \delta_L \sup_{\hat{\phi} \in \mathbb{R}^n} \frac{\|\hat{\phi}\|_1}{\sqrt{\langle \hat{\phi}, \hat{\phi} \rangle_M}}$$

A similar result holds for the term involving f . The statement now follows from a standard equivalence between finite-dimensional norms, $\|\hat{\phi}\|_1 \leq \sqrt{n} \|\hat{\phi}\|_2$, and the fact that \hat{M} is symmetric, which implies that $\langle \hat{\phi}, \hat{\phi} \rangle_M = \hat{\phi}^T \hat{M} \hat{\phi} \geq m \hat{\phi}^T \hat{\phi} = m \|\hat{\phi}\|_2^2$. \square

To summarise, Theorem 2 and Lemma 3 give us a formula for the projection performance p which states that

$$p^2 \leq 1 + \frac{n}{m} \frac{(\delta_L + \delta_f)^2}{\alpha_2^2}.$$

How large or small δ_f is will depend on the behaviour of f , if i.e. $f = \text{const}$ then $\delta_f = 0$. Both δ_f and δ_L are always finite even though L is possibly unbounded.

We comment on the best-approximation error ε_0 for two choices for the subspace D which will reappear again later in the paper.

Full partition Let \mathcal{S} be fully partitioned into disjoint sets A_1, \dots, A_{n+1} with centres x_1, \dots, x_{n+1} and such that $A_{n+1} := A$, and define $\chi_i := \chi_{A_i}$. These χ_i satisfy the assumptions (S1) and (S2) discussed in Section 3.1. By definition we can bound ε_0 by any interpolation $I\phi \in D$ of the solution ϕ :

$$\varepsilon_0 \leq \|\phi - I\phi\|_\mu.$$

As interpolation, we choose $I\phi(x) = \sum_i c_i \chi_i(x)$ where $c_i = \frac{1}{\|\chi_i\|_1} \int_{A_i} \phi(x) d\mu$. If the A_i are cubes of length h and ϕ is twice continuous differentiable, then using standard techniques one can show that ε_0 is linear in h , see e.g. [Bra12].

Incomplete partition Suppose the potential $V(x)$ has $n+1$ deep minima x_1, \dots, x_{n+1} . Let C_1, \dots, C_{n+1} be convex 'core' sets around x_1, \dots, x_{n+1} and such that $A = C_{n+1}$. We write $C = \cup_{i=1}^{n+1} C_i$ and $T = \mathcal{S} \setminus C$ and introduce $\tau_C = \inf\{t \geq 0 : X_t \in C\}$. We take χ_i to be the committor function associated to the set C_i , that is

$$\chi_i(x) = \mathbf{P}(X_{\tau_C} \in C_i | X_0 = x). \quad (23)$$

These χ_i satisfy the assumptions (S1) and (S2). Since we do not have an order parameter h controlling the resolution of the discretization, standard PDE techniques for bounding ε_0 fail. Indeed, typically we will have very few basis functions compared to a grid-like discretization. The following Lemma gives a bound on ε_0 .

Theorem 3. *Let Q be the orthogonal projection onto the subspace D spanned by the committor functions (23), and let ϕ be the solution of (19). Then we have*

$$\varepsilon_0 = \|Q^\perp \phi\|_\mu \leq \|P^\perp \phi\|_\mu + \mu(T)^{1/2} [\kappa \|f\|_\infty + 2\|P^\perp \phi\|_\infty]$$

where $\|\cdot\| = \|\cdot\|_\mu$, $\kappa = \sup_{x \in T} \mathbf{E}_x[\tau_{\mathcal{S} \setminus T}]$, and P is the orthogonal projection onto the subspace $V_c = \{v \in L^2(\mathcal{S}, \mu), v = \text{const on every } C_i\} \subset L^2(\mathcal{S}, \mu)$.

Proof. In Appendix B □

In theorem 3, κ is the maximum expected time of hitting the metastable set from outside (which is short). Note further that $P^\perp \phi = 0$ on T . The errors $\|P^\perp \phi\|_\mu$ and $\|P^\perp \phi\|_\infty$ measure how constant the solution ϕ is on the core sets. Theorem 3 gives us excellent control over ε_0 , and together with theorem 2 we have full control over the discretization error ϵ for the case of incomplete partitions. These error bounds are along the lines of MSM projection error bounds [SNS10], [DSS12], and to the best of the authors' knowledge they are new.

Remark 4. *It would be nice to have an error estimate also for the value function. In general such an estimate is difficult to get, because of the nonlinear logarithmic transformation $W = -\epsilon \log \phi$ involved. However we know that ϕ and its discrete approximation are both uniformly bounded and bounded away from zero. Hence the logarithmic transformation is uniformly Lipschitz continuous on its domain, which implies that the L^2 error bounds holds for the value function with an additional prefactor given by the Lipschitz constant squared; for a related argument see [ZLPH13]*

3.2. Interpretation in terms of a Markov decision problem

We derive an interpretation of the discretized equation (22) in terms of a MJP. We introduce the diagonal matrix Λ with entries $\Lambda_{ii} = \sum_j F_{ij}$ (zero otherwise) and the full matrix $G = K - \epsilon^{-1}(F - \Lambda)$, and rearrange (22) as follows:

$$\sum_{j=1}^{n+1} (G_{ij} - \epsilon^{-1}\Lambda_{ij}) \hat{\phi}_j = 0, \quad i \in \{1, \dots, n\} \quad (24)$$

$$\hat{\phi}_{n+1} = 1,$$

This equation can be given a stochastic interpretation. To this end let us introduce the vector $\pi \in \mathbb{R}^{n+1}$ with nonnegative entries $\pi_i = \langle \chi_i, \mathbf{1} \rangle$ and notice that $\sum_i \pi_i = 1$ follows immediately from the fact that the basis functions χ_i form a partition of unity, i.e. $\sum_i \chi_i = \mathbf{1}$. This implies that π is a probability distribution on the discrete state space $\hat{\mathcal{S}} = \{1, \dots, n+1\}$. We summarise properties of the matrices K , F and G :

Lemma 4. *Let K , G , F and π be as above.*

(i) *K is a generator matrix (i.e. K is a real-valued square matrix with row sum zero and positive off-diagonal entries) with stationary distribution π that satisfies detailed balance*

$$\pi_i K_{ij} = \pi_j K_{ji}, \quad i, j \in \hat{\mathcal{S}}$$

(ii) *$F \geq 0$ (entry-wise) with $\pi_i F_{ij} = \pi_j F_{ji}$ for all $i, j \in \hat{\mathcal{S}}$.*

(iii) *G has row sum zero and satisfies $\pi^T G = 0$ and $\pi_i G_{ij} = \pi_j G_{ji}$ for all $i, j \in \hat{\mathcal{S}}$.*

(iv) *There exists a (possibly ϵ -dependent) constant $0 < C < \infty$ such that $G_{ij} \geq 0$ for all $i \neq j$ if $\|f\|_\infty \leq C$. In this case equation (24) admits a unique and strictly positive solution $\hat{\phi} > 0$.*

Proof. (i) follows from $\sum_i \chi_i(x) = \mathbf{1}$ and reversibility of L : We have $\sum_i \pi(i) K_{ij} = \sum_i \langle \chi_i, L \chi_j \rangle_\mu = \langle L \mathbf{1}, \chi_j \rangle_\mu = 0$ and $\pi(i) K_{ij} = \langle \chi_i, L \chi_j \rangle_\mu = \langle L \chi_i, \chi_j \rangle_\mu = \pi(j) K_{ji}$. (ii) follows from $f(x)$ being real and positive for all x . As for (iii), G has row sum zero by (i) and the definition of Λ . $\pi(i) G_{ij} = \pi(j) G_{ji}$ follows from (i), (ii) and the fact that Λ is diagonal, and $\pi^T G = 0$ follows directly. For (iv), rewrite (24) as the $n \times n$ -system $\bar{G}_\lambda \bar{\phi} = g$ where \bar{G}_λ is the first n rows and columns of $G_\lambda := -G + \epsilon^{-1}\Lambda$, $\bar{\phi} = (\bar{\phi}, 1)^T$ and $-g$ is the vector of the first n entries of the $(n+1)$ st row of G_λ . Choose C such

that $\epsilon^{-1}\langle \chi_i, f\chi_j \rangle_\mu \leq \langle \chi_i, L\chi_j \rangle_\mu$ for all $i \neq j$. Then $g > 0$ and \bar{G}_λ is a non-singular M -matrix and thus inverse monotone [BP79], that is from $\bar{G}_\lambda \bar{\phi} = g$ and $g > 0$ follows $\bar{\phi} > 0$. \square

It follows that if the running costs f are such that (iv) in Lemma 4 holds, then G is a generator matrix of a MJP that we shall denote by $(\hat{X}_t)_{t \geq 0}$, and by lemma 1, (24) has a unique and positive solution of the form

$$\hat{\phi}(i) = \mathbf{E} \left[\exp \left(-\epsilon^{-1} \int_0^{\tau_A} \hat{f}(\hat{X}_s) ds \right) \middle| \hat{X}_0 = i \right]$$

with $\hat{f}(i) = \Lambda_{ii}$ and $\tau_A = \inf\{t \geq 0 | \hat{X}_t = i + 1\}$. In fact (24) can be interpreted as the backward Kolmogorov equation for $\hat{\phi}$. Moreover, the logarithmic transformation $\hat{W} = -\epsilon \log \hat{\phi}$ is well-defined and can be interpreted as the value function of the Markov decision problem (16)–(17), that is, we seek to minimize

$$\hat{J}(v; i) = \mathbf{E} \left[\int_0^{\tau_A} \left(\hat{f}(\hat{X}_s^v) + k^v(\hat{X}_s^v) \right) ds \middle| \hat{X}_0^v = i \right]$$

over Markov control strategies $v : \hat{\mathcal{S}} \rightarrow (0, \infty)$ with the costs

$$\hat{f}(i) = \Lambda_{ii}, \quad k^v(i) = \epsilon \sum_{j \neq i} G_{ij} \left\{ \frac{v(j)}{v(i)} \left[\log \frac{v(j)}{v(i)} - 1 \right] + 1 \right\}.$$

This completes the construction of the discrete control problem. We now analyse it in detail for the two choices of projection subspace D introduced before.

Full partitions We partition \mathcal{S} into disjoint sets A_i that we take to be rectangular with centres x_i , we let $S_{ij} = A_i \cup A_j$ and h_{ij} be the line joining x_i and x_j , see Figure 3. Let $m(A_i)$, $m(S_{ij})$ and $m(h_{ij})$ be the Lebesgue volumes of the cells A_i , surfaces S_{ij} and lines h_{ij} respectively, and let $\bar{x}_{ij} = S_{ij} \cap h_{ij}$.

We show in Appendix C that the matrix K_{ij} has then components

$$K_{ij} \approx \frac{1}{\Delta_{ij}} e^{-\epsilon^{-1}(V(\bar{x}_{ij}) - V(x_i))}, \quad \Delta_{ij}^{-1} = \epsilon \frac{m(S_{ij})}{m(h_{ij})m(A_i)} \quad (25)$$

if i and j are neighbours ($K_{ij} = 0$ otherwise). K is the generator of a MJP on the cells A_i and coincides with the so-called *finite volume approximation* of L discussed in [LMHS11]. The approximations we use in Appendix C to calculate the integrals coincide with the ones used in [LMHS11]. The invariant distribution of K is

$$\pi_i := \|\chi_i\|_1 = \int_{A_i} d\mu \approx m(A_i) e^{-\beta V(x_i)}. \quad (26)$$

Since $\chi_i \cap \chi_j = \emptyset$ for $i \neq j$, F is diagonal, and we obtain the running costs

$$\hat{f}(i) = \frac{1}{\pi_i} \int_{A_i} f(x) \mu(x) dx = \mathbf{E}_\mu[f(X_t) | X_t \in A_i] \quad (27)$$

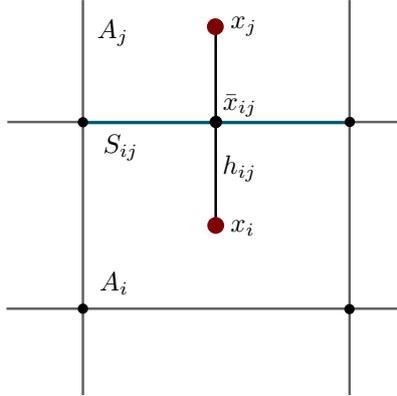


Figure 3: The mesh for the full partition.

by simply averaging $f(x)$ over the cell A_i . (27) is also a sampling formula for $\hat{f}(i)$. It follows directly that $G = K$, and in particular assumption (iv) of Lemma 4 holds for any f .

K and π can be computed from the potential V and the geometry of the mesh. In fact we can also derive a relation to standard Markov state modelling. Let P^τ be the MSM transition matrix with lag time τ associated to the partition $\{A_1, \dots, A_{n+1}\}$, that is:

$$P_{ij}^\tau = \frac{1}{\|\chi_i\|_1} \langle \chi_i, T_\tau \chi_j \rangle_\mu, \quad T_\tau = \exp(\tau L).$$

Then, by bounded convergence,

$$\lim_{\tau \rightarrow 0} \frac{1}{\tau} (P_{ij}^\tau - \delta_{ij}) = \lim_{\tau \rightarrow 0} \frac{1}{\pi_i} \langle \chi_i, \frac{1}{\tau} (T_\tau - \mathbf{1}) \chi_j \rangle = \frac{1}{\pi_i} \langle \chi_i, L \chi_j \rangle = K_{ij}, \quad (28)$$

thus K is the generator of the semigroup of transition matrices P^τ . For finite lag time τ the transition matrices P^τ can be sampled from long realizations of the original dynamics. This introduces a sampling error which depends on details of the partition, the available sampling data and the existence of rare transitions in the system. We do not address the sampling error in this paper, see e.g. [PWS⁺11, Roe08]. In view of (28) we could in principle sample K by sampling P^τ for very small τ , this is difficult however due to recrossing problems, see e.g. [CES⁺11].

Recovering Markov Chain approximations In Appendix D we show that, if \mathcal{S} is one-dimensional² and the cells are intervals of length h , we can write the nonzero off-diagonal components of the generator of the controlled MJP as

²These assumptions are mostly for notational convenience. The proof should be straightforward to generalise.

$$\begin{aligned}
G_{i,i\pm 1}^v &= \frac{1}{h^2} \left(\epsilon \mp \frac{h}{2} (\nabla V(x_i) - \alpha_v(i)) + \mathcal{O}(h^2) \right), \\
\alpha_v(i) &= \frac{\epsilon}{h} (\log v(i+1) - \log v(i-1))
\end{aligned} \tag{29}$$

where as usual $G_{ii}^v = -\sum_{j \neq i} G_{ij}^v$. We also show that the running costs of strategy v can be written as

$$k^v(i) = \frac{1}{4} \alpha_v^2(i) + \mathcal{O}(h). \tag{30}$$

This may be compared to a well-known discretization of continuous optimal control problems known as the Markov chain approximation (MCA); see [KD92]. The MCA discretization may be obtained by replacing derivatives with finite differences in the continuous control problem (7). The result is

$$\min_{\tilde{\alpha}(i) \in \mathbb{R}} \left[(\tilde{G}^{\tilde{\alpha}} \tilde{W})(i) + \frac{1}{4} \tilde{\alpha}^2(i) + \tilde{f}(i) \right] = 0 \text{ for } i \in \{1, \dots, n\}, \quad \hat{W}(n+1) = 0 \tag{31}$$

which is a Bellman equation for the MCA optimal cost \tilde{W} with strategies $\tilde{\alpha} \in \mathbb{R}^n$ and average running costs $\tilde{f}(i)$. The nonzero components of the MCA generator³ $\tilde{G}^{\tilde{\alpha}}$ corresponding to the strategy vector $\tilde{\alpha}$ are

$$\tilde{G}_{i,i\pm 1}^{\tilde{\alpha}} = \frac{1}{h^2} \left(\epsilon - \frac{h}{2} (\nabla V(i) - \tilde{\alpha}_i) \right), \quad \tilde{G}_{ii} = -\sum_{j \neq i} \tilde{G}_{ij}. \tag{32}$$

To compare both control problems, we need to be able to compare strategies. For our MJP control problem, strategies v were positive functions on $\hat{\mathcal{S}}$, but with $w = \log v$ we can think of U as \mathbb{R}^{n+1} . For the MCA approximation, $\tilde{U} = \mathbb{R}^n$. (29) gives a mapping $z : U \rightarrow \tilde{U}$ with $z(v) = \alpha_v$. It can be shown that z is onto and can therefore be used to map strategies. Now, comparing (31) and (32) with (16) and (29) gives $G^v = \tilde{G}^{\tilde{\alpha}} (1 + \mathcal{O}(h^2))$ if we set $\tilde{\alpha} = z(v) = \alpha_v$, and the Bellman equations are equal up to first order in h if strategies are mapped accordingly. Moreover, optimal strategies have the same functional dependence on optimal costs:

$$\begin{aligned}
\tilde{\alpha}^* &= -\frac{1}{h} \left(\tilde{W}(i+1) - \tilde{W}(i-1) \right) \\
\alpha_{v^*} &= \frac{\epsilon}{h} (\log v^*(i+1) - \log v^*(i-1)) = -\frac{1}{h} \left(\hat{W}(i+1) - \hat{W}(i-1) \right).
\end{aligned}$$

In the limit $h \rightarrow 0$, our discretization therefore coincides with the MCA. Convergence theory for MCAs [KD92] states that the discrete control problem (31) converges for

³In the literature, one usually considers the matrix $I + \tilde{G}^{\tilde{\alpha}}$ and interprets it as a transition matrix for a Markov chain. To be able to compare with our approach, we instead interpret $\tilde{G}^{\tilde{\alpha}}$ as a generator, which is equivalent.

$h \rightarrow 0$ to the continuous problem (7) and that $\tilde{W} \rightarrow W$ and $\tilde{\alpha} \rightarrow \alpha$ converge weakly in V . Therefore we can deduce that $\hat{W} \rightarrow W$ and $\alpha_v^* \rightarrow \alpha$ converge weakly in V as $h \rightarrow 0$.

Incomplete partitions We use a core set partition of \mathcal{S} as introduced in Section 3.1. The projection onto the committor basis χ_i also allows for a stochastic interpretation in terms of the forward and backward milestoning process \tilde{X}_t^\pm , which we define in the following way: $\tilde{X}_t^+ = i$ if the process X_t visits the core set C_i next, and $\tilde{X}_t^- = i$ if X_t came from C_i last. With this definition, the discrete costs can be written as

$$\hat{f}(i) = \frac{1}{\pi_i} \langle \chi_i, f \sum_j \chi_j \rangle = \int \nu_i(x) f(x) dx = \mathbf{E}_\mu \left[f(X_t) \middle| \tilde{X}_t^- = i \right] \quad (33)$$

where $\nu_i(x) = \pi_i^{-1} \chi_i(x) \mu(x) = \mathbf{P}(X_t = x | \tilde{X}_t^- = i)$ is the probability density of finding the system in state x given that it came last from i . Hence $\hat{f}(i)$ is the average costs conditioned on the information $\tilde{X}_t^- = i$, i.e. X_t came last from A_i , which is the natural extension to the full partition case where $\hat{f}(i)$ was the average costs conditioned on the information that $X_t \in A_i$.

The matrix $K_{ij} = \pi_i^{-1} \langle \chi_i, L \chi_j \rangle$ is reversible with stationary distribution

$$\pi_i = \langle \chi_i, \mathbf{1} \rangle = \mathbf{P}_\mu(\tilde{X}_t^- = i)$$

and is related to so called *core MSMs*. To see this, define the core MSM transition matrix P^τ with components $P_{ij}^\tau = \mathbf{P}(\tilde{X}_{t+\tau}^+ = j | \tilde{X}_t^- = i)$, and the mass matrix M with components $M_{ij} = \mathbf{P}(\tilde{X}_t^+ = j | \tilde{X}_t^- = i)$. Then, it is not hard to show that for reversible processes we have $P_{ij}^\tau = \pi_i^{-1} \langle \chi_i, T^\tau \chi_j \rangle_\mu$ and $M_{ij} = \pi_i^{-1} \langle \chi_i, \chi_j \rangle_\mu$ so that

$$K = \frac{1}{\pi_i} \langle \chi_i, L \chi_j \rangle_\mu = \lim_{\tau \rightarrow 0} \frac{1}{\tau} (P^\tau - M).$$

Formally, K is the generator of the P^τ , but these do not form a semigroup since $M \neq \mathbf{1}$, and therefore we cannot interpret K directly as e.g. the generator of \tilde{X}_t^- . Nevertheless, the entries of K are the transition rates between the core sets as defined in transition path theory [VE06]. We can obtain P^τ and M from sampling as in the full partition case. The difference is that if the core sets are chosen as the metastable states of the system, P^τ can be sampled for all lag times τ , and K can be sampled directly. See [SNS10], [DSS10] and [SNL⁺11] for more details on the construction and sampling of core MSMs. In Appendix E we show that F can also be sampled using

$$F_{ij} = \mathbf{E}_\mu \left[f(X_t) \chi_{\{\tilde{X}_t^+ = j\}} \middle| \tilde{X}_t^- = i \right] \quad (34)$$

Therefore, as in the construction of core MSMs, we do not need to compute committor functions explicitly. Note however that $G \neq K$, there is a reweighting due to the overlap of the χ_i 's which causes F to be non-diagonal. This reweighting is the surprising bit of this discretization. From Lemma 4 we see however that G and K are both reversible with stationary distribution π . Finally, note that if the cost function

$f(x)$ doesn't satisfy $\|f\|_\infty \leq C$ from (iv) in Lemma 4, G will not even be a generator matrix. In this case (22) still has a solution $\hat{\phi}$ which is the best-approximation to ϕ , but this solution may not be unique, it may not satisfy $\hat{\phi} > 0$, and we have no interpretation as a discrete control problem.

4. Numerical Results

We will present two examples to illustrate the approximation of LQ-type stochastic control problems based on a sparse Galerkin approximation using MSMs.

4.1. 1D triple well potential

To begin with we study diffusion in the triple well potential which is presented in Figure 4a. This potential has three minima at approximately $x_{0/1} = \pm 3.4$ and $x_2 = 0$. We choose the three core sets $C_i = [x_i - \delta, x_i + \delta]$ around the minima with $\delta = 0.2$. We choose $C_0 = A$ as the target set and the running cost $f = \sigma = \text{const}$, such that the control goal is to steer the particle into C_0 in minimum time.

In Figure 4a the potential V and effective potential U are shown for $\epsilon = 0.5$ and $\sigma = 0.08$ (solid lines), cf. equation (9). One can observe that the optimal control effectively lifts the second and third well up which means that the optimal control will drive the system into C_0 very quickly. The reference computations here have been carried out using a full partition FEM discretization of (12) with a lattice spacing of $h = 0.01$. Now we study the MJP approximation constructed via the committor functions shown in Figure 4b. These span a three-dimensional subspace, but due to the boundary conditions the subspace D_0 of the method is actually two-dimensional. The dashed line in Figure 4a gives the approximation to U calculated by solving (24). We can observe extremely good approximation quality, even in the transition region. In Figure 4c the optimal control $u^*(x)$ (solid line) and its approximation $\hat{u}^* = -2\nabla\hat{W}$ (dashed line) are shown. The core sets are shown in blue. We can observe jumps in \hat{u}^* at the left boundaries of the core sets. This is to be expected and comes from the fact that the committor functions are not smooth at the boundaries of the core sets, but only continuous.

Next we construct a core MSM to sample the matrices K and F . 100 trajectories of length $T = 20000$ were used to build the MSM. In Figure 4d, W and its estimate using the core MSM is shown for $\epsilon = 0.5$ and different values of σ . Each of the 100 trajectories has seen about four transitions. For comparison, a direct sampling estimate of W using the same data is shown (green). The direct sampling estimate suffers from a large bias and variance and is practically useless. In contrast, the MSM estimator for W performs well for all considered values of σ . The constant C which ensures $\hat{\phi} > 0$ when $\sigma \leq C$ is approximately 0.2 in this case. This seems restrictive but still allows to capture all interesting information about ϕ and W .

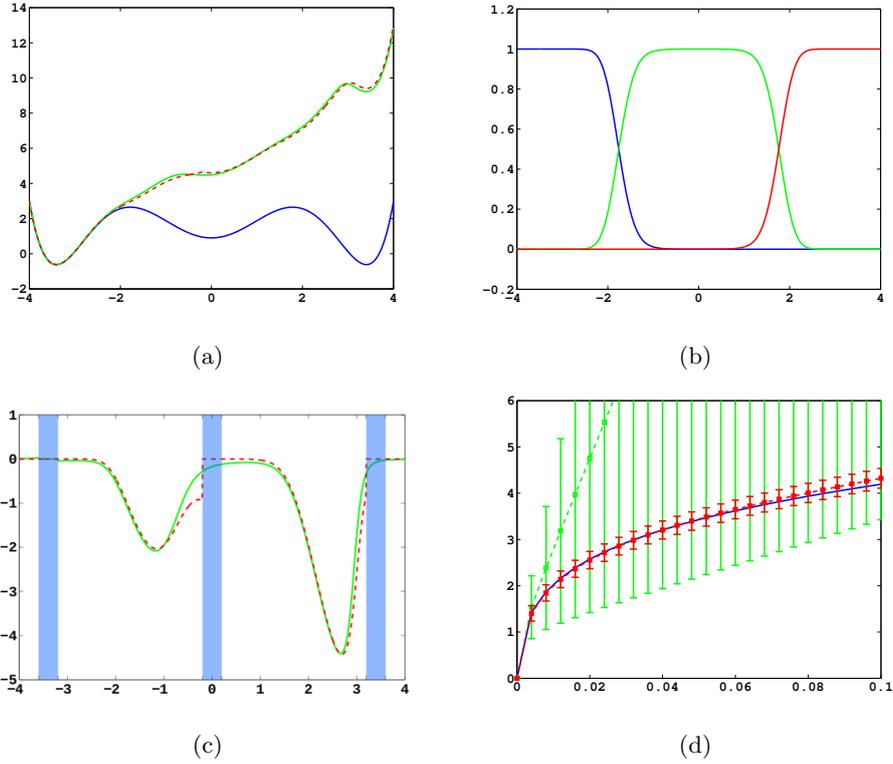


Figure 4: Three well potential example for $\epsilon = 0.5$ and $\sigma = 0.08$. (a) Potential $V(x)$ (blue), effective potential $U = V + 2W$ (green) and approximation of U with committors (dashed red). (b) The three committors $\chi_1(x)$, $\chi_2(x)$ and $\chi_3(x)$. (c) The optimal control $\alpha^*(x)$ (solid line) and its approximation (dashed line). Core sets are shown in blue. (d) Optimal cost W for $\epsilon = 0.5$ as a function of σ . Blue: Exact solution. Red: Core MSM estimate. Green: Direct sampling estimate.

4.2. Alanine dipeptide

As a second, non-trivial example we study the α - β conformational transition in Alanine dipeptide (ADP), a well-studied test system for molecular dynamics applications. We use a $1\mu s$ long molecular dynamics trajectory simulated in a box of 256 (explicit) water molecules using the CHARMM27 force field. The conformational dynamics is monitored as usual via the backbone dihedral angles ϕ and ψ . The data was first presented in [SNL⁺11]. In Figure 5, a cartoon of the molecule is shown.

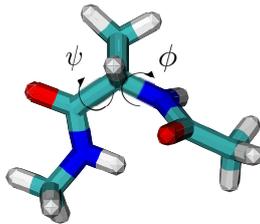


Figure 5: Alanine dipeptide.

Validation of the MSM approximation We construct a full partition MSM with 250 clusters using k -means clustering. The cluster algorithm uses the Cartesian coordinates of the ADP configurations as input data, ignoring the ADP velocities and the solvent molecules. Our first test is study the effect of the approximation of the generator matrix K by the sampled transition matrix P^τ according to $\tau^{-1}(P^\tau - 1)$; see (28). To obtain a robust estimate of K , we first focus on the mean first passage time (MFPT) $t(x) = \mathbf{E}_x[\tau_\alpha]$ where τ_α is the first hitting time of the α conformation, which we define as a ball C_α with radius $r = 45$ around the known minimum $(\phi_\alpha, \psi_\alpha) = (-80, -60)$ of the free energy landscape in (ϕ, ψ) . The MFPT satisfies the matrix equation

$$K\hat{t} = -1 \text{ outside } C_\alpha, \quad \hat{t} = 0 \text{ in } C_\alpha$$

which we study with K replaced by $\tau^{-1}(P^\tau - 1)$. In Figure 6a, the results are shown for $\tau = 5ps$, we can identify the β -structure as the red cloud of clusters where $t(x)$ is approximately constant. In 6b, $\hat{t}_{\beta\alpha} = \mathbf{E}(\hat{t}(i)|i \in \beta)$ is shown as a function of τ . We observe a linear behaviour for large τ which is due to the linear error introduced in the replacement of K with $\tau^{-1}(P^\tau - 1)$ and a nonlinear drop for small τ which is due to non-Markovianity. Our best guess is therefore a linear interpolation to $\tau = 0$, which is indicated by the solid line. The result is $\hat{t}_{\beta\alpha}^0 = 35.5ps$. As a comparison the reference value $\hat{t}_{\beta\alpha}^{\text{ref}} = 36.1ps$ from [SNL⁺11] is shown as a dashed line, that was computed therein as an inverse rate, using the slowest *implied time scale (ITS)* and information about the equilibrium weights of the α and β structure. We see very good agreement, which indicates that the strategy of linearly interpolating lag time dependent results to $\tau = 0$ is robust.

Controlled transition to the α -helical structure Next we consider an optimal control problem for steering the molecule into the α -structure. We choose as the target region $A = C_\alpha$ and define running costs in the (ϕ, ψ) variables as $f(\phi, \psi) = f_0 + f_1 \|\psi - \psi_\alpha\|^2$ where f_0 and f_1 are constants and $\|\cdot\|$ is a simple metric on the torus. We choose $f_0 = 0.01$ and $f_1 = 0.001$, which represents a mild penalty for being away from the target region. We discretize this control problem using the same partition as for the MSM construction above. The matrix K is again replaced by $\tau^{-1}(P^\tau - 1)$, the matrix F is diagonal and can be sampled straightforwardly. The resulting generator matrix G^{v^*} of the optimally controlled process can be used to compute the MFPT $\hat{t}^{(v^*)}$ of the controlled process according to the matrix equation

$$G^{v^*} \hat{t}^{(v^*)} = -1 \text{ outside } C_\alpha, \quad \hat{t}^{(v^*)} = 0 \text{ in } C_\alpha.$$

The results will again depend on the lag time τ . Figure 6c shows the results for $\tau = 5ps$, while 6d shows the MFPT for different lag times and a linear interpolation to $\tau = 0$. We observe that the control leads to a speedup of the MFPT by 1–2 orders of magnitude. A larger speedup could easily be achieved by increasing the relative weight of f , compared to the quadratic penalization of the force.

Figures 6e and 6f show the optimal cost \hat{W} and optimal strategy v^* for this problem. The optimal control v^* is best understood in terms of the jump rates

$$G_{ij}^{v^*} = \frac{G_{ij} v^*(j)}{v^*(i)}.$$

If $v^*(i)$ is low, the controller accelerates jumps out of state i while slowing down jumps into state i , and vice versa if $v^*(i)$ is high. The red cloud in Figure 6f actually has value 1, in accordance with the boundary conditions for v^* .

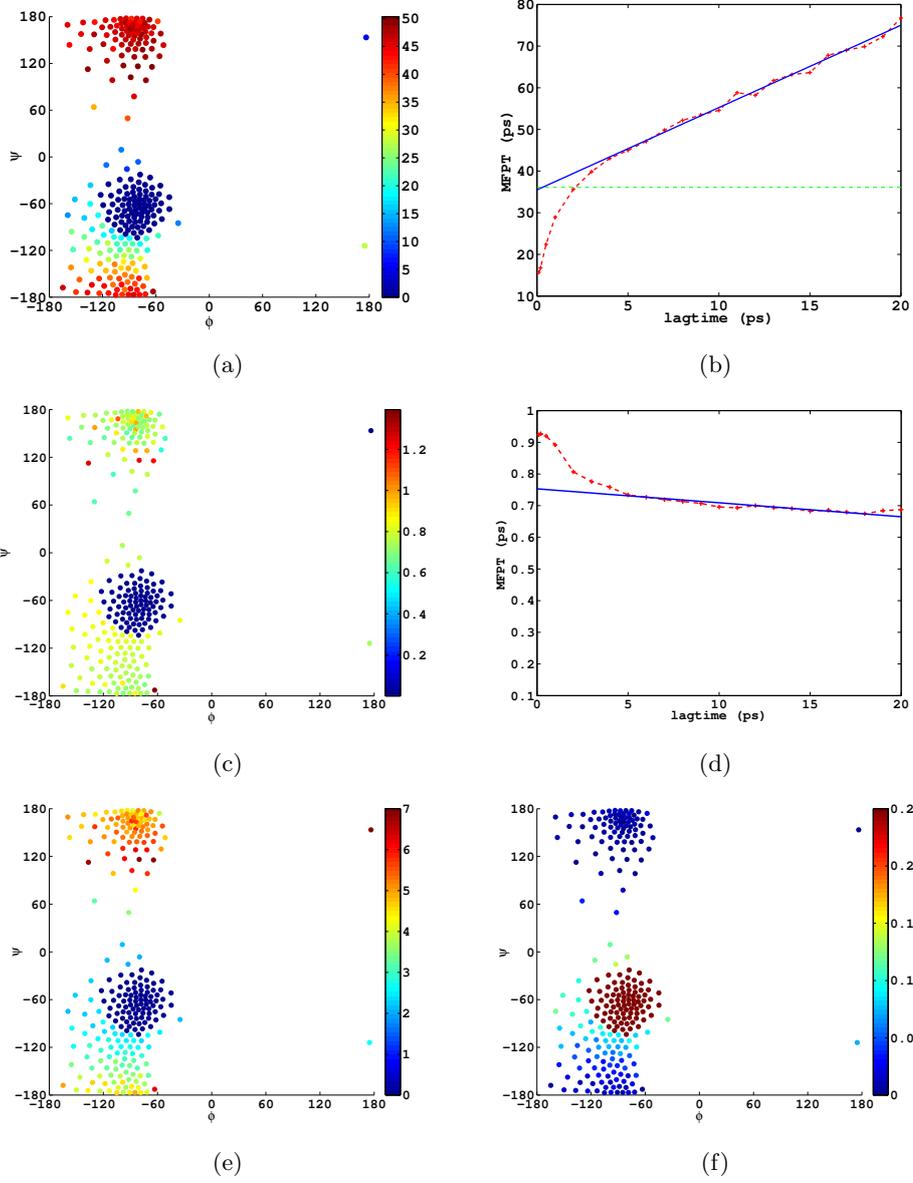


Figure 6: Top: (a) MFPT in ps from β to α in ϕ - ψ space for $\tau = 5ps$. (b) MFPT as a function of τ (dashed line) and linear interpolation to $\tau = 0$ (solid line). Green dashed line: Reference computed via slowest ITS. Middle: (c) and (d) same as (a) and (b), but for the controlled process. Bottom: Optimal cost (e) and optimal strategy (f) for the controlled process with lag time $\tau = 5ps$.

5. Conclusions

We have developed a Galerkin projection method that leads to an approximation of certain optimal control problems for reversible diffusions by Markov decision problems. The approach is based on the dual formulation of the optimal control problem in terms of a linear boundary value problem that can be discretized in a straightforward way. In this article we propose a discretization that preserves reversibility and the generator form of the linear equations, i.e., the discretization of the infinitesimal generator of the original diffusion process can be interpreted as the infinitesimal generator of a reversible Markov jump process (MJP). The discretized linear boundary value problem admits again a dual formulation in terms of a Markov decision problem.

Two special cases were discussed in detail: a Galerkin discretization based on a uniform box partition of state space and characteristic functions, that was found to agree with the known Markov chain approximation to first order in the size of the boxes, and a sparse approximation that uses the basis of committor functions of metastable sets of the dynamics; the latter does not require that the metastable sets partition the state space, hence the method can be applied to high-dimensional problems as they appear, e.g., in molecular dynamics. The committor functions in this case need not be known explicitly, as it is possible to sample the generator matrices and the discrete cost functions by a Monte-Carlo method, similarly to what is done in the Markov state modelling approach to protein folding. We could prove an L^2 error estimate for the Galerkin scheme, moreover the discretization was shown to preserve basic structural elements of the continuous problem, such as duality, reversibility or properties of the invariant measure. Our numerical results showed very good performance of the incomplete partition discretization on a simple toy example and a high-dimensional molecular dynamics problem, even with only a few basis functions, which is in line with the theoretical error bounds presented in this paper.

While we addressed the discretization error in this paper in great detail, we did not address the sampling error. In particular, for large systems our construction requires the coefficients of the MJP and therefore the transition rates between all metastable states as an input. This is not fully satisfactory. We believe that the optimal control framework presented here should be linked with Monte-Carlo methods for rare events, e.g., [HS12, DSW12], that exploit the same duality between optimal control and sampling to devise efficient importance sampling strategies as we did so as to reduce the sampling error. Moreover it would be desirable to use the MJP approach in a purely data-driven framework, e.g., for single molecule experiments or other optimal control applications in which a detailed microscopic model may not be feasible or may not be available. All this is ongoing work.

Acknowledgements

The authors thank Marco Sarich and Christof Schütte for helpful discussions and Frank Noé for providing the Alanine dipeptide data. The research was funded by the DFG Research Centre Matheon. RB holds a scholarship of the Berlin Mathematical School.

A. Proof of Theorem 2

Here we give the proof of Theorem 2 from Section 3.1. For ease of notation, let $\|\cdot\| = \|\cdot\|_\mu$.

Let ϕ be the solution to (20), and write $\phi = Q\phi + \phi_\perp = \phi_{||} + \phi_\perp$ with $\phi_\perp \in D^\perp$. The first step is to show that $\|\phi - \phi_{||}\| = \inf_{\psi \in D} \|\phi - \psi\|$, i.e. the infimum in the definition of ε_0 is attained at $\phi_{||}$. But this is clear since for any $\psi \in D$, by orthogonality we have

$$\|\phi - \psi\|^2 = \|\phi_{||} - \psi + \phi_\perp\|^2 = \|\phi_{||} - \psi\|^2 + \|\phi_\perp\|^2$$

which attains its minimum of $\varepsilon_0^2 = \|\phi_\perp\|^2$ for $\psi = \phi_{||}$. By (20), $\phi_{||}$ solves the equation

$$\mathcal{B}(\phi, \psi) = \mathcal{B}(\phi_{||}, \psi) + \mathcal{B}(\phi_\perp, \psi) = 0 \quad \forall \psi \in D,$$

and if we write $\phi_{||} = \sum_{i=1}^n \hat{\phi}_i^* \chi_i + 1\chi_{n+1}$ with n unknown coefficients $\hat{\phi}_i^*$ (note that a general element of D is of this form), this takes the matrix form

$$\hat{B}\hat{\phi}^* - c = F,$$

where in components we have $\hat{B}_{ij} = \mathcal{B}(\chi_i, \chi_j)$, $c_i = -\mathcal{B}(\phi_\perp, \chi_i) = -\langle \phi_\perp, B\chi_i \rangle_\mu$ and $F_i = -\langle \chi_i, B\chi_{n+1} \rangle_\mu$. On the other hand, the Galerkin solution $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i$ satisfies $\hat{B}\hat{\phi} = F$ by 21, hence we obtain

$$\hat{B}(\hat{\phi}^* - \hat{\phi}) = c. \tag{35}$$

Now we can write

$$\begin{aligned} \varepsilon^2 &= \|\phi_{||} + \phi_\perp - \hat{\phi}\|^2 = \|\phi_{||} - \hat{\phi}\|^2 + \|\phi_\perp\|^2 \\ &= \left\langle \sum_i (\hat{\phi}_i^* - \hat{\phi}_i) \chi_i, \sum_j (\hat{\phi}_j^* - \hat{\phi}_j) \chi_j \right\rangle_\mu + \varepsilon_0^2 \\ &= (\hat{\phi}^* - \hat{\phi})^T \hat{M} (\hat{\phi}^* - \hat{\phi}) + \varepsilon_0^2 \end{aligned}$$

where $\hat{M}_{ij} = \langle \chi_i, \chi_j \rangle_\mu$. The scalar product $\langle \cdot, \cdot \rangle_\mu$ on $D_0 \subset V$ induces a natural scalar product on \mathbb{R}^n by the isomorphism $\hat{\phi} \mapsto \sum_i \hat{\phi}_i \chi_i$:

$$\left\langle \sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}'_j \chi_j \right\rangle_\mu = \hat{\phi}^T \hat{M} \hat{\phi}' =: \langle \hat{\phi}, \hat{\phi}' \rangle_M$$

The error ε^2 is exactly ε_0^2 plus the distance between Galerkin solution and best approximation measured in this scalar product. There is also a natural bilinear form inherited from \mathcal{B} on \mathbb{R}^n :

$$\mathcal{B} \left(\sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}'_j \chi_j \right) = \hat{\phi}^T \hat{B} \hat{\phi}' = \langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi}' \rangle_M$$

The Matrix $\hat{M}^{-1} \hat{B}$ is symmetric since $\mathcal{B}(\cdot, \cdot)$ is symmetric. Moreover, since $\mathcal{B}(\cdot, \cdot)$ is elliptic,

$$\langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M = A \left(\sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}_j \chi_j \right) \geq \alpha_2 \left\langle \sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}_j \chi_j \right\rangle_\mu = \alpha_2 \langle \hat{\phi}, \hat{\phi} \rangle_M \quad (36)$$

In particular, $\hat{M}^{-1} \hat{B}$ is positive, hence it has a positive and symmetric square root $\hat{S}^2 = \hat{M}^{-1} \hat{B}$. Now, for any $\hat{\phi} \in \mathbb{R}^n$ it holds by virtue of (36),

$$\begin{aligned} \langle \hat{\phi}, \hat{\phi} \rangle_M &\leq \frac{1}{\alpha_2} \langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M = \frac{1}{\alpha_2} \langle \hat{S} \hat{\phi}, \hat{S} \hat{\phi} \rangle_M \\ &\leq \frac{1}{\alpha_2^2} \langle \hat{S} \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{S} \hat{\phi} \rangle_M = \frac{1}{\alpha_2^2} \langle \hat{M}^{-1} \hat{B} \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M. \end{aligned} \quad (37)$$

Now we apply the inequality (37) to $\hat{\phi}^* - \hat{\phi}$ and use (35):

$$\epsilon^2 \leq \epsilon_0^2 + \frac{1}{\alpha_2^2} \langle \hat{M}^{-1} c, \hat{M}^{-1} c \rangle_M. \quad (38)$$

Now for some final simplifications, note that the orthogonal projection Q onto D_0 can be written as

$$Q\psi = \sum_{i,j=1}^n \hat{M}_{ij}^{-1} \langle \chi_j, \psi \rangle_\mu \chi_i.$$

Using this we can write

$$\begin{aligned} \langle \hat{M}^{-1} c, \hat{M}^{-1} c \rangle_M &= \sum_{ij} c_i \hat{M}_{ij}^{-1} c_j = \sum_{ij} \langle \chi_i, B\phi_\perp \rangle_\mu \hat{M}_{ij}^{-1} \langle \chi_j, B\phi_\perp \rangle_\mu \\ &= \left\langle \sum_{ij} \hat{M}_{ij}^{-1} \langle \chi_j, B\phi_\perp \rangle_\mu \chi_i, B\phi_\perp \right\rangle_\mu = \langle QB\phi_\perp, B\phi_\perp \rangle_\mu \\ &= \langle QB\phi_\perp, QB\phi_\perp \rangle_\mu \end{aligned}$$

To arrive at the final result, notice that

$$\begin{aligned}
\langle QB\phi_\perp, QB\phi_\perp \rangle_\mu &\leq \left(\sup_{\phi'_\perp \in D^\perp} \frac{\langle QB\phi'_\perp, QB\phi'_\perp \rangle_\mu}{\langle \phi'_\perp, \phi'_\perp \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
&= \left(\sup_{\phi'_\perp \in D^\perp} \frac{\langle QBQ^\perp \phi'_\perp, QBQ^\perp \phi'_\perp \rangle_\mu}{\langle \phi'_\perp, \phi'_\perp \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
&\leq \left(\sup_{\phi' \in V} \frac{\langle QBQ^\perp \phi', QBQ^\perp \phi' \rangle_\mu}{\langle \phi', \phi' \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
&= \|QBQ^\perp\|^2 \langle \phi_\perp, \phi_\perp \rangle_\mu
\end{aligned}$$

Plugging these inequalities into (38) and dividing by ε_0^2 completes the proof. \square

B. Best-approximation error bound

In this appendix, we prove lemma 3:

$$\varepsilon_0 = \|Q^\perp \phi\|_\mu \leq \|P^\perp \phi\|_\mu + \mu(T)^{1/2} [\kappa \|f\|_\infty + 2\|P^\perp \phi\|_\infty].$$

Recall that $\kappa = \sup_{x \in T} \mathbf{E}_x[\tau_{\mathcal{S} \setminus T}]$ and P is the orthogonal projection onto the subspace $V_c = \{v \in L^2(\mathcal{S}, \mu), v = \text{const on every } C_i\} \subset L^2(\mathcal{S}, \mu)$. Note that $P^\perp \phi = 0$ on C . The errors $\|P^\perp \phi\|$ and $\|P^\perp \phi\|_\infty$ measure how constant the solution ϕ is on the core sets. We write $\|\cdot\| = \|\cdot\|_\mu$ throughout the proof for convenience.

Proof. The proof closely follows the proof of theorem (12) in [Sar11]. The first step of the proof is to realize that the committor subspace D where Q projects onto can be written as $D = \{v \in L^2(\mathcal{S}, \mu), v = \text{const on every } C_i, Lv = 0 \text{ on } C\}$. To see this, note that the values v takes on the C_i can be used as boundary values for the Dirichlet problem $Lv = 0$ on T . A linear combination of committor functions is obviously a solution to this problem. But the solution to the Dirichlet problem must be unique, otherwise one can construct a contradiction to the uniqueness of the invariant distribution, see [Sar11].

By definition we have $\|Q^\perp \phi\| \leq \|\phi - I\phi\|$ for every interpolation $I\phi \in D$ of ϕ . With the definition of P from above, we will take $q = I\phi$ such that

$$Lq = 0 \text{ on } T, \quad q = P\phi \text{ on } \mathcal{S} \setminus T. \quad (39)$$

Now $D \subset V$, therefore $q \in V_c$ and $Pq = q$. Therefore (39) is equivalent to

$$PLPq = 0 \text{ on } T, \quad q = P\phi \text{ on } \mathcal{S} \setminus T. \quad (40)$$

Now define $e := P\phi - q$. Then we have

$$PLPe = PLP(P\phi - q) = PLP\phi - PLPq = PL\phi - PLP^\perp \phi - PLPq$$

and by (40) and since $L\phi = f\phi$ on $\mathcal{S} \setminus A \supset T$, we have

$$PLPe = Pf\phi - PLP^\perp\phi \text{ on } T, \quad e = 0 \text{ on } \mathcal{S} \setminus T. \quad (41)$$

Therefore, $e \in E_\Theta = \{v \in L^2(\mathcal{S}, \mu), v = 0 \text{ on } \mathcal{S} \setminus T\}$ and with Θ being the orthogonal projection onto E_Θ , e has to fulfil

$$\Theta PLP\Theta e = \Theta Pf\phi - \Theta PLP^\perp\phi.$$

Since $\Theta P = P\Theta = \Theta$, this can be written as

$$Re := \Theta L\Theta e = \Theta f\phi - \Theta LP^\perp\phi.$$

The operator $R = \Theta L\Theta$ is invertible on E_Θ : If this wasn't the case, there would be a nontrivial solution v to

$$Lv = 0 \text{ on } T, \quad v = 0 \text{ on } \mathcal{S} \setminus T.$$

But the solution to this boundary value problem is again unique, and hence there is only the trivial solution. This gives

$$e = R^{-1}\Theta f\phi - R^{-1}\Theta LP^\perp\phi, \quad (42)$$

and $\|R^{-1}\| = \frac{1}{|\lambda_0|}$ where λ_0 is the principal eigenvalue of R . Due to an estimate by Varadhan we have

$$\frac{1}{|\lambda_0|} \leq \sup_{x \in T} \mathbf{E}_x[\tau_{\mathcal{S} \setminus T}] =: \kappa,$$

see e.g. [Bov09]. To complete the derivation we need to focus on the second term in (42). Since R^{-1} is an operator on E_Θ , we can write it as $R^{-1}\Theta LP^\perp\phi =: \Theta g$, where the function Θg solves

$$\Theta L\Theta g = R\Theta g = \Theta LP^\perp\phi \Leftrightarrow \Theta L[\Theta g - P^\perp\phi] = 0$$

by the definition of R and Θg . Therefore $w := \Theta g - P^\perp\phi$ solves the boundary value problem

$$Lw = 0 \text{ on } T, \quad w = -P^\perp\phi \text{ on } \mathcal{S} \setminus T \quad (43)$$

which implies that $\|w\|_\infty \leq \|P^\perp\phi\|_\infty$, this follows from Dynkin's formula or Lemma 3 in [Sar11]. Finally,

$$\|\Theta g\| \leq \mu(T)^{1/2} \|\Theta g\|_\infty \leq \mu(T)^{1/2} (\|P^\perp\phi\|_\infty + \|w\|_\infty) \leq 2\mu(T)^{1/2} \|P^\perp\phi\|_\infty$$

holds by the triangle inequality and the above considerations. Now focus on the first term in (42). Note that by the maximum principle, ϕ achieves its maximum of 1 on the boundary of $A^c \supset T$, therefore $\max_{x \in T} |\phi(x)| \leq 1$. Then we have

$$\|\Theta f\phi\| \leq \mu(T)^{1/2} \|f\|_\infty \max_{x \in T} |\phi(x)| \leq \mu(T)^{1/2} \|f\|_\infty.$$

Now putting everything together, we arrive at

$$\begin{aligned}\|e\| &\leq \|R^{-1}\|\|\Theta f\phi\| + \|R^{-1}\Theta LP^\perp\phi\| \\ &\leq \kappa\|\Theta f\phi\| + \|\Theta g\| \\ &\leq \mu(T)^{1/2} [\kappa\|f\|_\infty + 2\|P^\perp\phi\|_\infty].\end{aligned}$$

Finally, note that by the triangle inequality

$$\|Q^\perp\phi\| \leq \|\phi - q\| \leq \|\phi - P\phi\| + \|P\phi - q\| = \|P^\perp\phi\| + \|e\|$$

which completes the proof. \square

C. Finite-volume approximation

In this section we show (25), confirming that the Galerkin projection of L onto step functions gives the finite volume approximation discussed in [LMHS11]. Recall the definitions of S_{ij} , h_{ij} and A_i given in figure (3). We use the divergence representation $L\phi = e^{\beta V}\nabla \cdot (e^{-\beta V}\nabla\phi)$ with $\beta = \epsilon^{-1}$ and calculate

$$\begin{aligned}\langle \chi_i, L\chi_j \rangle_\mu &= \beta^{-1} \int_{\mathcal{S}} \chi_i e^{\beta V} \nabla \cdot (e^{-\beta V} \nabla \chi_j) e^{-\beta V} dx \\ &= \beta^{-1} \int_{A_i} \nabla \cdot (e^{-\beta V} \nabla \chi_j) dx \\ &= \beta^{-1} \int_{\partial A_i} e^{-\beta V} (\nabla \chi_j) \cdot \nu ds\end{aligned}$$

where ν is the surface normal vector field of ∂A_i . We write the integral over ∂A_i as a sum over surface integrals over $S_{ij'}$ where j' ranges over the set $\{i_l\}$ of neighbours of i and approximate the surface integrals by a point evaluation of the integrand at the midpoint $\bar{x}_{ij'}$ times the area of $S_{ij'}$. That gives

$$\langle \chi_i, L\chi_j \rangle_\mu \approx \beta^{-1} \sum_{j' \in \{i_l\}} m(S_{ij'}) e^{-\beta V(\bar{x}_{ij'})} (\nabla \chi_j \cdot \nu)|_{x=\bar{x}_{ij'}}$$

Now we can approximate the directional derivative of χ_j using a two-sided finite difference:

$$\nabla \chi_j \cdot \nu \Big|_{x=\bar{x}_{ij'}} = \frac{\nabla \chi_j \cdot h_{i,j'}}{m(h_{i,j'})} \Big|_{x=\bar{x}_{ij'}} \approx \frac{\chi_j(x_{j'}) - \chi_j(x_i)}{m(h_{ij'})} = \frac{\delta_{jj'} - 0}{m(h_{ij'})}.$$

Hence in the sum over neighbours of i , only j survives. Now we put everything together:

$$\langle \chi_i, L\chi_j \rangle_\mu \approx \beta^{-1} \frac{m(S_{ij})}{m(h_{ij})} e^{-\beta V(\bar{x}_{ij})}.$$

Finally, we divide by π_i using (26):

$$K_{ij} = \frac{1}{\pi_i} \langle \chi_i, L\chi_j \rangle_\mu \approx \frac{1}{\Delta_{ij}} e^{-\beta(V(\bar{x}_{ij})-V(x_i))}, \quad \Delta_{ij}^{-1} = \beta^{-1} \frac{m(S_{ij})}{m(h_{ij})m(A_i)}.$$

which confirms (25).

D. Markov chain approximations

We now show (29). Let i and j be nearest neighbours, and $\beta = \epsilon^{-1}$. For a regular d -dimensional lattice with lattice spacing h ,

$$\Delta_{ij} = \beta \frac{m(h_{ij})m(\Omega_i)}{m(S_{ij})} = \beta \frac{hh^d}{h^{(d-1)}} = \beta h^2.$$

Therefore, G as given by (25) simplifies to

$$G_{ij} = \frac{1}{\beta h^2} e^{-\beta(V(\bar{x}_{ij})-V(x_i))}.$$

We introduce the function $\hat{W}_v(i) = -\beta^{-1} \log v(i)$. Then, for neighbours i, j ,

$$G_{ij}^v = \frac{1}{\beta h^2} e^{-\beta(V(\bar{x}_{ij})-V(x_i)+\hat{W}_v(j)-\hat{W}_v(i))}.$$

Now we specialise to the one-dimensional case, thus $j = i \pm 1$. We write $V(\bar{x}_{i,i\pm 1}) - V(x_i) = \pm \frac{h}{2} \nabla V(x_i) + \mathcal{O}(h^2)$. Expanding the exponential gives

$$\begin{aligned} G_{i,i\pm 1}^v &= \frac{1}{\beta h^2} \left(1 - \frac{\beta h}{2} (\pm \nabla V(x_i)) - \beta (\hat{W}_v(i \pm 1) - \hat{W}_v(i)) + \mathcal{O}(h^2) \right) \\ &= \frac{1}{h^2} \left(\beta^{-1} - \frac{h}{2} \left(\pm \nabla V(x_i) + 2 \frac{\hat{W}_v(i \pm 1) - \hat{W}_v(i)}{h} \right) + \mathcal{O}(h^2) \right) \\ &= \frac{1}{h^2} \left(\beta^{-1} \mp \frac{h}{2} (\nabla V(x_i) - \alpha_v^\pm(i)) + \mathcal{O}(h^2) \right) \end{aligned}$$

with the definition

$$\alpha_v^\pm(i) := \pm \left(-2 \frac{\hat{W}_v(i \pm 1) - \hat{W}_v(i)}{h} \right).$$

Now consider the difference between $\alpha_v^+(i)$ and $\alpha_v^-(i)$:

$$\alpha_v^+(i) - \alpha_v^-(i) = -2h \frac{\hat{W}_v(i+1) - 2\hat{W}_v(i) + \hat{W}_v(i-1)}{h^2}$$

Assuming that \hat{W}_v converges to a twice differentiable function, the difference between $\alpha_v^+(i)$ and $\alpha_v^-(i)$ is of order h . In other words we may write $\alpha_v^\pm(i) = \alpha_v(i) + \mathcal{O}(h)$ where

$$\alpha_v(i) = \frac{1}{2}(\alpha_v^+(i) + \alpha_v^-(i)) = \frac{1}{h\beta} (\log v(i+1) - \log v(i-1)).$$

Then

$$G_{i,i\pm 1}^v = \frac{1}{h^2} \left(\beta^{-1} \mp \frac{h}{2} (\nabla V(x_i) - \alpha_v(i)) + \mathcal{O}(h^2) \right)$$

which confirms (29). Now we show (30). We use the representation (18) for $k^v(i)$:

$$\begin{aligned} \beta k^v(i) &= \sum_{j \neq i} G_{ij} \left\{ \frac{v(j)}{v(i)} \left[\log \frac{v(j)}{v(i)} - 1 \right] + 1 \right\} \\ &= G_{i,i+1} \frac{v(i+1)}{v(i)} \left(\log \frac{v(i+1)}{v(i)} - 1 \right) + G_{i,i-1} \frac{v(i-1)}{v(i)} \left(\log \frac{v(i-1)}{v(i)} - 1 \right) - G_{ii} \end{aligned}$$

Now we write this in terms of the shorthands $\alpha^\pm := \pm h^{-1} \log \frac{v(i\pm 1)}{v(i)}$. Notice that $\alpha^\pm = \frac{\beta}{2} \alpha_v^\pm(i)$ and use the formula for G above:

$$\begin{aligned} \beta k^v(i) &= \frac{1}{\beta h^2} \left[e^{h\alpha^+} (h\alpha^+ - 1) - e^{-h\alpha^-} (h\alpha^- + 1) + 2 \right] \\ &\quad - \frac{\nabla V(x_i)}{2h} \left[e^{h\alpha^+} (h\alpha^+ - 1) + e^{-h\alpha^-} (h\alpha^- + 1) \right] + \mathcal{O}(h) \\ &= \frac{1}{\beta h^2} [-1 - 1 + 2] + \frac{1}{\beta h} [\alpha^+ - \alpha^- - \alpha^+ + \alpha^-] + \frac{1}{2\beta} [(\alpha^+)^2 + (\alpha^-)^2] \\ &\quad - \frac{\nabla V(x_i)}{2h} [-1 + 1] - \frac{\nabla V(i)}{2} [\alpha^+ - \alpha^+ + \alpha^- - \alpha^-] + \mathcal{O}(h) \\ &= \frac{1}{2\beta} [(\alpha^+)^2 + (\alpha^-)^2] + \mathcal{O}(h) \\ &= \frac{\beta}{4} \alpha_v^2(i) + \mathcal{O}(h) = \frac{\beta}{4} \alpha_v^2(i) + \mathcal{O}(h). \end{aligned}$$

This confirms (18). In the second step, we have used Taylor expansions of $e^{h\alpha^\pm}$ up to second order. In the last step, we have used $\alpha_v^\pm(i) = \alpha_v(i) + \mathcal{O}(h)$.

E. Sampling of the discretized running cost

We show the sampling formula (34) for F :

$$F_{ij} = \mathbf{E}_\mu \left[f(X_t) \chi_{\{\tilde{X}_t^+ = j\}} \middle| \tilde{X}_t^- = i \right].$$

Recall that since the dynamics is reversible, $\chi_i(x) = \mathbf{P}(\tilde{X}_t^\pm = i | X_t = x)$ with \tilde{X}_t^\pm being the forward and backward milestoning processes defined in Section 3.2. Then,

$$\begin{aligned}
\hat{F}_{ij} &= \int f(x) \chi_j(x) \chi_i(x) \mu(x) dx = \int f(x) \mathbf{P}(\tilde{X}_t^+ = j | X_t = x) \mathbf{P}(\tilde{X}_t^- = i, X_t = x) dx \\
&= \int f(x) \mathbf{P}(\tilde{X}_t^+ = j, \tilde{X}_t^- = i, X_t = x) dx \\
&= \int f(x) \mathbf{P}(\tilde{X}_t^+ = j, \tilde{X}_t^- = i | X_t = x) \mathbf{P}(X_t = x) dx \\
&= \int f(x) \mathbf{E} \left(\chi_{\{\tilde{X}_t^+ = j, \tilde{X}_t^- = i\}} \middle| X_t = x \right) \mathbf{P}(X_t = x) dx \\
&= \mathbf{E}_\mu \left[f(X_t) \chi_{\{\tilde{X}_t^+ = j, \tilde{X}_t^- = i\}} \right].
\end{aligned}$$

This completes the proof.

References

- [AK11] Erik Asplund and Thorsten Klüner. Optimal control of open quantum systems applied to the photochemistry of surfaces. *Phys. Rev. Lett.*, 106:140404, 2011.
- [AMP⁺12] Steffen Arnrich, Alexander Mielke, Mark A. Peletier, Giuseppe Savaré, and Marco Veneroni. Passing to the limit in a wasserstein gradient flow: from diffusion to reaction. *Calc. Var. Partial Diff.*, 44(3-4):419–454, 2012.
- [Ber11] N. Berglund. Kramers’ law: Validity, derivations and generalisations. *ArXiv e-prints*, June 2011.
- [BH01] M. Boulbrachene and M. Haiour. The finite element approximation of hamilton-jacobi-bellman equations. *Comput. Math. Appl.*, 41(7-8):993 – 1007, 2001.
- [BJ07] G. Barles and E. Jakobsen. Error bounds for monotone approximation schemes for parabolic Hamilton-Jacobi-Bellman equations. *Math. Comput.*, 76:1861–1893, 2007.
- [BMZ09] O. Bokanowski, S. Maroso, and H. Zidani. Some convergence results for howard’s algorithm. *SIAM J. Numer. Anal.*, 47(4):3001–3026, 2009.
- [Bov09] A. Bovier. *Methods of Contemporary Statistical Mechanics*. Metastability. Springer, 2009.
- [BP79] A. Berman and R.J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, 1979.

- [Bra12] M. Braack. *Finite Elemente*. Lecture Notes, Universität Kiel, 2012.
- [BS91] G. Barles and P.E. Souganidis. Convergence of approximation schemes for fully nonlinear equations. *Asymptotic Anal.*, 4:271–283, 1991.
- [CES⁺11] J. D. Chodera, P. J. Elms, W. C. Swope, J.-H. Prinz, S. Marqusee, C. Bustamante, F. Noé, and V. S. Pande. A robust approach to estimating rates from time-correlation functions. *arXiv: 1108.2304*, August 2011.
- [DD70] Jr. Douglas, Jim and Todd Dupont. Galerkin methods for parabolic equations. *SIAM J. Numer. Anal.*, 7(4):575–626, 1970.
- [DSS10] N. Djurdjevac, M. Sarich, and Ch. Schütte. On Markov state models for metastable processes. *Proceedings of the ICM Hyderabad, India*, 2010.
- [DSS12] N. Djurdjevac, M. Sarich, and Ch. Schütte. Estimating the eigenvalue error of Markov state models. *Multiscale Model. Simul.*, 10 (1):61–81, 2012.
- [DSW12] P. Dupuis, K. Spiliopoulos, and H. Wang. Importance sampling for multiscale diffusions. *Multiscale Model. Simul.*, 10(1):1–27, 2012.
- [DW04] Paul Dupuis and Hui Wang. Importance sampling, large deviations, and differential games. *Stochastics and Stochastic Reports*, 76(6):481–508, 2004.
- [FE04] A.K. Faradjian and R. Elber. Computing time scales from reaction coordinates by milestoning. *J. Chem. Phys.*, 120:10880–10889, 2004.
- [Fle77] Wendell H. Fleming. Exit probabilities and optimal stochastic control. *Appl. Math. Optim.*, 4:329–346, 1977.
- [FM95] Wendell H. Fleming and William M. McEneaney. Risk-sensitive control on an infinite time horizon. *SIAM J. Control Optim.*, 33:1881–1915, 1995.
- [FS06] Wendell H. Fleming and Halil M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer, 2006.
- [FW84] M.I. Freidlin and A.D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, New York, 1984.
- [Gar96] C.W. Gardiner. *Handbook of Stochastic Methods: For Physics, Chemistry and the Natural Sciences*. Springer, 1996.
- [GS75] I.I. Gikhman and A.V. Skorokhod. *The Theory of Stochastic Processes II*. Springer-Verlag, New York, 1975.

- [HBS⁺13] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and Ch. Schütte. Characterization of rare events in molecular dynamics. *To appear in Entropy*, 2013.
- [Hei95] Ph. Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Trans. Model. Comput. Simul.*, 5:43–85, 1995.
- [HS12] C. Hartmann and C. Schütte. Efficient rare event simulation by optimal nonequilibrium forcing. *Journal of Statistical Mechanics: Theory and Experiment*, 11:4, November 2012.
- [Jam92] Matthew James. Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Math. Control Signals Systems*, 5:401–417, 1992.
- [Kap05] Hilbert J. Kappen. Path integrals and symmetry breaking for optimal control theory. *J. Stat. Mech. Theor. Exp.*, 2005(11):P11011, 2005.
- [KD92] H.J. Kushner and P.G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, 1992.
- [KV01] K. Kunisch and S. Volkwein. Galerkin proper orthogonal decomposition methods for parabolic problems. *Numer. Math.*, 90(1):117–148, 2001.
- [LMHS11] J.C. Latorre, Ph. Metzner, C. Hartmann, and Ch. Schütte. A structure-preserving numerical discretization of reversible diffusions. *Commun. Math. Sci.*, 9:1051–1072, 2011.
- [Øks03] Bernt K. Øksendal. *Stochastic Differential Equations: An Introduction With Applications*. Springer, 2003.
- [PBB10] Vijay S. Pande, Kyle Beauchamp, and Gregory R. Bowman. Everything you wanted to know about markov state models but were afraid to ask. *Methods*, 52(1):99 – 105, 2010.
- [PK02] José P. Palao and Ronnie Kosloff. Quantum computing by an optimal control algorithm for unitary transformations. *Phys. Rev. Lett.*, 89:188301, 2002.
- [PMR96] Paolo Pra, Lorenzo Meneghini, and Wolfgang J. Runggaldier. Connections between stochastic control and dynamic games. *Mathematics of Control, Signals and Systems*, 9:303–326, 1996.
- [PSV12] M. Peletier, G. Savaré, and M. Veneroni. Chemical reactions as γ -limit of diffusion. *SIAM Review*, 54(2):327–352, 2012.
- [PWS⁺11] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Fischbach, M. Held, J. D. Chodera, Ch. Schütte, and F. Noé. Markov models of molecular kinetics: Generation and validation. *J. Chem. Phys.*, 134:174105, 2011.

- [RdVRMK00] Herschel Rabitz, Regina de Vivie-Riedle, Marcus Motzkus, and Karl Kompa. Whither the future of controlling quantum phenomena? *Science*, 288(5467):824–828, 2000.
- [Roe08] Susanna Roebnitz. *Statistical Error Estimation and Grid-free Hierarchical Refinement in Conformation Dynamics*. PhD thesis, FU Berlin, 2008.
- [Sar11] M. Sarich. *Projected Transfer Operators*. PhD thesis, FU Berlin, 2011.
- [She82] Shuenn Jyi Sheu. *Stochastic control and its application to large deviation theory*. PhD dissertation thesis, Brown University, 1982.
- [She85] S. Sheu. Stochastic control and exit probabilities of jump processes. *SIAM Journal on Control and Optimization*, 23(2):306–328, 1985.
- [SNL⁺11] Ch. Schütte, F. Noe, J. Lu, M. Sarich, and E. Vanden-Eijnden. Markov state models based on milestoning. *J. Chem. Phys.*, 134 (20):204105, 2011.
- [SNS10] M. Sarich, F. Noe, and Ch. Schütte. On the approximation quality of Markov state models. *Multiscale Model. Simul.*, 8 (4):1154–1177, 2010.
- [SR95] Asmussen S. and R. Rubinstein. Steady-state rare events simulation in queueing models and its complexity properties. In J. Dhashalow, editor, *Advances in Queueing*, pages 429–462. CRC Press, Boca Raton, 1995.
- [Sta04] Henrik Stapelfeldt. Laser aligned molecules: Applications in physics and chemistry. *Physica Scripta*, 2004(T110):132–136, 2004.
- [Ste10] Andreas Steinbrecher. Optimal control of robot guided laser material treatment. In Alistair D. Fitt, John Norbury, Hilary Ockendon, and Eddie Wilson, editors, *Progress in Industrial Mathematics at ECMI 2008*, pages 505–511. Springer Berlin Heidelberg, 2010.
- [SWH12] Ch. Schütte, S. Winkelmann, and C. Hartmann. Optimal control of molecular dynamics using Markov state models. *Math. Program. (Series B)*, 134:259–282, 2012.
- [Tod09] E. Todorov. Efficient computation of optimal actions. *Proc. Natl. Acad. Sci. USA*, 106(28):11478–11483, 2009.
- [VE06] E. Vanden-Eijnden. Transition path theory. *Lect. Notes Phys.*, 703:439–478, 2006.
- [ZLPH13] W. Zhang, J.C. Latorre, G.A. Pavliotis, and C. Hartmann. Optimal control of multiscale systems using reduced-order models. *submitted to J. Comput. Dynamics*, 2013.