university of groningen

University Medical Center Groningen

University of Groningen

Application of adaptive multilevel splitting to high-dimensional dynamical systems

Baars, S.; Castellana, D.; Wubs, F.W.; Dijkstra, H.A.

Published in:
Journal of computational physics

DOI:
10.1016/j.jcp.2020.109876

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

Link to publication in University of Groningen/UMCG research database

Citation for published version (APA):
Baars, S., Castellana, D., Wubs, F. W., & Dijkstra, H. A. (2021). Application of adaptive multilevel splitting to high-dimensional dynamical systems. *Journal of computational physics*, *424*, [109876]. https://doi.org/10.1016/j.jcp.2020.109876

Copyright
Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the “Taverne” license. More information can be found on the University of Groningen website: https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment.

Take-down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): http://www.rug.nl/research/portal. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Download date: 19-11-2022

Journal of Computational Physics 424 (2021) 109876

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Application of adaptive multilevel splitting to high-dimensional dynamical systems

S. Baars [a,*], D. Castellana [b], F.W. Wubs [a], H.A. Dijkstra [b,c]

[a] Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, P.O. Box 407, 9700 AK Groningen, the Netherlands
[b] Institute for Marine and Atmospheric Research Utrecht, Department of Physics, Utrecht University, Princetonplein 5, 3584 CC Utrecht, the Netherlands
[c] Centre for Complex Systems Studies, Utrecht University, Leuvenlaan 4 3584 CE Utrecht, the Netherlands

ARTICLE INFO

Article history:
Received 13 December 2019
Received in revised form 13 August 2020
Accepted 25 September 2020
Available online 30 September 2020

Keywords:
Rare transitions
Multilevel splitting
Model order reduction
Stochastic dynamical systems
Ocean circulation

ABSTRACT

Stochastic nonlinear dynamical systems can undergo rapid transitions relative to the change in their forcing, for example due to the occurrence of multiple equilibrium solutions for a specific interval of parameters. In this paper, we modify one of the methods developed to compute probabilities of such transitions, Trajectory-Adaptive Multilevel Sampling (TAMS), to be able to apply it to high-dimensional systems. The key innovation is a projected time-stepping approach, which leads to a strong reduction in computational costs, in particular memory usage. The performance of this new implementation of TAMS is studied through an example of the collapse of the Atlantic Ocean Circulation.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

The Atlantic Meridional Overturning Circulation (AMOC) is an important component of the climate system because of its associated meridional heat transport. It is well known that the AMOC is sensitive to surface freshwater flux perturbations [1]. Freshening of the surface waters in the Nordic and Labrador Seas diminishes the production of deepwater that feeds the deep southward branch of the AMOC. The weakening of the AMOC leads to reduced northward salt transport freshening the northern North Atlantic and amplifying the original freshwater perturbation. In a hierarchy of ocean models, transitions in AMOC flows are found due to the occurrence of multiple equilibria (see e.g. chapter 6 in [2]). However, it is currently unknown whether the present-day AMOC is in a multiple equilibrium regime, and if so, what the probability is that it will undergo a collapse within the next 100 or 1000 years [3].

Probabilities of rare events, and in particular of transitions between different equilibrium states of a certain stochastic system, are often very difficult to compute. In the framework of Large Deviation Theory, there are several results available [4]. When the deterministic system can be described in terms of a potential, the transition rate for stochastic transitions induced by white additive noise is given by the Eyring–Kramers formula [5,6]. Extensions of this formula for non-gradient systems have been found [7], and although these can be very useful for determining the location of possible transition paths [8], directly computing transition probabilities using such formulas for high-dimensional systems is presently not feasible. Furthermore, these results allow to compute only transition rates, i.e. the probability that a transition occurs in a unit time. A more interesting quantity is the probability that the system undergoes a transition within a certain time interval.

* Corresponding author.
E-mail addresses: s.baars@rug.nl (S. Baars), d.castellana@uu.nl (D. Castellana), f.w.wubs@rug.nl (F.W. Wubs), h.a.dijkstra@uu.nl (H.A. Dijkstra).

https://doi.org/10.1016/j.jcp.2020.109876
0021-9991/© 2020 Elsevier Inc. All rights reserved.

The connection between these two quantities is not straightforward, unless specific assumptions are made on the noise generating the transitions: for example, if the process can be assumed to be Poisson, then the probability of undergoing a transition within a certain time follows the cumulative distribution function of the exponential distribution with the rate parameter being the transition rate [9].

From a totally different perspective, one could think of computing transition probabilities with standard Monte Carlo methods. This can be done by computing a large number of trajectories, and then counting the number of observed transitions. This approach is, however, not suitable for high-dimensional systems, especially when the probabilities are small. Therefore, specialized methods, such as Genealogical Particle Analysis (GPA) [10–12], Adaptive Multilevel Splitting (AMS) [13–15], and Trajectory-Adaptive Multilevel Sampling (TAMS) [9] have been developed. The idea behind these methods is similar: the algorithms try to push trajectories in the direction of the transition to make sure transitions actually occur, also for smaller sample sizes, while still being able to compute the transition probability. Rare transitions in a barotropic β-plane model have been studied by means of AMS in [8]. Transitions in the two-dimensional barotropic quasi-geostrophic equations were also studied in [16,17], but in that case only the most likely transition path was determined by means of the minimum action method [18–21]. To our knowledge, actual transition probabilities have never been computed for high-dimensional problems on the scale of our target problem, which is a global ocean model with a resolution of at least 2 degrees ($\sim 1.4 \cdot 10^6$ degrees of freedom). In this paper, however, we only present and validate the methodology suitable for the target problem, but do not actually compute transition probabilities for this problem.

Methods for finding transition probabilities have two main limitations. The first is that they always require some form of time integration, which is very expensive for high-dimensional systems. This is especially the case when an implicit time-stepping method is used, since then in every time step a linear system has to be solved. The other is that for methods based on AMS, a large number of states at different time steps have to be stored in memory in order to keep the computational cost minimal. We therefore propose a projected time-stepping method, for which in every time step, only a small projected linear system has to be solved, and for which only the projected states have to be stored, which are much smaller dimensional than the original states. The vectors that are used for the projection are the empirical orthogonal functions (EOFs) that are commonly used often easily computed for many geophysical models. The most dominant EOFs can be found by computing the eigenvectors belonging to the largest eigenvalues of the covariance matrices at both stable steady states. These vectors can be obtained from basis vectors of the low-rank solution of a generalized Lyapunov equation [22]. A theoretical background is given in [23].

The projected time-stepping method is related to the Karhunen–Loève transform in the context of stochastic processes [24] or the proper orthogonal decomposition in the context of fluid dynamics [25,26]. The difference is that we do not use an ensemble, but instead use both stable steady states and covariance matrices at those steady states to directly compute the basis vectors.

Model order reduction techniques such as the projection method described above have been applied in many fields when studying transitions [27–32], but to our knowledge not with the multilevel splitting methods that we need to compute transition rates in non-gradient systems. Note that there are many other available model reduction techniques that could be used with the methodology described in this paper such as temporal EOFs [29] and machine learning [32].

In this paper, TAMS, its optimization and the projected time-stepping are described in section 2. This is followed by section 3 where we apply it to a spatially two-dimensional model of the AMOC, governed by a set of stochastic partial differential equations. Summary and discussion are provided in section 4.

2. Methodology

Consider a system of stochastic differential algebraic equations (SDAEs) of the form

$$M(\mathbf{p})\mathrm{d}\mathbf{X}_t = F(\mathbf{X}_t;\mathbf{p})\mathrm{d}t + g(\mathbf{X}_t;\mathbf{p})\mathrm{d}\mathbf{W}_t \tag{1}$$

where $\mathbf{p} \in \mathbb{R}^{n_p}$ indicates a vector of parameters, $M \in \mathbb{R}^{n\times n}$ is usually referred to as the mass matrix, $\mathbf{X}_t \in \mathbb{R}^n$, $F : \mathbb{R}^n \to \mathbb{R}^n$, $g \in \mathbb{R}^{n\times n_w}$ and $\mathbf{W}_t \in \mathbb{R}^{n_w}$ is a vector of n_w independent standard Brownian motions [33]. Assume that the corresponding deterministic system

$$M(\mathbf{p})\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = F(\mathbf{x};\mathbf{p}) \tag{2}$$

has two stable steady states $\bar{\mathbf{x}}_A$ and $\bar{\mathbf{x}}_B$ for fixed values of parameters $\mathbf{p}$ and define two sets, A and B, which respectively contain $\bar{\mathbf{x}}_A$ and $\bar{\mathbf{x}}_B$. Our starting point will be the state $\bar{\mathbf{x}}_A$. Once we perturb it with noise, the state will wander around the equilibrium, with a motion depending on the characteristics of the noise. If the perturbation is strong enough, the system may undergo a transition and end up in the other equilibrium state $\bar{\mathbf{x}}_B$. We would like to compute the probability that such a transition occurs within a certain interval of time.

We compute bifurcation diagrams for (2) using the pseudo-arclength continuation method [34], in which branches of steady states $(\mathbf{x}(s), \mathbf{p}(s))$ are parameterized by an arclength parameter s and Newton's method is used on an augmented system to converge onto the steady states. Pseudo-arclength continuation has the advantage that it can also compute the unstable steady state $\bar{\mathbf{x}}_C$ that is located between $\bar{\mathbf{x}}_A$ and $\bar{\mathbf{x}}_B$.

2.1. Trajectory-adaptive multilevel sampling

The idea behind the algorithm is based on the fact that, once the trajectories leave A, there are more trajectories that come back to A than there are ones that reach B. A method that makes use of this is called the Adaptive Multilevel Splitting method (AMS) [13,14]. It was inspired by multilevel splitting methods which date back to [35] and [36]. The multilevel splitting methods are all based on the same idea of discarding trajectories that do not reach B and splitting (or branching) trajectories that are closer to B. This makes it so that the probability that a trajectory reaches B keeps increasing, which is why this method is more efficient than a naive method (such as Monte Carlo).

How close a vector $\mathbf{x}$ is to B is defined by a so-called reaction coordinate, or score function, which is a smooth one-dimensional function

$$\phi : \mathbb{R}^n \to \mathbb{R}. \tag{3}$$

In this paper, we assume that the reaction coordinate should satisfy

$$|\nabla\phi(\mathbf{x})| \neq 0, \ \forall x \in \mathbb{R}^n \backslash (A \cup B), \tag{4a}$$

$$A \subset \{\mathbf{x} \in \mathbb{R}^n : \phi(\mathbf{x}) < z_{\min}\}, \tag{4b}$$

$$B \subset \{\mathbf{x} \in \mathbb{R}^n : \phi(\mathbf{x}) > z_{\max}\}, \tag{4c}$$

where $z_{\min} < z_{\max}$ are two given real numbers. These properties were slightly adapted from [37]. For multi-dimensional problems, we propose some additional properties to make sure the method actually converges towards B, and not somewhere completely different which has the same value of ϕ. These properties are

$$\{\mathbf{x} \in \mathbb{R}^n : \phi(\mathbf{x}) = \inf\{\phi(\mathbf{y}) : \mathbf{y} \in \mathbb{R}^n\}\} \subset A, \tag{5a}$$

$$\{\mathbf{x} \in \mathbb{R}^n : \phi(\mathbf{x}) = \sup\{\phi(\mathbf{y}) : \mathbf{y} \in \mathbb{R}^n\}\} \subset B. \tag{5b}$$

Since the gradient $\nabla\phi$ is not allowed to be zero outside of A and B, this means that the reaction coordinate is always increasing towards B and decreasing towards A.

The efficiency of multilevel splitting methods is based on the choice of the reaction coordinate. With α indicating the transition probability, N the number of trajectories and μ_α, σ_α the sample mean and sample variance of the estimator of α, the relative error $\epsilon_\alpha = \sigma_\alpha/\mu_\alpha$ of standard Monte Carlo methods converges with $\mathcal{O}(\sqrt{1/(\alpha N)})$ [38], whereas multilevel splitting methods show behavior between $\mathcal{O}(\sqrt{\log(1/\alpha)/N})$ (optimal) and $\mathcal{O}(\sqrt{1/(\alpha N)})$ (worst case) [39,9]. Choosing a better reaction coordinate means that we get closer to the optimal convergence behavior. The optimal behavior, however, is only attained when the optimal reaction coordinate, which is also referred to as the committor, is used, in which case the proportionality constant is equal to one [14]. The committor can be obtained by solving a backward Fokker–Planck equation, but in practice this can never be done for high-dimensional systems [40].

In this paper, we define our reaction coordinate to be

$$\phi(\mathbf{x}) = \eta - \eta e^{-\gamma d_A^2} + (1-\eta)e^{-\gamma d_B^2}, \tag{6}$$

where $d_A = \|\mathbf{x} - \bar{\mathbf{x}}_A\|_2/\|\bar{\mathbf{x}}_A - \bar{\mathbf{x}}_B\|_2$ is the normalized distance between $\mathbf{x}$ and $\bar{\mathbf{x}}_A$, and $d_B = \|\mathbf{x} - \bar{\mathbf{x}}_B\|_2/\|\bar{\mathbf{x}}_A - \bar{\mathbf{x}}_B\|_2$ is the normalized distance between $\mathbf{x}$ and $\bar{\mathbf{x}}_B$ and γ is a real positive constant which we choose to be 8. To give equal weight to both stable steady states, a good choice for η would be 0.5. When the unstable state $\bar{\mathbf{x}}_C$ associated with the saddle-node bifurcation is known, however, a good choice would be $\eta = \|\bar{\mathbf{x}}_C - \bar{\mathbf{x}}_A\|_2/\|\bar{\mathbf{x}}_B - \bar{\mathbf{x}}_A\|_2$. This is the normalized distance between the unstable steady state $\bar{\mathbf{x}}_C$ and the stable steady state $\bar{\mathbf{x}}_A$. Using this value of η makes sure that more weight is given to the stable steady state farthest away from the unstable steady state, meaning that a trajectory is less likely to keep circling around it.

AMS can be used to first compute the mean first passage time. The transition probability can then be obtained from the mean first passage time under the assumptions that the transitions are a Poisson process and that they are instantaneous [9].

Trajectory-Adaptive Multilevel Sampling (TAMS) [9] is a slightly altered version of AMS, in which, instead of first computing the mean first passage time, the transition probability is computed directly when given a maximum time $T_{\max}$, without requiring any of the aforementioned assumptions [41]. In TAMS, there is an optional maximum number of iterations $k_{\max}$ that can be set, which allows us to stop even before all trajectories have reached B. We now explain step by step how to compute the transition probability using TAMS.

1. First generate N independent trajectories $(\mathbf{x})^{(1)}, \ldots, (\mathbf{x})^{(N)}$ with time step Δt that start in A until $t = T_{\max}$ is reached, or the trajectory ends up in B. Here $(\mathbf{x})^{(i)} = (\mathbf{x}_0^{(i)}, \mathbf{x}_1^{(i)}, \ldots, \mathbf{x}_{N_t}^{(i)})$ is a trajectory of length $N_t = T_{\max}/\Delta t$.
2. Each trajectory $(\mathbf{x})^{(i)}$, $i = 1, \ldots, N$, has a certain maximum value of the reaction coordinate (or maximum distance from A), which we define to be Q_i. Set $k = 1$, $w_0 = 1$.
3. Take $L = \{j : Q_j = \inf\{Q_i : i \in \{1, \ldots, N\}\}\}$, which are the indices for which the maximum value of the reaction coordinate is minimal, and take ℓ_k to be the number of elements in L. Since the trajectories $\{(\mathbf{x})^{(j)} : j \in L\}$ have the smallest value of Q_i, all other trajectories have some j for which $\phi(\mathbf{x}_j) \geq Q_l$ for all $l \in L$.

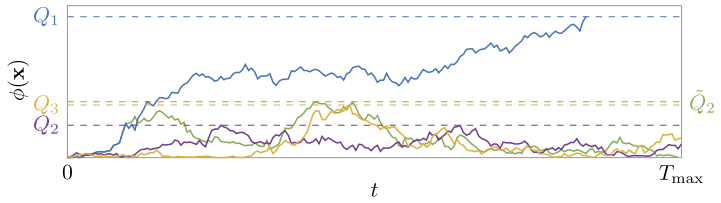

Fig. 1. Example run of TAMS with $N = 3$ applied to a two-dimensional double well potential, for which we take $F(\mathbf{x}) = -\nabla V(\mathbf{x})$ with $V(x, y) = \frac{1}{4}x^4 - \frac{1}{2}x^2 + y^2$ and $g(\mathbf{x}) = 0.4$. The second trajectory (purple) has the lowest maximum value of the reaction coordinate Q_2 after the first iteration. This trajectory is then branched from a random other trajectory, in this case the first trajectory (blue). The new trajectory (green) has a new maximum value of the reaction coordinate $\tilde{Q}_2$, which is larger than Q_2. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

4. Set $w_k = \left(1 - \frac{\ell_k}{N}\right) w_{k-1}$. For all $l \in L$, repeat steps 5-7.
5. Select a trajectory $(\mathbf{x})^{(r)}$ uniformly at random with r from $\{1, \dots, N\} \backslash L$, and set $(\tilde{\mathbf{x}})^{(l)} = (\mathbf{x}_0^{(r)}, \dots, \mathbf{x}_{j_{\min}}^{(r)})$, where $j_{\min}$ is the smallest value for which $\phi(\mathbf{x}_{j_{\min}}^{(r)}) \geq Q_l$.
6. Generate the rest of the trajectory starting from $\mathbf{x}_{j_{\min}}^{(r)}$ until again you reach either $t = T_{\max}$ or B. This trajectory has a new maximum value of the reaction coordinate $\tilde{Q}_l$, which is always greater than or equal to Q_l.
7. Set $(\mathbf{x})^{(l)} = (\tilde{\mathbf{x}})^{(l)}$ and $Q_l = \tilde{Q}_l$.
8. Repeat steps 3-7 with $k = k + 1$ until $Q_i \geq z_{\max}$, $\forall\, i = 1, \dots, N$ or $k = k_{\max}$.

An illustration of the workings of TAMS is given in Fig. 1.

The weights w_k, computed in each step, represent the probability of a trajectory reaching iteration $k + 1$. For instance, consider the case in which we start with 100 trajectories, and we have 2 trajectories for which the maximum value of the reaction coordinate is minimal. Since these two trajectories are eliminated, the probability of a trajectory reaching iteration 2 is $1 - 2/100 = 98/100$. We repeat this process, multiplying the probabilities that we find in every step. This gives us an unbiased estimator of the transition probability

$$\hat{\alpha}_N = \frac{N_B w_K}{N} = \frac{N_B}{N} \prod_{k=0}^{K} \left(1 - \frac{\ell_k}{N}\right), \tag{7}$$

where K is the number of iterations it took for all trajectories to converge, and N_B is the number of trajectories that reached B [9].

2.2. Optimizations

Say we have a problem of dimension $n = 10^4$ with a time step $\Delta t = 10^{-2}$, a maximum time $T = 10^3$, and $N = 10^3$ trajectories, which is not at all unreasonable for the problems that we want to solve. In this case we would need at least 7.28 TB of memory to store all of the states that we compute in double precision, which is a very large amount of memory. Fortunately, some simple optimizations exist to help with this.

First, it is important to observe that the only place where we actually use a state is when we branch a trajectory. Other than in this place, we only use the times to compute the quantities that we need. This means that we can discard any state $\mathbf{x}$ for which $\phi(\mathbf{x}) < Q_l$, since these will never be used for branching. Discarding these states can be done for instance every time a trajectory is branched, every so many iterations, or even in every TAMS iteration.

Secondly, since we only use the first $\mathbf{x}$ for which $\phi(\mathbf{x}) \geq Q_l$, we only have to store $\{\mathbf{x}_j \in (\mathbf{x}) \ : \ \phi(\mathbf{x}_j) > \sup\{\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_{j-1})\}\}$, where $(\mathbf{x})$ denotes a full trajectory. This means that if we iteratively determine the value of Q, we only store a state $\mathbf{x}$ if $\phi(\mathbf{x}) > Q$. In addition, one could also limit the number of states that we store by only storing one state per interval of ϕ. So for instance if $z_{\min} = 0.1$, $z_{\max} = 0.9$ and we only store a state with intervals of 0.005, at most 160 states would be stored.

Alternatively, one can implement a checkpointing strategy, where only a small subset of the states is stored, but all seeds of the random number generator, along with the corresponding values of $\phi(\mathbf{x})$ are stored. In that way, any state that is required for the branching process can be recomputed without actually storing it in memory. In practice it is common to store the state that is being branched from, meaning that only one state per trajectory has to be stored. Of course this comes at the cost of having to recompute $\mathbf{x}_{j_{\min}}$, starting from the checkpoint, in every iteration of TAMS. The storage in memory can be compressed even further by retaining only one checkpoint per group of trajectories that have a common ancestor, again at a slightly higher computational cost.

Optimizations can also be done in terms of parallelization. Since all trajectories in the first step are computed independently, parallelization of the first step is trivial. Later steps, however, are more difficult to parallelize due to the branching process. An example of a parallelized AMS method is given in [14].

2.3. Accuracy of the method

Generally, one realization of the algorithm is not enough to get an accurate answer. Also, from just one realization, we can not tell how accurate the answer actually is. Therefore, one computes M realizations of the algorithm, and uses those to compute both a more accurate answer, and an estimate of the error. From this we get a sample of M probability estimates $\hat{\alpha}_N$, which we can use to compute the sample mean

$$\mu_\alpha = \frac{1}{M}\sum_{m=1}^{M} \hat{\alpha}_{N,m} \tag{8}$$

and the sample variance

$$\sigma_\alpha^2 = \frac{1}{M-1}\sum_{m=1}^{M} \left(\hat{\alpha}_{N,m} - \mu_\alpha\right)^2. \tag{9}$$

An estimate of the relative error is given by

$$\epsilon_\alpha = \frac{\sigma_\alpha}{\mu_\alpha}. \tag{10}$$

We noted before that the relative error shows behavior between $\mathcal{O}(\sqrt{\log(1/\alpha)/N})$ (optimal) and $\mathcal{O}(\sqrt{1/(\alpha N)})$ (worst case).

It is known that when the committor is used as a reaction coordinate, the variance on the estimation of α tends to $\sqrt{|\log\alpha|}/\sqrt{N}$ for large N [13,14]. This means that we can compute a compensated variance

$$\sigma_0 = \sigma_\alpha \frac{\sqrt{N}}{\mu_\alpha\sqrt{|\log\mu_\alpha|}} \tag{11}$$

which can also be seen as the relative error divided by its expected optimal behavior. Since behavior is optimal when the committor is used, in that case σ_0 would be equal to 1. This quantity can be used as an indicator of the quality of the reaction coordinate.

2.4. Projected time-stepping in TAMS

In AMS and TAMS, but also in GPA, one is free to choose any stochastic time stepping method. One such method is the stochastic theta method with $\theta \neq 0$ [42], which, although convergence of the error of the rare event algorithm generally goes with $\sqrt{\Delta t}$, may be of benefit when large time steps are used [14]. In case $\theta \neq 0$, multiple linear systems have to be solved in every time step. Solving these linear systems is usually the computationally most expensive part of TAMS, even if the linear system is sparse, and an iterative solver with a good preconditioner is used.

At time step $j+1$ of the stochastic theta method performed on (1) with fixed $\mathbf{p}$, we take

$$\hat{F}(\mathbf{x}) = M\mathbf{x}_j - M\mathbf{x} + (1-\theta)\Delta t F(\mathbf{x}_j) + \theta\Delta t F(\mathbf{x}) + g(\mathbf{x}_j)\,\Delta\mathbf{W}_j = 0 \tag{12}$$

with Jacobian matrix

$$\hat{J}(\mathbf{x}) = \theta\Delta t J(\mathbf{x}) - M, \tag{13}$$

where J is the Jacobian matrix of F and M is the mass matrix. Newton's method for one implicit time step would normally look like this

1: $\mathbf{y}^{(0)} = \mathbf{x}_j$
2: **for** $k = 1, 2, \ldots$ until convergence **do**
3: **solve** $\hat{J}(\mathbf{y}^{(k-1)})\Delta\mathbf{y}^{(k)} = -\hat{F}(\mathbf{y}^{(k-1)})$
4: $\mathbf{y}^{(k)} = \mathbf{y}^{(k-1)} + \Delta\mathbf{y}^{(k)}$
5: $\mathbf{x}_{j+1} = \mathbf{y}^{(k)}$

A small potential optimization that can be made is computing $\hat{J}(\mathbf{x}_j)$ beforehand, and using this instead of $\hat{J}(\mathbf{y}^{(k-1)})$ (the chord method). This will make it such that Newton's method does not converge quadratically, but saves the time of computing the Jacobian matrix multiple times.

Instead of solving with this large sparse matrix $\hat{J}(\mathbf{y}^{(k-1)})$ in every Newton iteration, we instead propose to solve with a smaller matrix obtained from a Galerkin type projection $V^T\hat{J}(\mathbf{x}_j)V$, where V consists of an orthogonal basis of the most

important directions that are used by the stochastic perturbation. These vectors are given by the empirical orthogonal functions (EOFs), or more precisely by the eigenvectors belonging to the largest eigenvalues of the covariance matrix. After linearization around the steady state $\bar{\mathbf{x}}_A$, the covariance matrix C_A may be obtained by solving a generalized Lyapunov equation of the form

$$J(\bar{\mathbf{x}}_A)C_A M^T + MC_A J(\bar{\mathbf{x}}_A)^T + B_A B_A^T = 0, \tag{14}$$

where $B_A \in \mathbb{R}^{n\times n_w}$ is the matrix describing the noise [33]. Instead of computing the full covariance matrix, a low-rank approximation of the form $C_A \approx V_A Y_A V_A^T$, which only contains the eigenvectors belonging to the largest eigenvalues, may be computed instead [22]. Here $V_A \in \mathbb{R}^{n\times n_A}$ and $Y_A \in \mathbb{R}^{n_A\times n_A}$ with $n_A \ll n$. The cost of computing such a low-rank approximation of the covariance matrix is negligible compared to the cost of computing transition probabilities.

We can now compute an orthogonal basis V of the form

$$V = \text{orth}\left([\bar{\mathbf{x}}_A, V_A, B_A, \bar{\mathbf{x}}_B, V_B, B_B]\right), \tag{15}$$

where $\bar{\mathbf{x}}_A$ and $\bar{\mathbf{x}}_B$ are the two stable steady states, V_A is the basis of the low-rank Lyapunov solution at $\bar{\mathbf{x}}_A$, and V_B is the basis of the low-rank Lyapunov solution at $\bar{\mathbf{x}}_B$, and B_A and B_B are the noise matrices at $\bar{\mathbf{x}}_A$ and $\bar{\mathbf{x}}_B$. The basis may be expanded further with information from around the saddle-node bifurcation if available. Note that we can not compute V_C in the same way as V_A and V_B because the steady state $\bar{\mathbf{x}}_C$ is unstable.

We can now replace the Newton iteration with

1: $\mathbf{y}^{(0)} = V^T\mathbf{x}_j$
2: $A = V^T\hat{J}(\mathbf{x}_j)V$
3: **for** $k = 1, 2, \ldots$ until convergence **do**
4: **solve** $A\Delta\mathbf{y}^{(k)} = -V^T\hat{F}(V\mathbf{y}^{(k-1)})$
5: $\mathbf{y}^{(k)} = \mathbf{y}^{(k-1)} + \Delta\mathbf{y}^{(k)}$
6: $\mathbf{x}_{j+1} = V\mathbf{y}^{(k)}$

If we apply this in TAMS, we see that between two consecutive time steps, we always apply $V^TV\mathbf{y}^{(k)}$. Now since V is orthogonal, this means that we might as well apply TAMS itself to $\mathbf{y}^{(k)}$ directly. Doing this actually solves the most significant problem we observe when applying TAMS to high-dimensional systems, which is the storage of the trajectories, since we now only have to store the trajectories restricted to the space spanned by V. Say we have a system of size 10000 and V consisting of 100 vectors, which is a realistic scenario, then this reduces the required storage by a factor 100.

3. Application to AMOC model

In order to study the sensitivity of the AMOC to freshwater anomalies, we use the spatially quasi two-dimensional model as described in [43]. In the model, there are two active tracers: temperature T and salinity S, which are related to the density ρ by a linear equation of state

$$\rho = \rho_0\left(1 - \alpha_T\left(T - T_0\right) + \alpha_S\left(S - S_0\right)\right), \tag{16}$$

where α_T and α_S are the thermal expansion and haline contraction coefficients, respectively, and ρ_0, T_0, and S_0 are reference quantities.

3.1. Model formulation

In order to eliminate longitudinal dependence from the problem, we consider a purely buoyancy-driven flow on a non-rotating Earth [43]. We furthermore assume that inertia can be neglected in the meridional momentum equation. The mixing of momentum and tracers due to eddies is parameterized by simple anisotropic diffusion. In this case, the zonal velocity as well as the longitudinal derivatives are zero and the equations for the meridional velocity v, vertical velocity w, pressure p, and the tracers T and S are given by

$$-\frac{1}{\rho_0 r_0}\frac{\partial p}{\partial\theta} + A_V\frac{\partial^2 v}{\partial z^2} + \frac{A_H}{r_0^2}\left(\frac{1}{\cos\theta}\frac{\partial}{\partial\theta}\left(\cos\theta\frac{\partial v}{\partial\theta}\right) + \left(1 - \tan^2\theta\right)v\right) = 0, \tag{17a}$$

$$-\frac{1}{\rho_0}\frac{\partial p}{\partial z} + g\left(\alpha_T T - \alpha_S S\right) = 0, \tag{17b}$$

$$\frac{1}{r_0\cos\theta}\frac{\partial(v\cos\theta)}{\partial\theta} + \frac{\partial w}{\partial z} = 0, \tag{17c}$$

Table 1
Fixed model parameters of the two-dimensional ocean model, following [44].

D	=	$4.0 \cdot 10^3$	m	H_m	=	$2.5 \cdot 10^2$	m
r_0	=	$6.371 \cdot 10^6$	m	T_0	=	15.0	°C
g	=	9.8	m s^{-2}	S_0	=	35.0	psu
A_H	=	$2.2 \cdot 10^{12}$	m^2 s^{-1}	α_T	=	$1.0 \cdot 10^{-4}$	K^{-1}
A_V	=	$1.0 \cdot 10^{-3}$	m^2 s^{-1}	α_S	=	$7.6 \cdot 10^{-4}$	psu^{-1}
K_H	=	$1.0 \cdot 10^3$	m^2 s^{-1}	ρ_0	=	$1.0 \cdot 10^3$	kg m^{-3}
K_V	=	$1.0 \cdot 10^{-4}$	m^2 s^{-1}	τ	=	75.0	days

$$\frac{\partial T}{\partial t} + \frac{v}{r_0}\frac{\partial T}{\partial \theta} + w\frac{\partial T}{\partial z} = \frac{K_H}{r_0^2 \cos\theta}\frac{\partial}{\partial \theta}\left(\cos\theta \frac{\partial T}{\partial \theta}\right) + K_V \frac{\partial^2 T}{\partial z^2} + \mathrm{CA}(T), \tag{17d}$$

$$\frac{\partial S}{\partial t} + \frac{v}{r_0}\frac{\partial S}{\partial \theta} + w\frac{\partial S}{\partial z} = \frac{K_H}{r_0^2 \cos\theta}\frac{\partial}{\partial \theta}\left(\cos\theta \frac{\partial S}{\partial \theta}\right) + K_V \frac{\partial^2 S}{\partial z^2} + \mathrm{CA}(S). \tag{17e}$$

Here t is time, θ latitude, z the vertical coordinate, r_0 the radius of Earth, g the acceleration due to gravity, A_H (A_V) the horizontal (vertical) eddy viscosity, and K_H (K_V) the horizontal (vertical) eddy diffusivity. The terms with CA represent convective adjustment contributions.

The equations are solved on an equatorially symmetric, flat-bottomed domain. The basin is bounded by latitudes $\theta = -\theta_N$ and $\theta = \theta_N = 60°$ and has depth D. Free-slip conditions apply at the lateral walls and at the bottom. Rigid lid conditions are assumed at the surface and atmospheric pressure is neglected. The wind stress is zero everywhere, and "mixed" boundary conditions apply for temperature and salinity,

$$K_V \frac{\partial T}{\partial z} = \frac{H_m}{\tau}\left(\bar{T}(\theta) - T\right), \tag{18a}$$

$$K_V \frac{\partial S}{\partial z} = S_0 F_s(\theta). \tag{18b}$$

This formulation implies that temperatures in the upper model layer (of depth H_m) are relaxed to a prescribed temperature profile $\bar{T}$ at a rate τ^{-1}, while salinity is forced by a net freshwater flux F_s, which is converted to an equivalent virtual salinity flux by multiplication with S_0. The specification of the CA terms is given in [43]. The numerical values of the fixed model parameters are summarized in Table 1.

3.2. Bifurcation diagram

For the deterministic model, we take the surface forcing similar to what was used in [44] as

$$\bar{T}(\theta) = 10.0 \cos\left(\pi\theta/\theta_N\right), \tag{19a}$$

$$F_s(\theta) = \bar{F}_s(\theta) = \mu F_0 \frac{\cos(\pi\theta/\theta_N)}{\cos(\theta)} + \beta F_0 F_p(\theta), \tag{19b}$$

where $\mu = 0.3$ is the strength of the mean freshwater forcing, $F_0 = 0.1$ m yr^{-1} is a reference freshwater flux, β is the strength of an anomalous freshwater flux pattern F_p, which is -1 in the area $[45°N, 60°N]$, and 0 elsewhere.

The equations are a subset of those of a more general ocean model [45], where also the longitude dimension is taken into account. These more general equations are discretized on a longitude-latitude-depth equidistant $n_x \times n_y \times n_z$ grid using a second-order conservative central difference scheme. An integral condition expressing the overall conservation of salt is also imposed, as the salinity equation is only determined up to an additive constant. The total number of degrees of freedom is $n = 6n_x n_y n_z$, as there are six unknowns per point. The standard spatial resolution used is $n_x = 4, n_y = 32, n_z = 16$ and the solutions of the [45] model are uniform in the zonal direction, with the zonal velocity $u = 0$ and hence satisfy (17a).

The bifurcation diagram of the deterministic model for parameters as in Table 1 was already calculated in [44] and is shown in Fig. 2(a). On the y-axis, the sum of the maximum (Ψ^+) and minimum (Ψ^-) values of the meridional overturning streamfunction Ψ is plotted, where Ψ is defined through

$$\frac{\partial \Psi}{\partial z} = v\cos\theta, \quad -\frac{1}{r_0 \cos\theta}\frac{\partial \Psi}{\partial \theta} = w. \tag{20}$$

For the calculation of the volume transports (in Sv, where 1 Sv = 10^6 m^3 s^{-1}), the basin is assumed to have a zonal width of 64°. The value of $\Psi^+ + \Psi^-$ is zero when the AMOC is symmetric with respect to the equator.

We are interested in transitions from the top branch to the bottom branch of the bifurcation diagram. We will investigate transitions at parameter values between $\beta_a = 0.190$ and $\beta_b = 0.220$. The stochastic transitions will be from points between a and b to points between a' and b' with the same parameter values. The pattern of the asymmetric meridional overturning streamfunction at location a with sinking in the northern part of the basin, which represents the current state of the AMOC [2], is shown in Fig. 2(b), as well as the meridional overturning streamfunction at location a' with sinking in the southern part of the basin. The saddle-node bifurcations are located at $\beta = 0.2320$ and $\beta = -0.3371$.

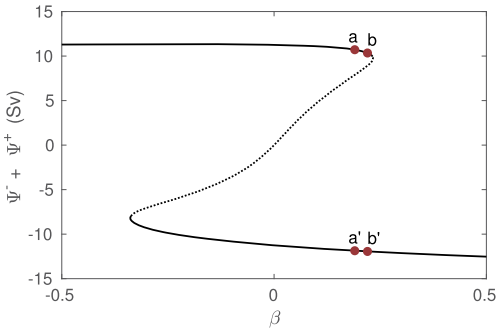

(a) Bifurcation diagram where a solid line is stable and dashed line is unstable.

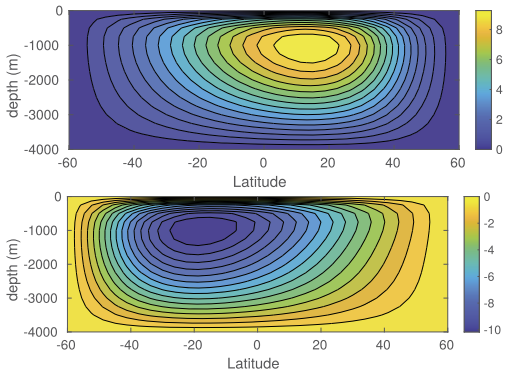

(b) The streamfunction at a and a' respectively.

Fig. 2. (a) Bifurcation diagram of the deterministic 2D AMOC model, with the forcing as in (19). (b) Meridional Overturning Streamfunction patterns at a, which represents the current state of the AMOC and a', in which the circulation is reversed, respectively, with contours in Sv.

3.3. Stochastic freshwater forcing

For the stochastic model, the freshwater forcing is chosen as

$$F_s(\theta,t) = (1+\sigma\zeta(\theta,t))\,\mu F_0 \frac{\cos(\pi\theta/\theta_n)}{\cos(\theta)} + \beta F_p(\theta), \tag{21}$$

where $\zeta(\theta,t)$ represents zero-mean white noise with a unit standard deviation, i.e., with $\mathbb{E}[\zeta(\theta,t)]=0$, $\mathbb{E}[\zeta(\theta,t)\zeta(\theta,s)] = \delta(\theta,t-s)$, $\mathbb{E}[\zeta(\theta,t)\zeta(\eta,t)] = \delta(\theta-\eta,t)$ and $\delta(\theta,t)$ being the Dirac delta function. The constant σ is the standard deviation of the noise which we set to $\sigma = 0.1$. The noise is additive and is only active in the freshwater component, only present at the surface, meridionally uncorrelated, and has magnitude σ of 10% of the background freshwater forcing amplitude at each latitude θ (see (21)).

3.4. Reaction coordinate

Since the pressure p is determined up to a constant, and therefore does not have a unique solution, computing the reaction coordinate by using the norm of the state vector as was done in (6) is not possible. We instead want to define the reaction coordinate in such a way that the distance between the meridional streamfunctions goes to zero, and therefore we decided to define a norm on the v-part of the state. This gives us a reaction coordinate of the form

$$\phi(\mathbf{x}) = \eta - \eta e^{-8d_A^2} + (1-\eta)e^{-8d_B^2}, \tag{22}$$

where, if $\|\cdot\|_v$ denotes the 2-norm on only the v-velocities, $d_A = \|\mathbf{x}-\bar{\mathbf{x}}_A\|_v / \|\bar{\mathbf{x}}_A - \bar{\mathbf{x}}_B\|_v$ is the normalized difference between $\mathbf{x}$ and $\bar{\mathbf{x}}_A$, and $d_B = \|\mathbf{x}-\bar{\mathbf{x}}_B\|_v / \|\bar{\mathbf{x}}_A - \bar{\mathbf{x}}_B\|_v$ is the normalized difference between $\mathbf{x}$ and $\bar{\mathbf{x}}_B$, and $\eta = \|\bar{\mathbf{x}}_C - \bar{\mathbf{x}}_A\|_v / \|\bar{\mathbf{x}}_B - \bar{\mathbf{x}}_A\|_v$ is the normalized difference between the unstable steady state $\bar{\mathbf{x}}_C$ and the stable steady state $\bar{\mathbf{x}}_A$. We assume that a trajectory has reached B if $\phi(\mathbf{x}) > z_{\max} = 0.75$, since at that point circulation has reversed and we are closer to $\bar{\mathbf{x}}_B$ as we are to the unstable steady state $\bar{\mathbf{x}}_C$, as can be seen in Fig. 2(a).

3.5. Results

In this section we report on the transition probabilities and compare the results from TAMS with the results from its projected variant. All optimizations as discussed in Section 2.2 were applied, with exception of the checkpointing strategy and the parallel computation of trajectories. The ocean model itself, however, does use a parallel domain decomposition. Computations are performed on Peregrine, the HPC cluster of the University of Groningen. Peregrine has nodes with 2 Intel Xeon E5 2680v3 CPUs (24 cores at 2.5 GHz) and each node has 128 GB of memory. For every instance of TAMS, four cores were used to allow for efficient parallelization. Due to the problem size and the way the convective adjustment is implemented, it did not make sense to use more than four cores per instance.

For our numerical experiments, we investigated the probability of a transition to a reversed circulation within 2000 years and ran TAMS 100 times with 300 trajectories. The stochastic theta method (12) was used with theta equals 0.5 and a time step of 2 years. Note that a (semi-)implicit method is required due to both the algebraic constraints and the large time step. The transition probability, relative error and compensated variance are shown in Fig. 3.

We first observe from Fig. 3(a) that the transition probability decreases rapidly when moving away from the bifurcation, and that the relative error increases as the transition probability gets smaller. We also see in Fig. 3(b), however, that the

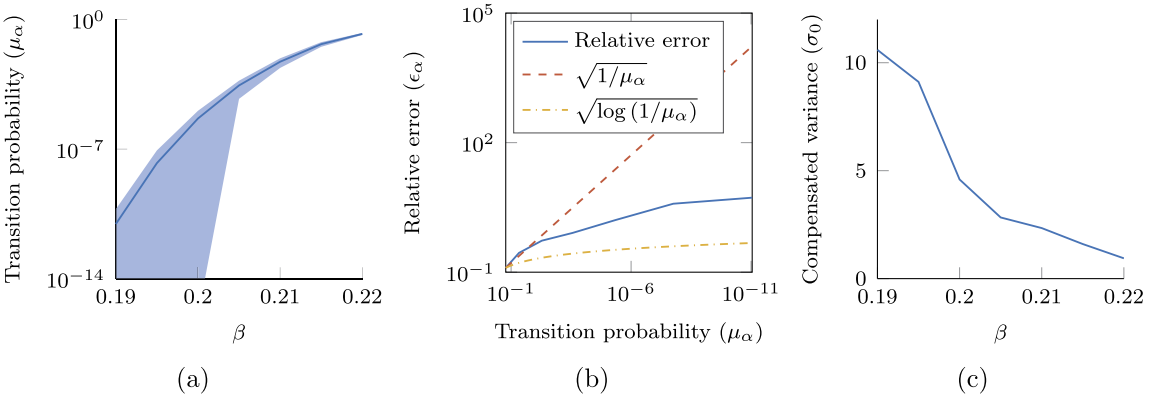

Fig. 3. Transition probability, relative error and compensated variance on the interval $[\beta_a, \beta_b]$ for $T_{\max} = 2000$ yr with the shaded area representing the standard deviation. To obtain the results we used 100 TAMS experiments with 300 trajectories each.

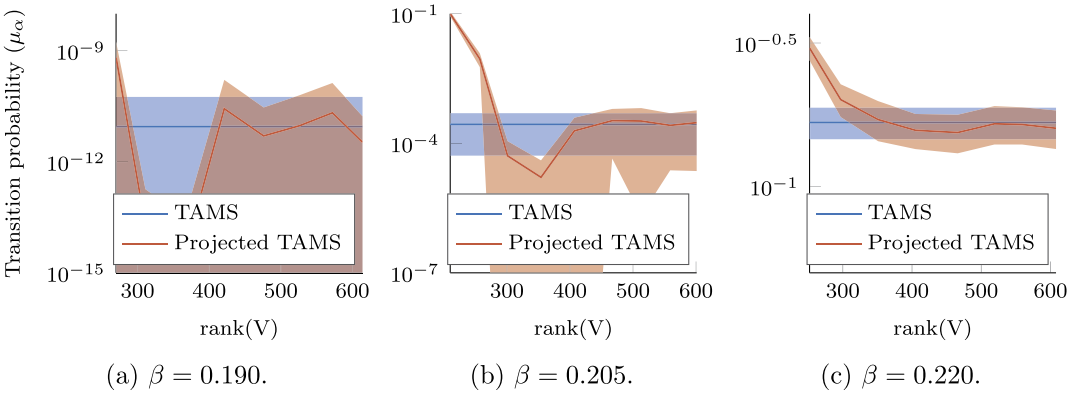

Fig. 4. Transition probability at different values of β and with increasing size of the basis V. The shaded area represents the standard deviation and $T_{\max} = 2000$ yr. To obtain the results we used 100 TAMS experiments with 300 trajectories each.

relative error shows much better behavior than the $\sqrt{1/\alpha}$ behavior we would get from a standard Monte Carlo method, as expected. From Fig. 3(c) we can conclude that the reaction coordinate can be further improved, since it gets further away from 1 as the transition probability decreases. The value of the compensated variance remains close to 1, however, in the region of interest where the probability of a transition within the next 2000 years is not negligible.

To determine how well the projected method performs, we look at accuracy at three different values of β: 0.190, 0.205 and 0.220, and with increasing sizes of the basis V. The basis is computed by the approximate solution of the generalized Lyapunov equation (14) where the relative residual $\|R\|_2/\|BB^T\|_2$ with $R = J(\bar{\mathbf{x}})VYV^TM^T + MVYV^TJ(\bar{\mathbf{x}})^T + BB^T$ is decreased to obtain a larger basis. The results are shown in Fig. 4. We see that for all values of β, the results are very close to the results obtained with the unprojected method when using a basis of rank 400 or larger, at the cost of a slightly larger error. This is remarkable, because it means that even for small probabilities (i.e. $\beta = 0.190$), we do not need a larger basis.

A common a priori method of determining how large the basis should be is by looking at the leading eigenvalues of the corresponding matrices. In our case these are the covariance matrices at two branches. As can be seen in Fig. 5, we never observe a sudden decrease of the eigenvalues, which would indicate that a sufficient amount has been determined. Indeed, we found that we are able to use the same basis size for all probabilities that we compute. This motivates to investigate whether, for a larger application, we would be able to determine the proper basis size once by looking at a figure such as Fig. 4(c), and then use that basis size for all further experiments. However, this is outside the scope of this paper. Note that computing Fig. 4(c) is still relatively cheap when the probability is large.

In Table 2 we show the computational cost of the unprojected method compared to the projected method for different ranks of the basis V. The values we show were determined at $\beta = 0.220$, but these are actually independent of the value of β. Here we observe that for a basis of rank 400, which we deemed sufficient, the projected method is twice as fast and uses 97% less memory.

4. Summary and discussion

The Adaptive Multilevel Splitting (AMS) method and its Trajectory variant (TAMS) are very promising numerical techniques to study rare events and in particular transitions between different equilibria of stochastic dynamical systems.

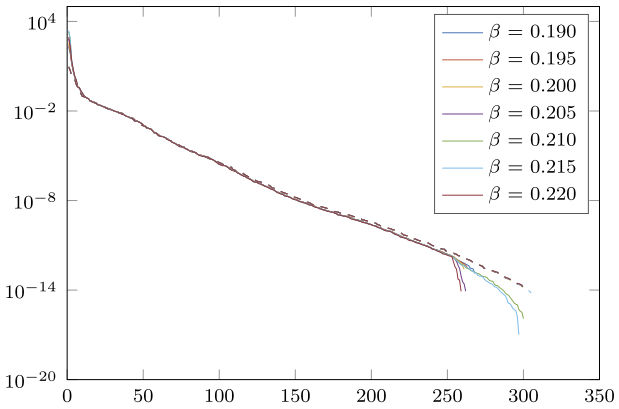

Fig. 5. Leading eigenvalues of the covariance matrix at the top (solid) and bottom (dashed) branch as determined by the Lyapunov solver RAILS [22] with a relative residual tolerance of 10^{-10}.

Table 2
Computational time per time step and memory cost per state vector of the unprojected method (first row) and the projected method at increasing sizes of V.

rank(V)	$t(s)$	Mem (kB)
–	0.32	96.00 kB
252	0.12	1.97 kB
297	0.13	2.32 kB
351	0.14	2.74 kB
405	0.17	3.16 kB
466	0.19	3.64 kB
519	0.22	4.05 kB
559	0.24	4.37 kB
608	0.26	4.75 kB

However, in their original form they cannot be applied to high-dimensional dynamical systems, i.e. as derived from three-dimensional ocean models. Therefore, we implemented a projected time-stepping method, which can be applied to any of the methods in the Generalized Adaptive Multilevel Splitting framework [46]. This approach solves the problem in a reduced space, which is obtained from a Galerkin-like projection with the low-rank solution of a generalized Lyapunov equation. The results obtained with the projected time-stepping approach are very similar to the ones obtained with standard TAMS. The main benefit of the projected time-stepping method is the reduced computational cost, especially in terms of memory usage, where only 3% of the amount of memory that is required for standard TAMS is used. Since memory is the limiting factor in many applications, this reduction, along with all other optimizations such as the checkpointing strategy we mentioned, will allow a larger application potential, without necessarily negatively impacting the computational time.

The benefit in terms of memory usage is present even if an explicit time stepping method is used. Downsides of using the projection in combination with an explicit method, however, are a higher computational cost due to the projection that has to be performed in every time step and the requirement of having access to the Jacobian matrix to be able to compute the low-rank approximation of the covariance matrix. In case of implicit time stepping, however, the Jacobian matrix is usually already available, and solving the projected system is often cheaper than solving with the full Jacobian matrix.

As an application, we computed transition probabilities for a quasi two-dimensional model of the Atlantic Meridional Overturning Circulation (AMOC). The model consists of a system of stochastic partial differential equations, representing the equations for mass, momentum, temperature and salinity in the oceanic basin, resulting (when discretized) in a dynamical system with a total of 12288 degrees of freedom. While this is a highly idealized model, it captures the essential feedback (the salt-advection feedback) which is responsible for the sensitivity of the AMOC to freshwater perturbations. Although we investigated only a few cases, the transition probability of the AMOC in the two-dimensional model decreases drastically when moving away from the saddle-node bifurcation. This indicates that, in this model, a stochastic transition is not very likely to occur unless the AMOC is close to critical conditions.

It would be interesting to see if the methodology can be extended to three-dimensional ocean models for which one can calculate bifurcation diagrams. In [22], we found for a primitive three-dimensional example where the noise is fully uncorrelated (worst case), that the size of the space that comes from the generalized Lyapunov equation increases with $\mathcal{O}(n)$. Therefore it is possible that for three-dimensional ocean models the cost of solving with the projected Jacobian matrix becomes dominant, and eventually more expensive than computations with the original Jacobian matrix, but the memory cost will still be at most 3% of the original TAMS. Whether we can suffice with predicting the required size of the basis from simulations with a high transition probability is still an open question.

Additionally, more sophisticated noise models may be used, such as correlated additive and multiplicative (CAM) noise, which is commonly used in climate science. The methodology from this paper can also be used for multiplicative types of noise, but in that case the generalized Lyapunov equations have to be extended with a bilinear term. We refer to the resulting equations as extended generalized Lyapunov equations. These may be solved by combining the methodology from [47] and [22].

Something that still needs further investigation is the reaction coordinate. In this paper, we just chose a reaction coordinate that works reasonably well for the AMOC application, but the fact that the compensated variance moves away from 1 as the transition probability becomes smaller suggests that it is possible to find more efficient reaction coordinates. It has also been shown theoretically for systems with a few degrees of freedom that paths that go through $\bar{\mathbf{x}}_C$ should be favored as they more closely represent the committor [40]. Further incorporating terms of the form $\|x - x_C\|$ may help with this [14,15]. It would be favorable if it were possible to find a reaction coordinate that works well for a wide variety of high-dimensional dynamical systems, and work on this is in progress.

CRediT authorship contribution statement

S. Baars: Conceptualization, Investigation, Methodology, Software, Writing - original draft, Writing - review & editing. **D. Castellana:** Conceptualization, Methodology, Writing - original draft, Writing - review & editing. **F.W. Wubs:** Conceptualization, Funding acquisition, Methodology, Supervision, Writing - original draft, Writing - review & editing. **H.A. Dijkstra:** Conceptualization, Funding acquisition, Methodology, Supervision, Writing - original draft, Writing - review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to thank Jeroen Wouters and Daan Crommelin for the useful discussions. We would also like to thank the reviewers for their very helpful and constructive comments, which greatly improved the paper. This work is part of the Mathematics of Planet Earth research program with project number 657.014.007, which is financed by the Netherlands Organisation for Scientific Research (NWO) (SB and FW), the SMCM project of the Netherlands eScience Center (NLeSC) with project number 027.017.G02 (SB, FW and HD) and the European Union's Horizon 2020 research and innovation program for the ITN CRITICS under Grant Agreement Number 643073 (DC and HD).

References

[1] H. Stommel, Thermohaline convection with two stable regimes of flow, Tellus B 13 (2) (1961), https://doi.org/10.3402/tellusb.v13i2.12985.
[2] H.A. Dijkstra, Nonlinear Physical Oceanography, Springer, Netherlands, 2005.
[3] D. Castellana, S. Baars, F.W. Wubs, H.A. Dijkstra, Transition probabilities of noise-induced transitions of the Atlantic Ocean circulation, Sci. Rep. 9 (1) (2019), https://doi.org/10.1038/s41598-019-56435-6.
[4] M.I. Freidlin, A.D. Wentzell, Random Perturbations, Springer US, 1984.
[5] H. Eyring, The activated complex in chemical reactions, J. Chem. Phys. 3 (2) (1935) 107–115, https://doi.org/10.1063/1.1749604.
[6] H.A. Kramers, Brownian motion in a field of force and the diffusion model of chemical reactions, Physica 7 (4) (1940) 284–304, https://doi.org/10.1016/s0031-8914(40)90098-2.
[7] F. Bouchet, J. Reygner, Generalisation of the Eyring–Kramers transition rate formula to irreversible diffusion processes, Ann. Henri Poincaré 17 (12) (2016) 3499–3532, https://doi.org/10.1007/s00023-016-0507-4.
[8] F. Bouchet, J. Rolland, E. Simonnet, Rare event algorithm links transitions in turbulent flows with activated nucleations, Phys. Rev. Lett. 122 (7) (2019), https://doi.org/10.1103/physrevlett.122.074502.
[9] T. Lestang, F. Ragone, C.-E. Bréhier, C. Herbert, F. Bouchet, Computing return times or return periods with rare event algorithms, J. Stat. Mech. Theory Exp. 2018 (4) (2018) 043213, https://doi.org/10.1088/1742-5468/aab856.
[10] P.D. Moral, J. Garnier, Genealogical particle analysis of rare events, Ann. Appl. Probab. 15 (4) (2005) 2496–2534, https://doi.org/10.1214/105051605000000566.
[11] P.D. Moral, Mean Field Simulation for Monte Carlo Integration, Chapman and Hall/CRC, 2013.
[12] J. Wouters, F. Bouchet, Rare event computation in deterministic chaotic systems using genealogical particle analysis, J. Phys. A, Math. Theor. 49 (37) (2016) 374002, https://doi.org/10.1088/1751-8113/49/37/374002.
[13] F. Cérou, A. Guyader, Adaptive multilevel splitting for rare event analysis, Stoch. Anal. Appl. 25 (2) (2007) 417–443, https://doi.org/10.1080/07362990601139628.
[14] J. Rolland, E. Simonnet, Statistical behaviour of adaptive multilevel splitting algorithms in simple models, J. Comput. Phys. 283 (2015) 541–558, https://doi.org/10.1016/j.jcp.2014.12.009.
[15] J. Rolland, F. Bouchet, E. Simonnet, Computing transition rates for the 1-D stochastic Ginzburg–Landau–Allen–Cahn equation for finite-amplitude noise with a rare event algorithm, J. Stat. Phys. 162 (2) (2015) 277–311, https://doi.org/10.1007/s10955-015-1417-4.
[16] F. Bouchet, J. Laurie, O. Zaboronski, Langevin dynamics, large deviations and instantons for the quasi-geostrophic model and two-dimensional Euler equations, J. Stat. Phys. 156 (6) (2014) 1066–1092, https://doi.org/10.1007/s10955-014-1052-5.
[17] J. Laurie, F. Bouchet, Computation of rare transitions in the barotropic quasi-geostrophic equations, New J. Phys. 17 (1) (2015) 015009, https://doi.org/10.1088/1367-2630/17/1/015009.
[18] W. E, W. Ren, E. Vanden-Eijnden, Minimum action method for the study of rare events, Commun. Pure Appl. Math. 57 (5) (2004) 637–656, https://doi.org/10.1002/cpa.20005.

[19] E. Vanden-Eijnden, M. Heymann, The geometric minimum action method for computing minimum energy paths, J. Chem. Phys. 128 (6) (2008) 061103, https://doi.org/10.1063/1.2833040.
[20] X. Zhou, W. Ren, W. E, Adaptive minimum action method for the study of rare events, J. Chem. Phys. 128 (10) (2008) 104111, https://doi.org/10.1063/1.2830717.
[21] T. Grafke, T. Schäfer, E. Vanden-Eijnden, Long term effects of small random perturbations on dynamical systems: theoretical and computational tools, in: Recent Progress and Modern Challenges in Applied Mathematics, Modeling and Computational Science, 2017, pp. 17–55.
[22] S. Baars, J.P. Viebahn, T.E. Mulder, C. Kuehn, F.W. Wubs, H.A. Dijkstra, Continuation of probability density functions using a generalized Lyapunov approach, J. Comput. Phys. 336 (2017) 627–643, https://doi.org/10.1016/j.jcp.2017.02.021.
[23] C. Hartmann, Balanced model reduction of partially observed Langevin equations: an averaging principle, Math. Comput. Model. Dyn. Syst. 17 (5) (2011) 463–490, https://doi.org/10.1080/13873954.2011.576517.
[24] M. Loève, Probability Theory: Foundations, Random Sequences, Van Nostrand, 1955.
[25] J.L. Lumley, The structure of inhomogeneous turbulent flows, in: Atmospheric Turbulence and Wave Propagation, 1967, pp. 166–178.
[26] W. Cazemier, R.W.C.P. Verstappen, A.E.P. Veldman, Proper orthogonal decomposition and low-dimensional models for driven cavity flows, Phys. Fluids 10 (7) (1998) 1685–1699, https://doi.org/10.1063/1.869686.
[27] C. Hartmann, C. Schütte, A constrained hybrid Monte-Carlo algorithm and the problem of calculating the free energy in several variables, Z. Angew. Math. Mech. 85 (10) (2005) 700–710, https://doi.org/10.1002/zamm.200410218.
[28] R. Yvinec, M.R. D'Orsogna, T. Chou, First passage times in homogeneous nucleation and self-assembly, J. Chem. Phys. 137 (24) (2012) 244107, https://doi.org/10.1063/1.4772598.
[29] D. Mukhin, D. Kondrashov, E. Loskutov, A. Gavrilov, A. Feigin, M. Ghil, Predicting critical transitions in ENSO models. Part II: spatially dependent models, J. Climate 28 (5) (2015) 1962–1976, https://doi.org/10.1175/jcli-d-14-00240.1.
[30] C. Hartmann, C. Schütte, W. Zhang, Model reduction algorithms for optimal control and importance sampling of diffusions, Nonlinearity 29 (8) (2016) 2298–2326, https://doi.org/10.1088/0951-7715/29/8/2298.
[31] M.A. Mohamad, W. Cousins, T.P. Sapsis, A probabilistic decomposition-synthesis method for the quantification of rare events due to internal instabilities, J. Comput. Phys. 322 (2016) 288–308, https://doi.org/10.1016/j.jcp.2016.06.047.
[32] C.X. Hernández, H.K. Wayment-Steele, M.M. Sultan, B.E. Husic, V.S. Pande, Variational encoding of complex dynamics, Phys. Rev. E 97 (6) (2018), https://doi.org/10.1103/physreve.97.062412.
[33] C.W. Gardiner, Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences, Springer, Berlin Heidelberg, 1985.
[34] H.B. Keller, Numerical solution of bifurcation and nonlinear eigenvalue problems, in: P.H. Rabinowitz (Ed.), Applications of Bifurcation Theory, Academic Press, New York, U.S.A., 1977, pp. 359–384.
[35] H. Kahn, T.E. Harris, Estimation of particle transmission by random sampling, Natl. Bur. Stand., Appl. Math. Ser. 12 (1951) 27–30.
[36] M.N. Rosenbluth, A.W. Rosenbluth, Monte Carlo calculation of the average extension of molecular chains, J. Chem. Phys. 23 (2) (1955) 356–359, https://doi.org/10.1063/1.1741967.
[37] F. Cérou, A. Guyader, T. Lelièvre, D. Pommier, A multiple replica approach to simulate reactive trajectories, J. Chem. Phys. 134 (5) (2011) 054108, https://doi.org/10.1063/1.3518708.
[38] G. Rubino, B. Tuffin, Rare Event Simulation Using Monte Carlo Methods, John Wiley & Sons, Ltd., 2009.
[39] E. Simonnet, Combinatorial analysis of the adaptive last particle method, Stat. Comput. 26 (1–2) (2014) 211–230, https://doi.org/10.1007/s11222-014-9489-6.
[40] P. Metzner, C. Schütte, E. Vanden-Eijnden, Illustration of transition path theory on a collection of simple examples, J. Chem. Phys. 125 (8) (2006) 084110, https://doi.org/10.1063/1.2335447.
[41] S. Baars, Numerical methods for studying transition probabilities in stochastic ocean-climate models, Ph.D. thesis, University of Groningen, Jun 2019.
[42] P.E. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Springer, Berlin Heidelberg, 1992.
[43] M. den Toom, H.A. Dijkstra, F.W. Wubs, Spurious multiple equilibria introduced by convective adjustment, Ocean Model. 38 (1–2) (2011) 126–137, https://doi.org/10.1016/j.ocemod.2011.02.009.
[44] M. van der Mheen, H.A. Dijkstra, A. Gozolchiani, M. den Toom, Q. Feng, J. Kurths, E. Hernandez-Garcia, Interaction network based early warning indicators for the Atlantic MOC collapse, Geophys. Res. Lett. 40 (11) (2013) 2714–2719, https://doi.org/10.1002/grl.50515.
[45] A. de Niet, F. Wubs, A.T. van Scheltinga, H.A. Dijkstra, A tailored solver for bifurcation analysis of ocean-climate models, J. Comput. Phys. 227 (1) (2007) 654–679, https://doi.org/10.1016/j.jcp.2007.08.006.
[46] C.-E. Bréhier, M. Gazeau, L. Goudenège, T. Lelièvre, M. Rousset, Unbiasedness of some generalized adaptive multilevel splitting algorithms, Ann. Appl. Probab. 26 (6) (2016) 3559–3601, https://doi.org/10.1214/16-aap1185.
[47] S.D. Shank, V. Simoncini, D.B. Szyld, Efficient low-rank solution of generalized Lyapunov equations, Numer. Math. 134 (2) (2015) 327–342, https://doi.org/10.1007/s00211-015-0777-7.