UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

EFFICIENT COMPUTATIONAL APPROACHES FOR TREATMENT OF TRANSFORMED PATH INTEGRALS

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

MATTHEW VON GONTEN
Norman, Oklahoma
2018

EFFICIENT COMPUTATIONAL APPROACHES FOR TREATMEMT
OF TRANSFORMED PATH INTEGRALS

A THESIS APPROVED FOR THE
SCHOOL OF AEROSPACE AND MECHANICAL ENGINEERING

BY

Dr. Prakash Vedula, Chair

Dr. Jivtesh Garg

Dr. Dimitrios Papavassiliou

Acknowledgments

I would like to begin by thanking Dr. Prakash Vedula for making this thesis possible. From the outset, he encouraged me to set my goals high, with the declaration that I would write a thesis so good I could pursue a PhD at any university I wanted. He turned my attention to a line of research that was new and innovative and gave me an opportunity to be a part of it, and throughout my journey he has offered guidance and support when I have most needed it, while encouraging me to discover and internalize things for myself whenever possible. I am proud to have worked with Dr. Vedula, and I believe I have become a more thoughtful and ambitious engineer because of him.

I would also like to thank Gnana Subramaniam, my mentor and predecessor who has helped me to build off of his own work in the Transformed Path Integral. Whenever I struggled to understand a concept from his research, he was more than happy to explain, and when I needed a fresh pair of eyes to help solve a problem or fix a bug, he was always willing to help. My respect for Gnana has only grown over the past two years as I have come to better understand and appreciate the brilliance of of his work and, more importantly, the strength of his character. I depart from this experience having gained a valued friend.

I extend my thanks to Dr. Jivtesh Garg and Dr. Dimitrios Papavassiliou for their interest in my work that led them to serve on my thesis topic committee. I am grateful to Dr. Farrokh Mistree for rekindling my desire to pursue an advanced degree in my Senior year, and for continuing to encourage me in the two years that followed. To the Graduate Student Community, I give thanks for making my graduate student experience beyond academics a fulfilling one. To the administration of the college of Aerospace and Mechanical Engineering, I am grateful for helping me to stay on track and better navigate the red tape. And to all the friends living tens, hundreds, or thousands of miles away who

never stopped asking how I was doing and offering their support, thank you as well.

Finally, I would like to thank my family for everything they have done to help me get this far. For all the care packages that gave me something good to eat when I was living off of scrambled eggs. For all the phone calls making sure I was getting enough sleep. For all the times they listened to me go on about my research even though they couldn't understand it. For respecting my space when I needed to focus and when I desperately needed to relax. For always giving me a home to return to. I will never be able to thank you enough for all you've done for me, but I will do my utmost to live up to the faith you've always had in me.

Table of Contents

List of Tables

List of Figures

Abstract

In this thesis, efficient computational approaches for treatment of transformed path integrals (TPIs) are proposed. The TPI-based method allows us to calculate the time evolution probability density functions (PDFs) using a short time propagator matrix that accounts for the transition probability in a transformed domain. A grid-based implementation of the TPI, in contrast to the conventional fixed-grid implementation of a path integral (PI), allows the propagation of the PDF to be performed on a dynamically adaptive grid parametrized by the mean and covariance of the PDF. TPI-based methods generate PDFs from all possible paths within the transformed space, and while these methods are found to be highly effective at capturing tail information in systems with large drifts, diffusions, and concentrations, they can become somewhat computationally expensive when applied to systems that must be represented by large numbers of data points. The purpose of this thesis is to develop computationally efficient TPI-based methods that largely preserve the accuracy and other desirable features of the original TPI method.

The first proposed method, referred to as the bandlimited TPI (BL-TPI) method, takes advantage of the fact that the transition probability is often concentrated around a set of peaks, with one natural peak occurring for each source state. This allows us to consider sparse matrix representations of the transition probability matrix operator and consider a region of importance about the peak transition probability curve for consideration in PDF propagation while neglecting all values outside of this region. With the use of sparse matrix tools, the BL-TPI enables us to perform PDF propagation using far fewer operations than the standard implementation. In the second proposed method, a TPI implementation based on the Symmetric Fast Gauss Transform (SFGT) is proposed. This method utilizes a Taylor series expansion of the Gaussian kernel in the propagator matrix to reduce the convolution operation for the PDF to an infinite sum of moments. This allows us to

perform calculations involving source and target terms separately, eliminating their convolution and in the process potentially reducing the associated computational complexity.

In order to demonstrate the effectiveness of the proposed approaches, comparisons with the standard TPI implementation are performed for canonical problems in one-dimensional and multi-dimensional state spaces. The results from the BL-TPI method appear promising and indicate that the method is applicable to a wide range of cases. In contrast, the effectiveness of the SFGT approach is found to be inherently conditional, and the computational cost of this method can exceed that of the standard TPI method in many cases.

1 Introduction

1.1 Stochastic Differential Equations and the Fokker-Planck Equation

A stochastic process refers to any process whose outcome is dependent on a series of variables that can be considered random due to internal and external factors that are uncontrollable or unpredictable. The study of these processes and the stochastic differential equations (SDEs) used to solve them are paramount to many fields of science, engineering, and mathematics, with applications ranging from particle physics, to thermodynamics, to structural and fluid mechanics, to vehicles navigating complex and turbulent environments. This wide range of uses for SDEs gives tremendous value to the development of efficient and accurate methods for solving them.

The state of a stochastic system at a given time can be modeled using a probability density function (PDF), and by assuming that the random excitation processes of the system occur in Wiener increments that follow the Markov property, the time evolution of the PDF that characterizes the system can in many cases be accurately predicted. In such cases, the time evolution of a PDF is governed by the Fokker-Planck equation (FPE), a second order differential equation which accounts for both the system dynamics and the random excitations through a drift vector function and a diffusion coefficient matrix, respectively. No general analytical solution to the FPE currently exists, so outside of a few specific circumstances the time evolution of a PDF must be either analytically approximated or numerically estimated.

Among the most commonly used numerical solutions are Monte Carlo (MC) methods, which are a class of numerical methods that construct an approximate PDF from a series of random samples. MC methods are particularly useful for representing infor-

mation about statistical properties of the PDF as a whole, such as mean and covariance. However, they often suffer from sampling errors that tend to result in loss of accuracy in representations of tail regions of the PDF, such regions representing extreme but not impossible events. Grid-based Finite Difference (FD) and Finite Element (FE) methods can more precisely capture tail information of the PDF over small time intervals and few iterations, but tend to be computationally expensive. Furthermore, such methods are traditionally implemented on a fixed grid which may not be able to accurately capture the transient behavior of the PDF for systems that are expanding, contracting, or translating.

1.2 Introduction to the Transformed Path Integral

The Path Integral (PI) solution to the Fokker-Planck equation utilizes a short time propagator to yield a solution that is exact in the short time limit. In conventional fixed grid implementations of the PI, the short time propagator is realized as a propagator matrix comprised of the transition probabilities between all source and target states. This propagator matrix is then used to find the new state PDF at the current time using a weighted sum of these "paths". Early formulations of Path Integrals were developed by Norbert Wiener as a theoretical tool to address Brownian motion, and have been found to have many appealing features, such as preserving the non-negativity of PDFs. Furthermore, in the short time limit as the time increment approaches zero, the PDF propagation is an exact solution of the FPE; however, similar to the aforementioned finite difference and finite element methods, traditional grid-based PI implementations struggle with the limitations of fixed grid formulations, and consequently have not found widespread practical usage.

To address these limitations, the Transformed Path Integral (TPI) method was developed as a robust alternative to the fixed grid PI (e.g., Subramaniam et al., 2017). This formulation of the propagator matrix utilizes a dynamic state space transformation that

continuously shifts and distorts the computational domain according to the mean and covariance of the probability state. The TPI inherits several of the desirable features of the PI, while through the aforementioned transformation maintaining an efficient resolution of a relevant state space. The latter feature allows the TPI to avert many common problems associated with conventional fixed grid approaches caused by large drift dynamical systems and large diffusion white noise excitations, as well as to ensure sufficient resolution to accurately capture PDF peak information in contracting systems that necessitate finer grid spacing over time.

1.3 Objective and Structure

The goal of this study is to develop, test, and compare new formulations based on the TPI method; to determine their respective computational efficiencies compared to that of the standard TPI method; to develop guidelines for their application; and to identify their limitations and areas for further research. Though it has many appealing features, the standard TPI implementation can be made more computationally efficient. In the standard implementation, the the convolution operation in the PDF propagation is reduced to a matrix-vector multiplication. This involves $O(N^2)$ operations in a computational grid with N grid points. In this study, we propose two classes of methods for improving the efficiency of the TPI formulation. The first of these methods, bandlimiting, takes advantage of the concentration of transition probability values in a relatively small region about a set of peaks, while the second method, based on the symmetric fast Gauss transform (SFGT), manipulates the Gaussian nature of the kernel in the propagator matrix to split the source and target terms in the convolution operation. For either method the number of operations reduces to $O(kN)$ for some $k << N$. In the BL-TPI method, k corresponds to bandwidth, while in the SFGT-TPI method, it corresponds to the size of the Taylor series

summation; in both cases it can be adjusted based on the desired accuracy.

In Chapter 2, we examine the mathematical formulation of Transformed Path Integrals and develop new TPI-based formulations that incorporate principles of bandlimiting and the Symmetric Fast Gauss Transform. In Chapter 3, we perform an analysis of the BL-TPI by comparing its performance for different bandwidths when applied to canonical SDOF systems. We then extend this analysis to MDOF systems to arrive at recommendations for an optimal bandwidth and conditions for optimal performance of the BL-TPI method. In Chapter 4, we analyze the SFGT TPI formulation and develop a framework for estimating the number of moments needed to produce accurate results within specific parameters and, conversely, to establish approximate boundaries within which the solution from the SFGT TPI method may be considered accurate. We then estimate the number of moments needed to accurately solve canonical stochastic systems, and measure the accuracy of the PDF propagation over a single time step for these systems within different boundary precisions. In Chapter 5 we study and compare the relative performances of the BL-TPI and the SFGT TPI to develop recommendations based on the results of this study. Chapter 6 presents our conclusions and Chapter 7 outlines the areas for future research to build upon these findings.

2 Mathematical Formulation

In this chapter, the definition of the Transformed Path Integral is outlined and its structure and properties are examined as a basis for bandlimiting and Symmetric Fast Gauss Transform formulations. Expressions for the peak states in the transition probability matrix and for an optimal bandwidth about this set of states are obtained for single degree of freedom (SDOF) systems, allowing us to develop bandlimited implementations of the TPI method. Different approaches for extending the bandlimited TPI (BL-TPI) formulations into multiple degree of freedom (MDOF) dynamical systems are also presented. Additionally, the formulation for the computational procedure symmetric fast Gauss transform (SFGT) is examined. New SFGT-based TPI implementations for both SDOF and MDOF systems are developed, while use of the Gauss-Hermite quadrature to improve efficiency is also illustrated.

2.1 Transformed Path Integral

A dynamical system subjected to random excitations may be modeled by an Itô stochastic differential equation, as shown below

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t), t)dt + \mathbf{A}(\mathbf{x}(t), t)d\mathbf{w} \tag{1}$$

in which $\mathbf{x} \in \mathbb{R}^{N_s \times 1}$, $\mathbf{f} \in \mathbb{R}^{N_s \times 1}$, and $\mathbf{A} \in \mathbb{R}^{N_s \times N_w}$ represent the state vector, the nonrandom drift, and the diffusion coefficient matrix, respectively, while $d\mathbf{w}(t) \in \mathbb{R}^{N_w \times 1}$ represents the Wiener increment at time t. These increments form a set of vectors with the distribution $\mathbf{w}(t) - \mathbf{w}(s) \sim N(\mu, \sigma^2)$ where $\mu = 0$ and $\sigma^2 = t - s$. When applied to

SDOF systems, Eq. (1) reduces to

$$dx(t) = f(x(t), t)dt + A(x(t), t)dw \tag{2}$$

The evolution of the state probability for the process in Eq. (1) may be expressed in terms of a transition probability density function which represents the conditional probability of systems reaching a target state given a source state. The transition probability for a system with N discrete source and target states can be expressed as a $N \times N$ matrix. In the short-time limit, i.e. for a small time increment dt, the matrix represents the short-time propagator for the system. For the dynamical system described by Eq. (1), each element of the short-time propagator matrix is given by

$$\begin{aligned} p(\mathbf{x}, t|\mathbf{x}', t') = {} & (4\pi dt)^{-N_s/2} \left\{ \det \left[\mathbf{G}(\mathbf{x}', t') \right] \right\}^{-1/2} \\ & \times \exp \left\{ -\frac{1}{4dt} \left[\mathbf{x} - \mathbf{x}' - \mathbf{f}(\mathbf{x}', t')dt \right]^T \left[\mathbf{G}(\mathbf{x}', t') \right]^{-1} \left[\mathbf{x} - \mathbf{x}' - \mathbf{f}(\mathbf{x}', t')dt \right] \right\} \end{aligned} \tag{3}$$

where $dt = t - t'$ and $\mathbf{G} = \frac{1}{2}(\mathbf{A}\mathbf{A}^T)$. In the case of SDOF systems, Eq. (3) reduces to

$$p(x, t|x', t') = \frac{1}{2\pi A^2(x, t)dt} \times \exp \left\{ -\frac{\left[x - x' - f(x', t')dt \right]^2}{2A^2(x', t')dt} \right\} \tag{4}$$

With the above expressions for the short-time propagator matrix, the PDF propagation is performed via the Chapman-Kolmogorov equation.

$$p(\mathbf{x}, t) = \int p(\mathbf{x}, t|\mathbf{x}', t') \times p(\mathbf{x}', t')d\mathbf{x}' \tag{5}$$

$$p(x, t) = \int p(x, t|x', t') \times p(x', t')dx' \tag{6}$$

Grid-based numerical implementations of Eqs. (5) and (6) reduce the convolution

operation to a matrix-vector multiplication. While fixed grid implementations of the path integral method are accurate, numerically stable, and maintain non-negativity of probability, such implementations face challenges in accurately describing the transient behavior of PDFs in systems with large drift, diffusion, or concentration of probability.

The transformed path integral approach offers a robust alternative to the conventional fixed grid implementation of path integrals. The TPI method is able to better address the challenges due to drift, diffusion, or concentration of the PDF through a dynamic transformation of the state space based on the mean and covariance of the state variables. This dynamic transformation is given by

$$z_t = (x_t - \mu_t)/\sigma_t \tag{7}$$

for SDOF systems, where x_t and z_t are the absolute coordinates and transformed coordinates at time t, respectively, and μ_t and σ_t^2 are the mean and covariance of x_t. The equivalent transformation for MDOF systems is given by

$$\mathbf{z}_t = \mathbf{R}_t^{-1}(\mathbf{x}_t - \boldsymbol{\mu}_t) \tag{8}$$

where $\boldsymbol{\Sigma} \equiv \mathbf{R}_t\mathbf{R}_t^T$ denotes the covariance matrix of $\mathbf{x}_t$ at time t.

Based on the transformation in Eq. (7), we obtain the new formulation of the short-time propagator in the transformed space for SDOF systems:

$$p(z,t|z',t') = \frac{\sigma}{\sqrt{2\pi\tilde{A}^2(z',t')dt}} \times \exp\left\{-\frac{\left[\sigma z - \sigma' z' - \delta\tilde{f}(z',t')dt\right]^2}{2\tilde{A}^2(z',t')dt}\right\} \tag{9}$$

where $\tilde{f}(z_t,t) = f(\sigma_t z_t + \mu_t, t)$, $\tilde{A}(z_t,t) = A(\sigma_t z_t + \mu_t, t)$, and $\delta f(x_t,t) = f(x_t,t) -$

$\langle f(x_t,t)\rangle_{x_t}$. Similarly, the formulation for MDOF systems is derived from Eq. (8) as

$$p(\mathbf{z},t|\mathbf{z}',t') = (4\pi dt)^{-N_s/2}\left\{\det\left[\mathbf{R}^{-1}\tilde{\mathbf{G}}(\mathbf{z}',t')\mathbf{R}^{-T}\right]\right\}^{-1/2} \times \exp\left\{-\frac{1}{4dt}\left[\mathbf{z}_\mathbf{e}^T\tilde{\mathbf{G}}^{-1}(\mathbf{z}',t')\mathbf{z}_\mathbf{e}\right]\right\} \quad (10)$$

where $\mathbf{z_e} = \mathbf{Rz} - \mathbf{R}'\mathbf{z}' - \delta\tilde{\mathbf{f}}(\mathbf{z}',t')dt$, $\tilde{\mathbf{f}}(\mathbf{z},t) = \mathbf{f}(\mathbf{Rz}+\boldsymbol{\mu},t)$, and $\tilde{\mathbf{G}}(\mathbf{z},t) = \mathbf{G}(\mathbf{Rz}+\boldsymbol{\mu},t)$.

The mean and covariance are updated using the equations

$$\mu_{t+dt} = \mu_t + \langle\tilde{f}(z_t,t)\rangle_{z_t}dt \quad (11)$$

$$\sigma^2_{t+dt} = \sigma^2_t + \langle\tilde{A}^2(z_t,t)\rangle_{z_t}dt + 2\sigma_t\langle\delta\tilde{f}(z_t,t)z_t\rangle_{z_t}dt + \langle[\delta\tilde{f}(z_t,t)]^2\rangle_{z_t}dt^2 \quad (12)$$

for SDOF systems and

$$\boldsymbol{\mu}_{t+dt} = \boldsymbol{\mu}_t + \langle\tilde{\mathbf{f}}(\mathbf{z}_t,t)\rangle_{\mathbf{z}_t}dt \quad (13)$$

$$\boldsymbol{\Sigma}_{t+dt} = \boldsymbol{\Sigma}_t + 2\langle\tilde{\mathbf{G}}(\mathbf{z}_t,t)\rangle_{\mathbf{z}_t}dt + \langle[\mathbf{R}_t\mathbf{z}_t]\left[\delta\tilde{\mathbf{f}}(\mathbf{z}_t,t)\right]^T + \left[\delta\tilde{\mathbf{f}}(\mathbf{z}_t,t)\right][\mathbf{R}_t\mathbf{z}_t]^T\rangle_{\mathbf{z}_t}dt + \langle\left[\delta\tilde{\mathbf{f}}(\mathbf{z}_t,t)\right]\left[\delta\tilde{\mathbf{f}}(\mathbf{z}_t,t)\right]^T\rangle_{\mathbf{z}_t}dt^2 \quad (14)$$

for MDOF systems.

The update equation between times t and $t+dt$ for a discrete representation of the PDF is given as

$$\mathbb{P}_{t+dt} = \mathbb{B}_{t,dt}\mathbb{P}_t \quad (15)$$

where $\mathbb{P}_t$ is a vector comprised of the PDF values corresponding to the grid points in the transformed space at time t and $\mathbb{B}$ is the short-time propagator matrix derived from Eqs. (9) and (10), for SDOF and MDOF systems respectively, for all source and target states.

The mean and covariance update equations (Eqs. (11), (12), (13), and (14)) are similarly discretized as weighted averages.

In order to better maintain the stability and accuracy of the TPI, it is desirable to ensure that the propagator matrix $\mathbb{B}$ obeys the zeroth moment condition of a Markov transition matrix

$$\int p(\mathbf{z}, t|\mathbf{z}', t')d\mathbf{z} = 1 \tag{16}$$

through the following normalization process:

$$B_{i,j}^{new} = \frac{B_{i,j}^{old}}{\sum_i B_{ij}(t, dt)\Delta z_i} \tag{17}$$

where Δz_i is the transformed mesh size. Likewise, the zeroth moment condition can be preserved in the updated $\mathbb{P}_{t+dt}$ through a similar normalization.

$$p_i^{new}(t + dt) = \frac{p_i^{old}(t + dt)}{\sum_i p_i(t)\Delta z_i} \tag{18}$$

Enforcing these zeroth moment properties through normalization allows us to maintain greater accuracy in numerical implementations.

2.2 Bandlimiting

In signal theory, bandlimiting is the practice of restricting a signal's representation to a specific domain, beyond which all values are considered negligible and assumed to be zero. To apply bandlimiting to the propagator matrix of the TPI, we recognize that for every source state, there is a single most probable target state about which all significant probability is concentrated. By developing a function of the source state, we obtain a set of peaks hereafter referred to as the peak transition probability set (PTPS). We then

designate a bandwidth based on the Gaussian distribution of the Wiener process, which we use to reduce the computation and storage demands of the propagator matrix by only obtaining the matrix values within that specified bandwidth about the PTPS.

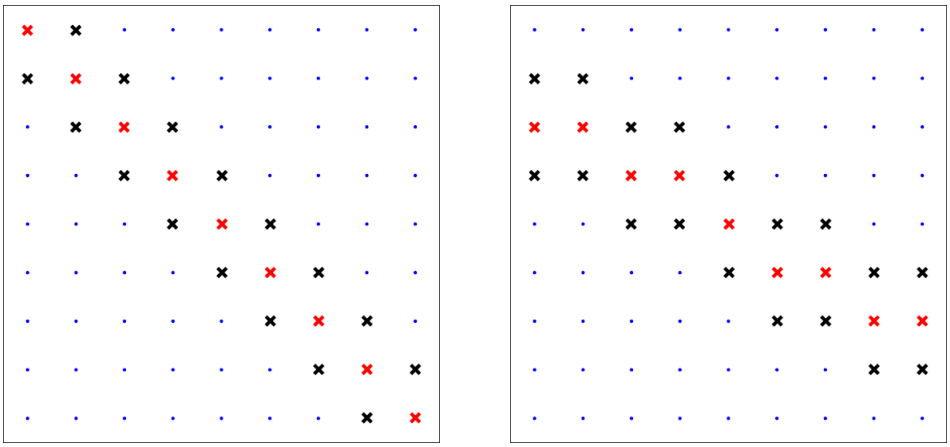

Figure 1: Model of a bandlimited propagator matrix for simple diffusion on a fixed grid (left) and transformed grid (right), with peak states in red.

2.2.1 Peak Transition Probability Curve

Each element B_{ij} in the short-time propagator matrix $\mathbb{B}$ represents the probability of the system reaching the ith target state from the jth source state in the transformed space. Since random excitation is modeled as a Gaussian Wiener process, there is a single peak target state for each source state. These peaks corresponds to the states obtained from the deterministic path in Eqs. (1) and (2), i.e., by setting the random excitation term to zero and, in the context of the transformed path integral method, incorporating the dynamic transformation in Eqs. (7) and (8). The PTPS can also be obtained by setting the gradient of $p(z,t|z',t')$ in Eqs. (9) and (10) to zero and isolating the target state z. In terms of transformed coordinates, the expressions for the peak states in SDOF and MDOF systems are, respectively, given by

$$z_{peak} = \frac{\sigma' z' + \delta \tilde{f}(z',t)dt}{\sigma} \tag{19}$$

$$\mathbf{z}_{peak} = \mathbf{R}^{-1}\left[\mathbf{R}'\mathbf{z}' + \delta\tilde{\mathbf{f}}(\mathbf{z}', t)\right] \tag{20}$$

The peak for each discrete source state can then be rounded to the nearest indexed target state. The resulting curve is the basis for developing an efficient bandlimited propagator matrix.

2.2.2 Diffusion Range

Because a Wiener process is characterized by Gaussian distribution, a reasonable range surrounding the PTPS can be calculated in terms of the strength A of the diffusion process in absolute space and $\tilde{A}$ in the transformed space, corresponding to a single standard deviation of its Gaussian distribution. In SDOF systems, any independent Wiener increment is characterized by the distribution $w(t+dt) - w(t) \sim N(0, dt)$. Consequently, a single standard deviation of the transformed random excitation term $\tilde{A}dw$ is given as

$$\sigma_{noise} = \tilde{A}\sqrt{dt} \tag{21}$$

By multiplying the factor N_σ, hereafter referred to as the "bandwidth strength", by the diffusion range, the semi-bandwidth—that is, the distance between the PTPS and the outer boundary of the bandlimited space, equal to half of the total bandwidth—can be obtained. Given a bandwidth strength N_σ, the bandlimited propagator matrix represents all transition probability within N_σ standard deviations of the Gaussian Wiener process about the PTPS.

In MDOF systems, there are multiple approaches to defining a suitable range about the PTPS. We begin by obtaining the principal semi-axes $\tilde{\mathbf{u}}$ of the ellipse represented by the transformed diffusion coefficient matrix $\tilde{\mathbf{A}}$, which can be accomplished using the principal axis theorem. If $\tilde{\mathbf{A}}$ is a diagonal matrix then these axes align with the axes of the

evaluated space and $\tilde{\mathbf{A}}\sqrt{dt}$ can be used as a basis for the bandwidth along each axis, but if that is not the case then an approximation is needed.

One approach is to apply the length of the maximum principal semi-axis $|\tilde{\mathbf{u}}_{max}|\sqrt{dt}$ as the basis for all axes, though this approach inevitably overestimates anisotropic diffusion processes. As an alternative, for each ith degree of freedom we recommend using the maximum component vector $\tilde{\mathbf{u}}_{i,max}$ among all principal semi-axes along the corresponding cardinal axis.

$$\sigma_{i,noise} = \tilde{\mathbf{u}}_{i,max}\sqrt{dt} \tag{22}$$

Similar to the SDOF case, this value must be multiplied by the user's chosen bandwidth strength N_σ and then converted to a number of indexed terms in both directions along each degree of freedom. This can be obtained from the relation $d\mathbf{x} = \mathbf{R}d\mathbf{z}$.

Due to the dynamic nature of the transformation used in TPI methods, the diffusion range must be recalculated at each iteration to keep it scaled to the transformed space. This increases the long-term efficiency of the TPI method when applied to expanding systems, while ensuring that accuracy is not lost over time in concentrating systems. In systems with state-invariant diffusion processes, to which we restrict our research in this study, only one range needs to be obtained at each discrete time step.

2.2.3 Bandlimiting for SDOF Case

Using the PTPS and semi-bandwidth values obtained in the previous two sections, it is possible to set boundaries beyond which all values in the short-time propagator matrix $\mathbb{B}$ are considered zero. In doing so we obtain a sparse matrix $\mathbb{C}$ as an approximation of the full propagator matrix $\mathbb{B}$, thereby reducing the computational cost and storage demands and potentially increasing the speed of the TPI implementation. By combining the upper and lower limits of the transformed data points z with the upper and lower limits of the

diffusion range about the PTPS, a set of boundaries can be obtained within which to gather data for a sparse matrix.

From these, the contents of the matrix $\mathbb{C}$ can be obtained and $\mathbb{C}$ can be substituted for $\mathbb{B}$ using sparse matrix software tools. This reduces the number of operations required to calculate all data points from $O(N^2)$ to $O(bN)$, where b is the total bandwidth in terms of indexed values. In the MDOF case, there is a bandwidth b_i for each individual ith degree of freedom, such that the total bandwidth is $b = \Pi_{i=1}^{n}(b_i)$ for an n-dimensional diffusion process. Because a sparse matrix indexes each data point in $\mathbb{C}$ by its coordinates and value, i.e. (i, j, C_{ij}), the total storage space required for a sparse matrix is $O(3bN)$. Therefore, if $3b < N$, then the storage space required by the BL-TPI is less than that of the standard TPI method.

2.2.4 Bandlimiting for MDOF Case

In order to extend the BL-TPI method to systems with $n > 1$ degrees of freedom, we examine two general approaches: bulk bandlimiting and operator splitting. The bulk bandlimiting method, which evaluates every data point within an n-dimensional space developed about the PTPS, is the more computationally expensive method but is able to capture more complex and detailed probability data. Operator splitting is examined as an alternative method that separates the MDOF BL-TPI implementation into a sequence of SDOF operations.

2.2.4.1 Bulk Bandlimiting

The most accurate approach to bandlimiting for MDOF systems is to use the diffusion range σ_{noise} along each axis to generate an n-orthotope or n-ellipsoid about each state in the PTPS. While an ellipsoid is theoretically more efficient, its boundaries are more complex and difficult to accurately and effectively code. Therefore, for the purpose of

this study, we will be developing an orthotope space about the PTPS whose lengths along each axis are the corresponding bandwidths.

2.2.4.2 Operator Splitting

An alternative to bulk bandlimiting considered in this thesis is the method of operator splitting. Operator splitting approximation exploits the property $y_1 = e^{(L_1+L_2)t}y_0 = e^{L_1 t}e^{L_2 t}y_0$ to separate a single convoluted system into a set of SDOF systems that operate along each degree of freedom. By approximating the MDOF short-time propagator matrix $\mathbb{B}$ as a sequence of SDOF propagations, an approximation of the time-transformed PDF can be obtained. In a 2D problem, the simplest form of operator splitting, hereafter referred to as "simple splitting", is implemented through the following steps:

$$\tilde{p}_{t'} = e^{L_{z_1} dt} p_{t'} \tag{23}$$

$$p_t = e^{L_{z_2} dt} \tilde{p}_{t'} \tag{24}$$

Alternative forms of the operator splitting process include Strang splitting, derived from the identity $e^{(L_1+L_2)dt}y_0 = e^{(\frac{1}{2}L_1+L_2+\frac{1}{2}L_1)dt}y_0$ and implemented through the steps

$$\tilde{p}_{t'} = e^{L_{z_1} \frac{dt}{2}} p_{t'} \tag{25}$$

$$\overline{p}_{t'} = e^{L_{z_2} dt} \tilde{p}_{t'} \tag{26}$$

$$p_t = e^{L_{z_1} \frac{dt}{2}} \overline{p}_{t'} \tag{27}$$

and high-order splitting, derived from the identity $e^{(L_1+L_2)dt}y_0 = e^{(L_2+L_1)dt}y_0$, implemented through the steps

$$\tilde{p}_{t',1} = e^{L_{z_1} dt} p_{t'} \tag{28}$$

$$p_{t,1} = e^{L_{z_2} dt} \tilde{p}_{t',1} \tag{29}$$

$$\tilde{p}_{t',2} = e^{L_{z_2} dt} p_{t'} \tag{30}$$

$$p_{t,2} = e^{L_{z_1} dt} \tilde{p}_{t',2} \tag{31}$$

$$p_t = \frac{1}{2}(p_{t,1} + p_{t,2}) \tag{32}$$

Expanding these methods to systems with three or more degrees of freedom is beyond the scope of this study, but variants of the operator splitting process for higher dimensions may be developed by substituting $e^{(L_n + L_{n+1})dt}$ for $e^{L_n dt}$.

The accuracy of the operator splitting process can be improved if the update equations for the mean and variance (Eqs. (13) and (14)) are performed between operations. This can be accomplished by reducing all but the nth rows of $\tilde{\mathbf{f}}(\mathbf{z}_t, t)$ and $\tilde{\mathbf{G}}(\mathbf{z}_t, t)$ to zero when operating along the nth axis. The overall accuracy of the operator splitting method is also found to improve when the normalization process, Eq. (18), is performed between operations.

2.3 Symmetric Fast Gauss Transform

The Symmetric Fast Gauss Transform (SFGT) is a method in the family of Fast Multipole Methods intended to reduce the computational complexity of convolution operations like path integrals by separating convoluted terms (e.g., Di Paola et al., 2009). This method is obtained from the path integral solution using the following key substitution:

$$\exp\left\{-\frac{(x-x')^2}{l^2}\right\} = \exp\left\{-\left(\frac{x-s}{l}\right)^2\right\} \exp\left\{-\left(\frac{x'-s}{l}\right)^2\right\} \times \sum_{k=0}^{\infty} \frac{2^k}{k!} \left\{\frac{x-s}{l}\right\}^k \left\{\frac{x'-s}{l}\right\}^k \tag{33}$$

This relationship can be proven by recognizing that $(x-x')^2 = [(x-s)-(x'-s)]^2 = (x-s)^2 + (x'-s)^2 - 2(x-s)(x'-s)$ and applying Taylor series expansion to the exponential function $\exp\left\{2\frac{(x-s)}{l}\frac{(x'-s)}{l}\right\}$. This enables us to separate the convolution of x and x'.

2.3.1 SDOF Transformed SFGT Formulation

The purpose of an SFGT path integral formulation is to separate source and target terms. For single degree of freedom systems, this formulation can be obtained by substituting $g(x',t') = x' - f(x',t')dt$ for x', setting $s = 0$ and $l = A\sqrt{2dt}$, and applying the relation in Eq. (33) to the PI formulation, Eqs. (4) and (6).

$$p(x,t) = \frac{1}{l\sqrt{\pi}} \sum_{k=0}^{\infty} \frac{2^k}{k!} \exp\left\{-\left(\frac{x}{l}\right)^2\right\} \left\{\frac{x}{l}\right\}^k \times \int \exp\left\{-\left(\frac{g(x',t')}{l}\right)^2\right\} \left\{\frac{g(x',t')}{l}\right\}^k p(x',t')dx' \quad (34)$$

In order to create a TPI formulation of the SFGT, Eq. (34) can be revised to iterate in the transformed space. By incorporating the relation in Eq. (7), a new TPI SFGT formulation is obtained

$$p(z,t) = \frac{1}{\tilde{l}\sqrt{\pi}} \sum_{k=0}^{\infty} \frac{2^k}{k!} \exp\left\{-\left(\frac{z}{\tilde{l}}\right)^2\right\} \left\{\frac{z}{\tilde{l}}\right\}^k \times \int \exp\left\{-\left(\frac{\tilde{g}(z',t')}{\tilde{l}}\right)^2\right\} \left\{\frac{\tilde{g}(z',t')}{\tilde{l}}\right\}^k p(z',t')dz' \quad (35)$$

where $\tilde{l} = l/\sigma$ and $\tilde{g}(z',t') = (\sigma' z' - \delta\tilde{f}(z',t')dt)/\sigma$.

Because the convolution of z and z' is separated, the integral component of the Eq. (35) can be calculated separately. Given a finite summation size of K, the integral com-

ponent of is of computational complexity $O(KN)$ for a space with N possible states, and the subsequent application of the obtained results to the summation is of equal complexity. The total computational complexity of the SFGT formula in SDOF cases is therefore $O(2KN)$. A larger K value results in greater accuracy, such that as K approaches infinity, the solution to the SFGT formulation approaches that of the TPI.

2.3.2 MDOF Transformed SFGT

The SFGT formulation in Eq. (35) can be extended to MDOF systems using the MDOF equivalent of Eq. (33). In this context, $\mathbf{M} \equiv \mathbf{L}\mathbf{L}^T$ and is therefore always symmetrical. Because of this symmetry, $\mathbf{x}^T\mathbf{M}^{-1}\mathbf{x}' = \mathbf{x}'^T\mathbf{M}^{-1}\mathbf{x}$, resulting in the following equation:

$$\begin{aligned}\exp\left\{-(\mathbf{x}-\mathbf{x}')^T\mathbf{M}^{-1}(\mathbf{x}-\mathbf{x}')\right\} &= \exp\left\{-(\mathbf{x}-\mathbf{s})^T\mathbf{M}^{-1}(\mathbf{x}-\mathbf{s})\right\} \\ \times \exp\left\{-(\mathbf{x}'-\mathbf{s})^T\mathbf{M}^{-1}(\mathbf{x}'-\mathbf{s})\right\} &\times \sum_{k=0}^{\infty}\frac{2^k}{k!}\left\{\mathbf{L}^{-1}(\mathbf{x}-\mathbf{s})\right\}^{[k]T}\left\{\mathbf{L}^{-1}(\mathbf{x}'-\mathbf{s})\right\}^{[k]}\end{aligned} \tag{36}$$

Here $\mathbf{x}^{[k]}$ represents the Kronecker power, which is the k-fold Kronecker product of $\mathbf{x}$, i.e., $\mathbf{x} \otimes \mathbf{x} \otimes ... \otimes \mathbf{x}$.

Setting $\mathbf{s} = 0$, $\mathbf{L} = \mathbf{A}\sqrt{2dt}$ and substituting $\mathbf{G}_x(\mathbf{x}', t') = \mathbf{x}' - f(\mathbf{x}', t')dt$ for $\mathbf{x}'$ allows us to obtain the following formulation for the fixed grid MDOF SFGT path integral implementation.

$$\begin{aligned}p(\mathbf{x}, t) = \{\pi\}^{-N_s/2}\left\{\det(\mathbf{M})\right\}^{-1/2} &\times \sum_{k=0}^{\infty}\frac{2^k}{k!}\exp\left\{-\mathbf{x}^T\mathbf{M}^{-1}\mathbf{x}\right\}\left\{\mathbf{L}^{-1}\mathbf{x}\right\}^{[k]T} \\ \times \int \exp\left\{-\mathbf{G}_x(\mathbf{x}', t')^T\mathbf{M}^{-1}\mathbf{G}_x(\mathbf{x}', t')\right\}&\left\{\mathbf{L}^{-1}\mathbf{G}_x(\mathbf{x}', t')\right\}^{[k]} p(\mathbf{x}', t')\Pi_{n=1}^{N_s}(d\mathbf{x}'_n)\end{aligned} \tag{37}$$

Applying the transformation in Eq. (8) to this relation allows us to obtain the MDOF SFGT TPI formulation.

$$p(\mathbf{z},t) = \{\pi\}^{-N_s/2}\left\{\det(\tilde{\mathbf{M}})\right\}^{-1/2} \times \sum_{k=0}^{\infty} \frac{2^k}{k!} \exp\left\{-\mathbf{z}^T\tilde{\mathbf{M}}^{-1}\mathbf{z}\right\}\left\{\tilde{\mathbf{L}}^{-1}\mathbf{z}\right\}^{[k]T}$$
$$\times \int \exp\left\{-\tilde{\mathbf{G}}_z(\mathbf{z}',t')^T\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{G}}_z(\mathbf{z}',t')\right\}\left\{\tilde{\mathbf{L}}^{-1}\tilde{\mathbf{G}}_z(\mathbf{z}',t')\right\}^{[k]} p(\mathbf{z}',t')\Pi_{n=1}^{N_s}(d\mathbf{z}'_n) \quad (38)$$

The inclusion of a Kronecker product significantly increases the computational complexity of the SFGT method in MDOF systems. Given a summation size of K and a dimensionality of n, the complexity can be approximated from the maximum Kronecker power n^K to be $O(2n^K KN)$ for a space with N possible states.

2.3.3 Gauss-Hermite Quadrature

A prospective method for further reducing the cost of the TPI method that is compatible with the SFGT formulation is through the use of the Gauss-Hermite quadrature, which is a method for approximating an integral using a relatively small number of weighted data points. The Gauss-Hermite quadrature can be applied to integrals of the form

$$I = \int_{-\infty}^{\infty} \exp\left\{-x^2\right\} f(x)dx \quad (39)$$

to create the following discrete approximation

$$\hat{I} \approx \sum_{i=1}^{n} w_i f(\zeta_i) \quad (40)$$

Here ζ_i are the roots of the nth order Physicist's Hermite polynomial and $w_i = \frac{2^{n-1}n!\sqrt{\pi}}{n^2[H_{n-1}(\zeta_i)]^2}$. For any nth order Hermite polynomial there exist n roots and weights. The

Physicist's Hermite polynomial is given as

$$H_n(z) = (-1)^n \exp\left\{z^2\right\} \frac{d^n}{dz^n} \exp\left\{-z^2\right\} \tag{41}$$

In order to incorporate the Gauss-Hermite quadrature into the SFGT formulation, Eq. (35) must be rewritten in order for the integral component to match the form of the left side of Eq. (39). We rewrite it as follows:

$$p(z,t) = \frac{1}{\tilde{l}\sqrt{\pi}} \sum_{k=0}^{\infty} A_k(z) B_k \tag{42}$$

$$A_k(z) = \frac{2^k}{k!} \exp\left\{-\left(\frac{z}{\tilde{l}}\right)^2\right\} \left\{\frac{z}{\tilde{l}}\right\}^k \tag{43}$$

$$B_k = \int \exp\left\{-z'^2\right\} \exp\left\{z'^2 - \left(\frac{\tilde{g}(z',t')}{\tilde{l}}\right)^2\right\} \left\{\frac{\tilde{g}(z',t')}{\tilde{l}}\right\}^k p(z',t') dz' \tag{44}$$

In this form, the source term component B_k can be discretized. The discrete approximation of Eq. (44) is given as

$$B_k \approx \sum_{i=1}^{n} w_i f(\zeta_i) \tag{45}$$

$$f(\zeta_i) = \exp\left\{\zeta_i'^2 - \left(\frac{\tilde{g}(\zeta_i',t')}{\tilde{l}}\right)^2\right\} \left\{\frac{\tilde{g}(\zeta_i',t')}{\tilde{l}}\right\}^k p(\zeta_i',t') \tag{46}$$

The Gauss-Hermite quadrature formulation can then be rewritten in terms of ζ_i.

$$p(\zeta_i,t) = \frac{1}{\tilde{l}\sqrt{\pi}} \sum_{k=0}^{\infty} A_k(\zeta_i) B_k \tag{47}$$

Because there is no uniform grid size dz, the resulting PDF can instead be normalized

using the corresponding weights and abscissas.

$$p_i^{new}(t+dt) = \frac{p_i^{old}(t+dt)}{\sum_i w_i \exp\left\{\zeta_i^2\right\} p_i^{old}(t+dt)} \tag{48}$$

In order to generalize the Gauss-Hermite quadrature to address MDOF systems, we consider the multivariate integral of N_s dimensions

$$I = \int ... \int \exp\left\{-\sum_{i=1}^{N_s}(x_i^2)\right\} f(x_1, ..., x_{N_s}) \prod_{i=1}^{N_s}(dx_i) \tag{49}$$

This integral can be discretely approximated as

$$\hat{I} \approx \sum_{i_1=1}^{n} ... \sum_{i_{N_s}=1}^{n} \prod_{j=1}^{N_s}(w_j) f(\zeta_{i_1}, ..., \zeta_{i_{N_s}}) \tag{50}$$

Eq. (38) can then be rewritten in the form of Eq. (49).

$$p(\mathbf{z}, t) = \{\pi\}^{-N_s/2} \left\{\det(\tilde{\mathbf{M}})\right\}^{-1/2} \times \sum_{k=0}^{\infty} \mathbf{A}_k(\mathbf{z}) \mathbf{B}_k \tag{51}$$

$$\mathbf{A}_k(\mathbf{z}) = \frac{2^k}{k!} \exp\left\{-\mathbf{z}^T \tilde{\mathbf{M}}^{-1} \mathbf{z}\right\} \left\{\tilde{\mathbf{L}}^{-1} \mathbf{z}\right\}^{[k]T} \tag{52}$$

$$\mathbf{B}_k = \int ... \int \exp\left\{-\Sigma_{n=1}^{N_s}(z_n^2)\right\} \times \exp\left\{\Sigma_{n=1}^{N_s}(z_n^2) - \tilde{\mathbf{G}}_z(\mathbf{z}', t')^T \tilde{\mathbf{M}}^{-1} \tilde{\mathbf{G}}_z(\mathbf{z}', t')\right\} \times \left\{\tilde{\mathbf{L}}^{-1} \tilde{\mathbf{G}}_z(\mathbf{z}', t')\right\}^{[k]} p(\mathbf{z}', t') \prod_{n=1}^{N_s}(d\mathbf{z}'_n) \tag{53}$$

This takes advantage of the scalar nature of $\tilde{\mathbf{G}}_z(\mathbf{z}', t')^T \tilde{\mathbf{M}}^{-1} \tilde{\mathbf{G}}_z(\mathbf{z}', t')$. From this, $\mathbf{B}_k$ can be discretized as

$$\mathbf{B}_k \approx \sum_{i_1=1}^{n} ... \sum_{i_{N_s}=1}^{n} \prod_{j=1}^{N_s}(w_j) \mathbf{f}(\boldsymbol{\zeta}'_i, t') \tag{54}$$

$$\mathbf{f}(\boldsymbol{\zeta}_i', t') = \exp\left\{\sum_{n=1}^{N_s}(z_n^2) - \tilde{\mathbf{G}}_z(\mathbf{z}', t')^T \tilde{\mathbf{M}}^{-1} \tilde{\mathbf{G}}_z(\mathbf{z}', t')\right\} \times \left\{\tilde{\mathbf{L}}^{-1}\tilde{\mathbf{G}}_z(\mathbf{z}', t')\right\}^{[k]} p(\mathbf{z}', t') \prod_{n=1}^{N_s}(d\mathbf{z}_n') \quad (55)$$

Finally, the MDOF formulation for the Gauss-Hermite quadrature can be obtained in terms of $\boldsymbol{\zeta}_i$.

$$p(\boldsymbol{\zeta}_i, t) = \{\pi\}^{-N_s/2}\left\{\det(\tilde{\mathbf{M}})\right\}^{-1/2} \times \sum_{k=0}^{\infty} \mathbf{A}_k(\boldsymbol{\zeta}_i)\mathbf{B}_k \quad (56)$$

3 Bandlimiting Results and Analysis

In this chapter we demonstrate the effectiveness of the bandlimited transformed path integral (BL-TPI) method by applying it to a set of canonical systems and performing error analysis in comparison to the standard TPI implementation. In order to develop a recommended bandwidth, in Examples 1-4 we compare the performance of the BL-TPI method at multiple bandwidths in single degree of freedom (SDOF) systems. Through these examples, we demonstrate the ability of the BL-TPI method to solve systems characterized by pure diffusion, constant drift, the Ornstein-Uhlenbeck process, and nonlinear drift processes. We compare the PDFs of the TPI and BL-TPI methods—p_0 and p respectively—using the root mean squares (RMS) method of error analysis, in which N denotes the number of data points.

$$\epsilon_{RMS} = \sqrt{\frac{1}{N}\sum_{n=1}^{N}(p_0(n) - p(n))^2} \tag{57}$$

We then examine a series of canonical multiple degree of freedom (MDOF) systems in Examples 5-8, in which we examine the effectiveness of the simple, Strang, and high-order methods of operator splitting, the bulk bandlimiting method, and the standard TPI. Through comparative error analysis, we develop recommendations for which bandlimiting method to use. Through these examples, we find that while the bulk BL-TPI produces consistently accurate results for all systems, it is also the most computationally expensive method of MDOF bandlimiting; in systems that meet certain conditions, which we discuss later, we find that operator splitting formulations of the BL-TPI can match the accuracy of the bulk formulation, and therefore recommend their usage over the bulk formulation under such conditions.

3.1 Single Degree of Freedom Systems

We begin by observing the effectiveness of the bandlimited TPI method in one dimensional state spaces in order to demonstrate the accuracy of this method and to develop a recommended bandwidth. It is worth noting that at higher bandwidths, the PDF output of the BL-TPI method approaches that of the standard TPI, not the exact solution of the system. Consequently, once the error of the standard TPI is closely matched, there is little benefit to extending the bandwidth further.

3.1.1 Example 1: Pure Diffusion in 1D

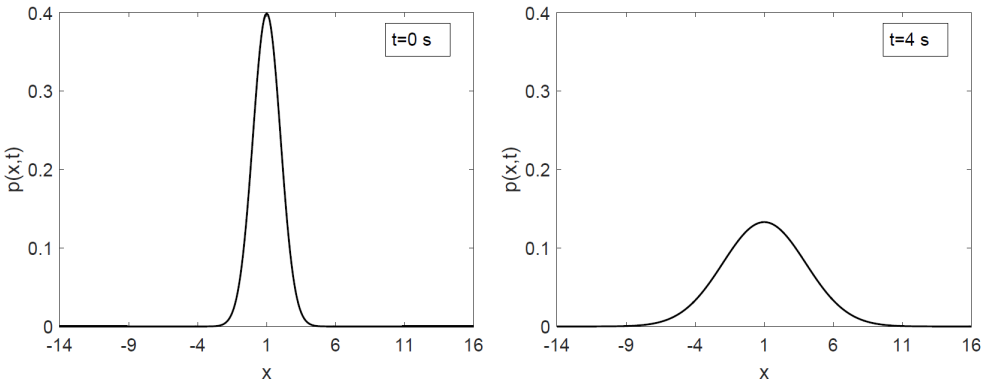

Figure 2: PDF of pure diffusion in the in absolute space at $t = 0\ s$ (left) and $t = 4\ s$ (right), as determined with the BL-TPI at a strength of 5.

Systems of pure diffusion—that is, systems experiencing zero nonrandom drift—are rudimentary examples of the time transformation of a PDF, as well as a useful baseline for demonstrating the effectiveness of the TPI method. In an Itô stochastic system with zero drift, the peak target state is always identical to its source state in absolute space. In the case of the standard PI formulation, this behavior manifests in the short-time propagator matrix as a perfectly diagonal PTPS.

By contrast, the principal diagonal of the short-time propagator matrix of the TPI

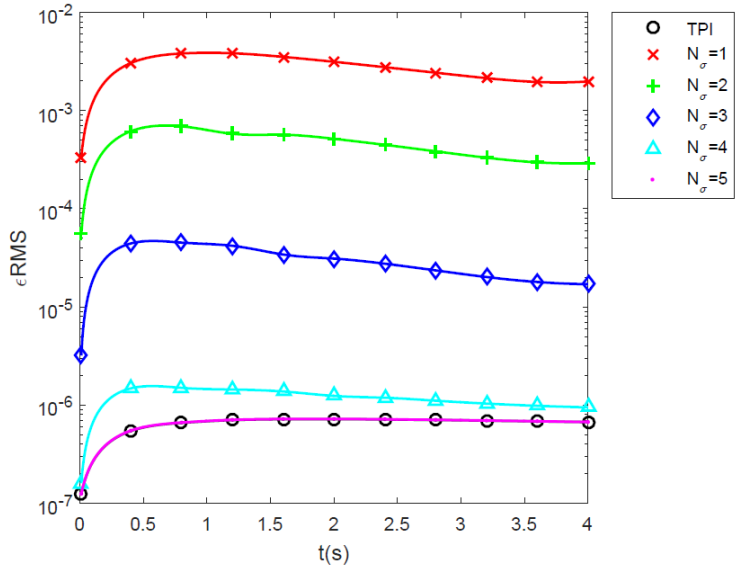

Figure 3: Comparison of the time evolution of RMS error relative to the complete TPI at different bandlimiting strengths N_σ in a system under pure diffusion.

method does not necessarily represent zero movement along its principal diagonal. This is due to the dynamic nature of the transformation, resulting in the source positions and target positions existing in different transformed spaces. In a simple diffusion case, the function for the peak states simplifies to $z_{peak} = \sigma' z' / \sigma$.

In this example, we evaluate the simple diffusion process governed by the following stochastic equation

$$dx_t = A dw_t \tag{58}$$

where the strength of the Wiener process is set to $A = \sqrt{2}$. A time increment of $dt = 0.01\ s$ is iterated over the time period of 0 s to 4 s. The initial Gaussian distribution $p(x, t_0) \sim N(\mu_0, \sigma_0^2)$ has a mean $\mu_0 = 1$ and a covariance $\sigma_0^2 = 1$, and the initial grid extends from -4 to 6 in absolute space over 201 grid points.

Because the drift and diffusion are both constant and therefore independent of the PDF and any error therein, there is zero error in the update equations for the mean and

variance (Eqs. (11) and (12)), and therefore loss of resolution is not a concern. The error that occurs even at $N_\sigma = 1$ is so small that a visual comparison with the complete TPI is not useful; however, a comparison of the time evolution of RMS error demonstrates the increase in accuracy that results from using a greater bandwidth.

Fig. 3.1.1 shows error analysis of the BL-TPI method with bandwidth strengths from 1 to 5. There is a clear decrease in error that accompanies each increase in bandwidth, and at a strength of 5 the error curve of the BL-TPI method is nearly identical to that of the standard TPI; $t = 4\ s$, the difference in error between the BL-TPI, 6.74×10^{-7}, and the standard TPI, 6.72×10^{-7}, becomes insignificant. Furthermore, all error curves indicate convergence, indicating that the BL-TPI method is a numerically stable approach.

3.1.2 Example 2: Constant Drift and Diffusion

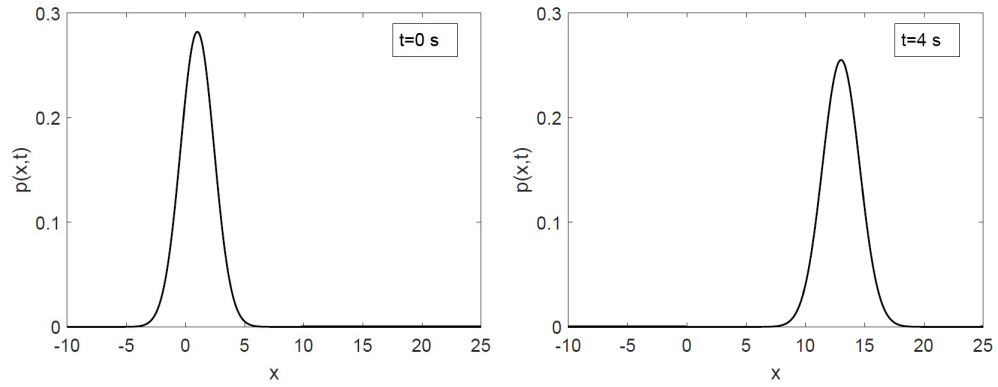

Figure 4: PDF of constant drift and diffusion in absolute space at $t = 0\ s$ (left) and $t = 4\ s$ (right), as determined with the BL-TPI at a strength of 5.

One of the most useful features of the TPI method is its ability to account for drift in order to limit the tail loss that occurs outside the boundaries of the grid. When applied to a system with constant drift, the transformed grid of the TPI shifts continuously to evaluate the translating system with the same accuracy and effectiveness as a stationary system of pure diffusion. Consequently, the BL-TPI method can also be expected to demonstrate

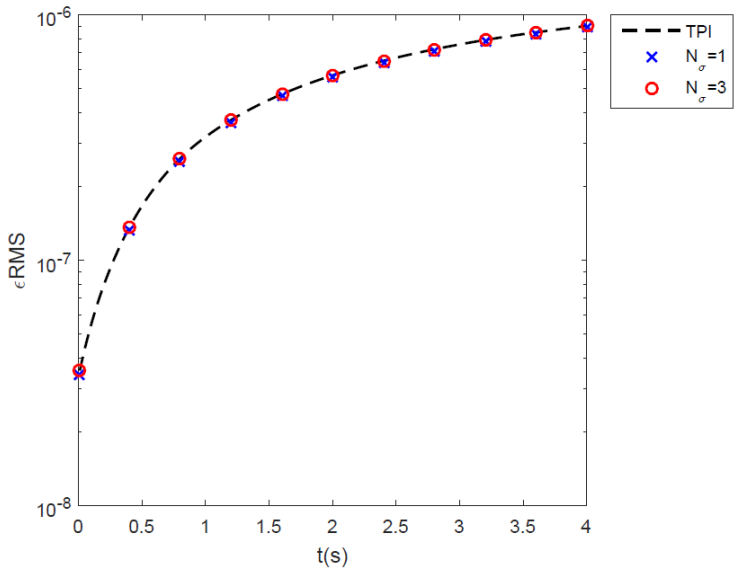

Figure 5: Comparison of the time evolution of RMS error at different bandlimiting strengths N_σ in a system under constant drift and diffusion.

similar accuracy to that seen in Example 1 when applied to a system of constant drift.

For this example, we evaluate a system governed by the equation

$$dx_t = \kappa dt + A dw_t \tag{59}$$

Here the drift is given as $\kappa = 3$ and the white noise strength is given as $A = \frac{1}{3}$. This process is iterated in increments of $dt = 0.01\ s$ over a period from $t = 0\ s$ to $t = 4\ s$. The initial Gaussian distribution is set to $p(x, t_0) \sim N(1, 2)$, and the number of grid points is 201, as in Example 1. We begin with bounds of $[1 - 5\sqrt{2}, 1 + 5\sqrt{2}]$ in absolute space which, given an initial covariance of $\sigma^2 = 2$, becomes $[-5; 5]$ in the transformed space.

Like in the case of pure diffusion, both the mean and variance update equations are independent of the PDF and therefore error-free. Unlike in Example 1, however, the results of the bandlimited TPI are very close to those of the complete TPI even at a bandwidth strength of 1. Transient RMS error analysis demonstrates a similar pattern of an initial

rise in error, followed by a trend toward convergence. Notably because a relatively small white noise strength is used in this example, the accuracy of the BL-TPI is very high even at a bandwidth strength of 1. Additionally, close examination reveals that the errors at bandwidth strengths 1 and 2 are exactly identical, at 9.01×10^{-7}, while the errors for all higher bandwidths from 3 to 5 are also identical at 8.88×10^{-7}. This is a product of the reliance of this bandlimiting method on rounding, such that for systems with very small diffusion terms, different bandwidth strengths can sometimes round to the same bandwidth.

3.1.3 Example 3: OU Process

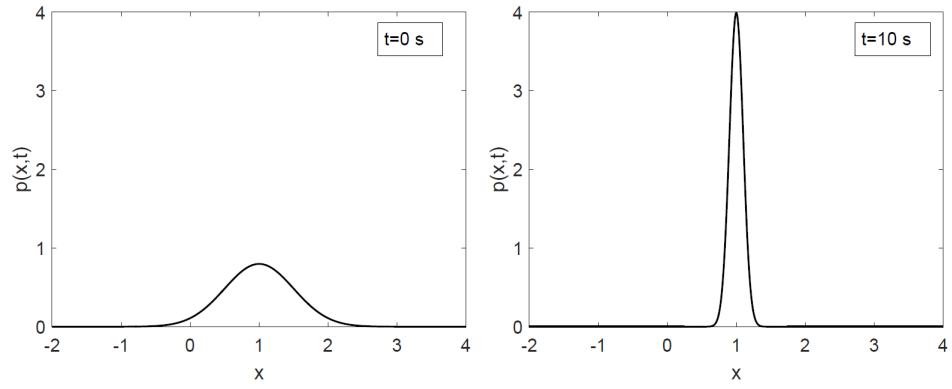

Figure 6: PDF of the OU process in absolute space at $t = 0\ s$ (left) and approaching the steady state solution at $t = 10\ s$ (right), as determined with the BL-TPI at a strength of 5.

The Ornstein-Uhlenbeck (OU) process is a common and simple form of nonuniform linear drift characterized by a single fixed point and the concentration of the PDF toward that point. The TPI is equipped to accurately solve problems of this nature because it is able to maintain an effective resolution as the concentration necessitates a finer grid. When applied to an OU process, the bandwidth about the PTPS is expected to expand over time because the diffusion coefficient, held constant, becomes larger relative to the shrinking transformed space.

The differential equation of the OU process is given as

$$dx_t = -\beta(x-\theta)dt + Adw_t \tag{60}$$

Here $\beta = \frac{1}{2}$ and $\theta = 1$, and the white noise strength is $A = 0.1$. We choose to run an initial distribution of $p(x,t_0) \sim N(1,0.25)$ in a space bounded by $[-2,4]$ and divided into 241 grid points.

The OU process is known to approach an analytical steady state solution of the Gaussian distribution $p(x,t_\infty) \sim N(\mu, A^2/2\beta) = N(1,0.01)$, which can be compared to the result of the TPI after 10 seconds in this example. The complete TPI produces an error of 2.42×10^{-3}, and even at a bandwidth strength of 1, the BL-TPI method yields a similar error of 2.96×10^{-3}. All subsequent errors at higher strengths approach the exact accuracy of the standard TPI method more closely, and at $N_\sigma = 5$ the BL-TPI error of 2.42×10^{-3} is identical to that of the standard TPI within three significant figures. It should be noted that while a smaller bandwidth is needed to converge to the error of the standard TPI in this example, that error is much larger in this example than in Examples 1 and 2.

3.1.4 Example 4: Nonlinear Drift Process

We next observe the applicability of the BL-TPI method to systems that demonstrate complex nonlinear drift processes and facilitate transformations into non-Gaussian PDFs. To illustrate the effectiveness of this method, we apply the BL-TPI to the following process

$$dx_t = -\gamma x(x^2-\delta)dt + Adw_t \tag{61}$$

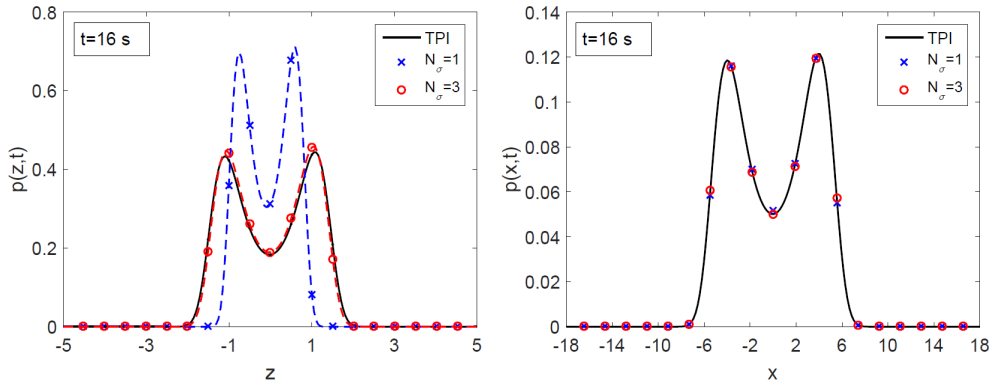

Figure 7: Comparison of PDFs of a nonlinear drift process at $t = 16\ s$ in the transformed space (left), and in absolute space (right) for the standard TPI and the BL-TPI at $N_\sigma = 1$ and $N_\sigma = 3$.

where $\gamma = \frac{1}{16}$, $\delta = 16$, and the white noise strength is $A = 3$. The initial distribution is given as $p(x, t_0) \sim N(1, 2)$, and the process is iterated in increments of $dt = 0.01\ s$ over a time period of 16 seconds.

Unlike in previous examples, the mean and covariance update equations do not exactly reproduce the results of the TPI when applied to this system. As a consequence, at small bandwidths the difference in shape in the transformed space between the BL-TPI and the standard TPI is very pronounced, as seen in Figure 6. Interestingly, when comparing the PDFs in absolute space, the error is found to be quite small even when the mean and covariance diverge significantly. This suggests that error in the update equations does not directly impact the overall accuracy of the BL-TPI method, beyond the error associated with smaller bandwidths in general. Nonetheless, it is still desirable to minimize error in the mean and covariance calculations in order to avoid complications due to loss of tail information and resolution. As seen in Figure 3.1.4, at a bandlimiting strength of $N_\sigma = 1$ the PDF is shown to have concentrated significantly by $t = 16$, potentially straining its resolution. By contrast, at $N_\sigma = 3$ the deviations in mean and covariance are much smaller and the resulting transformed space retains dimensions very close to the

results of the complete TPI.

To the author's knowledge at the time of writing, no transient solutions to this example exist. The steady state solution, however, is known and given by

$$p(x, t_\infty) = \frac{1}{\eta} \exp\left\{-\gamma \frac{x^2}{A^2}\left(\frac{x^2}{2} - \delta\right)\right\} \tag{62}$$

where η is the normalization factor. At $t = 16\ s$ the TPI is accurate to this solution with an RMS error of $\epsilon_{RMS} = 6.45 \times 10^{-4}$. At a bandwidth strength of 1, the bandlimited TPI is found to be reasonably accurate to this solution, at $\epsilon_{RMS} = 9.40 \times 10^{-4}$, despite the apparent loss of resolution in the transformed space. As documented in Table 1 in the following section, the error at a bandwidth strength of 4 is surprisingly smaller than those of both the standard TPI and the BL-TPI at a strength of 5. This phenomenon demonstrates that as the accuracy of the BL-TPI converges to that of the complete TPI, it can potentially err slightly toward the analytical solution rather than away from it.

3.1.5 Error Analysis

Table 1: Time and error of the standard TPI vs. the bandlimited TPI at different bandlimiting strengths in 1D.

	Pure Diffusion (4 s)		OU Process (10 s)		Nonlinear Drift (16 s)	
	T_{comp} (s)	ϵ_{RMS}	T_{comp} (s)	ϵ_{RMS}	T_{comp} (s)	ϵ_{RMS}
TPI	1.36×10^{0}	6.72×10^{-7}	1.36×10^{0}	2.42×10^{-3}	4.97×10^{0}	6.45×10^{-4}
$N_\sigma = 1$	7.28×10^{-1}	1.98×10^{-3}	7.30×10^{-1}	2.96×10^{-3}	2.21×10^{0}	9.40×10^{-4}
$N_\sigma = 2$	7.39×10^{-1}	2.92×10^{-4}	7.38×10^{-1}	2.81×10^{-3}	2.25×10^{0}	8.23×10^{-4}
$N_\sigma = 3$	7.45×10^{-1}	1.71×10^{-5}	7.45×10^{-1}	2.52×10^{-3}	2.33×10^{0}	6.47×10^{-4}
$N_\sigma = 4$	7.58×10^{-1}	9.49×10^{-7}	7.56×10^{-1}	2.43×10^{-3}	2.35×10^{0}	6.44×10^{-4}
$N_\sigma = 5$	7.64×10^{-1}	6.74×10^{-7}	7.63×10^{-1}	2.42×10^{-3}	2.42×10^{0}	6.45×10^{-4}

As shown in Table 1, when applied to SDOF examples, the BL-TPI becomes approx-

imately twice as fast as the complete TPI. The number of data points in all examples is between 200 and 250. Due to the Gaussian nature of the Wiener process, the bandwidth required for a system is directly proportional to the white noise strength of that system. Therefore it can be expected that for larger data sets with relatively small white noise strengths, the amount of time saved will be greater, especially when applied to expanding systems over long periods of time.

Interestingly, the effect of increasing or decreasing the bandlimiting strength on the computation time appears minimal, though this may be due to a limitation of the software and code used for this study rather than a fundamental property of the process. It must be emphasized that the exact speed of the BL-TPI method is beyond the scope of this study. The effectiveness of the BL-TPI implementation will largely depend on how cleanly and skillfully it is programmed.

While increasing the bandlimiting strength of this formulation is shown to reduce the error of the results, it also increases the time and storage space requirements to run a simulation, and as such we wish to establish a recommended bandlimiting strength. At $N_\sigma = 5$ the error of the BL-TPI becomes nearly identical to that of the complete TPI in all examples, though in some systems the difference in error is very small at strengths of as little as $N_\sigma = 3$. Whether the error at a bandlimiting strength of 3 or 4 is acceptable or not will depend on one's individual criteria and judgment; based on the results of this study, we believe that a strength of 5 can be generally recommended as a reliable convention.

3.2 Multiple Degree of Freedom Systems

In MDOF systems, we examine two general approaches to combining the bandlimiting ranges developed for each axis: operator splitting and bulk bandlimiting. For the purpose of this study we focus on three operator splitting processes: simple splitting, Strang

splitting, and high-order splitting. While operator splitting methods are faster and less computationally expensive compared to bulk bandlimiting, their effectiveness appears to be highly situational. When specific conditions are met, we find that bandlimiting with operator splitting can match the accuracy of bulk bandlimiting and the conventional TPI exactly. When these conditions are not met, the resulting error is found in some cases to be significantly greater than that of the standard TPI, and to not converge toward the exact outcome of the TPI at higher bandwidth strengths.

3.2.1 Example 5: Isotropic Diffusion in 2D

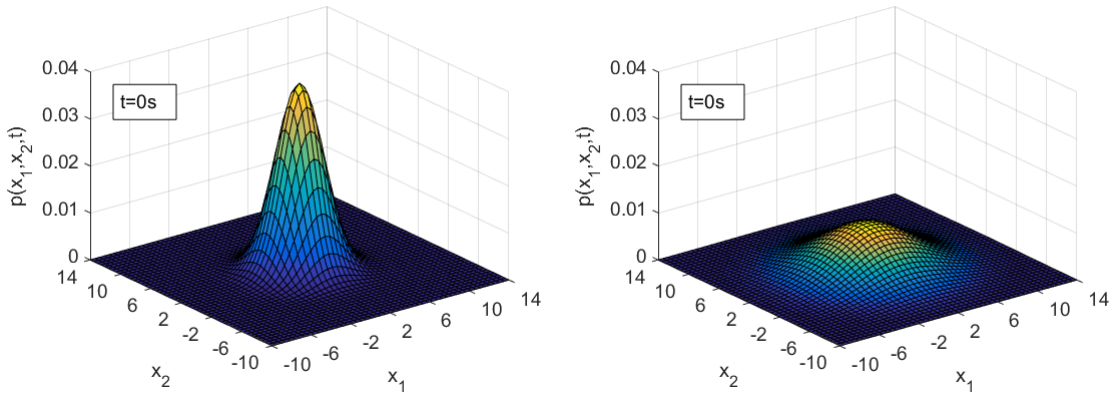

Figure 8: PDF of isotropic 2D pure diffusion in absolute space at $t = 0\ s$ (left), and at $t = 3\ s$ (right), as determined with the bulk bandlimited TPI at a strength of 5.

As with SDOF systems previously, we begin by examining pure diffusion in 2D as a rudimentary example of both the standard TPI and the bandlimited TPI. We model a 2D system under pure diffusion using the following governing equation:

$$d\mathbf{x}_t = \mathbf{A}d\mathbf{w}_t \tag{63}$$

For this example, we set the diffusion coefficient matrix $\mathbf{G} = \frac{1}{2}(\mathbf{A}\mathbf{A}^T) = [2,0;0,2]$. We apply this to an initial PDF with Gaussian distribution characterized by a mean of

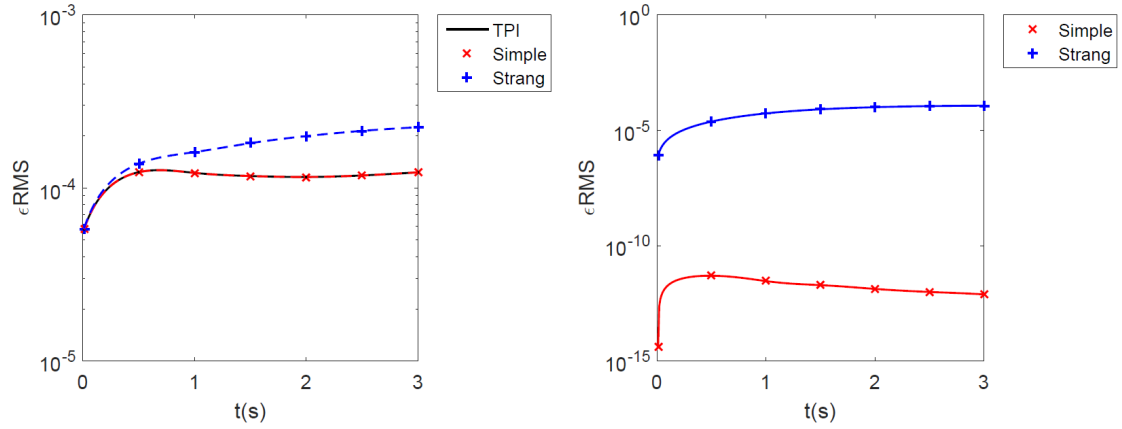

Figure 9: Comparison of the time evolution of RMS error in isotropic 2D pure diffusion calculated via different bandlimiting methods relative to the analytical solution (left) and relative to the TPI (right). In this example, high-order splitting and bulk bandlimiting are functionally identical to simple splitting.

$\boldsymbol{\mu}_0 = [2, 2]^T$, a diagonal covariance matrix $\boldsymbol{\Sigma}_0 = [4, 0; 0, 4]$, a domain of $|-4, 8|$ along both axes, and a 51×51 grid, over time increments of $dt = 0.01\ s$ for a period of 3 seconds.

Like in the SDOF case, the update equations for the mean and variance of this system are independent of the PDF and therefore error-free. This system can also be solved analytically as a function of time, allowing us to perform transient error analysis. We find that at large bandwidths, the solutions given by simple splitting, high-order splitting, bulk bandlimiting, and the complete TPI have the same error within at least three significant figures, at 1.13×10^{-4}. Only Strang splitting diverges noticeably from this outcome, at 2.16×10^{-4}.

A more nuanced comparison between BL-TPI methods comes from performing error analysis on the respective PDFs of the four bandlimiting methods—that is, the three operator splitting processes and bulk bandlimiting—relative to that of the standard TPI. Error taken by way of comparison to the standard TPI solution is hereafter referred to as "relative error", so as to distinguish from error taken through comparison to the given

analytical solution, or "absolute error".

Simple splitting, high-order splitting, and bulk bandlimiting all display excellent accuracy that is, again, identical within at least three significant figures, at 1.13×10^{-13}. Strang splitting is noticeably less accurate, at 1.17×10^{-4}, and though the difference in absolute accuracy is not disqualifying, simple splitting is both cheaper and more accurate. The difference in ability to exactly reproduce the outcome of the TPI can be explained by the fact that that 2D Gaussian Wiener process, such as in this system, can be exactly modeled as two independent 1D Wiener processes, one along each axis. Because of this, operating along each axis individually, as in operator splitting, is in this case mathematically identical to operating along all axes simultaneously, as in the case of bulk bandlimiting. This also explains why Strang splitting, which performs two half-step operations along one axis for each time step, yields a different PDF from the other two operator splitting processes, which only perform full-step operations.

3.2.2 Example 6: Anisotropic Diffusion in 2D

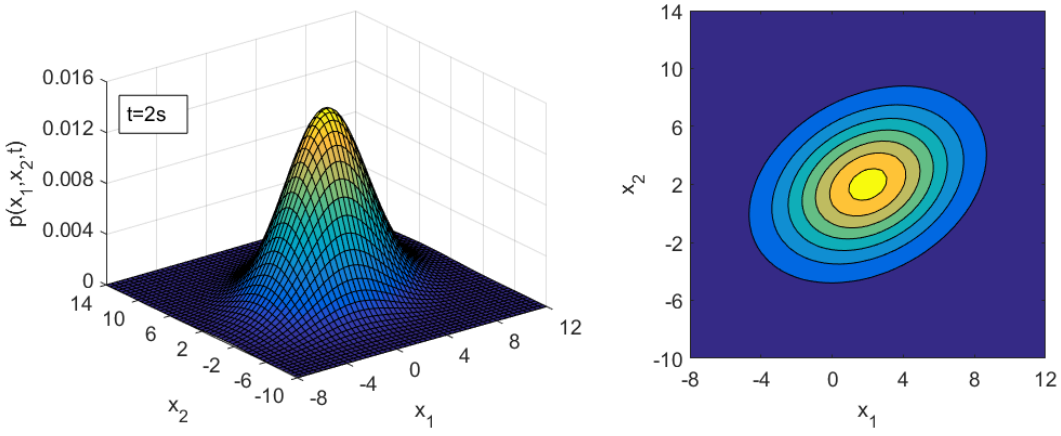

Figure 10: Surface plot (left) and contour plot (right) of a PDF of anisotropic 2D pure diffusion at t=2s, as determined with the bulk bandlimited TPI at a strength of 5.

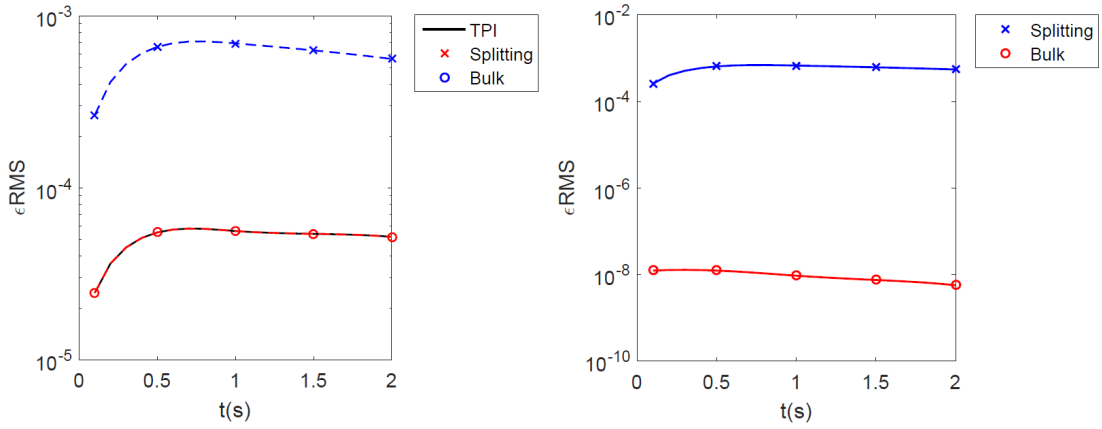

Figure 11: Comparison of the time evolution of RMS error in anisotropic 2D pure diffusion calculated via different bandlimiting methods relative to the analytical solution (left) and relative to the TPI (right). In this example, the difference in error between operator splitting methods is very small so only simple splitting is shown.

To determine the effect of an anisotropic noise term on the accuracy of the BL-TPI, and specifically on operator splitting formulations thereof, we examine a diffusion process characterized by a non-diagonal diffusion coefficient matrix—that is, a diffusion process with principal axes that do not correspond to the axes of the state space. For this example we use the same initial conditions as Example 5, applied to a system with a diffusion coefficient matrix of $\mathbf{G} = \frac{1}{2}(\mathbf{A}\mathbf{A}^T) = [3, \frac{\sqrt{3}}{2}; \frac{\sqrt{3}}{2}, 1]$ and iterated for 2 seconds at $dt = 0.1$ s. Interestingly, the errors in simple, Strang, and high-order splitting are very close to one another, at 5.62×10^{-4}, 5.58×10^{-4}, and 5.57×10^{-4}, respectively; the phenomenon of Strang splitting displaying error that diverges significantly from those of the other operator splitting processes is absent. Additionally, the errors of all operator splitting processes are an order of magnitude greater than that of bulk bandlimiting and the standard TPI, at 5.19×10^{-5}.

While the error of the bandlimited TPI appears numerically stable, as shown in Figure 10, and therefore might still be acceptable depending on one's individual needs, such an increase in error between the 1D and 2D cases indicates that when applied to systems

characterized by non-diagonal transformed diffusion coefficient matrices $\tilde{\mathbf{A}}$, the operator splitting method in its current form is not a promising option. Although there are no canonical problems with which to confirm this, we also surmise that a similar problem can surface when a process with a diagonal diffusion coefficient matrix is applied to a PDF with a distribution characterized by a non-diagonal covariance matrix. We therefore find that bandlimiting is most accurate when the principal axes of the diffusion process of the system and the distribution of the PDF are aligned.

3.2.3 Example 7: Uncoupled Nonlinear Drift in 2D

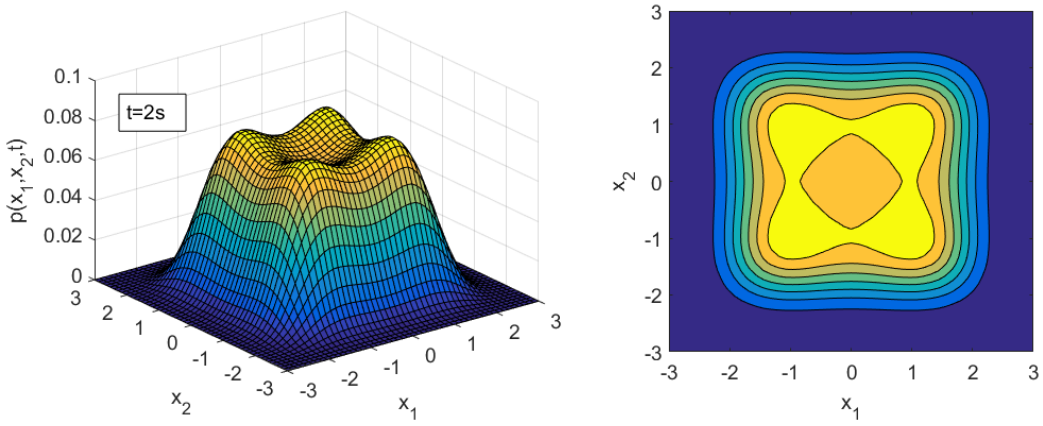

Figure 12: Surface plot (left) and contour plot (right) of a PDF of an uncoupled nonlinear drift process at t=2s, as determined with the bulk bandlimited TPI at a strength of 5.

Setting aside the diffusion coefficient matrix, it is also necessary that the bandlimited TPI is able to solve stochastic processes with complex nonlinear drift components. In this example we evaluate a stochastic system experiencing nonlinear drift.

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t)dt + \mathbf{A}d\mathbf{w}_t \tag{64}$$

This system is characterized by the drift process $\mathbf{f}(\mathbf{x}) = [x_1 - x_1^3; x_2 - x_2^3]$ and the

diffusion coefficient matrix $\mathbf{G} = \mathbf{A}\mathbf{A}^T/2 = [2, 0; 0, 2]$. The initial multivariate Gaussian distribution is given as $p(\mathbf{x}, t_0) \sim N(0, [0.25, 0; 0, 0.25])$, the initial domain is |-4,4| along each axis in absolute space, and the process is iterated over 2 seconds at a time step of $dt = 0.01\ s$. This process is known to result in a non-Gaussian distribution over time, which the TPI method is capable of effectively capturing, as seen in Figure 11. The analytical steady state solution to this process is given as follows:

$$p(x_1, x_2, t_\infty) = \frac{1}{\eta} \exp\left\{ -\begin{bmatrix} x_1^2 & x_2^2 \end{bmatrix} \mathbf{A}^{-1}\mathbf{A}^{-T} \left(\frac{1}{2} \begin{bmatrix} x_1^2 \\ x_2^2 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right) \right\} \tag{65}$$

As in Example 5, the simple splitting, high-order splitting, and bulk bandlimiting methods are all able to accurately reproduce the results of the standard TPI method with excellent accuracy, at a relative error of 2.04×10^{-9} and an absolute error of 4.01×10^{-4}. While the bulk bandlimiting formulation has been thus far found to be the most consistently reliable MDOF bandlimiting formulation, these results indicate that operator splitting TPI formulations are well-suited to solving systems with both uncoupled drift processes and diagonal diffusion coefficient matrices.

Also similar to Example 5, Strang splitting is shown to diverge significantly, but here, Strang splitting errs noticeably toward the analytical solution rather than away from it, resulting in a smaller absolute error than the TPI. While such outcomes are possible when the relative error of the Strang splitting method is of comparable magnitude to the absolute error of the standard TPI method, there is no known way to anticipate when this outcome will be favorable, especially when evaluating systems without known analytical solutions.

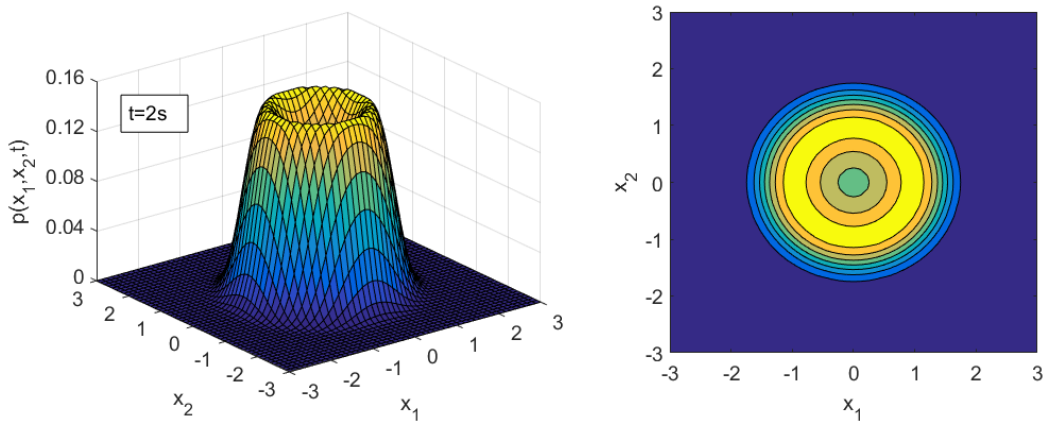

Figure 13: Surface plot (left) and contour plot (right) of a PDF of a coupled nonlinear drift process at t=2s, as determined with the bulk bandlimited TPI at a strength of 5.

3.2.4 Example 8: Coupled Nonlinear Drift in 2D

Another limitation of operator splitting methods surfaces when solving systems with coupled drift processes. In Eq. (64), we set $\mathbf{f}(\mathbf{x}) = [-\gamma x_1(x_1^2 + x_2^2 - c^2); -\gamma x_2(x_1^2 + x_2^2 - c^2)]$ where $\gamma = c = 1$, and $\mathbf{G} = \mathbf{A}\mathbf{A}^T/2 = [0.5, 0; 0, 0.5]$. Given an initial Gaussian distribution of $p(\mathbf{x}, t_0) \sim N(0, [4/49, 0; 0, 4/49])$ and an initial domain of |-1,1| along each axis in absolute space, this process is iterated in $dt = 0.01\ s$ increments up to 2 seconds. The system simulated in this example converges to the following steady state solution, where η is the normalization factor and $G_{11} = G_{22} = g = 0.5$.

$$p(x_1, x_2, t_\infty) = \frac{1}{\eta} \exp\left\{-\frac{\gamma}{4g}\left(x_1^2 + x_2^2\right)\left(x_1^2 + x_2^2 - 2c^2\right)\right\} \tag{66}$$

The error of the standard TPI method given these parameters is 6.57×10^{-4}. While bulk bandlimiting method is able to match the TPI's output almost exactly with a relative error of 9.26×10^{-10}, simple, Strang, and high-order operator splitting all err favorably toward the exact solution by $20 - 30\%$, as indicated in Table 2. It is noteworthy that here, unlike in Example 6, all three operator splitting processes have relative errors less than

the absolute error of the standard TPI. Based on these results, operator splitting appears to be an imperfect but viable alternative to the complete TPI for solving coupled nonlinear systems with diagonal noise terms. Still, further research is needed to determine how well this holds for systems with many degrees of freedom.

3.2.5 Error Analysis

Table 2: Time and error of the complete TPI vs. the bandlimited TPI with operator splitting and the bulk bandlimited TPI in 2D. A bandwidth strength of $N_\sigma = 5$ is used for all bandlimiting processes.

	Isotropic Diffusion (3 s)			Anisotropic Diffusion (2 s)		
	T_{comp} (s)	ϵ_{RMS}	$\epsilon_{RMS,TPI}$	T_{comp} (s)	ϵ_{RMS}	$\epsilon_{RMS,TPI}$
TPI	1.59×10^2	1.13×10^{-4}	0	8.92×10^1	5.19×10^{-5}	0
Simple	1.32×10^2	1.13×10^{-4}	7.76×10^{-13}	8.12×10^1	5.62×10^{-4}	5.49×10^{-4}
Strang	2.10×10^2	2.16×10^{-4}	1.17×10^{-4}	1.28×10^2	5.58×10^{-4}	5.45×10^{-4}
H.O.	2.93×10^2	1.13×10^{-4}	7.76×10^{-13}	1.72×10^2	5.57×10^{-4}	5.44×10^{-4}
Bulk	3.48×10^2	1.13×10^{-4}	7.76×10^{-13}	2.01×10^2	5.19×10^{-5}	5.66×10^{-9}
	Uncoupled Nonlinear Drift (2 s)			Coupled Nonlinear Drift (2 s)		
	T_{comp} (s)	ϵ_{RMS}	$\epsilon_{RMS,TPI}$	T_{comp} (s)	ϵ_{RMS}	$\epsilon_{RMS,TPI}$
TPI	1.59×10^2	4.01×10^{-4}	0	8.92×10^1	6.57×10^{-4}	0
Simple	1.41×10^2	4.01×10^{-4}	2.04×10^{-9}	8.12×10^1	5.32×10^{-4}	1.77×10^{-4}
Strang	2.10×10^2	3.17×10^{-4}	1.35×10^{-4}	1.28×10^2	5.03×10^{-4}	5.42×10^{-4}
H.O.	2.93×10^2	4.01×10^{-4}	2.04×10^{-9}	1.72×10^2	5.32×10^{-4}	1.77×10^{-4}
Bulk	3.48×10^2	4.01×10^{-4}	2.04×10^{-9}	2.01×10^2	6.57×10^{-4}	9.26×10^{-10}

As indicated in Table 2, despite all BL-TPI formulations requiring requiring far fewer calculations and less storage space than the standard TPI method, in practice only the simple splitting formulation is reliably faster than the standard TPI, and in all cases the difference is slight. As in the SDOF case, we believe that this can be attributed to the limitations of the prototype code used to test these methods rather than an inherent flaw of the methods themselves, so the time data in Table 2 should not be considered definitive.

It can be observed that in all cases, the accuracy of the simple splitting method is equal or nearly equal that of the high-order splitting method. We also observe that Strang splitting tends to have the greatest relative error of all operator splitting processes, except in systems with non-diagonal diffusion coefficient matrices, which we have found to be ill-suited to operator splitting methods in general. Therefore, in most cases where operator splitting is viable, simple splitting can be considered the best option, as any benefits of the Strang or high-order methods do not justify the accompanying increase in cost. Particularly, in systems with both diagonal noise terms and uncoupled drift processes, the accuracy of simple splitting matches that of bulk bandlimiting close to exactly.

In systems with non-diagonal diffusion coefficient matrices, however, while the bulk bandlimiting method remains effective if somewhat slow, the operator splitting methods of bandlimiting become significantly less reliable; in Example 6, the difference in error between bulk bandlimiting and operator splitting formulations was found to be an order of magnitude. This difference in error can be attributed to the fact that principal axes of such diffusion processes do not coincide with the axes along which operators are split. While all examples in this study use isotropic initial conditions, note that the diffusion coefficient matrix must be diagonal in the transformed space; when applied to PDFs with non-diagonal covariance matrices, isotropic diffusion can be expected to err in ways similar to the behavior seen in Example 6.

The errors of the operator splitting methods are also known to diverge from that of the standard TPI method when applied to systems with coupled drift processes, but unlike in anisotropic diffusion cases, here the error is comparable to that of the TPI and may even err favorably toward the exact solution. Still, additional research is needed to determine how this error manifests in similar problems with three or more degrees of freedom.

Overall, whereas Examples 1-4 indicate that bandlimiting is a viable alternative to

the standard TPI method for solving SDOF systems, Examples 5-8 demonstrate that the bulk bandlimiting method is the most consistently reliable MDOF bandlimiting variant, but that in cases with both uncoupled noise processes and diffusion processes whose principal axes align with the distribution of the PDF, simple operator splitting can be substituted flawlessly. Simple splitting might also be a viable alternative in some systems with coupled drift processes, though further research is necessary.

4 Symmetric Fast Gauss Transform Analysis

The symmetric fast Gauss transform (SFGT) formulation, as an infinite sum, approaches exact recreation of the transformed path integral (TPI) solution, but in practice it must be finite. For the SFGT to have practical value, it must generate a sufficiently accurate approximation of the TPI from a sufficiently small number of terms in the summation. As the SFGT implementation involves the multiplication and division of very large numbers that exceed software limitations at a high number of terms k_{max}, in the context of this study k_{max} cannot exceed 400 when applied to single degree of freedom (SDOF) systems, a limitation which becomes more stringent for multiple degree of freedom (MDOF) systems. This is found to be insufficient for many problems that are solvable with the conventional TPI formulation, including the examples given in Chapter 3. Due to the severe loss of tail information that results from these software limitations, experimental studies of the accuracy of the SFGT formulation over multiple time steps are unfeasible in this study.

In order to meaningfully study the accuracy and effectiveness of the SFGT approach, this chapter instead focuses on mathematical analysis of the formulation of this method to estimate the number of terms k_{max} needed to produce an accurate solution within the desired number of standard deviations z_{max}. To test the usefulness of this estimate, the error of the SFGT implementation over a single time step is measured within the boundaries $[-z_{max}, z_{max}]$ using multiple different parameters for calculating k_{max}. The goal of this research is to develop a recommended formula for the minimum number of terms needed to produce results that can be considered "sufficiently accurate".

4.1 Single Degree of Freedom Systems

We begin by examining the behaviors of the individual terms in the SDOF SFGT summation. Because the SFGT formulation approaches an exact recreation of the TPI solution, a sufficiently close approximation can be reasonably assumed to retain the stability of the TPI in the long run, assuming the number of terms is adjusted accordingly. Given this assumption, and given that we are iterating over only a single time step, we do not anticipate the resulting error data to differ in any noteworthy way between examples with different drift processes.

We therefore focus our analysis on the system of pure diffusion used in Example 1 of the previous chapter

$$dx_t = Adw_t$$

where $A = \sqrt{2}$, applied to a system with an initial distribution of $N \sim (1, 1)$ over 4 seconds in increments of $dt = 0.01\ s$. In symmetrical systems such as this one, odd-numbered moments in the SFGT summation are zero. Individual moments of the summation over a single iteration are plotted against the TPI solution, and the results are shown in Figure 13.

As indicated by this figure, for the SDOF case the $k = 0$ moment of the SFGT manifests as a single Gaussian curve that peaks at $z = 0$, while all subsequent moments manifest as two Gaussian curves that meet at and are symmetrical about $z = 0$. As k increases, the peaks of these curves move further from zero in either direction, while also generally decreasing in amplitude, though incidental increases may occur in non-Gaussian PDF curves.

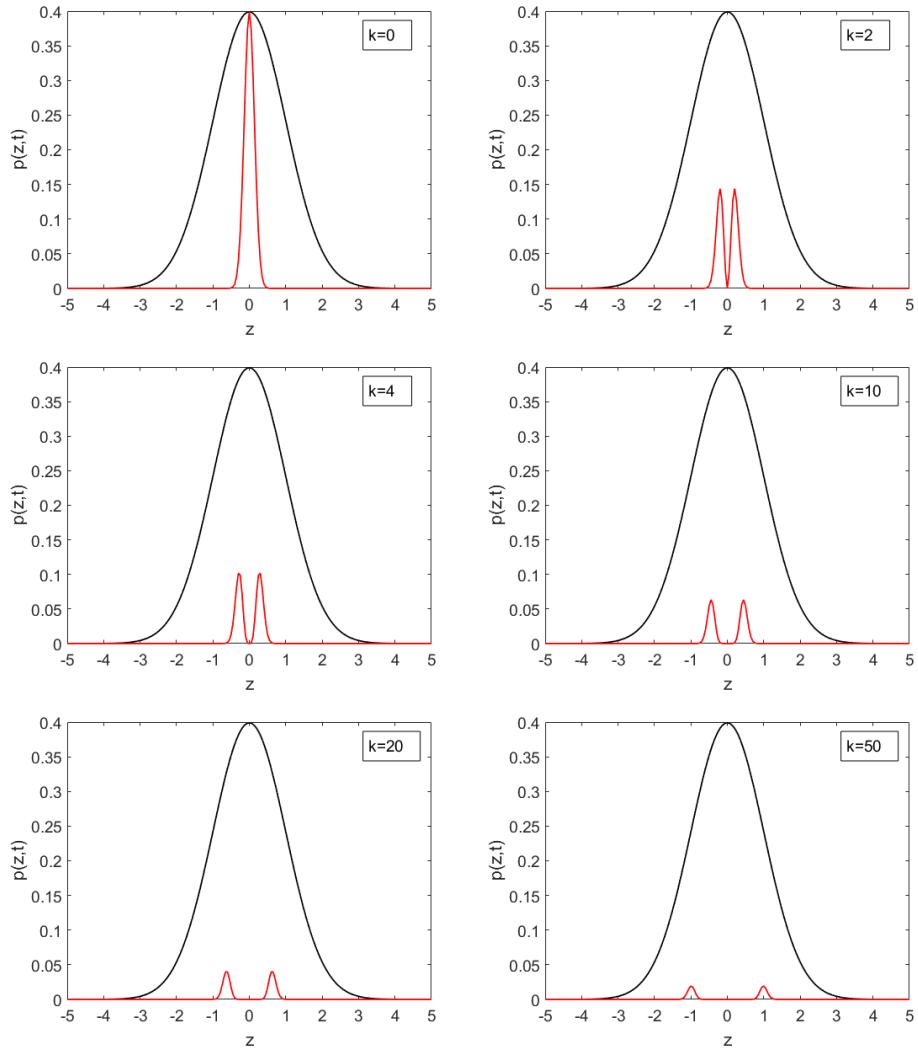

Figure 14: Comparison of the TPI solution to moments k=0 (top left), k=2 (top right), k=4 (middle left), k=10 (middle right), k=20 (bottom left), and k=50 (bottom right) of the SFGT for pure diffusion.

4.1.1 Relationship between Number of Moments and Accuracy

An equation for the locations of the peaks of the kth moment of the SFGT can be obtained from the gradient with respect to z of the SFGT formulation, that is Eq. (35). By setting the gradient equal to zero, the positions of the peaks of each moment in the

summation can be determined.

$$z_{peak} = \pm\frac{A}{\sigma}\sqrt{kdt} \tag{67}$$

Note that σ here specifically refers to the projected variance at time t, distinct from the initial variance at time t'. The square root of Eq. (12) may be substituted in order to rewrite this formula in terms of σ'. Similarly, we find that the variance of the singular Gaussian curve at $k = 0$ is equal to $\sigma_k = \frac{A}{\sigma}\sqrt{dt}$ while the variances of the two separate Gaussian curves for all subsequent terms are universally equal to $\sigma_k = \frac{A}{2\sigma}\sqrt{dt}$. These properties hold mathematically for the Gauss-Hermite quadrature formulation as well, as the formulation shares the separation of source and target terms of the SFGT formulation.

For the peak corresponding to any given moment k, the subsequent peak at $k+2$ (recalling that odd-numbered moments behave differently from even-numbered moments) tends to be both close and similar in magnitude. This suggests that statistically significant probability contributions can be made by the $k+2$th moment to the space bounded by $[-z_{peak}, z_{peak}]$ where z_{peak} corresponds to the kth moment. To address this property, we propose that the SFGT iteration may be considered accurate within the boundaries $[-(z_{peak} - N_k\sigma_k), (z_{peak} - N_k\sigma_k)]$, where N_k is a factor that can be adjusted according to the desired precision, hereafter referred to as the peak offset factor. All subsequent probability contributions to this bounded space by additional moments $k+n$ where $n \geq 1$ can be considered negligible.

Substituting $z_{peak} = z_{max} + N_k\sigma_k$, where $\pm z_{max}$ is the range within which the PDF is considered accurate in standard deviations, Eq. (67) can be rewritten in terms of k_{max} to calculate the required number of terms to yield an accurate solution within a desired semi-span z_{max}.

$$z_{max} = \frac{A}{\sigma}\sqrt{k_{max}dt} - N_k\sigma_k \tag{68}$$

$$k_{max} = \frac{(z_{max} + N_k\sigma_k)^2\sigma^2}{A^2 dt} \tag{69}$$

Consequently, when an insufficient number of moments z_{max} is used such that k_{max} is less than the desired semi-span, the PDF given over single time step adheres very closely to that of the standard TPI method at data points close to zero, but drops sharply near the z_{max} position corresponding to k_{max}. For example, in the pure diffusion problem described previously, the projected covariance after a single time step is $\sigma^2 = 1.02$. Given a desired z_{max} of 5 and an N_k of 1 yields a k_{max} of nearly 1350 terms, far exceeding the previously stated limit of 400. Conversely, at $k_{max} = 400$, the corresponding z_{max} given a peak offset factor of $N_k = 1$ is 2.66, beyond which the solution ceases to be accurate as seen in Figure 14.

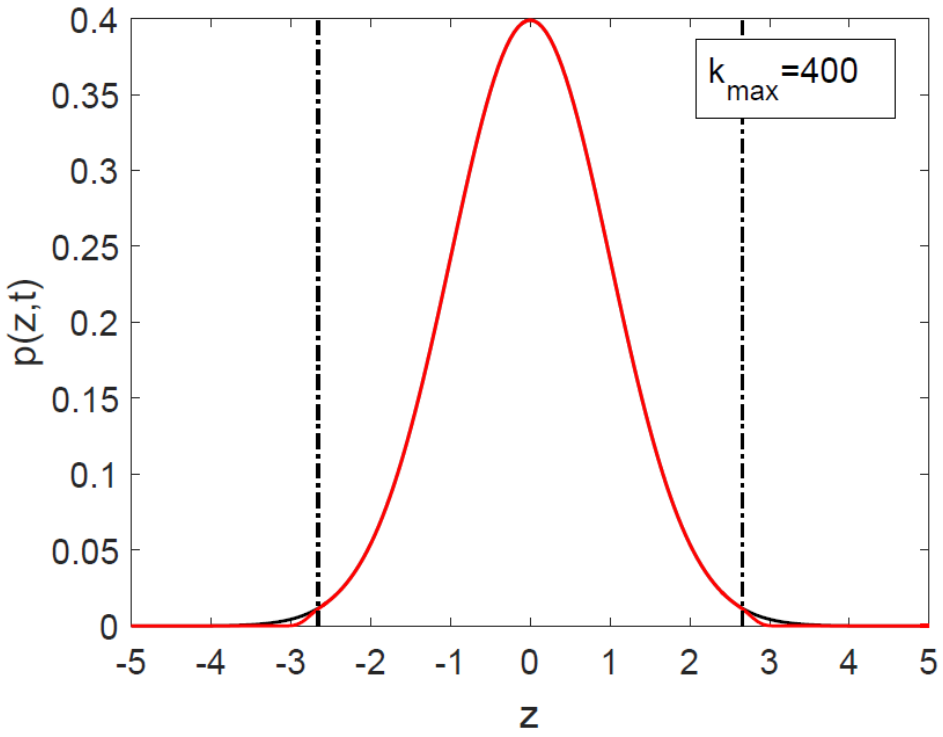

Figure 15: Comparison of the TPI solution to the SFGT at $k_{max} = 400$ for pure diffusion, shown to diverge at approximately ± 2.66.

For systems with constant white noise strengths, this presents a significant dilemma. For a system with a small white noise strength A relative to the variance σ of the PDF to which it is applied, a large number of moments k is needed to avoid significant loss of tail information. One way to mitigate this problem is to iterate over a large time step size, but this may compromise the accuracy of the results when applied to systems with sensitive, high-gradient drift processes.

Conversely, a system with a large white noise strength may initially warrant a relatively small number of moments, but systems with large diffusion processes tend to expand. This means that the variance of the PDF become progressively larger relative to the white noise strength, gradually requiring more moments in order to accurately reproduce the TPI in later iterations. In concentrating systems such as the OU process this is not the case, but the white noise strengths of such systems tend to be very small relative to their initial variances, resulting in a large initial hurdle. Due to shared target components of their formulations, these limitations extend to the Gauss-Hermite quadrature formulation of the SFGT as well.

We note that the behaviors described are not emergent properties of the TPI formulation of the SFGT method, but are inherited from the fixed grid SFGT. In the SFGT formulation of the path integral solution, rather than an equal number of terms covering a different amount of transformed space in later time steps, the same amount of absolute space is covered given the same number of terms, but distribution of the PDF itself changes, and with it the amount of coverage needed. Therefore, the limitations of the SFGT described thus far cannot be sidestepped using the fixed grid PI solution.

Importantly, PDFs within the boundaries $[-z_{max}, z_{max}]$ should not be considered reliable over multiple time steps, as the inaccurate portions of the PDF outside of these boundaries at one discrete point in time can contribute inaccurate information to the space

inside the boundaries when generating a PDF for the next point in time.

4.1.2 Error Analysis and Computational Requirements

While the symmetric fast Gauss transform cannot be accurately iterated over many time steps within the limitations of this study, both the accuracy of the SFGT method within a single time step and the evolution of its required number of terms can be evaluated. We first return to the conditions for 1D pure diffusion used in Example 1 and examine the accuracy of this process over a single time step within boundaries obtained from different N_k values using RMS error analysis, Eq. (57). In this case, p is the probability of the PDF yielded by the SFGT implementation and p_0 is that of the standard TPI, which p approaches exactly as k_{max} approaches infinity. For the purpose of this comparison, neither PDF is normalized, as the total probability of the PDF iterated via the SFGT in each example is significantly less than 1 and the normalization process would distort the results.

Table 3: Accuracy of pure diffusion after a single time step $dt = 0.01\ s$ using $k_{max} = 200$ and $k_{max} = 400$, within the boundaries determined using peak offset factors N_k from 0 to 3.

	$k_{max} = 200$		$k_{max} = 400$	
	z_{max}	ϵ_{RMS}	z_{max}	ϵ_{RMS} (s)
$N_k = 0$	± 1.98	3.15×10^{-3}	± 2.80	4.79×10^{-4}
$N_k = 1$	± 1.84	2.39×10^{-4}	± 2.66	5.49×10^{-5}
$N_k = 2$	± 1.70	1.30×10^{-5}	± 2.52	8.38×10^{-7}
$N_k = 3$	± 1.56	1.48×10^{-7}	± 2.38	1.22×10^{-7}

It is found that at both $k_{max} = 400$ and $k_{max} = 200$, the accuracy is greater within the smaller boundaries created by setting larger peak offset factors. At a smaller number of terms, this increase is more abrupt, which is likely due to the relatively large amplitude

of the peaks and rate of increase in the peak displacement z_{peak}. As seen in Table 3, the same offset factor yields PDF with a smaller error at a greater distance z_{max} and number of terms k_{max}. This can be attributed to the size of the probability measurements generally being smaller at greater distances from the mean, as well as the lower amplitudes of each additional moment further in the series. Based on this trend, we surmise that a sufficiently large z_{max} may be considered equal to z_{peak}.

To confirm the viability of this approximation, we modify Example 1 by increasing the time step $dt = 0.01\ s$ to $0.1\ s$. In this version of the problem, $\sigma^2 = 1.2$, but all else remains the same. Given these conditions, covering a total range of $z_{max} = 5$ requires a k_{max} of 233 for peak offset factor of $N_k = 3$, which yields an error of $\epsilon_{RMS} = 1.17 \times 10^{-7}$. Similarly, at a peak offset factor of $N_k = 0$, with a k_{max} of 150, the PDF error is nearly identical, at $\epsilon_{RMS} = 1.20 \times 10^{-7}$. This indicates that at large ranges, $[-z_{peak}, z_{peak}]$ may be considered an accurate set of boundaries for the PDF with no need for a peak offset factor.

While running full simulations over large numbers of time steps is not feasible within this study, it is possible to determine the computational requirements for such simulations based on the conventional TPI implementation. Setting a z_{max} of 5 and assuming that $N_k = 0$ is acceptable at this peak distance, we determine the required k_{max} for the initial and final time steps of Examples 1, 3, and 4 from Chapter 3.

Table 4: Initial and final required number of terms k_{max} (rounded up) for the SFGT TPI to be accurate within a range of $z_{max} = 5$ with a peak offset factor of $N_k = 0$.

	Pure Diffusion (4 s)		OU Process (10 s)		Nonlinear Drift (16 s)	
	σ^2	k_{max}	σ^2	k_{max}	σ^2	k_{max}
Initial	1.020	1275	0.2476	61900	2.108	586
Final	9.020	11275	0.03891	9728	13.36	3712

As predicted, systems that initially require relatively few moments tend to expand, while the system with the largest initial requirement is the OU process, which is contracting. For a discrete PDF comprised of N terms, the propagator matrix for the standard TPI for each problem is a $N \times N$ matrix, whereas the SFGT implementation requires the calculation of $N \times (k_{max} + 1)$ terms, and while the processes by which the terms in each method are calculated and developed differ, as a general rule the average k_{max} should not be larger than N, or else the SFGT offers little benefit. For a PDF with $N = 201$ terms, the minimum k_{max} far exceeds N in all examples given. Additionally, because high k values rely on the multiplication and division of very large numbers, software limitations may limit the possible k_{max} that can be used, as is the case in this study. Consequently, the SFGT in its current form appears to be ill-suited to solving SDOF systems that require high precision and must be iterated over small time steps.

4.1.3 SDOF Gauss-Hermite Quadrature

Because Eqs. (67), (68), and (69) can be identically derived from the formulation for the Gauss-Hermite quadrature, the limitations of the SFGT as outlined in the previous section persist in the quadrature formulation. However, the Gauss-Hermite quadrature implementation still has the potential to improve on the SFGT substantially. Because this method approximates N data points as n weighted abscissas, the number of terms that must be calculated in the Quadrature SFGT becomes approximately $n \times (k_{max} + 1)$, or n/N times that of the standard SFGT. In order to cover a range of $z_{max} = 5$ in the transformed space, n must be at least 18; to ensure that there is an abscissa at the exact center $z = 0$ of the space we instead use the odd number $n = 19$, as opposed to $N = 201$.

Returning to the pure diffusion system of Example 1, we again apply a white noise strength of $A = \sqrt{2}$, a time step size of $dt = 0.01$, and an initial distribution of $p(x, t_0) =$

$N(1,1)$. At $k_{max} = 200$, assigning a peak offset factor of $N_k = 3$, z_{max} is found to be ± 1.56. The middle 7 roots of the $19th$ Hermite polynomial fall into the range of [-1.560,1.560], which appears consistent with the results of iterating the Gauss-Hermite quadrature over a single time step, as indicated in Figure 15.

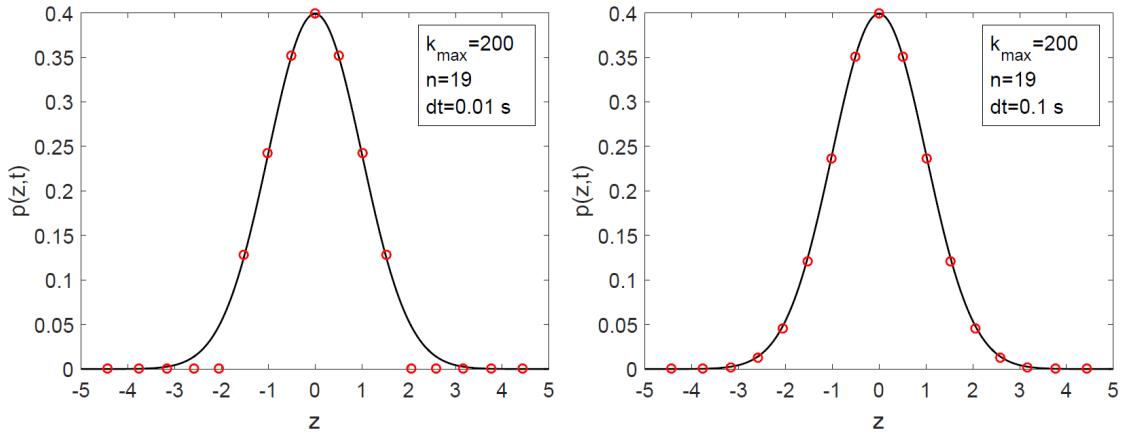

Figure 16: Comparison of the TPI solution to the Gauss-Hermite quadrature at $n = 19$ and $k_{max} = 200$ iterated over a single time step of $dt = 0.01s$ (left) and $dt = 0.1s$ (right).

The error of this method, however, is found to be significantly greater than that of the standard SFGT, at $\epsilon_{RMS} = 2.149 \times 10^{-3}$. Additional experimentation reveals that this error can be somewhat mitigated by increasing the number of abscissas. At $n = 29$, the error reduces to $\epsilon_{RMS} = 8.078 \times 10^{-4}$, though this is still relatively large. Furthermore, increasing the time step size to $dt = 0.1s$ such that $k_{max} > 5$ does not greatly improve the overall accuracy of the quadrature formulation, yielding an error of $\epsilon_{RMS} = 2.205 \times 10^{-3}$ at $n = 19$ and $\epsilon_{RMS} = 2.109 \times 10^{-3}$ at $n = 29$. From this error analysis it can be confirmed that the Gauss-Hermite quadrature does not converge to the exact TPI solution, and based on these results this implementation of the SFGT does not seem promising.

4.2 Multiple Degree of Freedom Systems

Similar to the SDOF case, we begin our analysis of the MDOF formulation by examining the behavior of individual terms in the SFGT sum, for which we return to Example 5 from Chapter 3. Due to additional technical limitations that occur in the MDOF formulation, we increase the time step to $dt = 0.2\ s$ to make the difference more pronounced. In the 2D case, we find that the $0th$ moment displays Gaussian distribution about the mean and origin, while subsequent moments take on the shapes of "rings" that move outward at higher k terms, with the peak values of each ring following an elliptical pattern about the origin.

4.2.1 Relationship between Number of Moments and Accuracy

An MDOF equivalent to Eq. (67) can be developed in order to obtain a vector $\mathbf{z}_{peak}$ containing the maximum from zero of the "peak ellipse" along each axis. Here $\mathbf{k}$ is a $N_s \times 1$ column vector where each term is equal to k.

$$\mathbf{z}_{peak} = \pm \mathbf{R}^{-1}\mathbf{A}\sqrt{\mathbf{k}dt} \tag{70}$$

The individual components of the vector $\mathbf{z}_{peak}$ correspond to the principal axes of the peak ellipse. Similar to the SDOF case, the update equation for $\mathbf{R}$, Eq. (14), can be incorporated into the above relation to rewrite it in terms of $\mathbf{R}'$. The moment at $k = 0$ is a simple Gaussian distribution of variance $\mathbf{R}_k = \mathbf{R}^{-1}\mathbf{A}\sqrt{dt}$ centered at the origin. All subsequent moments concentrate about an elliptical set of peaks rather than a singular point, forming a series of probability rings. The "thickness" of each of these rings—that is, the distance toward and away from the origin within which probability contributions of these rings can be considered significant—corresponds to the ellipse $\mathbf{R}_k = \frac{1}{2}\mathbf{R}^{-1}\mathbf{A}\sqrt{dt}$.

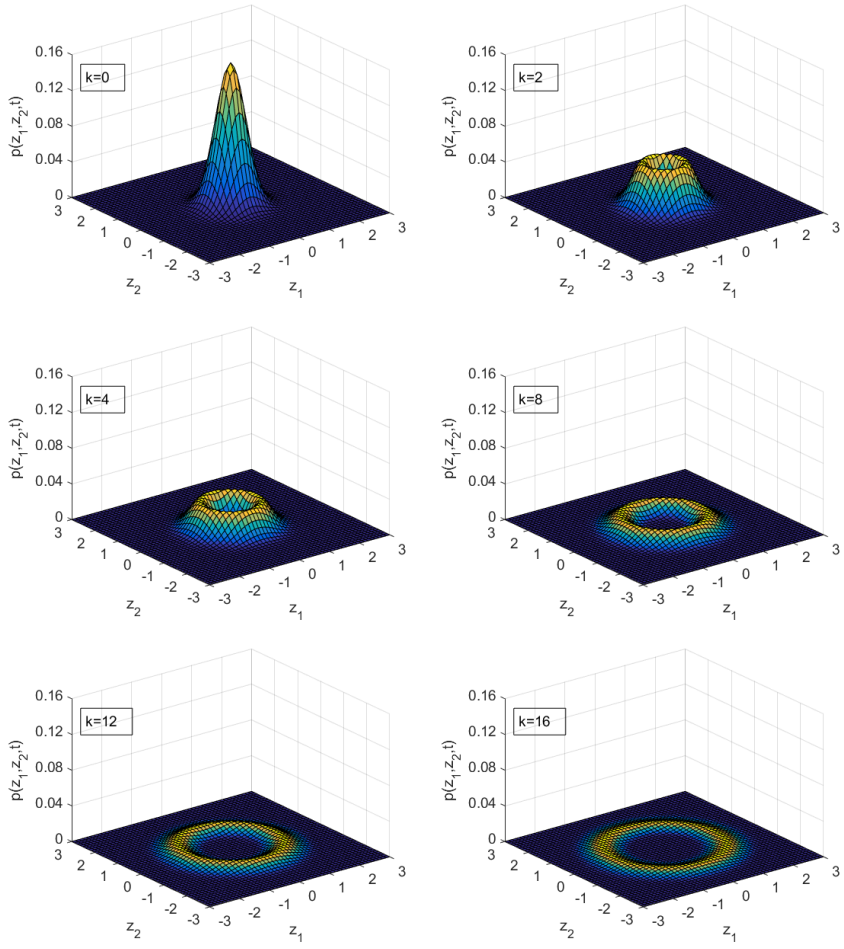

Figure 17: Moments k=0 (top left), k=2 (top center), k=4 (top right), k=8 (bottom left), k=12 (bottom center), and k=16 (bottom right) of the SFGT for 2D pure diffusion.

Echoing our SDOF analysis, we substitute $\mathbf{z}_{peak} = \mathbf{z}_{max} + \mathbf{R}_k\mathbf{N}_k$, $\mathbf{N}_k$ being a column vector whose terms are all equal to the desired precision N_k.

$$\mathbf{z}_{max} = \mathbf{R}^{-1}\mathbf{A}\sqrt{\mathbf{k}_{max}dt} - \mathbf{R}_k\mathbf{N}_k \tag{71}$$

$$\sqrt{\mathbf{k}_{max}} = \frac{1}{\sqrt{dt}}\mathbf{A}^{-1}\left[\mathbf{R}(\mathbf{z}_{max} + \mathbf{R}_k\mathbf{N}_k)\right] \tag{72}$$

Since a summation cannot be multiple sizes at once, we defer to the maximum component of the resulting vector $\mathbf{k}_{max}$.

4.2.2 Error Analysis and Computational Requirements

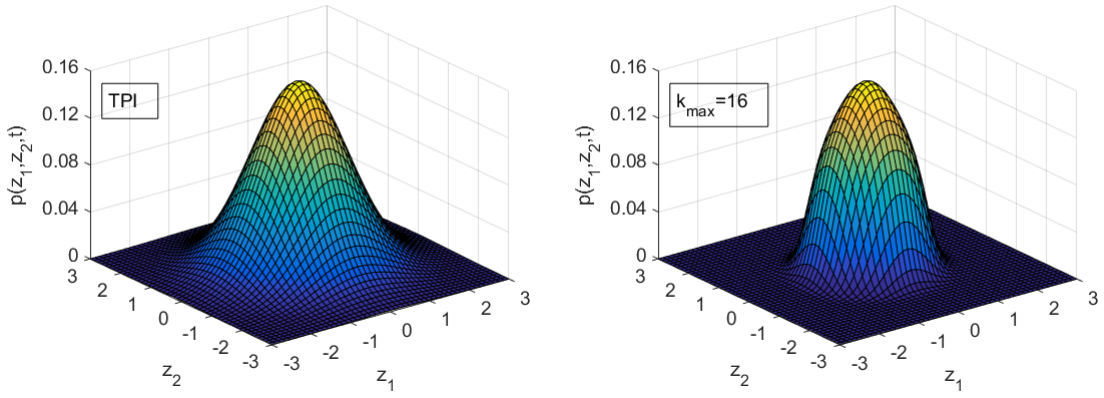

Figure 18: Comparison of the PDF for pure diffusion of a single time step $dt = 0.1\ s$ taken with the conventional TPI (left) and the SFGT at $k_{max} = 16$ (right).

Due to the dependence of the MDOF SFGT on repeated Kronecker multiplication, the computational requirements for higher k terms become very large. By $k = 20$, a 2×1 vector to the Kronecker power of 20 becomes a 1048576×1 vector, and two such vectors are needed to calculate each data point. We therefore limit our results to $k_{max} = 16$, which we apply to the pure diffusion case from Example 5 with an increased time step of $dt = 0.1$. The characteristic behavior of the 1D SFGT is found to persist in the 2D case, in which the data points nearest the mean follow the TPI outcome quite closely but the PDF's accuracy drops off abruptly past a certain distance. At $k = 16$, $\mathbf{z}_{peak} = [1.26; 1.26]$ and $\mathbf{R}_k = 0.158\mathbf{I}$, which gives us the constraints needed to develop an accurate range for this simulation. The error is taken using offset factors from 0 to 3 and recorded in Table 5. As in the 1D case, the error is found to steadily improve toward the center of the evaluated space.

Though they are not feasible within this experiment, we also examine the computational requirements to produce accurate results for 2D problems, using Examples 5, 7, and 8 from Chapter 3. Like in the 1D case, there is a tendency for the required number of

Table 5: Accuracy of 2D pure diffusion after a single time step $dt = 0.1\ s$ using $k_{max} = 16$ within the boundaries determined using peak offset factors N_k from 0 to 3.

	$k_{max} = 16$	
	z_{max}	ϵ_{RMS}
$N_k = 0$	± 1.26	4.50×10^{-4}
$N_k = 1$	± 1.10	9.38×10^{-5}
$N_k = 2$	± 0.94	1.35×10^{-5}
$N_k = 3$	± 0.79	1.01×10^{-6}

terms to expand over time.

Table 6: Initial and final required number of terms k_{max} (rounded up) for the 2D SFGT TPI to be accurate within a range of $z_{max} = 3$ with a peak offset factor of $N_k = 0$.

	Pure Diffusion (3 s)		Uncoupled Drift (2 s)		Coupled Drift (2 s)	
	$\mathbf{\Sigma}$	k_{target}	$\mathbf{\Sigma}$	k_{target}	$\mathbf{\Sigma}$	k_{target}
Initial	$4.040\mathbf{I}$	909	$0.2913\mathbf{I}$	66	$0.09274\mathbf{I}$	84
Final	$16.04\mathbf{I}$	3609	$1.287\mathbf{I}$	290	$0.6444\mathbf{I}$	580

While the prohibitive limitations of this study seems to have damning implications for the SFGT, this behavior can be largely attributed to the use of the Kronecker product. If a workaround is developed in future research that reduces the sizes of the vectors used in this process, this method could become far more viable. As formulated and tested in this paper, however, the SFGT has too many problems to recommend, especially when applied to MDOF systems.

5 Summary and Comparison

Through our experimentation and analysis, we find that of the two methods proposed, bandlimiting is the more universally applicable approach, as evidenced by our success in recreating a wide variety of simulations of stochastic differential equations, in both SDOF and MDOF cases. For every problem we considered in this study, at least one bandlimiting formulation is able to match the results of the standard TPI with a high degree of accuracy. While some bandlimited simulations in this experiment are found to run slower than the standard TPI, this can be attributed to the use of prototype code. Theoretically, a well-coded bandlimiting program should be faster than the standard TPI implementation in most if not all cases. Additionally, the propagator matrix developed for a bandlimited TPI requires greatly reduced storage space compared to that of the standard TPI, increasing the upper size limit of simulations that can be run given the same software capacity.

When applied to MDOF systems, bulk bandlimiting is found to be the most consistently accurate form of bandlimiting, but in systems with diagonal noise terms and uncoupled drift processes, operator splitting is equally accurate. In the case of a coupled drift process, the accuracy of operator splitting is still found to be acceptable, but further research is needed to determine how this will hold when applied to systems with three or more degrees of freedom. In a system with a non-diagonal diffusion term, the accuracy of the operator splitting methods is significantly worse than that of the standard TPI or bulk bandlimiting and therefore its usage is not recommended.

While the symmetric fast Gauss transform, like bandlimiting, is theoretically capable of matching the accuracy of the standard TPI, unlike bandlimiting it is not only possible but common for the computational demands of the SFGT method to exceed those of the standard TPI before the threshold of sufficient accuracy is reached. This formulation is

found to demand a large number of moments to yield an accurate result for all examples, and in many cases the number of moments required tends to increase over time. Systems for which this requirement instead decreases, such as OU processes, tend to have prohibitive initial requirements and converge at quantities that are still too large to be reasonable.

This also poses a problem because the SFGT implementation depends on multiplying and dividing very large numbers—specifically exponents and factorials—which may for high k values overtake the limits of the computers being used. These problems can be mitigated by using a larger time step, which greatly reduces the size of the required sum, but doing so may be detrimental to the accuracy of the results when evaluating sensitive systems. Another problem arises when the SFGT is applied to MDOF problems, due to its dependence on the Kronecker products of vectors, which further increases the storage and processing demands of the algorithm.

Despite these flaws, this study should not be construed as discrediting the SFGT implementation entirely. If an alternative to the Kronecker product is developed in future work, the demands of the MDOF SFGT could be greatly reduced, at which point its potential should be reevaluated. It is also worth noting that the SFGT is best-suited to solving systems with large white noise strengths that can be implemented over large time steps, which are also the systems for which bandlimiting offers the least benefit due to the relationship between white noise strength, time step size, and required bandwidth.

6 Conclusion

The goal of this thesis is to develop an alternate formulation of the Transformed Path Integral that, for an acceptably small trade-off in accuracy, retains its desirable features while improving on its computational efficiency. To that end, we propose and examine two general approaches to modifying the TPI. The first method, bandlimiting, takes advantage of the relative smallness of most terms in the propagator matrix in order to set boundaries outside of which all values are considered negligible. The second method, based on the Symmetric Fast Gauss Transform, manipulates the Gaussian kernel of the propagator matrix formulation to separate the source and target terms in the convolution operation. We find that both proposed methods are shown to approach the solution of the standard TPI, but bandlimiting is found to be the more dependable of the two approaches.

At a bandwidth strength of 5, which we propose as the generally recommended bandwidth, the bandlimited TPI is found to reproduce the TPI effectively in all examined cases. In MDOF systems, we find that the computationally inexpensive simple operator splitting method approaches the standard TPI solution for systems with uncoupled drift processes and diagonal white noise covariance matrices, and it is also reasonably accurate for systems with coupled drift processes, though further research is needed for the latter. In contrast, while bulk bandlimiting is the most expensive form of bandlimiting, it always approaches the standard TPI solution and can be recommended when the conditions for accurate operator splitting are not met

Our studies indicate that the SFGT method does not fare as well, because while it approaches an exact solution of the TPI, the summation size required by its Taylor series component causes it to far overtake the computational demands of the standard TPI. Furthermore, the large numbers involved in computing higher Taylor series terms exceed the limits of the software available for this study, and in cases of expanding systems, the

required Taylor series size increases over time, making the SFGT not only expensive but often unfeasible. This is further exacerbated in systems of two or more degrees of freedom, for which the SFGT depends on the repeated Kronecker multiplication of vectors, straining the limits of storage capacity at relatively small Taylor series summation sizes.

Because of this flaw, the SFGT implementation cannot be generally recommended, but in stable systems that can be iterated in large time steps, which require fewer Taylor series moments, it may prove beneficial. Furthermore, if an alternative to the use of Kronecker products can be developed in future research, the viability of the SFGT when applied to MDOF systems can be expected to greatly improve. A further modification to the SFGT, the Gauss-Hermite quadrature, is also considered, but its accuracy is found to be too poor to recommend its usage.

The TPI itself is a very new tool for numerical analysis, and since both bandlimiting and SFGT approaches are derived from its formulation, they have inherited its challenges. As the TPI method is further researched and developed to address issues relating to the curse of dimensionality and singular diffusion matrices, these alternative formulations will need to be updated and given new consideration in light of future developments.

Based on these findings, we recommend the use of bandlimiting to improve the efficiency of the TPI. Bandlimiting implementations greatly reduce the storage requirements involved in the TPI solution, and we believe that with optimal coding the speed can be significantly improved as well. Our findings also indicate that further research and improvements are necessary in the formulation of a TPI method based on the symmetric fast Gauss transform. We find that the computational requirements of the current method often exceed those of the standard TPI and in many cases tend to further increase over time.

7 Future Work

At the time of writing this thesis, the Transformed Path Integral is still a very new method that remains untested in many conditions. Notably, the stochastic systems for which the TPI has been proven effective have all had state-invariant, or uniform, diffusion processes, and only 1D and 2D state spaces have been tested, which this study reflects. Further research and testing are needed to better understand the stability of the TPI in higher dimensions and how well it holds up against the curse of dimensionality, as well as if and in what ways its behavior is affected by nonuniform diffusion. Pending such studies, the relative effectiveness of the SFGT and bandlimiting methods should be re-evaluated in that context.

With regard to bandlimiting, we find that for stochastic systems with white noise processes defined by non-diagonal matrices, operator splitting methods tend to be ineffective. One possible alteration to the current operator splitting bandlimiting formulation worth considering would be to develop a splitting process that operates along the principal axes of the diffusion process, rather than along the axes of the transformed space. If operator splitting can be made accurate for all systems regardless of noise or drift, them the relatively expensive bulk bandlimiting process could be made obsolete.

Additionally, we emphasize that the codes used to run the simulations in this study are essentially prototypes, and the recorded computation times reflect this fact. Through code optimization, we believe that the performance of both bandlimiting and SFGT methods can be further improved. On a related note, as floating point size limitations severely restrict the analysis of the symmetric fast Gauss transform-based method that is possible within this study, further studies that address this limitation are warranted. Developing a method to reduce the Kronecker power in the MDOF SFGT to something more manageable is also worth researching.

Software limitations aside, one concept worth considering for the improvement of the SFGT formulation is that of an evolving time step relative to the current transformed space. If the time step size is made to increase over time in proportion to the PDF and its resolution, this could potentially negate the increase over time in the required number of moments for expanding systems, while also decreasing the number of time steps needed to reach a certain time. If it is found that this proposed modification is not significantly detrimental to the accuracy of the results, such a method could also be applied to the standard and bandlimited TPI.

Finally, because the SFGT is the sum of a series of PDFs that each individually exhibit Gaussian behavior, it is possible to apply bandlimiting to the components of the Taylor series that comprises the SFGT. Notably, the formulas for locating the peaks of these Gaussian distributions have already been identified in Chapter 4. It may also be possible to apply the principles of bandlimiting to an SFGT-based TPI method; while such a method would not reduce the required summation size, it would reduce the computational cost of each element of the sum. We believe such a synthesis would be worth examining as SFGT-based methods and their potential uses are further researched.

References

[1] Abdelrahman, M.A.E., & Sohaly, M.A. 2018. "The development of the deterministic nonlinear PDEs in particle physicsto stochastic case." *Results in Physics,* **9**, 344-350.

[2] Brouwers, J.J.H. 2006. "Stochastic Processes in Mechanical Engineering." *Eindhoven University of Technology.*

[3] Di Paola, M., & Santoro, R. 2009. "Path integral solution handled by Fast Gauss Transform." *Probabilistic Engineering Mechanics,* **24**, 300-311.

[4] Drozdov, A.N., & Morillo, M. 1996. "Solution of nonlinear Fokker-Planck equations." *Physical Review E,* **54** (1), 931-937.

[5] Ermak, D.L., & Buckholz, H. 1980. "Numerical Integration of the Langevin Equation: Monte Carlo Simulation." *Journal of Computational Physics,* **35**, 169-182.

[6] Feynman, R.P. 1948. "Space-time approach to non-relativistic quantum mechanics." *Reviews of Modern Physics,* **20** (2), 367-387.

[7] Geiser, J., Tanoğlu, G., & Gücüyenen, N. 2011. "Higher order operator splitting methods via Zassenhaus product formula: Theory and applications." *Computers and Mathematics with Applications,* **62**, 1994-2015.

[8] Jäckel, P. 2005. "A note on multivariate Gauss-Hermite quadrature." *Selected documents by Peter Jäckel,* www.jaeckel.org.

[9] Jazwinski, A.H. 1970. *Stochastic Processes and Filtering Theory.* Academic Press, New York.

[10] Johnson, E.A., Wojtkiewicz, S.F., Bergman, L.A., & Spencer Jr., B.F. 1997. "Observations with regard to massively parallel computation for Monte Carlo simulation of stochastic dynamical systems." *International Journal of Non-Linear Mechanics,* **32** (4), 721-734.

[11] Kumar, P., & Narayanan, S. 2006. "Solution of Fokker-Planck equation by finite element and finite difference methods for nonlinear systems." *Sadhana,* **31** (4), 445-461.

[12] Landa, B., & Shkolnisky, Y. 2017. "Approximation scheme for essentially bandlimited and space-concentrated functions on a disk." *Applied and Computational Harmonic Analysis,* **43**, 381-403.

[13] Liu, Q., & Pierce, D.A. 1994. "A note on Gauss-Hermite Quadrature." textitBiometrika, **81** (3), 624-629, University of Oxford Press.

[14] Lyubarskii, Y., & Ortega-Cerdà, J. 2014. "Bandlimited Lipschitz functions." *Applied and Computational Harmonic Analysis,* **37**, 307-324.

[15] Proppe, C., Pradlwarter, H.J., & Schuëller, G.I. 2003. "Equivalent linearization and Monte Carlo simulation in stochastic dynamics." *Probabilistic Engineering Mechanics,* **18**, 1-15.

[16] Risken, H. 1989. "The Fokker-Planck Equation: Methods of Solution and Applications." *2nd edition, Springer Series in Synergetics,* Springer-Verlag, Berlin, Heidelberg, New York.

[17] Spencer Jr., B.F., & Bergman, L.A. 1993. "On the numerical solution of the Fokker-Planck equation for nonlinear stochastic systems." textitNonlinear Dynamics, **4** (4), 357-372.

[18] Strang, G. 2006. *Linear Algebra and its Applications (Fourth Edition).* Thomson Brooks/Cole, New York, 285.

[19] Strang, G. 1968. "On the construction and comparison of difference schemes." *SIAM Journal on Numerical Analysis,* **5** (3), 506-517.

[20] Subramaniam, G., & Vedula, P. 2017. "A transformed path integral approach for solution of the Fokker–Planck equation." *Journal of Computational Physics,* **346**, 49-70.

[21] Wiener, N. 1930. "Generalized Harmonic Analysis." *Acta Mathematica,* **55**, 117-258.