

A realizable filtered intrusive polynomial moment method

Graham Aلدredge, Martin Frank, Jonas Kusch,
Ryan McClarren

CRC Preprint 2021/23, May 2021

KARLSRUHE INSTITUTE OF TECHNOLOGY

CRC 1173



Participating universities



Universität Stuttgart

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN



Funded by

DFG

A Realizable Filtered Intrusive Polynomial Moment Method

Graham Alldredge^a, Martin Frank^b, Jonas Kusch^c, Ryan McClarren^d

^a*Berlin, gwak@posteo.net*

^b*Karlsruhe Institute of Technology, Karlsruhe, martin.frank@kit.edu*

^c*Karlsruhe Institute of Technology, Karlsruhe, jonas.kusch@kit.edu*

^d*University of Notre Dame, Notre Dame, rmclarr@nd.edu*

Abstract

Intrusive uncertainty quantification methods for hyperbolic problems exhibit spurious oscillations at shocks, which leads to a significant reduction of the overall approximation quality. Furthermore, a challenging task is to preserve hyperbolicity of the gPC moment system. An intrusive method which guarantees hyperbolicity is the intrusive polynomial moment (IPM) method, which performs the gPC expansion on the entropy variables. The method, while still being subject to oscillations, requires solving a convex optimization problem in every spatial cell and every time step.

The aim of this work is to mitigate oscillations in the IPM solution by applying filters. Filters reduce oscillations by damping high order gPC coefficients. Naive filtering, however, may lead to unrealizable moments, which means that the IPM optimization problem does not have a solution and the method breaks down. In this paper, we propose and analyze two separate strategies to guarantee the existence of a solution to the IPM problem. First, we propose a filter which maintains realizability by being constructed from an underlying Fokker-Planck equation. Second, we regularize the IPM optimization problem to be able to cope with non-realizable gPC coefficients. Consequently, standard filters can be applied to the regularized IPM method. We demonstrate numerical results for the two strategies by investigating the Euler equations with uncertain shock structures in one- and two-dimensional spatial settings. We are able to show a significant reduction of spurious oscillations by the proposed filters.

Keywords: uncertainty quantification, conservation laws, hyperbolic, intrusive, entropy, filtering, realizability

1. Introduction

Conservation laws describe the behavior of a large number of physical applications; these laws can be found in fields as diverse as fluid dynamics, radiation transport or plasma physics. Given the desire to use computational simulations to predict the behavior of such systems, quantifying the effect of stochastic uncertainties, such as measurement errors or modeling assumptions, is an important task. With this goal in mind we consider a general conservation law

$$\partial_t \mathbf{u}(t, \mathbf{x}, \boldsymbol{\xi}) + \nabla \cdot \mathbf{f}(\mathbf{u}(t, \mathbf{x}, \boldsymbol{\xi})) = \mathbf{0} \quad \text{in } D, \quad (1a)$$

$$\mathbf{u}(t = 0, \mathbf{x}, \boldsymbol{\xi}) = \mathbf{u}_{IC}(\mathbf{x}, \boldsymbol{\xi}) \quad (1b)$$

for the state variable $\mathbf{u} \in \mathbb{R}^m$ at time $t \in \mathbb{R}^+$ and spatial position $\mathbf{x} \in D \subseteq \mathbb{R}^d$ and also dependent on the random variable $\boldsymbol{\xi} \in \Theta \subseteq \mathbb{R}^p$ with probability density function $f_{\Xi}(\boldsymbol{\xi})$. Supplemented with adequate boundary conditions, these equations can be solved using various techniques [1]. In this work, we focus on

intrusive methods for the time evolution of stochastic quantities such as expected value or variance. These quantities are more generally the moments

$$\int_{\Theta} \varphi_i(\boldsymbol{\xi}) \mathbf{u}(t, \mathbf{x}, \boldsymbol{\xi}) f_{\Xi}(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad (2)$$

of the random vector \mathbf{u} with respect to the basis functions φ_i . In the sequel we use bracket notation for integration over Θ against f_{Ξ} :

$$\langle \cdot \rangle := \int_{\Theta} \cdot f_{\Xi}(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad (3)$$

and assume orthonormal basis functions. We use general polynomial chaos (gPC) polynomials [2, 3]. Hence the basis functions $\varphi_k : \Theta \rightarrow \mathbb{R}$ fulfill

$$\langle \varphi_k \varphi_{k'} \rangle = \delta_{k,k'} \quad \text{for all } k, k'. \quad (4)$$

For simplicity of exposition, let us start with a one-dimensional uncertainty ($p = 1$) and consider polynomial basis functions φ_i , where i corresponds to the degree of φ_i . Various intrusive methods can be used to approximate the moments, and they all come with certain advantages and shortcomings. Generally, intrusive methods specify a set of equations describing the time evolution of the moments, by first testing the original equations (1) with a finite set of the basis functions φ_i , say $i \in \{0, \dots, N\}$ leading to

$$\partial_t \langle \varphi_i \mathbf{u} \rangle + \nabla \cdot \langle \varphi_i \mathbf{f}(\mathbf{u}(t, \mathbf{x}, \cdot)) \rangle = 0 \quad \text{in } D, \quad i \in \{0, \dots, N\}. \quad (5)$$

This is not a closed system of equations because, due to the nonlinearity of f , $\langle \varphi_i \mathbf{f}(\mathbf{u}(t, \mathbf{x}, \cdot)) \rangle$ in general cannot be expressed in terms of the moments $\{\langle \varphi_0 \mathbf{u} \rangle, \dots, \langle \varphi_i \mathbf{u} \rangle\}$. The system is closed by providing an ansatz $\mathcal{U}(\hat{\mathbf{u}}_0, \dots, \hat{\mathbf{u}}_N) \simeq \mathbf{u}$ for the solution inside the flux, where $\hat{\mathbf{u}}_i \simeq \langle \varphi_i \mathbf{u} \rangle$ for each i . After inserting this closure into (5) and collecting the moments and basis functions into the matrix $\hat{\mathbf{u}} = (\hat{\mathbf{u}}_0, \dots, \hat{\mathbf{u}}_N)^T \in \mathbb{R}^{(N+1) \times m}$ and vector $\boldsymbol{\varphi} := (\varphi_0, \dots, \varphi_N)^T \in \mathbb{R}^{N+1}$, we have the approximate system

$$\partial_t \hat{\mathbf{u}} + \nabla \cdot \langle \boldsymbol{\varphi} \mathbf{f}(\mathcal{U}(\hat{\mathbf{u}})) \rangle = \mathbf{0} \quad \text{in } D, \quad (6a)$$

$$\hat{\mathbf{u}}(t = 0, \mathbf{x}) = \langle \boldsymbol{\varphi} \mathbf{u}_{\text{IC}}(\mathbf{x}, \cdot) \rangle. \quad (6b)$$

Note that a similar need to construct a closure also appears in fields such as kinetic theory, which is closely related to uncertainty quantification [4]. The choice of \mathcal{U} crucially affects properties of the resulting moment system. A straightforward closure choice makes use of the basis functions φ_i to represent the solution:

$$\mathcal{U}_k = \sum_{i=0}^N \hat{w}_{ik} \varphi_i \quad k \in \{1, \dots, m\},$$

where $\hat{w}_{ik} \in \mathbb{R}$ are expansion coefficients which are at this point free to choose and $\hat{\mathbf{w}} \in \mathbb{R}^{(N+1) \times m}$ is the matrix containing all the expansion coefficients. Assume that a given moment vector belongs to a function $\mathbf{u}_{\text{ex}} \in \mathbb{R}^m$. First, we pick $\hat{\mathbf{w}}$ such that the L^2 distance between this function and the polynomial approximation is minimized, i.e., we solve

$$(\hat{w}_{0k}, \dots, \hat{w}_{Nk}) = \arg \min_{\mathbf{w} \in \mathbb{R}^{N+1}} \mathcal{F}_k(\mathbf{w}), \quad \text{where } \mathcal{F}_k(\mathbf{w}) := \langle (\mathbf{w}^T \boldsymbol{\varphi} - \mathbf{u}_{\text{ex},k})^2 \rangle, \quad k \in \{1, \dots, m\}. \quad (7)$$

Since we are using orthonormal polynomials for basis functions, the resulting optimal expansion coefficients are simply the moments, i.e., $\hat{\mathbf{w}} \equiv \hat{\mathbf{u}}$. The resulting ansatz, which gives the stochastic-Galerkin (SG) closure [5] is

$$\mathcal{U}_{\text{SG}}(\hat{\mathbf{u}}) := \hat{\mathbf{u}}^T \boldsymbol{\varphi}.$$

Compared to other intrusive methods such as the intrusive polynomial moment (IPM) method [6] or the Roe transformation method [7], stochastic-Galerkin is computationally inexpensive, however suffers from a possible loss of hyperbolicity [6]. This necessitates manipulating the solution to be able to run the method in certain settings [8] or imposing a modification of the flux discretization [9]. Moreover, SG suffers from oscillations when the solution is not sufficiently smooth [10], which is a common issue in hyperbolic conservation laws where the solution tends to form discontinuous shocks [6]. These oscillations impair the solution quality and lead to physically incorrect solution values, which is why different strategies to mitigate oscillations have been developed. Oscillations dissolve when the number of chosen moments is increased [11, 12], which however adds computational costs. Additionally, amplifying the numerical diffusion mitigates oscillations [12], but at the same time leads to non-accurate solution approximations in deterministic regions. Furthermore, the use of multi-elements in the stochastic domain reduces artifacts which stem from oscillations [13, 14]. In [27] multi-elements are used for the IPM method. The strategy reduces numerical artifacts in the IPM solution while leading to a significant reduction in computational costs.

Another strategy for reducing oscillations in spectral methods is to apply a filter between time steps [15, 16]. Filters have for example been used in the context of kinetic theory [17, 18, 19, 20, 21] as well as spatial discretization methods [22, 23]. A filtered stochastic-Galerkin method has been proposed in [24], yielding a significant mitigation of oscillations. Just as the classical stochastic-Galerkin method, this filtered SG method can potentially lead to non-physical solution values, leading to a failure of the method. One possible strategy to derive a filter is to replace the SG cost function (7) by

$$\mathcal{F}^{(\lambda)}(\mathbf{w}_k; \hat{\mathbf{u}}_k) := \langle (\mathbf{w}_k^T \boldsymbol{\varphi} - \hat{\mathbf{u}}_k^T \boldsymbol{\varphi})^2 \rangle + \lambda T(\mathbf{w}_k), \quad (8)$$

for each $k \in \{1, \dots, m\}$, where $\hat{\mathbf{u}}_k$ is the k -th column of $\hat{\mathbf{u}}$ and T is a function which penalizes oscillations in the resulting ansatz $\mathbf{w}_j^T \boldsymbol{\varphi}$. The parameter λ is a user-determined filter strength. As an example, the L^2 filter for uniform distributions which has been investigated in [24] uses

$$T(\mathbf{w}) = \langle (\mathcal{L}(\mathbf{w}^T \boldsymbol{\varphi}))^2 \rangle, \quad \text{with } \mathcal{L}u := \frac{d}{d\xi} \left((1 - \xi^2) \frac{d}{d\xi} u \right). \quad (9)$$

Note that for more general distributions with gPC polynomials $P_i : \Theta \rightarrow \mathbb{R}$, the operator \mathcal{L} should be chosen to be the eigenoperator the chosen polynomial basis. I.e., we have $\mathcal{L}P_i = \mu_i P_i$, where μ_i are the corresponding eigenvalues. With this choice of T the minimizer of $\mathcal{F}^{(\lambda)}(\mathbf{w}_k; \hat{\mathbf{u}}_k)$ in (8) is given by the components $\bar{w}_{ik} := g_i(\lambda) \hat{u}_{ik}$, where for uniform distributions, the filter function g_i is

$$g_i(\lambda) := \frac{1}{1 + \lambda i^2 (i + 1)^2}. \quad (10)$$

Hence, the optimal expansion coefficients are dampened with increasing strength as the order i increases. A more sophisticated filter function takes the form

$$\tilde{g}_i(\lambda) := h(i/N)^{\lambda \Delta t}, \quad (11)$$

where Δt is the time step size of the time discretization chosen to numerically solve the SG moment system (6). Compared to (10), this filter adapts with the chosen time step size to ensure that increasing or decreasing Δt will not heavily weaken or amplify the effect of filtering. When $\Delta t \rightarrow 0$, the filtered equations tend towards a continuous limit [19]. Furthermore, a so-called filter order α can be introduced. This parameter aims at preserving the approximation order for sufficiently smooth functions, see e.g. [25]. One example for a filter of order of α is

$$h(\zeta) = \exp(c\zeta^\alpha) \quad (12)$$

where $c = \log(\varepsilon_M)$ and ε_M is the machine accuracy. This filter choice leads to the exponential filter [26]. Another choice, which is called the Erfc filter [25] uses

$$h(\zeta) = \frac{1}{2} \operatorname{erfc} \left(2\alpha^{1/2} (|\zeta| - 1/2) \right). \quad (13)$$

If we put the filter weights into the diagonal matrix $L(\lambda) := \text{diag}\{\tilde{g}_i(\lambda)\}_{i=0}^N$, we can write the solution as $\bar{\mathbf{u}} = L(\lambda)\hat{\mathbf{u}}$.

Within the full algorithm, the filter is applied as follows. First, a finite-volume method with forward-Euler time discretization for the moment equations (6) with the stochastic-Galerkin closure is given by

$$\hat{\mathbf{u}}_j^{n+1} = \hat{\mathbf{u}}_j^n - \frac{\Delta t}{\Delta x} (\mathbf{F}^*(\hat{\mathbf{u}}_j^n, \hat{\mathbf{u}}_{j+1}^n) - \mathbf{F}^*(\hat{\mathbf{u}}_{j-1}^n, \hat{\mathbf{u}}_j^n)), \quad (14)$$

where $\hat{\mathbf{u}}_j^n$ is the matrix of moment components for spatial cell j and time step n and \mathbf{F}^* is the numerical flux function (see Section 3.2 below for details). The filtered method is then given by

$$\bar{\mathbf{u}}_j^n = L(\lambda)\hat{\mathbf{u}}_j^n, \quad (15a)$$

$$\hat{\mathbf{u}}_j^{n+1} = \bar{\mathbf{u}}_j^n - \frac{\Delta t}{\Delta x} (\mathbf{F}^*(\bar{\mathbf{u}}_j^n, \bar{\mathbf{u}}_{j+1}^n) - \mathbf{F}^*(\bar{\mathbf{u}}_{j-1}^n, \bar{\mathbf{u}}_j^n)). \quad (15b)$$

We call this the filtered stochastic-Galerkin (fSG) method [24], and it gives accurate, non-oscillatory results if an adequate filter strength is chosen. Furthermore, the ability to analytically determine the minimizer of (8) allows low computational costs and therefore fast runtimes. Unfortunately, as with the original stochastic-Galerkin method, loss of hyperbolicity remains an issue, since the filter does not guarantee physically correct solution values. When for example taking a look at the Euler equations, the filter does not guarantee positivity of density and pressure. Hence, even though the filter tends to push the SG solution towards a physically feasible state, filtered stochastic-Galerkin can still generate physically incorrect solution values. This leads to a failure of the method, since various quantities such as the speed of sound will no longer be real numbers. One possibility of preserving physically correct solution bounds of the SG solution is to apply a hyperbolicity-preserving limiter to the SG method [8]. A strategy to combine the use of such a limiter as well as filters has been presented in [27].

A generalization of stochastic-Galerkin, which can preserve hyperbolicity without the use of a limiter is the intrusive polynomial moment (IPM) method [6]. The main idea is to close the system by minimizing a mathematical entropy of the conservation law (1) while fulfilling a moment constraint. For a given convex entropy $s : \mathbb{R}^m \rightarrow \mathbb{R}$ for the original problem (1), this constrained optimization problem reads

$$\mathcal{U}(\hat{\mathbf{u}}) = \arg \min_{\mathbf{u}} \langle s(\mathbf{u}) \rangle \quad \text{subject to } \hat{\mathbf{u}} = \langle \boldsymbol{\varphi} \mathbf{u} \rangle \quad (16)$$

where $\boldsymbol{\varphi} \mathbf{u}$ is an outer product, i.e., $[\boldsymbol{\varphi} \mathbf{u}]_{ik} = \varphi_i u_k$. Using duality, this constrained optimization problem can be transformed into the unconstrained finite-dimensional problem

$$\min_{\hat{\mathbf{w}} \in \mathbb{R}^{(N+1) \times m}} \left\{ \langle s_*(\hat{\mathbf{w}}^T \boldsymbol{\varphi}) \rangle - \hat{\mathbf{w}} \cdot \hat{\mathbf{u}} \right\}, \quad (17)$$

where $s_* : \mathbb{R}^m \rightarrow \mathbb{R}$ is the Legendre transformation of s and

$$\hat{\mathbf{w}} \cdot \hat{\mathbf{u}} = \sum_{i=0}^N \sum_{k=1}^m \hat{w}_{ik} \hat{u}_{ik}. \quad (18)$$

If we let $\hat{\mathbf{v}}(\hat{\mathbf{u}})$ denote the minimizer of the dual problem (17), the minimizer of the primal problem (16) is given by

$$\mathcal{U}(\hat{\mathbf{u}}) = (s')^{-1} (\hat{\mathbf{v}}(\hat{\mathbf{u}})^T \boldsymbol{\varphi}). \quad (19)$$

Here we use that $s'_*(v) = (s')^{-1}(v)$. It is important to point out that if the inverse of s' only takes values in a set $\mathcal{R}_{\mathbf{u}} \subset \mathbb{R}^m$, the presented optimization problem is only solvable if the moments $\hat{\mathbf{u}}$ lie inside the so-called realizable set

$$\mathcal{R} := \left\{ \hat{\mathbf{u}} \in \mathbb{R}^{(N+1) \times m} : \exists \mathbf{u} = \mathbf{u}(\xi) \in \mathcal{R}_{\mathbf{u}} \text{ for all } \xi \in \Theta \text{ such that } \hat{\mathbf{u}} = \langle \boldsymbol{\varphi} \mathbf{u} \rangle \right\}. \quad (20)$$

A detailed discussion of realizability in the context of uncertainty quantification can be found in [28].

By construction, the IPM system is hyperbolic, and, depending on the entropy chosen, the ansatz stays within certain physically relevant bounds [6]. These may be user-specified (through the choice of s) bounds which enforce the maximum-principle or the bounds of the domain of hyperbolicity (for example, for the Euler equations). While for scalar problems the solution is non-oscillatory [28], oscillations arise for systems like the Euler equations [24, 29], where the choice of admissible entropies is limited. Furthermore, since the optimization problem (16) can generally not be solved analytically, a numerical optimization method has to be used whenever the closure needs to be evaluated. This heavily increases computational costs, though these costs should be able to be alleviated by the use of multi-elements [27], high-order methods for the temporal-spatial discretization as well as by taking advantage of parallel scalability [30].

The objective of this paper is to develop a method that suppresses oscillations in the solution while guaranteeing hyperbolicity of the moment system. This goal is achieved by incorporating a filter into the IPM closure. To extend the IPM method to filters, first we simply apply the filter to the numerical solution of (6) (with the IPM closure) between time steps. But since an important property of the optimization problem is that it is infeasible when the given moment vector is not realizable, we must consider two options: either the filter must be chosen to preserve realizability or the optimization problem must be modified to become feasible for nonrealizable moment vectors. We propose two methods which follow these two paths. First, we design a filter which preserves realizability of the moments. Second, we change the IPM optimization problem to ensure a unique solution even if the moments are not realizable, allowing the use of common filters.

Parts of this work have been discussed in the PhD thesis [31]. To allow reproducing the numerical results of this work, the source code of the IPM solver implementation is publicly available [32] together with scripts which recompute all test cases from this paper. The implementation is an extension of the code framework presented in [29].

Before proceeding, we note that there are other ways to enforce the physical bounds of the solution and dampen oscillations in the closure. For example, one could enforce the physical bounds through constraints in the optimization problem instead of using the entropy s in the objective function. Also, instead of filtering the moment vector directly, oscillations could be dampened by penalizing them in the objective function. These alternatives also likely merit investigation, but we leave this for future work.

The rest of this paper is structured as follows: First we discuss how applying a standard filter to a realizable moment vector does not necessarily preserve its realizability and propose an alternate filter which does preserve realizability in Section 2. Then in the following Section 2.3, we apply the regularization technique from [33] to the optimization problem, thus making it feasible even for nonrealizable moment vectors and allowing the use of standard filters. We discuss our implementation in Sections 3 and present numerical results in Section 4. Section 5 draws conclusions and gives an outlook on future work.

2. Filtering and realizability

We consider two ways of incorporating a filter into the IPM method. First, we show that the filtered moments $L(\lambda)\hat{\mathbf{u}}$ are not guaranteed to fulfill the realizability condition (even when $\hat{\mathbf{u}}$ does) and therefore propose a modified filter which does preserve realizability. Second, we consider an alternative which modifies the IPM optimization problem (16) instead of the filter so that the problem is feasible even for non realizable moment vectors.

2.1. Effects of filtering on realizability

Before presenting the two novel approaches to include filters in the IPM method, we demonstrate the core problem that comes with this task. I.e., we demonstrate that the filtered moments $L(\lambda)\hat{\mathbf{u}}$ are not necessarily realizable, even when the moments $\hat{\mathbf{u}}$ are. Let us consider the scalar case, $m = 1$ and the entropy

$$s(u) = u \log(u). \tag{21}$$

Now, a solution to the IPM optimization problem (16) is only available if the moment vector lies in the realizable set \mathcal{R} , which we defined in (20). The realizable set which corresponds to the entropy (21) is

$$\mathcal{R} := \{\hat{\mathbf{u}} \in \mathbb{R}^{N+1} | \exists u : \Theta \rightarrow (0, \infty) \text{ such that } \hat{\mathbf{u}} = \langle \varphi u \rangle\}.$$

When $N = 2$ we have the more explicit form

$$\mathcal{R} = \{\hat{\mathbf{u}} \in \mathbb{R}^{N+1} | \hat{u}_0 > \hat{u}_2 > 0 \text{ and } \hat{u}_0 \hat{u}_2 > \hat{u}_1^2\}. \quad (22)$$

We refer to [34, 35, 36] for more information. In Figure 1a we see that $L(\lambda)\hat{\mathbf{u}}$ may be not realizable even when $\hat{\mathbf{u}}$ is. In this figure, we plot

$$L(\lambda)\mathcal{R}|_{\hat{u}_0=1} = \{L(\lambda)\hat{\mathbf{u}} | \hat{\mathbf{u}} \in \mathcal{R} \text{ and } \hat{u}_0 = 1\}$$

and see that for $\lambda \in \{0.05, 0.1, 0.2, 0.3\}$ we have $L(\lambda)\mathcal{R}|_{\hat{u}_0=1} \not\subset \mathcal{R}|_{\hat{u}_0=1}$, i.e., there are realizable moments which become nonrealizable after the filter is applied.

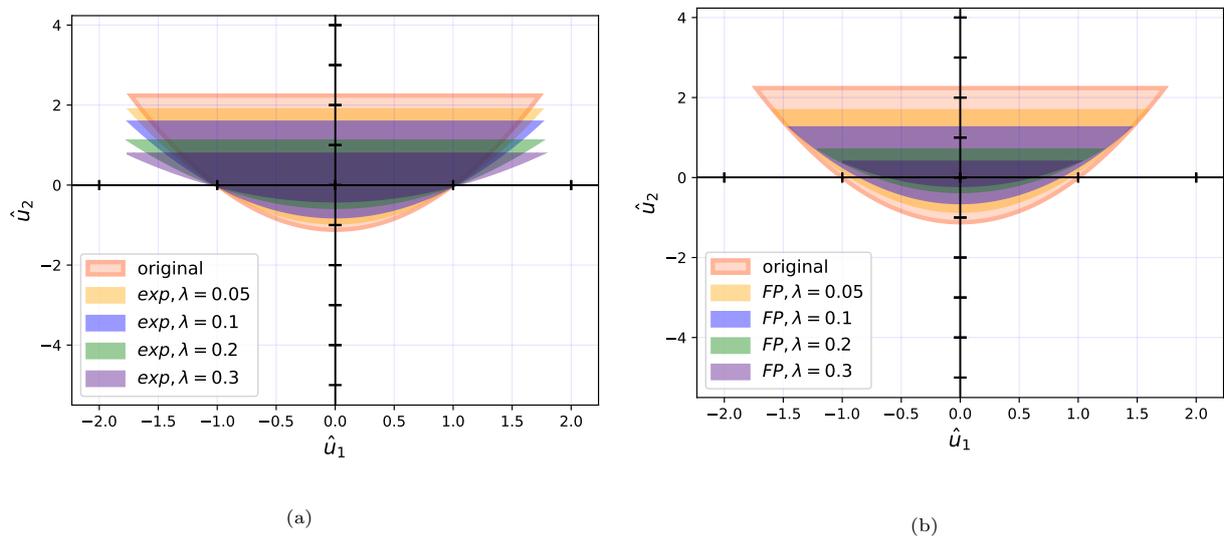


Figure 1: Images of the realizable set after application of filters. (a) Exponential filter with $\alpha = 7$ and $\Delta t = 0.1$, (b) Realizability-preserving Fokker-Planck filter

On the other hand, in Figure 1b we see the image of the realizable set \mathcal{R} after the application of a realizability-preserving filter, $F(\lambda)$, given below. Here one sees that $F(\lambda)\mathcal{R}|_{\hat{u}_0=1} \subset \mathcal{R}|_{\hat{u}_0=1}$ for each λ . In the following, we derive this realizability-preserving filter and discuss its properties.

2.2. A realizability-preserving filter

This section uses the Sturm-Liouville equation for orthogonal polynomials that we review in Appendix A for sake of completeness. Let us start by defining the filter for a scalar problem in the following theorem.

Theorem 1. *Assume that the chosen gPC basis is the solution to a Sturm-Liouville equation. I.e., a self-adjoint eigenoperator \mathcal{L} with eigenvalues μ_i for $i = 0, \dots, N$ exists. Furthermore, we assume that $\partial_t u = \mathcal{L}u$ fulfills a maximum principle. Let us define the filter function*

$$g_\lambda(i) = \exp(\mu_i \lambda). \quad (23)$$

Then, making use of the corresponding filter matrix which takes the form $F(\lambda) = \text{diag}\{\exp(\mu_i \lambda)\}_{i=0}^N$, the filtered moment vector $F(\lambda)\hat{\mathbf{u}} = \langle \boldsymbol{\varphi} u(\lambda, \cdot) \rangle$ lies in the realizable set

$$\mathcal{R} := \{ \hat{\mathbf{u}} \in \mathbb{R}^{N+1} \mid \exists u : \Theta \rightarrow (u_-, u_+) \text{ such that } \hat{\mathbf{u}} = \langle \boldsymbol{\varphi} u \rangle \},$$

for scalar solutions (i.e., $m = 1$) if the original moment vector $\hat{\mathbf{u}}$ is realizable.

Proof. The realizable moments, which we wish to filter can be written as $\hat{\mathbf{u}} = \langle \boldsymbol{\varphi} u_{\text{ex}} \rangle$, where $u_{\text{ex}}(\xi) \in (u_-, u_+)$. Consider the Fokker–Planck equation for $u = u(\lambda, \xi)$, where λ plays the role of time:

$$\partial_\lambda u = \mathcal{L}u \quad \text{with} \quad u(\lambda = 0, \xi) = u_{\text{ex}}(\xi).$$

The operator \mathcal{L} must be the eigenoperator (or Sturm–Liouville operator) to the chosen gPC basis. Since the solution to the Fokker–Planck equation fulfills the maximum principle and $u_{\text{ex}}(\xi) \in (u_-, u_+)$ for all $\xi \in \Theta$, we know $u(\lambda, \xi) \in (u_-, u_+)$ for all $\xi \in \Theta$ and $\lambda \geq 0$. Therefore the moments

$$\langle \boldsymbol{\varphi} u(\lambda, \cdot) \rangle \tag{24}$$

remain in the realizable set \mathcal{R} for all $\lambda \in [0, \infty)$. Furthermore, since \mathcal{L} is self-adjoint, these moments are straightforward to compute:

$$\partial_\lambda \langle \varphi_i u \rangle = \langle \varphi_i \mathcal{L}u \rangle = \langle (\mathcal{L}\varphi_i)u \rangle = \mu_i \langle \varphi_i u \rangle, \tag{25}$$

and thus

$$\langle \varphi_i u(\lambda, \cdot) \rangle = \exp(\mu_i \lambda) \langle \varphi_i u_{\text{ex}} \rangle. \tag{26}$$

Altogether with $\hat{\mathbf{u}} = \langle \boldsymbol{\varphi} u_{\text{ex}} \rangle$ and the definition of $F(\lambda) = \text{diag}\{\exp(\mu_i \lambda)\}_{i=0}^N$, we can write the filtered moment vector as $F(\lambda)\hat{\mathbf{u}} = \langle \boldsymbol{\varphi} u(\lambda, \cdot) \rangle$. The realizability of these moments is ensured by the maximum-principle satisfied by $u(\lambda, \xi)$. \square

When using uniform distributions, the chosen gPC basis are Legendre polynomials for which all assumptions of Theorem 1 hold. In this case, the eigenvalues are given by $\mu_i = -i(i+1)$, i.e. the filter matrix reads $F(\lambda) = \text{diag}\{\exp(-i(i+1)\lambda)\}_{i=0}^N$. Note that since $\mu_0 = 0$, the filtered zeroth order moment $\langle \varphi_0 u(\lambda, \cdot) \rangle$ is constant in λ and furthermore, as $\lambda \rightarrow \infty$, the solution $u(\lambda, \xi)$ becomes constant in ξ . Thus λ indeed acts like a filter-strength parameter.

Next we discuss how to extend this realizability-preserving filter to systems of conservation laws. First we define realizability more precisely for systems. Let $\mathcal{R}_{\mathbf{u}} \subseteq \mathbb{R}^m$ be the set of admissible states of the system. For example, above in the scalar case we took $\mathcal{R}_{\mathbf{u}} = (0, \infty)$ (although we could have just as easily taken $\mathcal{R}_{\mathbf{u}} = (u_{\min}, u_{\max})$, where u_{\min} and u_{\max} bound the initial condition, in order to enforce the maximum principle). In the case of the Euler equations, one would take

$$\mathcal{R}_{\mathbf{u}} = \{ (\rho, \rho \mathbf{v}, \rho E) \in \mathbb{R}^{d+2} : \rho > 0, p > 0 \} = \left\{ (\rho, \rho \mathbf{v}, \rho E) \in \mathbb{R}^{d+2} : \rho > 0, E > \frac{1}{2} \|\mathbf{v}\|^2 \right\} \tag{27}$$

which is the set of \mathbf{u} for which the Euler equations are hyperbolic. Here, $\rho \in \mathbb{R}$ denotes the density of the gas, $\mathbf{v} \in \mathbb{R}^d$ denotes the macroscopic gas velocity, $p \in \mathbb{R}$ is pressure and $\rho E \in \mathbb{R}$ is total energy. For a solution state \mathbf{u} which depends on the uncertainty ξ , let $\hat{\mathbf{u}} = (\hat{\mathbf{u}}_0, \hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_N)$ be the vector of moments with respect to the basis functions $\boldsymbol{\varphi}$. We call $\hat{\mathbf{u}}$ realizable when it belongs to the set

$$\mathcal{R} := \left\{ \hat{\mathbf{u}} \in \mathbb{R}^{(N+1) \times m} : \exists \mathbf{u} = \mathbf{u}(\xi) \in \mathcal{R}_{\mathbf{u}} \text{ for all } \xi \in \Theta \text{ such that } \hat{\mathbf{u}} = \langle \boldsymbol{\varphi} \mathbf{u} \rangle \right\}. \tag{28}$$

We assume that the underlying admissible set has the form

$$\mathcal{R}_{\mathbf{u}} := \left\{ \mathbf{u} \in \mathbb{R}^m : \exists f = f(v) \in I \subset \mathbb{R} \text{ for all } v \in V \subseteq \mathbb{R}^d \text{ such that } \mathbf{u} = \int_V \boldsymbol{\psi}(v) f(v) dv \right\} \tag{29}$$

where ψ is a given set of m basis functions in v , and I is usually $(0, \infty)$ or $(0, 1)$. For conservation laws derived from kinetic equations, v plays the role of the velocity variable from the underlying kinetic equation. This is the case for the Euler equations (where the f is the Maxwellian, $\psi(v) = (1, v, \frac{1}{2}|v|^2)$, and $I = (0, \infty)$) and more generally for entropy-based moment closures for kinetic equations [37]. Under this assumption, \mathcal{R} can be written as

$$\mathcal{R} = \left\{ \hat{\mathbf{u}} \in \mathbb{R}^{(N+1) \times m} : \exists f = f(\xi, v) \in I \text{ for all } (\xi, v) \in \Theta \times V \right. \\ \left. \text{such that } \hat{\mathbf{u}} = \left\langle \varphi \int_V \psi(v)^T f(\cdot, v) dv \right\rangle^T \right\}. \quad (30)$$

Now, since we have defined realizability for systems, let us discuss how the derived filter behaves in this case:

Theorem 2. *Assume that all assumptions from Theorem 1 hold. Applying the filter $F(\lambda)$ from Theorem 1 component-wise for systems of conservation laws preserves realizability. Hence, if $\hat{\mathbf{u}}_i \in \mathbb{R}^m$ collects the moments of order i for all states $k = 1, \dots, m$, the filtered moments $e^{\mu_i \lambda} \hat{\mathbf{u}}_i$ for $i = 0, \dots, N$ remain realizable.*

Proof. To filter a $\hat{\mathbf{u}} \in \mathcal{R}$ without destroying its realizability, let f be any distribution from (30) and then apply \mathcal{L} to its ξ -dependence to define g :

$$\partial_\lambda g(\lambda, \xi, v) = \mathcal{L}g(\lambda, \xi, v), \quad g(0, \xi, v) = f(\xi, v). \quad (31)$$

Again, the maximum principle for \mathcal{L} ensures $g(\lambda, \xi, v) \in I$ for all $(\lambda, v, \xi) \in [0, \infty) \times V \times \Theta$. Thus

$$\hat{\mathbf{u}}(\lambda) := \left\langle \int_V \psi g(\lambda, \cdot, v) dv \varphi^T \right\rangle^T \in \mathcal{R} \quad (32)$$

for all $\lambda \in [0, \infty)$. A practical expression for $\hat{\mathbf{u}}$ is also straightforward to derive, since

$$\begin{aligned} \partial_\lambda \hat{\mathbf{u}}_i &= \left\langle \varphi_i \int_V \psi \partial_\lambda g(\lambda, \cdot, v) dv \right\rangle \\ &= \left\langle \varphi_i \int_V \psi \mathcal{L}g(\lambda, \cdot, v) dv \right\rangle \\ &= \left\langle \mathcal{L}\varphi_i \int_V \psi g(\lambda, \cdot, v) dv \right\rangle \\ &= \mu_i \left\langle \varphi_i \int_V \psi g(\lambda, \cdot, v) dv \right\rangle \\ &= \mu_i \hat{\mathbf{u}}_i, \end{aligned}$$

i.e., $\hat{\mathbf{u}}_i(\lambda) = e^{\mu_i \lambda} \hat{\mathbf{u}}_i$. (This also confirms that $\hat{\mathbf{u}}_i(\lambda)$ is independent of the choice of f in (31).) Thus we see that we can simply apply the filter $F(\lambda)$ component-wise for systems of conservation laws without destroying realizability. \square

Note that the Fokker–Planck filter function (23) is closely related to the exponential filter (12) being applied at order $\alpha = 2$. Unfortunately, in uncertainty quantification high order filters appear to be a desirable choice, since they commonly yield satisfactory variance approximations, whereas low order filters heavily dampen the variance [27]. Therefore, we now turn to allowing the use of standard filters (with potentially high order) by enabling the dual problem to cope with non-realizable moment vectors.

2.3. Regularization

When faced with the potential infeasibility of the filtered moments, instead of modifying the filter to preserve realizability, we can instead modify the optimization problem to guarantee feasibility even for nonrealizable moment vectors.

A modification which achieves this is the regularization proposed in [38, 33], in which the equality constraints are exchanged with a penalty term in the objective function. Specifically, the primal optimization problem (16) defining the ansatz is replaced by

$$\mathcal{U}_\eta(\hat{\mathbf{u}}) = \arg \min_{\mathbf{u}} \left\{ \langle s(\mathbf{u}) \rangle + \frac{1}{2\eta} \|\hat{\mathbf{u}} - \langle \boldsymbol{\varphi} \mathbf{u} \rangle\|^2 \right\}, \quad (34)$$

where $\eta \in (0, \infty)$ is the regularization parameter and

$$\|\hat{\mathbf{w}}\|^2 = \sum_{i=0}^N \sum_{k=1}^m \hat{w}_{ik}^2. \quad (35)$$

This regularized problem is feasible for a larger set of moment vectors $\hat{\mathbf{u}}$, in particular when Θ is compact, it is feasible for all moment vectors.

Just as with the equality-constrained problem, the dual problem to (34) is finite dimensional:

$$\hat{\mathbf{v}}_\eta(\hat{\mathbf{u}}) := \arg \min_{\hat{\mathbf{v}} \in \mathbb{R}^{m \times (N+1)}} \left\{ \langle s_*(\hat{\mathbf{v}}^T \boldsymbol{\varphi}) \rangle - \hat{\mathbf{v}} \cdot \hat{\mathbf{u}} + \frac{\eta}{2} \|\hat{\mathbf{v}}\|^2 \right\}, \quad (36)$$

and the form of the primal minimizer is the same:

$$\mathcal{U}_\eta(\hat{\mathbf{u}}) = u_{ME}(\hat{\mathbf{v}}_\eta(\hat{\mathbf{u}})^T \boldsymbol{\varphi}), \quad (37)$$

and the flux for the resulting system is

$$\mathbf{f}(\hat{\mathbf{u}}) = \langle \boldsymbol{\varphi} \hat{\mathbf{f}}(\mathcal{U}_\eta(\hat{\mathbf{u}})) \rangle. \quad (38)$$

This gives a hyperbolic system of conservation laws, and, when $\hat{\mathbf{v}}_\eta$ is defined for all $\hat{\mathbf{u}}$, we can implement (15) directly with a numerical flux \mathbf{F}^* based on (38).

One remaining question is how to choose the regularization parameter η . A discussion of this issue can be found in [33], where the authors choose η according to the Morozov discrepancy principle, when interpreting the numerical errors as noise. Hence, η is chosen such that the effect of regularization does not dominate the numerical error.

Before discussing the space and time discretization as well as implementation of the IPM scheme, let us take a look at the Fokker–Planck and exponential filter functions, which are given in equations (23) and (12). To get a better impression of how both filters dampen moments, we plot both filter functions in Figure 2. The chosen filter strength is $\lambda = 0.04$. The time step size Δt for the exponential filter is chosen such that the exponential filter at order two resembles the Fokker–Planck filter. It can be seen that Fokker–Planck filtering (with a scaled filter strength) behaves like a second order exponential filter. When increasing the order of the exponential filter, the moment order i at which the filter starts to dampen gPC coefficients increases. Note that this will especially affect the variance approximation, which (for a polynomial solution ansatz) is given by a sum over the squares of moments higher than zero.

3. Implementation

In this section we discuss the implementation of our numerical scheme for solving the filtered system (15).

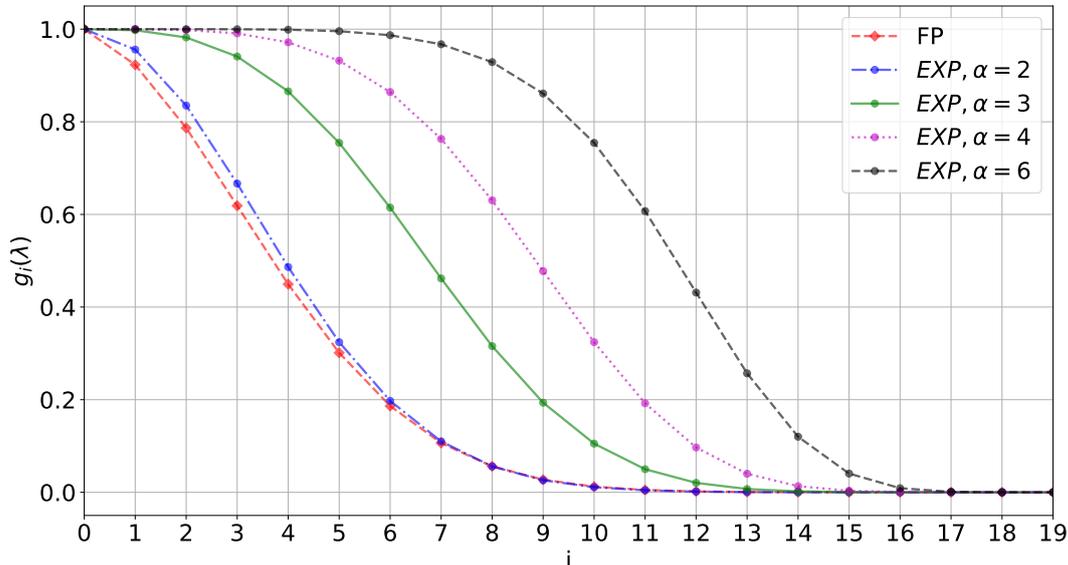


Figure 2: Fokker–Planck (FP) and exponential (EXP) filter functions for filter strength $\lambda = 0.04$. The exponential filter uses a time step parameter of $\Delta t = 12.5$ and varying orders α . The time step size is chosen such that the EXP filter for order 2 resembles FP filtering.

3.1. Solving the dual problem

The IPM method relies on solving the constrained optimization problem (16), or in the regularized case the unconstrained problem (34). Both of these problems are infinite-dimensional, so we solve their dual problems (17) and (36), which are finite-dimensional and unconstrained. In both cases we used Newton’s method stabilized with the standard backtracking linesearch. Our stopping criterion is the norm of the dual, i.e.,

$$\|\langle \varphi s'_*(\hat{\mathbf{v}}^T \varphi) \rangle - \hat{\mathbf{u}}\| < \tau \quad \text{or} \quad \|\langle \varphi s'_*(\hat{\mathbf{v}}^T \varphi) \rangle + \eta \hat{\mathbf{v}} - \hat{\mathbf{u}}\| < \tau \quad (39)$$

for the original and regularized cases respectively, where $\tau \in (0, \infty)$ is the user-specified tolerance parameter. The integrals needed to compute the dual objective function and its derivatives are computed with quadrature.

The main challenge in the implementation of the IPM method (and minimal entropy methods in general) is that the optimization problem is ill conditioned when the moment vector lies near the boundary of the realizable set (or in the regularized case, when the regularization parameter is small and the moment vector of the solution lies close to the boundary of the realizable set). In these cases the optimization algorithm needs many iterations to converge or may be unable to find a solution satisfying the stopping criterion. More information on this issue can for example be found in [28] for IPM and [39] for minimal entropy methods in kinetic theory.

3.2. Spatial and temporal discretization

Let us now write down the implementation of the regularized IPM method including filters. In order to discretize the closed moment system (6), we choose a spatial grid x_1, \dots, x_{N_x} and a time grid $0 = t_0, \dots, t_{N_t} = t_{\text{end}}$. We then represent the discretized solution by

$$\hat{u}_{ij}^n \simeq \frac{1}{\Delta x} \int_{x_{j-1/2}}^{x_{j+1/2}} \hat{u}_i(t_n, x) dx, \quad (40)$$

where i is the moment order and j is the spatial cell index. The moment vector is collected in $\hat{\mathbf{u}}_j^n = (\hat{u}_{ij}^n)_{i=0}^N$ for every spatial cell j and time step n . The corresponding dual state is denoted by

$$\hat{\mathbf{v}}_j^n := \hat{\mathbf{v}}(\hat{\mathbf{u}}_j^n),$$

when regularization is not used; when regularization is used, $\hat{\mathbf{v}}$ is replaced with $\hat{\mathbf{v}}_\eta$ (see (36)). Given a numerical flux $\mathbf{f}^*(\mathbf{u}_\ell, \mathbf{u}_r)$ for the deterministic system, a flux for the moment system is constructed by

$$\mathbf{F}^*(\hat{\mathbf{u}}(\hat{\mathbf{v}}_\ell), \hat{\mathbf{u}}(\hat{\mathbf{v}}_r)) := \mathbf{G}^*(\hat{\mathbf{v}}_\ell, \hat{\mathbf{v}}_r) := \langle \varphi \mathbf{f}^*(s'_*((\hat{\mathbf{v}}_\ell)^T \varphi), s'_*((\hat{\mathbf{v}}_r)^T \varphi)) \rangle. \quad (41)$$

We give this flux in terms of dual variables, since this more closely reflects the implementation. To compute the integrals in this expression one commonly uses a Gauss–Legendre quadrature rule, which equals the quadrature choice of the dual optimizer. This choice of the numerical flux is common in kinetic theory, where it is called the kinetic flux [40, 41, 42, 43]. It is frequently used in the context of uncertainty quantification, see e.g. [24, 29, 28, 44, 8, 45].

The basic strategy of the algorithm, as given in (15), is, in every time step, first to apply a filter to the moment vector in each spatial cell, and then perform an update in time starting from these filtered moment vectors. Some details of the implementation vary depending on which method we choose, so we present the methods separately in Algorithms 1 and 2.

In Algorithm 1, we note that after the dual variables are computed numerically for the filtered moment vector (using the Fokker–Planck filter), we compute in Line 7 the moment vector associated with these approximate dual variables. This is done to ensure the realizability of the moment vector at the next time step, $\hat{\mathbf{u}}_j^{n+1}$, see [28]. In Algorithm 2, this step is unnecessary.

Algorithm 1 Realizable Filtered IPM Method

```

1:  $\hat{\mathbf{u}}_j^0 \leftarrow$  setup Initial Conditions for all cells  $j$ 
2: choose filter strength  $\lambda \in (0, \infty)$ 
3: for  $n = 0$  to  $N_t$  do
4:   for  $j = 1$  to  $N_x$  do
5:      $\bar{\mathbf{u}}_j^n \leftarrow F(\lambda) \hat{\mathbf{u}}_j^n$ 
6:      $\tilde{\mathbf{v}}_j^n \leftarrow \hat{\mathbf{v}}(\bar{\mathbf{u}}_j^n)$  using Newton's method with gradient tolerance  $\tau$ 
7:      $\tilde{\mathbf{u}}_j^n \leftarrow \langle \varphi s'_*((\tilde{\mathbf{v}}_j^n)^T \varphi) \rangle$ 
8:   end for
9:   for  $j = 1$  to  $N_x$  do
10:     $\hat{\mathbf{u}}_j^{n+1} \leftarrow \bar{\mathbf{u}}_j^n - \frac{\Delta t}{\Delta x} (\mathbf{G}^*(\tilde{\mathbf{v}}_j^n, \tilde{\mathbf{v}}_{j+1}^n) - \mathbf{G}^*(\tilde{\mathbf{v}}_{j-1}^n, \tilde{\mathbf{v}}_j^n))$ 
11:   end for
12: end for

```

Algorithm 2 Regularized Filtered IPM Method

```

1:  $\hat{\mathbf{u}}_j^0 \leftarrow$  setup Initial Conditions for all cells  $j$ 
2: choose filter strength  $\lambda \in (0, \infty)$  and regularization parameter  $\eta \in (0, \infty)$ 
3: for  $n = 0$  to  $N_t$  do
4:   for  $j = 1$  to  $N_x$  do
5:      $\bar{\mathbf{u}}_j^n \leftarrow L(\lambda) \hat{\mathbf{u}}_j^n$ 
6:      $\tilde{\mathbf{v}}_{\eta,j}^n \leftarrow \hat{\mathbf{v}}_\eta(\bar{\mathbf{u}}_j^n)$  using Newton's method with gradient tolerance  $\tau$ 
7:   end for
8:   for  $j = 1$  to  $N_x$  do
9:      $\hat{\mathbf{u}}_j^{n+1} \leftarrow \bar{\mathbf{u}}_j^n - \frac{\Delta t}{\Delta x} (\mathbf{G}^*(\tilde{\mathbf{v}}_{\eta,j}^n, \tilde{\mathbf{v}}_{\eta,j+1}^n) - \mathbf{G}^*(\tilde{\mathbf{v}}_{\eta,j-1}^n, \tilde{\mathbf{v}}_{\eta,j}^n))$ 
10:   end for
11: end for

```

4. Results

In the following we will present numerical results for the two proposed filtering techniques. All results can be reproduced with the source code [32], which makes use of the intrusive UQ framework presented in [29].

4.1. Effects of the Regularization

Before moving to filtering, we study the effects of the regularization for different regularization strengths η . For this we investigate the following test case: We wish to solve the uncertain Sod shock tube problem, which is given by

$$\partial_t \begin{pmatrix} \rho \\ \rho v \\ \rho e \end{pmatrix} + \partial_x \begin{pmatrix} \rho v \\ \rho v^2 + p \\ v(\rho E + p) \end{pmatrix} = \mathbf{0},$$

with the initial conditions

$$\begin{aligned} \rho_{\text{IC}} &= \begin{cases} \rho_L & \text{if } x < x_{\text{interface}}(\xi) \\ \rho_R & \text{else} \end{cases} \\ (\rho v)_{\text{IC}} &= 0 \\ (\rho E)_{\text{IC}} &= \begin{cases} \rho_L E_L & \text{if } x < x_{\text{interface}}(\xi) \\ \rho_R E_R & \text{else} \end{cases} \end{aligned}$$

Here, ρ is the density, u is the velocity and e is the specific total energy. For a given heat capacity ratio γ which in our case has a value of 1.4, the pressure p can be determined from

$$p = (\gamma - 1)\rho \left(E - \frac{1}{2}v^2 \right).$$

The random interface position is given by $x_{\text{interface}}(\xi) = x_0 + \sigma\xi$ with a uniformly distributed random variable $\xi \sim U([-1, 1])$. We use Dirichlet boundary conditions at the left and right boundary. A reference solution is given by the analytic solution of the Sod shock tube test case. We approximate expected value and variance of the analytic solution by using a 100 point Gauss-Legendre quadrature rule. The remaining parameter values are

$[a, b] = [0, 1]$	range of spatial domain
$N_x = 2000$	number of spatial cells
$t_{\text{end}} = 0.14$	end time
$x_0 = 0.5, \sigma = 0.05$	interface position parameters
$\rho_L, p_L = 1.0, \rho_R = 0.125, p_R = 0.1$	initial states
$N = 10$	polynomial degree
$N_q = 30$	number of quadrature points
$\tau = 10^{-7}$	gradient tolerance for IPM

The IPM method uses the entropy

$$s(\rho, \rho v, \rho e) = \rho \ln \left(\rho^{-\gamma} \left(\rho e - \frac{(\rho v)^2}{2\rho} \right) \right).$$

One observes that this problem cannot be solved with SG or fSG, since negative densities ρ will show up already in the first time iteration. We now run this test case using IPM with different values for the regularization strength η . Hence, we run Algorithm 2 without applying the filter. Taking a look at the resulting expected value and variance of the density ρ in Figure 3, one sees that a big regularization

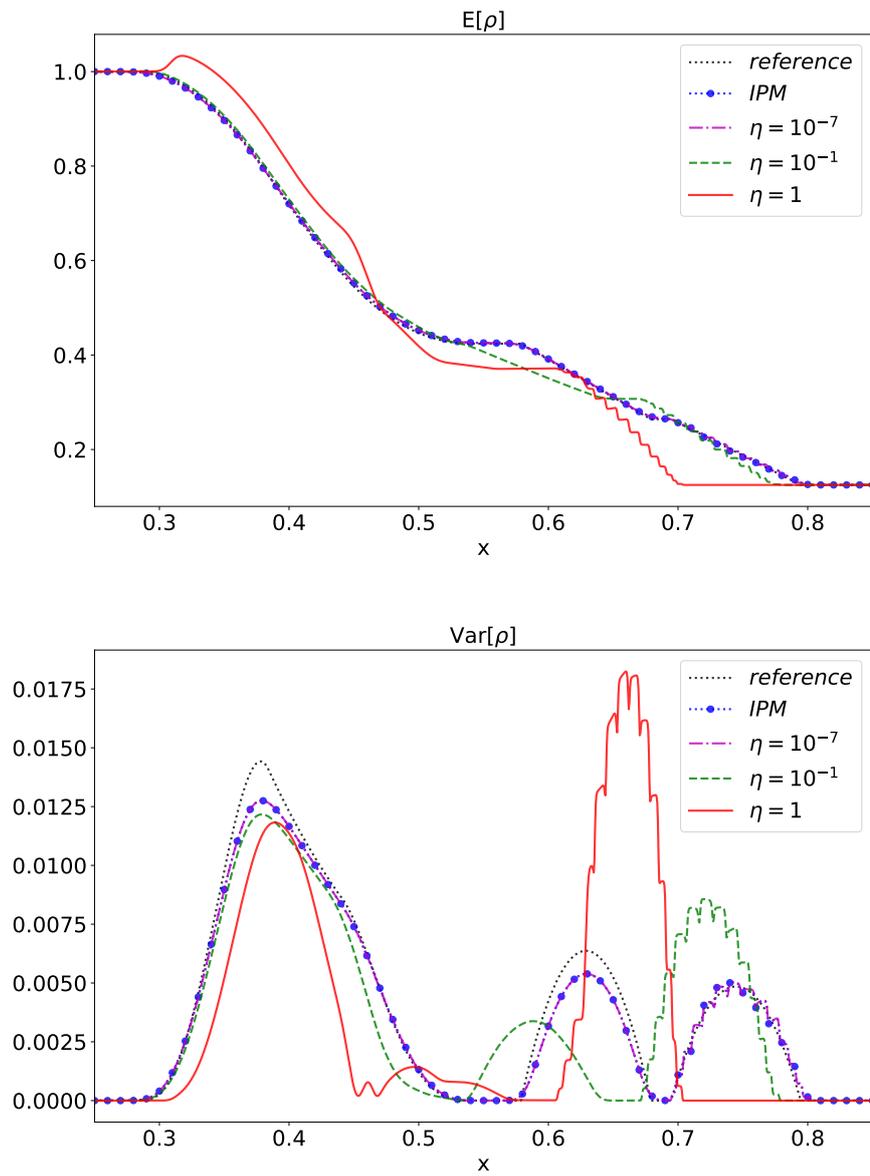


Figure 3: Expected value and variance of gas density for different regularization strengths.

parameter will heavily affect the solution. As the regularization strength decreases, the solution will approach the IPM solution and with $\eta = 10^{-7}$ the regularized solution shows good agreement with IPM. When the regularization parameter is not sufficiently small, we observe significant effects on higher order moments: When looking at the variance of the rarefaction wave and the contact discontinuity, one notices strong dampening. Note however that the variance in these regimes is very sensitive to the method used, see for example [24, 8]. It is important to point out that commonly, a regularization strength of $\eta = 10^{-7}$ is chosen and we compare against high values of the regularization strength to demonstrate the effects a regularization might have on the solution. For small regularization strength, the resulting numerical solution agrees well with the classical IPM solution.

The regularized solution with a small regularization parameter as well as the original IPM method show good approximation behavior in most parts of the solution. However, especially at the shock position, all methods will yield a non-physical, step-like approximation of the expected value and oscillatory results for the corresponding variance. Note that the step-like approximation of the IPM solution only becomes visible when further zooming into the shock region. For a deeper study of these artifacts we refer to [27].

This work focuses on obtaining better approximations at the shock position through filtering techniques. Note that the filter will further dampen the variance approximations at the rarefaction wave and the shock discontinuity. One idea to mitigate these effects that we leave for future studies is choosing an adaptive filter strength as done in [24].

4.2. Filtering for Sod's shock tube

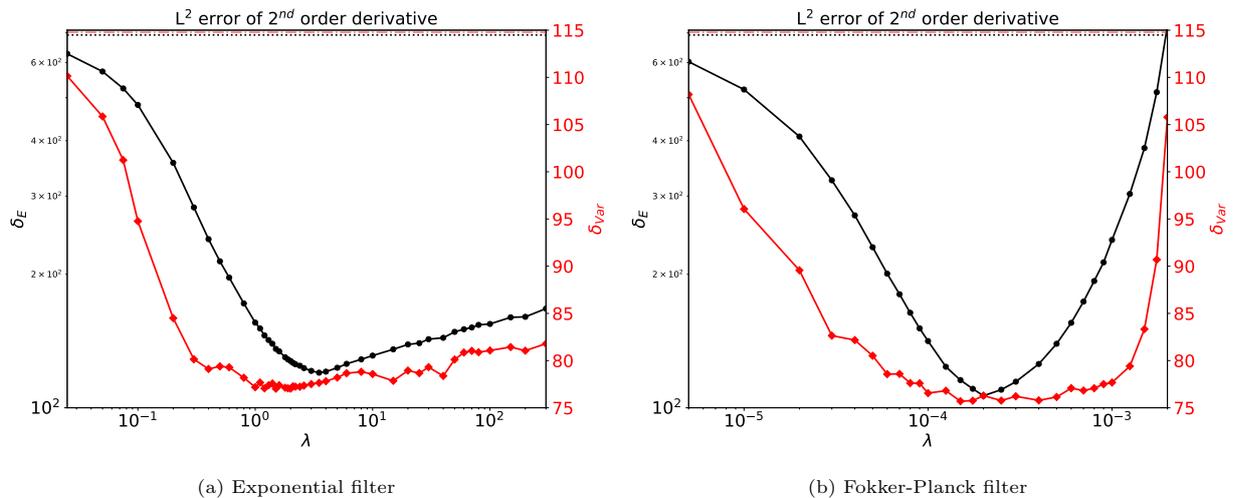


Figure 4: δ_E (in black with dots) and δ_{var} (in red with diamonds) for different filter strengths when using the exponential and Fokker-Planck filter. The exponential filter uses a regularization strength $\eta = 10^{-7}$ and order $\alpha = 10$. Values for δ_E and δ_{var} without filters are added as a straight line.

Let us start with studying the effects of the filter strength on the solution quality. As a measure for oscillations, let us investigate the error of second order derivatives. If the distance of a discrete numerical solution $u_\Delta := (u_1, \dots, u_{N_x})^T$ to a reference density $u_{\text{ex},\Delta} := (u_{\text{ex},1}, \dots, u_{\text{ex},N_x})^T$ is given by $e_\Delta := u_{\text{ex},\Delta} - u_\Delta$, we wish to determine

$$\delta_E := \sqrt{\sum_{j=1}^{N_x} \Delta x_j (\partial_{xx} \mathbb{E}[e_j])^2}$$

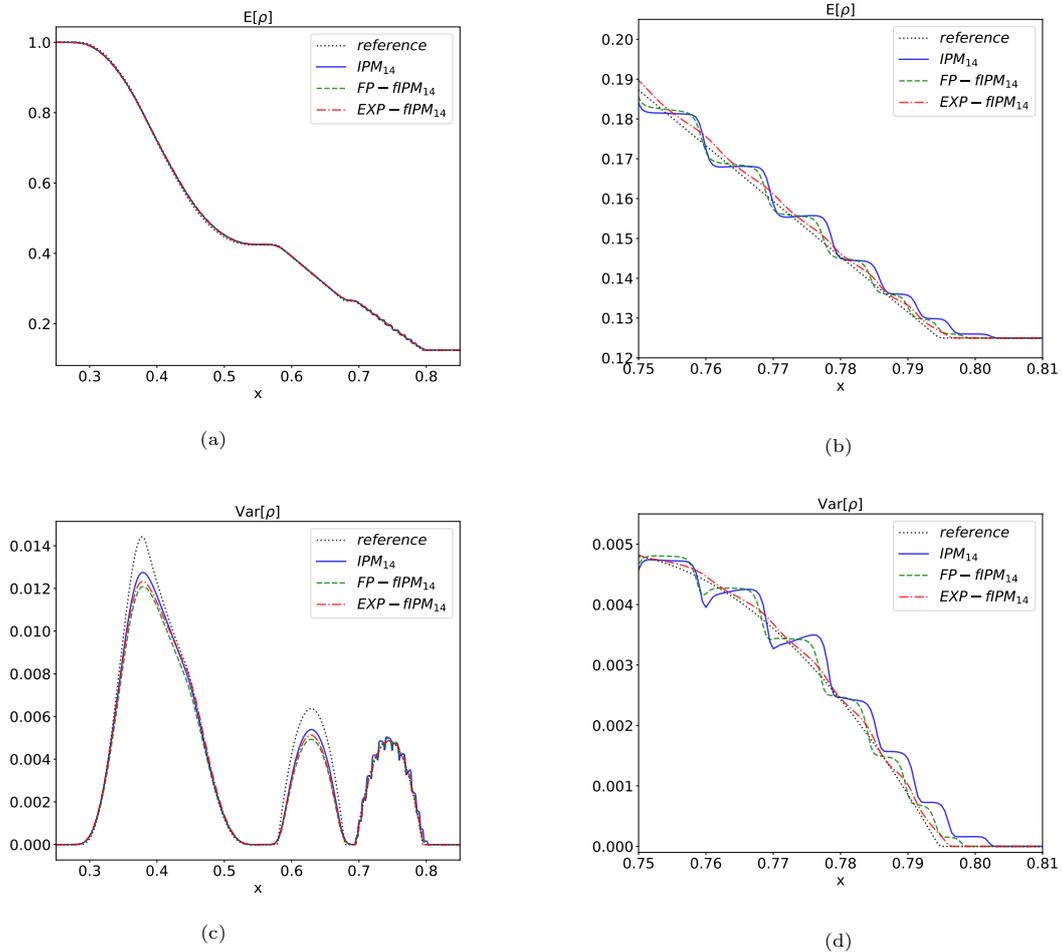


Figure 5: Expectation and variance computed with IPM and filtered IPM ($\lambda = 2, \alpha = 10$ for exponential filtering and $\lambda = 5 \cdot 10^{-5}$ for Fokker-Planck filtering) for Sod's shock tube. The exact expected value is depicted in red, the exact variance in blue. (a) Expected density. (b) Zoomed view on shock for expected density. (c) Variance of density. (d) Zoomed view on shock for variance of density.

and

$$\delta_{\text{Var}} := \sqrt{\sum_{j=1}^{N_x} \Delta x_j (\partial_{xx} \text{Var}[e_j])^2}.$$

The resulting curves for both variables when choosing varying filter strengths can be found In Figure 4. Here and in subsequent uses of the exponential and erfc filter, we solve issues resulting from realizability by picking a constant regularization parameter $\eta = 10^{-7}$. In all of our numerical experiments, this choice allowed solving the regularized dual problem. The filter order is here chosen to be $\alpha = 10$. Let us point out that the chosen measures δ_E and δ_{E} do not necessarily provide a comparison between different methods and simply allows an idea of a meaningful range for the filter strength. Furthermore, we evaluate both error measures only in the shock region $x \in [0.7, 0.8]$ to remove dampening effects from the rarefaction wave. It becomes clear that the optimal filter strength in terms of expectation and variance differs. Both filters reach an optimal variance value at a smaller filter strength compared to the optimal expectation. This means that while a bigger filter strength gives a good expectation approximation it heavily dampens out the variance,

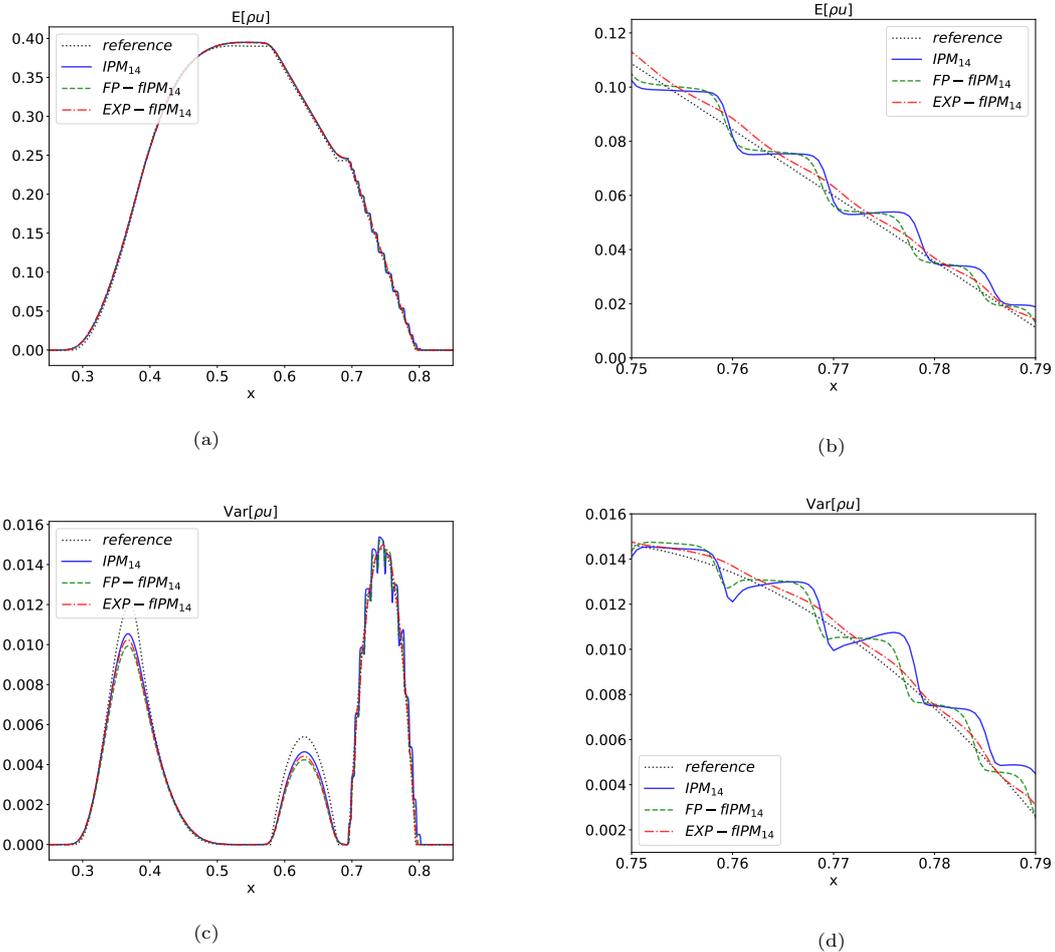


Figure 6: Expectation and variance computed with IPM and filtered IPM ($\lambda = 2$, $\alpha = 10$ for exponential filtering and $\lambda = 5 \cdot 10^{-5}$ for Fokker-Planck filtering) for Sod's shock tube. The exact expected value is depicted in red, the exact variance in blue. (a) Expected momentum. (b) Zoomed view on shock for expected momentum. (c) Variance of momentum. (d) Zoomed view on shock for variance of momentum.

i.e., yielding a significant error in the uncertainty. Therefore, let us in the following choose $\lambda = 2$ for the exponential filter. Fokker-Planck filtering seems to be optimal for values around $\lambda = 10^{-4}$ around the shock. However, this choice leads to heavy dampening of the variance at the rarefaction wave, which is why we slightly reduce this value to $\lambda = 5 \cdot 10^{-5}$. In this case, the Fokker-Planck filter shows a similar behaviour as the exponential filter at the rarefaction wave.

Note that we fixed the filter order $\alpha = 10$ for exponential filtering. Without going into detail, let us remark that increasing the order of the exponential filter mitigates the dampening effect at the rarefaction wave when increasing the filter strength. Thereby, a satisfactory value for λ , which improves the shock structure while preserving the remainder of the solution is easier to find. Choosing the order α too large will however require a significantly increased filter strength to obtain an improvement at the shock position. Since this behavior is difficult to capture by an error measure, we do not show a detailed parameter study as done for the filter strength. All remaining parameters are chosen as in section 4.1. Resulting expected values and variances are depicted in Figure 5. Again, the reference solution has been computed from the analytic solution of the Sod shock tube test case. As before, the expected value and variance of the analytic solution

are approximated by a 100 point Gauss-Legendre quadrature rule.

We first examine the shock position: The expected value and variance of the density again show heavy oscillations in the variance and a step-like profile at the shock if no filtering is applied. Using the exponential filter leads to a smooth linear connection between the left and right shock state for the expected value, which shows good agreement with the exact solution. The variance is slightly dampened, however its smooth profile nicely captures the main characteristics of the exact variance. This behavior can be examined in detail when zooming onto the shock position as done in Figures 5b and 5d. The Fokker-Planck filter appears to be less effective. Compared to standard IPM, the filter yields an improved approximation at the shock at the costs of dampening the variance at the rarefaction wave and contact discontinuity. When comparing exponential and Fokker-Planck filtering, one observes that exponential filtering yields a more satisfactory solution approximation, i.e., the shock region nicely captures the exact behavior while the variance at the rarefaction wave is only slightly damped. A similar conclusion can be drawn when looking at the expected momentum and its corresponding variance in Figure 6.

There are different strategies that can be chosen to mitigate non-physical dampening effects at the rarefaction wave. One example is using an adaptive filter strength. This approach can be combined with adaptivity in general [29]. In this work, however, we focus on mitigating oscillations at the shock and leave other open questions to future work.

4.3. Filtering for high density shock

In the following, we keep the same parameter setting, but change the initial condition to verify if the previous choice of the filtering parameters still yields a satisfactory result. The test case we investigate has an increased density state of $\rho_R = 0.8$ on the right hand side. The remaining shock states are left untouched. The resulting expected density and corresponding variance can be found in Figure 7.

Again, the reference solution has been computed with the analytic solution. Expected value and variance of the analytic solution are approximated by taking a 100 point Gauss-Legendre quadrature. Compared to the previous setting, the test case shows a large amount of density, which travels to the right hand side. In this test case, the rarefaction wave becomes less dominant. One observes that using the standard IPM method leads to heavy oscillations in the IPM solution of both, the expected value and the variance. Moreover, the variance approximation of IPM is heavily dampened out, which is most likely an artifact from the artificial viscosity of the finite volume method [46, Chapter 16.1]. Applying filters to this problem resolves oscillations. Especially the exponential filter captures the behavior of the exact variance nicely. Oscillations in the variance approximation are significantly mitigated, however the filter is not able to reduce the dampening effect the finite volume scheme has on the variance. The Fokker-Planck filter improves the approximation of the expected value. However, while oscillations in the variance are dampened, the Fokker-Planck filter further reduces the value of the variance approximation.

4.4. Filtering for Euler 2D

In the following, we will move to the two-dimensional Euler equations

$$\partial_t \begin{pmatrix} \rho \\ \rho v_1 \\ \rho v_2 \\ \rho E \end{pmatrix} + \partial_{x_1} \begin{pmatrix} \rho v_1 \\ \rho v_1^2 + p \\ \rho v_1 v_2 \\ v_1(\rho E + p) \end{pmatrix} + \partial_{x_2} \begin{pmatrix} \rho v_2 \\ \rho v_1 v_2 \\ \rho v_2^2 + p \\ v_2(\rho E + p) \end{pmatrix} = \mathbf{0},$$

where the pressure is given by

$$p = (\gamma - 1)\rho \left(E - \frac{1}{2}(v_1^2 + v_2^2) \right).$$

We will be interested in quantifying the uncertainty of a shock inside a two-dimensional nozzle. The nozzle geometry is composed of a chamber on the left, a throat in the center and the main nozzle region on the

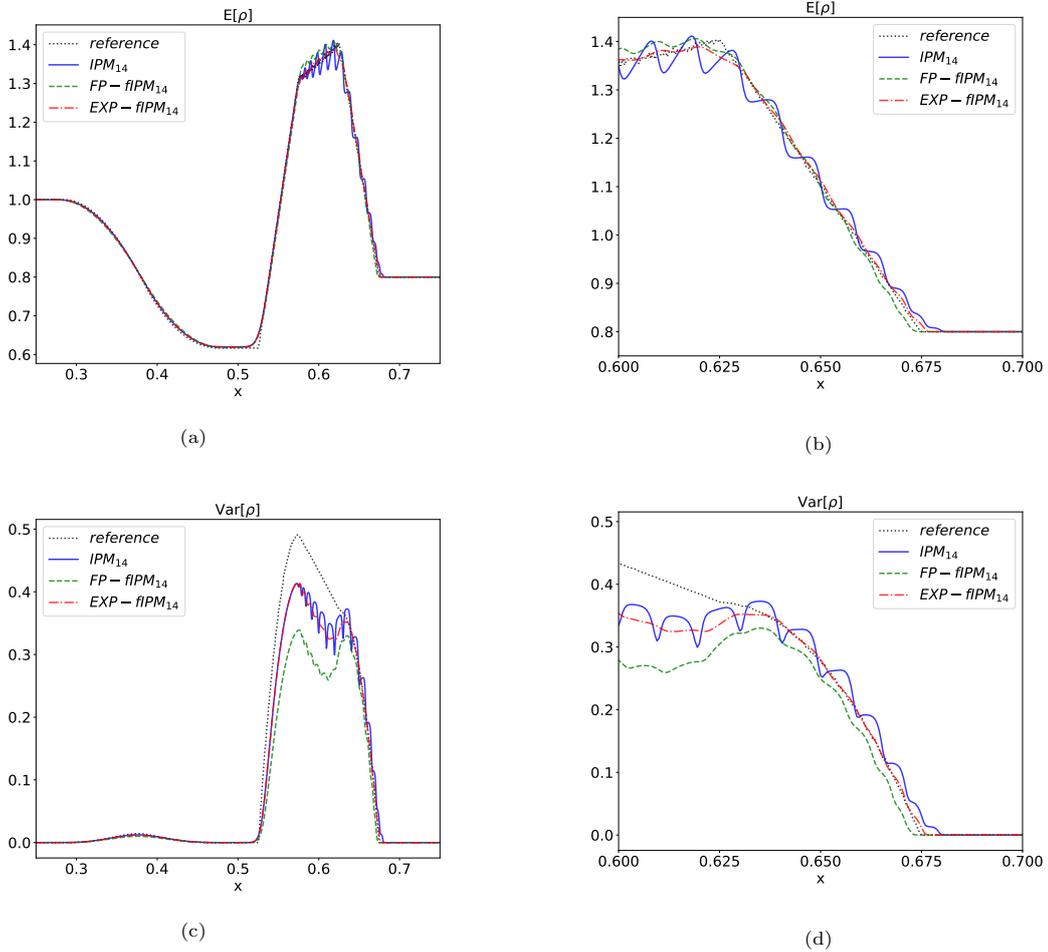


Figure 7: Expectation and variance for the high density shock computed with IPM and filtered IPM ($\lambda = 2, \alpha = 10$ for exponential filtering and $\lambda = 10^{-5}$ for Fokker-Planck filtering). The exact expected value is depicted in red, the exact variance in blue. (a) Expected density. (b) Zoomed view on shock for expected density. (c) Variance of density. (d) Zoomed view on shock for variance of density.

right. As in Sod's shock tube experiment, the gas is initially at rest with an uncertain shock occurring inside the domain. The initial condition is chosen to be

$$\rho_{\text{IC}} = \begin{cases} \rho_L & \text{if } x < x_{\text{interface}}(\xi) \\ \rho_R & \text{else} \end{cases}$$

$$(\rho \mathbf{v})_{\text{IC}} = \mathbf{0}$$

$$(\rho E)_{\text{IC}} = \begin{cases} \rho_L E_L & \text{if } x < x_{\text{interface}}(\xi) \\ \rho_R E_R & \text{else} \end{cases}$$

and we pick the shock states as in Section 4.3, i.e., the density on the right hand side is set to $\rho_R = 0.8$. The shock position $x_{\text{interface}}(\xi)$ is distributed uniformly inside the throat. Here we make the choice $x_{\text{interface}}(\xi) = x_{\text{th}} + \sigma \xi$, where x_{th} is the center x -coordinate of the throat. Since the throat has a length of 1, i.e., it ranges from $x_{\text{th}} - 0.5$ to $x_{\text{th}} + 0.5$, we choose $\sigma = 0.5$ and again use a uniformly distributed $\xi \sim U([-1, 1])$. A reference solution, depicted on the top of Figures 8 and 9, has been computed using Collocation (see [47])

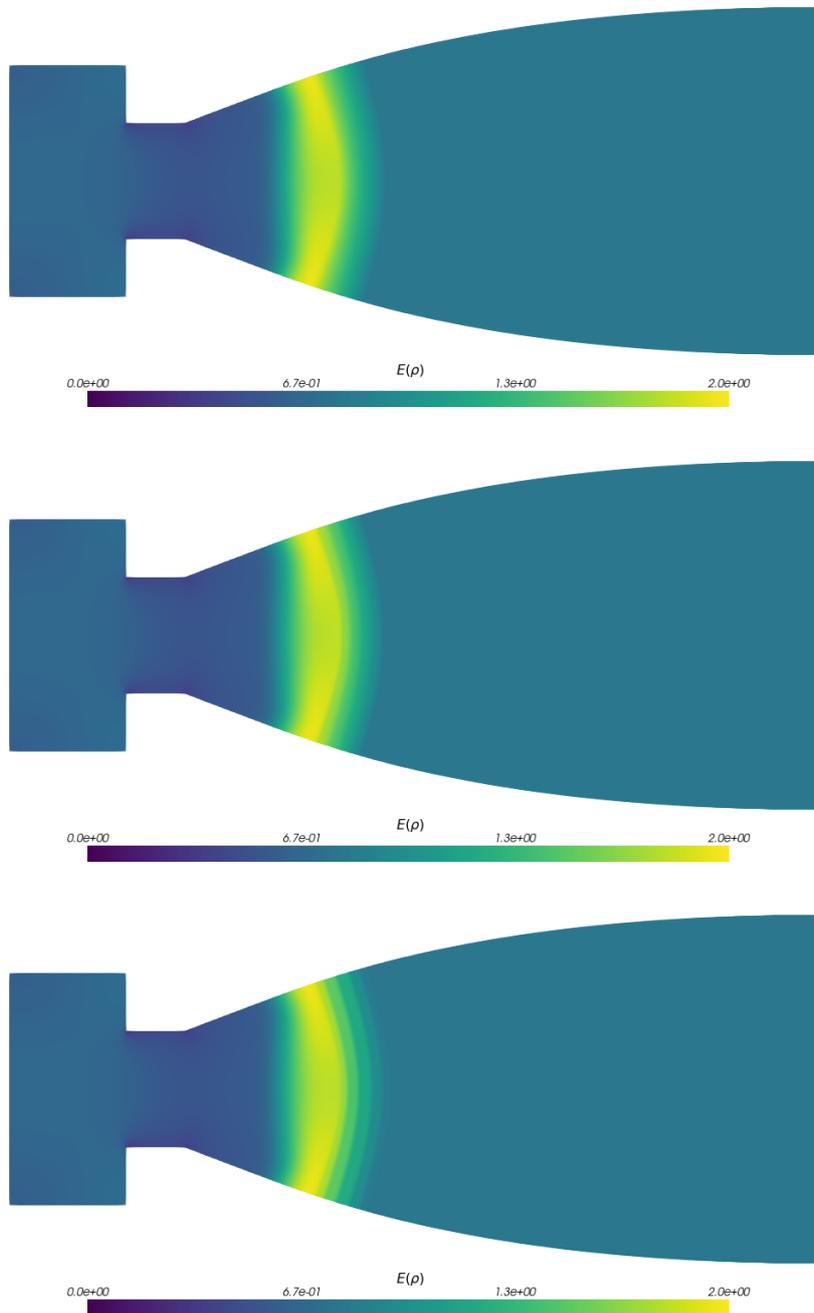


Figure 8: Expected density with different methods. From top to bottom: reference solution, fIPM, IPM.

with 100 Gauss-Legendre quadrature points. The solution shows an uncertain shock wave which travels to the right side of the nozzle.

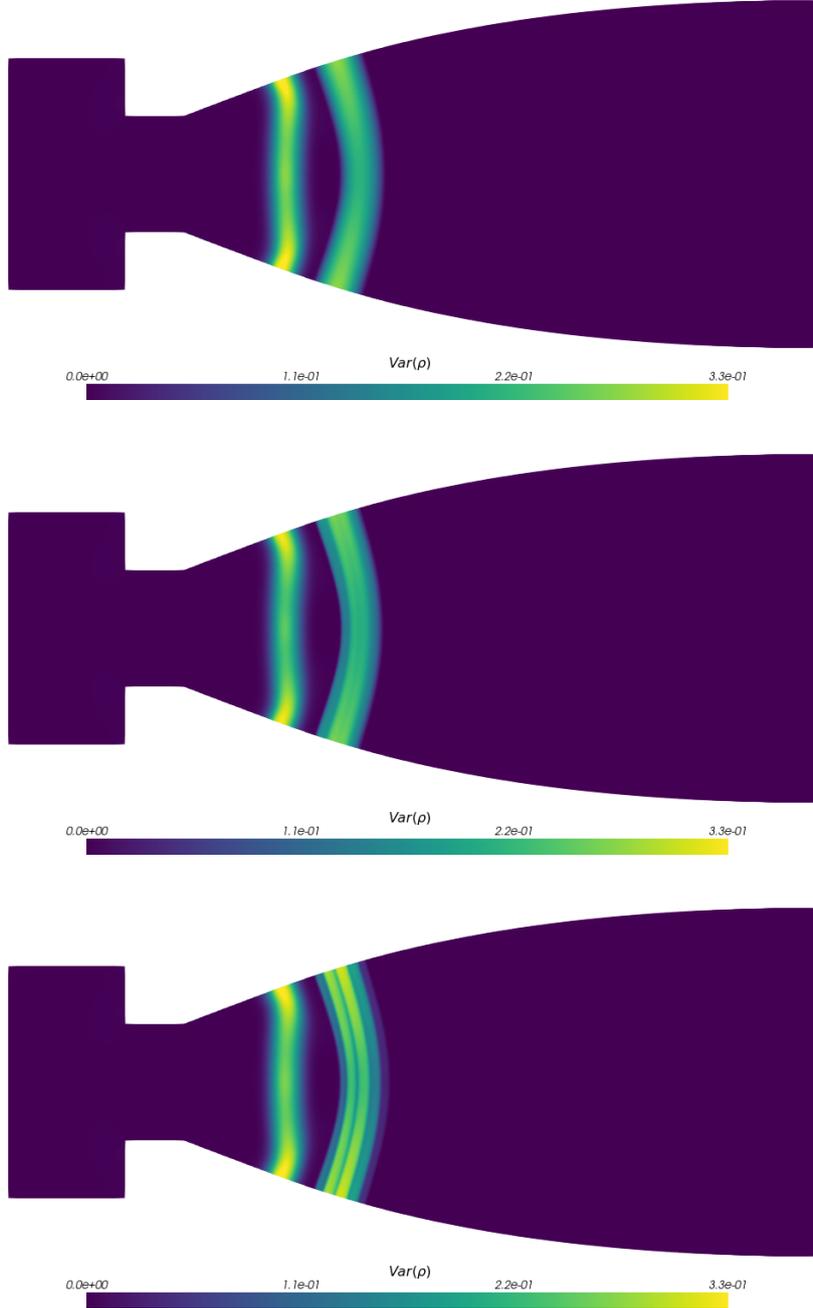


Figure 9: Variance of the density with different methods. From top to bottom: reference solution, fIPM, IPM.

In the two-dimensional setting IPM uses the entropy

$$s(\rho, \rho v_1, \rho v_2, \rho E) = -\rho \ln \left(\rho^{-\gamma} \left(\rho E - \frac{(\rho v_1)^2 + (\rho v_2)^2}{2\rho} \right) \right).$$

One can for example find the resulting IPM ansatz in [29, Appendix B]. The IPM solution, which is depicted in Figures 8 and 9 on the bottom again shows a step-like approximation of the expectation. Furthermore, it

leads to an oscillatory variance approximation at the shock position. Let us now apply the Erfc filter (13) to mitigate the spurious artifacts. Since this filter does not preserve realizability, we again need to add a regularization term with $\eta = 10^{-7}$, i.e., we choose Algorithm 2. The chosen filter strength is $\lambda = 0.9$ and we pick a filter order of $\alpha = 7$. It can be seen that the resulting approximation shows a satisfactory agreement with the reference solution. The filtered solution shows small artifacts in the shock region, however the overall approximation quality is significantly improved. The exponential filter (12) shows a similar result compared to the Erfc filter when choosing a filter strength of $\lambda = 0.045$ for the same order of $\alpha = 7$.

Now, let us study all filters used in this work, namely the exponential, Erfc and Fokker-Planck filters when changing the filter strength. In order to measure oscillations, we extend the measures δ_E and δ_{var} to two-dimensional spatial domains: If the distance of a discrete numerical density solution $\rho_\Delta := (\rho_1, \dots, \rho_{N_x})^T$ to a reference density $\rho_{\text{ex},\Delta} := (\rho_{\text{ex},1}, \dots, \rho_{\text{ex},N_x})^T$ is given by $e_\Delta := \rho_{\text{ex},\Delta} - \rho_\Delta$, we are interested in determining the error of second derivatives, which is

$$\delta_E := \sqrt{\sum_{j=1}^{N_x} \Delta x_j \left((\partial_{xx} \mathbb{E}[e_j])^2 + (\partial_{yy} \mathbb{E}[e_j])^2 \right)}$$

and

$$\delta_{\text{var}} := \sqrt{\sum_{j=1}^{N_x} \Delta x_j \left((\partial_{xx} \text{Var}[e_j])^2 + (\partial_{yy} \text{Var}[e_j])^2 \right)}.$$

With this error measure, we intend to capture the error behavior while measuring oscillations. Note that this measure only gives a coarse overview and minimal values in δ_E and δ_{var} do not always guarantee obtaining the optimal result when looking at the full expectation and variance plot. The resulting values of δ_E and δ_{var} for different filter strengths at a fixed filter order $\alpha = 7$ (when using the EXP and erfc filters) are depicted in Figure 10. Here, we added the values of δ_E and δ_{var} of the classical IPM method without filters as straight lines. Note that all filters require differently chosen filter strengths and we therefore study these filters on different grids for λ . For the order 7 filters, these grids are chosen such that the effects of filtering with each filter is plausible, i.e., the solution does not equal the standard IPM solution and the effects of uncertainty is not completely dampened out. The Fokker-Planck filter in Figure 10c is very sensitive with respect to the chosen filter strength and we can therefore observe the effect of filtering out the dependency on ξ for large filter strengths. Note that in this case, the variance has a value of zero. Hence, δ_{var} is simply the second derivative of the reference solution. Since the reference solution is smooth, this value is small compared to settings in which the filter is less dominant. However, the poor quality of the solution can be read from δ_E which heavily increases. At intermediate filter strengths, the Fokker-Planck filter does show improved results in terms of the chosen oscillation metric. It can be seen that both, the exponential filter in Figure 10a and the Erfc filter in Figure 10b mitigate oscillations compared to the classical IPM solution. However, the choice of the filtering strength affects the quality of the expected value and variance. While the optimal filter strength differs for the expected value and the variance, we observe that the minimum in the respective curves for δ_E and δ_{var} is reached at similar points. Therefore, if we choose the filter such that we obtain a good representation of the expected value, we can expect to also observe a satisfactory variance approximation.

5. Summary and Outlook

The main task of this work is to develop a method which mitigates oscillations that arise in the solutions of the IPM method. Applying a filter to the moments of IPM is challenging, since the filtered moments need to remain realizable and it can be shown that standard filters violate this property. In this work, we presented two different strategies to apply filtering to the IPM method. First, we proposed a new filter, which we called Fokker-Planck filter. It is based on an underlying Fokker-Planck equation, whose solutions obey the maximum-principle. We use this fact to guarantee realizability of filtered moment

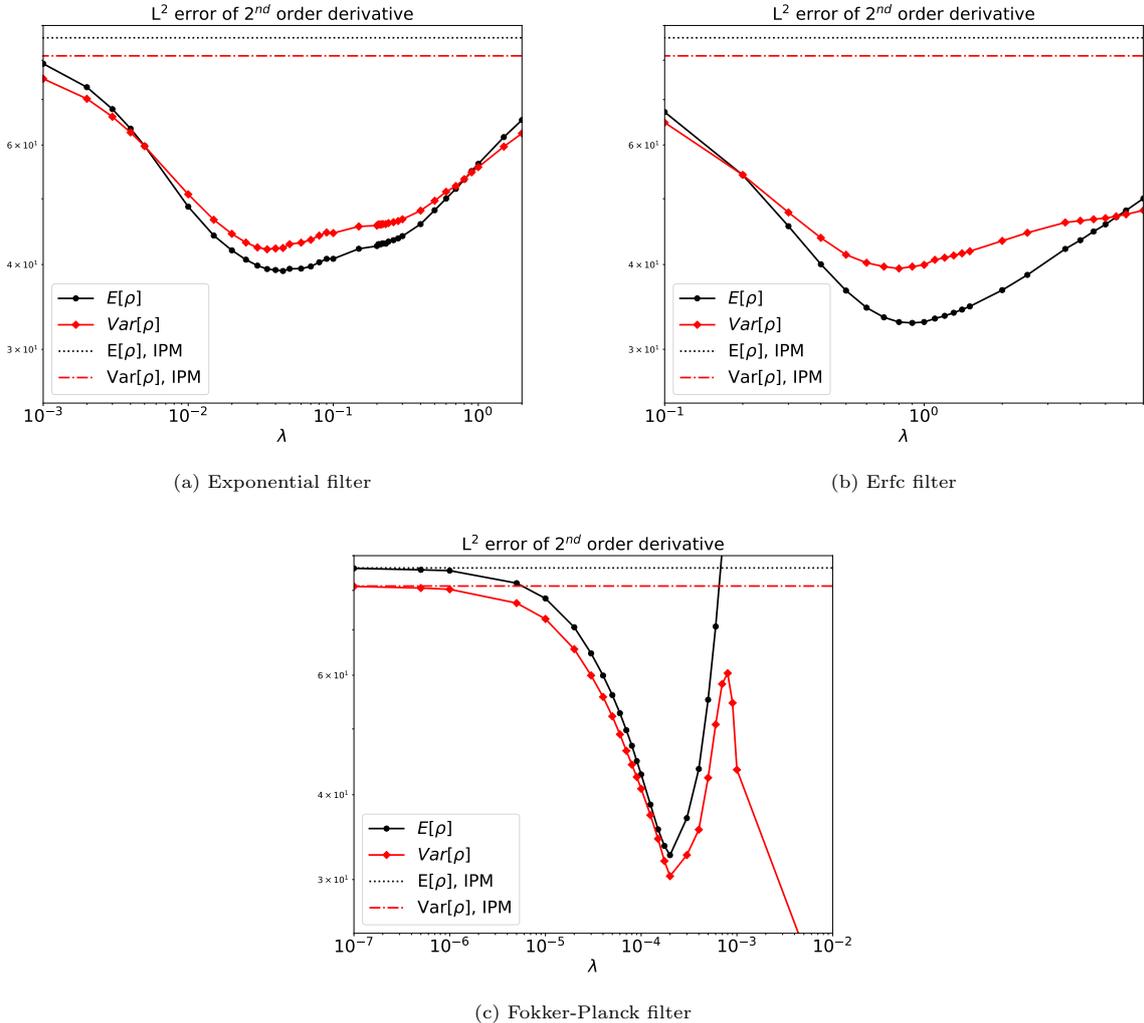


Figure 10: δ_E and δ_{Var} for different filter strengths when using the exponential, Erfc and Fokker-Planck filter. The exponential and erfc filters have a regularization strength $\eta = 10^{-7}$ and order $\alpha = 7$. Values for δ_E and δ_{Var} without filters are added as a straight line.

vectors. Furthermore, by investigating the filter on the kinetic level of the Euler equations, we are able to show realizability for applications in fluid dynamics. Second, we proposed to modify the IPM closure to be able to cope with non-realizable moments. This task is achieved by adding a regularization term to the IPM optimization problem. The regularization then allows choosing standard filters, which do not necessarily preserve realizability. We conduct numerical experiments to investigate the effectiveness of both approaches. It turns out that the second approach, i.e., using a regularization strategy yields the best results. I.e., the resulting approximations of expected value and variance show a significant reduction of artifact which resulted from spurious oscillations. Unfortunately, the Fokker-Planck filter, despite being able to improve the solution at the shock, heavily dampens the variance at the rarefaction wave.

Ideas to weaken this effect are the use of adaptive filter strengths, which yields the possibility to turn down the filter strength at the rarefaction wave and contact discontinuity. Note, that the classical IPM method without filters yields good results in these regions, meaning that one could only turn on filtering at the shock position. In this work, we focused on the exponential, Erfc and Fokker-Planck filter and one should

try out more filter functions. Note that the Lasso filter [24], which adaptively picks the filter strength in an automated fashion, unfortunately led to heavily dampened variance results. A method which is closely related to IPM is the use of Roe variables [7, 48]. Applying the proposed filtering and regularization strategies to Roe variables could potentially improve this class of methods. Furthermore, instead of regularizing the moment system, one could apply the regularization on the Euler equations as done in [33]. Such a system allows the use of non-physical quantities such as negative densities and pressures. In this case, one is able to apply the standard stochastic-Galerkin method to the regularized Euler equations without having to preserve positivity of certain physical quantities. The presented ideas are not restricted to uncertainty quantification as they can for example be used in kinetic theory and other research areas which require the construction of closures. Furthermore, one should further investigate the regularization, especially its effect on the run time. The observed run time speedup when increasing the regularization parameter holds the potential of speeding up standard IPM computations. However, one needs to keep in mind that the numerical experiments show strong solution manipulations when the regularization is chosen too big.

Acknowledgments

Funding: Jonas Kusch, Graham Alldredge and Martin Frank were supported by the German Research Foundation (DFG). Martin Frank was supported under grant FR 2841/6-1 and Graham Alldredge under AL 2030/1-1. Jonas Kusch was funded under Project-ID 258734477 - SFB 1173. Furthermore, Jonas Kusch would like to thank the German Academic Exchange Service (DAAD) for funding his research visit at the University of Notre Dame as well as the University of Notre Dame for hosting his stay during which large portions of this work were conducted.

References

- [1] Ryan G McClarren. *Uncertainty Quantification and Predictive Computational Science: A Foundation for Physical Scientists and Engineers*. Springer, 2018.
- [2] Norbert Wiener. The homogeneous chaos. *American Journal of Mathematics*, 60(4):897–936, 1938.
- [3] Dongbin Xiu and George Em Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing*, 24(2):619–644, 2002.
- [4] Jonas Kusch and Martin Frank. Intrusive methods in uncertainty quantification and their connection to kinetic theory. *International Journal of Advances in Engineering Sciences and Applied Mathematics*, pages 1–16, 2018.
- [5] Roger G Ghanem and Pol D Spanos. *Stochastic finite elements: a spectral approach*. Courier Corporation, 2003.
- [6] Gaël Poëtte, Bruno Després, and Didier Lucor. Uncertainty quantification for systems of conservation laws. *Journal of Computational Physics*, 228(7):2443–2467, 2009.
- [7] Per Pettersson, Gianluca Iaccarino, and Jan Nordström. A stochastic Galerkin method for the euler equations with roe variable transformation. *Journal of Computational Physics*, 257:481–500, 2014.
- [8] Louisa Schlachter and Florian Schneider. A hyperbolicity-preserving stochastic Galerkin approximation for uncertain hyperbolic systems of equations. *Journal of Computational Physics*, 375:80–98, 2018.
- [9] Dihan Dai, Yekaterina Epshteyn, and Akil Narayan. Hyperbolicity-preserving and well-balanced stochastic galerkin method for shallow water equations. *SIAM Journal on Scientific Computing*, 43(2):A929–A952, 2021.
- [10] OP Le Maître, OM Knio, HN Najm, and RG Ghanem. Uncertainty propagation using wiener–haar expansions. *Journal of computational Physics*, 197(1):28–57, 2004.
- [11] Per Pettersson, Gianluca Iaccarino, and Jan Nordström. Numerical analysis of the burgers’ equation in the presence of uncertainty. *Journal of Computational Physics*, 228(22):8394–8412, 2009.
- [12] Philipp Öffner, Jan Glaubitz, and Hendrik Ranocha. Stability of correction procedure via reconstruction with summation-by-parts operators for burgers’ equation using a polynomial chaos approach. *arXiv preprint arXiv:1703.03561*, 2017.
- [13] Xiaoliang Wan and George Em Karniadakis. Multi-Element Generalized Polynomial Chaos for Arbitrary Probability Measures. *SIAM J. Sci. Comput.*, 28(3):901–928, 2006.
- [14] J. Tryoen and A. Ern. Adaptive Anisotropic Spectral Stochastic Methods. 2010.
- [15] John P Boyd. *Chebyshev and Fourier spectral methods*. Courier Corporation, 2001.
- [16] Jan S Hesthaven, Sigal Gottlieb, and David Gottlieb. *Spectral methods for time-dependent problems*, volume 21. Cambridge University Press, 2007.
- [17] Ryan G. McClarren and Cory D. Hauck. Simulating radiative transfer with filtered spherical harmonics. *Physics Letters A*, 374(22):2290–2296, 2010.
- [18] Ryan G. McClarren and Cory D. Hauck. Robust and accurate filtered spherical harmonics expansions for radiative transfer. *Journal of Computational Physics*, 229(16):5597–5614, 2010.

- [19] David Radice, Ernazar Abdikamalov, Luciano Rezzolla, and Christian D Ott. A new spherical harmonics scheme for multi-dimensional radiation transport i. static matter configurations. *Journal of Computational Physics*, 242:648–669, 2013.
- [20] Vincent M. Laboure, Ryan G. McClarren, and Cory D. Hauck. Implicit filtered pn for high-energy density thermal radiation transport using discontinuous Galerkin finite elements. *Journal of Computational Physics*, 321:624–643, 2016.
- [21] Martin Frank, Cory Hauck, and Kerstin Kuepper. Convergence of filtered spherical harmonic equations for radiation transport. *Communications in Mathematical Sciences*, 14(5):1443–1465, 2016.
- [22] Gregor J Gassner and Andrea D Beck. On the accuracy of high-order discretizations for underresolved turbulence simulations. *Theoretical and Computational Fluid Dynamics*, 27(3-4):221–237, 2013.
- [23] Tomas Lundquist and Jan Nordström. Stable and accurate filtering procedures. *Journal of Scientific Computing*, 82(1):1–21, 2020.
- [24] Jonas Kusch, Ryan G. McClarren, and Martin Frank. Filtered stochastic galerkin methods for hyperbolic equations. *Journal of Computational Physics*, 403:109073, 2020.
- [25] J P Boyd. The erfc-log filter and the asymptotics of the Euler and Vandeven sequence accelerations. In *Proceedings of the Third International Conference on Spectral and High Order Methods*, pages 267–276. Houston Math. J, 1996.
- [26] Brian J Hoskins. Representation of the earth topography using spherical harmonics. *Monthly Weather Review*, 108(1):111–115, 1980.
- [27] Jonas Kusch and Louisa Schlachter. Oscillation mitigation of hyperbolicity-preserving intrusive uncertainty quantification methods for systems of conservation laws. *arXiv preprint arXiv:2008.07845*, 2020.
- [28] Jonas Kusch, Graham W Alldredge, and Martin Frank. Maximum-principle-satisfying second-order intrusive polynomial moment scheme. *SMAI-Journal of computational mathematics*, 5:23–51, 2019.
- [29] Jonas Kusch, Jannick Wolters, and Martin Frank. Intrusive acceleration strategies for uncertainty quantification for hyperbolic systems of conservation laws. *Journal of Computational Physics*, page 109698, 2020.
- [30] C. Kristopher Garrett, Cory Hauck, and Judith Hill. Optimization and large scale computation of an entropy-based moment closure. *Journal of Computational Physics*, 302:573–590, 2015.
- [31] Jonas Kusch. *Realizability-preserving discretization strategies for hyperbolic and kinetic equations with uncertainty*. PhD thesis, Karlsruhe Institute of Technology, May 2020.
- [32] Graham Alldredge, Martin Frank, Jonas Kusch, Ryan McClarren, and Jannick Wolters. *UQCreator test-cases for “A Filtered Intrusive Polynomial Moment Method”*, 2020. <https://git.scc.kit.edu/uqcreator/publication-a-filtered-ipm-method>.
- [33] Graham W. Alldredge, Martin Frank, and Cory D. Hauck. A regularized entropy-based moment method for kinetic equations. *SIAM Journal on Applied Mathematics*, 79(5):1627–1653, 2019.
- [34] Jonathan M Borwein and Adrian Stephen Lewis. Convergence of best entropy estimates. *SIAM Journal on Optimization*, 1(2):191–205, 1991.
- [35] James Alexander Shohat and Jacob David Tamarkin. *The problem of moments*. Number 1. American Mathematical Soc., 1943.
- [36] Raúl E Curto. Recursiveness, positivity and truncated moment problems. *Houston Journal of Mathematics*, 17:603–635, 1991.
- [37] C David Levermore. Moment closure hierarchies for kinetic theories. *Journal of statistical Physics*, 83(5-6):1021–1065, 1996.
- [38] A. Decarreau, D. Hilhorst, C. Lemaréchal, and J. Navaza. Dual methods in entropy maximization. Application to some problems in crystallography. *SIAM Journal on Optimization*, 2(2):173–197, 1992.
- [39] Graham W Alldredge, Cory D Hauck, and Andre L Tits. High-order entropy-based closures for linear transport in slab geometry ii: A computational study of the optimization problem. *SIAM Journal on Scientific Computing*, 34(4):B361–B391, 2012.
- [40] S Deshpande. Kinetic theory based new upwind methods for inviscid compressible flows. In *24th Aerospace Sciences Meeting*, page 275, 1986.
- [41] Amiram Harten, Peter D Lax, and Bram van Leer. On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM review*, 25(1):35–61, 1983.
- [42] Benoît Perthame. Boltzmann type schemes for gas dynamics and the entropy property. *SIAM Journal on Numerical Analysis*, 27(6):1405–1421, 1990.
- [43] Benoît Perthame. Second-order boltzmann schemes for compressible euler equations in one and two space dimensions. *SIAM Journal on Numerical Analysis*, 29(1):1–19, 1992.
- [44] Jonas Kusch and Martin Frank. An adaptive quadrature-based moment closure. *International Journal of Advances in Engineering Sciences and Applied Mathematics*, 11(3):174–186, 2019.
- [45] Tianbai Xiao and Martin Frank. A stochastic kinetic scheme for multi-scale plasma transport with uncertainty quantification. *arXiv preprint arXiv:2006.03477*, 2020.
- [46] Randall J LeVeque and Randall J Leveque. *Numerical methods for conservation laws*, volume 3. Springer, 1992.
- [47] Dongbin Xiu and Jan S Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM Journal on Scientific Computing*, 27(3):1118–1139, 2005.
- [48] Stephan Gerster and Michael Herty. Entropies and symmetrization of hyperbolic stochastic Galerkin formulations. *Comm. Computat. Phys.*, to appear, 2020.
- [49] Joseph M Powers and Mihir Sen. *Mathematical Methods in Engineering*. Cambridge University Press, 2015.

Appendix A. Sturm-Liouville equation of orthogonal polynomials

In the following, we provide a short summary of the Sturm-Liouville equation, following [49, Chapter 5]. The standard orthonormal polynomials $P_i : \Theta \rightarrow \mathbb{R}$ fulfill a Sturm-Liouville equation. This equation takes the form $\mathcal{L}P_i = \mu_i P_i$, where \mathcal{L} is a self-adjoint linear operator

$$\mathcal{L}P_i = \frac{1}{r(x)} \frac{d}{d\xi} \left(p(x) \frac{d}{d\xi} + q(x) \right).$$

I.e., the polynomials P_i are eigenfunctions to the operator \mathcal{L} with eigenvalues μ_i . Furthermore, we have that

$$\int_{\Theta} P_i(\xi) P_j(\xi) r(\xi) d\xi = \gamma_i \delta_{ij},$$

where γ_i denotes the weighted L^2 -norm. The *Legendre polynomials* which are used throughout this work fulfill this property, where

$$p(\xi) = 1 - \xi^2, \quad r(\xi) = 1, \quad q(\xi) = 0, \quad \mu_i = -i(i+1).$$

Further polynomials are

- *Chebyshev*: $p(\xi) = \sqrt{1 - \xi^2}$, $r(\xi) = \frac{1}{\sqrt{1 - \xi^2}}$, $q(\xi) = 0$, $\mu_i = -i^2$,
- *Hermite*: $p(\xi) = e^{-\xi^2}$, $r(\xi) = e^{-\xi^2}$, $q(\xi) = 0$, $\mu_i = -2i$,
- *Laguerre*: $p(\xi) = \xi e^{-\xi^2}$, $r(\xi) = e^{-\xi^2}$, $q(\xi) = 0$, $\mu_i = -i$.