

Legendre Fluids: A Unified Framework for Analytic Reduced Space Modeling and Rendering of Participating Media

Mohit Gupta and Srinivasa G. Narasimhan

Robotics Institute, Carnegie Mellon University

Abstract

In this paper, we present a unified framework for reduced space modeling and rendering of dynamic and non-homogenous participating media, like snow, smoke, dust and fog. The key idea is to represent the 3D spatial variation of the density, velocity and intensity fields of the media using the same analytic basis. In many situations, natural effects such as mist, outdoor smoke and dust are smooth (low frequency) phenomena, and can be compactly represented by a small number of coefficients of a Legendre polynomial basis. We derive analytic expressions for the derivative and integral operators in the Legendre coefficient space, as well as the triple product integrals of Legendre polynomials. These mathematical results allow us to solve both the Navier-Stokes equations for fluid flow and light transport equations for single scattering efficiently in the reduced Legendre space. Since our technique does not depend on volume grid resolution, we can achieve computational speedups as compared to spatial domain methods while having low memory and pre-computation requirements as compared to data-driven approaches. Also, analytic definition of derivatives and integral operators in the Legendre domain avoids the approximation errors inherent in spatial domain finite difference methods. We demonstrate many interesting visual effects resulting from particles immersed in fluids as well as volumetric scattering in non-homogenous and dynamic participating media, such as fog and mist.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction and Related Work

Participating media such as smoke, snow, dust, mist and fog exhibit a wide range of visual effects. Such media are characterized by their density, velocity and intensity fields that vary across both space and time. Accurate modeling of densities and velocities as well as rendering of intensities of these media is critical for achieving photo-realism in computer graphics. Also, many applications like games require interactive changes in lighting, view-point and the medium properties. For such applications, it is imperative to achieve these visual effects in real-time.

The first step in realizing this goal of visual realism requires modeling the time-varying density and velocity fields of participating media. The Navier-Stokes equations for incompressible fluid flow [CM90] provide a differential model for simulating the density and velocity fields. Explicit analytic solutions to Navier-Stokes equations are hard to obtain and hence, a number of works that employ numerical finite difference methods (FDM) have been proposed [FM96, FM97, Sta99, Sta01, FF01, FSJ01, NFJ02, SRF05]. Although simple to implement, such schemes require high spatial resolution to minimize the finite differencing numerical errors, plac-

ing serious demands on memory and compromising speed. Treuille et al [TMPS03] develop an approach for key-framing of fluid flows that alleviate the discretization errors. However, their approach becomes computationally prohibitive for large grid sizes. More recently, an interesting data-driven approach has been taken to simulate the velocity fields using a reduced dimensional PCA basis [TLP06]. This approach achieves considerable speed-ups and produces impressive results, but at the cost of high memory requirements and lengthy pre-computation. Furthermore, as the authors mention, it is unclear whether the approach generalizes to new fluid flows that are not represented in the pre-computation.

The second step towards the goal of creating the desired visual effects is rendering of participating media, which requires modeling the intensity fields resulting from volumetric scattering. Using the computed density field, the corresponding intensity field of the participating medium is then rendered by solving the light transport equation [Cha60]. Analogous to fluid modeling, many works that numerically solve the light transport equation based on FDMs have been proposed [KH84, Max94, Jen01, PM93, EP90, Sak90, LBC94, RT87]. As such, many of the issues related to numerical errors must be addressed here as well. While these methods



Figure 1: Legendre domain 3D fluid simulation and rendering: In this example, we have 3000 snow flakes being carried by a wind field (Legendre domain fluid simulation). We add mist to the scene using Legendre domain rendering for participating media. Notice further objects appearing brighter due to the air-light effect, and distant snow-flakes becoming invisible as the mist density is increased. The clear weather Christmas image was downloaded from www.survivinggrady.com/2005_12_01_archive.html

can produce impressive visual effects, they are too slow for interactive applications. Recent hardware-accelerated techniques [DYN02, REK*04, HL01] can significantly decrease the running times of numerical simulations, although they are specialized to particular phenomena.

In addition, note that the intensity fields depend on the illumination and viewing geometry as well as the scattering properties [NGD*06, HED05] of the participating medium. Moreover, the lighting, viewpoint and the densities of the medium may change with time. Thus, the pre-computations required are too prohibitive for data driven approaches to be applied to intensity fields. For the special case of homogeneous media, many previous analytic approaches [Max86, SRNN05, JMLH01, NN03] may be used to render the effects of scattering in real-time. However, homogeneous media are not the focus of our work.

The goal of this paper is fast modeling and rendering of dynamic and non-homogenous participating media, like smoke, dust and fog. The key idea is to represent the 3D spatial variation of the density, velocity and intensity fields using the same analytic basis. Jos Stam [Sta99] used Fourier basis to solve the diffusion and projection steps of the fluid simulation pipe-line over a domain with periodic boundary conditions. In this work, we use Legendre Polynomials [Cha60] as our basis functions. In many situations, natural effects such as mist, outdoor smoke and dust are smooth (low frequency) phenomena, and can be compactly represented by a small number of coefficients of a Legendre polynomial basis. In this work, we will focus on optically thin media where single scattering is the dominant form of light transport [SRNN05, NGD*06]. Under these conditions, the common Legendre polynomial basis for different fields allows us to analytically solve both the Navier-Stokes and light transport equations in the reduced Legendre space. It turns out that this solution requires us to analyze triple product integrals of Legendre polynomials and their sparsity [GN07], similar in spirit to the triple product wavelet integrals for relighting [NRH04].

Since all the fields are represented using Legendre polynomials, their derivatives (and integrals) can be computed analytically, thereby avoiding the numerical errors resulting from spatial finite differences approximation. Additionally, the compactness of the Legendre domain representations of natural effects makes modeling and rendering of such partic-

ipating media very fast. Depending on the number of coefficients required, we achieve computational speedups of one to three orders of magnitude as compared to spatial domain techniques. At the same time, only a few coefficients must be stored in memory as compared to the full 3D volumes that must be stored for the data-driven (eg., PCA) approaches.

The main contribution of this paper is a theoretical one: a unified framework for both fluid simulation and rendering in an analytic reduced space. We believe that this is an important first step towards bridging the gap between model reduction for fluid simulation and pre-computed radiance transfer for rendering. In addition, we believe that the mathematical results derived in the paper are general enough to find use in many computer graphics applications. We demonstrate several visual effects resulting from volumetric scattering in time-varying participating media, such as shadowgrams that are cast by the medium on a background, mixing of different gaseous media and airlight effects due to depth disparities in the scene. We also show fluid simulation results illustrating snow flakes (see Figure 1) and confetti immersed in turbulent wind fields, as well as smoke density fields evolving under the influence of user-defined forces.

For the purpose of deriving the unified reduced space simulation and rendering framework, we have made several limiting assumptions such as smooth (low frequency) phenomena, no objects within the medium, single scattering, orthographic viewing and distant lighting. In Section 8, we discuss these limitations in detail, along with future research directions for extending the technique to more general settings.

2. Physical Models for Participating Media

Dynamic and non-homogenous participating media can be characterized by density, velocity and intensity fields, that vary across both space and time. Whereas Navier-Stokes equations for incompressible fluid flow model the evolution of the density and velocity fields over time, the intensity fields are rendered using light transport equations. The time evolution of the velocity field \mathbf{u} is given by [CM90]:

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla) \mathbf{u} - \nu \nabla^2 \mathbf{u} + \nabla p + \mathbf{b}, \quad \text{s.t. } \nabla \cdot \mathbf{u} = 0, \quad (1)$$

where, ν is the kinematic viscosity, p is the pressure field and \mathbf{b} denotes the external forces (the notation used in the paper is

\mathbf{u}	velocity field
u_{\diamond}	\diamond – component of velocity field
r	density field
\mathbf{b}	external force field
b_{\diamond}	\diamond – component of force field
E^d	direct transmission intensity field
E^s	scattered intensity field
ν	kinematic viscosity
σ	extinction coefficient
β	scattering coefficient
θ	scattering angle
$\Omega(\theta)$	scattering phase function
ω^d	lighting direction
ω^s	viewing direction

Figure 2: Notation used in our paper. \diamond stands for either x , y or z .




Figure 3: The participating medium is illuminated by a distant light source and is viewed by an orthographic camera. Under the single scattering assumption, the intensity field within the medium volume can be split into two sets of light rays: the pre-scattering (direct transmission) intensity field $E^d(x,t)$ and post-scattering intensity field $E^s(x,t)$ (shown using red rays).

given in Figure 2). Following [Sta99, CM90], Equation 1 can be written as:

$$\frac{\partial \mathbf{u}}{\partial t} = \underbrace{P}_{\text{projection}} \left(\underbrace{-(\mathbf{u} \cdot \nabla) \mathbf{u}}_{\text{advection}} + \underbrace{\nu \nabla^2 \mathbf{u}}_{\text{diffusion}} + \underbrace{\mathbf{b}}_{\text{forces}} \right) \quad (2)$$

Here, P is a linear operator which projects a vector field to its divergence free component. Equation 2 can be resolved by splitting the right hand side into four sequential steps: (i) advection, (ii) diffusion, (iii) external forces and (iv) projection [Sta99]. Similarly, the time evolution of the density field r is given by:

$$\frac{\partial r}{\partial t} = \underbrace{-(\mathbf{u} \cdot \nabla) r}_{\text{advection}} - \underbrace{\kappa \nabla^2 r}_{\text{diffusion}} + \underbrace{-\alpha r}_{\text{dissipation}} + \underbrace{S_r}_{\text{source}}, \quad (3)$$

where, κ is the diffusion constant, α is the dissipation rate and S_r is the source term for density.

Using the density field r , we can render the intensity fields for any configuration of illumination and viewing geometry. In this work, we consider optically thin media where **single scattering** is the dominant form of light transport. Figure 3 shows an orthographic camera viewing a participating medium that is illuminated by a distant light source. Then, we can split the intensity fields into two components: the *pre-*

Orthogonality	$\int_{-1}^1 L_i(x) L_j(x) dx = \delta_{ij}$
Derivative	$L'_i(x) = \sum_k c_{ik} L_k(x)$
Integral	$\int L_i(x) dx = \sum_k b_{ik} L_k(x)$

Figure 4: Properties of Legendre Polynomials [Cha60].

scattering (direct transmission) intensity field $E^d(\mathbf{x}, t)$, and the *post-scattering* intensity field $E^s(\mathbf{x}, t)$. Mathematically, these intensity fields can be written as [Cha60]:

$$(\omega^d \cdot \nabla) E^d = -\sigma r \cdot E^d \quad (4)$$

$$(\omega^s \cdot \nabla) E^s = -\sigma r \cdot E^s + \beta r \cdot \Omega(\theta) \cdot E^d \quad (5)$$

Here, σ and β are the extinction and scattering coefficients respectively and $\Omega(\theta)$ is the phase function. When the camera is outside the medium, the acquired **2D image** of the medium is simply the boundary of the 3D intensity field E^s (Figure 3).

3. Compact Analytic Representation of Non-Homogenous Media

The key idea in this paper is to represent the 3D spatial variation of the density, velocity and intensity fields using the same analytic basis. We choose to use Legendre polynomials as basis functions. In many situations, natural effects such as mist, outdoor smoke and dust are smooth (low frequency) phenomena, and can be compactly represented by a small number of coefficients. Legendre polynomials are orthogonal, have global support (non-zero over the entire domain), and have analytic derivatives and integrals (Figure 4). As a result, they find wide application in mathematical physics literature in conjunction to solving differential equations [Cha60].

A function $f(x)$ can be represented as a linear combination of Legendre polynomials L_k of different orders $f(x) = \sum_k F_k L_k(x)$, where the Legendre domain coefficients $[F_k]$ can be computed analytically as:

$$F_k = \int_{-1}^1 f(x) L_k(x) dx. \quad (6)$$

In 3D, we represent a field $f(x, y, z)$ that is smooth in x -, y - and z -directions as:

$$f(x, y, z) = \sum_{ijk} F_{ijk} L_i(z) L_j(y) L_k(x). \quad (7)$$

For notational ease, Equation 7 is written as $f(x, y, z) \Leftrightarrow [F_{ijk}]$. The Legendre representations for the various fields are given in Figure 5.

4. Analytic Operators in Legendre Space

In this section, we derive the legendre space formulations for various operators and establish that they are compact, computationally efficient, and completely analytic in nature. For ease of exposition, we illustrate the concepts with 1D examples; analysis in 2D and 3D follows in an exactly similar manner.

4.1. Derivative Operator

Observe that spatial derivatives appear both in the Navier-Stokes and the light transport equations (2, 3, 4, 5) in the form of gradient and Laplacian operators. Using the property that derivative of a legendre polynomial can be expressed in terms of lower order legendre polynomials (Figure 4), we derive the

Field	Spatial \Leftrightarrow Legendre
Density Field	$r \Leftrightarrow [R]$
Velocity Field	$u_\diamond \Leftrightarrow [U_\diamond]$
Divergence free Velocity Field	$\hat{u}_\diamond \Leftrightarrow [\hat{U}_\diamond(t)]$
External Force Field	$b_+ \Leftrightarrow [B_+(t)]$
Direct Transmission Intensity Field	$E^d \Leftrightarrow [I^d]$
Scattered Intensity Field	$E^s \Leftrightarrow [I^s]$

Figure 5: Legendre representations of various fields, where \diamond stands for x, y or z . In Figures 5 and 6, sub-scripts and arguments have been dropped for brevity. For example, d and $[D]$ should be read as $d(x, y, z, t)$ and $[D_{ijk}(t)]$ respectively.

Operation	Operand	Result	Complexity
Derivative	$g \Leftrightarrow [G]$	$\frac{\partial}{\partial \diamond} g \Leftrightarrow D_\diamond \cdot [G]$	$O(K^2)$
Integral	$g \Leftrightarrow [G]$	$\int g d\diamond \Leftrightarrow \hat{I}_\diamond \cdot [G]$	$O(K^2)$
Product	$\begin{matrix} g \Leftrightarrow [G] \\ h \Leftrightarrow [H] \end{matrix}$	$g \cdot h \Leftrightarrow M^G \cdot [H]$	$O(K^3)$
Truncation	$[G]$	$[G^T] = T \cdot [G]$	$O(K^2)$
Legendre to Spatial	$[G]$	g	$O(NK)$

Figure 6: Legendre Space Operators (\diamond stands for x, y or z). N is the size of the spatial grid. K is the size of legendre coefficient representation.

derivative operator in legendre domain, which is **completely analytic**, and hence, devoid of the numerical errors resulting from the Finite Difference approximation:

$$\begin{aligned}
 f(x) &= \sum_i F_i L_i(x) \quad \Rightarrow \quad f'(x) = \sum_i F_i L'_i(x) \quad (8) \\
 \Rightarrow f'(x) &= \sum_k \left(\sum_i F_i * c_{ik} \right) L_k(x) \quad \dots \quad (\text{Figure 4}) \\
 &= \sum_k F'(k) L_k(x)
 \end{aligned}$$

where $F'(k) = \sum_i F_i * c_{ik}$. We can write this equation in matrix form, with $[F'_k]$ and $[F_i]$ as the coefficient vectors corresponding to the derivative and the original function respectively. The *derivative* operator (x -direction) in Legendre Domain is thus given by the matrix $D_x(i, k) = c_{ik}$:

$$[F'_k] = D_x * [F_i] \quad (9)$$

Derivatives in y and z and the integral operator can be defined likewise. Figure 6 lists all the legendre space operators that we derive, along with the corresponding time complexity. Given K as the size of legendre space representation $[F_i]$, the matrix-vector multiplication require $O(K^2)$ computations. Building the derivative and integral matrices is a one time operation, and takes $O(K^2)$ time.

4.2. Product Operator in Legendre Domain

The advection term in the Navier-Stokes equation (2, 3) as well as the single scattering equation for rendering (4, 5) entail multiplication of two fields to compute a third one. This motivates investigating the general problem of multiplying two functions, $h(x) = f(x) \cdot g(x)$, where both the functions and the result are represented in the Legendre Basis:

$$f(x) = \sum_j F_j L_j(x) \quad g(x) = \sum_k G_k L_k(x) \quad h(x) = \sum_i H_i L_i(x)$$

To compute the i^{th} basis coefficient for the result, we use orthogonality of Legendre Polynomials (see Figure 4)

$$\begin{aligned}
 H_i &= \int_{-1}^1 L_i(x) h(x) dx = \int_{-1}^1 L_i(x) f(x) g(x) dx \\
 &= \int_{-1}^1 L_i(x) \left(\sum_j F_j L_j(x) \right) \left(\sum_k G_k L_k(x) \right) dx \\
 &= \sum_{jk} F_j G_k T I_{ijk}
 \end{aligned}$$

where $T I_{ijk} = \int_{-1}^1 L_i(x) L_j(x) L_k(x) dx$ is the Legendre Polynomial triple product integral, and can be pre-computed a priori. As with the derivative and integral case, we can write the above equation in matrix form as follows:

$$[H_i] = M^G * [F_j] = M^F * [G_k] \quad (10)$$

where, $M^G(i, j) = \sum_k G_k T I_{ijk}$ and $M^F(i, k) = \sum_j F_j T I_{ijk}$. Given the size of legendre representations as K , the multiplication matrix has $O(K^2)$ entries. For each entry, $O(K)$ computations are required. Thus, we need $O(K^3)$ computations to build the multiplication matrix and $O(K^2)$ time for the matrix-vector multiplication. Therefore, total time complexity of legendre space multiplication is $O(K^3)$. However, we show that the 3D tensor TI is sparse using the **Legendre Polynomials Triple Product Integrals theorem** [GN07]. Using the theorem, we show that approximately $\frac{3}{4}$ of the entries of the TI tensor are exactly zero. We exploit this sparsity to achieve computational speed-ups in the advection and the rendering stages. Indeed, the time required to construct the multiplication matrix can be reduced by a factor of 4 in 1D and by $4^3 = 64$ in the 3D case.

Lower Order Approximation: Note that multiplying two polynomials of degree K each results in a polynomial of degree $2K$. Therefore, given two functions, each with Legendre representation of size K , the Legendre representation of the product will have size $2K$. For computational savings, it is desirable to keep the size of the Legendre representation constant. To this end, we devise a simple approximation scheme using the **Chebyshev Polynomials** to truncate a given Legendre representation from $2K$ terms to K terms, while keeping the approximation error low under the L_∞ norm [GN07]. We define the **Truncation Matrix Operator T in legendre space**, such that

$$\underbrace{[F'_i]^T}_{K \times 1} = \underbrace{T}_{K \times 2K} * \underbrace{[F_i]}_{2K \times 1}$$

where $[F_i]$ is the legendre representation of size $2K$, and $[F'_i]^T$ is the corresponding truncated representation of size K . As with derivatives and integrals, truncation requires a matrix multiplication with a time complexity of $O(K^2)$. Building the truncation matrix T is a one time operation requiring $O(K^2)$ operations.

5. Modeling in Legendre Domain

Using the Legendre representations for fields and the operators (derivative, multiplication, truncation), we solve the Navier-Stokes equations (2, 3) in the Legendre domain. For velocity simulation, we decompose Equation 2 into the 4 sequential steps of advection, diffusion, external forces and projection [Sta99]. Now we show how each of these steps can be simulated in the Legendre domain:

5.1. Advection

In the spatial domain, the conservation form of the advection equation is given by:

$$\frac{\partial}{\partial t} u_{\diamond} = -\nabla \cdot (\mathbf{u} u_{\diamond}) \quad (11)$$

$$= -\left(\frac{\partial}{\partial x} u_x u_{\diamond} + \frac{\partial}{\partial y} u_y u_{\diamond} + \frac{\partial}{\partial z} u_z u_{\diamond} \right) \quad (12)$$

Subscript \diamond denotes either x, y or z direction. This form implicitly assumes a divergence free velocity field, i.e. $\nabla \cdot \mathbf{u} = 0$.

Legendre space advection equation is then derived by substituting the legendre representations of the fields (Figure 5), along with the legendre space derivative and multiplication operators (Figure 6) in Equation 12:

$$\frac{\partial}{\partial t} [U_{\diamond}(t)] = -A \cdot [U_{\diamond}(t)] \quad (13)$$

$$\text{where } A = \underbrace{T}_{\text{truncation}} \cdot \left(\sum_{\diamond} \underbrace{D_{\diamond}}_{\text{derivative}} \cdot \underbrace{M^{U_{\diamond s}}}_{\text{multiplication}} \right)$$

$\sum_{\diamond}(\cdot)_{\diamond}$ is short-hand for $(\cdot)_x + (\cdot)_y + (\cdot)_z$. For example, $\sum_{\diamond} D_{\diamond} \cdot M^{U_{\diamond}}$ is expanded as $D_x \cdot M^{U_x} + D_y \cdot M^{U_y} + D_z \cdot M^{U_z}$. We update the legendre representations of the velocity field by computing the eigen decomposition of $A = A_E \cdot \Lambda \cdot A_E^{-1}$ [TLPO6]:

$$[U_{\diamond}(t + \Delta t)] = \left(A_E \cdot e^{\Delta t \cdot \Lambda} \cdot A_E^{-1} \right) \cdot [U_{\diamond}(t)] \quad (14)$$

A similar approach can be used to update the density field as well. Since it uses the multiplication operator, the time complexity of Legendre advection is $O(K^3)$ (Section 4), where K is the number of coefficients. In addition to the computational speed-up, using the completely analytic Legendre domain derivative operator reduces the numerical dissipation inherent in the FDM based approximations of the derivative operator (Figure 7).




Figure 7: Comparison between Legendre and Spatial domain advection (high intensities signify higher values of the field). Notice, that the field after advection in the spatial domain (c) has lower energy than the field resulting from analytic legendre domain advection (b). Spatial advection results in dissipation of energy due to discretization of the gradient operators. The grid size used for spatial advection was 500^2 , while 144 coefficients were used for legendre advection.

5.2. Diffusion

For the diffusion step, we solve the implicit form of diffusion equation:

$$\left(I_{N \times N} - \nu \Delta t \nabla^2 \right) u_{\diamond}(\mathbf{x}, t + \Delta t) = u_{\diamond}(\mathbf{x}, t) \quad (15)$$

where N is the total number of simulation grid voxels. The implicit form of diffusion equation is more stable than the explicit form. However, one drawback of the implicit form is that it requires solving a large system of linear equations. Fortunately, in our case, this issue is addressed by solving the diffusion equation in the reduced legendre space. Once again, we use the legendre representation of the fields and the operators (Figures 5 and 6) to obtain the **legendre space diffusion equation**:

$$\left(I_{K \times K} - \nu \Delta t D^2 \right) [U_{\diamond}(t + \Delta t)] = [U_{\diamond}(t)] \quad (16)$$

where, $D^2 = (D_x)^2 + (D_y)^2 + (D_z)^2$ is the legendre space **Laplacian** operator. Since we solve a $K \times K$ linear system, the time complexity of legendre space diffusion is $O(K^3)$. This is a considerable speed-up over solving the $(N \times N)$ system in spatial domain.

5.3. External Forces

External forces are handled by adding their legendre representation (Figure 5) to that of the velocity field:

$$[U_{\diamond}(t + \Delta t)] = [U_{\diamond}(t)] + [B_{\diamond}] \cdot \Delta t \quad (17)$$

5.4. Projection

This step ensures that the velocity field is divergence free, which is required to satisfy mass-conservation. For the projection step, we use the implicit definition of the projection operator P :

$$\nabla^2 q = \nabla \cdot \mathbf{u} \quad \hat{\mathbf{u}} = P\mathbf{u} = \mathbf{u} - \nabla q \quad (18)$$

This step requires solving the following Poisson system of equation for the scalar field q : $\nabla^2 q = \nabla \cdot \mathbf{u}$. $\hat{\mathbf{u}}$, the divergence free component of \mathbf{u} ($\nabla \cdot \hat{\mathbf{u}} = 0$), is then computed by subtracting the gradient of q from \mathbf{u} . The Poisson equation can be formulated as a linear system of equations by discretizing the ∇^2 operator in the spatial domain. Analogously, we can define P_L , the projection operator in the legendre space implicitly as follows:

$$D^2 \cdot [Q] = \sum_{\diamond} D_{\diamond} \cdot [U_{\diamond}(t)] \quad (19)$$

$$[\hat{U}_{\diamond}(t)] = P_L \cdot [U_{\diamond}(t)] = [U_{\diamond}(t)] - D_{\diamond} \cdot [Q] \quad (20)$$

Hence, in legendre space projection step, we need to solve the linear system of equations in the unknown vector $[Q]$ (Equation 19), requiring $O(K^3)$ time. As with diffusion, this is a considerable speed-up over solving the $(N \times N)$ linear system in spatial domain. **As an additional advantage**, using the analytic definitions of the derivative operators in all the simulation steps alleviates the numerical errors resulting from spatial finite difference approximations.

5.5. Density Dissipation

For density simulation, Equation 3 is solved in the legendre space. The advection, diffusion and source terms are handled in a way similar to velocity simulation. The dissipation term is then solved in the legendre space as follows:

$$(1 + \Delta t \alpha) \cdot [R(t + \Delta t)] = [R(t)] \quad (21)$$

Size of the Legendre representation: Figure 8 illustrates the time-evolution of 2D density and velocity fields for different sizes of Legendre representations. We start with the same low frequency density and velocity fields and apply the same forces throughout the 3 different Legendre domain simulations. We can observe that more coefficients allow for higher frequencies and vorticities in the density and velocity fields evolve. In Figures 9, 10 and 11, we also provide theoretical and empirical computational complexity of our framework as a function of the size of the Legendre representation (K). A user can use these as a guide for choosing the Legendre representation size that best addresses the demands (speed/ high frequency detail) of a particular application.

6. Rendering in Legendre Space

Rendering requires solving the light transport equations (4,5) in the Legendre domain using techniques similar to those used for the Navier-Stokes equations.

6.1. Direct Transmission intensity field

As earlier, substituting Legendre representations of various fields and Legendre operators (Figures 5 and 6) into Equation 4, we get:

$$\left(\sum_{\diamond} \omega_{\diamond}^d D_{\diamond} \right) \cdot [I^d] = -\sigma T \cdot M^R \cdot [I^d] \\ \Rightarrow L_{\omega^d} \cdot [I^d(t)] = \mathbf{0} \quad (22)$$

$$\text{where } L_{\omega^d} = \left(\sum_{\diamond} \omega_{\diamond}^d D_{\diamond} + \sigma T \cdot M^R \right).$$

6.2. Scattered intensity field

Similarly, we can project Equation 5 into the Legendre domain:

$$\left(\sum_{\diamond} \omega_{\diamond}^s D_{\diamond} \right) \cdot [I^s] = -\sigma T \cdot M^R \cdot [I^s] + \beta \Omega(\theta) \cdot T \cdot M^R \cdot [I^d] \\ \Rightarrow L_{\omega^s} \cdot [I^s] = \beta \Omega(\theta) \cdot T \cdot M^R \cdot [I^d] \quad (23)$$

$$\text{where } L_{\omega^s} = \left(\sum_{\diamond} \omega_{\diamond}^s D_{\diamond} + \sigma T \cdot M^R \right).$$

In the legendre space, both the light scattering equations are thus formulated as linear systems of equations in the unknowns $[I^d]$ (22) and $[I^s]$ (23). Along with the **boundary conditions**, which can be formulated as additional linear constraints, these systems can be solved in $O(K^3)$ time.

Imagine a camera observing the medium from the outside

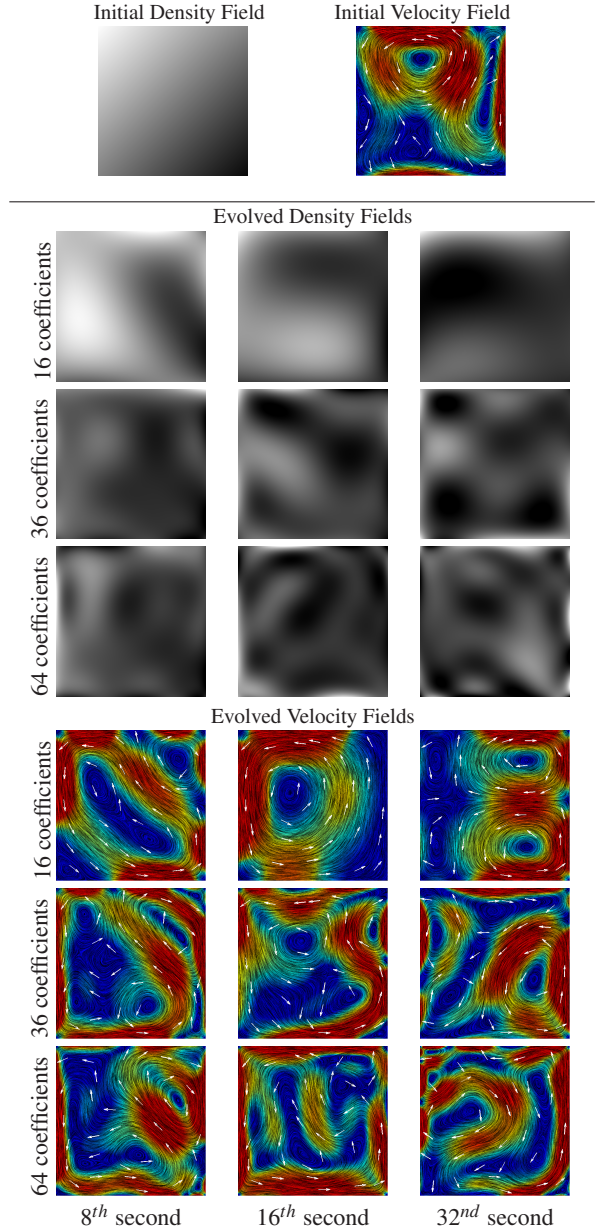


Figure 8: 2D Legendre domain Simulation results: Evolution of density and velocity for different number of Legendre coefficients. More coefficients allow higher frequencies and vorticities in the density and velocity fields.

(Figure 3). Then, the **image** recorded is given by the scattering intensity field E_s at the domain boundary:

$$E_s(x, y, z, t) = \sum_{ijk} I_{ijk}^s L_i(x) L_j(y) L_k(z) \quad (24)$$

If the image resolution is S , then time-complexity of image computation is $O(SK)$. Note that the image computation step is output-sensitive, and can easily be parallelized. Our Legendre domain modeling and rendering framework is summarized in Figure 9.

For every time step:	
• Update Velocity and Density Fields	
Advection (13) and Diffusion (16)	$O(K^3)$
Forces/Source (17)	$O(K^2)$
Projection(19, 20)	$O(K^3)$
• Update Intensity Fields	
Direct Transmission(22), Scattered (23)	$O(K^3)$
• Compute Image (24)	$O(SK)$

Figure 9: Legendre domain Rendering algorithm: K is the size of Legendre space representations and S is the image resolution.

		Number of Legendre Coefficients			
		16	36	64	144
Spatial Grid Size	200^2	500X	250X	75X	10X
	300^2	1250X	625X	187X	25X
	400^2	2500X	1250X	375X	50X
	500^2	5000X	2500X	750X	100X

Figure 10: Typical computational speed-ups for 2D simulation and rendering in Legendre domain as compared to the spatial domain.

7. Results

Our results show that Legendre polynomials can express a variety of interesting density and force distributions compactly, thereby letting the user manipulate the densities, velocities and forces globally to produce the desired effects.

Particles immersed in dynamic fluid media: Figure 12 and Figure 1 show simulations of 500 pieces of confetti and 3000 snow-flakes respectively being carried by a wind field simulated using 216 Legendre coefficients each. We can notice vorticities being created in the confetti example due to the turbulent behavior of the wind field. On the other hand, the snow flakes are carried by a more gentle, *breeze-like* wind. We encourage the reader to view the animation results in the supplementary video.

Simulation of smoke and advection of scattering albedos:

Figure 14 shows a vertically upwards axial impulse applied to a *vase shaped* smoke density field. Since the impulse is applied for a short duration, the density field dissolves towards the end of the simulation. For the first time, we also show advection of the optical properties of the medium (scattering albedos), in addition to the physical properties (densities and velocities), resulting in completely new colors and appearances as the medium evolves under external forces.

Single Scattering based rendering of participating media:

We demonstrate the visual effects of both relighting the medium under the single scattering model, and varying the viewpoint and scattering albedos, as the medium evolves under user defined forces (supplementary video). We also show interesting effects of shadowgrams that are cast by the medium on a background plane (Figures 13 and 14).

3D Visual effects resulting from volumetric scattering in non-homogenous and dynamic participating media:

In the examples of Figure 15 and Figure 1, we add non-homogenous mist to scenes with large depth variation. Notice how distant objects appear brighter due to the airlight [Kos24] effect. Reproducing such effects accurately, particularly for non-

		Number of Legendre Coefficients		
		64	125	216
Spatial Grid Size	20^3	100X	25X	8X
	30^3	1300X	325X	105X
	40^3	7500X	1875X	600X

Figure 11: Typical computational speed-ups for 3D simulation and rendering in Legendre domain as compared to the spatial domain.

homogenous media, is critical for achieving photo-realism while rendering 3D scenes. Finally, in Figure 16, we add non-homogenous and dynamic fog to a clear day fly-through of Swiss Alps.

Computational Speed-ups: Due to the compact representations of fields in the Legendre domain, we can achieve computational speed-ups of one to three orders of magnitude, depending on the number of Legendre coefficients (Figures 10 and 11). The comparison is made with our implementation of the Stable Fluids [Sta99] algorithm in the spatial domain. However, our technique places a restriction on the size of the simulation time-step; adding higher frequencies will require a progressively smaller time-step owing to stability considerations given by the CFL condition [FM97]. On the other hand, the Stable Fluids technique can support arbitrarily large time-steps. All our implementation was done in MATLAB on a 3.2GHz P-4 PC with 2 GB of RAM.




Figure 12: Legendre domain Simulation result: 500 pieces of confetti being carried by a turbulent wind field simulated using 216 Legendre coefficients.

8. Discussion of Limitations and Future Work

Our goal in this paper is fast rendering of non-homogenous and dynamic participating media. We achieve this by representing the spatio-temporally varying intensity (rendering), as well as density and velocity (simulation) fields in a reduced analytic Legendre space. This results in a single scattering based rendering technique for smooth non-homogenous and dynamic media, a significant improvement over similar techniques which make the severely limiting assumption of homogenous medium densities [SRNN05]. We believe this is the first work that provides a unified framework for both modeling and rendering in an analytic reduced space, and hope this can help bridge the gap between model reduction in fluids and pre-computed radiance transfer in rendering. However, the speed and analytic nature of the technique come at

the cost of its limited ability to handle high frequency fluid phenomena. Indeed, using only a global Legendre Polynomials basis offers limited local control and allows only for the **box-boundary conditions**, making it difficult to account for complex effects like local vorticities, turbulence and objects inside the medium. Also, being a **global sub-space method**, it offers low flexibility on the domain boundaries.


These limitations can be addressed by augmenting the global Legendre polynomials basis, which capture the majority of the energy of the fluid flow, with a local-support basis such as Haar-Wavelets or spatial voxels, thus accounting for the spatially sparse 'residual energy'. This is similar in spirit to adding local high frequency turbulence, or vorticities [FSJ01] to counter the dampening caused by the Stable Fluids semi-Lagrangian technique. Also, high frequency details in a particular dimension can be captured by keeping the full spatial representation and using Legendre expansion in the remaining directions. Using such *hybrid bases* can provide the desired local control in addition to computational speed-ups, and in our opinion, forms a very promising direction for future research. Since we also make assumptions of single scattering, orthographic viewing and distant lighting, extending our system to perspective viewer and more general, near-field lighting is another research direction worth exploring.

Acknowledgements

We would like to thank Adam Bargteil for helpful technical discussions and the anonymous reviewers for their comments and suggestions. This work was supported in parts by an NSF CAREER award #IIS-0643628, an NSF grant #CCF-0541307 and an ONR award #N00014-05-1-0188.

References

- [Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Oxford Univ. Press, 1960. 1, 2, 3
- [CM90] CHORIN A., MARSDEN J.: *A Mathematical Introduction to Fluid Mechanics*. Springer-Verlag. Texts in Applied Mathematics 4. Second Edition., New York, 1990. 1, 2, 3
- [DYN02] DOBASHI Y., YAMAMOTO T., NISHITA T.: Interactive rendering of atmospheric scattering effects using graphics hardware. In *Graphics Hardware Workshop* (2002), pp. 99–109. 2
- [EP90] EBERT D. S., PARENT R. E.: Rendering and animation of gaseous phenomena by combining fast volume and scanline a-buffer techniques. In *Proceedings of SIGGRAPH* (1990). 1
- [FF01] FOSTER N., FEDKIW R.: Practical animation of liquids. In *Proceedings of SIGGRAPH* (2001), pp. 23–30. 1
- [FM96] FOSTER N., METAXAS D.: Realistic animation of liquids. *Graph. Models Image Process.* 58, 5 (1996), 471–483. 1
- [FM97] FOSTER N., METAXAS D.: Modeling the motion of a hot, turbulent gas. In *Proceedings of SIGGRAPH* (1997). 1, 7
- [FSJ01] FEDKIW R., STAM J., JENSEN H. W.: Visual simulation of smoke. In *Proceedings of SIGGRAPH* (2001), pp. 15–22. 1, 8
- [GN07] GUPTA M., NARASIMHAN S.: *Legendre polynomials Triple Product Integral and lower-degree approximation of polynomials using Chebyshev polynomials*. Tech. Rep. CMU-RI-TR-07-22, Carnegie Mellon University, May 2007. 2, 4
- [HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: Acquisition of time-varying participating media. *ACM Trans. Graph.* 24, 3 (2005), 812–815. 2
- [HL01] HARRIS M., LASTRA A.: Real-time cloud rendering. In *Eurographics* (2001), pp. 76–84. 2
- [Jen01] JENSEN H. W.: *Realistic image synthesis using photon mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001. 1
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of SIGGRAPH* (2001), pp. 511–518. 2
- [KH84] KAJIYA J. T., HERZEN B. P. V.: Ray tracing volume densities. *SIGGRAPH Comput. Graph.* 18, 3 (1984), 165–174. 1
- [Kos24] KOSCHMIEDER H.: Theorie der horizontalen sichtweite. *beitr. In Phys. Freien Atm.* (1924), pp. 171–181. 7
- [LBC94] LANGUENOU E., BOUATOUCH K., CHELLE M.: Global illumination in presence of participation media with general properties. In *Fifth Eurographics Workshop on Rendering* (1994). 1
- [Max86] MAX N. L.: Atmospheric illumination and shadows. In *Proceedings of SIGGRAPH* (1986), pp. 117–124. 2
- [Max94] MAX N. L.: Efficient light propagation for multiple anisotropic volume scattering. In *Fifth Eurographics Workshop on Rendering* (1994), pp. 87–104. 1
- [NFJ02] NGUYEN D. Q., FEDKIW R., JENSEN H. W.: Physically based modeling and animation of fire. In *Proceedings of SIGGRAPH* (2002), pp. 721–728. 1
- [NGD*06] NARASIMHAN S. G., GUPTA M., DONNER C., RAMAMOORTHY R., NAYAR S. K., JENSEN H. W.: Acquiring scattering properties of participating media by dilution. *ACM Trans. Graph.* 25, 3 (2006), 1003–1012. 2
- [NN03] NARASIMHAN S. G., NAYAR S. K.: Shedding light on the weather. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (June 2003), vol. 1, pp. 665 – 672. 2
- [NRH04] NG R., RAMAMOORTHY R., HANRAHAN P.: Triple product wavelet integrals for all-frequency relighting. *ACM Trans. Graph.* 23, 3 (2004), 477–487. 2
- [PM93] PATTANAİK S., MUDUR S.: Computation of global illumination in a participating medium by monte carlo simulation. *Journal of Vis. and Computer Animation* 4, 3 (1993), 133–152. 1
- [REK*04] RILEY K., EBERT D., KRAUS M., TESSENDORF J., HANSEN C.: Efficient rendering of atmospheric phenomena. In *EuroGraphics Symposium on Rendering* (2004). 2
- [RT87] RUSHMEIER H. E., TORRANCE K. E.: The zonal method for calculating light intensities in the presence of a participating medium. In *Proceedings of SIGGRAPH* (1987), pp. 293–302. 1
- [Sak90] SAKAS G.: Fast rendering of arbitrary distributed volume densities. In *Eurographics* (1990), pp. 519–530. 1
- [SRF05] SELLE A., RASMUSSEN N., FEDKIW R.: A vortex particle method for smoke, water and explosions. *ACM Trans. Graph.* 24, 3 (2005), 910–914. 1
- [SRNN05] SUN B., RAMAMOORTHY R., NARASIMHAN S. G., NAYAR S. K.: A practical analytic single scattering model for real time rendering. *ACM Trans. Graph.* 24, 3 (2005). 2, 7
- [Sta99] STAM J.: Stable fluids. In *Proceedings of SIGGRAPH* (1999), pp. 121–128. 1, 2, 3, 5, 7
- [Sta01] STAM J.: A simple fluid solver based on the fft. *J. Graph. Tools* 6, 2 (2001), 43–52. 1
- [TLP06] TREUILLE A., LEWIS A., POPOVIC Z.: Model reduction for real-time fluids. *ACM Trans. Graph.* 25, 3 (2006). 1, 5
- [TMPS03] TREUILLE A., MCNAMARA A., POPOVIC Z., STAM J.: Keyframe control of smoke simulations. *ACM Trans. Graph.* 22, 3 (2003), 716–723. 1



Light Source and View Point Variation

Variation in Scattering Properties

Figure 13: 3D Legendre domain Rendering: Here we consider a smoke-cube illuminated by distant light source(s). The image is formed at an orthographic viewer observing the scene. Since the whole of our pipe-line is in the reduced Legendre domain, the user can control the view-point, lighting and scattering albedo interactively. Notice the varying shadow-gram patterns on the wall as the smoke evolves. The smoke and the shadow become darker as we decrease the albedo. Colored smoke and shadows can be created by varying the scattering properties differently across the color channels. This example required 64 coefficients for density and velocity, and 216 coefficients for intensity fields.




Figure 14: 3D Legendre domain simulation and advection of optical properties: 3D Simulation results for a vertically upwards axial impulse applied to a vase shaped smoke density field. Also, we advect the optical properties of the media (scattering albedos) along with the densities and velocities to create the effect of mixing of different media. This example required 216 Legendre coefficients for density and velocity fields (simulation) and 512 coefficients for intensity fields (rendering).



Clear Weather

Homogenous mist

Non-homogenous mist

Attenuation

Figure 15: Rendering of Non-homogenous participating media: Our technique can be used to render non-homogenous media as well under the single scattering model efficiently. Here we add mist to a clear weather scene (Images courtesy Google Earth). Non-homogenous density distributions, for example the high mist density over the lake provides for more realism as compared to homogenous mist. Also, notice how distant objects appear brighter due to the air-light effect, whereas distant objects appear darker in the attenuation-only image.




Figure 16: Snapshots from a fly-through of Swiss Alps with Non-homogenous and dynamic fog added (Images courtesy Google Earth). Images have been tone-mapped to high-lite the non-homogeneity of the medium. Complete video is included with the supplemental material.

