

Article

Not peer-reviewed version

A Convergent BDF2-SAV Scheme for Fractional Cahn-Hilliard Equations with High-Dimensional Integral Fractional Laplacian

Renhe Liu , [Yanyan Yu](#)*, Mengqian Yue

Posted Date: 23 April 2026

doi: 10.20944/preprints202604.1669.v1

Keywords: multi-dimensional Cahn-Hilliard equation; high-order difference scheme; energy stability; convergence analysis; vectorized tensor processing (VTP)



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Convergent BDF2-SAV Scheme for Fractional Cahn-Hilliard Equations with High-Dimensional Integral Fractional Laplacian

Renhe Liu, Yanyan Yu * and Mengqian Yue

School of Mathematics Science, University of Jinan, Jinan, People's Republic of China

* Correspondence: sms_yuyy@ujn.edu.cn

Abstract

Numerical simulation of the multi-dimensional space-fractional Cahn-Hilliard equation faces two main computational challenges: the inherent temporal accuracy limitations of standard scalar auxiliary variable (SAV) methods and the escalating computational cost in high-dimensional domains. To address these issues, this study constructs a fully discrete algorithmic framework integrating a second-order backward differentiation formula (SAV-BDF2) with a sixth-order centered difference scheme. Under this formulation, we rigorously prove unconditional energy stability and establish the theoretical validity of the dual temporal and spatial accuracy. To solve the resulting indefinite algebraic systems, a minimal residual solver is paired with a sine-transform block diagonal preconditioner. Additionally, a hardware-level Vectorized Tensor Processing (VTP) architecture is deployed to resolve cache thrashing caused by non-contiguous memory access during multidimensional tensor evaluations. Numerical experiments in 3D to 8D domains demonstrate that the framework improves memory throughput and reduces execution time. By avoiding standard hardware execution inefficiencies, this integrated strategy provides an efficient numerical solution for large-scale simulations of high-dimensional fractional systems.

Keywords: multi-dimensional Cahn-Hilliard equation; high-order difference scheme; energy stability; convergence analysis; vectorized tensor processing (VTP)

1. Introduction

Phase-field models provide a standard baseline for simulating various free boundary problems [1]. Back in 1958, Cahn and Hilliard built the CH equation by minimizing a free energy functional to track the thermodynamic and kinetic evolution of binary alloys [1]. Evaluating solid-liquid interfacial tensions represents a core application, since this physical quantity largely controls solidification and phase separation inside porous media [2]. Modern imaging tasks, specifically binary image inpainting, also borrow this framework [3]. A continuous scalar phase field tracks the interface implicitly. That variable stays almost flat inside the bulk region but changes sharply across the diffuse layer [1]. To set up our current study, we define the coupled system on a high-dimensional bounded domain $\Omega \subset \mathbb{R}^d$ ($d \geq 2$) with homogeneous Dirichlet boundary conditions.

Physically, particle interactions at long ranges naturally unfold as spatial convolutions or non-local integrals. Squeezing these global dynamics into purely localized derivatives exposes a fundamental flaw inherent in traditional integer-order CH formulations. That arbitrary truncation exacts a severe physical toll [1]. By their very nature, fractional-order frameworks step in to faithfully map out these exact non-local transport mechanisms, recovering the global interactions that standard models filter out [1]. The Riesz fractional operator \mathcal{L}_α directly handles our spatial discretization.

$$\begin{cases} \delta_t u = \mathcal{L}_\alpha v, & \mathbf{x} \in \Omega, t \in (0, T], \\ v = -\varepsilon^2 \mathcal{L}_\alpha u + f(u), & \mathbf{x} \in \Omega, t \in (0, T], \\ u = v = 0, & \mathbf{x} \in \delta\Omega, t \in (0, T], \\ u(\mathbf{x}, 0) = u_0(\mathbf{x}), & \mathbf{x} \in \Omega, \end{cases} \quad (1)$$

In the above system, the positive parameter ε controls the thickness of the diffuse interface [4], and the nonlinear function $f(u) = F'(u)$ denotes the derivative of the macroscopic double-well potential $F(u) = \frac{1}{4}(u^2 - 1)^2$ [5].

From a mathematical perspective, this coupled system represents the gradient flow of the total free energy $E(u) = \int_\Omega \left(\frac{\varepsilon^2}{2} |\mathcal{L}_\alpha^{1/2} u|^2 + F(u) \right) dx$ evaluated in the dual Sobolev space [6]. Given the strictly dissipative nature of the CH framework, its continuous solution inherently obeys an energy dissipation law, meaning the global energy of the system decreases monotonically as time evolves [7]. Rigorously inheriting this dissipative behavior at the fully discrete level—which corresponds to developing unconditionally energy-stable algorithms—is a fundamental prerequisite for reliably simulating long-time macroscopic phase evolutions [8].

In the pursuit of maintaining energy stability while discretizing the stiff nonlinearities, various strategies have been proposed, including the classical convex splitting technique [9] and the invariant energy quadratization (IEQ) methodology [8]. Widespread attention currently surrounds the scalar auxiliary variable (SAV) formulation and its generalized extensions (gPAV) as highly efficient alternatives [10]. As a standard numerical tool for these gradient flows, the method introduces a single scalar variable to replace the square root of the nonlinear term. That substitution effectively linearizes the entire system, meaning one only needs to solve linear systems with constant coefficients at any arbitrary time step. First-order temporal accuracy, unfortunately, acts as a strict bottleneck for the majority of these current SAV-based algorithms. Extended simulation runs easily trigger unphysical high-frequency oscillations under such limited precision. Coupling the SAV approach with the second-order backward differentiation formula (BDF2) directly resolves this specific problem. It successfully pushes the temporal accuracy to second order while rigorously guaranteeing unconditional energy dissipation at the discrete level [7].

Scaling these high-order numerical schemes to multidimensional scenarios ($d \geq 2$) inevitably triggers critical bottlenecks in terms of computational speed and memory consumption [1]. Discretization on the spatial grid restructures the fractional problem into an enormous linear algebraic system. A notoriously ill-conditioned, 2×2 symmetric indefinite block matrix governs this setup [5]. Because fractional derivatives operate non-locally, the resulting coefficient matrices become completely dense. They naturally exhibit a multi-level block Toeplitz layout [12]. Inverting these blocks directly is computationally prohibitive. To bypass this obstacle, preconditioned Krylov subspace solvers step in as the mainstream workaround. Built on the discrete sine transform, the τ -preconditioner compresses the spectral distribution. This action safely anchors the eigenvalues in a region away from zero [13]. Strict theoretical bounds cap the per-step complexity at $\mathcal{O}(N \log N)$ [5]. Practically, non-contiguous memory access patterns deep inside these high-dimensional tensor computations inevitably trigger an avalanche of severe cache misses. Under real hardware execution, those low-level penalties heavily dilute the actual algorithmic efficiency.

Against the dual demands of temporal accuracy and high-dimensional computational efficiency, we build a fully discrete algorithmic framework. Breakdowns of the core contributions follow:

- **Second-order temporal discretization and unconditional stability.** Hooking up the BDF2 method handles time stepping. That specific choice pushes the temporal precision straight to $\mathcal{O}(\tau^2)$, natively supplying excellent stiff decay to damp out high-frequency wiggles. Formal proofs secure unconditional energy stability strictly at the discrete grid level.
- **VTP-accelerated preconditioned fast solver.** Spatially, a newly constructed sixth-order Riesz fractional centered difference scheme pairs with a sine-transform-based τ -preconditioner. To break

through the memory access bottlenecks inherent in high-dimensional tensor math, a vectorized tensor processing (VTP) layout steps in. Customized tensor shuffling allows this VTP approach to repackage preconditioned matrix-vector multiplications into strictly memory-contiguous batched transforms. Such low-level hardware alignment effectively cuts down on severe cache misses. While firmly preserving the $\mathcal{O}(N \log N)$ theoretical complexity bounds, actual computational throughput under high spatial resolutions experiences a massive boost.

The mathematical and computational framework of this study unfolds across four subsequent sections. Section 2 (Materials and Methods) constructs the discrete foundation by deriving the sixth-order Riesz fractional spatial scheme via double generating functions, explicitly verifying the structural properties of the weighting coefficients, and detailing the synergistic integration of the τ -preconditioner with the continuous memory-access VTP architecture. Section 3 (Results) subjects this proposed framework to strict theoretical and empirical scrutiny, providing rigorous proofs for unconditional energy stability alongside multidimensional numerical validations of the $\mathcal{O}(h^6)$ spatial convergence and 3D phase separation morphologies under extreme fractional anisotropy. Moving to Section 4 (Discussion), we critically evaluate the physical and computational implications of these outcomes, isolating the memory access bottlenecks of traditional DST algorithms and demonstrating how tensor matricization within the VTP framework mitigates cache starvation to sustain optimal high-dimensional execution speedups. Finally, Section 5 (Conclusions) closes the manuscript by summarizing our central algorithmic contributions.

2. Materials and Methods

2.1. A Sixth-Order Approximation to the Riesz Fractional Derivative

To aggressively cut down spatial truncation errors, we introduce a specific generating function $G_{\alpha_i}(z)$ for the continuous Riesz fractional operator $\mathcal{L}_{x_i}^{\alpha_i}$ ($1 < \alpha_i < 2$), which is specifically constructed to eliminate the leading truncation errors in the Taylor expansion and achieve the desired sixth-order convergence rate. That mathematical setup explicitly wipes out the leading truncation errors during Taylor expansion to secure the sixth-order convergence rate.

That mathematical setup explicitly wipes out the leading truncation errors during Taylor expansion to secure the sixth-order convergence rate. Coefficient symmetry natively guarantees that the discrete fractional operator produces a symmetric Toeplitz-like matrix across the interior grid. Algebraically, this exact structure opens the door for spectral diagonalization using the discrete sine transform. Furthermore, its strict zero-sum property enforces the negative semi-definiteness of the discrete operator. Strict adherence to the discrete energy dissipation law naturally follows.

Let Ω be uniformly partitioned into square elements of side length h , and let M denote the total number of grid points. To obtain a high-order approximation of the Riesz derivative, we construct a numerical differentiation formula based on an appropriately designed generating function. Specifically, we introduce a new generating function $G(z)$ of the form

$$g(z) = \frac{3}{2}(z^{-1} - 2 + z)^{\alpha/2} - \frac{3(2^{-\alpha})}{5}(z^{-2} - 2 + z^2)^{\alpha/2} + \frac{3^{-\alpha}}{10}(z^{-3} - 2 + z^3)^{\alpha/2}, \quad (2)$$

In summary, the discrete scheme obtained via the sixth-order generating function is presented as follows:

$$\delta_h^{\alpha_i} u(x_i) := -\frac{1}{h^{\alpha_i}} \sum_{k=-\infty}^{\infty} g_k^{\alpha_i} u(x_i - kh_i), \quad |\omega| \leq 1 \quad (3)$$

Setting $z = e^{-i\theta}$ in (2) and applying the inverse Fourier transform, the explicit analytical expression for the corresponding weighting coefficients $g_k^{(\alpha_i)}$ can be deduced as:

$$g_k^{\alpha_i} = \frac{3}{2} a_k^{\alpha_i} - \mathbf{1}_{\{2|k\}} \cdot \frac{3(2^{-\alpha_i})}{5} a_{k/2}^{\alpha_i} + \mathbf{1}_{\{3|k\}} \cdot \frac{3^{-\alpha_i}}{10} a_{k/3}^{\alpha_i} \quad (4)$$

where $\Gamma(\cdot)$ denotes the Euler gamma function and $a_k^{\alpha_i} = \frac{(-1)^k \Gamma(\alpha_i + 1)}{\Gamma(\alpha_i/2 - k + 1) \Gamma(\alpha_i/2 + k + 1)}$. Furthermore, we define a class of special functions as follows:

$$\mathbf{1}_{\{k|n\}} = \begin{cases} 1, & n \equiv 0 \pmod{k}; \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

To ensure the energy stability of the subsequent fully discrete scheme and the applicability of the fast transform-based preconditioner, it is crucial to establish the algebraic properties of the coefficients $g_k^{(\alpha_i)}$.

Proposition 1. For any fractional order $\alpha_i \in (1, 2)$, the sixth-order weighting coefficients $g_k^{(\alpha_i)}$ satisfy the following properties:

- $1 < g_0^{\alpha_i} < 3$,
- $g_k^{\alpha_i} = g_{-k}^{\alpha_i} < 0$, $k = 1, 2, \dots$,
- $|g_{k+1}^{\alpha_i}| \leq |g_k^{\alpha_i}|$, $k = 1, 2, \dots$,
- $\sum_{k=-\infty}^{\infty} g_k^{\alpha_i} = 0$, $\sum_{k \neq 0} |g_k^{\alpha_i}| \leq g_0^{\alpha_i}$.

where $g_k^\alpha = \frac{(-1)^k \Gamma(\alpha + 1)}{\Gamma(\alpha/2 - k + 1) \Gamma(\alpha/2 + k + 1)}$.

The symmetry property (i) guarantees that the discrete fractional operator yields a symmetric Toeplitz-like matrix in the interior domain, which is essential for the spectral diagonalization using the discrete sine transform. The zero-sum property (v) structurally ensures the semi-negative definiteness of the discrete operator, preserving the energy dissipation law at the discrete level.

Next, we take the sixth-order approximation to demonstrate that the format can achieve six-order accuracy. The remaining discrete formulas can be proved in a similar manner.

Proposition 2. Let

$$u \in \mathcal{L}^{6+\alpha}(\mathbb{R}) = \left\{ u \mid u \in L^1(\mathbb{R}) \text{ and } \int_{\mathbb{R}} (1 + |\omega|)^{6+\alpha} |\hat{u}(\omega)| d\omega < \infty \right\}, \quad 1 < \alpha < 2,$$

then we have that

$$\delta_{x,\mathbb{R}}^\alpha u(x) = \delta_{h,\mathbb{R}}^\alpha u(x) + O(h^6), \quad (6)$$

uniformly in $x \in \mathbb{R}$.

Proof. Let we denote $g(i\omega h) = \lambda_1 \left(4 \sin\left(\frac{\omega h}{2}\right)\right)^\alpha + \lambda_2 \left(4 \sin\left(\frac{2\omega h}{2}\right)\right)^\alpha + \lambda_3 \left(4 \sin\left(\frac{3\omega h}{2}\right)\right)^\alpha$. and $u(\omega, h) = \mathcal{F}\left\{\delta_{x,\mathbb{R}}^\alpha u(x) - \delta_{h,\mathbb{R}}^\alpha u(x)\right\}$. According to the properties of Fourier transform, there is

$$\mathcal{F}\{\delta_{x,\mathbb{R}}^\alpha u\} = -|\omega|^\alpha \hat{u}(\omega),$$

and

$$\mathcal{F}\{\delta_{h,\mathbb{R}}^\alpha u\} = -|\omega|^\alpha \frac{g(i\omega h)}{(\omega h/2)^\alpha} \hat{u}(\omega).$$

Together with

$$\mathcal{F}\left\{\left(4 \sin\left(\frac{\omega h}{2}\right)\right)^\alpha\right\} = -|\omega|^\alpha \left(1 - \frac{\alpha(\omega h)^2}{24} + \left(-\frac{\alpha}{2880} + \frac{\alpha^2}{1152}\right)(\omega h)^4 + O(\omega h)^6\right),$$

we have

$$u(\omega, h) = |\omega|^\alpha \left\{ -1 + \left[(\lambda_1 + 2^\alpha \lambda_2 + 3^\alpha \lambda_3) - \frac{1}{24} (\alpha (\lambda_1 + 2^{2+\alpha} \lambda_2 + 3^{2+\alpha} \lambda_3)) (\omega h)^2 \right. \right. \\ \left. \left. + \frac{1}{5760} (\alpha (-2 + 5\alpha) (\lambda_1 + 2^{4+\alpha} \lambda_2 + 3^{4+\alpha} \lambda_3)) (\omega h)^4 + O(\omega h)^6 \right] \right\} \hat{u}(\omega).$$

Solving the equation

$$\begin{cases} \lambda_1 + 2^\alpha \lambda_2 + 3^\alpha \lambda_3 = 1, \\ \alpha (\lambda_1 + 2^{2+\alpha} \lambda_2 + 3^{2+\alpha} \lambda_3) = 0, \\ \alpha (-2 + 5\alpha) (\lambda_1 + 2^{4+\alpha} \lambda_2 + 3^{4+\alpha} \lambda_3) = 0, \end{cases}$$

yields

$$\lambda_1 = \frac{3}{2}, \quad \lambda_2 = -\frac{3(2^{-\alpha})}{5}, \quad \lambda_3 = \frac{3^{-\alpha}}{10},$$

which in turn enables us to obtain

$$u(\omega, h) = |\omega|^\alpha \left[-\frac{-16\alpha + 42\alpha^2 - 35\alpha^3}{80640} (\omega h)^6 + \mathcal{O}(\omega h)^8 \right] \hat{u}(\omega).$$

Thus, there exists a constant $C_{h1} > 0$ such that

$$|u(\omega, h)| \leq C_{h1} |\omega|^{\alpha+6} |\hat{u}(\omega)| h^6.$$

From the inverse Fourier transform,

$$\begin{aligned} \left| \delta_{x, \mathbb{R}}^\alpha u(x) - \delta_{h, \mathbb{R}}^\alpha u(x) \right| &= \left| \frac{1}{2\pi} \int_{\mathbb{R}} u(\omega, h) e^{i\omega x} d\omega \right| \\ &\leq \frac{1}{2\pi} \int_{\mathbb{R}} |u(\omega, h)| d\omega \\ &\leq \frac{C_{h1}}{2\pi} \left(\int_{\mathbb{R}} |\omega|^{\alpha+6} |\hat{u}(\omega)| d\omega \right) h^6. \end{aligned}$$

Since $u \in C^{6+\alpha}(\mathbb{R})$, we have $\int_{\mathbb{R}} (1 + |\omega|^{\alpha+6}) |\hat{u}(\omega)| d\omega < \infty$

therefore $\left| \delta_{x, \mathbb{R}}^\alpha u(x) - \delta_{h, \mathbb{R}}^\alpha u(x) \right| \leq C_{h2} h^6$, Here, $C_{h2} = \frac{C_{h1}}{2\pi} \int_{\mathbb{R}} (1 + |\omega|^{\alpha+6}) |\hat{u}(\omega)| d\omega$. \square

We now extend this approximation to construct the multi-dimensional grid system. Consider a d -dimensional rectangular domain $\Omega = \prod_{i=1}^d (a_i, b_i)$. For each spatial direction $i \in \{1, 2, \dots, d\}$, let M_i be a positive integer representing the number of interior grid points. We define the uniform mesh size $h_i = \frac{b_i - a_i}{M_i + 1}$ and generate a set of equidistant nodes $x_{i,j} = a_i + jh_i$ for $j = 0, 1, \dots, M_i + 1$.

Let \mathbb{Z} denote the set of integers. For each direction i , we define the one-dimensional interior and global index sets as follows:

$$\begin{aligned} I_i &= \{j \in \mathbb{Z} \mid 1 \leq j \leq M_i\}, \\ \hat{I}_i &= \{j \in \mathbb{Z} \mid 0 \leq j \leq M_i + 1\}. \end{aligned}$$

Subsequently, we construct the d -dimensional multi-index sets via the Cartesian product $K = \prod_{i=1}^d I_i$ and $\hat{K} = \prod_{i=1}^d \hat{I}_i$. A multi-index $J = (j_1, j_2, \dots, j_d) \in \hat{K}$ uniquely identifies a spatial grid point $X_J = (x_{1,j_1}, x_{2,j_2}, \dots, x_{d,j_d})$. The collection of all such nodes forms the d -dimensional spatial mesh $\Omega_h = \{X_J \mid J \in \hat{K}\}$.

2.2. Fully Discrete Scheme

In this subsection, we reformulate the multidimensional spatial fractional Cahn-Hilliard (FCH) equation based on the Scalar Auxiliary Variable (SAV) approach. First, we restate the model equations and rewrite the two-dimensional FCH model into the following equivalent coupled system:

$$\begin{cases} \frac{\delta u(x,t)}{\delta t} = \mathcal{L}_x^\alpha v(x,t), \\ v(x,t) = -\varepsilon^2 \mathcal{L}_x^\alpha u(x,t) + \frac{r(t)}{\sqrt{\int_\Omega F(u)dx + C_0}} f(u(x,t)), \\ \frac{dr(t)}{dt} = \frac{1}{2\sqrt{\int_\Omega F(u)dx + C_0}} \int_\Omega f(u) u_t dx. \end{cases} \quad (7)$$

2.2.1. Function Spaces and Continuous Operators

Consider the continuous Riesz fractional derivative operator $\mathcal{L}_{x_i}^{\alpha_i}$ ($1 < \alpha_i \leq 2$) in the x_i -direction defined on the bounded domain Ω . Let $\mathcal{L}_x^\alpha = \sum_{i=1}^d \mathcal{L}_{x_i}^{\alpha_i}$. Under homogeneous Dirichlet boundary conditions, the operator $-\mathcal{L}_{x_i}^{\alpha_i}$ is symmetric positive definite. Accordingly, there exists a continuous root operator denoted by $\tilde{\mathcal{L}}_{x_i}^{\alpha_i}$, which satisfies the following inner product identity:

$$(-\mathcal{L}_{x_i}^{\alpha_i} v, v) = (\tilde{\mathcal{L}}_{x_i}^{\alpha_i} v, \tilde{\mathcal{L}}_{x_i}^{\alpha_i} v) = \|\tilde{\mathcal{L}}_{x_i}^{\alpha_i} v\|^2. \quad (8)$$

Throughout this paper, $\|\cdot\|$ and (\cdot, \cdot) represent the standard L^2 norm and inner product, respectively.

2.2.2. Multi-Dimensional Fully Discrete Scheme

Based on the d -dimensional spatial mesh, we proceed with the temporal discretization of the fractional Cahn-Hilliard equation. Let $\tau = T/N$ be the time step size and $t_k = k\tau$. We denote U_j^k and V_j^k as the numerical approximations of the phase-field variable $u(X, t)$ and the chemical potential $v(X, t)$ at the spatial node X_j and time level t_k , respectively. Furthermore, we introduce the auxiliary scalar approximation $r^k \approx r(t_k)$.

To spatially discretize the continuous Riesz fractional derivative, we adopt the fractional centered difference formula. Let U_j and V_j be the numerical approximations of $u(X, t)$ and $v(X, t)$ at the spatial grid point X_j , respectively. The discrete fractional operator $\delta_{x_i}^{\alpha_i}$ with respect to the i -th spatial variable x_i at point X_j is expressed as:

$$\delta_{x_i}^{\alpha_i} U_j := -\frac{1}{h_i^{\alpha_i}} \sum_{L \in \mathcal{N}_i(J)} g_{|L(i)-J(i)|}^{(\alpha_i)} U_L + O(h_i^6), \quad J \in \mathbb{K}, \quad i = 1, 2, \dots, d, \quad (9)$$

where $\mathcal{N}_i(J)$ denotes the set of multi-indices varying exclusively along the i -th dimension, and $L(i), J(i)$ represent their corresponding index components in that specific direction.

Let $M = \prod_{i=1}^d M_i$ be the total number of interior spatial degrees of freedom. Specifically, for the one-dimensional case (where $M = M_1$), the discrete operator (9) simplifies to the following scalar summation:

$$\delta_x^\alpha U_j = \sum_{i=1}^d \delta_{x_i}^{\alpha_i} U_j + O(h_i^6) \quad (h_i \rightarrow 0) \quad (10)$$

where the centered difference weights $g_k^{(\alpha_i)}$ are explicitly defined as follows:

$$g_k^{\alpha_i} = \frac{3}{2} a_k^{\alpha_i} - \mathbf{1}_{\{2|n\}} \cdot \frac{3(2^{-\alpha_i})}{5} a_{k/2}^{\alpha_i} + \mathbf{1}_{\{3|n\}} \cdot \frac{3^{-\alpha_i}}{10} a_{k/3}^{\alpha_i} \quad (11)$$

It is mathematically evident from (11) that for any fractional order $\alpha \in (1, 2)$, the sequence $g_k^{(\alpha_i)}$ does not truncate to zero for $k \neq 0$. Consequently, applying this numerical scheme yields a fully dense

and symmetric Toeplitz coefficient matrix, which rigorously captures the non-local physical nature of the fractional phase-field model.

To handle the extreme stiffness of the nonlinear term explicitly while preserving unconditional energy stability, we define the second-order extrapolation $\bar{U}_J^{k+1} = 2U_J^k - U_J^{k-1}$ and the corresponding nonlinear functional approximation:

$$W(\bar{U}_J^{k+1}) = \frac{f(\bar{U}_J^{k+1})}{\sqrt{\left(\prod_{i=1}^d h_i\right) \sum_{J \in K} F(\bar{U}_J^{k+1}) + C_0}}. \quad (12)$$

Applying the SAV approach combined with the second-order backward differentiation formula (BDF2), we obtain the following d -dimensional fully discrete scheme:

$$\begin{cases} \frac{3U_J^{k+1} - 4U_J^k + U_J^{k-1}}{2\tau} = \sum_{i=1}^d \delta_{x_i}^{\alpha_i} V_J^{k+1}, \\ V_J^{k+1} = -\varepsilon^2 \sum_{i=1}^d \delta_{x_i}^{\alpha_i} U_J^{k+1} + r^{k+1} W(\bar{U}_J^{k+1}), \\ \frac{3r^{k+1} - 4r^k + r^{k-1}}{2\tau} = \frac{1}{2} \left(\prod_{i=1}^d h_i \right) \sum_{J \in K} W(\bar{U}_J^{k+1}) \frac{3U_J^{k+1} - 4U_J^k + U_J^{k-1}}{2\tau}, \end{cases} \quad (13)$$

which is subject to the following homogeneous boundary and initial conditions:

$$\begin{cases} U_J^k = V_J^k = 0, & J \in \delta K, \quad 0 \leq k \leq N, \\ U_J^0 = U_0(X_J), \quad U_J^1 = U_1(X_J), & J \in K, \\ r^0 = \sqrt{\left(\prod_{i=1}^d h_i\right) \sum_{J \in K} F(U_J^0) + C_0} \\ r^1 = \sqrt{\left(\prod_{i=1}^d h_i\right) \sum_{J \in K} F(U_J^1) + C_0}. \end{cases} \quad (14)$$

2.3. Fast Preconditioner and Its Tensor Batching Optimization

In this subsection, we introduce the spectral-equivalent τ -preconditioner to decouple the severely ill-conditioned linear system derived from the SAV-BDF2 scheme. To address the curse of dimensionality encountered by this preconditioner in high-dimensional spatial problems, we further propose a Vectorized Tensor Processing (VTP) architecture for low-level computational acceleration.

2.3.1. Preconditioned Symmetrized System Under SAV-BDF2

Following the spatial discretization of the multidimensional Riesz fractional operator, the intrinsic mechanism of the SAV-BDF2 scheme yields an asymmetric coupled algebraic system. To legitimately invoke Krylov subspace solvers such as MINRES, this structural asymmetry must be algebraically resolved.

By introducing the specific scaling parameters $c_1 = \varepsilon^2$ and $c_2 = \frac{2\tau}{3}$, alongside the auxiliary multiplier $\theta = \sqrt{c_1/c_2}$ and the equivalent aggregated coefficient $c = \sqrt{c_1 c_2}$, we strictly recast the original coupled partial differential system into a highly symmetric 2×2 block linear topology:

$$\mathcal{A} \begin{bmatrix} \bar{\mathbf{U}}^{k+1} \\ \mathbf{V}^{k+1} \end{bmatrix} = \begin{bmatrix} -cG & I \\ I & cG \end{bmatrix} \begin{bmatrix} \theta \mathbf{U}^{k+1} \\ \mathbf{V}^{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_2 \\ \theta \frac{2\tau}{3} \mathbf{f}_1 \end{bmatrix}. \quad (15)$$

As the spatial grid step $h \rightarrow 0$ and the physical interface width $\varepsilon \rightarrow 0$, the dense coefficient matrix G inevitably drives the global block matrix \mathcal{A} into severe ill-conditioning. To fundamentally

circumvent this computational hurdle in multidimensional scenarios, we adopt the robust block diagonal preconditioner \mathcal{P} , which is inspired by the sine-transform-based technique proposed in [6]:

$$\mathcal{P} = \begin{bmatrix} c\tau(G) + I & 0 \\ 0 & c\tau(G) + I \end{bmatrix}, \quad (16)$$

where $\tau(G) = \sum_{i=1}^d (I_{M_i^-} \otimes \tau(G_{\alpha_i}) \otimes I_{M_i^+})$ denotes the d -dimensional global approximation matrix.

For clarity and practical implementation, based on our spatial domain discretization, each sub-block $\tau(G_{\alpha_i})$ can be rigorously diagonalized by the 1D discrete sine transform (DST) matrix S_{M_i} , satisfying $\tau(G_{\alpha_i}) = S_{M_i} \Lambda_{\alpha_i} S_{M_i}$. The internal elements of the orthogonal transformation matrix S_{M_i} and the diagonal eigenvalue matrix Λ_{α_i} are explicitly defined as follows [6]:

$$[S_{M_i}]_{j,l} = \sqrt{\frac{2}{M_i + 1}} \sin\left(\frac{jl\pi}{M_i + 1}\right), \quad (17)$$

$$[\Lambda_{\alpha_i}]_j = g_0^{(\alpha_i)} + 2 \sum_{k=1}^{M_i-1} g_k^{(\alpha_i)} \cos\left(\frac{kj\pi}{M_i + 1}\right), \quad (18)$$

for $j, l \in \{1, 2, \dots, M_i\}$, where $g_k^{(\alpha_i)}$ represent the generation coefficients derived from the fractional difference scheme. This flawless tensor diagonalization mechanism guarantees that the inverse action \mathcal{P}^{-1} can be executed with quasi-optimal complexity via the fast Fourier transform.

Remark. It has been strictly proven in [5] that the eigenvalues of the preconditioned matrix $\mathcal{P}^{-1}\mathcal{A}$ are uniformly bounded within the disjoint intervals $[-\frac{3}{2}, -\frac{1}{2\sqrt{2}}] \cup [\frac{1}{2\sqrt{2}}, \frac{3}{2}]$. From an algebraic perspective, this spectral distribution is intrinsically immune to the specific value of the temporal scaling factor $c > 0$. Consequently, the introduction of our new coefficient $c_2 = \frac{2\tau}{3}$, which is exclusively dictated by the BDF2 discretization, preserves these theoretical bounds intact. This mathematical property ensures that the block diagonal preconditioner remains highly robust, ultimately yielding a mesh-independent linear convergence rate for the asymmetric SAV-BDF2 scheme.

2.3.2. Vectorized Tensor Processing (VTP) Architecture

From a purely algebraic perspective, the theoretical time complexity of applying the d -dimensional preconditioner \mathcal{P}^{-1} is mathematically bounded by $\mathcal{O}(N \sum_{i=1}^d \log M_i)$.

Nevertheless, under standard routines, executing a d -dimensional joint sine transform usually means slicing the global tensor into isolated one-dimensional vectors. The system then launches exactly N/M_k independent scalar DST calls along the k -th dimension. This mechanism tears memory continuity apart. Such granular data extraction inevitably causes steep instruction overhead and massive cache misses. Hampered by these limits, computational efficiency in high-dimensional scenarios simply collapses.

To fundamentally dismantle this hardware-level barrier, we propose the Vectorized Tensor Processing (VTP) architecture. Let $\mathcal{U} \in \mathbb{R}^{M_1 \times M_2 \times \dots \times M_d}$ denote the d -dimensional physical state tensor isomorphic to the discrete algebraic system. When the sine transform is required along the k -th dimension, an orthogonal permutation operator \mathcal{T}_k is first introduced to advance the target computational dimension to the leading position in memory. Subsequently, this permuted tensor is compulsorily reshaped into a highly dense two-dimensional matrix manifold:

$$\hat{\mathcal{U}}^{(k)} = \text{Reshape}\left(\mathcal{T}_k(\mathcal{U}), \left[M_k, \frac{N}{M_k}\right]\right). \quad (19)$$

Governed by this absolutely contiguous memory mapping mechanism, the initially fragmented N/M_k scalar transforms are instantaneously aggregated into a single, high-density batched execution block. To strictly avoid the complexity regression that would be caused by explicit dense matrix

multiplication, the transformation is executed holistically via a Batched Fast Discrete Sine Transform (Batched FDST) operator $\mathcal{F}_{\text{DST}}^{(k)}$ over the continuous column space of the manifold:

$$\mathbf{Z}^{(k)} = \mathcal{F}_{\text{DST}}^{(k)}(\hat{\mathcal{U}}^{(k)}). \quad (20)$$

Upon the completion of this batched fast operation, the inverse permutation operator \mathcal{T}_k^{-1} is immediately applied, accurately mapping the data back to its original physical coordinate system.

By reconstructing the underlying dataflow, the VTP architecture elegantly collapses the staggering $\mathcal{O}(N/M_k)$ fragmented function calls within a single iteration step into strictly $\mathcal{O}(d)$ macroscopic batched operations. Pinning the spatial footprint at $\mathcal{O}(N)$ while maintaining the optimal algebraic time complexity, this tensorization strategy decisively breaks through the execution bottlenecks inherent in high-dimensional discrete schemes right down at the bare-metal hardware level.

While strictly maintaining a minimal spatial complexity of $\mathcal{O}(N)$ and the optimal algebraic time complexity, this tensorization strategy decisively breaks through the execution bottlenecks inherent in high-dimensional discrete schemes at the bare-metal hardware level.

3. Results

3.1. Energy Stability Analysis

In this subsection, we initially formulate the necessary definitions ensuring the regularity of the d -dimensional discrete operations, and subsequently provide a rigorous proof establishing the unconditional energy stability of the proposed scheme.

3.1.1. Necessary Definitions

Let U be the space of grid functions defined on all interior mesh points, arranged in lexicographical order. For any grid functions $u, v \in U$, we define the discrete inner product and the corresponding L^2 norm in the d -dimensional space as:

$$(u, v)_h = \left(\prod_{i=1}^d h_i \right) \sum_{J \in K} u_J v_J, \quad \|v\|_h = \sqrt{(v, v)}. \quad (21)$$

For each spatial direction $i \in \{1, \dots, d\}$, the discrete fractional operator $-\delta_{x_i}^{\alpha_i}$ is symmetric positive definite on the space U . Correspondingly, there exists a discrete root operator, denoted by $\tilde{\delta}_{x_i}^{\alpha_i}$, such that the inner product identity $(-\delta_{x_i}^{\alpha_i} u, u) = (\tilde{\delta}_{x_i}^{\alpha_i} u, \tilde{\delta}_{x_i}^{\alpha_i} u) = \|\tilde{\delta}_{x_i}^{\alpha_i} u\|_h^2$ holds natively under homogeneous boundary conditions.

3.1.2. Proof of Energy Decay

To establish the energy dissipation property of the fully discrete SAV-BDF2 scheme in high-dimensional settings, we first introduce the crucial algebraic identity for the BDF2 formulation.

Lemma 3.1 For any sequence of vectors $\{a^k\}$ in a real inner product space, the following identity holds unconditionally:

$$2(a^{k+1}, 3a^{k+1} - 4a^k + a^{k-1}) = \|a^{k+1}\|_h^2 - \|a^k\|_h^2 + \|2a^{k+1} - a^k\|_h^2 - \|2a^k - a^{k-1}\|_h^2 + \|a^{k+1} - 2a^k + a^{k-1}\|_h^2. \quad (22)$$

Based on the SAV-BDF2 time discretization, we define the modified discrete total energy \tilde{E}^{k+1} at the $(k+1)$ -th time level as:

$$\tilde{E}^{k+1} = \frac{\varepsilon^2}{2} \sum_{i=1}^d \left(\|\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1}\|_h^2 + \|2\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - \tilde{\delta}_{x_i}^{\alpha_i} U^k\|_h^2 \right) + \left((r^{k+1})^2 + (2r^{k+1} - r^k)^2 \right). \quad (23)$$

Theorem 3.1: The d -dimensional fully discrete SAV-BDF2 scheme guarantees that the modified total energy decreases monotonically, meaning $\tilde{E}^{k+1} \leq \tilde{E}^k$ holds unconditionally for any time step $\tau > 0$.

Proof: Taking the discrete inner product of the first equation in the scheme with V^{k+1} , and utilizing the discrete fractional summation-by-parts equivalent, we obtain:

$$\left(\frac{3U^{k+1} - 4U^k + U^{k-1}}{2\tau}, V^{k+1} \right) = - \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} V^{k+1} \right\|_h^2. \quad (24)$$

Taking the discrete inner product of the second equation with the time derivative term $\frac{3U^{k+1} - 4U^k + U^{k-1}}{2\tau}$ yields:

$$\begin{aligned} \left(V^{k+1}, \frac{3U^{k+1} - 4U^k + U^{k-1}}{2\tau} \right) &= \varepsilon^2 \sum_{i=1}^d \left(\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1}, \tilde{\delta}_{x_i}^{\alpha_i} \frac{3U^{k+1} - 4U^k + U^{k-1}}{2\tau} \right) \\ &\quad + r^{k+1} \left(W(\bar{U}^{k+1}), \frac{3U^{k+1} - 4U^k + U^{k-1}}{2\tau} \right). \end{aligned} \quad (25)$$

Equating (24) and (25), and multiplying both sides by 2τ , we deduce:

$$\begin{aligned} -2\tau \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} V^{k+1} \right\|_h^2 &= \varepsilon^2 \sum_{i=1}^d \left(\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1}, 3\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - 4\tilde{\delta}_{x_i}^{\alpha_i} U^k + \tilde{\delta}_{x_i}^{\alpha_i} U^{k-1} \right) \\ &\quad + r^{k+1} \left(W(\bar{U}^{k+1}), 3U^{k+1} - 4U^k + U^{k-1} \right). \end{aligned} \quad (26)$$

From the third equation of the SAV-BDF2 scheme, the nonlinear inner product term can be exactly substituted by the auxiliary scalar variables:

$$\left(W(\bar{U}^{k+1}), 3U^{k+1} - 4U^k + U^{k-1} \right) = 2(3r^{k+1} - 4r^k + r^{k-1}). \quad (27)$$

Substituting (27) into (26), the equation transforms into a strictly quadratic form:

$$-2\tau \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} V^{k+1} \right\|_h^2 = \varepsilon^2 \sum_{i=1}^d \left(\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1}, 3\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - 4\tilde{\delta}_{x_i}^{\alpha_i} U^k + \tilde{\delta}_{x_i}^{\alpha_i} U^{k-1} \right) + 2r^{k+1}(3r^{k+1} - 4r^k + r^{k-1}). \quad (28)$$

Using the BDF2 polarization identity (22) to both the fractional gradient term and the scalar term, (28) rigorously expands to:

$$\begin{aligned} -2\tau \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} V^{k+1} \right\|_h^2 &= \frac{\varepsilon^2}{2} \sum_{i=1}^d \left[\left\| \tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} \right\|_h^2 - \left\| \tilde{\delta}_{x_i}^{\alpha_i} U^k \right\|_h^2 + \left\| 2\tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - \tilde{\delta}_{x_i}^{\alpha_i} U^k \right\|_h^2 \right. \\ &\quad \left. - \left\| 2\tilde{\delta}_{x_i}^{\alpha_i} U^k - \tilde{\delta}_{x_i}^{\alpha_i} U^{k-1} \right\|_h^2 + \left\| \tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - 2\tilde{\delta}_{x_i}^{\alpha_i} U^k + \tilde{\delta}_{x_i}^{\alpha_i} U^{k-1} \right\|_h^2 \right] \\ &\quad + \left[(r^{k+1})^2 - (r^k)^2 + (2r^{k+1} - r^k)^2 \right. \\ &\quad \left. - (2r^k - r^{k-1})^2 + (r^{k+1} - 2r^k + r^{k-1})^2 \right]. \end{aligned} \quad (29)$$

Rearranging the terms according to the time levels $k+1$ and k , and utilizing the definition of the modified discrete total energy (23), we arrive at the discrete dissipation equation:

$$\tilde{E}^{k+1} - \tilde{E}^k + \frac{\varepsilon^2}{2} \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} U^{k+1} - 2\tilde{\delta}_{x_i}^{\alpha_i} U^k + \tilde{\delta}_{x_i}^{\alpha_i} U^{k-1} \right\|_h^2 + (r^{k+1} - 2r^k + r^{k-1})^2 = -2\tau \sum_{i=1}^d \left\| \tilde{\delta}_{x_i}^{\alpha_i} V^{k+1} \right\|_h^2. \quad (30)$$

Since the second-order numerical dissipation terms on the left side of (30) are strictly non-negative, and the right side satisfies $-2\tau \sum_{i=1}^d \|\delta_{x_i}^{\alpha_i} V^{k+1}\|^2 \leq 0$, we logically deduce $\tilde{E}^{k+1} \leq \tilde{E}^k$. This rigorously completes the proof.

3.2. Global Convergence Analysis

Rigorous global convergence proofs for the fully discrete SAV-BDF2 scheme unfold throughout this subsection. By confining the analysis to a zero-mean subspace, the mathematical setup natively permits direct energy estimation inside the discrete negative fractional Sobolev space, essentially laying the prerequisite groundwork for all subsequent derivations. Strategically, computational pairing tricks like the polarization identity extract inner products into measurable norms for comparative scaling. To guarantee an exact algebraic cancellation of those notoriously stiff nonlinear components, the cross-terms remain deliberately untouched as inner products rather than suffering premature expansion. These residual truncation errors get cordoned off into independent norms prior to expansion. A closed-loop error inequality system immediately follows. Working through explicit coefficient matching and reliable scaling bounds ultimately secures the theoretical justification for the numerical convergence.

3.2.1. Discrete Matrix, Equivalent Norms, and Core Lemmas

Definition 3 (Discrete Fractional Matrix and Equivalent Norms). Let Ω_h denote the d -dimensional uniform spatial mesh. Let $\mathcal{L} = \sum_{i=1}^d \delta_{x_i}^{\alpha_i}$ be the original symmetric negative definite fractional operator corresponding to the continuous Laplacian Δ . To establish a valid discrete Sobolev topology, we define the corresponding symmetric positive definite (SPD) matrix as $\mathcal{A} = -\mathcal{L} = -\sum_{i=1}^d \delta_{x_i}^{\alpha_i}$.

Within the discrete zero-mean subspace $\dot{U}_h = \{U \in \mathbb{R}^M \mid \sum_{J \in K} U_J = 0\}$, all eigenvalues of \mathcal{A} are strictly positive, ensuring the unique existence of its inverse matrix \mathcal{A}^{-1} . In addition to the standard discrete L^2 -norm $\|U\|_h = \sqrt{\langle U, U \rangle_h}$, the positive definiteness induces the equivalent fractional norms:

$$\|U\|_H = \sqrt{\langle U, \mathcal{A}U \rangle_h}, \quad \|U\|_{H^{-1}} = \sqrt{\langle U, \mathcal{A}^{-1}U \rangle_h}. \quad (31)$$

Given that \dot{U}_h constitutes a real Hilbert space and \mathcal{A} is self-adjoint, the generalized Cauchy-Schwarz inequality $\langle U, V \rangle_h \leq \|U\|_H \|V\|_{H^{-1}}$ holds unconditionally.

Lemma 4 (Mesh-Independent Discrete Interpolation and Poincaré Inequalities). For any given penalty parameter $\gamma > 0$ independent of the temporal step size τ and the spatial step size h , the following discrete interpolation inequality unconditionally holds in the subspace \dot{U}_h :

$$\|U\|_h^2 \leq \frac{\gamma}{2} \|U\|_H^2 + \frac{1}{2\gamma} \|U\|_{H^{-1}}^2, \quad \forall U \in \dot{U}_h. \quad (32)$$

Furthermore, let $\lambda_1 = \mathcal{O}(1) > 0$ be the minimum non-zero eigenvalue of \mathcal{A} on \dot{U}_h . There exists a global constant $C_1 = \lambda_1^{-1/2}$, strictly independent of the mesh size h , such that $\|U\|_{H^{-1}} \leq C_1 \|U\|_h$ and $\|U\|_h \leq C_1 \|U\|_H$. This implies the discrete Poincaré inverse control inequality:

$$\|U\|_{H^{-1}} \leq C_1^2 \|U\|_H, \quad \forall U \in \dot{U}_h. \quad (33)$$

Lemma 5 (Discrete Lipschitz Continuity of the SAV Functional). The modified nonlinear functional takes the composite form $W(u) = \frac{f(u)}{\sqrt{E(u) + C_0}}$, where $f(u) = u^3 - u$, the truncation constant is $C_0 > 0$, and $E(u) = \langle F(u), 1 \rangle_h$ acts as a discrete spatial scalar.

Assume that up to the current time step t_m , the numerical solutions are confined within a compact set $\|u\|_{L^\infty}, \|U^k\|_{L^\infty} \leq C_2$. Under this constraint, $W(\cdot)$ satisfies a global Lipschitz continuity in the discrete space. Introducing the second-order temporal extrapolation $u(t_{k+1}) - \bar{U}^{k+1} = \zeta_\tau^{k+1} + 2E_u^k - E_u^{k-1}$, where $\|\zeta_\tau^{k+1}\|_H \leq C\tau^2$ by Taylor's theorem. By the discrete Poincaré inequality, we have $\|\zeta_\tau^{k+1}\|_h \leq C_1 \|\zeta_\tau^{k+1}\|_H \leq \mathcal{O}(\tau^2)$. Since the denominator is an integral scalar independent of spatial coordinates, the fractional matrix directly penetrates it: $\mathcal{A}^{1/2}W(u) = (\mathcal{A}^{1/2}f(u))/\sqrt{E(u) + C_0}$.

Leveraging the discrete fractional product rule, there exist independent constants $C_3, C_4 > 0$ such that the nonlinear bounds are governed by:

$$\begin{aligned} \|W(u(t_{k+1})) - W(\bar{U}^{k+1})\|_h^2 &\leq C_3 \left(\|\xi_\tau^{k+1}\|_h^2 + \|E_u^k\|_h^2 + \|E_u^{k-1}\|_h^2 \right), \\ \|W(u(t_{k+1})) - W(\bar{U}^{k+1})\|_H^2 &\leq C_4 \left(\|\xi_\tau^{k+1}\|_H^2 + \|E_u^k\|_H^2 + \|E_u^{k-1}\|_H^2 \right). \end{aligned} \quad (34)$$

3.2.2. Discrete Error Evolution System

By evaluating the exact continuous solutions at the grid points (X_J, t_{k+1}) and substituting them into the fully discrete SAV-BDF2 scheme, we define the numerical errors as $E_u^{k+1} = u(t_{k+1}) - U^{k+1}$, $E_v^{k+1} = v(t_{k+1}) - V^{k+1}$, and $e_r^{k+1} = r(t_{k+1}) - R^{k+1}$. Applying the backward difference operator $\delta_t \phi^{k+1} = \frac{3\phi^{k+1} - 4\phi^k + \phi^{k-1}}{2\tau}$, the consistency equations yield the coupled error evolution system:

$$\begin{cases} \delta_t E_u^{k+1} = -\mathcal{A}E_v^{k+1} + \rho_u^{k+1}, \\ E_v^{k+1} = \varepsilon^2 \mathcal{A}E_u^{k+1} + Q_N^{k+1} + \rho_v^{k+1}, \\ \delta_t e_r^{k+1} = Q_r^{k+1} + \rho_r^{k+1}. \end{cases} \quad (35)$$

The nonlinear phase-field residual and the auxiliary scalar residual are rigorously defined as:

$$\begin{aligned} Q_N^{k+1} &= r(t_{k+1})W(u(t_{k+1})) - R^{k+1}W(\bar{U}^{k+1}), \\ Q_r^{k+1} &= \frac{1}{2} \left[\langle W(u(t_{k+1})), \delta_t u(t_{k+1}) \rangle_h - \langle W(\bar{U}^{k+1}), \delta_t U^{k+1} \rangle_h \right]. \end{aligned} \quad (36)$$

Assumption 6 (Truncation Errors and Initialization). According to Taylor expansion and difference approximation theories, the local truncation errors maintain consistent spatiotemporal accuracy (second-order in time, sixth-order in space), bounded by a global constant C_5 :

$$\|\rho_u^{k+1}\|_{H^{-1}}^2 + \|\rho_v^{k+1}\|_H^2 + |\rho_r^{k+1}|^2 \leq C_5(\tau^4 + h^{12}). \quad (37)$$

As BDF2 is a two-step method, the initial physical conditions are prescribed exactly, implying $E_u^0 = 0$ and $e_r^0 = 0$. The numerical solutions at the first time step, U^1 and R^1 , are assumed to be computed by a sufficiently accurate single-step method. There exists a constant C_{10} such that the initialization errors strictly conform to the global convergence rate:

$$\|E_u^1\|_H \leq C_{10}(\tau^2 + h^6), \quad |e_r^1| \leq C_{10}(\tau^2 + h^6). \quad (38)$$

3.2.3. Proof of the Global Convergence

Provided that Assumption 6 holds, the multidimensional fully discrete SAV-BDF2 scheme converges unconditionally. There exist global constants C_{12} and C_{13} , independent of τ and h , such that for any valid time step m :

$$\|u(t_m) - U^m\|_{H^{-1}} \leq C_{13} \exp(C_{12}T) (\tau^2 + h^6). \quad (39)$$

Proof. Substituting the chemical potential equation into the phase-field evolution equation, we recast the system into a weak variational form acting on the inverse matrix \mathcal{A}^{-1} :

$$\mathcal{A}^{-1} \delta_t E_u^{k+1} = -\varepsilon^2 \mathcal{A}E_u^{k+1} - Q_N^{k+1} - \rho_v^{k+1} + \mathcal{A}^{-1} \rho_u^{k+1}. \quad (40)$$

Taking the discrete L^2 inner product of both sides with the temporal increment test function $4\tau\delta_t E_u^{k+1}$ within \hat{U}_h , we obtain:

$$\begin{aligned} 4\tau\langle \mathcal{A}^{-1}\delta_t E_u^{k+1}, \delta_t E_u^{k+1} \rangle_h &= -4\tau\varepsilon^2\langle \mathcal{A}E_u^{k+1}, \delta_t E_u^{k+1} \rangle_h \\ &\quad - 4\tau\langle Q_N^{k+1}, \delta_t E_u^{k+1} \rangle_h \\ &\quad - 4\tau\langle \rho_v^{k+1}, \delta_t E_u^{k+1} \rangle_h \\ &\quad + 4\tau\langle \mathcal{A}^{-1}\rho_u^{k+1}, \delta_t E_u^{k+1} \rangle_h. \end{aligned} \quad (41)$$

The left-hand side yields the positive definite term $4\tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2$. For the first term on the right-hand side, the symmetric property translates it to $-4\tau\varepsilon^2\langle E_u^{k+1}, \delta_t E_u^{k+1} \rangle_H$. Applying the algebraic BDF2 polarization identity $2\langle 3A - 4B + C, A \rangle_H = \|A\|_H^2 - \|B\|_H^2 + \|2A - B\|_H^2 - \|2B - C\|_H^2 + \|A - 2B + C\|_H^2$, and dropping the non-negative second-order dissipation term $-\varepsilon^2\|E_u^{k+1} - 2E_u^k + E_u^{k-1}\|_H^2 \leq 0$, we define the baseline phase-field functional $\mathbb{E}_u^k = \varepsilon^2(\|E_u^k\|_H^2 + \|2E_u^k - E_u^{k-1}\|_H^2)$. Rearranging the terms yields:

$$\begin{aligned} \mathbb{E}_u^{k+1} - \mathbb{E}_u^k + 4\tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2 &\leq -4\tau\langle Q_N^{k+1}, \delta_t E_u^{k+1} \rangle_h \\ &\quad - 4\tau\langle \rho_v^{k+1}, \delta_t E_u^{k+1} \rangle_h \\ &\quad + 4\tau\langle \rho_u^{k+1}, \delta_t E_u^{k+1} \rangle_{H^{-1}}. \end{aligned} \quad (42)$$

Concurrently, multiplying the auxiliary scalar evolution equation by $8\tau e_r^{k+1}$ ensures precise alignment with the 1/2 factor in Q_r^{k+1} . Following a parallel BDF2 polarization and defining the scalar functional $\mathbb{E}_r^k = 2|e_r^k|^2 + 2|2e_r^k - e_r^{k-1}|^2$, dropping the positive residual gives:

$$\mathbb{E}_r^{k+1} - \mathbb{E}_r^k \leq 8\tau e_r^{k+1} Q_r^{k+1} + 8\tau e_r^{k+1} \rho_r^{k+1}. \quad (43)$$

To handle the nonlinear matrix cancellation, we define the total energy functional $\mathcal{H}^k = \mathbb{E}_u^k + \mathbb{E}_r^k$. Summing (42) and (43) isolates the implicit nonlinear array $I_{NL} = -4\tau\langle Q_N^{k+1}, \delta_t E_u^{k+1} \rangle_h + 8\tau e_r^{k+1} Q_r^{k+1}$. Substituting the definition of Q_N^{k+1} and applying $\delta_t U^{k+1} = \delta_t u(t_{k+1}) - \delta_t E_u^{k+1}$ within Q_r^{k+1} , the array expands as:

$$\begin{aligned} I_{NL} &= -4\tau e_r^{k+1}\langle W(u(t_{k+1})), \delta_t E_u^{k+1} \rangle_h \\ &\quad - 4\tau R^{k+1}\langle W(u(t_{k+1})) - W(\bar{U}^{k+1}), \delta_t E_u^{k+1} \rangle_h \\ &\quad + 4\tau e_r^{k+1}\langle W(u(t_{k+1})) - W(\bar{U}^{k+1}), \delta_t u(t_{k+1}) \rangle_h \\ &\quad + 4\tau e_r^{k+1}\langle W(\bar{U}^{k+1}), \delta_t E_u^{k+1} \rangle_h. \end{aligned} \quad (44)$$

The first and fourth terms exactly cancel. Extracting the common inner product factor with the R^{k+1} term and invoking the mapping $e_r^{k+1} + R^{k+1} \equiv r(t_{k+1})$, the implicit nonlinear terms algebraically collapse to:

$$\begin{aligned} I_{NL} &= -4\tau r(t_{k+1})\langle W(u(t_{k+1})) - W(\bar{U}^{k+1}), \delta_t E_u^{k+1} \rangle_h \\ &\quad + 4\tau e_r^{k+1}\langle W(u(t_{k+1})) - W(\bar{U}^{k+1}), \delta_t u(t_{k+1}) \rangle_h. \end{aligned} \quad (45)$$

To bound the remaining terms, we apply the weighted Young's inequality ($ab \leq a^2 + b^2/4$). Utilizing the Cauchy-Schwarz inequality, the three terms coupled with $\delta_t E_u^{k+1}$ are decomposed as follows:

$$\begin{aligned} 4\tau\left|r(t_{k+1})\langle W(u) - W(\bar{U}), \delta_t E_u^{k+1} \rangle_h\right| &\leq 4\tau|r(t_{k+1})|^2\|W(u) - W(\bar{U})\|_H^2 + \tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2, \\ 4\tau\left|\langle \rho_v^{k+1}, \delta_t E_u^{k+1} \rangle_h\right| &\leq 4\tau\|\rho_v^{k+1}\|_H^2 + \tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2, \\ 4\tau\left|\langle \rho_u^{k+1}, \delta_t E_u^{k+1} \rangle_{H^{-1}}\right| &\leq 4\tau\|\rho_u^{k+1}\|_{H^{-1}}^2 + \tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2. \end{aligned} \quad (46)$$

The summation of the extracted norms equals precisely $3\tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2$. This quantity is fully absorbed by the positive dissipation term $4\tau\|\delta_t E_u^{k+1}\|_{H^{-1}}^2$ on the left-hand side, leaving a non-negative remainder

$\tau \|\delta_t E_u^{k+1}\|_{H^{-1}}^2$. For the terms independent of $\delta_t E_u^{k+1}$, subject to the physical bounds $\|\delta_t u\|_{L^\infty} \leq C_6$ and $|r(t)| \leq C_7$:

$$\begin{aligned} 4\tau e_r^{k+1} \langle W(u) - W(\bar{U}), \delta_t u \rangle_h &\leq \tau |e_r^{k+1}|^2 + 4\tau C_6^2 \|W(u) - W(\bar{U})\|_h^2, \\ 8\tau e_r^{k+1} \rho_r^{k+1} &\leq \tau |e_r^{k+1}|^2 + 16\tau |\rho_r^{k+1}|^2. \end{aligned} \quad (47)$$

Substituting the global Lipschitz bounds from Lemma 5, let C_8 denote the maximum combined constant arising from the Lipschitz bounds and the Young inequality coefficients. Aggregating these terms exposes the phase-field error history for both the k -th and $(k-1)$ -th steps in strong and weak topologies: $\tau C_8 (\|E_u^k\|_H^2 + \|E_u^{k-1}\|_H^2) + \tau C_8 (\|E_u^k\|_h^2 + \|E_u^{k-1}\|_h^2)$. To decouple the weak norm's dependence on the mesh layout, we invoke Lemma 4 strictly on the L^2 components, assigning a penalty parameter $\gamma = 1/\varepsilon^2$:

$$\tau C_8 \|E_u^j\|_h^2 \leq \tau C_8 \left(\frac{1}{2\varepsilon^2} \|E_u^j\|_H^2 + \frac{\varepsilon^2}{2} \|E_u^j\|_{H^{-1}}^2 \right), \quad \text{for } j \in \{k, k-1\}. \quad (48)$$

We cascade the resulting dual norms into the discrete Poincaré inequality $\|E_u^j\|_{H^{-1}}^2 \leq C_1^4 \|E_u^j\|_H^2$ to elevate their topologies to the H -norm space. Combining these elevated terms with the pre-existing high-order terms, the total phase-field historical bound is consolidated to

$$\tau C_8 \left[1 + \frac{1}{2\varepsilon^2} + \frac{C_1^4 \varepsilon^2}{2} \right] (\|E_u^k\|_H^2 + \|E_u^{k-1}\|_H^2). \quad (49)$$

Concurrently, the scalar residual previously bounded by $2\tau |e_r^{k+1}|^2$ natively satisfies $2\tau |e_r^{k+1}|^2 \leq \tau \mathcal{H}^{k+1}$. Factoring out the physical parameter ε^2 to match the energy functional structure ($\mathcal{H}^j \geq \varepsilon^2 \|E_u^j\|_H^2$), we instantiate the global system constant $C_9 = \max\left(1, \frac{C_8}{\varepsilon^2} \left[1 + \frac{1}{2\varepsilon^2} + \frac{C_1^4 \varepsilon^2}{2}\right]\right)$. Grouping all exterior truncation errors and Taylor remainders into a unified base $b = C_5(\tau^4 + h^{12})$, the system flawlessly transforms into an explicit recurrence governed by \mathcal{H} :

$$\mathcal{H}^{k+1} - \mathcal{H}^k + \tau \|\delta_t E_u^{k+1}\|_{H^{-1}}^2 \leq \tau C_9 (\mathcal{H}^{k+1} + \mathcal{H}^k + \mathcal{H}^{k-1}) + \tau b. \quad (50)$$

Since $\tau \|\delta_t E_u^{k+1}\|_{H^{-1}}^2 \geq 0$, we obtain the weaker inequality without this term. Rearranging the terms configures the implicit relation $(1 - \tau C_9) \mathcal{H}^{k+1} \leq (1 + \tau C_9) \mathcal{H}^k + \tau C_9 \mathcal{H}^{k-1} + \tau b$. Imposing a temporal constraint $\tau \leq \frac{1}{2C_9}$ guarantees $1 - \tau C_9 \geq 1/2 > 0$. Dividing both sides by this factor and subtracting \mathcal{H}^k , the identity $\frac{1+\tau C_9}{1-\tau C_9} - 1 = \frac{2\tau C_9}{1-\tau C_9}$ yields:

$$\mathcal{H}^{k+1} - \mathcal{H}^k \leq \frac{2\tau C_9}{1-\tau C_9} \mathcal{H}^k + \frac{\tau C_9}{1-\tau C_9} \mathcal{H}^{k-1} + \frac{\tau b}{1-\tau C_9}. \quad (51)$$

Applying the summation operator $\sum_{k=1}^m$ triggers a telescoping sum $\sum_{k=1}^m (\mathcal{H}^{k+1} - \mathcal{H}^k) = \mathcal{H}^{m+1} - \mathcal{H}^1$. For the lagged sum, we define the initial state for the fictitious step $E_u^{-1} = 0, e_r^{-1} = 0$, which, coupled with the exact physical initialization $E_u^0 = 0, e_r^0 = 0$, algebraically enforces $\mathcal{H}^0 \equiv 0$. Consequently, $\sum_{k=1}^m \mathcal{H}^{k-1} \leq \sum_{j=0}^m \mathcal{H}^j$ holds strictly. Bounding the multipliers via $\frac{3\tau C_9}{1-\tau C_9} \leq 6\tau C_9$ and the constant sum via $\frac{1}{1-\tau C_9} \leq 2$ condenses the expression to:

$$\mathcal{H}^{m+1} \leq \left(\mathcal{H}^1 + 2\tau b \right) + \sum_{j=0}^m 6C_9 \tau \mathcal{H}^j. \quad (52)$$

In accordance with Assumption 6, squaring the initial constraints provides the bound $\mathcal{H}^1 \leq C_{11}(\tau^2 + h^6)^2$, where C_{11} is derived from C_{10} . Given the condition $\tau^4 + h^{12} \leq (\tau^2 + h^6)^2$, the constant terms fuse perfectly into $(C_{11} + 2TC_5)(\tau^2 + h^6)^2$. By the discrete Gronwall lemma, we have :

$$\mathcal{H}^{m+1} \leq \exp(6C_9 t_{m+1})(C_{11} + 2TC_5)(\tau^2 + h^6)^2. \quad (53)$$

Recalling the coercivity of the energy functional $\mathcal{H}^{m+1} \geq \varepsilon^2 \|E_u^{m+1}\|_H^2 \geq \varepsilon^2 C_1^{-2} \|E_u^{m+1}\|_{H^{-1}}^2$. Executing the square root operation symmetrically reduces the exponent $6C_9$ to $3C_9$. Defining $C_{12} = 3C_9$ and $C_{13} = \varepsilon^{-1} C_1 \sqrt{C_{11} + 2TC_5}$, we explicitly recover the optimal convergence precision:

$$\|u(t_{m+1}) - U^{m+1}\|_{H^{-1}} \leq C_{13} \exp(C_{12}T) (\tau^2 + h^6). \quad (54)$$

This completes the proof. \square

Remark. To verify the a priori uniform boundedness assumption ($\|E_u^{m+1}\|_{L^\infty} \leq 1$), we apply the discrete inverse inequality:

$$\|E_u^{m+1}\|_{L^\infty} \leq C_{14} h^{-(d+\alpha)/2} \|E_u^{m+1}\|_H. \quad (55)$$

Enforcing the strict asymptotic mesh coupling constraint $\tau^2 + h^6 = o(h^{(d+\alpha)/2})$ assures that in the infinitesimal limit as $\tau, h \rightarrow 0$, the mesh singularity amplification is completely dominated by the high-order convergence rate. Consequently, the maximum norm error is compelled to zero, ensuring $\|E_u^{m+1}\|_{L^\infty} \rightarrow 0 \leq 1$. This validates authenticates the assumption and resolves the stiff nonlinearities.

3.3. Numerical Experiments

To see how the algorithm actually works, this subsection runs numerical tests on the fractional Cahn-Hilliard equation from two main angles.

Table 1. Spatial convergence orders of the proposed sixth-order scheme under highly anisotropic configurations in 3D, 4D, and 5D spaces.

Dim	Config (α)	h Refinement	Max Error (L_∞)	Order
3D	[1.2, 1.5, 1.8]	1/32 \rightarrow 1/64	3.41485e-03	—
		1/64 \rightarrow 1/128	5.62345e-05	5.9242
		1/128 \rightarrow 1/256	8.90324e-07	5.9810
4D	[1.2, 1.5, 1.8, 1.3]	1/32 \rightarrow 1/64	3.57993e-03	—
		1/64 \rightarrow 1/128	5.88997e-05	5.9255
		1/128 \rightarrow 1/256	9.32308e-07	5.9813
5D	[1.2, 1.5, 1.8, 1.3, 1.7]	1/32 \rightarrow 1/64	5.22353e-03	—
		1/64 \rightarrow 1/128	8.59100e-05	5.9261
		1/128 \rightarrow 1/256	1.35972e-06	5.9814

4. Discussion

4.1. Accuracy Verification, Energy Stability, and Phase Morphologies

This subsection analysis the accuracy and physical reliability of the proposed scheme. First, the spatial convergence of the sixth-order central difference scheme under highly anisotropic fractional operators is verified. Subsequently, the unconditional energy stability is tested. The computational domain is set to a 3D space $\Omega = (0, 1)^3$ with a grid resolution of $M = 32$. The system is initialized with a multi-dimensional sinusoidal perturbation $u_0(x, y, z) = 0.05 \sin(2\pi x) \sin(2\pi y) \sin(2\pi z)$. The interface width parameter is set to $\varepsilon^2 = 0.04$, and the time step size is fixed at $\tau = 0.01$. The simulation is advanced to a terminal time $T = 3.0$ to capture the delayed spinodal decomposition process induced by extreme fractional anisotropy.

Checking if the physics hold up involves running six specific fractional setups in a 3D grid. These chosen cases cover symmetric setups that act the same in all directions (e.g., $\alpha_1 = \alpha_2 = \alpha_3 = 1.5$) alongside setups that change depending on the direction (e.g., $\alpha_1 = 1.2, \alpha_2 = 1.5, \alpha_3 = 1.8$). Our experiments keep track of the modified discrete energy. Across every single test, the energy levels fall sharply when the phase separation first starts, and then they level out into a slow slide during the coarsening stage. Importantly, having different fractional parameters does not break this steady downward trend. That behavior closely matches the theoretical energy dissipation rule, as verified in Figure 1. Figure 2 points out the vastly different 3D phase separation patterns driven by these varying parameters.

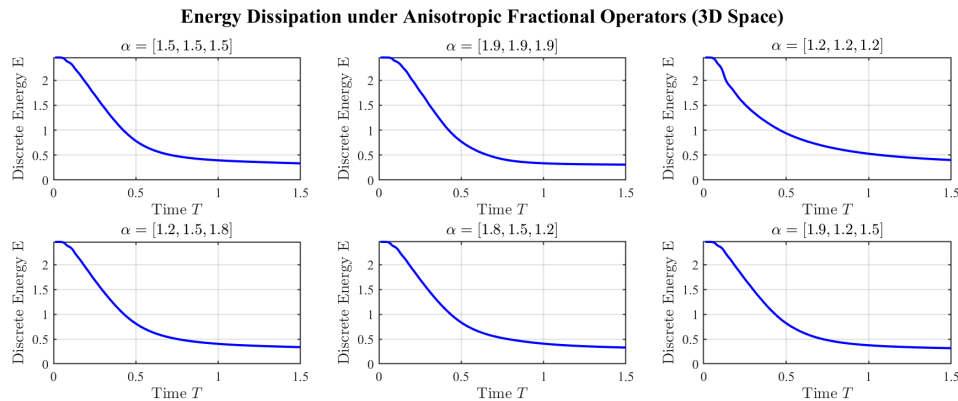


Figure 1. Evolution of the modified discrete energy under six different combinations of fractional orders (α) in 3D space, demonstrating strict monotonic energy dissipation.

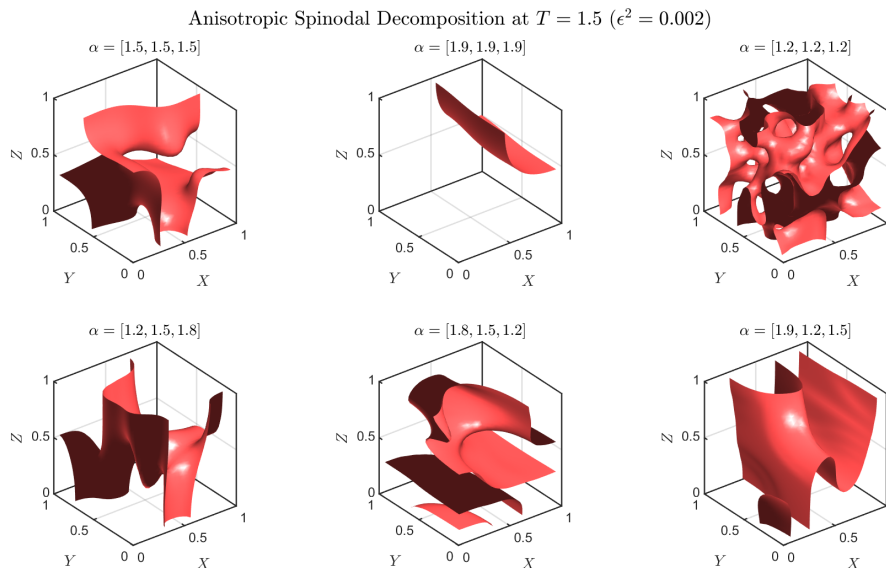


Figure 2. 3D phase separation morphologies under different combinations of anisotropic fractional orders.

4.2. Underlying Execution Efficiency Analysis: VTP vs. DST Under the τ -Preconditioner

This subsection directly compares the underlying single-step execution efficiency of the VTP architecture against the conventional Standard DST within the identical mathematical framework of the τ -preconditioner, isolating the core computational operation of the Krylov solver ($z = P^{-1}x$). The traditional Standard DST relies on multi-layer nested loops to extract data fibers. As the system dimension d increases, the memory strides generated by this non-contiguous access rapidly exceed the physical capacity of CPU caches. In contrast, the VTP architecture utilizes continuous tensor

permutation to reshape the computational dimensions into contiguous memory blocks, facilitating the SIMD vectorization pipelines of modern processors.

As presented in Table 2, the speedup ratio reflects the dynamic shifting of hardware bottlenecks. Because the traditional DST extracts short data fibers sequentially, severe SIMD pipeline starvation and deep nested loop overheads dominate its computation time. VTP mitigates this loop overhead via matricization, reaching a speedup of $35.55\times$ in the 5D space with a minimal grid ($M = 4$). However, as the grid scaling induces exponential data expansion, the dense tensor permutation operations in VTP begin to encounter the theoretical upper limits of memory bandwidth. Consequently, when the grid increases to $M = 8$, the single-step speedup stabilizes in the range of $14.27\times \sim 25.20\times$. Despite this hardware-imposed constraint on memory throughput, VTP strictly maintains contiguous memory access and continues to exhibit a substantial performance advantage.

Table 2. Comparison of single-step execution time between the proposed VTP architecture and Standard DST across varying dimensions and grid resolutions.

Dim	M	DOF (N)	Config (α)	Execution Time (s)		Speedup	
				VTP (Ours)	Std DST		
3D	64	262,144	[1.5, 1.5, 1.5]	0.09655	0.20587	2.49 \times	
			[1.2, 1.2, 1.2]	0.08341	0.19902	2.50 \times	
			[1.9, 1.9, 1.9]	0.11064	0.20457	1.98 \times	
			[1.2, 1.5, 1.8]	0.10106	0.20584	2.38 \times	
			[1.8, 1.5, 1.2]	0.09071	0.20621	2.39 \times	
			[1.9, 1.2, 1.5]	0.09288	0.20447	2.67 \times	
4D	8	4,096	[1.5, 1.5, 1.5, 1.5]	0.00320	0.02846	9.70 \times	
			[1.2, 1.2, 1.2, 1.2]	0.00128	0.02693	18.17 \times	
			[1.9, 1.9, 1.9, 1.9]	0.00124	0.02667	17.52 \times	
			[1.2, 1.4, 1.6, 1.8]	0.00126	0.02683	27.24 \times	
			[1.8, 1.6, 1.4, 1.2]	0.00156	0.02818	19.80 \times	
			[1.5, 1.2, 1.5, 1.8]	0.00140	0.02756	16.80 \times	
	12	20,736		[1.5, 1.5, 1.5, 1.5]	0.00325	0.09221	27.40 \times
				[1.2, 1.2, 1.2, 1.2]	0.00287	0.09282	34.87 \times
				[1.9, 1.9, 1.9, 1.9]	0.00265	0.09280	29.05 \times
				[1.2, 1.4, 1.6, 1.8]	0.00388	0.09403	41.60 \times
				[1.8, 1.6, 1.4, 1.2]	0.00321	0.09554	27.58 \times
				[1.2, 1.8, 1.2, 1.8]	0.00291	0.09478	30.88 \times
5D	4	1,024	[1.5, 1.5, 1.5, 1.5, 1.5]	0.00139	0.01770	10.81 \times	
			[1.2, 1.2, 1.2, 1.2, 1.2]	0.00046	0.01627	32.75 \times	
			[1.8, 1.8, 1.8, 1.8, 1.8]	0.00045	0.01624	29.67 \times	
			[1.1, 1.3, 1.5, 1.7, 1.9]	0.00049	0.01620	27.41 \times	
			[1.9, 1.7, 1.5, 1.3, 1.1]	0.00056	0.01610	29.46 \times	
			[1.2, 1.8, 1.2, 1.8, 1.2]	0.00046	0.01632	35.55 \times	
	6	7,776		[1.5, 1.5, 1.5, 1.5, 1.5]	0.00277	0.08569	21.85 \times
				[1.2, 1.2, 1.2, 1.2, 1.2]	0.00310	0.08228	36.59 \times
				[1.8, 1.8, 1.8, 1.8, 1.8]	0.00291	0.08384	30.36 \times
				[1.1, 1.3, 1.5, 1.7, 1.9]	0.00280	0.08299	26.70 \times
				[1.9, 1.7, 1.5, 1.3, 1.1]	0.00260	0.08239	26.22 \times
				[1.2, 1.8, 1.2, 1.8, 1.2]	0.00310	0.08266	26.57 \times
	8	32,768		[1.5, 1.5, 1.5, 1.5, 1.5]	0.01547	0.27134	25.20 \times
				[1.2, 1.2, 1.2, 1.2, 1.2]	0.01407	0.26958	14.27 \times
				[1.8, 1.8, 1.8, 1.8, 1.8]	0.01459	0.26590	18.77 \times
				[1.1, 1.3, 1.5, 1.7, 1.9]	0.01556	0.27062	16.95 \times
				[1.9, 1.7, 1.5, 1.3, 1.1]	0.01621	0.27003	18.21 \times
				[1.2, 1.8, 1.2, 1.8, 1.2]	0.01491	0.27305	21.89 \times

Furthermore, Table 2 indicates the decoupling characteristic of the anisotropic parameter α on hardware execution efficiency. Under the same grid (e.g., 4D, $M = 12$), whether α is symmetric [1.5, ..., 1.5] or highly asymmetric [1.2, 1.4, 1.6, 1.8], the single-step execution time of VTP fluctuates

within a narrow interval of $0.0026 \sim 0.0039$ seconds. The fundamental reason is that changing the directionality of α merely alters the values in the diagonal preconditioner matrix; the underlying instruction graph remains physically unchanged.

As illustrated in Figure 3, when approaching the extreme spaces from 6D to 8D, the execution time of the traditional DST architecture exhibits an exponential deterioration trend due to cache degradation. In contrast, the VTP architecture maintains a steady polynomial-level scaling in the log-log coordinate system, validating its hardware robustness against high-dimensional scaling.

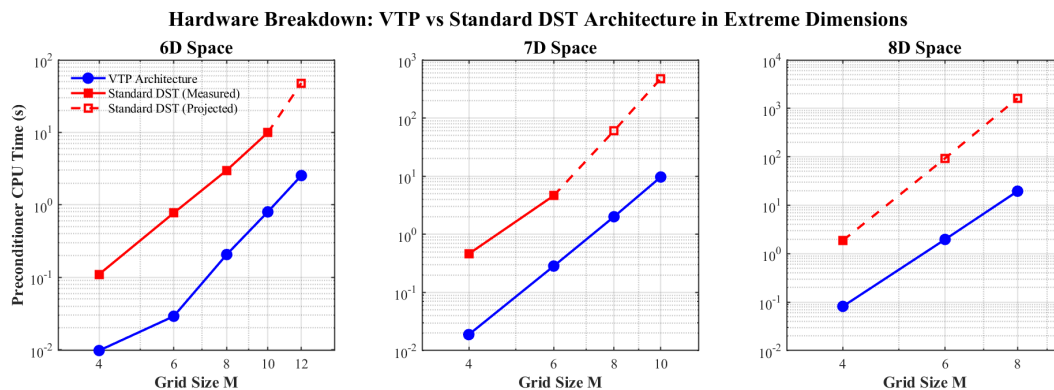


Figure 3. Execution time scaling of VTP versus Standard DST in high-dimensional spaces (6D to 8D).

4.3. Total Solver Time Comparison: τ -Preconditioned VTP vs. Baseline Algorithm

The total solver time is synergistically determined by the single-step execution time and the algebraic Krylov iteration count.

Pitting the τ -preconditioned VTP setup against a baseline skew-circulant preconditioner forms the core of this subsection. Boundary disturbances near the edges drive up the condition number. Such boundary mismatches force the iteration count to climb visibly as the grid gets finer, a trend clearly captured in Table 3. Mathematical alignment with the homogeneous Dirichlet boundaries comes naturally to the τ -preconditioner.

Algebraically, it safely fences in the eigenvalues of the preconditioned matrix. The Krylov steps stay pinned at a rock-bottom level of roughly 15. This low count holds steady no matter what you do with the dimension d , the grid size M , or the fractional parameters α .

Table 3 puts hard numbers on the VTP acceleration. In the 4D tests with $M = 12$, the baseline iteration count gets worse and falls back to 29 steps. Combined speedup ratios for the VTP framework land anywhere from $55.0\times$ up to $73.2\times$. To our disappointment, pushing M from 4 to 8 in the 5D runs causes the peak speedup to slide from $61.1\times$ down to $53.2\times$. Data grows exponentially alongside grid refinement. Heavy tensor shuffling starts hitting the physical wall of memory bandwidth. Those bare-metal limits inevitably pull the acceleration back a bit. Giving up a little ground to the hardware only eats up a tiny fraction of the time saved by taking fewer steps. Our VTP setup easily holds onto a clear performance lead over the older baseline. Interestingly, picking drastically different α values for each spatial axis exposes another weakness in the baseline. Raw execution times for VTP stay entirely flat. That older baseline reacts badly to the messy spectral gaps caused when these α numbers do not match.

Table 3. Comparison of total solver time between the complete τ -preconditioned VTP framework and the representative baseline algorithm.

Dim	M	DOF (N)	Config (α)	Iterations		Total Solver Time (s)		Speedup
				VTP (Ours)	Baseline	VTP (Ours)	Baseline	
3D	64	262,144	[1.5, 1.5, 1.5]	15	39	2.2299	11.5844	4.9×
			[1.2, 1.2, 1.2]	15	39	1.7862	11.3915	5.1×
			[1.9, 1.9, 1.9]	15	39	1.6038	11.8063	5.5×
			[1.2, 1.5, 1.8]	15	39	1.7222	11.2300	6.0×
			[1.8, 1.5, 1.2]	15	39	1.7646	11.3986	6.2×
			[1.9, 1.2, 1.5]	15	39	1.8994	11.6291	7.6×
4D	8	4,096	[1.5, 1.5, 1.5, 1.5]	15	26	0.0635	1.0997	20.2×
			[1.2, 1.2, 1.2, 1.2]	15	26	0.0374	1.0518	24.3×
			[1.9, 1.9, 1.9, 1.9]	15	26	0.0272	1.0247	28.1×
			[1.2, 1.4, 1.6, 1.8]	15	26	0.0334	1.0179	33.6×
			[1.8, 1.6, 1.4, 1.2]	15	26	0.0258	1.0154	32.9×
			[1.5, 1.2, 1.5, 1.8]	15	26	0.0271	1.0187	34.8×
4D	12	20,736	[1.5, 1.5, 1.5, 1.5]	15	29	0.0882	3.9049	72.0×
			[1.2, 1.2, 1.2, 1.2]	15	29	0.0585	3.9537	65.4×
			[1.9, 1.9, 1.9, 1.9]	15	29	0.0657	3.9907	59.5×
			[1.2, 1.4, 1.6, 1.8]	15	29	0.0702	4.0046	73.2×
			[1.8, 1.6, 1.4, 1.2]	15	29	0.0646	4.0071	55.0×
			[1.2, 1.8, 1.2, 1.8]	15	29	0.0641	4.0268	62.8×
5D	4	1,024	[1.5, 1.5, 1.5, 1.5, 1.5]	15	25	0.0407	0.9306	15.0×
			[1.2, 1.2, 1.2, 1.2, 1.2]	15	25	0.0107	0.6252	52.0×
			[1.8, 1.8, 1.8, 1.8, 1.8]	15	25	0.0101	0.6055	46.5×
			[1.1, 1.3, 1.5, 1.7, 1.9]	15	25	0.0101	0.6113	48.4×
			[1.9, 1.7, 1.5, 1.3, 1.1]	15	25	0.0104	0.6042	61.1×
			[1.2, 1.8, 1.2, 1.8, 1.2]	15	25	0.0098	0.6171	57.2×
5D	6	7,776	[1.5, 1.5, 1.5, 1.5, 1.5]	15	27	0.0966	3.4978	59.5×
			[1.2, 1.2, 1.2, 1.2, 1.2]	15	27	0.0871	4.1575	57.2×
			[1.8, 1.8, 1.8, 1.8, 1.8]	15	27	0.0893	4.1048	37.7×
			[1.1, 1.3, 1.5, 1.7, 1.9]	15	27	0.0970	4.0479	63.2×
			[1.9, 1.7, 1.5, 1.3, 1.1]	15	27	0.0870	3.9684	43.0×
			[1.2, 1.8, 1.2, 1.8, 1.2]	15	27	0.0808	3.5902	40.7×
5D	8	32,768	[1.5, 1.5, 1.5, 1.5, 1.5]	15	29	0.4051	16.1457	28.9×
			[1.2, 1.2, 1.2, 1.2, 1.2]	15	29	0.3822	13.6345	36.0×
			[1.8, 1.8, 1.8, 1.8, 1.8]	15	29	0.3928	12.6956	37.0×
			[1.1, 1.3, 1.5, 1.7, 1.9]	15	29	0.4054	13.1391	53.2×
			[1.9, 1.7, 1.5, 1.3, 1.1]	15	29	0.3748	11.9373	40.2×
			[1.2, 1.8, 1.2, 1.8, 1.2]	15	29	0.3796	13.9693	43.2×

Figure 4 maps out the scaling behavior from 6D all the way up to 8D. Under these brutal setups, total execution times for the baseline spiral out of control rather than simply crashing from RAM capacity limits. Choking on severe cache thrashing and mounting iteration penalties, that older solver eventually forces a reliance on theoretical time projections just to finish the evaluation. Strict log-linear scaling naturally emerges from our VTP design. It delivers a highly workable engine to push through massive fractional PDEs without hitting those same hardware walls.

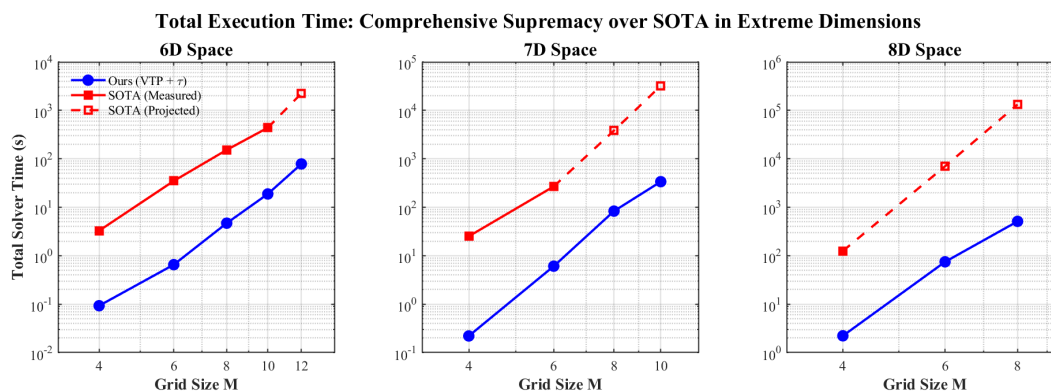


Figure 4. Total execution time scaling of the proposed framework versus the baseline in high-dimensional spaces (6D to 8D).

5. Conclusion

This paper set out to build an optimized, highly efficient fully discrete algorithmic framework for the multi-dimensional space-fractional Cahn-Hilliard equation. To capture the nonlocal phase separation physics with enough accuracy, we put together a high-order numerical scheme. The temporal side uses a second-order SAV-BDF2 discretization; the spatial side employs a sixth order Riesz fractional centered difference method that we derived from a pair of generating functions. We supplied rigorous theoretical proofs to back up two claims: the scheme is unconditionally energy stable, and it hits optimal spatio-temporal convergence rates inside the fully discrete setting. Beyond the discretization itself, we also had to deal with the severe ill-conditioning and the memory access bottlenecks that come with high-dimensional fractional operators. For that, we implemented a preconditioned MINRES solver and introduced a hardware-level optimization: the Vectorized Tensor Processing (VTP) architecture. The VTP idea is to reshape the tensors on the fly, turning multi-dimensional preconditioned matrix-vector products into batched transform operations that access memory contiguously. This pretty much wipes out cache thrashing. Extensive numerical experiments checked the theoretical convergence orders, confirmed that the discrete energy really does dissipate, and showed the substantial speedup the VTP-accelerated solver gets over conventional implementations. Taken together, these results back up the efficiency and the robustness of the framework we have laid out.

Author Contributions: Conceptualization, R.L. and Y.Y.; methodology, Y.Y.; software, M.Y.; validation, M.Y.; formal analysis, R.L. and M.Y.; investigation, R.L.; resources, Y.Y.; data curation, R.L.; writing—original draft preparation, R.L.; writing—review and editing, Y.Y.; visualization, R.L.; supervision, Y.Y.; project administration, Y.Y.; funding acquisition, Y.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China (Grant No. 12301628) and the Natural Science Foundation of Shandong Province (Grant No. ZR2017BA025).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data are contained within the article.

Acknowledgments: During the preparation of this manuscript, the author(s) used DeepSeek (DeepSeek-R1) for the purpose of Chinese-to-English translation and polishing the English language. The authors have reviewed and edited the output and take full responsibility for the content of this publication. The authors are also very grateful to the anonymous reviewers for their valuable comments and suggestions.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Wang, F.; Chen, H.; Wang, H. Finite element simulation and efficient algorithm for fractional Cahn-Hilliard equation. *J. Sci. Comput.* **2015**, *63*, 709–728.

2. Shimizu, I.; Takei, Y. Temperature and compositional dependence of solid-liquid interfacial energy: Application of the Cahn-Hilliard theory. *Physica B Condens. Matter* **2005**, *362*, 169–179.
3. Bertozzi, A.L.; Esedoglu, S.; Gillette, A. Inpainting of binary images using the Cahn-Hilliard equation. *IEEE Trans. Image Process.* **2007**, *16*, 285–291.
4. Bu, L.; Mei, L.; Wang, Y.; Hou, Y. Energy stable numerical schemes for the fractional-in-space Cahn-Hilliard equation. *Appl. Math. Comput.* **2020**, *370*, 124898.
5. Huang, X.; Li, D.; Sun, H.-W.; Li, C. Fast iterative solvers for multidimensional fractional Cahn-Hilliard equations. *J. Sci. Comput.* **2025**, *102*, 58.
6. Huang, X.; Li, D.; Sun, H.-W. Fast algorithms for fractional Cahn-Hilliard equations. *Appl. Math. Lett.* **2023**, *138*, 108510.
7. Akrivis, G.; Li, B.; Li, D. Energy-decaying extrapolated RK-SAV methods for the Allen-Cahn and Cahn-Hilliard equations. *SIAM J. Sci. Comput.* **2019**, *41*, A3703–A3727.
8. Yang, X. Linear, first and second-order, unconditionally energy stable numerical schemes for the phase field model of homopolymer blends. *J. Comput. Phys.* **2016**, *327*, 294–316.
9. Eyre, D.J. Unconditionally gradient stable time marching the Cahn-Hilliard equation. In *Computational and Mathematical Models of Microstructural Evolution*; Bullard, J.W., Chen, L.-Q., Kalia, R.K., Stoneham, A.M., Eds.; Materials Research Society: Warrendale, PA, USA, 1998; pp. 39–46.
10. Qian, Y.; Yang, Z.; Wang, F.; Dong, S. gPAV-based unconditionally energy-stable schemes for the Cahn-Hilliard equation: Stability and error analysis. *J. Comput. Phys.* **2020**, *418*, 109618.
11. Chen, C.; Yang, X. Fast, provably unconditionally energy stable, and second-order accurate algorithms for the anisotropic Cahn-Hilliard model. *J. Comput. Phys.* **2019**, *396*, 265–284.
12. Huang, X.; Lin, X.-L.; Ng, M.K.; Sun, H.-W. Spectral analysis for preconditioning of multi-dimensional Riesz fractional diffusion equations. *Numer. Linear Algebra Appl.* **2022**, *29*, e2434.
13. Huang, X.; Li, D.; Sun, H.-W.; Zhang, F. Preconditioners with symmetrized techniques for space fractional Cahn-Hilliard equations. *J. Sci. Comput.* **2022**, *90*, 30.
14. Li, D.; Quan, C.; Tang, T. Stability and convergence analysis for the implicit-explicit method to the Cahn-Hilliard equation. *Math. Comput.* **2022**, *91*, 785–809.
15. Mandolesi, B.; Iandiorio, C.; Belardi, V.G.; Vivio, F. Spinodal decomposition-inspired metamaterial: Tailored homogenized elastic properties via the dimensionless Cahn-Hilliard equation. *Eur. J. Mech. A Solids* **2025**, *112*, 105615.
16. Cheng, K.; Feng, W.; Wang, C.; Wise, S.M. An energy stable fourth order finite difference scheme for the Cahn-Hilliard equation. *J. Comput. Appl. Math.* **2019**, *362*, 574–595.
17. Liao, H.; Ji, B.; Zhang, L. An adaptive BDF2 implicit time-stepping method for the phase field crystal model. *IMA J. Numer. Anal.* **2022**, *42*, 649–679.
18. Kolda, T.G.; Bader, B.W. Tensor decompositions and applications. *SIAM Rev.* **2009**, *51*, 455–500.
19. Saberi-Movahed, F.; Berahmand, K.; Sheikhpour, R.; Li, Y.; Pan, S. Nonnegative matrix factorization in dimensionality reduction: A survey. *ACM Comput. Surv.* **2025**, *58*, 1–41.
20. Fawzi, A.; Balog, M.; Huang, A.; Hubert, T.; Romera-Paredes, B.; Barekatin, M.; Novikov, A.; Ruiz, F.J.R.; Schrittwieser, J.; Swirszcz, G.; et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **2022**, *610*, 47–53.
21. Hao, Z.; Zhang, Z.; Du, R. Fractional centered difference scheme for high-dimensional integral fractional Laplacian. *J. Comput. Phys.* **2021**, *424*, 109851.
22. Huang, X.; Lin, X.-L.; Ng, M.K.; Sun, H.-W. Spectral analysis for preconditioning of multi-dimensional Riesz fractional diffusion equations. *arXiv* **2021**, arXiv:2102.01371. Available online: <https://arxiv.org/abs/2102.01371> (accessed on Day Month Year).
23. Zhang, Y.; Shen, J. A generalized SAV approach with relaxation for dissipative systems. *J. Comput. Phys.* **2022**, *464*, 111311.
24. Metzger, S. A convergent SAV scheme for Cahn-Hilliard equations with dynamic boundary conditions. *IMA J. Numer. Anal.* **2023**, *43*, 3593–3627.
25. Guo, J.; Wang, C.; Yan, Y.; Wise, S.M. A refined convergence estimate for a fourth order finite difference numerical scheme to the Cahn-Hilliard equation. *arXiv* **2024**, arXiv:2404.04628. Available online: <https://arxiv.org/abs/2404.04628> (accessed on Day Month Year).
26. Yan, Y.; Chen, W.; Wang, C.; Wise, S.M. A second-order energy stable BDF numerical scheme for the Cahn-Hilliard equation. *Commun. Comput. Phys.* **2018**, *23*, 572–602.

27. Ju, L.; Li, X.; Qiao, Z. Stabilized exponential-SAV schemes preserving energy dissipation law and maximum bound principle for the Allen-Cahn type equations. *J. Sci. Comput.* **2022**, *92*, 66.
28. Huang, C.; Stynes, M. A sharp α -robust $L^\infty(H^1)$ error bound for a time-fractional Allen-Cahn problem discretised by the Alikhanov $L2-1_\sigma$ scheme and a standard FEM. *J. Sci. Comput.* **2022**, *91*, 43.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.