A general framework for high accuracy solutions to energy gradient flows from material science models

Andrew Christlieb Keith Promislow Brian Wetton

Mathematics Department
University of British Columbia
www.math.ubc.ca/∼wetton

SUSTech Seminar, May 25, 2017

Institute of Applied Mathematics

University of British Columbia

- Faculty participation from many departments.
- Interdisciplinary graduate programme.

Overview

- Introduction to energy gradient flows
- Numerical approximation
- Preconditioned conjugate gradient solution
- Performance on several examples
- Comparison to splitting methods
- (Higher order time stepping, GPU implementation)
- JCP **257** 193-215 (2014)

Introduction to energy gradient flows

Allen Cahn equation

$u(x,t)$, 2π-periodic in x solves

$$u_t = \epsilon^2 u_{xx} - u^3 + u$$

Allen and Cahn, Acta Metall 1979

- for discussion, consider $\epsilon = 0$
- A-C is then an autonomous ODE with fixed points $u = \pm 1$ (stable) and $u = 0$ (unstable) at each space location
- solutions tend to $u = \pm 1$ in O(1) time: spinodal evolution
- with $\epsilon > 0$ there is an interface of width O(ϵ) that is formed between the two phases

Introduction to energy gradient flows

Allen Cahn equation (cont.)

$$u_t = \epsilon^2 u_{xx} - u^3 + u$$

- Steady state solution

$$u = \tanh\left(\frac{x - x_0}{2\epsilon}\right)$$

- in higher dimensions, x_0 is replaced by the curve between the phases $u = \pm 1$ and $x - x_0$ is replaced by a normal distance to the curve
- in this case, the solution is approximate and the interface will move in a slow time scale: ripening evolution
- higher order asymptotic terms can determine a motion law for the interface: gamma limit
- For 2D A-C curves move with curvature motion as $\epsilon \to 0$ in an $O(\epsilon^{-2})$ time scale.

Introduction to energy gradient flows

Allen Cahn equation (cont.)

$$\begin{aligned} u_t &= \epsilon^2 u_{xx} - W'(u) \\ W(u) &= \frac{1}{4}(u^2-1)^2 \end{aligned}$$

- This equation is gradient flow on the energy

$$\mathcal{E} = \int_0^{2\pi} \left(\epsilon^2 u_x^2 + W(u)\right) dx$$

- This leads to a symmetric Jacobian matrix for the implicit time steps of the discretization
- Carr and Pego (1989) The linearization of a ripening state with M transition layers has M exponentially small (in ϵ) eigenvalues. The rest are negative and bounded away from zero.

Introduction to energy gradient flows

Cahn-Hilliard equation

$$u_t = -\epsilon^2 u_{xxxx} + (u^3 - u)_{xx}$$

Cahn and Hilliard, J Chem Phys 1958

- Same steady state solution as A-C

$$u = \tanh\left(\frac{x - x_0}{2\epsilon}\right)$$

- gradient flow on the same energy $\mathcal{E}$ but in the H_{-1} norm that has inner product

$$(u, v)_{H_{-1}} := (u, \Delta^{-1} v)$$

- the gamma limit is nonlocal, Mullins-Sekerka flow
- for a dilute phase, a later ripening evolution is known as Ostwald ripening
- C-H conserves the mass of the two phases

Introduction to energy gradient flows

1D Cahn-Hilliard solution energy

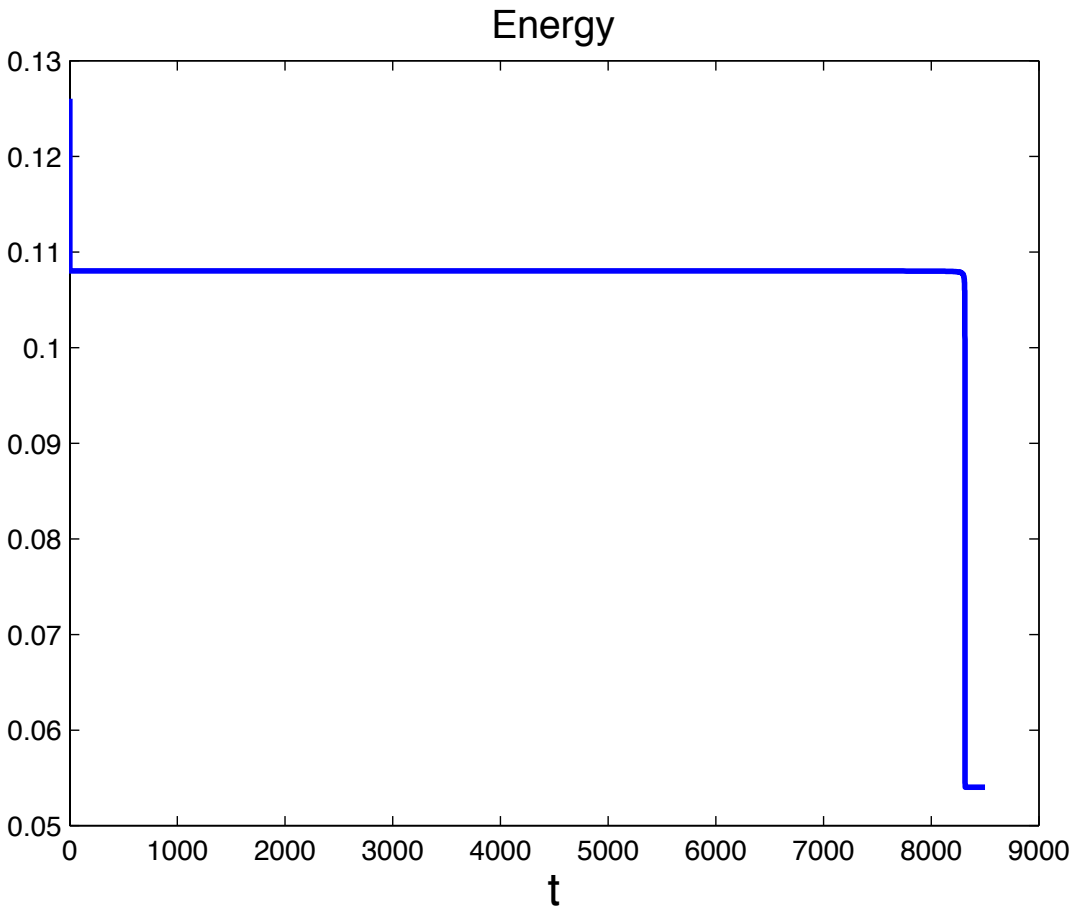

Numerical Approximation of Energy Gradient Flows

Goals

1. (Relatively) fast and accurate method
2. Easily adapted to different models:
   - Vector models
   - Additional terms and well shapes
   - Different PDE order

Spectral approximation in space, adaptive implicit time stepping

Note: I am considering Cahn-Hillard as the materials science model, not as a method for tracking material interfaces.

Conventional Wisdom

Much of it false

- Fully implicit time stepping always requires a small time step (false).
- The recommended energy stable (convex splitting) time stepping:
  - Any size time step can be taken: the implicit problem is convex (can even be linear) and any step is guaranteed to reduce energy (true).
  - Large time steps are very inaccurate.

Numerical Approximation

Spectral approximation in space

$$u_t = -\epsilon^2 u_{xxxx} + (u^3 - u)_{xx}$$

Dual form of the approximation:

$$\begin{aligned} u(jh,t) &\approx U_j(t), \quad j = 1, \ldots N \\ u(x,t) &\approx \sum_{\alpha=-N/2+1}^{N/2} \hat{U}_\alpha(t) e^{i\alpha x} \end{aligned}$$

- N is the number of spatial grid points
- $h = 2\pi/N$ is the grid spacing
- $\hat{\mathbf{U}}$ is the DFT of $\mathbf{U}$, $\hat{\mathbf{U}} = \mathcal{F}\mathbf{U}$

Approximation of RHS above (Λ_α is diag$(-\alpha^2)$):

$$-\mathcal{F}^{-1}\Lambda_\alpha\mathcal{F}(\epsilon^2\mathcal{F}^{-1}\Lambda_\alpha\mathcal{F}\mathbf{U} - \mathbf{U}^{<3>} + \mathbf{U})$$

Numerical Approximation

Implicit time stepping

$$u_t = -\epsilon^2 u_{xxxx} + (u^3 - u)_{xx}$$

Fully discrete approximation

$$u(jh, t_m) \approx U_j^m, \; j = 1, \ldots N \text{ and } j = 0, \ldots M$$

with time steps $k_m = t_m - t_{m-1}$

- explicit FE predictor (here $\Delta_h = \mathcal{F}^{-1}\Lambda_\alpha \mathcal{F}$):

$$\mathbf{U}^* = \mathbf{U}^{m-1} - k_m \Delta_h \left[\epsilon^2 \Delta_h \mathbf{U}^{m-1} - W'(\mathbf{U}^{m-1})\right]$$

- implicit BE step

$$\mathbf{G}(\mathbf{U}^m) := \mathbf{U}^m + k_m \Delta_h \left[\epsilon^2 \Delta_h \mathbf{U}^m - W'(\mathbf{U}^m)\right] - \mathbf{U}^{m-1} = \mathbf{0}$$

Numerical Approximation

Adaptive time stepping

- Prescribe a tolerance δ for the error for *each* time step
- For the BE step the error is approximately $\sigma = \|u_{tt}\| k_m^2/2$
- $\sigma \approx \|\mathbf{U}^m - \mathbf{U}^*\|/2$
- Reject the step and repeat with time step $k_m/2$ if
  - if $\sigma > \delta$ *or*
  - Newton iterations fail to converge *or*
  - $\mathcal{E}$ increases
- Otherwise accept the step and take

$$k_{m+1} = 0.8 k_m \sqrt{\frac{\delta}{\sigma}}$$

Numerical Approximation

PCG solution of the implicit system

$$\mathbf{G}(\mathbf{U}^m) := \mathbf{U}^m + k_m \Delta_h \left[\epsilon^2 \Delta_h \mathbf{U}^m - W'(\mathbf{U}^m)\right] - \mathbf{U}^{m-1} = \mathbf{0}$$

- Newton's method with symmetric (in H^{-1}) Jacobian matrix

$$\mathcal{J} = I + k_m \Delta_h (\epsilon^2 \Delta_h - \Lambda_2)$$

where Λ_2 is the diagonal matrix with entries

$$W''(U_j^{(r)}) = 3[U_j^{(r)}]^2 - 1.$$

- Symmetric preconditioner (Scott Maclachlan and Zhengfu Xu)

$$\mathcal{Q} = I + k_m \Delta_h (\epsilon^2 \Delta_h - 2k_m I)$$

- $\mathcal{J}$ and $\mathcal{Q}^{-1}$ are dense but can multiply by these matrices efficiently

Numerical Approximation

Basic numerical tests

1D Cahn-Hilliard model, $\epsilon = 0.18$, fixed time step computations to time $t = 0.2$.

$E_k = \|U_k - U_{k/2}\|$, results for $N = 128$:

k	E_k
2e-4	1.32e-5
1e-4	6.6e-6
5e-5	3.3e-6

$E_k = \|U_N - U_{2N}\|$, results for $k = 1e-4$:

N	E_N	E_N for $\epsilon = 0.09$
32	2.0e-3	0.139
64	9.3e-7	4.4e-3
128	9.0e-13	1.3e-6

First order convergence in time, spectral in $N = O(1/\epsilon)$.

Numerical Approximation

Tests of adaptive time stepping

1D Cahn-Hilliard, $\epsilon = 0.18$, $N = 128$ to time $t = 8500$

δ	time steps	ripening time	total CG
1e-4	848	8180	19,105
1e-5	2580 (3.04)	8273	39,942 (2.09)
1e-6	8072 (3.13)	8304	87,563 (2.19)
1e-7	25446 (3.15)	8314	227,799 (2.60)

- Confirmation of adaptive time stepping strategy ($\sqrt{10} \approx 3.16$)
- Solver improves as $k \to 0$
- Ripening times can be approximated accurately
- solver iterations are independent of N

Numerical Approximation

Tests of adaptive time stepping (cont.)

1D Cahn-Hilliard, $\epsilon = 0.18$, $N = 128$, $\delta = 1e-4$ to time $t = 8500$

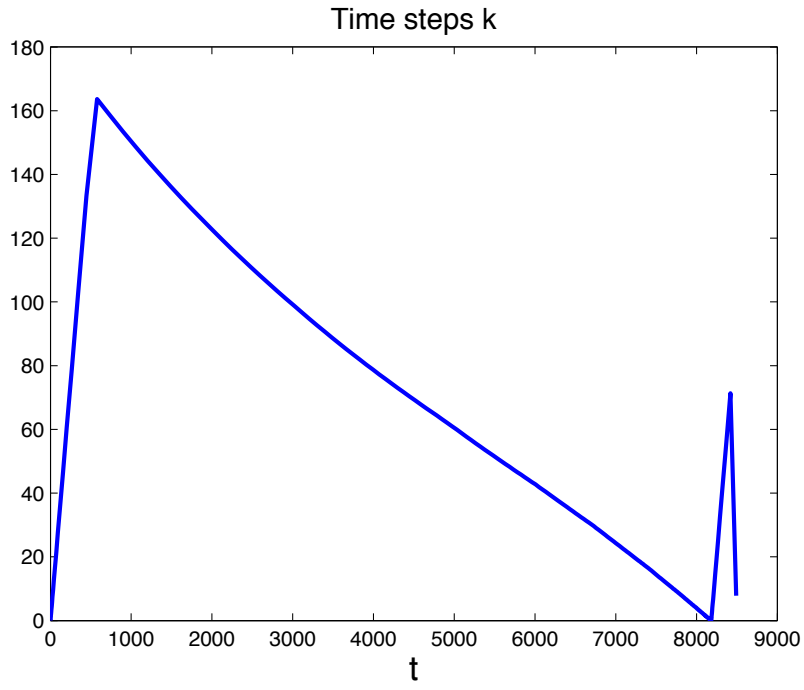

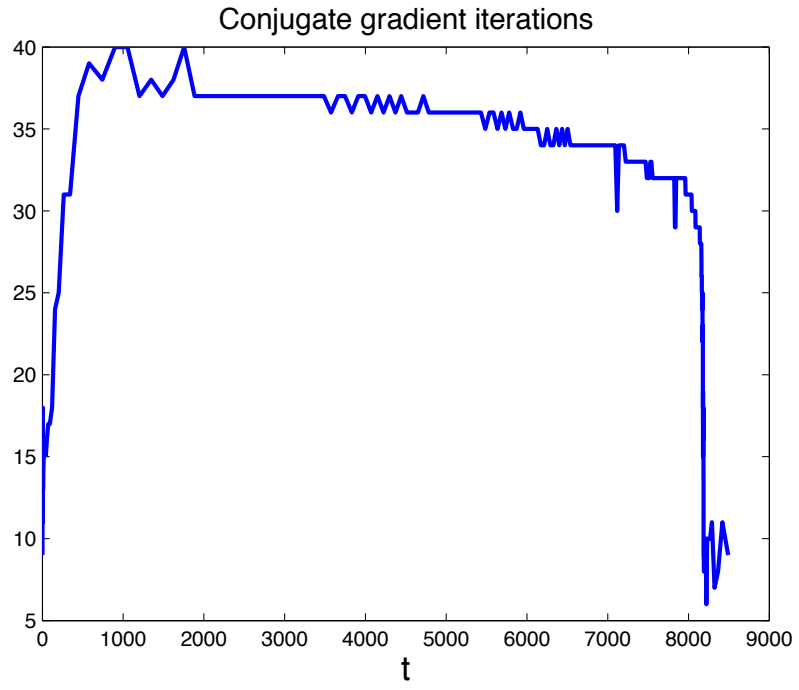

Preconditioned conjugate gradient solver

spectra

$$\begin{aligned}\mathcal{J}_{CH} &= I + k_m\epsilon^2\Delta_h\Delta_h - k_m\Delta_h\Lambda_2 \\ \mathcal{Q}_{CH} &= I + k_m\epsilon^2\Delta_h\Delta_h - 2k_m\Delta_h.\end{aligned}$$

Spectra of $\mathcal{A} = \mathcal{Q}^{-1}\mathcal{J}$ determines the performance of the CG iterations, order $\sqrt{\kappa(\mathcal{A})}$ iterations to a given tolerance

k	$\epsilon = 0.16$	$\epsilon = 0.08$	$\epsilon = 0.04$
1	14.3	23.7	43.5
2	27.2	44.8	82.6
4	52.4	87.7	160
8	103	172	316

Suggests condition number $O(k/\epsilon)$ in the ripening regime ($O(k)$ independent of ϵ for AC), confirmed by formal asymptotics.

Preconditioned conjugate gradient solver

spectra (cont.)

1D C-H, ripening state $\epsilon = 0.18$, $k = 10$

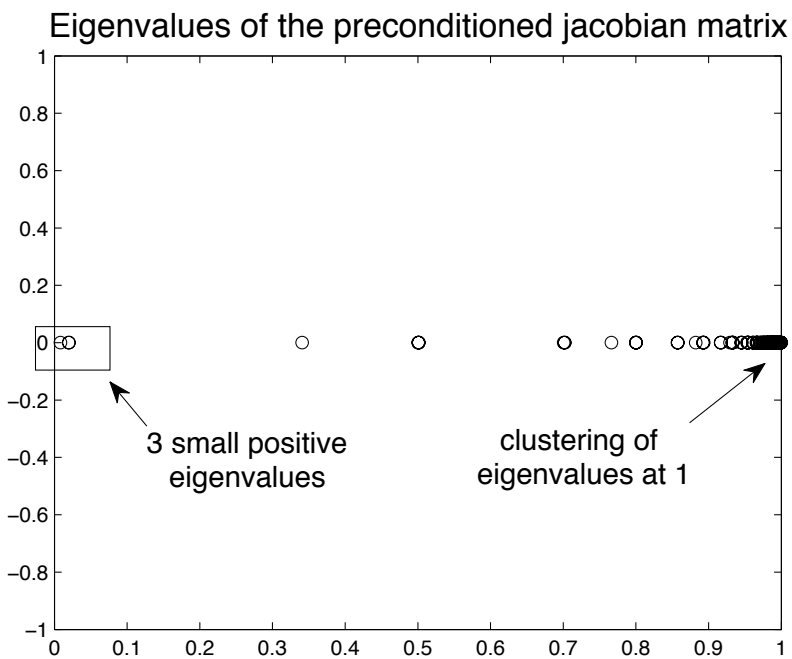

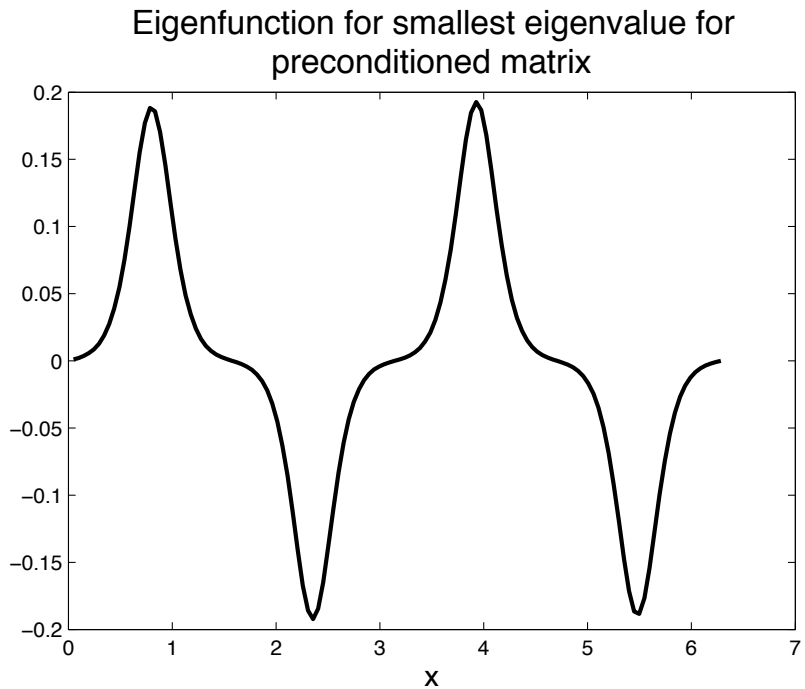

Theorem on the rank of a modified distance matrix for the asymptotics for the Cahn-Hilliard case proved in Cheng, Li, Shirokoff, W, Journal of Statistical Physics, **166**, 1029-1035 (2017).

Preconditioned conjugate gradient solver

Asymptotics for Allen-Cahn case (easier)

Eigenvalues σ and eigenvectors ψ of $\mathcal{A} = \mathcal{Q}^{-1}\mathcal{J}$

$$(I - k\mathcal{L})\psi = \sigma[I - k(\mathcal{L} - 3(u^2 - 1))]\psi$$

- Recall M dimensional $\mathcal{V}$, span of exponentially small eigenvalues of $\mathcal{L}$.
- Consider asymptotics in powers of (large) k, $\sigma = \beta/k$.

$$\begin{aligned} O(k): &\quad \mathcal{L}\psi = 0 \\ O(1): &\quad \psi = 3\beta(1 - u^2)\psi \quad \text{in } \mathcal{V} \text{ component} \end{aligned}$$

Satisfying the $O(1)$ term above gives $\beta \approx 0.41$, matches numerically computed eigenvalues $\approx 0.41/k$.

Performance on other models

Sixth order model

$$u_t = \Delta \left(\epsilon^2 \Delta - W''(u) + \epsilon^2 \eta\right) \left(\epsilon^2 \Delta u - W'(u)\right)$$

where $\eta > 0$ (promotes the formation of phase interface)

$$\begin{aligned}
\mathcal{J}_6 &= I - k_m \epsilon^4 \Delta_h \Delta_h (\Delta_h + \eta I) + k_m \epsilon^2 \Delta_h \Delta_h \Lambda_2 \\
&\quad + k_m \epsilon^2 \Delta_h (\Lambda_L \Lambda_3 + \Lambda_2 \Delta_h - \eta \Lambda_2) \\
&\quad + k_m \Delta_h (\Lambda_2^2 + \Lambda_1 \Lambda_3) \\
\mathcal{Q}_6 &= I - k_m \epsilon^4 \Delta_h \Delta h \Delta_h + k_m (\eta \epsilon^4 - 4\epsilon^2) \Delta_h \Delta_h \\
&\quad - k_m (4 - 2\epsilon^2 \eta) \Delta_h
\end{aligned}$$

where here Λ_i, $i = 1, 2, 3$ is the diagonal matrix with entries

$$\frac{d^i W}{du^i}(U_j^{(r)})$$

and Λ_L is the diagonal matrix with entries $\Delta_h U^{(r)}$.

Performance on other models

Sixth order model: Model of amphiphilic materials

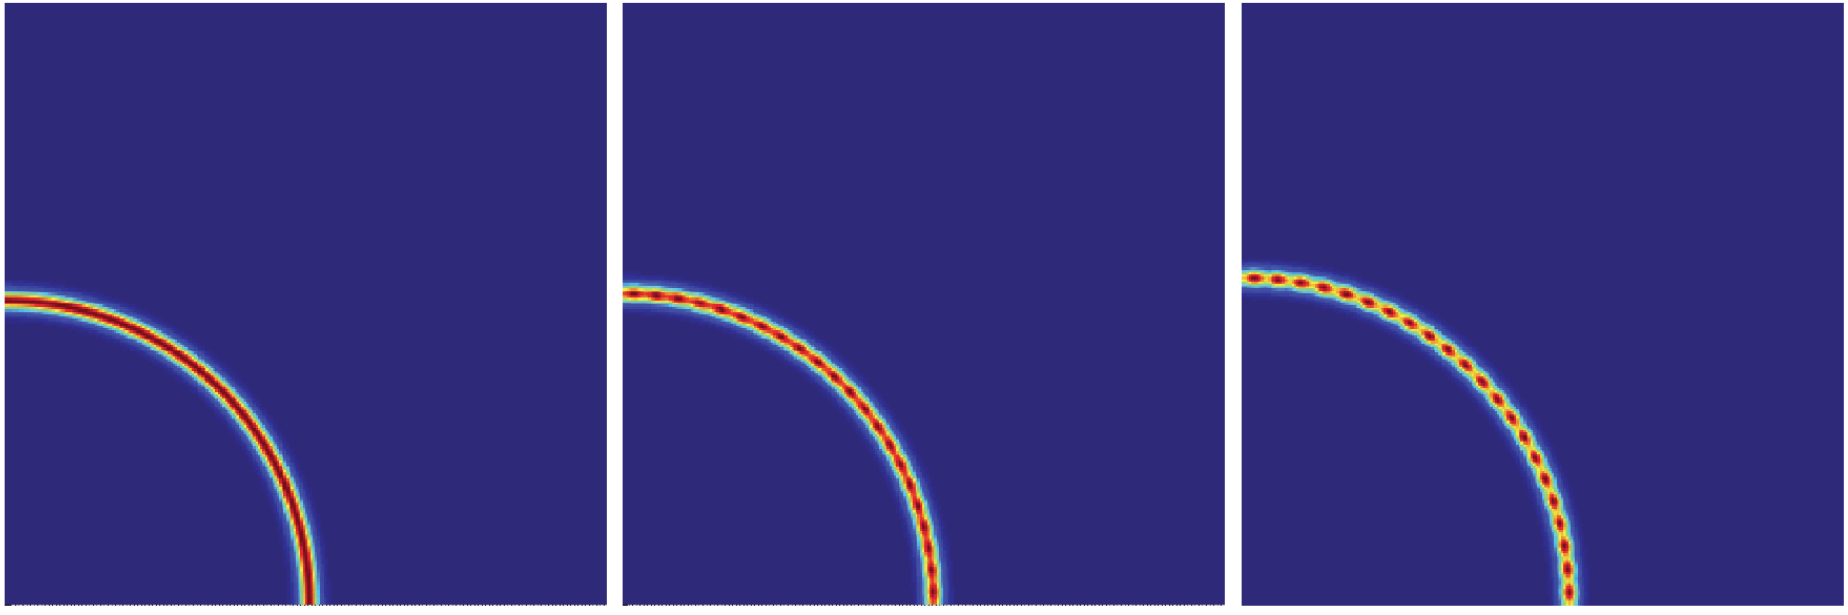

Doelman et. al., SIMATH 46, 3640-3677 (2014)

Performance on other models

Vector model

Vector $\mathbf{u} = (u, v)$

$$\mathbf{u}_t = -\epsilon^2 \Delta\Delta\mathbf{u} + \Delta\nabla_{\mathbf{u}} W(\mathbf{u}) \tag{1}$$

where here

$$W(\mathbf{u}) = \prod_{i=1}^{3} |\mathbf{u} - \mathbf{u}_i|^2 \tag{2}$$

and $\mathbf{u}_i$ are the points in the (u, v) plane that correspond to the cube roots of unity. This is a volume preserving model that forms symmetric triple junctions between three phases.

Comparison to splitting methods

Eyer's method in Allen-Cahn Framework

our implicit BE step

$$\mathbf{G}(\mathbf{U}^m) := \mathbf{U}^m - k_m \left[\epsilon^2 \Delta_h \mathbf{U}^m - \mathbf{U}^{m,<3>} + \mathbf{U}^m\right] - \mathbf{U}^{m-1} = \mathbf{0}$$

Convex-concave splitting method Eyer, 1998

$$\mathbf{G}_E(\mathbf{U}^m) := \mathbf{U}^m - k_m \left[\epsilon^2 \Delta_h \mathbf{U}^m - \mathbf{U}^{m,<3>}\right] - (1 + k_m)\mathbf{U}^{m-1} = \mathbf{0}$$

The splitting technique has some desirable properties

- Unique solution for any k that decreases energy $\mathcal{E}$
- Condition number of PCG iterations independent of ϵ and k (MSU fixed point method)

But suffers from poor accuracy for large time steps. Local error size $\lambda^2 k^2$ for backward Euler compared to λk^2 for Eyre's splitting where λ can be exponentially small in ϵ (1D problems).

Comparison to splitting methods

Large errors from splitting: simple explanation

Consider the Dahlquist test problem for $u(t)$:

$$\dot{u} = -\gamma u$$

with $\gamma \ll 1$ and $\gamma = \alpha - \beta$ with α, β size $O(1)$.
Now compare over one time step:

Exact: $e^{-\gamma k} \approx 1 - \gamma k + \frac{1}{2}\gamma^2 k^2$

BE: $1/(1+\gamma k) \approx 1 - \gamma k + \gamma^2 k^2$, error $\gamma^2 k^2/2$.

Eyer: $(1+\beta k)/(1-\alpha k) \approx 1 - \gamma k + \gamma\alpha k^2$, error $\gamma\alpha k^2$.

Controversial statement: All the split step schemes I have looked at suffer from this loss of accuracy and hence are unsuitable for accurate simulation of meta-stable dynamics.

Comparison to splitting methods

Splitting methods are inaccurate

Our implicit BE step applied to 1D C-H model (reprise)

δ	time steps	ripening time	total CG
1e-4	848	8180	19,105
1e-5	2580 (3.04)	8273	39,942 (2.09)
1e-6	8072 (3.13)	8304	87,563 (2.19)
1e-7	25446 (3.15)	8314	227,799 (2.60)

Eyer's splitting method (similar poor performance from other IMEX methods)

δ	time steps	ripening time	total CG
1e-4	70,517	13147	1,039,676
1e-5	202,549 (2.87)	9582	2,368,051 (2.27)
1e-6	618,431(3.06)	8695	5,205,739 (2.19)

Comparison to splitting methods

Splitting methods are inaccurate (cont.)

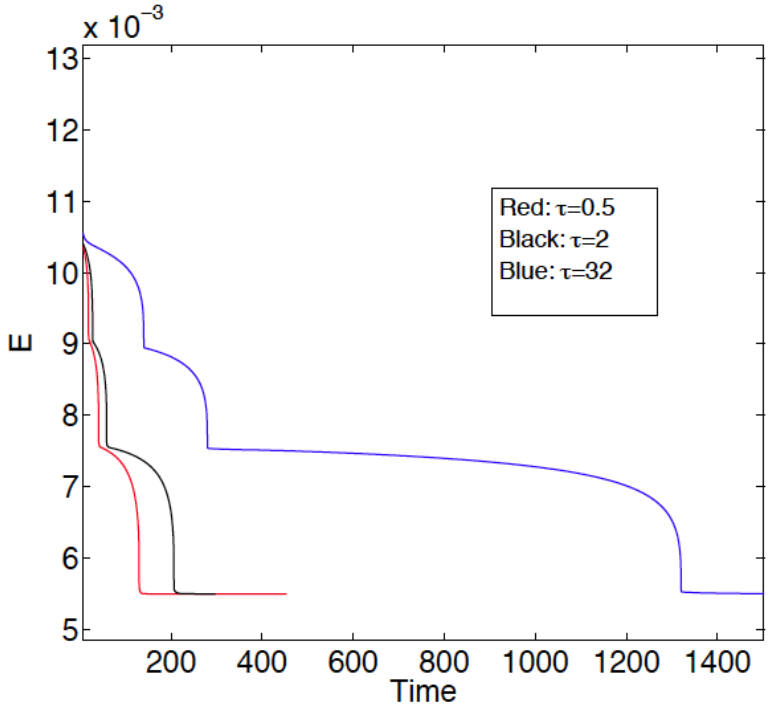

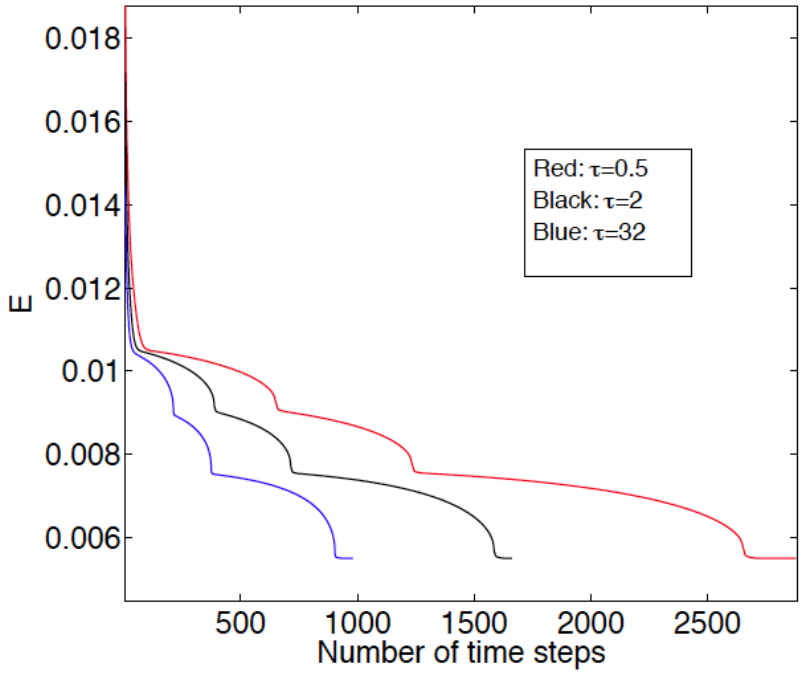

Christlieb et. al., Commun. Math. Sci. **11**, 345-360 (2013)

Summary

1. General framework for solving energy gradient problems from materials science
2. Reasonably efficient, scales reasonably with ϵ
   - $N = O(1/\epsilon)$, necessary for a uniform grid approximation
   - $O(1/\sqrt{\epsilon})$ solver iterations per time step (for C-H)
3. Easy to adapt to new problems
4. Lacks theoretical guarantees but in practice much more efficient than widely used splitting methods

Additional:

- Higher order BDF time stepping (Mark Willoughby, MSc thesis)
- GPU implementation (Jaylan Jones, MSU PhD student)

3D GPU implementation

Method fits very well into GPU computational framework (FFT, diagonal multiplication), but there are limitations on the size of the problem (128^3 with unit from last year).

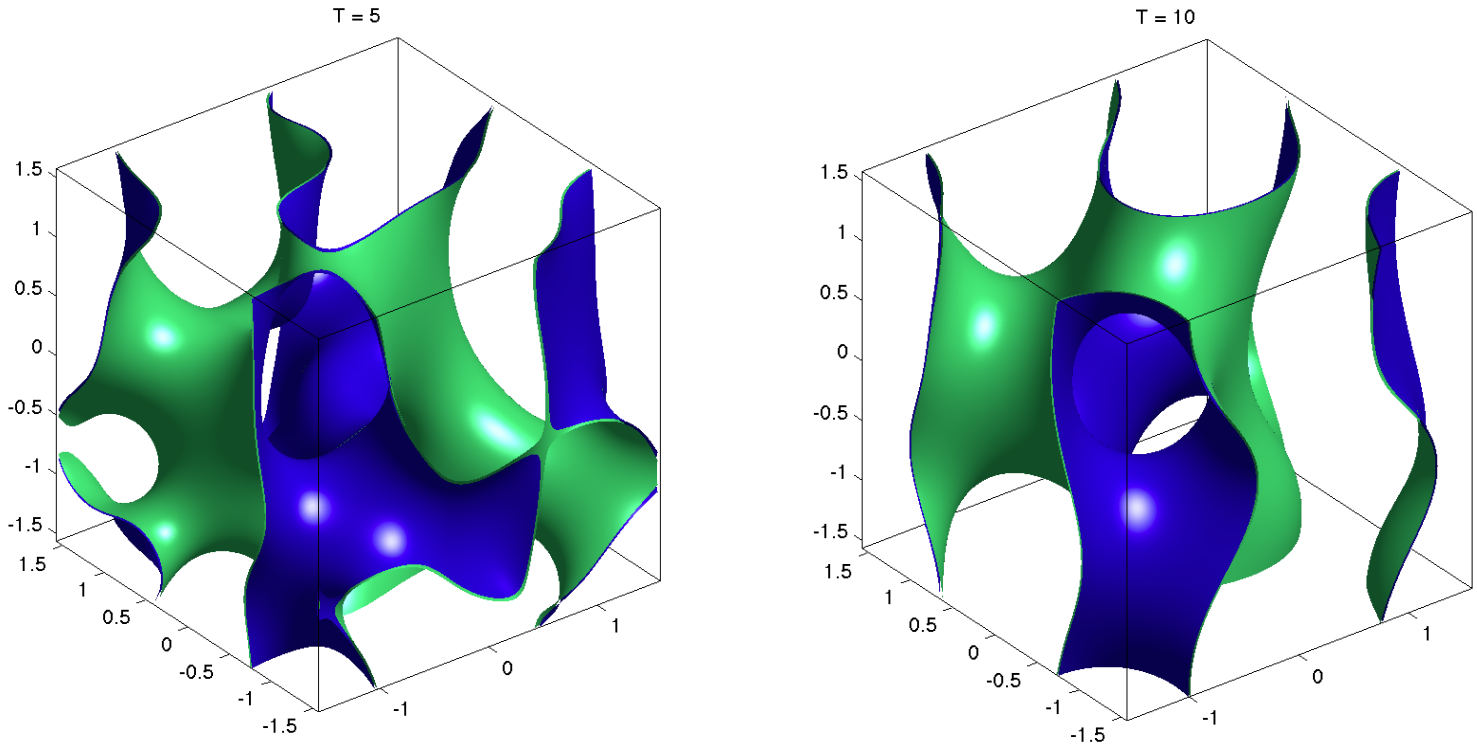

Speedup of 6.5 times to two Quad-core Intel Xeon E5620 processors.