Accelerating Large-scale Phase Field Simulation with GPU

Jian Zhang

Computer Network Information Center(CNIC),
Chinese Academy of Sciences

中国科学院
计算机网络信息中心
Computer Network Information Center,
Chinese Academy of Sciences

Outline

Background

- Phase Field Model
- Large Scale Smiulations

Compute intensive large time step algorithm

- CETD Schemes
- Localized Exponential Integration

Acceleration on heterogeneous platform

- GPU
- Sunway TaihuLight, MIC

Summary

Background

中国科学院
计算机网络信息中心
Computer Network Information Center,
Chinese Academy of Sciences

Micro-structures in Materials

Micro-structures: meso-scale morphological patterns

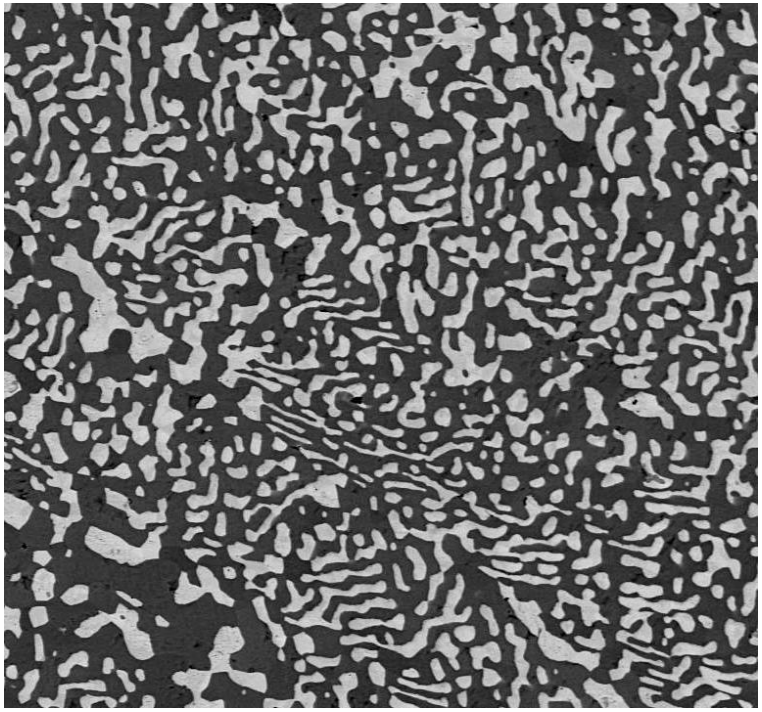

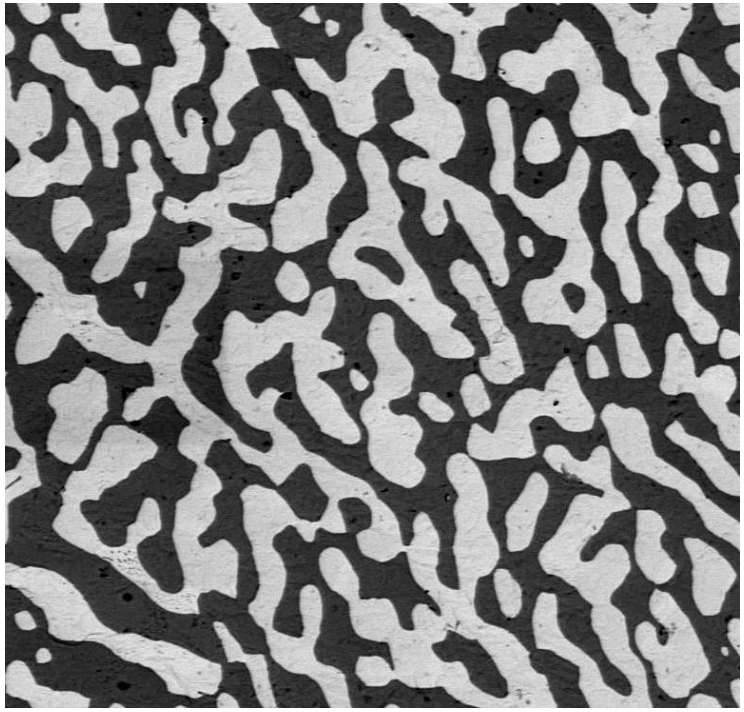

Micro-structures in Materials

Fatigue Failure

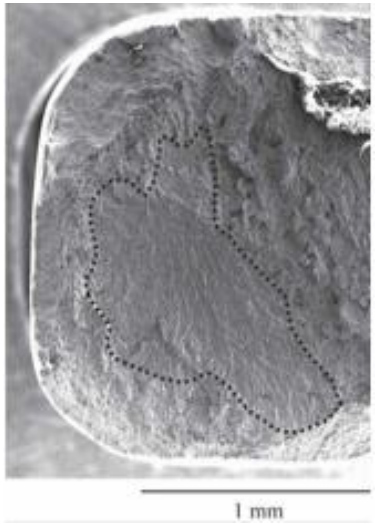

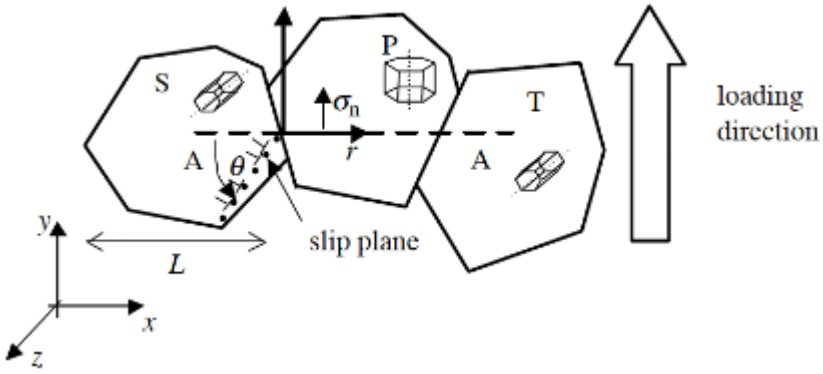

loading direction

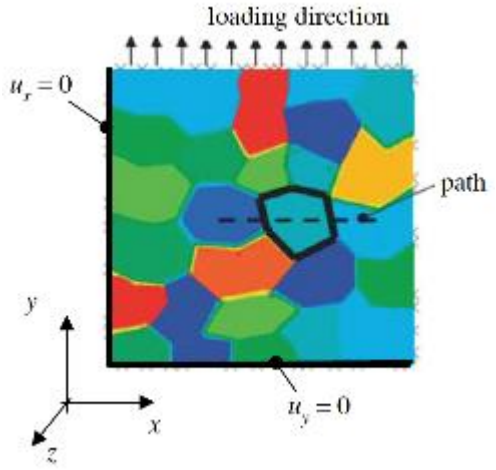

相场模型

Given a free energy functional $E(\phi)$, the gradient flow in L^2:

$$\frac{\partial \phi}{\partial t} = -\frac{\partial E(\phi)}{\partial \phi};$$

or the gradient flow in H^{-1}:

$$\frac{\partial \phi}{\partial t} = \Delta \frac{\delta E(\phi)}{\delta \phi}.$$

It is easy to see that they satisfy, respectively the following energy laws:

$$\frac{d}{dt}E(\phi) = -\|\frac{\delta E(\phi)}{\delta \phi}\|^2,$$

or

$$\frac{d}{dt}E(\phi) = -\|\nabla \frac{\delta E(\phi)}{\delta \phi}\|^2.$$

- 相场、成分场等
- 梯度流系统

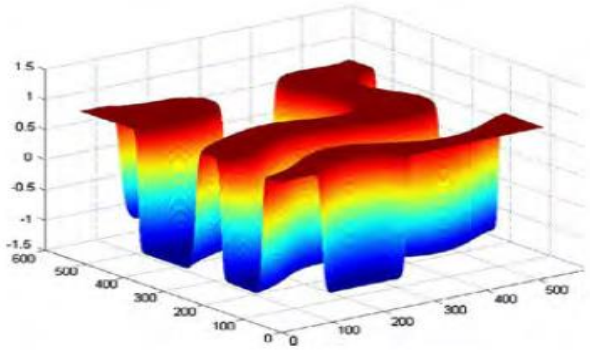

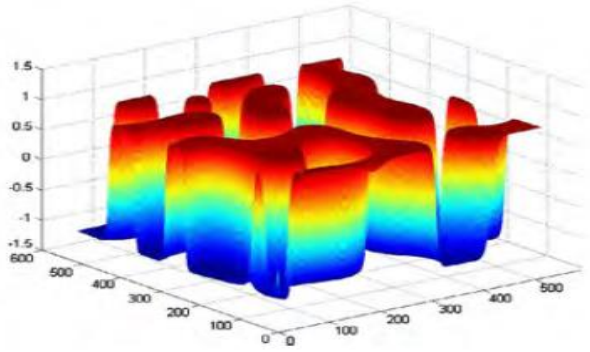

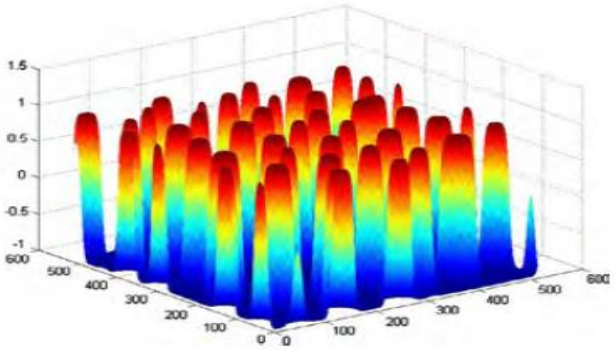

Phase Field Model

Given a free energy functional $E(\phi)$, the gradient flow in L^2:

$$\frac{\partial \phi}{\partial t} = -\frac{\partial E(\phi)}{\partial \phi};$$

or the gradient flow in H^{-1}:

$$\frac{\partial \phi}{\partial t} = \Delta \frac{\delta E(\phi)}{\delta \phi}.$$

It is easy to see that they satisfy, respectively the following energy laws:

$$\frac{d}{dt}E(\phi) = -\|\frac{\delta E(\phi)}{\delta \phi}\|^2,$$

or

$$\frac{d}{dt}E(\phi) = -\|\nabla \frac{\delta E(\phi)}{\delta \phi}\|^2.$$

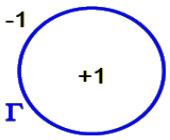

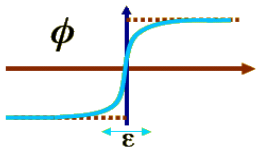

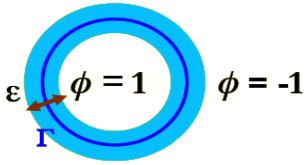

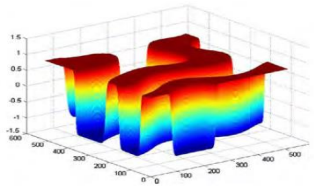

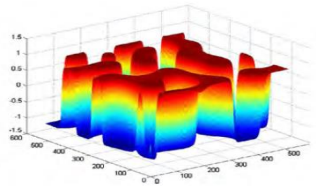

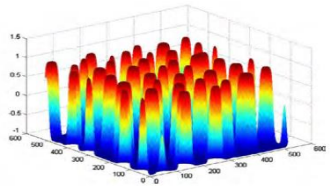

Explicit time marching- small time step

Allen-Cahn (AC) equation

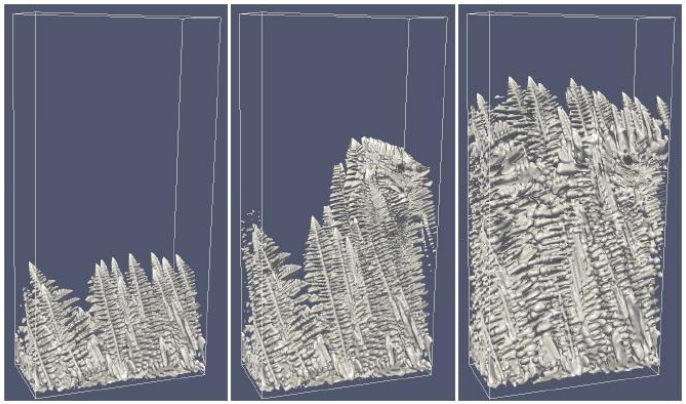

Takashi Shimokawabe et. al. , SC2011.
4 ×10^9 cells, TSUBAME 2.0.

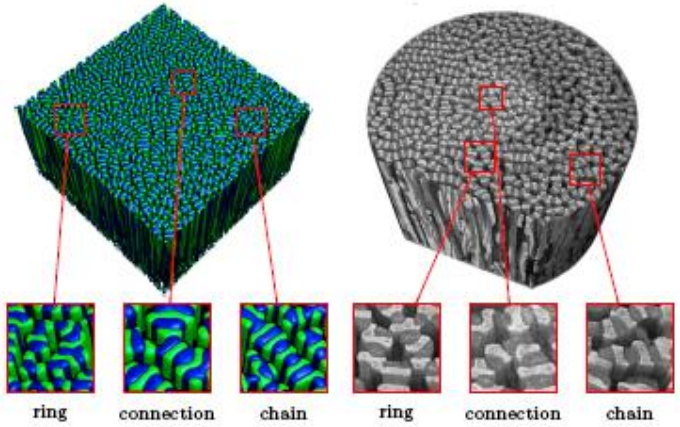

Martin Bauer et. al., SC2015. 8 ×10^9 cells.
SuperMUC, Hornet and JUQUEEN.

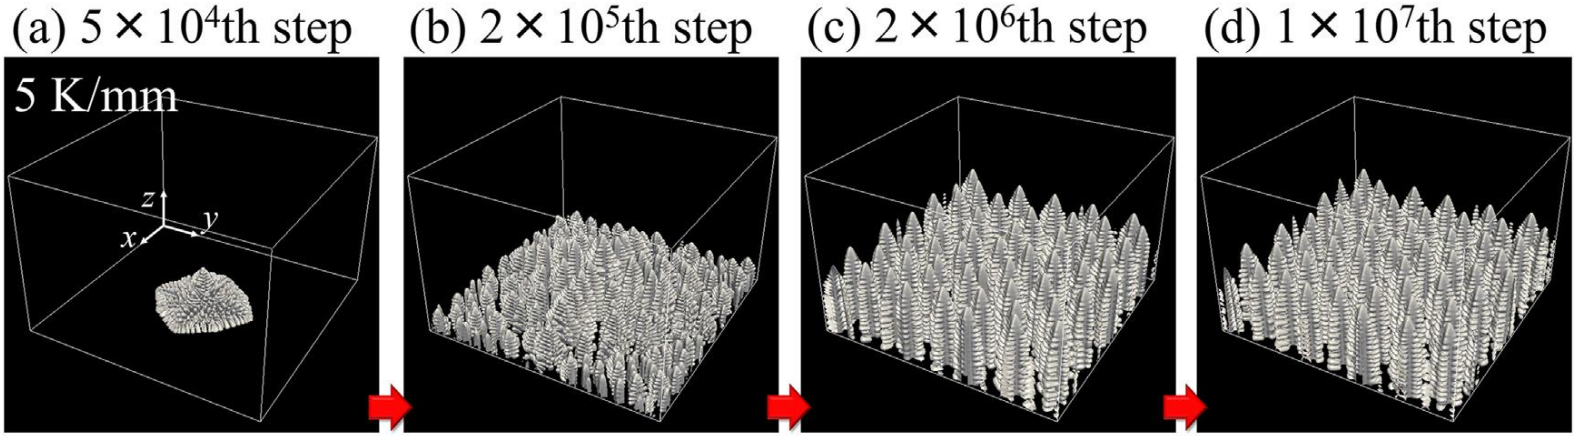

Tomohiro Takaki et. al. Acta Materialia, 2016. 4 ×10^9 cells, TSUBAME 2.5.

Energy stability

To fix the idea, we let $E(\phi) = \int_\Omega [\frac{1}{2}|\nabla\phi|^2 + \frac{1}{\eta^2}F(\phi)]dx$, where $F(\phi)$ is a general nonlinear free energy, and consider the gradient flow in H^{-1}:

$$\phi_t = \nabla \cdot \nabla \frac{\delta E}{\delta \phi}, \qquad \partial_n w|_{\partial\Omega} = 0;$$

$$w = \frac{\delta E}{\delta \phi} = -\Delta\phi \quad + \frac{1}{\eta^2} f(\phi), \qquad \partial_n \phi|_{\partial\Omega} = 0,$$

which satisfies the energy law:

$$\frac{\partial}{\partial t}\int_\Omega \left(\frac{1}{2}|\nabla\phi|^2 + \frac{1}{\eta^2}F(\phi)\right) = -\int_\Omega |\nabla(-\Delta\phi + \frac{1}{\eta^2}f(\phi))|^2.$$

Goal: Design simple, efficient and accurate numerical schemes that satisfy a discrete energy law.

Large Scale Phase Field Simulations

AC equation, explicit time marching

Small time step-size
Integration scheme design, easy
Stencil computing
performance ~ 25% peak

Large scale simulation ~10 billion cells

CH equation, implicit time marching

Large time step-size
Integration scheme design, hard
Multi-level preconditioner-solver
performance < 10% peak

Large scale simulation ~ 0.1 billion cells

The limited resolution in 3D simulations(CH) constitutes bottlenecks in validating predictions based on the phase field approach.

Accurate large-time-step marching scheme, scalability, efficiency

Compute intensive large time step algorithm

中国科学院
计算机网络信息中心
Computer Network Information Center,
Chinese Academy of Sciences

Exponential Time Differencing (ETD)

$$u_t = Lu + N(u,t).$$

$$u(t_{n+1}) = e^{L\Delta t}u(t_n) + e^{L\Delta t}\int_0^{\Delta t} e^{-Ls}N(u(t_n+s), t_n+s)ds.$$

exact integration polynomial approx.

Stable large time step-size

Exact integration & proper splitting of L and N

High order accuracy

Multi-step, prediction-correction, Runge-Kutta

Second order ETD scheme

$$\begin{cases} \frac{\partial u}{\partial t} = \nabla \cdot \big(M(u)\nabla(-\gamma\Delta u + u^3 - u)\big), & \mathbf{x} \in \Omega, t \in [t_0, t_0 + T], \\ u|_{t=t_0} = u_0, & \mathbf{x} \in \Omega, \end{cases}$$

$$\frac{\partial u}{\partial t} = -C(\gamma\Delta^2 u) - \nabla \cdot \big((M(u) - C)\nabla(\gamma\Delta u)\big) + \nabla \cdot \big(M(u)\nabla(u^3 - u)\big).$$

$$C \geq C - M((\mathbf{U})_{i,j}) \geq 0 \qquad C = \max_{i,j} M((\mathbf{U})_{i,j}).$$

$$P_y^{-1} ⓨ P_x^{-1} ⓧ \eta^\star = P_y^{-1} ⓨ P_x^{-1} ⓧ \eta_n \odot (e^{\odot})^{-H_1 \Delta t} - P_y^{-1} ⓨ P_x^{-1} ⓧ (Lf(\eta_n) - Lk_1\eta_n) \odot S^1_{r,s}$$

$$P_y^{-1} ⓨ P_x^{-1} ⓧ \eta_{n+1} = P_y^{-1} ⓨ P_x^{-1} ⓧ \eta_n \odot (e^{\odot})^{-H_2 \Delta t} - P_y^{-1} ⓨ P_x^{-1} ⓧ \left(\frac{Lf(\eta_n) - Lk_2\eta_n}{2}\right) \odot S^2_{r,s} - P_y^{-1} ⓨ P_x^{-1} ⓧ \left(\frac{Lf(\eta_\star) - Lk_2\eta_\star}{2}\right) \odot S^2_{r,s}$$

$$(H_1)_{ij} = -L[\epsilon(d_i^x + d_j^y) - k_1], \quad (H_2)_{ij} = -L[\epsilon(d_i^x + d_j^y) - k_2], \quad S^i_{r,s} = \frac{1 - (e^{\odot})^{-H_i \Delta t}}{H_i}$$

Unconditionally Energy Stable

Time Integration Accuracy

High order accuracy in time is important for simulating coarsening dynamics with large-time-step schemes.

Time step-size can be 10-100X than 1st order implicit schemes; > 4 orders of magnitude larger than explicit Euler scheme.

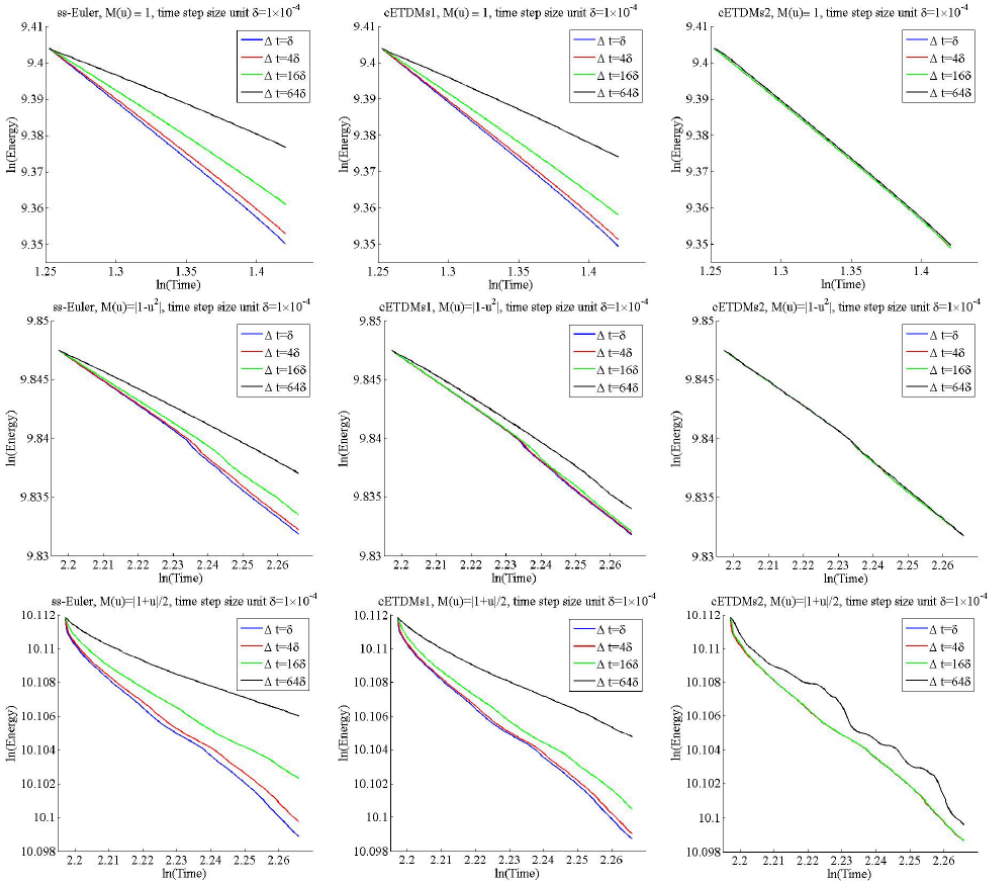

1st order stabilized semi-implicit Euler

1st order cETD

2nd order cETD

Extensive numerical experiments can be found in

"Fast and accurate algorithms for simulating coarsening dynamics of Cahn–Hilliard equations", Computational Materials Science, 108 (2015), pp 272-282.

Example

$$F_c = \int_V \mathrm{d}\mathbf{r}\left[\frac{1}{2}\alpha|\nabla c|^2 + \frac{1}{2}\beta|\nabla\eta|^2 + f(c,\eta)\right],$$

$$f(c,\eta) = \tfrac{1}{2}A(c-c')^2 - \tfrac{1}{2}B\eta^2 - \tfrac{1}{4}C\eta^4 + \tfrac{1}{6}D\eta^6 + \tfrac{1}{2}Gc\eta^2$$

$$\frac{\partial c(\mathbf{r},t)}{\partial t} = M\nabla^2\frac{\delta F}{\delta c(\mathbf{r},t)},$$

$$\frac{\partial \eta(\mathbf{r},t)}{\partial t} = -L\frac{\delta F}{\delta \eta(\mathbf{r},t)},$$

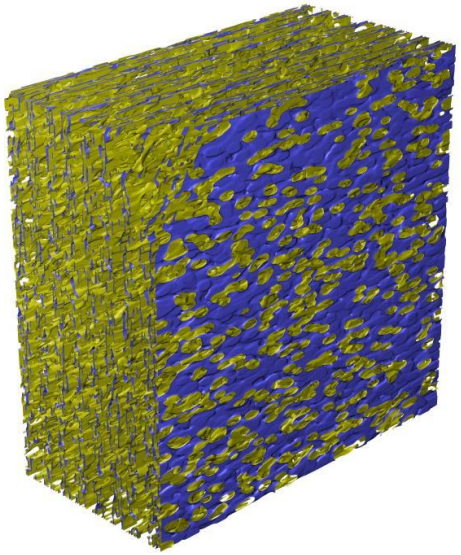

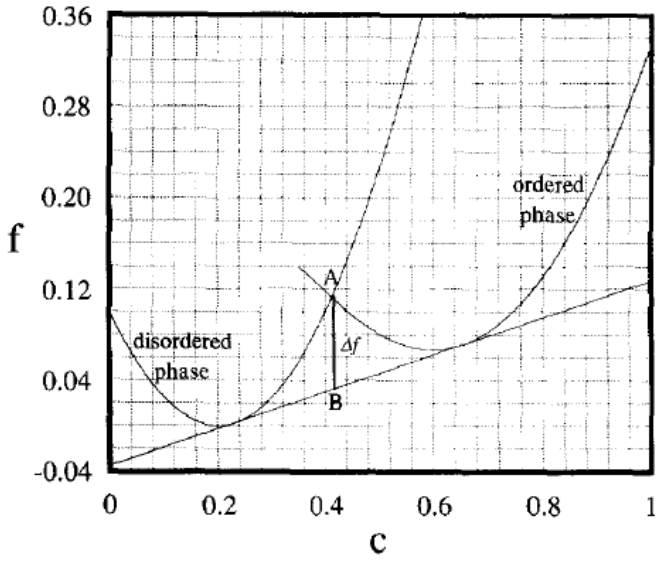

Fig. 1. Specific free energy vs composition curves for both the ordered and disordered solid solutions calculated according to equation (3) with $\eta = \eta_0(c)$ and $A = 62.5$, $B = 15.0$, $C = 12.0$, $D = 6.25$ and $G = 25.0$. See text for explanations.

Example

$$\frac{\partial \eta}{\partial t} = -L\left(\begin{array}{l} h'(\eta)\left(F^{v}(c_v,c_i)-F^{m}(c_v,c_i)\right)-k_\eta{}^2\nabla^2\eta-\nabla\left(k_\eta{}^2\right)\cdot\nabla\eta- \\ \dfrac{\partial\left((\nabla\eta)^2\left(k_\eta\right)\left(\left(k_{\eta\varphi}{}'\right)\left(-\dfrac{\eta_y}{(\eta_x)^2+(\eta_y)^2}\right)+\left(k_{\eta\theta}{}'\right)\left(\dfrac{\eta_x\eta_z}{(\nabla\eta)^2\sqrt{(\eta_x)^2+(\eta_y)^2}}\right)\right)\right)}{\partial x}- \\ \dfrac{\partial\left((\nabla\eta)^2\left(k_\eta\right)\left(\left(k_{\eta\varphi}{}'\right)\left(\dfrac{\eta_x}{(\eta_x)^2+(\eta_y)^2}\right)+\left(k_{\eta\theta}{}'\right)\left(\dfrac{\eta_y\eta_z}{(\nabla\eta)^2\sqrt{(\eta_x)^2+(\eta_y)^2}}\right)\right)\right)}{\partial y}- \\ \dfrac{\partial\left((\nabla\eta)^2\left(k_\eta\right)\left(\left(k_{\eta\theta}{}'\right)\left(-\dfrac{\sqrt{(\eta_x)^2+(\eta_y)^2}}{(\nabla\eta)^2}\right)\right)\right)}{\partial z} \end{array}\right)$$

$$\frac{\partial c_v}{\partial t} = \nabla\cdot\left(\left(M_v\right)\nabla\left(\left(1-h(\eta)\right)\frac{\partial\left(F^{m}(c_v,c_i)\right)}{\partial c_v}+h(\eta)\frac{\partial\left(F^{v}(c_v,c_i)\right)}{\partial c_v}\right)\right)-\gamma_0\eta^2 c_v c_i$$

$$\frac{\partial c_i}{\partial t} = \nabla\cdot\left(\left(M_i\right)\nabla\left(\left(1-h(\eta)\right)\frac{\partial\left(F^{m}(c_v,c_i)\right)}{\partial c_i}+h(\eta)\frac{\partial\left(F^{v}(c_v,c_i)\right)}{\partial c_i}\right)\right)-\gamma_0\eta^2 c_v c_i$$

$$F^{m}(c_v,c_i) = F_A + E_v^f c_v + E_i^f c_i + k_B T[c_v\ln(c_v)+c_i\ln(c_i)+(1-c_v-c_i)\ln(1-c_v-c_i)]$$

$$F^{v}(c_v,c_i) = k_B T((c_v-1)^2+c_i^2)$$

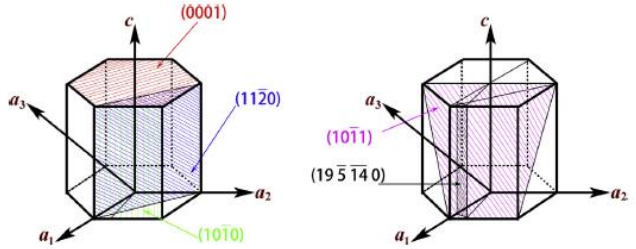

Fig. 2. An illustration of five representative planes in an hcp lattice.

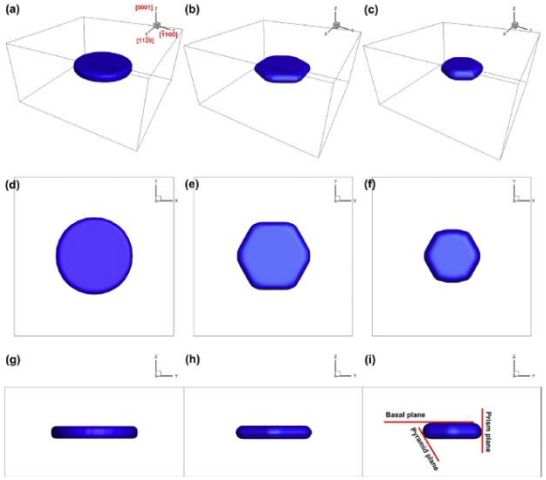

Localization

ETD

$$U = \begin{pmatrix} \bullet \\ \bullet \\ \vdots \\ \bullet \end{pmatrix}_{N_x \times N_y \times N_z} \qquad e^{L\Delta t} = \begin{pmatrix} \bullet & \bullet & \cdots & \bullet \\ \bullet & \bullet & \cdots & \bullet \\ \vdots & \vdots & \ddots & \vdots \\ \bullet & \bullet & \cdots & \bullet \end{pmatrix}_{(N_x \times N_y \times N_z) \times (N_x \times N_y \times N_z)}$$

M. Hochbruck and A. Ostermann, "Exponential integrators," Acta Numerica, vol. 19, pp. 209–286, 2010.

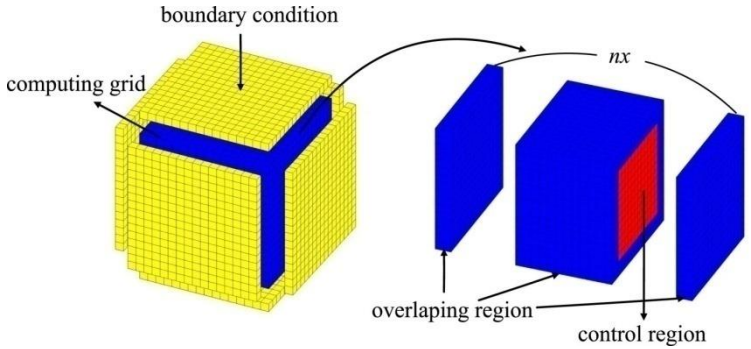

compact ETD

$$U = \quad e^{A\Delta t} = \begin{pmatrix} \bullet & \bullet & \cdots & \bullet \\ \bullet & \bullet & \cdots & \bullet \\ \vdots & \vdots & \ddots & \vdots \\ \bullet & \bullet & \cdots & \bullet \end{pmatrix}_{N_x \times N_x}$$

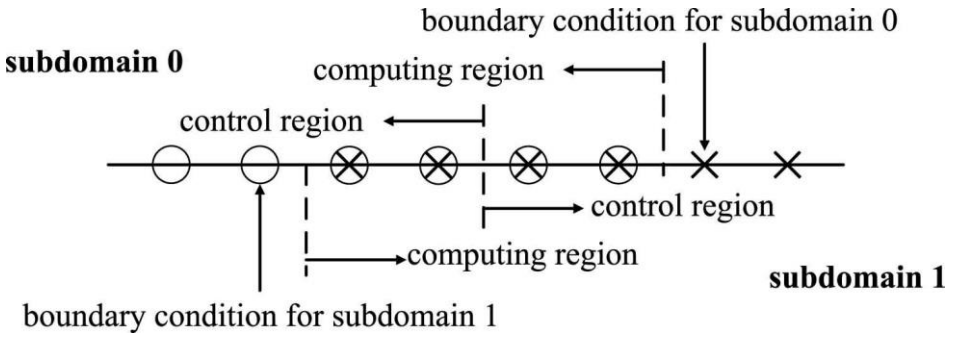

Efficient direct subdomain integration based on FD spatial discretization

+

subdomain coupling techniques
overlapping
BC & discretization

large time step-size, stable and accurate, compute intensive

GPU Acceleration

中国科学院
计算机网络信息中心
Computer Network Information Center,
Chinese Academy of Sciences

MPI Communication

26 adjacent subdomains

Twice per step

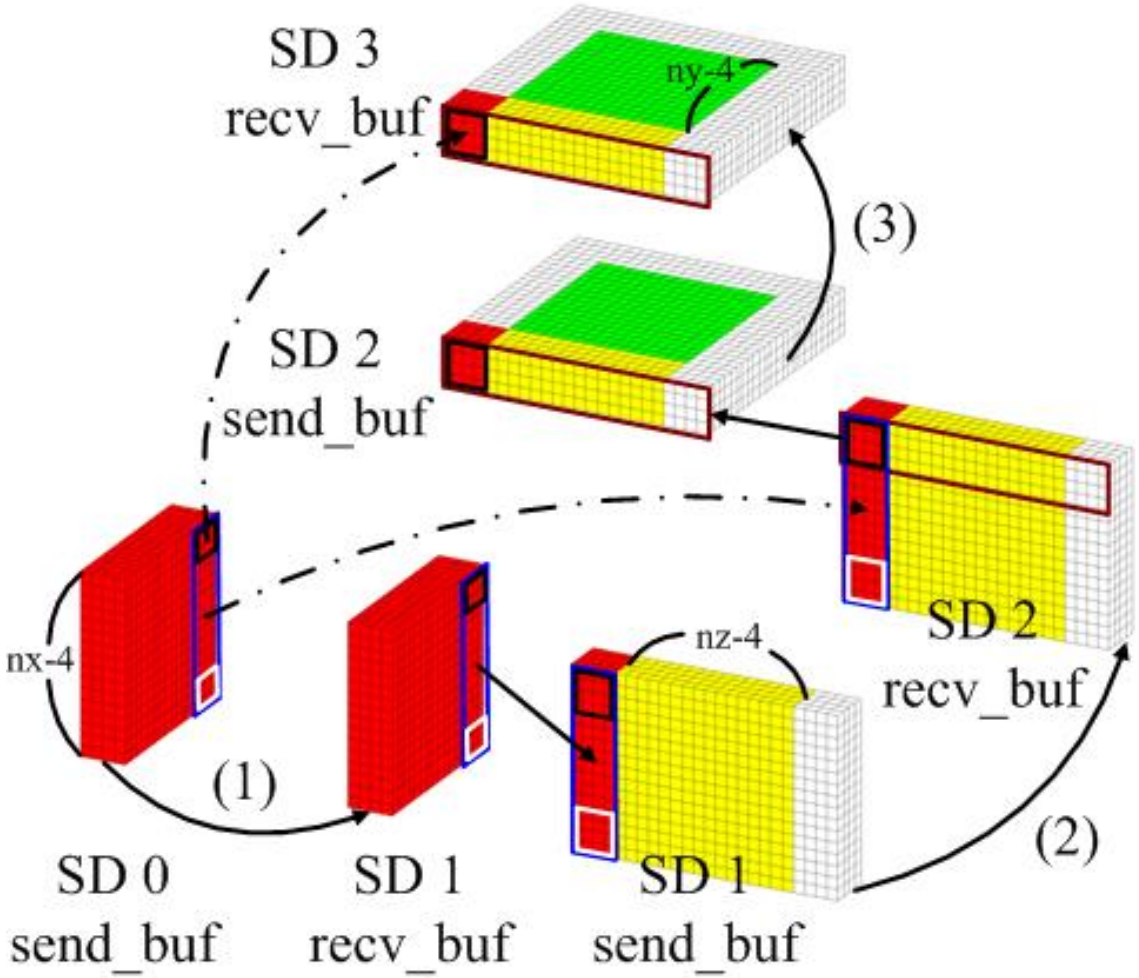

3-round scheme

Simulation setup

P100-PCIe-12GB： 4.7T=4812.8GFlops；540GB

subdomain： 768*768*384=0.2109G Points,

216 subdomain = 45G points; 20,000~50,000 time steps, average step size ~ 10,000X vs. explicit

Subdomain divided into 192*192*192 blocks when calculating matrix exponentials

~ perform 32 tensor dot production simultaneously

2.45TFlops/step

Performance

Between subdomain: 73ms (pack,copy,MPI)

Tensor dot production: 2.42T@3.19T/sec, 759ms, ~ 66% peak

Stencil & pointwise: 47ms

Overall performance: DP 2,787GFlops ~ 58% peak, ~880ms/step

Explicit FD scheme

Stencil： 12.8GFlops/step @ 40% peak ~6.2ms/step

10,000 steps= 62 sec

ETD is 70X faster!

Other Platforms

Sunway TaihuLight

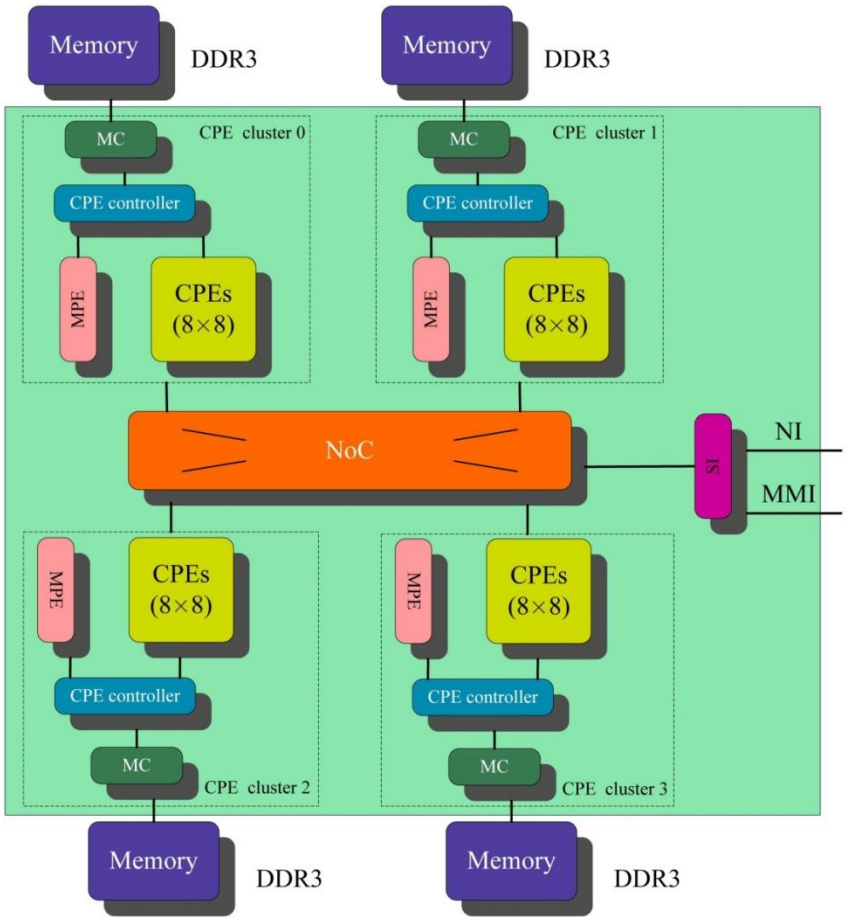

40,960 SW26010 many-core processors;

260 cores, divided into 4 core groups (CGs),

1 MPE + 64 CPEs

8GB main memory for each CG

64KB SPM for each CPE

MPI recommended among CGs
DMA available SPM ⟷ main memory

Performance Analysis

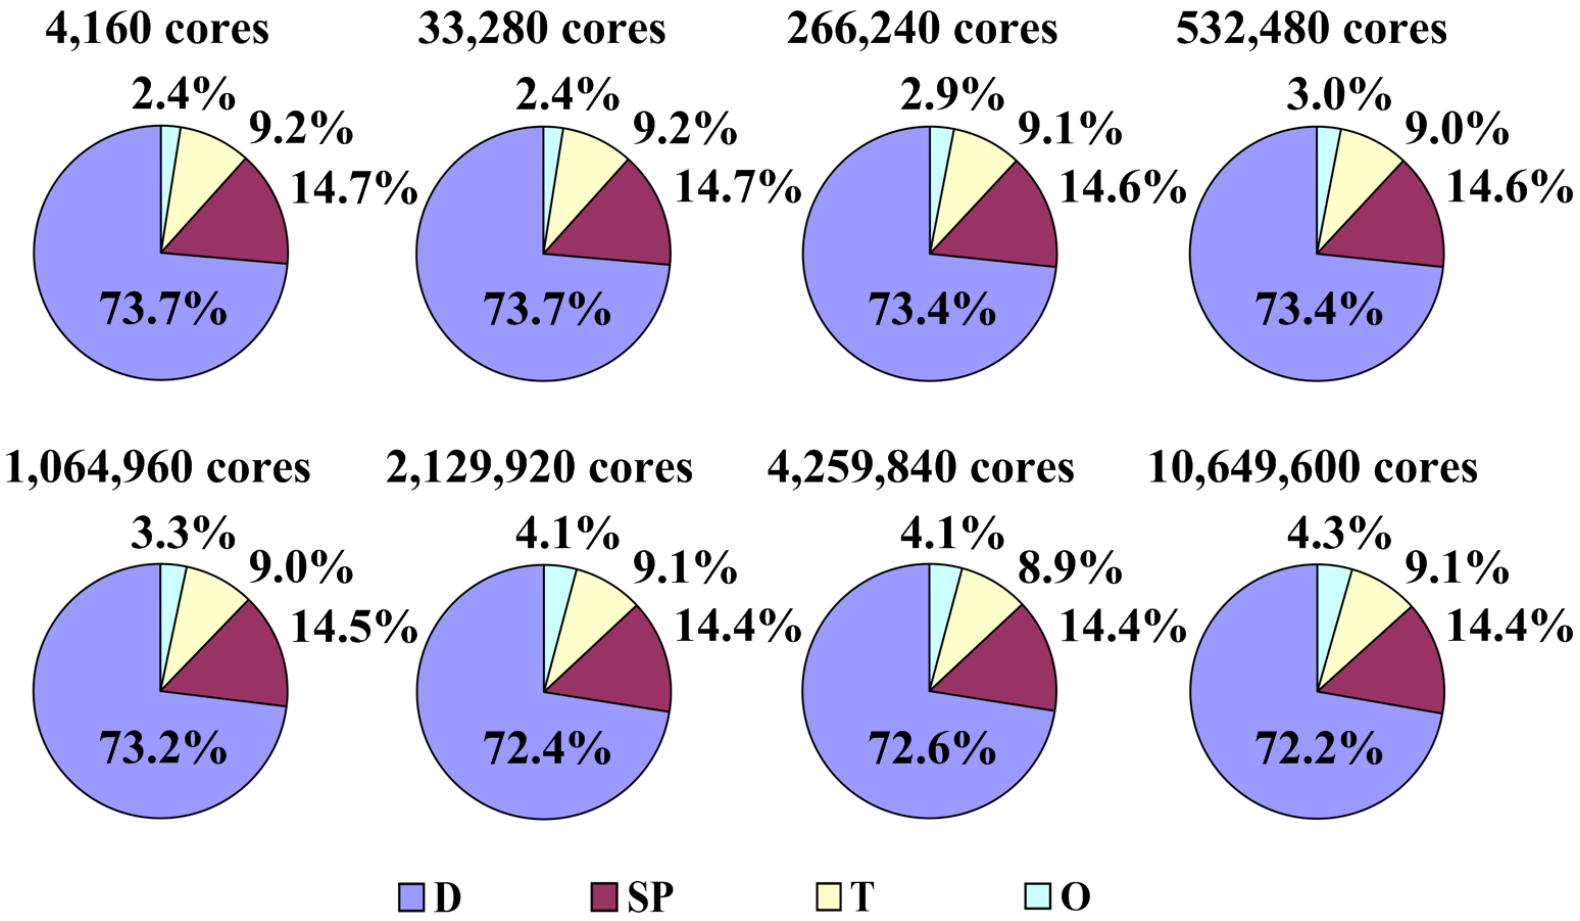

DGEMM: 457.2 and 408.5 GFlops, 60% and 53% peak
Aggregate DMA BW in T and SP: ~ 22GB/s
Overall : 316.1 to 324.5 Gflops, 41%-42% peak

Summary

中国科学院
计算机网络信息中心
Computer Network Information Center,
Chinese Academy of Sciences

Summary

A promising algorithm for a variety of architectures
Large time step, scalable, compute intensive

Idea applicable to other stiff evolution equations
fluid dynamics, structure-fluid interaction…

Thank you!