

Parallel Fast Fourier Transforms and Their Application to Particle Simulation

Michael Pippig

Department of Mathematics
Chemnitz University of Technology

September 21, 2011

supported by BMBF grant 01IH08001B

Table of Contents

- 1 Motivation
- 2 Parallel Fast Fourier Transform
- 3 Parallel Nonequispaced Fast Fourier Transforms
- 4 Parallel Fast Ewald Summation




- 1 **Motivation**
- 2 Parallel Fast Fourier Transform
- 3 Parallel Nonequispaced Fast Fourier Transforms
- 4 Parallel Fast Ewald Summation

Motivation

Coulomb Interaction in Particle Systems - $\mathcal{O}(M^2)$

$$\phi(\mathbf{x}_j) = \sum_{l=1}^{M'} \frac{q_l}{\|\mathbf{x}_j - \mathbf{x}_l\|_2}, \quad j = 1, \dots, M$$

$$\tilde{\phi}(\mathbf{x}_j) = \sum_{\mathbf{r} \in \mathbb{Z}^3} \sum_{l=1}^{M'} \frac{q_l}{\|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2}, \quad j = 1, \dots, M$$



Motivation

Coulomb Interaction in Particle Systems - $\mathcal{O}(M^2)$

$$\phi(\mathbf{x}_j) = \sum_{l=1}^{M'} \frac{q_l}{\|\mathbf{x}_j - \mathbf{x}_l\|_2}, \quad j = 1, \dots, M$$

$$\tilde{\phi}(\mathbf{x}_j) = \sum_{\mathbf{r} \in \mathbb{Z}^3} \sum_{l=1}^{M'} \frac{q_l}{\|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2}, \quad j = 1, \dots, M$$

Fast Algorithms

Ewald Sum	P ³ MG	Tree Codes	FMM	P ³ M	Fast Sum
1921	1977	1986	1987	1987	2004
$\mathcal{O}(M^{3/2})$	$\mathcal{O}(M)$	$\mathcal{O}(M \log M)$	$\mathcal{O}(M)$	$\mathcal{O}(M \log M)$	$\mathcal{O}(M \log M)$
Ewald	Brandt Hackbusch Trottenberg	Barnes Hut	Greengard Rokhlin	Hockney Eastwood	Nieslony Potts Steidl

Outline

- 1 Motivation
- 2 Parallel Fast Fourier Transform**
- 3 Parallel Nonequispaced Fast Fourier Transforms
- 4 Parallel Fast Ewald Summation

Fast Fourier Transform

Task of 3d-DFT

Consider a three-dimensional dataset of $n_0 \times n_1 \times n_2$ complex numbers $\hat{g}_{k_0 k_1 k_2} \in \mathbb{C}$. Compute


$$\begin{aligned} g_{l_0 l_1 l_2} &:= \sum_{k_0=0}^{n_0-1} \sum_{k_1=0}^{n_1-1} \sum_{k_2=0}^{n_2-1} \hat{g}_{k_0 k_1 k_2} e^{-2\pi i(k_2 \frac{l_2}{n_2} + k_1 \frac{l_1}{n_1} + k_0 \frac{l_0}{n_0})} \\ &= \sum_{k_0=0}^{n_0-1} \left(\sum_{k_1=0}^{n_1-1} \sum_{k_2=0}^{n_2-1} \hat{g}_{k_0 k_1 k_2} e^{-2\pi i(k_2 \frac{l_2}{n_2} + k_1 \frac{l_1}{n_1})} \right) e^{-2\pi i k_0 \frac{l_0}{n_0}} \end{aligned}$$

for all $l_t = 0, \dots, n_t - 1$ ($t = 0, 1, 2$).

Realized by 3d-FFT ($n_0 = n_1 = n_2 = n$)

$\Rightarrow \mathcal{O}(n^3 \log n)$ instead of $\mathcal{O}(n^6)$

One-Dimensional Data Distribution



p - number of processors

Features of FFTW [Frigo, Johnson]

- open source
- easy interface
- arbitrary size
- d -dim. FFT
- in place FFT
- high performance
- many transforms
- communicator
- adjust planning
- collect wisdom


Maximum Number of Processes p_{\max}^{1D} ($n_0 = n_1 = n_2 = n$)

$$p_{\max}^{1D} = n$$

FFTW combines portable performance and good usability, but is not scalable enough.

Two-Dimensional Data Distribution

[Ding, Eleftheriou et al. 03, Plimpton, Pekurovsky - P3DFFT]



$p_0 \times p_1$ - size of process grid

Maximum Number of Processes p_{\max}^{2D}
($n_0 = n_1 = n_2 = n$)

$$p_{\max}^{2D} = n^2$$

n	$p_{\max}^{1D} = n$	$p_{\max}^{2D} = n^2$
64	64	4096
128	128	16384
256	256	65536
512	512	262144
1024	1024	1048576

Comparison of PFFT and P3DFFT

P3DFFT Unique Features [Pekurovsky]

- support of different serial FFT vendors


Common Features

- open source
- high scalability
- portability
- multiple precisions
- Fortran interface
- C interface
- r2c FFT
- ghost cell support

PFFT Unique Features


- c2c FFT
- completely in place FFT
- FFTW like interface (basic, advanced and guru)
- adjustable blocksize
- separate communicator
- accumulated wisdom
- change of planning effort without recompilation
- d -dimensional parallel FFT
- over- and downsampling

Ghost Cell Support




Oversampled & Downsampled FFT


Without Library Support




PFFT Library Support



Scaling PFFT of Size 512^3 on BlueGene/P



Scaling PFFT of Size 512^3 on BlueGene/P



Outline

- 1 Motivation
- 2 Parallel Fast Fourier Transform
- 3 Parallel Nonequispaced Fast Fourier Transforms**
- 4 Parallel Fast Ewald Summation

Nonequispaced Discrete Fourier Transform

Task of 3d-DFT and 3d-NDFT

For $\hat{f}_{k_0 k_1 k_2} \in \mathbb{C}$ compute

$$f_{l_0 l_1 l_2} = \sum_{k_0=0}^{N-1} \sum_{k_1=0}^{N-1} \sum_{k_2=0}^{N-1} \hat{f}_{k_0 k_1 k_2} e^{-2\pi i(k_2 \frac{l_2}{N} + k_1 \frac{l_1}{N} + k_0 \frac{l_0}{N})} \quad (\text{DFT})$$

for all $l_0, l_1, l_2 \in \{0, \dots, N-1\}$ ($\Rightarrow \frac{l_0}{N}, \frac{l_1}{N}, \frac{l_2}{N} \in [0, 1)$), and

$$f_j = \sum_{k_0=0}^{N-1} \sum_{k_1=0}^{N-1} \sum_{k_2=0}^{N-1} \hat{f}_{k_0 k_1 k_2} e^{-2\pi i(k_2 z_j + k_1 y_j + k_0 x_j)} \quad (\text{NDFT})$$

for $x_j, y_j, z_j \in [0, 1)$, $j = 1, \dots, M$.

Nonequispaced Discrete Fourier Transforms

Task of 3d-NDFT and 3d-NDFT^H

For $\hat{f}_{k_0 k_1 k_2} \in \mathbb{C}$ compute

$$f_j = \sum_{k_0=0}^{N-1} \sum_{k_1=0}^{N-1} \sum_{k_2=0}^{N-1} \hat{f}_{k_0 k_1 k_2} e^{-2\pi i(k_2 z_j + k_1 y_j + k_0 x_j)} \quad (\text{NDFT})$$

for $x_j, y_j, z_j \in [0, 1)$, $j = 1, \dots, M$, and for $h_j \in \mathbb{C}$ compute

$$\hat{h}_{k_0 k_1 k_2} = \sum_{j=1}^M h_j e^{+2\pi i(k_2 z_j + k_1 y_j + k_0 x_j)} \quad (\text{NDFT}^H)$$

for $z_j, y_j, x_j \in [0, 1)$, $k_0, k_1, k_2 \in \{0, \dots, N-1\}$.

Nonequispaced Fast Fourier Transforms

Matrix-Vector-Notation of NDFT and NDFT^H

For $\hat{\mathbf{f}} \in \mathbb{C}^{N^3}$ and $\mathbf{h} \in \mathbb{C}^M$ compute

$$\mathbf{f} = \mathbf{A}\hat{\mathbf{f}} \in \mathbb{C}^M, \quad (\text{NDFT})$$

$$\hat{\mathbf{h}} = \mathbf{A}^H \mathbf{h} \in \mathbb{C}^{N^3}, \quad (\text{NDFT}^H)$$

where $\mathbf{A} = \left(e^{-2\pi i(k_2 z_j + k_1 y_j + k_0 x_j)} \right)_{j, k_0, k_1, k_2} \in \mathbb{C}^{M \times N^3}$.


NFFT [Dutt, Rohklin 93, Beylkin 95, Steidl 96, ...]

$$\mathbf{A} \approx \mathbf{BFD}, \quad \mathbf{A}^H \approx \mathbf{DF}^H \mathbf{B}^T$$


- $\mathbf{D} \in \mathbb{R}^{N^3 \times N^3}$ diagonal matrix
- $\mathbf{F} \in \mathbb{C}^{n^3 \times N^3}$ truncated Fourier matrix ($n \geq N$)
- $\mathbf{B} \in \mathbb{R}^{M \times n^3}$ sparse matrix

$\Rightarrow \mathcal{O}(N^3 \log N + \log^3(\frac{1}{\epsilon})M)$ instead of $\mathcal{O}(N^3 M)$

Scaling PNFFT of Size 128^3 on BlueGene/P



Scaling PNFFT of Size 128^3 on BlueGene/P



Outline

- 1 Motivation
- 2 Parallel Fast Fourier Transform
- 3 Parallel Nonequispaced Fast Fourier Transforms
- 4 Parallel Fast Ewald Summation**

Coulomb Interaction in Periodic Particle Systems

Calculation of the Potential

$$\tilde{\phi}(\mathbf{x}_j) = \sum_{\mathbf{r} \in \mathbb{Z}^3} \sum_{l=1}^M \frac{q_l}{\|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2}, \quad j = 1, \dots, M$$

- conditionally convergent
- computational work is $\mathcal{O}(M^2)$

Ewald Splitting

$$\frac{1}{r} = \frac{1 - f(r)}{r} + \frac{f(r)}{r} \quad \Rightarrow \quad \tilde{\phi} = \tilde{\phi}^{\text{real}} + \tilde{\phi}^{\text{reci}}$$




- $\frac{1-f(r)}{r}$ converges fast in real space
- $\frac{f(r)}{r}$ converges fast in reciprocal space and $\lim_{r \rightarrow 0} \frac{f(r)}{r}$ exists

Ewald Splitting with Error Function

Error Function

$$f(r) = \text{erf}(\alpha r) = \frac{2}{\sqrt{\pi}} \int_0^r e^{-t^2} dt, \quad \lim_{r \rightarrow 0} \frac{\text{erf}(\alpha r)}{r} = \frac{2\alpha}{\sqrt{\pi}}$$

$$\frac{1}{r} = \frac{\text{erfc}(\alpha r)}{r} + \frac{\text{erf}(\alpha r)}{r}$$



Ewald Summation

Nearfield Approximation - $\mathcal{O}(M^2 R^3) \xrightarrow{\alpha_{\text{opt}}} \mathcal{O}(M^{3/2})$

$$\begin{aligned}\tilde{\phi}^{\text{real}}(\mathbf{x}_j) &= -2q_j \frac{\alpha}{\sqrt{\pi}} + \sum_{\mathbf{r} \in \mathbb{Z}^3} \sum_{l=1}^M q_l \frac{\text{erfc}(\alpha \|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2)}{\|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2} \\ &\approx -2q_j \frac{\alpha}{\sqrt{\pi}} + \sum_{\mathbf{r} \in I_{\mathbf{R}}} \sum_{l=1}^M q_l \frac{\text{erfc}(\alpha \|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2)}{\|\mathbf{x}_j - \mathbf{x}_l + \mathbf{r}\|_2}\end{aligned}$$

Farfield Approximation - $\mathcal{O}(MN^3) \xrightarrow{\alpha_{\text{opt}}} \mathcal{O}(M^{3/2})$

$$\begin{aligned}\tilde{\phi}^{\text{reci}}(\mathbf{x}_j) &= \frac{1}{\pi} \sum_{\mathbf{k} \in \mathbb{Z}^3 \setminus \{\mathbf{0}\}} \frac{e^{-\pi^2 \|\mathbf{k}\|_2^2 / \alpha^2}}{\|\mathbf{k}\|_2^2} \sum_{l=1}^M q_l e^{-2\pi i \mathbf{k}(\mathbf{x}_j - \mathbf{x}_l)} \\ &\approx \frac{1}{\pi} \sum_{\mathbf{k} \in I_{\mathbf{N}} \setminus \{\mathbf{0}\}} \frac{e^{-\pi^2 \|\mathbf{k}\|_2^2 / \alpha^2}}{\|\mathbf{k}\|_2^2} \left(\sum_{l=1}^M q_l e^{+2\pi i \mathbf{k} \mathbf{x}_l} \right) e^{-2\pi i \mathbf{k} \mathbf{x}_j}\end{aligned}$$

Fast Ewald Summation [Hedman, Laaksonen 2006]

Farfield Approximation - $\mathcal{O}(MN^3)$

$$\tilde{\phi}^{\text{reci}}(\mathbf{x}_j) \approx \frac{1}{\pi} \sum_{\mathbf{k} \in I_N \setminus \{\mathbf{0}\}} \frac{e^{-\pi^2 \|\mathbf{k}\|_2^2 / \alpha^2}}{\|\mathbf{k}\|_2^2} \left(\sum_{l=1}^M q_l e^{+2\pi i \mathbf{k} \mathbf{x}_l} \right) e^{-2\pi i \mathbf{k} \mathbf{x}_j}$$

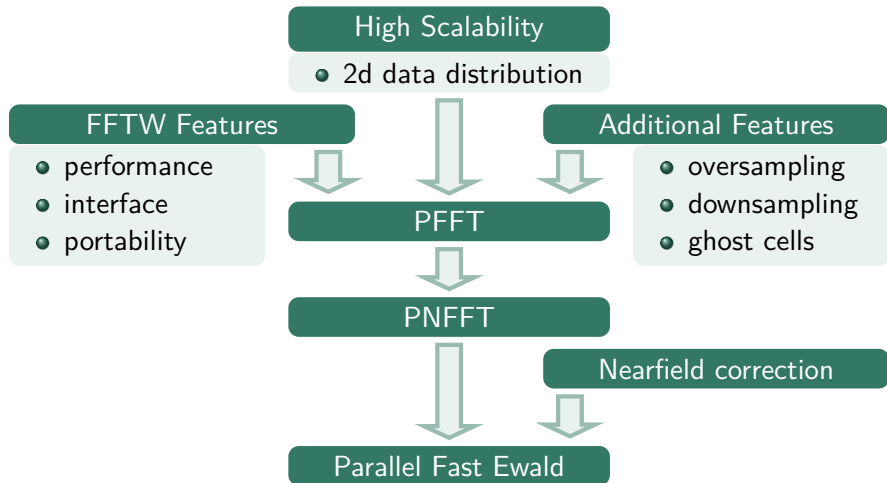
Calculate by NFFT^H - $\mathcal{O}(\log^3(\frac{1}{\epsilon})M + N^3 \log N)$

$$S(\mathbf{k}) = \sum_{l=1}^M q_l e^{+2\pi i \mathbf{k} \mathbf{x}_l}$$

Calculate by NFFT - $\mathcal{O}(\log^3(\frac{1}{\epsilon})M + N^3 \log N)$

$$\sum_{\mathbf{k} \in I_N \setminus \{\mathbf{0}\}} \left(\frac{1}{\pi} \frac{e^{-\pi^2 \|\mathbf{k}\|_2^2 / \alpha^2}}{\|\mathbf{k}\|_2^2} S(\mathbf{k}) \right) e^{-2\pi i \mathbf{k} \mathbf{x}_j}$$

Summary



PNFFT & PFFT Software Library and Papers

Available at

<http://www.tu-chemnitz.de/~mpip>