Massively parallel computation of nonequispaced fast Fourier transforms

Michael Pippig, Daniel Potts

Department of Mathematics
Chemnitz University of Technology

June 30, 2012

supported by BMBF grant 01IH08001B

Table of Contents

Outline

1. Parallel Nonequispaced Fast Fourier Transform

2. Application: Fast Summation

Discrete Fourier Transforms

Task of 3d-DFT (Discrete Fourier Transform)

For $\hat{f}_{\boldsymbol{k}} \in \mathbb{C}$ compute

$$f_{\boldsymbol{l}} = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}} \hat{f}_{\boldsymbol{k}} \, \mathrm{e}^{-2\pi \mathrm{i}(k_0 \frac{l_0}{N} + k_1 \frac{l_1}{N} + k_2 \frac{l_2}{N})}$$

for all $\boldsymbol{l} \in I_{\boldsymbol{N}} := \{0, \dots, N-1\}^3$ $\left(\Rightarrow \frac{l_0}{N}, \frac{l_1}{N}, \frac{l_2}{N} \in [0,1)\right)$.

Task of 3d-NDFT (Nonequispaced DFT)

For $\hat{f}_{\boldsymbol{k}} \in \mathbb{C}$ compute

$$f_j = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}} \hat{f}_{\boldsymbol{k}} \, \mathrm{e}^{-2\pi \mathrm{i}(k_0 x_j + k_1 y_j + k_2 z_j)}$$

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

Window Function

Gaussian

$$\varphi(x) = \frac{1}{\sqrt{(\pi b)}} \mathrm{e}^{-\frac{(nx)^2}{b}}, \qquad \hat{\varphi}_k = \frac{1}{n} \mathrm{e}^{-b\left(\frac{\pi k}{n}\right)^2}$$

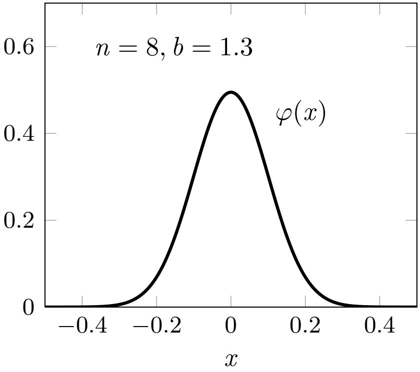

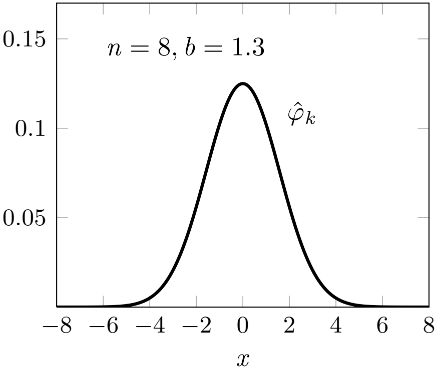

Window Function

Gaussian

$$\varphi(x) = \frac{1}{\sqrt{(\pi b)}} \mathrm{e}^{-\frac{(nx)^2}{b}}, \qquad \hat{\varphi}_k = \frac{1}{n} \mathrm{e}^{-b\left(\frac{\pi k}{n}\right)^2}$$

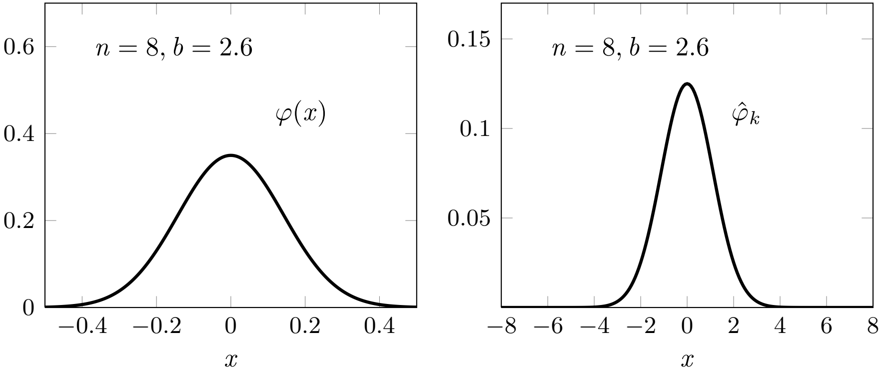

Window Function

Gaussian

$$\varphi(x) = \frac{1}{\sqrt{(\pi b)}} \mathrm{e}^{-\frac{(nx)^2}{b}}, \qquad \hat{\varphi}_k = \frac{1}{n} \mathrm{e}^{-b\left(\frac{\pi k}{n}\right)^2}$$

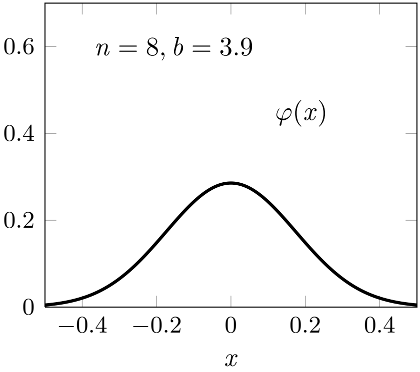

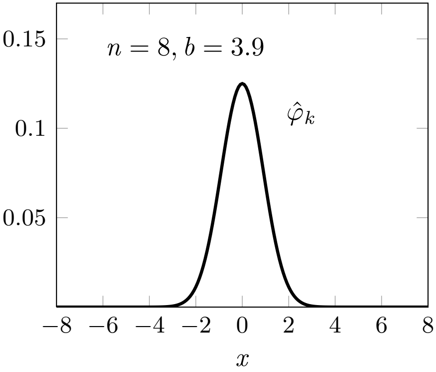

Nonequispaced Fast Fourier Transform

1. Deconvolution Step $\mathcal{O}(N)$

$$\hat{g}_{\boldsymbol{k}} = \frac{1}{|\mathcal{I}_{\boldsymbol{n}}|} \cdot \frac{\hat{f}_{\boldsymbol{k}}}{\hat{\varphi}_{k_0}\hat{\varphi}_{k_1}\hat{\varphi}_{k_2}}, \quad \boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}$$

2. Oversampled FFT $\mathcal{O}(N^3 \log N)$

$$g_{\boldsymbol{l}} = \sum_{\boldsymbol{k}\in\mathcal{I}_{\boldsymbol{N}}} \hat{g}_{\boldsymbol{k}}\, \mathrm{e}^{-2\pi\mathrm{i}(k_0\frac{l_0}{n}+k_1\frac{l_1}{n}+k_2\frac{l_2}{n})}, \quad \boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}$$

3. Convolution Step $\mathcal{O}(|\log\varepsilon|^3 M)$

$$f_j \approx \sum_{\boldsymbol{l}\in\mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right)\varphi\left(y_j - \tfrac{l_1}{n}\right)\varphi\left(z_j - \tfrac{l_2}{n}\right) g_{\boldsymbol{l}}, \quad j = 1, \dots, M$$

Nonequispaced Fast Fourier Transforms

Matrix-Vector-Notation of NDFT and adjont NDFT

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

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

$$\hat{\boldsymbol{h}} = \boldsymbol{A}^{\mathsf{H}}\boldsymbol{h} \in \mathbb{C}^{N^3}, \qquad \text{(adjoint NDFT)}$$

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

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

$$\boldsymbol{A} \approx \boldsymbol{C}\boldsymbol{F}\boldsymbol{D}, \qquad \boldsymbol{A}^{\mathsf{H}} \approx \boldsymbol{D}\boldsymbol{F}^{\mathsf{H}}\boldsymbol{C}^{\top}$$

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

$\Rightarrow \mathcal{O}(N^3 \log N + |\log \varepsilon|^3 M)$ instead of $\mathcal{O}(N^3 M)$

Parallel Deconvolution Step

1. Deconvolution Step $\mathcal{O}(N)$

$$\hat{g}_{\boldsymbol{k}} = \frac{1}{|\mathcal{I}_{\boldsymbol{n}}|} \cdot \frac{\hat{f}_{\boldsymbol{k}}}{\hat{\varphi}_{k_0}\hat{\varphi}_{k_1}\hat{\varphi}_{k_2}}\,, \quad \boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}$$

2. Oversampled FFT $\mathcal{O}(N^3 \log N)$

$$g_{\boldsymbol{l}} = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}} \hat{g}_{\boldsymbol{k}}\, \mathrm{e}^{-2\pi\mathrm{i}(k_0 \frac{l_0}{n} + k_1 \frac{l_1}{n} + k_2 \frac{l_2}{n})}\,, \quad \boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}$$

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$f_j \approx \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right) \varphi\left(y_j - \tfrac{l_1}{n}\right) \varphi\left(z_j - \tfrac{l_2}{n}\right) g_{\boldsymbol{l}}\,, \quad j = 1, \ldots, M$$

Parallel Fast Fourier Transform

1. Deconvolution Step $\mathcal{O}(N)$

$$\hat{g}_{\boldsymbol{k}} = \frac{1}{|\mathcal{I}_{\boldsymbol{n}}|} \cdot \frac{\hat{f}_{\boldsymbol{k}}}{\hat{\varphi}_{k_0}\hat{\varphi}_{k_1}\hat{\varphi}_{k_2}}, \quad \boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}$$

2. Oversampled FFT $\mathcal{O}(N^3 \log N)$

$$g_{\boldsymbol{l}} = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}} \hat{g}_{\boldsymbol{k}} \, \mathrm{e}^{-2\pi\mathrm{i}(k_0 \frac{l_0}{n} + k_1 \frac{l_1}{n} + k_2 \frac{l_2}{n})}, \quad \boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}$$

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$f_j \approx \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right) \varphi\left(y_j - \tfrac{l_1}{n}\right) \varphi\left(z_j - \tfrac{l_2}{n}\right) g_{\boldsymbol{l}}, \quad j = 1, \ldots, M$$

Highly Scalable Parallel FFT

FFTW

[Frigo, Johnson 2005]

Features of FFTW

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

Highly Scalable Parallel FFT

FFTW

[Frigo, Johnson 2005]

1d Data Decomposition

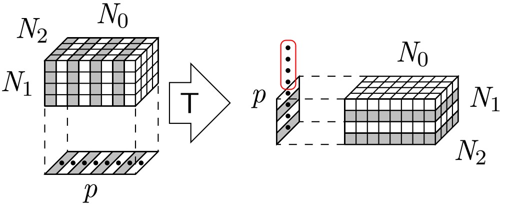

Maximum Number of Processes $p_{\max}^{\mathbf{1D}}$ ($N_0 = N_1 = N_2 = N$)

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

Highly Scalable Parallel FFT

FFTW

[Frigo, Johnson 2005]

1d Data Decomposition

FFTW_MPI

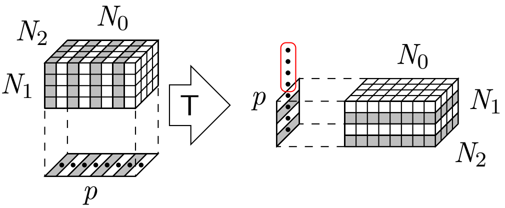

Maximum Number of Processes $p^{\mathbf{1D}}_{\max}$ $(N_0 = N_1 = N_2 = N)$

$$p^{1D}_{\max} = N$$

Highly Scalable Parallel FFT

FFTW

[Frigo, Johnson 2005]

1d Data Decomposition

FFTW_MPI

2d Data Decomposition

[Ding 1995]

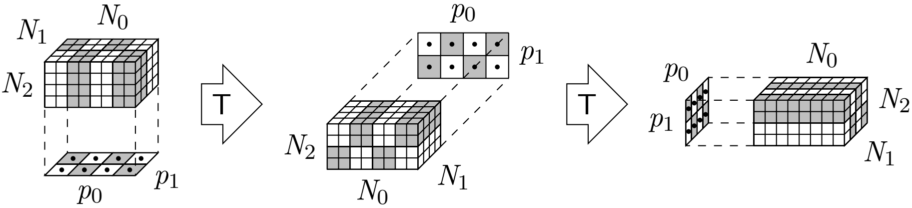

Highly Scalable Parallel FFT

FFTW

[Frigo, Johnson 2005]

1d Data Decomposition

FFTW_MPI

2d Data Decomposition

[Ding 1995]

Maximum Number of Processes $p^{\mathbf{2D}}_{\max}$ $(N_0 = N_1 = N_2 = N)$

$$p^{\mathrm{2D}}_{\max} = N^2$$

N	$p^{\mathrm{1D}}_{\max} = N$	$p^{\mathrm{2D}}_{\max} = N^2$
64	64	4096
128	128	16384
256	256	65536
512	512	262144
1024	1024	1048576

Highly Scalable Parallel FFT

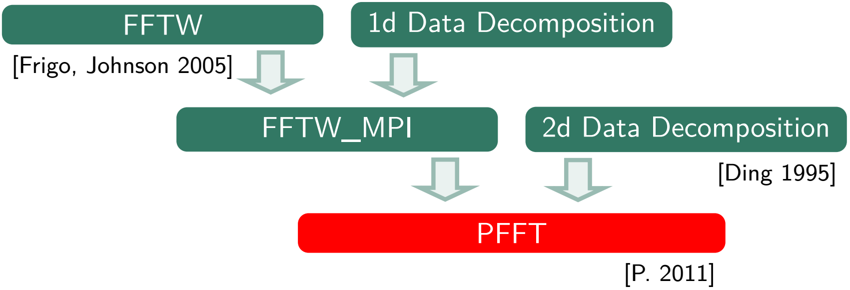

Features of PFFT

- open source
- high scalability
- portability
- c2c, r2c FFT
- FFTW like interface
- completely in place FFT
- d-dimensional parallel FFT
- ghost cell support

Scaling Parallel FFT of Size 512^3 on BlueGene/P

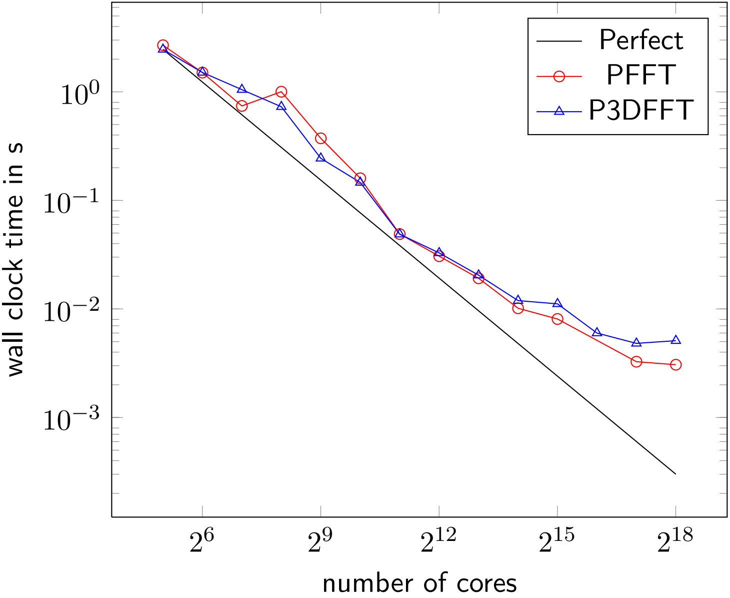

Parallel Oversampled FFT

Without Library Support

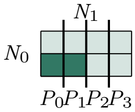

Parallel Oversampled FFT

Without Library Support

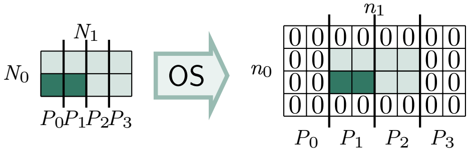

Parallel Oversampled FFT

Without Library Support

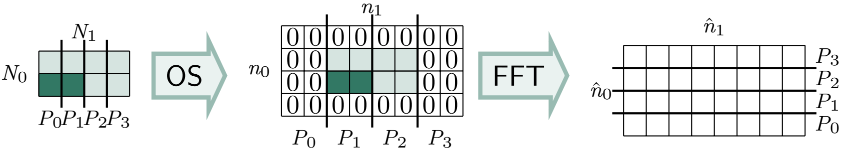

Parallel Oversampled FFT

Without Library Support

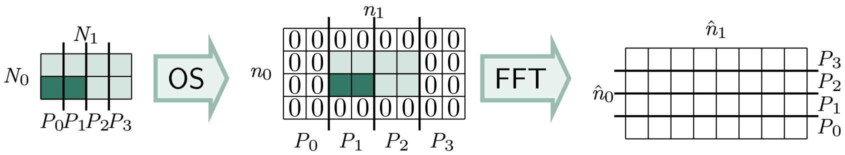

PFFT Library Support

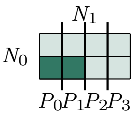

Parallel Oversampled FFT

Without Library Support

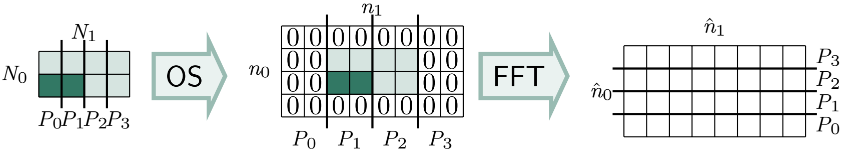

PFFT Library Support

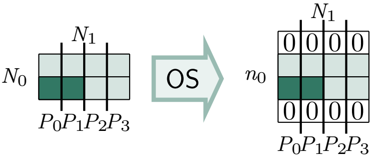

Parallel Oversampled FFT

Without Library Support

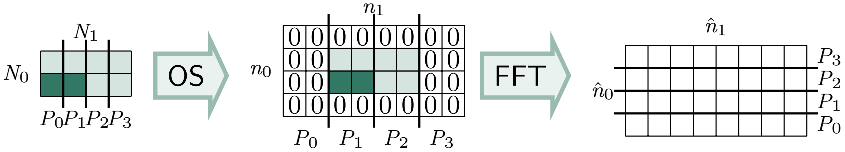

PFFT Library Support

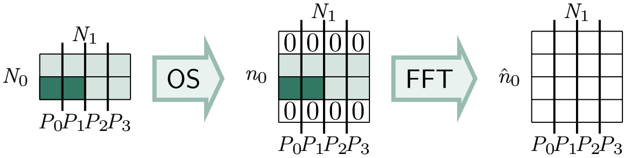

Parallel Oversampled FFT

Without Library Support

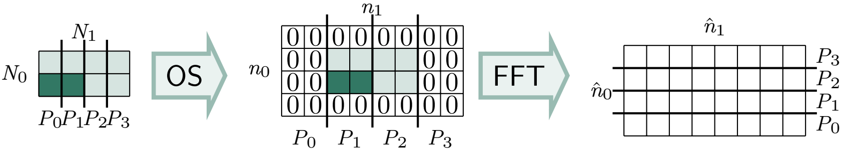

PFFT Library Support

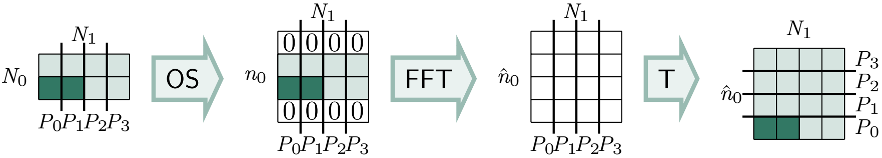

Parallel Oversampled FFT

Without Library Support

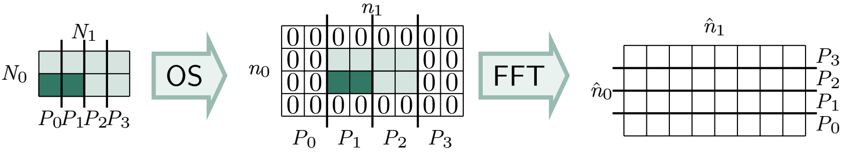

PFFT Library Support

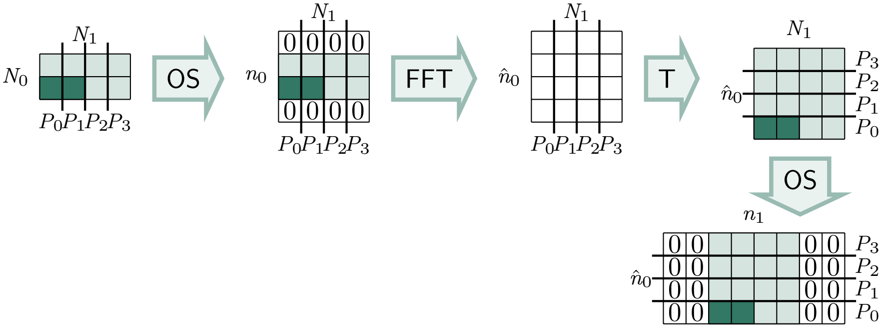

Parallel Oversampled FFT

Without Library Support

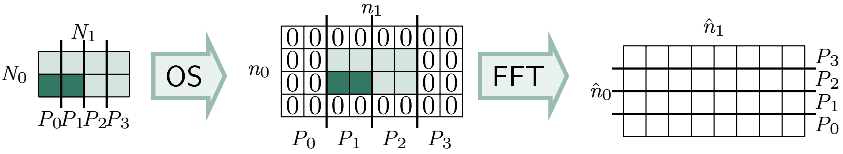

PFFT Library Support

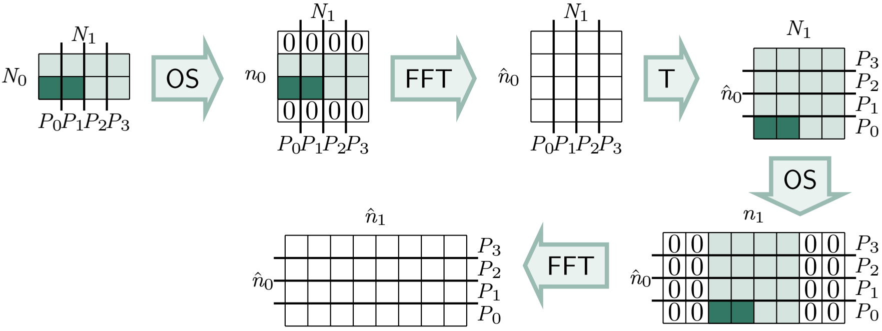

Parallel Convolution Step

1. Deconvolution Step $\mathcal{O}(N)$

$$\hat{g}_{\boldsymbol{k}} = \frac{1}{|\mathcal{I}_{\boldsymbol{n}}|} \cdot \frac{\hat{f}_{\boldsymbol{k}}}{\hat{\varphi}_{k_0} \hat{\varphi}_{k_1} \hat{\varphi}_{k_2}}, \quad \boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}$$

2. Oversampled FFT $\mathcal{O}(N^3 \log N)$

$$g_{\boldsymbol{l}} = \sum_{\boldsymbol{k} \in \mathcal{I}_{\boldsymbol{N}}} \hat{g}_{\boldsymbol{k}} \, \mathrm{e}^{-2\pi \mathrm{i}\left(k_0 \frac{l_0}{n} + k_1 \frac{l_1}{n} + k_2 \frac{l_2}{n}\right)}, \quad \boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}$$

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$f_j \approx \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right) \varphi\left(y_j - \tfrac{l_1}{n}\right) \varphi\left(z_j - \tfrac{l_2}{n}\right) g_{\boldsymbol{l}}, \quad j = 1, \ldots, M$$

Parallel Convolution Step

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$s_j = \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right) \varphi\left(y_j - \tfrac{l_1}{n}\right) \varphi\left(z_j - \tfrac{l_1}{n}\right) g_{\boldsymbol{l}}\,, \quad j = 1, \ldots, M$$

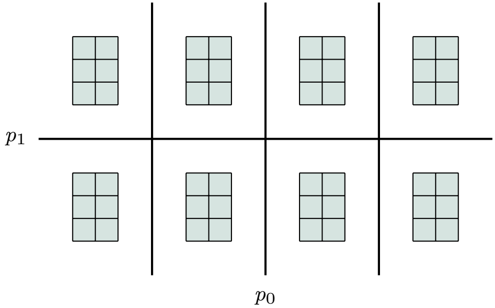

Parallel Convolution Step

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$s_j = \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \frac{l_0}{n}\right) \varphi\left(y_j - \frac{l_1}{n}\right) \varphi\left(z_j - \frac{l_1}{n}\right) g_{\boldsymbol{l}}\,, \quad j = 1, \ldots, M$$

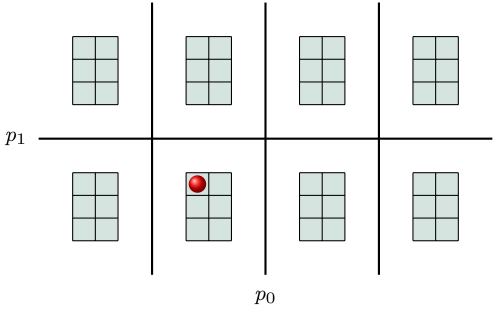

Parallel Convolution Step

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$s_j = \sum_{l \in \mathcal{I}_n} \varphi\left(x_j - \frac{l_0}{n}\right) \varphi\left(y_j - \frac{l_1}{n}\right) \varphi\left(z_j - \frac{l_1}{n}\right) g_l\,, \quad j = 1, \ldots, M$$

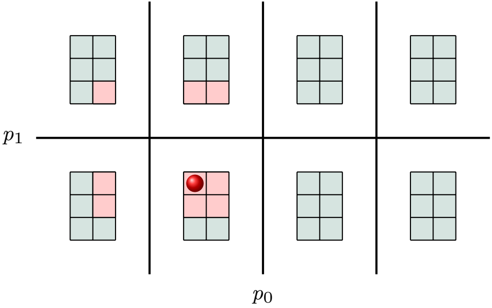

Parallel Convolution Step

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$s_j = \sum_{\boldsymbol{l} \in \mathcal{I}_n} \varphi\left(x_j - \frac{l_0}{n}\right) \varphi\left(y_j - \frac{l_1}{n}\right) \varphi\left(z_j - \frac{l_1}{n}\right) g_{\boldsymbol{l}}\,, \quad j = 1, \ldots, M$$

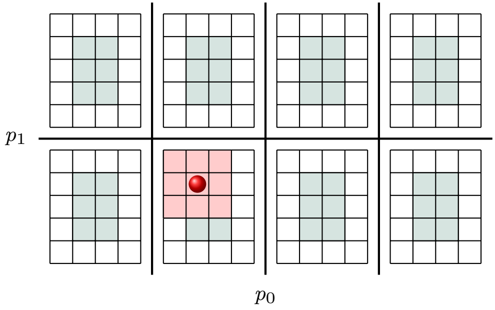

Parallel Convolution Step

3. Convolution Step $\mathcal{O}(|\log \varepsilon|^3 M)$

$$s_j = \sum_{\boldsymbol{l} \in \mathcal{I}_{\boldsymbol{n}}} \varphi\left(x_j - \tfrac{l_0}{n}\right) \varphi\left(y_j - \tfrac{l_1}{n}\right) \varphi\left(z_j - \tfrac{l_1}{n}\right) g_{\boldsymbol{l}}\,, \quad j = 1, \ldots, M$$

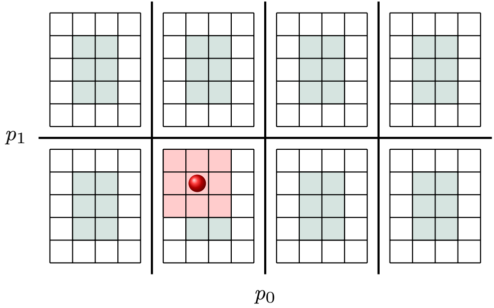

The PFFT software library offers very flexible ghost cell support.

Scaling PNFFT of Size 256^3 on BlueGene/P

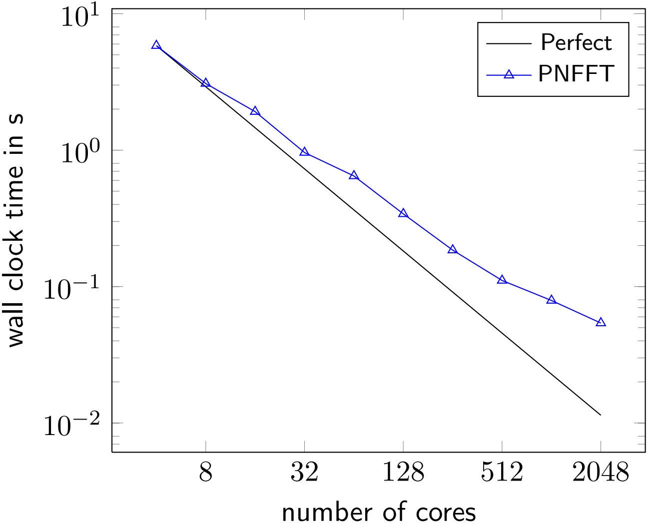

NFFT parameters: $N = 256, n = 288, m = 4, M = 103680$

Scaling PNFFT of Size 256^3 on BlueGene/P

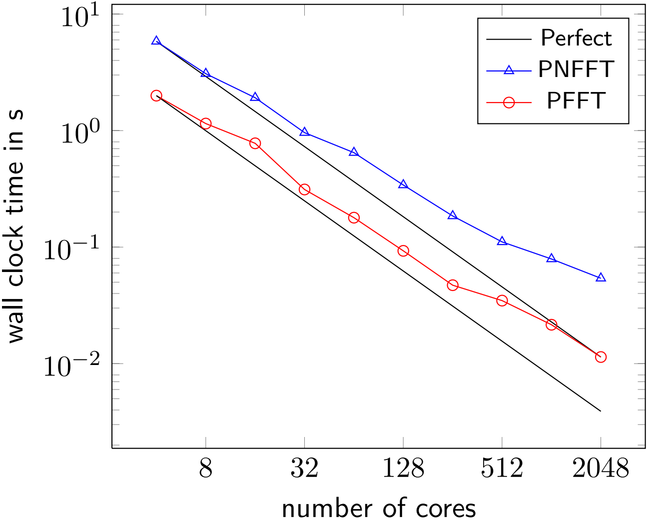

NFFT parameters: $N = 256, n = 288, m = 4, M = 103680$

Outline

1. Parallel Nonequispaced Fast Fourier Transform
2. Application: Fast Summation

Coulomb Interaction in Open Particle Systems

Calculation of the Potentials - $\mathcal{O}(M^2)$

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

Split Kernel into Nearfield and Farfield

$$\frac{1}{r} = \frac{1}{r} - R(r) + R(r)$$

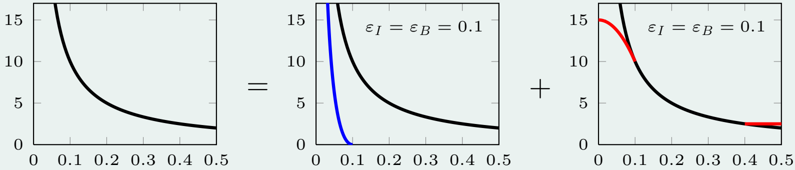

Approximate Farfield by Fourier Series

$$R(\|\boldsymbol{x}\|_2) \approx \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k}\boldsymbol{x}}$$

Fast Summation [Potts, Steidl 2004]

Nearfield Approximation - $\mathcal{O}(\nu M)$

$$\phi^{\text{near}}(\boldsymbol{x}_j) = -q_j R(0) + \sum_{l \in I_j}{}' q_l \left(\frac{1}{\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2} - R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \right)$$

$$I_j = \left\{ l = 1, \ldots, M : \|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2 < \varepsilon_I \right\}, \quad \nu := \max_j |I_j|$$

Farfield Approximation - $\mathcal{O}(MN^3) \to \mathcal{O}(\log^3(\frac{1}{\varepsilon})M + N^3 \log N)$

$$\phi^{\text{far}}(\boldsymbol{x}_j) = \sum_{l=1}^{M} q_l R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \approx \sum_{l=1}^{M} q_l \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k}(\boldsymbol{x}_j - \boldsymbol{x}_l)}$$

$$= \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \left(\sum_{l=1}^{M} q_l \mathrm{e}^{+2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_l} \right) \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_j}$$

Matrix decomposition: $\boldsymbol{C}^{\text{near}} + \boldsymbol{A} \, \mathrm{diag}(\hat{R}_{\boldsymbol{k}}) \, \boldsymbol{A}^{\mathsf{H}}$

Fast Summation [Potts, Steidl 2004]

Nearfield Approximation - $\mathcal{O}(\nu M)$

$$\phi^{\text{near}}(\boldsymbol{x}_j) = -q_j R(0) + \sum_{l\in I_j}{}' q_l \left(\frac{1}{\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2} - R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \right)$$

$$I_j = \left\{ l = 1, \ldots, M : \|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2 < \varepsilon_I \right\}, \quad \nu := \max_j |I_j|$$

Farfield Approximation - $\mathcal{O}(MN^3) \to \mathcal{O}(\log^3(\frac{1}{\varepsilon})M + N^3 \log N)$

$$\phi^{\text{far}}(\boldsymbol{x}_j) = \sum_{l=1}^{M} q_l R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \approx \sum_{l=1}^{M} q_l \sum_{\boldsymbol{k}\in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \mathrm{e}^{-2\pi\mathrm{i}\boldsymbol{k}(\boldsymbol{x}_j - \boldsymbol{x}_l)}$$

$$= \sum_{\boldsymbol{k}\in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \left(\sum_{l=1}^{M} q_l \mathrm{e}^{+2\pi\mathrm{i}\boldsymbol{k}\boldsymbol{x}_l} \right) \mathrm{e}^{-2\pi\mathrm{i}\boldsymbol{k}\boldsymbol{x}_j}$$

adjoint NFFT

Matrix decomposition: $\boldsymbol{C}^{\text{near}} + \boldsymbol{A}\, \mathrm{diag}(\hat{R}_{\boldsymbol{k}})\, \boldsymbol{A}^{\mathsf{H}}$

Fast Summation [Potts, Steidl 2004]

Nearfield Approximation - $\mathcal{O}(\nu M)$

$$\phi^{\text{near}}(\boldsymbol{x}_j) = -q_j R(0) + \sum_{l \in I_j}{}' q_l \left(\frac{1}{\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2} - R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \right)$$

$$I_j = \left\{ l = 1, \ldots, M : \|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2 < \varepsilon_I \right\}, \quad \nu := \max_j |I_j|$$

Farfield Approximation - $\mathcal{O}(MN^3) \to \mathcal{O}(\log^3(\frac{1}{\varepsilon})M + N^3 \log N)$

$$\phi^{\text{far}}(\boldsymbol{x}_j) = \sum_{l=1}^{M} q_l R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \approx \sum_{l=1}^{M} q_l \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k}(\boldsymbol{x}_j - \boldsymbol{x}_l)}$$

$$= \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \left(\sum_{l=1}^{M} q_l \mathrm{e}^{+2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_l} \right) \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_j}$$

convolution in Fourier space

Matrix decomposition: $\boldsymbol{C}^{\text{near}} + \boldsymbol{A}\, \mathrm{diag}(\hat{R}_{\boldsymbol{k}})\, \boldsymbol{A}^{\mathsf{H}}$

Fast Summation [Potts, Steidl 2004]

Nearfield Approximation - $\mathcal{O}(\nu M)$

$$\phi^{\text{near}}(\boldsymbol{x}_j) = -q_j R(0) + \sum_{l \in I_j}{}' q_l \left(\frac{1}{\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2} - R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \right)$$

$$I_j = \left\{ l = 1, \ldots, M : \|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2 < \varepsilon_I \right\}, \quad \nu := \max_j |I_j|$$

Farfield Approximation - $\mathcal{O}(MN^3) \to \mathcal{O}(\log^3(\frac{1}{\varepsilon}) M + N^3 \log N)$

$$\phi^{\text{far}}(\boldsymbol{x}_j) = \sum_{l=1}^{M} q_l R(\|\boldsymbol{x}_j - \boldsymbol{x}_l\|_2) \approx \sum_{l=1}^{M} q_l \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k}(\boldsymbol{x}_j - \boldsymbol{x}_l)}$$

$$= \sum_{\boldsymbol{k} \in I_{\boldsymbol{N}}} \hat{R}_{\boldsymbol{k}} \left(\sum_{l=1}^{M} q_l \mathrm{e}^{+2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_l} \right) \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{k} \boldsymbol{x}_j} \qquad \text{NFFT}$$

Matrix decomposition: $\boldsymbol{C}^{\text{near}} + \boldsymbol{A} \, \mathrm{diag}(\hat{R}_{\boldsymbol{k}}) \, \boldsymbol{A}^{\mathsf{H}}$

Scaling Parallel Fast Summation on BlueGene/P

silica melt with 103680 particles: RMS-force error 2.03×10^{-5}

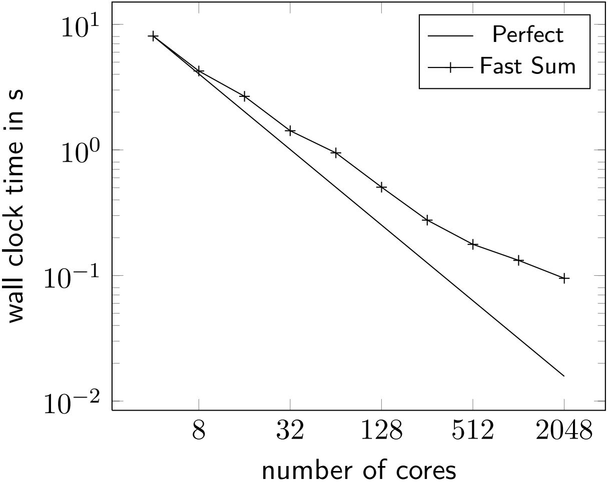

Parameters: $N = 256, n = 288, m = 4, \varepsilon_I = \varepsilon_B = 0.016$

Scaling Parallel Fast Summation on BlueGene/P

silica melt with 103680 particles: RMS-force error 2.03×10^{-5}

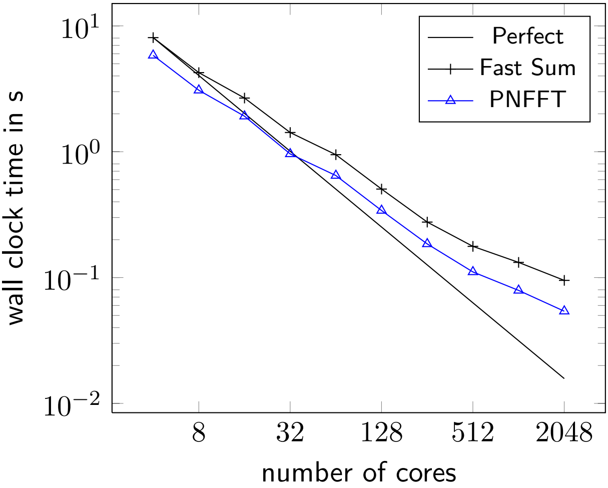

Parameters: $N = 256, n = 288, m = 4, \varepsilon_I = \varepsilon_B = 0.016$

Scaling Parallel Fast Summation on BlueGene/P

silica melt with 103680 particles: RMS-force error 2.03×10^{-5}

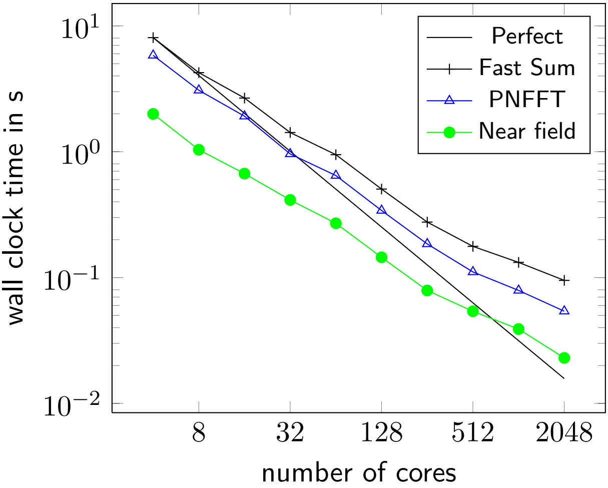

Parameters: $N = 256, n = 288, m = 4, \varepsilon_I = \varepsilon_B = 0.016$

Scaling Parallel Fast Summation on BlueGene/P

silica melt with 103680 particles: RMS-force error 2.03×10^{-5}

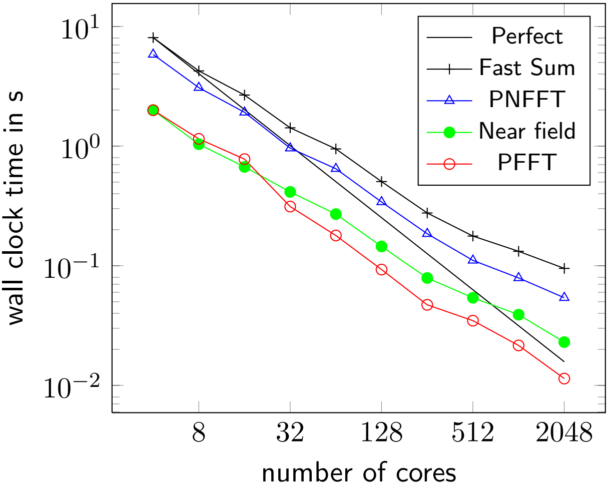

Parameters: $N = 256, n = 288, m = 4, \varepsilon_I = \varepsilon_B = 0.016$

Summary

Parallel FFT

$\boldsymbol{F}$ $\boldsymbol{F}^{\mathsf{H}}$

Summary

Parallel FFT

Window Convolution

Parallel NFFT

$\boldsymbol{C}\,\boldsymbol{F}\,\boldsymbol{D}$ $\quad$ $\boldsymbol{D}\boldsymbol{F}^{\mathsf{H}}\,\boldsymbol{C}^{\top}$

Summary

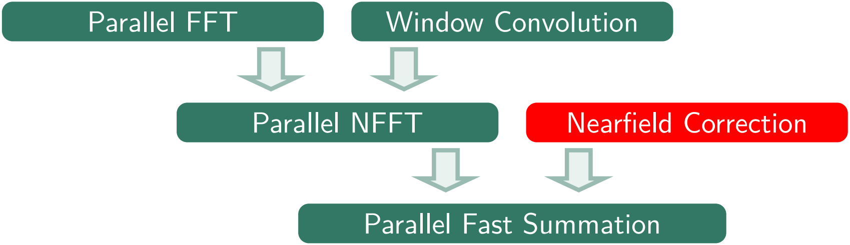

$$\boldsymbol{C}\,\boldsymbol{F}\,\boldsymbol{D}\,\mathrm{diag}(\hat{R}_{\boldsymbol{k}})\;\boldsymbol{D}\boldsymbol{F}^{\mathsf{H}}\,\boldsymbol{C}^{\top} + \boldsymbol{C}^{\mathrm{near}}$$

Summary

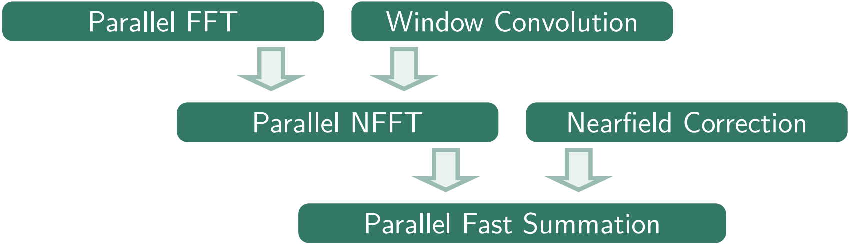

$$\boldsymbol{C}\,\boldsymbol{F}\,\boldsymbol{D}\,\mathrm{diag}(\hat{R}_{\boldsymbol{k}})\ \boldsymbol{D}\boldsymbol{F}^{\mathsf{H}}\,\boldsymbol{C}^{\top} + \boldsymbol{C}^{\mathrm{near}}$$

PFFT & PNFFT Software Library and Papers

Available at

`http://www.tu-chemnitz.de/~mpip`