A Fast Spectral Sum-of-Gaussians Method for Electrostatic Summation in Quasi-2D Systems

Xuanzhao Gao

Joint work with Shidong Jiang, Jiuyang Liang, Zhenli Xu, and Qi Zhou

2025-8-18

Hong Kong University of Science and Technology

Molecular dynamics simulation

Molecular dynamics simulation is a widely used method based on theoretical modeling and computer simulation, which plays a key role in many fields, including materials science, biophysics, and drug design.

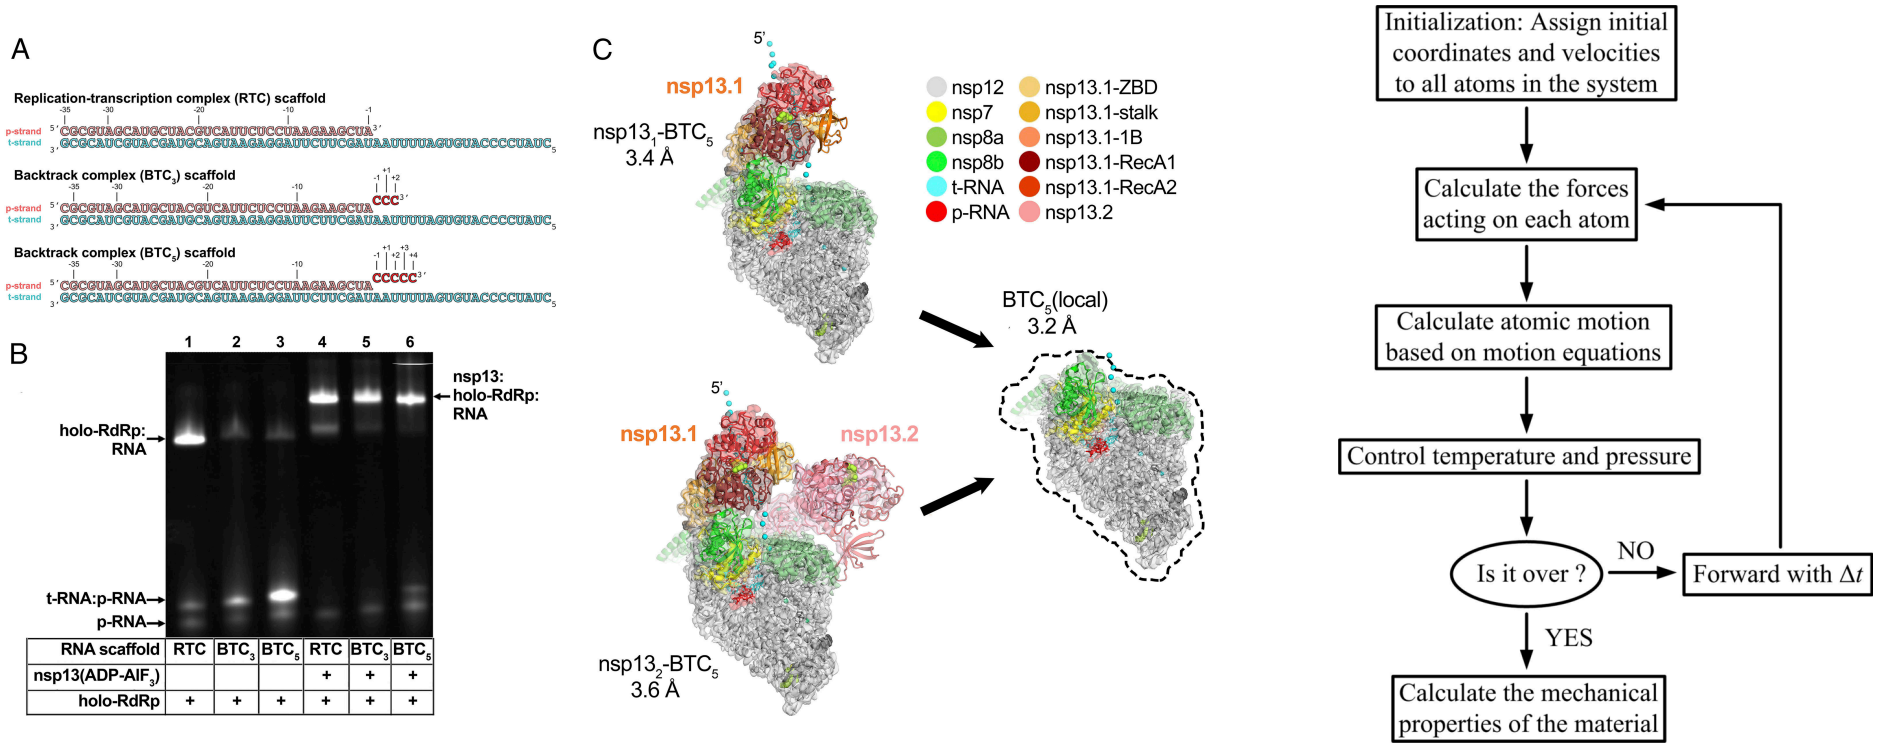

Figure 1: (Left) Study of the structure of SARS-CoV-2 replication-transcription complex by molecular dynamics simulation (Malone et al., 2021). (Right) Illustration of the workflow of molecular dynamics simulation.

Quasi-2D charged systems

Quasi-2D systems (Mazars, 2011) are macroscopic in x and y, but microscopic in z. They are widely found in nature and engineering, e.g., cell membranes, electrolytes near surfaces, and ultrathin polymer films.

They are typically modeled as infinite planar layers and treated as doubly periodic in simulations. We consider Q2D systems with point charges interacting via the Coulomb potential.

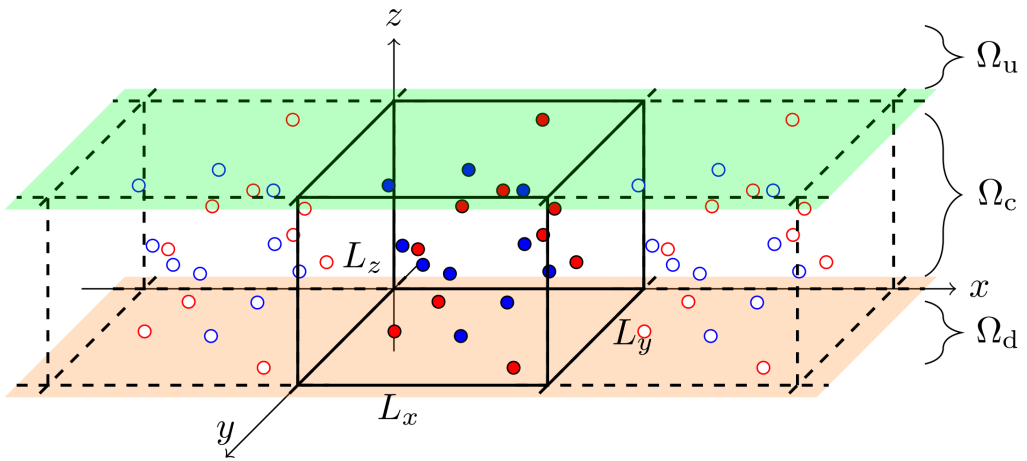

Figure 2: Illustration of a quasi-2D charged system.

Coulomb interaction

The bottleneck of simulating Q2D charged systems is the Coulomb interaction.

The electrostatic interaction in Q2D systems is given as follows:

$$U = \frac{1}{2} \sum_{i,j=1}^{N} \sum_{\boldsymbol{m}} q_i q_j G\big(\boldsymbol{r}_i, \boldsymbol{r}_j + L_{\boldsymbol{m}}\big),$$

where $G(\boldsymbol{r}, \boldsymbol{r}')$ is the Green's function, $L_{\boldsymbol{m}} = \big(L_x m_x, L_y m_y, 0\big)$.

If the system is homogeneous, the Green's function is given by:

$$G(\boldsymbol{r}, \boldsymbol{r}') = \frac{1}{|\boldsymbol{r} - \boldsymbol{r}'|}$$

It decays slowly and is singular at $r = 0$, which makes such summation computationally expensive.

Algorithms for Q2D charged systems

Methods have been developed to accelerate the calculation of Coulomb interaction in Q2D systems.

- **Ewald2D** (Parry, 1975): based on the Ewald splitting, $O(N^2)$ complexity.

To reduce the complexity, most methods rely on the following strategies:

- **Fourier spectral method** (Maxian et al., 2021; Yuan et al., 2021): based on Ewald splitting and fast Fourier transform (FFT), with $O(N \log N)$ complexity. Example: PPPM2D, ICM-PPPM.
- **Fast multipole methods** (Greengard & Rokhlin, 1987; Liang et al., 2020): accelerated by hierarchical low-rank compression, adaptive and with $O(N)$ complexity. Example: HSMA, periodic FMM.
- **Random batch Ewald** (Jin et al., 2021): based on Ewald splitting and random batch sampling, stochastic and with $O(N)$ complexity, efficient parallelization.

Is the problem fully solved?

Ewald splitting

$$\frac{1}{r} \approx \underbrace{\left(\frac{1-\operatorname{erf}(\alpha r)}{r}\right)\mathbb{1}_{r<r_c}}_{\text{near-field, real space}} + \underbrace{\frac{\operatorname{erf}(\alpha r)}{r}}_{\text{far-field, Fourier space}}$$

The near-field contribution is computed by a real space truncation, and the far-field part is a periodic summation of a smooth kernel, which can be summed up in Fourier space.

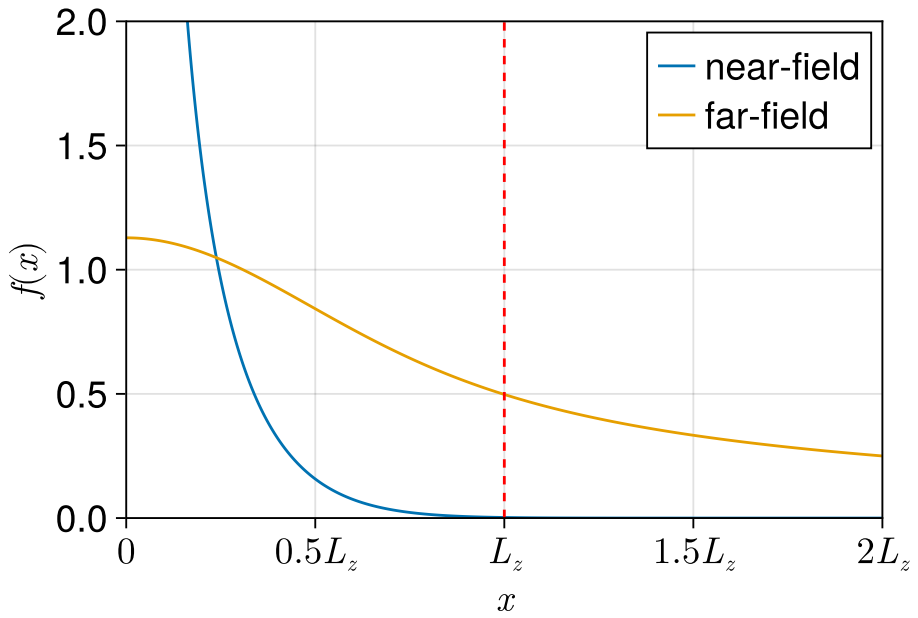

Figure 3: Near-field and far-field parts of the Coulomb kernel.

For triply periodic systems, the long range part can be calculated by standard Fourier spectral solvers (Darden et al., 1993):

$$U_{\text{long}}^{\text{3D}} = \frac{1}{2V} \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{k} q_i q_j \frac{e^{-\frac{k^2}{4\alpha^2}}}{k^2} e^{ik(\boldsymbol{r}_i - \boldsymbol{r}_j)}$$

For the Q2D systems, a common strategy is to **periodize the system in z**, approximating the origin system with a 3D periodic one.

However, Q2D systems are naturally **multi-scaled**, i.e., $L_z \ll L_x, L_y$, which leads to large prefactors compared to 3D-PBC solvers (Gao et al., 2025).

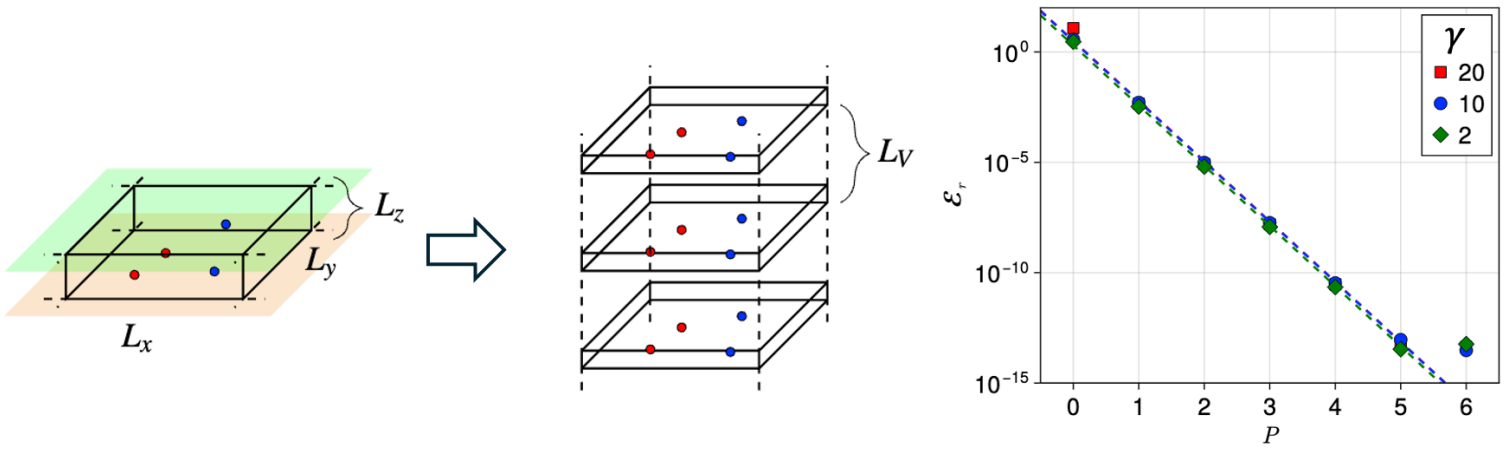

Figure 4: Zero-padding for FFT based methods. $\gamma = \frac{L_x}{L_z}$, $P := \frac{L_V}{L_x}$.

The sum-of-Gaussians approximation

The bilateral series approximation (Beylkin & Monzón, 2010) of the Coulomb kernel:

$$\frac{1}{r} \approx \frac{2\log b}{\sqrt{2\pi\sigma^2}} \sum_{l=-\infty}^{\infty} \frac{1}{b^l} e^{-\frac{r^2}{(\sqrt{2}b^l\sigma)^2}}, \text{ with } \mathcal{E}_r < 2\sqrt{2}e^{-\frac{\pi^2}{2\log b}}, r > 0$$

With the SOG approximation, we further split the potential into three parts:

$$\frac{1}{r} \approx \underbrace{\left(\frac{1}{r} - \sum_{l=0}^{M} w_l e^{-\frac{r^2}{s_l^2}}\right) \mathbb{1}_{r<r_c}}_{\text{near-field}} + \underbrace{\sum_{l=0}^{m} w_l e^{-\frac{r^2}{s_l^2}}}_{\text{mid-range}} + \underbrace{\sum_{l=m+1}^{M} w_l e^{-\frac{r^2}{s_l^2}}}_{\text{long-range}}$$

where m is selected so that $s_m < \eta L_z < s_{m+1}$, where η is $O(1)$ constant.

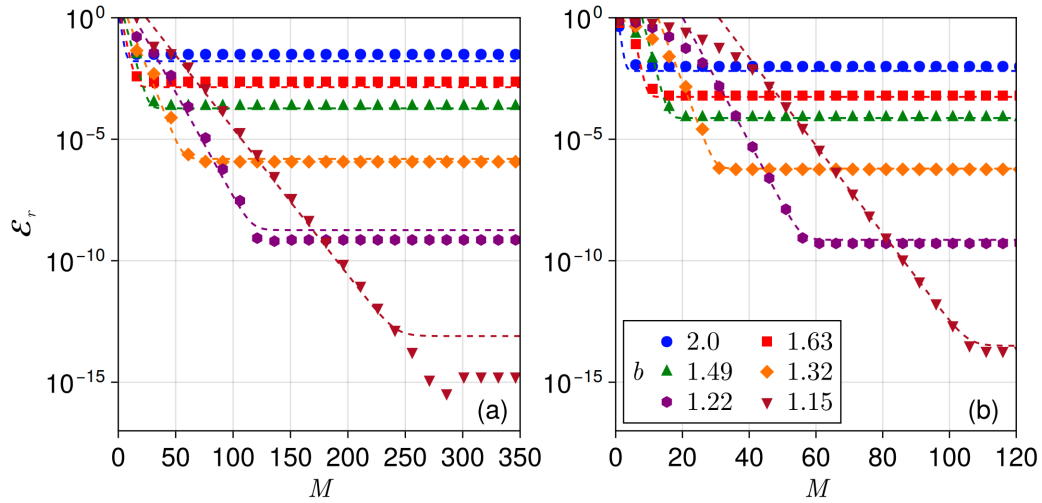

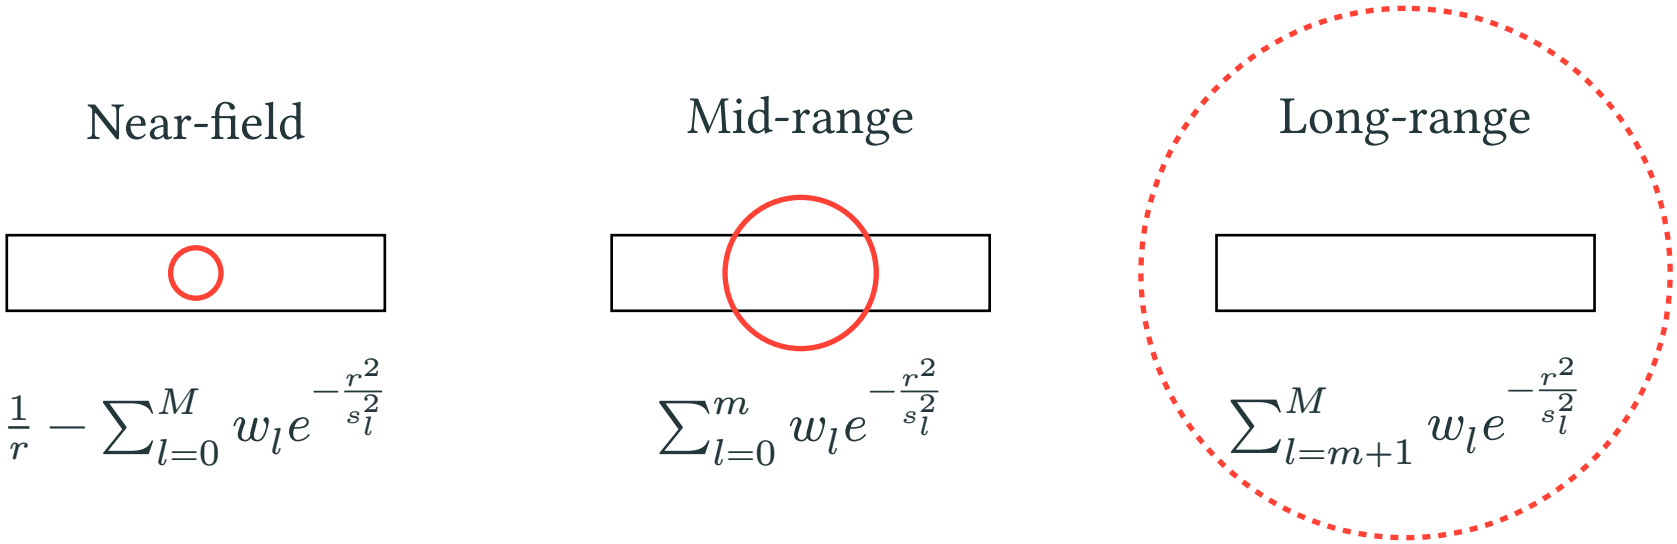

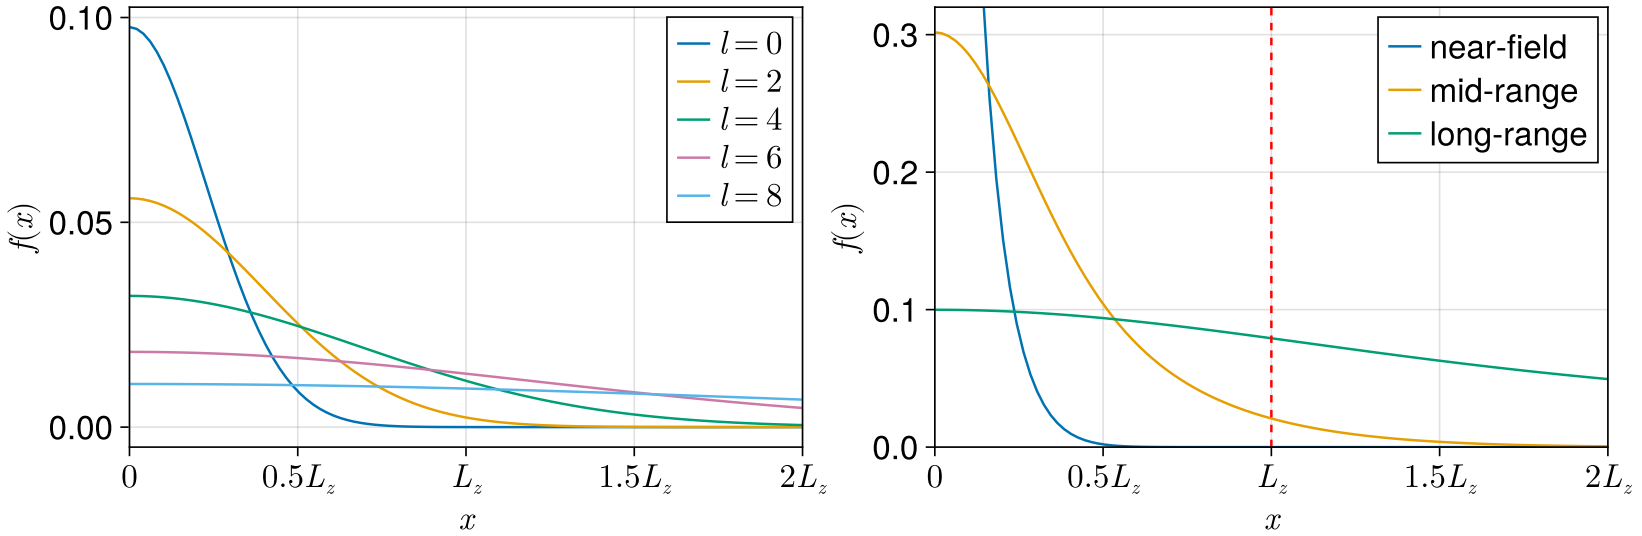

Figure 5: Mid-range part and long-range part of the potential, $\eta \approx 0.6$.

Near-field potential

The weight of the narrowest Gaussian is modified to be

$$w_0 = \omega \frac{2 \log b}{\sqrt{2\pi\sigma^2}}$$

to enforce the C^0 and C^1 continuity (Predescu et al., 2020) of the near-field potential at $r = r_c$, which is important for MD simulations (Shamshirgar et al., 2019).

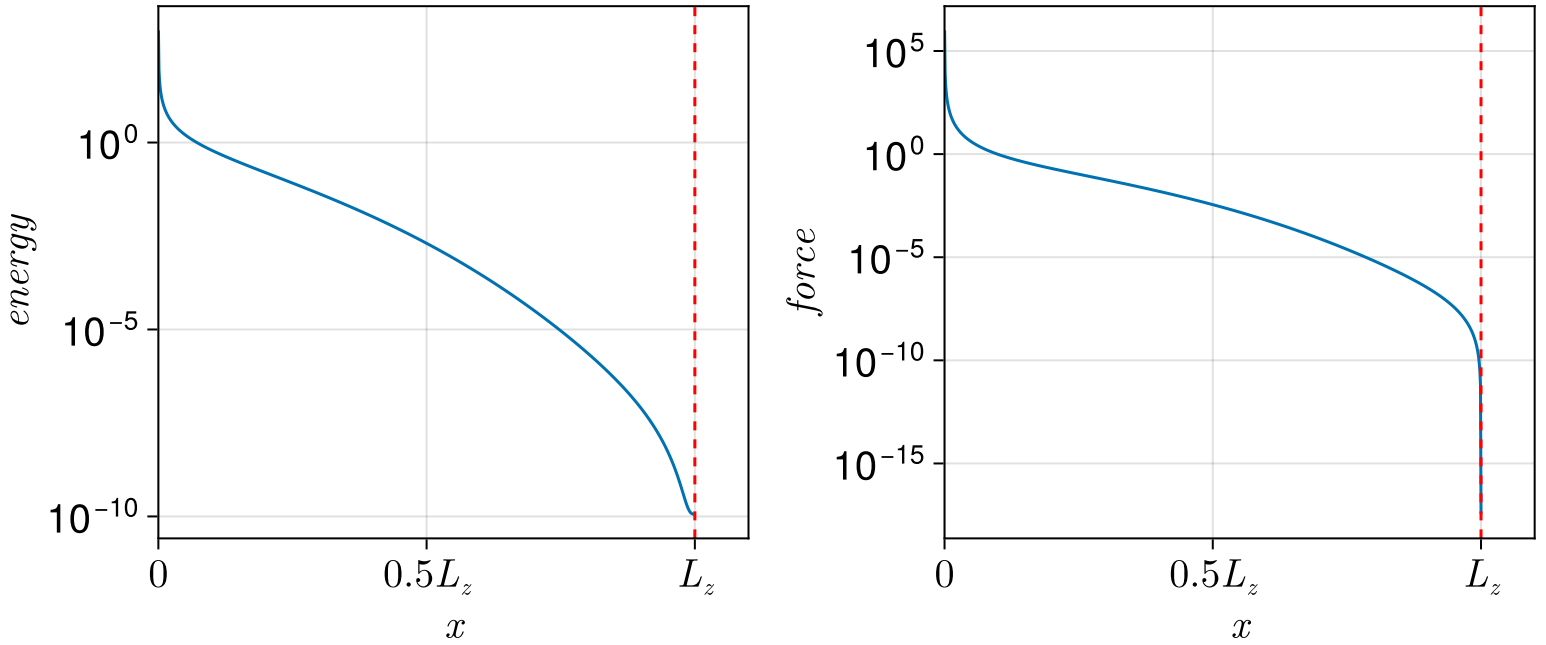

Mid-range potential

$$\Phi^l_{\text{mid}}(\vec{r}) = \sum_{\vec{n}} \sum_{j=1}^{N} q_j w_l e^{-\frac{\left(\vec{r}-\vec{r}_j+\vec{n}\circ\vec{L}\right)^2}{s_l^2}}, \quad s_l < \eta L_z$$

Since the mid-range Gaussians decay rapidly to zero for $z > L_z$, the mid-range potential is computed by a standard Fourier spectral solver (ChebParticleMesh.jl) with little zero padding ($\lambda_z < 2$ for double precision).

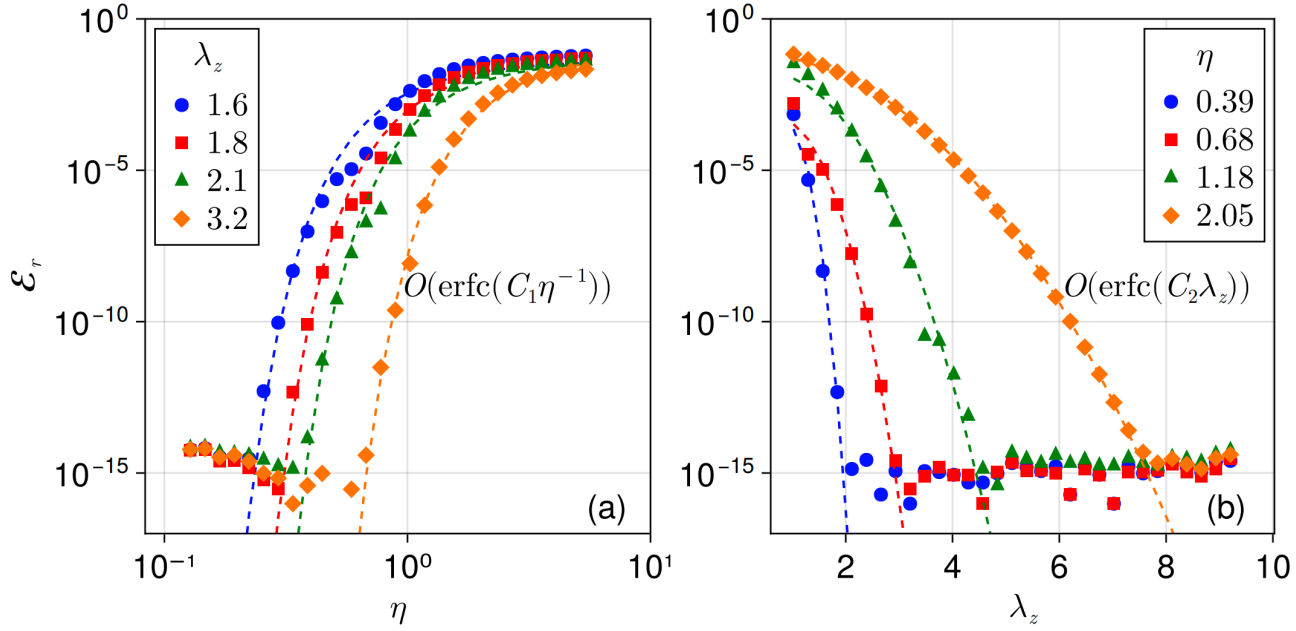

Figure 6: Relation between the padding ratio λ_z, η and the error.

Long-range potential

$$\Phi^l_{\text{long}}(\vec{r}) = \sum_{\vec{n}} \sum_{j=1}^{N} q_j w_l e^{-\frac{\left(\vec{r}-\vec{r}_j+\vec{n}\circ\vec{L}\right)^2}{s_l^2}}, \quad s_l > \eta L_z$$

The long-range potential is computed by a Fourier-Chebyshev solver.

The extremely smooth long-range Gaussians are interpolated on the Chebyshev proxy points in z; **only $O(1)$ Chebyshev points are required**.

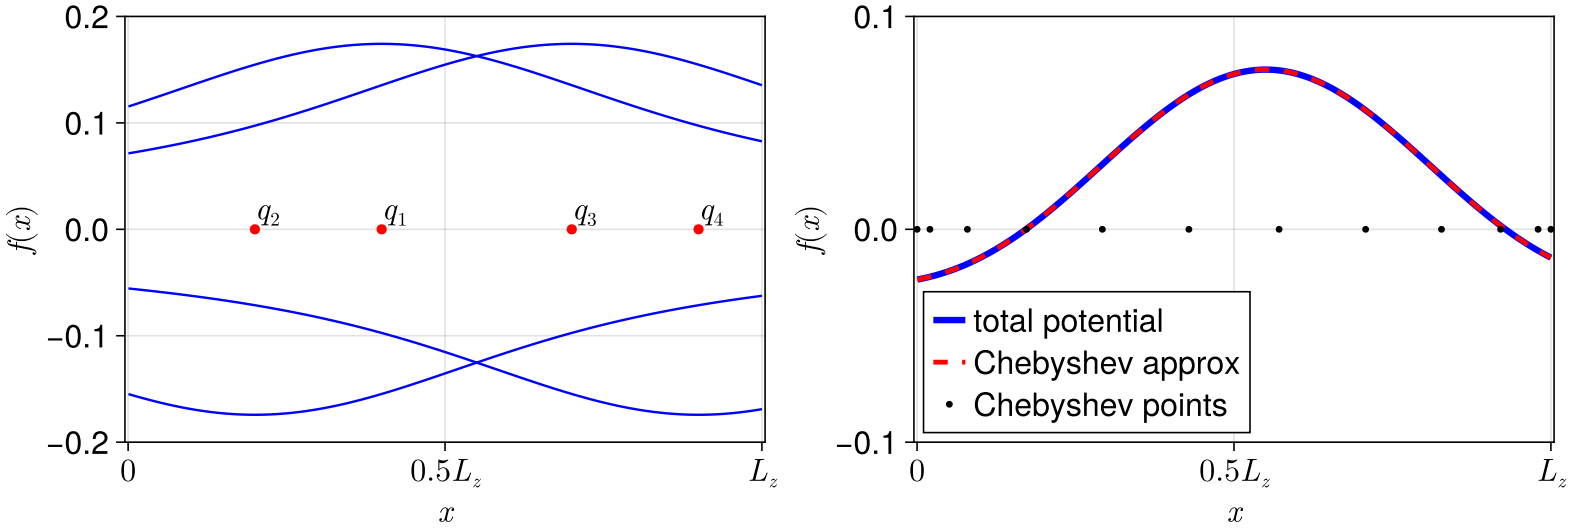

Then 2D non-uniform FFT (Greengard & Lee, 2004) can be used to evaluate the potential on a tensor-product grid, with cost of $O(N \log N)$.

Complexity

Using DFT for long-range potential, the complexity is

$$O(\underbrace{4\pi r_c^3 \rho_r N}_{\text{near-field}} + \underbrace{\mathcal{P}_x \mathcal{P}_y \mathcal{P}_z N + \frac{\lambda_z\left(1+\frac{\delta}{L_z}\right)}{r_c^3 \rho_r} N \log N}_{\text{mid-range}} + \underbrace{\frac{P L_x L_y}{\eta^2 L_z^2} N}_{\text{long-range}})$$

where $\mathcal{P}_x, \mathcal{P}_y, \mathcal{P}_z$ are the window supports, λ_z is the padding ratio, δ is the extended length of the box in the free direction to accommodate the support of the window function, P is the number of Chebyshev points. By taking $r_c \sim O(1)$ and assume $L_z \sim O\left(\sqrt{L_x L_y}\right)$, the complexity is $O(N \log N)$.

Using 2D-NUFFT for long-range potential, the complexity is

$$O(\underbrace{4\pi r_c^3 \rho_r N}_{\text{near-field}} + \underbrace{\mathcal{P}_x \mathcal{P}_y \mathcal{P}_z N + \frac{\lambda_z\left(1+\frac{\delta}{L_z}\right)}{r_c^3 \rho_r} N \log N}_{\text{mid-range}} + \underbrace{\mathcal{P}_x \mathcal{P}_y P N + \frac{P}{\rho_r L_z^3 \eta^2} N \log N}_{\text{long-range}})$$

which is needed when $L_z \ll L_x, L_y$, the total complexity is also $O(N \log N)$.

Numerical results

The method is benchmarked on the following systems:

- Cubic systems with fixed aspect ratio equals to 1.
- Strongly confined systems with fixed L_z, aspect ratio up to $10^{3.5}$

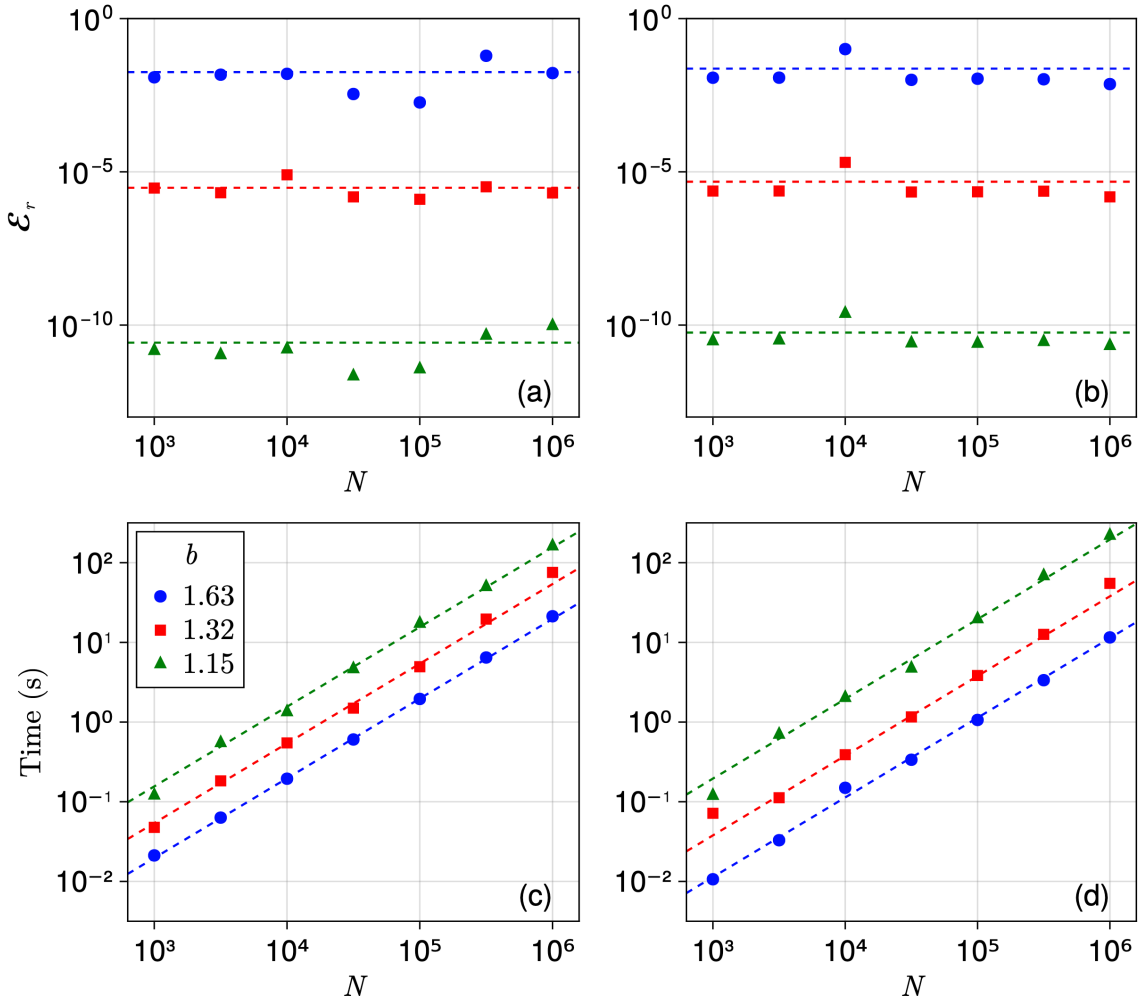

Figure 7: Error and time cost for the SOG method in the (a,c) cubic and (b,d) strongly confined systems.

Summary

A fast and accurate solver for Q2D charged systems is developed based on the sum-of-Gaussian approximation of the Coulomb kernel and the kernel splitting technique.

The Coulomb kernel is splitted into three parts:

- Near-field: solved by real space truncation
- Mid-range: solved by the Fourier spectral method with little zero padding
- Long-range: solved by the Fourier-Chebyshev method with $O(1)$ number of Chebyshev points

The solver **addresses challenges arising from singularities and strong confinement**, and has the following advantages:

- spectrally accurate with rigorous error analysis (Liang et al., 2023)
- requires little or no zero padding for high-aspect-ratio rectangular boxes
- no upsampling in the gridding and gathering steps
- all computations occur within the fundamental cell
- simple to implement and parallelize for large-scale MD simulations

Limitation: non-adaptive; complexity is $O(N \log N)$ rather than $O(N)$.

Acknowledgments

Prof. Zecheng Gan
HKUST(GZ)

Prof. Zhenli Xu
SJTU

Prof. Shidong Jiang
Flatiron Institute

Jiuyang Liang
Flatiron Institute

Qi Zhou
SJTU

Links

Source code and preprint are available on GitHub and arXiv.

FastSpecSoG.jl, available on GitHub

arXiv preprint: 2412.04595

Thank you for your attention!

Appendix

The U-series and its derivative

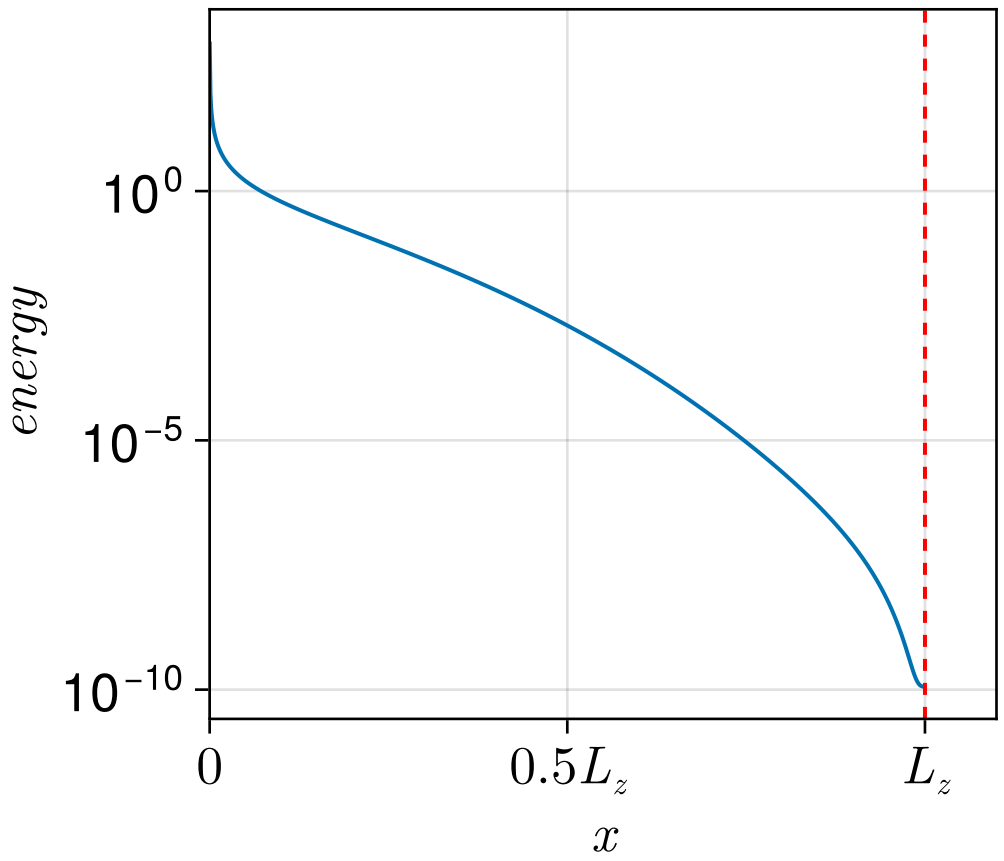

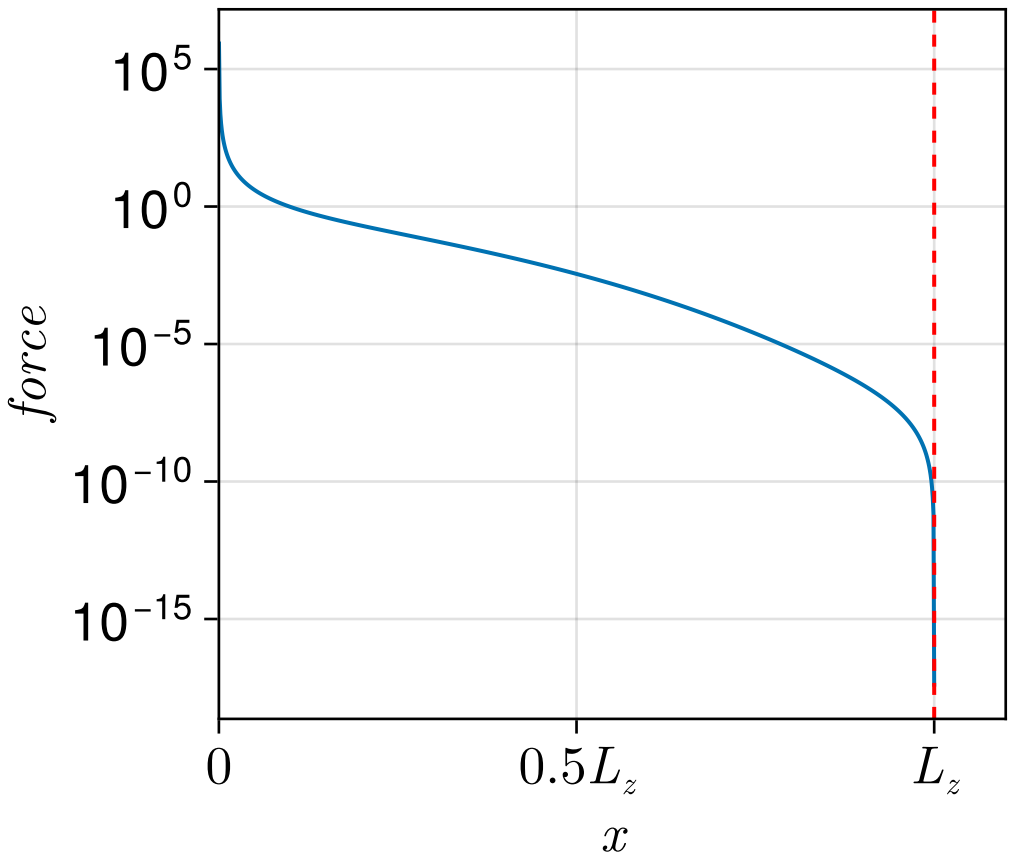

Figure 8: The U-series and its derivative, $r_c = 10.0$.

b	r_0	ω	Energy		Force	
			Error	M	Error	M
2	1.9892536839080267	0.9944464927622323	3.12×10^{-2}	16	9.93×10^{-3}	11
1.62976708826776469	2.7520026668023417	1.0078069793438068	2.33×10^{-3}	31	6.21×10^{-4}	16
1.48783512395703226	3.7554672283554990	0.9919117057598183	2.29×10^{-4}	46	7.98×10^{-5}	26
1.32070036405934420	4.3914554711638349	1.0018891411481198	1.18×10^{-6}	76	5.76×10^{-7}	41
1.21812525709410644	5.6355288151271085	1.0009014615603334	7.14×10^{-10}	166	5.14×10^{-10}	71
1.14878150173321925	7.2956245490719404	1.0000368348358225	1.30×10^{-15}	271	1.98×10^{-14}	116

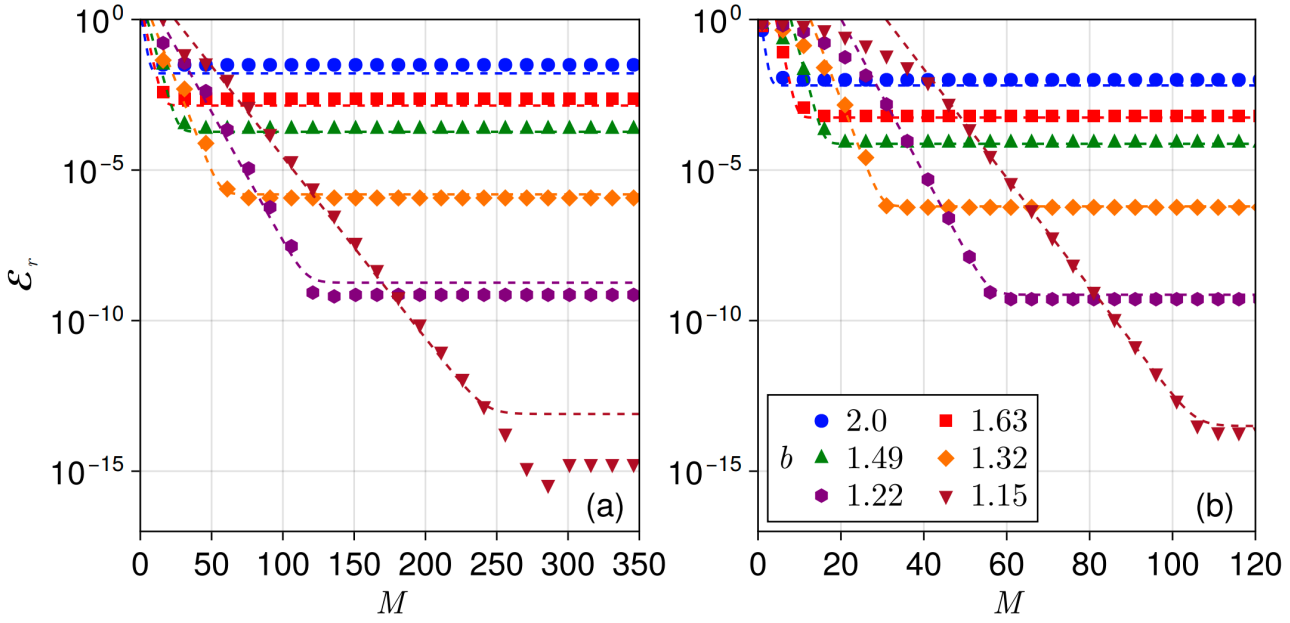

Figure 9: U-series parameters and the error

Window function

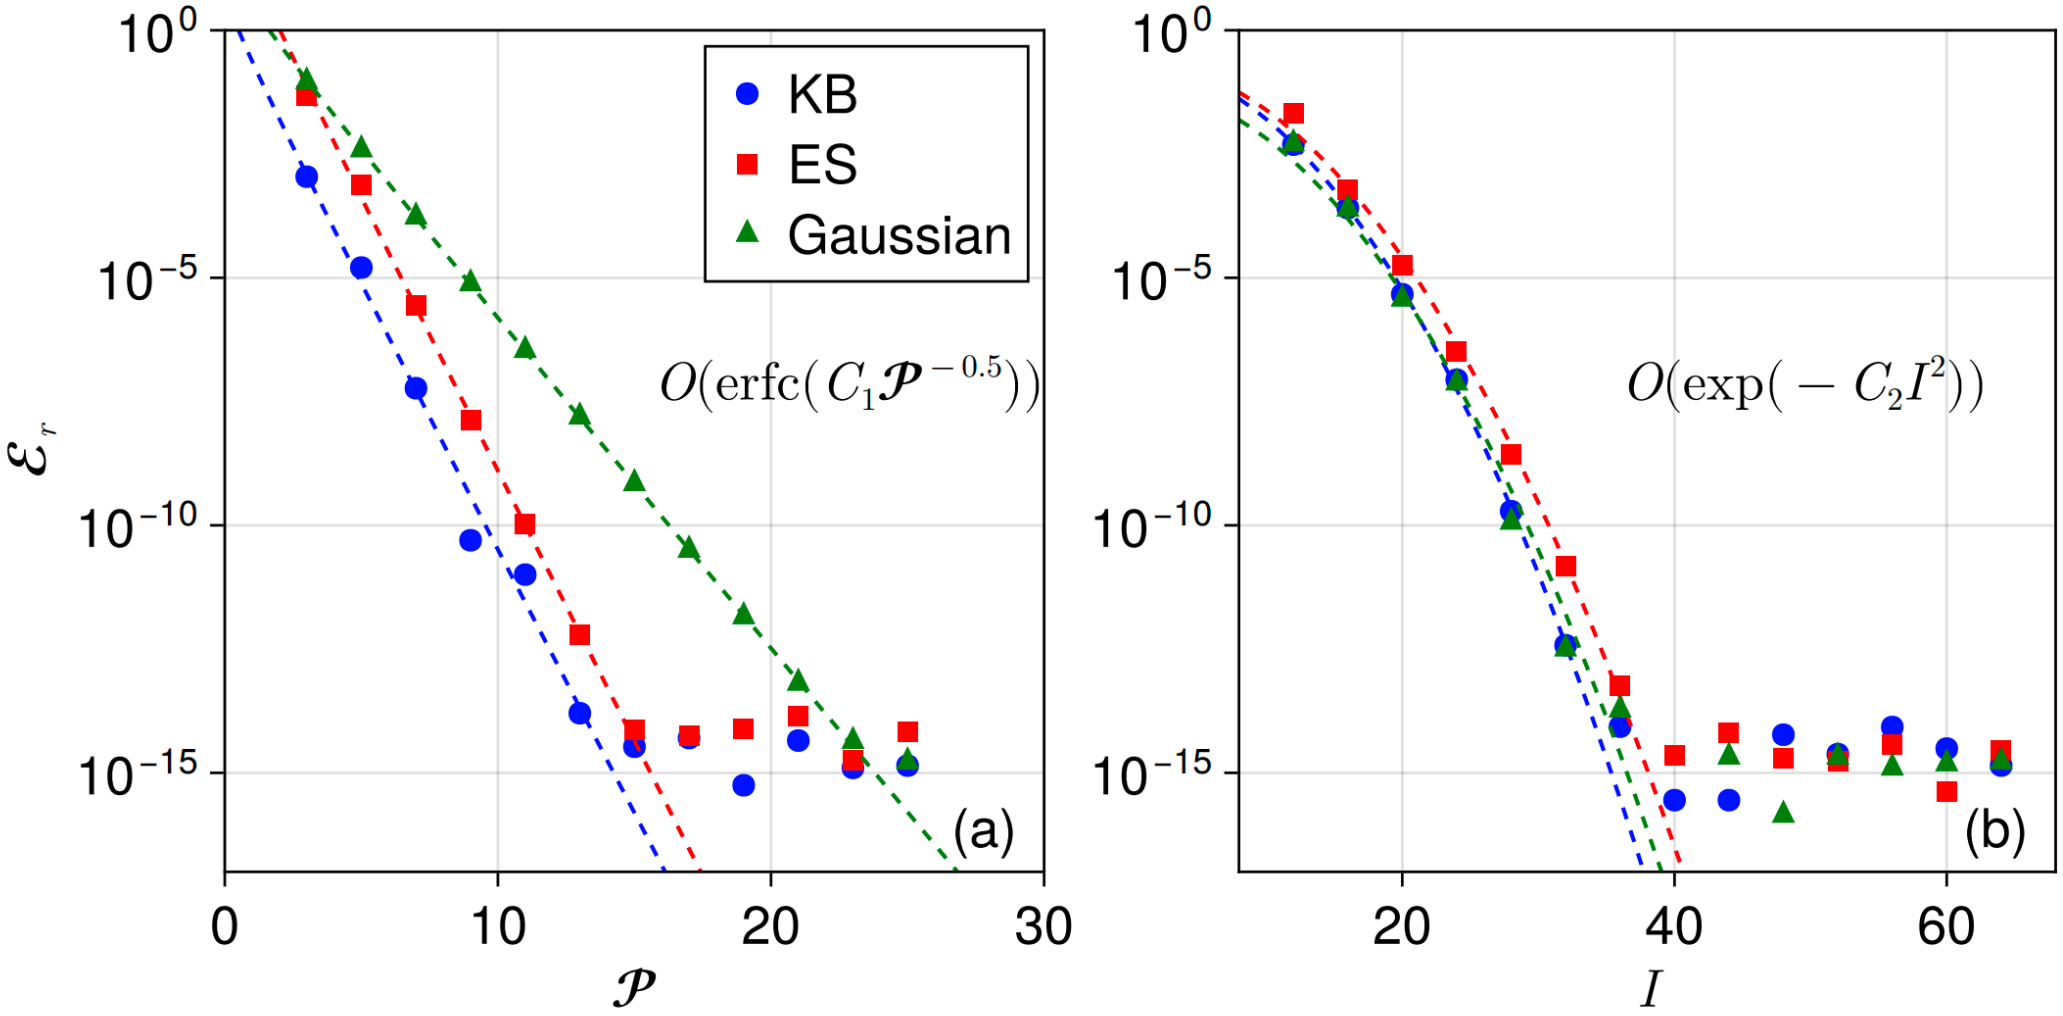

Figure 10: Different window functions.

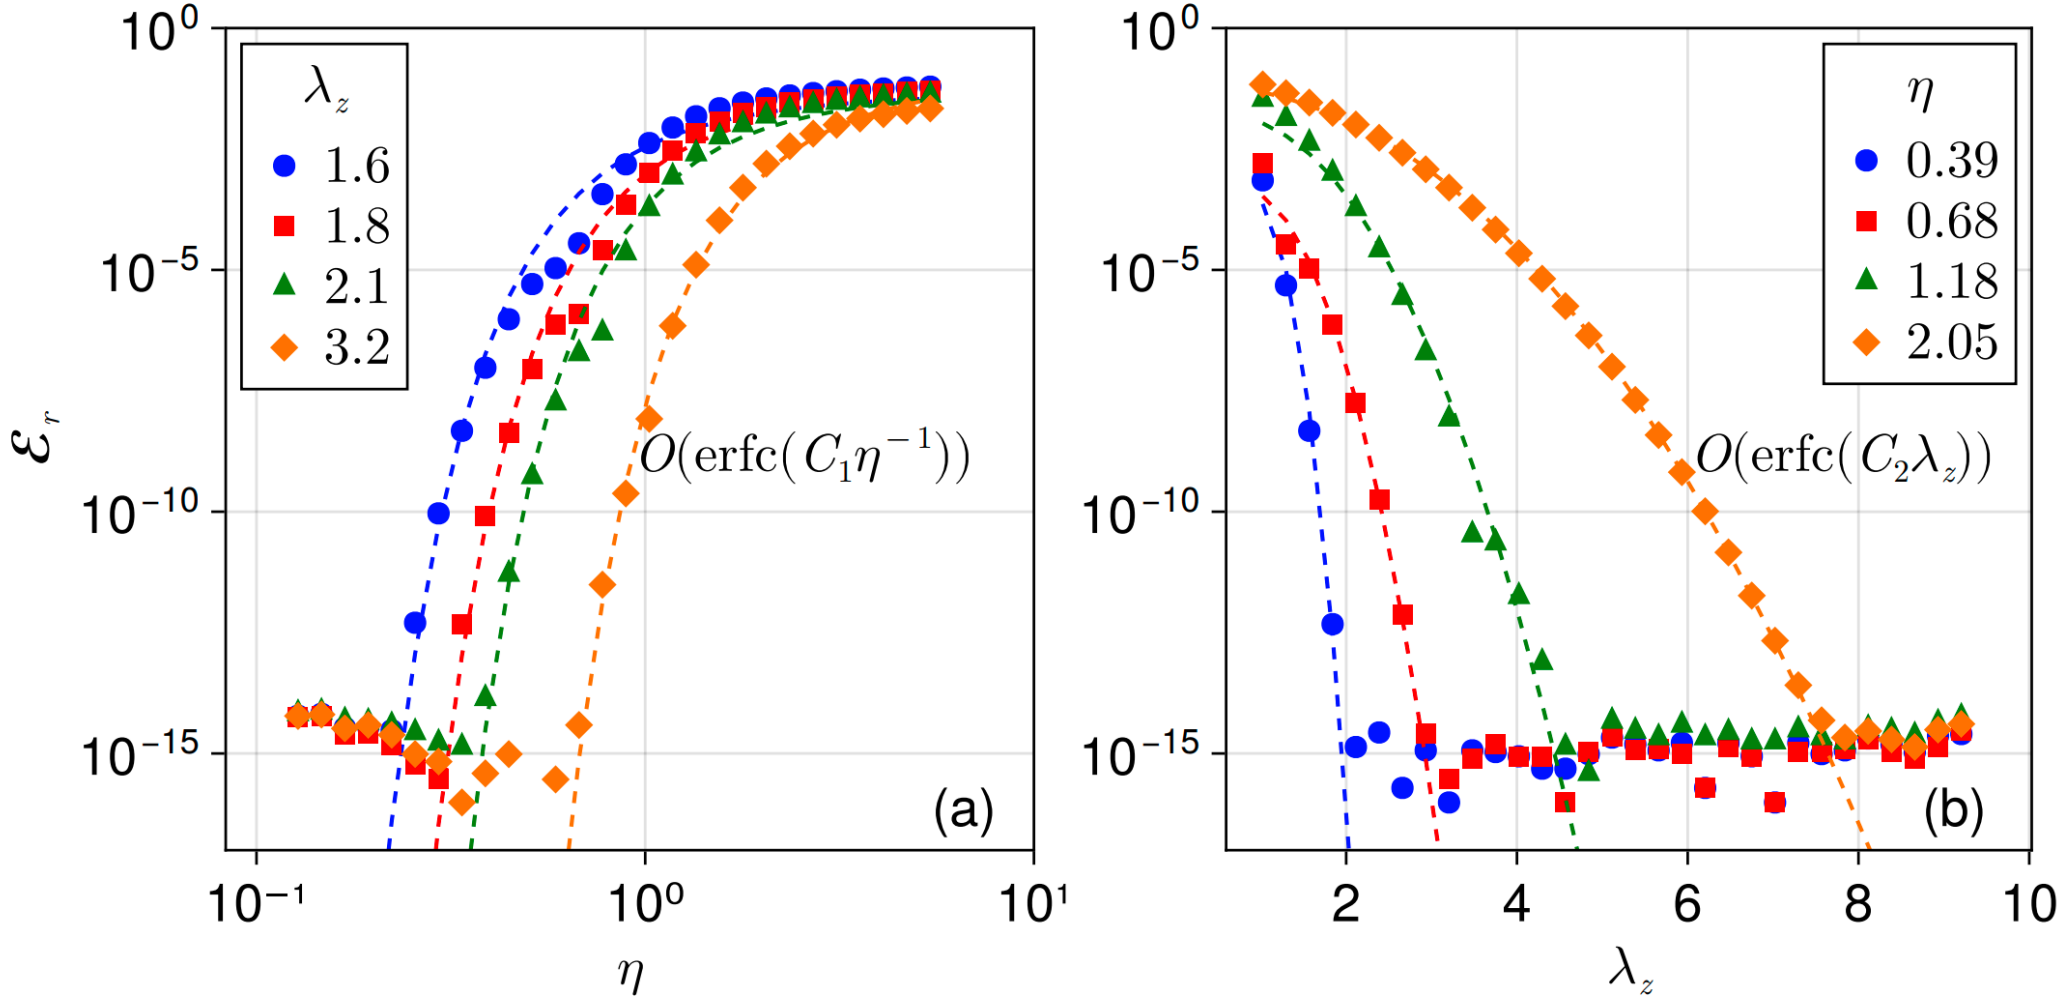

Figure 11: Zero-padding.

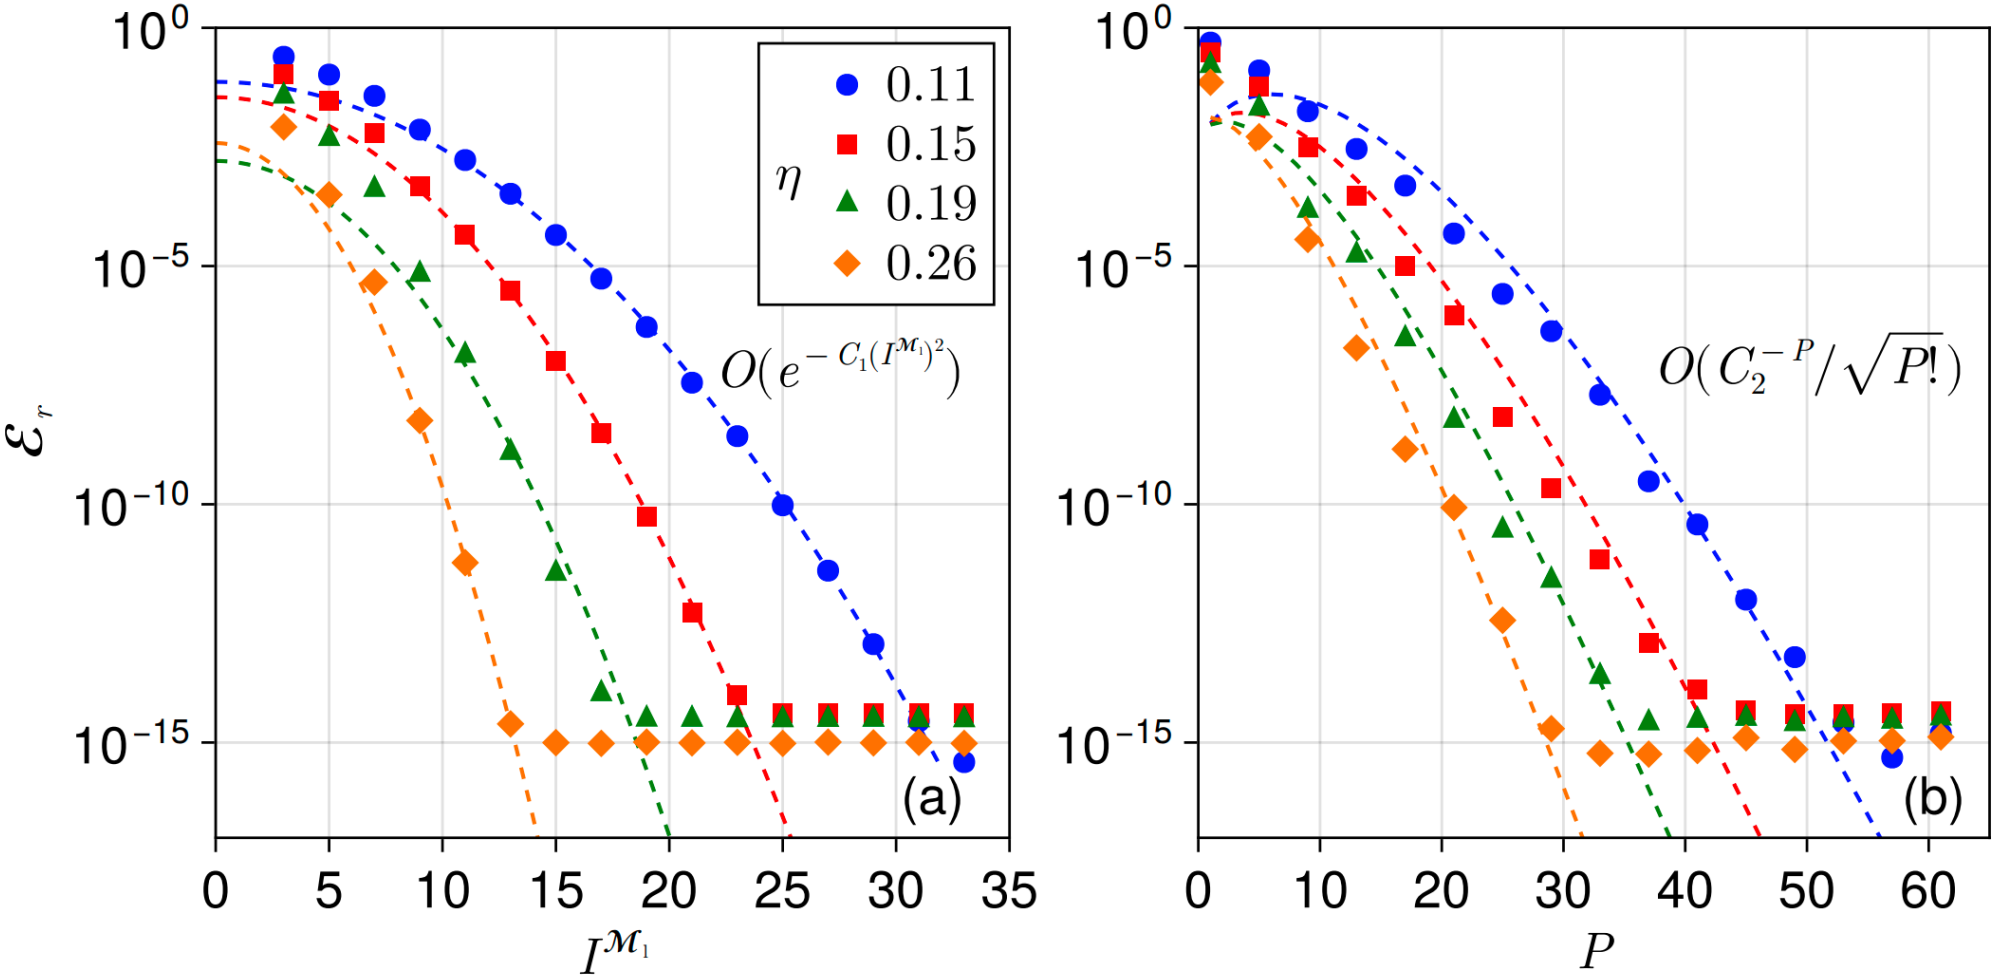

Figure 12: Accuracy of the Fourier-Chebyshev solver.

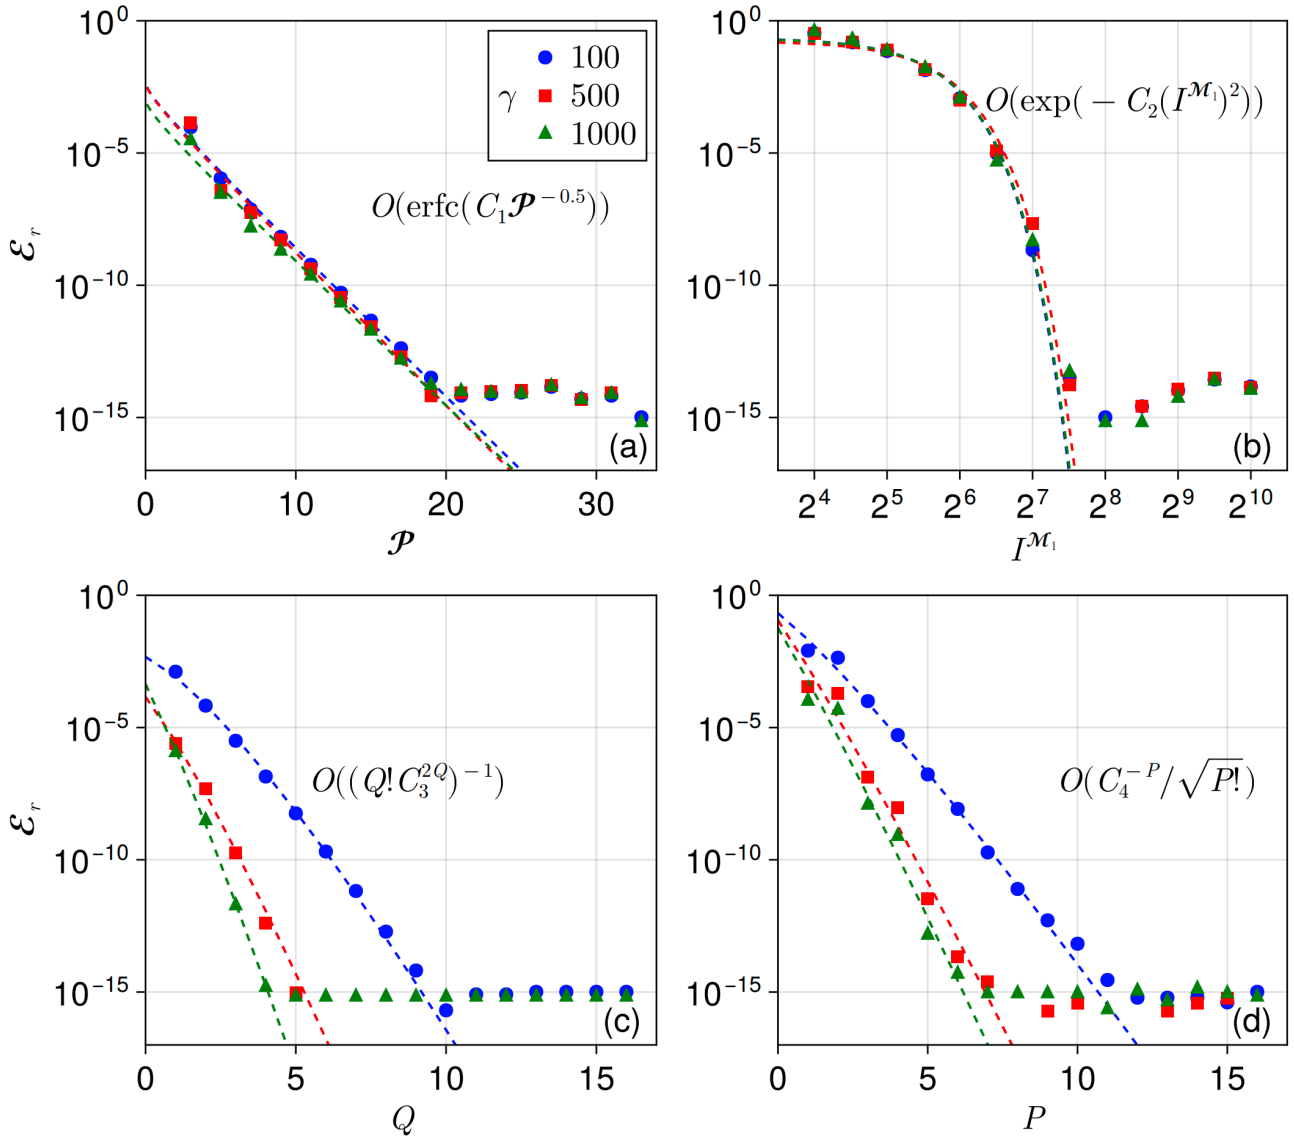

Figure 13: Strongly confined systems.

References

Beylkin, G., & Monzón, L. (2010). Approximation by exponential sums revisited. *Applied and Computational Harmonic Analysis*, *28*(2), 131–149.

Darden, T., York, D., & Pedersen, L. (1993). Particle mesh Ewald: An N*c* ·log(N) method for Ewald sums in large systems. *J. Chem. Phys.*, *98*(12), 10089–10092.

Gao, X., Zhou, Q., Gan, Z., & Liang, J. (2025). Accurate Error Estimates and Optimal Parameter Selection in Ewald Summation for Dielectrically Confined Coulomb Systems. *Journal of Chemical Theory and Computation*, *21*(12), 5890–5904. https://doi.org/10.1021/acs.jctc.5c00438

Greengard, L., & Lee, J.-Y. (2004). Accelerating the nonuniform fast Fourier transform. *SIAM Rev.*, *46*, 443–454.

Greengard, L., & Rokhlin, V. (1987). A fast algorithm for particle simulations. *J. Comput. Phys.*, *73*(2), 325–348.

Jin, S., Li, L., Xu, Z., & Zhao, Y. (2021). A random batch Ewald method for particle systems with Coulomb interactions. *SIAM J. Sci. Comput.*, *43*(4), B937–B960.

Liang, J., Xu, Z., & Zhou, Q. (2023,). *Error estimate of the u-series method for molecular dynamics simulations*. https://arxiv.org/abs/2305.05369

References

Liang, J., Yuan, J., Luijten, E., & Xu, Z. (2020). Harmonic surface mapping algorithm for molecular dynamics simulations of particle systems with planar dielectric interfaces. *J. Chem. Phys.*, *152*(13), 134109.

Malone, B., Chen, J., Wang, Q., Llewellyn, E., Choi, Y. J., Olinares, P. D. B., Cao, X., Hernandez, C., Eng, E. T., Chait, B. T., Shaw, D. E., Landick, R., Darst, S. A., & Campbell, E. A. (2021). Structural basis for backtracking by the SARS-CoV-2 replication–transcription complex. *Proceedings of the National Academy of Sciences*, *118*(19), e2102516118.

Maxian, O., Peláez, R. P., Greengard, L., & Donev, A. (2021). A fast spectral method for electrostatics in doubly periodic slit channels. *J. Chem. Phys.*, *154*(20), 204107.

Mazars, M. (2011). Long ranged interactions in computer simulations and for quasi-2D systems. *Phys. Rep.*, *500*(2–3), 43–116.

Parry, D. E. (1975). The electrostatic potential in the surface region of an ionic crystal. *Surf. Sci.*, *49*(2), 433–440. https://doi.org/10.1016/0039-6028(75)90362-3

Predescu, C., Lerer, A. K., Lippert, R. A., Towles, B., Grossman, J., Dirks, R. M., & Shaw, D. E. (2020). The u-series: A separable decomposition for electrostatics computation with improved accuracy. *J. Chem. Phys.*, *152*(8), 84113.

References

Shamshirgar, D. S., Yokota, R., Tornberg, A.-K., & Hess, B. (2019). Regularizing the fast multipole method for use in molecular simulation. *The Journal of Chemical Physics*, *151*(23), 234113.

Yuan, J., Antila, H. S., & Luijten, E. (2021). Particle–particle particle–mesh algorithm for electrolytes between charged dielectric interfaces. *J. Chem. Phys.*, *154*(9), 94115.