




ELSEVIER

Contents lists available at ScienceDirect

Computer Physics Communications

journal homepage: www.elsevier.com/locate/cpc

Computer Programs in Physics

Floating-point-consistent cross-verification methodology for reproducible and interoperable DDA solvers with fair benchmarking

Clément Argentin ^a, Patrick C. Chaumet ^b, Michel Gross^c, Maxim A. Yurkin ^{a,*}^a Université Rouen Normandie, INSA Rouen Normandie, CNRS, CORIA UMR 6614, Rouen, 76000, France^b Institut Fresnel, Aix Marseille Univ, CNRS, Centrale Marseille, Marseille, 13013, France^c Université Montpellier II, CNRS, L2C UMR 5221, Montpellier, 34090, France

ARTICLE INFO

Communicated by Dr. W Jong

Keywords:

Light scattering
Discrete dipole approximation
Reproducibility
Benchmark
Cross-verification

ABSTRACT

The discrete dipole approximation (DDA) is a widely used and versatile numerical method for solving electromagnetic scattering by arbitrarily shaped objects. Despite its popularity, quantitative comparisons between independent implementations remain challenging due to differences in linear-system conventions, solver settings, and default numerical parameters. In this work, we introduce a unified software-assisted methodology for cross-verification and benchmarking of three major open-source DDA solvers: DDSCAT, ADDA, and IFDDA. We demonstrate how machine-precision agreement can be achieved across implementations by aligning all free parameters and provide practical equivalence tables enabling reproducible and interoperable simulations. Using this methodology, we perform systematic CPU and GPU performance comparisons covering OpenMP, MPI, and CUDA/OpenCL parallelization. Beyond benchmarking, our approach serves as a practical guide for configuring consistent DDA simulations and for understanding how precision, solver choice, and hardware architecture affect runtime, scalability, and accuracy in computational light-scattering studies. The software package also supports regression testing and bitwise reproducibility verification for future code releases.

PROGRAM SUMMARY

Program Title: dda-bench*CPC Library link to program files:* <https://doi.org/10.17632/vzfy32vfn7.1>*Developer's repository link:* <https://doi.org/10.5281/zenodo.18836855>*Licensing provisions:* GPLv3*Programming language:* Python

Nature of problem: Independent implementations of the discrete dipole approximation (DDA), such as DDSCAT, ADDA, and IFDDA, are widely used for electromagnetic scattering by arbitrarily shaped objects. However, quantitative cross-code comparisons and fair performance benchmarks remain difficult because the codes may differ not only in linear-system conventions, but also in default numerical parameters, solver settings, polarizability and interaction formulations, and output conventions or units. Small mismatches can modify iterative convergence histories and lead to discrepancies that mask genuine algorithmic or hardware effects. Users and developers therefore lack a practical and reproducible way to (i) verify that two codes are solving the same numerical problem, (ii) assess the number of matching digits that should be expected when specific parameters differ, and (iii) benchmark runtimes without confounding accuracy differences.

Solution method: The presented software provides a lightweight, command-line-based framework to compare numerical DDA implementations through their executables. The user specifies, in a dedicated input file, the command-lines to be tested for each code or code version. The Python wrapper automatically executes the specified commands, collects the generated output files, and extracts the physical quantities selected by the user for comparison. The extracted data are converted automatically into consistent definitions and units when necessary. Agreement between simulations is quantified using a matching-digits metric, which directly reflects floating-point consistency. This approach enables straightforward comparison between different DDA codes or between multiple versions of the same solver, thus supporting regression testing, cross-verification, and fair benchmarking once equivalent configurations are enforced.

* Corresponding author.

E-mail addresses: clement.argentin@coria.fr (C. Argentin), patrick.chaumet@univ-amu.fr (P.C. Chaumet), gross.lkb1@gmail.com (M. Gross), yurkin@gmail.com (M.A. Yurkin).<https://doi.org/10.1016/j.cpc.2026.110172>

Received 5 March 2026; Received in revised form 16 April 2026; Accepted 21 April 2026

Available online 30 April 2026

0010-4655/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Additional comments including restrictions and unusual features: The software is limited to DDA implementations that provide a command-line interface and assumes that the user has already installed and compiled the corresponding codes. The framework itself acts as a Python wrapper that operates on the generated executables. This design choice facilitates the practical use of the framework as a lightweight Python package that can be installed and imported easily, while allowing users to track and compare different releases of the same code and to freely choose which executable versions are used in a given comparison.

1. Introduction

Light scattering by particles is a fundamental process governing a wide range of optical phenomena [1]. Quantitative modeling of scattering, absorption, and near-field effects is essential for interpreting experiments and designing optical systems. Analytical solutions of Maxwell's equations are limited to a few idealized geometries (e.g., spheres or spheroids), making numerical approaches indispensable for realistic systems of arbitrary shape, composition, and environment.

Among general numerical techniques, several broad families of methods exist: finite-difference time-domain (FDTD) [2], finite-element (FEM) [3], boundary-element or surface-integral equation (BE-M/SIE) [4], and volume-integral equation (VIE) [5] ones. Within the latter category, the discrete dipole approximation (DDA) is commonly used due to its conceptual simplicity, modularity, and favorable scaling with particle size and wavelength [6,7]. Introduced in Ref. [8], and further developed in Refs. [9,10], DDA discretizes the scatterer into an array of polarizable voxels (dipoles) and computes their induced polarizations by solving a large linear system derived from the volume-integral form of the Maxwell's equations. Once the voxel polarizations are known, all scattering observables, including cross sections, Mueller matrices, and field distributions, can be computed in both near- and far-field regions. DDA is employed for arbitrarily shaped particles, and complex optical environments, across disciplines from astrophysics [9] and atmospheric science [11] to nanotechnology [12] and biomedical applications [13]. Moreover, as it is a "numerically exact" method, DDA is often used as a reference to validate or correct approximate analytical models [14,15].

Over the past three decades, several open-source DDA codes have been developed, each optimized for different architectures and applications. The most widely used are DDSCAT [10], ADDA [16], and IFDDA [17]. An extensive list of publicly available DDA implementations and related tools is available online [18,19].

Despite their shared theoretical foundation, comparing the numerical performance and accuracy of different DDA implementations remains nontrivial. Discrepancies often arise from subtle differences in polarizability formulations, iterative solvers, or linear system conventions. As a result, cross-code benchmarking [20] is challenging to interpret, since the codes have different accuracy and different computational performance.

The aim of this paper is therefore twofold. First, we introduce a unified software-assisted methodology for floating-point-consistent cross-verification between independent DDA implementations. By harmonizing all free and hidden parameters, we demonstrate that solvers can agree up to machine precision in double precision (i.e. 14–15 correct digits). Focusing on DDSCAT, ADDA, and IFDDA, we provide practical equivalence tables that enable reproducible and interoperable DDA simulations. Because these implementations rely on different linear-system conventions, and FFT backends, strict bitwise reproducibility is not achievable across codes, as even tiny differences in floating-point operation ordering lead to round-off divergence. Such limitations are well-known concerns in high-performance scientific computing, particularly for parallel numerical simulations [21,22].

Second, using this consistent numerical baseline, we conduct a fair benchmarking study across CPUs and GPUs, covering OpenMP and MPI parallelization on CPUs and CUDA/OpenCL acceleration on GPUs. Performance differences observed in this study therefore reflect algorithmic and hardware architectural effects rather than configuration inconsistencies.

Beyond cross-verification, the software package also supports comparisons between different versions of the same implementation. In that regime, true bitwise reproducibility is achievable, enabling continuous integration and continuous delivery (CI/CD) regression tests and version-to-version consistency checks. For instance, this will complement the automatic tests already present in the ADDA workflow.

Finally, this paper aims to serve as a reference for best practices in setting up DDA simulations, and understanding the impact of implementation choices. Section 2 provides a comparative overview of the three DDA codes and their internal features; Section 3 acts as a methodological guide and a technical prerequisite for the performance benchmarks presented in Sections 4 and 5.

2. DDA implementations and workflows

The main objective of this section is to provide a concise and self-consistent description of the DDA workflow, introducing its key internal components to facilitate the discussion in later sections. We also summarize distinctive features of each considered implementation so that this paper serves as a practical comparative overview of the three main open-source DDA codes.

All three implementations: DDSCAT, ADDA, and IFDDA, aim to simulate the interaction of electromagnetic waves with arbitrarily shaped particles or objects. Both DDSCAT and ADDA are primarily designed for general light-scattering problems, and particularly suited for databases and large-scale parametric studies. By contrast, IFDDA focuses on applications involving complex optical setups, including multilayered substrates and microscopy configurations. It can generate observable quantities such as intensity images making it especially suitable for modeling and visualization tasks in applied or experimental optics [23].

All three codes implement similar DDA workflows but differ in numerical details and hardware optimization strategies. Each supports several Krylov-space iterative solvers (of the conjugate-gradient, CG, family), which are used for solving the large sparse linear systems. As noted in Ref. [24], it is generally not possible to identify a single solver that performs best in all scenarios, since performance depends strongly on particle morphology and refractive index. In Ref. [24], the QMR solver (used in IFDDA) is implemented in its non-symmetrized form (as in DDSCAT); when symmetrized, as done in ADDA, the number of matrix-vector products is reduced by a factor of two, resulting in similar convergence rates and execution times as other CG-family solvers. Consequently, solver performance comparisons between codes must take these implementation details into account. For consistency, all performance results presented in Sections 4 and 5 are therefore based on the BiCGStab solver.

DDA has undergone several refinements in polarizability models and interaction formulations [7,25]. The classical approach, i.e., point-dipole interaction (POI) combined with the lattice dispersion relation (LDR) polarizability [10], remains widely used, while more modern formulations such as the filtered coupled dipoles (FCD) [26,27] and the integrated Green's tensor (IGT) [28,29] have been developed to improve both accuracy and iterative convergence [7]. Fortunately, all three codes support at least one of these advanced formulations. FCD and IGT come with their own polarizability prescriptions, while POI support many others [25]: Clausius-Mossotti (CM), radiative-reaction (RR) correction, digitized Green's function (DGF / GB), Lakhtakia (LAK), and lattice dispersion relation (LDR) or corrected lattice dispersion relation (CLDR).

Table 1
Overview of selected features in DDSCAT, ADDA, and IFDDA.^a

Feature / Code	DDSCAT	ADDA	IFDDA
Application focus	general	general	multi-layered substrates and microscopy
Community and contributions	open source (website)	open source (GitHub)	open source (GitLab)
User interface	CLI-oriented (configuration file) with optional GUI	CLI-oriented with optional GUI	GUI-oriented with optional CLI
Language	Fortran 90	C99	Fortran 77
Hardware support	CPU only	CPU and GPU (OpenCL)	CPU and GPU (CUDA)
Parallelization	OpenMP + MPI (partial) ^b	MPI (full)	OpenMP
Orientation averaging	Simpson rule	Romberg integration	–
Precision	single (default) or double	double (default)	double (default) or single
Incident wave types	plane wave only	plane wave and others	plane wave and others
FFT algorithms	GPFA, Intel MKL	GPFA, FFTW	Singleton, FFTW
DDA formulations	POI, FCD	POI, FCD, IGT	POI, FCD, IGT
Polarizability models	LDR, CLDR, FCD	CM, RR, LDR, CLDR, FCD, LAK, DGF	CM, RR, LDR, FCD, LAK, GB
Initial field	zero only	zero and others	zero and others
Preconditioner	left (Jacobi) or none	none	left, right, both, (Chan) or none
Volume correction	yes	optional (yes/no)	no
Iterative solver	BCGS2, BiCGStab, GPBiCG, QMR, PETRKP	BCGS2, BiCG, BiCGStab, CGNR, CSYM, QMR, QMR2	GPBiCG, GPBiCGStab, BiCGStab, CG, QMR, IDRS and many others
Unique features	periodic target	Bessel beams, EELS/CL, sparse mode, cuboid voxels, homogeneous substrate (with 3D FFT)	multi-layered substrates, microscopy, fully anisotropic target

^a Used abbreviations are explained in the text.

^b Only for different particle orientations (orientation averaging) in contrast to distributed-memory parallelization of a single simulation in ADDA.

The incident-field models differ across implementations. DDSCAT primarily supports plane-wave illumination. ADDA additionally provides (approximate) Gaussian beams and point-dipole excitation, accepts user-specified arbitrary fields from file, and in its master branch provides a variety of built-in Bessel beams [30]. Development branches further implement electron energy-loss spectroscopy (EELS) and cathodoluminescence (CL) for particles in arbitrary host medium [31]. IFDDA supports physical Gaussian beams and a wider set of structured illuminations (circularly polarized, multi-plane-wave, speckle, confocal, and arbitrary imported fields), and allows targets to be embedded in multilayer substrates. It can also generate microscopy images (e.g., bright-field, dark-field, phase-contrast, holography) from the computed fields.

Iterative convergence may be accelerated by advanced initial fields and appropriate preconditioners. IFDDA implements several advanced initial guesses such as uGu that can be combined with Chan-type preconditioning [32] to reduce iteration counts [33]. Recently, an implementation of a three-level circulant preconditioner extended this work and has been shown to further reduce iteration counts [34]. ADDA provides, in addition to the classic zero and incident field, a Wentzel–Kramers–Brillouin (WKB) initial field option, while DDSCAT provides the zero-field initialization with a possible left Jacobi preconditioning for BiCGStab solver. It should be noted that although these initial fields may be beneficial in some cases, they require initialisation time to be calculated and can therefore increase calculation time if used outside their limits. A recent study [35] compared the initial WKB field with the uGu formulation [36] and showed that their performance was broadly equivalent. Beyond rigorous but computationally expensive DDA, IFDDA also offers approximate direct methods that are effective for multilayer transmission-reflection problems [37].

All implementations are portable across operating systems and require few external dependencies. By default, DDSCAT uses the self-contained GPFA FFT algorithm [38], which requires only a Fortran compiler (e.g., `gfortran`), making it highly portable. ADDA can also be compiled with GPFA (named `Temperton`), ensuring compatibility without external libraries (except for `gcc` and `gfortran` compilers). However, as shown in Section 4.2, GPFA underperforms compared to optimized FFT libraries such as Intel MKL routine `DFTI` and `FFTW` [39]. The same applies to IFDDA with the `Singleton` algorithm [40]. Therefore, the recommended configurations are DDSCAT MKL, ADDA FFTW, and IFDDA FFTW. GPU modes introduce further dependencies: ADDA requires `OpenCL` and `cFFTW`, while IFDDA relies on `CUDA`. ADDA also supports package managers such as `Spack` [41] or `Nix`, facilitating portable and reproducible

installations in contained environments and high-performance computing (HPC) clusters.

HPC has strongly influenced the evolution of the DDA codes. DDSCAT and IFDDA are optimized for shared-memory architectures using `OpenMP`, whereas ADDA employs `MPI` to leverage distributed-memory clusters. Recently, GPU accelerators have become standard for accelerating matrix–vector operations. In this context, IFDDA uses `CUDA`, the most mature and widely optimized GPU framework, but limited to `NVIDIA` hardware, while ADDA uses `OpenCL` to ensure cross-vendor portability across `NVIDIA`, `AMD`, and `Intel` GPUs, at the cost of lower performance.

Floating-point precision is another important design choice. Double precision ensures better numerical stability but increases memory traffic and runtime, whereas single precision offers higher arithmetic throughput and lower memory requirements, especially beneficial for GPUs, where available memory is often limited. As demonstrated in Sections 4.3 and 5.3, single precision can yield substantial speedups with acceptable error levels for well-conditioned problems. However, convergence of the iterative solver may degrade for some cases, illustrating the trade-offs between performance, portability, and numerical accuracy inherent to the modern DDA implementations.

We used the latest releases (or commits) of each code: DDSCAT v7.3.4, IFDDA v1.0.27, and ADDA v.1.5.0-alpha3 (commit `b03d648`). Minor, non-intrusive modifications were applied to enable machine-precision cross-comparisons in our benchmark tests; all such changes are documented in Appendix A. Build configurations (compilers, flags, libraries) are reported alongside the benchmarks to ensure reproducibility.

3. Accuracy, interoperability, and guidelines

This section serves as both a methodological guide and a technical prerequisite for the performance benchmarks. We first analyze the main sources of numerical error in DDA simulations and clarify what “accuracy” means in this context. We then outline the interoperability methodology used to align the parameters of different DDA codes, providing practical guidance for achieving consistent, reproducible, and high-precision results.

3.1. Understanding accuracy in DDA simulations

The accuracy of a DDA simulation depends on several interacting factors. DDA is fully deterministic, and its errors can therefore be

decomposed into well-defined categories. Understanding these sources of error is essential for interpreting the observed differences between simulation results and for assessing the level of agreement achievable between different implementations.

- **Iterative convergence error.** The iterative solver used to obtain the internal electric field must reach a specified tolerance η , defined as the relative norm of the residual (see Eq. (16)). The corresponding relative error in the field and derived quantities (e.g., extinction cross section) is typically η , which is generally negligible in comparison to the method error below (for the used default values of $\eta = 10^{-4}$ or 10^{-5}). Still, it may become an issue if η is increased in an effort to accelerate computations [31].
- **Method error.** This represents the intrinsic accuracy of DDA compared to the exact electromagnetic solution. It is composed of two main contributions:
 - *Discretization error* which appears from approximating the target by a finite number of voxels N . Accuracy improves with increasing discretization grid, e.g., the number of voxels along the x -axis n_x , or equivalently, with smaller voxel spacing d [42].
 - *Shape error* due to mismatch between the discretized geometry and the actual particle shape. Apart from grid refinement, this error may be decreased by the weighted discretization [43].
 The required or expected method error largely depends on application. Typically, 1% error is considered satisfactory for integral scattering quantities, while up to 10% can be fine for angle-resolved one. We stress, however, that better accuracy is perfectly possible, and relative errors down to 10^{-7} have been demonstrated [44].
- **Round-off error.** Finite-precision arithmetic leads to accumulation of rounding errors over many iterations. When the number of matrix-vector products and iteration count N_{it} are large, these small numerical fluctuations (about 10^{-7} or 10^{-16} in single and double precision, respectively) accumulate and ultimately limit the number of matching digits to that set by the convergence threshold η .

In practice, when comparing different implementations or numerical setups, two typical levels of agreement can be identified:

- **Machine-precision agreement.** When all simulation parameters are aligned across codes, one can achieve near machine-level agreement, typically 10–15 matching digits in double precision. This defines the upper limit of reproducibility for deterministic DDA solvers.
- **η -level agreement.** When comparing results across codes with identical physical settings but different ways to solve the resulting linear system, the agreement is limited by the solver tolerance η , which bounds the numerical accuracy of the solution. In practice, this leads to agreement up to approximately $-\log_{10}(\eta)$ significant digits.

Understanding these different sources of error provides a foundation for the cross-code verification methodology described in the following subsection. In practice, machine-level agreement is only achievable when the same free parameters are used (see Section 3.4), identical precision modes are selected, and linear system conventions (see Section 3.2) agree. When any of these conditions differ, one can only hope to reach the η -level agreement (see Section 3.3), but in some cases only agreement within the method error is obtained. Note also, that η -level agreement is sufficient for most physical applications and can be considered sufficiently good for software testing, since η can be varied in a wide dynamic range.

3.2. Different linear systems (ADDA, IFDDA, DDSCAT)

In this section, we aim to identify the conditions under which the linear systems used in different DDA implementations are not just equivalent (which is always the case), but lead to exactly the same convergence trajectory of standard iterative methods. This marks the difference between η -level and machine-precision agreement in the computed quantities.

In the following, we express all quantities using the SI units. All permittivities and refractive indices are expressed relative to the vacuum background. The relative permittivity tensor is defined as

$$\bar{\epsilon}(\mathbf{r}) \stackrel{\text{def}}{=} \begin{cases} \bar{\mathbf{I}}, & \mathbf{r} \in V_{\text{int}}, \\ \bar{\epsilon}_p(\mathbf{r})/\epsilon_0, & \mathbf{r} \in V_{\text{ext}}, \end{cases} \quad (1)$$

where ϵ_0 is the vacuum permittivity, $\bar{\epsilon}_p(\mathbf{r})$ is the permittivity of the particle, V_{int} denotes the particle volume, and V_{ext} denotes the exterior region. The complex relative refractive index and electric susceptibility are defined as

$$\bar{\mathbf{m}}(\mathbf{r}) = \sqrt{\bar{\epsilon}(\mathbf{r})}, \quad \bar{\chi}(\mathbf{r}) \stackrel{\text{def}}{=} \bar{\epsilon}(\mathbf{r}) - \bar{\mathbf{I}}. \quad (2)$$

The wavenumber in vacuum is defined as $k = \omega\sqrt{\epsilon_0\mu_0}$, where μ_0 is the vacuum permeability. The free-space dyadic Green's function is given by:

$$\bar{\mathbf{G}}(\mathbf{r}, \mathbf{r}') = \frac{\exp(ikR)}{4\pi R} \left[\left(\bar{\mathbf{I}} - \frac{\mathbf{R} \otimes \mathbf{R}}{R^2} \right) + \frac{ikR-1}{k^2 R^2} \left(\bar{\mathbf{I}} - 3 \frac{\mathbf{R} \otimes \mathbf{R}}{R^2} \right) \right], \quad (3)$$

with $\mathbf{R} = \mathbf{r} - \mathbf{r}'$, $R = \|\mathbf{R}\|$, and the dyad $\mathbf{R} \otimes \mathbf{R}$.

When referring to code-specific linear systems, we use superscripts to distinguish between systems: \mathbf{A}^p for the standard polarization form [9] (DDSCAT [10]), \mathbf{A}^x for the symmetrized form [7] (ADDA [16]), and \mathbf{A}^E for the internal field form [6] (IFDDA [17]).

3.2.1. Standard DDA (polarization form, e.g., DDSCAT)

In SI units, the linear system solved in the standard polarization formulation [7], after discretization, is:

$$\bar{\alpha}_i^{-1} \mathbf{p}_i - k^2 \sum_{j \neq i} \bar{\mathbf{G}}_{ij} \mathbf{p}_j = \mathbf{E}_i^{\text{inc}} \quad (4)$$

where $\bar{\mathbf{G}}$ is the free-space dyadic Green's tensor defined in Eq. (3), and $\bar{\alpha}$ is the polarizability tensor. Note that the voxel moments \mathbf{p}_i and polarizabilities $\bar{\alpha}_i$ are scaled by a factor of $1/\epsilon_0$ (see Eq. (10) in Ref. [7]) compared to their conventional SI definitions to ensure consistency with the linear form in the CSG convention. The latter is different from Eq. (4) only by the lack of k^2 factor (see Eq. (20) in Ref. [25]).

This can be written in compact matrix form as:

$$\mathbf{A}^p \mathbf{p} = \mathbf{E}^{\text{inc}} \quad (5)$$

where the matrix elements \mathbf{A}_{ij}^p are defined as:

$$\mathbf{A}_{ij}^p = \begin{cases} \bar{\alpha}_i^{-1}, & i = j \\ -k^2 \bar{\mathbf{G}}_{ij}, & i \neq j \end{cases} \quad (6)$$

This matrix \mathbf{A}^p is complex symmetric for arbitrary complex symmetric $\bar{\alpha}_i$.

3.2.2. Symmetrized-form (change of variable, e.g., ADDA)

Change of variable based on the square root of the polarizability tensor [7,45]:

$$\mathbf{b}_i = \bar{\beta}_i \mathbf{E}_i^{\text{inc}}, \quad \mathbf{x}_i = \bar{\beta}_i^{-1} \mathbf{p}_i \quad (7)$$

where $\bar{\beta}_i = \bar{\alpha}_i^{1/2}$. Replacing into the standard DDA form leads to a transformed system:

$$\mathbf{A}^x \mathbf{x} = \mathbf{b} \quad (8)$$

with matrix elements given by:

$$\mathbf{A}_{ij}^x = \begin{cases} \bar{\mathbf{I}}, & i = j \\ -k^2 \bar{\beta}_i \bar{\mathbf{G}}_{ij} \bar{\beta}_j, & i \neq j \end{cases} \quad (9)$$

This matrix \mathbf{A}^x is complex symmetric for arbitrary complex symmetric $\bar{\alpha}_i$ and has unit diagonal (i.e, it is Jacobi-preconditioned).

3.2.3. Internal field form (e.g., IFDDA)

The local internal electric field satisfies in SI units:

$$\mathbf{E}_i^\ell = \mathbf{E}_i^{\text{inc}} + \sum_{j \neq i} k^2 \bar{\mathbf{G}}_{ij} \bar{\alpha}_j \mathbf{E}_j^\ell \quad (10)$$

The structure of the linear system is different by a factor of k^2 compared to Ref. [6]-CGS form, due to the transformation of $\bar{\mathbf{G}}$. The polarizabilities $\bar{\alpha}_i$ are scaled by a factor $1/\epsilon_0$ relative to their strict SI definitions (as discussed previously). Additionally, while in Ref. [6] the susceptibility is defined as $\bar{\chi} = \frac{1}{4\pi}(\bar{\epsilon}_r - \mathbf{I})$ and uses a Green's tensor without the $1/4\pi$ prefactor, we instead adopt the standard SI definition $\bar{\chi} = \bar{\epsilon}_r - \mathbf{I}$ and include the $1/4\pi$ factor directly in our Green's tensor (see Eq. (3)).

Here \mathbf{E}_i^ℓ is the local field, also called microscopic field [46] (according to DDSCAT, see Sec. 25 of the manual [47]), or excited field [25] \mathbf{E}^{exc} , and corresponds to the total field \mathbf{E}_i without the field induced by the voxel on itself. A general formula, in case of non-magnetic scatterer ($\mu = 1$), linking these two fields is [25]:

$$\mathbf{E}_i^\ell = [\bar{\mathbf{I}} + (\bar{\mathbf{L}}_i - \bar{\mathbf{M}}_i)\bar{\chi}_i]\mathbf{E}_i. \quad (11)$$

Here $\bar{\mathbf{L}}_i$ is the self-term dyadic, which equals to $\bar{\mathbf{I}}/3$ for cubical (or spherical) voxels, and $\bar{\mathbf{M}}_i$ is the finite-size correction, which choice determines the specific polarizability prescription. For instance, assumption of $\bar{\mathbf{M}}_i = 0$ corresponds to the so-called weak form of the DDA [48], leading to:

$$\mathbf{E}_i^\ell = \mathbf{E}_i \frac{\bar{\epsilon} + 2\bar{\mathbf{I}}}{3}. \quad (12)$$

and thus to Clausius–Mossotti polarizability.

Note that all three codes return the total field as output, while IFDDA also provides the local field. However, the initialization of an iterative solver may differ. While IFDDA always defines the initial field in terms of the local field, ADDA can define in terms of the local or total field depending on the `-init_field` mode (see Sec. 12.1 of its manual [49]). Moreover, IFDDA scales this field using the Hegedüs trick [50], which consists in applying an optimal scalar factor to the initial field so as to minimize the initial residual norm.

This leads in matrix notation to:

$$(\mathbf{I} - \mathbf{GD}_\alpha)\mathbf{E}^\ell = \mathbf{E}^{\text{inc}}. \quad (13)$$

Here $\mathbf{D}_\alpha = \text{diag}(\bar{\alpha}_j)$ is a block-diagonal matrix of polarizabilities. To be consistent with the two previous forms, we can define a matrix $\mathbf{A}^E = \mathbf{I} - \mathbf{GD}_\alpha$, where the elements of the matrix are:

$$\bar{A}_{ij}^E = \begin{cases} \mathbf{I}, & i = j \\ -k^2 \bar{\mathbf{G}}_{ij} \bar{\alpha}_j, & i \neq j \end{cases} \quad (14)$$

This matrix \mathbf{A}^E always has unit diagonal but is complex symmetric only when all $\bar{\alpha}_i$ are identical scalars.

3.2.4. Relationship between \mathbf{A}^x , \mathbf{A}^E , \mathbf{A}^p

For a homogeneous and isotropic target, the polarizability reduces to a scalar $\bar{\alpha}_i = \alpha \mathbf{I}$ for all voxels. In this special case, the three matrices reduce to

$$\mathbf{A}^x = \mathbf{A}^E = \alpha \mathbf{A}^p, \quad (15)$$

which differ only by a constant scaling. However, the linear systems also differ by the choice of primary unknown (\mathbf{E}_i^ℓ , $\mathbf{p} = \bar{\alpha} \mathbf{E}_i^\ell$, and $\mathbf{x} = \bar{\beta}^{-1} \mathbf{E}_i^\ell$) and source terms ($\mathbf{E}_i^{\text{inc}}$ and $\mathbf{b} = \bar{\beta} \mathbf{E}_i^{\text{inc}}$), with $\bar{\beta} = \bar{\alpha}^{1/2}$. Thus, when evaluating the residual:

$$r = \frac{\|\mathbf{b} - \mathbf{Ax}\|}{\|\mathbf{b}\|} \quad (16)$$

of the systems, the scalar factor cancels out, resulting in identical solutions up to machine precision, with only minimal differences caused by round-offs and the order of operations.

For anisotropic or inhomogeneous materials, however, the polarizability is no longer a scalar and may vary from voxel to voxel, which

breaks commutativity of \mathbf{G} and \mathbf{D}_α (or \mathbf{D}_β). As a result, \mathbf{A}^x , \mathbf{A}^E , and \mathbf{A}^p are no longer related by simple scalings, and their residuals cannot be reconciled merely by transformations of the primary unknowns or source terms. Nevertheless, all systems aim to approximate the same physical scattering problem, and since the DDA is a deterministic method, only one exact solution exists. The differences between the computed solutions are therefore dominated by the solver tolerance, resulting in η -level agreement.

ADDA supports anisotropic materials but only under the assumption that the dielectric tensor of the target is diagonal in the chosen frame. DDSCAT, on the other hand, can handle off-diagonal components by constructing a rotated dielectric tensor: the user provides the diagonal elements (in the local dielectric frame) along with rotation angles, and DDSCAT computes the off-diagonal components (see Sec. 28 of the manual [47]). Still, the dielectric tensor has to be symmetric. By contrast, IFDDA allows for a fully general dielectric tensor, including off-diagonal elements and lack of symmetry. However, any anisotropic materials in this code can only be combined with Clausius–Mossotti or radiative reaction correction [9] polarizabilities. Naturally, all three codes support inhomogeneous materials.

3.3. Impact of free parameters on simulation agreement

Even though DDA is a deterministic method, small differences in free parameters can strongly affect the numerical agreement between independent runs. To quantify this sensitivity, we varied individual parameters between two otherwise identical DDA simulations and compared the resulting extinction efficiencies and internal fields. The number of matching digits between the two results serves as a practical metric of agreement, where 10–15 digits correspond to full machine precision for doubles. Table 2 summarizes the expected level of agreement for each parameter category when only that parameter differs across codes or configurations.

As shown in Table 2, discrepancies in the considered scattering problem immediately break reproducibility, yielding, in general, zero matching digits. Differences in the DDA formulation or computational grid typically yield one or two significant digits, where the number of digits is quantified such that a relative error of 10^{-x} corresponds to x digits of agreement. Next, solver-related parameters (choice of algorithm, stopping criterion η , or initial field) generally control the agreement at the level of η . For example, with $\eta = 10^{-4}$, agreement is typically limited to four digits. Importantly, this may even correspond to full machine-precision agreement, provided that the solver tolerance is tightened to 10^{-16} in double precision.

3.4. Default parameters and unified command-line setup

To enable floating-point-consistent cross-verification between independent DDA codes, all relevant numerical and physical parameters must be harmonized. Although IFDDA, ADDA, and DDSCAT share a similar theoretical foundation, their default configurations differ significantly in terms of target definition, polarization handling, numerical conventions, and solver parameters. Table 3 summarizes the main default values for each code. These differences must be explicitly resolved to achieve consistent and reproducible simulations.

While all three implementations employ a plane-wave incident field by default, differences remain in polarization handling and coordinate conventions. Both ADDA and DDSCAT can compute two orthogonal polarizations simultaneously, whereas IFDDA currently performs one polarization per run. The propagation and polarization directions also differ by default. Moreover, polarizability models and solver choices differ substantially, leading to variations in iterative convergence and accuracy if not explicitly aligned.

To facilitate systematic benchmarking, we developed a Python wrapper enabling DDSCAT to be executed directly from the command line [51]. Using it we can formulate minimal command-line examples

Table 2

Expected agreement (in terms of matching digits) when varying individual parameters across DDA codes.

Parameter category	Description	Matching digits
Studied system	particle shape, refractive index, wavelength	0
Grid	number and size of voxels	1–2 ^a
DDA formulation	polarizability model and interaction formulation (e.g., LDR, FCD, IGT)	1–2
Solver parameters	type of iterative solver and convergence threshold η	η -level
Initial field	initial guess for the iterative solver (e.g., zero, WKB, or preconditioned)	η -level

^a This depicts the typical value of method error (discussed in Section 3.1), which can also be significantly smaller.**Table 3**

Default parameters across DDA codes. Differences must be resolved to achieve machine-precision equivalence.

Parameter	IFDDA	ADDA	DDSCAT
Shape	sphere	sphere	16-sphere aggregate
Size	$R = 100$ nm	$D \approx 6.734552 \mu\text{m}$ ^a	$a_{\text{eff}} = 0.25198 \mu\text{m}$
Incident field	plane wave	plane wave	plane wave
Propagation direction	\hat{z}	\hat{z}	\hat{x}
Polarization direction	\hat{y}	\hat{y}	\hat{y}
Target orientation	zxz-convention	zyz-convention	xzx-convention
Wavelength λ	632.8 nm	$2\pi \mu\text{m}$	$0.6 \mu\text{m}$
Refractive index	$\varepsilon = 1.1$	$m = 1.5$	$m = 1.33 + 0.01i$
Polarizability	RR	LDR	CLDR
Grid resolution n_x	10	16	24
Solver	GPBiCG1	QMR	PBCGS2
Convergence threshold η	10^{-4}	10^{-5}	10^{-5}
Initial field	Born	automatic	zero field
Volume correction	none	enabled	enabled
Scattering formula	Draine's formulation [9]	Draine's formulation	Draine's formulation
Unit system	SI outputs (with code equations in CGS)	CGS	CGS

^a Equal to $(6N/\pi)^{1/3}\lambda/(10|m|)$ for default values of λ and m , while $N = 2176$ for a sphere with the default value of n_x .

constructed with consistent physical and numerical parameters, yielding agreement up to machine precision between the three codes:

```

$./adda -lambda 632.8 -size 200 -init_field zero
-grid 15
$./ifdda -polarizability LS -methodeit QMRCLA
-epsmulti 2.25 0.0 -nnnr 15 -ninitest 0
-tolinit 1.d-5 -object sphere 99.5540512465
$ddscatcli -CMDSOL QMRCCG -CALPHA LATTDR
-MEM_ALLOW "15 15 15" -CSHAPE ELLIPSOID -SHPAR
"15 15 15" -DIEL "diel/m1.50_0.00"
-WAVELENGTHS "632.8 632.8 1 'LIN'" -AEFF "100
100 1 'LIN'" -run

```

Note that it is still up to the user to manually create or configure the input file for the refractive index in DDSCAT. In the above and following command lines, we use the DDSCAT style convention to name them according to their value.

A table with all the equivalent parameters and their corresponding accuracy agreement is provided in Appendix C. For all command-line examples, we exploit the fact that DDSCAT and ADDA allow for arbitrary physical units for all lengths (consistently for all input and output) to express the input wavelength and particle size in nanometers to match the convention used by IFDDA. This choice is made solely to improve readability and usability and does not affect the iterative convergence process. Moreover, the equivalence between the codes and the scale-invariance rule [52] in general indicates that the same flexibility exists for IFDDA, although some care should be exercised in interpreting the output. For instance, the resulting cross sections would not have the units of m^2 if the input unit is not nm.

4. Performance comparison: CPU benchmark

This section presents the performance benchmarks conducted across the three DDA implementations: IFDDA, ADDA, and DDSCAT. CPU benchmarks are performed for all three codes, while GPU benchmarks are

restricted to ADDA and IFDDA. The main goal is to evaluate runtime, memory footprint, and precision effect under identical physical and numerical conditions. Notably, DDSCAT allows single-precision execution on CPU, while IFDDA supports it on GPU, providing additional insight into performance–accuracy trade-offs.

To ensure fair comparisons, all benchmarks were configured to achieve machine-precision agreement across the three implementations. This requirement imposes tight constraints on the simulation setup, as not all combinations of parameters or solver options are available in every code. We therefore focus on a single, well defined case, a homogeneous cubic particle of ice with $kD = 30$ (D is the edge size), solved using the BiCGStab iterative solver and the FCD interaction and polarizability formulation. This configuration is representative of typical DDA workloads and was also used in Ref. [20].

As mentioned in Section 2, BiCGStab is chosen to avoid the two times difference in algorithmic complexity between QMR implementations in different codes. However, BiCGStab implementations also differ slightly in their stopping procedures. The original algorithm [53] can stop in the middle of an iteration, thereby saving one matrix-vector product. This feature is implemented in ADDA, but not in the PIMZ [54] version used in DDSCAT and IFDDA. Whether this happens is quasi-random and depends on the problem parameters. The following test case was explicitly tuned to avoid this issue. Other scattering problems may instead lead only to η -level of agreement and to extra computational-time difference for a single matrix-vector product. Hence, unfortunately, there is currently no ideal iterative solver for robust comparison across all three codes. However, such options exist for pair-wise comparisons (see Table C.1).

Due to floating-point rounding errors, the convergence history differs slightly between the codes and hardware backends, leading to different iteration counts for a fixed η . This degrades floating-point consistency at the η -level. To remove this issue, we explicitly tuned η for each implementation so that all CPU runs terminated after the same number of iterations (54 in this case), yielding 10–12 matching digits in the extinction cross section and efficiency (see Table C.3).

An alternative and more robust approach to enforce identical iteration counts would be to limit the maximum number of iterations

using options such as `-maxiter` (in ADDA), `-MXITER` (in DDSCAT), or `-nlim` (in IFDDA). However, while ADDA still computes physical observables when this limit is reached, both IFDDA and DDSCAT terminate execution, preventing a comparison. We therefore adopt threshold tuning as a practical compromise. However, we suspect such tuning to become impractical beyond 100 iterations.

Since our comparison focuses on runtime and solver performance rather than angular quantities, we reduced the angular sampling: in DDSCAT we set `ETASCA = 10` as recommended by the manual and used `IORTH = 1` to compute a single polarization; for ADDA this corresponds to `-ntheta 10`. The exact command lines used were:

```

$./ifdda -object cube 2387.3241463784303 -lambda
500 -epsmulti 1.723969 0.0 -ninitest 0 -nnnr
150 -tolinit 9.46237161365793d-5 -methodeit
BIGGSTAB -polarizability FG
$./adda -shape box 1 1 -size 2387.3241463784303
-lambda 500 -m 1.313 0.0 -init_field zero
-grid 150 -eps 4.024 -iter biggstab -pol fcd
-int fcd -scat dr -ntheta 10 -maxiter 54
$ddscatcli -CSHAPE RCTGLPRSM -AEFF
"1480.9777061418503 1480.9777061418503 1
'LIN'" -WAVELENGTHS "500 500 1 'LIN'" -DIEL
"diel/m1.313_0.00" -MEM_ALLOW "150 150 150"
-SHPAR "150 150 150" -TOL 9.46237161365793e-5
-CMDSOL PBCGST -CALPHA FLTRCD -ETASCA 10
-IORTH 1 -NPLANES 0 -NRFLD 0 -run

```

The memory footprint grows linearly with the number of voxels N , so it can be extrapolated straightforwardly for each code. By contrast, the reported runtimes and speedups in this chapter are problem-size dependent. As discussed in Refs. [7,55], the time per iteration does not scale exactly linearly with N (e.g., FFT introduces $N \log N$ behavior) and depends sensitively on the FFT implementation, numerical precision, and paradigm of parallelization. These effects are also hardware dependent.

A comprehensive sweep over N is beyond our present study. Consequently, the speedups reported here characterize the selected test cases and hardware, and should be interpreted or extrapolated with caution to different problem sizes.

4.1. Computational setup

The CPU architectures used for benchmarking are summarized in Table 4. They correspond respectively to a single node of a high-performance Cray cluster and to a modern laptop processor. This contrast allows us to illustrate both the performance potential and the practical limitations that can be expected on HPC hardware versus everyday desktop environments.

All cluster executions were submitted using the `SBATCH --exclusive` option of SLURM to eliminate performance variability due to shared-node contention. Each simulation was repeated three times to compute the mean and sample standard deviation, assessing run-to-run variability.

Table 4

CPU platforms used in benchmarks. Cache sizes are aggregated *per socket*. Memory bandwidths are theoretical peaks *per socket* (decimal GB/s).

	AMD EPYC 9654 (Genoa)	Intel Core Ultra 7 165H
Sockets	2	1
Cores / Threads	96 / 192 (per socket)	16 / 22 (8 E + 6 P cores)
Base / Turbo (GHz)	2.4 / 3.7	1.4 / 5
L1 / L1D	3 MiB / 3 MiB	896 KiB / 544 KiB
L2 / L3	96 MiB / 384 MiB	18 MiB / 24 MiB
Peak memory BW	460.8 GB/s (DDR5-4800)	120 GB/s (LPDDR5x-7467)
RAM (node)	768 GiB (24 × 32 GiB)	2 × 32 GiB

All three codes were compiled using the same toolchain to minimize environment-related bias. Compiler version had negligible impact: switching from GCC 11.2 to GCC 12.2 produced identical results within measurement noise. On the cluster, we used the following software stack: GCC 12.2 with OpenMP 4.5, MPICH 3.4a2, FFTW 3.3.10, and Intel oneMKL 2023.0, while on the laptop, we used GCC 11.4 with OpenMP 4.5, Open MPI 4.1.2, FFTW 3.3.8, and Intel oneMKL 2025.0. ADDA was compiled with the `-Ofast` optimization flag, IFDDA with `-O3`, and DDSCAT with `-O2`, following their respective build defaults.

The following subsections present the CPU results while the GPU benchmark is reported separately in another section. Each test reports total runtime, speedup factors, memory footprints, and precision effects.

4.2. Time performance

We tested several grids and FFT factorizations and found that grid sizes of $n_x = 150$ and $n_x = 250$ allow all codes to show their best. For all CPU benchmarks, the number of MPI processes was chosen to follow ADDA recommended practice for parallel FFTs, i.e., this number divides n_x . Figs. 1 and 2 report wall-clock times and speedups on the AMD EPYC 9654 node and the Intel Core Ultra 7 165H processor, respectively. Error bars denote the sample standard deviation and are shown only when the corresponding coefficient of variation exceeds 5%.

In this section, we report only the total wall-clock time, since it is dominated by the iterative solver phase, which is itself dominated by matrix-vector products, i.e., FFTs. By contrast, for GPU benchmarks, the solver phase is the portion offloaded to the device and, thus, its time may become comparable to that of the parts remaining on the CPU.

Consistent with the DDSCAT manual [47], MKL outperforms the legacy GPFA backend on modern CPUs, especially on Intel processors, as MKL is heavily optimized for Intel microarchitectures. As a result, DDSCAT MKL achieves roughly twice the performance of its GPFA build on the Intel laptop, while the improvement is more modest on the AMD EPYC node. This makes DDSCAT MKL competitive with ADDA FFTW and IFDDA FFTW in single-core runs on the Intel laptop for the larger grid. On the AMD node, however, ADDA FFTW and IFDDA FFTW remain significantly faster, outperforming DDSCAT MKL by factors of roughly 2.5 and 2.0 for $n_x = 150$ and $n_x = 250$, respectively. These differences are expected. Indeed, while DDSCAT performs a full 3D FFT transform in one run [56], ADDA reduces the total FFT workload by decomposing the 3D convolution into a sequence of 1D FFTs with intermediate transpositions. This 3D-to-1D strategy, introduced in Ref. [57], almost halves the number of involved 1D FFTs by explicit usage of zero-padding inherent to the convolution, sustaining high parallel efficiency and optimizing both memory usage and execution time. During the course of this work, IFDDA was also modified to adopt a similar 3D-to-1D implementation as ADDA, but optimized primarily for execution time (see Section 4.4 for further details on memory differences), resulting in performance gains of approximately 10% here. Related improvements also appear in OpenDDA [55]. Finally, the FFTW library, featuring hardware-adaptive planning, shows optimal performance over all tested hardware, in agreement with previous inter-code comparisons [20].

For multi-core CPU runs, DDSCAT exhibits essentially flat scaling, as its FFT routines, which constitute the largest fraction of the total runtime, are not parallelized with OpenMP. DDSCAT parallelization is primarily designed for orientation-averaging (outside of our scope), where much larger speedups would be observed. By contrast, both ADDA and IFDDA scale reasonably well on the AMD node, achieving speedups up to 24× (for 125 cores) and 7× for $n_x = 250$. On this architecture, ADDA FFTW remains more than 5× faster than IFDDA FFTW for $n_x = 150$ and more than 3× faster for $n_x = 250$ (for 125 cores). Moreover, we observed that variability is greater for OpenMP than for MPI. On the Intel laptop, however, IFDDA displays better runtimes and scaling for $n_x = 250$ reaching speedup up to 1.5×.

To explain these behaviors it is necessary to discuss saturation effects. As visible in Fig. 1, ADDA and IFDDA exhibits nearly linear scaling

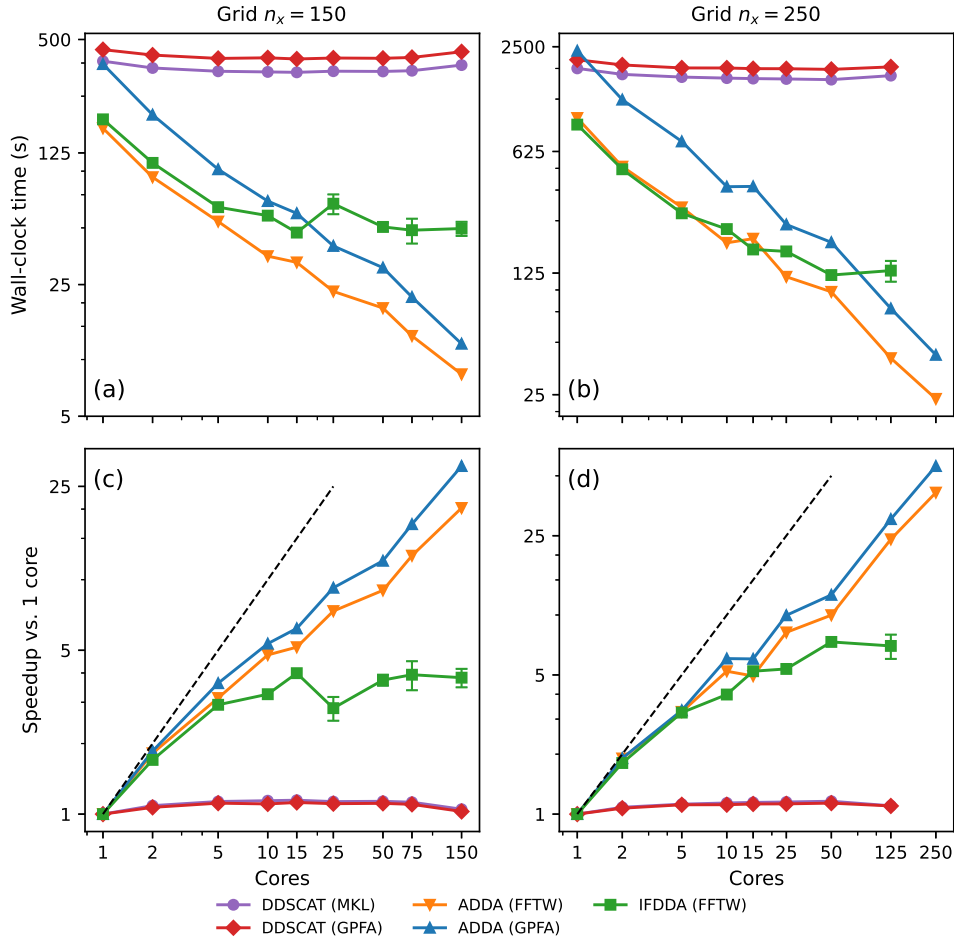


Fig. 1. Wall-clock time (top row) and speedup relative to the single-core reference (bottom row), for grid sizes $n_x = 150$ (left) and $n_x = 250$ (right) on the AMD EPYC 9654 node. Log-log scale is used, and the dashed curves indicate ideal linear scaling. Error bars, denoting the sample standard deviation, are shown when it is larger than 5% of the value.

up to 2–3 cores, after which deviations from ideal scaling appear. This saturation arises from Amdahl’s law and memory-bandwidth limits, as more threads compete for a fixed number of memory channels, communication costs progressively decrease parallel efficiency. Similar trends were reported in a recent comparative study of FFTW backends using OpenMP and MPI [58]. The different saturation points between $n_x = 150$ and $n_x = 250$ highlight the grid-dependent scaling behavior of 3D FFTs. Larger grids maintain high parallel efficiency for longer, as their computational intensity grows faster than synchronization and communication overheads.

On the Intel laptop, these effects appear much sooner. The available memory bandwidth is significantly lower and the cache hierarchy is smaller (see Table 4), causing the working set of the 3D FFT to overflow to main memory. As a result, speedups saturate well before 10 cores despite slightly better single-core performance due to higher clock frequency. This also explains why ADDA GPFA exhibits better scaling than ADDA FFTW on the AMD node, while the opposite trend is observed on the Intel laptop (although the wall time of FFTW is superior in all cases). Specifically, the limited memory bandwidth and cache capacity make the adaptive nature of FFTW more beneficial during the parallel execution on the Intel laptop. This highlights the strong dependence of FFT performance on both the underlying hardware architecture and the algorithmic optimization strategy.

Finally, note that the wall-clock time of ADDA FFTW for $n_x = 250$ at 250 cores was obtained using two AMD nodes, as a single EPYC 9654 node provides only 192 cores. This was done to illustrate that ADDA MPI

implementation continues to scale across nodes. The runtime can decrease further when using more nodes with fewer cores per node, indicating that latency effects remain a limiting factor. A more systematic study, including a saturation predictor based on grid size and hardware characteristics, would be valuable for future work.

To conclude, users should prefer FFTW or MKL over GPFA backends and, for large-scale problems, use the MPI execution of ADDA.

4.3. Precision effects

A practical advantage of DDSCAT is that it supports single precision, which reduces the array footprints by a factor of two. For $n_x = 150$ case, single precision also reduces the per-iteration time since the iteration counts matched the double precision one.

Switching from double to single precision yielded a time reduction of approximately 30% with GPFA and 60% with MKL on the laptop, and 10% and 40% respectively on the cluster (data not shown). Moreover, accuracy remained good, as results match the η -level digits.

Single precision can therefore be faster than IFDDA and ADDA on a single core of the laptop for this grid (115 seconds with MKL). Overall, single precision in DDSCAT offers substantial savings in memory (factor of two) and runtime (up to 60% with MKL) with negligible effect on accuracy for this case. However, the latter is not guaranteed, especially with an increasing number of iterations. Thus, more sophisticated approaches to use the double precision for only a subset of the variables have been proposed [20]. Users should therefore be

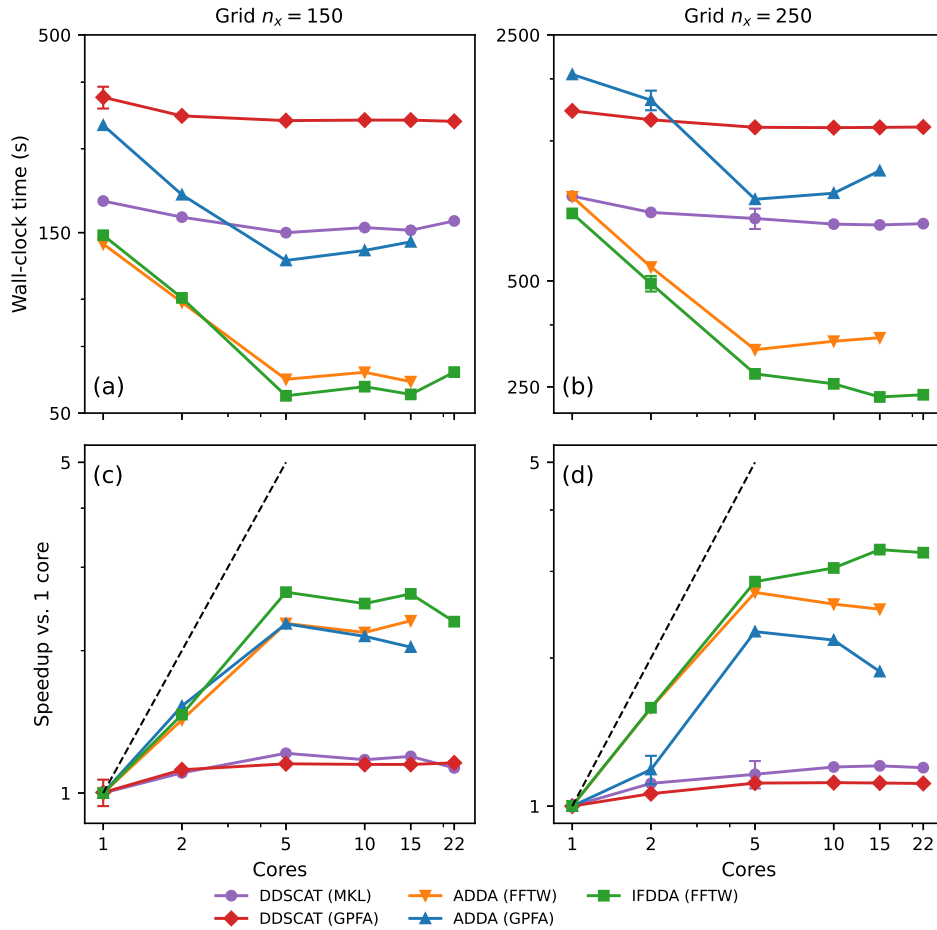


Fig. 2. Wall-clock time (top row) and speedup relative to the single-core reference (bottom row), for grid sizes $n_x = 150$ (left) and $n_x = 250$ (right) on the Intel Core Ultra 7 165H processor. Log-log scale is used, and the dashed curves indicate ideal linear scaling. Error bars, denoting the sample standard deviation, are shown when it is larger than 5% of the value.

Table 5

Node-level physical memory [GiB] for grid size $n_x = 150$. Totals are RSS per node; for MPI (ADDA), computed as the average of RSS memory from each rank, times the number of ranks (from `sacct`).

Method	1	25	50	75	150
ADDA	2.5	3.0	4.2	5.4	12.5
IFDDA	6.5	6.5	6.5	6.5	6.6
DDSCAT	6.3	6.3	6.3	6.3	6.3

cautious when trading performance gains against potentially reduced accuracy.

4.4. Memory performance

Table 5 reports memory footprints using binary prefixes as recommended by the International Electrotechnical Commission (IEC): GiB denotes gibibytes ($1 \text{ GiB} = 2^{30} \text{ B}$). The only exception is memory bandwidth in Table 4, which uses decimal GB. We report physical memory only, i.e., the resident set size (RSS). In our experience the virtual memory size (VM/VSZ) is not a useful capacity indicator for these workloads: VM is always \geq RSS because it includes the full address space (reserved but untouched pages) and the complete size of shared libraries mapped into the process; by contrast, RSS counts only the pages that are actually resident in RAM.

For codes that use OpenMP (IFDDA and DDSCAT), threads share a single process and therefore one copy of the large working arrays; RSS

remains essentially flat as the thread count increases (see Table 5). For an MPI program, however, an increase is observed with the number of ranks: although ADDA uses about 2.5 less memory at one rank, the aggregate memory can exceed the OpenMP codes at higher core counts. This behavior is primarily due to MPI. Compared to shared memory, MPI launches P separate processes. Each rank keeps its own copy of certain auxiliary MPI variables, leading to a noticeable memory increase with the number of ranks. This MPI-related overhead is clearly visible in moderate-size runs but becomes negligible for larger grids, where the properly distributed working arrays dominate. A detailed analysis of the MPI memory scaling is provided in Appendix D.

The lower initial memory usage in ADDA originates from how the interaction matrix is stored in Fourier space. Indeed, in FFT-accelerated DDA formulations, only the six independent components of the dyadic Green's tensor are required due to symmetry; these components are transformed to Fourier space and stored on an extended zero-padded domain of twice the size in each spatial direction. In the classical approach [56], all six components are stored over the full circulant domain. ADDA reduces this cost by exploiting mirror symmetries in two directions, while not using symmetry along the third axis to facilitate MPI parallelization, thereby reducing the storage requirement by a factor of four compared to the classical formulation. Further details are provided in Ref. [55], which also shows how OpenDDA exploits symmetry in all three spatial directions to reach the theoretical minimum storage.

On the full node (192 cores, 714 GiB RAM) we were able to run ADDA up to $n_x = 960$. Moreover, using fewer cores per node and moving to multiple nodes allow even larger grids. For example, a recent large-scale

Table 6

GPU platforms used in the benchmarks. Peak floating-point throughput and memory bandwidth values are vendor-reported, while cache sizes are taken from third-party hardware databases. FP64 throughput on Ada GPUs is 1/64 of FP32 peak. SM denotes “Streaming Multiprocessors”.

	A100 (80 GiB SXM4)	H200 (141 GiB SXM)	RTX 6000 Ada	RTX 2000 Ada
FP64 (TFLOP/s)	9.7	34	1.4	0.2
FP32 (TFLOP/s)	19.5	67	91.1	12
Mem. BW (GB/s)	2,039	4,890	960	256
L1 / SM (KiB)	192	256	128	128
L2 (MiB)	40	50	96	12
SM	108	132	142	22

ADDA run was reported in Ref. [59], reaching $n_x = 3168$ for a cube with size parameter $kD = 2000$ and refractive index $m = 1.02$, using 44 nodes (3168 CPU cores) with a peak memory footprint of 27.6 TiB and an MPI overhead of only 0.4 TiB (i.e., 1.2%).

5. Performance comparison: GPU benchmark

The GPU modes of both IFDDA (v1.0.27) and ADDA (v.1.5.0-alpha3, commit b03d648) are still under active development. Beyond version control, the ongoing development activity is also evident from the public ADDA issue tracker and recent benchmark discussions on the official discussion group. It is also worth noting the recent `fork`¹ by one of the authors, which replaces the original OCL code with a CUDA implementation (not used in this work). The timings and memory footprints reported here therefore characterize the current implementations and test configurations rather than hard limits. Thus, future releases will significantly change these results.

5.1. Computational setup

The GPU benchmarks were performed on four NVIDIA accelerator devices spanning both cluster- and workstation-architectures. These include the A100 (Ampere), H200 (Hopper), RTX 6000 Ada, and RTX 2000 Ada. Their main architectural characteristics are summarized in Table 6.

Recent NVIDIA GPUs integrate tensor cores units specialized in matrix-multiply-accumulate (MMA) that accelerate general matrix multiplication (GEMM) operations. They deliver very high throughput in mixed/low precision but also in standard precision (FP64 tensor core) on recent A100/H100 GPUs. However, standard FFT libraries (FFTW, cuFFT, cFFFT) do not use them by default. Research prototypes that reformulate FFTs to be tensor cores-friendly [60] reported 1.1–3× speedups over cuFFT on NVIDIA A100, but we have not considered them.

Each GPU was paired with the host processor available on the corresponding system: dual-socket AMD EPYC 7543 nodes for the A100, dual-socket Intel Xeon 8558 nodes for the H200, an Intel Core Ultra 7 165H laptop for the RTX 2000 Ada, and an Intel Xeon w5-3525 workstation for the RTX 6000 Ada. Ideally, all GPUs would be benchmarked under the same host CPU to isolate accelerator performance. However, this was not possible due to hardware availability. Nevertheless, this does not affect the validity of our comparison, as for each run we separately record both the GPU-resident solver time and the total wall-clock time. This allows us to separate the contribution of the GPU from that of the host processor and to assess the impact of both GPU and CPU hardware on the overall runtime.

The same numerical configurations as in the CPU section were used to ensure direct cross-platform comparability. While the GPU double-precision runs also terminated in 54 iterations, the GPU single-precision runs with IFDDA converged in 53 iterations instead of 54. This difference arises because the step between successive residuals becomes too large to be tuned in the same manner as in double precision. Nevertheless,

this does not affect the present comparison, since single-precision accuracy is inherently limited to 7–8 significant digits, and a difference of one iteration has a negligible impact on the measured execution times reported in Section 5.

As in the CPU section, the reported timings correspond to the mean of three identical runs. Furthermore, all codes were compiled using the same library versions and optimization flags, with CUDA 13.0 for IFDDA and OpenCL 3.0 for ADDA.

5.2. Time performance

For DDA solvers, the dominant cost per iteration is the FFT stage. Despite many optimization efforts over the years [55,61,62], FFTs remain the bottleneck of the iterative solver, with the remainder of the time spent mostly in BLAS (Basic Linear Algebra Subprograms) operations. In IFDDA, the entire solver loop is offloaded to the GPU, which minimizes host–device communication. By contrast, ADDA offloads only the FFTs (matrix-vector products), while the linear-algebra operations remain on the CPU, introducing additional transfers. ADDA also provides an experimental OCL_BLAS mode, which offloads the BLAS operations to the GPU via the `clBLAS` library, thereby both accelerating the BLAS operations and removing the corresponding transfers. As of the current ADDA release, this mode supports only the BiCG solver. Since our focus in this study is on cross-code comparisons and floating-point-consistent reproducibility, we did not include it in the main benchmark. However, we report its performance in Appendix E, where it achieves multiple speedups, demonstrating the benefit of a fully GPU-resident implementation.

We report in Fig. 3 solver-only time (iterative phase) and total wall-time (end-to-end) for grid sizes $n_x = 150$ and $n_x = 250$ on four NVIDIA GPUs (see Table 6), except on the NVIDIA RTX 2000 Ada, where the IFDDA computation fails due to an out-of-memory error, and two CPU settings (1 and 10 cores). IFDDA is run in both double (DP) and single precision (SP), while ADDA is reported only for 1 core because its default GPU mode is currently mutually exclusive with the MPI. For ADDA, we also show the time for the FFTs, as a part of the solver. In contrast to Fig. 2, the sample standard deviation is not shown here for clarity. The corresponding coefficients of variation remain below 5% in all cases (and typically around 1%), except for the solver mean runtime of the IFDDA single-precision run for $n_x = 150$ on the NVIDIA RTX 6000 Ada, where it reaches 6%.

The overall GPU behavior of IFDDA DP can be interpreted within the roofline model [63], which relates arithmetic intensity (FLOP/B) to the machine balance (peak FLOP/s per peak BW). The arithmetic intensity of 3D FFTs is low and depends on problem size; accordingly, DDA iterations lie predominantly in the memory-bound regime [64], with possible transitions into the compute-bound or latency-bound regimes [63]. For a fully device-resident solver such as IFDDA, one therefore expects the solver time to scale approximately with the inverse memory bandwidth. The measured runtimes follow this trend qualitatively, as GPUs with higher bandwidth deliver shorter solver times, H200 > A100 > RTX 6000 Ada > RTX 2000 Ada. However, quantitative description does not seem feasible, as synchronization overheads, cache and

¹ <https://github.com/michelgross34/adda/tree/cuda-port>.

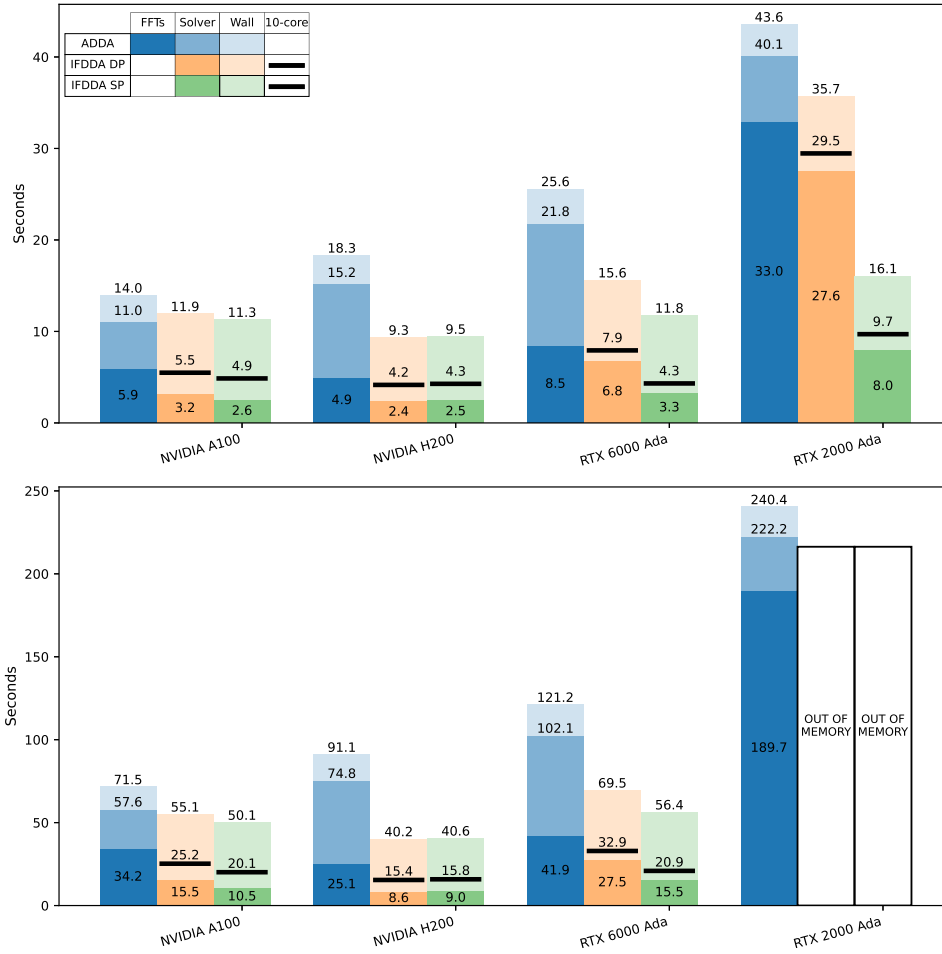


Fig. 3. GPU timings for DDA codes. Bars are grouped by code and stacked to show FFT time, solver time, and 1-core wall-time. Horizontal black lines indicate 10-core wall-times for IFDDA (DP/SP). Grid size $n_x = 150$ (top row); $n_x = 250$ (bottom row).

memory-path, and FP64/FP32 ratio vary significantly between devices. These factors introduce an additive latency term, causing a loss of performance. Consequently, the H200 does not reach the full improvement predicted by its bandwidth advantage over the A100 or workstation GPUs.

For ADDA, the trend resembles that of IFDDA DP except for the H200, where the solver and total runtimes are higher than on the A100. Since only the FFTs are executed on the GPU, each solver iteration requires multiple host-device transfers and synchronizations, making the performance sensitive to system-level latency. Evidence for this interpretation is provided by the OCL_BLAS mode (see Fig. E.1), where offloading the full solver to the GPU reverses the trend back to the expected one. This indicates that the performance inversion observed for ADDA arises from latency costs in the hybrid CPU-GPU workflow, rather than from GPU-side throughput. We emphasize that it cannot be explained by any hardware specification, as both the memory hierarchy and CPUs associated with the H200 system are nominally superior (see Table 6). Systematic investigation of this fact is left for the future.

In the single-core case, IFDDA DP systematically outperforms ADDA, even when comparing only the GPU-offloaded part, which would in principle be expected to favor ADDA since only FFTs are performed. This behavior can arise from several factors. First, the previously discussed CPU-GPU latency and synchronization overheads. Second, on NVIDIA hardware, cuFFT (used in IFDDA) is expected to outperform c1FFT [65]. Finally, ADDA relies on data transpositions associated with FFT execution plans, using empirically chosen block sizes, which may be suboptimal for current GPU architectures. Taken together, these effects make it dif-

ficult to draw definitive conclusions or ratios from absolute runtimes alone. Similar considerations apply to total wall-clock times, where differences in CPU performance also contribute.

Nevertheless, by considering the relative speedup achieved by the OCL_BLAS mode on the solver part for the BiCG solver (see Appendix E) and extrapolating this improvement to the BiCGstab case, one may expect ADDA to achieve performance comparable to that of IFDDA when the entire solver loop is offloaded to the GPU. This interpretation is further supported by independent implementations. A MATLAB implementation, MPDDA [62], which offloads the complete iterative solver to the GPU, reports better runtimes than the default ADDA OpenCL mode, although this difference may partly arise from the use of Apple’s OpenCL FFT backend rather than the standard c1FFT library. More recently, a Python implementation, CPDDA [66], based on a similar strategy and relying on CuPy, reports performance surpassing MPDDA at smaller lattice spacings. Altogether, these results support the benefit of offloading the complete iterative solver to the GPU.

Finally, IFDDA can further reduce wall-clock time by increasing the number of CPU threads, even in GPU mode, since its non-FFT components benefit from OpenMP parallelism. For instance, using 10 CPU cores yields a 2x reduction in the total runtime on the A100.

5.3. Precision effects

For bandwidth-bound problems such as FFTs, switching from double to single precision ideally halves the memory traffic (8 to 4 bytes), potentially yielding a 2x speedup. For compute-bound regimes, the

achievable ratio is instead limited by the device FP64:FP32 throughput, and for latency-bound regimes essentially no speedup is expected, since the dominant cost lies in waiting for data rather than processing it.

The measured ratios therefore differ across GPUs depending on their architectural balance between bandwidth, FP64 capability, and overheads. On the A100, FP64 performance is relatively strong (and half of FP32) and memory bandwidth is high. For moderate grid sizes ($n_x = 150$ – 250) the kernels do not fully saturate the memory subsystem, and the observed speedups of 1.4–1.5 \times reflect a mixture of bandwidth and latency effects. For larger grids (not shown here), we observed the ratio of 1.8.

On the RTX 6000 Ada, the FP64 throughput is extremely low (of 1/64 FP32), and memory bandwidth is modest. This makes the latency delays less important, which explains larger DP/SP ratio of 1.8–2.1 still in the bandwidth-regime but closer to a compute-bound regime for $n_x = 150$. This trend is even more pronounced on the RTX 2000 Ada, whose FP64 capability is still lower. There, the DP solver is clearly more compute-bound, manifested by the DP/SP ratio of 3.5, between bandwidth- and computations-based values.

The H200 presents a different behaviour. Despite sharing the same FP64:FP32 peak ratio (1/2) as the A100 and offering both higher FP64 throughput and memory bandwidth, the IFDDA DP and SP solver times on the H200 are nearly identical or even higher for SP. This indicates a latency-bound regime. Unlike the ADDA case discussed in [Appendix E](#), this behavior cannot originate from CPU–GPU transfers, since IFDDA is fully device-resident and the comparison is purely between double and single precision. A plausible explanation is that, on the H200, the double-precision solver already operates near a practical performance limit set not by bandwidth but by internal latency overheads on the FFT kernels side.

Overall, the results show that, on workstation GPUs with limited FP64 throughput and small memory bandwidth, the choice of precision strongly affects both solver and total time, and significant gains can be obtained by using single precision when iterative convergence permits. On high-bandwidth cluster GPUs, however, the benefits of switching to single precision may decrease or even vanish if the solver approaches an latency-limited regime.

5.4. GPU vs CPU

This section compares the best GPU runtimes obtained on the NVIDIA A100 with the best CPU results measured on the AMD EPYC node. Because the available core counts differ between the two grid sizes, the CPU reference is defined as the fastest configuration yielding the smallest wall-clock time in the CPU benchmarks: 75 cores for $n_x = 150$ in the case of IFDDA, 150 cores for $n_x = 150$ in the case of ADDA, and 50 and 125 cores respectively for $n_x = 250$, in order to avoid taking advantage of ADDA multi-node implementation.

For $n_x = 150$, the best CPU runtime for ADDA is 8.3 s, whereas the A100 requires 14 s. Thus, ADDA in OpenCL mode is about 1.7 \times slower than the corresponding multi-core CPU execution. By contrast, IFDDA benefits from GPU acceleration, as the double-precision GPU time is 4.1 \times faster, and the single-precision time yields a 4.3 \times speedup.

At $n_x = 250$, ADDA becomes 1.8 \times slower on the GPU than on the 125-core CPU run, whereas IFDDA DP and IFDDA SP achieve 2.2 \times and 2.4 \times speedups, respectively.

For the ADDA OCL_BLAS mode on the A100, we observe a 1.3 \times speedup relative to the 125-core CPU run (computed for BiCG solver). However, for very large grids (e.g., $n_x = 400$), a pure CPU MPI run becomes faster, even though the GPU-based solver itself is 15–20% faster. In this regime, the performance is limited by the initialization stage, which is not parallelized in the OCL mode and therefore dominates the overall runtime. Introducing MPI parallelism in the OCL initialization stage, or ideally offloading it to the GPU, would remove this limitation,

Table 7

Peak GPU device memory (GiB) for two problem sizes (i.e., $n_x = 150$ and 250).

Method	$n_x = 150$	$n_x = 250$
ADDA	1.5	6.5
IFDDA SP	2.9	13.4
IFDDA DP	5.8	28.5

at the expense of a higher memory footprint and a reduced maximum attainable grid size.

Finally, SLURM energy accounting on the AMD node shows that a typical large MPI run of ADDA (75 cores) consumes about 700 kJ, whereas the single-GPU A100 run requires only 36 kJ. Even when the GPU is slower in wall time, its energy efficiency is therefore much higher.

Overall, GPUs provide substantial benefits for fully GPU-resident solvers (IFDDA and ADDA OCL_BLAS), especially when single precision is usable. For ADDA, GPU mode is competitive for small and medium grids, but CPU-based MPI remains necessary for the largest cases. Memory capacity is the critical factor for the latter: even modern accelerators offer far less RAM than multi-socket CPU nodes, constraining feasible grid sizes.

5.5. Memory performance

From [Table 7](#), IFDDA has a larger footprint than ADDA, as IFDDA SP and IFDDA DP require about 2 and 4 \times the RAM of ADDA, respectively. This ratio is higher than what we observe on CPU memory ([Section 4.4](#) reports 2.5 \times for a single rank) due to ADDA offloading only the FFT data on GPU compared to IFDDA.

Precision has, in trend, the expected first-order effect, i.e., switching from DP to SP reduces per-element storage by a factor of 2. Thus, when the iterative convergence permits single precision, the 2 \times memory reduction can be decisive for feasibility on GPU memory-limited hardware.

Practically, the grid $n_x = 250$ IFDDA DP peak (28.5 GiB) requires a 40 GiB GPU card (e.g., RTX 6000 Ada), whereas IFDDA SP (13.4 GiB) fits on 16 GiB workstations, and ADDA (6.5 GiB) remains broadly portable. On the A100 80 GiB, our measurements indicate a feasible maximum around $n_x \approx 350$ for IFDDA DP, $n_x \approx 450$ for IFDDA SP, and $n_x \approx 530$ for ADDA. With ADDA OCL_BLAS mode, the limit is slightly reduced to $n_x \approx 510$ (for BiCG). Note that ADDA provides the `-opt mem` option, which enforces a memory-conservative execution mode by processing the computational domain one layer at a time along the x direction. This reduces the GPU memory footprint and can therefore slightly increase the maximum feasible grid size. For example, for $n_x = 530$, the maximum GPU memory usage decreases from 78.2 GiB to 71.7 GiB. In addition to reducing memory consumption, this option may also improve runtime stability, as operating close to the GPU memory limit can lead to inefficiencies and performance degradation. The overall performance impact of this trade-off warrants further investigation.

6. Conclusion

We presented a unified benchmarking and interoperability software-assisted methodology for the three most widely used implementations of DDA (DDSCAT, ADDA, and IFDDA). By aligning all physical, numerical, and solver parameters, we demonstrated that the three codes can reach machine-precision agreement on identical test cases. Although applied here to only three solvers, the software package itself is fully general, as any DDA implementation can be integrated into the same cross-verification and benchmarking methodology. This establishes that modern DDA solvers are both numerically mature and internally consistent, and it enables performance comparisons that are decoupled from accuracy differences.

Our CPU benchmarks show that FFT strategy is the dominant factor in performance, as MKL and FFTW substantially outperform legacy GPPA algorithms. Among the tested codes, ADDA achieves the best scalability thanks to its 3D–1D FFT decomposition and optimized MPI implementation.

GPU benchmarks confirm that offloading the entire solver phase to the device, as implemented in IFDDA (and in ADDA experimental OCL_BLAS mode), provides the strongest acceleration. By contrast, ADDA default OpenCL mode offloads only the FFTs while keeping the solver loop on the CPU, making its performance sensitive to CPU–GPU-transfer latency. GPU performance also depends primarily on architectural factors such as memory bandwidth for cluster-type cards, while peak FLOP/s can have an impact on workstation cards.

Overall, this study demonstrates that numerical interoperability between DDA codes is not only achievable but also highly informative. Once physical and numerical parameters are harmonized, the remaining performance differences reveal genuine architectural and implementation effects. Our methodology thus enables reproducible, cross-platform verification and provides fair comparisons between independent DDA implementations.

Finally, we note that during the preparation and publication of this work, newer versions of IFDDA were released, introducing modifications that affect certain aspects of the performance analysis, in particular the memory footprint. While these updates do not impact the overall conclusions of this study, they may lead to different quantitative results in future benchmarks.

CRedit authorship contribution statement

Clément Argentin: Writing – original draft, Visualization, Validation, Software, Investigation, Data curation, Conceptualization; **Patrick C. Chaumet:** Writing – review & editing, Validation, Software, Investigation; **Michel Gross:** Writing – review & editing, Software; **Maxim A. Yurkin:** Writing – review & editing, Validation, Supervision, Software, Investigation, Funding acquisition, Conceptualization.

Data and code availability

The exact version of the main cross-comparison code `dda-bench` used in this paper is archived on Zenodo and available at <https://doi.org/10.5281/zenodo.18836855>. The Python wrapper `ddscatcli` is available at <https://doi.org/10.5281/zenodo.18619006>.

All data and scripts used to produce the benchmark results and figures in this study are publicly available at Zenodo: <https://doi.org/10.5281/zenodo.18847515>. To reproduce the software setup, patches must be applied to the exact base versions documented in the repository README.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the Normandy Region (project RAD-DAERO). Part of the simulations were performed using computing resources of CRIANN (Normandy, France). The authors also thank Bruce Draine and Piotr Flatau for their comments on the DDSCAT code.

Appendix A. Minor code modifications

To ensure numerical consistency and fair performance comparison across DDSCAT, ADDA, and IFDDA, several minor edits were made to

DDSCAT (version 7.3.4). These modifications, listed below, improve precision handling and solver consistency but do not affect physical results. Complete patch files are available in the Zenodo repository (see Data availability, folder `patches/`). We have also submitted these files to the DDSCAT authors for potential inclusion in the next release. Similar minor changes to IFDDA (during the course of this study) are not listed here, since they were already included in the latest release.

A.1. Precision of numbers in output

In order to test floating-point consistency in double precision, output quantities must be printed with 16 significant digits. This required minor modifications to the output routines of ADDA and DDSCAT.

A.2. Modifications of DDSCAT

Precision handling in LDR coefficients. The literal precision of the Lattice Dispersion Relation (LDR) coefficients was increased to ensure correct rounding in double precision (see Appendix B). This eliminates rounding differences observed at the 7–8 significant-digit level between DDSCAT and the other implementations.

Precision-consistent in trigonometric routines. In the cosine–sine integral subroutine, numerical constants were extended to full precision, tolerance thresholds were made precision-dependent, and complex variables were explicitly typed to prevent down-casting. This change does not affect single-precision results, but avoids loss of precision in double ones.

Solver restart policy. BiCGStab restarts were disabled to align the solver iteration counts and timing with those of ADDA and IFDDA, which do not use restarts.

Preconditioner and stopping criterion. By default, DDSCAT applies a left Jacobi preconditioner and a true-residual stopping criterion, resulting in an additional matrix–vector product per iteration. For benchmarking equivalence, we disabled preconditioning and used a relative-residual criterion, consistent with the other DDA codes. This change avoids the extra matrix–vector operation and ensures the same linear operator is solved in all cases. Note also that the resulting behavior is the same as that already used for other iterative solvers in DDSCAT.

Appendix B. Evaluation of the LDR coefficients

The following summarizes the steps required to evaluate the lattice dispersion relation (LDR) coefficients with arbitrary numerical precision. We use the equivalent representations of Eqs. (3.19) and (A.4) from Ref. [67],

$$g(\beta, x) \stackrel{\text{def}}{=} \sum_{n=-\infty}^{\infty} e^{-x(\beta+n)^2} = e^{-x\beta^2} \theta_3(ix\beta, e^{-x}) = \sqrt{\frac{\pi}{x}} \theta_3\left(\pi\beta, e^{-\pi^2/x}\right), \quad (\text{B.1})$$

where the elliptic theta functions ensure rapid convergence for $x \geq \pi$ and $x \leq \pi$, respectively [68]. In particular, they imply

$$g(0, x) = \sqrt{\frac{t}{\pi}} g(0, t), \quad t \stackrel{\text{def}}{=} \frac{\pi^2}{x}. \quad (\text{B.2})$$

The coefficient c_1 is defined by Eq. (4.2) of Ref. [67]:

$$c_1 = \frac{2}{3} \int_0^{\infty} dx \left[g^3(0, x) - \left(\frac{\pi}{x}\right)^{3/2} - 1 \right]. \quad (\text{B.3})$$

This integral is split at $x = \pi$, and the change of variables $x \rightarrow t = \pi^2/x$ is applied to the interval $(0, \pi)$. Using Eq. (B.2), this yields

$$c_1 = \frac{2}{3} \int_{\pi}^{\infty} dx \left(1 + \sqrt{\frac{\pi}{x}} \right) [g^3(0, x) - 1] - 2\pi. \quad (\text{B.4})$$

The cubic power of g can be expressed as a lattice sum over all nonzero integer vectors (Eq. (A.3) of Ref. [67]),

$$g^3(0, x) - 1 = \sum_{\mathbf{n} \neq 0} e^{-x|\mathbf{n}|^2}. \quad (\text{B.5})$$

Thus, defining $u = \pi|\mathbf{n}|^2$, one can obtain

$$c_1 = \pi \left\{ \frac{2}{3} \sum_{\mathbf{n} \neq 0} \left[\frac{e^{-u}}{u} + \sqrt{\frac{\pi}{u}} \operatorname{erfc}(\sqrt{u}) \right] - 2 \right\}. \quad (\text{B.6})$$

This can also be written in terms of incomplete gamma functions $\Gamma(1, u)$ and $\Gamma(1/2, u)$, making the whole approach of evaluating the lattice sums similar to that in Ref. [69].

The second coefficient is defined by Eq. (4.3) of Ref. [67]:

$$c_2 = \frac{1}{2} \int_0^\infty dx g(0, x) \frac{\partial g(0, x)}{\partial x} \frac{\partial^2 g(\beta, x)}{\partial \beta^2} \Big|_{\beta=0}. \quad (\text{B.7})$$

Derivatives are evaluated using Eqs. (B.1) and (B.2), yielding

$$\frac{\partial^2 g(\beta, x)}{\partial \beta^2} = -2x \left[g(\beta, x) + 2x \frac{\partial g(\beta, x)}{\partial x} \right], \quad (\text{B.8})$$

$$2x \frac{\partial g(0, x)}{\partial x} = -\sqrt{\frac{1}{\pi}} \left[g(0, t) + 2t \frac{\partial g(0, t)}{\partial t} \right]. \quad (\text{B.9})$$

Using these identities, separating the integral at $x = \pi$, and transforming g back into the sum, Eq. (B.7) can be rewritten as

$$c_2 = \int_\pi^\infty dx \left(1 + \sqrt{\frac{\pi}{x}} \right) \sum_{\mathbf{n} \neq 0} e^{-x|\mathbf{n}|^2} \left[x n_x^2 - 2x^2 n_y^2 n_z^2 \right]. \quad (\text{B.10})$$

The integration of each term in the sum leads to incomplete gamma functions $\Gamma(3/2, u)$, $\Gamma(2, u)$, $\Gamma(5/2, u)$, and $\Gamma(3, u)$, which admit closed-form expressions in terms of $\exp(-u)$ and $\operatorname{erfc}(\sqrt{u})$ [70]. Due to the cubic symmetry, n_x^2 can be replaced by $|\mathbf{n}|^2/3$, while $n_y^2 n_z^2$ by $|\mathbf{n}|^4 v/9$, where

$$v \stackrel{\text{def}}{=} \frac{3(n_x^2 n_y^2 + n_y^2 n_z^2 + n_z^2 n_x^2)}{|\mathbf{n}|^4}, \quad 0 \leq v \leq 1. \quad (\text{B.11})$$

Finally, one obtains

$$c_2 = \frac{\pi}{3} \sum_{\mathbf{n} \neq 0} \left\{ \frac{e^{-u}}{u} \left[1 + 2u - \frac{v}{3} (4 + 7u + 4u^2) \right] + \left(\frac{1-v}{2} \right) \sqrt{\frac{\pi}{u}} \operatorname{erfc}(\sqrt{u}) \right\}, \quad (\text{B.12})$$

with $u = \pi|\mathbf{n}|^2$.

The third coefficient satisfies the exact identity

$$c_1 - 4c_2 + 2c_3 = -\frac{4}{3} \int_0^\infty dx \frac{\partial}{\partial x} \left\{ x \left[g^3(0, x) - \left(\frac{\pi}{x} \right)^{3/2} - 1 \right] \right\} = 0, \quad (\text{B.13})$$

which follows directly from their definitions in Ref. [67] and the fact that the function under the derivative vanishes at both ends of the integration domain. Therefore, c_3 does not need to be evaluated independently. This identity has also been indirectly proven in Appendix A2 of Ref. [71].

When evaluated with sufficient precision to ensure accuracy up to quad (128-bit) arithmetic (using truncation at $|\mathbf{n}| \leq 5$ for c_1 and $|\mathbf{n}| \leq 5.2$ for c_2), the final LDR coefficients (in CGS units) are

$$b_1 = -\frac{c_1}{\pi} = 1.8915316529870796511106114030718259, \quad (\text{B.14})$$

$$b_2 = -\frac{c_2}{\pi} = -0.16484691508771947306079362778185226, \quad (\text{B.15})$$

$$b_3 = \frac{3c_2 + c_3}{\pi} = 1.7700004019321371908592738404451742. \quad (\text{B.16})$$

Importantly, we use the sign convention employed in the ADDA code [16], opposite to the original paper [67], to facilitate comparison with other DDA formulations [25]. When expressed in SI units, the coefficients b_1 and b_3 should be additionally divided by 4π . These values have been used in ADDA since 2022 (commit 887a63c) and they satisfy

$$b_1 - 10b_2 - 2b_3 = 0 \quad (\text{B.17})$$

due to Eq. (B.13).

As a final note, the above approach may also be used to calculate LDR coefficients for rectangular cuboid lattices [29,72].

Appendix C. Equivalence of parameters

Performing equivalent simulations with ADDA, IFDDA, and DDSCAT requires matching the command-line parameters level by level, following the accuracy hierarchy introduced in Table 2. In practice, one starts from the executables `adda` and `ifdda`, and the Python wrapper `ddscatcli` [51], then progressively enforces the equivalences listed in Table C.1 ignoring all lines that are not relevant (not used by any of the codes). The number of matching digits is then determined by the first non-matching line. If all lines in the table are matched, the machine-precision level is reached (10–15 digits).

Table C.1 is designed to show the correspondence between parameters, rather than to provide a list of specific numerical values that would enforce strict equivalence across codes, since achieving this requires additional care discussed below. We also provide a set of cross-code equivalent command-line examples in the Python package [73] (in the file `DDA_commands`) as concrete working examples. Note that even with the `ddscatcli` wrapper, the user must still create or modify the input file for the refractive index if it is not available in the DDSCAT examples folder.

A few conventions are essential for interpreting Table C.1. First, particle size and shape definitions differ across codes. In particular, DDSCAT defines the target size through an equivalent radius a_{eff} . For a given volume V , this corresponds to

$$a_{\text{eff}} = \left(\frac{3V}{4\pi} \right)^{1/3}. \quad (\text{C.1})$$

For shapes affected by volume correction, a direct one-to-one mapping of the nominal radius or diameter does not preserve the exact physical volume across codes. Consequently, the particle size must be adjusted in either IFDDA or DDSCAT to ensure volume equivalence (ADDA can use `-no_vol_cor` to switch). For inhomogeneous materials (e.g., coated spheres), arbitrary shapes in IFDDA (which are treated as box-embedded), and for the diagonally anisotropic materials (IFDDA requires CM or RR polarizability), the maximum achievable agreement is limited to the η -level, even when all other parameters are matched. Indeed, in these cases, the linear systems cannot be made strictly equivalent across codes (see Section 3.2). The bisphere is treated as an inhomogeneous material in both IFDDA and DDSCAT and therefore requires the corresponding material specification. In this work, identical dielectric functions are assigned to all components to represent a homogeneous bisphere and achieve full digit agreement between the codes. For coated spheres, DDSCAT further requires three dielectric components when using `-SHPAR 1`, following the anisotropic-shell configuration. Note that for arbitrary shapes, DDSCAT and ADDA share a common geometry file format (original to DDSCAT, based on integer voxel positions), enabling direct reuse between the two codes. IFDDA, by contrast, employs a distinct format based on real voxel coordinates. Additionally, in contrast to DDSCAT and ADDA, IFDDA require explicit grid size in this file, in addition to the corresponding input parameter used for memory allocation for both IFDDA and DDSCAT (`-nnnr` and `-MEM_ALLOW`). Moreover, this grid parameter corresponds to the maximum particle extent in IFDDA, whereas ADDA expects one along the x direction (i.e. n_x). Finally, as discussed in Section 3.4, the physical units used for wavelength and particle size are arbitrary as long as the relevant ratios are preserved and the outputs are rescaled consistently.

In addition, the incident propagation direction and the target orientation are treated differently between the codes (although the same unit of degrees is assumed for input angles). Both DDSCAT and ADDA build the voxel grid in a target-fixed reference frame (Target Frame, TF), while the incident beam is described in the laboratory frame (Lab Frame, LF).

Table C.1
Command-line equivalences between IFDDA, ADDA, and DDSCAT, organized by accuracy-critical levels.

Level	Parameter	IFDDA	ADDA	DDSCAT	Digits
Shape	Sphere	-object sphere <R>	-shape sphere -size <D>	-CSHAPE ELLIPSOID -SHPAR "<Nx> <Nx> <Nx>" -AEFF "<aeff> <aeff> 1 'LIN'"	0
	Cube	-object cube <D>	-shape box 1 1 -size <D>	-CSHAPE RCTGLPRSM -SHPAR "<Nx> <Nx> <Nx>" -AEFF "..."	
	Cuboid	-object cuboid1 <Dx> <Dy> <Dz>	-shape box <Dy/Dx> <Dz/Dx> -size <Dx>	-CSHAPE RCTGLPRSM -SHPAR "<Nx> <Ny> <Nz>" -AEFF "..."	
	Cylinder	-object cylinder <R> <h>	-shape cylinder <h/D> -size <D>	-CHSAPE CYLNDRPBC -SHPAR "<Nx> <Nx> 3 0 0 1" -AEFF "..."	
	Ellipsoid	-object ellipsoid <Rx> <Ry> <Rz>	-shape ellipsoid <y/x> <z/x> -size <Dx>	-CHSAPE ELLIPSOID -SHPAR "<Nx> <Ny> <Nz>" -AEFF "..."	
	Bisphere	-object nspheres 2 <R> <R>	-shape bisphere 1 -size <Dx>	-CSHAPE ELLIPSO_2 -SHPAR "<Nx> <Ny> <Nz>" -AEFF "..." ^a	
	Coated sphere	-object concentricsphere 2 <Rin> <Rout>	-shape coated <Din/D> -size <D>	-CSHAPE ONIONSHEL -SHPAR "1 <D> <Din/D>" -AEFF "..."	
	Arbitrary	-object arbitrary <filename>	-shape read <filename> -size <Dx>	-CSHAPE FROM_FILE -AEFF "..."	
Particle	Orientation	-orient <phi> <theta> <psi>	-orient <alpha> <beta> <gamma>	-ROT_BETA "<beta> <beta> 1" -ROT_THETA "<Theta> <Theta> 1" -ROT_PHI "<Phi> <Phi> 1"	
Material	Isotropic	-epsmulti <eps_r> <eps_i>	-m <m_r> <m_i>	-DIEL "filename"	0
	Diagonal anisotropic	-trope <eps1_r> <eps1_i> ... <eps3_r> <eps3_i>	-anisotr -m <m1_r> <m1_i> <m2_r> <m2_i> <m3_r> <m3_i>	-NCOMP 3 --diels "filename1" "filename2" "filename3"	
	Inhomogeneous	-epsmulti <...> (per material)	-m <...> (per material)	-NCOMP <Ncomp> --diels "... (per material)"	
Incident beam	Linear (or \hat{y})	-beam pwavelinear 0 0	-beam plane -sym enf -scat_matr none	-IORTH 1 -POL_E01 "(0,0) (1,0) (0,0)"	0
	Linear (\perp or \hat{x})	-beam pwavelinear 1 0	natively unsupported ^b	-IORTH 1 -POL_E01 "(0,0) (0,0) (-1,0)"	
	Linear (arbitrary)	-beam pwavelinear <pola> 0	natively unsupported ^b	-IORTH 1 -POL_E01 "(0,0) (<s>,0) (-<p>,0)"	
	Two linear polarizations	—	-beam plane	-IORTH 2 -POL_E01 "(0,0) (1,0) (0,0)"	
	Circular (right)	-beam pwavecircular 1	natively unsupported ^b	-IORTH 1 -POL_E01 "(0,0) (1,0) (0,-1)"	
	Circular (left)	-beam pwavecircular -1	natively unsupported ^b	-IORTH 1 -POL_E01 "(0,0) (1,0) (0,1)"	
	Propagation direction	-prop <theta> <phi>	-prop <x> <y> <z>	—	
	Wavelength	-lambda <value>	-lambda <value>	-WAVELENGTHS "<value> <value> 1 'LIN'"	
Grid	Grid size	-nnnr <N>	-grid <N>	-MEM_ALLOW "<Nx> <Ny> <Nz>" -SHPAR "..." ^c	1-2
Formulation	Clausius-Mossotti	-polarizability CM	-pol cm	—	1-2
	Radiative reaction	-polarizability RR	-pol rrc	—	
	Digitized Green's function	-polarizability GB	-pol dgf	—	
	Lakhtakia	-polarizability LA	-pol lak	—	
	LDR (standard)	-polarizability LS	-pol ldr	-CALPHA LATTDR	
	LDR (isotropized)	-polarizability LR	-pol ldr avgpol	—	
	Integrated Green's tensor	-igt <value>	-int igt <value>	—	
	Filtered Green	-polarizability FG	-pol fcd -int fcd	-CALPHA FLTRCD	
	POI	-igt 0	-int poi	—	
	Scattering formula	—	-scat dr	—	
Solver	QMR	-methodeit QMRCLA	-iter qmr	-CMDSOL QMRCCG	η
	BiCGStab	-methodeit BICGSTAB	-iter bicgstab	-CMDSOL PBCGST	
	GPBiCG	-methodeit GPBiCG1	—	-CMDSOL GPBiCG	
	BCGS2	—	-iter bcgs2	-CMDSOL PBCGS2	
	Threshold η	-tolinit 1.d-<n>	-eps <n>	-TOL 1.d-<n>	
	Initial guess	-ninitest 0	-init_field zero	—	

^a Here, -SHPAR defines the grid size of one sphere.

^b Not directly available in ADDA; see the text for available workarounds.

^c -SHPAR option defines both geometry of the shape (aspect ratios and, sometimes, the main axes) and the level of the discretization (n_x , etc.). The accuracy associated with grid is relevant if only the discretization parameters are varied keeping all aspect ratios fixed.

Table C.2

Command-line equivalences between IFDDA, ADDA, and DDSCAT, defining the calculated output quantities.

Parameter	IFDDA	ADDA	DDSCAT
Field	-near_field 0 1 0 -save_data 2	-store_int_field	-NRFLD 1
Polarization	—	-store_dip_pol	-NRFLD 1
Force	-nforce 1 1 0 0	-Cpr	-CMDTRQ DOTORQ
Torque	-nforce 1 1 1 1	—	-CMDTRQ DOTORQ
Mueller matrix ^a	—	-scat_matr muel -ntheta 180	-IORTH 2 -CMDFRM LFRAME -NPLANES 1 -PLANE "0 0 180 1" -NSMELTS 16 -SMELTS_LIST "11 12 13 14 21 22 23 24 31 32 33 34 41 42 43 44"

^a Corresponds to the default ADDA options for a single scattering plane with particle symmetric with respect to the xz -plane.

To rotate the particle without rebuilding the grid, both codes use a passive approach, i.e., instead of actively rotating the dipole coordinates, they rotate the incident propagation and polarization vectors (and scattering directions) into the TF. Thus, rotating the target is completely equivalent to that of the incident beam – the corresponding matrices are inverse of each other.

Despite this common strategy, the two codes differ by conventions and default axis choices. ADDA employs a passive Euler rotation matrix in the zyz convention and the default $\hat{\mathbf{k}}_0 = \hat{\mathbf{z}}$, while DDSCAT— xzx convention [74] and hard-coded $\hat{\mathbf{k}}_0 = \hat{\mathbf{x}}$. Here, $\hat{\mathbf{k}}_0$ denotes the incident propagation in the LF to discriminate it from $\hat{\mathbf{k}}$ in the TF that is used in the DDA equations. Moreover, for some shapes (e.g. ellipsoids or SPHERES_N), DDSCAT allows one to interchange or rotate the target axes (denoted by $\hat{\mathbf{a}}_i$), which corresponds to preliminary active rotation of the particle — we do not investigate this feature and assume the default $\hat{\mathbf{a}}_i$.

While any combination of $\hat{\mathbf{k}}_0$ and Euler angles in ADDA can be matched by DDSCAT, one-to-one correspondence takes place only if $\hat{\mathbf{k}}_0$ is fixed. Most naturally for the latter is to use $\hat{\mathbf{k}}_0 = \hat{\mathbf{x}}$, corresponding to -prop 1 0 0 in ADDA, which allows us to disentangle equivalence conditions for propagation direction and orientation. The latter then boils down to equality of Euler-rotation matrices with different conventions

$$\mathbf{R}_{zyz}(\alpha, \beta, \gamma) = \mathbf{R}_{xzx}(\Phi, \Theta, \beta_D), \quad (\text{C.2})$$

where we denote the third Euler angle in DDSCAT as β_D to discriminate β used in its manual from that of ADDA. As expected, default orientations (all zero angles) correspond to each other, but otherwise the transformation of angles expressions involve direct and inverse trigonometric functions.

By contrast, IFDDA applies an active rotation, i.e., the voxel coordinates are defined in the LF and then rotated to be tested against the particle shape in the TF. This approach is physically equivalent to the passive formulation, but it does not reuse the same voxel grid for different orientations. Effectively, the DDA equations are solved in the LF rather than in the TF. Thus, exact agreement between IFDDA and other codes is only expected for rotations that map the grid lattice onto itself, i.e., for all angles being multiples of 90° . Otherwise, the agreement is limited to method error.

The propagation direction in IFDDA can be freely specified as in ADDA, but using spherical angles instead of Cartesian components of $\hat{\mathbf{k}}_0$. Thus, -prop 90 0 in IFDDA corresponds to $\hat{\mathbf{k}}_0 = \hat{\mathbf{x}}$ in DDSCAT, in which case, the rotation matrix $\mathbf{R}_{xzx}(\psi, \theta, \varphi)$ should be equal to the ones in Eq. (C.2) (all three using different conventions). Angles θ, φ are as used in IFDDA manual; they should not be confused with angles defining incident or scattering directions.

In the LF, DDSCAT uses $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$ as basic polarization vectors, while ADDA and IFDDA employ $\hat{\mathbf{y}}$ and $\hat{\mathbf{x}}$ with the default $\hat{\mathbf{k}}_0 = \hat{\mathbf{z}}$. When the propagation vector is rotated into $\hat{\mathbf{x}}$, the latter two polarizations become $\hat{\mathbf{y}}$ and $-\hat{\mathbf{z}}$, i.e., the same as in DDSCAT apart from the minus. This minus sign has no effect for most of the simulated quantities, except for internal fields. As discussed below, $\hat{\mathbf{y}}$ defines the default scattering plane in both ADDA and DDSCAT for both of the discussed directions of the incident field. Thus, we denote it as \parallel to this plane and the complementary polarization as \perp . This designation remains meaningful for any $\hat{\mathbf{k}}_0$.

Although DDSCAT is the most flexible code in specifying the single incident polarization by a full complex vector (which can match any setting of IFDDA and ADDA), the second polarization is automatically obtained from the first (by default, it is $\hat{\mathbf{y}} \rightarrow \hat{\mathbf{z}}$). That means that -IORTH 2 for DDSCAT would use the same two polarizations as ADDA, but the second one will have an inverse sign. Alternatively, one may use -IORTH 2 -POL_E01 "(0,0) (0,0) (-1,0)" to use exactly the same polarizations, but in reverse order.

IFDDA supports circular polarizations, as well as arbitrary linear ones. The latter is controlled by a single real parameter pol_a (with $0 \leq \text{pol}_a \leq 1$), which sets the relative weight of two basic polarizations. This can be mapped to DDSCAT by specifying the complex incident polarization unit vector $\hat{\mathbf{e}}_{01}$ as

$$\hat{\mathbf{e}}_{01} = (0, s, -p), \quad s \stackrel{\text{def}}{=} \sqrt{1 - \text{pol}_a}, \quad p \stackrel{\text{def}}{=} \sqrt{\text{pol}_a}, \quad (\text{C.3})$$

ADDA is designed mostly for simulating two incident polarizations to obtain the Mueller matrix. To simulate a single one (and not waste computational resources), one can trick ADDA into considering the scatterer symmetric, but then produced Mueller matrix will generally be incorrect (and better discarded). Thus, addition of -sym enf -scat_matr none is a workaround to simulate only \parallel polarization. To rotate this single linear polarization, there are two possible workarounds. First, it can be emulated by additional rotation of the particle. However, no simple expressions exist for arbitrary particle orientation and $\hat{\mathbf{k}}_0$. Still, a few special case can be named: for $\hat{\mathbf{k}}_0 = \hat{\mathbf{z}}$, one only need to increment Euler angle $\gamma \rightarrow \gamma + \arcsin p$, while for $\hat{\mathbf{k}}_0 = \hat{\mathbf{x}}$ and default particle orientation we have $\alpha = -\pi/2$, $\beta = \arcsin p$, $\gamma = \pi/2$. The second option is through the general Bessel beams implemented in ADDA, by setting both order and half-cone angle to zero: -beam besse1M 0 0 0 0 <p> <s>, although it is incompatible with calculation of radiation forces. Each of the workarounds should be combined with the above symmetry trick, otherwise the behavior will be similar to -IORTH 2 in DDSCAT.

For circular polarizations, we assume the definition of IFDDA, in agreement with Ref. [75], specifically, $\hat{\mathbf{e}}_{R,L} \stackrel{\text{def}}{=} (\hat{\mathbf{y}} \pm i\hat{\mathbf{x}})/\sqrt{2}$ for right- and left-circular polarizations, respectively, of a beam propagating along $\hat{\mathbf{z}}$. That is essentially opposite to the ones given in DDSCAT manual, which are equivalent to $\hat{\mathbf{e}}_L$ and $i\hat{\mathbf{e}}_R$, respectively, after rotation of $\hat{\mathbf{k}}_0$. In ADDA, such polarizations can be emulated through the Bessel beam, for instance, -beam besse1M 0 0 0 0 0 <is2> 0 0 <±is2> 0 for $\hat{\mathbf{e}}_{R,L}$, respectively, where $\text{is2} = 1/\sqrt{2}$. Same as above, such emulation is not suitable for radiation forces. Moreover, computation of two polarizations with ADDA is completely meaningless in this case, since those two polarizations will be equivalent (in contrast to -IORTH 2 in DDSCAT). Finally, the QMR solver (as well as QMR2 and BiCG) breaks down for circular polarizations, since the pseudo-norm of incident field vanishes [49]; this can be resolved by switching to BiCGStab.

For the isotropized LDR polarizability, ADDA and IFDDA use different conventions. ADDA averages the term S only over incident polarizations keeping the propagation direction fixed (so the result depends on the latter, see Eq. (46) in Ref. [7]), whereas IFDDA performs full averaging leading to $\langle S \rangle = 1/5$. Therefore, the incident propagation angles must be chosen carefully to reproduce the same value of S . An example is

$\hat{\mathbf{k}} = (\sqrt{(5 - \sqrt{5})/10}, 0, \sqrt{(5 + \sqrt{5})/10})$, which corresponds to angle $\theta = (1/2) \arctan 2$ in IFDDA.

Several additional numerical aspects are worth noting. The option `-scat fin` of ADDA affects the value of C_{abs} only for some combination of other parameters (like absorbing particles with the LDR polarizability). By contrast, enabling `igt_so` reduces the accuracy of C_{ext} to 1–2 significant digits compared to full IGT, as for other variations of formulations. However, most differences in the formulation class are $\mathcal{O}[(kd)^2]$, which decreases with particle size (for fixed discretization). The notable exception is the interaction formulations, which remain significant even for particles much smaller than the wavelength. Finally, the equivalence of the BiCGstab solver is problem dependent, due to the last residual stopping criterion.

Table C.2, by contrast, describes equivalent command-line parameters to produce additional scattering quantities. They can be optionally added on top of the free parameters listed in Table C.1, as they do not affect other output of the codes. When these additional combinations are used exactly as listed, these quantities inherit the accuracy level of the corresponding configuration listed in Table C.1. Otherwise, the expected accuracy level drops to 0, as it may incur comparison of different types of fields or different components of forces.

ADDA and IFDDA compute the radiation force (and torque in IFDDA) exerted on each dipole [76], whereas DDSCAT provides only total integrated values obtained from numerical integration over scattering directions [77]. Moreover, in DDSCAT the radiation force \mathbf{Q}_{pr} must be post-processed using momentum conservation (i.e. $Q_{\text{ext}} \hat{\mathbf{k}} - Q_{\text{sca}} \mathbf{g}$, where \mathbf{g} is the asymmetry parameter). As a result, while machine-precision agreement is achievable between ADDA and IFDDA, comparisons involving DDSCAT are typically limited to about 4–5 significant digits due to numerical integration errors, even when using millions of scattering directions. Moreover, the above relation is valid only for plane incident waves, although the same limitation is currently present in ADDA for dipole-wise calculation of forces.

To compare radiation forces between IFDDA and ADDA, additional post-processing is required, since IFDDA outputs the total time-averaged force \mathbf{F}_{rad} (in N), while ADDA outputs the radiation-pressure cross section C_{pr} (in nm^2 in our case) and the corresponding efficiency Q_{pr} . The cross section and the total force are related by

$$\mathbf{F}_{\text{rad}} = C_{\text{pr}} |\mathbf{E}_0|^2 \frac{\epsilon_0}{2}. \quad (\text{C.4})$$

For the radiation torques, the process is similar since IFDDA outputs the total time-averaged torque Γ_{rad} (in $\text{N} \cdot \text{m}$), while DDSCAT outputs the radiation-torque efficiency Q_{trq} . They are related by

$$\Gamma_{\text{rad}} = Q_{\text{trq}} \pi a_{\text{eff}}^2 |\mathbf{E}_0|^2 \frac{\epsilon_0}{2k}. \quad (\text{C.5})$$

Regarding the internal field, IFDDA and DDSCAT compute the macroscopic field by multiplying the local/excited field by, respectively, radiative-reaction correction and the Clausius–Mossotti factor, whereas ADDA rigorously computes the internal field according to the selected polarizability (see Eq. (11)). Therefore, machine-precision agreement is achievable only for these two polarizabilities (and not between all three codes); for other choices, the expected level of agreement drops to that obtained when the polarizations are not matched (1–2 digits). Moreover, DDSCAT does not support the Clausius–Mossotti polarizability (see Table 1), so full machine precision is not achievable for it as well. One can, in principle, compare instead voxel polarizations between DDSCAT and ADDA, but this requires additional efforts: the source code of DDSCAT need first to be modified to save the internal polarizations multiplied by the voxel volume d^3 (to match the ADDA convention).

Finally, it is also possible to compare Mueller matrices between ADDA and DDSCAT. In the latter code (with `-CMDFRM LFRAME`), the scattering direction is parameterized by (θ, ϕ) , where $\phi = 0$ corresponds to the reference scattering plane coinciding with the $(\hat{\mathbf{x}}_{\text{LF}}, \hat{\mathbf{y}}_{\text{LF}})$ plane. Note that these θ and ϕ are distinct from the particle-orientation angles Θ and

Table C.3

Matching digits for C_{ext} and Q_{ext} between ADDA, IFDDA, and DDSCAT.

Grid size	Code	C_{ext} [nm^2]	Q_{ext}
$n_x = 150$	ADDA	24729875.277 382743	3.58901283529 54913
	IFDDA	24729875.277 422981	–
	DDSCAT	–	3.58901283529 78903
$n_x = 250$	ADDA	24730049.76 7323192	3.5890381587 809035
	IFDDA	24730049.76 6945954	–
	DDSCAT	–	3.5890381587 500482

Φ). In ADDA, the scattering plane is also set by the incident basis vectors and corresponds to the $(\hat{\mathbf{k}}, \hat{\mathbf{e}}_y)$ plane. Therefore, once these vectors are mapped between the two formulations (see above), both codes share the same scattering-plane convention, enabling direct comparisons of Mueller matrices.

Comparison of the amplitude scattering matrices would be more complicated, since DDSCAT defines it with respect to fixed incident polarizations [9] rather than to the ones rotated together with the scattering plane [75]. Moreover, DDSCAT does not produce the original complex elements of this matrix, but rather real quadratic combinations that are further used to compute the Mueller matrix.

Table C.3 provides an example of the matching digits obtained in this work (see the command-lines in Section 4). Note that ADDA reports the extinction cross section C_{ext} in nm^2 (when nm is assumed for input quantities), so the IFDDA values, originally expressed in m^2 , were converted to nm^2 before comparison. By contrast, DDSCAT is compared only through the dimensionless efficiency Q_{ext} .

D. MPI memory scaling model

In MPI programs, each rank maintains its own process space and communication buffers, which leads to a node-level memory overhead that scales both linearly and quadratically with the number of ranks P . We model the total resident set size (RSS) as

$$\text{RSS}(N, P) \approx D_{\text{ADDA}}(N) + C_{\text{RSS}} P + \alpha_{\text{RSS}} P^2, \quad (\text{D.1})$$

where $D_{\text{ADDA}}(N)$ is the application-dependent working set predicted by Eq. (5) of the ADDA manual [49], C_{RSS} is the resident portion of the constant per-rank libraries runtime environment used, and α_{RSS} quantifies the memory usage of MPI communication pools. Preliminary fits on our cluster suggest $C_{\text{RSS}} \approx 20 \text{ MiB}$ and $\alpha_{\text{RSS}} \approx 0.3 \text{ MiB}$, but they are expected to vary depending on hardware and MPI implementation.

The quadratic term is most pronounced in single-node runs with many ranks; for multi-node configurations with fewer ranks per node, the growth approaches linearity. At large grid sizes, $D_{\text{ADDA}}(N)$ dominates and MPI overhead becomes negligible compared to the overall working set. This seems to correspond to the natural design goal of MPI implementations, that most of the node memory should remain accessible to the application itself.

E. Fully offloaded ADDA solver (OCL_BLAS)

Fig. E.1 reports the FFTs time, solver time, and single-core wall-clock time on the A100, H200, and RTX 2000 Ada GPUs, for both $n_x = 150$ and $n_x = 250$ grids, comparing the OCL_BLAS mode of ADDA (BiCG solver) with the default OpenCL backend using the same BiCG solver. This ensures a fair comparison with OCL_BLAS and allows a meaningful reference to the original ADDA BiCGstab results (also duplicated here from Fig. 3). Unfortunately, timing inside the OpenCL code has not yet been implemented; thus, the FFT timing is not available for OCL_BLAS mode.

On all three GPUs and for both grid sizes, OCL_BLAS systematically outperforms the default ADDA OpenCL mode, except on the NVIDIA RTX 2000 Ada for $n_x = 250$, where the computation fails due to an out-of-memory error (limit is around $n_x = 230$). The improvement is most pronounced on the H200, where solver time is reduced by factors of approximately 9 ($n_x = 150$) and 6.5 ($n_x = 250$). This confirms that, in the

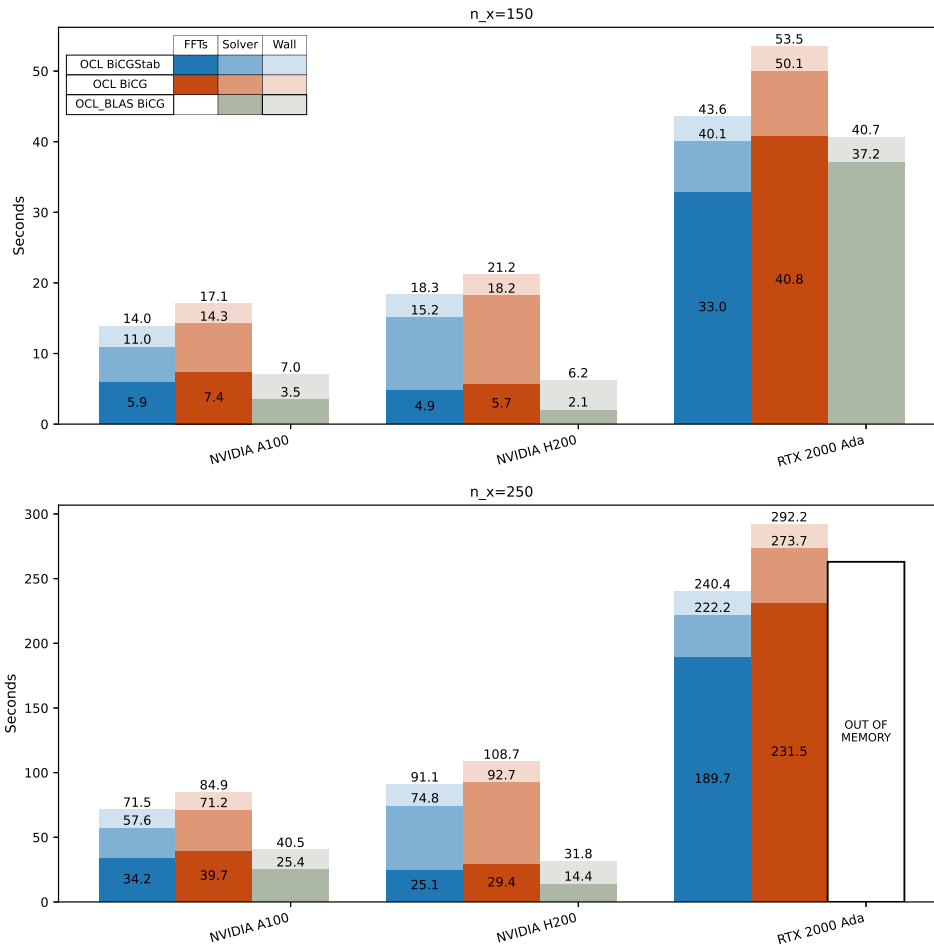


Fig. E.1. GPU timings for ADDA and ADDA OCL_BLAS mode. Bars are grouped by mode and solver type and stacked to show FFT time, solver time, and 1-core wall-time. Grid size $n_x = 150$ (top row); $n_x = 250$ (bottom row). Blue-column results (BICGStab) are the same as in Fig. 3.

standard OpenCL mode, performance is limited primarily by CPU–GPU-transfer latency rather than by GPU compute throughput.

The trends observed here closely mirror those of IFDDA: switching from A100 to H200 increases the OCL_BLAS solver performance by a factor of approximately 1.8, consistent with the expected scaling trend based on GPU memory bandwidth.

References

- [1] M.I. Mishchenko, *Electromagnetic Scattering by Particles and Particle Groups: An Introduction*, Cambridge University Press, Cambridge, UK, 2014. <https://doi.org/10.1017/CBO9781139019064>
- [2] A. Taflove, S.C. Hagness, *Computational Electromagnetics: The Finite-Difference Time-Domain Method*, Artech House, 2005.
- [3] P. Monk, *Finite Element Methods for Maxwell's Equations*, Oxford university press, 2003. <https://doi.org/10.1093/acprof:oso/9780198508885.001.0001>
- [4] R.F. Harrington, *Field Computation by Moment Methods*, Wiley-IEEE Press, 1993.
- [5] M.A. Yurkin, M.I. Mishchenko, Volume integral equation for electromagnetic scattering: Rigorous derivation and analysis for a set of multilayered particles with piecewise-smooth boundaries in a passive host medium, *Phys. Rev. A* 97 (2018) 043824. <https://doi.org/10.1103/PhysRevA.97.043824>
- [6] P.C. Chaumet, The discrete dipole approximation: a review, *Mathematics* 10 (2022) 3049. <https://doi.org/10.3390/math10173049>
- [7] M.A. Yurkin, Discrete dipole approximation, in: M.P. Menguc, M. Francoeur (Eds.), *Light, Plasmonics and Particles*, Elsevier, Amsterdam, 2023, pp. 167–198. <https://doi.org/10.1016/B978-0-323-99901-4.00020-2>
- [8] E.M. Purcell, C.R. Pennypacker, Scattering and adsorption of light by nonspherical dielectric grains, *Astrophys. J.* 186 (1973) 705–714. <https://doi.org/10.1086/152538>
- [9] B.T. Draine, The discrete dipole approximation and its application to interstellar graphite grains, *Astrophys. J.* 333 (1988) 848–872. <https://doi.org/10.1086/166795>
- [10] B.T. Draine, P.J. Flatau, Discrete-dipole approximation for scattering calculations, *J. Opt. Soc. Am. A* 11 (1994) 1491–1499. <https://doi.org/10.1364/JOSAA.11.001491>
- [11] C.M. Sorensen, J. Yon, F. Liu, J. Maughan, W.R. Heinson, M.J. Berg, Light scattering and absorption by fractal aggregates including soot, *J. Quant. Spectrosc. Radiat. Transf.* 217 (2018) 459–473. <https://doi.org/10.1016/j.jqsrt.2018.05.016>
- [12] V. Amendola, Surface plasmon resonance of silver and gold nanoparticles in the proximity of graphene studied using the discrete dipole approximation method, *Phys. Chem. Chem. Phys.* 18 (2016) 2230–2241. <https://doi.org/10.1039/C5CP06121K>
- [13] M.A. Yurkin, K.A. Semyanov, P.A. Tarasov, A.V. Chernyshev, A.G. Hoekstra, V.P. Maltsev, Experimental and theoretical study of light scattering by individual mature red blood cells by use of scanning flow cytometry and a discrete dipole approximation, *Appl. Opt.* 44 (2005) 5249–5256. <https://doi.org/10.1364/AO.44.005249>
- [14] J. Yon, C. Rozé, T. Girasole, A. Coppalle, L. Mèès, Extension of RDG-FA for scattering prediction of aggregates of soot taking into account interactions of large monomers, *Part. Part. Syst. Charact.* 25 (2008) 54–67. <https://doi.org/10.1002/ppsc.200700011>
- [15] C. Argentin, M.J. Berg, M. Mazur, R. Ceolato, J. Yon, Electromagnetic coupling and determination of the structure factor of fractal aggregates, *J. Quant. Spectrosc. Radiat. Transf.* 296 (2023) 108451. <https://doi.org/10.1016/j.jqsrt.2022.108451>
- [16] M.A. Yurkin, A.G. Hoekstra, The discrete-dipole-approximation code ADDA: capabilities and known limitations, *J. Quant. Spectrosc. Radiat. Transf.* 112 (2011) 2234–2247. <https://doi.org/10.1016/j.jqsrt.2011.01.031>
- [17] P.C. Chaumet, D. Sentenac, G. Maire, M. Rasedujjaman, T. Zhang, A. Sentenac, IFDDA, an easy-to-use code for simulating the field scattered by 3D inhomogeneous objects in a stratified medium: tutorial, *J. Opt. Soc. Am. A* 38 (2021) 1841–1852. <https://doi.org/10.1364/JOSAA.432685>
- [18] ADDA-team, ADDA Wiki page on other related DDA codes, 2025 (accessed 25 November 2025), <https://github.com/adda-team/adda/wiki/Links#other-related-dda-codes>.
- [19] Discrete dipole approximation codes Wikipedia page, 2025 (accessed 25 November 2025), https://en.wikipedia.org/wiki/Discrete_dipole_approximation_codes.
- [20] A. Penttilä, E. Zubko, K. Lumme, K. Muinonen, M.A. Yurkin, B.T. Draine, J. Rahola, A.G. Hoekstra, Y. Shkuratov, Comparison between discrete dipole

- implementations and exact techniques, *J. Quant. Spectrosc. Radiat. Transf.* 106 (2007) 417–436. <https://doi.org/10.1016/j.jqsrt.2007.01.026>
- [21] M. Wiesenberger, L. Einkemmer, M. Held, A. Gutierrez-Milla, X. Saez, R. Iakymchuk, Reproducibility, accuracy and performance of the Feltor code and library on parallel computer architectures, *Comput. Phys. Commun.* 238 (2019) 145–156. <https://doi.org/10.1016/j.cpc.2018.12.006>
- [22] C. Collange, D. Defour, S. Graillat, R. Iakymchuk, Numerical reproducibility for the parallel reduction on multi- and many-core architectures, *Parallel Comput.* 49 (2015) 83–97. <https://doi.org/10.1016/j.parco.2015.09.001>
- [23] L. Durdevic, A. Relano Gínés, A. Roueff, G. Blivet, G. Baffou, Biomass measurements of single neurites in vitro using optical wavefront microscopy, *Biomed. Opt. Express* 13 (2022) 6550–6560. <https://doi.org/10.1364/BOE.471284>
- [24] P.C. Chaumet, A comparative study of efficient iterative solvers for the discrete dipole approximation, *J. Quant. Spectrosc. Radiat. Transf.* 312 (2024) 108816. <https://doi.org/10.1016/j.jqsrt.2023.108816>
- [25] M.A. Yurkin, A.G. Hoekstra, The discrete dipole approximation: an overview and recent developments, *J. Quant. Spectrosc. Radiat. Transf.* 106 (2007) 558–589. <https://doi.org/10.1016/j.jqsrt.2007.01.034>
- [26] N.B. Pillier, O.J.F. Martin, Increasing the performance of the coupled-dipole approximation: a spectral approach, *IEEE Trans. Antennas Propag.* 46 (1998) 1126–1137. <https://doi.org/10.1109/8.718567>
- [27] M.A. Yurkin, M. Min, A.G. Hoekstra, Application of the discrete dipole approximation to very large refractive indices: filtered coupled dipoles revived, *Phys. Rev. E* 82 (2010) 036703. <https://doi.org/10.1103/PhysRevE.82.036703>
- [28] P.C. Chaumet, A. Sentenac, A. Rahmani, Coupled dipole method for scatterers with large permittivity, *Phys. Rev. E* 70 (2004) 036606. <https://doi.org/10.1103/PhysRevE.70.036606>
- [29] D.A. Smuneev, P.C. Chaumet, M.A. Yurkin, Rectangular dipoles in the discrete dipole approximation, *J. Quant. Spectrosc. Radiat. Transf.* 156 (2015) 67–79. <https://doi.org/10.1016/j.jqsrt.2015.01.019>
- [30] S.A. Glukhova, M.A. Yurkin, Vector Bessel beams: general classification and scattering simulations, *Phys. Rev. A* 106 (2022) 033508. <https://doi.org/10.1103/PhysRevA.106.033508>
- [31] A.A. Kichigin, M.A. Yurkin, Simulating electron energy-loss spectroscopy and cathodoluminescence for particles in arbitrary host medium using the discrete dipole approximation, *J. Phys. Chem. C* 127 (2023) 4154–4167. <https://doi.org/10.1021/acs.jpcc.2c06813>
- [32] T.F. Chan, An optimal circulant preconditioner for Toeplitz systems, *SIAM J. Sci. Stat. Comput.* 9 (1988) 766–771. <https://doi.org/10.1137/0909051>
- [33] P.C. Chaumet, G. Maire, A. Sentenac, Accelerating the discrete dipole approximation by initializing with a scalar solution and using a circulant preconditioning, *J. Quant. Spectrosc. Radiat. Transf.* 298 (2023) 108505. <https://doi.org/10.1016/j.jqsrt.2023.108505>
- [34] S. Lanier, Learning to precondition: reinforcement learning enhanced three-level circulant preconditioning for the discrete dipole approximation, *J. Quant. Spectrosc. Radiat. Transf.* 350 (2026) 109741. <https://doi.org/10.1016/j.jqsrt.2025.109741>
- [35] C. Argentin, M.A. Yurkin, Accelerating iterative solvers in the discrete dipole approximation using dedicated initial guesses, in: *The 21th Electromagnetic and Light Scattering Conference, University of Messina, Milazzo, Italy, 2025*, pp. 50. URL <https://hal.science/hal-05142441>
- [36] P.C. Chaumet, G. Maire, A. Sentenac, Scalar approximation of Maxwell equations: derivation and accuracy, *J. Opt. Soc. Am. A* 39 (2022) 1462–1467. <https://doi.org/10.1364/JOSAA.462034>
- [37] P.C. Chaumet, G. Maire, A. Sentenac, Accuracy of the multilayer Born approximation in transmission and reflection, *J. Opt. Soc. Am. A* 42 (2025) 413–421. <https://doi.org/10.1364/JOSAA.549504>
- [38] C. Temperton, A generalized prime factor FFT algorithm for any $N = 2^p 3^q 5^r$, *SIAM J. Sci. Stat. Comput.* 13 (1992) 676–686. <https://doi.org/10.1137/0913039>
- [39] M. Frigo, S.G. Johnson, FFTW: an adaptive software architecture for the FFT, in: *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 3, 1998, pp. 1381–1384. <https://doi.org/10.1109/ICASSP.1998.681704>
- [40] R.C. Singleton, On computing the fast Fourier transform, *Commun. ACM* 10 (1967) 647–654. <https://doi.org/10.1145/363717.363771>
- [41] T. Gamblin, M. LeGendre, M.R. Collette, G.L. Lee, A. Moody, B.R. de Supinski, S. Futral, The Spack package manager: bringing order to HPC software chaos, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2015, pp. 1–12. <https://doi.org/10.1145/2807591.2807623>
- [42] M.A. Yurkin, V.P. Maltsev, A.G. Hoekstra, Convergence of the discrete dipole approximation. I. Theoretical analysis, *J. Opt. Soc. Am. A* 23 (2006) 2578–2591. <https://doi.org/10.1364/JOSAA.23.002578>
- [43] N.B. Pillier, Influence of the edge meshes on the accuracy of the coupled-dipole approximation, *Opt. Lett.* 22 (1997) 1674–1676. <https://doi.org/10.1364/OL.22.001674>
- [44] M.A. Yurkin, M. Kahnert, Light scattering by a cube: accuracy limits of the discrete dipole approximation and the T-matrix method, *J. Quant. Spectrosc. Radiat. Transf.* 123 (2013) 176–183. <https://doi.org/10.1016/j.jqsrt.2012.10.001>
- [45] M. Yurkin, Erratum to "Discrete dipole approximation" [Light, Plasmonics and Particles, ed. Menguc M.P. and Francoeur M.,] (2025) 167–198. <https://doi.org/10.13140/RG.2.2.18469.61920>
- [46] J.D. Jackson, *Classical Electrodynamics*, Wiley, New York, 3rd ed., 1998.
- [47] B.T. Draine, P.J. Flatau, User Guide for the Discrete Dipole Approximation Code DDSCAT 7.3, 2013. [arXiv:1305.6497](https://arxiv.org/abs/1305.6497)
- [48] A. Lakhtakia, Strong and weak forms of the method of moments and the coupled dipole method for scattering of time-harmonic electromagnetic fields, *Int. J. Mod. Phys. C* 3 (1992). <https://doi.org/10.1142/S0129183192000385>
- [49] M.A. Yurkin, A.G. Hoekstra, *User Manual for the Discrete Dipole Approximation Code ADDA 1.4.0*, 2020. <https://github.com/adda-team/adda/raw/v1.4.0/doc/manual.pdf>
- [50] Z. Strakoš, J. Liesen, On numerical stability in large scale linear algebraic computations, *ZAMM* 85 (2005) 307–325. <https://doi.org/10.1002/zamm.200410185>
- [51] C. Argentin, ddsccatcli: v1.0.4, 2026, (Zenodo). URL <https://github.com/Argencle/ddsccatcli>. <https://doi.org/10.5281/zenodo.18619006>
- [52] M.I. Mishchenko, Scale invariance rule in electromagnetic scattering, *J. Quant. Spectrosc. Radiat. Transf.* 101 (2006) 411–415. <https://doi.org/10.1016/j.jqsrt.2006.02.047>
- [53] H.A. van der Vorst, Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 13 (1992) 631–644. <https://doi.org/10.1137/0913035>
- [54] R.D.d. Cunha, T. Hopkins, The parallel iterative methods (PIM) package for the solution of systems of linear equations on parallel computers, *Appl. Numer. Math.* 19 (1995) 33–50. [https://doi.org/10.1016/0168-9274\(95\)00017-0](https://doi.org/10.1016/0168-9274(95)00017-0)
- [55] J.M. Donald, A. Golden, S.G. Jennings, Opendda: a novel high-performance computational framework for the discrete dipole approximation, *Int. J. High Perform. Comput. Appl.* 23 (2009) 42–61. <https://doi.org/10.1177/1094342008097>
- [56] J.J. Goodman, B.T. Draine, P.J. Flatau, Application of fast-fourier-transform techniques to the discrete-dipole approximation, *Opt. Lett.* 16 (1991) 1198–1200. <https://doi.org/10.1364/OL.16.001198>
- [57] A.G. Hoekstra, M.D. Grimminck, P.M.A. Slood, Large Scale simulations of elastic light scattering by a fast discrete dipole approximation, *Int. J. Mod. Phys. C* 09 (1998) 87–102. <https://doi.org/10.1142/S012918319800008X>
- [58] A. Strack, C. Taylor, D. Pflüger, Parallel FFTW on RISC-V: a comparative study including openMP, MPI, and HPX, 2025. [arXiv:2506.08653](https://arxiv.org/abs/2506.08653)
- [59] M.A. Yurkin, D.A. Smuneev, S.A. Glukhova, A.A. Kichigin, A.E. Moskalensky, K.G. Inzhevatkin, C. Argentin, P. Bouillon, Recent developments of the ADDA code, in: *The 21th Electromagnetic and Light Scattering Conference, University of Messina, Milazzo, Italy, 2025*, pp. 36. <https://hal.science/hal-05142352>
- [60] B. Li, S. Cheng, J. Lin, tcfft: accelerating half-precision FFT through tensor cores (2021). [arXiv:2104.11471](https://arxiv.org/abs/2104.11471)
- [61] B.E. Barrowes, F.L. Teixeira, J.A. Kong, Fast algorithm for matrix–vector multiply of asymmetric multilevel block-Toeplitz matrices in 3-D scattering, *Microwave Opt. Technol. Lett.* 31 (2001) 28–32. <https://doi.org/10.1002/mop.1348>
- [62] M. Shabaninezhad, M.G. Awan, G. Ramakrishna, MATLAB package for discrete dipole approximation by graphics processing unit: fast fourier transform and bi-conjugate gradient, *J. Quant. Spectrosc. Radiat. Transf.* 262 (2021) 107501. <https://doi.org/10.1016/j.jqsrt.2020.107501>
- [63] S. Williams, A. Waterman, D. Patterson, Roofline: an insightful visual performance model for multicore architectures, *Commun. ACM* 52 (2009) 65–76. <https://doi.org/10.1145/1498765.149878>
- [64] A. Ayala, S. Tomov, M. Stoyanov, A. Haidar, J. Dongarra, Performance analysis of parallel FFT on large multi-GPU systems, in: *2022 IEEE International Parallel and Distributed Processing Symposium Workshops, IEEE, 2022*, pp. 372–381. <https://doi.org/10.1109/IPDPSW55747.2022.00072>
- [65] P. Steinbach, M. Werner, gearshift – The FFT benchmark suite for heterogeneous platforms, in: *High Performance Computing: 32nd International Conference, Springer-Verlag, Berlin, Heidelberg, 2017*, p. 199–216. https://doi.org/10.1007/978-3-319-58667-0_11
- [66] D. Xu, P. Tiersun, S. Li, M. Wang, L. Jiang, CPDDA: a python package for discrete dipole approximation accelerated by CuPy, *Nanomaterials* 15 (2025) 500. <https://doi.org/10.3390/nano15070500>
- [67] B.T. Draine, J. Goodman, Beyond Clausius-Mossotti-Wave propagation on a polarizable point lattice and the discrete dipole approximation, *Astrophys. J.* 405 (1993) 685–697. <https://doi.org/10.1086/172396>
- [68] J.D. Fenton, R.S. Gardiner-Garden, Rapidly-convergent methods for evaluating elliptic integrals and theta and elliptic functions, *J. Aust. Math. Soc. B* 24 (1982) 47–58. <https://doi.org/10.1017/S033427000003301>
- [69] B.R.A. Nijboer, F.W. De Wette, The internal field in dipole lattices, *Physica* 24 (1958) 422–431. [https://doi.org/10.1016/S0031-8914\(58\)95803-8](https://doi.org/10.1016/S0031-8914(58)95803-8)
- [70] NIST Digital Library of Mathematical Functions, Section 8.4: Incomplete Gamma and Related Functions, (<https://dlmf.nist.gov/8.4>). F.W.J. Olver, A.B. Olde Daalhuis, D.W. Lozier, B.I. Schneider, R.F. Boisvert, C.W. Clark, B.R. Miller, B.V. Saunders, H.S. Cohl, and M.A. McClain, eds. Release 1.0.24 (2019-09-15).
- [71] D.W. Mackowski, Discrete dipole moment method for calculation of the T matrix for nonspherical particles, *J. Opt. Soc. Am. A* 19 (2002) 881–893. <https://doi.org/10.1364/JOSAA.19.000881>
- [72] D. Gutkowicz-Krusin, B.T. Draine, Propagation of electromagnetic waves on a rectangular lattice of polarizable points, 2004. [arXiv:astro-ph/0403082](https://arxiv.org/abs/astro-ph/0403082)
- [73] C. Argentin, dda-bench: v0.1.1, 2026, (Zenodo). URL <https://github.com/Argencle/dda-bench>. <https://doi.org/10.5281/zenodo.18836855>
- [74] M.A. Yurkin, et al., Fair evaluation of orientation-averaging techniques in light-scattering simulations: Comment on “Evaluation of Higher-Order Quadrature Schemes in Improving Computational Efficiency for Orientation-Averaged Single-Scattering Properties of Nonspherical Ice Particles” by Fenni et al, *J. Geophys. Res. Atmos.* 128 (2023). <https://doi.org/10.1029/2021JD036088>
- [75] C.F. Bohren, D.R. Huffman, *Absorption and Scattering of Light by Small Particles*, John Wiley & Sons, 1983.
- [76] P.C. Chaumet, M. Nieto-Vesperinas, Time-averaged total force on a dipolar sphere in an electromagnetic field, *Opt. Lett.* 25 (2000) 1065–1067. <https://doi.org/10.1364/OL.25.001065>
- [77] B.T. Draine, J.C. Weingartner, Radiative torques on interstellar grains. I. Superthermal spin-up, *Astrophys. J.* 470 (1996) 551. <https://doi.org/10.1086/177887>