

Acceleration techniques for semiclassical Maxwell–Bloch systems: An application to discrete quantum dot ensembles[☆]

C. Glosser, E. Lu, T.J. Bertus, C. Piermarocchi ^{*}, B. Shanker

Department of Physics & Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA

Department of Electrical & Computer Engineering, Michigan State University, 428 South Shaw Lane, East Lansing, MI 48824, USA



ARTICLE INFO

Article history:

Received 23 October 2019

Received in revised form 17 April 2020

Accepted 20 July 2020

Available online 31 July 2020

Keywords:

Maxwell–Bloch equations

Quantum dots

Adaptive integral method

Integral equation

ABSTRACT

The solution to Maxwell–Bloch systems using an integral-equation-based framework has proven effective at capturing collective features of laser-driven and radiation-coupled quantum dots, such as light localization and modifications of Rabi oscillations. Importantly, it enables observation of the dynamics of each quantum dot in large ensembles in a rigorous, error-controlled, and self-consistent way without resorting to spatial averaging. Indeed, this approach has demonstrated convergence in ensembles containing up to 10^4 interacting quantum dots (Glosser et al., 2017). Scaling beyond 10^4 quantum dots tests the limit of computational horsepower, however, due to the $\mathcal{O}(N_t N_s^2)$ scaling (where N_t and N_s denote the number of temporal and spatial degrees of freedom). In this work, we present an algorithm that reduces the cost of analysis to $\mathcal{O}(N_t N_s \log^2 N_s)$. While the foundations of this approach rely on well-known particle–particle/particle–mesh and adaptive integral methods, we add refinements specific to transient systems and systems with multiple spatial and temporal derivatives. Accordingly, we offer numerical results that validate the accuracy, effectiveness and utility of this approach in analyzing the dynamics of large ensembles of quantum dots.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

The computational simulation of the nonlinear propagation of laser pulses through materials presents formidable challenges, particularly in materials containing dispersed quantum dots or nanoparticles that have strong light–matter coupling. One phenomenon, Rabi oscillations, demonstrates nonlinear behavior that can arise from such coupling. These oscillations have a long history of study in single quantum dots [1–3], though understanding the collective Rabi dynamics of quantum dot *ensembles* requires a careful analysis of secondary emissions that couple quantum dots to produce many-body collective effects. Experiments on novel systems based on perovskite nanocrystals have recently demonstrated these effects [4]; accordingly, researchers have a significant interest in examining these effects from theoretical, numerical, and experimental perspectives [5–8] given their potential to in developing novel composite materials with enhanced optical properties.

Typical theoretical and computational analyses use variations of the Maxwell–Bloch equations [9] to describe the collective

behavior of ensembles of optically active centers in which classical radiation fields couple a quantum description of each center. To this end, methods such as homogenization [10,11], differential-equation-based methods [12–16], and, more recently, integral-equation-based methods [7] all describe the dynamics of coupled quantum dot systems, each with a differing level of fidelity.

Approaches that do not rely on homogenization use coupled discrete methods to solve the classical Maxwell equations and local time evolution techniques to solve the Bloch equations for each quantum dot. Spatial homogenization, on the other hand, describes the near and far radiation characteristics of a quantum dot assuming homogeneous background material properties [17]. This approach has limited validity as it does not account for strong interactions between particles in each other's nearfield, a shortcoming exacerbated by the non-linear regimes considered here. Differential equation methods [18] to solve the Maxwell system have included finite-difference, time-domain finite-element, and discontinuous Galerkin methods, though all succumb to various numerical inaccuracies inherent to discretization. The inaccuracies most pertinent to simulation of quantum dot systems extend from the need to include point dipole sources in the simulation and capture near field effects that behave as $1/r^3$ (where r denotes the distance between centers). Accurately recovering these fields has numerous challenges and one needs dense discretization in the vicinity of quantum dots.

[☆] The review of this paper was arranged by Prof. Stephan Fritzsche.

^{*} Corresponding author at: Department of Physics & Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA.

E-mail address: piermaro@msu.edu (C. Piermarocchi).

together with equivalent/soft sources to accurately capture these effects [15]. Additionally, static null spaces that grow linearly with time present another challenge with conventional time-domain finite-element techniques [18]. On the positive side, these methods offer a high degree of flexibility and can accommodate different background linear bulk materials.

Our approach [7] differs significantly in that we make use of an integral equation-based formulation that employs a retarded potential to compute fields radiated by every quantum dot. This approach does not rely on a particular discretization and only depends on the number of quantum dots under investigation. Demonstrations of the accuracy of this approach appear in [7], though this method faces a fundamental bottleneck: the required memory and computational time both scale as $\mathcal{O}(N_t N_s^2)$ which becomes prohibitive for extended ensembles. A traditional acceleration technique that readily adapts to these equations makes use of a rotating frame approximation to reduce the number of time steps by a factor of ~ 1000 . This approximation exploits the narrowband nature of the nonlinearity which enables the use of envelope functions [16]. Physically, this arises due to the large difference between the characteristic Rabi energy associated with light-matter coupling and the optical transition energy. Computationally, this enables time step sizes much larger than the inverse of the laser frequency. While this approximation offers significant acceleration, the key bottleneck remains the cost of evaluating retarded potentials, which scales *quadratically* with the number of quantum dots. Thankfully, there exists extensive literature on reducing this cost that we examine next.

Two algorithms typically see use in accelerating the evaluation of retarded potentials: the Plane-Wave Time-Domain method (PWTD) and the Adaptive Integral Method (AIM, a variation of particle-particle/particle-mesh techniques) [19,20]. Unfortunately, neither of these methods apply directly to the problem at hand. To set the stage for discussion, assume that one needs to evaluate the radiated electric field $\mathbf{E}_R(\mathbf{r}, t) = \mathfrak{F}\{\mathbf{P}(\mathbf{r}, t)\}$ [21] due to a polarization density $\mathbf{P}(\mathbf{r}, t)$ via $\mathfrak{F}\{\mathbf{P}(\mathbf{r}, t)\} = \mathfrak{L}\{g(\mathbf{r}, t)\} * \mathbf{P}(\mathbf{r}, t)$ where $g(\mathbf{r}, t) = \delta(t - |\mathbf{r}|/c)/(4\pi|\mathbf{r}|)$ denotes the retarded potential, $\mathfrak{L}\{g(\mathbf{r}, t)\} = -\mu_0(\partial_t^2 \mathbf{I} - c^2 \nabla \nabla \cdot)g(\mathbf{r}, t)$ denotes a dyadic differential operator, and \mathbf{I} denotes the identity dyad.

PWTD exploits the properties of radiated fields due to quiescent sources that occupy a bounded spatial domain. These fields have a bandlimit (in momentum space) that PWTD leverages to reconstruct them to arbitrary precision using a tree-based approach; see [19] and the references therein. However, the crux of this methodology lies in the fact that spatial variation scales with the temporal one (times c). Unfortunately, while the correlation between spatial and temporal scales holds in the fixed frame, it does not in the rotating frame.

AIM, on the other hand, relies on moments around a uniform grid (independent of temporal variation) to reconstruct sources and their resulting field distributions. This idea – using grids for computing translationally-invariant functions by exploiting an underlying block Toeplitz structure – has seen extensive use in molecular dynamics simulations to evaluate Laplace, Helmholtz, and wave equation kernels [20,22,23]. In each of these cases, one evaluates the space (and time) convolution with a scalar quantity without the operator $\mathfrak{L}\{\cdot\}$ required herein. Unfortunately, the nature of $\mathfrak{L}\{\cdot\}$ determines the number of potentials that need evaluation. As we will see in the ensuing sections, the $\mathfrak{L}\{\cdot\}$ used for quantum dots contains numerous dyadic terms, each with different orders of temporal derivatives. As a result, a naïve application of AIM to each term of the expression quickly becomes untenable and we seek to develop a more expedient technique.

This paper has two principal contributions: (i) development and demonstration of techniques that overcome computational complexity (memory and CPU costs) associated with evaluation

of integral equation operators, and (ii) demonstration of these algorithms to examine optical systems containing quantum dots. In developing these algorithms, we examine their runtime and accuracy and, more importantly, show that evaluating $\mathfrak{L}\{g(\mathbf{r}, t)\} * \mathbf{P}(\mathbf{r}, t)$ incurs approximately the same cost as evaluating $g(\mathbf{r}, t) * \mathbf{P}(\mathbf{r}, t)$.

We organize the rest of this paper as follows: in Section 2 we define the problem, and provide the means to a solution in Section 3. Section 4 develops the AIM method for the Maxwell–Bloch problem and outlines its computational complexity. Next, in Section 5, we present a number of results that verify the claims of accuracy, complexity, and applicability of this method to a collection of quantum dots. Finally, in Section 6 we summarize the contributions of the paper and outline future research avenues.

2. Formulation

Consider a domain Ω that contains N_s randomly distributed quantum dots. A time-varying electromagnetic field of central frequency ω impinges on Ω and excites each quantum dot. We wish to develop the means to study the evolution of these quantum dots in response to both the incident excitation as well as radiation produced by other quantum dots. Toward this end, we employ a semi-classical approach to understand the response of each quantum dot to the incident field that comprises the laser field as well as fields radiated by other quantum dots (computed classically). In what follows, we provide a brief description of the requisite formulation for completeness; [7] provides a more detailed description.

As we assume dipolar transitions govern the response of each quantum dot to the exciting field, we write the time-dependence of a given quantum dot's density matrix, $\hat{\rho}(t)$, as

$$\frac{d\hat{\rho}}{dt} = \frac{-i}{\hbar} [\hat{\mathcal{H}}(t), \hat{\rho}] - \hat{\mathcal{D}}[\hat{\rho}]. \quad (1)$$

For two-level systems, $\hat{\rho}(t)$ denotes a two-by-two matrix with three unique unknowns (ρ_{00} and the real and imaginary parts of ρ_{01}), $\hat{\mathcal{H}}(t)$ represents a local Hamiltonian that governs the internal two-level structure of the quantum dot as well as its interaction with an external electromagnetic field, and $\hat{\mathcal{D}}$ describes the effects of spontaneous emission in the quantum dot. Explicitly,

$$\hat{\mathcal{H}}(t) = \begin{pmatrix} 0 & \hbar\chi(t) \\ \hbar\chi^*(t) & \hbar\omega_0 \end{pmatrix} \quad (2a)$$

$$\hat{\mathcal{D}}[\hat{\rho}] = \begin{pmatrix} (\rho_{00} - 1)/T_1 & \rho_{01}/T_2 \\ \rho_{10}/T_2 & \rho_{11}/T_1 \end{pmatrix}, \quad (2b)$$

where $\chi(t) \equiv \mathbf{d} \cdot \hat{\mathbf{E}}(\mathbf{r}, t)/\hbar$, $\mathbf{d} \equiv \langle 1|\hat{\mathbf{e}}\mathbf{r}|0\rangle$, and the kets represent the highest valence and lowest conduction states of the quantum dot under consideration. Finally, T_1 and T_2 characterize the spontaneous emission and decoherence times in the two-level system [24]. Unlike plasmonic resonances in metallic nanoparticles [25], atom-like excitonic resonances characterize semiconductor quantum dots. Therefore, we expect the spontaneous emission here to outweigh thermal noise at room temperature as thermal noise becomes insignificant at optical frequencies.

We compute the semi-classical interaction between quantum dots assuming coherent fields and negligible quantum statistical effects. Such assumptions imply classical electromagnetic interactions while preserving the two-level structure of individual quantum dots. To this end, we write the total electric field at any point in space and time as $\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_L(\mathbf{r}, t) + \mathfrak{F}\{\mathbf{P}(\mathbf{r}, t)\}$ where $\mathbf{E}_L(\mathbf{r}, t)$ denotes the incident laser field, $\mathbf{P}(\mathbf{r}, t)$ gives a polarization distribution arising from the off-diagonal elements (coherences)

of $\hat{\rho}$, and

$$\begin{aligned} \tilde{\mathcal{F}}\{\mathbf{P}(\mathbf{r}, t)\} &\equiv -\mu_0(\partial_t^2 \mathbf{I} - c^2 \nabla \nabla) \mathbf{g}(\mathbf{r}, t) * \mathbf{P}(\mathbf{r}, t) \\ &\equiv \frac{-1}{4\pi\epsilon_0} \int (\mathbf{I} - \bar{\mathbf{r}} \otimes \bar{\mathbf{r}}) \cdot \frac{\partial_t^2 \mathbf{P}(\mathbf{r}', t_R)}{c^2 |\mathbf{r} - \mathbf{r}'|} \\ &\quad + (\mathbf{I} - 3\bar{\mathbf{r}} \otimes \bar{\mathbf{r}}) \cdot \left(\frac{\partial_t \mathbf{P}(\mathbf{r}', t_R)}{c |\mathbf{r} - \mathbf{r}'|^2} + \frac{\mathbf{P}(\mathbf{r}', t_R)}{|\mathbf{r} - \mathbf{r}'|^3} \right) d^3 \mathbf{r}' \end{aligned} \quad (3)$$

(see [26, section §72]). In the above expression, $\bar{\mathbf{r}} \equiv (\mathbf{r} - \mathbf{r}')/|\mathbf{r} - \mathbf{r}'|$, \otimes represents the tensor product (i.e. $(\mathbf{a} \otimes \mathbf{b})_{ij} = a_i b_j$), $t_R \equiv t - |\mathbf{r} - \mathbf{r}'|/c$.¹ Thus, in a system composed of multiple quantum dots, Eq. (3) couples the evolution of each quantum dot by way of the off-diagonal matrix elements appearing in Eq. (2a). Note that this approach does not require an instantaneous dipole-dipole Coulomb term between (charge-neutral) quantum dots; the interactions between structures occur only via the electric field which propagates through space with finite velocity (see [30, sections A_{IV} and C_{IV}] for in-depth discussions of this point).

In the systems under consideration here, ω_0 lies in the optical frequency band (~ 1500 meV/h). Consequently, integrating Eq. (1) directly to resolve the Rabi dynamics that occur on the order of 1 ps quickly becomes computationally infeasible. Introducing $\tilde{\rho} = \hat{U} \hat{\rho} \hat{U}^\dagger$ where $\hat{U} = \text{diag}(1, e^{i\omega_0 t})$, we may instead write Eq. (1) as

$$\frac{d\tilde{\rho}}{dt} = \frac{-i}{\hbar} [\hat{U} \hat{H} \hat{U}^\dagger - i\hbar \hat{V}, \tilde{\rho}] - \hat{\mathcal{D}}[\tilde{\rho}], \quad \hat{V} \equiv \hat{U} \frac{d\hat{U}^\dagger}{dt}, \quad (4)$$

which contains only terms proportional to $e^{i(\omega_0 \pm \omega)t}$ if $\mathbf{E}(t) \sim \tilde{\mathbf{E}}(t) \cos(\omega t)$. Consequently, we ignore the high-frequency quantities (corresponding to $\omega_0 + \omega$) under the assumption that such terms will integrate to zero in solving Eq. (4) over appreciable timescales—an approximation known as the rotating wave approximation [24]. One can then construct efficient numerical strategies for solving Eq. (4). A similar transformation applies to the source distribution $\mathbf{P}(\mathbf{r}, t)$; by assuming $\mathbf{P}(\mathbf{r}, t) = \tilde{\mathbf{P}}(\mathbf{r}, t) e^{i\omega_0 t}$ in Eq. (3) the radiated field envelope becomes

$$\begin{aligned} \tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\} &\equiv \frac{-1}{4\pi\epsilon_0} \int (\mathbf{I} - \bar{\mathbf{r}} \otimes \bar{\mathbf{r}}) \\ &\quad \left(\partial_t^2 \tilde{\mathbf{P}}(\mathbf{r}', t_R) + 2i\omega \partial_t \tilde{\mathbf{P}}(\mathbf{r}', t_R) - \omega^2 \tilde{\mathbf{P}}(\mathbf{r}', t_R) \right) e^{-i\omega |\mathbf{r} - \mathbf{r}'|/c} \\ &\quad + \frac{c^2 |\mathbf{r} - \mathbf{r}'|}{(\mathbf{I} - 3\bar{\mathbf{r}} \otimes \bar{\mathbf{r}}) \cdot \frac{(\partial_t \tilde{\mathbf{P}}(\mathbf{r}', t_R) + i\omega \tilde{\mathbf{P}}(\mathbf{r}', t_R)) e^{-i\omega |\mathbf{r} - \mathbf{r}'|/c}}{c |\mathbf{r} - \mathbf{r}'|^2} \\ &\quad + (\mathbf{I} - 3\bar{\mathbf{r}} \otimes \bar{\mathbf{r}}) \cdot \frac{\tilde{\mathbf{P}}(\mathbf{r}', t_R) e^{-i\omega |\mathbf{r} - \mathbf{r}'|/c}}{|\mathbf{r} - \mathbf{r}'|^3} d^3 \mathbf{r}'. \end{aligned} \quad (5)$$

(Note that Eq. (5) critically maintains the high-frequency phase relationship between sources oscillating at ω via the factors of $e^{-i\omega |\mathbf{r} - \mathbf{r}'|/c}$ that appear.) Accordingly, we define

$$\tilde{\mathbf{E}}(\mathbf{r}, t) \equiv \tilde{\mathbf{E}}_L(\mathbf{r}, t) + \tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\} \quad (6)$$

and the evolution of the ensemble relies on a self-consistent solution of Eqs. (4) and (5). As evident from Eq. (5), this comprises a large number of costly potential integrals arising from the number of dyadic components and number of time derivatives. As such, we now turn our attention to an efficient computational infrastructure for ameliorating this cost.

¹ The convolution of this (dyadic) Green's function with $\mathbf{P}(\mathbf{r}, t)$ provides a solution to Maxwell's equations that innately embeds the radiation boundary condition. This obviates the need for absorbing boundary layers that would effectively behave as a source of blackbody thermal noise [27]. See [21,28,29] for more details about both time- and frequency-domain formulations.

3. Discrete solution

Solving Eqs. (4) and (6) self-consistently proceeds via the following steps: (i) represent the time varying behavior of the polarization (ii) using (6), evaluate $\tilde{\mathbf{E}}(\mathbf{r}, t)$ at a given time step, and (iii) use a predictor corrector approach to evaluate $\tilde{\rho}$ via (4). Representing $\mathbf{P}(\mathbf{r}, t)$ in terms of space and time basis functions such that

$$\tilde{\mathbf{P}}(\mathbf{r}, t) \approx \sum_{\ell=0}^{N_s-1} \sum_{m=0}^{N_t-1} \tilde{\mathcal{A}}_\ell^{(m)} \mathbf{s}_\ell(\mathbf{r}) T(t - m \Delta t), \quad (7)$$

$\tilde{\mathcal{A}}_\ell^{(m)} = \tilde{\rho}_{\ell,01}(m \Delta t)$ gives the polarization associated with the ℓ^{th} quantum dot at the m^{th} time step, and Δt denotes a fixed time interval chosen to accurately sample the dynamics of the physical quantities involved. Both $\mathbf{s}_\ell(\mathbf{r})$ and $T(t)$ have finite support and $T(t)$ obeys (discrete) causality (i.e. $T(t) = 0$ if $t < -\Delta t$). In particular, we consider shifted Lagrange polynomials for the $T(t)$ and assume dipolar transitions in the quantum dots allowing for $\mathbf{s}_\ell(\mathbf{r}) = \mathbf{d}_\ell \delta(\mathbf{r} - \mathbf{r}_\ell)$, though this analysis readily extends to accommodate any similar set of functions [31,32].

Substituting Eq. (7) into Eq. (5) and projecting the resulting fields onto $\delta(t - m \Delta t) \mathbf{s}_\ell(\mathbf{r})$ for all $\ell \in \{0, 1, \dots, N_s - 1\}$, we obtain

$$\tilde{\mathcal{E}}^{(m)} = \tilde{\mathcal{E}}_{\text{inc}}^{(m)} + \sum_{m'=0}^m \tilde{\mathcal{F}}^{(m-m')} \cdot \tilde{\mathcal{A}}^{(m')} \quad (8)$$

where

$$\tilde{\mathcal{E}}_\ell^{(m)} \equiv \langle \mathbf{s}_\ell(\mathbf{r}), \tilde{\mathbf{E}}(\mathbf{r}, m \Delta t) \rangle; \quad 0 \leq \ell < N_s \quad (9a)$$

$$\tilde{\mathcal{E}}_{\text{inc}, \ell}^{(m)} \equiv \langle \mathbf{s}_\ell(\mathbf{r}), \tilde{\mathbf{E}}_{\text{inc}}(\mathbf{r}, m \Delta t) \rangle; \quad 0 \leq \ell < N_s \quad (9b)$$

and $\tilde{\mathcal{F}}^{(k)}$ gives a sparse matrix of dimension $N_s \times N_s$ such that

$$\tilde{\mathcal{F}}_{\ell\ell'}^{(k)} \equiv \langle \mathbf{s}_\ell(\mathbf{r}), \tilde{\mathcal{F}}\{\mathbf{s}_{\ell'}(\mathbf{r}) T(k \Delta t)\} \rangle. \quad (9c)$$

(In the above equations, $\langle \cdot, \cdot \rangle$ denotes the inner product between functions.) Due to the finite support of the 3-D retarded potential, $\tilde{\mathcal{F}}^{(k)}$ has a sparse, lower-triangular, and banded structure. Note, Eq. (8) equivalently represents a discrete convolution.

A self-consistent solution to Eqs. (4) and (6) then has the following prescription for any time step: (i) determine $\tilde{\mathcal{A}}_\ell^{(m)} = \tilde{\rho}_{\ell,01}(m \Delta t)$ from the known history of the system, (ii) compute $\tilde{\mathcal{E}}_\ell^{(m)}$ using Eq. (8), (iii) find $\partial_t \tilde{\rho}_{\ell,01}(m \Delta t)$ using Eq. (4), and (iv) correct $\tilde{\rho}_{\ell,01}(m \Delta t)$ and iterate steps (ii) through (iv) until converged. The time complexity of the entire algorithm follows naturally from the above description; as discussed earlier, the cost of evaluating Eq. (8) scales as $\mathcal{O}(N_t N_s^2)$ while the cost of solving Eq. (4) for every quantum dot scales as $\mathcal{O}(N_t N_s)$. As a result, the bottleneck arises from the discrete convolution/field evaluation at every time step, and we address strategies to ameliorate this cost in the next section.

4. Acceleration via fast Fourier transforms

As a kernel-independent accelerator, TD-AIM forms the basis of our approach to reducing the computational complexity. Unfortunately, we cannot directly apply existing methodologies due to overhead induced by the multiplicity of terms in Eq. (5). In what follows, we develop a variation of TD-AIM that relies on propagating the convolution of the retarded potential with the source function and *local* evaluation of spatial and temporal derivatives.

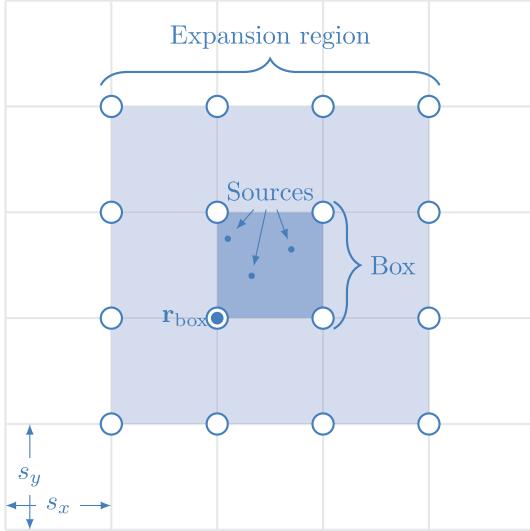


Fig. 1. Illustration of the grid structure and related terminology. All of the sources within a box (shown as the central shaded square) map to the same set of expansion points (shown as open circles) indexed relative to \mathbf{r}_{box} .

4.1. Algorithmic details

Here, we give algorithmic steps for accelerating the evaluation of the convolution in Eq. (8). Essentially, to effect a sub-quadratic calculation of Eq. (8), we approximate the discrete convolution with $\tilde{\mathcal{F}}^{(k)}$ as a sum of near- and far-field contributions. The near-field matrix elements follow directly from Eq. (9c): sources within a prescribed distance threshold interact “directly” so as to avoid incurring unreasonable approximation error between adjacent basis functions. Sources beyond this threshold, however, interact via auxiliary spatial basis functions that reside at the vertices of a regular Cartesian grid. These auxiliary sources recover $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$ at large distances (by design) and have two salient advantages: (i) they compress the interaction matrix by representing sources within the same spatial region in terms of the same auxiliary set (Fig. 1), and (ii) they impose a Toeplitz structure on the resulting interaction matrix that lends itself to efficient diagonalization through application of a multidimensional FFT. In devising a methodology to recover $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$ efficiently, we start with the observation that $\tilde{\mathcal{F}}\{\mathbf{P}(\mathbf{r}, t)\} = \tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}e^{i\omega t} = \mathcal{L}\{g(\mathbf{r}, t)\} * \mathbf{P}(\mathbf{r}, t)$. It follows that

$$\begin{aligned} \tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}e^{i\omega t} &= -\mu_0(\partial_t^2 \mathbf{I} - c^2 \nabla \nabla \cdot)g(\mathbf{r}, t) * \tilde{\mathbf{P}}(\mathbf{r}, t)e^{i\omega t} \\ &= -\mu_0(\partial_t^2 \mathbf{I} - c^2 \nabla \nabla \cdot) \int \tilde{\mathbf{P}}(\mathbf{r}', t_R) e^{i\omega t} \frac{e^{-i\omega |\mathbf{r} - \mathbf{r}'|/c}}{4\pi |\mathbf{r} - \mathbf{r}'|} d^3 \mathbf{r}' \quad (10) \\ &= -\mu_0(\partial_t^2 \mathbf{I} - c^2 \nabla \nabla \cdot) \tilde{\Phi}(\mathbf{r}, t) e^{i\omega t} \\ &\equiv \mathcal{M}\{\tilde{\Phi}(\mathbf{r}, t)\} e^{i\omega t}. \end{aligned}$$

Eq. (10) formulates $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$ in terms of spatiotemporal operators acting on a propagated field, $\tilde{\Phi}(\mathbf{r}, t)$. With no loss in generality, we assume that one can derive a discrete equation for the evaluation of $\tilde{\Phi}(\mathbf{r}, t)$ in the same manner as Eq. (8). Indeed, using Eq. (7), to evaluate

$$\tilde{\mathcal{K}}_{\ell}^{(m)} \equiv \langle \delta(t - m \Delta t) \mathbf{s}_{\ell}(\mathbf{r}), \tilde{\Phi}(\mathbf{r}, t) \rangle; \quad 0 \leq \ell < N_s \quad (11)$$

yields

$$\tilde{\mathcal{K}}^{(m)} = \sum_{m'=0}^m \tilde{\mathcal{Z}}^{(m-m')} \cdot \tilde{\mathcal{A}}^{(m')}, \quad (12)$$

where

$$\tilde{\mathcal{Z}}_{\ell\ell'}^{(k)} \equiv \frac{1}{4\pi |\mathbf{r}_{\ell} - \mathbf{r}_{\ell'}|} \exp\left(-i \frac{\omega |\mathbf{r}_{\ell} - \mathbf{r}_{\ell'}|}{c} T(k \Delta t)\right). \quad (13)$$

Our prescription for acceleration of the evaluation of Eq. (8) via FFT embeds the computational domain in a spatial grid of uniformly-spaced points along each Cartesian direction. We denote the grid spacing Δs ; in what follows, we refer to a domain of size Δs as a box. As we will demonstrate, three parameters determine the accuracy of the approximation: the grid spacing, Δs , the expansion order of the boxes, M , and a nearfield parameter, γ . The expansion order M determines the number of equivalent grid points used to represent each $\mathbf{s}_{\ell}(\mathbf{r})$ as $(M + 1)^3$, and the nearfield parameter controls the minimum grid distance beyond which interactions proceed via FFTs (Fig. 2); sources within this distance interact directly via Eq. (9c).

We effect the discrete convolution in Eq. (8) with the following strategy: Every box in the computational domain contains a number of adjacent neighbors; we label the interactions between the quantum dots in these domains as “direct” and compute them using a subset of $\tilde{\mathcal{F}}^{(k)}$ with corrections. We compute the remainder of the interactions – so-called “far” interactions – via TD-AIM. As noted earlier, the number of terms in Eq. (5) makes using conventional TD-AIM unwieldy for computing far interactions. We take a different route: we evaluate $\tilde{\mathcal{K}}^{(m)}$ (and its equivalent for time derivatives of $\tilde{\Phi}$) via FFT and then compose $\tilde{\mathcal{E}}^{(m)}$ via the operator $\mathcal{M}\{\cdot\}$.

Specifically, we can write

$$\tilde{\mathcal{K}}_{\text{FFT}}^{(m)} = \sum_{m'=0}^m \tilde{\mathcal{Z}}_{\text{FFT}}^{(m-m')} \cdot \tilde{\mathcal{A}}^{(m')}, \quad (14)$$

where each element of $\tilde{\mathcal{Z}}_{\text{FFT}}^{(m)}$ arises as a product of three matrices: one that effects a mapping from a source point to the uniform grid points, one to evaluate Eq. (13) (with \mathbf{r}_{ℓ} and $\mathbf{r}_{\ell'}$ denoting grid points), and finally one that projects these potentials back to the observation point. As such, at every time step, our algorithm proceeds as follows:

- 1. Projection on to the uniform grid:** At time step m , project each of the $\mathbf{s}_{\ell}(\mathbf{r}) \tilde{\mathcal{A}}_{\ell}^{(m)}$ onto the auxiliary sources that reside on a set of grid points. We denote the distribution of auxiliary sources on the grid by $\tilde{\mathbf{P}}_{\text{aux}}(\mathbf{u}, t)$. The cost of this operation scales linearly with the number of dots in the system.
- 2. Effect the convolution in Eq. (8) between auxiliary sources:** Having imposed a regular structure on $\tilde{\mathbf{P}}_{\text{aux}}(\mathbf{u}, t)$, we may efficiently diagonalize $\tilde{\mathcal{Z}}$ – the interactions between grid points – with (up to four-dimensional) blocked FFTs. In effect this evaluates the discrete convolution $\tilde{\mathcal{K}}_{\text{FFT}}^{(m)}$ at every point \mathbf{u} in the grid. (Note, to properly recover $\tilde{\mathcal{E}}^{(m)}$ we must also evaluate $\partial_t^i \tilde{\mathcal{K}}^{(m)}$ for $i \in \{0, 1, 2\}$. We trivially accomplish this by using $\partial_t^i T(k \Delta t)$ in Eq. (13).) The cost of this operation scales as $\mathcal{O}(N_t N_s \log^2 N_s)$.
- 3. Projection back from the grid:** We recover the total field at $\mathbf{s}_{\ell}(\mathbf{r})$ by projecting the $\tilde{\mathcal{K}}_{\text{FFT}}^{(m)}$ back onto the $\mathbf{s}_{\ell}(\mathbf{r})$ under the action of $\mathcal{M}\{\cdot\}$. These projection matrices account for zeroth- and second-order spatial derivatives in the operator. As before, the cost of this operation scales linearly with the number of quantum dots.
- 4. Correction of near fields:** $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$ evaluated through the grid remains accurate only at large (many Δs) distances. To correct this, we subtract the contribution to the fields from close sources and replace them with Eq. (9c) (Fig. 3 gives a schematic illustration of this correction.) Finally, this cost scales linearly with the number of quantum dots.

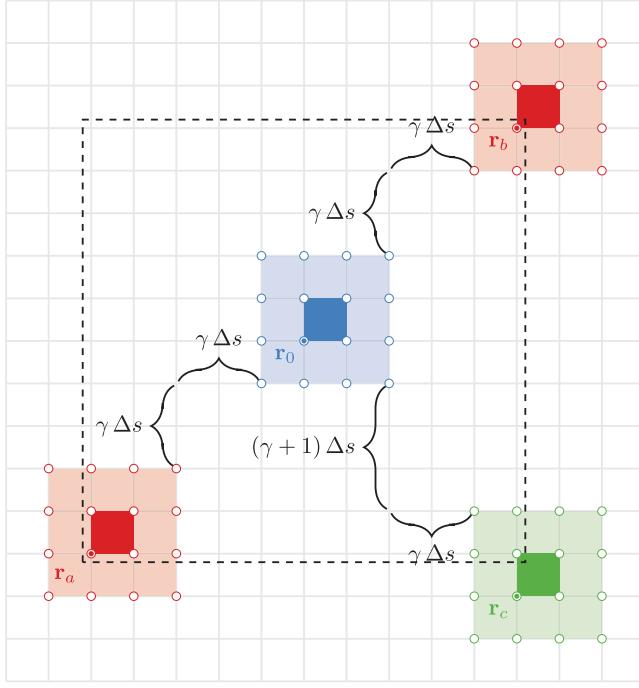


Fig. 2. Illustration of the nearfield criterion for a third order expansion $M = 3$ corresponding to $(3+1)^2 = 16$ grid points in this 2D illustration, and $\gamma = 2$. The dashed line indicates the complete nearfield of the box associated with \mathbf{r}_0 —i.e. all boxes that have an expansion point within $\gamma \Delta s$ (infinity norm) of the expansion around \mathbf{r}_0 . Consequently, all of the $\mathbf{s}_\ell(\mathbf{r})$ within the central dark square have a pairwise interaction with the $\mathbf{s}_{\ell'}(\mathbf{r})$ inside the dashed box.

Next, we describe the construction of auxiliary matrices – Λ^\dagger that maps from sources to the auxiliary grid and $\Lambda_{\mathfrak{M}}$ that maps from the auxiliary grid back to sources (under the action of \mathfrak{M}) – that play a critical role in steps 1 and 3 above.

4.1.1. Auxiliary matrices

The construction of projection matrices Λ^\dagger and $\Lambda_{\mathfrak{F}}$ critically underpins the above process. These operators map fields onto the uniform grid and back, though the operator $\Lambda_{\mathfrak{F}}$ differs slightly from Λ^\dagger as it accounts for all the spatial derivatives contained within \mathfrak{M} . To start, we represent the primary $\mathbf{s}_\ell(\mathbf{r})$ basis functions as a weighted sum of δ -functions on the surrounding grid points using

$$\psi_\ell(\mathbf{r}) \approx \sum_{\mathbf{u} \in C_\ell} \Lambda_{\ell\mathbf{u}}^\dagger \delta(\mathbf{r} - \mathbf{u}). \quad (15)$$

Here, $\psi_\ell(\mathbf{r}) \in \{\mathbf{s}_\ell(\mathbf{r}) \cdot \hat{\mathbf{x}}, \mathbf{s}_\ell(\mathbf{r}) \cdot \hat{\mathbf{y}}, \mathbf{s}_\ell(\mathbf{r}) \cdot \hat{\mathbf{z}}\}$ and C_ℓ denotes the collection of grid points within the expansion region of $\mathbf{s}_\ell(\mathbf{r})$ (Fig. 1). For an expansion of order M , this sum contains $(M+1)^3$ terms corresponding to the $(M+1)^3$ grid points nearest to $\mathbf{s}_\ell(\mathbf{r})$. Consequently, the $\Lambda_{\ell\mathbf{u}}^\dagger$ matrices contain few nonzero elements and we have elected to use a moment-matching scheme to capture the $(M+1)^3$ multipole moments of $\mathbf{s}_\ell(\mathbf{r})$ according to

$$\int (x-x_0)^{m_x} (y-y_0)^{m_y} (z-z_0)^{m_z} \left[\psi_\ell(\mathbf{r}) - \sum_{\mathbf{u} \in C_\ell} \Lambda_{\ell\mathbf{u}}^\dagger \delta(\mathbf{r} - \mathbf{u}) \right] d^3\mathbf{r} = 0. \quad (16)$$

In this expression, $0 \leq m_x, m_y, m_z \leq M$ and $\mathbf{r}_0 \equiv x_0\hat{\mathbf{x}} + y_0\hat{\mathbf{y}} + z_0\hat{\mathbf{z}}$ denotes the origin about which we compute the multipoles. To determine the $\Lambda_{\ell\mathbf{u}}^\dagger$, we solve the least-squares system

$$\sum_{\mathbf{u} \in C_\ell} W_{\mathbf{m}\mathbf{u}} \Lambda_{\ell\mathbf{u}}^\dagger = Q_{\ell\mathbf{m}} \quad (17)$$

where

$$W_{\mathbf{m}\mathbf{u}} = (u_x - x_0)^{m_x} (u_y - y_0)^{m_y} (u_z - z_0)^{m_z} \quad (18a)$$

$$Q_{\ell\mathbf{m}} = \int \psi_\ell(\mathbf{r}) (x - x_0)^{m_x} (y - y_0)^{m_y} (z - z_0)^{m_z} d^3\mathbf{r}, \quad (18b)$$

$\mathbf{u} \in C_\ell$, and \mathbf{m} denotes the multi-index $\mathbf{m} = \{m_x, m_y, m_z\}$. With an infinite precision calculation, the choice of $\mathbf{r}_0 = x_0\hat{\mathbf{x}} + y_0\hat{\mathbf{y}} + z_0\hat{\mathbf{z}}$ merely defines an origin for the polynomial expansion system. To minimize numerical issues, we choose \mathbf{r}_0 at the center of $\mathbf{s}_\ell(\mathbf{r})$ and we employ low-order projection/Vandermonde matrices as done in [22,33,34].

4.2. Details and relevant nuances

Analogous to

$$\tilde{\mathcal{K}}_{\text{FFT}}^{(m)} = \sum_{m'=0}^m \left[\Lambda \tilde{\mathcal{Z}}^{(m-m')} \Lambda^\dagger \right] \cdot \tilde{\mathcal{A}}^{(m)}, \quad (19)$$

it follows from the above that one may reconstruct

$$\tilde{\mathcal{E}}_{\text{FFT}}^{(m)} = \sum_{m'=0}^m \left[\Lambda_{\mathfrak{M}} \left(\begin{array}{c} \partial_t^0 \tilde{\mathcal{Z}}^{(m-m')} \\ \partial_t^1 \tilde{\mathcal{Z}}^{(m-m')} \\ \partial_t^2 \tilde{\mathcal{Z}}^{(m-m')} \end{array} \right) \Lambda^\dagger \right] \cdot \tilde{\mathcal{A}}^{(m)}. \quad (20)$$

In this equation, $\tilde{\mathcal{Z}}$ captures interactions between points in the auxiliary grid. One can go on to embed this within blocked space-time FFTs [34] to obtain advantageous algorithmic scaling. However, for brevity we restrict ourselves to simpler spatial-only acceleration. In the construction of $\tilde{\mathcal{E}}^{(m)}$ from $\tilde{\Phi}(\mathbf{r}, t)$ (sampled on the grid) under \mathfrak{M} , we require time derivatives of $\tilde{\Phi}(\mathbf{r}, t)$ which we accomplish by differentiating $T(t)$ in Eq. (13). In addition to time derivatives, \mathfrak{M} requires spatial derivatives (taken with respect to the observation point). As Λ effects a polynomial reconstruction of fields at any observation point, one can trivially compute spatial derivatives as long as the order of this reconstruction exceeds the number of derivatives required. As a result, differentiating Eq. (16) removes the high-order moments in Eq. (18b). Accordingly, $\Lambda_{\mathfrak{M}}$ captures these derivatives so as to correctly reproduce the behavior of $\mathfrak{M}\{\cdot\}$. It follows that

$$\tilde{\mathcal{E}}^{(m)} \approx \tilde{\mathcal{E}}_{\text{inc}}^{(m)} + \tilde{\mathcal{E}}_{\text{direct}}^{(m)} + \tilde{\mathcal{E}}_{\text{FFT}}^{(m)} \quad (21)$$

where

$$\tilde{\mathcal{E}}_{\text{direct}}^{(m)} = \begin{cases} \sum_{m'=0}^m \tilde{\mathcal{F}}_{\text{direct}}^{(m-m')} \cdot \tilde{\mathcal{A}}^{(m)} - \tilde{\mathcal{E}}_{\text{FFT}}^{(m)} & R_{\ell\ell'} \leq \gamma \\ 0 & \text{otherwise,} \end{cases} \quad (22)$$

$$\tilde{\mathcal{F}}_{\text{direct}, \ell\ell'}^{(m)} = \begin{cases} \mathcal{F}_{\ell\ell'}^{(m)} & R_{\ell\ell'} \leq \gamma \\ 0 & \text{otherwise.} \end{cases}$$

Here, γ serves as adjustable input parameters to control the accuracy of the simulation and $R_{\ell\ell'}$ gives the minimum distance (in integral units of the grid spacing) between the expansion regions enclosing $\mathbf{s}_\ell(\mathbf{r})$ and $\mathbf{s}_{\ell'}(\mathbf{r})$ (Fig. 2) via

$$R_{\ell\ell'}^{\text{grid}} = \min\{\|\mathbf{u} - \mathbf{u}'\|_\infty \mid \mathbf{u} \in C_\ell, \mathbf{u}' \in C_{\ell'}\}. \quad (23)$$

4.3. Convergence analysis

Next, we present a succinct analysis of convergence. While such analyses arise in different contexts [22], the analysis herein approaches it from an interpolation perspective and enables one to obtain an error bound on the overall operator. With no loss of generality, consider two point particles located at x_{src} and x_{obs} . A time-independent Green's function, $g(x_{\text{obs}} - x_{\text{src}})$, describes the interaction between the two particles and we wish to construct a polynomial approximation of $g(x - x_{\text{src}})$ for x in the vicinity of x_{obs} as in Fig. 4.

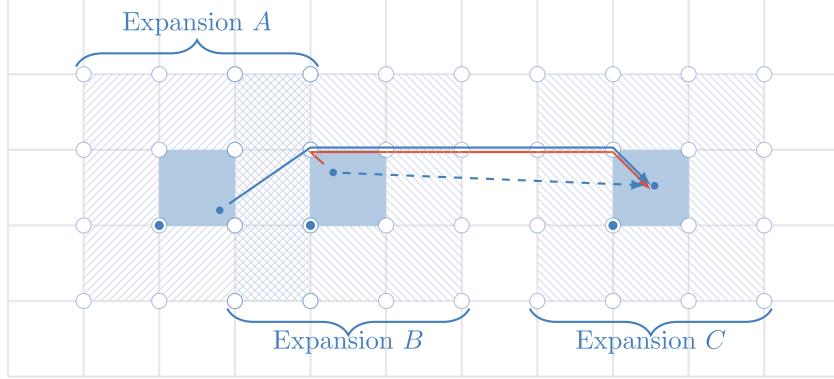


Fig. 3. Illustration of nearfield corrections between close boxes. The expansions of boxes A and B overlap, but only box B lies in the nearfield of box C for $\gamma = 2$. As the grid-based propagation strategy only remains accurate for distant source/observer pairs, we remove the interaction “through the grid” between the BC pair (red line) and replace it with a more accurate “direct” interaction (dashed blue line). The AC pair requires no such treatment as they have well-separated expansion regions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

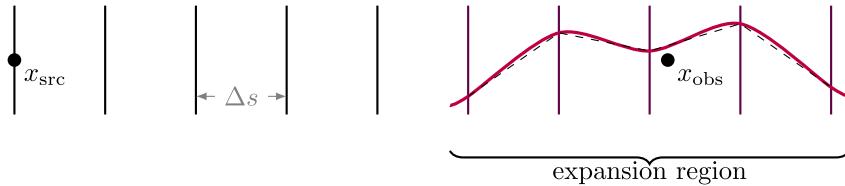


Fig. 4. Polynomial interpolation of $g(x - x_{\text{src}})$ near x_{obs} . Here, the red curve represents the actual $g(x - x_{\text{src}})$ and the dashed black line its approximation. Evaluating the m th-order approximation requires samples of the signal at $m + 1$ grid points surrounding x_{obs} .

To construct an interpolation polynomial over the expansion region of order M , we define a polynomial coordinate $x_p = (x - x_0)/\Delta s$ such that $x_p^{\min} \leq x_p \leq x_p^{\min} + M$ where $x_p^{\min} \equiv -\lfloor M/2 \rfloor$. Consequently, the expansion points about x_{obs} correspond to $x_p \in \{-\lfloor M/2 \rfloor, -\lfloor M/2 \rfloor + 1, -\lfloor M/2 \rfloor + 2, \dots\}$ with the 0th order expansion point, x_0 , equivalent to $x_p = 0$. Such a coordinate system defines the Vandermonde’s linear equation $\sum_j V_{ij} w_j = g_i$ for the weights of an interpolating polynomial where

$$V_{ij} = (x_p^{\min} + i)^j \quad (24a)$$

$$g_i = g((x_0 - x_{\text{src}}) + (x_p^{\min} + i) \Delta s) \quad (24b)$$

and $0 \leq i, j \leq M$. Approximating $g(x - x_{\text{src}})$ at x_{obs} then becomes a matter of evaluating this polynomial at $x_p = (x_{\text{obs}} - x_0)/\Delta s$, i.e.

$$\begin{aligned} g(x_{\text{obs}} - x_{\text{src}}) &= g\left((x_0 - x_{\text{src}}) + \left(\frac{x_{\text{obs}} - x_0}{\Delta s}\right) \Delta s\right) \\ &\approx \sum_{i=0}^M w_i \left(\frac{x_{\text{obs}} - x_0}{\Delta s}\right)^i. \end{aligned} \quad (25)$$

Accordingly, the polynomial approximation to $g(x_{\text{obs}} - x_{\text{src}})$ contains terms of order $\mathcal{O}(\Delta s^{-M})$ and we can expect the approximation error to scale as $\mathcal{O}(\Delta s^{-(M+1)})$. Moreover, this motivates using the approximation to calculate interactions involving differential operators; applying an n th-order derivative reduces the polynomial order by n , thus the error scales like $\mathcal{O}(\Delta s^{-(M+1)+n})$. The preceding analysis generalizes to three dimensions.

5. Numerical results

Next, we present a number of results using the methodologies developed thus far. We seek to demonstrate controllable accuracy of the proposed scheme, the cost complexity in both computation time and memory, and finally some exemplar simulations of quantum dot systems.

5.1. Accuracy

To start, we examine error incurred in our approach in evaluating the space-time convolution in (6). To isolate the errors incurred, our experiment proceeds as follows. We set up two domains with sufficient separation such that the interactions between these occur only via AIM. Each domain contains 64 randomly distributed quantum dots, we prescribe the temporal variation of the polarization of each quantum dot, and we measure the total radiated field at each quantum dot. Finally, we fix the temporal interpolation basis order at 3 and the polarization of each quantum dot varies as

$$P(t) = e^{-\frac{(t-t_0)^2}{2\sigma^2}}. \quad (26)$$

The simulation runs for 1024 time steps of size $\Delta t = 0.1$ ps, the width of the Gaussian $\sigma = 1024 \Delta t / 12$ and its center $t_0 = 1024 \Delta t / 2$. This approach admits a readily available analytic solution via Eq. (3) which we measure against the AIM solution. For this, we calculate the ℓ_2 norm differences between the two solutions as a function of AIM grid size for different expansion orders to validate the error behavior described in Section 4. Fig. 6 gives geometric parameters and results; as shown by the figure, we observe excellent convergence.

Next, we examine errors incurred when conducting a similar experiment in the rotating frame. All quantum dots begin in the ground state $(\rho_{00}, \rho_{01})|_{t=0} = (1, 0)$, and their density matrix elements evolve according to Eq. (4). The dipole moment of each quantum dot aligns with the laser field, given by

$$\tilde{\mathbf{E}}(\mathbf{r}, t) = \tilde{E}_0 \hat{\mathbf{x}} e^{-\frac{(\mathbf{k}\mathbf{r} - \omega(t-t_0))^2}{2\sigma^2}}. \quad (27)$$

We use a third order expansion with AIM spacing $\Delta s = 5 \times 10^{-3} \lambda$, $\gamma = 1$, and 1000 time steps of size $\Delta t = 0.01$ ps (Table 1 gives additional simulation parameters.) As before, we compare results from AIM (Fig. 7) to those obtained using the direct method, as it permits us to normalize against the error in using temporal basis sets.

Table 1

Dynamic simulation parameters; e and a_0 denote the elementary charge and Bohr radius. The decoherence times here, while shorter than those typical of optical resonance experiments, afford a shorter computational time but preserve dynamical emission phenomena.

Quantity	Symbol	Value
Speed of light	c	$300 \mu\text{m ps}^{-1}$
Transition frequency	ω_0	$1500 \text{ meV}/\hbar$
Transition dipole moment (magnitude)	$ d $	$10 ea_0$
Decoherence times	T_1, T_2	$10 \text{ ps}, 20 \text{ ps}$
Laser frequency	ω	$1500 \text{ meV}/\hbar$
Laser wavevector	$ \mathbf{k} $	$7.6016 \mu\text{m}^{-1}$
Laser wavelength	λ	827 nm
Laser peak shift	t_0	5 ps
Pulse width	σ/ω	1 ps
Pulse area	-	π

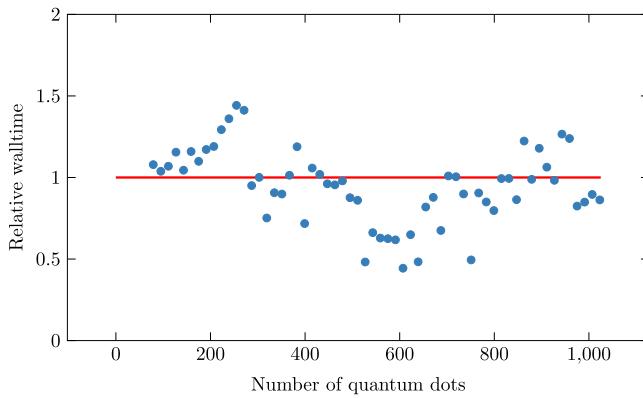


Fig. 5. Simulation time of $g(\mathbf{r}, t) * \tilde{\mathbf{P}}(\mathbf{r}, t)$ relative to $\mathcal{L}\{g(\mathbf{r}, t)\} * \tilde{\mathbf{P}}(\mathbf{r}, t)$ (i.e. $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$) for various system sizes. This indicates that evaluating fields (with multiple derivatives) has almost no overhead relative to evaluating simple retarded potentials.

5.2. Cost of evaluation of higher order spatial derivatives

Fig. 5 shows wall time results for the calculation of $g(\mathbf{r}, t) * \tilde{\mathbf{P}}(\mathbf{r}, t)$ relative to $\tilde{\mathcal{F}}\{\tilde{\mathbf{P}}(\mathbf{r}, t)\}$ – i.e. a “simple” scalar propagator that evaluates a potential relative to a complex one involving dyadics and derivatives (alternatively, evaluating the field) – for various system sizes. Each experiment uses the same configuration of sources (arranged linearly at consistent density) and system parameters and we time only the time stepping procedure assuming pre-filled matrices. We attribute the correlated variation in **Fig. 5** to AIM – the efficiency of the grid-based acceleration scheme accurately depends on the geometry/density of sources – though we note both propagators appear to take roughly the same amount of computational effort to evaluate. This indicates that our modified TD-AIM formulation can accommodate *any* propagation kernel involving arbitrary spatiotemporal derivatives with little-to-no additional computational overhead.

5.3. Complexity

Next, we present a set of experiments that demonstrate the $O(N_s \log(N_s))$ complexity scaling of AIM. For this, we perform simulations in both the fixed frame with prescribed polarizations, and the rotating frame with full Liouville equation dynamics. To ensure proper examination of computational complexity, we start with a box of side length $6 \Delta s$ (chosen to minimize the number of nearfield pairs), and filled with quantum dots at random locations. We obtain each successive value of N_s by doubling the side length and in effect, increasing the number of quantum dots by a factor of eight. We use a third order expansion $M = 3$ with

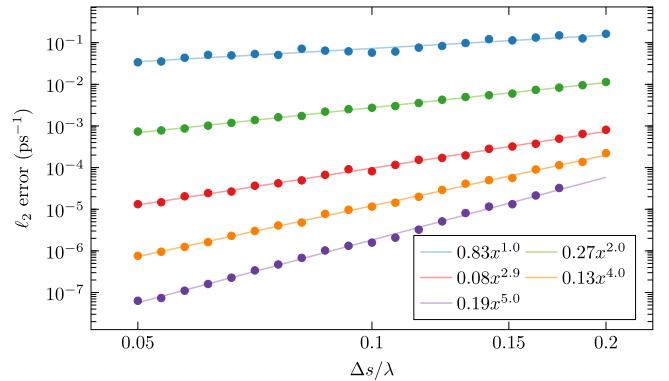


Fig. 6. l_2 error of the Rabi frequency magnitude $|\chi|$ with respect to grid spacing for expansion orders M 2 through 6 and $\gamma = 1$, using source and observer boxes of volume λ^3 separated by $\Delta r = 2\lambda(\hat{\mathbf{x}} + \hat{\mathbf{y}} + \hat{\mathbf{z}})$, each containing 64 randomly generated quantum dots. For an expansion order M one expects the overall error to scale as $O(\Delta s^{M-1})$, consistent with the results above.

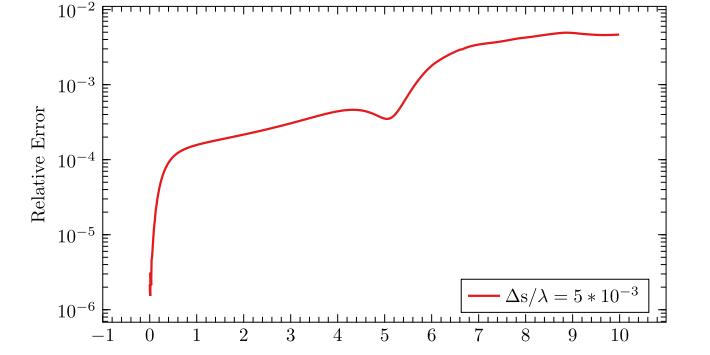
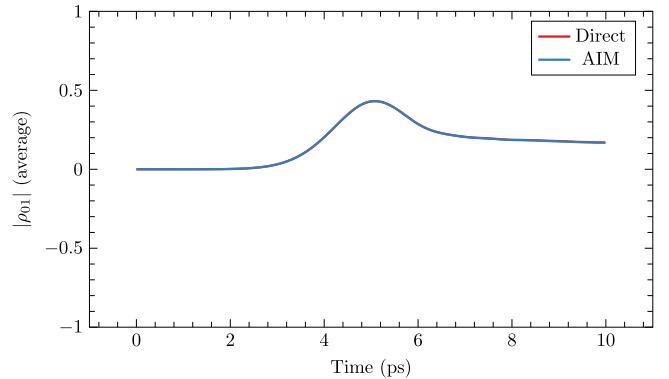


Fig. 7. (Top) $|\tilde{\rho}_{01}|$ for a system of 128 quantum dots in a cube of length $\lambda/10$ (chosen to promote observable coupling effects through very close quantum dots), computed using the direct algorithm and AIM. (Bottom) l_2 relative error of AIM algorithm against the direct algorithm for this simulation.

AIM spacings $\Delta s = \lambda/400$ and $\Delta s = \lambda/10$ for the fixed and rotating wave cases, respectively. Time steps mirror those used in Section 5.1. **Fig. 8** gives runtimes for both cases, demonstrating that the two FFT-accelerated simulations outpace their direct counterparts near $N_s = 1000$ and $N_s = 2000$, respectively.

5.4. Large scale physical simulations

The largest system simulated in [7] without TD-AIM consists of 10 000 quantum dots randomly distributed in a cylinder of radius $0.2 \mu\text{m}$ and length $4 \mu\text{m}$. **Fig. 9** shows an equivalent simulation

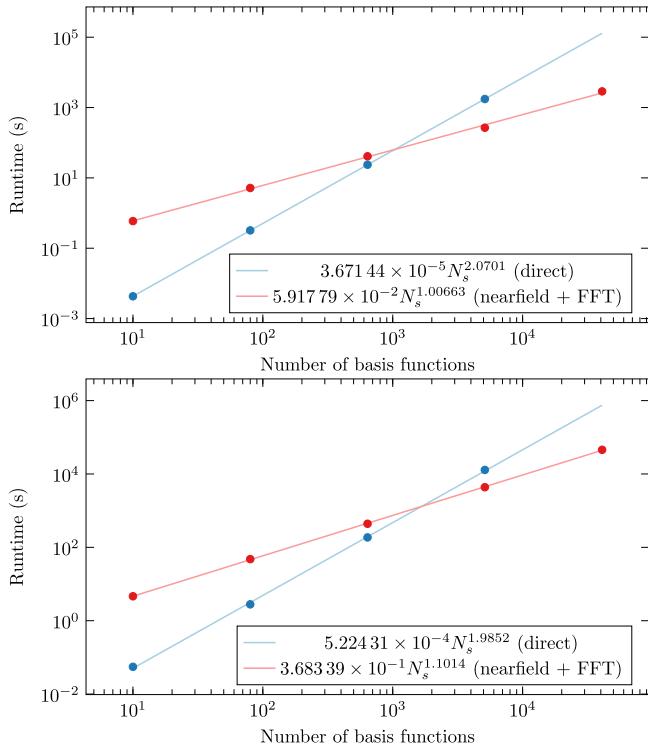


Fig. 8. FFT runtime (excluding setup time) using a third-order expansion. (Top) 1024 time steps with $\Delta s = \lambda/400$ and prescribed polarizations in the fixed frame. (Bottom) 1000 time steps with $\Delta s = \lambda/10$ and Liouville-dynamics polarization in the rotating frame. Both cases have a quasi-quadratic scaling in the direct calculation, whereas the FFT-accelerated calculation performs slightly worse than linear.

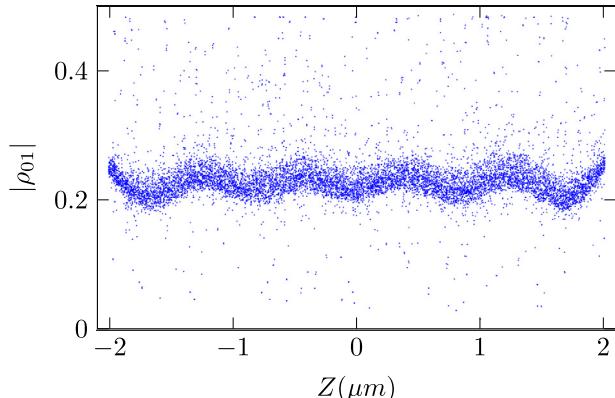


Fig. 9. $\hat{\mathbf{z}}$ -distribution of polarization $|\tilde{\rho}_{01}|$ for a 10 000-dot cylindrical simulation, replicating the parameters in [7]. The AIM calculation recovers the oscillatory long-range pattern that we obtained using a direct calculation [7].

with TD-AIM that reproduces features arising from quantum dot interactions; we conduct a very similar experiment here. The figure shows the polarization of each quantum dot in the cylinder as a function of their \mathbf{z} -coordinate (the axis of the cylinder), under the effect of a resonant π pulse. Each of the quantum dots has an identical (fixed) dipole moment (see [7] for the details of the simulation parameters). Note how the secondary radiation produces random shifts in the polarization due to short-range effects in the local neighborhood of each quantum dot. In addition, the simulation shows an oscillation of the polarization due to long-range collective effects. This oscillation reflects the role

Table 2
AIM parameters for the simulation of Section 5.4.

Quantity	Symbol	Value
Simulation time step	Δt	0.02 ps
AIM spacing	Δs	$0.040\lambda = 33.06 \text{ nm}$
AIM expansion order	M	5
Nearfield border parameter	γ	1
Transverse domain length	-	$16 \Delta s = 529 \text{ nm}$
Longitudinal domain length	-	$1500 \Delta s = 49.59 \mu\text{m}$

of boundary conditions in the confinement of the macroscopic electric field in the system.

The algorithm introduced in this paper facilitates simulations of much larger systems. In Figs. 10 and 11 we examine the response of a system of 100 000 quantum dots – randomly distributed throughout a cuboid – to an applied laser pulse traveling along $\hat{\mathbf{z}}$. The transition dipole moment of each quantum dot has a fixed magnitude but random orientation. Tables 1 and 2 list simulation parameters.

Fig. 10 displays a color map of $|\tilde{\rho}_{01}|$ as an indicator of the polarization $|\mathbf{P}|$ of each quantum dot at different time steps after the pulse peak. The figure shows only quantum dots located in a central segment of about $4 \mu\text{m}$ of the entire cuboid. The random orientation of the dipole moments creates a variation in the amplitude of the polarization with quantum dots whose dipole moments (anti-)align with the laser field having greatest amplitude. In addition, despite each quantum dot resonantly coupling to the pulse, inhomogeneity arises due to the inter-dot coupling. These simulations can resolve inhomogeneities at the microscopic level, taking into account the orientation of the transition dipole moment of each quantum dot, as well as the effect of local secondary fields.

To visualize long-range effects, Fig. 11 shows $|\tilde{\rho}_{01}|$ as a function of the \mathbf{z} coordinate of each quantum dot, corresponding to the color plots of Fig. 10. Here we show the entire cuboid having sides of $20 \mu\text{m}$. In contrast to the results of Fig. 9, we do not observe the oscillatory behavior due to confinement since the length of the system far exceeds the radiation wavelength. Moreover, we observe a dispersion of the polarization due to the random orientation of the transition dipoles. Since the strength of the coupling scales with $\mathbf{E} \cdot \mathbf{d} = \cos(\theta)$, the distribution peaks at the value of $|\tilde{\rho}_{01}|$ when $\theta = 0$ or $\theta = \pi$, with a tail corresponding to all the intermediate values. Only a few quantum dots, for which the secondary fields constructively interfere, have a polarization larger than the peak value. Finally, note how the value of the peak polarization slightly increases from left to right due to pulse propagation.

Furthermore, we calculate the inverse participation ratio (IPR) of the quantum dot polarization:

$$\text{IPR}(t) = \frac{\sum_{\ell} |\tilde{\rho}_{01}(t)|^4}{(\sum_{\ell} |\tilde{\rho}_{01}(t)|^2)^2} \quad (28)$$

with results shown in Fig. 12. This quantity ranges from $1/N$ to 1, and collectively measures the spatial localization of the polarization, with $1/N$ corresponding to a completely delocalized spatial distribution, and 1 to the case of the polarization completely localized on a single site. For comparison, we also include an IPR plot for the case of uniform (pulse-aligned) dipoles. In the uniform case, all quantum dots participate equally until the onset of the pulse peak, whereupon inter-dot coupling leads certain quantum dots to retain their polarization longer than neighbors. This contrasts the non-uniform case, which exhibits localization of polarization to quantum dots that align with the laser pulse.

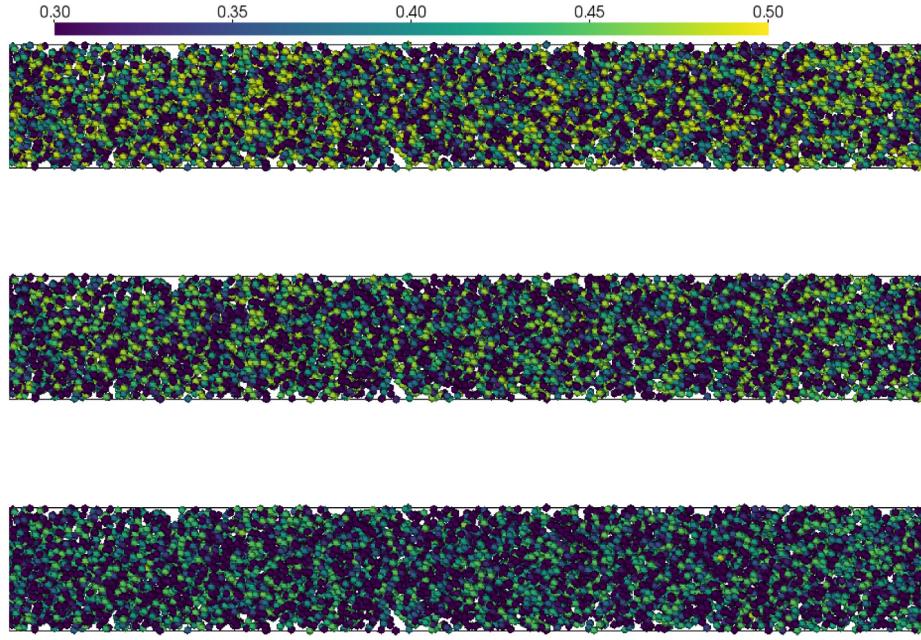


Fig. 10. Coloration of $|\tilde{\rho}_{01}|$ as an indicator of $|\tilde{P}|$ at $t_1 = 6.0$ ps (top), $t_2 = 7.0$ ps (middle), $t_3 = 8.0$ ps (bottom) relative to the peak of a 1 ps-wide pulse, for a system of 100 000 quantum dots.

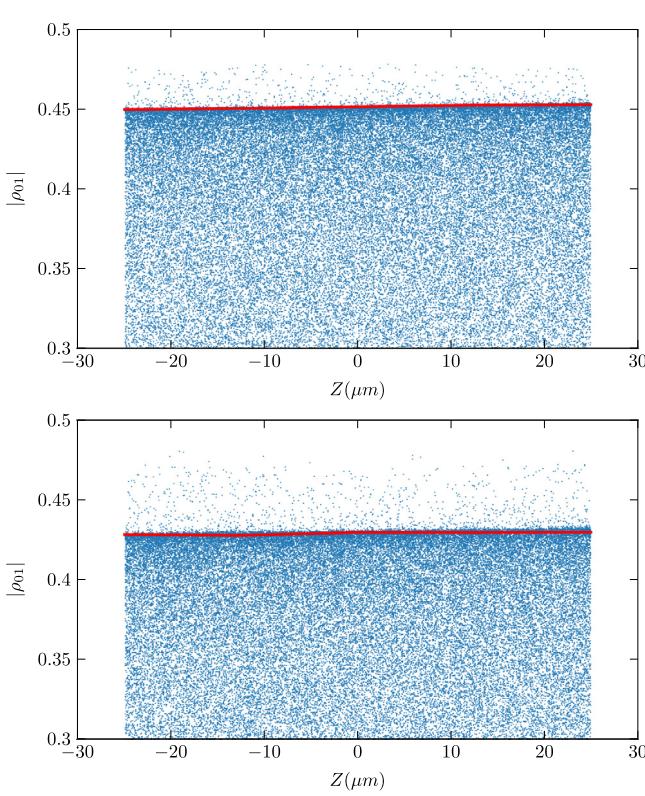


Fig. 11. Scatterplots of $|\tilde{\rho}_{01}|$ corresponding to the bottom two color plots of Fig. 10. There exists a single preferred polarization, represented by the linear region of greatest density, arising from quantum dots whose transition dipole moments (anti-)align with the laser field. Radiative coupling produces polarizations that exceed this value. The changes in the red line reflect pulse propagation.

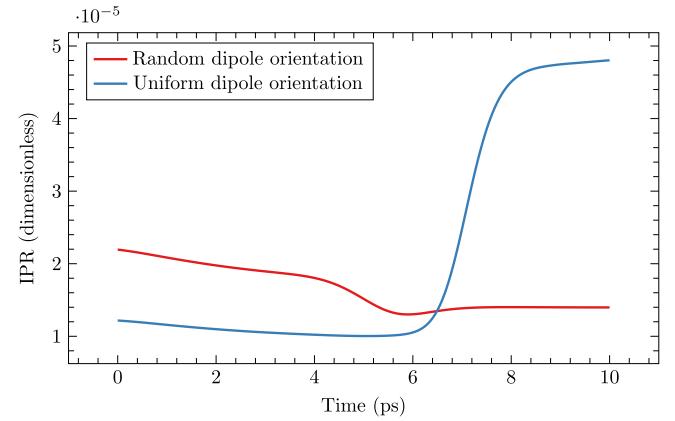


Fig. 12. Inverse Participation Ratio (IPR) for the system of Fig. 10 (red) and a similar system of 100 000 quantum dots with uniform dipole orientations (blue). The IPR for the uniform dipole orientation dominates at later times and approaches $5 \cdot 10^{-5}$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6. Conclusions

Here we have presented novel variations to TD-AIM that enables analysis of large ensembles of quantum dots, and we provide an implementation of this algorithm at [35]. We have discussed numerous features of the approach, including accuracy, convergence, and complexity. The latter for prescribed and fixed polarization, as well as when the polarization evolves. Additionally, we have validated the approach against “direct” simulations that use no acceleration techniques, demonstrating a reduction in overall complexity from $\mathcal{O}(N_t N_s^2)$ to $\mathcal{O}(N_t N_s \log N_s)$. Finally, we have used the accelerated approach to simulate a system with 100 000 quantum dots. We observe results identical to direct solutions, thus these techniques can reliably simulate much larger

systems. The next phase of our research focuses on additional capabilities to both the underlying physics model as well as the computational infrastructure, specifically, parallelization along the lines of those for both particle-particle, particle-mesh [33] and AIM [34].

Declaration of competing interest

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

We gratefully acknowledge the HPCC facility at Michigan State University for their support of this work. We also acknowledge support from the National Science Foundation under grants NSF ECCS-1408115 and OAC-1835267.

References

- [1] T.H. Stievater, X. Li, D.G. Steel, D. Gammon, D.S. Katzer, D. Park, C. Piermarocchi, L.J. Sham, Phys. Rev. Lett. 87 (2001) 133603, <http://dx.doi.org/10.1103/PhysRevLett.87.133603>, URL <http://link.aps.org/doi/10.1103/PhysRevLett.87.133603>.
- [2] H. Kamada, H. Gotoh, J. Temmyo, T. Takagahara, H. Ando, Phys. Rev. Lett. 87 (2001) 246401, <http://dx.doi.org/10.1103/PhysRevLett.87.246401>, URL <https://link.aps.org/doi/10.1103/PhysRevLett.87.246401>.
- [3] H. Htoon, T. Takagahara, D. Kulik, O. Baklenov, A.L. Holmes, C.K. Shih, Phys. Rev. Lett. 88 (2002) 087401, <http://dx.doi.org/10.1103/PhysRevLett.88.087401>, URL <https://link.aps.org/doi/10.1103/PhysRevLett.88.087401>.
- [4] G. Rainò, M.A. Becker, M.I. Bodnarchuk, R.F. Mahrt, M.V. Kovalenko, T. Stöferle, Nature 563 (7733) (2018) 671–675, <http://dx.doi.org/10.1038/s41586-018-0683-0>.
- [5] G.Y. Slepyan, S.A. Maksimenko, A. Hoffmann, D. Bimberg, Phys. Rev. A 66 (2002) 063804, <http://dx.doi.org/10.1103/PhysRevA.66.063804>, URL <http://link.aps.org/doi/10.1103/PhysRevA.66.063804>.
- [6] G.Y. Slepyan, A. Magyarov, S.A. Maksimenko, A. Hoffmann, D. Bimberg, Phys. Rev. B 70 (2004) 045320, <http://dx.doi.org/10.1103/PhysRevB.70.045320>, URL <http://link.aps.org/doi/10.1103/PhysRevB.70.045320>.
- [7] C. Glosser, B. Shanker, C. Piermarocchi, Phys. Rev. A 96 (3) (2017) 033816, <http://dx.doi.org/10.1103/PhysRevA.96.033816>, URL <https://link.aps.org/doi/10.1103/PhysRevA.96.033816>.
- [8] K. Asakura, Y. Mitsumori, H. Kosaka, K. Edamatsu, K. Akahane, N. Yamamoto, M. Sasaki, N. Ohtani, Phys. Rev. B 87 (2013) 241301, <http://dx.doi.org/10.1103/PhysRevB.87.241301>, URL <https://link.aps.org/doi/10.1103/PhysRevB.87.241301>.
- [9] M. Gross, S. Haroche, Phys. Rep. 93 (5) (1982) 301–396, [http://dx.doi.org/10.1016/0370-1573\(82\)90102-8](http://dx.doi.org/10.1016/0370-1573(82)90102-8), URL <http://www.sciencedirect.com/science/article/pii/0370157382901028>.
- [10] S.L. McCall, E.L. Hahn, Phys. Rev. 183 (1969) 457–485, <http://dx.doi.org/10.1103/PhysRev.183.457>.
- [11] N.E. Rehler, J.H. Eberly, Phys. Rev. A 3 (1971) 1735–1751, <http://dx.doi.org/10.1103/PhysRevA.3.1735>, URL <http://link.aps.org/doi/10.1103/PhysRevA.3.1735>.
- [12] N. Bachelard, R. Carminati, P. Sebbah, C. Vanneste, Phys. Rev. A 91 (2015) 043810, <http://dx.doi.org/10.1103/PhysRevA.91.043810>, URL <http://link.aps.org/doi/10.1103/PhysRevA.91.043810>.
- [13] A. Fratalocchi, C. Conti, G. Ruocco, Opt. Express 16 (12) (2008) 8342–8349, <http://dx.doi.org/10.1364/OE.16.008342>.
- [14] C. Vanneste, P. Sebbah, Phys. Rev. Lett. 87 (2001) 183903, <http://dx.doi.org/10.1103/PhysRevLett.87.183903>, URL <http://link.aps.org/doi/10.1103/PhysRevLett.87.183903>.
- [15] A. Baczevski, Integral Equation and Discontinuous Galerkin Methods for the Analysis of Light-Matter Interaction (Ph.D. thesis), Michigan State University, 2013.
- [16] D. Stynk, Mathematical Modeling of Quantum Dots with Generalized Envelope Functions Approximations and Coupled Partial Differential Equations (Ph.D. thesis), Wilfrid Laurier University, 2009.
- [17] V.V. Temnov, U. Woggon, Opt. Express 17 (7) (2009) 5774–5782, <http://dx.doi.org/10.1364/OE.17.005774>.
- [18] J.M. Jin, The Finite Element Method in Electromagnetics, Wiley, New York, 2002.
- [19] B. Shanker, A.A. Ergin, M. Lu, E. Michielssen, IEEE Trans. Antennas and Propagation 51 (3) (2003) 628–641, <http://dx.doi.org/10.1109/TAP.2003.809054>.
- [20] A.E. Yilmaz, J.-M. Jin, E. Michielssen, IEEE Trans. Antennas and Propagation 52 (10) (2004) 2692–2708, <http://dx.doi.org/10.1109/TAP.2004.834399>.
- [21] G. Kobidze, J. Gao, B. Shanker, E. Michielssen, IEEE Trans. Antennas and Propagation 53 (3) (2005) 1215–1226, <http://dx.doi.org/10.1109/TAP.2004.841295>.
- [22] E. Bleszynski, M. Bleszynski, T. Jaroszewicz, Radio Sci. 31 (5) (1996) 1225–1251, <http://dx.doi.org/10.1029/96RS02504>.
- [23] D.C. Rapaport, The Art of Molecular Dynamics Simulation, Cambridge University Press, 2004.
- [24] L. Allen, J. Eberly, Optical Resonance and Two-level Atoms, in: Dover books on physics and chemistry, Dover, 1975.
- [25] J.S. Totero Gongora, A.E. Miroshnichenko, Y.S. Kivshar, A. Fratalocchi, Laser Photonics Rev. 10 (3) (2016) 432–440, <http://dx.doi.org/10.1002/lpor.201500239>.
- [26] L. Landau, E. Lifshitz, The Classical Theory of Fields, in: Course of Theoretical Physics, vol. 2, Elsevier Science, 2013.
- [27] J. Andreassen, H. Cao, A. Taflove, P. Kumar, C.-q. Cao, Phys. Rev. A 77 (2) (2008) 023810, <http://dx.doi.org/10.1103/PhysRevA.77.023810>.
- [28] J. Jin, The Finite Element Method in Electromagnetics, third ed., Wiley-IEEE Press, 2014.
- [29] A. Peterson, S. Ray, R. Mittra, I. Antennas, P. Society, Computational Methods for Electromagnetics, in: IEEE/OUP Series on Electromagnetic Wave Theory, IEEE Press, 1998.
- [30] C. Cohen-Tannoudji, J. Dupont-Roc, G. Grynberg, Photons and Atoms: Introduction to Quantum Electrodynamics, Wiley Online Library, 1989.
- [31] A.J. Pray, N.V. Nair, B. Shanker, IEEE Trans. Antennas and Propagation 60 (8) (2012) 3772–3781, <http://dx.doi.org/10.1109/TAP.2012.2201101>.
- [32] B. Shanker, A.A. Ergin, K. Aygun, E. Michielssen, IEEE Trans. Antennas and Propagation 48 (7) (2000) 1064–1074, <http://dx.doi.org/10.1109/8.876325>.
- [33] T. Theuns, Comput. Phys. Comm. 78 (3) (1994) 238–246, [http://dx.doi.org/10.1016/0010-4655\(94\)90002-7](http://dx.doi.org/10.1016/0010-4655(94)90002-7).
- [34] A.E. Yilmaz, J.-M. Jin, E. Michielssen, IEEE Antennas and Propagation Society International Symposium. Digest. Held in conjunction with: USNC/CNC/URSI North American Radio Sci. Meeting (Cat. No.03CH37450), vol. 3, 2003, pp. 543–546, <http://dx.doi.org/10.1109/APS.2003.1219905>.
- [35] C. Glosser, cglosser/QuEST: v0.2.1, Zenodo, 2018, <http://dx.doi.org/10.5281/zenodo.1246090>.