

Long-time simulations of nonlinear Schrödinger-type equations using step size exceeding threshold of numerical instability

T.I. Lakoba*

Department of Mathematics and Statistics, 16 Colchester Ave.,
University of Vermont, Burlington, VT 05401, USA

December 10, 2016

Abstract

We propose an exponential time-differencing method based on the leapfrog scheme for numerical integration of the generalized nonlinear Schrödinger-type equations. The key advantage of the proposed method over the widely used Fourier split-step method is that in the new method, numerical instability at high wavenumbers is strongly suppressed. This allows one to use time steps that considerably exceed the instability threshold, which leads to a proportional reduction of the computational time. Moreover, we introduce a technique that eliminates numerical instability at low-to-moderate wavenumbers that is common for methods based on the leapfrog scheme. We illustrate the performance of the proposed method with examples from two applications areas: deterministic wave turbulence and solitary waves.

Keywords: Numerical instability, Nonlinear wave equations, Statistics of nonlinear waves.

*tlakoba@uvm.edu, 1 (802) 656-2610

1 Introduction

The Fourier split-step method (SSM) is, perhaps, the most popular numerical method of solving the generalized nonlinear Schrödinger equation (NLS),

$$iu_t + u_{xx} + (\gamma|u|^2 + V(x))u = 0, \quad (1)$$

and similar equations, which arise in diverse applications. For a survey of various methods for this equation see, e.g., recent reviews [1, 2] and references therein. The SSM is explicit and hence easy to implement; it is relatively fast as it requires $O(N \log N)$ operations per time step, where N is the number of grid points; it also conserves (to numerical precision) the L_2 -norm of the solution. While it does not conserve the Hamiltonian (or some spatially discretized form of it), the SSM only causes it to slightly fluctuate about the initial value rather than systematically deviate from the latter, as, e.g., any explicit Runge–Kutta method would do. Such a quasi-preservation of conserved quantities by the SSM makes it well suited for simulation of long-time evolution.

Let us note that in such long-time simulations, in contrast to short-time ones, one rarely focuses on high accuracy of the obtained solution. For example, when simulating oscillations of a pulse in a potential well, it is reasonable to ask for the period and amplitude of those oscillations with accuracy, say, 0.1 %, but not for the exact location of the pulse after hundreds of such periods. Other actively researched areas concerned with long-time evolution are wave turbulence and formation of waves of extremely high amplitude (rogue waves); see, e.g., reviews [3]–[5]. In those studies one is interested in *statistics* of, e.g., waves' height and speed, but not in precise values of the solution; see, e.g., [6]. Accuracy of about 0.1 % or even lower may be quite appropriate for such statistical simulations.

Since simulations of long-time evolution require a considerable amount of physical time, one may be tempted to increase the numerical time step, given that a relatively low accuracy would be acceptable. However, for the SSM, the upper bound on the time step Δt is often set by the condition of numerical stability:

$$\Delta t < \Delta x^2 / \pi, \quad (2)$$

where Δx is the spatial mesh size. Condition (2) was first obtained in [7] for the pure NLS and later rederived in numerous studies. While this numerical instability (NI) may often be too weak to affect the solution over short time intervals, it will destroy or strongly degrade the solution over long times.

In this paper we propose an alternative to the SSM that allows one to use time steps that can be several times¹ greater than the threshold (2). This leads to a proportional reduction of the computational time. The main application of our technique is in studies of wave turbulence and rogue wave formation, where non-localized solution is simulated in a large domain with periodic boundary conditions. We will demonstrate with examples that the accuracy of statistical information obtained by the proposed technique is acceptable for such applications. Another advantage of

¹at least three–four times in one spatial dimension and even more in higher dimensions

the newly proposed method is that it can be straightforwardly applied to conservative or nearly conservative evolution equations with any type of the nonlinear term. This is in contrast to the SSM, whose applicability hinges on one’s ability to efficiently solve the nonlinear “substep” of the evolution.

Before we present our main idea, let us explain why two well-known methods to suppress NI, and hence to use a large time step, will not work for the above applications. One such method is filtering of high- k (k is wavenumber) components of the solution where NI occurs. However, filtering leads to energy dissipation, which affects both the L_2 -norm and the Hamiltonian and thus may considerably alter the evolution over *long times*. This will occur even though the “share” of numerically unstable high- k Fourier harmonics in the solution is typically very small, less than 0.001 % or so. The reason is that in the NLS or similar equations, nonlinearity results in energy “flowing” from low to high wavenumbers during the evolution; the inverse “cascade” also takes place; see, e.g., [5, 6] and references therein. Thus, numerical dissipation (i.e., filtering) at high k will affect the low- k , i.e. smooth, part of the solution.

The other way to greatly increase the NI threshold is to use either a linearly implicit method (as, e.g., in [8]) or a SSM with a finite-difference, instead of spectral, discretization of the u_{xx} term in (1). This has been shown to modify the stability condition (2) to be

$$\Delta t_{\text{thresh}} = O(\Delta x), \quad (3)$$

where the numerical coefficient on the right-hand side (r.h.s.) of (3) depends on the details of the method as well as on the nature of the simulated solution, e.g., plane wave [7] or soliton [9]. However, both linearly implicit methods and finite-difference discretizations, unless the latter are of very high order (see, e.g., [10]), truthfully mimic the dispersion only for $k = O(1)$. Already for *moderately* high $k = O(1/\sqrt{\Delta x}) \ll k_{\text{max}} = \pi/\Delta x$, which still represent physically significant components of the solution, a typical finite-difference discretization significantly distorts the dispersion relation²; see, e.g., [11]. Thus, it cannot be used to simulate phenomena like rogue wave formation, where high- k components play an important role.

In this paper we propose an exponential time differencing (ETD) method that uses the leap-frog (LF) scheme for integration of the nonlinear and potential-dependent terms in (1). We will refer to this method as ETD-LF. We will first demonstrate that the ETD-LF has a considerably weaker high- k NI than the SSM. However, this alone does not allow one to use this method for long-time simulations. Since the method is based on the LF scheme, it suffers from a low- k NI. The main contribution of the present work is a demonstration that a properly modified ETD-LF is capable of simulating long-time evolution of the generalized NLS with time steps exceeding the NI threshold severalfold compared to the SSM. The accuracy achieved with it appears to be acceptable in most statistical studies.

The main part of this work is organized as follows. In Section 2 we introduce the ETD-LF and a related method. In Section 3 we analyze their low- k and high- k NI. In particular, we show that

²It is, in fact, this distortion that raises the NI threshold from (2) to (3).

the high- k NI of the ETD-LF is considerably weaker than that of the SSM. In Section 4 we discuss modifications of the ETD-LF that address these NIs. In Sections 5 and 6 we report numerical examples that demonstrate performance of the modified ETD-LF in simulations of deterministic wave turbulence in one and two spatial dimensions, respectively. In Section 7 we summarize this work. Appendices A and B contain auxiliary results on high- k NI of the ETD-LF. Appendix A also demonstrates that the ETD-LF is effective in simulations of solitary waves.

2 Exponential time differencing and integrating factor LF methods

The idea behind the ETD and integrating factor (IF) methods was originally formulated for ordinary differential equations [12]–[14] and later applied to partial differential equations [15], [16]. Below we illustrate it for equations of the form

$$iu_t + \mathcal{L}u + \mathcal{N} = 0, \quad (4)$$

where \mathcal{L} is a linear operator with spatially constant coefficients and \mathcal{N} includes all other terms. The generalized NLS (1) is a special case of (4). Let \mathbf{F} and \mathbf{F}^{-1} be operators of Fourier transform and its inverse:

$$\widehat{u}(k, t) \equiv \mathbf{F}[u] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} u(x, t) e^{-ikx} dx, \quad u(x, t) = \mathbf{F}^{-1}[\widehat{u}] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \widehat{u}(k, t) e^{ikx} dx.$$

In what follows we will label the Fourier transform of any quantity with an over-hat; in particular, $\widehat{\mathcal{L}} \equiv \widehat{\mathcal{L}}(k)$ will denote the Fourier symbol of operator \mathcal{L} . Taking the Fourier transform of (4) yields

$$i\widehat{u}_t + \widehat{\mathcal{L}}\widehat{u} + \widehat{\mathcal{N}} = 0. \quad (5)$$

Solving this from t_1 to t_2 as a linear inhomogeneous equation yields:

$$e^{-i\widehat{\mathcal{L}}t_2} \widehat{u}(t_2) - e^{-i\widehat{\mathcal{L}}t_1} \widehat{u}(t_1) = \int_{t_1}^{t_2} e^{-i\widehat{\mathcal{L}}t'} i\widehat{\mathcal{N}}(t') dt', \quad (6)$$

where we have suppressed the obvious k -dependence of variables. Thus, the linear term in (4), which is often numerically stiff, is accounted for exactly by (6). One may think that this may eliminate the numerical stiffness due to \mathcal{L} in methods based on (6). However, as we will show in Section 3.2, the stiffness may still remain depending on the numerical implementation of the integral term in (6).

This numerical implementation depends on two things: (i) what finite difference discretization is used to approximate the integral, and (ii) what is assumed about the time evolution of $\widehat{\mathcal{N}}(t)$. Regarding (i), in this work we use the LF discretization. This choice is based on the well-known fact that the LF scheme preserves the L_2 -norm when applied to linear conservative equations. While the generalized NLS is not linear, we still expect — and later will confirm by examples —

that an LF-based method nearly preserves the L_2 -norm and the Hamiltonian of its solution³. In fact, a number of previous studies have used LF-based methods to solve the (generalized) NLS [17, 18] and other conservative nonlinear equations [19, 8]. Let us also note that the LF scheme is a member of the so called “geometric integrators” family of numerical methods (see, e.g., [20]), and such methods are known (see, e.g., [21]) to be well suited for integration of Hamiltonian equations. Finally, the LF scheme is explicit, which makes its implementation simple and hence competitive with that of the SSM.

Assumption (ii) in the previous paragraph distinguishes between the ETD and IF methods (see, e.g., [16]). For the ETD, one assumes that $\widehat{\mathcal{N}}(t)$ in (6) varies on a time scale of order one (i.e., much greater than $O(\Delta t)$) for all k , including those where $|\widehat{\mathcal{L}}| \gg 1$. On the other hand, the factor $\exp[-i\widehat{\mathcal{L}}t]$ is integrated exactly. Then, applying the LF scheme to (6), one obtains the ETD-LF method:

$$e^{-i\widehat{\mathcal{L}}t_{n+1}} \widehat{u}(t_{n+1}) - e^{-i\widehat{\mathcal{L}}t_{n-1}} \widehat{u}(t_{n-1}) = \frac{e^{-i\widehat{\mathcal{L}}t_{n+1}} - e^{-i\widehat{\mathcal{L}}t_{n-1}}}{-i\widehat{\mathcal{L}}} i\widehat{\mathcal{N}}(t_n). \quad (7)$$

In the case of the NLS (1), $\widehat{\mathcal{L}} = -k^2$ and $\widehat{\mathcal{N}} = \mathbf{F}[(\gamma|u|^2 + V(x))u]$; then with $t_n = n\Delta t$, (7) becomes:

$$e^{ik^2\Delta t} \widehat{u}(t_{n+1}) - e^{-ik^2\Delta t} \widehat{u}(t_{n-1}) = 2i\Delta t \operatorname{sinc}(k^2\Delta t) \widehat{\mathcal{N}}(t_n). \quad (8)$$

For the IF methods, one assumes that the entire integrand in (6), i.e. $e^{-i\widehat{\mathcal{L}}t} i\widehat{\mathcal{N}}(t)$, varies on a time scale of order one. Then, applying the LF scheme to (6), one obtains the IF-LF method:

$$e^{-i\widehat{\mathcal{L}}t_{n+1}} \widehat{u}(t_{n+1}) - e^{-i\widehat{\mathcal{L}}t_{n-1}} \widehat{u}(t_{n-1}) = 2\Delta t e^{-i\widehat{\mathcal{L}}t_n} i\widehat{\mathcal{N}}(t_n). \quad (9)$$

For the NLS (1) this simplifies to:

$$e^{ik^2\Delta t} \widehat{u}(t_{n+1}) - e^{-ik^2\Delta t} \widehat{u}(t_{n-1}) = 2i\Delta t \widehat{\mathcal{N}}(t_n). \quad (10)$$

One can see that the only difference between (8) and (10) is in the factor $\operatorname{sinc}(k^2\Delta t)$. In the next section we will show that this factor is responsible for the difference in the stability properties of the ETD-LF and IF-LF.

Both these methods have the second-order accuracy in time, inheriting this property from the LF scheme⁴. As the first step of either method, one can, in principle, use any first-order method without compromising the overall second-order accuracy. However, we found by experimentation that it benefits the accuracy to compute the solution at $t_1 = \Delta t$ by a second-order method. We used the second-order Runge–Kutta ETD method, ETD-2RK, from [16]:

$$c_1 = \Delta t e^{-ik^2\Delta t} \operatorname{sinc}(k^2\Delta t/2), \quad c_2 = (\Delta t - c_1)/(ik^2\Delta t); \quad (11a)$$

$$\widehat{a} = \widehat{u}_0 e^{-ik^2\Delta t} + c_1 i\widehat{\mathcal{N}}_0, \quad a = \mathbf{F}^{-1}[\widehat{a}]; \quad (11b)$$

³Here “nearly” means that these quantities slightly fluctuate about, but do not systematically deviate from, their initial values.

⁴There are no known “geometric integration” *explicit* methods of order higher than two, apart from fractional-step (also known as split-step) methods; see, e.g., [22]. Thus, one cannot increase the order of the ETD-LF (or the IF-LF) without making it implicit and thereby significantly slower, which would undo its main advantage over the SSM.

$$\widehat{u}_1 = \widehat{a} + c_2 i(\widehat{\mathcal{N}}_a - \widehat{\mathcal{N}}_0), \quad u_1 = \mathbf{F}^{-1}[\widehat{u}_1]; \quad (11c)$$

where u_0 is the initial condition and $\widehat{\mathcal{N}}_0$ and $\widehat{\mathcal{N}}_a$ stand for $\widehat{\mathcal{N}}$ evaluated at u_0 and a , respectively. The Fourier transform (including that for the computation of spatial derivatives) was implemented by the discrete (fast) Fourier transform for all the methods considered in this paper.

While the ETD and IF methods, and, of course, the LF scheme, have been known for a long time (see, e.g., recent studies [24]–[30] and references therein), we are not aware of any work where their combination (7)–(10) has been applied to the NLS-type equations.

3 Numerical instability of ETD-LF and IF-LF

We will obtain an equation for the numerical error and use it to infer stability of the methods derived in the previous section. We will employ the same technique as in [23] (see below) to analyze NI of the SSM on the background of the soliton of the NLS. The advantages of this technique over the von Neumann analysis are: (i) it does not require the background solution to be a plane wave or a constant in space, nor does it limit the equation to have spatially constant coefficients; (ii) it may provide intuitive understanding of the mechanism of the instability of the numerical method (see Section 3.1). The limitations of this technique are: (i) the equation is specific to the error *within a certain spectral range* rather than for all Fourier harmonics at once; (ii) it is valid only for a sufficiently small Δt , so that one can approximate the difference equation with a differential one. In Sections 3.1 and 3.2 we will obtain equations for the low- k and high- k numerical errors of *the NLS*; that is, we focus on specific implementations (8) and (10) of the general methods (7) and (9).

3.1 Low- k numerical instability of ETD-LF and IF-LF

We assume $\Delta t \ll 1$ and $k = O(1)$ in (8) and (10); the coefficients of $\widehat{\mathcal{N}}$ are always assumed to be of order one. We represent the numerical solution as

$$u(x, t_n) \equiv u_n(x) = u_b(x, t_n) + \tilde{u}_n(x), \quad |\tilde{u}_n| \ll |u_b|, \quad (12)$$

where u_b mimics the actual (“physical”), or background, solution, and \tilde{u}_n is the numerical error. Substitution of the Fourier transform of (12) into (8) and (10) yields:

$$e^{ik^2\Delta t} \widehat{\tilde{u}}_{n+1} - e^{-ik^2\Delta t} \widehat{\tilde{u}}_{n-1} = 2i\Delta t \widehat{\mathcal{G}}[k^2\Delta t] \widehat{\mathcal{M}}(u_b, \tilde{u}_n), \quad (13a)$$

where $\widehat{\mathcal{G}} = \text{sinc}(k^2\Delta t)$ for the ETD-LF, $\widehat{\mathcal{G}} \equiv 1$ for the IF-LF, and $\widehat{\mathcal{M}}$ is the linearization of $\widehat{\mathcal{N}}$:

$$\widehat{\mathcal{M}} = \mathbf{F} [2\gamma|u_b|^2\tilde{u}_n + u_b^2\tilde{u}_n^* + V(x)\tilde{u}_n]. \quad (13b)$$

The zeroth-order in Δt solution of (13) satisfies: $\widehat{\tilde{u}}_{n+1} = \widehat{\tilde{u}}_{n-1}$, which admits two possibilities:

$$\widehat{\tilde{u}}_{n+1}^{(\pm)} = \pm \widehat{\tilde{u}}_n^{(\pm)}. \quad (14)$$

This is well known for any method based on the LF scheme, which involves three time levels. The quantity $\tilde{u}^{(+)}$ will be referred to as the “physical” error. Indeed, it exhibits a continuous evolution: $\tilde{u}_{n+1}^{(+)} - \tilde{u}_n^{(+)} = O(\Delta t) \cdot |\tilde{u}^{(+)}|$ (see the r.h.s. of (13a)) and thus mimics the physical perturbation of the background solution u_b . On the other hand, $\tilde{u}^{(-)}$ is a parasitic error, i.e., a numerical artifact (its time evolution is *not* continuous: $\tilde{u}_{n+1}^{(-)} - \tilde{u}_n^{(-)} = O(1) \cdot |\tilde{u}^{(-)}|$). The parasitic error is harmless unless it grows and contaminates the solution $u(x, t)$. The evolution equation for the numerical error is obtained by expanding the l.h.s. of (13a) into the Maclaurin series while assuming that $(\pm 1)^n \hat{u}^{(\pm)}(t_n)$ are *continuous* functions of time, so that

$$\hat{u}^{(\pm)}(t_{n+1}) = \pm \left(\hat{u}^{(\pm)}(t_n) + \Delta t \hat{u}_t^{(\pm)}(t_n) + \frac{\Delta t^2}{2} \hat{u}_{tt}^{(\pm)}(t_n) + \dots \right). \quad (15)$$

Neglecting terms $O(\Delta t^3)$, one finds from (13a):

$$\hat{u}_t^{(\pm)} + ik^2 \hat{u}^{(\pm)} = \pm \widehat{\mathcal{M}}(u_b, \hat{u}^{(\pm)}); \quad (16)$$

here we have used that, for both the ETD-LF and IF-LF, $\widehat{\mathcal{G}}[k^2 \Delta t]$ can be replaced by unity for $k^2 \Delta t \ll 1$. Taking the inverse Fourier transform of (16) yields (for $V(x) \equiv 0$, which we set in the remaining derivations for simplicity)⁵:

$$i\tilde{u}_t^{(\pm)} + \tilde{u}_{xx}^{(\pm)} \pm \gamma(2|u_b|^2 \tilde{u}^{(\pm)} + u_b^2 \tilde{u}^{(\pm)*}) = 0. \quad (17)$$

Not surprisingly, this is (up to the (\pm) in front of γ) the linearized NLS; the accuracy with which it has been obtained is consistent with the LF being a second-order scheme.

The evolution of the numerical error depends on the sign of γ (and, of course, on the background solution u_b). For the focusing NLS ($\gamma > 0$) and u_b being a plane wave of amplitude A , small perturbations with $|k| \in (0, 2\gamma A^2)$ are known to undergo modulational instability (see, e.g., [31]). The evolution of the “physical” part of the error, $\tilde{u}^{(+)}$, is governed by the linearized NLS (with the $(+)$ in front of γ) and hence will be indistinguishable from the evolution of any “physical” perturbation that may exist in the system. On the other hand, the parasitic error, $\tilde{u}^{(-)}$, obeys the evolution governed by Eq. (17) with the $(-)$ in front of γ . This is equivalent to a “physical” perturbation of a plane wave in the *defocusing* case, where perturbations are known to be stable. Thus, for $\gamma > 0$, the numerical low- k error will be dominated by $\tilde{u}^{(+)}$ and therefore will closely approximate the behavior of a “physical” perturbation. When u_b is not a plane wave, the conclusion will be qualitatively the same. This is because $\tilde{u}^{(+)}$ and $\tilde{u}^{(-)}$ will obey the evolutions governed by the focusing and defocusing NLS, respectively, and the former is known to exhibit much more volatile dynamics than the latter (we also will demonstrate this by an example in Section 5).

For the defocusing NLS ($\gamma < 0$), the “physical” error $\tilde{u}^{(+)}$ will obey the modulationally stable dynamics, while Fourier harmonics with $k \in (0, k_{\text{MI}})$ of the parasitic error $\tilde{u}^{(-)}$ will be modulationally unstable. That is, the parasitic error will be governed by the “more volatile” focusing NLS

⁵All conclusions of this paper remain valid also for $V(x) \not\equiv 0$ as long as $\max |V(x)| \Delta t \ll 1$. Thus, for $V(x) \sim x^2$, they may be altered, although the ETD- and IF-LF methods themselves will remain applicable for sufficiently small Δt .

and hence will, eventually, overcome the “physical” solution. Thus, both ETD- and IF-LF are *numerically unstable* for simulating a defocusing NLS. The underlying reason for this low- k NI is the existence of the parasitic error $\tilde{u}^{(-)}$. Let us note that a similar NI was explained in [32] by the von Neumann analysis for the method proposed in [17]. However, as we have remarked earlier, the von Neumann analysis did not reveal a connection between the NI observed there and the “physical”, i.e. modulational, instability of the parasitic error. Our method, based on Eq. (17), predicts that *any* numerical method based on the LF scheme, not just that of [17], will suffer from this NI for the defocusing NLS.

To conclude this subsection, let us note that the presence of a parasitic error $\tilde{u}^{(-)}$ is also responsible for a *nonlinear* NI in the *focusing* NLS. We did not undertake its analysis here since this would have taken us in a direction completely unrelated to the main thrust of this work. We merely note that analysis of such (or closely related) nonlinear NI, for a different equation, was done in [33]. This nonlinear NI occurs for Fourier harmonics with relatively low k , which are within the spectral band of the “bulk” of the solution. In Section 4.1 we will give an example of this NI. Its manifestation is rather universal [33]: once those parasitic harmonics reach a certain critical size (which may take several hundreds of time units), they undergo an explosive growth, and the numerical solution blows up within just a few time units. In contrast, for the *defocusing* NLS, addressed in the previous paragraph, the growth of $\tilde{u}^{(-)}$ is more gradual, but the parasitic error takes a smaller time to develop from the noise level to a size where it destroys the solution. For the parameters considered in this work, this time was several tens to about a hundred of time units.

3.2 High- k numerical instability of ETD-LF and IF-LF

Here we assume that the spectrum of numerical error \tilde{u}_n is concentrated near $|k| = k_0 \gg 1$. As in [23], we will show that high- k NI can occur only for $k_0 \approx k_{\pi m}$, where

$$k_{\pi m}^2 \Delta t = \pi m, \quad m = 1, 2, \dots \quad (18)$$

Denote $k_0^2 \Delta t \equiv \phi$; then in the zeroth order in Δt , \tilde{u} satisfies $e^{i\phi} \widehat{u}_{n+1} = e^{-i\phi} \widehat{u}_{n-1}$, which admits two solutions:

$$\widehat{u}_{n+1}^{(\pm)} = \pm e^{-i\phi} \widehat{u}_n^{(\pm)}. \quad (19)$$

In the next approximation, we allow $e^{i\phi} \widehat{u}^{(\pm)}$ to be slow functions of time, just as in (15). For \tilde{u} with a sufficiently narrow spectrum whose wavenumbers satisfy

$$|k^2 - k_0^2| \Delta t \ll 1, \quad (20)$$

one obtains, similarly to (16):

$$\widehat{u}_t^{(\pm)} + i(k^2 - k_0^2) \widehat{u}^{(\pm)} = \pm \widehat{\mathcal{G}}[\phi + (k^2 - k_0^2) \Delta t] \widehat{\mathcal{M}}(u_b, \widehat{u}^{(\pm)}). \quad (21)$$

Here we have assumed, just as when approximating the difference equation (13a) by the differential equation (16), that the term on the r.h.s. of (21) multiplied by Δt is sufficiently small. For future

reference we note that this assumption may be violated for some high- k harmonics at those times when the spectrum expands considerably.

Before moving on, let us explain why $k_0 \approx k_{\pi m}$ is necessary for the occurrence of NI [23]. The r.h.s. of (21) contains two terms varying with time steps as $e^{-i\phi n}$ and $e^{i\phi n}$, respectively; see (13b) and (19). An instability, similar to the modulational instability, can occur only when these terms are in synchronism with each other, which requires $e^{2i\phi} = 1$. This condition is equivalent to (18).

The above considerations imply that the threshold for high- k NI is

$$\Delta t_{\text{thresh}} = \pi/k_{\text{max}}^2, \quad (22)$$

where k_{max} is defined after (3). Note that (22) is consistent with condition (2).

Continuing with (21), one can easily take its inverse Fourier transform when $\widehat{\mathcal{G}} \equiv 1$, i.e. for the IF-LF method. The result is:

$$i\tilde{u}_t^{(\pm)} + (\tilde{u}_{xx}^{(\pm)} + k_{\pi m}^2 \tilde{u}^{(\pm)}) \pm \gamma(2|u_b|^2 \tilde{u}^{(\pm)} + u_b^2 \tilde{u}^{(\pm)*}) = 0, \quad (23)$$

which is *the same* (apart from the (\pm) in front of γ) as the equation for the high- k numerical error of the SSM; see Eq. (A2) in [23]. As shown there, the role of the γ -term is to enable interaction of harmonics with $k \approx k_{\pi m}$ and $k \approx -k_{\pi m}$ with each other, which is the mechanism of high- k NI. Importantly, unlike in the low- k case, the NI resulting from that interaction is *independent* of the sign of γ (see Eq. (3.25) in [23]). Therefore, for any background solution u_b whose spectrum is negligibly small near $|k| \sim k_\pi$ (see below), such as a soliton or plane wave, the NI of the IF-LF and SSM must be very similar. In Appendix A we demonstrate this for u_b being the soliton. On the other hand, for solutions whose spectrum expands considerably past k_π , as in the examples in Section 5, we have found that NI for the IF-LF is stronger than that for the SSM, as illustrated in Fig. 1(a). We comment on a possible reason for that in Appendix A.

For the ETD-LF, i.e. when $\widehat{\mathcal{G}}[k^2 \Delta t] = \text{sinc}(k^2 \Delta t) \equiv \text{sinc}(\pi m + (k^2 - k_{\pi m}^2) \Delta t)$, the inverse Fourier transform of (21) cannot be found exactly. However, note that $\text{sinc}(\pi m) = 0$ for $m = 1, 2, \dots$ (see Fig. 2) and therefore $\widehat{\mathcal{G}}[k^2 \Delta t] = O(|k^2 - k_{\pi m}^2| \Delta t) \ll 1$. A counterpart of (23) for the ETD-LF is then:

$$i\tilde{u}_t^{(\pm)} + (\tilde{u}_{xx}^{(\pm)} + k_{\pi m}^2 \tilde{u}^{(\pm)}) + O(|k^2 - k_{\pi m}^2| \Delta t) = 0. \quad (24)$$

In particular, for $|k - k_{\pi m}| = O(1)$, the last term in (24) is $O(\sqrt{\Delta t}) \ll 1$. This shows that the high- k NI of the ETD-LF is strongly suppressed compared with that of the IF-LF and the SSM; see Fig. 1(a). Let us note that the spikes seen in the spectrum of the ETD-LF solution do *not* grow exponentially; this is confirmed by Fig. 1(b). As for the IF-LF, they are created and augmented at times when the spectrum expands past the corresponding $k_{\pi m}$ (see Appendix A). With each subsequent expansion they grow (on average), but do so much slower than exponentially. As we will discuss in Section 4.2 and demonstrate in Sections 5 and 6, those spectral distortions affect the solution only insignificantly.

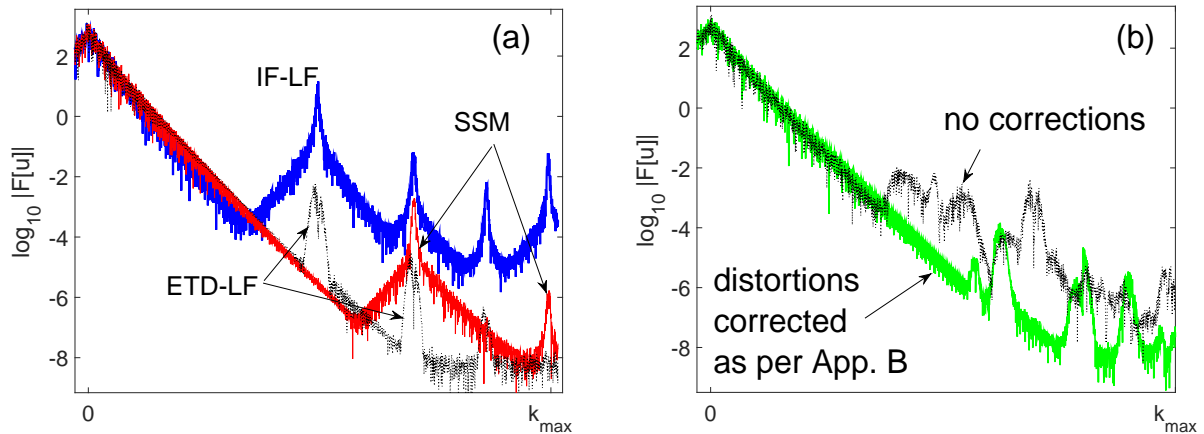


Figure 1: (Color online) Fourier spectra of the numerical solution of (1) with $V(x) \equiv 0$ and the initial condition (32). Here and below in this paper, we show only the part with $k > 0$, since the spectra are symmetric relative to $k = 0$. The parameters in all displayed cases are the same as in Section 5.1, except that $\Delta t = 1.8 \cdot 10^{-3}$ and $t = 20$ for panel (a) and $t = 10,000$ for panel (b). Panel (a) compares IF-LF (thicker solid, blue), SSM (thinner solid, red), and ETD-LF (dotted, black). No stabilization (see Section 4) is used for IF-LF and ETD-LF. The spikes in the spectra of IF-LF and SSM are the exponentially growing numerically unstable modes. (For example, the IF-LF solution is destroyed by the NI before $t = 50$.) Panel (b) shows the solution obtained with ETD-LF at $t = 10,000$ for the same parameters as in (a). Dotted (black) line is for the case where no amelioration of spectral distortions near $|k| = k_\pi$ and $k_{2\pi}$ was used; solid (green) line is when such amelioration techniques, described in Appendix B (see also Section 4.2) were used. Both simulations in panel (b) had to use stabilization (28) every $t_{\text{stabilize}} = 10$ time units to overcome the low- k NI described in Sections 3.1 and 4.1.




Figure 2: Functions depicting the strength of the nonlinear interaction, $\widehat{\mathcal{G}}$, for the ETD-LF (solid) and for its modified version used in Section 6 (dashed). The dotted lines are guides for the eye.

In Appendix A we also show that the NI does not occur in the ETD-LF for the soliton background solution. In particular, in this case the ETD-LF appears to be numerically stable for Δt

exceeding the NI threshold by more than ten times.

4 Suppression of low- k numerical instability and high- k distortions of ETD-LF

Since the high- k NI of the ETD-LF is much weaker than that of the IF-LF, we have chosen the former method as a “platform” to modify. Our goal is to suppress the low- k NI, as well as ameliorate distortions in the spectrum in the vicinity of $k_{\pi m}$, described in the previous section. We will show later that suppression of the low- k NI is critically important to render the ETD-LF usable for long-time simulations. On the other hand, amelioration of the high- k distortions will prove not to be significant to improve performance of the ETD-LF with large time steps.

4.1 Suppression of low- k numerical instability

As we noted in Section 3.1, the low- k NIs are due to the parasitic component $\tilde{u}^{(-)}$ of the numerical error. Therefore, if $\tilde{u}^{(-)}$ is suppressed, so will be the NI. The growth rate of the *linear* low- k NI for the defocusing NLS is $O(1)$: see (17). For the *nonlinear* low- k NI for the focusing NLS, the growth rate is even smaller; see the end of Section 3.1. Then, it will suffice to suppress $\tilde{u}^{(-)}$ every $t_{\text{stabilize}} = O(1)$ time units to keep it from growing in the long term. The most common way to do so, which has long been known for LF-based methods, is to simply discard solution u_{n+1} while keeping u_n and then restart the simulation from scratch. However, this does not eliminate $\tilde{u}^{(-)}$ that has already been created, but merely makes its accumulation slower because $\tilde{u}^{(-)}$'s created at each restart add up incoherently.

We have chosen another approach, where $\tilde{u}^{(-)}$ is actually eliminated (with the same accuracy with which (19) is valid) at each $t_{\text{stabilize}}$ time unit. The simplest suppression scheme is:

$$\widehat{v}_n = (\widehat{u}_n + \widehat{u}_{n-1})/2, \quad \widehat{v}_{n+1} = (\widehat{u}_{n+1} + \widehat{u}_n)/2, \quad (25)$$

where $\widehat{v}_n = \widehat{u}_n e^{ik^2 \Delta t_n}$; \widehat{v}_n and \widehat{v}_{n+1} are then used for the computation of the solution at the $(n+2)$ th time level. The averaging of the two successive time levels in (25) suppresses $\tilde{u}^{(-)}$, as seen from (14). However, this introduces dissipation. Indeed,

$$\bar{u}(t_{n+1}) = u(t_n + 0.5\Delta t) + \frac{\Delta t^2}{2}(u(t_n + 0.5\Delta t))_{tt} + \dots, \quad (26)$$

and then this step modifies the simulated equation, (1), as:

$$iu_t + u_{xx} + (\gamma|u|^2 + V(x))u = O(\Delta t^2) \cdot iu_{tt}, \quad (27)$$

which reveals dissipation with rate $O(\Delta t^2)$. Over long simulation times this may be significant. Therefore, we employed averaging over several time levels:

$$\widehat{v}_n = \frac{11}{16}\widehat{u}_n + \frac{15}{64}(\widehat{u}_{n-1} + \widehat{u}_{n+1}) - \frac{3}{32}(\widehat{u}_{n-2} + \widehat{u}_{n+2}) + \frac{1}{64}(\widehat{u}_{n-3} + \widehat{u}_{n+3}), \quad (28)$$

and similarly for \widehat{v}_{n+1} . This produces dissipation at a rate $O(\Delta t^6)$, which we found not to affect our results. We will refer to (28) as a stabilization step. A value of $t_{\text{stabilize}}$ can be established by a quick experimentation; we found the results not to be sensitive to it within a reasonably wide range.

Let us mention that time instances when stabilization needs to be applied can instead be determined automatically. Given (14) or (19), one can write the numerical solution as

$$u_n = u_b(t_n) + \tilde{u}_n^{(+)} + (-1)^n \tilde{u}_n^{(-)}, \quad (29)$$

where $u_b(t_n)$ is the background solution at t_n and the last two terms are the two components of the numerical error. Both u_b and $\tilde{u}^{(\pm)}$ vary little over one time step. Consider now two quantities:

$$\Delta_1 u \equiv u_{n+1} - u_n \approx (u_b)_t \Delta t + \tilde{u}_t^{(+)} \Delta t - 2(-1)^n \tilde{u}^{(-)} \approx (u_b)_t \Delta t - 2(-1)^n \tilde{u}^{(-)}, \quad (30a)$$

$$\Delta_2 u \equiv u_{n+1} - 2u_n + u_{n-1} \approx (u_b)_{tt} \Delta t^2 - 4(-1)^n \tilde{u}^{(-)}; \quad (30b)$$

here all functions on the r.h.s. are evaluated at t_n , and we have neglected derivatives of $\tilde{u}^{(+)}$ compared to the corresponding derivatives of u_b . Since $\Delta t \ll 1$ and initially one has

$$|\tilde{u}^{(-)}| \ll |(u_b)_t| \Delta t, \quad (31a)$$

then, also initially,

$$\|\Delta_1 u\| = C_0 \|\Delta_2 u\|, \quad C_0 \gg 1, \quad (31b)$$

where $\|\dots\|$ is the L_2 -norm. If $\|\tilde{u}^{(-)}\|$ begins to grow, that will eventually increase the ratio $\|\Delta_2 u\|/\|\Delta_1 u\|$. One can monitor that ratio (say, every $t = 1$ units) and prescribe that once that ratio has exceeded C_0 by a specified factor (say, five), a stabilization step (28) is to be applied. Once the decision to apply stabilization has been made, one saves the solution at seven successive time levels and at the seventh, applies (28). In our simulations, however, we did not need to implement such a monitoring because finding a working value of $t_{\text{stabilize}}$ was easy.

Concluding the discussion of the low- k NI in the focusing NLS, in Fig. 3 we illustrate its development when suppression of $\tilde{u}^{(-)}$ is *not* applied. The simulated setup is that considered in Section 5.2 with similar parameters and the time step is $\Delta t < \Delta t_{\text{thresh}}$, so that no high- k NI or spectral distortions could occur. The time evolution of $\|\Delta_1 u\|$ and $\|\Delta_2 u\|$ (Fig. 3(a)) indicates that near $t = 200$, the parasitic error component $\tilde{u}^{(-)}$ begins to grow. Figures 3(b,c) show the spectra of $\Delta_1 u$ and $\Delta_2 u$ at $t = 1$ and $t = 250$, which confirm the above conclusion. Figure 3(d) shows the spectra of the solution u at $t = 365$ when stabilization is and is not applied. The feature that indicates the low- k NI is the broader central part compared to that of the undistorted (i.e., stabilized) solution. In the x -space, such a broader spectrum is manifested by a relatively small-period ripple in u ; compare Figs. 3(e,f). Finally, the solution blows up (due to the NI) around $t = 373$, i.e. only a few time units after the broader central part of the spectrum of u had developed to become visible on a logarithmic scale.

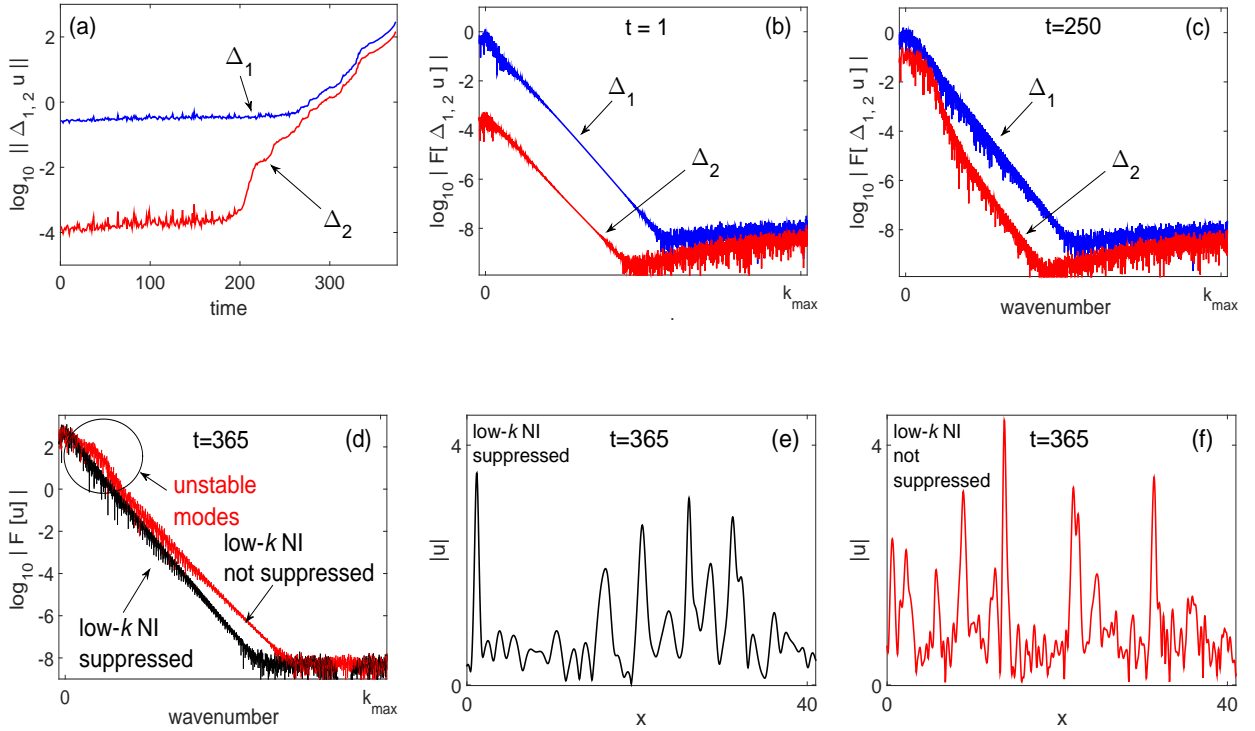


Figure 3: See text for explanation. Panels (e,f) show only part of the computational domain.

4.2 Amelioration of high- k distortions

We have tried several approaches to suppress distortions in the solution's spectrum near $\pm k_\pi$ and $\pm k_{2\pi}$. For example, we multiplied the coefficient $\widehat{\mathcal{G}} = \text{sinc}(k^2 \Delta t)$ in (8) (see also (13a)) by empirically chosen factors that reduced the value of that coefficient in a broader vicinity of $|k| = k_{\pi m}$. The result of applying two approaches (different from the one just described) which were most successful in combatting the spectral distortions is shown in Fig. 1(b). The distortions near $\pm k_\pi$ are entirely removed, while near $k_{2\pi}$ their size is considerably reduced. Despite that, we have found that they did not noticeably alter the accuracy of the statistical quantities (see Section 5) computed by the ETD-LF. For this reason we did not use *any* such correctional techniques in the examples reported in Section 5. For the same reason, we describe the techniques employed to “clean up” the spectrum in Fig. 1(b) in Appendix B rather than in the main text.

5 Numerical experiments in one spatial dimension

We will report results for four cases. The first two are different examples of the generalized focusing NLS (1) while the other two illustrate robustness of our modified ETD-LF method in situations where the LF scheme alone has been known to be prone to low- k NI. In all examples, the simulation time was $t = 3000$, and the initial condition was taken as

$$u(x, 0) = 1 + 0.24 \cos(15\Delta k x + 0.7) + 0.33 \sin(21\Delta k x + 0.2) + 0.42 \cos(31\Delta k x) + \xi(x), \quad (32)$$

where $\xi(x) \sim 10^{-10}$ is white noise added to define the noise floor. To avert the low- k NI, the stabilization by Eq. (28) was applied every $t_{\text{stabilize}} = 10$ time units in Sections 5.1 – 5.3. In Section 5.4, where we considered a defocusing NLS, we used a smaller value, $t_{\text{stabilize}} = 1$. This was necessary because, as we mentioned at the end of Section 3.1, the parasitic error $\tilde{u}^{(-)}$ grows initially faster for the defocusing than for the focusing NLS. Since the above values of $t_{\text{stabilize}}$ correspond to once per every several hundreds or even thousands of time steps, applying the stabilization step did not affect the computational time perceptibly.

For each example we report the following statistical quantities: the probability density function (PDF) of $|u|^2$; the PDF of $|u_x|^2$; the correlation function in space, $C_x(\delta x) = \langle u^*(x, t)u(x + \delta x, t) \rangle / \langle |u(x, t)|^2 \rangle$; and the correlation function in time, $C_t(\delta t) = \langle u^*(x, t)u(x, t + \delta t) \rangle / \langle |u(x, t)|^2 \rangle$. Note that having information from the PDF of both $|u|^2$ and $|u_x|^2$ allows one to conclude about the distribution of the solution's energy between the potential and kinetic parts (these are given by, respectively, the second and first terms in (35c) below); see also the last two sentences in the caption to Fig. 10.

The PDFs were computed as follows. Starting at some time moment in a late stage of the simulation, we recorded values of $|u|^2$ and $|u_x|^2$ at every grid point. This was done $n_{\text{av}, x}$ times every $t_{\text{record}, x}$ time units. We used the following respective values for $n_{\text{av}, x}$ and $t_{\text{record}, x}$: 10^3 and 0.5 (Section 5.1); 400 and 0.25 (Section 5.2); 10^3 and 0.05 (Section 5.3); 500 and 2 (Section 5.4). The values for $t_{\text{record}, x}$ were chosen so that $|C_t(t_{\text{record}, x})| \ll |C_t(0)|$, i.e. the records were only weakly correlated in time, and then the values of $n_{\text{av}, x}$ ensured that the system's properties remained statistically constant over the averaging period. The PDFs are then computed as histograms normalized in the standard way:

$$\int_0^{y_{\text{max}}} \text{PDF}(y) dy = 1.$$

The x -correlation function was obtained by a similar averaging procedure.

For the t -correlation function, the averaging had to be done differently. At instances t_j in a late stage of the simulation, where $j = 1, \dots, n_{\text{av}, t}$, we recorded the field $u(x, t_j)$ for all x . We then computed the product $P_j = \sum_{x_i} u^*(x_i, t_j)u(x_i, t_j + \delta t) / \sum_{x_i} |u(x_i, t_j)|^2$ at later time instances ($t_j + \delta t$). Finally, we took the average over all t_j :

$$C_t(\delta t) = \frac{1}{n_{\text{av}, t}} \sum_{j=1}^{n_{\text{av}, t}} P_j \equiv \frac{1}{n_{\text{av}, t}} \sum_{j=1}^{n_{\text{av}, t}} \frac{\sum_{x_i} u^*(x_i, t_j)u(x_i, t_j + \delta t)}{\sum_{x_i} |u(x_i, t_j)|^2}. \quad (33)$$

The instances t_j were taken sufficiently far apart so that the field is substantially uncorrelated: $C_t(t_{j+1} - t_j) \approx 0$, and we used $n_{\text{av}, t} = 100$.

We will now address the issue of benchmarking, i.e. what method(s) with what spectral and temporal resolution can be used to provide a reference solution. We will first address this issue in general, and provide specific details for each of the four examples in the respective subsections below.

The candidate methods were the second-order (symmetrized) SSM and the ETD-4RK. The latter is the ETD method based on the 4th-order Runge–Kutta scheme, proposed in [16]. It has been shown to produce highly accurate results when applied to the NLS and other nonlinear wave equations (see, e.g., [26, 27, 29]). Let us, however, clarify three aspects about the ETD-4RK. First, it is almost four times slower than the SSM and ETD-LF (whose computational speeds are approximately the same). Second, even though for a given Δt , the ETD-4RK is much more accurate than the other two methods, its speed cannot be increased considerably by taking larger time steps. The reason is that the ETD-4RK has the same NI threshold Δt_{thresh} as the IF-LF and SSM. This was revealed by the von Neumann analysis for a plane-wave background solution, which we do not present here. Even though the growth rate of harmonics with $|k| \approx k_\pi$ was found to be about an order of magnitude less for the ETD-4RK than for the SSM, the high- k NI still destroys the solution over long times considered in the examples in Sections 5.1–5.4 and 6.1, 6.2. A similar comment pertains to 4th-order SSMs, whose stability was analyzed in [23] and which we did not use in this paper. Third, unlike the ETD-LF and SSM, the ETD-4RK does *not* nearly preserve the L_2 -norm and the Hamiltonian (see Section 2); rather, those quantities computed for the ETD-4RK solution exhibit systematic drift. However, we have found that for sufficiently small time steps, such as those used below, that drift up to $t \sim 3000$ may be less than average fluctuations of the conserved quantities computed by the ETD-LF and SSM.

For all benchmarking simulations we used a smaller time than for the “main” ones: $t = 1000$ instead of $t = 3000$. This had to be done because benchmarking simulations with the highest resolution for the ETD-4RK took quite a long time even for $t = 1000$: more than a week in Section 5.1 and more than three days in Section 5.2 (on a personal computer running Matlab).

5.1 Pure focusing NLS

Here we set $V(x) \equiv 0$, $\gamma = 2$, and $L = 48\pi$ in the NLS (1).

5.1.1 Benchmarking simulations

To compare accuracy of the ETD-4RK, ETD-LF, and SSM, we varied both the number of grid points N , i.e. spectral resolution, and Δt . For $N = 2^{12} \equiv N_0$, one has $k_{\text{max}} = 85.3$, and the NI threshold is $\Delta t_{\text{thresh}}(N_0) \equiv \Delta t_0 = 4.3 \cdot 10^{-4}$.

First, we have verified that all four statistical quantities — the two PDFs and the two correlation functions — obtained by the ETD-4RK remain practically the same when N was increased to $N = 2N_0$, $3N_0$, and $6N_0$ and Δt was decreased according to (2):

$$\Delta t(N) = 2 \cdot 10^{-4} / (N/N_0)^2 \lesssim (\Delta t_0/2) / (N/N_0)^2. \quad (34)$$

Thus, results for any of the above N 's can be chosen as benchmark ones; in Fig. 4 they are shown by the thick black line.

Second, we repeated this for the ETD-LF and SSM with $\Delta t = \Delta t(N)/2$. The ETD-LF results for $N = N_0$ are already quite close to the benchmark ones: see thin solid lines in Fig. 4. They improve




Figure 4: Benchmarking simulations with ETD-4RK (thick solid), ETD-LF (thin solid), SSM (dashed); see more explanations in the text. In (c), $\text{Im}[C_x(\delta x)]$ is almost zero for all methods and hence is not plotted.

only a little for $N = 2N_0$ and $N = 3N_0$, whereas for $N = 6N_0$ they become indistinguishable from the benchmark results. The SSM's PDFs for $N = N_0$ are quite far from the benchmark ones, although the correlation functions are not considerably worse than those obtained by the ETD-LF; see dashed lines in Fig. 4. For $N = 2N_0$, all four of the SSM's results are improved and become about as close to the benchmark results as the ETD-LF's ones with $N = N_0$ to $3N_0$. For $N = 3N_0$, the SSM's results become noticeably closer to the benchmark ones than they are for $N = 2N_0$. Thus, the SSM's performance improves more gradually with N than the ETD-LF's performance. Finally, for $N = 6N_0$ the SSM's results become indistinguishable from the benchmark ones.

In Fig. 5 we illustrate how the three methods handle conserved quantities of the pure NLS:

$$I_2 = \int_0^L |u|^2 dx \quad (\text{mass}); \quad (35a)$$

$$J_2 = i \int_0^L (u^* u_x - u_x^* u) dx \quad (\text{momentum}); \quad (35b)$$

$$I_4 = \int_0^L (-|u_x|^2 + (\gamma/2)|u|^4) dx \quad (\text{Hamiltonian}); \quad (35c)$$

$$I_6 = \int_0^L (-|u_{xx}|^2 + (\gamma^2/2)|u|^6 - (\gamma/2)(|u^2|_x)^2 - 3\gamma|u_x|^2|u|^2) dx \quad (\text{higher-order law [34]}). \quad (35d)$$

As we have mentioned earlier, the systematic drift of the conserved quantities computed by the ETD-4RK is much smaller than their fluctuations in the ETD-LF and SSM, and hence can be ignored.




Figure 5: Time evolution of $R_j(t) = (I_j(t) - I_j(0))/I_j(0)$ (solid: $j = 2$; dashed: $j = 4$; dash-dotted: $j = 6$) and of $J_2(t) - J_2(0)$ (dotted), computed by the ETD-4RK (a), ETD-LF (b), and SSM (c) for $N = N_0$ and $\Delta t = \Delta t_0/2$. The lines for I_6 and J_2 are labeled in panel (a) for clarity. The lines for I_2 and J_2 are identically zero in panel (c).

As a side note, let us mention that these conserved quantities, or more specifically, their errors R_j , defined in the caption to Fig. 5, can be used to confirm the second-order accuracy of the ETD-LF even for $\Delta t > \Delta t_{\text{thresh}}$ ⁶. Let us define

$$R_{j, \text{av}} = \frac{1}{T_{\text{av}}} \int_{1000-T_{\text{av}}}^{1000} \sqrt{R_j(t)^2} dt, \quad j = 2, 4, 6, \quad (36)$$

where $T_{\text{av}} = 500$ for the ETD-LF and SSM and $T_{\text{av}} = 100$ for the ETD-4RK (because its data exhibit a systematic drift). Figure 6 demonstrates that these quantities sufficiently closely (within 20%) follow the second-order (for the ETD-LF and SSM) and fourth-order (for the ETD-4RK) dependencies on the time step. It should be mentioned that the fact that the conserved quantities computed by the SSM are less accurate than those computed by the ETD-LF (for the specific example in this subsection), is of secondary importance. Of primary importance is the fact that the ETD-LF can operate with a four-time greater time step and hence save its user the computational time.

Third and last, we set $N = N_0$ and investigated how much one was to reduce Δt in order to have the SSM approximate the benchmark PDFs sufficiently accurately. We found that for $\Delta t = \Delta t_0/10$, the PDFs obtained by the SSM achieve approximately the same accuracy as the PDFs obtained by the ETD-LF for $N = N_0$ and $\Delta t = \Delta t_0/2$ (Fig. 4(a,b)). Thus, for the pure NLS with the above parameters, it will take the SSM about five times longer to obtain results with accuracy similar to

⁶The solution itself of the NLS cannot be used for this purpose. Indeed, since the system exhibits deterministic chaos, any two solutions that are initially different, no matter how slightly, will separate by an amount $O(1)$ over the long times considered in this study.




Figure 6: Dependence of the time-averaged errors (38) on the time step for $N = N_0$. Circles/solid lines, stars/dashed lines, and triangles/dashed-dotted lines pertain to the ETD-4RK, ETD-LF, and SSM, respectively. The thin dashed line underneath the ETD-LF data has the slope of 2. The data for the ETD-4RK and SSM for $\Delta t > \Delta t_0$ are not shown because the numerical solution there is unstable and hence is dominated by the unphysical, high- k noise. Data for the ETD-LF for $\Delta t \geq 8\Delta t_0$ are not shown, as in this case the error $R_{6,av}$ becomes greater than 1%, which is beyond the accuracy accepted in this study. The $R_{2,av}$ -data for the SSM are not shown because this quantity for the SSM is identically zero.

that of the ETD-LF. Perhaps more surprisingly, even the ETD-4RK will be faster than the SSM in this particular case.

Based on the above, we chose the ETD-4RK with $\Delta t \lesssim \Delta t_0/2$ as the benchmark method in this subsection.

5.1.2 Results

Here we simulated the pure NLS (1) with parameters listed in Section 5.1.1, and we used $t = 3000$ and $N = N_0 \equiv 2^{12}$. As noted there, the maximum resolved wavenumber and the NI threshold are $k_{\max} \approx 85.3$ and $\Delta t_{\text{thresh}} = 4.3 \cdot 10^{-4}$. The corresponding computational spectral window is rather tight, in the sense that during the moments when the spectrum⁷ expanded the most, it rised above the noise floor (see (32)) by about an order of magnitude at the edges of the window (as opposed to not exceeding that floor). Let us emphasize that we made this choice intentionally, because for a wider spectral window, the time step would have to decrease according to (2) if one is to exceed the NI threshold by a fixed factor. A decreased Δt leads to increased accuracy of the results. This is further aided by the fact that for a smaller Δt , the wavenumbers $\pm k_{\pi m}$ are farther from the center of the spectrum and hence the distortions around $\pm k_{\pi m}$ will be weaker (see Fig. 1(b)). We will mention an example later of this subsection. Let us note that we have not observed any effect of spectral aliasing on the quantities that we measured.

In Fig. 7 we compare the statistical quantities for the benchmark method (ETD-4RK with $\Delta t = 2 \cdot 10^{-4}$) and the ETD-LF with $\Delta t = 1.7 \cdot 10^{-3} \approx 4\Delta t_{\text{thresh}}$. The spectrum of the ETD-LF solution looks qualitatively similar to that in Fig. 1(b) (dotted line). Despite such distortions of the

⁷of the benchmark solution

spectrum, the statistical quantities obtained by the ETD-LF match fairly well with the benchmark ones. If one uses $\Delta t = 9 \cdot 10^{-4} \gtrsim 2\Delta t_{\text{thresh}}$ with the ETD-LF, then $C_x(\delta x)$ becomes almost indistinguishable from the benchmark one on the scale of Fig. 7(c), while in $C_t(\delta t)$ conspicuous deviations begin to occur past $\delta t = 2$, similarly to the green curve in Fig. 7(d) (see below for its meaning). Decreasing the time step by an additional factor of two, to $\Delta t = 4 \cdot 10^{-4}$, makes $C_t(\delta t)$ quite accurate for the entire range shown in that figure. On the other hand, both PDFs require a much smaller Δt , about $\Delta t_{\text{thresh}}/4$, to become indistinguishable (on a logarithmic scale) from the corresponding benchmark curves.




Figure 7: Comparison of the four statistical quantities for the pure NLS (Section 5.1.2), obtained by the ETD-4RK with $\Delta t = 2 \cdot 10^{-4}$ (thick solid) and ETD-LF with $\Delta t = 1.7 \cdot 10^{-3}$ (thin solid). In addition, in (d), the green lines show the result obtained by the ETD-LF with $N = 3 \cdot 2^{11}$ and $\Delta t = 1.7 \cdot 10^{-3} / (1.5)^2$ (see (34)).

Now recall that the spectral window for which we have shown the above results is rather tight. In fact, it took us some experimentation, where we had gradually narrowed the spectral window and determined when the results would change significantly. Instead, one may choose to stay on the safe side and use a 50-percent wider window (with $N = 3 \cdot 2^{11}$ grid points), while still using the time step four times greater than the adjusted NI threshold as per (34). Then $C_x(\delta x)$ becomes almost indistinguishable from the benchmark result, while $C_t(\delta t)$ improves considerably: see the

green curves in Fig. 7(d). The PDFs, however, do not visibly change.

Finally, we ran a ETD-4RK simulation with $N = 2^{12}$ and $\Delta t = 4 \cdot 10^{-4}$, which is just below the NI threshold. This appears to affect the PDF($|u_x|^2$) and the two correlation functions rather insignificantly. The PDF($|u|^2$) also remains almost the same up to a value 10^{-1} and then begins to exceed the benchmark curve; at $|u|^2 = 20$ it exceeds it by almost an order of magnitude.

The above leads to the following comparison of computational speeds.

- If one is willing to accept a PDF($|u|^2$) with a somewhat inaccurate “tail”, as described in the previous paragraph, and also the ETD-LF results with $\Delta t = 1.7 \cdot 10^{-3}$ and $N = 2^{12}$, then the speed gain of the ETD-LF over the ETD-4RK is about $4 \times (3.5 \dots 4) \approx 15$ times. Here the former factor is due to $1.7 \cdot 10^{-3} / \Delta t_{\text{thresh}} \approx 4$ and the latter one is due to the ETD-LF being faster than the ETD-4RK.
- If it is essential to resolve the “tail” of $C_x(\delta x)$ accurately, then one needs to use the ETD-LF with a twice smaller step size, $\Delta t \approx 2\Delta t_{\text{thresh}}$. In this case, the speed gain is 7 to 8 times.
- If it is essential to resolve the fine oscillatory structure of the “tail” of $C_t(\delta t)$, then further reduction of Δt for the ETD-LF is required, and its speed gain reduces to the mere speed gain of the ETD-LF over ETD-4RK with the same time step (i.e. 3.5 to 4).
- If it is essential to resolve the PDFs, and especially PDF($|u|^2$), very accurately, then the ETD-LF provides no to little gain over the ETD-4RK, since the former would require a 3–4 times smaller time step than the latter.
- Finally, a speed gain of the ETD-LF over the SSM would amount only to the excess of Δt over Δt_{thresh} , as both methods have about the same speed for the same Δt . If only the correlation functions need to be approximated, then the speed gain of the ETD-LF over the SSM is two to four times. If, however, one needs to resolve the PDFs, and the accuracy of the ETD-LF results shown in Fig. 7(a,b) is sufficient, then the speed gain is comparable or greater than that over the ETD-4RK. This follows from our observation in Section 5.1.1 that the SSM requires $\Delta t \sim \Delta t_{\text{thresh}}/10$ to resolve the PDF($|u|^2$) with the same accuracy as the ETD-LF with $\Delta t \sim \Delta t_{\text{thresh}}$.

5.2 Focusing NLS with potential

We simulated (1) with $V(x) \equiv 0.25 \sin(6x + \varphi)$ and $\gamma = 2$ in a domain of length $L = 32\pi$ using $N = 2^{12}$ grid points. The corresponding $k_{\text{max}} \approx 128$ and hence $\Delta t_{\text{thresh}} = 1.9 \cdot 10^{-4}$. The following considerations led us to choose the above parameters.

As far as the choice of $V(x)$ goes, we wanted the effect caused by the potential to satisfy two conditions. On one hand, it had to make the statistical quantities to be substantially different from those in the pure NLS case. From our study [35], we knew that a correlated multiplicative noise in the NLS considerably increases the probability that the amplitude of the solution of (1) can

reach high values. While our potential is deterministic and time-independent, the same mechanism of solution enhancement has appeared to also apply for it. Thus, after a quick experimentation, we chose a periodic potential with wavelength close to 1. On the other hand, we had to keep its amplitude, 0.25, to be relatively small; otherwise the generation of high-amplitude events was so strong that it would require a much wider computational spectrum, and hence a much smaller time step, for which simulations would have become prohibitively time-consuming (see below).

The need to include the phase φ into $V(x)$ will be explained when we discuss the benchmarking simulations below.

The computational spectral window was, again, chosen to be tight: the solution’s spectrum at $|k| = k_{\max}$, at the moments of its greatest expansion, exceeded the noise floor by almost three orders of magnitude. However, such bursts were also isolated and very rare — just a few over $t = 3000$, and the solution’s spectrum at the edges of the spectral window stayed below the noise floor most of the time. As we have noted earlier, using a wider computational spectral window, and thus a smaller time step, would only improve the results of the ETD-LF whose time step exceeds the NI threshold by a fixed factor.

5.2.1 Benchmarking simulations

In benchmarking simulations ($t = 1000$), we set $\varphi = 0$ in $V(x)$ and observed that results obtained by any of the methods — ETD-4RK, ETD-LF, and SSM — could depend on the time step value even for sufficiently small time steps: $\Delta t = \Delta t_{\text{thresh}}/2$ and $\Delta t = \Delta t_{\text{thresh}}/4$. This should be contrasted with the situation for the pure NLS, where the ETD-4RK results with such time steps were almost indistinguishable from one another. The greatest difference occurred in the time correlation function, whose values after the first “lobe” could differ by up to 100%. Widening the computational spectrum by using $N = 3 \cdot 2^{11}$ instead of 2^{12} grid points for the same $L = 32\pi$ (and reducing Δt accordingly) would not improve the situation.

Our explanation of this inconsistency among simulations was this. The dynamics of the solution of the NLS, both with and without potential, is chaotic. Hence, using different time steps and/or numerical methods *should* be expected to lead to solutions that at a long time may be completely different. However, the statistical quantities would remain close, as long as averaging of the solution takes place during its evolution. In the pure NLS case, the solution is uniform in space, and hence the averaging occurs over *both space and time*. On the contrary, in the presence of the potential, averaging over space occurs less efficiently, because the solution is not uniform in space.

To make the space averaging more thorough, one could increase the spatial domain. We, however, adopted a different strategy. Since the initial condition (32) is the same in all simulations, one generates different solutions by simply shifting the potential. Therefore, for each method and each Δt , we ran eight simulations with

$$V(x) = 0.25 \sin(6x - 2j\pi/8), \quad j = 0, 1, \dots, 7, \quad (37)$$

and then averaged all statistical quantities over these eight cases. Using considerably more φ -values

in the interval $[0, 2\pi)$ would have improved the spatial averaging. However, it would have also been time-prohibitive, considering that the ETD-4RK simulation with $\Delta t = \Delta t_{\text{thresh}}/4$ for $t = 3000$, reported below, took more than 10 days. For the same reason, we did not use simulations with $N = 3 \cdot 2^{11}$ in our comparison of results.

Thus, for the benchmarking, we verified that at $t = 1000$, the PDFs and the spatial correlation function obtained by the ETD-4RK, ETD-LF, and SSM with $\Delta t = \Delta t_{\text{thresh}}/2$ and $\Delta t = \Delta t_{\text{thresh}}/4$ were all consistent among themselves after averaging over the eight values of φ . The time correlation functions were still not very close even after the averaging, but at least they showed the same trend as δt increased. Unlike the pure NLS case, here the SSM results were as accurate as the ETD-4RK ones, and so the SSM is the main competitor of the ETD-LF in terms of speed.

5.2.2 Results

Figure 8 shows the comparison of the four statistical quantities obtained by the ETD-4RK with $\Delta t = \Delta t_{\text{thresh}}/2$ and $\Delta t = \Delta t_{\text{thresh}}/4$, by the ETD-LF with $\Delta t = 6.0 \cdot 10^{-4} > 3\Delta t_{\text{thresh}}$, and by the SSM with $\Delta t = \Delta t_{\text{thresh}}$. We presented ETD-4RK results for two different Δt 's in order to illustrate that the values of the time-correlation function deviate from each other even in these cases, which presumably would be more accurate than the lower-order methods (ETD-LF and SSM) with greater Δt 's. Moreover, the deviations between the values of $C_t(\delta t)$ for the two ETD-4RK cases are comparable in size to those for the ETD-LF and SSM. Therefore, we conclude that the ETD-LF with a step size exceeding the NI threshold by more than three times is almost as accurate as the SSM with $\Delta t = \Delta t_{\text{thresh}}$. Thus, the ETD-LF here provides a three-fold speedup compared to the SSM.

Evolution of conserved quantities: I_2 , given by (35a), and a modified form of the Hamiltonian (35c):

$$I_{4,\text{mod}} = \int_0^L (-|u_x|^2 + (\gamma/2)|u|^4 + V(x)|u|^2) dx, \quad (38)$$

obtained by the ETD-4RK with $\Delta t = \Delta t_{\text{thresh}}/4$, is similar to that shown in Fig. 5(a). The relative fluctuations of I_2 and $I_{4,\text{mod}}$ do not exceed 0.001% at $t = 3000$ in this case. Relative fluctuations of these quantities obtained by the ETD-LF with $\Delta t \gtrsim 3\Delta t_{\text{thresh}}$ grow up to 0.05% and 1%, respectively. For the SSM with $\Delta t \lesssim \Delta t_{\text{thresh}}$, the relative fluctuations of $I_{4,\text{mod}}$ increased to 1%; the quantity I_2 is conserved exactly.

5.3 Focusing NLS with linear pumping and nonlinear damping

The purpose of this example is to illustrate that the ETD-LF with the stabilization step, such as (28), is suitable for long-time simulation of problems involving dissipation. This may not be obvious given the well-known fact that the LF scheme itself is unstable for such problems. However, the instability of the LF scheme is caused by the parasitic error $\tilde{u}^{(-)}$, discussed in Section 3.1. The stabilization step suppresses growth of such an error and thereby makes the ETD-LF stable.




Figure 8: (Color online) Comparison of the four statistical quantities for the NLS with potential (Section 5.2), averaged over eight values of φ (see (37)). Black: ETD-4RK (thinner — $\Delta t = \Delta t_{\text{thresh}}/4$; thicker — $\Delta t = \Delta t_{\text{thresh}}/2$); green: ETD-LF; red: SSM. In panels (a)–(c) some of the curves appear indistinguishable. Note that all of the horizontal scales are different from the corresponding scales in Fig. 7.

We simulated (1) with $V(x) = -i\alpha$ and $\gamma = 2 + i\alpha$, where we set $\alpha = 0.05$. A similar equation was considered, e.g., in [36]. Other simulation parameters were as in Section 5.1: $L = 48\pi$ and $N = 2^{12}$, so that $\Delta t_{\text{thresh}} = 4.3 \cdot 10^{-4}$. The spectra of the solution at $t = 1000$ and $t = 3000$ are shown in Fig. 9(a), and a sample of the waveform is shown in Fig. 9(b). One can see that the given combination of the linear pumping and nonlinear damping terms leads to the development of a highly oscillatory solution, whose frequency increases with time. One can also see that for the parameters chosen for our example, the computational spectral window is sufficiently tight in the sense discussed in the previous subsections.

From benchmarking simulations, performed for $t = 1000$, we have found that the four statistical quantities obtained by the ETD-4RK with $\Delta t \lesssim \Delta t_{\text{thresh}}$ and $\Delta t \lesssim \Delta t_{\text{thresh}}/2$ are almost indistinguishable. Therefore, they can be used as benchmark results. In fact, we have also found that the results obtained by the SSM with $\Delta t \lesssim \Delta t_{\text{thresh}}$ are very similar to the ETD-4RK results. In Fig. 10 we present the four statistical quantities obtained by the ETD-4RK with $\Delta t \lesssim \Delta t_{\text{thresh}}/2$,




Figure 9: Spectra (a) and waveforms (b) of the solution of the NLS with linear pumping and nonlinear damping (Section 5.3) at different times.

SSM with $\Delta t \lesssim \Delta t_{\text{thresh}}$, and ETD-LF with $\Delta t \sim 3.5\Delta t_{\text{thresh}}$. The latter results are seen to be very close to those obtained by the methods with $\Delta t < \Delta t_{\text{thresh}}$. Thus, the computational time of the ETD-LF is approximately 3.5 times less than that of its closest competitor, which in this case is the SSM.

5.4 Defocusing NLS with potential

The purpose of this example is to illustrate that the ETD-LF with a stabilization step, such as (28), is suitable for long-time simulation of yet another type of problems where a LF-based method *without* stabilization would not be applicable. In Section 3.1 we explained that the parasitic error $\tilde{u}^{(-)}$, produced by the ETD-LF when simulating the defocusing NLS, leads to low- k NI. The stabilization step with $t_{\text{stabilize}} = 1$ has been found to suppress the parasitic error, which made the method suitable for simulation of the defocusing NLS.

We simulated (1) with $\gamma = -2$ and

$$V(x) = 2 \sin(4x - 2j\pi/8), \quad j = 0, 1, \dots, 7, \quad (39)$$

where considering several values of the phase of the potential served the same purpose as in Section 5.2 (see also below). In the defocusing NLS the dynamics of the solution is relatively “laminar”, compared to the “volatile” dynamics in the focusing NLS. As a consequence, the spectrum of the solution is much narrower in the defocusing case. We had found by experimentation that presence of a potential such as (39) makes the spectrum a few tens of percent broader, and hence we included it in our example. The sample spectrum and waveform of the solution are shown in Fig. 11. Other simulation parameters were as follows: $L = 100\pi$ and $N = 2^{12}$, so that $\Delta t_{\text{thresh}} = 1.9 \cdot 10^{-3}$. Unlike in the previous examples, the computational spectral window is not tight for the solution obtained by benchmark methods, although it is still tight for the solution obtained by the ETD-LF with $\Delta t > \Delta t_{\text{thresh}}$.




Figure 10: (Color online) Comparison of the four statistical quantities for the NLS with linear pumping and nonlinear damping (Section 5.3): ETD-4RK with $\Delta t = 2.0 \cdot 10^{-4}$ (thick solid); ETD-LF with $\Delta t = 1.5 \cdot 10^{-3}$ (thin solid); SSM with $\Delta t = 4.0 \cdot 10^{-4}$ (dashed). In panels (a)–(c) some of the curves appear indistinguishable. Note that the range of $|u_x|^2$ is much greater than that of $|u|^2$. This indicates a highly oscillatory solution, in agreement with Fig. 9(b).

As in Section 5.3, benchmarking simulations at $t = 1000$ revealed that the four statistical quantities, including $C_t(\delta t)$, obtained by the ETD-4RK with $\Delta t \lesssim \Delta t_{\text{thresh}}$ and $\Delta t \lesssim \Delta t_{\text{thresh}}/2$ are almost indistinguishable. In particular, the averaging of the solution over space, facilitated by taking multiple shifted replicas of the potential as per (39), occurs much more efficiently for the defocusing NLS than for the focusing one, considered in Section 5.2. Therefore, below we use the ETD-4RK results with $\Delta t \lesssim \Delta t_{\text{thresh}}/2$ as the benchmark.

In Fig. 12 we present the four statistical quantities obtained by the ETD-4RK with $\Delta t \lesssim \Delta t_{\text{thresh}}/2$, the SSM with $\Delta t \lesssim \Delta t_{\text{thresh}}$, and the ETD-LF with $\Delta t \sim 4\Delta t_{\text{thresh}}$. The latter results, except for $C_t(\delta t)$, are seen to be very close to the benchmark and SSM ones. We also verified that the time correlation function obtained by the ETD-LF with $\Delta t \sim 3\Delta t_{\text{thresh}}$ lies approximately midway between that obtained with $\Delta t \sim 4\Delta t_{\text{thresh}}$ and the benchmark one. Thus, the computational time of the ETD-LD is approximately 3 – 4 times (depending on the desired accuracy) less than that of its closest competitor, which, as in Section 5.3, is the SSM. It is possible that the ETD-LF with a

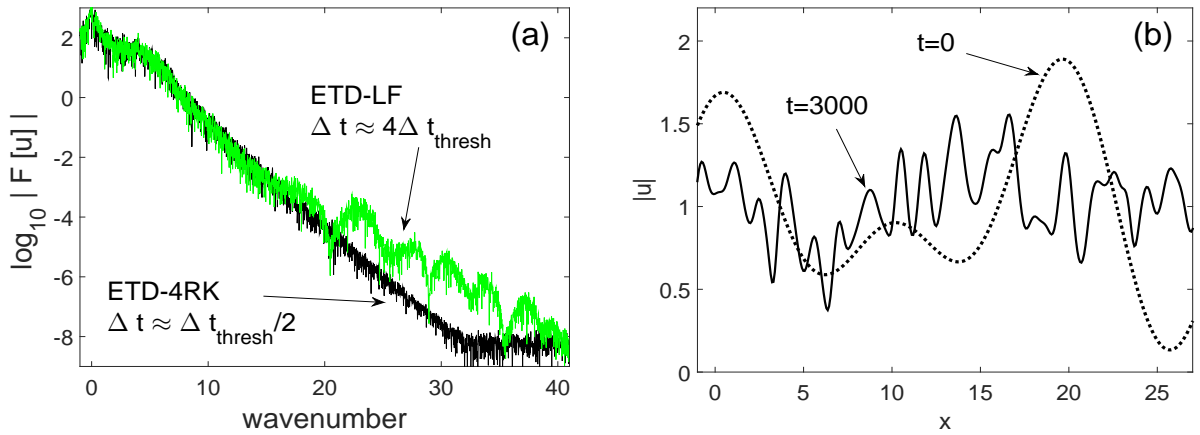


Figure 11: (a): Spectra for the ETD-4RK with $\Delta t = 0.9 \cdot 10^{-3}$ and ETD-LF with $\Delta t = 7.5 \cdot 10^{-3}$ at $t = 3000$; and (b): waveforms (for the ETD-4RK) of the solution of the defocusing NLS with potential (Section 5.4).

still greater time step can also provide reasonably accurate results. However, we did not push such a comparison because Δt_{thresh} in this case is already quite large; then using $\Delta t > 4\Delta t_{\text{thresh}}$ will likely result in loss of accuracy, even though the NI may still be suppressed.

Evolution of conserved quantities I_2 and $I_{4, \text{mod}}$ is similar to that reported in Section 5.2, except that their relative fluctuations are more than an order of magnitude smaller than their counterparts in that subsection.

6 Numerical experiments in two spatial dimensions

In higher spatial dimensions, simulations typically take longer than in one dimension (even though some of the simulations considered in Sections 5.1 and 5.2 took several days). Therefore, speeding up higher-dimensional simulations may be even more essential than one-dimensional ones. Below we will first argue that in $D > 1$ dimensions, the gain in the computational speed provided by the ETD-LF over the SSM can be up to a factor of D higher than in one dimension. Then we will present two numerical examples supporting this claim. In the process, we will propose a simple modification of the ETD-LF which will enable the method to better preserve the conserved quantities of the (generalized) NLS.

We will explain the above claim about the computational speed gain using $D = 2$ as an example. The explanation hinges on a few assumptions. First, for simplicity, we will assume that the maximum wavenumbers along both dimensions are equal: $k_{\text{max}, x} = k_{\text{max}, y} \equiv k_{\text{max}}$; this is when the upper bound, D , for the speed gain can be achieved (see below). Thus, the spectral computational domain is a square

$$|k_x| \leq k_{\text{max}}, \quad |k_y| \leq k_{\text{max}}. \quad (40)$$

Second, we will assume that this domain is tight in the sense defined at the beginning of Section 5.1.2 and as was the case in all examples in Section 5 and as will be in both examples considered

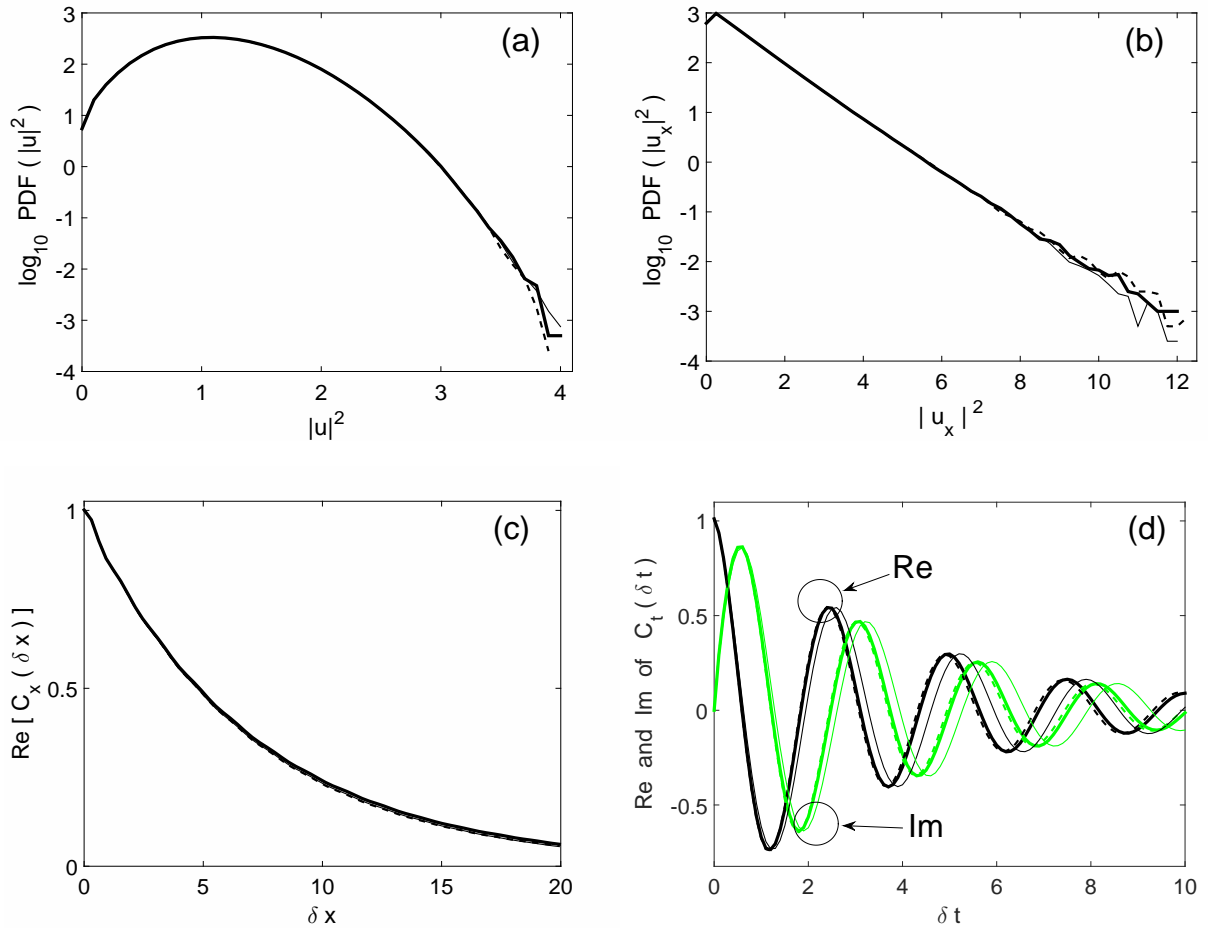


Figure 12: Comparison of the four statistical quantities for the defocusing NLS with potential (Section 5.4). The line styles are similar to those in Fig. 10: ETD-4RK with $\Delta t = 0.9 \cdot 10^{-3}$ (thick solid); ETD-LF with $\Delta t = 7.5 \cdot 10^{-3}$ (thin solid); SSM with $\Delta t = 1.8 \cdot 10^{-3}$ (dashed). In panels (a)–(c) some of the curves appear indistinguishable. Note that the horizontal scales are different from the corresponding scales in Figs. 7, 8, and 10.

below. Third, we will assume that the solution’s spectrum has an (at least approximate) circular symmetry; this indeed occurs in many applications (see the numerical examples below) and does *not* imply any symmetry in the spatial domain. Finally, our comparison between $D = 1$ and $D = 2$ (or, more generally, an arbitrary D) will pertain to such solutions for which the ratio of their spectral width to k_{\max} is approximately the same for all values of D .

The last three assumptions have the following important corollary: The accuracy of the numerical solution is determined by the ratio (k_π/k_{\max}) . Here $k_\pi = \sqrt{\pi/\Delta t}$ (see (18)) is the approximate lowest location of wavenumbers where conspicuous distortions of the ETD-LF solution’s spectrum occur: see Figs. 1(b) and 11(a). To justify the above corollary, note that the denominator, i.e., k_{\max} , in that ratio, is *not* the largest magnitude of the wavenumber vector in the computational domain (40), but rather the radius of the largest circle inscribed in that domain. Further note that accurate resolution of Fourier harmonics outside that circle,

$$k_x^2 + k_y^2 = k_{\max}^2, \quad (41)$$

would not increase the accuracy of the solution, because not all of them are contained inside the domain (40). Therefore, recalling the assumed circular symmetry of the solution's spectrum, one concludes that distortions in those harmonics do not affect the numerical solution's accuracy (as long as those harmonics do not increase exponentially in time, which would signify an NI).

Let σ_D be the gain in the computational speed of the ETD-LF over the SSM in D dimensions. We want to demonstrate that $\sigma_D/\sigma_1 = D$. As in most examples in Section 5, one has $\sigma_D = \Delta t/\Delta t_{\text{thresh},D}$, where Δt is a time step yielding a certain accuracy (or, equivalently, a certain amount of distortion in the numerical solution) and $\Delta t_{\text{thresh},D}$ is the NI threshold in D dimensions. Given the above expression for k_π , one has $\sigma_D = (\pi/k_\pi^2)/\Delta t_{\text{thresh},D}$. Next, in two dimensions, the NI threshold is given, instead of (22), by

$$\Delta t_{\text{thresh},D} = \pi/(k_{\text{max},x}^2 + k_{\text{max},y}^2) = \pi/(2k_{\text{max}}^2) \equiv \pi/(D k_{\text{max}}^2), \quad (42)$$

because the numerically unstable harmonics first occur at wavenumbers with $(|k_x|, |k_y|) = (k_{\text{max}}, k_{\text{max}})$.

Using the above information, one deduces the computational speed gain in D dimensions compared to one dimension:

$$\frac{\sigma_D}{\sigma_1} = \frac{(\pi/k_\pi^2) / (\pi/(D k_{\text{max}}^2))}{(\pi/k_\pi^2) / (\pi/k_{\text{max}}^2)} = D, \quad (43)$$

which proves our claim. Note that above we have essentially used one of our earlier assumptions that the same Δt produces similar distortions to the numerical solution in D and one dimensions.

Recall that in the one-dimensional examples in Section 5 we observed a typical speed gain of the ETD-LF over SSM to be a factor of three to four. Therefore, in two dimensions we expect the gain to be a factor of six to eight. The two examples considered below support this prediction. They both refer to the generalized NLS

$$iu_t + (u_{xx} + u_{yy}) + G(|u|^2)u = 0, \quad (44)$$

where two different forms of the function G are considered in Sections 6.1 and 6.2.

6.1 Defocusing cubic nonlinearity

It is well known that in the focusing cubic NLS (44), with $G(|u|^2) = \gamma|u|^2$ and $\gamma > 0$, the solution undergoes a collapse in finite time. Therefore, as our first example, we considered the case $\gamma = -2$. The initial condition was taken to be similar to (32):

$$\begin{aligned} u(x, y, 0) = & 1 + 0.24 \cos(15\Delta k x + 0.7) + 0.33 \sin(21\Delta k x + 0.2) + 0.42 \cos(31\Delta k x) + \xi(x) \\ & + 0.27 \cos(16\Delta k y - 0.7) + 0.36 \sin(22\Delta k y - 0.6) + 0.39 \cos(30\Delta k y + 0.9). \end{aligned} \quad (45)$$

Its exact form, however, is not important for the conclusions to be drawn, as long as it presents some sufficiently deeply modulated, pseudo-random profile. We chose the spatial domain to be a square $[-L/2, L/2] \times [-L/2, L/2]$ with $L = 12\pi$ and simulated (44) until $t = 300$. We have empirically found that the solution's spectrum gradually broadens, and to resolve it up to $t = 300$ we needed

$N = 2^{10}$ grid points along each dimension. (In this case the longest of the simulations reported below, which used the ETD-4RK, took approximately one week.) The corresponding solution and its spectrum at $t = 300$ are shown in Fig. 13⁸. Note that $k_{\max} = 85.3$, so that according to (42), $\Delta t_{\text{thresh},2} = 2.2 \cdot 10^{-4}$. When using the ETD- and IF-LF, we employed the stabilization step (28) with $t_{\text{stabilize}} = 1$, as in Section 5.4.

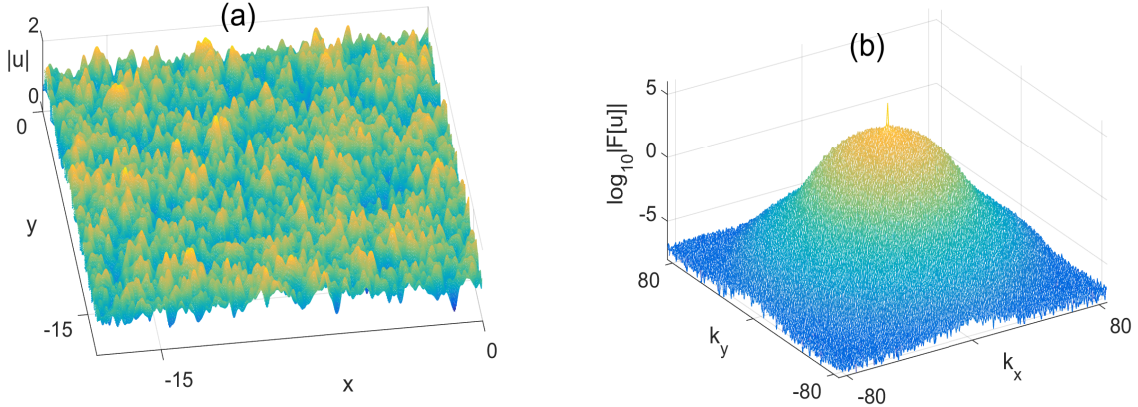


Figure 13: Magnitude (a) and spectrum (b) of the solution of (44) with $G(|u|^2) = -2|u|^2$ at $t = 300$; other parameters are listed in Section 6.1. The solution is uniform in space, and therefore only quarter of the spatial computational domain is shown.

The statistical quantities: the two PDFs and two correlation functions, were computed exactly as described in Section 5, except that the summation was performed over both dimensions of the spatial grid. The data for the averaging were collected over the last $t = 50$ units of the evolution. Just for simplicity, we chose to continue computing the x -correlation function $C_x(\delta x)$, as opposed to the more natural circularly symmetric correlation function $C_r(\delta r)$, where $\delta r \equiv \sqrt{\delta x^2 + \delta y^2}$.

We monitored the conserved quantities of the numerical solution: mass, x - and y -components of the momentum, and the Hamiltonian, defined as in (35)(a-c), except the integration is now in two dimensions. When we applied the SSM or IF-LF with $\Delta t = 2 \cdot 10^{-4} \lesssim \Delta t_{\text{thresh},2}$, those quantities fluctuated around the mean zero by less than $10^{-3}\%$. For the ETD-4RK with the same Δt , they exhibited a systematic drift, but it was less than $10^{-5}\%$ over the considered time, and hence could be ignored. However, for the ETD-LF with $\Delta t = 1.25 \cdot 10^{-3}$ (see below), the relative change of the Hamiltonian is over 1%, and, most disturbingly, it is a systematic drift (unlike in all of the one-dimensional examples in Section 5): see the dashed line in Fig. 14(f). The rate of the drift decreased approximately quadratically with Δt , but still remained conspicuous even for $\Delta t \lesssim \Delta t_{\text{thresh},2}$. The relative changes of the mass and momentum were one–two orders of magnitude less than the corresponding change of the Hamiltonian, but also exhibited a systematic

⁸We also carried out simulations of (44) with a potential term, $V(x,y)u$, and found the solution and its spectrum to be qualitatively similar to those obtained for $V(x,y) \equiv 0$. We could not, however, perform the averaging of the results over the potential, as described in Section 5.2.1, because it would have been time prohibitive given our computational resources (two personal computers running Matlab).

drift. By varying $t_{\text{stabilize}}$, we verified that this drift is *not* due to the stabilization step (28).

We were unable to find an explanation for this drift. (The most puzzling was the fact that it occurs in two dimensions, but not in one.) However, given that this drift is absent in the IF-LF, we have found a means to substantially reduce it in the ETD-LF. Namely, we conjectured that it was due to the strength of the nonlinear interaction, $\widehat{\mathcal{G}} = \text{sinc}((k_x^2 + k_y^2)\Delta t)$, being substantially less than unity (see Section 3.1 and, specifically, the solid line in Fig. 2) in those parts of the computational spectral domain where the solution’s Fourier harmonics still contained “enough” energy. Now note that instead of the above expression for $\widehat{\mathcal{G}}$, one could have used an *empirical* expression

$$\widehat{\mathcal{G}}_{\text{mod}} = \text{sinc}(\pi [(k_x^2 + k_y^2)\Delta t / \pi]^m), \quad m > 1. \quad (46)$$

In Fig. 2 we show the trace (in one dimension) of this expression for $m = 3$; it is seen to be close to unity over a significantly larger (compared to $m = 1$) part of the spectral domain where the solution has most of its content. Yet, $\widehat{\mathcal{G}}_{\text{mod}}$ vanishes at all $k_{\pi l}$, $l = 1, 2, \dots$ and hence guarantees suppression of the high- k NI.

In the simulations with ETD-LF reported below, we used $\widehat{\mathcal{G}}_{\text{mod}}$ with $m = 3$ instead of $\text{sinc}((k_x^2 + k_y^2)\Delta t)$ in (8). For the same Δt , this substantially reduced the amount of drift of conserved quantities: compare the dashed and thin solid lines in Fig. 14(f).

We compared the PDFs and correlation functions obtained by the ETD-LF with $\Delta t = 1.25 \cdot 10^{-3} \approx 6\Delta t_{\text{thresh},2}$ and with $\Delta t = 1.6 \cdot 10^{-3} \lesssim 8\Delta t_{\text{thresh},2}$ with the same quantities obtained by the ETD-4RK and SSM with $\Delta t = 2 \cdot 10^{-4} \lesssim \Delta t_{\text{thresh},2}$; see Fig. 14(a–e). The PDFs are seen to agree quite well for all cases. The spatial correlation functions obtained by the SSM, ETD-LF with $\Delta t = 1.25 \cdot 10^{-3}$, and ETD-LF with $\Delta t = 1.6 \cdot 10^{-3}$ differ from that obtained by the ETD-4RK by about 1, 2, and 3% for sufficiently large δx . Since the temporal correlation functions in all four cases appear to be very close in Fig. 14(d), we compared (in panel (e)) the deviations of these functions from the “asymptotic”⁹ functional form of $C_t(\delta t)$. Parameters of that form: $C_{t,\text{asympt}}(\delta t) = 0.82 \exp[i(2.9\delta t + 0.005)]$, were found by an inspection of the $C_t(\delta t)$ obtained by the ETD-4RK close to $\delta t = 4$. From Fig. 14(e) we see that the temporal correlation function obtained by the ETD-LF with $\Delta t \approx 6\Delta t_{\text{thresh},2}$ stays within about 2% of the “benchmark” one (obtained by the ETD-4RK) up to $\delta t \approx 2$; as one increases Δt up to $8\Delta t_{\text{thresh},2}$, larger deviations are observed as early as for $\delta t \approx 1$.

Let us also mention that the statistical quantities obtained by the “original” ETD-LF (i.e., using (46) with $m = 1$) with $\Delta t \approx 6\Delta t_{\text{thresh},2}$ were within 1% from those obtained by the modified ETD-LF with the same time step. This may be consistent with the fact (Fig. 14(f)) that the Hamiltonian in the original ETD-LF drifts only by about 1%. However, in longer simulations, deviations can become greater, and then the use of the modified ETD-LF over the original one may become imperative.

⁹In passing, let us note that the presence of nonvanishing correlations in space in time, as seen in panels (c) and (d), indicates that the solution has a hidden “plane-wave background” that oscillates in time as a whole. The presence of this background is also evident from the spectral peak at $k_x = k_y = 0$ in Fig. 13(b).

To summarize the results of this numerical experiment, we have shown that if one is willing to tolerate moderately small (up to 3%) errors in the statistical quantities, one can reduce the computational time by a factor of 8, by using the ETD-LF with a time step that is 8 times greater than the NI threshold for the SSM. Using the ETD-LF with a smaller Δt and thereby achieving a speed gain of the factor of 6 instead of 8, one can reduce the error to about 2%. For reference, the SSM simulation in this section took approximately 40 hours on a personal computer running Matlab; the more accurate ETD-4RK simulation required more than four times as long.

6.2 Focusing saturable nonlinearity

As our second example, we considered the generalized NLS (44) with focusing, but saturable nonlinearity: $G(|u|^2) = 2|u|^2/(1+|u|^2)$. Note that for large $|u|^2$, this function saturates at the constant value 2, and this arrests a collapse. Somewhat surprisingly, the solution's spectrum in the saturable focusing case turned out to be narrower than in the cubic defocusing case; see Fig. 15(b). This allowed us to use a smaller spectral computational domain and hence a larger time step and a longer computational time.

The initial condition and the spatial computational domain are the same as in Section 6.1. We used $N = 3 \cdot 2^8$ grid points along each dimension, which resulted in $k_{\max} = 63.8$ and $\Delta t_{\text{thresh}, 2} = 3.8 \cdot 10^{-4}$. We computed the solution until $t = 500$. When using the ETD-LF, we employed the stabilization step (28) with $t_{\text{stabilize}} = 5$. Computation of the statistical quantities was performed exactly as in Section 6.1. The expression for $C_{t, \text{asympt}}(\delta t)$ is now $0.43 \exp[i(1.7\delta t - 0.3)]$.

We used the ETD-4RK and SSM with $\Delta t = 3.7 \cdot 10^{-4} \lesssim \Delta t_{\text{thresh}, 2}$ and the modified (as per (46) with $m = 3$) ETD-LF with $\Delta t = 2.5 \cdot 10^{-3} \gtrsim 6\Delta t_{\text{thresh}, 2}$ and $\Delta t = 3 \cdot 10^{-3} \approx 8\Delta t_{\text{thresh}, 2}$. Panels (a)–(e) in Fig. 16 demonstrate that the agreement between the ETD-LD results with $\Delta t \gtrsim 6\Delta t_{\text{thresh}, 2}$ and the ETD-4RK ones is the same or even slightly better than between the SSM and ETD-4RK results (see panel (e)). For $\Delta t \approx 8\Delta t_{\text{thresh}, 2}$, the ETD-LF and SSM results are approximately equidistant from the ETD-4RK ones. The overall better performance of the ETD-LF in this example, compared to that considered in Section 6.1, appears to be due to the fact that the “bulk” of the solution's spectrum is narrower relative to the spectral domain's width¹⁰ in Fig. 15(b) than in Fig. 13(b). The same reason appears to be behind a much smaller drift of the Hamiltonian: compare Figs. 14(f) and 16(f).

To summarize the results of this numerical experiment, the ETD-LF provided a gain in the computational speed between factors of 6 and 8 compared to the SSM, for a similar accuracy.

7 Conclusions

We proposed to use an ETD method based on the LF scheme to integrate the generalized NLS (1) or (44). The equation of the method is given by (8), or, more generally, by (7). Adding a

¹⁰Even though in both cases, the domains are tight in the sense defined in Section 5.1.2.

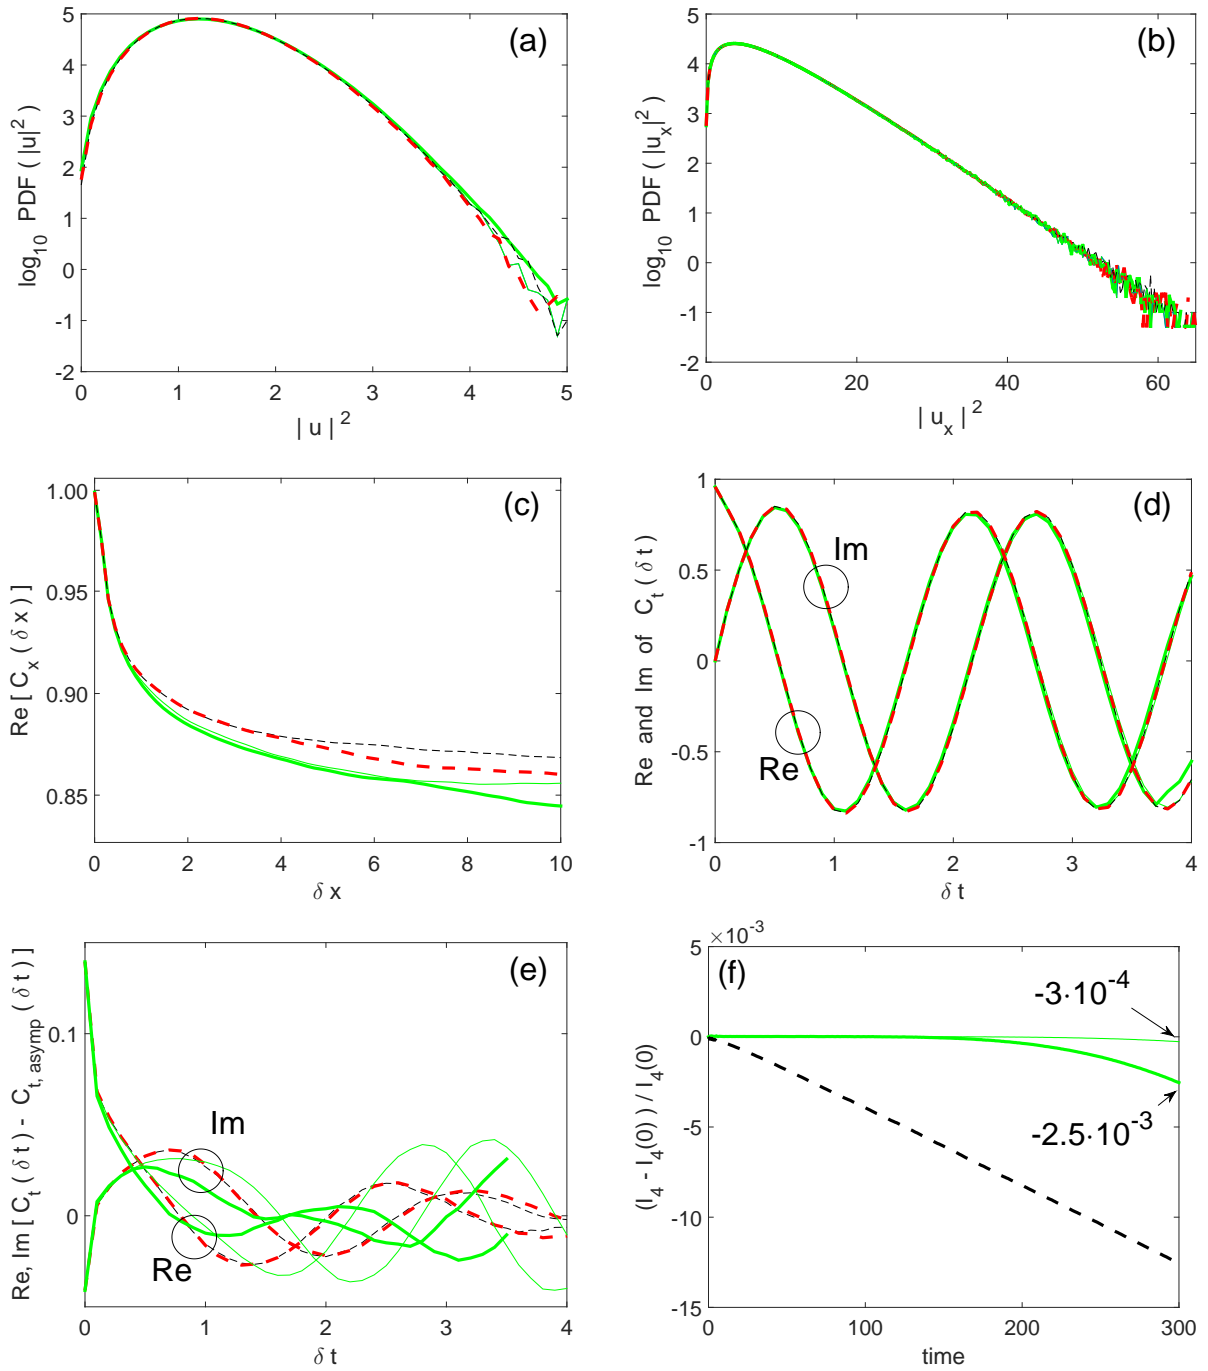


Figure 14: (Color online) Panels (a)–(e): Comparison of the four statistical quantities for the defocusing cubic NLS (Section 6.1). ETD-4RK (black, thin dashed) and SSM (red, thick dashed) with $\Delta t = 2 \cdot 10^{-4}$; modified (as per (46) with $m = 3$) ETD-LF with $\Delta t = 1.25 \cdot 10^{-3}$ and with $\Delta t = 1.6 \cdot 10^{-3}$ (green solid, thin and thick, respectively). Note the vertical scales in panels (c) and (e). In panel (d), a small “dent” in the thick green curve around $\delta t = 3.6$ is a computational artifact. Panel (f): Relative change of the Hamiltonian computed by the original ETD-LF (dashed) with $\Delta t = 1.25 \cdot 10^{-3}$ and the modified ETD-LF with $\Delta t = 1.25 \cdot 10^{-3}$ (thin solid) and $\Delta t = 1.6 \cdot 10^{-3}$ (thick solid).

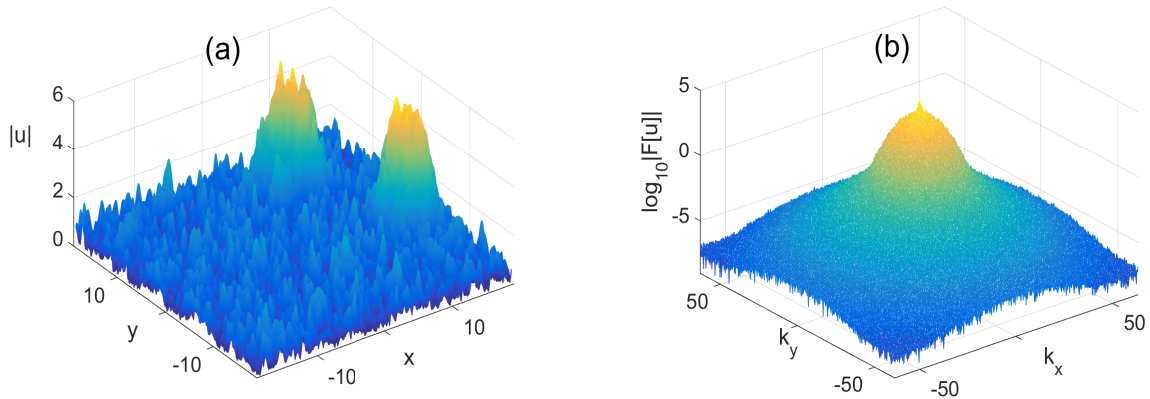


Figure 15: Magnitude (a) and spectrum (b) of the solution of (44) with $G(|u|^2) = 2|u|^2/(1 + |u|^2)$ at $t = 500$; other parameters are listed in Section 6.2.

stabilization step, such as (28), that suppresses a parasitic numerical error, we eliminated NI of Fourier harmonics with wavenumbers $k = O(1)$, which would otherwise occur in any method based on the LF scheme.

We have predicted and demonstrated that the NI of high- k harmonics is eliminated (or strongly suppressed) by the use of the ETD, as opposed to, say, IF, method; compare (8) with (10). This suppression of the NI has allowed us to use time steps that exceeded the NI threshold 3 to 4 times and still obtain sufficiently accurate characteristics of the *statistical* behavior of the solution. In the examples considered in Sections 5.2 – 5.4 this resulted in the gain of the computational speed (compared to the SSM) by a factor of 3 to 4; the exact values are stated at the end of the respective subsections). In the case of the pure NLS, considered in Section 5.1, the speed gain could be substantially higher, depending on the goal of the study. In the two-dimensional examples considered in Section 6, the speed gain was also higher: a factor between 6 and 8. We also argued there that this is the generic case: the speed gain in D -dimensional simulations can be a factor of D greater than that in one-dimensional simulations. A summary of typical CPU times and of respective computational gains provided by the ETD-LF over its closest competitors for the examples considered in Sections 5 and 6 is given in Table 1.

Rephrasing a statement from the previous paragraph, let us emphasize that the speed gain by the ETD-LF was *not* due to its order of accuracy being higher than those of other methods. The ETD-LF’s order is two, the same as the order of the SSM used in this work. The order of the ETD-4RK is four. However, one could not use larger time steps in the ETD-4RK to speed it up because this would have resulted in a high- k NI. The same pertains to the SSM or its higher-order versions¹¹. In contrast, the time steps in the ETD-LF *can* be increased far beyond the NI threshold, with the limitation being imposed only by the accuracy of the obtained solution, but not by the NI. It was an important point of this study to demonstrate that this accuracy is *sufficient* for the ETD-LF to yield practically the same statistical information about the solution as the other, much

¹¹It was shown in [23] that fourth-order versions of the SSM have the same NI threshold as the second-order one.

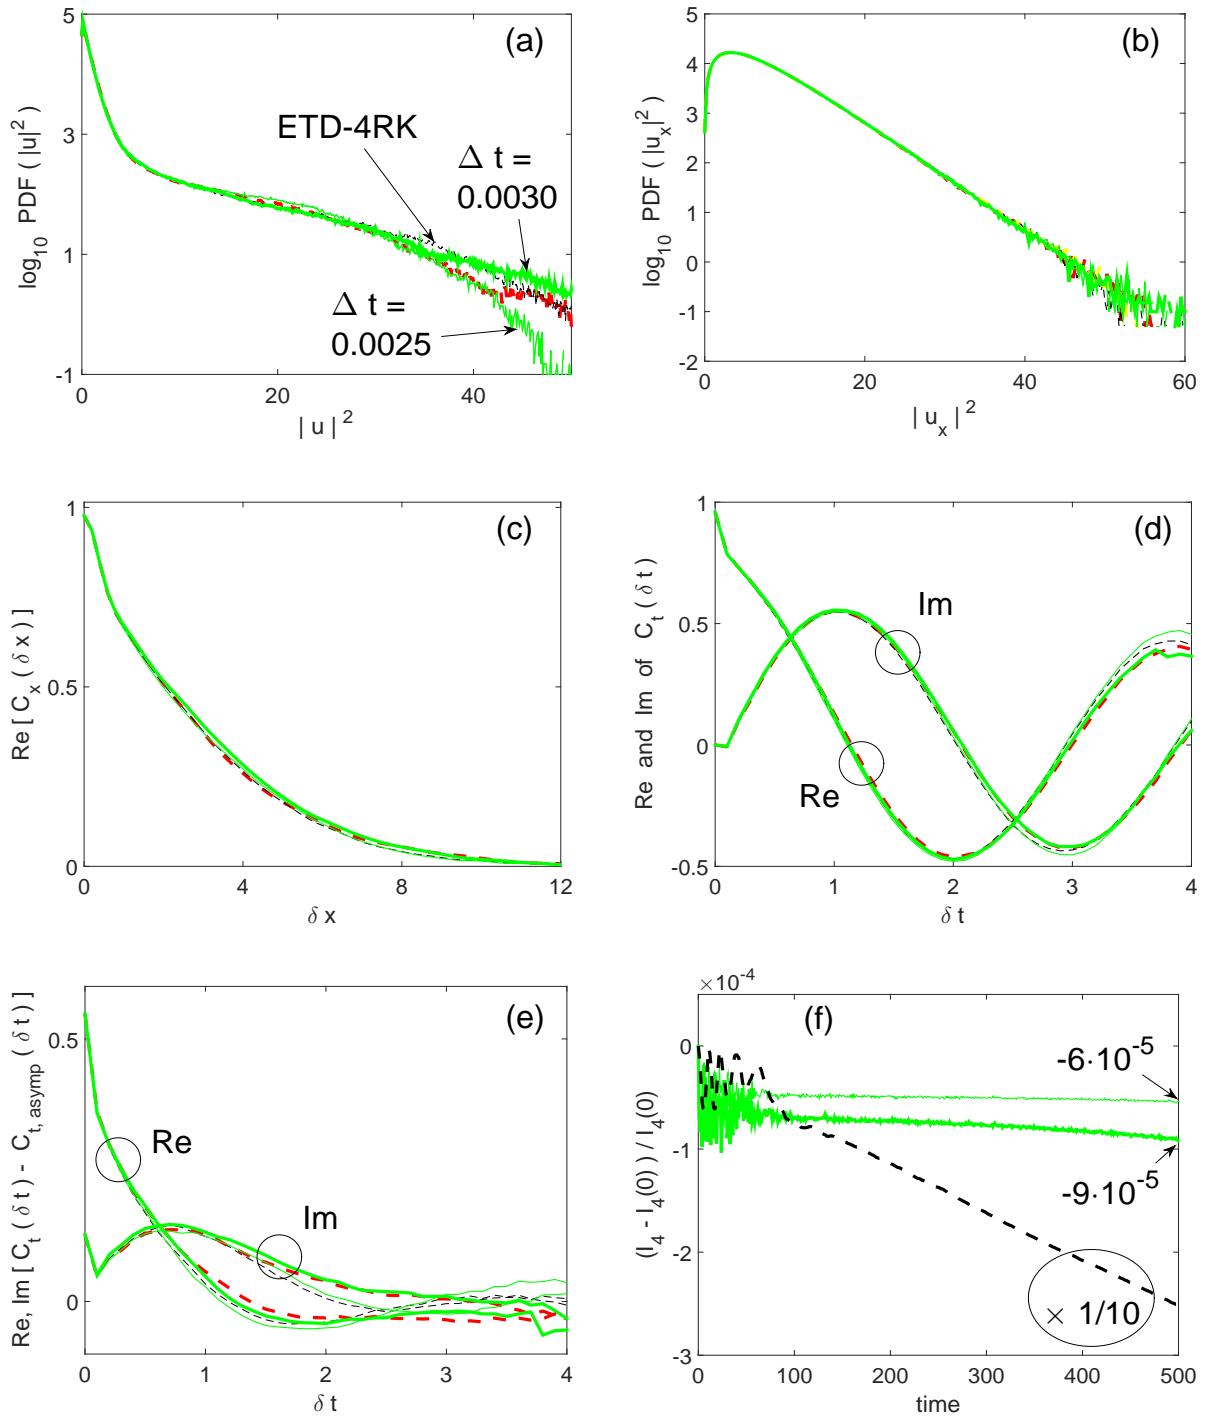


Figure 16: (Color online) Panels (a)–(e): Comparison of the four statistical quantities for the focusing saturable NLS (Section 6.2). ETD-4RK (black, thin dashed) and SSM (red, thick dashed) with $\Delta t = 3.7 \cdot 10^{-4}$; modified (as per (46) with $m = 3$) ETD-LF with $\Delta t = 2.5 \cdot 10^{-3}$ and with $\Delta t = 3 \cdot 10^{-3}$ (green solid, thin and thick, respectively). Panel (f): Relative change of the Hamiltonian computed by the original ETD-LF (dashed; but see next sentence) with $\Delta t = 2.5 \cdot 10^{-3}$ and the modified ETD-LF with $\Delta t = 2.5 \cdot 10^{-3}$ (thin solid) and $\Delta t = 3 \cdot 10^{-3}$ (thick solid). The curve for the original ETD-LF had to be scaled down by a factor of 10 in order to facilitate comparison within the same vertical axis.

Example in Sec.	CPU time for ETD-LF	Name and CPU time of closest competitor	Speed-up factor
5.1 ^a	0.77	ETD-4RK 5.7	7.4
5.2	26	SSM 79	3.0
5.3	0.25	SSM 0.83	3.3
5.4 ^b	0.067	SSM 0.20	3.0
6.1 ^b	6.3	SSM 40	6.3
6.2 ^c	3.0	SSM 23	7.7

Table 1: Computational times (in hours) of the ETD-LF and its closest competitors for the examples considered in Sections 5 and 6. All simulations were performed in Matlab 2014b on a personal computer with a Intel(R) Core(TM) i5-3570 3.4-GHz CPU and 4 GB RAM, running a 64-bit Windows operating system. Given a considerable variation of processor speeds over different platforms, we found it appropriate to present only two significant figures of the data.

^aETD-LF data in this line pertain to the 2nd bulleted item at the end of Sec. 5.1.

^bETD-LF data in this line pertain to the more accurate of the two simulations.

^cETD-LF data in this line pertain to the *less* accurate of the two simulations, as its result appear to be as accurate as in the previous examples.

slower, methods.

Another advantage of the ETD-LF (as well as of the IF-LF) over the SSM is that it can be straightforwardly applied to *any* nonlinear term \mathcal{N} in Eq. (4). The algorithm in that case is exactly the same as for the purely cubic nonlinearity. On the contrary, to apply the SSM, one needs to first work out an analytical expression for the solution in the absence of the $\mathcal{L}u$ -term. (For example, this had to be done in Sections 5.3 and 6.2.) For many types of \mathcal{N} that are more complicated than those considered in this work (e.g., for the term $u(|u|^2)_x$, describing the Raman effect in optical fibers), such an analytical expression does not exist; this would impede the implementation of the “nonlinear substep” of the SSM but would not deter the ETD-LF.

Finally, we note that ETD-LF-like methods may be used to speed up simulations of other types of non-dissipative (or weakly dissipative) equations, such as the Korteweg–de Vries or related equations, where statistical information about the solution, rather than the solution’s precise value, is sought (see, e.g., [37]).

Acknowledgement

This work was supported in part by the NSF grant DMS-1217006. I thank the anonymous referees whose constructive comments helped improve this paper.

Appendix A: Comparison of high- k NI of IF-LF and SSM

We will first illustrate that the high- k NI of these two methods is the same as long as the spectrum of the background solution has negligible content in the vicinity of the resonant wavenumbers $k_{\pi m}$, defined in (18). Then we will present a heuristic argument why the NI of the IF-LF appears to be stronger than that of the SSM when the solution's spectrum expands to and beyond $k_{\pi m}$.

Figure 17 shows a typical spectrum of a soliton initial condition,

$$u(x, 0) = \text{sech}(x) + \xi(x), \quad (47)$$

whose evolution is computed by the IF-LF (Fig. 17(a)) or ETD-LF (Fig. 17(b)) with a time step exceeding the NI threshold. The spectrum computed by the SSM is indistinguishable from that computed by the IF-LF. The white noise $\xi(x)$ of magnitude of order 10^{-10} is added in (47) to facilitate controllable observation of NI; in all simulations $\xi(x)$ was the same. The two spikes in the vicinity of k_{π} in Fig. 17(a) are the unstable modes. In Table 2 we list the locations and magnitudes of these spikes for the SSM as one varies Δt . The corresponding locations for the IF-LF were found to be exactly the same as for the SSM, while the magnitudes were within $\pm 5\%$. This demonstrates the close similarity between the high- k NIs of the two methods.

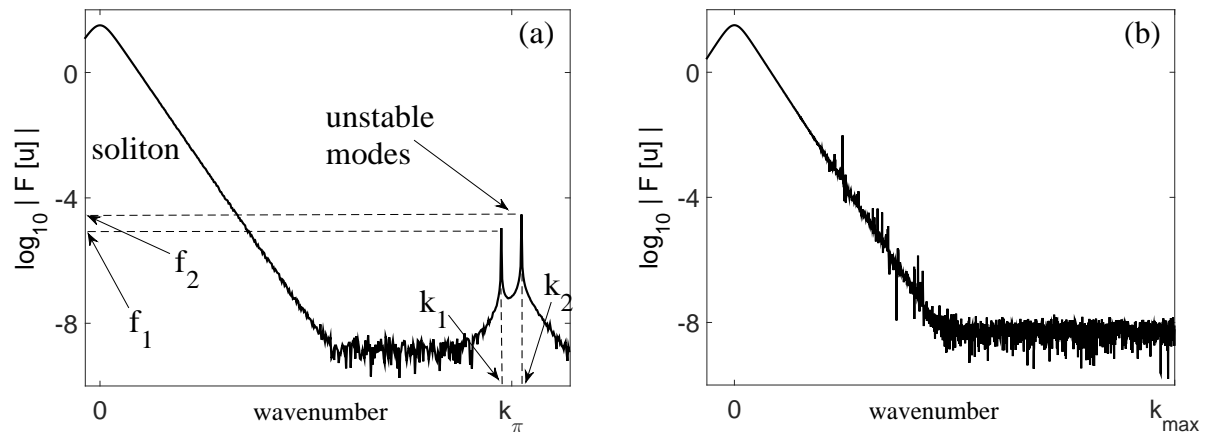


Figure 17: Spectra of the numerical solution of (1) with $\gamma = 2$, $V(x) \equiv 0$, and the near-soliton initial condition (47). The domain length $L = 128\pi$ and the number of grid points $N = 2^{12}$ are as in the experiment reported in [23]. Panel (a) shows a typical spectrum obtained at $t = 1000$ with the IF-LF with $\Delta t_{\text{thresh}} < \Delta t < 2\Delta t_{\text{thresh}}$. (For some Δt , there may be more than two unstable modes around k_{π} ; see Table 2.) Panel (b) shows the spectrum obtained by the ETD-LF at $t = 10,000$ (i.e. 10 times longer than in (a)) with $\Delta t = 0.05$. Stabilization, discussed in Section 4, did not need to be applied for either the IF- or ETD-LF.

Figure 17(b) shows that the NI remains suppressed in the solution obtained by the ETD-LF over a very long time even when Δt exceeds the threshold more than 15-fold. Note that unlike in Fig. 17(a), here the spikes do *not* correspond to numerically unstable modes. These spikes occur in the first few time units of the evolution (probably due to a discretization error) and then remain almost unchanged over a very long time.

Δt	$k_1, k_2,$ etc.	$f_1, f_2,$ etc.
0.0041	27.03, 28.28	-6.1, -6.2
0.0042	26.89, 27.77	-5.1, -5.2
0.0043	26.56, 27.45	-5.3, -6.0
0.0044	26.08, 27.31	-6.3, -6.1
	25.80, 27.58	-6.9, -6.8
0.0045	25.88, 26.92	-5.6, -6.0
	25.56, 27.22	-7.1, -7.2

Table 2: Locations and magnitudes of the unstable modes for the SSM at $t = 1000$; other parameters are listed in the caption to Fig. 17. For reference, $\Delta t_{\text{thresh}} \approx 0.0031$. Notations k_1 , etc. are explained in Fig. 17(a). For the last two values of Δt there are four unstable modes near k_π [23]; the locations and magnitude for the secondary pair of modes are listed below those for the first pair.

We will now present a heuristic explanation of why the unstable modes in the IF-LF method appear to grow much faster than in the SSM (see Fig. 1(a)). In Fig. 18 we provide evidence that this “extra” growth is not exponential. Rather, it occurs only when the spectrum expands so much that near $\pm k_{\pi m}$ it considerably exceeds the noise floor. Figure 18 shows that the magnitude of unstable modes near k_π has increased by almost two orders of magnitude precisely at those short time intervals of duration ~ 0.1 when the solution’s spectrum has expanded beyond k_π . After the spectrum has shrunk (Fig. 18(c)), these modes retain their abnormally increased amplitude. They proceed on growing slowly¹² until the next instance of spectrum expansion beyond k_π occurs.




Figure 18: Spectra (on the \log_{10} -scale) of the solution of the pure NLS obtained by the IF-LF with $\Delta t = 3.2 \cdot 10^{-4} > \Delta t_{\text{thresh}}$. Parameters are similar to those in Section 5.1, but $L = 80\pi$, $N = 2^{13}$. Shown are successive stages of the spectrum expansion beyond k_π (a,b) and subsequent shrinking (c).

In contrast, the SSM does not produce such quickly growing modes; numerically unstable modes in its solution grow at an approximately constant rate and independently of the behavior of the

¹²In the specific example used in Figs. 1 and 18, the NI growth rate is on the order of $0.01 \dots 0.1$, whereas spectrum expansion occurs with frequency of order once per less than ten time units.

spectrum. This indicates that a possible reason for the quickly growing modes in the IF-LF is the presence of *two* solutions of its three-level difference equation (8)¹³. Indeed, the SSM, being a two-level method, has only one solution of its difference scheme. An indirect confirmation of our conjecture comes from the solution by the ETD-LF, which, like the IF-LF, is a three-level method but, unlike the IF-LF, has no or a greatly reduced high- k NI. Figure 1 shows that the ETD-LF solution also develops spectral peaks around $|k| = k_{\pi m}$, and these peaks grow only when the spectrum expands beyond their locations. However, we have been unable to quantitatively explain the formation of these peaks around $|k| = k_{\pi m}$ during spectrum expansions.

Appendix B: Reduction of spectral distortions near $|k| = k_{\pi}$ and $k_{2\pi}$

We have observed that the distortions in questions appeared only near $k_{\pi m}$ but not, for example, near $k_{\pi/2}$ (where $k_{\pi/2}^2 \Delta t = \pi/2$). Therefore, one expects that distortions for $|k| \approx k_{\pi}$ will be reduced if one somehow could use a twice as small time step, $\Delta t_{\text{new}} = \Delta t/2$, to resolve the dynamics of *those specific Fourier modes* (because then $k_{\pi}^2 \Delta t_{\text{new}} = \pi/2$). This can be done by using, instead of (8), a formula

$$e^{ik^2 \Delta t/2} \hat{u}(t_{n+1}) - e^{-ik^2 \Delta t/2} \hat{u}(t_n) = i\Delta t \operatorname{sinc}(k^2 \Delta t/2) \hat{\mathcal{N}}(t_{n+1/2}), \quad (48a)$$

where the term on the r.h.s. is found by the standard linear extrapolation:

$$\hat{\mathcal{N}}(t_{n+1/2}) \equiv \frac{3}{2} \hat{\mathcal{N}}(t_n) - \frac{1}{2} \hat{\mathcal{N}}(t_{n-1}). \quad (48b)$$

For the result shown in Fig. 1(b) with the green solid line, this technique was applied to harmonics with $|k| \in [k_{\pi} - 12, k_{\pi} + 10]$ (for reference, $k_{\pi} = 41.8$ there).

We have empirically found that this idea does *not* work if extended to $|k| \approx k_{2\pi}$, where one would use $\Delta t_{\text{new}} = \Delta t/4$. The problem appeared to be in a linear NI that set in when the term $\hat{\mathcal{N}}(t_{n+3/4})$, required for an extension of (48a), was computed by an extrapolation analogous to (48b).

Therefore we used a different idea for $|k| \approx k_{2\pi}$. Those harmonics do not need to be resolved accurately, because they contribute little to the solution. Rather, one simply needs to arrest their abnormal growth (see the spectrum with “no corrections” in Fig. 1(b)). A well-known way to gain stability at the expense of accuracy is to use a linearly implicit method:

$$\hat{u}(t_{n+1}) - \hat{u}(t_n) = (-ik^2 \Delta t/2)(\hat{u}(t_{n+1}) + \hat{u}(t_n)) + i\Delta t \hat{\mathcal{N}}(t_{n+1/2}), \quad (49)$$

where $\hat{\mathcal{N}}(t_{n+1/2})$ is computed by (48b). We have empirically found that the best results are obtained when (49) is applied every third step. The solid green line in Fig. 1(b) shows the result of doing so for harmonics with $|k| \in (k_{\pi} + 10, k_{2\pi} + 9]$.

¹³Note that approximation (19) is not applicable for modes with those k where the spectrum varies by several orders of magnitude during its expansion; in particular, this occurs for $|k| \gtrsim k_{\pi}$. For such modes, the size of the nonlinear term in (8) may be much greater than that of the terms on the l.h.s. when the spectrum expands.

References

- [1] I. Alonso-Mallo, A. Duran, N. Reguera, Simulation of coherent structures in nonlinear Schrödinger-type equations, *J. Comp. Phys.* 229 (2010) 8180–8198.
- [2] X. Antoine, W. Bao, C. Besse, Computational methods for the dynamics of the nonlinear Schrödinger/Gross–Pitaevskii equations, *Comp. Phys. Comm.* 184 (2013) 2621–2633.
- [3] N. Akhmediev, J.M. Dudley, D.R. Solli, S.K. Turitsyn, Recent progress in investigating optical rogue waves, *J. Opt.* 15 (2013) 060201.
- [4] M. Onorato, S. Residori, U. Bortolozzo, A. Montinad, F.T. Arecchi, Rogue waves and their generating mechanisms in different physical contexts, *Phys. Rep.* 528 (2013) 47–89.
- [5] V.E. Zakharov, F. Dias, A. Pushkarev, One-dimensional wave turbulence, *Phys. Rep.* 398 (2004) 1–65.
- [6] A.O. Korotkevich, Influence of the condensate and inverse cascade on the direct cascade in wave turbulence, *Math. Comp. Simul.* 82 (2012) 1228–1238.
- [7] J.A.C. Weideman and B.M. Herbst, Split-step methods for the solution of the nonlinear Schrödinger equation, *SIAM J. Numer. Anal.* 23 (1986) 485–507.
- [8] T.F. Chan, T. Kerkhoven, Fourier methods with extended stability intervals for the Korteweg–de Vries equation, *SIAM J. Numer. Anal.* 22 (1985) 441–454.
- [9] T.I. Lakoba, Instability of the finite-difference split-step method applied to the nonlinear Schrödinger equation. I. standing soliton, *Num. Meth. Part. Diff. Eqs.* 32 (2016) 1002–1023.
- [10] T. Kremp, W. Freude, Fast split-step wavelet collocation method for WDM system parameter optimization, *J. Lightwave Technol.* 23 (2005) 1491–1502.
- [11] T.I. Lakoba, Instability of the split-step method for a signal with nonzero central frequency, *J. Opt. Soc. Am. B* 30 (2013) 3260–3271.
- [12] J. Certaine, The solution of ordinary differential equations with large time constants, in: A. Ralston, H.S. Wilf. (Eds.), *Mathematical methods for digital computers*, Wiley, New York, 1960, pp. 128–132.
- [13] D.A. Pope, An exponential method of numerical integration of ordinary differential equations, *Commun. ACM* 6 (1963) 491–493.
- [14] J.D. Lawson, Generalized Runge–Kutta processes for stable systems with large Lipschitz constants, *SIAM J. Numer. Anal.* 4 (1967) 372–380.
- [15] P.A. Milewski, E. Tabak, A pseudospectral procedure for the solution of nonlinear wave equations with examples from free-surface flow, *SIAM J. Sci. Comp.* 21 (1999) 1102–1114.

- [16] S.M. Cox, P.C. Matthews, Exponential time differencing for stiff systems, *J. Comp. Phys.* 176 (2002) 430–455.
- [17] Z. Fei, V.M. Perez-Garcia, L. Vazquez, Numerical simulation of nonlinear Schrödinger systems: a new conservative scheme, *Appl. Math. Comp.* 71 (1995) 165–177.
- [18] Y. Zhang, W. Bao, Dynamics of the center of mass in rotating Bose–Einstein condensates, *Appl. Numer. Math.* 57 (2007) 697–709.
- [19] B. Fornberg, G.B. Whitham, A numerical and theoretical study of certain nonlinear wave phenomena, *Phil. Trans. Roy. Soc.* 289 (1978) 373–404.
- [20] E. Hairer, C. Lubich, G. Wanner, Geometric numerical integration illustrated by the Störmer–Verlet method, *Acta Numerica* 12 (2003) 399–450.
- [21] A.L. Islas, D.A. Karpeev, C.M. Schober, Geometric integrators for the nonlinear Schrödinger equation, *J. Comp. Phys.* 173 (2001) 116–148.
- [22] H. Yoshida, Construction of higher order symplectic integrators, *Phys. Lett. A* 150 (1990) 262–268.
- [23] T.I. Lakoba, Instability analysis of the split-step Fourier method on the background of a soliton of the nonlinear Schrödinger equation, *Num. Meth. Part. Diff. Eqs.* 28 (2012) 641–669.
- [24] S. Zhao, J. Ovadia, X. Liu, Y.-T. Zhang, Q. Nie, Operator splitting integration factor methods for stiff reaction–diffusion–advection systems, *J. Comput. Phys.* 230 (2011) 5996–6009.
- [25] T. Jiang, Y.-T. Zhang, Krylov implicit integration factor WENO methods for semilinear and fully nonlinear advection-diffusion-reaction equations, *J. Comput. Phys.* 253 (2013) 368–388.
- [26] A.Y. Suhov, An accurate polynomial approximation of exponential integrators, *J. Sci. Comput.* 60 (2014) 684–698.
- [27] C.-Y. Jung, T. B. Nguyen, Semi-analytical time differencing methods for stiff problems, *J. Sci. Comput.* 63 (2015) 355–373.
- [28] L. Ju, J. Zhang, L. Zhu, Q. Du, Fast explicit integration factor methods for semilinear parabolic equations, *J. Sci. Comput.* 62 (2015) 431–455.
- [29] L. Zhu, L. Ju, W. Zhao, Fast high-order compact exponential time differencing Runge-Kutta methods for second-order semilinear parabolic equations, *J. Sci. Comput.* 67 (2016) 1043–1065.
- [30] D. Lu, Y.-T. Zhang, Krylov integration factor method on sparse grids for high spatial dimension convection-diffusion equations, *J. Sci. Comput.* (2016) DOI 10.1007/s10915-016-0216-7.
- [31] G.P. Agrawal, *Nonlinear fiber optics*, 3rd Ed., Academic Press, San Diego, 2001, Sec. 5.1.

- [32] M. Dahlby, B. Owren, Plane wave stability of some conservative schemes for the cubic Schrödinger equation, *ESAIM: M2AN* 43 (2009) 677–687.
- [33] W.L. Briggs, A.C. Newell, T. Searie, Focusing: A mechanism for instability of nonlinear finite difference equations, *J. Comp. Phys.* 51 (1983) 83–106.
- [34] V.E. Zakharov, A.B. Shabat, Exact theory of two-dimensional self-focusing and one-dimensional self-modulation of waves in nonlinear media, *Sov. Phys. JETP* 34 (1972) 62–69.
- [35] T.I. Lakoba, Effect of noise on extreme events probability in a one-dimensional nonlinear Schrödinger equation, *Phys. Lett. A* 379 (2015) 1821-1827.
- [36] P.M. Lushnikov, N. Vladimirova, Non-Gaussian statistics of multiple filamentation, *Opt. Lett.* 35 (2010) 1965–1967.
- [37] D. Dutykh, E. Pelinovsky, Numerical simulation of a solitonic gas in KdV and KdV–BBM equations, *Phys. Lett. A* 378 (2014) 3102-3110.