

Open Boundaries for the Nonlinear Schrödinger Equation

A. Soffer and C. Stucchio

February 2, 2008

Abstract

We present a new algorithm, the Time Dependent Phase Space Filter (TDPSF) which is used to solve time dependent Nonlinear Schrödinger Equations (NLS). The algorithm consists of solving the NLS on a box with periodic boundary conditions (by any algorithm). Periodically in time we decompose the solution into a family of coherent states. Coherent states which are outgoing are deleted, while those which are not are kept, reducing the problem of reflected (wrapped) waves. Numerical results are given, and rigorous error estimates are described.

The TDPSF is compatible with spectral methods for solving the interior problem. The TDPSF also fails gracefully, in the sense that the algorithm notifies the user when the result is incorrect. We are aware of no other method with this capability.

1 Introduction and Definitions

Consider a semilinear Schrödinger equation on \mathbb{R}^{N+1}

$$i\partial_t\psi(x, t) = -(1/2)\Delta\psi(x, t) + g(t, \vec{x}, \psi(\vec{x}, t))\psi(\vec{x}, t) \quad (1.1)$$

where $g(t, \vec{x}, \cdot)$ is some semilinear perturbation. For instance, $g(t, \vec{x}, \cdot)$ could be $V(\vec{x}, t) + f(|\psi(\vec{x}, t)|^2)$ for some smooth function f and (spatially) localized potential $V(\vec{x}, t)$. Abusing terminology, we refer to $g(t, \vec{x}, \cdot)$ as the nonlinearity even when it is linear.

We assume the initial condition and nonlinearity are such that the nonlinearity remains localized inside some box $[-L_{\text{NL}}, L_{\text{NL}}]^N$. Outside this region the solution is assumed to behave like a free wave.

One very common method of solving such a problem is domain truncation. That is, one solves the PDE (1.1) numerically on a region $[-L, L]^N$. On the finite domain, boundary conditions must be specified. Dirichlet or Neumann boundaries introduce spurious reflections, while periodic boundaries allow outgoing waves to wrap around the computational domain. This causes the numerical solution to become incorrect after a time $T \approx L/k_{\text{max}}$, where k_{max} is the “maximal velocity” (the highest relevant spatial frequency) of the solution.

There are two major approaches to dealing with this problem. The Dirichlet-to-Neumann method consists of attempting to use the exact solution as a boundary condition. This works rather well for the wave equation [1, 2, 9, 11, 13], but it is fraught with problems for dispersive waves and only limited progress has been made [16, 19, 20, 25]. Dirichlet-to-Neumann methods also prevent the use of fast spectral methods for solving the interior problem because spectral methods based on the FFT naturally impose periodic boundary conditions. A spectral solver for the interior problem is highly desirable since dispersive waves are difficult to calculate by other methods (see Remark 3.1).

The other approach is to add a dissipative term which is localized on a buffer region $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$. The dissipative term can be an absorbing potential [17] or the more sophisticated PML [3]. This dissipates outgoing waves, but also dissipates incoming waves located near the boundary. This can introduce spurious dissipation (see Section 3.2) which does not occur for the TDPSF.

1.1 Our Approach

We propose a new method. Suppose we want to approximate the solution in the space H^s (with s large enough to make (1.1) is well posed). We make the assumption that outside some region $[-L_F, L_F]^N$, $\psi(x, t)$ behaves like a free wave, i.e. $\psi(x, t) \approx e^{i(1/2)\Delta t}\psi_+(x)$. The size of the computational box L_{int} must be larger than L_F .

We assume the nonlinearity is nearly localized both in position and momentum. That is, we assume the existence of $L_{\text{NL}}, k_{\text{max,NL}}$ so that:

$$\left\| [1 - \chi_{[-L_{\text{NL}}, L_{\text{NL}}]^N}(\vec{x})]g(t, \vec{x}, \psi(x, t))\psi(x, t) \right\|_{H^s} \approx 0 \quad (1.2a)$$

$$\left\| [1 - \chi_{[-k_{\text{max,NL}}, k_{\text{max,NL}}]^N}(\vec{k})]g(t, \vec{x}, \psi(x, t))\psi(x, t) \right\|_{H^s} \approx 0 \quad (1.2b)$$

For the rigorous proof, we assume that the nonlinearity is Lipschitz in H^s , though this is probably unnecessary. In particular, the nonlinearity could take the form of a time dependent short range potential $V(\vec{x}, t)\psi(\vec{x}, t)$ or a local nonlinearity $f(|\psi(\vec{x}, t)|^{2\sigma})\psi(\vec{x}, t)$ with $f(z)$ a bounded function.

We also assume that the solution remains localized in frequency, that is the H^s norm of $\hat{\psi}(\vec{k}, t)$ is small outside $[-k_{\text{max}}, k_{\text{max}}]^N$ for some large number k_{max} (the maximal momentum of the problem, which we assume exists).

Our algorithm is as follows. We assume the initial data is localized on a region $[-L_{\text{int}}, L_{\text{int}}]^N$. We solve (1.1) on the box $[-(L_{\text{int}} + w), L_{\text{int}} + w]^N$ on the time interval $[0, T_{\text{step}}]$. We insist that $L_{\text{int}} > L_F, L_{\text{NL}}$ so that the solution behaves like a free wave on $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$.

By making T_{step} smaller than $w/3k_{\text{max}}$, we can ensure that $\psi(\vec{x}, T_{\text{step}})$ is mostly localized inside box $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N$. Since very little mass has reached the boundary, the error is nearly zero [24, Thm. 4.1].

We then decompose $\psi(\vec{x}, T_{\text{step}})$ into a sum of Gaussians (indexed by $\vec{a}, \vec{b} \in \mathbb{Z}^N$, with lattice spacing x_0 in position, k_0 in momentum and standard deviation

σ):

$$\psi(x, T_{\text{step}}) \approx \sum_{\substack{|\vec{a}x_0|_\infty \leq L_{\text{int}} + w \\ |\vec{b}k_0|_\infty \leq k_{\text{max}}}} \psi_{(\vec{a}, \vec{b})} \pi^{-N/4} \sigma^{-N/2} e^{ik_0 \vec{b} \cdot \vec{x}} e^{-|\vec{x} - \vec{a}x_0|_2^2 / 2\sigma^2}$$

We then examine the Gaussians near the boundary, $|\vec{a}x_0|_\infty \geq L_{\text{int}}$, and determine whether they are leaving the box or not under the free flow ($|\vec{x}|_p$ denotes the l^p norm on finite vectors). By assumption, $e^{i(1/2)\Delta t}$ is a sufficiently accurate approximation to the true solution for this part of the solution. We apply this approximation to each framelet:

$$\begin{aligned} & e^{i(1/2)\Delta t} \pi^{-N/4} \sigma^{-N/2} e^{ik_0 \vec{b} \cdot \vec{x}} e^{-|\vec{x} - \vec{a}x_0|_2^2 / 2\sigma^2} \\ &= \frac{\exp\left(i\vec{b}k_0 \cdot (\vec{x} - (\vec{b}k_0/2)t - \vec{a}x_0)\right)}{\pi^{N/4} \sigma^{N/2} (1 + it/\sigma^2)^{N/2}} \exp\left(\frac{-|\vec{x} - \vec{b}k_0 t - \vec{a}x_0|_2^2}{2\sigma^2(1 + it/\sigma^2)}\right) \quad (1.3) \end{aligned}$$

Under the free flow, a framelet moves along the trajectory $\vec{a}x_0 + \vec{b}k_0 t$ while spreading about its center. If a given Gaussian is leaving the box, we delete it (set $\psi_{(\vec{a}, \vec{b})} = 0$). Some Gaussians spread more quickly than their center of mass moves, and we do not present here an algorithm to deal with these Gaussians.

After this filtering operation, the only Gaussians remaining are either inside the box $[-L_{\text{int}}, L_{\text{int}}]^N$, or inside the box $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N$ but moving towards $[-L_{\text{int}}, L_{\text{int}}]^N$. Thus, it is safe to propagate for a time T_{step} , since what remains will not hit the boundaries before this time. We do the same at time $2T_{\text{step}}, 3T_{\text{step}}$, etc.

The main drawback to this approach is that some Gaussians are ambiguous. Consider a Gaussian with velocity 0: $\pi^{-N/4} \sigma^{-N/2} e^{-|\vec{x} - \vec{a}x_0|_2^2 / 2\sigma^2}$. Suppose also that $\vec{a}x_0$ is located inside $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$. By examining (1.3), we observe that this framelet is spreading out laterally both into the box and outward. If we delete it, we have removed waves which should have returned. If we fail to delete it, then that part of the framelet which is spreading outwards will wrap around and cause an error.

This happens only for Gaussians which are moving slowly relative to the box. Thus, we impose one additional assumption – we assume that Gaussians moving so slowly that they move both ways do not occur in the solution. The algorithm to deal with low frequency waves is more involved and we relegate this to a future work [23].

1.2 Error Bounds

We prove rigorous error bounds for the TDPSF algorithm in [24], as well as stating explicitly the assumptions and defining what we mean by “ \approx ”. Although the error bound is too long to state here, we describe briefly the form it takes.

For a general time-stepping algorithm (with periodic boundaries and no filtering), the error bound would take the following form:

$$\begin{aligned} \sup_{t \in [0, T_{\max}]} \|\mathcal{U}(t)\psi_0(x) - \Psi(x, t)\|_{H_b^s} &\leq \text{BoundaryError}(T_{\max}) \\ &+ \text{HighFrequency}(T_{\max}) + \text{LowFrequency}(T_{\max}) \\ &+ \text{NonlocalNonlinearity}(T_{\max}) + \text{Instability}(T_{\max}) \quad (1.4) \end{aligned}$$

The term $\text{BoundaryError}(T_{\max})$ encompasses errors due to waves wrapping/reflecting from the boundaries of the box. For many problems, this is the dominant error term. It is directly proportional to the mass which would have (if we were solving the problem on \mathbb{R}^N) radiated outside the box $[-L_{\text{int}}, L_{\text{int}}]^N$.

The $\text{HighFrequency}(T_{\max})$ part stems from waves with momenta too high to be resolved by the discretization. The term $\text{LowFrequency}(T_{\max})$ encompasses errors due to waves with wavelength that is long in comparison to the box. The term $\text{NonlocalNonlinearity}(T_{\max})$ stems from that fraction of the nonlinearity itself which is located outside the box. The $\text{Instability}(T_{\max})$ stems from the possibility that the dynamics of the solution itself might amplify the other errors dramatically (e.g. in strongly nonlinear problems).

Our algorithm reduces the term $\text{BoundaryError}(T_{\max})$. The other errors are assumed to be small, since we are concerned only with the boundary error.

The method of proof is a direct calculation. We calculate the errors made in the filtering step as well as the errors made while propagating in between filtering. We then add up the error over time, taking into account instabilities of the system being simulated. The error takes the form described in (1.4). The $\text{BoundaryError}(T_{\max})$ term can be reduced by increasing w , and the error behaves like $\text{BoundaryError}(T_{\max}) = O(e^{-Cw}T_{\max})$. In some cases (problems which are asymptotically complete and have no bound states), we believe that this error can be proved to be time independent.

The main drawback of our algorithm is that it does not provide us the ability to filter waves for which the wavelength is longer than the buffer region. This is a due to the Heisenberg uncertainty principle. This problem will be addressed by a novel multiscale argument in [23]. The problem is shared with most methods of absorbing boundaries. The TDPSF algorithm also alerts the user when it fails, unlike all other methods we are aware of.

2 The Algorithm

2.1 The Windowed Fourier Transform

The discrete windowed Fourier transform frame is a way of localizing a function $f(\vec{x})$ in the phase space. For the big picture, see [6, 7, 8]; see also [24, Section 3] for technical details needed by the TDPSF.

The WFT is an expansion of a function $f(x)$ into the following ‘‘basis’’

(actually a frame) for $L^2(\mathbb{R}^N)$:

$$\phi_{(\vec{a}, \vec{b})}(\vec{x}) = \pi^{-N/4} \sigma^{-N/2} e^{ik_0 \vec{b} \cdot \vec{x}} e^{-|\vec{x} - \vec{a}x_0|^2 / 2\sigma^2}, \quad \vec{a}, \vec{b} \in \mathbb{Z}^N \times \mathbb{Z}^N \quad (2.1)$$

The lattice spacing in position and momentum, x_0 , k_0 , and the standard deviation σ must positive real numbers. To be a frame, it must hold that $x_0 k_0 < 2\pi$.

This family of functions is not a basis in the usual sense, but a frame. A family of functions $\{\phi_j(x)\}_{j \in J}$ is a frame in the Hilbert space H if there exists an $0 < A_F \leq B_F < \infty$ so that:

$$A_F \|f(x)\|_H \leq \left(\sum_{j \in J} |\langle f | \phi_j(x) \rangle|^2 \right)^{1/2} \leq B_F \|f(x)\|_H$$

A frame is an overcomplete, non-orthonormal basis with a numerically stable reconstruction procedure. We can then decompose any $f \in L^2(\mathbb{R}^N)$ as:

$$f(\vec{x}) = \sum_{(\vec{a}, \vec{b}) \in \mathbb{Z}^N \times \mathbb{Z}^N} f_{(\vec{a}, \vec{b})} \phi_{(\vec{a}, \vec{b})}(\vec{x}) \quad (2.2)$$

$$f_{(\vec{a}, \vec{b})} = \langle f(\vec{x}) | \tilde{\phi}_{(\vec{a}, \vec{b})}(\vec{x}) \rangle \quad (2.3)$$

The functions $\tilde{\phi}_{(\vec{a}, \vec{b})}(\vec{x})$ are given by $\tilde{\phi}_{(\vec{a}, \vec{b})}(\vec{x}) = e^{ik_0 \vec{b} \cdot \vec{x}} \tilde{g}(\vec{x} - \vec{a}x_0)$ with $\tilde{g}(\vec{x}) \in L^2(\mathbb{R}^N)$ (see [6, 7] for procedures to calculate $\tilde{g}(\vec{x})$).

In [24] we extend the analysis of [7] and show that if $x_0 k_0 = 2\pi/\mathbf{M}$ for $\mathbf{M} \in 2\mathbb{N}$, $g(\vec{x})$ decays exponentially in $|\vec{x}|_1$ and $\hat{g}(\vec{k})$ decays exponentially in $|\vec{k}|_1$:

$$|\partial_x^{\vec{\alpha}} \tilde{g}(\vec{x})| \leq \mathbf{g}(x_0, k_0, N, \vec{\alpha}) e^{-\mathbf{r}(x_0, k_0) |\vec{x}|_1} \quad (2.4)$$

$\mathbf{g}(x_0, k_0, N, \vec{\alpha})$ is a constant (explicitly bounded in [24]) as is $\mathbf{r}(x_0, k_0) = x_0 \mathbf{M} / 8\pi\sigma$. (2.4) holds for $\hat{g}(\vec{k})$, but with x_0 and k_0 swapped and σ^{-1} replacing σ .

2.1.1 Phase Space Localization

The WFT allow us to define a concrete realization of phase space. From here onward, we will consider $\mathbb{Z}^N \times \mathbb{Z}^N$ to be a discrete realization of phase space. The vector $(\vec{a}, \vec{b}) \in \mathbb{Z}^N \times \mathbb{Z}^N$ will represent the point at $\vec{a}x_0$ in position, and $\vec{b}k_0$ in momentum.

With this in mind, we can now construct phase space localization operators very simply. For a set $F \in \mathbb{Z}^N \times \mathbb{Z}^N$, we define \mathcal{P}_F by:

$$\mathcal{P}_F \psi(x) = \sum_{(\vec{a}, \vec{b}) \in F} \psi_{(\vec{a}, \vec{b})} \phi_{(\vec{a}, \vec{b})}(\vec{x}) = \sum_{(\vec{a}, \vec{b}) \in F} \langle \tilde{\phi}_{(\vec{a}, \vec{b})}(\vec{x}) | \psi(x) \rangle \phi_{(\vec{a}, \vec{b})}(\vec{x}) \quad (2.5)$$

It can be shown [24, 5, 6] that phase space localization in terms of the WFT corresponds closely (though not exactly) to the usual phase space localization. We state one theorem (proved in [24], based on a result by Daubechies from [5, 6]) as an example of this:

Theorem 2.1 *Let $B_X = [-X, X]^N$, $B_K = [-K, K]^N$ for $X, K < \infty$. Suppose also that $x_0 k_0 = 2\pi/\mathbf{M}$ for $\mathbf{M} \in 2\mathbb{N}$. Then there exists C , \mathbf{X} and \mathbf{K} so that if $X' = X - \mathbf{X}$, $K' = K - \mathbf{K}$, then:*

$$\begin{aligned} & \left\| f(x) - \mathcal{P}_{B_{X'} \times B_{K'}} f(x) \right\|_{H^s} \\ & \leq C \left(\left\| (1 - P_{B_X; x_0}^s(\vec{x})) f(\vec{x}) \right\|_{H^s} + \left\| (1 - P_{B_K; k_0}^0(\vec{k})) f(\vec{x}) \right\|_{H^s} + \epsilon \|f\|_{H^s} \right) \end{aligned} \quad (2.6)$$

The constant C depends on $\epsilon, s, x_0, k_0, \sigma, X$ and K , as do \mathbf{X} and \mathbf{K} . C is given explicitly in [24].

2.1.2 Computation of the WFT Coefficients and Phase Space Projections: How to do it, and how hard it is

We now present the algorithm for computing the framelet coefficients. The algorithm consists of computing the products $f(\vec{x})\tilde{g}(\vec{x} - \vec{a}x_0)$, followed by Fourier transforming the results. Due to the spatial decay of $\tilde{g}(\vec{x})$, we can truncate the domain to the box $[-L_\epsilon, L_\epsilon]^N$ with minimal error.

Note that $\tilde{g}(\vec{x})$ can be calculated as accurately as necessary [6, 7], and we will not discuss this here.

Algorithm 2.1 Calculation of Windowed Fourier Transforms

This algorithm calculates, for a function $f(x)$, the WFT coefficients $f_{(\vec{a}, \vec{b})}$ for $\vec{a} \in A \subseteq \mathbb{Z}^N$. We assume that the frequencies are bounded above by k_{\max} . The lattice spacing in frequency, k_0 is taken to be $2\pi/L_\epsilon$; taking it to be any larger than this yields only logarithmic improvements in computational complexity.

1. Let $A \subset \mathbb{Z}^N$ be some set of position coordinates.
2. For each $\vec{a} \in A$ multiply $f(\vec{x})\tilde{g}(\vec{x} - \vec{a}x_0)$ for $\vec{x} \in [-L_\epsilon, L_\epsilon]^N + \vec{a}x_0$ only.
3. Calculate the Fast Fourier Transform of $f(\vec{x})\tilde{g}(\vec{x} - \vec{a}x_0)$ on this region. For each \vec{a} , the resulting function is $f_{(\vec{a}, \vec{b})}$.

We observe that this algorithm is local in space. This means that if A is finite, then the computational cost is proportional to:

$$|A| k_{\max}^N |\text{supp } \tilde{g}(\vec{x})|^N \log(k_{\max} |\text{supp } \tilde{g}(\vec{x})|) = |A| 2^N k_{\max}^N L_\epsilon^N$$

The reason for this complexity is as follows. For each $\vec{a} \in A$, we need to compute an FFT. The FFT is computed only on the ‘‘support’’ (after truncating) of $\tilde{g}(\vec{x})$ and the lattice spacing in this region is $2\pi/k_{\max} = O(1/k_{\max})$. Thus, there are $O(|\text{supp } \tilde{g}(\vec{x})|^N k_{\max}^N) = M$ data points in this region. The FFT has computational complexity $M \log(M)$. In addition, we need to compute $|A|$ of these FFT’s.

Note that this algorithm can be parallelized very easily, simply by having different processors compute the FFT’s for different $\vec{a}x_0$.

Remark 2.2 As an example, consider the case when $A = \{\vec{a} \in \mathbb{Z}^N : \vec{a}x_0 \in [-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N\}$. Then the size of A is proportional to $|\{[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N\}|/x_0^N$. If $L_{\text{int}} \gg w$, then this is of order L_{int}^{N-1} , and the computational complexity is $O(L_{\text{int}}^{N-1}k_{\text{max}}^N \log(L_{\text{int}}k_{\text{max}}))$.

We consider this case since this is what the TDPSF requires.

Phase space projections can also be computed. Let $F \subset \mathbb{Z}^N \times \mathbb{Z}^N$ be finite. Let $A = \{\vec{a} \in \mathbb{Z}^N : \exists \vec{b} \in \mathbb{Z}, (\vec{a}, \vec{b}) \in F\}$. Then we provide the phase space projection algorithm:

Algorithm 2.2 *Phase Space Projection Algorithm*

This algorithm computes the phase space projection onto a region of phase space F .

1. Compute $A = \{\vec{a} : (\vec{a}, \vec{b}) \in F\}$. Assume this is finite.
2. Compute $f_{(\vec{a}, \vec{b})}$ for $\vec{a} \in A$ as in Algorithm 2.1.
3. Define a new function $f^t : \mathbb{Z}^N \times \mathbb{Z}^N \rightarrow \mathbb{C}$ by:

$$\begin{aligned} f_{(\vec{a}, \vec{b})}^t &= f_{(\vec{a}, \vec{b})}, & (\vec{a}, \vec{b}) \in F \\ f_{(\vec{a}, \vec{b})}^t &= 0, & (\vec{a}, \vec{b}) \notin F \end{aligned}$$

4. Compute the inverse WFT of f^t . The result approximates $\mathcal{P}_F f(x)$, with errors caused by the truncation of $\tilde{g}(\vec{x})$ in Algorithm 2.1.

Clearly the computational complexity of Algorithm 2.2 is of the same order as that of Algorithm 2.1.

2.2 Propagation with Periodic Boundaries: FFT/Split Step Algorithm

The TDPSF is a filtering algorithm, which is built on top of another propagation method which solves (1.1) on a box with periodic boundaries. The exact manner in which this is done is irrelevant for our purposes, provided it is sufficiently accurate. However, it is important to note that the spectral propagation method is significantly better than the usual FDTD methods, and that compatibility with spectral methods is an important feature of the TDPSF.

Fix a grid spacing δx and timestep δt . The computational grid is simply the region $[-L_{\text{comp}}, L_{\text{comp}}]^N$ sampled at the lattice with spacing δx . This corresponds to a lattice spacing in momentum of $2\pi/L_{\text{comp}}$, with maximal momentum $2\pi/\delta x$. A common rule of thumb is that if the problem has a maximal momentum k_{max} , then $\delta x = 4\pi/k_{\text{max}}$.

Algorithm 2.3 *Split Step Propagation Algorithm*

This algorithm approximates the propagation of (1.1) on the region $[-L_{\text{comp}}, L_{\text{comp}}]^N$. The accuracy is $O(\delta t^2)$ time and $O(\delta x^\omega)$ (“spectral accuracy”) in space. It is $O(\delta t^3)$ in time if (1.1) is linear.

First, the discrete operator $e^{i(1/2)\Delta t} f(x)$ is defined by computing the FFT of $f(x)$, multiplying by $e^{ik^2 t}$ and then computing the inverse FFT of the result.

With this in mind, the algorithm is as follows (taking $\psi_0(x)$ as initial condition, and assuming t is a multiple of δt):

1. Define $\psi_{1/2}(x) = e^{i(1/2)\Delta\delta t/2}\psi_0(x)$.

2. For $j = 0, \dots, t/\delta t - 2$, define:

$$\psi_{j+1/2}(x) = e^{i(1/2)\Delta\delta t} e^{-ig(j\delta t, \vec{x}, \psi_{j+1/2}(x))\delta t} \psi_{j+1/2}(x)$$

3. Finally, define $\psi_{t/\delta t}(x) \approx \psi(x, t)$ by:

$$\psi_{t/\delta t}(x) = e^{i(1/2)\Delta(\delta t/2)} e^{-ig(j\delta t, \vec{x}, \psi_{j+1/2}(x))\delta t} e^{i(1/2)\Delta\delta t} \psi_{t/\delta t-3/2}(x)$$

This is a numerical realization of the Trotter-Kato product formula:

$$\mathcal{U}(t) \approx e^{i(1/2)\Delta(-\delta t/2)} \left[\prod_{j=1}^{t/\delta t} e^{i(1/2)\Delta\delta t} e^{-ig(j\delta t, \vec{x}, \psi(x, j\delta t))\delta t} \right] e^{i(1/2)\Delta\delta t/2} \quad (2.7)$$

The computational complexity of Algorithm 2.3 is

$$O((L_{\text{int}} + w)^N k_{\text{max}}^N \ln[(L_{\text{int}} + w)k_{\text{max}}])$$

per timestep (since the number of data points is of order $(L_{\text{int}} + w)k_{\text{max}}$). The number of timesteps is $T_{\text{max}}/\delta t$. Thus, split step propagation has time complexity $O[L_{\text{int}}^N k_{\text{max}}^N \log(L_{\text{int}} k_{\text{max}})(T_{\text{max}}/\delta t)]$. Algorithm 2.3 is by now a textbook result [4], so we will not discuss it further.

2.3 The TDPSF Algorithm

The TDPSF algorithm consists of solving (1.1) on the box $[-L_{\text{comp}}, L_{\text{comp}}]^N$ with periodic boundary conditions, using Algorithm 2.3 or any other method. At times $T_{\text{step}}, 2T_{\text{step}}$, etc, we subtract from the solution the projection onto Gaussians which are outgoing by means of Algorithm 2.2.

Outgoing Gaussians are defined to be those $\phi_{(\vec{a}, \vec{b})}(\vec{x})$ with $\vec{a}x_0 \in [-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$ and which are moving strictly outward under the free flow; that is, those which are leaving $[-L_{\text{int}}, L_{\text{int}}]^N$ soon (before a time T_{step} has passed) and which will never return.

Incoming Gaussians are those which will not leave $[-L_{\text{comp}}, L_{\text{comp}}]^N$ before time T_{step} has passed.

There is also a set of ambiguous gaussians. These are gaussians with velocity sufficiently low so that they spread about their center faster than they

move. They are ambiguous because they are both leaving $[-L_{\text{comp}}, L_{\text{comp}}]^N$ and returning to $[-L_{\text{int}}, L_{\text{int}}]^N$ due to their spread. An example of this would be $e^{-(x-L_{\text{int}}-w/2)^2/2\sigma^2}$. This gaussian is located inside the buffer region, and is spreading laterally with no motion whatsoever.

The TDPSF algorithm does not know what to do with these gaussians, which correspond to low frequency waves. However, unlike other algorithms (e.g. Dirichlet-Neumann, PML or absorbing potential), the TDPSF does know when such waves are causing a problem. Thus, incorrect results are never returned.

Algorithm 2.4 *Propagation algorithm* This is the main propagation algorithm. In this algorithm, we consider σ , x_0 , k_0 to be fixed. L_{int} and T_{max} are also considered fixed. The initial data is considered fixed, and localized inside $[-L_{\text{int}}, L_{\text{int}}]^N$. The approximation will be denoted by $\Psi(x, t)$. Also, fix a small tolerance $\epsilon > 0$.

1. Before beginning, precalculate the set of framelets which are outgoing, and those which are ambiguous.
2. Define $\Psi(x, t)$ iteratively as follows. Loop over $n = 0 \dots T_{\text{max}}/T_{\text{step}}$. In what follows, the propagator $\mathcal{U}(t)$ is calculated by Algorithm 2.3.

(a) For $t \in [(n-1)T_{\text{step}}, nT_{\text{step}})$, define

$$\psi(x, t) = \mathcal{U}(t - (n-1)T_{\text{step}})\psi(x, (n-1)T_{\text{step}})$$

(b) For $t = nT_{\text{step}}$, define:

$$\Psi(x, nT_{\text{step}}) = (1 - \mathcal{P}_{\text{OUT}})\mathcal{U}(T_{\text{step}})\psi(x, (n-1)T_{\text{step}})$$

(c) At each integer multiple of T_{step} , compute

$$\|\mathcal{P}_{\text{AMB}}\mathcal{U}(T_{\text{step}})\psi(x, (n-1)T_{\text{step}})\|_{H^s}$$

where AMB is the set of ambiguous framelets. If this quantity is greater than ϵ , stop the program and notify the user.

Remark 2.3 The TDPSF fails gracefully in the following respect. If the TDPSF is unable to filter due to slowly moving waves, the program fails and the user is notified. Absorbing potentials, the PML and Dirichlet-Neumann maps do not have this property.

Remark 2.4 In order to calculate \mathcal{P}_{OUT} , we must calculate the WFT coefficients of the solution in the region $[-L_{\text{comp}}, L_{\text{comp}}]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$. To calculate \mathcal{P}_{AMB} , we need the WFT coefficients from the same region. Thus, only one WFT is needed for steps (b) and (c). The cost of computing the H^s norm of the ambiguous framelets is cheaper than the cost of a WFT; this is merely a sum rather than a set of FFT's.

Thus, all that remains is to do is explain which framelets are outgoing, incoming and ambiguous.

For a given (\vec{a}, \vec{b}) , this can be determined easily, by a direct calculation based on (1.3). This is done in [24] for $s = 0, 1$, and other cases can be done with the help of Maple/Mathematica.

The end result of the calculation is the following. Define $R_{\vec{b}}(t)$ by

$$R_{\vec{b}}(t) = \sqrt{\sigma^2 + t^2/\sigma^2} (\Gamma^{-1}(N/2, 2\epsilon^2\pi^{N/2}/|S^{N-1}|))^{1/2} \quad (2.8a)$$

if we measure the error in $L^2(\mathbb{R}^N)$ or

$$R_{\vec{b}}(t) = \sqrt{\sigma^2 + t^2/\sigma^2} \max \left\{ \left[\Gamma^{-1} \left(N/2, \frac{\epsilon^2\pi^{N/2}}{2|S^{N-1}|(1+|\vec{b}k_0|_2^2)} \right) \right]^{1/2}, \left[\Gamma^{-1} \left((N+2)/2, \frac{\epsilon^2\sigma^2\pi^{N/2}}{2|S^{N-1}|} \right) \right]^{1/2} \right\} \quad (2.8b)$$

if we measure the error in $H^1(\mathbb{R}^N)$.

The function $\Gamma^{-1}(a, x)$ is the inverse function of $x \mapsto \Gamma(a, x)$, with $\Gamma(a, x)$ the complementary incomplete Gamma function. This definition implies that all the mass (up to a tolerance ϵ) of $e^{i(1/2)\Delta t} \phi_{(\vec{a}, \vec{b})}(\vec{x})$ is, at time t , contained in a ball of radius $R_{\vec{b}}$ about the point $\vec{a}x_0 + \vec{b}k_0t$. That is:

$$\left\| e^{i(1/2)\Delta t} \pi^{-N/4} \sigma^{-N/2} e^{ik_0\vec{b}\cdot\vec{x}} e^{-|\vec{x}-\vec{a}x_0|_2^2/2\sigma^2} \right\|_{H^{0,1}(B(t)^c)} \leq \epsilon \quad (2.9a)$$

$$B(t) = \{x : |\vec{a}x_0 + \vec{b}k_0t - \vec{x}|_2 \leq R_{\vec{b}}(t)\} \quad (2.9b)$$

Thus, if $B(t)$ does not intersect $[-L_{\text{int}}, L_{\text{int}}]^N$, this framelet is strictly bad.

Similar calculations can be done in H^s , although $R_{\vec{b}}$ must be redefined (and now depends on \vec{b}). Thus, given the result of this calculation, we define the set of outgoing Gaussian's to be the set of $(\vec{a}, \vec{b}) \in \mathbb{Z}^N \times \mathbb{Z}^N$ such that:

1. $\vec{a}x_0 \in [-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$
2. $B(t)$ does not intersect $[-L_{\text{int}}, L_{\text{int}}]^N$ for any $t < T_{\text{max}}$.

A simple criterion to determine whether a given framelet is outgoing is to check whether $\vec{a}x_0 \in [-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$, and whether $\vec{b}_n k_0 \geq 2 \ln \epsilon / \sigma$. Here, \vec{b}_n is the component of $\vec{b}k_0$ pointing in the direction normal to $[-L_{\text{int}}, L_{\text{int}}]^N$. Conversely, this means that to filter waves with frequency k_{min} (or higher) making an error ϵ , then $\sigma \geq \ln \epsilon / k_{\text{min}}$. This determines the width of the buffer region. This appears to be somewhat larger than what a PML requires, although we were unable to obtain high accuracy using current implementations of the PML (see Remark 3.1).

The set of ambiguous framelets is the set of framelets for which $B(t)$ intersects $[-L_{\text{int}}, L_{\text{int}}]^N$, but also intersects $\mathbb{R}^N \setminus [-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N$. That is to say, these are framelets with both incoming and outgoing components.

A calculation based on (1.3) shows that these consist of framelets with either low velocities, or framelets which are located on the boundary but are moving tangentially to the boundary.

To assess the time complexity of the TDPSF algorithm, we recall remark 2.2, which says that when $L_{\text{int}} \gg w$, the cost of computing \mathcal{P}_{OUT} (by Algorithm 2.2) is only $O(L_{\text{int}}^{N-1} k_{\text{max}}^N \log(L_{\text{int}} k_{\text{max}}))$. Thus, in this case, we find that the added complexity of TDPSF propagation over the Split Step algorithm is of lower order than the Split Step itself. However the constant is significantly larger in our experiments, so on a small grid the TDPSF propagator is slower than FFT/Split Step propagation.

We discuss some possible improvements on this algorithm in section 5.

3 Numerical Examples

In this section we discuss the results of our numerical tests.

The TDPSF algorithm is built in the program Kitty. Kitty is written in Python, with C extensions, calling the external libraries FFTW [12], Numarray and Matplotlib. The external programs gnuplot, ImageMagick and gifsicle were also used for making graphs/movies.

Kitty is licensed under the GPL. It is very much a work in progress, and has little documentation and minimal user interface (an end-user version is currently in progress). Various test cases, spanning many types of parameters, are also available for download from the author's webpage, <http://math.rutgers.edu/~stucchio>.

3.1 Simple Tests: Free Schrödinger Equation

The standard method for testing absorbing boundaries is simply to throw coherent states (which are well localized in frequency) at the boundary and compute the difference between the approximate result and the true solution. This is a useful test, although it by no means completely characterizes the errors, as we discuss in Section 3.2.

Our implementation is as follows. The free Schrödinger equation (i.e. $g(t, \vec{x}, \psi(x, t)) = 0$) was solved on $[-102.4, 102.4]$ with TDPSF boundaries on the regions $[-102.4, -88]$ and $[88, 102.4]$. The lattice spacing was $\delta x = 0.1$ in space, 5×10^{-4} in time, and the TDPSF filtering interval was $T_{\text{step}} = 2 \times 10^{-3}$. The initial condition was taken to be $\psi(x, 0) = e^{-x^2/4} e^{ivx}$ with v ranging from 1 to 25. The simulation was run from $t = 0$ out to $T_{\text{max}} = 3 \cdot 51.2/v$.

After the solution was given sufficient time to exit the computational domain, the error in the region $[-88, 88]$ was measured (comparing with the exact result (2.9)). The error as a function of velocity is graphed in Figure 1. The TDPSF was used with $\sigma = 1, 2, 4$. The results are comparable to the complex absorbing potential $V(x) = -25ie^{-(x-25.6)^2/16}$, which is also shown in Figure 1. The width of the complex potential was chosen so that its spatial extent is comparable to the width of the TDPSF used. Altering the height of the potential changed only position of the dip in Figure 1, and achieved little other benefit.

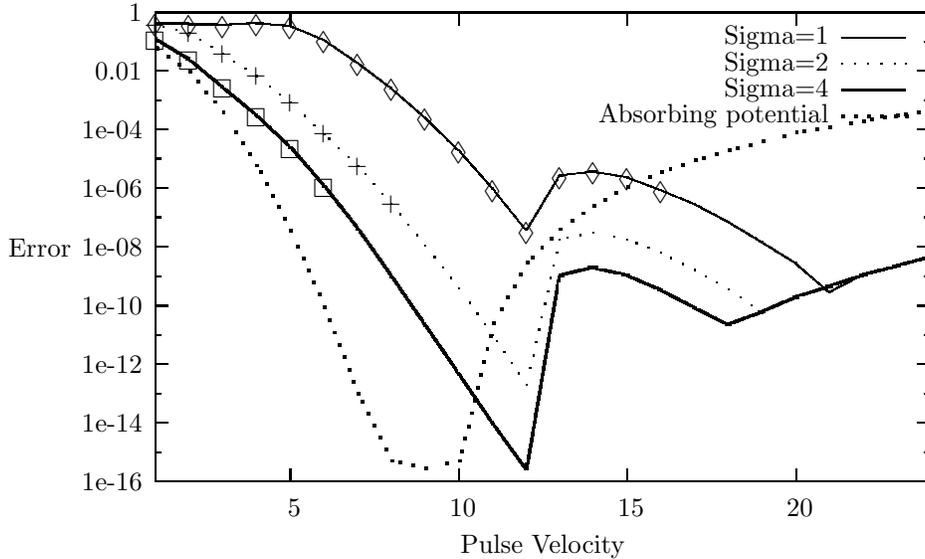


Figure 1: A graph of error vs the velocity of an outgoing pulse. The boxes, diamonds and “+” signs indicate TDPSF error reports.

In step 2c of Algorithm 2.4, the TDPSF checks whether it is making an error. The above simulations were run with tolerance 10^{-6} , and the simulations for which the TDPSF reported an error are indicated in Figure 1. Each time the TDPSF failed (and 4 times when it was successful), it notified the user of the error. The absorbing potential was unaware of it’s own failures.

This particular example demonstrates no major advantage of the TDPSF over the absorbing potential apart from the awareness of errors when they occur. The TDPSF works better for high frequencies, but not low. The advantage of the TDPSF is not that it successfully dissipates outgoing waves, but that it does not dissipate incoming waves. An example of this will be demonstrated in the next section.

Remark 3.1 Using the programs from [15] (available at <http://www.math.unm.edu/~hagstrom/>), we ran a similar set of simulations to the ones above using the PML for boundary absorption. Since FDTD methods rather than spectral methods were used, we took $\delta x = 6.7 \times 10^{-4}$, $\delta t = 2.4 \times 10^{-6}$, and a PML with a thickness of 500 lattice points¹. The error was of order $10^{-2} - 10^{-3}$, probably due to the use of FDTD instead of spectral methods for solving the interior problem.

¹The computational region was $[-20, 20]$ with 30,000 lattice points. Solving on $[-100, 100]$ would require $5 \times 30,000 = 150,000$ lattice points (the TDPSF/spectral solver uses 512 lattice points for the region $[-20, 20]$).

3.2 Harder Tests: Long Range Potentials

The main problem associated with an absorbing potential or PML is that not all waves located near the boundary are outgoing. The problem is that some waves are incoming, and should not be dissipated.

Consider the following linear Schrödinger equation (with $(\vec{x}, t) \in \mathbb{R}^{2+1}$):

$$\begin{aligned} i\partial_t\psi(x, t) &= \left[-(1/2)\Delta - \frac{15}{0.05|\vec{x}|_2^2 + 1} \right] \psi(x, t) \\ \psi(x, 0) &= e^{i7x_2} e^{-|\vec{x}|_2^2/20} + e^{i4x_1} e^{-|\vec{x}|_2^2/20} \end{aligned} \quad (3.1)$$

The initial condition consists of two coherent states of equal mass, one with velocity 4 and one with velocity 7. The notable fact about this particular potential is that the fast Gaussian has enough kinetic energy to (mostly) escape from the binding potential, while the slow Gaussian does not. The slow Gaussian moves toward the boundary, turns around and returns.

The problem with the absorbing potential approach is that it does not distinguish between incoming and outgoing waves. It dissipates everything on the boundary including the waves that should have returned. We believe this should also occur for the PML (at least the PML found in [15]).

We ran three simulations of (3.1). The first was performed using the TDPSF with $\sigma = 2.0$, $x_0 = 0.8$, $k_0 = 2\pi/12.8 = 0.491$ and $T_{\text{step}} = 0.1$. The region of computation was $[-25.6, 25.6]^2$, with lattice spacing $\delta x = 0.2$ in space and $\delta t = 0.025$ in time. The region of interest was $[-10, 10]^2$. Note that even if the region of interest were larger (as in Section 3.1), the width of the TDPSF region would not grow². An identical simulation was performed (on the same region) with an absorbing potential

$$V_1(\vec{x}) = -20ie^{-(\vec{x}_1 \pm 25.6)^2/36} - 20ie^{-(\vec{x}_2 \pm 25.6)^2/36}$$

The third was solved with periodic boundary conditions on $[-204.8, 204.8]^2$ (all other parameters were the same). This boundary is sufficiently distant so that the outgoing waves cannot return to the origin for a time $408.6/7.0 \approx 58$. Letting $\psi_t(x, t)$ be the solution with the TDPSF boundary, $\psi_a(x, t)$ be the solution with the absorbing boundary and $\psi_d(x, t)$ be the solution with the distant boundary, we measured the relative error:

$$E_{t,a}(t) = \frac{\|\psi_{t,a}(x, t) - \psi_d(x, t)\|_{L^2([-10,10]^2)}}{\|\psi_0(x)\|_{L^2(\mathbb{R}^2)}}$$

Figure 2 plots the error. In total, the error for the absorbing potential reaches 7.0% by $t = 58$, while the TDPSF makes an error of 5.0%. The reason the error is so large for the TDPSF is that there is a concentration of mass near

²We could have performed this test on the box $[-204.8, 204.8]^2$, with the TDPSF region still having width 15.6. But comparisons with an accurate simulation on a larger box would be impossible.

$k = 0$ near the boundary. This poses a problem for the absorbing potential, as Figure 2 shows.

However, the problem with the absorbing potential is even more serious than it appears at first glance. The true solution $\psi(x, t)$ can be written as $\psi(x, t) = \sum_j e^{-iE_j t} \psi_j \phi_j(x) + \psi_d(x, t)$, with $\phi_j(x)$ the bound states and $\psi_d(x, t)$ the dispersive part of $\psi(x, t)$. According to our simulations, by approximately $t = 30$, $\psi_d(x, t)$ appears to have dispersed. Therefore, after this point, all that should remain are bound states.

This implies that $M(t) = \|\psi(x, t)\|_{L^2([-10, 10]^2)}$ should remain approximately constant after $t = 30$. This behavior can be seen in the distant boundaries simulation as well as the TDPSF simulation. This does not occur for the absorbing potential. Figure 3 graphs $M(t)$ for $t \in [0, 60]$ (in the distant boundaries case) and $t \in [0, 120]$ otherwise. The absorbing potential appears to be dissipating, although we know this is impossible. For larger times, even the qualitative behavior of the absorbing potential simulation is wrong.

The reason the TDPSF performs better than the complex potential is that it distinguishes outgoing waves from incoming waves. The TDPSF only removes waves which sit on the boundary and are also clearly outgoing. Low frequency waves on the boundary, which return to the computation region are incorrectly dissipated by the absorbing potential, but (correctly) ignored by the TDPSF. This is why the TDPSF gives the correct long time behavior.

This problem is more dangerous than this example suggests. Although in this example we can determine that the dissipation is artificial, we cannot always do so. Similar problems occur often in atomic physics, but with exponentially decaying resonances as well as bound states (corresponding to complex E_j). In this case, one cannot distinguish the dissipation caused by the absorbing potential from that caused by the dynamics, and the measured decay rate will be incorrect. Since the object of such a simulation is often the measurement of the decay rate, this is a serious problem.

3.3 Soliton Filtering

Consider the nonlinearity, $g(t, \vec{x}, \psi(\vec{x}, t)) = -|\psi(\vec{x}, t)|^2$. It is desirable to construct a numerical algorithm which filters outgoing solitons as well as free waves. Although the TDPSF was not designed to filter solitons, it turns out to work well for solitons which are not moving slowly. This is however a coincidence, and cannot be expected to hold in general. A referee pointed out that it will fail for the KdV equation.

The reason for this is that an outgoing soliton with sufficiently high velocity is localized in phase space on outgoing waves. Consider a simple soliton, $\phi(x, v, t) = 2^{-1/2} e^{i(vx + (1-v^2)t)} / \cosh((x - vt))$. A simple calculation shows that for $|k - v| \gg 1$, $\hat{\phi}(k, v, t) \sim e^{-|k-v|}$. Thus, for $k \gg k_{\min}$, this shows that the framelet coefficients of $\phi(x, v, t)$ which are moving too slowly to resolve have exponentially small mass. This shows that under the free flow this soliton is strictly outgoing.

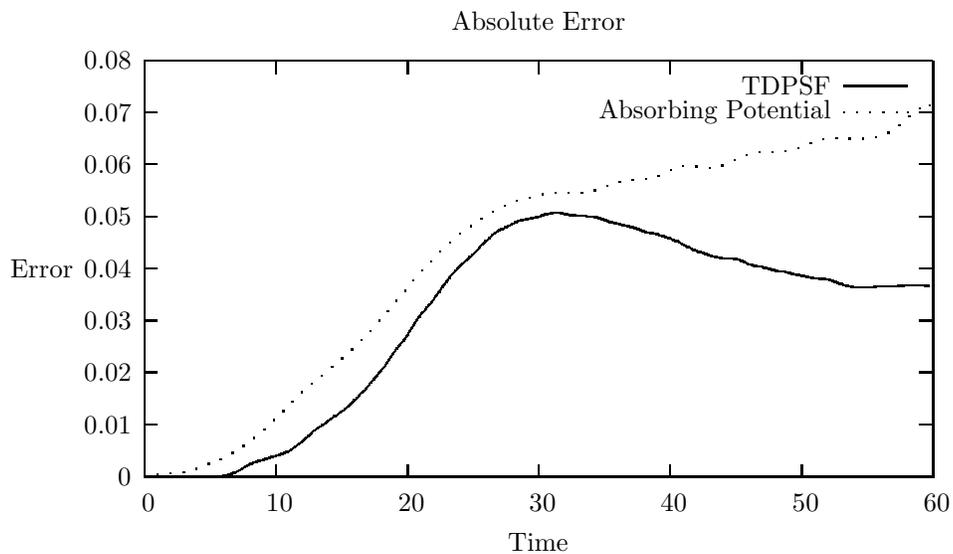


Figure 2: A graph of the error in $L^2([-10, 10]^2)$ for the absorbing potential and TDPSF up to $t = 58$.

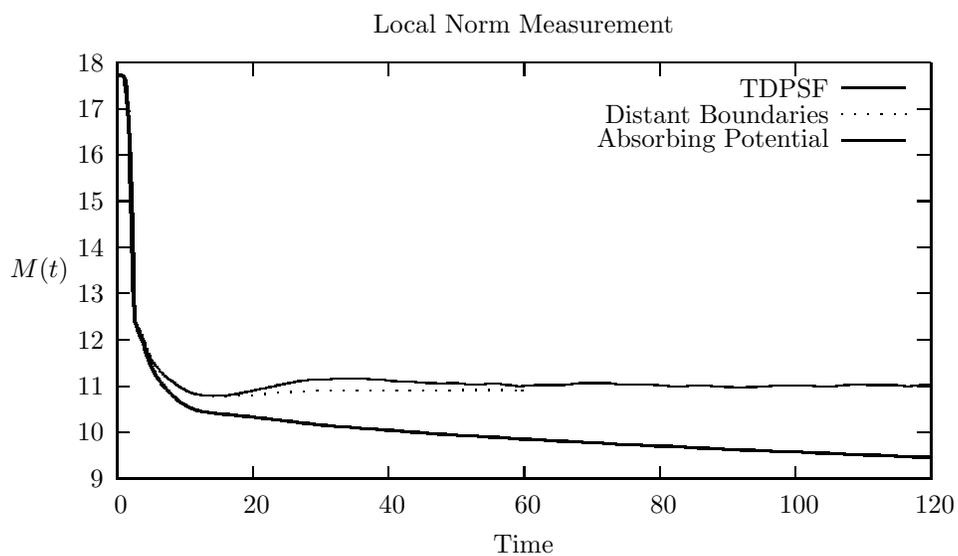


Figure 3: A graph of $M(t) = \|\psi(x, t)\|_{L^2([-10, 10]^2)}$ vs t . The distant boundary simulation is invalid at time $t = 58$, due to the return of outgoing pulse.

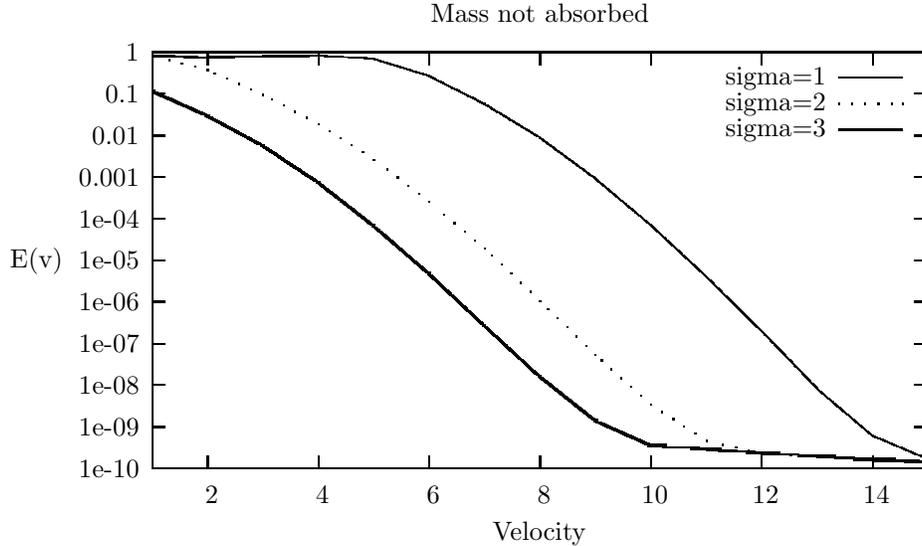


Figure 4: A plot of the error (defined as in (3.2)) as a function of velocity. Note the exponential improvement in accuracy with velocity.

The soliton is also leaving the box under the full flow $\mathcal{U}(t)$. Although $e^{i(1/2)\Delta t}$ and $\mathcal{U}(t)$ move the soliton very differently (one dispersively, one coherently), they both move it out of the box and in nearly the same direction. For this reason, we expect the TDPSF is to filter soliton solutions correctly.

We ran numerical tests to demonstrate this as follows. We solved (1.1) with $g(t, \vec{x}, \psi(\vec{x}, t)) = -|\psi(\vec{x}, t)|^2$ on the region $[-25.6, 25.6]$ with $\delta x = 0.05$, $\delta t = 0.002/v$ and $T_{\text{step}} = 0.008/v$ (the timestep's are scaled with the velocity to speed up the simulations). In this simulation, $L_{\text{int}} = 12.0$ and $w = 13.6$. The initial condition was taken to be $\psi(x, 0) = 2^{-1/2} e^{ivx} / \cosh(x)$ for $v = 1..15$.

The TDPSF was used with $T_{\text{step}} = 0.08/v$, $x_0 = 0.20$, $k_0 = 2\pi/3.2$, and $\sigma = 1, 2, 3$. We measured the following quantity:

$$E(v) = \sup_{t < 200/v} \frac{\|\psi(x, t) - \psi_{ex}(x, t)\|_{L^2([-12, 12])}}{\|\psi_{ex}(x, 0)\|_{L^2(\mathbb{R})}} \quad (3.2)$$

The function $\psi_{ex}(x, t)$ is the exact solution. The result of this experiment is plotted in Figure 3.3. The time $200/v$ was chosen since it is more than enough time for errors to return to the region $[-10, 10]$. We believe the error floor near 10^{-10} visible in Figure 3.3 is due to truncation errors in the calculation of the WFT, and errors due to time-stepping in solving the interior problem.

Remark 3.2 The paper [26] proposes an alternative method of absorbing boundaries (namely the paradifferential strategy), based on a novel method of ap-

proximating the Dirichlet-to-Neumann operator. A similar numerical test was performed for those boundary conditions. For a soliton at velocity 15, Szeftel obtained $E(15) = 0.08$ at best. For comparison, we obtain $E(15) = 1.69 \times 10^{-10}$ for $\sigma = 1$ and $E(v) = 1.40 \times 10^{-10}$ for $\sigma = 3$.

Our tests are not directly comparable. Our region of interest was $[-12, 12]$ with TDPSF region on $[-25.6, -12]$ and $[12, 25.6]$, as opposed to Szeftel who used $[-5, 5]$ (although we used 256 spatial lattice points rather than 200, as used in [26]). We used FFT/Split Step propagation for the interior problem, as opposed to the finite differences of [26]. Additionally, storage of the history on the boundary was required in [26], unlike the TDPSF.

It is quite surprising that the TDPSF beats the Dirichlet-to-Neumann boundary by such a large margin because [26] takes the nonlinearity into account while the TDPSF assumes the nonlinearity is zero on the boundary.

4 Comparison to Other Methods

4.1 Dirichlet-to-Neumann Boundaries

The Dirichlet-to-Neumann map was originally constructed by Engquist and Majda in [11, 9] (see also [1, 2]). Their guiding principle was that near the boundary, the physical optics approximation to wave flow is sufficiently accurate to filter off the outgoing waves. The TDPSF is a direct analogue of this - the Gaussian framelet elements behave (under the free flow) like classically free particles. We use a different method to filter, but the guiding principle is the same.

Modern approaches attempt to construct the exact solution on the boundary and then impose it as a boundary condition. In principle, this is the best possible approach, although in practice it is difficult to implement for dispersive equations. We briefly mention two major approaches that we are aware of, and remark that only one [26] even attempts to deal with nonlinear equations.

An additional problem is that this approach precludes the use of spectral methods (e.g. Algorithm 2.3) to solve the interior problem, reducing accuracy of the solution on the interior for dispersive equations. The FFT naturally imposes periodic boundaries rather than Dirichlet-to-Neumann, thus requiring the use of FDTD or other local methods.

4.1.1 Current Constructions

To deal with the free Schrödinger equation (no nonlinearity or potential), Lubich and Schädlé [16, 19] constructed an approximation to the exact integral kernel of the Dirichlet-Neumann operator based on a piecewise exponential (in time) approximation. This approach appears to work nicely for the free Schrödinger equation, although it is uncertain that it could be applied to the full Dirichlet-to-Neumann operator of a nonlinear equation or long range problem.

We are aware of only one fully nonlinear Dirichlet-to-Neumann operator, constructed by J. Szeftel in [25]. Szeftel uses a modified version of the paradiifferential calculus (see references in [25]) in 1 space dimension to deal with

smooth nonlinearities and potentials. He proves local well posedness of (1.1) of the boundary conditions assuming H^6 regularity of the initial data.

However, extensions to \mathbb{R}^N appear highly nontrivial. The assumptions are significantly stronger than ours, and there are no error bounds. The numerical experiments look promising and the results appear accurate for radiative problems (see also Remark 3.2).

4.2 Absorbing Potentials/ PML

4.2.1 Absorbing Potentials

Absorbing (complex) potentials, described in [17], are the current “industry standard”. One can add a dissipative term $-ia(x)\psi(\vec{x}, t)$ to the right side of (1.1) and solve it on the region $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N$. The function $a(x)$ is a positive function supported in $[-(L_{\text{int}} + w), (L_{\text{int}} + w)]^N \setminus [-L_{\text{int}}, L_{\text{int}}]^N$. This has the effect of (partially) absorbing waves which have left the domain of interest, although it might create spurious reflections and dissipation. This approach is the mainstay of absorbing boundaries, being simple to program and compatible with spectral methods.

The potential $a(x)$ must be tuned to the given problem. Given $k_{\text{min}}, k_{\text{max}}$, one must select the height and width of the absorber so that it kills most of the wave between k_{min} and k_{max} . Waves with momentum lower than k_{min} are mostly reflected, and waves with momentum higher than k_{max} are mostly transmitted and wrap around the computational domain.

Heuristic calculations and numerical experiments suggest that the absorber must have width proportional to $Ck_{\text{max}} \ln(\epsilon)/k_{\text{min}}$, with C depending on the specific shape of the potential. The TDPSF works on a layer of width $C \ln(\epsilon)/k_{\text{min}}$, which is smaller by a factor of k_{max} . Spurious dissipation is an additional problem with absorbing potentials, as illustrated in Section 3.2.

4.2.2 Perfectly Matched Layers

Perfectly Matched Layers (PML) are a variation on this approach, first proposed in [3] for Maxwell’s equations. In [15] they are designed and tested for the 1 dimensional free Schrödinger equation, with reasonable results.

In the special case of the Schrödinger equation³, the PML consists of adding a term $-ia(x)(1/2)\Delta\psi(\vec{x}, t)$ to the right side of (1.1) instead of merely $-ia(\vec{x})\psi(\vec{x}, t)$. If $a(\vec{x})$ is chosen carefully, one can completely eliminate reflections at the interface (the boundary of $\text{supp} a(\vec{x})$) for certain frequencies. It also has the property of dissipating high frequency waves more strongly than low frequency ones, thereby requiring a boundary layer of size only $C \ln(\epsilon)/k_{\text{min}}$.

As discussed in remark 3.1, numerical tests using the PML of [15] achieved errors of only 10^{-3} . The PML is likely to have the same problem as complex

³In general, one constructs the PML by exterior complex scaling, but for the Schrödinger equation it takes this simple form.

absorbing potentials with spurious dissipation. Some PML methods are unstable [18]. The PML method for the Schrödinger equation is still very much undeveloped, so a more detailed comparison is difficult to make.

5 Conclusion

We have described in this work a new method, the Time Dependent Phase Space Filter, of filtering outgoing waves for the nonlinear Schrödinger equation. Unlike absorbing potentials and (most likely) PML methods, the TDPSF filters only those regions of phase space containing outgoing waves. Dissipative terms on the boundary filter all waves, including waves which should have returned to the computational region from the boundary (see Section 3.2).

Our method is easier to construct than the Dirichlet-Neumann map, being based on standard techniques of signal processing rather than complicated pseudo/para-differential calculus. It is local in time, whereas the Dirichlet-Neumann map is history dependent. The TDPSF is also compatible with spectral methods, unlike the Dirichlet-to-Neumann map.

Unlike all other methods we are aware of, the TDPSF fails gracefully, in the sense that when it cannot filter the outgoing waves it notifies the user rather than reporting an incorrect calculation as correct.

The main problem with the method is the problem of filtering low frequency outgoing waves, a problem shared with all other methods. We are currently working on an extension of the TDPSF [23] which will reduce the cost of filtering from $C \ln(\epsilon)/k_{\min}$ to $C \ln(\epsilon) \ln(k_{\min})$.

Acknowledgements: We thank O. Costin and M. Kiessling for useful discussions, and R. Falk for reading the manuscript. This work was supported by NSF grant DMS01-00490.

References

- [1] Alvin Bayliss and Eli Turkel. Radiation boundary conditions for wave-like equations. *Comm. Pure Appl. Math.*, 33(6):707–725, 1980.
- [2] Alvin Bayliss and Eli Turkel. Outflow boundary conditions for fluid dynamics. *SIAM J. Sci. Statist. Comput.*, 3(2):250–259, 1982.
- [3] Jean-Pierre Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.*, 114(2):185–200, 1994.
- [4] John P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover Publications, Mineola, NY, 2001.
- [5] Ingrid Daubechies. Time-frequency localization operators: a geometric phase space approach. *IEEE Trans. Inform. Theory*, 34(4):605–612, 1988.

- [6] Ingrid Daubechies. *Ten lectures on wavelets*, volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
- [7] Ingrid Daubechies and A. Grossmann. Frames in the Bargmann space of entire functions. *Comm. Pure Appl. Math.*, 41(2):151–164, 1988.
- [8] Ingrid Daubechies, A. Grossmann, and Y. Meyer. Painless nonorthogonal expansions. *J. Math. Phys.*, 27(5):1271–1283, 1986.
- [9] Björn Engquist and Andrew Majda. Absorbing boundary conditions for numerical simulation of waves. *Proc. Nat. Acad. Sci. U.S.A.*, 74(5):1765–1766, 1977.
- [10] Bjorn Engquist and Andrew Majda. Absorbing boundary conditions for the numerical simulation of waves. *Math. Comp.*, 31(139):629–651, 1977.
- [11] Björn Engquist and Andrew Majda. Radiation boundary conditions for acoustic and elastic wave calculations. *Comm. Pure Appl. Math.*, 32(3):314–358, 1979.
- [12] Matteo Frigo and Steven G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. special issue on "Program Generation, Optimization, and Platform Adaptation".
- [13] Dan Givoli and Beny Neta. High-order non-reflecting boundary scheme for time-dependent waves. *J. Comput. Phys.*, 186(1):24–46, 2003.
- [14] Thomas Hagstrom. Radiation boundary conditions for the numerical simulation of waves. In *Acta numerica, 1999*, volume 8 of *Acta Numer.*, pages 47–106. Cambridge Univ. Press, Cambridge, 1999.
- [15] Thomas Hagstrom. New results on absorbing layers and radiation boundary conditions. In *Topics in computational wave propagation*, volume 31 of *Lect. Notes Comput. Sci. Eng.*, pages 1–42. Springer, Berlin, 2003.
- [16] Christian Lubich and Achim Schädle. Fast convolution for nonreflecting boundary conditions. *SIAM J. Sci. Comput.*, 24(1):161–182 (electronic), 2002.
- [17] Daniel Neuhauser and Micheal Baer. The time-dependent schrodinger equation: Application of absorbing boundary conditions. *J. Chem. Phys.*, 90(8), 1989.
- [18] D. Gottlieb S. Abarbanel and J.S. Hesthaven. Long time behavior of the perfectly matched layer equations in computational electromagnetics. *Journal of Scientific Computing*, 17(1), 2002.
- [19] Achim Schädle. Non-reflecting boundary condition for a Schrödinger-type equation. In *Mathematical and numerical aspects of wave propagation (Santiago de Compostela, 2000)*, pages 621–625. SIAM, Philadelphia, PA, 2000.

- [20] Achim Schädle. Non-reflecting boundary conditions for the two-dimensional Schrödinger equation. *Wave Motion*, 35(2):181–188, 2002.
- [21] I. M. Sigal and A. Soffer. Asymptotic completeness of short-range many-body systems. *Bull. Amer. Math. Soc. (N.S.)*, 14(1):107–110, 1986.
- [22] I. M. Sigal and A. Soffer. The N -particle scattering problem: asymptotic completeness for short-range systems. *Ann. of Math. (2)*, 126(1):35–108, 1987.
- [23] A. Soffer and C. Stucchio. Multiscale resolution of shortwave-longwave interactions in time dependent dispersive waves. *in preparation*.
- [24] A. Soffer and C. Stucchio. Time dependent phase space filters: Nonreflecting boundaries for semilinear schrodinger equations. *preprint*, 2006.
- [25] Jérémie Szeftel. Design of absorbing boundary conditions for Schrödinger equations in \mathbb{R}^d . *SIAM J. Numer. Anal.*, 42(4):1527–1551 (electronic), 2004.
- [26] Jeremie Szeftel. Absorbing boundary conditions for nonlinear schrodinger equations. *Numerische Mathematik*, (104):103–127, 2006.