

**A Preconditioned Newton-Krylov Method for Computing
Steady-State Pulse Solutions of Mode-Locked Lasers**

by

Jonathan R. Birge

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of
Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY


February 2008

© Massachusetts Institute of Technology 2008. All rights reserved.

Author *Jonathan R. Birge*
School of Engineering
January 18, 2008

Certified by *Jacob K. White*
Jacob K. White
Professor of Electrical Engineering
Thesis Supervisor

Accepted by *Jaime Peraire*
Professor of Aeronautics and Astronautics
Codirector, Computation for Design and Optimization Program



ARCHIVES

0001 0002 0003 0004

0005 0006 0007 0008

A Preconditioned Newton-Krylov Method for Computing Steady-State Pulse Solutions of Mode-Locked Lasers

by

Jonathan R. Birge

Submitted to the School of Engineering
on January 18, 2008, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

We solve the periodic boundary value problem for a mode-locked laser cavity using a specially preconditioned matrix-implicit Newton-Krylov solver. Solutions are obtained at least an order of magnitude faster than with dynamic simulation, the standard method. Our method is demonstrated experimentally on a one-dimensional temporal model of an eight femtosecond mode-locked laser operating in the dispersion-managed soliton regime.

Our solver is applicable to finding the steady-state solution of any nonlinear optical cavity with moderate self phase modulation, such as those of solid state lasers, and requires only a model for the round-trip action of the cavity. We conclude by proposing avenues of future work to improve the method's convergence and expand its applicability to lasers with higher degrees of cavity nonlinearity. Our approach can be extended to spatio-temporal cavity models, potentially allowing for the first feasible simulation of the full dynamics of Kerr-lens mode locking.

Thesis Supervisor: Jacob K. White

Title: Professor of Electrical Engineering

Acknowledgments

First, I would like to thank Dr. Peraire and Dr. Freund for starting the CDO program. It has been invaluable in my research, and is a great asset to MIT. I'd also like to thank all of the professors who taught core classes in the CDO program. When I came to MIT, I expected the main focus of the faculty to be their research, and it was a very pleasant surprise to find professors who are so dedicated to teaching.

I'd like to express my appreciation to Dr. White for being nice enough to be my thesis advisor while on sabbatical, and to Dr. Peraire for being willing to read my thesis on short notice.

Laura Koller deserves a tremendous deal of gratitude for the job she does as the administrator of CDO. She consistently goes beyond the call of duty to help, and looks out for students in a way that is truly extraordinary. I'm certain I would have not graduated without her help.

Finally, I'd like to thank my Ph.D. advisor, Dr. Kärtner, for being so generous in allowing me to pursue the CDO degree in parallel with my doctoral research.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction	13
1.1	Background	13
1.2	Contribution	15
1.3	Organization	17
2	Physical Model	19
2.1	Nonlinear wave equation	19
2.2	Split-step method	23
2.2.1	Derivation	23
2.2.2	Convergence	24
2.3	Dispersion managed soliton mode-locking	25
2.4	Laser cavity numerical model	27
2.4.1	Dispersion	27
2.4.2	Gain Material	28
2.4.3	Fast Saturable Absorber	28
2.5	Titanium:sapphire test model	29
3	Newton-Krylov Solver	31
3.1	Problem Statement	31
3.2	Jacobian Properties	32
3.3	Diagonal Preconditioner	33
3.4	Krylov subspace solver	35
3.5	Theoretical Convergence	37

4	Results and Summary	39
4.1	Accuracy	39
4.2	Empirical convergence	39
4.3	Future Work	42
4.3.1	Better preconditioning for high SPM	42
4.3.2	Alternative linear solvers	43
4.3.3	Line search implementation	43
4.3.4	Phase normalization handling	43
4.3.5	Reduced basis sets	43
4.3.6	Parallelization	44
4.3.7	Extension to full spatio-temporal model	44
A	Code Appendix	49
A.1	Cavity round trip function	49
A.1.1	Nonlinear propagation	51
A.1.2	Phase normalization	51
A.2	Solver	52
A.3	Preconditioner	55

List of Figures

2-1	An example of a Kerr lens mode-locked laser (top) and a schematic of the 1D model for the cavity used as our test problem. (GVD: group velocity dispersion; SA: saturable absorber; SPM: self-phase modulation.)	26
2-2	Illustration of the pulseshaping mechanism of a dispersion managed soliton laser over one round-trip. Adapted from [13].	27
2-3	Evolution of model laser starting from noise, shown on a log scale. The field intensity as a function of time is on the left, with the power spectral density shown on the right. The pulse does not exactly follow our moving window because of nonlinear effects shifting the spectrum from our assumed center frequency.	30
2-4	Convergence of dynamic cavity evolution. Once the initial transients die, the convergence of the pulse shaping mechanism is linear.	30
3-1	Jacobian of model shooting problem (3.14) near stationary point: (a) log magnitude, (b) phase.	33
3-2	SVD spectrum for Newton problem before (left) and after (right) preconditioning. Note the different scales.	35
4-1	Left: comparison of the amplitude of the final solutions obtained by our method (dots) and standard dynamical evolution (solid line) for our test mode-locked laser model. Right: absolute value of difference between the two solutions.	40

4-2	Comparison of convergence between our method (left curve) and standard dynamical evolution (right curve) for our test mode-locked laser model. All cavity evaluations are included, including those used to solve the linear Newton problems.	40
4-3	Convergence map of Newton-Krylov solver for Gaussian starting iterates of various widths and amplitudes for a ten femtosecond laser model (chosen for its quick convergence). The darker colors represent fewer steps, with the outer red region starting points that did not ever converge. The vertical axis is the amplitude and the horizontal the width of the starting Gaussian. The actual solution to the model is best approximated by the dot at roughly (2500, 10). This map suggests that the method will converge for a wide range of starting guesses, but that starting with energetic short pulses is a good strategy in the absence of any information about the true solution.	41
4-4	The sequence of residuals produced by the direct solver as a function of the number of round trip evaluations, log scaled.	42

List of Tables

2.1	Model parameters for Titanium:sapphire laser	29
-----	--	----

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Background

The first demonstration of the laser in 1960 [14] began an era of unprecedented control over light. In its simplest form, a laser consists of a gain medium within a low loss optical cavity. In theory, the single resonant frequency which experiences the most net gain is amplified in a laser to the point where the gain medium is saturated to match the cavity loss and only that one mode can survive. In practice, various nonidealities such as inhomogeneous broadening and hole-burning conspire against true single-frequency operation, such that most lasers actually operate with a small cluster of frequencies. However, even when multiple frequencies lase in a standard laser, the modes have randomly changing relative phases and are spaced unevenly in the frequency domain; the only effect is to broaden the spectrum and reduce the coherence length of the laser, typically to something on the order of centimeters (i.e., millions of cycles). For most purposes, therefore, conventional lasers can be considered monochromatic.

However, shortly after the realization of the first (nearly) monochromatic lasers, researchers realized that adding a strong enough nonlinear filter to the cavity could couple the cavity modes, causing multiple frequencies to not only simultaneously lase, but to “line up” in frequency and lock in phase so that they form a uniform comb of frequencies, producing a train of short pulses. This mode of pulsed laser operation is termed mode-locking, and is a prime example of nonlinear self-organization [11]. Given that carrier frequencies of light are on the order of PHz, pulses far shorter than those attainable with electronics can be created with negligible relative bandwidth.

The simplest nonlinear effect that can be used to promote mode-locking is so-called “slow” saturable absorption, whereby a material is placed in the cavity that attenuates light at a rate inverse to the light intensity. It is termed slow because the absorption changes due to real changes in the population of electronic levels in the material, which have an intrinsic recovery time that limits the time scale of the pulses produced. Nonetheless, in 1966, DeMaria [9] used a saturable absorber dye placed in the cavity of a Nd:glass laser to produce pulses on the order of 10 ps, so short that they could not be measured by standard electronic detection. Thus began the field of ultrafast optics, and a long path towards shorter and shorter optical pulses from mode-locked lasers. The culmination of this trajectory was the development in the 1990s of “dispersion managed soliton” lasers, which rely on fast saturable absorption through the nearly instantaneous Kerr lensing effect [7], explained briefly in Section 2.3. Today, such lasers are capable of producing pulses so short that they contain less than two cycles of the optical field, and have bandwidths spanning hundreds of THz [15].

Mode-locked lasers produce light at remarkable physical extremes of both time and intensity, and are thus able to manipulate matter in unique ways [17]. The few-cycle pulses produced by the best commercially available lasers are some of the shortest electromagnetic events ever created—nearly at the the very limit of what is physically possible—and efforts are already underway at several laboratories to produce single-cycle pulses. Femtosecond pulses allow the probing of physical phenomena on a commensurate time scale, sufficient to resolve electronic relaxation processes in molecules or bonding dynamics [8], or control molecular quantum states to guide reactions [3]. With high energy pulses below three cycles or so, effects due to the *absolute* optical phase of the pulse start to appear for the first time, and techniques have emerged [23] to stabilize the carrier phase of such pulses relative to their envelopes (called CEP stabilization). One hugely promising application of such exquisitely controlled light fields is the field acceleration of outer electrons of noble gases to create UV and X-ray pulses, a process known as high harmonic generation (HHG) [18]. Soft X-ray pulses as low as 130 attoseconds [21] have been produced with this technique, which has the potential to create the first lab-scale source of coherent X-rays for use in molecular imaging. Another application of CEP stabilized pulses is in the frequency domain, where the stable frequency comb can be used as an optical “clock work” to convert optical frequencies to radio frequencies, allowing one to literally count optical cycles using

standard RF electronics [23].

A corollary of the short time span within which the energy is localized in a mode-locked laser is that tremendous intensities are created. The peak focused output coming directly from a mode-locked lasers is already on the order of 10^{12} W/cm, which is on the scale of the binding energy of outer electrons in molecules and crystals. As such, the fields associated with the laser pulse are high enough to produce significant nonlinear polarization responses in most materials.¹ For example, a significant portion of a beam of infrared light (at about 800 nm) from an unamplified mode-locked Ti:sa laser can be efficiently converted to blue (400 nm) by the use of a very thin nonlinear crystal. The peak intensities of commercially available *amplified* femtosecond lasers are over 10^{15} W/cm, sufficient to ionize matter. One novel application of such intense pulses involves creating a plasma in the interior of a transparent material, using the resulting localized heat source to write waveguides embedded in the interior [6].

1.2 Contribution

The effective numerical solution of the steady-state solution of a nonlinear cavity is essential to the design and study of mode-locked lasers [13], as well as the modeling of field enhancement cavities, such as those used in high harmonic generation. The standard method for tackling such problems is to develop a numerical model of a cavity round trip, and then use the model to explicitly simulate the dynamic operation of the cavity in question until convergence is reached at some precision [1]. While this has the advantage of demonstrating self-starting and solution stability, dynamics simulation is rather poor when viewed as a numerical algorithm. Transients in cavities inherently decay exponentially, meaning that dynamic simulation exhibits linear numerical convergence to the final solution. For example, a solid-state laser operating in the dispersion-managed soliton regime can take many thousands of round trips to converge.

In contrast, our algorithm converges quadratically to the stable solution, typically requiring the evaluation of less than a hundred round trips to converge to within numerical precision, often two to three orders of magnitude faster than with dynamic simulation [4]. This speed up is achieved by directly solving the periodic boundary value problem

¹In a classical sense, the electrons travel far enough in their light induced oscillations to “sense” the structure of the molecule or crystal and deviate from a linear path.

for the nonlinear cavity using a Newton-Raphson algorithm. At each Newton step we use a matrix-implicit, preconditioned Krylov subspace method to approximately solve the linearized problem. The preconditioning is critical to both enabling the quadratic convergence of the overall Newton iterations, as well as the efficiency of solving the linearized system as each step.

While matrix-free Krylov shooting methods have been previously applied to periodically driven RF systems [22, 16], this is the first application of such methods to a *passive* nonlinear system that exhibits periodicity as a result of nonlinear self-organization. Thus, in addition to the domain specific issue of preconditioning, this application is unique in that our period length is unknown at the outset and must be solved for along with the steady-state field.

As a matrix-implicit method, our approach only needs access to a function that computes the action of the nonlinear cavity on an arbitrary input field. As such, it can often be applied to existing problems with little additional effort from the user. In basic terms, our solver operates by sending a series of “trial” perturbations through the cavity, using the observed results to make a series of increasingly accurate estimates of the final solution. It often takes less than 100 trial cavity evaluations to obtain a solution to machine precision, whereas the natural laser dynamics might take thousands of cavity round trips to converge to the same level. This enables several new opportunities for design and analysis of mode-locked lasers, such as putting the laser model inside an optimization loop.

1.3 Organization

The structure of this thesis is as follows:

- Chapter 2: We derive the relevant mathematical model governing nonlinear wave propagation. The theory behind mode-locking is briefly introduced and the model laser system to be used throughout is described.
- Chapter 3: We present our direct solution approach, including a brief introduction to the Krylov subspace iterative linear system solver used, the ORTHOMIN method, which is based on the Generalized Conjugate Residual (GCR) algorithm.
- Chapter 4: The results of the application of our solver to the model laser are shown for various laser parameters. We also summarize the results and conclude with a suggested list of future work to improve the algorithm.
- MATLAB source code for the major components of the simulation and solver are provided in the Appendix.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Physical Model

2.1 Nonlinear wave equation

We begin by stating Maxwell's equations in differential form for the case of a nonmagnetic material that is source-free. In Gaussian units¹ the curls of the electric and magnetic fields are related by [5]

$$\nabla \times \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{H}}{\partial t} \quad (2.1a)$$

$$\nabla \times \mathbf{H} = \frac{1}{c} \frac{\partial \mathbf{D}}{\partial t}, \quad (2.1b)$$

where c is the vacuum speed of light and \mathbf{D} is the electric displacement vector. Even though there are no free charges, the motion of paired charges in the material still matters and so we must introduce a separate vector \mathbf{D} , which can be viewed as taking into account the possibility of paired charge motion (hence its appearance as a time derivative). All of the interesting action in optics really occurs in the charge dynamics of the material, and thus in some sense the most important equation is actually that relating \mathbf{D} to \mathbf{E} , known as the constitutive relation:

$$\mathbf{D} = \mathbf{E} + 4\pi\mathbf{P}(\mathbf{E}), \quad (2.2)$$

where \mathbf{P} is the polarization vector, which represents the material dipole moment at a given point in space. The polarization is driven by the electric field, and this is what allows for

¹Gaussian units are defined such that the electric and magnetic fields are related in terms of the speed of light in vacuum. This eliminates the notion of a vacuum permittivity by making \mathbf{E} and \mathbf{D} equivalent in vacuum. Similarly, vacuum permeability is gone, and $\mathbf{H} = \mathbf{B}$ in a nonmagnetic material.

nontrivial solutions, especially when the relation between \mathbf{P} and \mathbf{E} is nonlinear, as it is in our case.

To derive a wave equation, we take the curl of (2.1a) and substitute in (2.1b) to eliminate the magnetic field, leaving

$$\nabla \times \nabla \times \mathbf{E} + \frac{1}{c^2} \frac{\partial^2 \mathbf{D}}{\partial t^2} = 0. \quad (2.3)$$

We can rewrite the first term of this by using the vector identity $\nabla \times \nabla \times \mathbf{E} = \nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E}$ to give us

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} + \frac{1}{c^2} \frac{\partial^2 \mathbf{D}}{\partial t^2} = 0. \quad (2.4)$$

In this thesis, we are concerned with the temporal propagation of a laser beam inside a material. Spatial effects will either be ignored or approximated by other means. We thus assume a transversely polarized infinite planewave propagating in the z direction in a homogeneous medium. This allows us to ignore the divergence of \mathbf{E} in the first term of (2.4) and collapse the Laplacian to a spatial derivative in z . Furthermore, in the lasers we seek to model, the polarization is invariant so we may simplify (2.4) to the inhomogeneous scalar wave equation,

$$\frac{\partial^2}{\partial z^2} E(z, t) - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} D(z, t) = 0. \quad (2.5)$$

Returning to the question of how to model the material polarizability, we can assume that the vast majority of the polarization is linear with the electric field. Furthermore, given that the materials used are not centrosymmetric, we can neglect even-order nonlinearities. We thus assume the only appreciable nonlinear polarization is third-order with the electric field. Furthermore, we are only interested in self-modulation, so that we consider only those third-order terms which result in the polarization being driven at the same frequency as the electric field. Given a local monochromatic field written in phasor form, $Ee^{i\omega t} + \text{c.c.}$, this means that we consider the nonlinear contribution to the polarization phasor to be given by

$$P_{\text{NL}} = \chi^{(3)} E^* E E = \chi^{(3)} |E|^2 E. \quad (2.6)$$

In treating the polarizability $\chi^{(3)}$ as a constant scalar (in general, it is a fourth-rank tensor and a function of frequency) we imply an isotropic instantaneous nonlinearity. This is not entirely physical, of course, but it has proven in practice to be a reasonable approximation that captures the salient effects of soliton mode-locking. The nonlinear polarization thus

looks like the standard linear polarizability scaled by the amplitude of the field. As such, this effect can be roughly viewed as an “intensity dependent permittivity.” We will thus proceed with our derivation of the wave equation by treating D as resulting from a linear filter operating on E plus a perturbative nonlinear term, such that they are related in the frequency domain by

$$\tilde{D}(z, \omega) = [\epsilon(\omega) + \Delta\epsilon_{\text{NL}}] \tilde{E}(z, \omega), \quad (2.7)$$

where we will treat the nonlinear term as a constant for the time being. We define the transform of a function $G(z, t)$ as

$$\tilde{G}(z, \omega) = \int_{-\infty}^{\infty} dt G(z, t) e^{-i\omega t}. \quad (2.8)$$

Taking the Fourier transform of (2.5) and substituting in (2.7) gives the harmonic wave equation

$$\frac{\partial^2}{\partial z^2} \tilde{E}(z, \omega) + [\epsilon(\omega) + \Delta\epsilon_{\text{NL}}] \frac{\omega^2}{c^2} \tilde{E}(z, \omega) = 0,$$

where we have implicitly allowed the intensity $|E|^2$ to pass through the derivative despite it being a function of time. This is valid so long as $\partial_t |E|^2 \ll \partial_t E$, not an entirely unreasonable assumption until the envelope starts to approach a single-cycle. The quantity in front of the second term in (2.1) has units of wavenumber, and it is simpler to just write

$$\frac{\partial^2}{\partial z^2} \tilde{E}(z, \omega) + [k^2(\omega) + \Delta k_{\text{NL}}^2] \tilde{E}(z, \omega) = 0. \quad (2.9)$$

To consider pulse propagation, we assume our electric field can be written as a slowly varying envelope function modulating a monochromatic optical carrier at frequency ω_0 ,

$$E(z, t) = A(z, t) e^{i(k_0 z - \omega_0 t)} + \text{c.c.},$$

where k_0 is the carrier wave number. The analytic (positive frequency) transform of $E(z, t)$ is related to the transform of $A(z, t)$ by

$$\begin{aligned} \tilde{E}(z, \omega) &= \tilde{A}(z, \omega - \omega_0) e^{ik_0 z} + \tilde{A}^*(z, \omega + \omega_0) e^{-ik_0 z} \\ &\approx \tilde{A}(z, \omega - \omega_0) e^{ik_0 z}, \end{aligned} \quad (2.10)$$

where the approximation is obtained by noting that the envelope cannot, by construction,

have significant components at optical frequencies. Finally, we make one more assumption, which is somewhat more subtle. We neglect any backwards traveling waves, and presume that we will follow the forward traveling wave (2.1) in a forward moving frame such that little change occurs in z . Thus, we may neglect all but the lowest order derivatives in z that operate on the envelope.² Substituting (2.10) into (2.9) and keeping only the first-order spatial derivative yields

$$\frac{\partial}{\partial z} \tilde{A}(z, \omega) - i \left(\frac{k^2(\omega) + \Delta k_{\text{NL}}^2 - k_0^2}{2k_0} \right) \tilde{A}(z, \omega) = 0. \quad (2.11)$$

If we let $k(\omega_0) = k_0$, then $k(\omega)$ will not deviate much from k_0 over the pulse bandwidth. We can thus write the quantity in brackets in terms of a series expansion of $k(\omega)$,

$$\begin{aligned} \left(\frac{k^2(\omega) + \Delta k_{\text{NL}}^2 - k_0^2}{2k_0} \right) &\approx k(\omega) + \Delta k_{\text{NL}} - k_0 \\ &= -ig_0 + \Delta k_{\text{NL}} + \frac{1}{v_g(\omega_0)}(\omega - \omega_0) + \sum_{n=2}^{\infty} \frac{k_n}{n!}(\omega - \omega_0)^n. \end{aligned} \quad (2.12)$$

We have specifically labeled a small signal gain term $g_0 \equiv ik_0$ and a group velocity $v_g(\omega_0) \equiv 1/k_1$. In addition to the simplifying assumptions mentioned earlier regarding the nonlinearity, it is also now apparent that our approximate treatment of the nonlinearity implicitly restricts the nonlinearity to only affecting the phase velocity, having no effect on the group velocity.

Rewriting (2.11) in terms of (2.12), we have

$$\frac{\partial \tilde{A}}{\partial z} = g_0 \tilde{A} + i\Delta k_{\text{NL}} \tilde{A} + \frac{i}{v_g}(\omega - \omega_0) \tilde{A} + i \sum_{n=2}^{\infty} \frac{k_n}{n!}(\omega - \omega_0)^n \tilde{A}. \quad (2.13)$$

Taking the inverse Fourier transform of (2.13) and applying the derivative theorem, the powers of ω become temporal derivatives, and the time domain wave equation is

$$\frac{\partial A}{\partial z} = g_0 A + i\gamma |A|^2 A - \frac{i}{v_g} \frac{\partial A}{\partial t} - i \sum_{n=2}^{\infty} \frac{k_n}{n!} \frac{\partial^n A}{\partial t^n}, \quad (2.14)$$

where we've expressed the nonlinear wavenumber perturbation in terms of the intensity,

²This is often referred to as the slowly varying envelope approximation. However, were this really the approximating we were making, we would also have ignored higher-order time derivatives, which have been kept. In fact, it turns out that ignoring higher-order spatial derivatives, as we have done, has the sole effect of limiting us to waves traveling in one direction, and is actually an exact wave equation otherwise.

as per (2.6), and an effective parameter γ known as the nonlinear coefficient. Given that, in practice, we are actually dealing with focused beams and not the ideal plane waves assumed here, γ can be regarded as an empirically determined parameter that includes spatial effects. The nonlinear term acts in quadrature, and simply advances the phase of the field in proportion to the local intensity. The effect due to this term is thus termed self phase modulation (SPM).

Finally, we complete the derivation by rewriting the equation in a frame that is moving with the group velocity, by performing a change of variables into a “local” time $T \equiv t - z/v_g$, and writing the envelope in terms of the local field $u(z, T)$. The result of this transformation is a tilting of the solution in the z - t plane, such that time derivatives are unchanged, but spatial derivatives of A become mixed time and space derivatives in u . Application of the chain rule shows that

$$\frac{\partial A}{\partial z} = \frac{\partial u}{\partial z} - \frac{1}{v_g} \frac{\partial u}{\partial T}.$$

Going to the moving frame thus eliminates the group delay delay term, as we’d hope, and so our final propagation equation is

$$\boxed{\frac{\partial u}{\partial z} = g_0 u + i\gamma |u|^2 u - i \sum_{n=2}^{\infty} \frac{k_n}{n!} \frac{\partial^n u}{\partial T^n}.} \quad (2.15)$$

All series coefficients k_n are assumed to be complex to account for a spectrally arbitrary gain (or loss) mechanism. This equation is known as the generalized nonlinear Schrödinger equation (GNLSE).

2.2 Split-step method

2.2.1 Derivation

The GNLSE can be numerically integrated efficiently by a pseudospectral method known as the split-step method [1]. To begin with, we express (2.15) in terms of two operators,

$$\frac{\partial}{\partial z} u(z, T) = (\hat{D} + \hat{N}) u(z, T), \quad (2.16)$$

where \hat{D} represents all the linear terms (i.e. the gain/loss and dispersion) and \hat{N} is the SPM operator,

$$\hat{D} = g_0 - i \sum_{n=2}^{\infty} \frac{k_n}{n!} \frac{\partial^n}{\partial T^n}, \quad (2.17)$$

$$\hat{N} = i\gamma|u|^2. \quad (2.18)$$

The key to the split-step method is the recognition that the dispersion operator is diagonal in the Fourier domain, and the SPM operator is diagonal (and purely imaginary) in the time domain. Thus, each can be propagated *exactly* in their respective domains with a single exponential. We can efficiently compute an approximate propagation over a distance h by first handling the dispersion alone, and then transforming to the time domain and having the nonlinearity act,

$$u(z+h, T) \approx \exp(h\hat{N}) \exp(h\hat{D})u(z, T). \quad (2.19)$$

This single step requires only two FFT operations, as the exponential of a diagonal matrix is a vector operation. As will be proven in the following section, the accuracy of the iteration can be improved by symmetrizing it to yield the following iteration

$$u(z+h, T) \approx \exp\left(\frac{h}{2}\hat{D}\right) \exp(h\hat{N}) \exp\left(\frac{h}{2}\hat{D}\right) u(z, T). \quad (2.20)$$

Other than requiring a single extra dispersion propagation at the end of a computation, this requires no more computation than (2.19), and is thus the scheme used in practice (see Appendix A.1.1).

2.2.2 Convergence

To ascertain the convergence of the symmetric iteration in (2.19), we first consider the exact solution to (2.16) in terms of the operator exponential,

$$u(z+h, T) = \exp[h(\hat{D} + \hat{N})]u(z, T), \quad (2.21)$$

where \hat{N} is assumed to be invariant in z . The fact that \hat{N} is nonlinear and does not commute with \hat{D} implies that it will, in fact, vary with space. However, for small nonlinearities

this will be a negligible effect compared to the error caused by the fact that the split-step method treats everything as commuting operators, and it is this error we are analyzing.³

We can rewrite the product of the exponentials in the split-step iteration (2.20) by breaking the center time-domain step in to two equal pieces and applying the Baker-Hausdorff lemma [2] to each pair to yield

$$\begin{aligned} \exp\left(\frac{\hbar}{2}\hat{D}\right) \exp(\hbar\hat{N}) \exp\left(\frac{\hbar}{2}\hat{D}\right) = \\ \exp\left(\frac{\hbar}{2}\hat{D} + \frac{\hbar}{2}\hat{N} + \frac{1}{2}\frac{\hbar^2}{4}[\hat{D}, \hat{N}] + O[\hbar^3]\right) \\ \times \exp\left(\frac{\hbar}{2}\hat{N} + \frac{\hbar}{2}\hat{D} + \frac{1}{2}\frac{\hbar^2}{4}[\hat{N}, \hat{D}] + O[\hbar^3]\right). \end{aligned} \quad (2.22)$$

The above shows that the error in each half step is $O[\hbar^2]$ due to the fact that the two operators do not commute. Applying the lemma a second time to the right hand side of the above allows the commutators to cancel, since $[a, b] + [b, a] = 0$, leaving us with only third-order terms⁴ in excess of the exact solution,

$$\exp\left(\frac{\hbar}{2}\hat{D}\right) \exp(\hbar\hat{N}) \exp\left(\frac{\hbar}{2}\hat{D}\right) = \exp\left(\hbar\hat{D} + \hbar\hat{N} + O[\hbar^3]\right) \quad (2.23)$$

$$= \exp\left(\hbar\hat{D} + \hbar\hat{N}\right) + O[\hbar^3]. \quad (2.24)$$

Comparing the above with (2.21), we can see that the symmetric split-step method converges as $O[\hbar^3]$.

Another advantage of the split-step approach is that it allows us to handle each effect in the most natural basis. As such, dispersion and gain are not, in actuality, expressed in terms of series coefficients as in (2.15), but are instead simply represented in spectral form as the complex elements of the diagonal matrix representation of \hat{D} .

2.3 Dispersion managed soliton mode-locking

A dispersion managed soliton laser consists of three primary sections: two negative (anomalous) dispersion regions surrounding a central amplifying region with positive dispersion

³When actually performing the split-step method, of course, we allow \hat{N} to vary at each step.

⁴The third order terms not explicitly shown in (2.22) involve commutators of commutators, and thus the switching of the order of operators which allowed the second-order terms to cancel does not cause the third-order terms to cancel.




Figure 2-1: An example of a Kerr lens mode-locked laser (top) and a schematic of the 1D model for the cavity used as our test problem. (GVD: group velocity dispersion; SA: saturable absorber; SPM: self-phase modulation.)

and self-phase modulation [7], Fig. 2-1. Briefly, the effect of the alternating dispersion is to create a negatively dispersed field as the pulse enters the gain material. When the negatively chirped pulse experiences the SPM of the gain medium, it causes the spectral bandwidth to compress during its travel through the first half of the crystal. This has the effect of squeezing the spectrum to fit the gain bandwidth, thus allowing the laser to support a greater bandwidth than its gain bandwidth would otherwise allow. The discovery of this mechanism was the key breakthrough that allowed femtosecond lasers to operate well below ten femtoseconds.

The dispersion of the whole cavity is balanced such that the dispersion is exactly reversed in sign by the gain material. In a modern laser, negative dispersion is typically achieved by dispersion compensating mirrors which are engineered to take longer to reflect “red” light than “blue” light. The second half of the crystal expands the spectrum and reverses the sign of the chirp. This is illustrated in Fig. 2-2.

In addition, the cavity also must contain a fast saturable absorber of some sort, which supports the initiation of mode-locking and stabilizes the pulses. The shortest pulses are created by lasers which utilize a saturable absorption mechanism which acts instantaneously, with no appreciable recovery time. This is achieved in practice through an effect known as Kerr lensing, whereby the spatial self focusing of the beam causes brighter in-




Figure 2-2: Illustration of the pulseshaping mechanism of a dispersion managed soliton laser over one round-trip. Adapted from [13].

tensities to focus themselves into the highly pumped region of the gain material, thereby experiencing greater gain.

2.4 Laser cavity numerical model

In this thesis, we'll consider a dispersion managed soliton laser as a test case, considering only the time domain and neglecting any spatial effects (such as self-focusing or diffraction). Thus, we consider only a complex analytic pulse envelope $u(z, T)$, as discussed in Section 2.1. The cavity model we will use is shown in Fig. 2-1.

2.4.1 Dispersion

In our model, the net cavity dispersion is expressed in terms of second- and third-order series coefficients in ω , written as D_2 and D_3 , respectively. Where there are no nonlinear elements, dispersive elements (e.g. air, mirrors) may be treated trivially in lumped form. In the Fourier domain this operation is written as

$$U'(\omega) = \exp [-i (D_2\omega^2 + D_3\omega^3)] U(\omega), \quad (2.25)$$

where U is the field before and U' after the dispersive element. It would be trivial to handle arbitrary spectral dispersion profiles in the simulation, but for the sake of simplicity in this test model, we limit ourselves to two series terms.

2.4.2 Gain Material

We model the gain material via the GNLSE (2.15), handling propagation using the aforementioned split step. For the levels of nonlinearity typically encountered in a mode-locked laser, a discretization of 30 steps is generally sufficient. The gain is presumed to saturate as a function of the total intracavity power P , given by

$$P = \frac{1}{T_R} \int dT |u|^2, \quad (2.26)$$

with T_R the round trip time. The effect of the gain over a spatial step Δz is handled in the Fourier domain by

$$U(z + \Delta z, \omega) = \exp \left(\frac{g_0}{1 + P/P_{\text{sat}}} \frac{\Delta \omega}{\omega + \Delta \omega} \frac{\Delta z}{\ell} \right) U(z, \omega), \quad (2.27)$$

with g_0 the small signal gain, and $\Delta \omega$ the gain bandwidth, P_{sat} the saturation power.

Given the focusing that is occurring in the gain medium, the nonlinear parameter γ is technically a function of z . However, to simplify things we specify the nonlinearity in terms of an empirically determined net nonlinear phase per unit of intensity per pass, known as the SPM coefficient δ . Nonetheless, we still consider the SPM in distributed form, acting throughout the gain crystal. The effect of self phase modulation on the envelope is then given by

$$u(z + \Delta z, T) = \exp \left[i \delta |u(z, T)|^2 \frac{\Delta z}{\ell} \right] u(z, T). \quad (2.28)$$

The split-step method is thus implemented by alternating between (2.27) and (2.28).

2.4.3 Fast Saturable Absorber

The full spatio-temporal Kerr lensing mechanism is a prohibitively complex process to model (especially given the standard simulation methods which this thesis aims to improve upon) and thus saturable absorption is not modeled physically, but phenomenologically, using a simple lumped model given by

$$u'(T) = \exp \left(\frac{-q}{1 + |u(t)|^2 / I_{\text{sat}}} \right) u(T), \quad (2.29)$$

Table 2.1: Model parameters for Titanium:sapphire laser

Parameter	Value
T_R	$1/100 \text{ MHz}^{-1}$
$D_{\text{net}}^{(2)}$	-0.25 fs^2
g_0	0.1
$D_{\text{gain}}^{(2)}$	232 fs^2
$D_{\text{gain}}^{(3)}$	160 fs^3
$\Delta\omega$	60 THz
P_{sat}	1 W
δ	20 rad/GW^{-1}
q	-0.02
I_{sat}	0.3 MW
l_{OC}	0.02

where q is the unsaturated absorption and I_{sat} is the saturation intensity. Both parameters are determined empirically by comparing model results with actual lasers.

2.5 Titanium:sapphire test model

To test our model and solver, we implement a simulation of a typical Titanium:sapphire mode-locked laser. The model parameters we use as a baseline are presented in Table 2.1. For sake of convenience, dispersion is specified in terms of gain material dispersion and the net dispersion of an entire round-trip. We discretize the field as a 256 element vector representing a window of 200 femtoseconds, evaluated at a z slice right before the output coupler (to simulated the pulse as exiting the laser cavity).

In Fig. 2-3, we show the evolution of the cavity over 2000 round-trips, starting from random noise (representing spontaneous emission). The resulting pulse is slightly asymmetric and has a duration (as measured by its full width half maximum) of 8.56 fs. The evolution of the residual, as defined by the “energy” of the difference between the input and output of the cavity relative to the energy of the pulse, is shown in Fig. 2-4. After a period of oscillation as transients die down, the simulation converges linearly overall toward the steady state solution, as expected.




Figure 2-3: Evolution of model laser starting from noise, shown on a log scale. The field intensity as a function of time is on the left, with the power spectral density shown on the right. The pulse does not exactly follow our moving window because of nonlinear effects shifting the spectrum from our assumed center frequency.

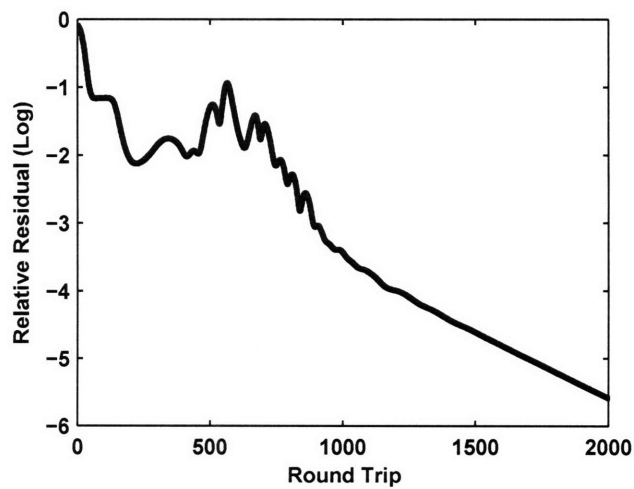


Figure 2-4: Convergence of dynamic cavity evolution. Once the initial transients die, the convergence of the pulse shaping mechanism is linear.

Chapter 3

Newton-Krylov Solver

In the context of this chapter, we define computational complexity in terms of round-trip evaluations (i.e. calls to the cavity simulation function). As demonstrated in Section 2.5, the standard method of dynamic simulation converges linearly, requiring thousands of propagations through the cavity. In fact, some lasers require 50,000 or more round trips to converge to significant precision. It is generally the rule that mode-locked lasers perform optimally in regions of marginal stability, further limiting the efficacy of dynamic simulation as a solution method.

By treating the solution of the stationary cavity mode as a nonlinear problem to be solved using Newton's method, we hope to converge directly to the solution quadratically. We have no way to efficiently compute the Jacobian of the cavity, however, and so we use a matrix-free shooting method [22], modified to the specifics of our problem. In many cases, Newton-based shooting methods do not require preconditioning [16], though we find that in our case, the shooting problem is poorly conditioned and does not converge efficiently as it stands. We derive a diagonal preconditioner, which improves the problem conditioning by a couple orders of magnitude.

3.1 Problem Statement

We regard the cavity as a nonlinear operator $\mathbf{g}(\mathbf{u})$ acting on a vector representing the Fourier coefficients of the field in our temporal window. We operate in the basis of the Fourier modes because they are eigenvectors of dispersion and gain, and thus we can trivially invert those operators of the cavity.

It's important to note that \mathbf{u} need only be sufficient to represent the solution. We can project this vector in and out of a higher dimensional space to perform the actual computation. In fact, given the “breathing” nature of dispersion managed soliton lasers, for example, it may be necessary to propagate using a much larger temporal window than is necessary to describe the solution.

To solve the problem, we seek an “eigenfunction” \mathbf{u} such that

$$\mathbf{f}(\mathbf{u}) \equiv e^{-i\phi(\mathbf{u})} \mathbf{g}(\mathbf{u}) = \mathbf{u}, \quad (3.1)$$

where $\phi(\mathbf{u})$ is a function that takes care of normalizing out a constant and linear phase,

$$\phi_i(\mathbf{u}) = \exp [-i(\Delta\phi(\mathbf{u}) - \Delta v_g(\mathbf{u})\omega_i)]. \quad (3.2)$$

Due to the presence of nonlinearity, we cannot predict ahead of time what the actual periodicity of the cavity will be, and need to account for a potential perturbation Δv_g to the group delay already assumed in Section 2.1. We also normalize the overall phase of our solution by such that the DC component of the field experiences no phase change. More details can be found in Appendix A.1.2. The core of our method is the repeated solution of linear Newton subproblems given by

$$[\mathbf{J}_f(\mathbf{u}) - \mathbf{I}] (\mathbf{u}_{k+1} - \mathbf{u}_k) = \mathbf{f}(\mathbf{u}_k), \quad (3.3)$$

with the cavity Jacobian \mathbf{J}_f defined by

$$(\mathbf{J}_f)_{ij} \equiv \frac{\partial f_i(\mathbf{u})}{\partial u_j}. \quad (3.4)$$

The bulk of the work in this thesis deals with the efficient solution of this problem.

3.2 Jacobian Properties

Unfortunately, it turns out that the Jacobian of the full problem $\mathbf{J} = \mathbf{J}_f - \mathbf{I}$ is poorly conditioned and has a full eigenspectrum. Thus, a naive application of Newton's method to (3.1) fails due to numerical truncation error in approximating the Jacobian of the cavity. Furthermore, even when it succeeds, the computation of the full Jacobian involves the




Figure 3-1: Jacobian of model shooting problem (3.14) near stationary point: (a) log magnitude, (b) phase.

evaluation of hundreds of round trips, negating much of the computational savings.

In Fig. 3-1 we show the Jacobian of our test problem at a stationary point. The Jacobian has the following salient properties, which will be relevant to a solution of the associated system:

1. Complex. This will affect our choice of solver.
2. Nearly diagonal. (Note that the plot in Fig. 3-1 is on a logarithmic scale). This suggests we can effectively precondition the system using a diagonal matrix.
3. The nondiagonal part is close to Hermitian. Thus we will have to use a generalized solver, but we may be able to treat the matrix as orthogonal in certain cases.

3.3 Diagonal Preconditioner

To derive a diagonal left preconditioner, we begin by approximating the cavity as composed off all Fourier diagonal elements followed by all time diagonal elements, approximating the cavity with a single nonsymmetric split-step propagation,

$$\mathbf{f}(\mathbf{u}) \approx \mathbf{N}(\mathbf{u})\mathbf{D}(\mathbf{u})\mathbf{u}. \quad (3.5)$$

As discussed in Chapter 2, everything in the cavity is diagonal in the Fourier domain except for the saturable absorber and SPM, which form the full matrix \mathbf{N} . However, we can actually compute the diagonal of \mathbf{N} rather efficiently, and use this to approximate

the contribution of SPM and saturable absorption to the diagonal of the Fourier domain Jacobian. The matrix representation of the time diagonal components can be written as

$$\mathbf{N} = \mathbf{F}^\dagger \lambda_{\mathbf{N}} \mathbf{F}, \quad (3.6)$$

where \mathbf{F} is the discrete Fourier transform matrix and $\lambda^{\mathbf{N}}$ is the diagonal matrix of the non-linear operator in the time domain. A diagonal matrix in the time domain is banded in the Fourier domain, and therefore the diagonal of (3.6) is constant. Since the trace of a matrix is equal to its spectral trace, the diagonal must therefore be

$$N_{ii} = \frac{\text{Tr}\{\mathbf{N}\}}{n} \quad (3.7)$$

$$= \frac{\text{Tr}\{\lambda_{\mathbf{N}}\}}{n} \quad (3.8)$$

$$\equiv \sum_k \exp(n_k) / n \quad (3.9)$$

$$\approx \exp\left[\sum_k n_k / n\right], \quad (3.10)$$

where n is the dimension of our system, and n_k are the diagonal elements of the nonlinear operator, expressed in “small signal” form for simplicity and to match how the computations are done in practice. The whole cavity function can thus be approximated by a diagonal operator

$$\mathbf{f}_i(\mathbf{u}) \approx N_{ii} D_{ii} \quad (3.11)$$

$$\approx \exp\left[\sum_k n_k(\mathbf{u}) + d_i(\mathbf{u})\right] u_i, \quad (3.12)$$

with d_i the small signal elements of the Fourier domain operators (i.e. dispersion and gain). If all of the operators are perturbative, in the sense that

$$\frac{d}{du} e^{g(u)} u \approx e^{g(u)},$$

then the cavity Jacobian \mathbf{J}_f can be approximated by the diagonal matrix

$$\mathbf{B} = \text{diag} \left\{ \exp \left[\sum_k n_k(\mathbf{u}) + \mathbf{d}(\mathbf{u}) \right] \right\}. \quad (3.13)$$




Figure 3-2: SVD spectrum for Newton problem before (left) and after (right) preconditioning. Note the different scales.

Ideally, we'd use a preconditioner than includes more than the diagonal, but this is sufficient for most regimes (with nonlinear phase shifts much less than π) as the Jacobian is itself highly diagonal. With the preconditioner, the linear Newton subproblem (3.3) becomes

$$\boxed{[\mathbf{B}(\mathbf{u}_k) - \mathbf{I}]^{-1} [\mathbf{J}_f(\mathbf{u}_k) - \mathbf{I}] (\mathbf{u}_{k+1} - \mathbf{u}_k) = -[\mathbf{B}(\mathbf{u}_k) - \mathbf{I}]^{-1} \mathbf{f}(\mathbf{u}_k).} \quad (3.14)$$

The inversion of the full system preconditioner $\mathbf{B} - \mathbf{I}$ is trivial for a diagonal \mathbf{B} .

Our simple preconditioner (3.13) results in significant reduction in the spectrum of the linear subproblem. Fig. 3-2 shows the singular value spectrum for the Newton problem for our soliton laser at its stationary point, both before and after preconditioning. The conditioning improves by roughly a factor of 100 with most of the spectrum collapsing to unity, demonstrating that our simple diagonal preconditioner nearly inverts the system.

3.4 Krylov subspace solver

The fact that we do not have direct access to the analytic Jacobian of our problem—as well as the fact that our problem is amenable to preconditioning—suggests we solve (3.14) using a matrix-implicit Krylov subspace iterative solver. Given a problem $\mathbf{A}\mathbf{x} = \mathbf{b}$, such methods search for a solution within the space spanned by the set of vectors produced by repeated multiplication of the system matrix \mathbf{A} with the residual. Assuming a null initial iterate, this gives

$$\mathcal{K}_m(\mathbf{A}) = \text{span} \{ \mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \dots, \mathbf{A}^{m-1}\mathbf{b} \}. \quad (3.15)$$

Put another way, the Krylov subspace is the space of all vectors that can be written as $\mathbf{x} = p\{\mathbf{A}\}\mathbf{b}$, where p is a polynomial of degree less than or equal to $m - 1$. Essentially, we seek to invert a matrix with a polynomial of the same matrix. The power of the Krylov subspace comes from three facets: (a) it turns out to be a very efficient space in which to search given clustered eigenvalues; (b) such a space arises naturally when iteratively solving a matrix by projection methods; and (c) the solution can be found by computing the operation of a series of vectors through the system, obviating the need to know the full Jacobian.

Among Krylov methods, the most promising for our problem are of the class called *optimal* Krylov solvers, which find solutions \mathbf{x} given by

$$\arg \min_{\mathbf{x} \in \mathcal{K}_m} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2. \quad (3.16)$$

Given the properties of the Jacobian enumerated in Section 3.2, our two main options are GMRES [20] or GCR [10]. Both methods involve finding orthogonal basis sets over which to solve (3.16). Given the generality of the system, both involve storing up to m vectors. The advantage of GMRES is that it often exhibits superior stability, and requires less space [19]. However, space is not a concern for the size of our problem, and our method will only be worthwhile if we are able to precondition well enough such that only tens of iterations will be needed to solve each linear system. We thus chose GCR for its ability to efficiently compute the residual and updated solution at each iteration. This allows us to minimize the number of round-trip function calls needed.

GCR is a variational method that involves incrementally generating an $\mathbf{A}^\dagger \mathbf{A}$ -orthogonal basis using a Gram-Schmidt orthogonalization. The GCR algorithm, slightly modified for complex systems, follows.

Algorithm 1 *Complex Generalized Conjugate Residual (GCR)*

1. Compute $\mathbf{r}_0 = \mathbf{b}$. Set the initial search vector $\mathbf{p}_0 = \mathbf{r}_0$ and $x_0 = 0$.
2. For $j = 0, 1, \dots$ until $\|\mathbf{r}_j\|_2 < \epsilon$, Do
3. $\alpha_j = \frac{\Re(\mathbf{r}_j^\dagger \mathbf{A}\mathbf{p}_j)}{\|\mathbf{A}\mathbf{p}_j\|_2^2}$
4. $\mathbf{x}_{j+1} = \mathbf{x}_j + \alpha_j \mathbf{p}_j$

5. $\mathbf{r}_{j+1} = \mathbf{r}_j - \alpha_j \mathbf{A} \mathbf{p}_j$
6. Compute $\beta_{ij} = \frac{\Re[(\mathbf{A} \mathbf{r}_{j+1})^\dagger \mathbf{A} \mathbf{p}_i]}{\|\mathbf{A} \mathbf{p}_i\|_2^2}$ for $i = 0, 1, \dots, j$
7. $\mathbf{p}_{j+1} = \mathbf{r}_{j+1} - \sum_{i=0}^j \beta_{ij} \mathbf{p}_i$
8. *End Do*

When implementing GCR for nearly symmetric systems, we can cheat on the orthogonalization somewhat (line 6) by only orthogonalizing the current search direction relative to the last s directions. This method, known as ORTHOMIN(s) [12], turns out to work well for our problem for $s = 15$, and improves both speed and convergence by limiting round-off errors.

When solving for the Newton step using ORTHOMIN(s), we are essentially sending a series of trial perturbations through our cavity, and using the information gleaned from the perturbed output to make an optimal guess as to the best direction in which to move towards a stationary point. Since we only ever need to know the action of the system (3.14) on a single trial vector, we can avoid having to compute the full cavity Jacobian by approximating all system matrix-vector products $\mathbf{A} \mathbf{p}$ in the ORTHOMIN(s) program by the forward difference

$$(\mathbf{B} - \mathbf{I})^{-1}(\mathbf{J}_f - \mathbf{I})\mathbf{p} \approx (\mathbf{B} - \mathbf{I})^{-1} \left(\frac{\mathbf{f}(\mathbf{u}) - \mathbf{f}(\mathbf{u} + d\mathbf{p}/\|\mathbf{p}\|_2)}{d} - \mathbf{p} \right), \quad (3.17)$$

with d an appropriate size: small enough to yield an accurate finite difference approximation of the derivative, and yet large enough to avoid significant round-off error.

3.5 Theoretical Convergence

If we are close enough to a solution, the outer loop should converge quadratically. This is not guaranteed, of course, even if the function is well behaved, as we are using an iterative solver. However, the number of outer steps will be small in any event, assuming convergence occurs. The greater concern is how many ORTHOMIN(15) steps will be required in each inner loop.

The convergence of optimal Krylov methods can be shown to be a function of the distribution of the eigenvalues. Roughly speaking, the convergence will be better the lower the condition number, and the more clustered the eigenvalues [19]. Given how close our system is to symmetric, the distribution of eigenvalues can be approximated by the singular values. As we've already shown in Fig. 3-2, the preconditioning greatly improves both the conditioning and spectral distribution of the system. Based on the fact that roughly 90% of the singular values are clustered very close to one, we'd expect that the subproblems should converge roughly an order of magnitude faster than would be required for a full matrix. For our model problem, this implies each inner loop should only require on the order of ten round trip evaluations. Assuming the Newton steps converge in less than ten iterations, we should thus expect to solve the cavity using on the order of 100 round trip evaluations.

Whether or not this is much of an improvement over naive simulation depends, of course, on the convergence of the natural system. As a rule of thumb, the more dispersion in the system, the slower it naturally converges. On the other hand, the more nonlinear the cavity, the faster it will converge. The utility of this algorithm then, hinges on the granularity of the cavity relative to the pulse shaping mechanisms. On one extreme, weakly nonlinear solid-state lasers can take 50,000 round trips to converge, with our algorithm only taking around 30 round trips. At the other end of the spectrum, fiber lasers are capable of significant pulse shaping over one pass, and they tend to converge within hundreds of round trips, offering little opportunity for improvement given the expected slow convergence of our inner loop problems in the presence of high nonlinearity.

Chapter 4

Results and Summary

4.1 Accuracy

Given that the convergence test of the algorithm is the same convergence test for the standard method (i.e. that the cavity reproduces the solution to within the some level of precision) and given that our method uses the same cavity model, the accuracy of our method is not really an issue and is limited by the cavity model. However, it is technically possible that our method could converge to a quasi-stable cavity mode that would not be found by a dynamic simulation. In Fig. 4-1 we compare the final solution found by our method with that found by dynamic simulation. As expected, they are roughly similar to within the convergence criteria of 10^{-8} .

4.2 Empirical convergence

As with all Newton solvers, the convergence of the outer loop is dependent on the starting point. As such, our solver is best used to refine an initial rough guess to high precision. However, the guess can be off by a significant amount (as illustrated in Fig. 4-3) and in some cases can converge starting from noise. Nonetheless, the better the guess the more reliable the convergence. However, this is not inconsistent with the two main intended uses of this algorithm: (a) to compute the results of many different cavity parameters, where the output of one simulation may be used as a seed to the next; and (b) in an optimization loop, where the parameters evaluated will be highly correlated with those from preceding computations. If a solution starting from nothing is desired, the best approach is




Figure 4-1: Left: comparison of the amplitude of the final solutions obtained by our method (dots) and standard dynamical evolution (solid line) for our test mode-locked laser model. Right: absolute value of difference between the two solutions.

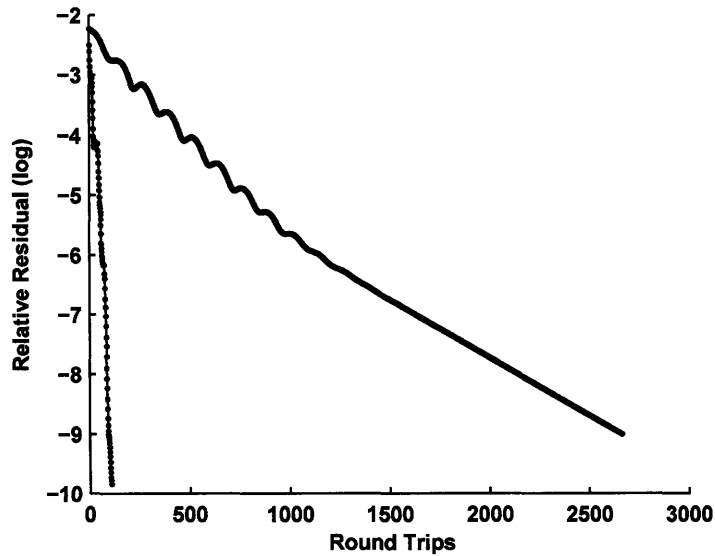


Figure 4-2: Comparison of convergence between our method (left curve) and standard dynamical evolution (right curve) for our test mode-locked laser model. All cavity evaluations are included, including those used to solve the linear Newton problems.

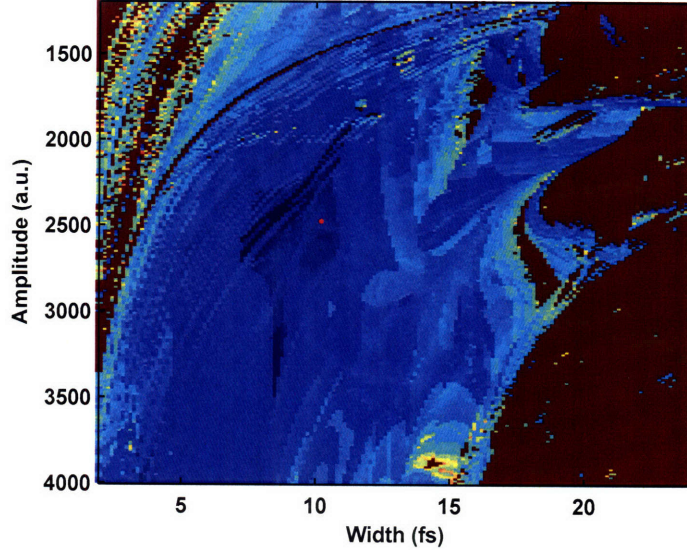


Figure 4-3: Convergence map of Newton-Krylov solver for Gaussian starting iterates of various widths and amplitudes for a ten femtosecond laser model (chosen for its quick convergence). The darker colors represent fewer steps, with the outer red region starting points that did not ever converge. The vertical axis is the amplitude and the horizontal the width of the starting Gaussian. The actual solution to the model is best approximated by the dot at roughly (2500, 10). This map suggests that the method will converge for a wide range of starting guesses, but that starting with energetic short pulses is a good strategy in the absence of any information about the true solution.

to run the dynamic simulation for a hundred round trips or so, and then let the algorithm take over once the pulse has evolved sufficiently.

This was the approach we took to test the algorithm's convergence relative to dynamic simulation. In Fig. 4-2, we compare the convergence of our algorithm to that obtained with standard simulation. Both the Newton-Krylov solver and the cavity dynamic simulation were seeded with a very rough starting pulse obtained from iterating the cavity 200 times on a starting impulse. Continuing with the dynamic simulation required over 2500 round trips to converge to within 10^{-9} . Our method was more than 34 times faster, and was able to converge to 10^{-10} while requiring only 76 cavity round-trip evaluations. The quadratic convergence of the outer Newton process is apparent in Fig. 4-2.

To provide a visualization of the kind of paths taken by the Newton-Krylov solver, we plot the exact sequence of trial pulses sent through the cavity solver in terms of the log of their residuals, Fig. 4-4. Note that quite early on the rough pulse shape is found and further refinement simply involves small perturbations and scaling. It thus makes sense

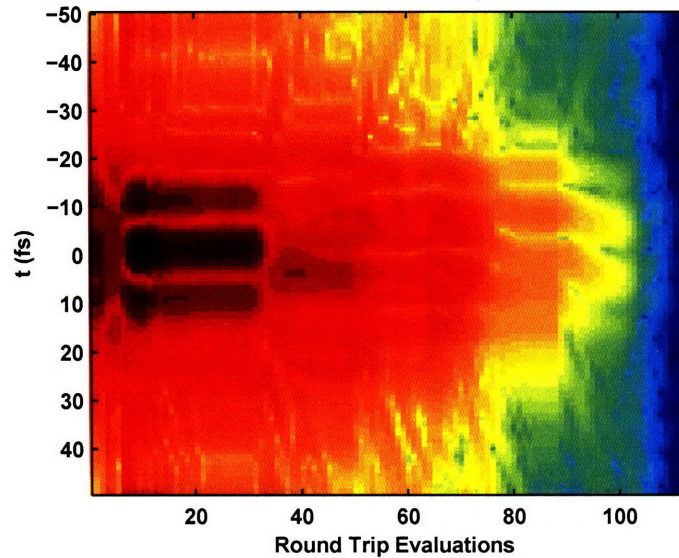


Figure 4-4: The sequence of residuals produced by the direct solver as a function of the number of round trip evaluations, log scaled.

that searching along paths in the direction of the residual would work well. The observed fact that the solution tends to overshoot during early steps suggests that convergence could be improved markedly were care taken in determining the Newton step length.

4.3 Future Work

The method appears to work quite well up to roughly 0.1π radians of nonlinear phase per round trip. This is good enough to deal with all but the most extreme solid-state lasers. There are several avenues of improvement that could be pursued to enhance the applicability and convergence of the algorithm, however. In addition, there is potential for this algorithm to be applied to more general cavity models incorporating spatial effects.

4.3.1 Better preconditioning for high SPM

The reason for the failure at high levels of nonlinearity is most likely the fact that we are using a diagonal preconditioner in the Fourier domain, which ignores all but a constant factor contribution from the nonlinearity. Given that the computational cost of the algorithm is almost completely dominated by the costly propagation through the gain medium, especially given the small vectors needed to represent our solutions, it might be worthwhile

to do a more involved preconditioning. We could easily compute several stripes of the banded matrix representing the lumped nonlinearity, and its inversion would likely be negligible compared to the cavity round trip evaluations.

4.3.2 Alternative linear solvers

So far we've only tried GCR and ORTHOMIN. Given that we experienced improvement by moving to ORTHOMIN it's likely we are, in fact, running into numerical stability problems. Despite our hope otherwise given the small number of iterations, the numerical stability issues suggest we may have chosen the wrong solver. It may be productive to investigate other Krylov solvers, such as ORTHODIR and GMRES.

4.3.3 Line search implementation

One admittedly egregious omission on my part was to not implement a proper line search algorithm for the Newton steps. It's possible that our realm of convergence could be greatly improved by improving the Newton steps. At the very least the result shown in Fig. 4-4 and discussed in Section 4.2 suggest that a line search could significantly improve the rate of convergence.

4.3.4 Phase normalization handling

Our solution to the problem of having no *a priori* knowledge of the cavity phase or group velocity (see the discussion in Section 3.1) was the addition of an ad hoc normalization function to our cavity round trip. Undoubtedly, this affects the convergence of our Krylov method to some extent as it significantly breaks the symmetry of the system. There may be better ways to handle this novel facet of our problem, such as solving for the ideal normalizing phases in a separate step, and then taken them as fixed during the solution of the Newton subproblems.

4.3.5 Reduced basis sets

As mentioned earlier, there is no reason why we must compute a round trip in the same basis that we represent our search. However, we have not yet taken advantage of the

potentially significant performance improvements via this route. In general, the representation that is optimal for accurately computing a round trip will be of significantly higher dimension than that needed to sufficiently represent the solution. Not only do we waste effort operating on wider temporal windows than we need (which end up filled with zeros) but we also hurt our convergence with the resulting poorly conditioned system.

4.3.6 Parallelization

The approach we've taken here has a potentially significant unrealized benefit over dynamic simulation: the ability to be parallelized. Dynamic simulation is inherently a serial process, whereas the solution to our linear subproblems can take advantage of parallel linear numerical techniques. At the extreme where we have at least as many processors as dimensions in our field, we can solve each linear problem in the time it takes to compute two round trips. This would allow another order of magnitude advantage over standard simulation.

4.3.7 Extension to full spatio-temporal model

An unfortunate truth is that while we can phenomenologically model Kerr lens mode-locked lasers sufficiently well to capture their salient features, we cannot simulate them well enough to use simulation as a primary design tool. Designers of these lasers thus currently have to include significant margins in the laser as built, and getting them to work is largely a matter of experience and trial-and-error. A significant aspect of this is our lack of quantitative understanding of Kerr lensing, driven by our lack of ability to effectively model it in a cavity.

The method presented in this thesis could be applied in a straightforward way to the simulation of a spatio-temporal cavity model. Preconditioning would work similarly, with the Jacobian for the linear cavity elements being diagonal and computable analytically if we operate in a basis composed of temporal and spatial frequency modes. Were an extension of this method to work with similar effectiveness on a spatial model, it might allow us to refine a quantitative model of Kerr lens mode locking by matching the output of the simulation with spatial measurements from actual lasers. Such an "optimization" would require a fast cavity solver.

The ability to fully model a nonlinear cavity would provide a significant benefit to the field of ultrafast optics, as it would allow for the precise engineering of mode-locked lasers without need for the trial-and-error tweaking involved today in the development of a laser. An even more auspicious goal would be the development of a model sufficiently accurate to act as a testbed for new laser development and research. Laser development is currently at the point where models do not adequately predict the operation of the shortest pulsed lasers. Providing a computational method capable of accurately modeling the full physics could revolutionize the development of mode-locked lasers, which currently require the construction of prototypes costing hundreds of thousands of dollars to build.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

- [1] G. Agrawal. *Nonlinear Fiber Optics*. Academic Press, 2001.
- [2] A. Arvanitoyeorgos. *An Introduction to Lie Groups and the Geometry of Homogeneous Spaces*. Americal Mathematical Society, 2003.
- [3] A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, V. Seyfried, M. Strehle, and G. Gerber. Control of chemical reactions by feedback-optimized phase-shaped femtosecond laser pulses. *Science*, 282:919–922, 1998.
- [4] J. R. Birge and F. X. Kärtner. A preconditioned newton-krylov method for computing stationary pulse solutions of mode-locked lasers. In *CLEO Proceedings*, page CTuC7, Baltimore, 2007. OSA.
- [5] R. W. Boyd. *Nonlinear Optics*. Academic Press, San Diego, 2nd edition, 2003.
- [6] J. W. Chan, T. Huser, S. Risbud, and D. M. Krol. Structural changes in fused silica after exposure to focused femtosecond laser pulses. *Opt. Lett.*, 26:1726–1728, 2001.
- [7] Y. Chen, F. X. Kärtner, U. Morgner, S. H. Cho, H. A. Haus, H. G. Fujimoto, and E. P. Ippen. Dispersion managed mode locking. *JOSA B*, 16(11):1999–2004, 1999.
- [8] M. Dantus, M. J. Rosker, and A. H. Zewail. Real-time femtosecond probing of “transition states” in chemical reactions. *J. Chem. Phys.*, 87(4):2395–2397, 1987.
- [9] A. J. DeMaria, D. A. Stetser, and H. Heynau. Self mode-locking of lasers with saturable absorbers. *Applied Physics Letters*, 8(7):174–176, 1966.
- [10] S. C. Eisenstat, H. C. Elman, and M. H Schultz. Variational iterative methods for nonsymmetric systems of linear equations. *SIAM Sci. Stat. Comput.*, 20(2):345–357, 1983.
- [11] A. Gordon and B. Fischer. Phase transition theory of many-mode ordering and pulse formation in lasers. *Phys. Rev. Lett.*, 89(10):103901, 2002.
- [12] K.C. Jea and D. M. Young. Generalized conjugate gradient acceleration of nonsymmetrizable iterative methods. *Lin. Alg. Appl.*, 34:159–194, 1980.
- [13] F. X. Kärtner, editor. *Few-Cycle Laser Pulse Generation and its Applications*. Springer, 2004.
- [14] T. H. Maimon. Stimulated optical radiation in ruby. *Nature*, 187:493–494, 1960.

- [15] U. Morgner, F. X. Kärtner, S. H. Cho, Y. Chen, H. A. Haus, J. G. Fujimoto, and E. P. Ippen. Sub-two-cycle pulses from a kerr-lens mode-locked ti:sapphire laser. *Opt. Lett.*, 24:411–413, 1999.
- [16] O. Nastov, R. Telichevesky, K. Kundert, and J. White. Fundamentals of fast simulation algorithms for rf circuits. *Proc. IEEE*, 95(3):600–620, March 2007.
- [17] J. Ouellette. Femtosecond lasers prepare to break out of the laboratory. *Physics Today*, page 36, January 2008.
- [18] P. M. Paul, E. S. Toma, P. Breger, G. Mullot, F. Augé, Ph. Balcou, H. G. Muller, and P. Agostini. Observation of a train of attosecond pulse from high harmonic generation. *Science*, 292:1689–1692, 2001.
- [19] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, 2nd edition, 2003.
- [20] Y. Saad and M. H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986.
- [21] G. Sansone, E. Benedetti, F. Calegari, L. Avaldi, R. Flammini, L. Poletto, P. Villoresi, C. Altucci, R. Velotta, S. Stagira, S. De Silvestri, and M. Nisoli. Isolated single-cycle attosecond pulses. *Science*, 314:443–446, 2006.
- [22] R. Telichevesky, K. Kundert, and J. White. Efficient steady-state analysis based on matrix-free krylov-subspace methods. In *Proc. ACM IEEE Des. Auto. Conf.*, pages 480–484, Santa Clara, June 1995.
- [23] T. Udem, R. Holzwarth, and T. W. Hänsch. Optical frequency metrology. *Nature*, 416:233–237, 2002.

Appendix A

Code Appendix

A.1 Cavity round trip function

```
function [Uout, Jdiag, phi0out, phi1out] = ...
    solitoncavitystep(f, Uin, cavparams, norm)
% solitoncavitystep Propagate soliton-like laser cavity for m steps.
% Assumes length(ts) is an even number (best if a power of 2).
% Optionally returns the diagonal of the Jacobian in Jdiag, for use in
% preconditioning solvers.
%
% This is a fast version of the solitoncavityfd function that is meant to
% be called from other functions, notable cavnewtonfd. It uses column
% vectors and only computes one step.
%
% Cavity parameters. %
TR = cavparams.TR; % cavity roundtrip time (fs from MHz)
D2net = cavparams.D2net; % GDD of whole cavity
D3net = cavparams.D3net; % TODO Fix this
g0 = cavparams.g0; % peak small signal gain for one RT
Wg = cavparams.Wg; % gain bandwidth (PHz)
Psat = cavparams.Psat; % gain saturation power (W)
q = cavparams.q; % saturable absorber gain per RT
Isat = cavparams.Isat; % saturation intensity of absorber (W)
oc = sqrt(cavparams.l); % output coupler gain spectrum
if ~isfield(cavparams, 'nstep')
    nstep = 32;
else
    nstep = cavparams.nstep;
end
if nargin > 3
    if strcmp(norm, 'none')
        normed = false;
    else
        normed = true;
    end
else
    normed = true;
end
```

```

end

% Memory allocation and precalculations. %
n = length(f);
w = 2*pi*f;
dt = 1/f(2)/n; % f(2) is df
A = dt/TR; % average power integral scaling
D2 = D2net - cavparams.D2nl; % bulk dispersion ex crystal
D3 = D3net - cavparams.D3nl;
phibulk = -j*(D2*w.^2/2 + D3*w.^3/6); % bulk dispersion phase
H4 = exp(phibulk/4); % transfer function of FT domain parts
H2 = exp(phibulk/2); % transfer function of FT domain parts

% Calculation. %

% DCM
U = H4.*Uin;

% Nonlinear gain material with SA
u = ifft(U);
P = (u'*u)*A;
nlparams.gsp = g0*lorentznorm(f, Wg)/(1 + P/Psat)/4; % ss gain
nlparams.d = cavparams.d/4; % SPM coefficient of nl material (1/W)
nlparams.D2 = cavparams.D2nl/4; % GDD of nl material (fs^2)
nlparams.D3 = cavparams.D3nl/4;
U = nlprop(f, fft(u), nlparams, floor(nstep/2));
u = ifft(U);
Iu = real(u .* conj(u)); % intensity (W)
u = exp(q./(1 + Iu/Isat)/2) .* u;
U = nlprop(f, fft(u), nlparams, floor(nstep/2));

% DCM, two passes
U = H2.*U;

% Nonlinear gain material with SA
U = nlprop(f, U, nlparams, floor(nstep/2));
u = ifft(U);
Iu = real(u .* conj(u)); % intensity (W)
u = exp(q./(1 + Iu/Isat)/2) .* u;
U = nlprop(f, fft(u), nlparams, floor(nstep/2));

% DCM
U = H4.*U;

% OC
U = oc.*U;

% Phase normalization and output.
if normed
    Uout = phasenorm(f, Uin, U);
else
    Uout = U;
end

```

```

% Optional output.
if nargout > 1
    phinet = -j*(D3net*w.^3 + D2net*w.^2);
    Jdiag = oc.*exp(phinet + 4*nlparams.gsp + q./(1 + Iu/Isat));
end
if nargout > 2 && normed
    phi0out = phi0 - phi0in;
    phi1out = phi1 - phi1in;
end

```

A.1.1 Nonlinear propagation

```

function Uout = nlprop(f, Uin, nlparams, m)
% nlpropfd Propagate NLSE through material broken into m steps.

% Cavity parameters. %
D2 = nlparams.D2; % GDD (fs^2)
D3 = nlparams.D3;
d = nlparams.d; % SPM coefficient (1/W)
gsp = nlparams.gsp; % total small signal gain

% Memory allocation and precalculations. %
w = 2*pi*f;
phi = -j*(D2*w.^2/2 + D3*w.^3/6);
H = exp((gsp + phi)/m);
H2 = exp((gsp + phi)/m/2);

% Calculate m steps. %
U = H2.*Uin;
for k = 1:m,
    % Time domain:
    u = ifft(U);
    Iu = real(u .* conj(u)); % intensity (W)
    u = exp(-j*d/m*Iu) .* u;
    U = fft(u);

    % Frequency domain:
    if k == m
        Uout = H2.*U;
    else
        U = H.*U;
    end
end
end

```

A.1.2 Phase normalization

```

function Uout = phasenorm(f, Uin, U)
% Minimize the change in phase between Uin and Uout by adding trivial
% phases to U.

phi0in = angle(Uin(1));
phi1in = (angle(Uin(2)) - angle(Uin(end)))/2/f(2);

```

```

phi0 = angle(U(1));
phi1 = (angle(U(2)) - angle(U(end)))/2/f(2); % 2nd order FD
dphi = (phi1-phi1in)*f + (phi0-phi0in);
Uout = U.*exp(-j*dphi);

```

10

A.2 Solver

```

function [ufinal, converged, Fnorms, us] = ...
    cavitysolver(ts, u0, cav, tol, maxiter, ngcr)

% Initialization and parameters.
a = 1e1; % perturbation used in finite difference
s = 15;
diags = true;
n = length(ts);
if nargin < 6
    ngcr = 100; % max iterations
    autogcr = true;
else
    autogcr = false;
end
if nargin < 5
    maxiter = 10;
elseif nargin < 6
    ngcr = 32; % max GCR iterations
    autogcr = true;
else
    autogcr = false;
end
if nargout > 2
    Fnorms = zeros(1,2);
end
if nargout > 3
    us = zeros(n,3);
    us(:,1) = u0;
end
dt = ts(2) - ts(1); % sample period (fs)
tp = n*dt; % window time (fs)
f = [(0:n/2-1), -(n/2:-1:1)].'/tp; % FFT frequencies (PHz)
k = 0;
fn = 0;
p = zeros(n, ngcr);
Jp = zeros(n, ngcr);

% *** Newton Iterations. ***
U = fft(u0(:));
Ucav = solitoncavitystep(f, U, cav); fn = fn + 1;
Jdiag = (cavityprecond(f, U, cav) + cavityprecond(f, Ucav, cav))/2;
Binv = diag(1./(1 - Jdiag));
BinvF = Binv*U - Binv*Ucav; % what we want to set to zero
F = U - Ucav;
Fnorm = norm(F)/norm(U);

```

10

20

30

40

```

if nargout > 2
    Fnorms(end+1) = Fnorm;
end
kiter = 0;
converg = false;
done = false;

% Diagnostics.
if diags
    nl = 45;
    fprintf([repmat('-', 1, nl) '\n'])
    fprintf('iter\tevals\tdU| (log)\t\t|F| (log)\n')
    fprintf([repmat('-', 1, nl) '\n'])
    fprintf('%d\t\t%d\t\t%f\t\t%f\n', k, fn-1, 0, log10(Fnorm))
end

while ~done % *** Newton ***
    m = 0;

    r = -BinVF; % null starting vector
    r0 = r;
    dU = zeros(n,1);
    gcrdone = false;
    while (m < ngcr) && ~gcrdone, % *** GCR ***
        m = m + 1;
        kiter = kiter + 1;

        p(:,m) = r; % use residual as search direction

        % Compute Approximate BinV*J*p using finite differences.
        d = a/norm(p(:,m));
        dUp = d*p(:,m);
        Ucavdp = solitoncavitystep(f, U + dUp, cav); fn = fn + 1;
        Jp(:,m) = BinV*(p(:,m) + (Ucav - Ucavdp)/d);

        % Make the new Jp vector orthogonal to the s most recent Jp vectors. (Try doing exactly as shown in Saad,
        for j = max(1,m-1-s):m-1,
            beta = real(Jp(:,m)' * Jp(:,j)); % projection
            p(:,m) = p(:,m) - beta*p(:,j); % subtract out orthogonal parts
            Jp(:,m) = Jp(:,m) - beta*Jp(:,j); % "
        end

        % Make the orthogonal Jp vector of unit length.
        Jpnorm = norm(Jp(:,m),2);
        Jp(:,m) = Jp(:,m)/Jpnorm;
        p(:,m) = p(:,m)/Jpnorm;

        % Determine the optimal amount to change x in the p direction
        % by projecting r onto Mp
        alpha = real(r' * Jp(:,m));

        % Update x and r
        dU = dU + alpha*p(:,m);
        oldrnorm = norm(r);

```

```

r = r - alpha*Jp(:,m);

% Automatically terminate.
if autogcr
    % relr = norm(r)/norm(r0);
    % if relr < 0.5e-1 % 0.2e-1
    % gcrdone = true;
    % end
    if (oldrnorm - norm(r))/oldrnorm < 0.5e-2
        gcrdone = true;
    end
end
110

if nargout > 2
    Utest = solitoncavitystep(f, U + dU, cav);
    Fnorms(kiter) = norm(U + dU - Utest)/norm(U);
end
if nargout > 3
    us(:,kiter+1) = ifft(U + dU);
end
120
end % *** GCR ***

U = U + dU;
dUnorm = norm(dU);

k = k + 1;

Ucav = solitoncavitystep(f, U, cav); fn = fn + 1;
Jdiag = (cavityprecond(f, U, cav) + cavityprecond(f, Ucav, cav))/2;
Binv = diag(1./(1 - Jdiag));
130

BinvF = Binv*U - Binv*Ucav;
F = (U - Ucav);
Fnorm = norm(F)/norm(U);

converg = Fnorm < tol;
done = converg || (k > maxiter);

if nargout > 2
    Fnorms(kiter+1) = Fnorm;
end
140

if diags
    fprintf('%d\t\t%d\t\t%f\t\t%f\n', k, fn-1, log10(dUnorm), log10(Fnorm))
end
end % *** Newton ***
fprintf([repmat('=', 1, nl) '\n'])

if nargout > 1
    if converg
        converged = fn;
    else
        converged = 0;
    end
end
150

```

end

ufinal = ifft(U);

A.3 Preconditioner

```
function d = cavityprecond(f, U, cav)
% CAVITYPRECOND Compute diagonal preconditioner
% Compute diagonal preconditioner for the cavity given by cav, around the
% point given by the Fourier component vector U, at the frequencies given
% in f.
%
% It would be slightly more efficient to compute this in
% solitoncavitystep, but is put here for the sake of

n = length(U);
dt = 1/f(2)/n;
w = 2*pi*f;
u = ifft(U);
oc = sqrt(cav.l);
P = (u' * u) * dt / cav.TR;
lu = real(u .* conj(u));
dgsp = cav.g0 * lorentznorm(f, cav.Wg) / (1 + P / cav.Psat);
phinet = -j * (cav.D2net * w.^2 / 2 + cav.D3net * w.^3 / 6);
d = oc * exp(mean(j * cav.d * lu + cav.q ./ (1 + lu / cav.lsat)) + phinet + dgsp);
```

10