Title

Computing the Weak Limit of KdV

Authors

McLaughlin, D W
Strain, J A

Publication Date

1992-07-01

LBL-32990
UC-405
Preprint

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Physics Division

Mathematics Department

To be submitted for publication

Computing the Weak Limit of KdV

D.W. McLaughlin and J.A. Strain

July 1992

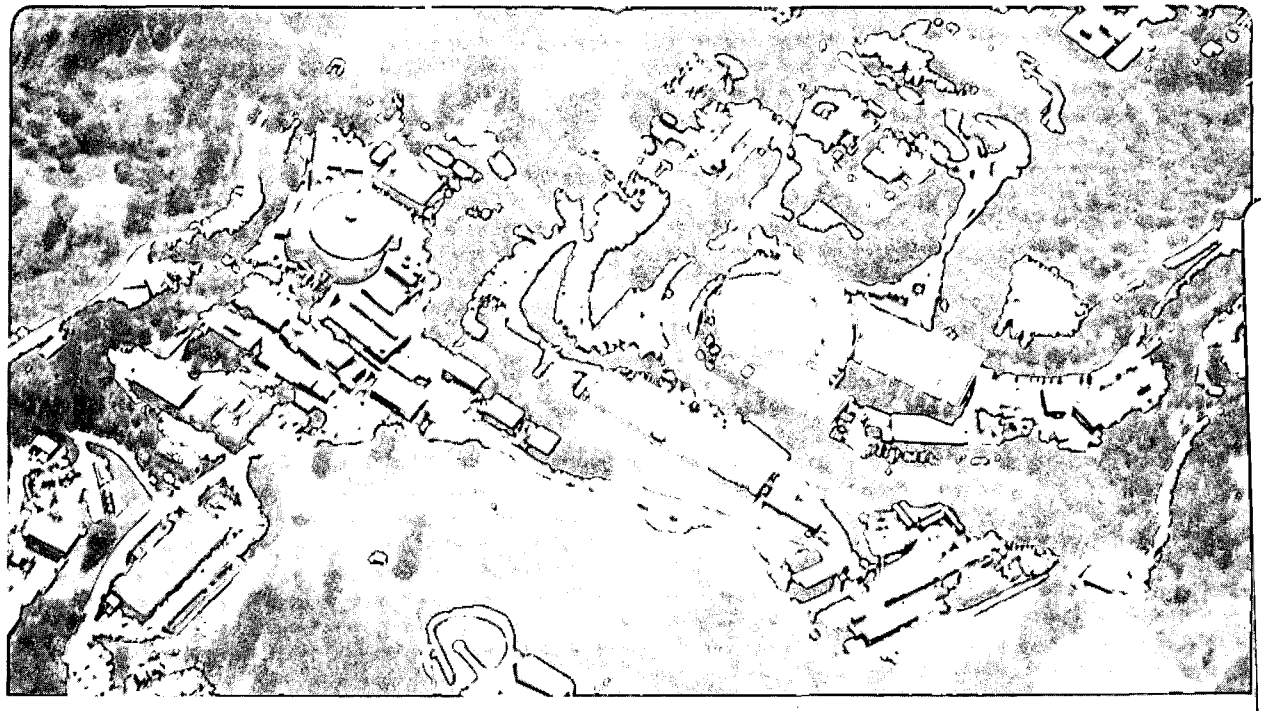

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

LBL-32990

COMPUTING THE WEAK LIMIT OF KdV

David W. McLaughlin[1]
Department of Mathematics
and
Program in Applied and Computational Mathematics
Princeton University
Princeton, NJ 08544

and

John A. Strain[2]
Department of Mathematics
and
Program in Applied and Computational Mathematics
Princeton University
Princeton, NJ 08544
and
Department of Mathematics and Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

July 1992

[1] Funded in part by AFOSR-90-0161 and NSF DMS 8922717 A01.

[2] Research supported in part by a NSF Mathematical Sciences Postdoctoral Research Fellowship and in part by the Applied Mathematical Sciences Subprogram of the Office of Energy Research, U.S. Department of Energy under Contract DE-AC03-76SF00098 while the author was visiting Lawrence Berkeley Laboratory.

Abstract

The solution of the KdV equation with single-minimum initial data has a zero-dispersion limit characterized by Lax and Levermore as the solution of an infinite-dimensional constrained quadratic minimization problem. An adaptive numerical method for computing the weak limit from this characterization is constructed and validated. The method is then used to study the weak limit. Initial simple experiments confirm theoretical predictions, while experiments with more complicated data display multiphase behavior considerably beyond the scope of current theoretical analyses. The method computes accurate weak limits with multiphase structures sufficiently complex to provide useful test cases for the calibration of numerical averaging algorithms.

Contents

List of Figures

List of Tables

1 Introduction

The Korteweg-deVries equation

$$u_t - 6uu_x + \epsilon^2 u_{xxx} = 0 \qquad x \in \mathbf{R} \tag{1.1}$$

$$u(x,0) = v(x) \tag{1.2}$$

appears as an asymptotic model in physical systems ranging from water waves to plasma physics. This ubiquity is partly explained by the fact that it is the simplest model for dispersive regularization of nonlinear hyperbolic conservation laws. In the absence of the regularizing term ($\epsilon = 0$), equation (1.1) reduces to a scalar conservation law, whose solutions steepen and develop infinite derivatives in finite time. Regularizing terms, such as $\epsilon^2 u_{xxx}$, prevent the development of such singularities. The most common regularization is Burgers' equation

$$u_t - 6uu_x - \epsilon u_{xx} = 0, \tag{1.3}$$

in which dissipation is modeled by the ϵu_{xx} term. This term smooths the shock into a transition front of spatial width $O(\epsilon)$. In the limit as $\epsilon \to 0$, this front converges strongly to the discontinuous shock profile which itself is a weak solution of the conservation law.

When the physical situation is dispersive rather than dissipative, the simplest regularization is the KdV equation (1.1), with its third order term $\epsilon^2 u_{xxx}$. This dispersive term also prevents the development of singularities, but in an entirely different manner than the dissipation in Burgers' equation. In the KdV situation, the steepening profile is smoothed by the onset of short wavelength [$O(\epsilon^1)$] finite amplitude [$O(\epsilon^0)$] oscillations [GP74, FW78]. The front becomes a rapidly oscillating wavetrain, which converges weakly but not strongly to a limit $\bar{u}$. This limit is not a weak solution of the scalar conservation law; rather, $\bar{u}$ and the local wavenumbers and frequencies of the oscillating wavetrain solve a coupled system of several hyperbolic conservation laws.

To contrast these different regularizations, we have used a numerical method described in Section 3.9 to solve both the Burgers and KdV equations at moderately small values of ϵ. Figure 1 compares the short wavelength, finite amplitude oscillations of the KdV case to the smooth profile of the Burgers solution. Figure 1 also exhibits the weak limit $\bar{u}$ of KdV as $\epsilon \to 0$, as computed by the methods of this paper.

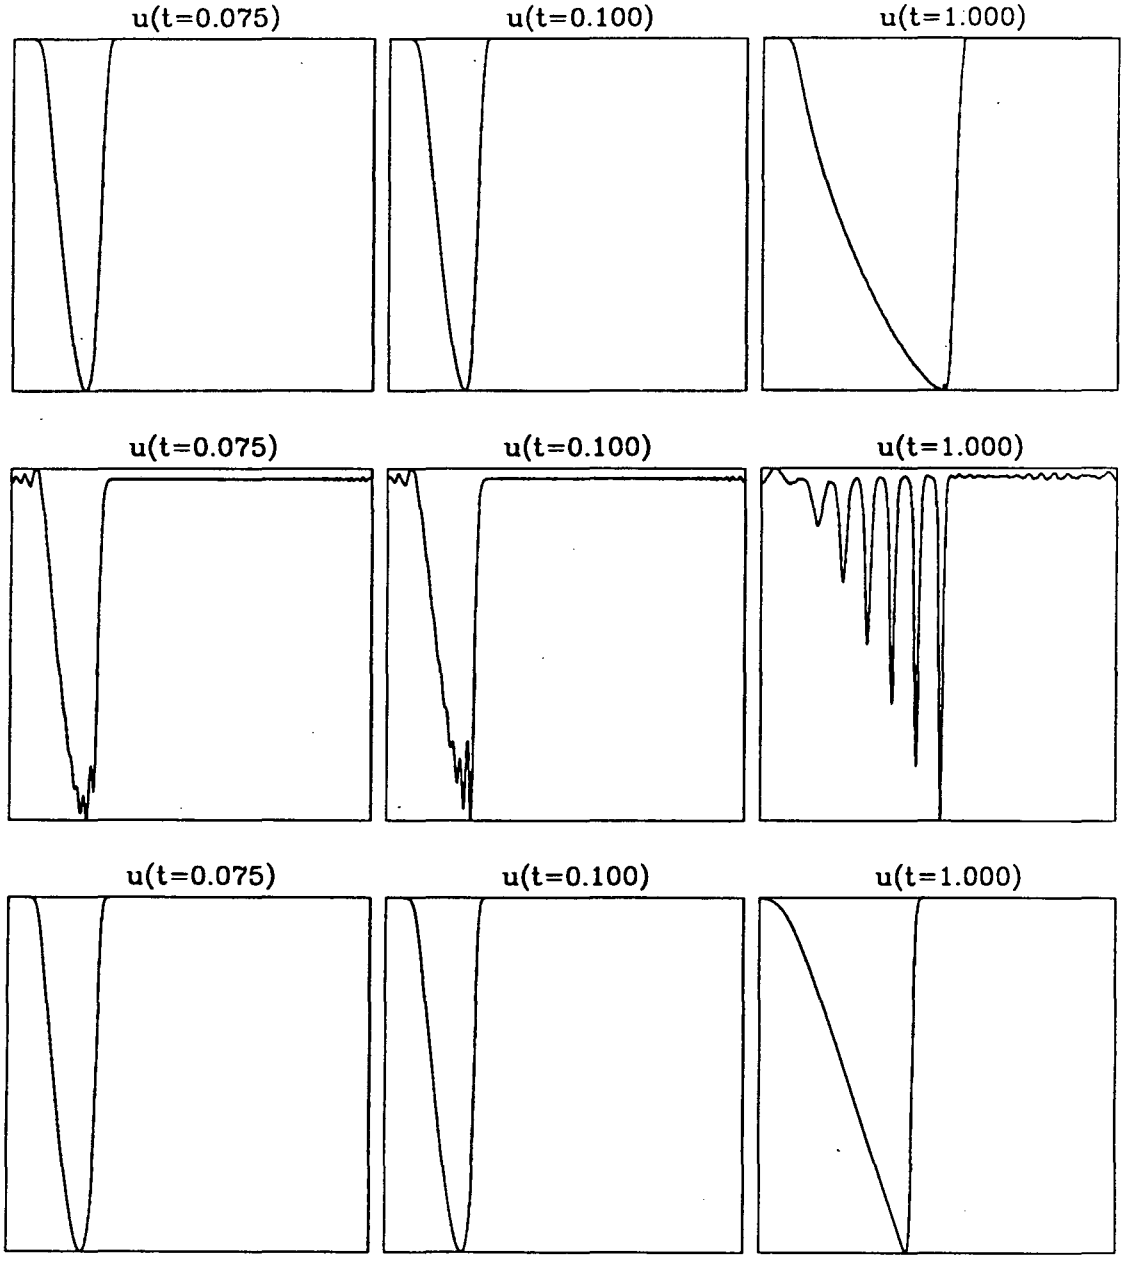

Figure 1: Weak limit, KdV solution and Burgers profile with $\epsilon = 0.08$. The interval is $[-2, 10]$ and each plot is scaled to fit the solution. The breaktime for this solution is $t_b = 1/12 = 0.8666....$

Oscillations make KdV difficult to solve numerically when ϵ is small, and computing the weak limit $\bar{u}$ via numerical averaging seems even more difficult.

This paper is devoted to the numerical study of the structure of $\bar{u}$. Our basic tool is the characterization of $\bar{u}$ due to Lax and Levermore [LL83]. They used the inverse scattering solution of KdV due to Gardner, Greene, Kruskal and Miura [GGKM67] to solve (1.1) almost exactly, then took the limit $\epsilon \to 0$ analytically. This produced an infinite-dimensional constrained quadratic minimization problem, for each (x,t), for a function $f(k;x,t)$ which mediates between $\bar{u}(x,t)$ and the spectral parameter $k \in [0,1]$ of the Schrödinger operator with potential v. The coefficients and constraints of the minimization problem are computable functions of v; this data fixes the minimization problem which has a unique solution $f^*(k;x,t)$; the weak limit $\bar{u}(x,t)$ is then calculated from $f^*(k;x,t)$. These results are summarized in Section 2.

This suggests a natural numerical method for computing $\bar{u}$ at any point (x,t) in spacetime: simply solve the minimization problem numerically. This method is attractively direct and parallel. It offers insight into the folding structure and spectral information associated with $\bar{u}$, information which would be difficult if not impossible to obtain by computing u and letting $\epsilon \to 0$. It eliminates the oscillations analytically, rather than averaging over them numerically.

We have designed and implemented such a method, using a finite element Rayleigh-Ritz method with an adaptive grid. Singular quadrature methods, fast adaptive tabulation, and numerical smoothing and differentiation also play roles in our method. It is described in Section 3, where we also describe how we construct initial data v and a simple method for computing solutions of (1.1) and (1.3) with $\epsilon > 0$.

In Section 4, we validate the numerical method by presenting numerical results which demonstrate its accuracy and efficiency. It turns out to be reasonably and controllably accurate, and surprisingly efficient for large-scale computation. It would also be an extremely natural candidate for parallel computation.

In Section 5, we present large-scale calculations which probe the structure of the weak limit. First, with simple cases, our experiments confirm analytical predictions of the Lax-Levermore theory. Turning to data with several inflection points, we then compute multiphase behavior beyond the scope of current theoretical analyses. In particular, several phases separated by sharp spacetime boundaries develop and interact in interesting fashion.

In these cases, the weak limit has visible macroscopic structure associated with the multiphase regions. In these regions, our simulations of KdV with small dispersion show wavetrains with oscillatory microstructure, presenting a challenge to numerical averaging methods for dispersive limit problems.

The paper is arranged so that the non-numerical reader need not follow Sections 3 and 4 in order to understand Sections 5 and 6, while the reader familiar with the Lax-Levermore method need study Section 2 only for the notation.

2 The Lax-Levermore Method

We now summarize the results of [LL83], in notation convenient for our purposes. A recent survey is provided in [LLV92]. We begin by stating the minimization problem which forms the main analytical tool of this paper, then discuss its derivation and meaning.

2.1 The minimization problem

We consider the KdV equation (1.1) with initial condition (1.2), under the following assumptions on v;

(a) v is smooth,

(b) $v(x) \leq 0$,

(c) v has a single strict minimum at $x = 0$, where $v(0) = -1$ and $v'' > 0$,

(d) v has compact support.

This last requirement eliminates an additional numerical error due to truncating the support of v without significantly decreasing generality.

Under these assumptions, Lax and Levermore constructed v_ϵ converging strongly to v in $L^2(\mathbf{R})$ such that the solution u_ϵ of (1.1) with initial data v_ϵ has a weak limit $\bar{u}$. Furthermore, $\bar{u}$ is computable from

$$\bar{u}(x,t) = \partial_x^2 Q^*(x,t), \tag{2.1}$$

where $Q^*(x,t)$ is the minimum value of the quadratic functional

$$\begin{aligned} Q(f;x,t) &= (a,f) + \frac{1}{2}(f,Kf) \qquad (2.2)\\ &= \int_0^1 a(k;x,t)f(k)dk + \frac{1}{2}\int_0^1\int_0^1 K(k,k')f(k)f(k')dkdk', \end{aligned}$$

taken over integrable f subject to the constraints

$$0 \leq f(k) \leq \varphi(k) \qquad 0 \leq k \leq 1. \tag{2.3}$$

Here K, φ and a are given by

$$K(k,k') = \frac{-1}{\pi^2} \log\left(\frac{k-k'}{k+k'}\right)^2$$

$$\varphi(k) = \int_{x_-(k)}^{x_+(k)} \frac{k}{\sqrt{-v(x)-k^2}}\, dx \tag{2.4}$$

$$a(k;x,t) = \frac{4}{\pi}\left(xk - 4k^3t - \vartheta(k)\right) \tag{2.5}$$

$$\vartheta(k) = kx_+(k) + \int_{x_+(k)}^{\infty} k - \sqrt{k^2+v(x)}dx \tag{2.6}$$

and the "turning" points $x_- \leq x_+$ are the solutions of $v(x) = -k^2$, defined on the interval $0 \leq k \leq 1$. Note that Q is linear in x and t. This infinite-dimensional minimization problem has a unique solution f^*, because K defines a compact positive definite operator and thus Q is convex. The minimizer f^* satisfies the Euler-Lagrange conditions corresponding to Q.

2.2 Inverse scattering and WKB

This characterization of $\bar{u}$ is derived by using the exact solution of KdV by inverse scattering [GGKM67], which relates the solution u of KdV to the spectrum of the Schrödinger equation

$$-\epsilon^2\partial_x^2\psi + u\psi = E\psi \qquad x \in \mathbf{R}. \tag{2.7}$$

The zero-dispersion limit of KdV becomes the semiclassical limit $\epsilon \to 0$ of the Schrödinger equation, in which the WKB theory [Mer70] applies. In the WKB limit, the negativity of v makes the reflection coefficient exponentially small. Thus Lax and Levermore construct reflectionless potentials $v_\epsilon \to v$ in L^2, for which the Kay-Moses explicit solution of the inverse scattering problem [KM56] applies, and the continuous spectrum can be ignored. Then u_ϵ becomes a large sum, analogous to a partition function in statistical mechanics, and the leading behavior of the sum as $\epsilon \to 0$ is dominated by its largest term, via a kind of Laplace's method in function space. This largest term leads to the minimization problem.

Lax and Levermore go on to transform the minimization problem into a Riemann-Hilbert problem, under further assumptions on the solution, and

solve the Riemann-Hilbert problem in the long-time limit. Thus they find, for example, that

$$\bar{u} \sim \frac{-1}{2\pi t}\sqrt{\varphi(x/4t)} \qquad \text{as} \quad t\to\infty \tag{2.8}$$

when $0 < x/t < 4(\min v(x))^2$. This fixes the speed of the leading edge of the oscillatory region at $4(\min v(x))^2$. We do not make use of the Riemann-Hilbert problem, nor any further assumptions on the solution.

2.3 Structure of the minimizer

The structure of the minimizer is interesting: Because of the constraint $0 \le f \le \varphi$, the interval $[0,1]$ can be divided into two sets: the "active set" A where $f = 0$ or $f = \varphi$ (a constraint is "active") and the inactive set I where $0 < f < \varphi$. On the inactive set, only Q matters and f satisfies the Euler-Lagrange equation for Q. On the active set, the gradient of Q points out of the admissible region $0 \le f \le \varphi$, and f touches a constraint. The active set corresponds to a "gap" in the spectrum of the Schrödinger equation (2.7), or to a "phase" in the Whitham or Flaschka-Forrest-McLaughlin(FFM) theories [Whi74, FFM80].

The structure of the minimizer changes dramatically at the "breaktime"

$$t_b = \frac{1}{6\max_x v'(x)} \tag{2.9}$$

where the classical solution of the dispersionless equation

$$u_t - 6uu_x = 0 \tag{2.10}$$

becomes multivalued ("breaks"). Before the breaktime, the weak limit is actually strong and the active set is a single k-interval $[\sqrt{-\bar{u}(x,t)}, 1]$. This corresponds to a spectrum $[-\sqrt{-\bar{u}(x,t)}, +\infty]$ in the energy variable $E = -k^2$. After t_b, oscillations in u prevent strong convergence and the active set splits into two or more intervals. In the FFM theory, the gaps in the spectrum are occupied by "phases" which describe the oscillations in u. The number of dynamical variables in the FFM theory thus increases at each breaking. The number of gaps is the *genus* of the problem. Weak limits are no longer solutions of (2.10) when the genus is nonzero; Lax and Levermore show that as $t\to\infty$, the genus becomes unity over large areas of space, so single-phase behavior dominates at large times.

2.4 Generalizations

Finally, we remark that the Lax-Levermore theory can be extended to more general initial data v. The restriction of nonpositivity has been removed by Venakides [Ven85], while Lax and Levermore sketched how to handle v with several minima. Instead of a single function f, we get a function f_j and a constraint φ_j for each potential well in v. Again a minimization problem can be found, but now Q is a functional of $f_1, f_2 \ldots, f_n$:

$$Q(f_1 \ldots, f_n; x, t) = \sum_{j=1}^{n} (a_j, f_j) + \frac{1}{2}(f, Kf)$$

with

$$f = \sum_{j=1}^{n} f_j$$

and the constraints are $0 \leq f_j \leq \varphi_j$ for $j = 1, 2, \ldots, n$. Here φ_j is the φ function corresponding to the jth potential well. This seems at first glance greatly to increase the complexity of the problem. However, this functional is linear in all but one variable, so if we introduce f as a new variable then we can compute the minimizing $(f_1, \ldots, f_n)$ by solving $n - 1$ very simple constrained linear minimization problems, and then only one truly quadratic problem. Numerically, this is a straightforward generalization of the numerical method presented in this paper.

3 The Numerical Method

We now develop a numerical method based on Lax-Levermore theory; thus we solve the minimization problem to calculate Q^* and differentiate numerically to obtain $\bar{u}$. This approach to $\bar{u}$ avoids the difficulty of calculating u and letting $\epsilon \to 0$; the oscillations have already been eliminated analytically. However, many other technical complications arise; singularities, singular integrals, numerical differentiation and efficiency are major concerns. We describe how to overcome each difficulty in turn. We also discuss the selection of initial data v and the solution of KdV and Burgers with $\epsilon > 0$.

3.1 A useful example

First, we try to develop some intuition by writing down a function

$$v(x) = \min(0, x^2 - 1)$$

for which the coefficients and constraints can be evaluated analytically. A routine calculation shows that

$$\varphi(k) = \pi k \tag{3.1}$$

and

$$\vartheta(k) = \frac{1}{2}\left(k + (1-k^2)\log\frac{1+k}{\sqrt{1-k^2}}\right), \tag{3.2}$$

with $\vartheta(1) = 1/2$. Note that $\varphi(0) = 0$; thus we must seek our minimum among functions f with $f(0) = 0$. The logarithmic singularity in ϑ for this example is typical. However, this φ is untypically smooth; later examples will involve much more singular φ's. At $t = 0$, the exact minimizer f can be computed analytically;

$$f^*(k; x, 0) = \begin{cases} \pi k & x \le -\sqrt{1-k^2} \\ k\left(\frac{\pi}{2} - \sin^{-1}\left(\frac{x}{\sqrt{1-k^2}}\right)\right) & |x| \le \sqrt{1-k^2} \\ 0 & x \ge \sqrt{1-k^2}. \end{cases} \tag{3.3}$$

Note that f^* has a square-root singularity at $x = \pm\sqrt{1-k^2}$. This singular behavior occurs whenever k enters or leaves the active set, which in this example is $A = \{k \in [0,1] \mid k > \sqrt{1-x^2}\}$ for $|x| \le 1$ and $A = [0,1]$ for

$|x| > 1$. Thus there can be several moving singularities in f^*, occurring at unpredictable locations in $[0,1]$. This difficulty must be handled with care in designing numerical methods for minimizing Q.

The weak limit $\bar{u}$ itself is also singular; it develops a jump discontinuity at the breaktime.

3.2 Discretization

Now we consider our central task, the minimization of Q over the set S of functions $f \in L^1(0,1)$ with $0 \le f \le \varphi$ and $f(0) = 0$. This problem is infinite-dimensional, so we approximate it by a finite-dimensional problem using the Rayleigh-Ritz method; we choose an N-dimensional subset $S_N \to S$ as $N \to \infty$, and approximate f^* by the minimizer f^N of Q over S_N. This approach is particularly effective because we want Q^* as well as f^*; since f^N minimizes Q exactly over a subspace, the error $Q(f^*) - Q(f^N)$ is proportional to $(f^* - f^N)^2$, which is much smaller than $f^* - f^N$ when f^N is close to f^*. The choice of S_N is made on a balance of computational convenience and good approximation properties. Let

$$G: \quad 0 = k_0 < k_1 < \ldots < k_N = 1$$

be a grid and let S_N be the space of functions f which are linear between grid points and satisfy $0 \le f_j \le \varphi_j$ and $f_0 = 0$. Here $f_j = f(k_j)$ and $\varphi_j = \varphi(k_j)$. (Note that $f(0) = 0$ since $\varphi(0) = 0$.) This choice is computationally convenient because (f, Kf) can be evaluated exactly for $f \in S_N$, and it has excellent approximation properties if the grid is chosen with proper regard to the singularities of f^*.

Given this grid, we can compute matrix elements K_{ij} such that

$$(f, Kf) = \sum_{i=1}^{N} \sum_{j=1}^{N} K_{ij} f_i f_j = f^T K f$$

exactly, for every $f \in S_N$. Moreover, we can find $a \in \mathbf{R}^N$ such that

$$(a, f) = \sum_{i=1}^{N} a_i f_i = a^T f,$$

up to an arbitrarily small error depending on the numerical evaluation of ϑ. The computation of (a, f) is discussed further in Section 3.7, since it depends on our method for evaluating ϑ.

3.3 Quadratic programming

Given a grid G and the associated values of K, φ and a, we now have a finite-dimensional quadratic program to solve:

$$\begin{cases} \text{minimize} & Q(f) \equiv a^T f + \frac{1}{2} f^T K f \\ \text{subject to} & 0 \le f_i \le \varphi_i, \qquad 1 \le i \le N. \end{cases} \tag{3.4}$$

The matrix K is positive definite since it is just the positive operator K restricted to S_N, so the solution f^N of (3.4) is completely determined by the *feasibility* requirement

$$0 \le f_j \le \varphi_j \tag{3.5}$$

and the *variational conditions*

$$\begin{aligned} f_j = 0 &\Rightarrow \nabla Q(f) > 0 \\ f_j = \varphi_j &\Rightarrow \nabla Q(f) < 0 \\ 0 < f_j < \varphi_j &\Rightarrow \nabla Q(f) = 0. \end{aligned}$$

The "active set method" [Fle80] for solving (3.4) is based on these two conditions. If we knew the active set $A \cup B$ of indices j where $f_j = 0$ (for A) or $f_j = \varphi_j$ (for B), then f solves the "equality constraint" problem of minimizing Q subject to $f_j = 0$ for $j \in A$ and $f_j = \varphi_j$ for $j \in B$. Since Q is quadratic, f can be found by solving a linear system

$$\hat{K}\hat{f} + \hat{a} = 0 \tag{3.6}$$

and setting

$$f_j = \begin{cases} 0 & j \in A \\ \varphi_j & j \in B \\ \hat{f}_j & \text{otherwise} \end{cases} \tag{3.7}$$

Here $\hat{K}$ and $\hat{a}$ are obtained from K and a by striking out every row and column with an index in $A \cup B$.

We don't know the active set, however, so we find it iteratively. Given an iterate f^0 and the corresponding A and B, we solve (3.6) for the inactive component $\hat{f}$ of a new iterate f. There are now two cases to consider. First, $\hat{f}$ may be feasible: $0 \le f_j \le \varphi_j$ for $j \notin A \cup B$. If this happens, then we must

check whether the active components of ∇Q have the right signs. If they do, then we are done and (3.7) gives us the minimizer. Otherwise, we remove from the active set the index j for which $\nabla Q(f)_j$ is furthest from feasiblfrom feasible, and repeat the iteration. Second, $\hat{f}$ may not be feasible. In this case, we start from the previous iterate f^0 and proceed as far as we can in the direction of f until a constraint is reached. Then we add the newly active constraint or constraints to the active set and repeat the iteration. Eventually, the iteration terminates at the minimizer. It costs $O(N^3)$ work per step, if we solve (3.6) by the standard technique of LU-decomposition [PFT86], and in the worst possible case it takes N steps; thus its worst-case time requirement is $O(N^4)$. However, there exist $O(N^2)$ techniques for solving (3.6) when only a few indices are added to the active set, and a good initial guess can result in a solution after only a few steps. Thus it can be as little as $O(N^2)$ in general to solve a quadratic program. This is still superlinear, so it is clearly a good idea to keep the number of grid points N as small as possible by choosing S_N appropriately.

3.4 Grid construction

We can now compute the minimizer associated to a given grid G. The next question is how to determine G. We construct G to fit f^* iteratively, by solving a sequence of approximate problems on finer and finer grids. The iteration involves two ideas; first, we speed up the determination of the active set by recursively constructing a finer and finer grid, and second, we reduce the number of degrees of freedom tremendously by refining the grid only selectively where it needs it.

We begin with the error bound for linear interpolation:

$$|f(k) - \tilde{f}(k)| \leq \frac{1}{4}h^2|f''(\tilde{k})|. \tag{3.8}$$

Here $\tilde{f}$ is the linear interpolant of f on an interval of length h containing both k and $\tilde{k}$. For our minimizer f^*, we know that there will be square root singularities, so we want to choose h small near the singularities where f'' is large, to maintain uniformly small error with the smallest possible number of grid points.

We don't know f^*, so we construct the grid iteratively. Starting with a uniform grid of, say, N points, we compute a minimizer f^N. Then we can use

f^N to approximate the error on each grid interval and bisect those intervals where the approximate error is large. This gives a finer grid, selectively refined near the singularities, on which we solve (3.4) again. The process is repeated until the grid is as fine as desired.

One further step must be added to this process in practice; smoothing the mesh between refinements. After intervals are bisected, the mesh size varies dramatically from one interval to the next. The computed minimizer f^N tends to reflect this with numerical kinks and irregularities where the mesh size jumps. To smooth over this difficulty, we smooth the mesh between each refinement with several steps of the following process: Take each grid point and move it to the midpoint of its two neighbors. This produces a smooth mesh, with no extraneous kinks in f^N, at the price of slightly increasing the refined area of each grid.

This process is far more efficient than using a fine uniform grid, for two reasons. First, the adaptive grid concentrates the N degrees of freedom where they do the most good, near the singularities. Second, the iterative structure gives us a highly accurate starting point for each new quadratic program, and this reduces the cost of solving the quadratic program by a factor of N. Numerical results demonstrating a large speedup are given in Section 4.3.

3.5 Computation of φ and ϑ

Now we can solve the minimization problem; to formulate it, however, we need to evaluate the constraint function

$$\varphi(k) = \int_{x_-(k)}^{x_+(k)} \frac{k}{\sqrt{-v(x) - k^2}}\, dx \tag{3.9}$$

and the coefficient

$$\vartheta(k) = kx_+(k) + \int_{x_+(k)}^{\infty} k - \sqrt{k^2 + v(x)}dx. \tag{3.10}$$

If v has compact support, then the integral in ϑ is actually over a finite interval $[x_+(k), R]$ and

$$\vartheta(k) = kR - \int_{x_+(k)}^{R} \sqrt{k^2 + v(x)}dx.$$

The integral defining φ is difficult to evaluate when k is close to 1, because x_- and x_+ coalesce around a singularity in the integrand. As $k \to 1$, the roots $x_\pm \to 0$ since $x = 0$ is a strict minimum of v and $v(0) = -1$. Thus we can Taylor expand v about $x = 0$. A routine calculation then gives

$$\varphi(k) \sim \sqrt{\frac{2}{v''(0)}} \quad \pi k \qquad \text{as} \quad k \to 1;$$

recall that we assumed $v''(0) > 0$. In particular, $\varphi(1) = \sqrt{2}\pi/\sqrt{v''(0)}$.

For other values of k, we calculate φ and ϑ by numerical quadrature. We use bisection [PFT86] to find $x_\pm$ to the desired precision, then integrate using adaptive singular quadrature rules from Quadpack [PdKUK83] to handle the square-root singularities of the integrands in (3.9) and (3.5). This procedure is expensive, but produces excellent and controllable accuracy.

3.6 Adaptive tabulation

We need many values of φ and ϑ in the course of a calculation; for example, if we use 200 points in x and 50 steps in t, and it takes 100 points in k for each evaluation of $\bar{u}$, then we need well over one *million* φ and ϑ evaluations. This could well be the major expense in a large calculation. φ and ϑ depend only on the initial data and not on x and t, so it seems unnecessary to re-evaluate them constantly, but we need their values at unpredictable k-grid locations for each x and t, so they cannot simply be evaluated once and for all.

A standard method [PFT86] for dealing with this situation is *tabulation and interpolation*: we lay down a uniform grid on $[0, 1]$, evaluate φ and ϑ at the grid points, and store the resulting table of values. When we need $\varphi(k)$, we approximate it by interpolating between tabulated values. In the present case, however, it turns out that φ and ϑ have singularities, so uniform tabulation fails to work. Millions of points are needed to get uniform three-digit accuracy.

We overcome this difficulty by *adaptive* tabulation; we tabulate φ and ϑ on a nonuniform grid of points such that linear interpolation between grid points gives their values to any user-specified tolerance ϵ_T. Of course, ϵ_T must be larger than the error in numerical evaluation of φ and ϑ to begin with, or we'll just be resolving numerical noise.

To construct such a grid, say for φ, we begin with a coarse grid of say twenty points and φ values at each grid point. We form a *stack* storing all

the intervals of the grid, and process each interval in the stack by comparing φ and its linear interpolant at the midpoint. If the error exceeds ϵ_T, we split the interval in half and store each of the two new half-intervals at the end of the stack. Otherwise, we proceed to the next interval in the stack. When we reach the end of the stack, we have a grid on which φ is very likely to be approximated within ϵ by linear interpolation. The grid for ϑ is constructed in precisely the same way.

This adaptive tabulation scheme reduces the CPU time required for large calculations by orders of magnitude, at a very modest expense in storage and initialization. Figure 2 shows the tabulation constructed for the example of Section 3.1, with the adaptive tabulation points shown as tick marks, with tolerance $\epsilon_T = 10^{-4}$. This example is discussed further in Section 4.3.

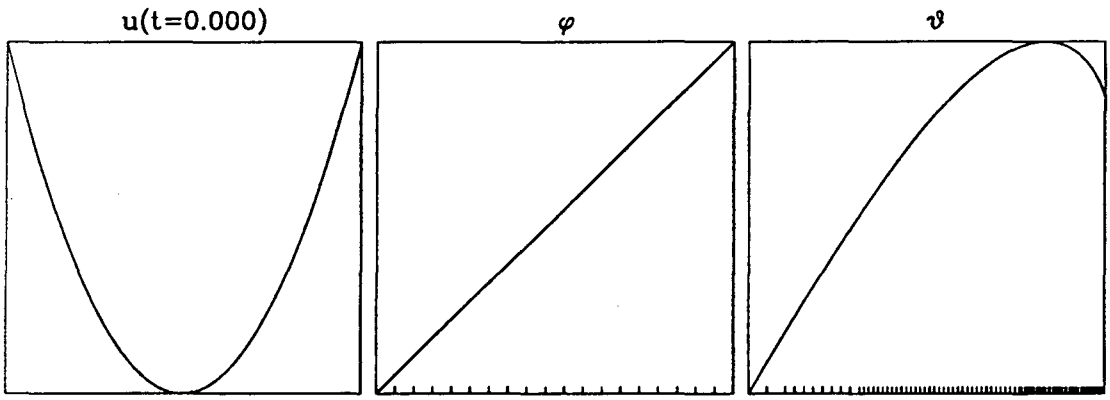

Figure 2: Initial data $v(x)$ and spectral functions φ and ϑ for Example A, with the adaptive tabulation points indicated by tick marks and $\epsilon_t = 10^{-4}$. The interval for $v(x)$ is $[-1, 1]$ and for the others is $[0, 1]$; all plots are scaled to the maximum of the function plotted.

Note that with this piecewise linear approximation for ϑ, we can compute (a, f) with error $\leq \epsilon_T$ by piecewise Gaussian quadrature. In other words, we interleave the grids where a (given by (2.5)) and f are piecewise polynomial functions of k (within ϵ_T) to get a single grid on $[0, 1]$ on which af is a piecewise quartic function of k (within ϵ_T); on each interval of this grid, three-point Gauss-Legendre quadrature integrates af exactly since q-point Gauss-Legendre quadrature is exact for polynomials up to degree $2q - 1$.

3.7 Numerical differentiation

We can now evaluate $Q^*(x,t)$ at any point with user-specified accuracy. The final challenge is now to compute

$$\bar{u} = \partial_x^2 Q^*(x,t) \tag{3.11}$$

by numerically differentiating Q^* twice. This is difficult for two reasons; $\bar{u}$ is not smooth in general, so neither is Q^*, and also the values of $Q^*(x,t)$ are computed independently at different spacetime points and therefore have uncorrelated errors. Hence a standard difference formula like

$$f_j'' = \frac{1}{h^2}(f_{j+1} - 2f_j + f_{j-1}) \qquad f_j = f(jh) \tag{3.12}$$

will produce noisy results; the numerical errors effectively get divided by h^2. We experimented with many methods, and eventually have settled, at least for the time being, on removal of noise with cosine smoothing (see below) in real space and then differencing with either widely-spaced centered differencing or ENO methods.

Cosine smoothing filters noise by computing averaged values

$$\tilde{f}_j = \frac{1}{4}(f_{j+1} + 2f_j + f_{j-1}) \qquad 2 \le j \le N-1, \tag{3.13}$$

on an N-point uniform mesh. The process can be repeated several times if necessary; it can be carried out with the FFT, but works better in real space when the data are not periodic. After smoothing, the standard centered difference formula (3.12) works well if $h = 1/N$ is not too small. Otherwise, we use ENO differencing [HOEC87]. The idea of ENO (an acronym for "essentially non-oscillatory") is to approximate $f''(x)$ by $P''(x)$, where P is a dth-degree polynomial interpolating f at a stencil of $d+1$ uniformly spaced grid points containing x in its interior. The distinguishing feature of ENO is that it slides the stencil left or right based on higher differences of f, in order to avoid the catastrophic loss of accuracy that goes with differencing across discontinuities. Thus it produces better results than centered differencing when jumps are present.

We found that d passes of cosine smoothing combined well with dth-degree ENO differencing, presumably because the stencils match.

Another promising method which we have not yet implemented is LSENO, for "Least-squares ENO." Here, we adopt the ENO sliding stencil but rather than interpolating, we fit a least-squares polynomial through the stencil. This smooths and differentiates simultaneously.

Remark: The Lax-Levermore theory offers an alternative formula for $\bar{u}$, namely

$$\bar{u}(x,t) = \int_0^1 k f_x^*(k;x,t)dk$$

where f_x^* is the derivative of f^* with respect to x. Given the active set, f_x^* can be computed directly without iteration and $\bar{u}$ can be evaluated by integration, so this seems more attractive than (3.11) at first glance. However, f_x^* is more singular than f^*, as (3.3) shows, hence harder to compute to high accuracy. The minimum value Q^* is much easier to compute accurately, because it is a minimum value. Hence the first variation of Q vanishes, so f^* computed with error ϵ in L^1 gives Q^* with error $O(\epsilon^2)$, much smaller than $O(\epsilon)$. This helps overcome the noise introduced by numerical differentiation.

3.8 Initial data

Our algorithm as formulated above and as coded works for essentially arbitrary single-well initial data $v(x)$. We want v negative, vanishing outside a compact set, with a single strict minimum at $x = 0$, where $v(0) = -1$ and $v''(0) > 0$. For purposes of exploring this infinite-dimensional space of experimental inputs, we must construct control parameters for a fairly broad class of initial data. The formula (2.9) for the breaktime suggests that we can produce interesting and complicated weak limits if we can control the inflection points of the solution; we can produce more interesting weak limits by adding more inflection points.

We adopt the following strategy for controlling v. We take v to be a quintic polynomial on each of a succession of intervals $[x_j, x_{j+1}]$, for $j = 1, 2, \ldots, M-1$. At each breakpoint x_j, we specify $v(x_j) = v_j$, $v'(x_j) = v'_j$, and $v''(x_j) = v''_j$, with $v_1 = v_M = 0$ for continuity. Then we set $v = 0$ for $x \le x_1$ and $x \ge x_M$. This structure is easy to control and gives a C^2 function, except possibly at the ends where we may allow v to be merely continuous. Furthermore, any continuous function can be uniformly well approximated by a function of this class. (The special data $v(x) = \min(0, x^2 - 1)$ of Example A falls conveniently into this class as well.) It is easy to evaluate

v and its derivatives since it is locally a polynomial. The integrals of v, which we need to measure the error in Q^* at $t = 0$, can be computed *exactly* using q-point Gauss-Legendre quadrature on each interval $[x_j, x_{j+1}]$. We need $q = 3$ for the first integral, since q-point Gauss-Legendre quadrature integrates polynomials of degree $2q - 1$ exactly, and $q = 4$ for the second integral.

3.9 An algorithm for KdV with $\epsilon > 0$.

At several points in this paper, we display the solutions of KdV and Burgers with $\epsilon > 0$, computed numerically. We now describe the algorithm we use for integrating the initial value problem

$$u_t - f(u)_x + \epsilon^2 u_{xxx} = 0 \qquad x \in \mathbf{R} \tag{3.14}$$

$$u(x, 0) = v(x) \tag{3.15}$$

with given f and a moderate to small dispersion coefficient ϵ. Other methods for solving this problem are described, for example, in [BDK86, CK85, FW78]. Our basic approach is to integrate out the linear constant-coefficient third order term by variation of parameters, then solve the remaining first-order hyperbolic problem numerically by a pseudospectral method. This approach allows us to use an explicit time-stepping method without an unnecessarily restrictive Courant condition.

This method is most conveniently implemented for periodic data v; since our data is compactly supported, we extend it to periodic data with a period sufficiently large that the imposed periodicity doesn't affect the results.

To integrate out the third-order term, we let A be the operator $-\epsilon^2 \partial_x^3$, with periodic boundary conditions. We introduce a new variable w by

$$u(\cdot, t) = e^{tA} w(\cdot, t).$$

Then w satisfies the integrodifferential equation

$$w_t = e^{-tA} f(e^{tA} w)_x.$$

To discretize the w equation, we used Fourier collocation in space and fourth-order Runge-Kutta in time. (A Fourier method is natural since e^{tA} is most naturally computed by Fourier series. We experimented with several

other explicit time-stepping methods, and settled on fourth-order Runge-Kutta because of its accuracy and stability properties.)

We validated our computations by testing the method on the soliton solution given in (5.6) with $k = 1$ and $\epsilon = 0.08$. These parameter values were chosen to make the soliton width comparable to the microscale in our numerical experiments in Section 5. Using 512 grid points on the interval $[-2, 2]$ and 5000 time steps, we achieved pointwise four-digit accuracy at $t = 1$, when the soliton has moved through a full period, indicating that our calculations are reliable. The L^2 norm and the energy

$$H = \int \frac{\epsilon^2}{2} u_x^2 + u^3 \, dx$$

of u (which are conserved for exact solutions), were conserved to at least four digits. The computation required about fifty minutes on a Sparcstation 1. A convergence study was carried out for each of the calculations presented in Section 5, and the solutions shown have converged to at least graphical accuracy. The L^2 norm and energy were conserved to several digits.

An exactly similar approach allows us to solve the Burgers equation (1.3), by defining the operator A differently. This is an easier calculation, because diffusion smoothes the solution where dispersion produces oscillations.

4 Validation

We now describe the numerical parameters of the method and demonstrate its accuracy and efficiency. Thus we validate the extensive computations presented in Section 5.

4.1 Parameters

The method requires several numerical parameters. First, there are the tolerances which control the accuracy of the φ and ϑ evaluation:

ϵ_r: the tolerance for rootfinding in the evaluation of $x_\pm$,

ϵ_i: the tolerance for numerical quadrature of φ and ϑ,

ϵ_T: the tolerance for tabulation of φ and ϑ.

These parameters cannot be varied independently. The tabulation tolerance must be substantially larger than the quadrature error, or else we are making an accurate table of the quadrature error rather than the function. Thus we set $\epsilon_T = 10\epsilon_i$ once and for all. The singularity in the integrand means that the roots $x_\pm$ must be found to much higher accuracy than ϵ_i, and must be inside rather than outside the singularity. Thus we fixed $\epsilon_r = 10^{-12}$. This tiny tolerance does not cost too much, for two reasons. Bisection costs $\log \epsilon$ work to get accuracy ϵ, so a few more digits cost only a little bit more CPU time, and also we evaluate φ and ϑ by rootfinding and quadrature only when we construct the initial tabulation.

Thus the error in φ and ϑ is controlled by the single parameter ϵ_T, the tabulation tolerance. We were able to achieve up to seven-digit accuracy with reasonable table sizes and running times. It turns out that φ and ϑ need to be evaluated to several more digits than f^*, or else f^* develops kinks. Thus our runs were mostly made with ϵ_T between 10^{-4} and 10^{-8}.

Next there are the parameters which control the k-grid on which f^* is computed:

N the number of points in the initial k-grid for each point (x, t),

ϵ_k the refinement tolerance for the k-grid,

n_L the number of levels of refinement permitted in the k-grid,

n_s the number of smoothings between k-grid refinements.

Clearly we will achieve at best error ϵ_k in f^*. This will lead, however, to error in Q^* which is on the order of $\epsilon_k^2 + \epsilon_T$, much smaller than ϵ_k, because Q^* is a minimum so the first variation vanishes. (The constraints contribute an additional ϵ_T because f^* is accurate at best of order ϵ_T where it touches φ.) The error in f^* can be reduced by increasing N or n_L, or by decreasing ϵ_k or ϵ_T.

Finally, there is the grid spacing h_x in the x variable, which controls the accuracy of $\bar{u}$ given Q^*. We also need to specify the type of differencing (ENO or centered) and the number of smoothing passes.

We study the performance of the method on two examples with different initial data $v(x)$:

A has $v(x) = \min(x^2 - 1, 0)$, so we know exact analytical formulas (Equations (3.1), (3.2) and (3.3)) for $f^*(k; x, 0)$, $\varphi(k)$ and $\vartheta(k)$, in addition to $Q^*(x, 0)$ and $\bar{u}(x, 0) = v(x)$, which we know for any $v(x)$. Figure 2 shows v, φ and ϑ, with the tabulation points shown as tick marks for $\epsilon_T = 10^{-4}$.

B has a more complicated piecewise quintic $v(x)$, for which we do not know f^*, φ and ϑ, and we study only Q^* and $\bar{u}$ at $t = 0$. Figure 3 shows v and the corresponding functions φ and ϑ; note the spike in φ and the clustering of the tick marks at the spike. The breakpoints, values and first two derivatives of $v(x)$ are given in Table 1.

4.2 Accuracy

We now discuss the errors in φ, ϑ, and f^* for Example A at $t = 0$, and the errors in Q^* and $\bar{u}$ for both examples A and B at $t = 0$.

The errors in φ and ϑ, measured in the max-norm

$$|e|_\infty = \max_j |e(k_j)|$$

over the set of midpoints k_j of intervals in the tabulation, are less than the tolerance ϵ_T by construction. Errors measured at random points typically are no larger than $3\epsilon_T$ even near a singularity.

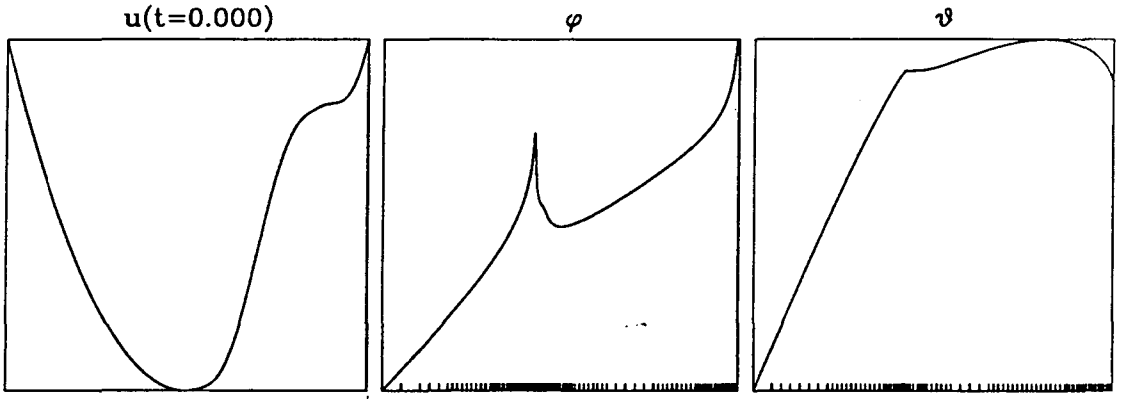

Figure 3: Initial data $v(x)$ and spectral functions φ and ϑ for Example B, with the adaptive tabulation points indicated by tick marks and $\epsilon_t = 10^{-4}$.

j	x_j	$v(x_j)$	$v'(x_j)$	$v''(x_j)$
1	-1	0	-2	2
2	0	-1	0	2
3	0.1	-0.99	0.2	2
4	0.7	-0.2	0.3	1
5	1	0	2	2

Table 1: Breakpoints, values and derivatives for the piecewise quintic $v(x)$ of Example B.

For Example A, we show the errors in f^* in the norms

$$f_{\infty,1} = \max_i \frac{\sum_j (k_{j+1} - k_{j-1}) |f^*(k_j; x_i, 0) - f^N(k_j; x_i, 0)|}{\sum_j (k_{j+1} - k_{j-1}) |f^*(k_j; x_i, 0)|}$$

and

$$f_{\infty,\infty} = \max_i \frac{\max_j |f^*(k_j; x_i, 0) - f^N(k_j; x_i, 0)|}{\max_j |f^*(k_j; x_i, 0)|},$$

and the maximum errors in Q^* defined by

$$Q_\infty = \frac{\max_i |Q^*(x_i, 0) - Q^N(x_i, 0)|}{\max_i |Q^*(x_i, 0)|}.$$

Here k_j are the points of the adaptive k-grid, x_i are the points of the uniform x-grid, f^* is the exact minimizer and f^N is the computed approximation to f^*. Similarly Q^* is the exact minimum and Q^N the computed minimum. Table 3 gives results from several refinement strategies. Case a) shows the result of *uniform* refinement with $N = 5, 10, 20, \ldots, 320$ points. The maximum L^1 errors $f_{\infty,1}$ show a rough first-order convergence, while $f_{\infty,\infty}$ is more erratic, as we would expect. By contrast, the error in Q^* is clearly second-order, in accordance with the variational structure of the problem. Cases b) through e) show similar errors, and will be discussed in the next section. In these examples, we took $\epsilon_T = 10^{-8}$ and $h_x = 0.1$ on the interval $-1 \le x \le 1$.

Figure 4 shows the computed and exact minimizers for an adaptive grid solution, the fifth line of Case c), from $x = -1$ to $x = 1$ in steps of 0.2. The grid points are indicated by tick marks along the lower edge, and the constraint φ is shown dashed. To graphical accuracy, the errors are almost invisible.

Next we study the results of various differencing techniques; centered and ENO with and without smoothing. Table 4 shows the errors in $\bar{u}(x, 0)$ computed by various methods for Example A, and Table 5 does the same for Example B. Here we used enough points in the adaptive k-grid to get six-figure accuracy in Q^*, so we are looking almost entirely at differencing error. The subscripts denote relative L^1 and L^∞ errors in $\bar{u}$. As the number n_x of grid points increases, the errors decrease at first, then the noise in the Q^* values takes over and the errors increase again. Thus there is an optimum grid size for a given accuracy in Q^*. (The ENO results for Example A are anomalous because the exact solution is quadratic so third-order ENO is exact. Hence the error is entirely due to noise in this case.) This is illustrated by Figure 5, which shows three types of differencing on Example A and the exact solution for comparison. Figure 6 shows similar results for Example B.

4.3 Efficiency

There are three areas in which our method attempts to be more efficient than a straightforward implementation. These are adaptive tabulation, recursive grid construction and selective refinement. The latter two are the main elements of our adaptive grid strategy.

First, we consider the efficiency of adaptive tabulation. Table 2 shows the number of points n, the mesh ratio m (largest grid interval divided by

smallest) and the construction time t in seconds for φ and ϑ for Examples A and B.

If φ and ϑ were smooth functions, the error in linear interpolation on a uniform n-point grid would be $O(n^{-2})$, so error ϵ_T could be achieved with $n = O(\epsilon_T^{-1/2})$ points. The purpose of adaptive tabulation is to maintain this relationship even for nonsmooth functions. Table 2 shows that decreasing ϵ_T by a factor of 10 increases n_φ and n_ϑ by a factor of roughly $\sqrt{10} \simeq 3.2$, even though φ and ϑ are both singular. (The first column for Example A is anomalous because φ is actually linear, so linear interpolation is exact.) The CPU time increases slightly faster than this, because we are evaluating the functions to greater precision as well as at more points. The large mesh ratios m_φ and m_ϑ show that millions of points would be needed to achieve the same accuracy with a uniform mesh.

Example	ϵ_T	n_φ	m_φ	t_φ	n_ϑ	m_ϑ	t_ϑ
A	10^0	20	1	0.26	20	1	0.16
	10^{-1}	20	1	0.28	20	1	0.18
	10^{-2}	20	1	0.31	21	2	0.24
	10^{-3}	20	1	0.28	33	16	0.44
	10^{-4}	20	1	0.31	97	256	1.66
	10^{-5}	20	1	0.29	303	1024	5.63
	10^{-6}	20	1	0.28	955	4096	18.69
B	10^0	20	1	0.26	20	1	0.16
	10^{-1}	20	1	0.30	20	1	0.19
	10^{-2}	31	32	0.68	22	4	0.23
	10^{-3}	74	128	2.50	34	32	0.54
	10^{-4}	216	512	11.12	105	256	2.37
	10^{-5}	683	1024	51.35	323	2048	9.27
	10^{-6}	2165	1024	200.47	1031	8192	34.83

Table 2: Number of grid points n, mesh ratio $m = \max h / \min h$, and CPU time t required for the adaptive tabulation of singular functions φ and ϑ to accuracy ϵ_T.

Adaptive tabulation requires some initial investment of time to form the table of values for each function, but pays off in dramatically reducing the

cost of each later evaluation. With six-figure accuracy, for example, we have to spend three minutes of CPU time to form the φ table for Example B, but each evaluation by table lookup instead of by numerical integration costs 0.000011 seconds instead of 0.046 seconds, a speedup of order 4000. In a large calculation with say a million evaluations, this would save eleven hours of CPU time.

Next, we demonstrate the savings in CPU time produced by recursive grid construction. For this, we compare case a) of Table 3 to case b). In case a), f^* is computed with a single uniform grid, and the time required grows like $O(N^3)$ or worse for large N. In case b), we use the same grid, but we construct it recursively starting with a 5-point grid and bisecting each interval. This is much faster for large N, because we use each active set to start the next iteration. Hence we need only a few solves of the linear system (3.6) to solve the quadratic program. The time required now increases only like $O(N^2)$ for large N, and the code runs about twenty times faster for $N = 320$, obtaining six-figure accuracy in fifteen minutes rather than five hours per space-time point.

The other half of our grid refinement strategy, selective refinement, is demonstrated in case c) of Table 3. With selective refinement, only intervals where the error estimate is larger than ϵ_k are bisected. Now the CPU time required to attain error ϵ_k is a more reasonable measure of efficiency than the time versus N, since N varies from one point to another. The combined effect is to produce six-figure accuracy in ninety seconds rather than five hours, a speedup of more than two orders of magnitude.

The combined effect of these three speedups is a factor of several thousand for six-figure accuracy in Q^*, reducing the time required for evaluating Q^* at many points from months to hours.

Case	N	ϵ_k	n_L	$f_{\infty,1}$	$f_{\infty,\infty}$	Q_∞	t
	5	0.00E+00	1	0.51E+00	0.37E+00	0.43E-03	11.88
	10	0.00E+00	1	0.15E+00	0.27E+00	0.12E-03	12.51
	20	0.00E+00	1	0.54E-01	0.24E+00	0.24E-04	14.64
a	40	0.00E+00	1	0.22E-01	0.25E+00	0.93E-05	26.49
	80	0.00E+00	1	0.59E-02	0.13E+00	0.11E-05	122.65
	160	0.00E+00	1	0.47E-02	0.21E+00	0.62E-06	1238.18
	320	0.00E+00	1	0.11E-02	0.10E+00	0.80E-07	17179.04
	5	0.00E+00	1	0.51E+00	0.37E+00	0.43E-03	11.85
	5	0.00E+00	2	0.15E+00	0.27E+00	0.12E-03	23.68
	5	0.00E+00	3	0.54E-01	0.24E+00	0.24E-04	37.02
b	5	0.00E+00	4	0.22E-01	0.25E+00	0.93E-05	55.89
	5	0.00E+00	5	0.59E-02	0.13E+00	0.11E-05	97.65
	5	0.00E+00	6	0.47E-02	0.21E+00	0.62E-06	247.05
	5	0.00E+00	7	0.11E-02	0.10E+00	0.79E-07	879.71
	5	0.10E-01	1	0.51E+00	0.37E+00	0.43E-03	11.83
	5	0.10E-01	2	0.15E+00	0.27E+00	0.14E-03	22.36
	5	0.10E-01	3	0.54E-01	0.24E+00	0.24E-04	33.53
c	5	0.10E-01	4	0.41E-01	0.25E+00	0.93E-05	45.67
	5	0.10E-01	5	0.19E-01	0.16E+00	0.11E-05	58.96
	5	0.10E-01	6	0.91E-02	0.21E+00	0.68E-06	73.71
	5	0.10E-01	7	0.47E-02	0.91E-01	0.74E-07	89.24

Table 3: Scaled errors in f^* and Q^* in discrete L^1 and L^∞ norms, as functions of the initial grid size N, refinement tolerance ϵ_k, and maximum level of refinement n_L.

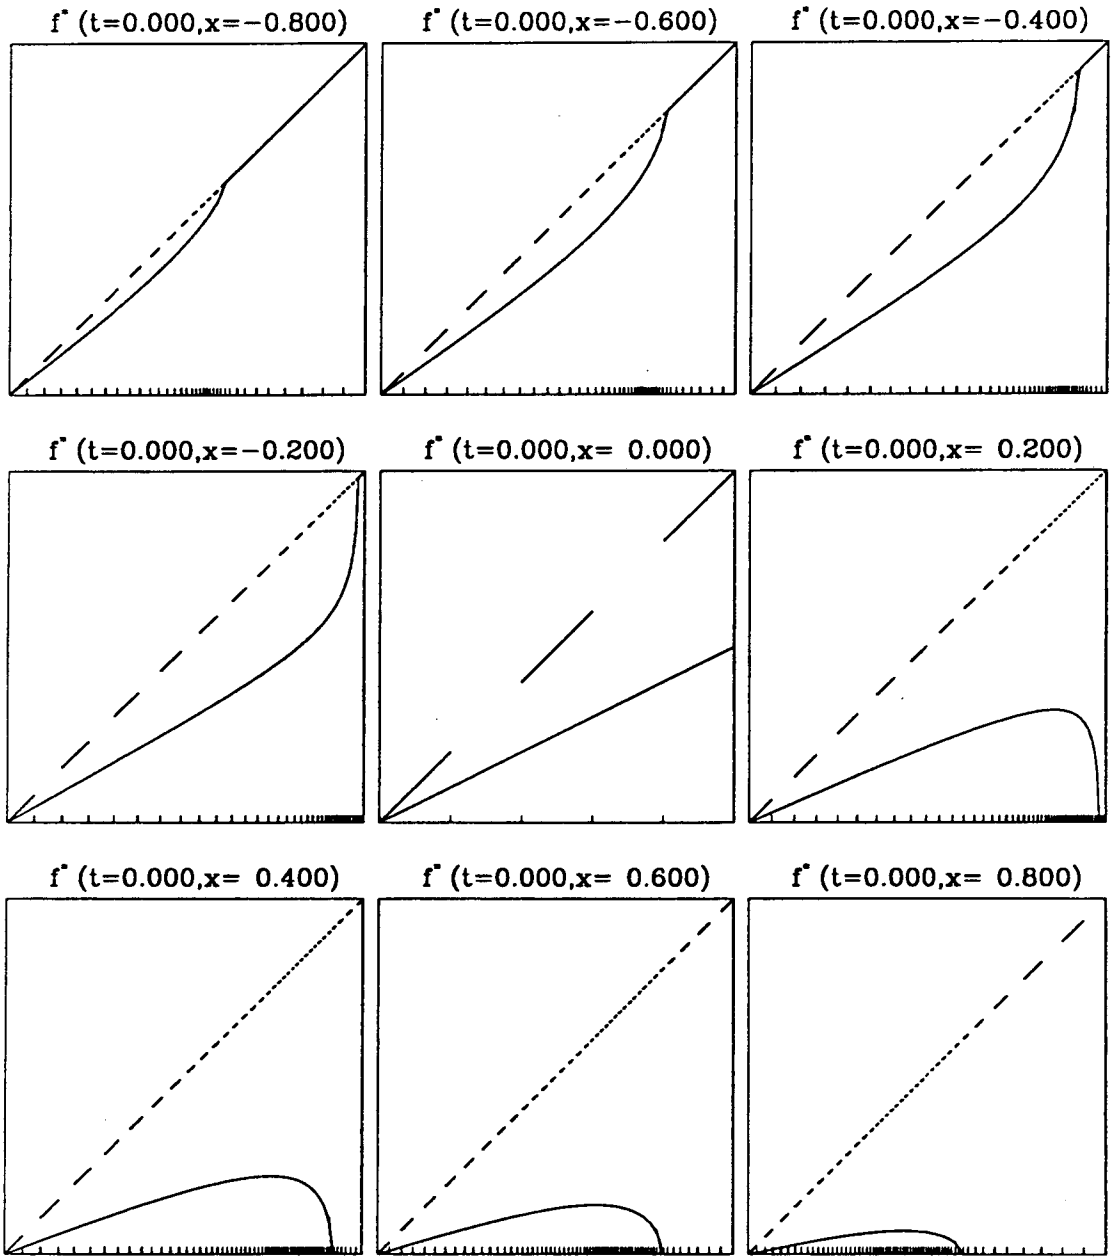

Figure 4: Exact and computed minimizers f^* for Example A at $t = 0$, computed with $-1 \leq x \leq 1$, $N = 5$, $n_L = 6$ and $\epsilon_k = 10^{-3}$.

Method	Order	n_x	$\bar{u}_1$	$\bar{u}_\infty$
Centered	2	10	$0.31E-01$	$0.74E-01$
	2	20	$0.51E-02$	$0.18E-01$
	2	40	$0.95E-03$	$0.46E-02$
	2	80	$0.20E-03$	$0.11E-02$
	2	160	$0.77E-04$	$0.46E-03$
	2	320	$0.16E-03$	$0.92E-03$
ENO	3	10	$0.25E-06$	$0.52E-06$
	3	20	$0.53E-06$	$0.13E-05$
	3	40	$0.78E-05$	$0.54E-04$
	3	80	$0.12E-04$	$0.95E-04$
	3	160	$0.52E-04$	$0.57E-03$
	3	320	$0.19E-03$	$0.25E-02$
Smoothed centered	2	10	$0.54E-01$	$0.58E-01$
	2	20	$0.12E-01$	$0.14E-01$
	2	40	$0.27E-02$	$0.35E-02$
	2	80	$0.65E-03$	$0.89E-03$
	2	160	$0.16E-03$	$0.22E-03$
	2	320	$0.72E-04$	$0.46E-03$
Smoothed ENO	3	10	$0.90E-01$	$0.65E-01$
	3	20	$0.22E-01$	$0.15E-01$
	3	40	$0.55E-02$	$0.38E-02$
	3	80	$0.14E-02$	$0.95E-03$
	3	160	$0.35E-03$	$0.27E-03$
	3	320	$0.96E-04$	$0.35E-03$

Table 4: Errors in $\bar{u}$ for Example A at $t = 0$, with various differencing methods and step sizes. The maximum error in any computed value of Q^* was less than 10^{-6}.

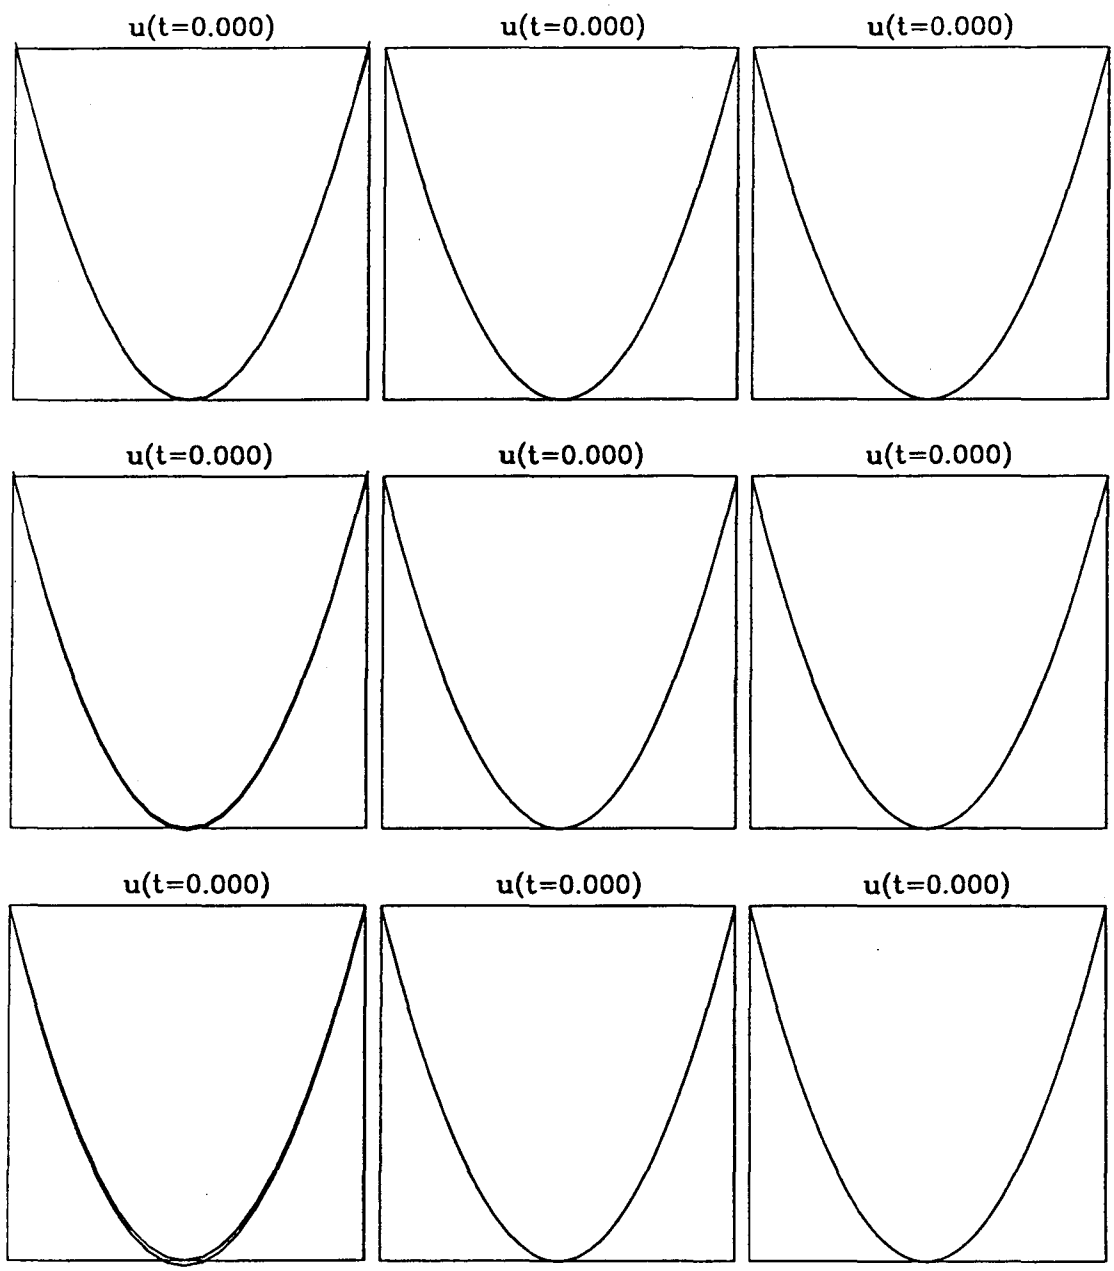

Figure 5: Exact and computed values of $\bar{u}$ for Example A at $t = 0$, with (from left to right) 20, 80 and 320 grid points and (from top to bottom) centered, smoothed centered, and smoothed third-order ENO differencing on $[-1, 1]$.

Method	Order	n_x	$\bar{u}_1$	$\bar{u}_\infty$
Centered	2	10	$0.42E-01$	$0.77E-01$
	2	20	$0.15E-01$	$0.97E-01$
	2	40	$0.28E-02$	$0.29E-01$
	2	80	$0.53E-03$	$0.56E-02$
	2	160	$0.31E-03$	$0.16E-02$
	2	320	$0.66E-03$	$0.42E-02$
ENO	3	10	$0.60E-01$	$0.27E+00$
	3	20	$0.17E-01$	$0.77E-01$
	3	40	$0.14E-02$	$0.83E-02$
	3	80	$0.36E-03$	$0.24E-02$
	3	160	$0.28E-03$	$0.18E-02$
	3	320	$0.73E-03$	$0.35E-02$
Smoothed centered	2	10	$0.73E-01$	$0.93E-01$
	2	20	$0.26E-01$	$0.54E-01$
	2	40	$0.67E-02$	$0.23E-01$
	2	80	$0.15E-02$	$0.49E-02$
	2	160	$0.38E-03$	$0.84E-03$
	2	320	$0.29E-03$	$0.13E-02$
Smoothed ENO	3	10	$0.12E+00$	$0.17E+00$
	3	20	$0.40E-01$	$0.60E-01$
	3	40	$0.12E-01$	$0.22E-01$
	3	80	$0.31E-02$	$0.68E-02$
	3	160	$0.76E-03$	$0.17E-02$
	3	320	$0.28E-03$	$0.79E-03$

Table 5: Errors in $\bar{u}$ for Example B at $t = 0$, with various differencing methods and step sizes. The maximum error in any computed value of Q^* was less than 2×10^{-6}.

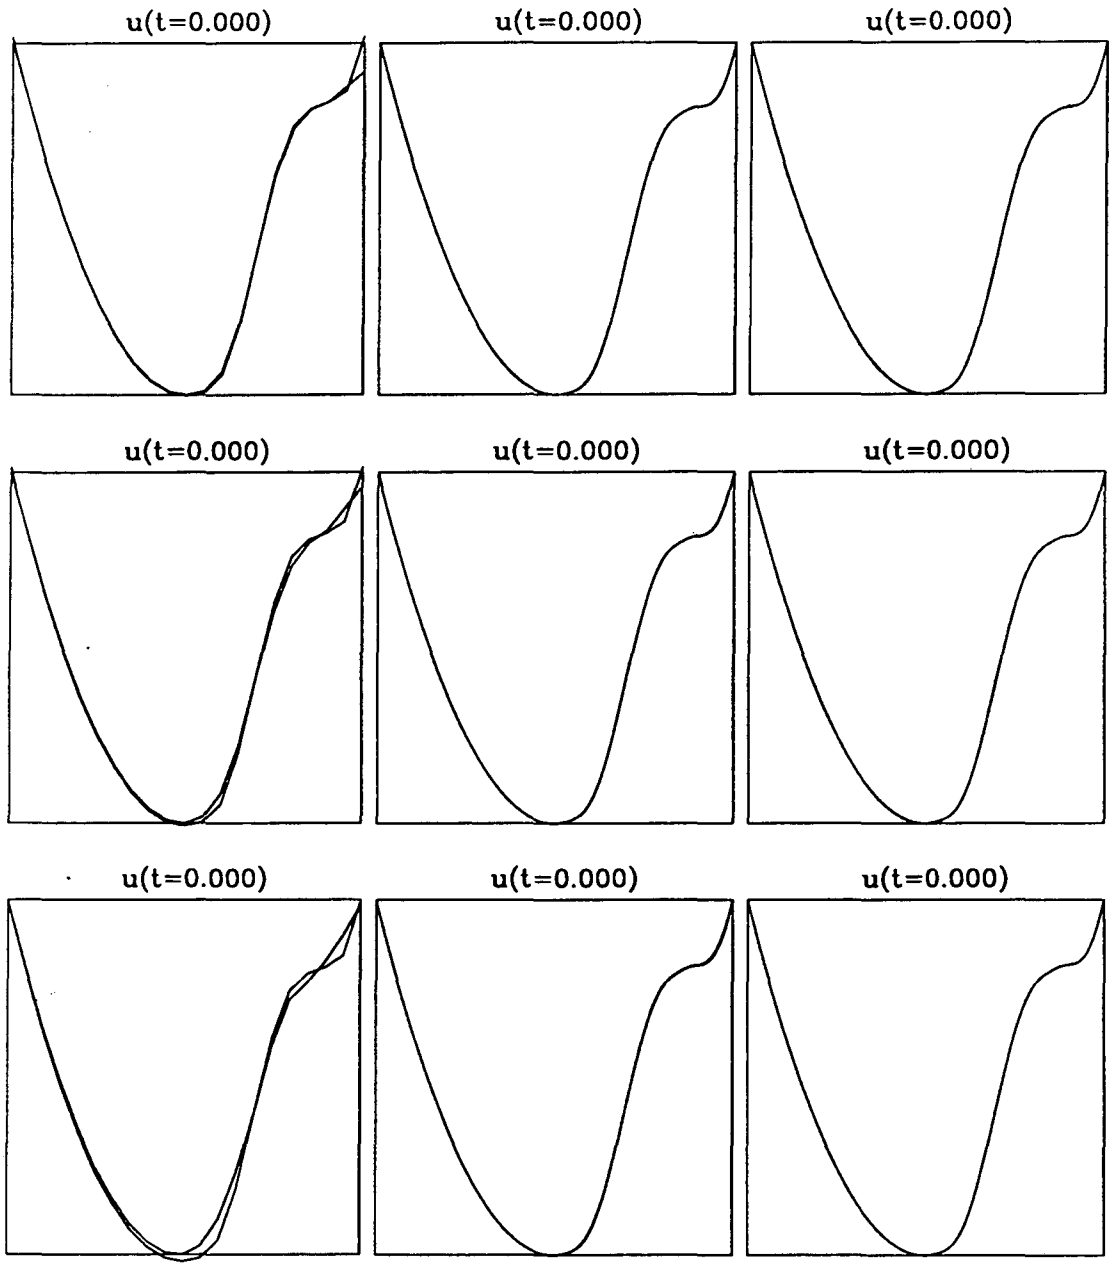

Figure 6: Exact and computed values of $\bar{u}$ for Example B at $t = 0$, with (from left to right) 20, 80 and 320 grid points and (from top to bottom) centered, smoothed centered, and smoothed third-order ENO differencing on $[-1, 1]$.

5 Structure of the Weak Limit

We now present large-scale computations which probe the detailed structure of the weak limit $\bar{u}$.

5.1 Parameters

We experimented with various combinations of the numerical parameters, seeking maximum accuracy in Q^* for a given time investment. We tried various tolerances for the computation of φ and ϑ, and eventually settled on on $\epsilon_r = 10^{-12}$, $\epsilon_i = 10^{-7}$ and $\epsilon_T = 10\epsilon_i$ as compromises between accuracy and speed. Values of ϵ_i much larger than this tended to produce spurious kinks in f^*, while smaller tolerances resulted in excessive CPU time requirements without improving accuracy.

We took $N = 20$ points equispaced on $[0, 1]$ as an initial grid for f^*, and set a refinement tolerance of $\epsilon_k = 10^{-3}$. Thus we expect to get at best one part in one thousand accuracy in f^* and at best one part in a million in Q^*. We found that four or five levels of grid refinement, with an equal number of grid smoothings between each refinement, produce sufficiently accurate results. The k-grid mesh ratio m for the more complex cases ranged as high as 20, indicating that the adaptive mesh strategy speeds up the calculation by several orders of magnitude for a fixed accuracy.

These computations were carried out on a Cray X-MP, more because of the large data files generated than because of the modest increase in speed obtainable with this not very vectorizable code. However, we emphasize again that the method is completely natural for massively parallel computing, because each value of Q^* solves a completely independent minimization problem.

Data for the first two examples A and B were described in Section 4.1, while Example C is somewhat more complicated.

5.2 Example A: Single-phase data

In Example A we consider extremely simple initial data:

$$v(x) = \min(0, x^2 - 1). \tag{5.1}$$

The corresponding spectral functions φ and ϑ are given in (3.1) and (3.2), and shown, for convenience, in Figure 7.

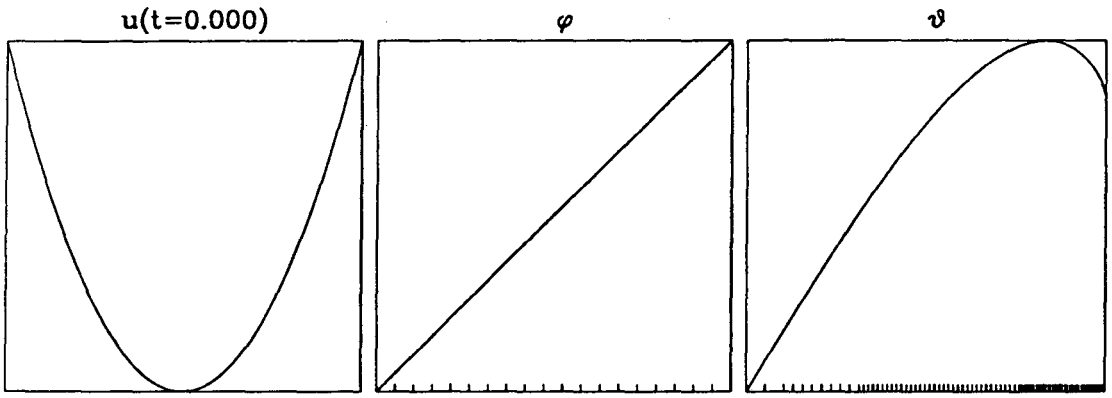

Figure 7: Initial data $v(x)$ and spectral functions φ and ϑ for Example A, with the adaptive tabulation points indicated by tick marks and $\epsilon_t = 10^{-4}$.

5.2.1 Orientation

For purposes of orientation, we first summarize some features of the zero dispersion limit which are known theoretically – emphasizing properties which are directly related to the quadratic data (5.1). Before the "breaktime" given by (2.9), the solution of KdV converges strongly to the classical solution u of the conservation law:

$$u_t - 6uu_x = 0 \tag{5.2}$$

$$u(x,0) = v(x). \tag{5.3}$$

Thus initially u steepens to the right until it "breaks" by developing an infinite derivative. Prior to breaking, all of the features of $\bar{u}$ can be deduced directly from the method of characteristics:

$$u(\xi(x,t),t) = v(x) \tag{5.4}$$

$$\xi(x,t) = x - 6tv(x). \tag{5.5}$$

In particular, the minimum value of $\bar{u}$ moves to the right at speed 6; breaktime and position are given by $(t_b = 1/12, x_b = 1)$. The point of first breaking

can be traced back to $x = 1$, the location at which the quadratic initial data abruptly changes concavity.

Breaking occurs when a tall, fast moving trailing wave in the left of the profile overtakes a shorter, slower moving component to the right. The manner in which the third order derivative in the KdV equation prevents the onset of a singularity in this conservative dispersive model is very different from a dissipative situation in which a small diffusive term eventually saturates steepening and produces a shock discontinuity in the zero-diffusion limit (see Figure 1). Here the third order dispersive term generates rapid oscillations, with $O(\epsilon)$ wavelengths and $O(1)$ amplitudes. These oscillations prevent the existence of a strong limit; after breaktime, the limit is only weak. This weak limit is described by a multisheeted surface over the $x - t$ plane which arises because, at and beyond breaktime, the overtaking process causes the function $\bar{u}$ to "fold," initially into a three sheeted surface. The projection of the folds onto the $x - t$ plane defines two curves which emanate from the breakpoint (x_b, t_b) and which separate that region of the $x - t$ plane in which the solution of KdV is oscillatory from that region in which oscillations are absent. These curves in the $x - t$ plane are analogous to caustics in the theory of linear dispersive waves [JLM91]. In that theory a three sheeted surface also arises in the description of the weak limit, which, in the linear case, can be viewed as a multivalued solution of the conservation law (5.2). Here, in the nonlinear case, the evolving three sheeted surface does *not* solve the conservation law (5.2); rather, it solves a coupled system of three conservation laws known as "Whitham's equations" [Whi74, Whi65]. The three state variables for this hyperbolic system are the three heights of the three sheeted surface, or equivalently, and more physically, the amplitude, local wave number, and local frequency of the oscillatory wave train. Formulas exist which map this three sheeted surface to the weak limit $\bar{u}$.

Lax and Levermore *prove* that, after breaktime, the weak limit is indeed described by a three sheeted surface which evolves according to Whitham's equations. For quadratic data (5.1), the $x - t$ plane splits into one region supporting rapid oscillations of the field, and a region of quiescent behavior. The space-time curve (caustic) which separates these two regions emanates from a "cusp" at (x_b, t_b). Asymptotically for large time, Lax and Levermore show that the leading caustic travels at speed 4. More detailed information about behavior near these caustic curves has recently been obtained by Tsarev [Tsa85], Tian [Tia92b], and Wright [Wri91]. Detailed surveys de-

scribing this general situation may be found in [LLV92] and in the conference proceedings [ELS91].

The speed 4 of the leading edge of the oscillatory region can be understood intuitively as follows: The oscillatory region is composed of densely packed solitons, and the leading edge is moving with the speed of the fastest moving soliton in the wave packet. Each of these solitons is of the form

$$u_j(x,t) = -2k_j^2 \mathrm{sech}^2\left(\frac{k_j(x - x_j - 4k_j^2 t)}{\epsilon}\right) \tag{5.6}$$

where k_j is one of the bound state eigenvalues of the Schrödinger equation with potential $v(x)$,

$$-\epsilon^2\psi_{xx} + v(x)\psi = -k^2\psi. \tag{5.7}$$

The largest eigenvalue is associated with the fastest moving soliton, and, for the data we study here and ϵ small, this largest eigenvalue will lie at the bottom of the potential $v(x)$ with value $k = 1 = \sqrt{-v(0)}$. Thus, the speed of the leading edge of the oscillatory region is expected to be $4k^2 = 4$, as the Lax-Levermore theory predicts and our numerical results confirm.

5.2.2 Results

We now compute the minimizer f^*, the minimum value Q^*, and the weak limit $\bar{u}$. In the calculation of the minimizer f^*, we keep track of the number of intervals in the active set, and thus count the number of phases at each space-time point. In this manner, we can depict the caustics bordering the oscillatory region. This information for the quadratic data (5.1) of Example A is depicted in the space-time diagram of Figure 8. Here the numerals at each point indicate the number of free endpoints of the active set, a slightly more informative quantity than the number of phases. The numerics clearly confirm the break point at $(x_b = 1, t_b = 1/12)$, the single phase nature of the oscillatory region, and the asymptotic speed of 4 for the leading edge of the oscillatory region.

Figure 9 shows the spatial profiles of the minimizer Q^* and the weak limit $\bar{u}(x,t)$ for times before, shortly after, and long after breaking. Notice that it is difficult to detect the effects of oscillatory microstructure in the wave itself directly from the weak limit, as the profile of $\bar{u}$ appears rather insensitive to

these underlying microscopic oscillations. (Of course, this oscillatory structure does affect the temporal evolution of $\bar{u}$ through the Whitham equations.) While continuous, $\bar{u}$ does lose smoothness at the location of the caustics, as Figure 9 shows.

For purposes of comparison, we have solved the KdV equation itself with $\epsilon = 0.16$, 0.08 and 0.04, using the method described in Section 3.9. In Figure 10 we show profiles of u_ϵ at the three times shown in Figure 9. Note in particular the regular structure in the oscillatory region, and the leading solitary wave at its front. This leading solitary wave has height 2, as is consistent with the soliton formula (5.6). On the other hand, this tall soliton is very narrow. In the weak limit, the height of $\bar{u}$ drops to 1 as is clear from the maximum principle for the weak limit [LL83]. Note also that before the breaktime, the oscillations vanish as $\epsilon \to 0$, while after t_b they remain $O(1)$ in amplitude, independent of ϵ. Also, Figure 1 displays a numerical solution of Burgers' equation

$$u_t - 6uu_x - \epsilon u_{xx} = 0, \tag{5.8}$$

with the same quadratic initial data and $\epsilon = 0.08$, in order to illustrate the striking differences between conservative and dissipative regularizations. Note that ϵ in Burgers' equation corresponds to ϵ^2 in KdV.

The pde space-time information contained in $\bar{u}$ is calculated from the minimizer f^* of the quadratic variational problem, which depends primarily upon spectral information. To understand the behavior of the minimizer and its connection to the weak limit, one must learn to interpret this spectral information. In Figure 11 we display the minimizer as a function of k at various space- time locations. In the figures, the "active" set of k values where f^* touches the constraints is the set of "gaps" in the spectrum of the local Schrödinger operator used for modulation theory [FFM80]. Prior to breaking, this local spectrum contains no gaps, and consists of one band of spectrum filling the k-interval $[0, \sqrt{-\bar{u}(x,t)}]$. Clearly, the right hand endpoint of this interval moves as a function of x and t. (In this picture, the continuous spectrum extends from $k = -\infty$ to $k = 0$ and a gap running from $k > 0$ to $+\infty$ doesn't count.)

After breaktime, in the oscillatory region, a gap $[0, k_0]$ opens in this local spectrum. This gap consists of the leftmost interval of the "active set" of k values, for which the constraint is active; thus, over the gap, the minimizer is locked to one of the two constraints $f^* = 0$ or $f^* = \varphi$. (One additional

feature, for this data with compact support, is that one of the endpoints of the gap remains locked to 0, as is predicted in [LL83]. In essence, this means that only two of the three sheets of the surface are free to evolve.) For large times, in the oscillatory region, the gap gets very large and pinches the band of spectrum to become very narrow. This narrow band of spectrum indicates the presence of spatially localized solitons in the wave. As one moves toward the leading edge of the oscillatory region, this narrow band of spectrum moves toward $k = 1$, which indicates the presence of a soliton at the leading edge which moves at speed 4 in spacetime.

5.2.3 Summary of Example A

In this first example, the Lax-Levermore predictions are certainly confirmed for simple quadratic data. At the leading edge of the oscillatory region, numerics and theory agree very well. The example also shows interesting algebraic behavior in the location of the trailing caustic. This algebraic behavior has been confirmed theoretically [Tia92a]. Finally, from this first example, one learns to interpret the spectral representation of the minimizer.

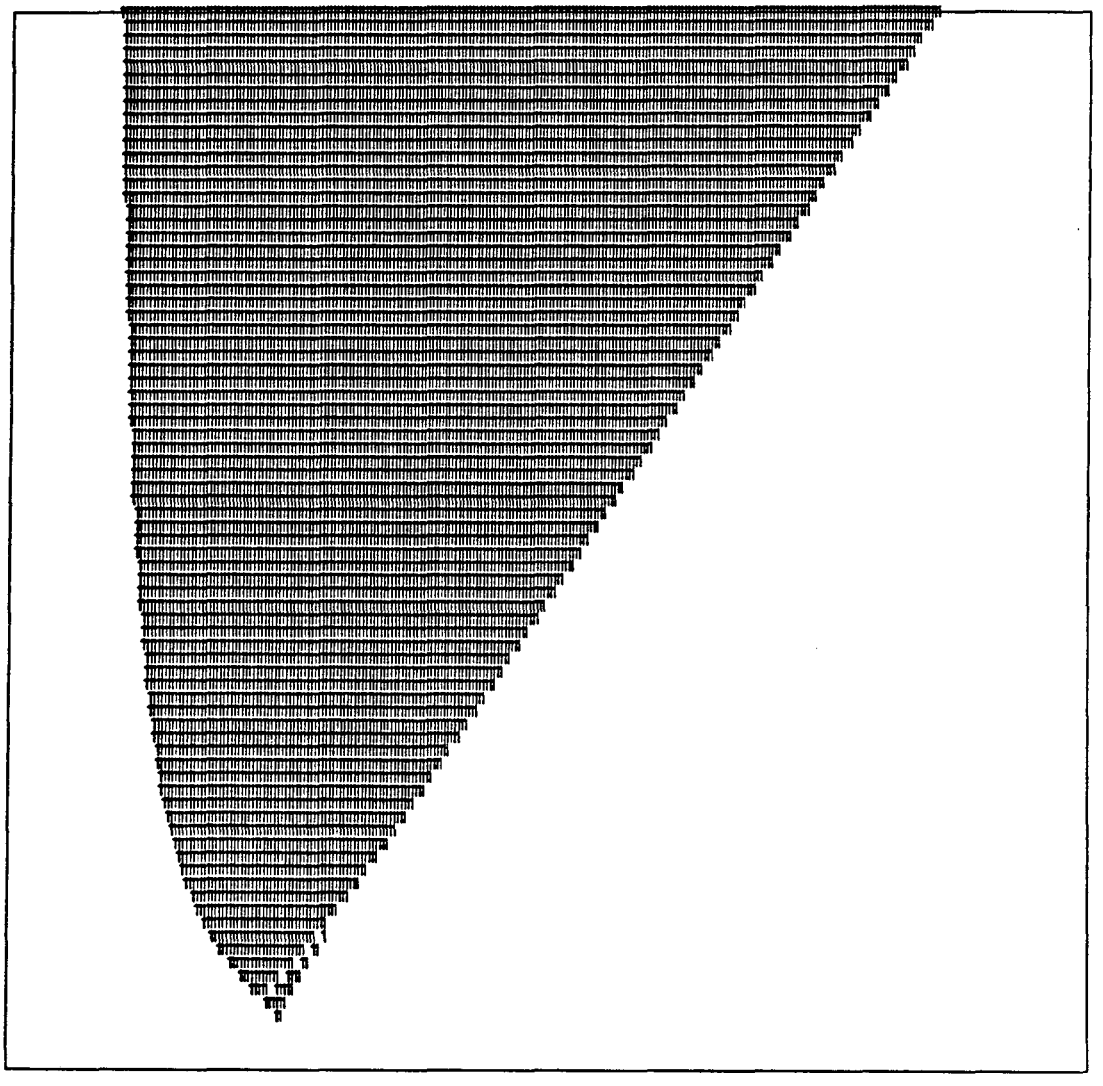

Figure 8: Space-time diagram for Example A. The domain is $[-2, 10] \times [0, 2]$.

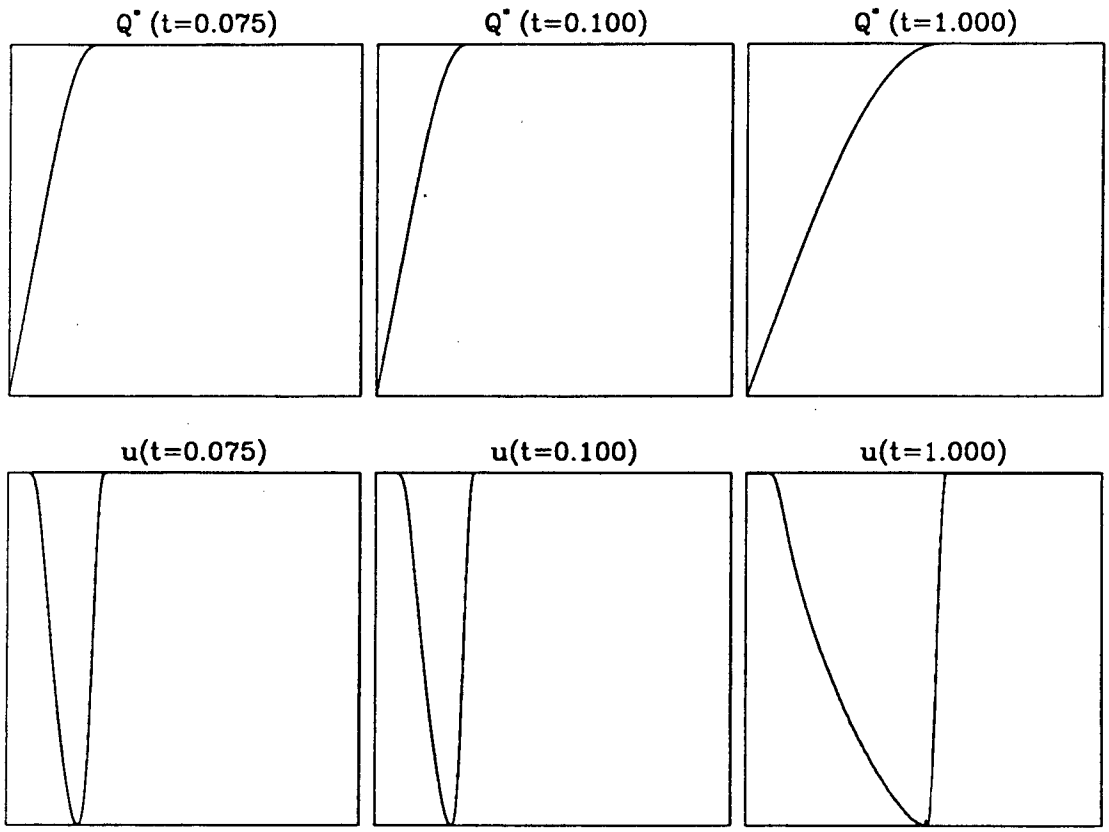

Figure 9: The minimum value Q^* and the weak limit $\bar{u}$ of Example A for selected values of t. The top row shows Q^*, while the second row is $\bar{u}$ computed with centered differencing on a 300-point grid on $[-2, 10]$.

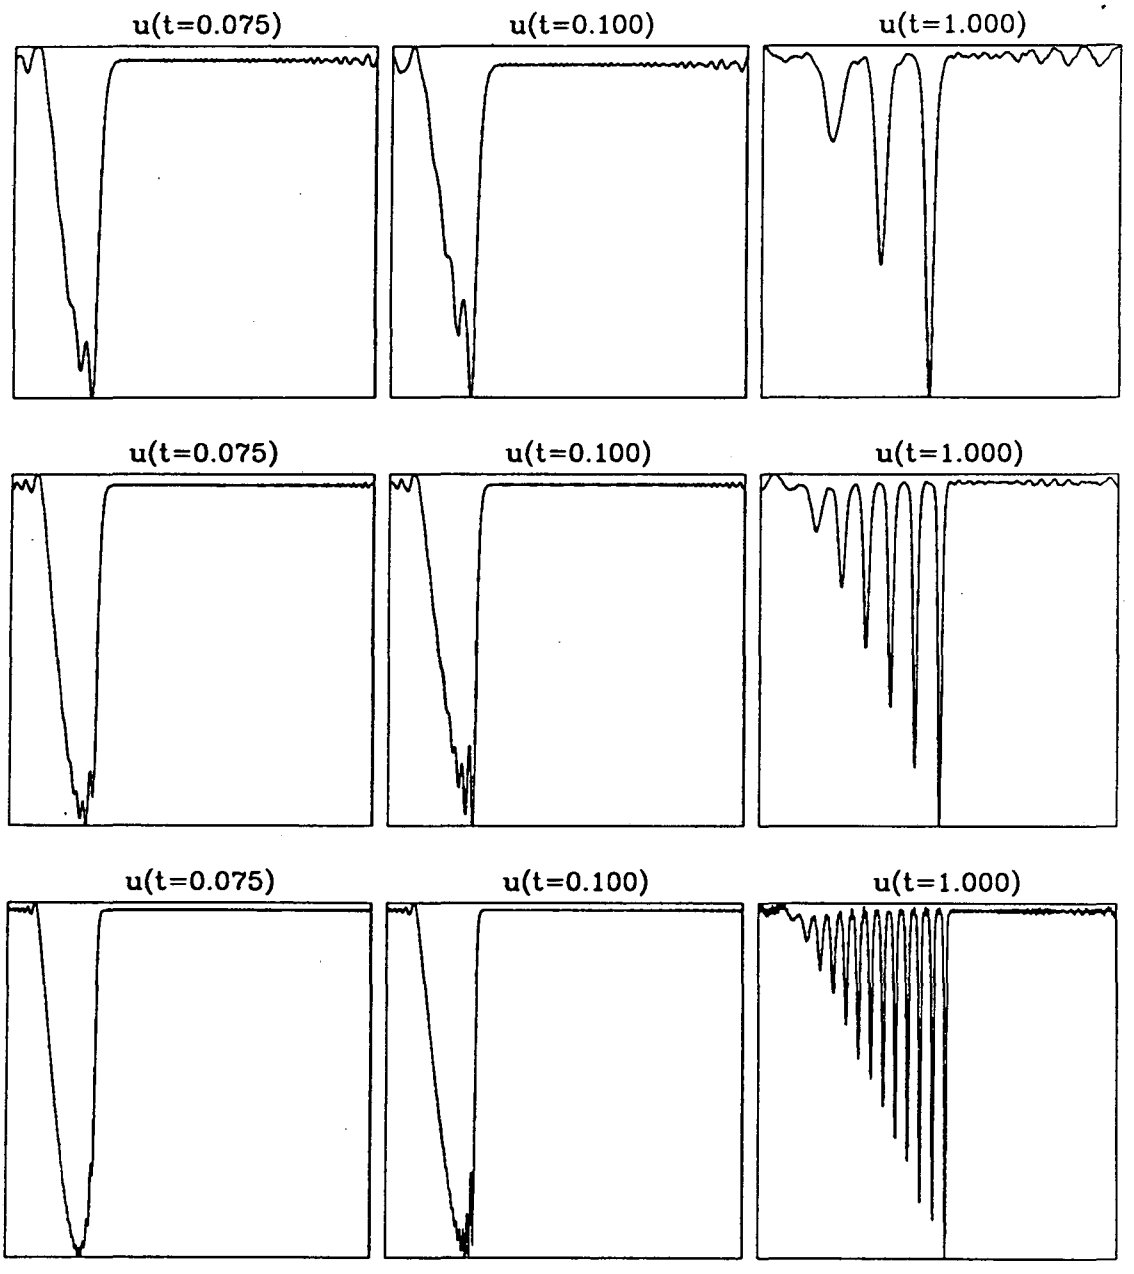

Figure 10: Solutions of KdV equation for Example A, with $\epsilon = 0.16$, 0.08 and 0.04 (from top to bottom). The domain of each plot is $[-2, 10]$.

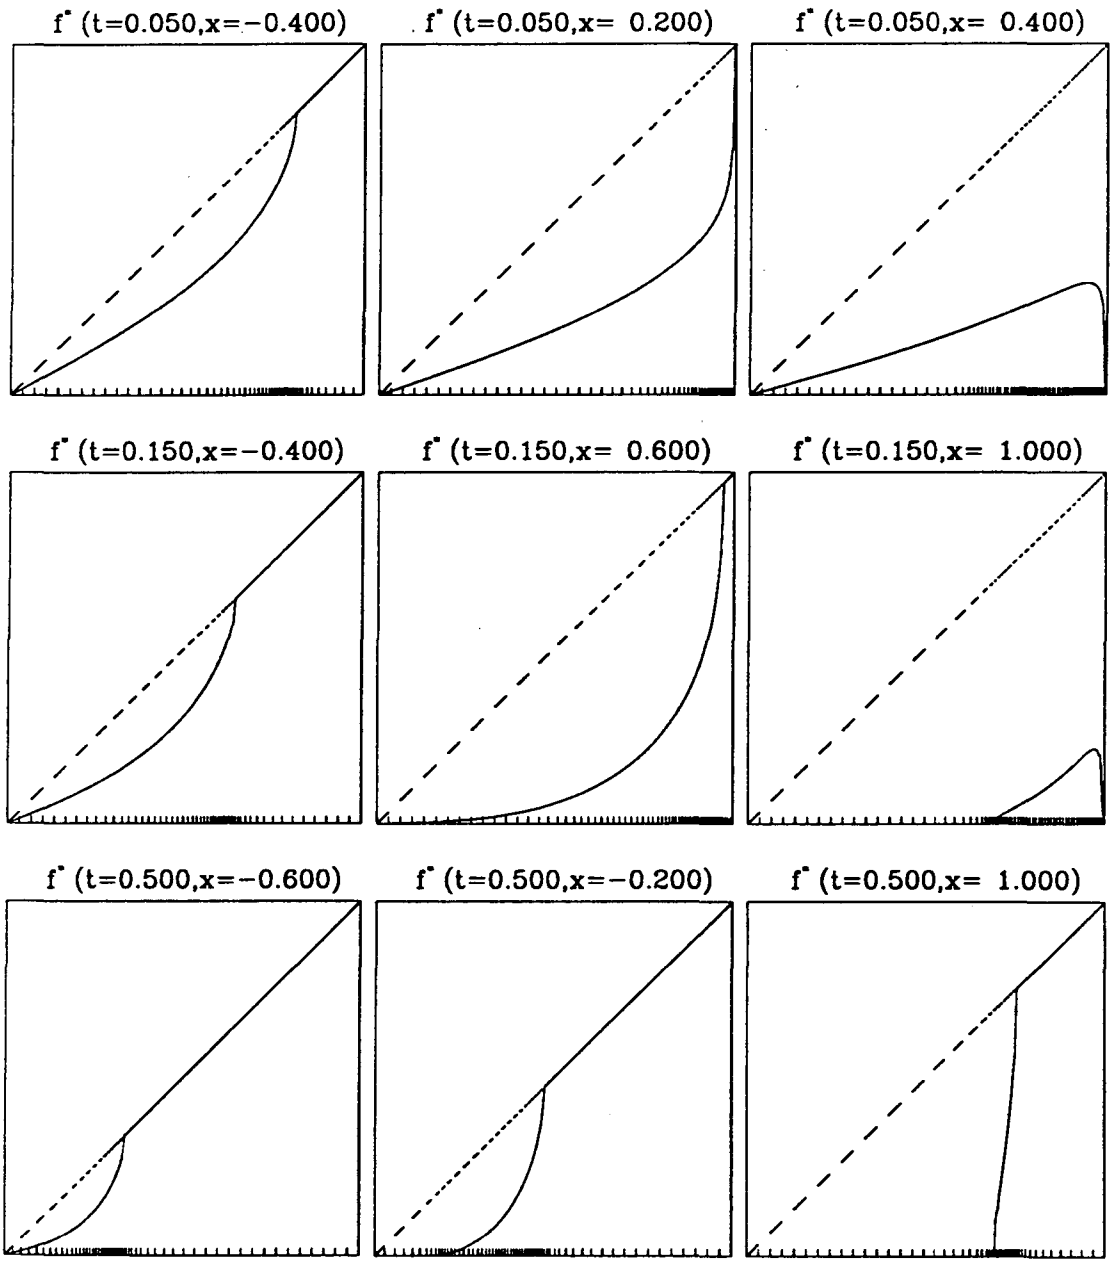

Figure 11: The minimizer f^* for Example A, shown for $t < t_b = 1/12$, $t > t_b$ and $t >> t_b$ on the interval $0 \leq k \leq 1$.

5.3 Example B: two-phase data

In this second example, we initialize data with several inflection points, a piecewise quintic function with breakpoints and values given in Table 1. Figure 3 shows v, φ and ϑ for this example. This data was chosen to generate a second breaking, with its associated two-phase behavior. In the two-phase regime, the weak limit $\bar{u}$ is analytically described with by a five-sheeted surface.

First, we display the spacetime diagram (Figure 12). Here we have printed the number of free end points of the active set at each point in spacetime; blank (0) is no phases, 1 is one phase and 2 and 3 correspond to two-phase regions. Note that asymptotically in time, the oscillations relax to single phase behavior, with the leading edge of the oscillatory region traveling to the right at speed 4. Each of these features confirms theoretical results in [LL83]. Also, note the extremely *sharp phase boundaries*, even in the transitions between one and two phase behavior. The "island" of 3's toward the top is probably connected to the rest, but the connector is skipped over by our uniform grid. For even longer time periods, it vanishes entirely as predicted.

Next, we display in Figure 13 the minimizer f^*. The "gaps" in the spectrum of the local Schrödinger operator corresponds to the active set of $k \in [0,1]$ where f^* is touching a constraint. Note that in the single phase regions the active set sometimes arises because of the constraint $f^* = 0$, and sometimes because $f^* = \varphi$. Also note that in some cases the endpoints of the gaps in the local spectrum (i.e., the heights of the surface) are locked to either 0 or 1. In these cases, the surface is more rigid than it is usually, with one or more of its heights locked and not free to move as a function of x and t. Also, note the minimizer associated with the 5-sheeted, 2-phase behavior, as displayed in the last row of Figure 13. Finally, for large time, the asymptotic behavior reduces to the single phase case, with a minimizer as in the preceding example A.

This example nicely illustrates an interesting feature connecting the structure of the minimizer $f^*(k)$ with inflection points in the initial data, a feature which may be viewed as a generalization to the multiphase case of a well known fact about breaking of a scalar conservation law. For the scalar case, breakpoints may be traced directly back to inflection points in the initial data. Here, inflection points in the initial data generate very sharp peaks

in its Abel transform, that is, in its WKB constraints. These sharp peaks are clearly visible in the data of Figure 3; their existence can be confirmed analytically by Taylor-expanding v about the inflection point and integrating the Abel transform (2.4) analytically. As the minimizer evolves, the presence of these sharp peaks in the constraints forces gaps to form (see Figure 13).

Next, we turn to the weak limit $\bar{u}$ itself, as displayed in Figure 14. In this figure, the profile $\bar{u}(t)$ is shown at several times. Comparison with Figure 12 shows that an initial breaking occurs at the front of the wave and generates single phase behavior. For this data, a two phase component later emerges, surrounded by two single phase components. As time increases, these oscillatory components slide toward the back of the wave. As it slides, the two phase component diminishes in spatial extent and finally disappears, leaving only a single phase component. While there is change in smoothness at the transition boundary curves, these changes are difficult to detect from observations of $\bar{u}(x,t)$ as a function of x for fixed t.

The subtle but apparent structure which the two-phase data produces in the weak limit (see Figure 14) is a challenge to more generally applicable numerical methods such as averaging small-ϵ solutions of KdV. To quickly illustrate the effect of oscillations on the weak limit, we next compute the solution of KdV for $\epsilon = 0.05$, display the spatial profiles (see Figure 15) and consider locally averaging over the oscillations to obtain the weak limit. The behavior of the weak limit $\bar{u}(x,t)$ should be contrasted with the actual oscillatory structure which is present in the wave at small, but positive, values of ϵ. We note that the expected *multiphase* microstructure in the spatial profile of u is not apparent at this value of ϵ, because the oscillations are certainly not fine enough to resolve the local two-soliton structure.

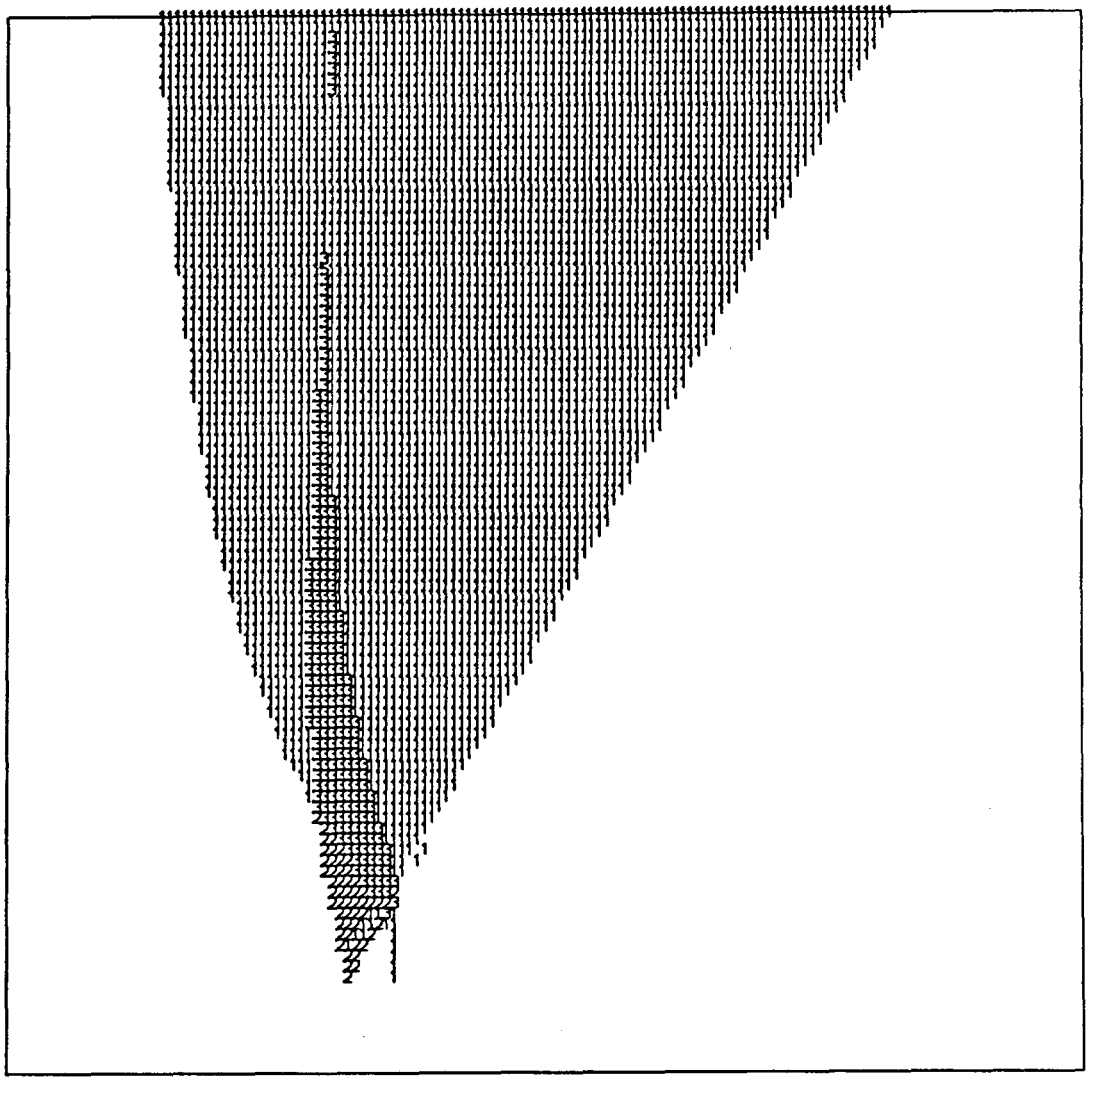

Figure 12: Space-time diagram for Example B. The domain is $[-1.5, 5.5] \times [0, 1]$.

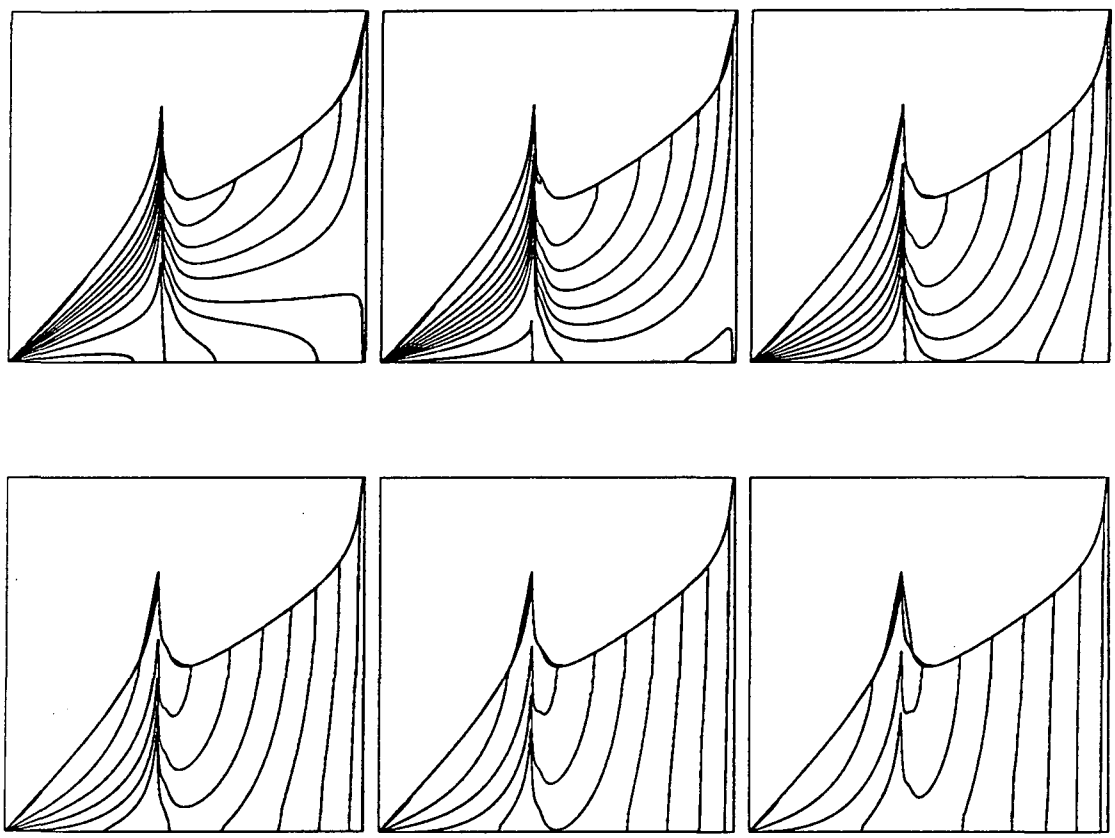

Figure 13: The minimizer f^* for Example B, shown for $t = 0$ to 0.5 in steps of 0.1 and for several x values at each t, from left to right and top to bottom.

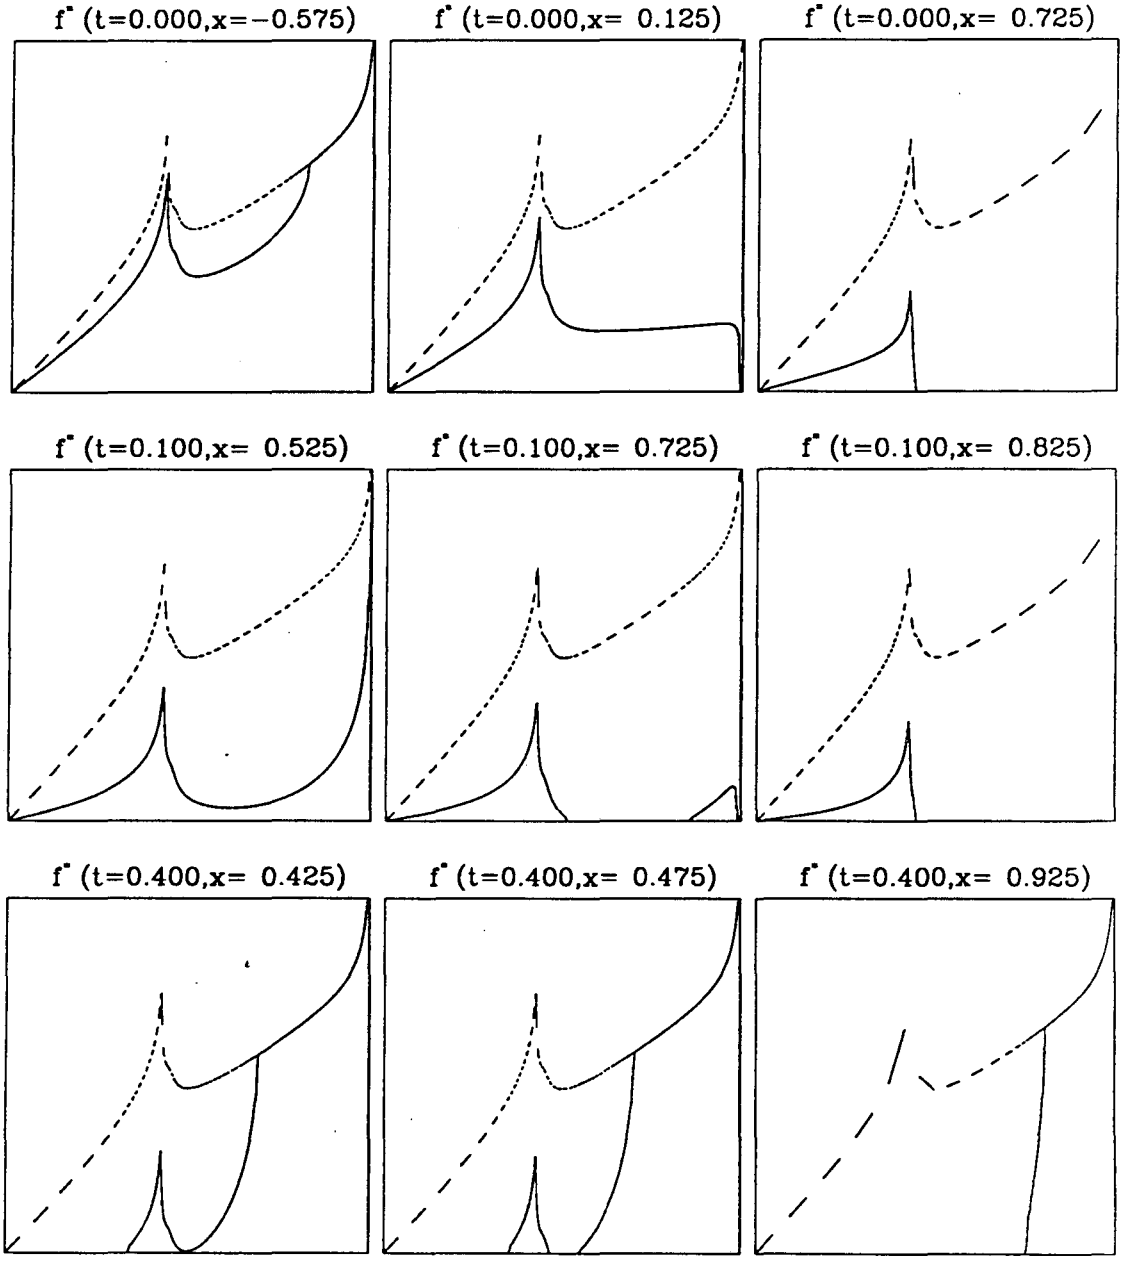

Figure 13: (continued) Selected minimizers from Example B, with the constraint shown dashed.

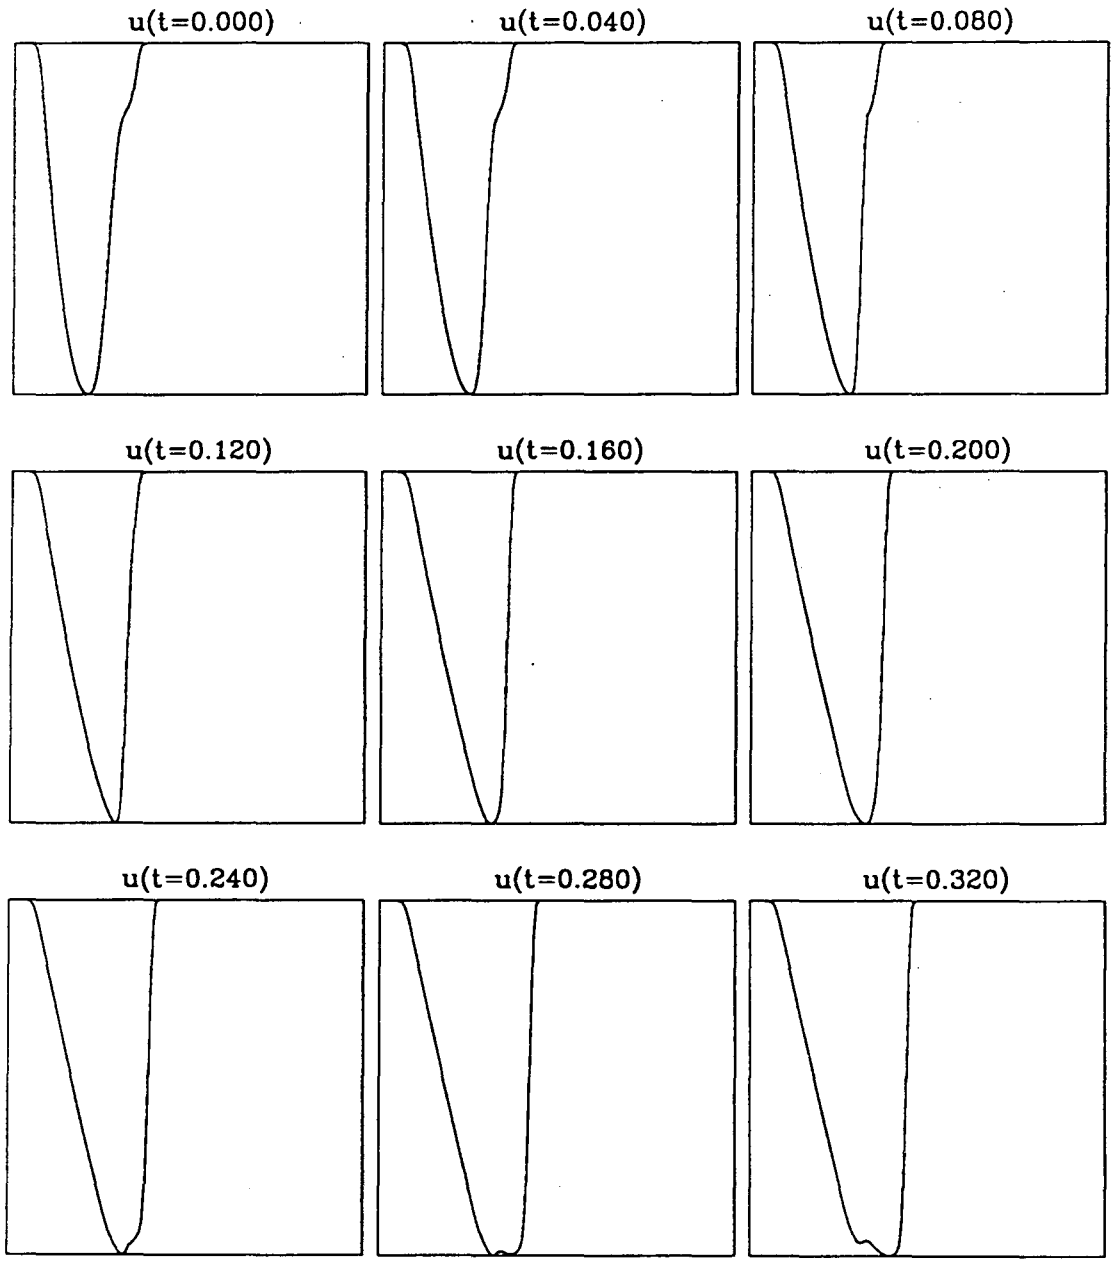

Figure 14: The weak limit $\bar{u}$ for selected values of t in Example B, computed with third-order smoothed ENO differencing on $[-1.5, 5.5]$.

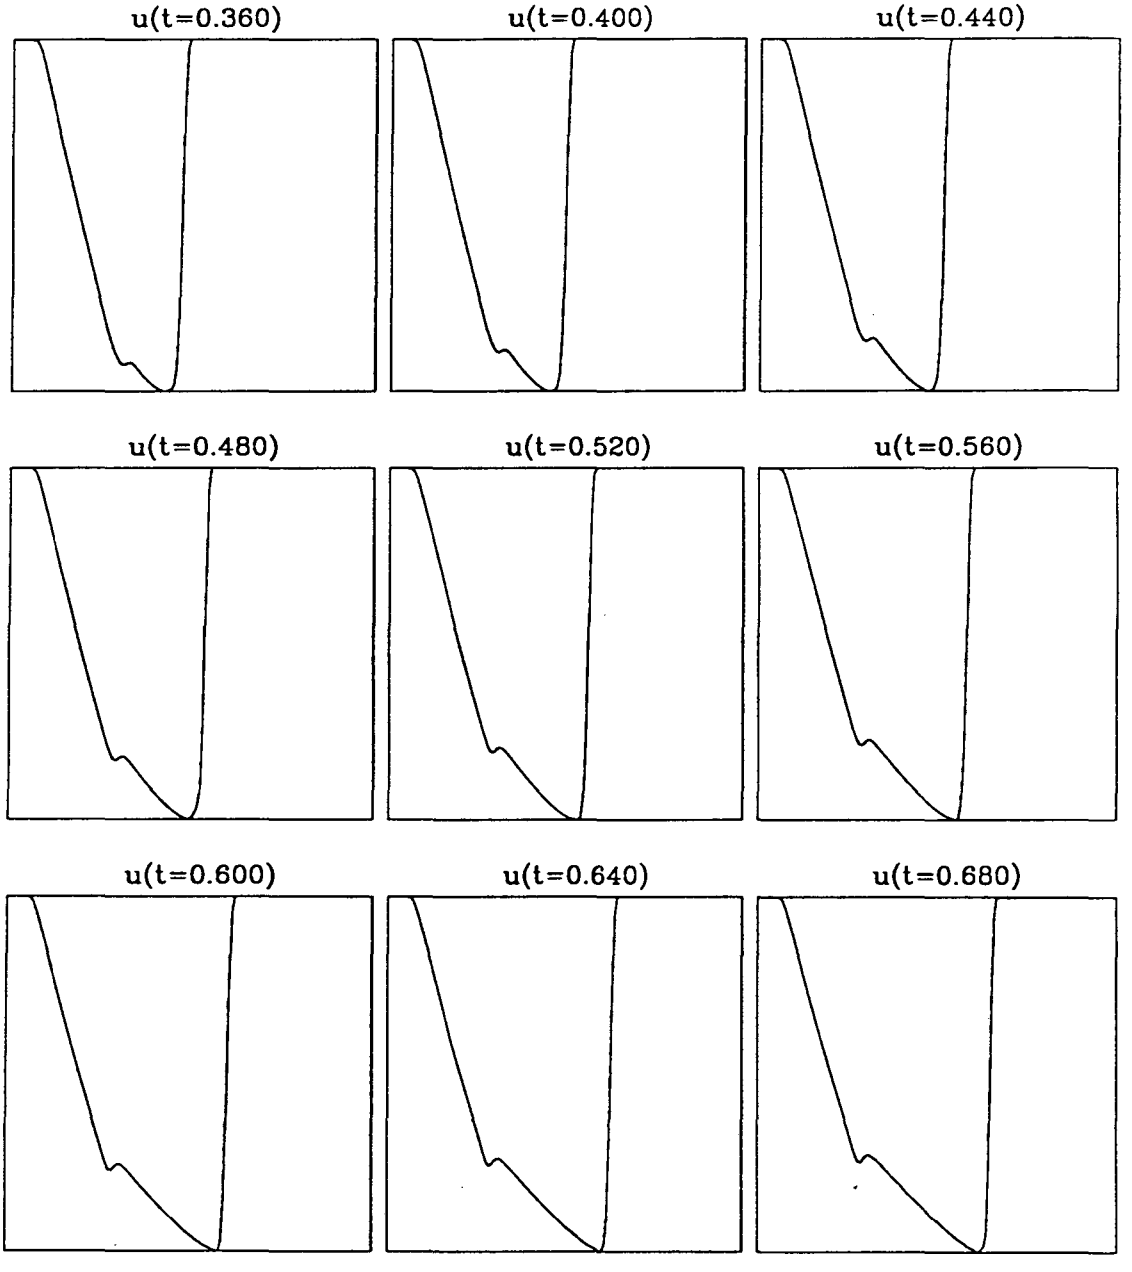

Figure 14: (continued)

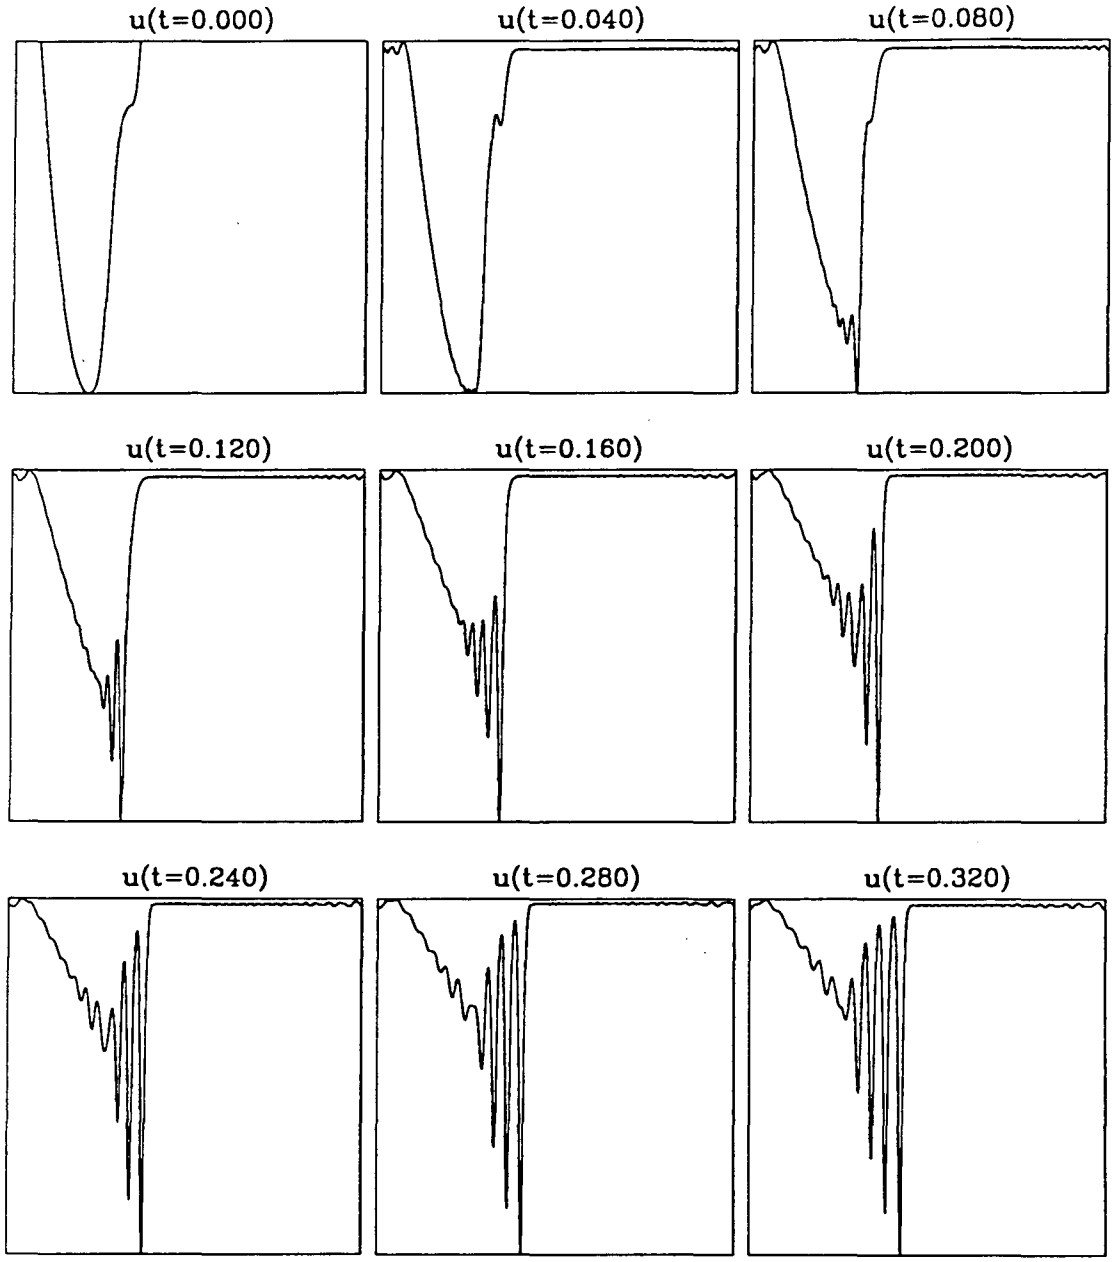

Figure 15: Solution of KdV equation for Example B with $\epsilon = 0.05$, computed on $[-1.5, 5.5]$.

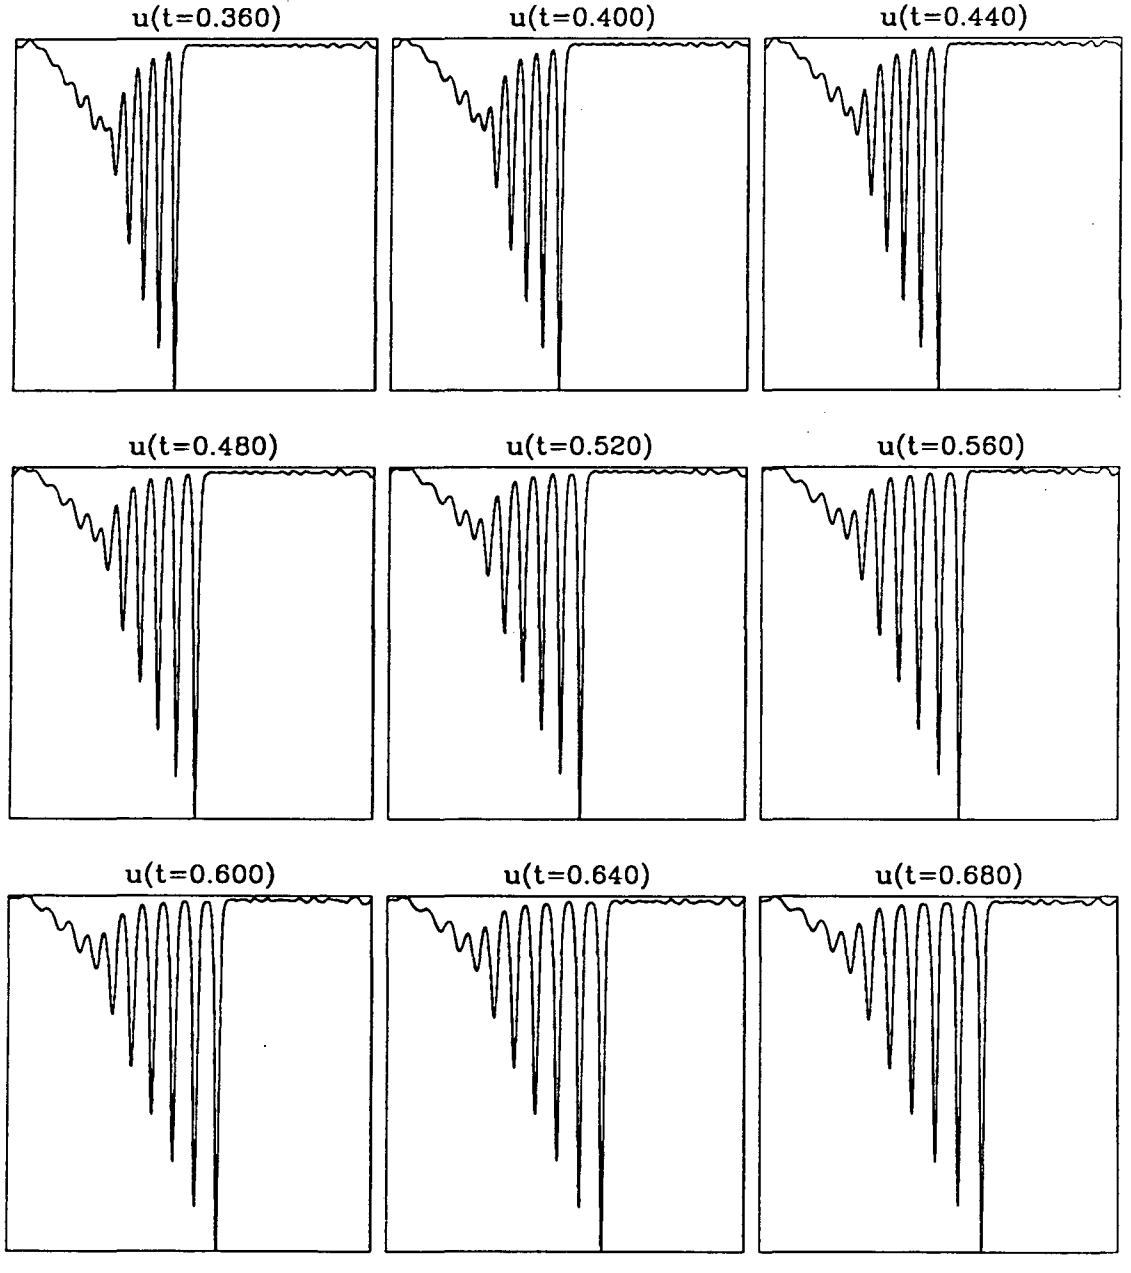

Figure 15: (continued)

5.4 Example C: three-phase data

Example B described phase boundaries between one-phase and two-phase components, phenomena which are just beyond current theoretical studies [Tia92b, Wri91]. In example C we proceed further, with data chosen to generate three-phase behavior. Specifically, we choose piecewise quintic data with breakpoints, values and derivatives given in Table 6. Figure 16 shows v, φ and ϑ for this example.

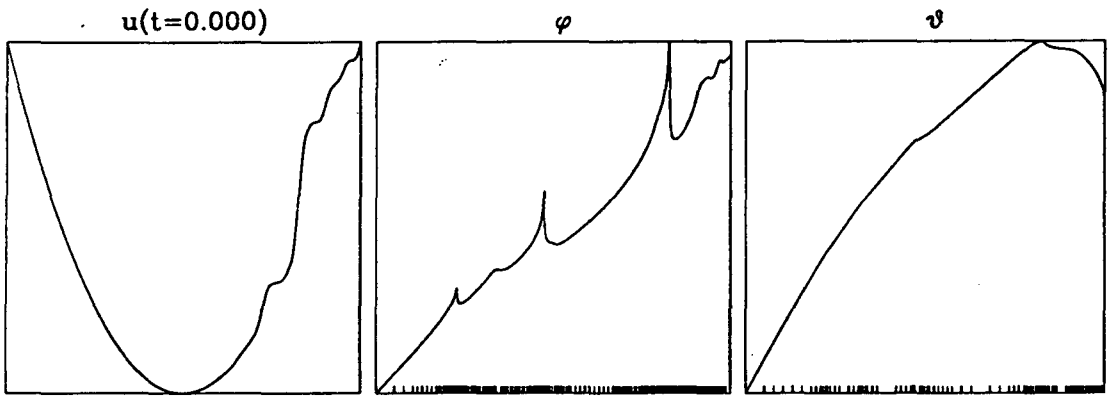

Figure 16: The initial data $v(x)$ and spectral functions φ and ϑ for Example C, with the adaptive tabulation points indicated by tick marks and $\epsilon_t = 10^{-4}$.

We first display the space-time diagram in Figure 17. The phase structure of this example is rather complicated, reaching three phases (a seven-sheeted surface marked with 4 and 5 in the figure). Still, one notes very sharp phase boundaries. Also, for this example, some of the phase regions are quite narrow in spatial extent, but last for long durations of time. Again, the long time behavior is single phase.

We organize the presentation of this example in the same order as in example B: (1) representative minimizers and local spectrum in Figure 18; (2) spatial profiles of the weak limit in Figure 19; (3) direct numerical simulations of KdV with $\epsilon = 0.05$ in Figure 20.

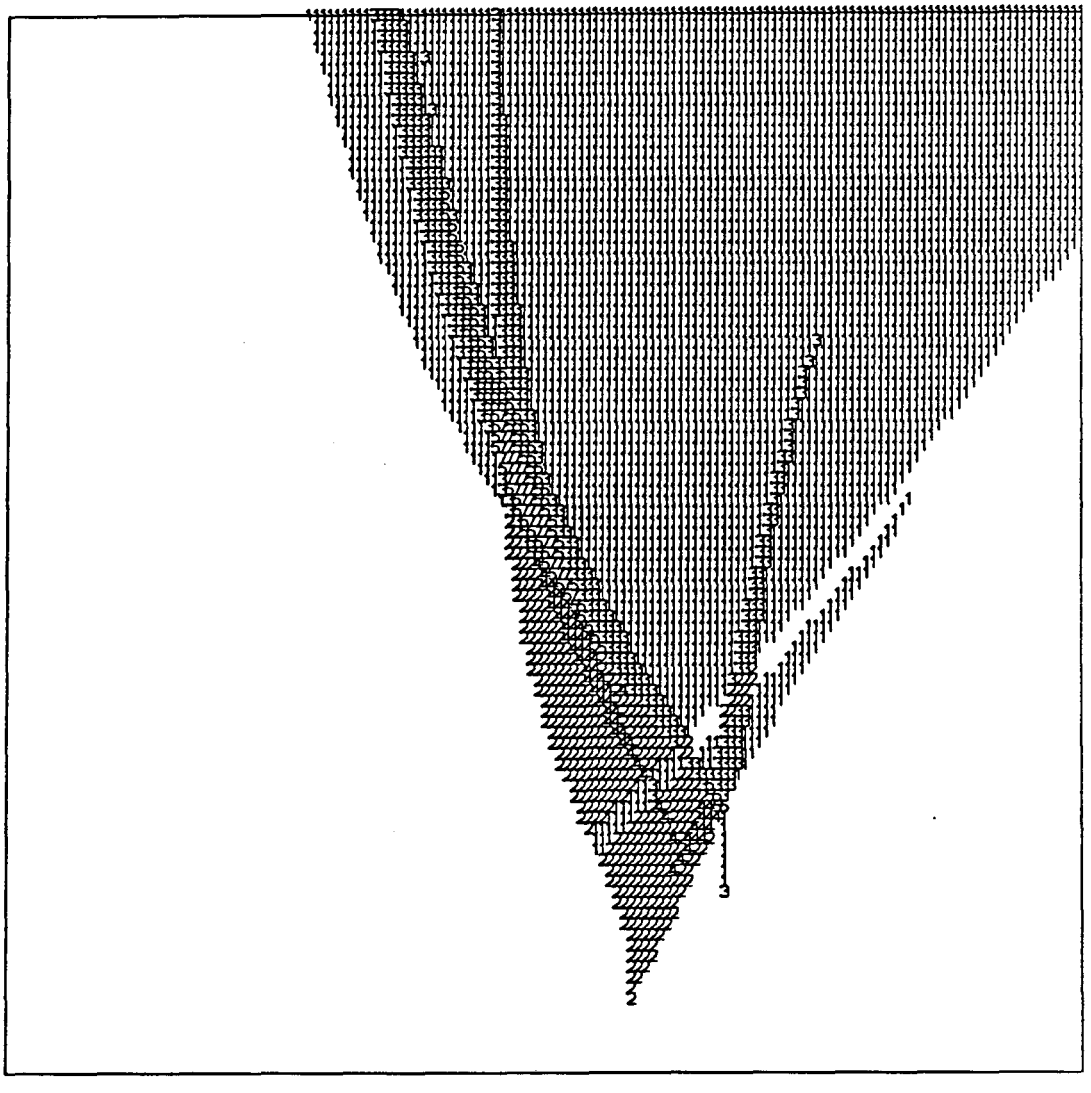

Figure 17: Space-time diagram for Example C. The domain is $[-1, 2] \times [0, 0.5]$.

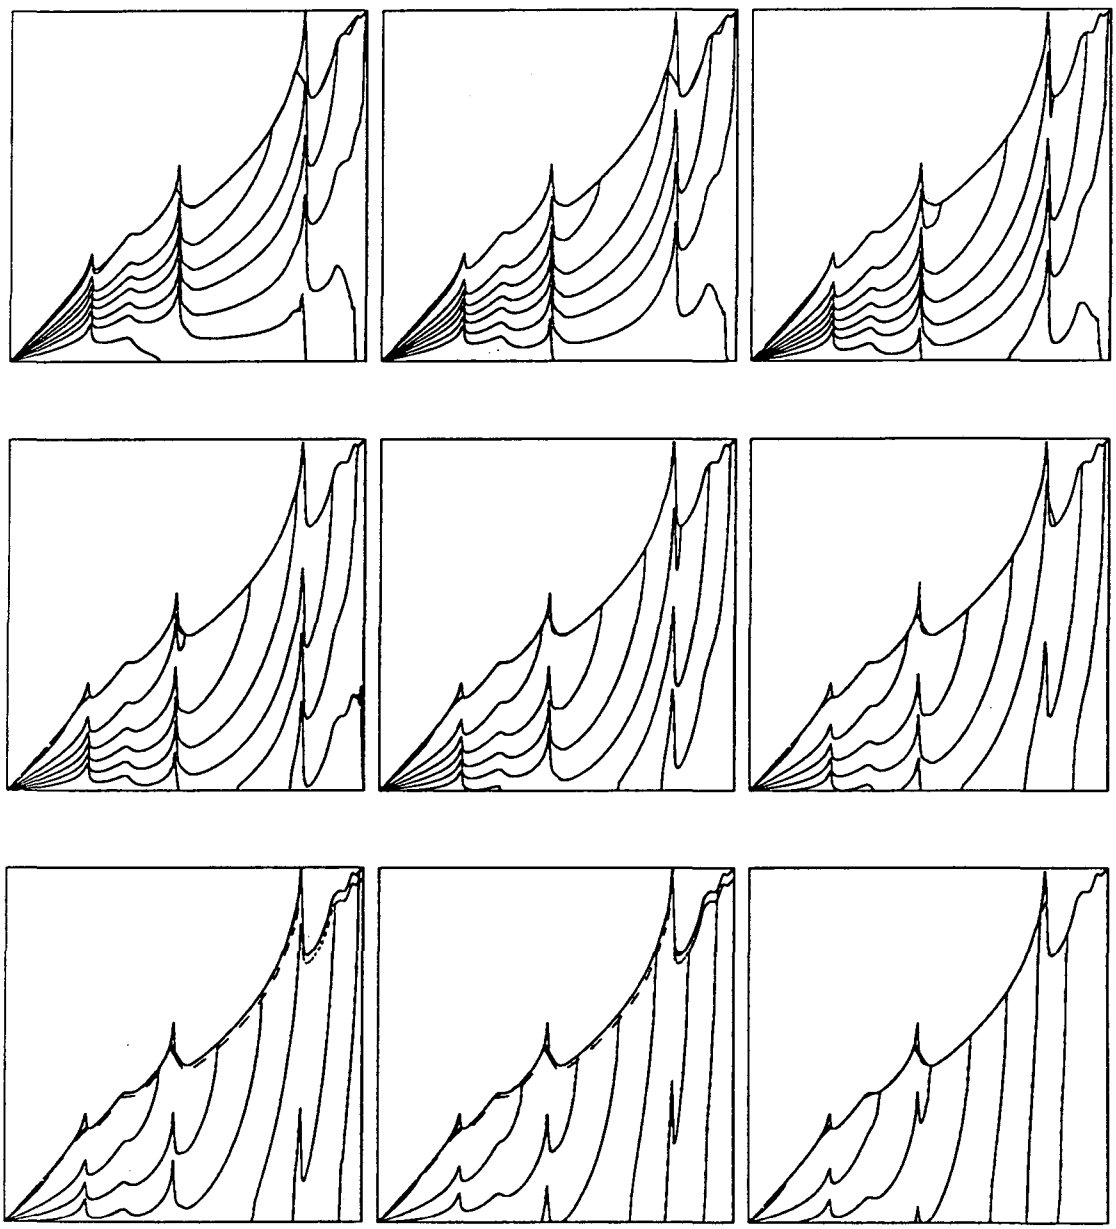

Figure 18: The minimizer f^* for Example C, shown for $t = 0$ to 0.3 in steps of 0.05 and for $t = 0.3$ to $t = 0.5$ in steps of 0.1, from left to right and top to bottom. The domain of each plot is $[-1, 2]$.

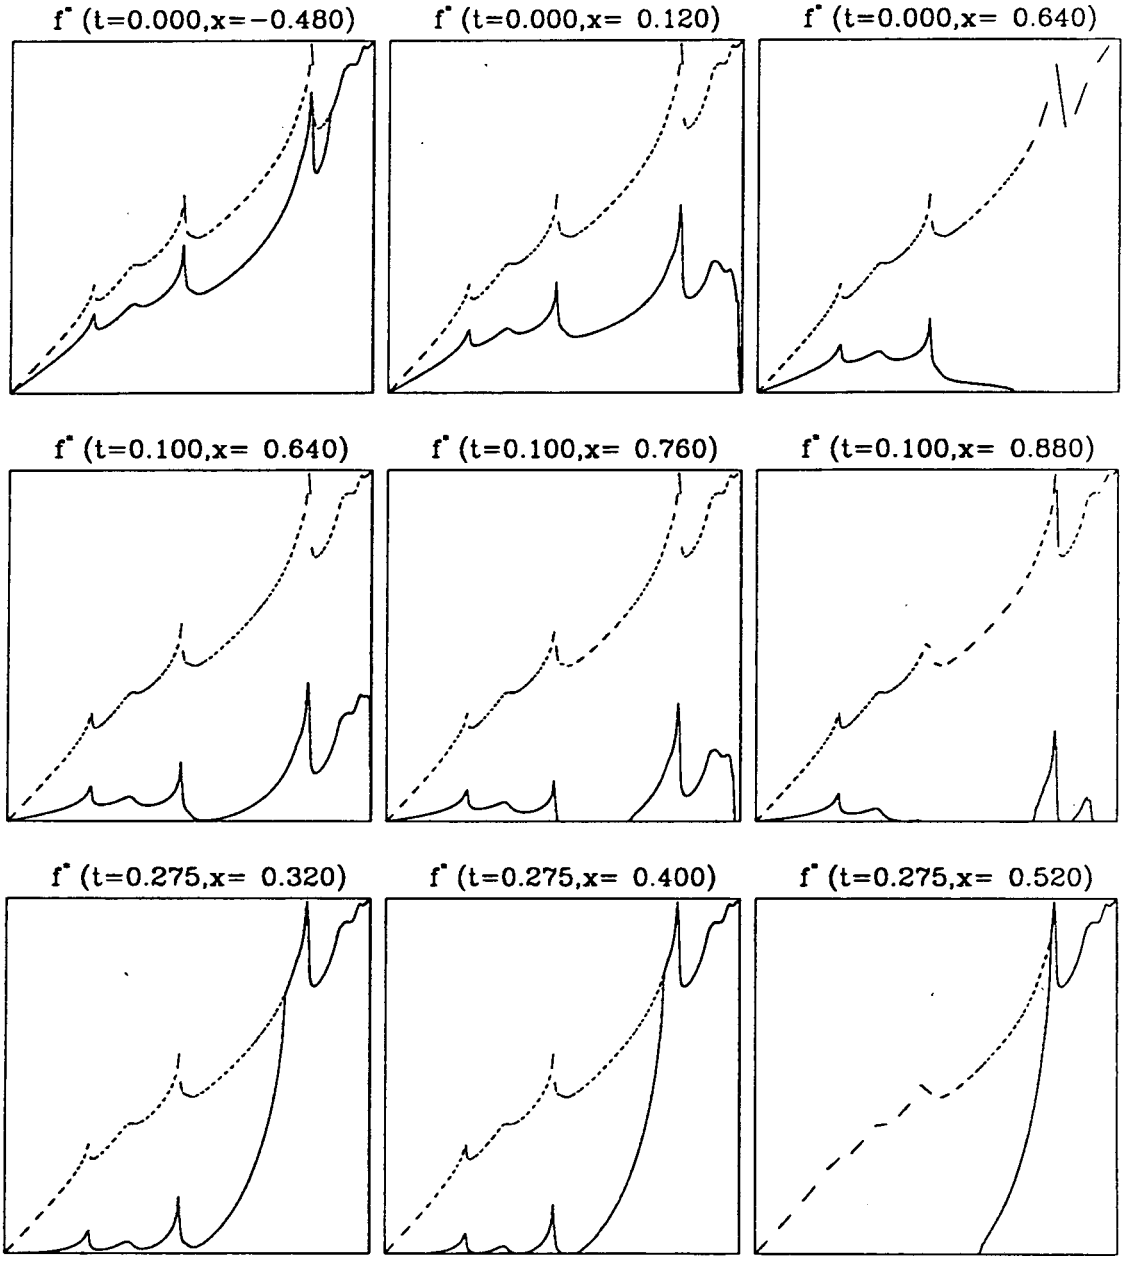

Figure 18: (continued) Selected minimizers from Example C. Note particularly the multiple breaking sequence in the bottom row.

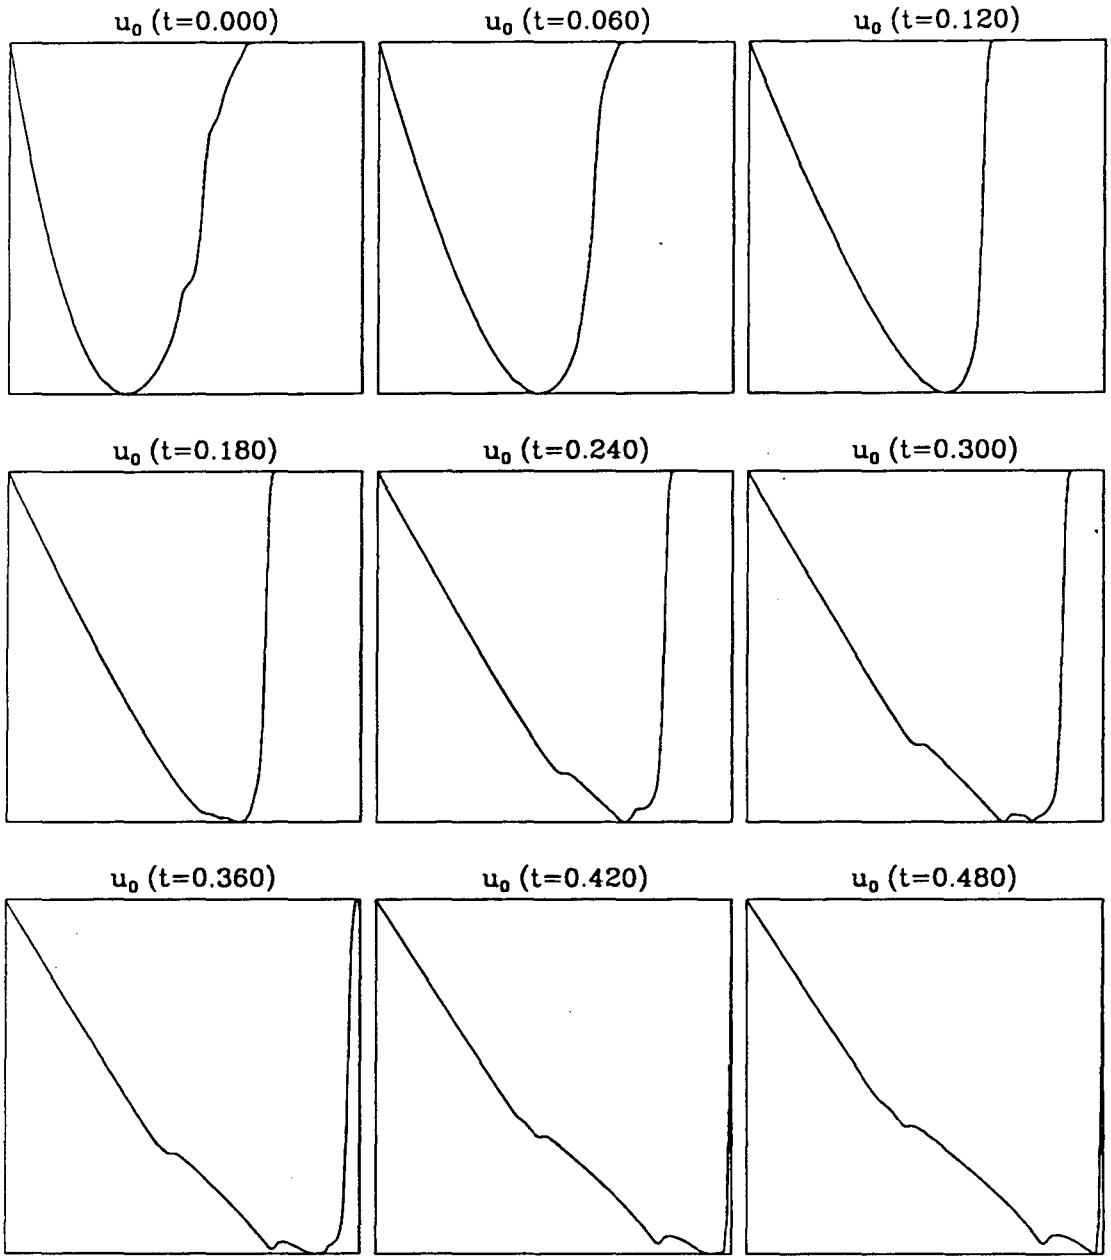

Figure 19: The weak limit $\bar{u}$ for selected values of t in Example C, computed with third-order smoothed ENO differencing. The domain of each plot is $[-1, 2]$.

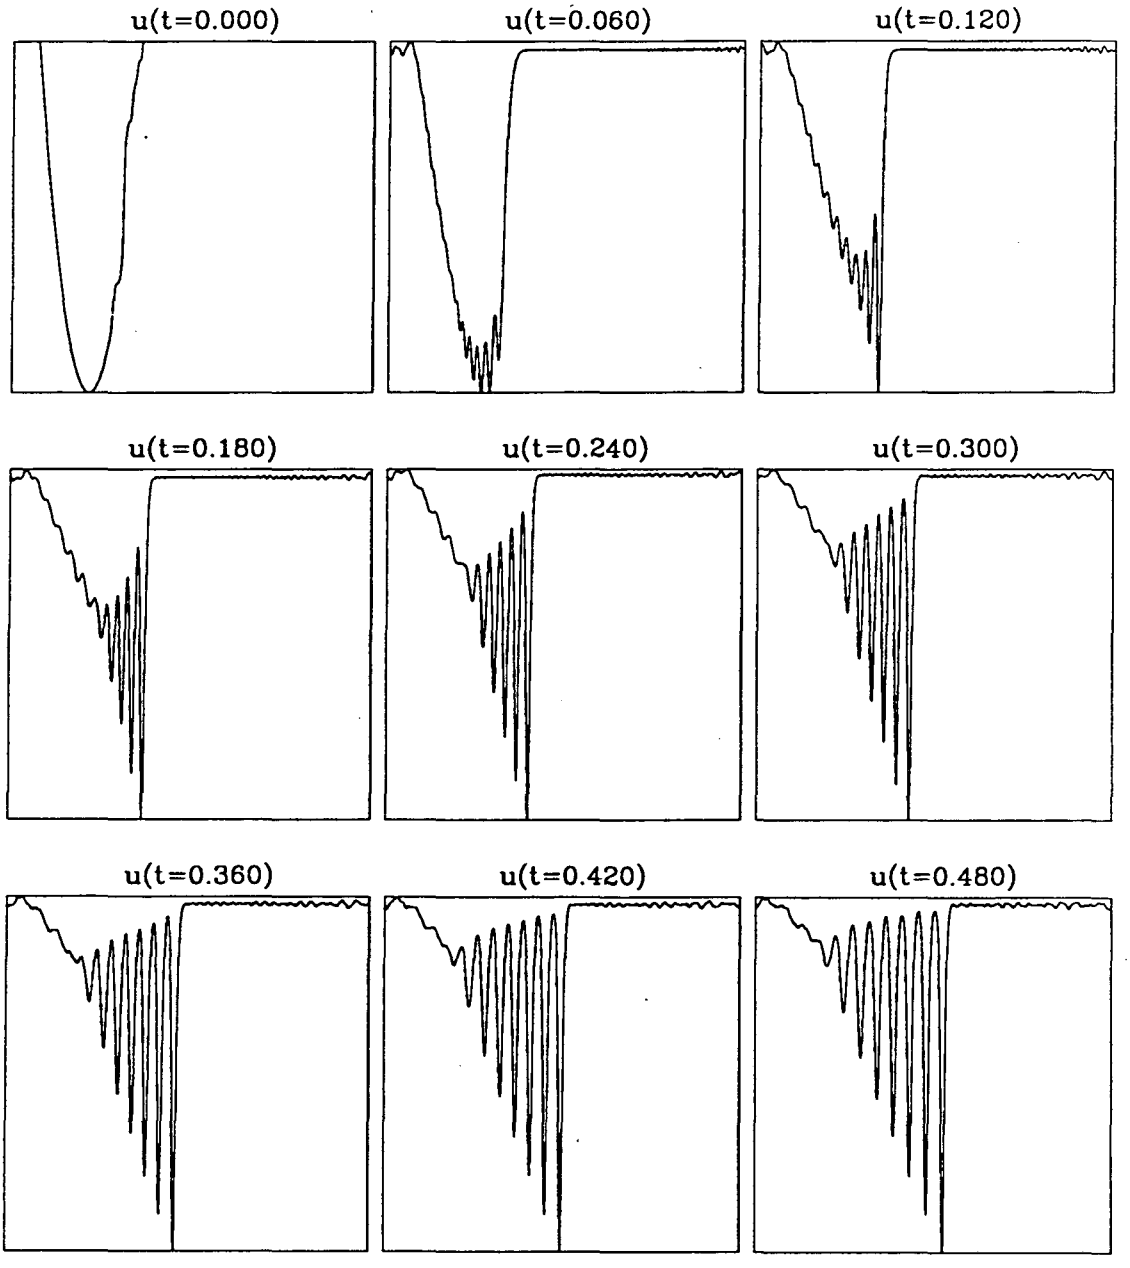

Figure 20: Solution of KdV equation for Example C with $\epsilon = 0.05$, on $[-1, 2]$.

j	x_j	$v(x_j)$	$v'(x_j)$	$v''(x_j)$
1	-1	0	-2	2
2	0	-1	0	2
3	0.1	-0.99	0.2	2
4	0.2	-0.96	0.4	2
5	0.3	-0.91	0.7	2
6	0.4	-0.84	0.8	2
7	0.5	-0.69	0.2	0
8	0.6	-0.64	1.3	2
9	0.7	-0.25	1	0
10	0.8	-0.18	1.8	2
11	0.9	-0.08	1.2	0
12	1	0	2	2

Table 6: Breakpoints, values and derivatives for the piecewise quintic $v(x)$ of Example C.

6 Conclusion

The generation and propagation of oscillations is an important natural phenomenon for nonlinear conservative waves. In the presence of small dispersion, these oscillations can be sufficiently intense to prevent the existence of a strong limit as the dispersion vanishes. In general nonlinear wave situations, it is difficult to develop a numerical procedure which accurately captures the weak limit by averaging over these oscillations [EH89]. On the other hand, for special integrable nonlinear waves, the Lax-Levermore formulation provides a natural characterization of the weak limit which reveals detailed structure far beyond that currently accessible to direct computational methods. Analytically, it is somewhat difficult to extract detailed quantitative information from the Lax-Levermore characterization. However, as a quadratic minimization problem, the Lax-Levermore characterization of the weak limit is amenable to numerical study. In this article we have developed an algorithm for this study for the KdV equation. This algorithm should be useful (1) for understanding the structure of the weak limit of the special integrable KdV problem, and (2) as a means to calibrate algorithms which compute the weak limit for more general problems.

The algorithm, as described in the text, is surprisingly efficient when properly programmed. In addition, since the quadratic minimization problem has *no coupling* between different space-time points, it is natural for massively parallel computation.

Even in these initial studies of the weak limit of the KdV equation, the algorithm enables us to investigate phenomena beyond those currently accessible to analytical methods. In particular, we mention multiple breaking and the associated structure of $\bar{u}$ in this regard. Our numerical experiments clearly demonstrate that different phases, which live in sharply defined regions of space time, evolve and interact in interesting ways. Thus distinct phases, which result from microscopic oscillations, have observable macroscopic consequences in the profile of $\bar{u}$. For such KdV studies, the quadratic minimization problem seems more efficient than a direct study of Whitham's equations, particularly in transition regions where the number of phases changes. Computing the weak limit by numerical averaging seems even more difficult, especially if one wishes to resolve the macroscopic consequences of the multiphase microstructure.

The algorithm is sufficiently efficient to consider a study of more interest-

ing initial data – data with many minima and even random data. For such studies, one would presumably need to implement the algorithm on a parallel machine, and consider ways to further improve its efficiency as well. For example, one might make use of its simple structure to implement wavelet technology [ABCR90]. These algorithms might reduce the number of operations required to solve the minimization problem with N points from $O(N^3)$ to $O(N)$. However, it is unclear whether they would actually be faster in practical situations. Also, better ways to handle the singular functions which appear could certainly be developed.

The variational methods implemented here can certainly be adapted to other soliton equations such as the Toda lattice and the (defocusing) nonlinear Schrödinger equation. However, they seem intrinsically restricted to integrable equations. It would be really exciting if a similar variational approach could be invented for the numerical study of oscillations in more general nonlinear waves!

References

[ABCR90] B. Alpert, G. Beylkin, R. Coifman, and V. Rokhlin. Wavelets for the fast solution of second-kind integral equations. Research Report YALEU/DCS/RR-837, Yale University Department of Computer Science, December 1990.

[BDK86] J. L. Bona, V. A. Dougalis, and O. A. Karakashian. Fully discrete Galerkin methods for the Korteweg-deVries equation. *Comp. and Maths with Appls.*, 12A:859–884, 1986.

[CK85] T. F. Chan and T. Kerkhoven. Fourier methods with extended stability intervals for the Korteweg-deVries equation. *SIAM Jour. Num. Analysis*, 22:441–454, 1985.

[EH89] B. Engquist and T. Y. Hou. Particle method approximation of oscillatory solutions to hyperbolic differential equations. *SIAM Jour. Num. Analysis*, 26:289–319, 1989.

[ELS91] N. Ercolani, D. Levermore, and D. Serre, editors. *Lyons Conference*, 1991.

[FFM80] H. Flaschka, M. G. Forest, and D. W. McLaughlin. Multiphase averaging and the inverse spectral solution of the Korteweg-de Vries equation. *Comm. Pure Appl. Math.*, 33:739–784, 1980.

[Fle80] R. Fletcher. *Practical Methods of Optimization.* John Wiley and Sons, 1980.

[FW78] B. Fornberg and G. B. Whitham. A numerical and theoretical study of certain nonlinear wave phenomena. *Phil. Trans. Roy. Soc. Lond. ser. A*, 289:373–404, 1978.

[GGKM67] C. S. Gardner, J. M. Greene, M. D. Kruskal, and R. M. Miura. Method for solving the Korteweg-deVries equation. *Phys. Rev. Lett.*, 19:1095–1097, 1967.

[GP74] A. V. Gurevich and L. P. Pitaevskii. Nonstationary structure of a collisionless shock wave. *Sov Phys JETP*, 38(2):291–297, 1974.

[HOEC87] A. Harten, S. Osher, B. Engquist, and S. R. Chakravarthy. Uniformly high order accurate essentially non-oscillatory schemes. *J. Comput. Phys.*, 71:231–257, 1987.

[JLM91] S. Jin, C. D. Levermore, and D. W. McLaughlin. Semi-classical limits of defocusing NLS. *Preprint*, 1991.

[KM56] I. Kay and H. E. Moses. Reflectionless transmission through dielectrics and scattering potentials. *J. Appl. Phys.*, 27:1503–1508, 1956.

[LL83] P. D. Lax and C. D. Levermore. The small dispersion limit of the Korteweg-de Vries equation I, II, III. *Comm. Pure Appl. Math.*, 36:253–290,571–593,809–830, 1983.

[LLV92] P. D. Lax, C. D. Levermore, and S. Venakides. The generation and propagation of oscillations in dispersive IVPs and their limiting behavior. In T. Fokas and V. E. Zakharov, editors, *Important developments in soliton theory 1980-1990.* Springer-Verlag, Berlin, 1992.

[Mer70] E. Merzbacher. *Quantum Mechanics.* John Wiley and Sons, 1970.

[PdKUK83] R. Piessens, E. deDoncker Kapenga, C. Uberhuber, and D. Kahaner. *Quadpack: a subroutine package for automatic integration.* Springer-Verlag, 1983.

[PFT86] H. Press, W. Flannery, and S. Teukolsky. *Numerical recipes: the art of scientific computing.* Cambridge University Press, 1986.

[Tia92a] F. R. Tian, 1992. Personal communication.

[Tia92b] F. R. Tian. Oscillations of the zero-dispersion limit of the Korteweg-deVries equation. *Comm. Pure and Appl. Math.*, to appear, 1992.

[Tsa85] S. P. Tsarev. Poisson brackets and one-dimensional hamiltonian systems of hydrodynamic type. *Soviet Math. Dokl*, 31:488–491, 1985.

[Ven85] S. Venakides. The zero dispersion limit of the Korteweg-deVries equation with nontrivial reflection coefficient. *Comm. Pure Appl. Math.*, 38:125–155, 1985.

[Whi65] G. B. Whitham. Non-linear dispersive waves. *Phil. Trans. Roy. Soc. Lond. ser. A*, 139:283–291, 1965.

[Whi74] G. B. Whitham. *Linear and Nonlinear Waves.* Wiley-Interscience, 1974.

[Wri91] O. C. Wright. *Korteweg-De Vries Zero Dispersion Limit: A Restricted Initial Value Problem.* PhD thesis, Princeton University, 1991.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
TECHNICAL INFORMATION DEPARTMENT
BERKELEY, CALIFORNIA 94720