Mat1062: Introductory Numerical Methods for PDE

Mary Pugh

January 6, 2009

1 Ownership

These notes are the joint property of Rob Almgren and Mary Pugh.

2 General classification

A partial differential equation (PDE) is an equation expressing a relationship among partial derivatives of an unknown function. If the function is u (a scalar or a vector), and it depends on variables $x, y, \ldots$ (it must depend on more than one variable to have *partial* derivatives), then the most general PDE would have the form

$$F\big(x, y, \ldots, u, u_x, u_y, \ldots, u_{xx}, u_{xy}, u_{yy}, \ldots, u_{xxx}, \ldots\big) = 0 \qquad (1)$$

where the arguments include all derivatives of u up to some finite order.

In addition to the PDE (1), we need to specify a *domain* $\Omega \subset \mathbb{R}^n$, where n is the number of independent variables, and *boundary conditions* that u and some appropriate combination of its derivatives should satisfy on the boundary of Ω. If one of the variables is time, then some part of the boundary conditions may be called *initial* or *terminal conditions.*

Example: Robert Merton (1969) considered the problem of "Lifetime portfolio selection under uncertainty." The independent variables are t, representing time throughout an person's life, and w, representing his or her total wealth. Then the optimal investment/consumption strategy is determined by the "utility" $u(w, t)$ which satisfies the PDE

$$u_t - \rho u + r w u_w + \frac{1-p}{p}\frac{1}{u_w^{p/(1-p)}} - \frac{\beta^2}{2}\frac{u_w^2}{u_{ww}} = 0,$$

on the domain $w > 0$ and $t < T$ where T is time of death (assumed known). Terminal conditions are given at $t = T$. The parameters ρ, r, etc. represent various financial coefficients.[1]

This PDE is highly nonlinear, since the first derivative u_w appears with various exponents, and the second derivative u_{ww} appears in a denominator; in addition, the solution has singular behavior near the boundaries. There are no good numerical methods for solving such an equation; Merton obtained analytical solutions by exploiting a scaling symmetry and reducing the problem to an ODE.

2.1 Simple forms

Fortunately, most problems arising in physical sciences (and even in finance) do not have such horrible structures. We will therefore define a more restricted class of problems than the general (1).

A PDE is *quasilinear* if the highest derivatives appear linearly. For the simplest case of two independent variables x, y (one of which may represent time), and restricting attention to second-order PDEs (at most two derivatives), a quasilinear equation may be written

$$\begin{aligned} a(x,y,u,u_x,u_y)\,u_{xx} \;+\; b(x,y,u,u_x,u_y)\,u_{xy} \;+\; c(x,y,u,u_x,u_y)\,u_{yy} \\ = \; F(x,y,u,u_x,u_y) \end{aligned}$$

where $a(\cdot), b(\cdot), c(\cdot), F(\cdot)$ may be arbitrary nonlinear functions.

It is *semilinear* if it is quasilinear, and the coefficient functions on the highest derivatives do not involve u or its lower derivatives:

$$a(x,y)\,u_{xx} \;+\; b(x,y)\,u_{xy} \;+\; c(x,y)\,u_{yy} \;=\; F(x,y,u,u_x,u_y)$$

It is *linear* if u appears only linearly:

$$a(x,y)\,u_{xx} \;+\; b(x,y)\,u_{xy} \;+\; c(x,y)\,u_{yy} \;=\; F(x,y).$$

A linear equation has the property that if u_1 and u_2 are solutions, then $u_1 + u_2$ is also a solution. Also, if μ is a number then μu_1 is also a solution.

---

[1] ρ is a utility discount rate, r is a riskless interest rate, p is a utility exponent, and β is a risk-adjusted excess return rate.

The qualitative behavior of a PDE is determined by how its highest derivatives enter the equation. For a second-order, quasilinear equation the coefficients on the second derivatives define a quadratic form

$$G(p,q) \;=\; a\,p^2 + b\,pq + c\,q^2 \;=\; \begin{pmatrix} p & q \end{pmatrix} Q \begin{pmatrix} p \\ q \end{pmatrix}, \qquad Q = \begin{pmatrix} a & \frac{1}{2}b \\ \frac{1}{2}b & c \end{pmatrix}$$

where p, q are only dummy variables. It is then natural to classify the equation depending on the type of the matrix Q:

- The PDE is *elliptic* if Q is positive definite or negative definite: $G(p,q)$ has the same sign for all nonzero p, q. The paradigmatic example of a linear elliptic equation is *Poisson's equation* with $G = I$:

  $$u_{xx} + u_{yy} \;=\; F(x,y)$$

  If $F = 0$ then it is *Laplace's equation.* Elliptic equations generally do not have a time variable, and come from equilibrium problems.

- The PDE is *hyperbolic* if Q is indefinite: $G(p,q)$ takes opposite signs for different values of p, q. The classic example is the *wave equation*

  $$u_{yy} - u_{xx} \;=\; 0.$$

  In this case, y would represent time. Hyperbolic equations tend to be associated with transport phenomena.

- The PDE is *parabolic* in the in-between case that Q is semidefinite: $G(p,q)$ is always either ≥ 0 or always ≤ 0, but does take zero values. The classic example is the *diffusion equation*

  $$u_y \;=\; u_{xx}.$$

  Again, y would represent time. Note that the first-order term u_y comes from the right-hand side $F(x, y, u, u_x, u_y)$.

These define broad categories that we often apply qualitatively to more general problems.

A vaguer but broader classification is *time-dependent* vs. *stationary*:

- Time-dependent problems are conceptually similar to ODEs in lots of dimensions (infinitely many). The time derivative or derivatives are generally determined as an explicit function of the space derivatives, and the problem can be marched forward in time. In addition to the accuracy of the discrete formulation, the main issue is *stability.*

- Stationary problems must be solved across the entire domain at once. The main issue in solving them is how to solve the large systems of simultaneous equations that result from the discretization. Problems of this form often appear as parts of other problems; for example, solving an implicit discretization of a parabolic equation is equivalent to solving an elliptic problem at each step.

Here are a few more examples of PDEs. They are single equations or systems, depending on whether the unknown u has one or several components.

Reaction-diffusion $u(\vec{x},t)$ is a real-valued function of space representing concentration of something or another. It solves the semilinear parabolic equation

$$u_t = \Delta u + f(u)$$

where $f(u)$ is a real-valued function giving the reaction rates in terms of local concentrations, and the Laplacian operator $\Delta = \nabla^2 = \partial_{x_1x_1} + \ldots \partial_{x_nx_n}$ if $\vec{x} \in \mathbb{R}^n$. In two and three dimensions, one would write $\Delta = \partial_{xx} + \partial_{yy}$ (2-D) or $\partial_{xx} + \partial_{yy} + \partial_{zz}$ (3-D).

Reaction diffusion equations can be written for more than one species. In this case, one would have $\vec{u}(\vec{x},t)$ which would be a real-valued vector; each component representing the concentration of a particular species. Their evolution would be modelled by a coupled system of reaction diffusion equations:

$$u_{j_t} = D_j \Delta u_j + f_j(\vec{u}) \qquad 1 \le j \le n. \tag{2}$$

Note that the reaction rate f_j can depend on all of the species and that each species can have a different diffusion rate D_j. Such reaction-diffusion systems can have very interesting dynamics.

Incompressible fluid dynamics The *incompressible Euler equations* for velocity $\vec{u}$ and pressure p are

$$\begin{aligned} \vec{u}_t + \vec{u} \cdot \nabla \vec{u} &= -\nabla p \\ \operatorname{div} \vec{u} &= 0. \end{aligned}$$

In 3-D, there are four components in all: $\vec{u} = (u_1, u_2, u_3)$ and p. Although this system does not fall within our classification above, we say that it is partly hyperbolic, by virtue of the transport behavior, and is partly elliptic, because of the need to solve for the pressure field (which is in equilibrium to acoustic waves). The appearance of the time derivative is deceptive: this is much more complicated system than a simple evolution problem.

The *Navier-Stokes equations* add a diffusion term:

$$\begin{aligned} \vec{u}_t + \vec{u} \cdot \nabla \vec{u} &= -\nabla p + \nu \Delta \vec{u} \\ \operatorname{div} \vec{u} &= 0, \end{aligned}$$

which makes the system parabolic as well.

Compressible fluid dynamics The *compressible Euler equations* are the first-order quasilinear conservation law

$$\vec{u}_t + \operatorname{div}\big(\vec{F}(\vec{u})\big) = 0, \tag{3}$$

where $\vec{F}(\vec{u})$ is a vector-valued function of the solution vector $\vec{u}$, whose components include velocity (three components if in 3-D), pressure, density, temperature, etc. With appropriate assumptions on $F(\cdot)$, these are hyperbolic problems even though they are first-order. As an example, in one dimension ($\operatorname{div} = \partial_x$), for a scalar u, taking the simplest nonlinear function $F(u) = \frac{1}{2}u^2$ gives the *inviscid Burgers' equation*

$$u_t + u u_x = 0.$$

We can also add a diffusion term to the right-hand side of (3) obtain the *compressible Navier-Stokes equation*, which has parabolic nature as well as hyperbolic.

Optimal glider flying One model of the optimal strategy for a glider aircraft flying in uncertain lift constructs a value function $u(x, z, \ell)$, where u is expected time to reach the goal, x is distance from the goal, z is altitude, and ℓ is local lift. This function solves the system of partial differential inequalities

$$\begin{aligned} u_x + H(u_z, \ell) + \frac{\ell}{\xi} u_\ell &\le \frac{\bar{\ell}^2}{\xi} u_{\ell\ell} \\ u_z(x, \ell, z) &\ge -\frac{1}{\ell - s_{\min}}, \end{aligned}$$

where $H()$ is a known function incorporating the glide performance of the aircraft, $\bar{\ell}$ is a root-mean-square lift strength, and ξ is a correlation length. Each inequality is an equality on part of the domain, and part of the problem is to find the dividing line between the two regions.

More examples can be found everywhere. I have given these just to emphasize that the classification elliptic/parabolic/hyperbolic must be interpreted qualitatively as well as quantitatively. A vast amount of theoretical work goes into classifying the behavior of solutions to such equations; in this course we shall discuss only the simplest phenomena and methods. We shall begin with parabolic equations to illustrate consistency, convergence, and stability, continue with elliptic equations to discuss solving large linear systems, and finally consider hyperbolic equations and systems, where the conservation properties of the physical problem must be reflected in the discretization.

3 Parabolic equations

Physical derivation

Suppose $u(x,t)$ is temperature of some medium, in one or more space dimensions. The internal energy per unit volume is some function $H(u)$ of the temperature. Suppose F is the vector flux of energy through the material, driven primarily by thermal gradients. By conservation of energy, the total energy inside Ω changes only due to flux in or out across the boundaries. If c is the specific heat per unit volume, then this may be written

$$\frac{d}{dt}\int_\Omega H\big(u(x,t)\big)\,dx \;=\; -\int_{\partial\Omega} F\cdot\vec{n}\,dA$$

where $\vec{n}$ is the unit outward normal, $\partial\Omega$ denotes the boundary surface of Ω, and dA is surface area. Integrating by parts (also known as Green's formula), we may write this as

$$\int_\Omega c\,u_t\,dV \;=\; -\int_\Omega \operatorname{div}(F)\,dV$$

where $c = dH/du$ is the specific heat per unit volume. Since this is true for every volume Ω, if the integrands are continuous then they must be equal pointwise, and we obtain the differential equation

$$cu_t + \operatorname{div}(F) = 0$$

which is satisfied at every point in Ω.

To complete the model, we must specify how F depends on u. In the conservation law (3) above, we assumed that F was a function of u itself, giving a first-order system. Now we shall let F depend on derivatives of u. The simplest model is the linear Fourier law, in which the flux F is simply proportional to the temperature gradient ∇u:

$$F(x, t, u, \nabla u) = -\kappa(x, t, u)\,\nabla u$$

where the coefficient of thermal conductivity κ in principle may depend on position, time, and the local density itself. We thus obtain the quasilinear parabolic equation

$$cu_t = \operatorname{div}\big(\kappa(x, t, u)\nabla u\big). \tag{4}$$

If c and κ are constant in space and time, and independent of u itself, then the nonlinear equation (4) reduces to the linear equation

$$u_t = D\,\Delta u, \tag{5}$$

where $D = \kappa/c$ is the *diffusivity*. The units of D are length^2/time. The same equation arises from other physical models; for example, u could be a chemical potential with a driving particle flux F which changes the particle density $H(u)$. In one dimension our model is

$$u_t = D\,u_{xx}. \tag{6}$$

We shall spend several weeks talking about numerical methods for equations (5,6). But before we begin, let us outline a few of their mathematical properties.

Properties of the diffusion equation

For simplicity, we focus on the one-dimensional version (6).

Conservation For any fixed interval $[a, b]$, a smooth solution of (6) will satisfy

$$\frac{d}{dt}\int_a^b u(x,t)\,dx \;=\; \int_a^b D u_{xx}(x,t)\,dx \;=\; D u_x(x,t)\Big|_{x=a}^{b}. \tag{7}$$

This does not say that "stuff" is conserved in the sense that it never changes, but that we can account for the total change in an interval by looking at what crosses the endpoints. Stuff doesn't appear or disappear from the interior because there are no sources or sinks in equation (6).

The analogue for the diffusion equation in higher dimensions (5) is as follows. Let $\Omega \in \mathbb{R}^n$ be a bounded set with a reasonably smooth boundary. Then

$$\begin{aligned}\frac{d}{dt}\int_\Omega u(\vec{x},t)\,d\vec{x} \;&= D\int_\Omega \nabla\cdot\nabla u(\vec{x},t)\,d\vec{x} = D\int_{\partial\Omega} \nabla u(\vec{x},t)\cdot\vec{n}\;dA\\ &= D\int_{\partial\Omega} \frac{\partial u}{\partial\vec{n}}(\vec{x},t)\;dA\end{aligned}$$

where $\vec{n}$ is the unit outward normal. And so the change in mass is equal to the total flux through the boundary of Ω.

Finally, these observations apply to the nonlinear diffusion equation (4):

$$\frac{d}{dt}\int_\Omega u(\vec{x},t)\,d\vec{x} = \frac{1}{c}\int_{\partial\Omega} \kappa(\vec{x},t,u)\frac{\partial u}{\partial\vec{n}}(\vec{x},t)\;dA$$

Maximum principle Suppose $u(x,t)$ has a smooth interior local maximum, when viewed as a function of x for a fixed t. Then the second derivative $u_{xx} \le 0$. Thus the time derivative u_t is also ≤ 0, and the *local maximum cannot move upwards*. Similarly, at a local minimum $u_{xx} \ge 0$ and so $u_t \ge 0$, and the *local minimum cannot move downwards*. If M is an upper bound at time t_0 then the graph of u cannot get any higher later on

$$\max_{x\in\mathbb{R}} u(x,t_0) \le M \quad\Longrightarrow\quad \max_{x\in\mathbb{R}} u(x,t)\;\le M, \qquad \text{for } t > t_0.$$

Similarly, if m is a lower bound at time t_0 then the graph of u cannot get any lower later on

$$m \le \max_{x\in\mathbb{R}} u(x,t_0) \quad\Longrightarrow\quad m \le \max_{x\in\mathbb{R}} u(x,t), \qquad \text{for } t > t_0. \tag{8}$$

The maximum principle was written to allow the possibility of initial data which is a constant function: $u(x,t_0) = M = m$ for all $x \in R$. Because such a function is a solution of the diffusion equation, we need to allow an inequality in the above. However, if the initial data is not a constant function then one can make the stronger statements

$$\max_{x\in\mathbb{R}} u(x,t_0) \le M \quad\Longrightarrow\quad \max_{x\in\mathbb{R}} u(x,t) \ < M, \qquad \text{for } t > t_0. \tag{9}$$

$$m \le \max_{x\in\mathbb{R}} u(x,t_0) \quad\Longrightarrow\quad m < \max_{x\in\mathbb{R}} u(x,t), \qquad \text{for } t > t_0. \tag{10}$$

One can easily prove (9) and (10) by representing the solution via convolution against the Green's function

$$u(\cdot,t) = u_0 * G(\cdot, t - t_0).$$

See Subsection 3.1.1 for more about the Green's function.

Note that above we are considering the diffusion equation on all of $\mathbb{R}$, allowing us to ignore what might happen at the boundary. One has a Green's function in $\mathbb{R}^n$ and so one can make the same statements about solutions on R^n. If one is considering the diffusion equation on a bounded domain, Ω, then one cannot make statements like (9) and (10) about the extremal values because it is possible that the maximum value, say, is achieved at the boundary.

In terms of physical problems for which u is a density it follows immediately from (8) that if the initial data is everywhere nonnegative, then the solution will never become negative at a later time.

In addition, because the diffusion equation is linear and has constant coefficients it follows that any space or time derivative of u also satisfies the diffusion equation. Therefore, since the derivatives do not get larger than their initial value, the solution only becomes smoother as it evolves in time. In fact the solution instantly becomes infinitely smooth. This justifies the invocation of derivatives above.

It will be very nice if the numerical scheme can preserve these properties. But to put this principle in perspective, let us describe an alternative source for a parabolic PDE, arising in pricing of financial derivatives.

Black-Scholes equation If V is the market price of a traded contract whose value depends on the price of an underlying asset S, then, under a suitable model for the dynamics of S, $V(S,t)$ must satisfy

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r-q)\, S\, \frac{\partial V}{\partial S} - rV = 0.$$

where σ, r, q are parameters of the model for S (volatility, interest rate, and dividend yield). This equation is linear, but has nonconstant coefficients depending on the independent variable S. By changing variables to $x = \log S$, the coefficients become constant, and a few more simple changes reduce it to the form (6), in which D is a ratio of volatility to interest rate.

Although this equation has the mathematical form as the physical model above, the underlying reasoning is very different. Conservation of V has no particular significance, nor does decrease of energy; instead a principle called *absence of arbitrage* is central. This suggests that different numerical methods may be appropriate for the Black-Scholes equation than for physical diffusion problems.

3.1 Special solutions

Above we have described fundamental properties of the structure of the diffusion equation, especially as it arises in physical problems. To gain some insight into the behavior of its solutions in general, let us consider two particular explicit solutions.

3.1.1 Green's function

Consider the function

$$G(x,t) \;=\; \frac{1}{\sqrt{4\pi D t}}\, e^{-x^2/4Dt}. \tag{11}$$

For $t > 0$, no denominators vanish, so this function is perfectly well defined and has derivatives of all orders; it describes a Gaussian pulse centered at the origin, of total mass one, and of width $\sqrt{2Dt}$. (Because the solution is not compacty supported we cannot just take the length of the support as the width of the solution. Here, the "width" is the square root of variance $\int x^2 G(x,t)\,dx$.) You may verify that this function satisfies (6). (To find G, look for solutions with the similarity form $u(x,t) = U(x/\sqrt{Dt})/\sqrt{Dt}$ for an unknown function $U(\xi)$, obtain the ODE $2U''(\xi) + \xi U'(\xi) + U(\xi) = 0$, find the solution $U(\xi) = A\exp(-\xi^2/4)$, and then choose A so mass is one.)

Note that

$$G(x,0) := \lim_{t\searrow 0} G(x,t) = \begin{cases} 0 & \text{if } x \neq 0 \\ \infty & \text{if } x = 0 \end{cases}$$

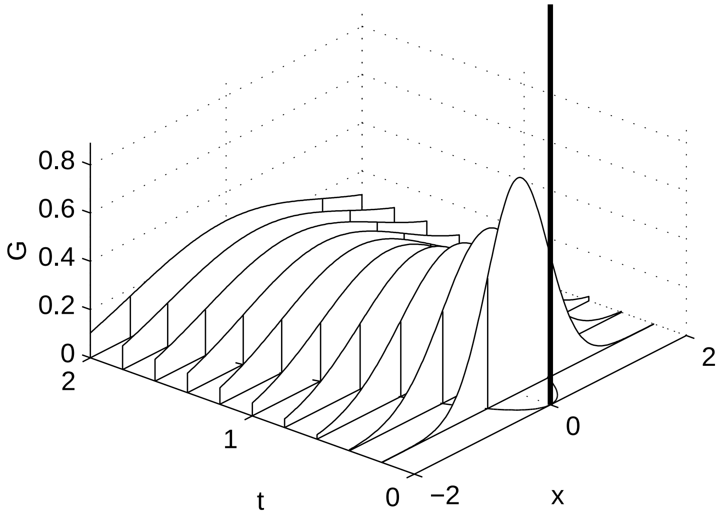

Figure 1: Green's function (11) for the diffusion equation $u_t = \frac{1}{2}u_{xx}$, showing a curve on which G is a fixed fraction of its height at the origin. The first few slices expand very rapidly away from $x = 0$, and then slow down.

is not a real-valued function. However, by extrapolating known properties of $G(x,t)$ for $t > 0$, we may say that $G(x,0)$ is some sort of unit mass, centered at the origin, but of width zero—this strange object is called a *Dirac delta distribution.* Of course, no physical object can really have zero width, so you may think of this solution as an "outer asymptotics," approximately valid on scales larger than the length scale of the initial data. I want to draw your attention to two aspects of this solution.

First, because our PDE is linear, we can solve the initial-value problem

$$u_t = Du_{xx} \quad \text{on } -\infty < x < \infty,\ t \geq 0. \qquad \text{with } u(x,0) = u_0(x),$$

by adding up the contributions from each part of the initial data to get the famous Green's formula:

$$u(x,t) = \int_{-\infty}^{\infty} G(x-y,t)\, u_0(y)\, dy.$$

Note that if $u_0(x) \geq 0$ for all x, then, $u(x,t) \geq 0$ for all (x,t) since $G > 0$. Furthermore, if $u_0 \geq 0$ is continuous and if $u_0 > 0$ *somewhere*, then

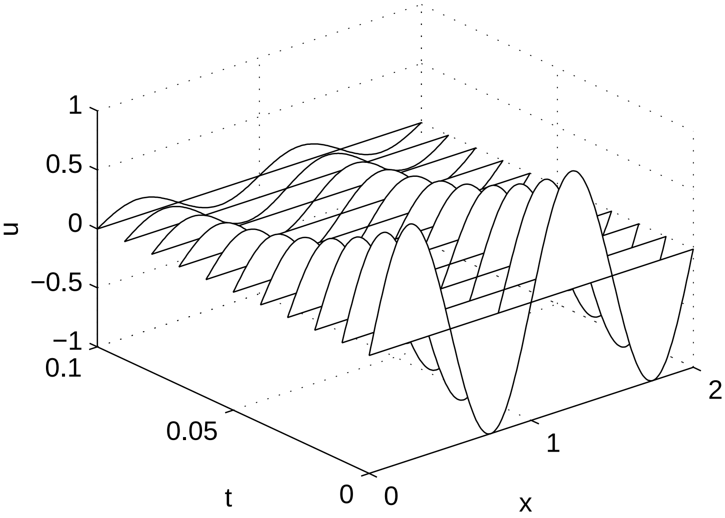

Figure 2: Periodic solutions (12) for the diffusion equation, with $k = 2\pi$ so the spatial period is 1. The time decay constant is $k^2/2 \approx 20$, so the decay is extremely rapid (note the time scale of the plot). By $t = 0.1$ the amplitude is only about 1/7 of its initial value.

$u(x,t) > 0$ for all x and for all $t > 0$ *no matter how small* t *is.* Thus initial data which is positive in a bounded set and zero elsewhere results in a solution which is instantaneously positive everywhere — the "speed of propagation" is infinite. However, for small t, the largest contributions to u at x come from values $u_0(y)$ for y near x.

Second, in (11), x appears only in the combination $x/\sqrt{Dt}$. This means that if we look at the physical location of any qualitative aspect of the solution—for example, the boundaries between which half of the solution mass is contained—then we will see x scaling proportional to $\sqrt{Dt}$. For small t, this x will move very rapidly, but for large t, the speed of advance will be slower and slower.

3.1.2 Eigenfunctions

The second class of special solutions takes sinusoidal initial data $u_0(x) = \sin kx$, where k is any real number. You may easily verify that the solution with this initial data is

$$u(x,t) \;=\; e^{-\sigma t}\sin kx, \qquad \sigma = Dk^2. \tag{12}$$

The formula $\sigma = Dk^2$ is called the *dispersion relation:* it gives us the time dependence of wave-like solutions in terms of their spatial wave number k. For any positive value of D, σ becomes very large as k increases, meaning that *high-frequency waves damp fast.* Describing this behavior in terms of wavelengths and periods rather than wave numbers and decay rates, the wavelength $\Lambda = 2\pi/k$ and the decay time $T = 1/\sigma$, so $\Lambda = 2\pi\sqrt{DT}$. This is the same length $\propto \sqrt{\text{time}}$ that we observed in the Green's function.

Since the PDE is linear, a general solution may be considered as a superposition of all these periodic waves, and we conclude that small-scale features of the initial data are very rapidly smoothed out.