

Notes for Introduction to Partial Differential Equations, Fall 2018

Esteban G. Tabak

1 Traffic Flow

We explore a simple mathematical model for traffic flow. This will be our introduction to systems of conservation laws, characteristics and shock waves.

There are various possible approaches to modeling traffic. First, we can distinguish between discrete and continuous models. In a discrete model, one follows individual cars, studying their position and speed as functions of time, and modeling their evolution in terms of interactions with the other cars, traffic lights, etc. In a continuous model, one looks at the road from far enough that the behavior of individual cars becomes unimportant, and deals instead with average concepts, such as the number of cars per mile and the cars' mean speed. This is analogous to the distinction, in the study of fluid motion, between the molecular and the continuous viewpoints: the former is more accurate for the description of behavior in the nano-scales, while the latter is far more practical for a macroscopic description.

A second distinction is between stochastic and deterministic models. In a stochastic model, the evolution of the traffic is not fully specified by the model, but depends on random variables with specified probability distributions, such as the probability that a car exceeds a certain speed, that a number of cars to cluster in a section of the road, etc. In a deterministic model, the evolution is dictated exclusively by the models laws.

These notes will describe traffic flow through a continuous, deterministic model. In the process, we will learn about a class of partial differential equations (hyperbolic conservation laws), and unveil some simple reasons for a number of mysteries that we observe daily on the road, such as the sudden formation of seemingly causeless traffic congestions, and their eventual dissipation into thin air.

Variables

Consider a segment of a road of uniform physical characteristics –width, state of the asphalt, etc.– and with no crossings. In describing the traffic through a road from a macroscopic perspective, a sensible choice for *independent* variables are the distance along the road, which we will denote x , and the time t .

Two *dependent* quantities of interest are the density $\rho(x, t)$, measuring the number of cars per unit length of the road at position x and time t , and the mean velocity of cars $U(x, t)$. Another quantity easier to measure than U –and, it will turn out, more fundamental from a modeling perspective– is the flux $Q(x, t)$, representing the number of cars crossing position x per unit time at time t .

These three variables are related by the identity

$$Q = \rho U, \tag{1}$$

which follows from the following considerations: The number of cars crossing a section at speed u in the time interval dt equals the number of cars with that speed in the spatial interval $dx = u dt$, given by $\rho u dt$. Hence, for a uniform car speed, (1) follows. To take the average over the various car speeds in the general case, we need to introduce the density in phase space

$$r(x, t, u),$$

so that

$$\int_{u_1}^{u_2} r(x, t, u) du$$

represents the amount of cars with speeds between u_1 and u_2 per unit length of the road at position x and time t . Then our prior density ρ is given by

$$\rho(x, t) = \int_0^\infty r(x, t, u) du$$

and the mean velocity U by

$$U = \frac{\int_0^\infty r(x, t, u) u du}{\int_0^\infty r(x, t, u) du}.$$

Similarly, we may introduce a flux $q(x, t, u)$ in phase space. Then

$$Q = \int_0^\infty q(x, t, u) du = \int_0^\infty r(x, t, u) u du = \rho U,$$

justifying (1).

Equations

Our “physics” will be limited to a simple, rather intuitive fact: that the number of cars in a segment of the road can change only driven by the difference between the number of cars entering the segment at one end and those leaving at the other. The differential form of the equation for car conservation is

$$\rho_t + Q_x = 0. \tag{2}$$

It follows from the integral identity

$$\int_{x_1}^{x_2} (\rho(x, t_2) - \rho(x, t_1)) dx = \int_{t_1}^{t_2} (Q(x_1, t) - Q(x_2, t)) dt, \quad (3)$$

which equates the change in the number of cars in a given segment of the road between times t_1 and t_2 to the difference between the cars entering and leaving the segment through its two ends at x_1 and x_2 . The derivation of (2) proceeds by invoking the fundamental theorem of calculus to rewrite (3) in the form

$$\int_{x_1}^{x_2} \int_{t_1}^{t_2} (\rho_t + Q_x) dt dx = 0.$$

Since this must hold for any values of $(x_{1,2}, t_{1,2})$, equation (2) follows whenever the integrand is continuous.

Equation (2), a single constraint for the two unknowns ρ and Q , cannot determine their evolution alone. In order to close the system, we need another constraint. The simplest one follows from the realization that the typical car speed depends on the level of traffic congestion, i.e. the car density. Turning this qualitative statement into a sharp functional relation, we may propose that

$$Q(x, t) = Q(\rho(x, t)), \quad (4)$$

where $Q(\rho)$ is a prescribed function that depends on the characteristics of the road, cars and drivers. Under this assumption, equation (2) becomes

$$\rho_t + Q'(\rho) \rho_x = 0, \quad (5)$$

a closed equation for the single variable ρ .

Since the mean speed U should clearly be a decreasing function of ρ , reaching zero at the fully congested level ρ_{max} , then the flux $Q(\rho)$, computed from (1), must be concave, with zeros at both $\rho = 0$ and $\rho = \rho_{max}$, and achieving its maximum value in between, say at $\rho = \rho^*$. Figure 1 plots one such function, corresponding to the simplest choice

$$Q = \rho(1 - \rho), \quad (6)$$

where we have picked the unit of length slightly larger than the mean length of cars, so that $\rho_{max} = 1$, and the unit of time such that the maximal speed of cars is one.

Characteristics

We will solve the partial differential equation (5) for $\rho(x, t)$ using the fact that there are certain lines in the (x, t) plane along which it becomes an ordinary differential equation. To see this, consider the time derivative of ρ along a prescribed curve $x = X(t)$:

$$\frac{d}{dt} \rho(X(t), t) = \rho_t + \frac{dX}{dt} \rho_x.$$

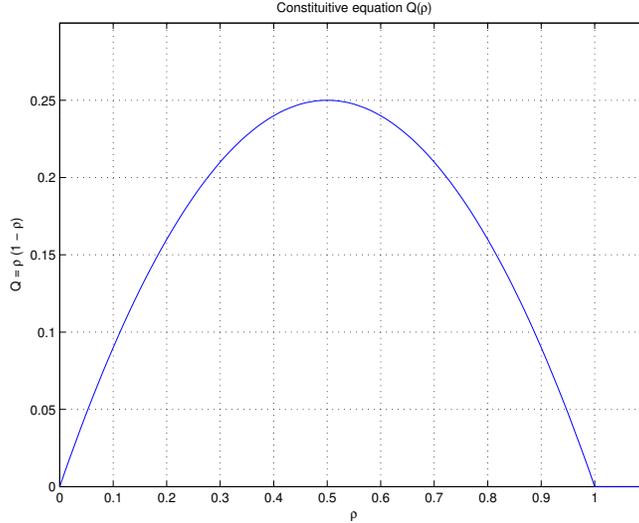


Figure 1: Simple constitutive relation for traffic flow

If we choose the curve $X(t)$ in such a way that

$$\frac{dX}{dt} = Q'(\rho(X(t), t)), \quad (7)$$

then (5) yields

$$\frac{d}{dt}\rho(X(t), t) = 0. \quad (8)$$

Equations (7, 8) are the *characteristic equations* for (5), and the family of curves $X(t, \alpha)$ satisfying (7) are its *characteristics* (Here α is any parameter, such as the initial value $X(t = 0)$, distinguishing the family members.) It follows from (8) that the density ρ is constant along characteristics, and then from (7) that the characteristic curves $X(t)$ are straight lines.

The characteristics are often described as directions along which information propagates. In our case, this information is just the value of the solution $\rho(x, t)$. It is worth noting that, for traffic flow, the characteristics are always slower than the cars:

$$\frac{dX}{dt} = \frac{dQ}{d\rho} = \frac{d(\rho U)}{d\rho} = U + \rho \frac{dU}{d\rho} \leq U,$$

since U is a decreasing function of ρ . This fact makes driving doable, since the information on the state of the traffic always comes from the cars ahead, not from those behind!

A consequence of the characteristics being straight is that the solution at any point (x, t) equals that at $(x - Q'(\rho(x, t)t, 0)$. This allows us to write the

following implicit solution to (5) subject to initial data $\rho(x, 0) = \rho_0(x)$:

$$\rho(x, t) = \rho_0(x - Q'(\rho(x, t))t) \quad (9)$$

Shocks

From the construction of characteristics in the section above, summarized in the characteristic equations (7, 8) or alternatively in the implicit solution (9), we see that the information yielding a solution $\rho(x, t)$ to (5) “travels” along straight lines in the (x, t) plane. These straight lines, however, are not generally all parallel, since $Q'(\rho)$ adopts different values for different values of ρ . Hence characteristics will eventually intersect, yielding contradictory information about the solution $\rho(x, t)$. What becomes of the solution after the first such intersection occurs?

The concavity of $Q(\rho)$ for traffic flow implies that Q' is in fact a decreasing function of ρ . Hence lower traffic densities will travel faster than higher ones. Consequently, any non-constant initial profile will deform as time progresses, with the areas of low density catching up with those of higher density ahead of them, and getting farther away from those behind them. This phenomenon is illustrated in figure 2, which displays various snap-shots of a numerical simulation of (5), with the constitutive relation given by (6) (The algorithm used will be described a few subsections below), and initial data

$$\rho_0(x) = \frac{1}{4}(1.5 + \sin(x - \pi)).$$

Notice that at approximately $t = 2$, the solution seems to deform to the point of developing a vertical slope at $x \approx 3.6$. This corresponds to the first crossing of characteristics. This breaking time and position could have been predicted from the implicit solution (9): differentiating this, one obtains

$$\rho_x(x, t) = \frac{\rho'_0(x_0)}{1 + \rho'_0(x_0) Q''(\rho_0(x_0))t},$$

with $x_0 = x - Q'(\rho(x, t))t$. Hence the slope $\rho_x(x, t)$ blows up at the minimum positive value of

$$t = \frac{-1}{\rho'_0(x_0) Q''(\rho_0(x_0))}.$$

Since Q'' is negative, blow up takes place at a location of positive slope, agreeing with the picture of a region of low density catching up with a more congested one. In our case, $Q'' = -2$ for all values of ρ , so blow up occurs at the point of maximal slope $\rho'_0 = 1/4$, achieved at $x_0 = \pi$. The corresponding time is $t = 2$, at position $x = \pi + Q'(\rho(\pi, 0)) * 2 = 3.64$, in agreement with the numerical results of figure (2).

After the blow-up time, the implicit solution given by (9) is multivalued; typically three characteristics converge at each point of the overlap region, each carrying a different value for ρ . The profile corresponding to the initial data

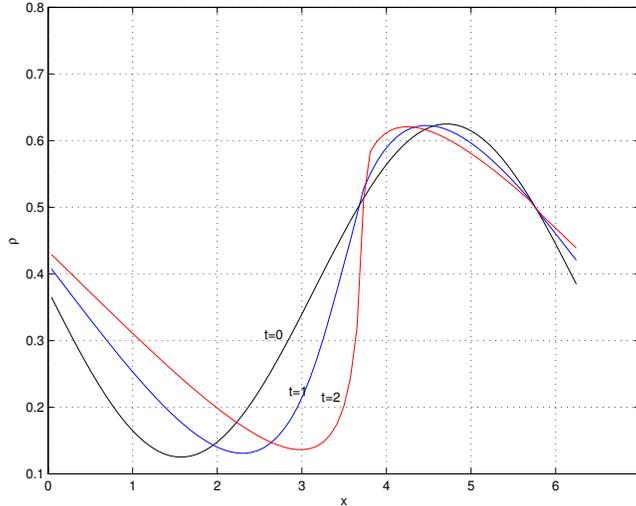


Figure 2: Snap shots of a numerical solution of the traffic flow equations, displaying nonlinear steepening up to a breaking point with vertical slope.

above, evaluated at $t = 4$, is presented in figure (3). Even though we shall see examples later where multivalued solutions of this kind make perfect physical sense, this is certainly not the case of traffic flow, where only one car density can exist at each point of the road. A possible solution is to discard two of the three solutions in the region of characteristic overlap, keeping only one. At the point where the solutions switches from a characteristic branch to another, the solution ρ will be discontinuous. An example of this construction is shown in figure (3) with a dotted line; the locus of discontinuity is called a shock wave, a name originating in gas dynamics. In traffic flow, it corresponds to a phenomenon that we frequently observe: a car driving relatively fast in light traffic is suddenly forced to slow down and enter a more congested segment of the road, with no apparent reason explaining the sharp transition between the two.

Yet it remains to determine where to position the shock, and more generally which meaning to give to a solution to (2) once it becomes discontinuous. The answer lies in the integral formulation (3). Notice that this simple-minded statement of car conservation does not depend on either ρ nor Q being continuous. Hence we may think of (3) as more fundamental than (2), and call *weak solutions* of (2) those functions, not necessarily continuous, that satisfy (3) for all choices of $x_{1,2}, t_{1,2}$.

Equation (3) can be generalized, switching from rectangles to arbitrary regions Ω in the (x, t) plane:

$$\int_{\delta\Omega} (\rho dx - Q dt) = 0, \quad (10)$$

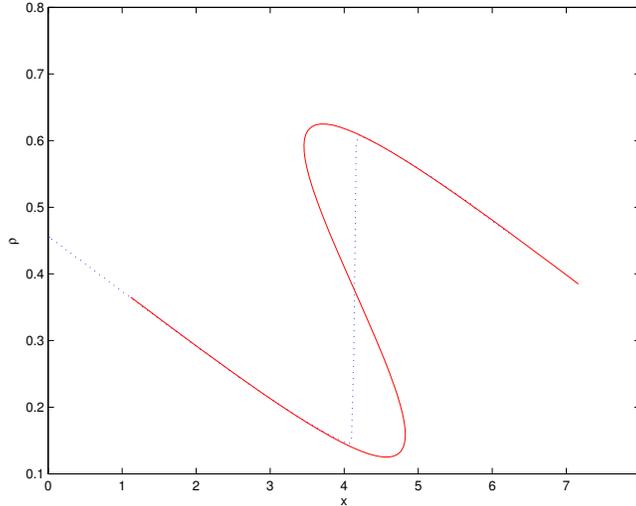


Figure 3: Multivalued solution along characteristics of the traffic flow equation, for time $t = 4$, well after breaking.

where $\delta\Omega$ is the closed boundary of the domain Ω . Equation (10) can be derived in many ways: physically, by noticing that it represents the generalized flux of cars across $\delta\Omega$, and so just states that cars are neither created nor destroyed; geometrically, by approximating the surface $\delta\Omega$ by a union of rectangles and applying (3); analytically, by invoking Green's theorem in conjunction with (2). To see that (10) really stands for car flux in (x, t) space, think of an observer moving at speed c . The number of cars that the moving observer would encounter during a small time interval Δt equals $\rho(c - U)\Delta t$, i.e.: the car density times the distance traveled with respect to the moving traffic. But this equals $\rho c \Delta t - \rho U \Delta t = \rho \Delta x - Q \Delta t$, justifying (10).

The integral formulation (10) allows us to compute the speed at which a shock wave propagates. To see this, consider the location of a shock in the (x, t) plane, and draw the boundary of a thin domain Ω , including a small enough segment of the shock that its speed $c = dx/dt$ and the values ρ^- and ρ^+ to its left and right can be approximated by constants (see figure(4)). When applied to this domain, (10) yields

$$c = \frac{dx}{dt} = \frac{Q(\rho^+) - Q(\rho^-)}{\rho^+ - \rho^-}. \quad (11)$$

This relation between the speed of the shock and the jumps of ρ and Q across it is called the Rankine–Hugoniot condition, or simply the jump condition across shocks.

Is the shock speed c all we need to compute weak solutions to (2) near shocks? The answer is positive if the characteristics on both sides of the shock converge

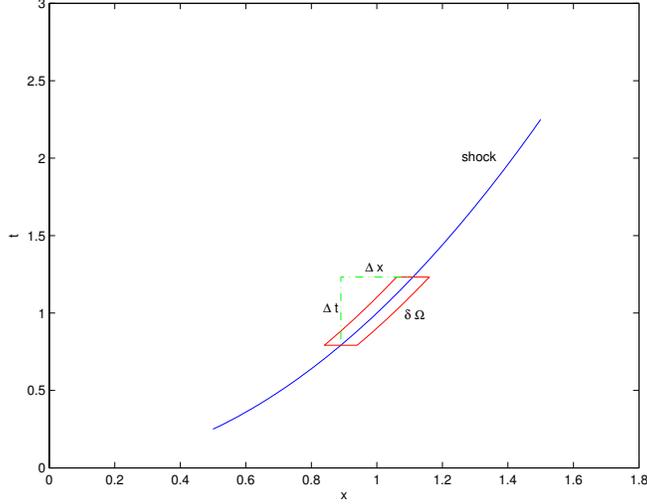


Figure 4: Control volume around a shock leading to the RH conditions

toward the shock, or at least do not emanate from it. For then the information about ρ^- and ρ^+ is updated by the characteristic flow on each smooth side of the shock, while the speed of the shock itself is given by (11). Notice that this convergence of characteristics is what gave rise to the shock to start with, with areas of lower car density catching up with more congested ones. Hence we should expect the condition

$$Q'(\rho^-) \geq \frac{Q(\rho^+) - Q(\rho^-)}{\rho^+ - \rho^-} \geq Q'(\rho^+) \quad (12)$$

to apply to realizable shocks. This is the *Lax condition* for shocks, which will re-appear below relating to issues of uniqueness and stability. The downward concavity of $Q(\rho)$ for traffic flow makes this condition equivalent to $\rho^- \leq \rho^+$, i.e., realizable shocks always have the more congested area ahead.

Notice that the jump condition (11) does not depend only on the differential equation (5), but also on the integral principle (10) from which it derives. To see that this is not determined uniquely by the differential equation alone, take any smooth function $F(\rho)$, and rewrite (5) in the form

$$F_t + G_x = 0, \quad (13)$$

where $G = \int F'(\rho) Q'(\rho) d\rho$. The shock speed for (13) is given by

$$c = \frac{G^+ - G^-}{F^+ - F^-} \neq \frac{Q^+ - Q^-}{\rho^+ - \rho^-}$$

(Verify that the inequality holds for functions F that are either strictly convex or concave.) Hence different forms of writing the same differential equation give

rise to different weak solutions, so it is crucial to know which is the conservation principle from which the equation has been derived.

The Riemann Problem

Equation (2) is an example of a *conservation law*, referring specifically in this case to the conservation of cars. Other phenomena modeled through systems of conservation laws include gas dynamics and river flow. For such systems, one of the basic building blocks for both theory and numerics is the Riemann problem: a set of initial conditions consisting of two constant states separated by a discontinuity. For traffic flow, this problem is given by equations (2) and (4), subject to the initial condition

$$\rho(x, 0) = \begin{cases} \rho^- & \text{for } x < 0 \\ \rho^+ & \text{for } x > 0 \end{cases} . \quad (14)$$

Notice that both the PDE (2) and the initial condition (14) are invariant under the stretching

$$x \rightarrow \lambda x, \quad t \rightarrow \lambda t.$$

This implies that if $\rho(x, t)$ is a solution to the Riemann problem, so is $\rho(\lambda x, \lambda t)$, for any value of λ . Hence for any two points (x_1, t_1) , (x_2, t_2) , such that $x_2/x_1 = t_2/t_1$, the solution ρ will adopt the same value. In other words,

$$\rho(x, t) = R(x/t).$$

If $R(\xi = x/t)$ is a smooth function, plugging it into (5) yields

$$(Q'(R) - \xi) R'(\xi) = 0,$$

with alternative solutions

$$R = \text{constant} \quad (15)$$

and

$$Q'(R) = \xi. \quad (16)$$

Since Q' is a decreasing function of R , the solution given by (16) is a decreasing function of $\xi = x/t$. Hence, if $\rho^- > \rho^+$, we can build a solution to the Riemann problem (14) by combining (15) and (16) in the following way:

$$\rho(x, t) = \begin{cases} \rho^- & \text{for } x/t \leq Q'(\rho^-) \\ Q'^{-1}(x/t) & \text{for } Q'(\rho^-) < x/t < Q'(\rho^+) \\ \rho^+ & \text{for } x/t \geq Q'(\rho^+) \end{cases} . \quad (17)$$

An example of such a solution, called a *rarefaction*, is plotted in figure 5a. Since in our example $Q'(R) = 1 - 2R$ is a linear function, the resulting rarefaction is linear in $\xi = x/t$ too.

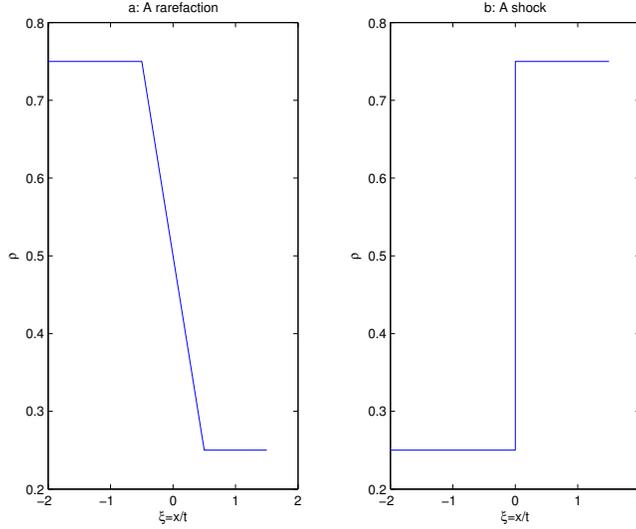


Figure 5: Solution to the Riemann problem for (2), with constitutive relation $Q = \rho(1 - \rho)$. a: A rarefaction. b: A shock wave.

If $\rho^- < \rho^+$, on the other hand, no such matching of solutions is possible. However, in this case, we may fit a shock wave between ρ^- and ρ^+ , with speed given by (11):

$$\rho(x, t) = \begin{cases} \rho^- & \text{for } x/t \leq \frac{Q(\rho^+) - Q(\rho^-)}{\rho^+ - \rho^-} \\ \rho^+ & \text{for } x/t > \frac{Q(\rho^+) - Q(\rho^-)}{\rho^+ - \rho^-} \end{cases} . \quad (18)$$

This discontinuous solution, plotted in figure 5b, completes the solution of the Riemann problem for traffic flow.

1.1 Stability of shocks and uniqueness of weak solutions

The attentive reader may be slightly troubled by the sudden end of the previous section. There we found two qualitatively different solutions to the Riemann problem, depending on the relative sizes of ρ^- and ρ^+ : a rarefaction and a shock wave. Clearly, the rarefaction only fits the case when $\rho^- > \rho^+$; there is no way to match it to the outer, constant solution otherwise. But couldn't the shock solution (18) apply to this case as well?

Indeed, there is nothing in the discontinuous solution (18) that requires that $\rho^- < \rho^+$. Hence we have a problem of non-uniqueness: both (17) and (18) are solutions to the Riemann problem (14), when $\rho^- > \rho^+$. In a way, by extending our definition of a solution to (2) to include weak solutions, i.e. piece-wise smooth functions satisfying the integral constraint (10), we gave ourselves too

much freedom: now we have more solutions than we need; certainly at least two for the Riemann problem with $\rho^- > \rho^+$. Is there a way (and a need) to eliminate all but one of them? We shall argue below that indeed there is a need and a way, and that the only valid solution to (14) with $\rho^- > \rho^+$ is in fact the rarefaction (17).

Let us start with some plausibility arguments. We must recall the reason why we were forced to generalize the notion of solution to (2): because different characteristics may catch-up with each other and cross, thus yielding multi-valued solutions. But, for this to happen, the characteristic starting behind needs to be faster than the one ahead; for traffic flow, this means that it needs to have a lower density ρ . Hence we should not expect a shock to form with $\rho^- > \rho^+$: it would not arise naturally from smooth initial data. This is a first argument in favor of rejecting shocks with the “wrong” ordering of ρ^- and ρ^+ .

Yet, in posing the Riemann problem (14), we just started with a discontinuity in the initial data, which could well have had $\rho^- > \rho^+$ (For instance, the initial condition corresponding to a traffic light that has just turned green has $\rho^- = \rho_{max}$ and $\rho^+ = 0$.) Is there any reason why such discontinuity, once there, could not survive as a shock? To see that there is one reason, recall our logic behind the introduction of the Lax condition (12): that the information required to evolve the smooth solutions on each side of the shock should be available; i.e., that the characteristics carrying this information should be converging toward the shock, not diverging from it. When this condition is violated, as is the case of (18) with $\rho^- > \rho^+$, it is not clear where the information comes from to keep the solution at the two sides of the shock at their constant values. This is a second argument in favor of rejecting the “rarefaction shock”.

Here comes a third, related argument, based on stability. When one writes a differential equation as a model for a physical process, it is often a desired property that it should lead to *well posed problems*: small changes in the data, or small *thermal* perturbations along the way, should not give rise to dramatically different solutions. Yet this is exactly the case of the shock solution (18) when $\rho^- > \rho^+$. For consider, instead of (14), the slightly perturbed initial condition

$$\rho(x, 0) = \begin{cases} \rho^- & \text{for } x < -\epsilon \\ R\left(\frac{x}{\epsilon}\right) & \text{for } -\epsilon < x < \epsilon \\ \rho^+ & \text{for } x > \epsilon \end{cases} , \quad (19)$$

where ϵ is an arbitrarily small number and $R(y)$ is a smooth, monotonic function with $R(\pm 1) = \rho^\pm$. This monotonically decreasing initial condition, will evolve along diverging characteristics and hence never break, approaching a profile arbitrarily close to the rarefaction in (17) as $\epsilon \rightarrow 0$. This is, therefore, the solution of choice, if one is to require well-posedness of the initial value problem.

An alternative wording of the argument above states that a shock wave not satisfying the Lax condition (12) is unstable to small perturbations: if the shock should be slightly smoothed at any time, it will immediately start expanding into a fully blown rarefaction. Shocks satisfying (12), on the other hand, recombine shortly after being pushed apart, due to the compressive effect of the converging characteristics.

A condition such as (12) that restores unicity to the weak formulation of a system of conservation laws is often denoted an *entropy condition*. The origin of this name is to be found in gas dynamics, where physical shocks create entropy, while the unstable, un-physical ones manage to dissipate it, contrary to the second principle of thermodynamics. In the context of our equation, one can show that any strictly concave function $F(\rho)$ works as an entropy: under suitable boundary conditions, the spatial integral of $\eta(x, t) = F(\rho(x, t))$ is constant while the solution $\rho(x, t)$ remains smooth, but grows in the presence of shocks satisfying the Lax condition and decays at discontinuities that violate it.

Problem: Prove this statement when $\rho(x, t)$ satisfies the periodic boundary conditions $\rho(-\pi, t) = \rho(\pi, t)$, modeling traffic in a circular road.

1.2 Numerics: Godunov's method

We now switch from theoretical considerations to the problem of simulating the traffic flow equation (2) numerically. Three main challenges associated with solving systems of conservation laws numerically are the following:

- The solutions one is after are typically not smooth, due to the presence of shock waves. Hence methods that approximate derivatives by assuming some local smoothness, such as finite differences or spectral collocation, are not well-suited to the task.
- The speed of propagation of shocks depends not only on the differential equations being solved, but also on the integral form from which they derive (which, as shown above, are not uniquely determined by the differential equations.) Hence the numerics needs to know about these integral principles.
- Information propagates along particular directions, the system's characteristics. If the numerics does not mimic this flow of information correctly, instabilities ensue.

One numerical method that addresses successfully all the issues above, is due to Godunov, who devised it originally in the context of gas dynamics. This is the method that has been used to produce the numerical results in figure (2), and is the one that we shall choose to describe and use in these notes.

A first ingredient of the method is that it is in *conservation form*, hence enforcing the integral formulation of the problem: the number of cars is preserved exactly by the numerics. The way this is achieved is through the introduction of a discrete spatial grid with mesh size Δx (that we shall for simplicity assume uniform), and a discretized dependent variable ρ_j^n approximating not quite $\rho(x_j, t_n)$, as in most numerical methods, but rather its average over a grid cell:

$$\rho_j^n \approx \frac{1}{\Delta x} \int_{x_j - \frac{\Delta x}{2}}^{x_j + \frac{\Delta x}{2}} \rho(x, t_n) dx .$$

Every time step t_n , ρ_j^n is updated by an equation of the form

$$\rho_j^{n+1} = \rho_j^n + \frac{\Delta t}{\Delta x} \left(Q_{j-1/2}^{n+\frac{1}{2}} - Q_{j+1/2}^{n+\frac{1}{2}} \right), \quad (20)$$

where $Q_{j+1/2}^{n+\frac{1}{2}}$ is an approximation to the flux $Q(x_j + \frac{\Delta x}{2}, t)$ averaged over the time interval $(t_n, t_n + \Delta t)$. Notice that, independently of the accuracy of this approximation to Q , the form of (20), leading to telescopic sums, makes the total mass between two sections $x_{j_1-\frac{1}{2}}$ and $x_{j_2+\frac{1}{2}}$,

$$M^n = \Delta x \sum_{j=j_1}^{j_2} \rho_j^n,$$

change only in response to the difference of fluxes between its two ends:

$$M^{n+1} - M^n = \Delta t \left(Q_{j_1-\frac{1}{2}}^{n+\frac{1}{2}} - Q_{j_2+\frac{1}{2}}^{n+\frac{1}{2}} \right),$$

thus satisfying a discrete version of the integral principle in (3).

Hence the only task left to complete the algorithm is to decide how to compute the approximate fluxes $Q_{j+1/2}^{n+\frac{1}{2}}$. What Godunov proposed at this stage is to compute the *exact* fluxes corresponding to an *approximate* initial profile $\rho(x, t_n)$ consistent with the current local numerical averages ρ_j^n . He chose the simplest such profile: a piecewise constant function, defined by

$$\rho(x, t_n) = \rho_j^n \quad \text{for } x_j - \frac{\Delta x}{2} < x < x_j + \frac{\Delta x}{2}$$

(see figure 6.)

In order to compute the fluxes Q at the cell interfaces, it is enough to notice that, locally, these are the solutions to Riemann problems: two initially constant states separated by a discontinuity. Since information travels along characteristics at a finite speed, non-neighboring cells will not interact with each other, provided that we pick Δt small enough that characteristics coming from two adjacent interfaces do not intersect. As we saw in the prior subsection, the solution to the Riemann problem will be constant along lines of constant slope x/t . For our procedure, we are interested in $Q(\rho(x/t = 0))$: the flux at the interface between cells.

This completes the description of the (first order) Godunov method, applied to single conservation laws. Its Matlab implementation *traffic.m* can be found in the class website. Figure 7 displays a numerical solution with periodic boundary conditions,

$$\rho(x + 2\pi, t) = \rho(x, t),$$

intended to represent a pattern that repeats periodically along the road. We see that an initial smooth disturbance to a uniform traffic profile, due perhaps to natural fluctuations in the traffic, breaks into a discontinuous profile, experimented by the drivers as a sudden traffic congestion. The intensity of this shock

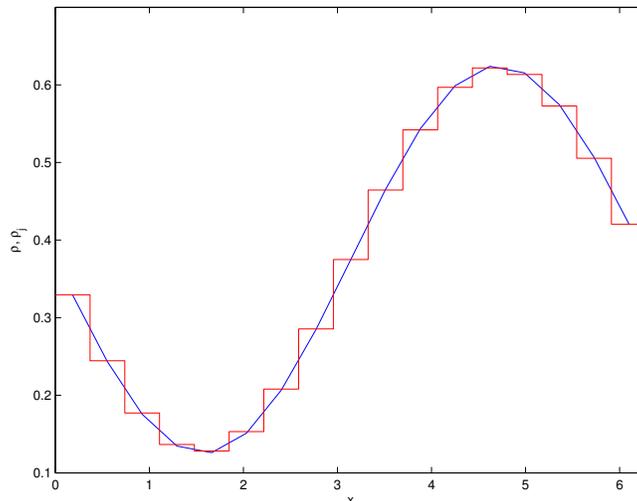


Figure 6: Approximation of ρ by a piecewise constant function

wave then gradually diminishes. This is due to the fact that shock waves that conserve the “mass” $\int \rho(x, t) dx$ dissipate the “energy” $\int [\rho(x, t)]^2 dx$. Then the saw-tooth shaped solution that we observe in the later snapshots of figure 7 decay, asymptotically reaching a uniform steady state with the same mean as the initial data.

Problem: Prove that shocks they preserve mass dissipate energy.

1.3 An empirical validation of the theory

The simple kinematic model for traffic flow described above illustrates the power of the theory of partial differential equations: using just a handful of intuitive facts about traffic –cars don’t suddenly appear or vanish; the lighter the traffic, the faster the cars– one develops a model with rich mathematical structure –characteristics, weak solutions, shock waves, entropy conditions– capable of explaining a mysterious fact of everyday driving: the sudden appearance of traffic congestions and their smooth dissolution into nothingness. Here’s another empirical fact that I have noticed while driving in the streets of New York and Buenos Aires and that the theory explains rather well:

You are in a long traffic queue in a two-way street –say Park Avenue–, waiting for the green light. When the light finally turns green, the first car in the queue moves, then the second, and so on: a wave propagates backward through the queue, telling the cars when to start. On the other hand, you see on your left the first cars moving the opposite way, as they start with the green light at the intersection and soon reach your point in the queue. My observation is –please confirm this with your own experience!– that the time at which the first car

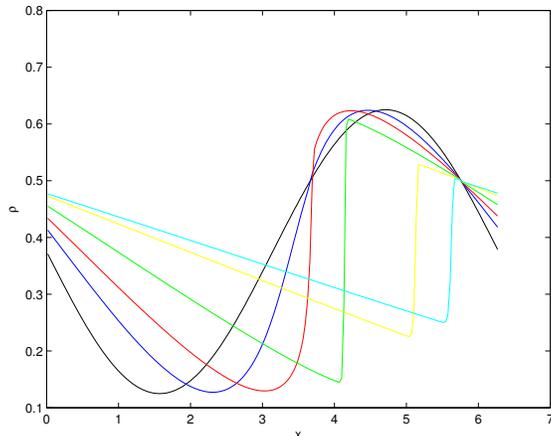


Figure 7: A numerical solution to the traffic flow problem, with periodic boundary solutions. The various colors correspond to snapshots of the solution at different times.

going the other way reaches your position in the queue agrees almost exactly with your starting time. Is there an explanation for this peculiar coincidence?

In the model developed above, the backward starting wave in the queue is the leading edge of a rarefaction wave, with speed given by $c = Q'(\rho_{max})$, where ρ_{max} is the car density at maximal congestion. The first cars on the opposite lane, on the other hand, move at speed $-U(0)$ –the minus sign because of our other way perspective, the zero density because they move into an area with no cars. Then our observation is equivalent to the statement

$$Q'(\rho_{max}) = -U(0),$$

linking the flow at the two extreme opposite conditions: maximal congestion and no traffic at all. Is there a reason why these two should agree?

Consider the simplest model for the speed of cars: maximal speed when the road is empty, decreasing linearly to zero at maximal congestion:

$$U(\rho) = U(0) \frac{\rho_{max} - \rho}{\rho_{max}}.$$

The corresponding flux function is given by

$$Q(\rho) = \rho U(\rho) = \rho U(0) \frac{\rho_{max} - \rho}{\rho_{max}},$$

with

$$Q'(\rho_{max}) = -U(0)$$

as the observation suggests. Then the observation is consistent with the simplest rule for driving speed as a function of traffic conditions. That it should apply to cities like New York and Buenos Aires, where people drive following complex, rather idiosyncratic rules, remains nonetheless a mystery!

2 More on Conservation Laws

We apply a model entirely similar to the traffic flow model of section 1 to river flow. In this context, some of the model's limitations become apparent, which motivates us to switch from a single equation to systems of conservation laws. For general systems, we study characteristics, Riemann invariants, simple waves and shocks.

2.1 Flood waves: a kinematic model

The analysis developed in section 1 has applications far beyond traffic flow. Here we extend it to long waves in rivers, where the principle of car conservation is replaced by volume conservation, under the assumption of negligible evaporation, infiltration and rain. Since the volume between two cross-sections of the river at positions x_1 and x_2 is given by

$$V = \int_{x_1}^{x_2} S(x, t) dx,$$

where S is the cross-sectional area of the river up to the free surface of the water, volume conservation takes the form

$$S_t + Q_x = 0, \tag{21}$$

entirely analogous to (2), with S replacing the car density ρ and $Q(x, t)$ representing the volume flow per unit time through the river's cross-section at position x and time t :

$$Q(x, t) = \int_{S(x,t)} u(x, y, z, t) dy dz.$$

Here u is the component of the fluid velocity normal to the cross-section. As for traffic flow, we have the kinematic constraint that

$$Q = SU,$$

where U is the fluid's mean velocity across the section:

$$U(x, t) = \frac{\int_{S(x,t)} u(x, y, z, t) dy dz}{S(x, t)}.$$

In analogy with traffic flow, we may close equation (21) invoking a relation between Q and S . Hydraulic engineers denote such a relation a *hydrological law*. This is customarily measured in various cross-sections of the world's main rivers: a vertical stick measures the water height h , a surrogate for the area S if the geometry of the river bed is known, while a variety of devices are used for measuring the water speed u at various points. Integrating these velocity measurements across S yields an estimate for Q .