Lecture 1-27

We now return to differential equations and start with a new textbook, the Boyce-DiPrima book mentioned in the syllabus and on the website. We will begin with the initial-value problem $y' = f(x,y), y(x_0) = y_0$ mentioned last quarter, where $f(x,y)$ is a continuous real-valued function on a closed rectangle $R = [a,b] \times [c,d]$ with $(x_0, y_0) \in (a,b) \times (c,d)$. We also assume that f is differentiable as a function of $y \in [c,d]$ for each fixed $x \in [a,b]$ and that the derivative f_y of f with respect to y is continuous on R. Then *this problem has a unique solution $y = g(x)$, whose graph exists up to the boundary of* $[a,b] \times [c,d]$. To prove this we begin (by way of motivation) with a sequence of numbers that nothing to do with differential equations. It is defined recursively by $a_1 = \sqrt{2}, a_{n+1} = \sqrt{2}^{a_n}$. What is $\lim_{n\to\infty} a_n$? This seems quite hard to work out until we hit on the observation that the limit L must satisfy $\sqrt{2}^L = L$, whence it follows that $L = 2$.

What can this sequence possibly have to do with differential equations? To answer this question we first reformulate the initial-value problem. Consider the *operator* F taking the continuous function g on $[a,b]$ to the *function* $Fg = y_0 + \int_{x_0}^x f(t, g(t))\,dt$. This function $F(g)$ satisfies $F(g)(x_0) = y_0$ and $F(g(x))' = f(x, g(x))$, by the Fundamental Theorem fo Calculus, so if we can find a *F-fixed point*, meaning a function $g(x)$ with $Fg(x) = g(x)$, then we will have solved our initial-value problem. Here is where the sequence comes into the picture: we argued that its limit is a fixed point of the function sending x to $\sqrt{2}^x$. In our function setting, if we start with an arbitrary function g_0 on $[a,b]$, say the constant function y_0, and define $g_n(x) = F^n(g_0)(x)$, then the limit of $g_n(x)$ as $n \to \infty$ should be the fixed point of F that we are looking for.

But why should this limit exist? To answer this question we need to define a notion of distance between two continuous functions G, H on a closed bounded interval I: we take this distance $d(G,H)$ to be $\max_{x\in I} |G(x) - H(x)|$ (which is finite). Then it is easy to check that $d(G,H) = 0$ if and only if $G = H$ and $d(G,H) \le d(G,K) + d(K,H)$ for any continuous function K on I. Returning now to the given function $f(x,y)$, both it and $f_y(x,y)$ are continuous on R and so there is an upper bound M for both $f(x,y)$ and $f_y(x,y)$ on this rectangle R (extending a theorem we proved about continuous functions of one real variable last quarter to two variables). Moreover, if $(x,y_1), (x,y_2)$ both belong to R, then by the mean-value theorem $f(x,y_2) - f(x,y_1) = (y_2 - y_1) f_y(x, y_3)$ for some y_3 between y_1 and y_2, whence $|f(x,y_2) - f(x,y_1| \le M|y_2 - y_1|$. Now choose a positive number h small enough that both $Mh = \alpha < 1$ and $S = [x_0 - h, x_0 + h] \times [y_0 - Mh, y_0 + Mh] \subset R$. Then any continuous function f_1 defined on $I = [x_0 - h, x_0 + h]$ such that $|f_1(x) - y_0| \le Mh$ for all $x \in I$ (i.e., whose graph stays in S) is such that $|F(f_1)(x) - y_0| = |\int_{x_0}^x f(t, f_1(t))\,dt| \le M|x - x_0| \le Mh$, so the graph of $F(f_1)$ continues to stay in S. Furthermore any two continuous functions f_1, f_2 whose graphs both stay in S are such that $d(F(f_1), F(f_2) = \max_x |\int_{x_0}^x f(t, f_1(t)) - f(t, f_2(t))|\,dt \le \alpha d(f_1, f_2)$; that is, applying F makes the distance between any two functions whose graphs lie in S at most α times what it was before.

Setting g_0 equal to the constant function y_0 on $[x_0 - h, x_0 + h]$ and g_n the result of applying F to g_0 n times, we are ready to investigate the sequence $\{g_n\}$ of functions. For any $x \in I$, the absolute value $|g_0(x) - g_1(x)|$ is at most the distance $d(g_0, g_1)$ between g_0 and g_1, the absolute value $|g_1(x) - g_2(x)|$ is at most $d(g_1, g_2) \le \alpha d(g_0, g_1)$, and so on; the upshot is that $g_n(x) - g_{n+1}(x)| \leq \alpha^n d(g_0, g_1)$, whence the sequence $g_0(x), g_1(x), \ldots$ is uniformly Cauchy and the sequence g_n of functions converges uniformly to a continuous limit g. Then by continuity $F(g) = \lim_{n\to\infty} g_{n+1} = g$, so g is indeed the fixed point we are looking for. Moreover, there cannot be more than one fixed point, for if g, h are both fixed by F, then $d(F(g), F(h)) = d(g, h)$; but we have already seen that $d(F(g), F(h)) \leq \alpha d(g, h) < d(g, h)$, a contradiction.

Thus there is indeed a unique solution to the equation $y' = f(x, y)$ with $y(x_0) = y_0$. If the graph of this solution does not exist all the way to the boundary of R, then let (x_0', y_0') be the rightmost point on this graph and set up a new initial-value problem with the same differential equation $y' = f(x, y)$ but now the initial condition $y(x_0') = y_0'$. The unique solution to this new problem agrees with the old one where both are defined, by uniqueness, so by putting these solutions together we get a new unique solution whose graph properly extends that of the old one. Iterating this process as many times as necessary we eventually get a unique solution whose graph exists all the way to the boundary of R, as desired.

A simple but instructive example of the above construction arises from the initial-value problem $y' = y, y(0) = 1$. Following the notation of the above proof, we get $g_0 = 1, g_1 = 1 + \int_0^t 1\, dx = 1 + t.g_2 = 1 + \int_0^x (1 + t)\, dt = 1 + t + t^2/2$, and so on; here the limit g of the g_i is the power series $\sum_{i=0}^{\infty} t^i/i!$, which you recognize as our old friend e^t. Note that in general we cannot predict where the graph of the unique solution leaves the boundary of the rectangle R; for example, given the initial-value problem $y' = y^2 + 1, y(0) = 0$, the unique solution $y = \tan x$ blows up at $x = \pm\pi/2$. It is by no means obvious from the differential equation $y' = y^2 + 1$ alone that the solution blows up for a finite x; all we know in advance is that the solution must be defined on some interval containing $x = 0$.