An Essay on ODEs from the MF Section 5.2 Perspective PhL 12.20.09

ODE Ordinary Differential Equation (1 variable)
BV Boundary Value
MF Morse & Feshbach 1953
D d/dx
PDE Partial Differential Equation (n variables)
WW Whittaker & Watson circa 1927

1. **Comparison of MF 5.2 to Stakgold Chapter 4.** First, let's discuss what this little essay is NOT about. In the MF approach of Section 5.2, we don't care if the operator L is self-adjoint, because we are not doing boundary value problems, meaning eigenvalue problems. In an EV problem, you want L = L* so you can swing L around in the inner product and be able to prove various facts such as the reality of eigenvalues and the orthogonality of eigenfunctions. In effect, you want L to be like a Hermitian Matrix in quantum mechanics. At a slightly more advanced level, if you compute the Green's Function (kernel of integral equation), and if it is Hilbert-Schmidt, then you know that the set of eigenfunctions forms a complete set, an example being Legendre functions on (-1,1). That is, the eigenfunctions form a basis in the L_2 Hilbert Space, which space's definition includes the endpoints of an interval (a,b).

None of these things is of interest in Section 5.2 of MF. We never talk here about any particular interval (a,b) and we never mention boundary conditions. There is no scalar product, there is no Hilbert Space. Our only concern is to be able to write down the "general solution" of the ODE, meaning we want to know the two independent functions which solve the ODE. The subjects mentioned in the above paragraph are treated in Chapter 4 of Stakgold volume 1, the spectral theory of L operators. In particular, the "general" self-adjoint L = -D(pD) + q is treated (see page 268) in both regular situations and in singular situations, all from the eigenfunction perspective. "Regular and singular" in this BV context is unrelated to "regular singular point" in our MF Section 5.2 context to be described below. For example, if an end of the interval (a,b) goes to infinity, the spectrum becomes continuous and we speak of that endpoint as being "singular" in Stakgold. But that does not mean the ODE has a singularity at ∞ (regular or irregular). It is the boundary value problem which is regular or singular, not a point in the variable plane that is ordinary or singular. And it is the endpoints of the interval (a,b) that make the regular/singular BV distinction, whereas in the MF theory we don't even have any endpoints.

Of course in the boundary value problem world, the functions which solve the problem are going to be the same functions (combinations of them) which are discovered in the MF Section 5.2 approach. In general, these will be "special functions" with people's names attached to them, like Bessel or Legendre or Mathieu or Lamé.

2. General Solution sought. So, what exactly is the MF Section 5.2 "world" we want to talk about? In this world, we are handed a second order differential operator $L = D^2 + pD + q$ and we are asked to solve the equation Lu=0 (the "homogeneous equation"), and maybe at some later time the equation Lu = f with some ("inhomogeneous") driving term f. We will only consider Lu=0 in this essay. We are looking for "the general solution" of Lu= 0 which we know is a linear combination of two independent functions (at most), coefficients often written A and B. There are no boundary conditions, no intervals. The argument is treated as a complex variable.

2A. Source of ODEs. Before we delve into this topic, we should ask where these ODE's might be coming from? In the MF perspective, focus is kept on a particular PDE called the Helmholtz equation which, if the constant k_1 is set to 0, becomes the Laplace equation. The focus is also on systems of separable coordinates. When you are able to separate a PDE in some coordinate system, you end up with an ODE in each variable, and each ODE contains the "separation constants" in some way. These are the ODE's that are of main interest to MF. They come from PDE's and then are formed by attempting a separation. Some or all of the "singularities" of these ODE's come from singularities in the way to coordinates are defined, locations MF call "concentration points" in loose language.

3. Forms of an ODE. Now back to our ODE Lu = 0 with

$$L = D^2 + pD + q \,. \qquad (*)$$

Here, p and q are the traditional symbols for general functions appearing in this equation. I have called this "the canonical form" to distinguish it from another form which is the Frobenius form noted below. Notice that these functions p and q are NOT the same as those which appear in the Stakgold "self-adjoint form". The Stakgold is this,

$$L = -D(pD) + q = -pD^2 - p'D + q = (-p)[D^2 + (p'/p)D - (q/p)]$$

If we avoid places where p=0, and if we are interested in Lu=0, then we can identify

$$p = (p'/p) \qquad q = -(q/p)$$

The Frobenius form is

$$L = (z-a)^2D^2 + (z-a)P_a D + Q_a = (z-a)^2 [D^2 + (P_a/(z-a)) D + Q_a/(z-a)^2]$$

and again, if we avoid z=a and care only about Lu=0, we can identify

$$p = (P_a/(z-a)) \qquad q = Q_a/(z-a)^2$$

This Frobenius form puts special attention on the point z = a. We could put attention on a different point b and rewrite (*) just as easily in this way

$$L = (z-b)^2D^2 + (z-b)P_b D + Q_b = (z-b)^2 [D^2 + (P_b/(z-b)) D + Q_b/(z-b)^2]$$
$$p = (P_b/(z-b)) \qquad q = Q_b/(z-b)^2$$

where we are always considering only Lu=0 so we think of L and factor x L as the same thing, as long as we make sure factor $\neq 0$ or ∞.

4. Statement of the MF Problem. Now, finally, we come to the core of our essay. Someone hands you an ODE Lu= 0 with L given by $L = D^2 + pD + q$. The variable is z, and we are interested in z being anywhere in the complex plane. There is no *interval* a,b, although we are often interested in a disk-shaped *area* of the plane in which some power series might converge. You stare at this ODE Lu=0. Your

problem is to find "the general solution" to this equation at all points in the z plane!!! Remember that you are given p(z) and q(z). It is certainly not obvious, upon naive consideration, how on earth you would find this general solution, even for some small region of z. What function, if you differentiate it once, multiply by some function p, add to the second derivative, will end up being -q(z) ? This sounds very "trial and error". It is not very obvious how you might "integrate" the ODE to get a solution, since it is not likely to involve a perfect differential. So the big question is this: how do you even begin to answer this question?

5. Aside on constant coefficients. There is a whole world of ODE study where the coefficient functions are constants, meaning they don't depend on z. One sees $L = a D^2 + bD + c$ and one considers $Lu=0$. Dividing by "a", which is non-singular everywhere, we get $L = D^2 + pD + q$ where p and q are constants. This kind of equation with constant coefficients, arises in many contexts in physics and electronics, to name two fields I know a little about. In our essay today, we are never thinking about functions p and q which are constants. In the constant case, we know pretty much all about solving $Lu=0$ using Laplace transforms, and even $Lu=f$ is pretty straightforward. It is the constant-coefficients case that "the student" usually first encounters in his or her educational path. The solutions are usually exponentials and trig functions combined with other elementary functions. The solutions are not what we call "special functions".

6. Frobenius Power Series Approach. So, where do we begin solving $Lu=0$ with $L = D^2 + pD + q$? The key starting point idea is a very ancient in mathematical terms: the power series. In some crude sense, we seek a solution u(z) as a projection onto powers z^n, knowing that such powers don't form a "basis" or a "complete set of functions" or anything like that. The main reason we consider power series is that we know what D and D^2 do to a power of z, and we hope to expand p and q in their own power series. Now when we say "a power series", we have in mind "a power series about some point z = a", though we often think generically about doing this at the point a = 0. In general, if we can find a power series solution of $Lu=0$ about some point z = a (we hope in fact to find two independent such power series solutions), we might expect this series to converge in a disk that can be enlarged until it hits some "singularity" of the "full" general solution function for the ODE, whatever that solution might be.

7. Where are the branch points of the general solution? So how does this power series approach get us anywhere? What we really want is the "full general solution", not some local power series solution.

Well, our first subject of interest in thinking about the full general solution is to try to figure out where that general solution has singularities -- places where the general solution blows up in some manner or has multiple values as you circulate around the singularity, indicating the presence of multiple Riemann sheets connected at a branch point. We digress momentarily on this latter case.

8. Comments on branch points. The function $f(z) = \sqrt{z}$ does not blow up at z = 0, it is 0 there. But we know there is a branch point at z = 0, and we have to draw a branch cut and break the z plane into two "sheets" in order to truly "understand" this function.

We shall refer to $f(z) = (z-a)^s$ as having a **"power branch point"** at z = a. The power s, assuming it is real, could be an irrational number, a rational number, or an integer.

If it is a positive integer, then there really is no branch point. If it is a negative integer, again there is no branch point and we have a pole of order s. In all other cases, we really do have a power branch point. All cases I have seen in the MF world involve real powers, so we shall simply ignore the case of imaginary or complex power s.

If s is irrational, there are an infinite number of Riemann sheets associated with the branch point, and if rational (aka a "fractional power"), there are a finite number of such sheets.

As far as I know, there is only one other type of branch point, the **logarithmic** one. The function f(z) = ln(z-a) has a log branch point at z = a, and such branch points always have an infinite number of sheets to worry about.

We can of course construct fancy functions like f(z) = ln[$(z-a)^s + 4$] which has fancier branch points, but really each branch point has its source in either a power or a log and these then get propagated as you build up a more complicated function.

9. Singularity Types. We can classify singularities of a complex function f(z) into three bins: (1) poles of finite order; (2) power or log branch points; (3) essential singularities. An **essential singularity** is isolated in the same way a pole (of any order) is isolated -- there is no branch point and cut. The usual example of a function with an essential singularity is f(z) = exp($A/(z-a)^s$) with Re(s)>0. As $z \to a$, the exponent $\to \infty$ and $f \to e^{\infty}$ which is some kind of exponential infinity that blows up worse than any pole of finite order. If you make a power series for such an f(z), it is a Laurent series including all negative powers, meaning it has poles of all finite orders.

I think this list of types of singularities is exhaustive. People talk about the "species" of an essential singularity which gets higher as s increases in our example.

My impression is that essential singularities are unpleasant things that you want to stay away from. If you have to have one, it is good if it can be had at $z = \infty$ where it is sort of out of sight, out of mind.

10. The Frobenius Theory and Classification. Now back to our main flow. We want to try to identify the singularities of the general solution of our ODE Lu = 0 where $L = D^2 + pD + q$. This is where what I might call "the Frobenius theory" comes to our aid. There are only three cases to think about:

(a) if p and q are both analytic at z = a, the two power series around z = a are analytic, there is no singularity at z = a in the general solution, the Frobenius exponents are r = 1 and 0. [These are the solutions, sometimes called indices, of the quadratic "indicial equation" obtained by balancing the leading powers when everything is power series expanded.] Any such point z = a is called an **ordinary point**. In general, these two power series solutions at an ordinary point z = a will converge in a disk that is as large as you want, providing that disk does not bump into a branch point, pole or an essential singularity of the general solution function. These in turn can be located according to the next two paragraphs.

(b) if p has a simple pole at z = a, and/or if q has a simple pole and/or a double pole at z = a, then in general there will be a power branch point at z = a. The two independent solutions of Lu=0 will, in general, have the form u(z) = $(z-a)^{r1}$ f(z) and u(z) = $(z-a)^{r2}$ g(z) where r_1 and r_2 are the Frobenius exponents (solutions of the indicial equation), and f and g are analytic power series about z = a. There are several "special cases" of interest. If $r_1 - r_2$ = an integer, including the case $r_1 = r_2$, then one solution will be of the form just stated (power branch point function times an analytic function) while the second solution will have a factor ln(z-a) and will thus have a log branch point.

A point z = a of this type is called a **regular singular point**.

(c) if p and or q has *any other* singularity, such as a pole of higher order than described in b, or such as a power or log branch point, or such as an essential singularity, then the general solution of the ODE will have an essential singularity at z = a. If you consider the two independent solutions near z = a, at least

one of these solutions will show that essential singularity, so the general solution will have one. A point z = a of this type is called an **irregular singular point**.

That then is the complete "Frobenius theory" which lets you identify the singular points of the general solution of an ODE of second order Lu = 0 where $L = D^2 + pD + q$. Given p and q, we can take out a piece of paper and immediately write down a list of the regular points (case b) and of the irregular points (case c). All other points are then ordinary points. When discussing the z plane, we always have to include the point $z = \infty$ in our discussion, and it is not quite as obvious how you determine whether $z = \infty$ is ordinary, regular or irregular, but there is a well-defined method for doing this.

11. What we now know. So at this point, given our ODE, we know at once the location and nature of all the singularities of the general solution of the ODE. That is an impressive thing to know, considering we started completely cold, knowing nothing at all. At each regular singular point z = a, we know the two associated exponents r_1 and r_2 (indices), so we know the nature of the power branch points (and possible special case log branch points) at each singular point. And we also know where the essential singularities are located. Notice that we therefore know ALL the branch points of the general solution and we know the nature of these branch points, so we have a complete knowledge of the Riemann sheet structure of our general solution.

12. A path not taken. So, what about our actual solutions to the ODE? We could go to some distant ordinary point and find the two power series solutions at that point and give the series function "names". Perhaps we would choose a point far from all singularities in order to get a large disk of convergence. I think that typically this is NOT what is done. I suspect that one really needs to know the solution in the region of singular points, so the approach just mentioned is not really useful.

13. Shuffling to get standard forms. Typically an ODE (of practical interest) has some small number of regular singular points, and at most one irregular singular point. It turns out that by changing variables and/or functional form [eg, let $u(z) = e^z v(z)$] , you can in general cause the singular points to be at "standard locations". The most standard location is $z = \infty$ (out of sight, out of mind), the next is z = 0, and the next is z = 1. Sometimes instead of z = 0 we use z = -1 as with Legendre functions. If you are given an ODE, say, with three regular singular points, you can message the ODE to cause these points to be at $0,1,\infty$. The general ODE with 3 regular singular points a,b,c is called the **Papperitz** or **Riemann P** equation, while the one with those points at $0,1,\infty$ is the **hypergeometric equation**.

By "shuffling" the ODE, it is also possible to "tune" the exponents to get "nice values" instead of ugly values. The "nicest" value an exponent can have at some regular singular point is r = 0, because then the series solution corresponding to that exponent is the entire solution, there is no multiplicative power factor with a branch point. For example, at the regular singular point z = 0, the just-mentioned hypergeometric equation has one exponent being 0, and the series solution that goes with that exponent is called F(a,b,c,z), known as the hypergeometric series. Although this thing is a series, we think of it as a well-defined and standard function (as further explained below), like the "sine function" or "the Bessel function". This thing F(a,b,c,z) is then an example of "a special function", which just means it is a new function that is not an elementary function and which is a solution of a certain ODE that is of interest. It turns out that F(a,b,c,z) = F(b,a,c,z) so people sometimes write F(a,b,c,z) as F(a,b;c;z) or F(a,b|c|z). In this same example, and in other examples, one can usually obtain the "other" independent ODE solution using what I have called "**the trick formula**" which I think is 5.2.6 in MF. For the hypergeometric equation case, this trick produces a power factor times another F function with different parameters but

the same argument. I quote MF just this one time regarding the general solution which is a linear combination of F(a,b,c,z) and that second solution, with constants A and B,

$$AF(a,b|c|z) + Bz^{1-c}F(b-c+1, a-c+1|\,2-c|\,z) \quad (5.2.45)$$

Here then is the "general solution" of the hypergeometric equation expanding around z = 0. You can see by inspection that the two exponents for the regular singular point z = 0 are r = 0 and r = 1-c. Since we know there is a regular singularity at z = 1 (the nearest other singular point), we know that the above solution has a radius of convergence of 1. The above solution gives us good knowledge of the ODE solution in a fairly large neighborhood of the singular point z = 0, but it is not very useful for the neighborhood of z = 1. If we need to know the general solution near z = 1, we have to find power series solutions at that point, which we know how to do. Those solutions, if they were of the variable (z-1), would also have a radius of convergence of 1, since their disk would hit the z = 0 singularity.

14. Naming and kinds of special functions. So far, then, we know all about the singularity structure of our ODE general solution, and we have solutions that are valid in certain disks of the complex plane. At least one of these solution functions is given a special function name, like F(a,b,c,z), or like $P_\nu^\mu(z)$ for Legendre. The "other" function is not uniquely determined since if you found $f_{\mathtt{other}}(z)$, then 3 $f_{\mathtt{first}}(z)$ + 5.6 $f_{\mathtt{other}}(z)$ would be a "just as viable" other independent function. Usually one first finds a function that is a useful Frobenius solution at some regular singular point (like $P_\nu^\mu(z)$ at z = 1, or like F(a,b,c,z) at z = 0) and gives that function a "special function" letter name (like P or F, with the parameters glued on in some manner) and a spelled-out language name (like Legendre, or Hypergeometric: usually the name is of one of the first people who dealt with this ODE and its solution). One usually refers to this function as "**the first kind**" solution. A choice is then made for the "other" solution and it is then referred to as the "**second kind**" solution, like the Legendre $Q_\nu^\mu(z)$ or the second F function shown in the quote above.

15. The idea of analytic continuation. We might be a little unhappy with our "spotty" method of solution of an ODE, where we have a set of little disks, each associated with a different series solutions surrounding a regular singular point. We can also have disks and solutions around ordinary points, as noted above.

So here is the idea, where we shall take the hypergeometric case as our example. We have a certain convergent series for F(a,b,c,z) -- we know the convergence radius is R =1 about z = 0. We then try to come up with a way to extend the definition of this function F(a,b,c,z) to values of z which lie outside the radius. Obviously, for such values, the series no longer converges, so what could this possibly mean? Suppose we could first find a region in the z plane where F(a,b,c,z) does converge, and where the two solutions about z = 1 also converge. Maybe we go to some point near z = 3/4 and put a little region around it. We know that in this region, both our z = 0 and both our z = 1 solutions are convergent series. But we also know that we can write each of the z = 0 solutions as a linear combination of the z = 1 solutions. So, suppose we find that (this is sometimes called a joining relation)

$$F(a,b,c,z) = K_1(a,b,c)\; f^{(1)}{}_{\mathtt{z=1}}(a,b,c,z) + K_2(a,b,c)\; f^{(2)}{}_{\mathtt{z=1}}(a,b,c,z)$$

where we have made up a little notation for our two independent z = 1 solutions. Now suppose we move out of our little dual "safe region" disk near z = 3/4. Suppose we move to the right say to a region near z = 1.3. We know for sure that the nominal series F(a,b,c,z) diverges there, but we also know that the two z =

1 series converge there. So the idea is to regard the RHS of the above equation as a the *meaning* of F(a,b,c,z) "function" in the neighborhood of z = 1.3. We have thus "analytically continued" the function F(a,b,c,z) outside its nominal convergence disk of radius 1. When we do this, we are then saying that the function F(a,b,c,z) has some kind of life of its own, beyond its "series representation" as a power series around z = 0. In fact, we regard F(a,b,c,z) as a more general object which just happens to have that series representation by which we first discovered it! We start referring to this thing as the hypergeometric "function" instead of the hypergeometric "series".

In this manner, it is possible to give F(a,b,c,z) a meaning in a neighborhood of any point in the z plane!! How this works out in detail involves the 20 Kummer relations listed on pages 106-107 of Bateman (vol 1). Each of these triple joining relations is like the one shown above, relating three solutions of the ODE. In addition to these 20 relations, there are 6 basic functions, one of which is F(a,b,c,z), called u_1 thru u_6, each of which has three alternate forms which in themselves provide some analytic continuation "capability". For example, we have

$$u_1 = F(a,b;c;z) = (1-z)^{-a} F(a,c-b;c; z/(z-1))$$

The left function has series convergence for $|z| < 1$, our usual disk, but the right form has convergence for the region $|z-0| < |z-1|$ which is all points closer to 0 than to 1, which is an entire half of the complex plane with a vertical line edge coming down through the point z = 1/2. So this one "alternate form" equation analytically continues F(a,b,c,z) to an entire half plane, all by itself! This formula is proved by a simple shuffling of the ODE, by the way. Another way these things are done is by writing something like F(a,b,c,z) as an integral (an "integral representation", as opposed to a "series representation") and then we can manipulate the integral perhaps doing parts integration or whatever to get another form which converges in some different z range.

There is a theorem which states that if two analytic functions agree on some finite region of the complex plane, but have different global convergence regions in the plane, then each function is the *unique* analytic continuation of the other outside its normal region of convergence. So there is never any issue of not finding the right analytic continuation. Of course you have to pay attention to the various branch cuts you have and make sure you know what sheet of each cut you are on. (This is something Maple worries about a lot, sometimes excessively).

16. What we now know? So here then is our final status. We are handed an ODE. We find its regular singular points and if it has an irregular singularity, we arrange to have that at $z = \infty$. (If it has multiple irregular singular points, we throw up our hands and think numerical.) We also arrange to have the regular singular points be at standard locations, and we arrange to have "nice" Frobenius exponents at these locations. We then pick one or two "standard" functions like F(a,b,c,z) which start off being specific convergent power series, but then we analytically continue these functions to all points in the z plane. Then we can say that this function and its second kind partner ARE the two independent solutions of our ODE for all z, and we are done! Problem solved, mission accomplished. Such a function is called "a special function". There are lots of books just about "special functions" and their properties. Note that the pair of independent solutions you find in this way is not unique, since you could linearly combine them to create a new pair of independent solutions.

17. Five Regular Singular Points and Confluence. It is possible to write down a "most general" ODE that has five regular singular points with one of these at $z = \infty$. This thing is called the **generalized Lamé equation**. It turns out that by considering confluences of these singular points in various ways, you can

generate pretty much all the ODE's of interest in theoretical physics. For example, if you let two pairs of regular points conflow (become the same), you end up with only three regular points, and this is the hypergeometric equation situation. If you let three points converge, they become a single irregular point, and then you have a situation with one irregular point (which you can then throw out to $z = \infty$) and two regular points, which is the Mathieu world. If you let two of the regular points of the hypergeometric equation conflow, you end up with an ODE having one irregular and one regular singular point, and this is the "confluent" hypergeometric equation with solutions F(a,c,z) . See WW Chap 10 for all the cases.

18. Curvilinear harmonics. As we noted earlier, if you study the Laplace equation in various separable curvilinear coordinate systems, you end up with certain ODE's as your separated equations. These are the ODE's that most interest MF in Section 5.2 of their tome. The solutions to these ODE's have lots of interesting names, depending on coordinate system, and the product of the three solution functions (which is then a solution of the full Laplace PDE) are called **harmonics** specific to that coordinate system. The triple product always has some kind of correlation between the separation constants. This correlation reaches its most complex form in the case of ellipsoidal coordinates and the corresponding ellipsoidal harmonics, which are regular Lamé functions.

19. Quantization of separation constants. In these curvilinear coordinate systems, it is true that for each separation ODE we do get a certain "interval" of interest, such as (-1,1) for $z = \cos\theta$ in spherical coordinates. As we noted at the very start above, in this little essay we are not dealing with intervals and boundary value problems, we are just talking general solutions. It is only when we start looking at specific intervals that we end up getting quantization of the separation constants, so this subject is beyond this scope of this essay. Usually the "harmonic" functions include the effect of this quantization, the most familiar being $\mathbf{r}^{\ell}$ and $\mathbf{r}^{-\ell-1}$ times $Y_{\ell m}(\theta,\varphi)$, where Y are the so-called "spherical harmonics" with $\ell = 0,1,2$ and $|m| \leq \ell$.

Appendix A. Clarification of Stakgold "regular" and "singular"

We have noted the Stakgold self-adjoint form to be this (functions are all real for Stak, and here I use non-italic p and q for the Stakgold functions)

$$L = (-p) \{ D^2 + [p'/p] D - q/p \} = (-p) \{ D^2 + p_1 D + q_1 \} \qquad Lu = 0$$

$$p_1 = d (\ln p)/dp = p'/p \qquad q_1 = -q/p$$

where I try to put this roughly in MF canonical form with canonical functions p_1 and q_1 . Imagine that we have some interval (a,b) on the real axis. The conditions Stakgold gives for a regular BV problem are (p 268) that his p (and p') and q be *continuous* in the interval and that p be positive. The positivity of p justifies our cancelling it out when we think of Lu = 0, so we then have

$$L = D^2 + p_1 D + q_1 \qquad p_1 = d (\ln p)/dp = p'/p \quad q_1 = -q/p$$

Now we can try to compare the MF and Stak singularity senses. All points within the interval would be ordinary points (to MF) as long as p_1 and q_1 were analytic in the interval. Certainly if we follow that Stak

rule that the functions be "continuous", we cannot allow p , p' or q to have even a single pole in the closed interval. That would pretty much rule out q_1 having any poles. I am pretty sure this rules out a pole in p_1 as well. And no doubt branch points are also ruled out. So I think that, basically, the Stak "regular" boundary value problem means that p and q are analytic "on the closed interval". Of course the whole Stak theory is for a real variable so "analytic" is a little hazy. But I think the claim would be this:

Claim #1: If you want to have a "regular BV problem" for the interval (a,b), the functions p and q should be analytic in the entire closed interval (a,b). This means that all points of the closed interval (a,b) should be ordinary points of the ODE. If the ODE has singular points of either regular or irregular type *outside* the interval, the BV problem is not affected. We only care about points in the interval (a,b) because only these points are involved in the scalar product.

Claim #2: For Stakgold, a "singular BV problem" occurs when one or both of the *endpoints* a, b are "singular". I think the functions p and q still have to be analytic on the entire open interval (a,b). As Stak says on page 295, an endpoint is singular in only two cases: (1) it is infinite, eg, $b=\infty$; (2) $p = 0$ there, eg, $p(a) = 0$. Recall that for the regular problem p is positive on the closed interval (a,b).

Observation: In either case, the regular *or* singular BV problem for (a,b), the solution function to the ODE on (a,b) with (homo) BC's will be analytic at every point in the *open* interval. I argue this since all interior interval points are analytic for p and q and are therefore ordinary, and the general theory above says that the general solution at ordinary points is always analytic. The upshot is that BV problem solutions f(x) are not going to have strange singularities within (a,b), though they might have singularities at the endpoints. For example, I would expect that the functions $P_\nu^\mu(z)$ *and* $Q_\nu^\mu(z)$ have no singularities in the open interval (-1,1). Each of these can be expressed in terms of powers of $(z \pm 1)$ times F(a,b,c; [1-z]/2) where the F converges in a disk of radius 2 centered at $z = 1$. For the Q function, there are two such terms you add together, see bottom of Bateman p 130. [Yes, you might draw the branch cuts along the interval (-1,1) but it is understood then that our region of interest is a disk of radius 2 between the two endpoints (-1,1), even if half this disk lies on one sheet and half on another. It is clearer if we just pull the branch cuts away and sent them radially outward from $z = 1$ and $z = -1$, for purposes of our current discussion.]