Convergence of solutions of p-adic differential equations

In this unit, we consider the radius of convergence of a local horizontal section of such a differential module. In particular, we define a fundamental invariant, the *generic radius of convergence*, and some related invariants, the *subsidiary radii*.

1 Radius of convergence on a disc

View $K\langle \alpha/t, t/\beta \rangle$ as a differential ring with derivation $d = \frac{d}{dt}$, the formal differentiation in the variable t. Note that this does indeed carry $K\langle \alpha/t, t/\beta \rangle$ into itself.

Proposition 1. *Any finite differential module over $K\langle \alpha/t, t/\beta \rangle$ is torsion-free, and hence free.*

Thus I could omit saying "free" in many statements, though I will continue to do so for emphasis.

Proof. Exercise. □

Let M be a finite free differential module over $K\langle t/\beta \rangle$. The fundamental theorem of p-adic ordinary differential equations, if it were true, would say that $M \otimes K\langle t/\rho \rangle$ admits a basis of horizontal elements (elements in the kernel of D) for any $\rho \in [0, \beta)$. Unfortunately, this is simply false, as we saw in the introduction; it already fails for the module $M = K\langle t/\beta \rangle$ with $D(x) = x$, when $\beta > p^{-1/(p-1)}$.

We define the *radius of convergence* for M around 0, denoted $R(M)$, as the supremum of the set of $\rho \in (0, \beta)$ such that $M \otimes K\langle t/\rho \rangle$ has a basis of horizontal elements. We will see a bit later (Corollary 5) that the set is nonempty, so this supremum makes sense.

Writing everything in terms of such a basis, we see that the only elements of any $M \otimes K\langle t/\rho \rangle$ that can be horizontal are the K-linear combinations of the basis elements; we call such linear combinations *local horizontal sections* of M. (As I may have done before, I use the word "section" because I am thinking geometrically, in terms of a vector bundle equipped with a connection.)

Here are some easy consequences of the definition; note the parallels with properties of the truncated spectral norm. In particular, the proof of (c) below may clarify the proof of the corresponding property of the truncated spectral norm.

Lemma 2. *Let M, M_1, M_2 be finite free differential modules over $K\langle t/\beta \rangle$.*

(a) If $0 \to M_1 \to M \to M_2 \to 0$ is exact, then

$$R(M) = \min\{R(M_1), R(M_2)\}.$$

(b) We have

$$R(M^\vee) = R(M).$$

(c) We have

$$R(M_1 \otimes M_2) \geq \min\{R(M_1), R(M_2)\},$$

with equality when $R(M_1) \neq R(M_2)$.

Proof. For (a), it is clear that $R(M) \leq \min\{R(M_1), R(M_2)\}$; we must check that equality holds. Choose $\lambda < \min\{R(M_1), R(M_2)\}$, so that $M_1 \otimes K\langle t/\lambda\rangle$ and $M_2 \otimes K\langle t/\lambda\rangle$ are both trivial. If we choose a basis of M compatible with the sequence, then the action of D will be block upper triangular nilpotent, and trivializing M amounts to antidifferentiating the entries in the nonzero block. We may not be able to perform this antidifferentiation in $K\langle t/\lambda\rangle$, but we can do it in $K\langle t/\lambda'\rangle$ for any $\lambda' < \lambda$. Since we can make λ and λ' as close to $\min\{R(M_1), R(M_2)\}$ as we like, we find $R(M) \geq \min\{R(M_1), R(M_2)\}$.

For (b), we obtain $R(M^\vee) \geq R(M)$ from the fact that if $M \otimes K\langle t/\lambda\rangle$ is trivial, then so is its dual $M^\vee \otimes K\langle t/\lambda\rangle$. Since M and $M^\vee$ enter symmetrically, we get $R(M^\vee) = R(M)$.

For (c), the inequality is clear from the fact that the tensor product of two trivial modules over $K\langle t/\lambda\rangle$ is also trivial. If $R(M_1) < R(M_2)$, then we have

$$\begin{aligned} R(M_1) &= \min\{R(M_1), R(M_2)\} \\ &\leq \min\{R(M_1 \otimes M_2), R(M_2^\vee)\} \\ &\leq R(M_1 \otimes M_2 \otimes M_2^\vee). \end{aligned}$$

Moreover, $M_2 \otimes M_2^\vee$ contains a trivial submodule (the trace), so $M_1 \otimes M_2 \otimes M_2^\vee$ contains a copy of M; hence by (a), $R(M_1 \otimes M_2 \otimes M_2^\vee) \leq R(M_1)$. We thus obtain a chain of inequalities leading to $R(M_1) \leq R(M_1)$; this forces the intermediate equality $R(M_1) = \min\{R(M_1 \otimes M_2), R(M_2^\vee)\}$. Since $R(M_1) \neq R(M_2) = R(M_2^\vee)$, we can only have $R(M_1) = R(M_1 \otimes M_2)$. □

Here is a simple example.

Proposition 3. *Let M be the differential module of rank 1 over $K\langle t/\beta\rangle$ defined by $D(v) = \lambda v$ with $\lambda \in K$. Then*

$$R(M) = \min\{\beta, |p|^{-1/(p-1)}|\lambda|^{-1}\}.$$

Proof. Exercise. □

In general, the radius of convergence is difficult to compute. To get a better handle on it, we introduce another invariant which looks more complicated to define but has much simpler behavior; this is the generic radius of convergence introduced below.

2 Generic radius of convergence

For $\rho > 0$, let F_ρ be the completion of $K(t)$ under the ρ-Gauss norm $|\cdot|_\rho$. Put $d = \frac{d}{dt}$ on $K(t)$; then d extends by continuity to F_ρ, and

$$|d|_{F_\rho} = \rho^{-1}, \qquad |d|_{\mathrm{sp},F_\rho} = \lim_{n\to\infty} |n!|^{1/n}\rho^{-1} = p^{-1/(p-1)}\rho^{-1}.$$

Let (V, D) be a finite differential module over F_ρ. We define the *generic radius of convergence* (or for short, the *generic radius*) of V to be

$$R(V) = p^{-1/(p-1)}|D|^{-1}_{\mathrm{tsp},V};$$

note that $R(V) > 0$. We will see later (Proposition 10) that this does indeed compute the radius of convergence of horizontal sections of V on a "generic disc".

In the interim, we note the following relationship with the usual radius of convergence. In the language of Dwork, this is a *transfer theorem*, because it transfers convergence information from one disc to another.

Theorem 4. *For any finite free differential module M over $K\langle t/\rho\rangle$, $R(M) \geq R(M \otimes F_\rho)$. That is, the radius of convergence is at least the generic radius.*

Proof. Suppose $\lambda < \rho$ and $\lambda < p^{-1/(p-1)}|D|^{-1}_{\mathrm{tsp},V}$. We claim that for any $x \in M$, the Taylor series

$$y = \sum_{i=0}^{\infty} \frac{(-t)^i}{i!} D^i(x) \tag{1}$$

converges under $|\cdot|_\lambda$. To see this, pick $\epsilon > 0$ such that $\lambda p^{1/(p-1)}(|D|_{\mathrm{sp},V} + \epsilon) < 1$; then there exists $c > 0$ such that $|D^i(x)| \leq c(|D|_{\mathrm{sp},V} + \epsilon)^i$ for all i. The i-th term of the sum defining y thus has norm at most $\lambda^i p^{i/(p-1)} c(|D|_{\mathrm{sp},V} + \epsilon)^i$, which tends to 0 as $i \to \infty$.

By differentiating the series expression, we find that

$$\begin{aligned} Dy &= \sum_{i=0}^{\infty} \frac{(-t)^i}{i!} D^{i+1}(x) + \sum_{i=1}^{\infty} \frac{-(-t)^{i-1}}{(i-1)!} D^i(x) \\ &= \sum_{i=0}^{\infty} \frac{(-t)^i}{i!} D^{i+1}(x) - \sum_{i=0}^{\infty} \frac{(-t)^i}{i!} D^{i+1}(x) = 0. \end{aligned}$$

That is, y is a horizontal section of $V \otimes K\langle t/\lambda\rangle$.

If we run this construction over a basis of M, we obtain horizontal sections of $V \otimes K\langle t/\lambda\rangle$ whose reductions modulo t form a basis; they thus form a basis themselves by Nakayama's lemma (and the fact that finite differential modules over a PID are free). This proves the claim. ∎

Corollary 5. *For any finite free differential module M over $K\langle t/\rho\rangle$, $R(M) > 0$.*

We can translate some basic properties of the truncated spectral radius into properties of generic radii, leading to the following analogue of Lemma 2. Alternatively, one can first check Proposition 10 and then simply invoke Lemma 2 itself around a generic point.

Lemma 6. *Let V, V_1, V_2 be finite differential modules over F_ρ.*

(a) For $0 \to V_1 \to V \to V_2 \to 0$ exact,

$$R(V) = \min\{R(V_1), R(V_2)\}.$$

(b) We have

$$R(V^\vee) = R(V).$$

(c) We have

$$R(V_1 \otimes V_2) \geq \min\{R(V_1), R(V_2)\},$$

with equality when $R(V_1) \neq R(V_2)$.

In some situations, it is more natural to consider the *intrinsic generic radius of convergence*, or for short the *intrinsic radius*, defined as

$$IR(V) = \rho^{-1} R(V) = \frac{|d|_{\mathrm{sp},F_\rho}}{|D|_{\mathrm{tsp},V}} \in (0,1].$$

For emphasis, I will sometimes refer to the usual generic radius of convergence as the *extrinsic* generic radius of convergence. The term “intrinsic” is used to connote a certain independence from scale; see Proposition 11. (It also is more natural from the point of view of the definition of truncated spectral norm given by Baldassarri and di Vizio, as described in a previous unit.)

3 Some examples in rank 1

An important class of examples is given as follows. For $\lambda \in K$, let V_λ be the differential module of rank 1 over F_ρ defined by $D(v) = \lambda t^{-1} v$.

Proposition 7. *We have $IR(V_\lambda) = 1$ if and only if $\lambda \in \mathbb{Z}_p$.*

Proof. Exercise. □

Proposition 8. *We have $V_\lambda \cong V_{\lambda'}$ if and only if $\lambda - \lambda' \in \mathbb{Z}$.*

Proof. Note that $V_\lambda \cong V_{\lambda'}$ if and only if $V_{\lambda-\lambda'}$ is trivial, so we may reduce to the case $\lambda' = 0$. By Proposition 7, V_λ is nontrivial whenever $\lambda \notin \mathbb{Z}_p$; by direct inspection, V_λ is trivial whenever $\lambda \in \mathbb{Z}$.

It remains to deduce a contradiction assuming that V_λ is trivial, $\lambda \in \mathbb{Z}_p$, and $\lambda \notin \mathbb{Z}$. There is no harm in enlarging K now, so we may assume that K contains a scalar of norm

ρ; by rescaling, we may reduce to the case $\rho = 1$. We now have $f \in F_1^\times$ such that $t\frac{df}{dt} = \lambda f$; by multiplying by an element of $K^\times$, we can force $|f|_1 = 1$.

Let λ_1 be an integer such that $\lambda \equiv \lambda_1 \pmod{p}$. Then

$$\left|\frac{d(ft^{-\lambda_1})}{dt}\right|_1 = |(\lambda - \lambda_1)ft^{-\lambda_1-1}|_1 \leq p^{-1}.$$

On the other hand, the map d induces the derivation $\frac{d}{dt}$ on the residue field $\kappa_K(t)$ of F_1, and likewise on $\kappa_K((t))$. In the latter, it is clear that the kernel of $\frac{d}{dt}$ is precisely the series with only exponents divisible by p. In particular, if we expand the residue of f as a power series around $t = 0$, it can only have exponents congruent to λ modulo p.

By considering the reduction of f modulo p^n and arguing similarly, we find that the image of f in $\kappa_K((t))$ has only exponents congruent to λ modulo $p^2, p^3, \dots$. But since $\lambda \notin \mathbb{Z}$, this means that the image of f in $\kappa_K((t))$ cannot have any terms at all, contradiction. $\square$

4 Open discs and annuli

Although we have been talking about closed discs so far, it is clearly quite natural to consider open discs. After all, if M is a finite free differential module over $K\langle t/\beta\rangle$ with radius of convergence M, that only guarantees that the local horizontal sections exist on the union of the closed discs of radii strictly less than 1.

We will stick to the following convention, which I'll explain just in the half-open case (since I can consider a fully open disc or annulus as the union of two half-opens pasted together along a common boundary). By a finite differential module M on the half-open annulus $\alpha \leq |t| < \beta$ (which becomes a disc if $\alpha = 0$), let us mean a sequence of finite free differential modules M_i over $K\langle \alpha/t, t/\beta_i\rangle$ with $\beta_1, \beta_2, \dots$ an increasing sequence with limit β, together with isomorphisms $M_{i+1} \otimes K\langle \alpha/t, t/\beta_i\rangle \cong M_i$. In geometric language, this corresponds to a locally free coherent sheaf on the corresponding rigid or Berkovich analytic space, equipped with a connection.

It is unambiguous to refer to the generic radius of convergence $R(M \otimes F_\rho)$ for $\rho \in [\alpha, \beta)$; simply restrict to some interval $[\alpha, \beta_i]$ containing ρ and make the definition there.

5 A cautionary note

Let M be a differential module over $K\langle t/\beta\rangle$ for which $IR(M \otimes F_\beta) = 1$. Then by Theorem 4, for any $\rho \in [0, \beta)$, $M \otimes K\langle t/\rho\rangle$ is trivial, so we have an isomorphism

$$M \otimes K\langle t/\rho\rangle \cong K\langle t/\rho\rangle^{\oplus n}$$

of differential modules.

Now let M be a differential module over $K\langle \alpha/t, t/\beta\rangle$ for which $IR(M \otimes F_\rho) = 1$ for all $\rho \in [\alpha, \beta]$. This most favorable situation was originally thought to be analogous to the situation of regular singularities in the complex setting. In particular, it was believed that for any $\alpha < \gamma \leq \delta < \beta$, it would be possible to write

$$M \otimes K\langle \gamma/t, t/\delta \rangle \cong M_{\lambda_1} \oplus \cdots \oplus M_{\lambda_n}$$

for some $\lambda_1, \ldots, \lambda_n \in \mathbb{Z}_p$, where M_λ is the differential module of rank defined by $D(v) = \lambda t^{-1} v$ (as in the previous section).

This hope was dashed when a counterexample was exhibited by Monsky; it is the rank 2 differential module associated to the differential polynomial $p(1-x)T^2 - xT - a$, where $a \in \mathbb{Z}_p$ is constructed so that

$$\liminf_{m \to +\infty} |a + m|^{1/m} < 1, \qquad \liminf_{m \to +\infty} |a - m|^{1/m} = 1. \tag{2}$$

(The existence of such a is left as an exercise, or see [DR77, §7.20].) I plan to expand on this in a further unit; in the meantime, see [DR77, §7] for further discussion.

What this suggests is the hypothesis on the intrinsic radius needs to be supplemented with some extra hypotheses in order to get the decomposition we want. We will see one such hypothesis later (the existence of a Frobenius structure); in the interim, see the notes for further discussion.

6 Geometric interpretation

You might be wondering why we call $R(V)$ the generic radius of convergence; here is the construction that explains the name. First, the geometric motivation. Say we have $f \in F_\rho$ given by a limit of a Cauchy sequence $f_1, f_2, \ldots$ in $K(t)$. For any complete extension L of K and any $t_\rho \in L$ of norm ρ, we can evaluate the $f_i(t_\rho)$ to get a convergent sequence provided that f_i does not have a pole in the disc $|t - t_\rho| < \rho$.

However, the poles of elements of $K(t)$ belong to F^{alg}. So suppose we take t_ρ to be a *generic point of norm ρ relative to F*, i.e., an element of some larger complete nonarchimedean field such that the disc $|t - t_\rho| < \rho$ contains no elements of F^{alg}. Then we can evaluate each of the $f_i(t_\rho)$ to get a convergent sequence, and thus evaluate $f(t_\rho)$.

Let us turn this geometric thinking into algebra. Let L be the completion of $K(t_\rho)$ for the ρ-Gauss norm. For any $\lambda \in (0, \rho)$, embed $K[t]$ into $L\langle (t - t_\rho)/\lambda \rangle$ by mapping t to $t_\rho + (t - t_\rho)$. Note that the image of f has constant term $f(t_\rho)$, whose norm is strictly greater than that of the remaining terms of the image. Thus each nonzero element of $K[t]$ maps to a unit in $L\langle (t - t_\rho)/\lambda \rangle$, so we can extend to isometric embeddings of $K(t)$ and F_ρ.

By tensoring with each $L\langle (t - \rho)/\lambda \rangle$, we obtain a finite differential module V' on the open disc of radius ρ centered at t_ρ. We can speak of the radius of convergence of this module using the previous definition.

We already made the previous observation, but let us now formalize it.

Lemma 9. *For any $f \in F_\rho$ and any $\lambda < \rho$, $|f|_\rho$ equals the λ-Gauss norm of f within $L\langle (t - t_\rho)/\lambda \rangle$.*

Proof. This holds because the constant term $f(t_\rho)$ is dominant. $\square$

Proposition 10. *The generic radius of convergence of V is equal to the radius of convergence of V'.*

Proof. Let G_λ be the completion of $L(t - t_\rho)$ for the λ-Gauss norm. If we now compute $|D|_{\mathrm{tsp},V}$ in terms of some basis and apply Lemma 9 to every matrix entry in the calculation, we get the same norms whether we work in F_ρ or G_λ. In other words,

$$|D|_{\mathrm{tsp},V\otimes G_\lambda} = \max\{|d|_{\mathrm{sp},G_\lambda}, |D|_{\mathrm{tsp},V}\} = \max\{p^{-1/(p-1)}\lambda^{-1}, |D|_{\mathrm{tsp},V}\}.$$

On one hand, this implies $R(V) \leq R(V')$ by applying Theorem 4 to $V \otimes L\langle(t - t_\rho)/\lambda\rangle$ for a sequence of values of λ converging to ρ.

On the other hand, pick any $\lambda < R(V \otimes L[\![t - t_\rho]\!])$; then $V \otimes G_\lambda$ is a trivial differential module, so the truncated spectral norm of D on it is $p^{-1/(p-1)}\lambda^{-1}$. We thus have

$$|D|_{\mathrm{tsp},V} \leq p^{-1/(p-1)}\lambda^{-1},$$

so $R(V) \geq \lambda$. This yields $R(V) \geq R(V')$. $\square$

Here is an example illustrating both the use of the geometric interpretation and a good transformation property of the intrinsic normalization.

Proposition 11. *Let m be a positive integer coprime to p, and let $f_m : F_\rho \to F_{\rho^m}$ be the map $t \mapsto t^m$. Then for any finite differential module V over F_ρ, $IR(V) = IR(V \otimes F_{\rho^m})$.*

Proof. This follows from the geometric interpretation plus the fact that

$$|t - t_\rho| < c\rho \Leftrightarrow |t^m - t_\rho^m| < c\rho^m \qquad (c \in (0, 1)), \tag{3}$$

whose proof is left as an exercise. $\square$

7 Subsidiary radii

Let V be a finite differential module over F_ρ. Let $V_1, \dots, V_m$ be the Jordan-Hölder constituents of V; that is, take a maximal filtration of V by differential submodules, and let $V_1, \dots, V_m$ be the successive quotients. (It is a standard algebra exercise that the resulting list does not depend on the filtration up to ordering.)

We define the multiset of *generic radii of subsidiary convergence*, or for short *subsidiary radii*, to be the multiset consisting of $R(V_i)$ with multiplicity $\dim V_i$ for $i = 1, \dots, m$. We also have *intrinsic (generic) radii of subsidiary convergence* obtained by multiplying the subsidiary radii by ρ^{-1}. You may want to think of the product of the subsidiary radii (or the intrinsic subsidiary radii) as an analogue of irregularity over $\mathbb{C}((z))$; this analogy will be further supported later.

It is not yet clear how to interpret the subsidiary radii as the radii of convergence of anything. We will give this interpretation in a later unit.

8 Notes

The notion of considering a "generic disc" as above originates in the work of Dwork [Dwo73]. Our definition of the generic radius of convergence is taken from Christol and Dwork [CD94]. The intrinsic radius of convergence was introduced in [Ked07], where it is called the "toric normalization"; it also figures prominently in the coordinate-free treatment of generic radii of convergence given in [BdV07].

A theory of exponents for differential modules on an annulus with intrinsic radius of convergence 1 everywhere was developed by Christol and Mebkhout [CM97, §4–5]; an alternate development was later given by Dwork [Dwo97] (see also [DGS94, §6]). The exponents are elements of the quotient $\mathbb{Z}_p/\mathbb{Z}$; this makes the construction somewhat complicated, as one must use *archimedean* considerations to identify a p-adic number, in a manner we will not elaborate further here. When the differences between exponents satisfy a p-adic non-Liouville condition (that is, they cannot be approximated unusually well by integers), one obtains a decomposition into modules of rank 1 [CM97, §6]. This is automatic in case of a Frobenius structure, as then the set of exponents is invariant under the operation $x \to x^p$ and so all of the exponents are forced to be rational numbers, whose differences are always non-Liouville. See [Loe97] (or a promised later unit) for a detailed exposition.

9 Exercises

1. Prove Proposition 1. (Hint: first prove that $K\langle \alpha/t, t/\beta \rangle$ has no nonzero differential ideals. Then given a finite differential module over $K\langle \alpha/t, t/\beta \rangle$, consider the annihilator of the torsion submodule.)

2. Exhibit an example showing that the cokernel of $\frac{d}{dt}$ on $K\langle \alpha/t, t/\beta \rangle$ is not spanned over K by t^{-1}. That is, antidifferentiation with respect to t is not well-defined.

3. Prove Proposition 3. Optional: give an explicit formula for $IR(V_\lambda)$ in terms of ρ and the minimum distance from λ to an integer.

4. Prove Proposition 7. (Hint: for (a), consider the cases $\lambda \in \mathbb{Z}_p$, $\lambda \in \mathfrak{o}_K - \mathbb{Z}_p$, and $\lambda \notin \mathfrak{o}_K$ separately. For (b), use (a) to reduce to the case $\lambda \in \mathbb{Z}_p$.)

5. Prove that there exists $a \in \mathbb{Z}_p$ satisfying (2).

6. Prove (3).