Differential algebra, functional transcendence, and model theory

Math 512, Fall 2017, UIC

James Freitag

October 15, 2017

0.1 Notational Issues and notes

I am using this area to record some notational issues, some which are needing solutions.

- I have the following issue: in differential algebra (every text essentially including Kaplansky, Ritt, Kolchin, etc), the notation $\{A\}$ is used for the differential radical ideal generated by A. This notation is bad for several reasons, but the most pressing is that we use $\{-\}$ often to denote and define certain sets (which are sometimes then also differential radical ideals in our setting, but sometimes not). The question is what to do - everyone uses $\{-\}$ for differential radical ideals, so that is a rock, and set theory is a hard place. What should we do?
- One option would be to work with a slightly different brace for the radical differential ideal. Something like:

  $$\{\!|-|\!\}$$

  from the `stix` package. I am leaning towards this as a permanent solution, but if there is an alternate suggestion, I'd be open to it. *This is the currently employed notation.* Dave Marker mentions: this notation is a bit of a pain on the board. I agree, but it might be ok to employ the notation only in print, since there is almost always more context for someone sitting in a lecture than turning through a book to a random section.
- Pacing notes: the first four lectures actually took five classroom periods, in large part due to the lecturer/author going slowly or getting stuck at a certain point.

Chapter 1

Course Description

Differential algebra is roughly the study of solution sets of algebraic differential equations. This course will focus on connections to algebraic geometry, functional transcendence, and model theory. There will be three general parts to the course:

The first part will consist of an introduction to differential algebra and the model theory of differential fields. In this part of the course we will cover topics such as the differential Nullestellensatz, connections to Jet spaces, quantifier elimination and elimination of imaginaries. The second part of the course will focus on differential Galois theory. In this part of the course, we will cover topics such as Picard-Vessiot theory, strongly normal extensions, and the inverse differential Galois theory problem. We will also cover applications to functional transcendence questions. The third part of the course will focus on using tools from geometric stability theory to prove differential algebraic results; again, we will have an eye toward applications in functional transcendence. Here, the functions in question will satisfy certain nonlinear differential equations for which there has not been (up to now) a Galois theory sufficient to prove the sort of functional transcendence results in which we are interested. Allowing sufficient time, we will apply these results toward topics around Mordell-Lang, Manin-Mumford, and Andre-Oort type problems.

There are plenty of interesting topics in differential algebraic equations which don't fit into the main theme of this course, but as you might have noticed from the description above, our common threads and guiding interests throughout each of the parts of these notes are questions around functional transcendence and algebraic relations between solutions of a differential equation. This is an exciting and active area of research, and our goal is to get close to some cutting edge results of this theme.

Chapter 2

Differential commutative algebra and differential algebraic geometry

2.1 Differential Rings

In this section we cover the basics of differential ring theory, building towards the Ritt-Raudenbusch basis theorem - the rings we deal with in differential algebra are usually non-Noetherian, but differential polynomial rings satisfy the ascending chain condition on radical differential ideals. This result was proven originally in the analytic context by Ritt [] and later worked out in the abstract setting by Raudenbusch [].

Lecture 1

Definition 2.1.1. A *differential ring* is a pair (R, δ) where R is a commutative ring with unit and $\delta : R \to R$ is a derivation of R, that is, an additive homomorphism such that

$$\delta(ab) = \delta(a)b + a\delta(b)$$

for all $a, b \in R$.

There are many familiar and natural consequences of the definition (basically derivations follow the sort of rules you expect and know well). Here are a few examples:

Exercise 2.1.2. Show that $\delta(a^n) = na^{n-1}\delta(a)$ for all $a \in R$.

Exercise 2.1.3. Show that for units $b \in R$, $\delta(\frac{a}{b}) = \frac{\delta(a)b - a\delta(b)}{b^2}$.

We will sometimes use the notation a' for $\delta(a)$. Generally, we will use lower case letters a, b for elements in a ring or field, and capital letters S, T for subsets. The letter n will always be a natural number.

Definition 2.1.4. We define the *ring of constants* of R:

$$R^\delta := \{a \in R \mid \delta(a) = 0\}.$$

Example 2.1.5. Here is one boring and three interesting examples of δ-rings.

- Let R be any commutative ring. Let $\delta(x) = 0$ for all $x \in R$. This is called the *trivial* derivation.
- The field of rational functions $\mathbb{C}(x)$ with $\delta = \frac{d}{dx}$. Here $R^\delta = \mathbb{C}$.
- The ring of analytic functions $\mathcal{O}(U)$ where $U \subseteq \mathbb{C}$ is open and connected with the usual derivation.
- The ring (actually field) of meromorphic functions $\mathcal{M}er(U)$ where $U \subseteq \mathbb{C}$ is open and connected with the usual derivation.

Exercise 2.1.6. Explain why $\mathcal{O}(U)$ is an integral domain, but the ring of $\mathcal{C}^\infty$ functions on U is not.

Example 2.1.7. Let Λ be a lattice in $\mathbb{C}$. Let

$$\wp(z) = \sum_{\lambda \in \Lambda \setminus \{0\}} \left(\frac{1}{(z-\lambda)^2} - \frac{1}{\lambda^2} \right)$$

be the Weierstrass $\wp$-function associated with Λ. There are g_2, g_3 such that $\delta(\wp(z)^2 = 4\wp(z)^3 - g_2 wp(z) - g_3$. Then $\mathbb{C}(\wp(z), \delta(\wp(z))$ is a differential ring (field).

Definition 2.1.8. A δ-morphism between δ-rings R and S is a morphism of rings $\phi : R \to S$ which commutes with δ.[1] R is a δ-subring of S if the inclusion map is a δ-morphism.

Exercise 2.1.9. Let L/K be a δ-field extension. Let $a \in L^\delta$ with a algebraic over K. Prove that a is algebraic over K^δ.

Exercise 2.1.10. Let $\eta_1, \dots, \eta_n$ be elements in δ-field, K. Show that $\eta_1, \dots, \eta_n$ are linearly dependent over K^δ if and only if the Wronskian of $\eta_1, \dots, \eta_n$ vanishes.[2]

Lemma 2.1.11. *Let R be a δ-ring and let $S \subseteq R$ be a multiplicatively closed subset. There is a δ-ring structure on $S^{-1}R$ so that*

$$R \to S^{-1}R$$

given by $a \mapsto \frac{a}{1}$ is a morphism of δ-rings.

Exercise 2.1.12. Prove the previous result. Use the quotient rule from earlier exercises to extend the derivation.

Definition 2.1.13. An ideal $I \subseteq R$ is called a differential ideal if $\delta(I) \subseteq I$.

[1] There is an abuse of notation since we are using δ for the derivation on R and on S. In practice, this really doesn't ever cause confusion, and we are not going to mention this abuse again.

[2] In fact one can generalize this problem to allow the coefficients not only in K^δ, but in any *δ-complete* δ-variety [4]

Exercise 2.1.14. Describe the δ-ideals of the ring $\mathbb{C}[x]$ with the derivation $\frac{d}{dx}$.

Definition 2.1.15. Let $S \subseteq R$. The δ-ideal generated by S is the smallest δ-ideal of R containing S and is denoted by $[S]$.[3]

Exercise 2.1.16. Let R be the ring of holomorphic functions on $\mathbb{C}^2$ in variables x, y, with the derivation $\delta = \frac{\partial}{\partial y}$. What are the generators (as an *ideal*) of $[xy] \subset R$.

Lemma 2.1.17. *Let R be a δ-ring and let I be a δ-ideal of R. There is a unique structure of a δ-ring on R/I so that $R \to R/I$ is a δ-homomorphism.*

Exercise 2.1.18. Prove the previous Lemma.

Recall the radical of an ideal $\sqrt{I} := \{x \in R \mid x^n \in I\}$. An ideal is *radical* if $x^n \in I$ for some $n \in \mathbb{N}$ implies that $x \in I$. That is, I is radical if $I = \sqrt{I}$.

Lecture 2

Lemma 2.1.19. *Let $I \subset R$ be a radical differential ideal. If $ab \in R$ then $ab' \in I$ and $a'b \in I$.*

Proof. Since I is a differential ideal, $(ab)' \in I$. Thus:

$$(ab)' \cdot (a'b) = (a'b)^2 + (ab')(a'b) \in I,$$

and the second term in the sum is in I by hypothesis. So, $(a'b)^2 \in I$, and since I is radical, we have $a'b \in I$. Finally, $ab' = (ab)' - a'b \in I$, completing the proof. □

Lemma 2.1.20. *Let $I \subset R$ be a radical differential ideal. Let $S \subset R$. Then $I(S) := \{x \in R \mid xS \subset I\}$ is a radical differential ideal of R.*

Proof. That $I(S)$ is an ideal follows readily from the definitions. To see that $I(S)$ is a differential ideal, one can apply Lemma 2.1.19.

Let $x^n \in I(S)$. Then for any $s \in S$, we have that $x^n s^n \in I$, but since I is radical, this implies that $xs \in I$, which implies that $x \in I(S)$. Thus, $I(S)$ is radical.

□

Intersections of radical ideals are radical (and we earlier noted this was true for differential ideals as well), so for any set S, there is a unique smallest radical differential ideal containing S, which we denote by $\{\!\{S\}\!\}$. We shall see, that when R is a Ritt algebra (differential ring containing $\mathbb{Q}$), the smallest differential radical ideal containing S is given by $\sqrt{[S]}$, but in general, it need not be the case that $\sqrt{[S]}$ is even a differential ideal. Here are two examples, each of a slightly different nature, illustrating this point:

[3] To see that $[S]$ exists, note that an intersection of δ-ideals is again a δ-ideal.

Example 2.1.21. Let $R = \mathbb{Z}[x]$ with $\delta = \frac{d}{dx}$. Let $I = (2, x^2)$. Note that $[I] = I$. If $\sqrt{I}$ is a δ-ideal, then since $x \in \sqrt{I}$, $\delta x \in \sqrt{I}$. But, then in the ring R/I, it must be that for some $n \in \mathbb{N}$, $(\delta x)^n = 0$, but this is nonsense, since $\delta x = 1$ in the ring R/I (since this is true in R, and R/I is a nontrivial ring).

The previous example starts in characteristic zero, but the effect of adding 2 to the ideal essentially transitions the example to positive characteristic. Such examples are pervasive in positive characteristic, and are an indication of some of the additional difficulties one encounters in positive characteristic differential algebra.

Lemma 2.1.22. *For any $a \in R$ and $S \subseteq R$, $a\{S\} \subseteq \{aS\}$.*

Proof. The set of all x with the property that $ax \in \{aS\}$ is, by Lemma 2.1.20, a differential radical ideal. But, it contains in particular S, so contains $\{S\}$. □

Lemma 2.1.23. *Let $S, T \subseteq R$. Then*

$$\{S\}\{T\} \subseteq \{ST\}.$$

Proof. The set of those x so that $x\{T\} \subseteq \{ST\}$ contains S by Lemma 2.1.22. It is a radical differential ideal by Lemma 2.1.20 and so it contains $\{S\}$ by definition. □

Lemma 2.1.24. *Let $T \subseteq R$ be a multiplicatively closed subset of a differential ring. Let I be a radical ideal which is maximal with respect to containment among all differential ideals which do not contain any element of T. Then I is prime.*

Proof. Suppose that $ab \in I$, but $a \notin I$ and $b \notin I$. Now $\{I, a\}$ and $\{I, b\}$ are radical ideals which properly contain I. So, there are elements $t_! \in \{I, a\}$ and $t_2 \in \{I, b\}$ which are also in the set T. Now $t_1 t_2 \in \{I, a\}\{I, b\}$ and lies in the set T. We also have, using Lemma 2.1.23 $\{I, a\}\{I, b\} \subseteq I$. This is a contradiction to our hypotheses. □

Theorem 2.1.25. *Any radical differential ideal can be written uniquely as an irredundant intersection of prime differential ideals.*

Proof. Let $I \subset R$ be a radical differential ideal and fix a single element $a \in R \setminus I$. We will be done as long as we can give a prime differential ideal, $\mathfrak{p}_x$ which contains I but does not contain x (because then taking the intersection over all such $\mathfrak{p}_x$, we obtain the ideal I). Take the multiplicative set consisting of the powers of x. This set does not intersect I, because I is a radical ideal and $x \notin I$. So, take $\mathfrak{p}_x$, a radical differential ideal which contains I and which is maximal with respect to containment among the set of all differential radical ideals which contain no element of T. Such an ideal $\mathfrak{p}_x$ is prime by Lemma 2.1.24. We leave the uniqueness as an exercise to the reader. □

Exercise 2.1.26. Show that any prime δ-ideal which is maximal with respect to inclusion is prime.

Definition 2.1.27. We call a δ-ring R a *Ritt algebra* if $\mathbb{Q} \subseteq R$.

Exercise 2.1.28. Show that any derivation δ on $\mathbb{Q}$ is trivial.

Lemma 2.1.29. *Let R be a Ritt algebra. Let I be a δ-ideal. If $a^n \in I$ then $(a')^{2n-1} \in I$.*

Proof. We are going to prove by induction on k that $a^{n-k}(a')^{2k-1} \in I$ for $k = 1, \dots, n$. Of course, when $k = n$, this is precisely the result we are trying to establish. First, the case when $k = 1$: $(a^n)' = na^{n-1}a' \in I$ so $a^{n-1}a' \in I$.

Now, suppose we have the result for some value of k.

$$\delta(a^{n-k}(a')^{2k-1}) = (n-k)a^{n-(k+1)}a'(a')^{2k-1} + a^{n-k}(2k-1)(a')^{2k-2}a''.$$

Multiplying by a', we obtain

$$a'\delta(a^{n-k}(a')^{2k-1}) = (n-k)a^{n-(k+1)}(a')^{2k+1} + a^{n-k}(2k-1)(a')^{2k-1}a''.$$

The second term in the sum is in I, by the inductive hypothesis, so the first term must be, establishing the result.

□

A nearly immediate consequence of the previous Lemma is:

Theorem 2.1.30. *Let R be a Ritt algebra. Let $S \subseteq R$. Then $\{S\} = \sqrt{[S]}$.*

Exercise 2.1.31. Let $R \subset S$ be a δ-rings.

1. Let $I \subseteq S$ be a radical δ-ideal such that $I \cap R$ is a prime δ-ideal, P. Show that there is some prime δ-ideal Q of S containing I such that $Q \cap R = P$. In this case, we say that Q is a prime δ-ideal *extending* P or Q *lies over* P.

2. Suppose that $P_1 \subset R$ is a prime δ-ideal. Show that there is some $Q_1 \subset S$, a prime δ-ideal such that $Q_1 \cap R = P_1$.

3. Let $I \subseteq S$ be a radical δ-ideal such that $I \cap R$ is a prime δ-ideal, P. Suppose that I is the intersection of all prime δ-ideals which extending P. Show that if $ab \in I$ with $a \in R$ and $b \in S$, then $a \in I$ or $b \in I$.

The third part of the above exercise gives a necessary condition for a radical δ-ideal to be the intersection of the prime ideals lying over a given prime ideal in a subring; the next result shows that the condition from the exercise is in fact sufficient.

Theorem 2.1.32. *Let $R \subset S$ be a δ-rings. Let $I \subseteq S$ be a radical δ-ideal such that $I \cap R$ is a prime δ-ideal, P. Suppose that if $ab \in I$ with $a \in R$ and $b \in S$, then $a \in I$ or $b \in I$. Show that I is the intersection of all prime δ-ideals lying over P.*

Proof. Take some $x \in S \setminus I$. We will obtain a prime δ-ideal $Q \supseteq I$ and lying over P such that $x \notin Q$. Let $T := \{ax^n \mid n \in \mathbb{N}, a \in R \setminus P\}$. The set T is multiplicative and disjoint from I, so by Lemma 2.1.24, we can take some prime δ-ideal Q which contains I and does not intersect T. Take some $a \in Q \cap R$. Then $ax \in Q$, and so it must be the case that $a \in P$. Thus, Q extends P. □

Lecture 3

Definition 2.1.33. Let R be a δ-ring. The δ-polynomial ring $R\{y_1, \ldots, y_n\}$ is the polynomial ring in infinitely many variables $y_i^{(j)}$ for $i = 1, \ldots, n$ and $j \in \mathbb{N}$ equipped with the derivation extending δ on R with $\delta y_i^{(j)} = y_i^{(j+1)}$.

Exercise 2.1.34. Let $\mathfrak{I}$ be a differential ideal of a differential polynomial ring over a differential field. Suppose I is generated (as a differential ideal) by linear differential polynomials. Show that either $\mathfrak{I} = R$ or $\mathfrak{I}$ is a prime ideal.

Up until now, our development has essentially mirrored that of classical commutative algebra, but starting with this lecture, things will begin to diverge. There are a number of reasons for this divergence. First and most pressing is one of practicality: the differential ideal generated by a single irreducible differential polynomial is not necessarily prime. This issue is related to what we will eventually see; the approach we take is related to deep and still unresolved difficulties related to the *Ritt problem*.

Exercise 2.1.35. Show that $[(x')^2 - 4x] \subseteq K\{x\}$ for K a δ-field is not prime. Show that $[(x')^2 - 4x^3]$ is prime.

Definition 2.1.36. Let R be a ring and S a multiplicative subset. Let $\mathfrak{A} \subset R$ be an ideal. The *saturation of* $\mathfrak{A}$ *by* S, denoted $\mathfrak{A} : S$ is the ideal given by the elements

$$\{f \in R \mid \exists s \in S \text{ such that } sf \in \mathfrak{A}\}.$$

Exercise 2.1.37. Show that $\mathfrak{A} : S$ is an ideal.

When $S = \{s^n \mid n \in \mathbb{N}\}$ is given by the collection of powers of a single element, we sometimes use the notation $\mathfrak{A} : s^\infty$ in place of $\mathfrak{A} : S$. If it is the case that $\mathfrak{A} = \mathfrak{A} : S$, then one says that $\mathfrak{A}$ is saturated with respect to S.

Lemma 2.1.38. *Let R be a δ-ring, $S \subset R$ a multiplicative subset, and $\mathfrak{A} \subset R$ a δ-ideal. Then $\mathfrak{A} : S$ is a δ-ideal.*

Proof. That $\mathfrak{A} : S$ is an ideal is the previous exercise, so it remains to show that $\mathfrak{A} : S$ is closed under differentiation. If $sa \in \mathfrak{A}$, then $s^2a \in \mathfrak{A}$, so $\delta(s^2a) = 2\delta(s)sa + s^2\delta(a) \in \mathfrak{A}$. It follows that $s^2D(a) \in \mathfrak{A}$, and so $D(a) \in \mathfrak{A} : S$. □

Exercise 2.1.39. Let $I \subset R$ be a differential ideal. Suppose that $a, b \in R$, $a + b = 1$, $a \cdot b \in I$. Show that $[a] \subseteq (a) + I$. Show $(a) + I$ is a differential ideal. Show

$$((a) + I) \cap ((b) + I) = ((a) + I) \cdot ((b) + I)\,.$$

In what follows, fix a differential ring, R; we will be considering the differential polynomial ring $R\{y_1, \ldots, y_n\}$.

Definition 2.1.40. A *ranking* on $\bar{y} = (y_1, \ldots, y_n)$ is a total ordering, $\preceq$, on the set $\Theta\bar{y} = \{\delta^j y_i \,|\, i, j \in \mathbb{N}\}$ such that for any $\theta \in \Theta = \{\delta^j | \, j \in N\}$ and any $u \in \Theta\bar{y}$, $u \prec \theta u$.

Definition 2.1.41. A ranking is *integrated* if for any $\theta_1 y_j, \theta_2 y_i$ there is $\theta \in \Theta$ such that $\theta_2 y_i \prec \theta\theta_1 y_j$. A ranking of order type $\omega, <$ is called *sequential*. A ranking is *orderly* if $\theta_1 y_j \prec \theta_2 y_i$ whenever θ_2 has higher order than θ_1.

Exercise 2.1.42. Let $\prec$ be a ranking.

1. Show that if $\prec$ is orderly, then $\prec$ is integrated. Is the reverse implication true?
2. Show that if $\prec$ is orderly, then $\prec$ is sequential. Is the reverse implication true?
3. Show that if $\prec$ is integrated, then $\prec$ is sequential. Is the reverse implication true?

Now fix a ranking $\prec$ on $\bar{y}$.

Definition 2.1.43. Let $f \in R\{\bar{y}\}$. The leader of f with respect to $\prec$ is the highest ranking element of $\Theta\bar{y}$ which appears in f. The leader of f will be denoted by u_f. Letting $d = deg_{u_f}(f)$, we can write:

$$f = \sum_{i=0}^{d} a_i u_f^i,$$

where the elements a_i are of rank strictly less than u_f. The element a_d is called the *initial* of f. The differential polynomial $\frac{\partial f}{\partial u_f} = \sum_{i=1}^{d} a_i i u_f^{i-1}$ is called the *separant* of f, and is denoted by S_f.

One can extend the ordering $\prec$ on the differential monomials to an pre-ordering on the entire ring $R\{\bar{y}\}$ as follows:

- All elements of R have smaller rank than any element of $R\{\bar{y}\} \setminus R$.
- If $u_f < u_g$ or if $u_f = u_g$ and $deg_{u_f}(f) < deg_{u_g}(g)$ then $f \prec g$.
- Any two elements not satisfying any of the previous conditions have the same rank.

Lemma 2.1.44. *For any $f \in R\{\bar{y}\} \setminus R$ and any $\theta \in \Theta$ of order larger than zero, $\theta f - s_f \theta u_f$ has rank lower than θu_f.*

Exercise 2.1.45. Prove the lemma.

Exercise 2.1.46. If $f \in R\{\bar{y}\}$ and $\delta f \in (f)$, then $f \in R[(\theta\bar{y})^p]$ where $p = char(R)$.

Definition 2.1.47. Fix $f, g \in R\{\bar{y}\}$ and $S \subset R\{\bar{y}\}$. We say that f is *partially reduced* with respect to g if no proper derivative of u_g appears in f. We say that f is *reduced* with respect to g if f is partially reduced and $deg_{u_g}(f) < deg_{u_g}(g)$. We say f is (partially) reduced with respect to S if f is (partially) reduced with respect to each element of S. We say S is autoreduced if each element of S is not in R and each element of S is reduced with respect to the remaining elements of S.

Lemma 2.1.48. *Every autoreduced set is finite.*

We will prove the previous Lemma as part of the next exercise.

Exercise 2.1.49. Consider the set $\mathbb{N}^m \times \{1, \ldots, n\}$ equipped with the partial order $(n_1, \ldots, n_m, k) <_\times (p_1, \ldots, p_m, j)$ if and only if $j = k$ and $n_i < p_i$ for each $i = 1, \ldots, m$. Show that any infinite sequence of elements in the set $\mathbb{N}^m \times \{1, \ldots, n\}$ has an infinite subsequence which is increasing with respect to our partial order (we say that this is a well-quasi-order). Deduce (from a special case of the previous result - pick a certain small value of m) that any autoreduced set in $R\{\bar{y}\}$ with respect to any ranking must be finite.

Lemma 2.1.50. *Let $S \subseteq R\{\bar{y}\}$ be autoreduced and let $f \in R\{\bar{y}\}$. There are $n_s \in \mathbb{N}$ for $s \in S$ and $\hat{f} \in R\{bary\}$ called the partial remainder of f with respect to S with the following properties:*

1. *$\hat{f}$ has rank less than or equal to f.*

2. *$\prod_{a \in S} S_a^{n_a} f - \hat{f} \in [S]$ and $\prod_{a \in S} S_a^{n_a} f - \hat{f}$ can be written as a linear combination (over $R\{\bar{y}\}$ of derivatives θa for $a \in S$ and θu_a is less than or equal to the leader of f.*

Proof. Define $H_S := \prod_{f \in S} I_f S_f$. Now, fix some autoreduced set S. Fix $f \in R\{\bar{y}\}$. We now define the *partial remainder of f, denoted $\hat{f}$ with respect to S* along with the natural numbers n_a for each $a \in S$ so that $\prod_{a \in S} S_a^{n_a} f - \hat{f} \in [S]$.

If f is partially reduced with respect to S, then $\hat{f} = f$ and $n_a = 0$ for all $a \in S$. Now, suppose that f involves some derivative of u_a for $a \in S$. Now we will give the definition inductively on the rank of the highest differential monomial v of f with the property that v is the derivative of a leader of an element of f. Let u_a denote the highest leader of which v is a derivative. Let θ be such that $\theta u_a = v$. Then by Lemma 2.1.44 $S_a v = g + \theta a$ with g of rank less than v. Let $e = deg_v(f)$, and write $f = \sum_{i<e} t_i v^i$ where t_i are free from v. Then $S_a^e f = \sum_{i<e} S_a^{e-i} t_i (S_a v)^i$, and

$$\sum_{i<e} S_a^{e-i} t_i (S_a v)^i \equiv \sum_{i<e} S_a^{e-i} t_i (g)^i \pmod{\theta a}.$$

Now, the differential polynomial $f_a = \sum_{0 \le i \le e} S_a^{e-i} t_i g^i$ involves only derivatives of u_a of rank lower than v and it is also clear that the rank of f_a of rank lower than or equal to f. So, by induction, the partial remainder of f_a and natural numbers m_s for $s \in S$ for f_a with the desired properties exist by induction. Define $\hat{f} = \hat{f}_a$, $n_a = m_a + e$, and for $s \in S \setminus \{a\}$, $n_s = m_s$. It is routine to verify that $\hat{f}$ and the given natural numbers verify the hypotheses of the Lemma. □

Exercise 2.1.51. Let $f_1 \ldots f_d \in R\{\bar{y}\}$ and $A \subseteq R\{\bar{y}\}$ an autoreduced set. Show that there are $g_1, \ldots, g_d$ and some natural numbers t_a for $a \in A$ such that $\prod_{a \in A} S_a^{t_a} \cdot f_i - g_i \in [A]$ for each $i = 1, \ldots, d$.

Exercise 2.1.52. (Clairaut's equation) Consider $f = y - xy' - \frac{1}{4}(y')^2 \in \mathbb{C}(x)\{y\}$ with $\delta = \frac{d}{dx}$ on $\mathbb{C}(x)$. Calculate $\delta(P)$. Show that neither S_P nor y'' are in $[P]$. Explain why $[P]$ is not prime.

Lecture 4

A slightly different procedure allows one to obtain the *remainder* of f with respect to an autoreduced set A. While the partial remainder is partially reduced with respect to the autoreduced set, the remainder will be reduced. We will not cover the full details of this algorithmic procedure, but it can be found in various sources, e.g. [14, Page 6]

Theorem 2.1.53. *Let $A = \{f_1, \dots, f_n\}$ be autoreduced with elements ordered by increasing ranks of their leaders and let $g \in R\{\bar{y}\}$. Then there are natural numbers $r_1, \dots, r_n$ and $s_1, \dots, s_n$ such that when a certain linear combination of the elements of A and their derivatives is subtracted from*

$$S_{f_1}^{s_1} \dots S_{f_n}^{s_n} I_{f_1}^{r_1} \dots I_{f_n}^{r_n} g,$$

the result, $\bar{f}$ is reduced with respect to A.

Exercise 2.1.54. Prove the previous result by first applying the procedure of partial reduction from the previous lecture and then doing purely algebraic operations to make the partial remainder reduced.

Corollary 2.1.55. *Let $A = \{f_1, \dots, f_n\}$ be autoreduced $g \in R\{\bar{y}\}$. Then there are natural numbers $r_1, \dots, r_n$ and $s_1, \dots, s_n$ and some differential polynomial g_0 which is reduced with respect to A such that*

$$S_{f_1}^{s_1} \dots S_{f_n}^{s_n} I_{f_1}^{r_1} \dots I_{f_n}^{r_n} g - g_0 \in [A].$$

Next, lets see that the reduction procedure shows that the colon ideal alleviates that annoying property that the differential ideal generated by a single irreducible differential polynomial is not necessarily prime.

Corollary 2.1.56. *Take K to be a Ritt algebra and I a nonzero δ-ideal in $K\{y\}$. Then there is some minimal rank differential polynomial p such that*

$$[p] \subset I \subset [p] : (S_P I_P)^\infty.$$

Further, if I is prime then $I = [p] : (S_P I_P)^\infty$.

Proof. Take p to be a δ-polynomial in I which has minimal rank among all such δ-polynomials. It is clear that $[p] \subset I$, and so we need only establish the other containment. Take some other $q \in I$. By Corollary 2.1.55, there are r, s such that $S_p^s I_p^r q - q_0 \in [p]$, where q_0 is of smaller rank than p. But since $S_p^s I_p^r q - q_0 \in [p]$, it must be that q_0 is an element of I, but as p has minimal rank already, it must be that $q_0 = 0$. Thus, q must be in the ideal $[p] : (S_P I_P)^\infty$.

To see the final statement of the Corollary, note that if I is prime, then since by our choice of p, S_p and I_p are not elements of I, if $g \in [p] : (S_P I_P)^\infty$, then $g \cdot S_P^s I_P^r \in [p] \subset I$, and so by primality, it must already be the case that $g \in I$. □

Exercise 2.1.57. Suppose that $p \in K\{y\}$ is an irreducible. Show that $[p] : (S_P I_P)^\infty$ is a prime ideal.

The polynomial p is an example of what is called a *characteristic set* of a differential ideal. In order to develop the notion in general, we need to extend our ranking to collections of autoreduced differential polynomials. Again, we have fixed a ranking $\prec$ on the differential monomials, which we extended to differential polynomials (at that point it is only a quasiorder). Given two autoreduced sets $A = \{f_1, \ldots, f_r\}$ and $B = \{g_1, \ldots, g_s\}$ whose elements are written in ascending order by the rank of their leaders, we say that A is of lower rank than B, writing $A \prec B$ if either:

1. For some $j \leq r, s$, we have that f_i and g_i are of the same rank whenever $i < j$ and $f_j \prec g_j$.
2. We have that $r > s$ and f_i and g_i are of the same rank when $i \leq s$.

Exercise 2.1.58. Let $\mathfrak{I} \subset K\{y_1, \ldots, y_n\}$ be a differential ideal. Show that the collection of autoreduced subsets A whose elements are in the ideal $\mathfrak{I}$ is well-(quasi-)ordered by rank.

Deduce that there are minimal rank autoreduced sets among this collection. Such sets are called *characteristic sets* of $\mathfrak{I}$.

Corollary 2.1.59. *Let $\mathfrak{p}$ be a prime differential ideal of $K\{\bar{y}\}$. Let $A = \{f_1, \ldots, f_r\}$ be a characteristic set of $\mathfrak{p}$. Then $\mathfrak{p} = [f_1, \ldots, f_r] : (S_{f_1} \ldots S_{f_r} I_{f_1} \ldots I_{f_r})^\infty$.*

The utility of working with colon ideals also comes from the next two results, which shows that various algebraic properties of the differential ideal can be deduced from a simpler algebraic counterpart. For any set of differential polynomials, A, we let H_A denote the product of the separants of the elements of A.

Lemma 2.1.60. *Let $A\{a_1, \ldots, a_k\}$ be an autoreduced set in $K\{y_1, \ldots, y_n\}$. Then any element $f \in K\{y_1, \ldots, y_n\}$ which is partially reduced with respect to A and is in the differential ideal $[A] : H_A^\infty$ is actually an element of the ideal $(A) : H_A^\infty$.*

Proof. Let $f \in [A] : H_A^\infty$ be partially reduced. Write $H_A^k f$ for some k as a linear combination (over $K\{\bar{y}\}$) of elements of A and their derivatives:

$$H_A^k f = \sum_{1 \leq i \leq r} f_i \theta_i A_{j_i} + \sum g_i h_i$$

where $\theta_i \in \Theta$ of positive order, $f_i \in K\{\bar{y}\}$, $A_{j_i} \in A$, $g_i \in K\{\bar{y}\}$ is partially reduced with respect to A and $h_i \in (A)$. Now, if the first sum in the above expression can be taken to be zero, then we will have established the result.

Suppose that this is not the case; and suppose that v is the largest rank element among $\theta_i u_{A_{j_i}}$ among the elements in the sum (without loss of generality, lets suppose that $v = \theta_r u_{A_{j_r}}$). Suppose that we take an expression as above in which v has lowest rank among all such expressions. Now, we rewrite the expression:

$$\sum_{1 \leq i \leq r-1} f_i \theta_i A_{j_i} + f_r \theta_r A_{j_r} + \sum g_i h_i,$$

where for $1 \leq i \leq r-1$, we have that $\theta_i A_{j_i} \prec \theta_r A_{j_r}$. In the above expression, note that $\theta_r A_{j_r} = S_{A_{j_r}} v + b$ where b is a differential polynomial of lower rank than v. So, we have

$$\sum_{1 \leq i \leq r-1} f_i \theta_i A_{j_i} + f_r \left(S_{A_{j_r}} v + b \right) + \sum g_i h_i.$$

Now, in the previous equation, treat the differential monomial v as a variable, and note that it does not appear on the left hand side, while on the right hand side it may appear in the f_i's. Substitute $\frac{-b}{S_{A_{j_r}}}$ for v, and after clearing the $S_{A_{j_r}}$ by multiplying both sides of the expression by a suitable power of H_A, we obtain an expression in which the largest ranking derivative of a leader of an element of A in the linear combination is strictly less than v, a contradiction to the minimality of the originally chosen expression. □

Lemma 2.1.61. *Let A be an autoreduced set in $K\{y_1, \ldots, y_n\}$. Then $[A] : H_A^\infty$ is prime if and only if $(A) : H_A^\infty$ is prime.*

Proof. Let $f, g \in K\{\bar{y}\}$ such that $fg \in [A] : H_A^\infty$. Let f_0 and g_0 be the remainders of f, g with respect to A, respectively. Then $f_0 g_0$ is partially reduced with respect to A and $f_0 g_0 \in [A] : H_A^\infty$, so by Lemma 2.1.60 we have that $f_0 g_0 \in (A) : H_A^\infty$. □

Lecture 5

Throughout, let K be a characteristic zero differential field, and $R = K\{\bar{y}\}$. One major difference $K\{\bar{y}\}$ and $K[\bar{y}]$ is that $K\{\bar{y}\}$ is non-Noetherian. It is even true that $K\{\bar{y}\}$ does not satisfy the ascending chain condition (every ascending chain is finite) on *differential ideals*:

Exercise 2.1.62. Give an example of a strictly ascending chain of differential ideals. (*Hint:* Consider using differential polynomials from the set $\{(\delta^n y)^2\}_{n \in \mathbb{N}}$.

Nevertheless, not everything is lost. We will show that $K\{\bar{y}\}$ has the ascending chain condition on differential radical ideals in this lecture. Though we are using the language of differential rings, this result answers a fundamental question which was asked long before the formalism of differential ring theory was invented. Given as system of infinitely many differential equations (in finitely many variables), is there a finite subsystem of differential equations whose vanishing implies the vanishing of the entire system? This sort of question goes back (at least) to the late 1800s (Drach, Picard).

Exercise 2.1.63. Let R be a δ-ring. Show that the following conditions are equivalent:

1. Let I in R be a δ-radical ideal. Then there is a finite set $A \subset I$ such that $I = \{\!\{A\}\!\}$. Such as set A is sometimes called a *basis of I*.

2. R satisfies the ascending chain condition on δ-radical ideals.

3. Any nonempty set of δ-radical ideals contains a maximal element with respect to inclusion.

Exercise 2.1.64. Let R be a δ-ring satisfying any of the three conditions of the previous exercise. Let $I \subset R$ be a radical δ-ideal. Show that there is a bijection between the radical δ-ideals of R which contain I and the radical δ-ideals of R/I. Let S be a multiplicative subset of R. Show that there is a bijection between the radical δ-ideals of $S^{-1}R$ and the radical δ-ideals of R which contain no element of S.

A δ-ring satisfying the ascending chain condition on δ-radical ideals is sometimes called *Rittian*.

Theorem 2.1.65. *(Ritt-Raudenbush basis theorem) Let R be a* Rittian *δ-ring which is a Ritt algebra. Then if I is a radical δ-ideal in $R\{\bar{y}\}$, then I has a finite basis.*

Proof. It is enough to show the statement when $n = 1$. Suppose that the theorem is false. Then there are radical differential ideals which have no finite basis; call this collection $\mathcal{B}$. We claim that there are in the collection $\mathcal{B}$ some elements which are maximal with respect to containment. To see this, it suffices to show that any chain of ideals in $\mathcal{B}$ has an upper bound in $\mathcal{B}$.

Take some infinite chain $I_1 \subset I_2 \subset I_3 \subset \ldots$. Now, the union $I = \cup_{i \in \mathbb{N}} I_i$ is a radical differential ideal, and we claim it must be in the collection $\mathcal{B}$. Suppose not; then if I has finite basis given by B, it must be that B is contained I_k for some $k \in \mathbb{N}$. But then $I = \{\!\{B\}\!\} \subset I_k \subset I$, and so the chain must have stabilized. Now, fix some maximal element of $\mathbb{B}$, $\mathfrak{p}$. We claim that $\mathfrak{p}$ is a prime differential ideal. Suppose not, and take some $a, b \in R\{y\}$ with $a, b \notin \mathfrak{p}$, but $ab \in \mathfrak{p}$. Both of the ideals $\{\!\{a, \mathfrak{p}\}\!\}$ and $\{\!\{b, \mathfrak{p}\}\!\}$, containing $\mathfrak{p}$ properly, must not be elements of $\mathcal{B}$, and thus must have finite bases.

We claim that $\{\!\{a, \mathfrak{p}\}\!\}$ has a basis of the form $\{a, f_1, \ldots, f_k\}$, where $f_i \in \mathfrak{p}$. Suppose that we have an arbitrary basis $\{c_1, \ldots, c_n\}$ of $\{\!\{a, \mathfrak{p}\}\!\}$. Then $c_i \in \{\!\{a, \mathfrak{p}\}\!\}$, so $c_i^{n_i} \in [a, \mathfrak{p}]$, for some n_i so there must be some $b_{i,1}, \ldots, b_{i,k_i} \in \mathfrak{p}$ such that $c_i^{n_i} \in [a, b_{i,1}, \ldots, b_{i,k_i}]$, and so $c_i \in \{\!\{a, b_{i,1}, \ldots, b_{i,k_i}\}\!\}$. Now, the collection $\cup_{i=1}^{n} \{b_{i,1}, \ldots, b_{i,n_i}\}$ together with a gives a basis of the desired form. Similarly, take $\{b, g_1, \ldots, g_m\}$ to be a basis of $\{\!\{b, \mathfrak{p}\}\!\}$. Now, note that

$$\{\!\{a, \mathfrak{p}\}\!\}\{\!\{b, \mathfrak{p}\}\!\} \subset \{\!\{C\}\!\},$$

where C is the set of products of elements fg where $f \in \{a, f_1, \ldots, f_k\}$ and $g \in \{b, g_1, \ldots, g_m\}$. But for arbitrary $d \in \mathfrak{p}$, we have $d^2 \in \mathfrak{p}$, and $d \in \{\!\{a, \mathfrak{p}\}\!\}$, $d \in \{\!\{b, \mathfrak{p}\}\!\}$ so it follows that $d \in \{\!\{C\}\!\}$ and thus $\{\!\{C\}\!\} = \mathfrak{p}$, that is C is a finite basis of $\mathfrak{p}$, a contradiction to our assumption that $\mathfrak{p}$ has no finite basis. Thus, it must be that $\mathfrak{p}$ is prime.

The ideal $\mathfrak{p} \cap R$ is radical, and so by our assumption has a finite basis as does $Q = \{\!\{\mathfrak{p} \cap R\}\!\} \subset R\{y\}$. Now we have $Q \subset \mathfrak{p}$. Take some element $f \in \mathfrak{p} \setminus Q$ which has minimal rank.

Exercise 2.1.66. Show that the initial and separant of f must not be in $\mathfrak{p}$. (*Hint:* if it is the case that either of these elements is in $\mathfrak{p}$, construct an element of $\mathfrak{p} \setminus Q$ with lower rank.)

Now we have that $\mathfrak{p} \subset \{\!\{\mathfrak{p}, S_f I_f\}\!\}$ is a proper containment, so $\{\!\{\mathfrak{p}, S_f I_f\}\!\}$ must not be in $\mathcal{B}$, and thus must have a finite basis. Now take $g \in \mathfrak{p}$. There are $n_1, n_2 \in \mathbb{N}$ an some $h \prec f$ so that $S_f^{n_1} I_f^{n_2} g - h \in [f]$ by Theorem 2.1.53. So, we have that $h \in \mathfrak{p}$, and has rank less than f. By our assumptions, then $h \in \{\!\{\mathfrak{p} \cap R\}\!\} \subset R\{y\}$. But now an arbitrary element $g \in \mathfrak{p}$ has the property that $S_f I_f g \in \{\!\{\{\!\{\mathfrak{p} \cap R\}\!\}, f\}\!\}$. Now take some $h_1, \ldots, h_r$ such that $\{\!\{S_f I_f, \mathfrak{p}\}\!\} = \{\!\{S_f I_f, h_1, \ldots, h_r\}\!\}$. Now we have

$$\mathfrak{p}^2 \subset \mathfrak{p}\{\!\{S_f I_f, \mathfrak{p}\}\!\} \subset \{\!\{S_f I_f \mathfrak{p}, \mathfrak{p}h_1, \ldots, \mathfrak{p}h_r\}\!\} \subset \{\!\{\{\!\{\mathfrak{p} \cap R\}\!\}, f, h_1, \ldots, h_r\}\!\}.$$

But now $\mathfrak{p}$ has a finite basis because $\{\!\{\mathfrak{p} \cap R\}\!\}$ does. □

Of main interest for us is the particular case of the previous theorem:

Corollary 2.1.67. *Let K be a differential field. Then every radical δ-ideal of $K\{\bar{y}\}$ has a finite basis.*

Exercise 2.1.68. Let K be a differential field. Let $I \subset K\{y\}$ be the δ-ideal generated by $\delta^i(y)\delta^j(y)$ for $(i, j) \in \mathbb{N}^2$. Show that $\{y\}$ is a basis for the ideal $\{\!\{I\}\!\}$.

Exercise 2.1.69. Given a differential field K, calculate the cardinality of the set of prime differential ideals of $K\{\bar{y}\}$ in terms of $|K|$.

Notice, in our earlier proof that every radical differential ideal is the intersection of prime differential ideals (Theorem 2.1.25), we appealed to Zorn's lemma. In the case that we are working with a Ring with the ACC on differential radical ideals, such an appeal is not required, since every such chain is finite. It is, therefore, reasonable to ask if there is an algorithmic procedure for calculating the prime decomposition of a radical differential ideal. Whether or not such a procedure exists in general is an open problem, now known as the *Ritt problem*. There is a detailed computational literature devoted to studying the effectiveness of various aspects of the Ritt problem, which we will now elaborate on. Formally speaking, here is the problem we are discussing:

Question 2.1.70. *Given a finite set of differential polynomials $B \subseteq R\{\bar{y}\}$, is there an algorithm which allows one to write $\{\!\{B\}\!\}$ as an intersection of $\mathfrak{p}_1 \cap \ldots \cap \mathfrak{p}_n$ of unique prime differential ideals with the property that for no i, j do we have $\mathfrak{p}_i \subset \mathfrak{p}_j$?*

One could split, in light of our approach via characteristic sets given above into two separate questions, the first of which is:

Question 2.1.71. *Can one find a finite collection of autoreduced sets $A_1, \ldots, A_n$ such that each is the characteristic set of a prime differential ideal which contains $\{\!\{B\}\!\}$ and each prime component of $\{\!\{B\}\!\}$, $\mathfrak{p}$, has a characteristic set from among the collection $A_1, \ldots, A_n$?*

In [8], the previous question is answered affirmatively, though in quite a terse manner. A slower and more constructive explanation is given in [15].

The previous question alone is not enough to answer the Ritt problem, but if additionally one could answer the next question, then an answer to the Ritt problem would follow:

Question 2.1.72. *Given an autoreduced set, A, is there an algorithm to decide if it is the characteristic set of some prime component of $\{B\}$?*

The previous question is open except in special cases. Because there is an algorithm to decide if the prime ideals associated with two characteristic sets are identical, the previous question is equivalent to:

Question 2.1.73. *Given charactisteric sets A_1, A_2 of prime differential ideals $\mathfrak{p}_1, \mathfrak{p}_2$, is there an algorithm to decide if $\mathfrak{p}_1 \subset \mathfrak{p}_2$?*

There has not been much progress on answering the question, but a variety of works have investigated the Ritt problem, for instance[6, 5]. The one notable exception in which there has been progress on the Ritt problem is given by Ritt's Low Power Theorem (see [8]), a topic to which we will return later.

2.2 Differential fields

Lecture 6

The next two lectures follow [12] almost verbatim.

For the remainder of the section, we let K be a differential field of characteristic zero. Suppose $\bar{a} \in L$, a differential field extending K. Then, by $K\langle \bar{a} \rangle$, we mean the differential field generated by $\bar{a}$ over K, the smallest differential subfield of L containing K and $\bar{a}$.

Exercise 2.2.1. Let $p(y) \in K\{y\}$ be an irreducible differential polynomial in a single variable. Show that if $g \in [p] : H_p^\infty$ then the order of g is at least that of p. Show that if the orders are the same, then p must divide g.

The next result shows how prime differential ideals (at least in one variable) behave under extension of the base field. This result is a special case of [8, Proposition 3, page 131], a result we will prove once we have the necessary terminology.

Lemma 2.2.2. *Let $f \in K\{y\}$ be irreducible of order n. Let f_1 be an irreducible factor of $f \in L\{y\}$. Then $I = [f] : H_f^\infty \subset K\{y\}$ is equal to $[f_1] : H_{f_1}^\infty \cap K\{y\}$.*

Proof. Say f is of order n. Then if f factors in $L\{x\}$ then we claim that any irreducible factor must be of order n. In general, any irreducible polynomial $p(y_1, \dots, y_n) \in K[y_1, \dots, y_n]$ in which y_i appears has the property that over L, any irreducible factor of p must have y_i appear. To see this note that each of the automorphism conjugates of a generic solution of $p(\bar{y})$ over K are by quanifier elimination in ACF precisely the generic solutions of the zero sets of the irreducible factors of $p(\bar{y})$ over K^{alg}. Each of these conjugates has the same type over K, and thus y_i must appear in the corresponding irreducible factor.

Thus f_1 is of order n. Let $g \in [f_1] : H_{f_1}^\infty$. Take $\bar{g}$ to be the partial remainder of g with respect to f. Then f_1 divides g_1. But now $g_1 \in K\{x\}$, and so it is the case that all of the conjugates of f_1 over K must also divide g_1 (each of them has coefficient tuple with the same type over K). But, this means that f must divide g_1, and this establishes that $[f_1] : H_{f_1}^{\infty} \cap K\{x\} \subset [f] : H_f^{\infty}$. We leave establishing the reverse containment as an exercise to the reader. □

Definition 2.2.3. The theory of *differentially closed fields of characteristic zero* is the theory in the language of differential rings given by:

1. ACF_0, the axioms for algebraically closed fields of characteristic zero.

2. δ is a derivation.

3. For any non-constant differential polynomials $f(x)$ and $g(x)$ in which the order of g is less than the order of f, there is some x such that $f(x) = 0 \wedge g(x) \neq 0$.

We will abbreviate the theory DCF.[4]

Exercise 2.2.4. Show that if $a \in L$ is algebraic over the field of constants of K, then $a \in L^{\delta}$.

Lemma 2.2.5. *Every differential field K of characteristic zero is contained in a differentially closed field.*

Proof. Let f be of order n and g be of order strictly less than n. Then take some irreducible factor f_1 of f which is of order n. Now, by Exercise 2.2.1, we have that $g \notin [f_1] : H_{f_1}^{\infty}$ and $[f_1] : H_{f_1}^{\infty}$ is prime. Take L to be the fraction field of the ring $K\{x\}/[f_1] : H_{f_1}^{\infty}$. Then taking a to be the image of x in the quotient map, we have that $f(a) = 0$ and $g(a) \neq 0$. Iterate this construction. □

Lemma 2.2.6. *Let K, L be ω-saturated models of DCF. Let $\bar{a} \in K, \bar{b} \in L$. Let $k = \mathbb{Q}\langle \bar{a} \rangle$ and $l = \mathbb{Q}\langle \bar{b} \rangle$. Let $\sigma : k \to l$ be an isomorphism such that $\sigma(\bar{a}) = \bar{b}$. For all $\alpha \in K$, there is an extension of σ to an isomorphism mapping $k\langle \alpha \rangle$ into L.*

Proof. Take some $\alpha \in K$. There are two cases to consider: 1) α is differentially algebraic over k and 2) α differentially transcendental over k. We first consider case 1). Let f belong to the ideal of differential polynomials over K which vanish at α such that f has minimal rank among all such differential polynomials. Suppose that the order of f is n and let g be the image of the differential polynomial under σ. Now take

$$\Gamma(v) := \{g(v) = 0\} \cup \{h(v) \neq 0 \,|\, h \text{ of order} < n \text{ in } l\{x\}\ \}.$$

Then $\Gamma(v)$ is finitely satisfiable in L, by Lemma 2.2.5. So, $\Gamma(v)$ is satisfiable in L, by saturation. Let β realize Γ, and extend the isomorphism by setting $\sigma(\alpha) = \beta$.

Case 2) is easier. Extend σ to α by sending α to any differential transcendental over l, which exists by the saturation hypothesis.

□

[4] Differentially closed fields of characteristic p also exist, and if we refer to that theory, we write DCF_p.

We will use a standard quantifier elimination test given in the next exercises, proofs of which could be found in various textbooks (e.g. [11]).

Exercise 2.2.7. Suppose that every quantifier free $\mathcal{L}$-formula $\theta(\bar{v}, w)$ has an associated quantifier free formula $\psi(\bar{v})$ such that $T \models \forall \bar{v}(\exists w \theta(\bar{v}, w) \leftrightarrow \psi(\bar{v})$. Then show that T has quantifier elimination.

Exercise 2.2.8. Suppose that $\mathcal{L}$ is a language with at least one constant. Let $\phi(\bar{v})$ be an $\mathcal{L}$-formula with free variables $\bar{v}$. Then the following are equivalent:

1. There is a quantifier free $\mathcal{L}$-formula $\psi(\bar{v})$ so that $T \models \forall \bar{v}(\phi(\bar{v}) \leftrightarrow psi(\bar{v})$.

2. If $\mathcal{M}, \mathcal{N}$ are models of T and $\mathcal{C}$ is a common substructure, then for all $\bar{a} \in \mathcal{C}$, $\mathcal{M} \models \phi(\bar{a})$ if and only if $\mathcal{N} \models \phi(\bar{a})$.

Theorem 2.2.9. *DCF has elimination of quantifiers.*

Proof. Fix $K, L \models$ DCF with k a subfield of K, L, $\bar{a} \in k$, and $b \in K$. Suppose $\phi(v, \bar{w})$ is quantifier free and suppose that $K \models \phi(b, \bar{a})$. In light of the two previous exercises, we need only show that there is some $d \in L$ so that $L \models \phi(d, \bar{a})$. The existence of such a d is invariant under elementary extensions, so we can assume L is ω-saturated. But then by Lemma 2.2.6 we have some $d \in L$ such that $k\langle b \rangle \cong k\langle d \rangle$. So, we must have $L \models \phi(d, \bar{a})$. □

Exercise 2.2.10. DCF is complete and model complete.

Lecture 7

The next result is the analog of Hilbert's Nullstellensatz for differential fields:

Corollary 2.2.11. *If k is a differential field and Σ is a finite collection of differential equations and inequations over k which has a solution in some differential field extension L extending k, then Σ has a solution in any differentially closed $K \supset k$.*

Exercise 2.2.12. Give a proof of the previous result using quantifier elimination.

Exercise 2.2.13. There is a bijection from the space of n-types over k to prime differential ideals induced by $p \mapsto I_p = \{f \in k\{x\} \mid f(x) = 0 \in p(x)\}$. Prove that DCF is ω-stable.

Let $\mathcal{M} \models T$ be saturated enough. Let p be a type over $\mathcal{M}$. B is a canonical base for p if for any σ an automorphism of $\mathcal{M}$, σ fixes p setwise if and only if σ fixes B pointwise.

Lemma 2.2.14. *If B is a canonical base for $\phi(v, a)$ then there is some formula $\psi(v, w)$ and some $b \in B$ so that $\phi(v, a) \leftrightarrow \psi(v, b)$ and for all $b \neq b' \in B$, $\phi(v, a) \not\leftrightarrow \psi(v, b)$.*

Proof. Let

$$\Gamma(v) = \{\psi(v) \mid \psi \text{ has parameters from } B \text{ and } \phi(v,a) \to \psi(v)\}.$$

We claim that $\Gamma(v) \to \phi(v,a)$. If not, then there is, by saturation, some $c \in \mathcal{M}$ such that $\Gamma(c)$ and $\neg\phi(c,a)$. Whenever the types of c and c' over B are the same, then again by saturation, there is some $\sigma \in Aut(\mathcal{M}/B)$ which sends c to c'. We also then have $\neg\phi(c',a)$, since B is assumed to be a canonical base, and so $tp(c/B) \to \neg\phi(v,a)$. By compactness, take some formula $\theta(v) \in tp(c/B)$ such that $\theta(v) \to \neg\phi(v,a)$. Now $\neg\theta(v) \in \Gamma(v)$, a contradiction. Thus, we must have $\Gamma(v) \to \phi(v,a)$. Again by compactness, we take some $\chi(v,b) \in \Gamma(v)$ such that $\chi(v,b) \to \phi(v,a)$. By construction, the reverse implication holds, since all of the elements of Γ were taken to be consequences of ϕ.

When b and b' have the same type over the empty set, there is an automorphism taking b to b'. Such an automorphism cannot, by the definition of canonical bases preserve the set of realizations of $\chi(x,b)$. This establishes the Lemma. □

Definition 2.2.15. A theory T has elimination of imaginaries if every formula $\phi(v,a)$ has a canonical base.

Lemma 2.2.16. *Suppose T has elimination of imaginaries and at least two constants. Let $\mathcal{M} \models T$ and let E be a $\emptyset$-definable equivalence relation on $\mathcal{M}^n$ then there is an $\emptyset$-definable $f : \mathcal{M}^n \to \mathcal{M}^m$ such that $E(x,y)$ if and only if $f(x) = f(y)$.*

Proof. By EI, and the previous lemma, given any formula $\phi(x,a)$, there is some formula, $\psi_a(v,w)$ and some unique b so that $\phi(v,a) \leftrightarrow \psi_a(v,b)$. As the value of the parameter a varies, the family of formulas covers all of the domain of the equivalene relation. The by compactness, finitely many formulas must suffice, call them $\psi_1(v,b_1), \ldots, = \psi_n(v,b_k)$. Now using standard coding, we can assume there is a single such formula, ψ. Let the map f be given by $a \mapsto b$ where b is such that $\phi(v,a) \leftrightarrow \psi(v,b)$.

□

We say that B is a canonical base for a set of types if every automorphism permutes the set of types if and only if it fixes the set B pointwise.

Lemma 2.2.17. *Let T be an ω-stable theory and let $\mathcal{M} \models T$ be saturated enough. If every set of conjugate complete types over $\mathcal{M}$ has a canonical base then T admits elimination of imaginaries.*

Proof. Given a formula $\phi(x,y)$, we define $E_\phi(y,z)$ to hold when $\forall x, \phi(x,y) \leftrightarrow \phi(y,z)$. We claim that for any automorphism σ of $\mathcal{M}$, σ fixes the set $\phi(x,a)$ (setwise) if and only if it preserves the E_ϕ-class of a. First, suppose the E_ϕ-class is preserved. Then $\sigma(a)$ is equivalent to a and $\{x \mid \phi(x,\sigma(a))\} = \{x \mid \phi(x,a)\}$. Now, suppose that σ fixes $\{x \mid \phi(x,a)\}$ setwise. Consider the image of a under σ. We claim that this image must be in the E_ϕ-class of a. To see this, note that

$$\sigma(\{x \mid \phi(x,a)\}) = \{\sigma(x) \mid \phi(x,a)\},$$

but by elementarity of the map σ, we have that

$$\{\sigma(x) \mid \phi(x,a)\} = \{x \mid \phi(\sigma^{-1}(x),a)\} = \{x \mid \phi(x,\sigma(a))\}.$$

Now, take $p_1, \ldots, p_n$ to be global types of maximal rank which contain $E_\phi(y,a)$. There are finitely many conjugacy classes of $p_1, \ldots, p_n$ and by hypothesis, for each conjugacy class, we can find a canonical base. Take A to be the canonical base of one conjugacy class, call it $p_1, \ldots, p_k$.

Any automorphism permutes the collection of types $p_1, \ldots, p_k$ if and only if it fixes A pointwise. Any automorphism permutes $p_1, \ldots, p_k$ if and only if it fixes the E_ϕ-class of a, because we know that the image of the E_ϕ-class of a under automorphism is the E_ϕ-class of some element b, preserving a single point in the class is sufficient. Now we can see that A must be a canonical base for $\phi(x,a)$. □

Definition 2.2.18. Let K be a field and let $I \subset K[\bar{x}]$ an ideal. Then $k \subseteq K$ is a field of definition for I if I is generated by elements in $k[\bar{x}]$.

Theorem 2.2.19. *Every ideal $I \subset K[\bar{x}]$ has a unique smallest field of definition, k. Any automorphism of K which fixes I fixes k pointwise.*

Proof. Let M be a basis of the monomials of $K[\bar{x}]/I$ thought of as a K vector space. Then any monomials in the ring $K[\bar{x}]$ can be written as a sum of elements of the form $am + g$ with $a \in K, m \in M, g \in I$. Write each monomial of this form:

$$u = \sum a_{u,i} m_i + g_i,$$

and let k be the subfield of K which is generated by the $a_{u,i}$. Now, take some $f \in K[\bar{x}]$, and write it as a sum of monomials:

$$\sum_u b_u u = \sum_u b_u (u - \sum_i a_{u,i} m_i) + \sum_u b_u (\sum_i a_{u,i} m_i) = \sum_u b_u (u - \sum_i a_{u,i} m_i) + \sum_i c_i m_i.$$

If we take $f \in I$, then the elements m_i form a basis of the ideal, and thus the c_i are zero. So, the elements $u - \sum a_u, i m_i$ are in I, and are a basis of $K[\bar{x}]/I$, thus generate the ideal I. Thus, k is a field of definition of I.

Take some other field of definition, call it F. Then let $f_1, \ldots f_r$ be a generating set for I with coefficients in the field F. Then for each monomial u, there must be $g_{u,1} \ldots g_{u,r} \in K[\bar{x}]$ such that $u - \sum a_{u,i} m_i = \sum g_{u,i} f_i$. But this is, a system of $K[\bar{x}]$-linear equations over $F[x]$. The system has a solution in K, and thus must have a solution already in F. But the m_i form a basis for $K[\bar{x}]/I$, and so if we have $u - \sum c_{u,i} m_i \in I$, it must be the case the the $c_{u,i}$ we chose are actually the earlier fixed $a_{u,i}$. But then we can see that $k \subseteq F$.

To see that any automorphism σ which fixes I as a set must fix k pointwise, apply σ to

$$u - \sum a_{u,i} m_i \in I$$

and notice that we also have

$$u - \sum \sigma(a_{u,i}) m_i \in \sigma(I) = I.$$

Now, by the same argument as before, we have that $\sigma(a_{u,i}) = a_{u,i}$, because the elements m_i form a basis of $K[\bar{x}]/I$.

□

Lemma 2.2.20. *Take $I_1, \ldots, I_n$ to be a collection of conjugate prime ideals in $K[\bar{x}]$. There is a subfield $k \subset K$ such that if $\sigma \in Aut(K)$, then $\sigma \in Aut(K/k)$ if and only if σ permutes $I_1, \ldots, I_n$.*

Proof. Let $I = \cap I_j$. Then I is a radical ideal, written as an intersection of prime ideals, and we know, by Theorem 2.1.25 that there is a unique irredundant such decomposition (apply 2.1.25 in the case that the derivation is trivial). So, take k to be, by Theorem 2.2.19, a field of definition of I. I is preserved by σ if and only if $I_1, \ldots, I_n$ are permuted by σ if and only if σ fixes k pointwise. □

Lemma 2.2.21. *Let $I_1, \ldots, I_n$ be a collection of conjugate prime differential ideals in $K\{x_1, \ldots, x_m\}$. Then there is some differential subfield $k \subseteq K$ such that any differential field automorphism σ of K fixes k pointwise if and only if σ permutes the collection $I_1, \ldots, I_n$.*

Proof. Let $J = \cap I_j$ Then J is a radical differential ideal. Let $f_1, \ldots, f_s$ be a basis for J, that is $J = \{\!\{f_1, \ldots, f_s\}\!\}$. There is some N such that f_i has order bounded by N for all $i = 1, \ldots, s$. Let $J_0 = K[\bar{x}, \ldots, \bar{x}^{(N)}]$. Take k to be the differential field generated by the field of definition of J_0. Any differential field automorphism of K fixes k pointwise if and only if it fixes the ideal J_0 if and only if it fixes the ideal J if and only if it permutes $I_1, \ldots, I_n$, again here using the uniqueness of primary decomposition, Theorem 2.1.25. □

Theorem 2.2.22. *The theories ACF and DCF have elimination of imaginaries.*

Proof. Combine Lemmas 2.2.17 with Lemmas 2.2.20 and 2.2.21 respectively. □

Poizat [13] page xxxi compares the elimination of imaginaries to the "...gliding of vultures above the high Himalayan peaks..." It always seemed more natural to me to compare imaginaries to the *shadows* of vultures, but Poizat's comment was mostly about travel, not analogy, I think.

Lecture 8

In this lecture, we will explain the basics of the Kolchin topology, which we are already familiar with indirectly - In this course, and in most works in the subject, it suffices to consider affine differential varieties over a differential field K - that is - the zero sets of collections of differential polynomials in finitely many variables with coefficients over K. There are, of course, more general ways to approach the subject - projective or even abstract differential varieties (similar to abstract varieties in the vein of Weil, etc., see [9]) or differential schemes [10]. These approaches might prove essential for certain problems, but at the present time, it seems most applications one might have pursue can be tackled using our more straightforward approach. That said, it is often the case that more sophisticated methods from algebraic geometry are can be usefully employed to solve problems of a differential algebraic nature (for instance, [3]). Some of the main difficulties with developing sophisticated tools in the arena of differential schemes is that compared to their algebraic counterparts, sheaves of regular differential functions on differential varieties are rather poorly behaved. For instance, given a differential ring R, there is a natural map from R to the ring of global sections of the corresponding affine differential scheme, however this map is not necessarily injective or surjective (it is an isomorphism in the algebraic case). This sort of difficulty (along with others) is a natural barrier to one building up things like sheaf theory in the differential setting in the same smooth manner in which one does in the algebraic setting - many such difficulties can be overcome by adding assumptions to the class of differential rings one considers (for instance the mentioned issue disappears when one assumes that R is a *Keigher ring*, see [16]). Nevertheless, as we have said, thus far, approaches utilizing more sophisticated tools from the algebraic setting have more often found ways to convert a differential problem to an algebraic one rather than developing the analogous tools from scratch in the differential setting [3] [1].

In the written version of this lecture, there are no proofs, but in class, we did many of the exercises which appear here (at this point, given that you know the previous lectures, the proofs just fall out, essentially).

Let $S \subset K\{x_1, \dots, x_n\}$ be a set of δ-polynomials. We denote, by $\mathbb{V}(S)$, the zero set of S in K^n:

$$\mathbb{V}(S) = \{\bar{a} \in K^n \mid \forall f \in S,\, f(a) = 0\}.$$

Proposition 2.2.23. *Let $S, T, S_i \subseteq K\{\bar{x}\}$. The map $\mathbb{V}$ from collections of differential polynomials over K in n variables to subsets of K^n has the following properties:*

1. $\mathbb{V}(0) = K^n$ *and* $\mathbb{V}(1) = \emptyset$.
2. *When* $S \subseteq T$, *we have* $\mathbb{V}(T) \subseteq \mathbb{V}(S)$.
3. $\mathbb{V}(S) = \mathbb{V}([S]) = \mathbb{V}(\{\!\{S\}\!\})$.
4. $\mathbb{V}(\cup S_i) = \mathbb{V}(\sum [S_i]) = \cap \mathbb{V}(S_i)$.
5. $\mathbb{V}([S] \cap [T]) = \mathbb{V}([S][T]) = \mathbb{V}(S) \cup \mathbb{V}(T)$.

Proof. We leave the proof of the previous proposition as an exercise. □

Sets of the form $\mathbb{V}(S)$ are thus for the closed sets of a topology over K^n, which we call the Kolchin topology.

Let $X \subseteq K^n$. We denote, by $\mathbb{I}(X)$, the collection of differential polynomials which vanish at each element of X:

$$\mathbb{I}(X) = \{f \in K\{x_1, \dots, x_n\} \mid \forall a \in X,\, f(a) = 0\}.$$

Lemma 2.2.24. $\mathbb{I}(X)$ *is a differential radical ideal of* $K\{x_1, \dots, x_n\}$. *When X is closed in the Kolchin topology, we call* $\mathbb{I}(X)$ *the defining ideal of X.*

Proof. Exercise. □

Exercise 2.2.25. Show that $\mathbb{I}(X \cup Y) = \mathbb{I}(X) \cap \mathbb{I}(Y)$. Show that whenever $X \subseteq Y$, we have $\mathbb{I}(Y) \subseteq \mathbb{I}(X)$.

We mentioned the following corollary of quantifier elimination for differentially closed fields in the previous lecture, but one can nicely formulate the differential Nullestellensatz in terms of the maps $\mathbb{V}$ and $\mathbb{I}$.

Theorem 2.2.26. *Let K be differentially closed. The maps $\mathbb{V}, \mathbb{I}$ give an inclusion reversing bijection between affine Kolchin closed subsets and differential radical ideals.*

Proof. □

Exercise 2.2.27. When $X \subset K^n$, we denote, by $\bar{X}$, the *Kolchin closure* of X, that is, the smallest Kolchin closed subset Y which contains X. Show that $\bar{X} = \mathbb{V}(\mathbb{I}(X))$. The Kolchin topology on K^n thus has closed subsets given by the zero sets of radical ideals.

An easy consequence of the basis theorem and the correspondence between closed subsets and differential radical ideals is the Noetherianity of the Kolchin topology:

Corollary 2.2.28. *Any descending chain of Kolchin closed sets in K^n stabilizes.*

Definition 2.2.29. We say that $X \subset K^n$ is irreducible if X can not be written as the union of two proper closed subsets (in the subspace topology induced by the Kolchin topology on X).

Exercise 2.2.30. Show that a nonempty open subset of an irreducible Kolchin-closed set is irreducible and dense.

Exercise 2.2.31. Show that X is irreducible if and only if $\bar{X}$ is irreducible.

Exercise 2.2.32. Show that $X \subset K^n$ is irreducible if and only if $\mathbb{I}(X)$ is irreducible.

Exercise 2.2.33. Let X be a nonempty closed subset of K^n. Show that X can be written uniquely as an irredundant finite union of irreducible closed subsets. The closed subsets in this union are called the *components* of X.

Exercise 2.2.34. Show that the components of X are maximal among irreducible closed subvarieties of X. Show that the corresponding ideals are minimal among prime differential ideals containing $\mathbb{I}(X)$.

When X is a Kolchin-closed subset of K^n, we call $K\{x_1, \ldots, x_n\}/\mathbb{I}(X)$ the *coordinate ring of X* (sometimes differential coordinate ring). We denote, by $K\langle X\rangle$, the ring of total fractions of X, that is, the localization of the coordinate ring with respect to all nonzero divisors.

Exercise 2.2.35. Show that $K\{X\}$ is an integral domain if and only if X is irreducible. In this case, we call $K\langle X\rangle$ the *differential function field* of X.

Exercise 2.2.36. Show that when $X = \cup X_i$ with X_i the irreducible components of X, then $K\langle X \rangle = \prod K\langle X_i \rangle$.

Definition 2.2.37. A morphism of δ-varieties is a map $\phi : X \to Y$, $X \subset \mathbb{A}^n, Y \subset \mathbb{A}^m$ such that $\phi(\bar{x}) = (\phi_1(\bar{x}), \ldots, \phi_m(\bar{x}))$ such that $\phi_i \in K\{X\}$ and $\phi(\bar{x}) \in Y$ for all $\bar{x} \in X$.

It would be possible to define the morphisms slightly differently to be everywhere defined differential rational maps - the two approaches are not equivalent as an everywhere defined differential rational function is not necessarily a differential polynomial (in fact this difficulty is one of the issues with differential scheme theory). Taking differential rational maps as morphisms is equivalent (by QE over a differentially closed field) to taking definable maps between definable sets as the category (which is often how we will work in these notes).

Exercise 2.2.38. Let K be a δ-closed field. Let F be the functor from affine δ-varieties over K to reduced δ-algebras over K obtained by sending a variety to its coordinate ring. Explain how F induces, for every map of δ-varieties, $X \to Y$, a dual map (going in the opposite direction) of the K-algebras given by the coordinate rings. Show that F is an antiequivalence of categories (equivalence of one category and the opposite of another).

Definition 2.2.39. A morphism $\phi : X \to Y$ is a closed embedding if it induces an isomophism of δ-varieties over K between X and $\phi(X)$.

Exercise 2.2.40. Show that $\phi : X \to Y$ is a closed embedding if and only if the dual morphism of coordinate rings is surjective.

Exercise 2.2.41. Let $\phi : X \to Y$ be a morphism of δ-varieties. Show that $\mathbb{V}(\text{Ker}(\phi^*))$ is equal to the Kolchin-closure of $\phi(X)$.

Exercise 2.2.42. Show that if X is irreducible, then $\phi(X)$ is irreducible.

Exercise 2.2.43. Show that morphisms are continuous maps in the Kolchin topology.

Chapter 3

Groups and Galois theory

Lecture 9 - Basics of differential algebraic groups

Definition 3.0.1. A δ-group over K is given by a δ-variety over K, $G \subset K^n$ equipped with a group operation and inverse which are given by δ-morphisms over K. δ-subgroups over K are subgroups which are cut out by δ-subvarieties.

The choice of whether the morphisms are given as in the previous section, essentially, by δ-polynomials or by δ-rational functions is again somewhat arbitrary here. One can make a nowhere vanishing δ-rational function into a δ-polynomial function on a different closed set by taking a different embedding. Various sources have taken the approach of having the group operations given by δ-rational functions, for instance [2].

Next, we give some examples of differential algebraic groups.

Example 3.0.2. Let $L(x)$ be a linear homogeneous differential polynomial over K, and let G be the subset of K given by $\{x \mid L(x) = 0\}$. G is a subgroup of the additive group of K, $\mathbb{G}_a(K)$.

Example 3.0.3. $SL_n(K)$, and $GL_n(K)$. Note that in order to make $GL_n(K)$ fit the definition above, one should take the natural embedding of $GL_n(K)$ into $SL_{n+1}(K)$ given by sending a matrix A to the matrix $\begin{pmatrix} A & 0 \\ 0 & \det(A)^{-1} \end{pmatrix}$.

Exercise 3.0.4. Show that the δ-variety consisting of elements of $GL_3(K)$ of the following form:

$$\begin{pmatrix} a & 0 & c \\ 0 & 1 & b \\ 0 & 0 & a \end{pmatrix}$$

where $\delta(a) = c$ is a δ-subgroup of GL_3.

One of the most important aspects of the theory of δ-groups separating the theory from the special case of algebraic groups is the fact that there is a natural homomorphism from the multiplicative group to the additive group given by the logarithmic derivative:

$$\text{dlog} : \mathbb{G}_m \to \mathbb{G}_a$$

given by $\text{dlog}(a) = \frac{\delta(a)}{a}$. In fact, as we will eventually see, there is an analogous map for any algebraic group.

Definition 3.0.5. Let G be a δ-group over K. Then we denote by λ_g the translation morphism, $x \to gx$. Similarly, we define $\rho_g(x) = xg^{-1}$ and $\iota_g(x) = gxg^{-1}$.

Sometimes techniques in the theory of δ-groups (or more generally stable groups) blur the lines between group theory and differential algebra - one of the reasons is that with respect to groups, the Kolchin topology is particularly well-behaved. Some of this good behavior can be seen as a consequence of the following lemma:

Lemma 3.0.6. *Let H be a (abstract) subgroup of a δ-group G over K. Then $\bar{H}$, the Kolchin-closure of H over K is a δ-subgroup of G.*

Proof. The group operations are given by morphisms, so we need only to verify that $\bar{H}$ is itself a subgroup of G. Take any $h \in H$. Then $H \subseteq h^{-1}\bar{H}$, and $h^{-1}\bar{H}$ is also a closed subset of G. Thus it follows that $\bar{H} \subseteq h^{-1}\bar{H}$ (note that this follows because $\bar{H}$ is contained in any Kolchin closed subset containing H). This implies that $h\bar{H} \subseteq \bar{H}$, for all $h \in H$. Thus $H \cdot \bar{H} \subseteq \bar{H}$. So, for all $\bar{h} \in \bar{H}$, we have that $H\bar{h} \subseteq \bar{H}$. Similar to our reasoning before, this implies that $\bar{H}\bar{h} \subseteq \bar{H}$ for all $\bar{h} \in \bar{H}$ (note that $\bar{H}\bar{h}$ is the closure of $H\bar{h}$, the latter of which is contained in $\bar{H}$ already). This shows that $\bar{H}$ is closed under multiplication.

The map $x \mapsto x^{-1}$ is a homeomorphism, and it preserves a dense subset, H of $\bar{H}$, so it preserves $\bar{H}$. □

The following exercises can be done with a rather elementary approach (though one could also use the theory of ω-stable groups after certain ranks are well understood in this context). For the elementary type of proof the maps from Definition 3.0.5 are useful. Throughout, let G be a δ-group.

Exercise 3.0.7. Let $U, V \subseteq G$ be open and dense. Show that $U \cdot V = G$.

Exercise 3.0.8. Show that there is a unique component G^0 of G which contains the identity of the group, and that it is a closed normal subgroup of G. We call G^0 the *connected component* of G.

Exercise 3.0.9. The irreducible components of G are given by the cosets of G^0.

Exercise 3.0.10. Let $\phi : G \to H$ a morphism of δ-groups over K. Prove that $\text{Ker}(\phi)$ is a normal δ-subgroup of G. Prove that the image of G in H is a δ-subgroup of H. Show that $\phi(G^0) = (\phi(G))^0$ (that is, the image of the connected component of G is the connected component of the image of G).

Lecture 10 - Basic Picard-Vessiot Theory

Let $L(x)$ be a linear homogeneous differential polynomial over K of order n.

Lemma 3.0.11. *Let $x_0, \ldots, x_n$ be solutions to $L(x) = 0$. Then $x_0, \ldots, x_n$ are linearly dependent over K^δ.*

Proof. The Wronskian vanishes. □

Lemma 3.0.12. *Let K be differentially closed. Then there are $x_1, \ldots, x_n$ independent solutions to $L(x) = 0$.*

Proof. The vanishing of the Wronskian of $x_1, \ldots, x_j$ is an order $j - 1$ differential equation, and by the axioms of DCF and induction on j there are solutions $x_1, \ldots, x_n$ to $L(x) = 0$ which are linearly independent over K^δ. □

So, we have the following result:

Theorem 3.0.13. *Let K be differentially closed. Then there are $x_1, \ldots, x_n \in K$ which are linearly independent over K^δ, and the solution set of $L(x) = 0$ is the span of $x_1, \ldots, x_n$ over K^δ.*

The elements $x_1, \ldots, x_n$ satisfying the Theorem is called a fundamental system of solution to $L(x) = 0$.

Definition 3.0.14. Let K/K_0 be an extension of differential fields. We call K a Picard-Vessiot extension of K_0 if there is some linear homogeneous differential polynomial L over K and $x_1, \ldots, x_n \in K$ such that $x_1, \ldots x_n$ is a fundamental set of solutions to $L(x) = 0$, $K = \langle x_1, \ldots, x_n \rangle$ and $K_0^\delta = K^\delta$.

Theorem 3.0.15. *Let K_0 be a δ-field with K_0^δ algebraically closed and suppose that $L(x) = 0$ is a linear homogeneous δ-polynomial over K_0. Then there is a PV-extension K of K_0 for L and K is unique up to isomorphism.*

Proof. □

Lecture 11 - Continuing Basic PV theory and Differential algebraic groups

I want to begin by covering a bit of material on differential algebraic groups which is not required to move forward in the main thread of the lectures, but is interesting enough to mention, in my view. The results here can be proven rather easily at this point via model theoretic methods or more basic methods, but I think that introducing *Hopf algebras* is worthwhile. The results in this section on differential algebraic groups have already been proved by more elementary methods, but there is at least one big theorem in differential algebra whose proof does seem to rely in a more serious manner on Hopf algebra techniques: Cassidy's classification of simple differential algebraic groups (in the partial differential case) [2]. We won't prove that result in this course, but it would be nice to have an example of some results which are a bit easier to prove but seem to use Hopf algebras in a deeper manner.

A K-Hopf algebra is a K-algebra A with equipped with three K-algebra maps:

- Comultiplication, $\Delta : A \to A \otimes_K A$
- Coinverse $S : A \to A$
- Counit $\epsilon : A \to K$

such that the following diagrams commute:

$$\begin{array}{ccc} A \otimes_K A \otimes_K A & \xleftarrow[id \otimes \Delta]{} & A \otimes_K A \\ \uparrow{\scriptstyle \Delta \otimes id} & & \uparrow{\scriptstyle \Delta} \\ A \otimes_K A & \xleftarrow[\Delta]{} & A \end{array}$$

$$\begin{array}{ccc} A & \xleftarrow[S \otimes id]{} & A \otimes_K A \\ \uparrow & & \uparrow{\scriptstyle \Delta} \\ K & \xleftarrow[\epsilon]{} & A \end{array}$$

$$\begin{array}{ccc} K \otimes_K A & \xleftarrow[\epsilon \otimes id]{} & A \otimes_K A \\ & {\scriptstyle \cong} \nwarrow & \uparrow{\scriptstyle \Delta} \\ & & A \end{array}$$

δ-K-Hopf algebras are δ-K-algebras which are also K-Hopf algebras. Morphisms of δ-K-Hopf algebras are morphisms of δ-K-algebras which commute with the Hopf algebra maps. The natural example of a δ-K-Hopf algebra comes from the coordinate ring of a δ-group over K. The maps come the dual maps of the group operation, inverse map, and the unit of the group. Take the simplest case of $G = \mathbb{G}_a$. Then $A = K\{G\} = K\{y\}$ and $\Delta(y) = y \otimes 1 + 1 \otimes y$, and $S(y) = -y$ and $\epsilon(y) = 0$.

We won't prove the next theorem (see Chapter 1 of [17], for instance - one needs to make sure that in the classical correspondence, everything works fine with demanding that the maps behave well with respect to the derivation, but all of this works out without trouble).

Theorem 3.0.16. *Let K be a δ-closed field. There is an antiequivalence of categories between affine δ-groups over K and commutative reduced K-Hopf algebras given by the functor $G \mapsto K\{G\}$. Each possible group scheme structure on an affine scheme corresponds to a K-Hopf algebra structure on its coordinate algebra.*

We've already seen that the Kolchin closure of an abstract group is a differential algebraic group (in fact our proof works in an arbitrary topological group), but we will reprove the theorem now in order to get used to Hopf algebras.

Proposition 3.0.17.

Proposition 3.0.18.

Now, we'll move on to differential Galois theory.

Lecture 12 - Some very general differential Galois theory

Let K/K_0 be a differential field extension, and $G = Aut(K/K_0)$ the differential Galois group. There are natural maps between intermediate field differential field extensions and subgroups G: given $H \leq G$, let H' be the differential field of elements fixed pointwise by each element of H. Given L, an intermediate differential field extension, let L' be the subgroup of G of elements which fix each element of L. We call a group (or a differential subfield of K) *closed* if $H = H''$ (or $L = L''$).

Define K to be *normal* over K_0 if for any $a \in K \setminus K_0$, we have some $\sigma \in Aut(K/K_0)$ such that $\sigma(a) \neq a$.

Exercise 3.0.19. Show that for any subgroup H (or differential subfield L) $H' = H'''$ (or $K' = K'''$).

It is easy to see that there is a bijective correspondence between the closed subgroups and closed differential subfields, but this leaves the important question of which subfields or subgroups are closed completely open. We will eventually answer the question in some special cases.

Theorem 3.0.20. *Let G be the differential Galois group of K/K_0.*

1. *If $H \lhd G$ then any automorphism of K/K_0 sends H' to itself.*

2. *If L is a differential subfield of K such that any automorphism of K/K_0 sends L to itself, then we have that L' is a normal subgroup.*

Proof. Let $\sigma \in Aut(K/K_0)$, and let $a \in H'$. We need to show that $\sigma(a) \in H'$. This is equivalent to $\tau(\sigma(a)) = \sigma(a)$ for all $\tau \in H$. This is equivalent to $\sigma^{-1}\tau\sigma(a) = a$. But, of course, this is true since H is normal.

To see that L' is normal, follow through the proof of the first part in reverse. Now, we have a homomorphism from G to the automorphism group of L over K_0 via restricting the automorphism to L. The kernel of this map is L' and the image is the collection of automorphisms of L/K_0 which extend to all of K. □

Exercise 3.0.21. Using the previous result, show that the closure of a normal subgroup of G is a normal subgroup.

Exercise 3.0.22. Explain why the δ-field extension associated with a normal subgroup of G is necessarily normal. Give an example showing that the converse is not true. The proof of Lemma 3.0.25 gives you an idea of what properties a counterexample must have.

Exercise 3.0.23. Check that nothing so far in this section depends on the structures in question being differential fields. That is, one could set all of this up in the setting of arbitrary first order structures, and the previous theorem would still be true.

Lemma 3.0.24. *Let L be a closed subfield. Let $H = L'$. Then the normalizer of H, $N_G(H)$, is the collection of all $\sigma \in G$ such that $\sigma(L) = L$.*

Proof. We gave the proof in lecture, but leave it as an exercise to the reader. □

Lemma 3.0.25. *Let L be a closed subfield which is normal over K_0. Let $H = L'$. Suppose that $N_G(H)$ is closed and every differential automorphism of L/K_0 can be extended to K. Then H must be normal (so $N_G(H) = G$) and $G/H = Aut(L/K_0)$.*

Proof. To prove H is normal, we can show that $N_G(H)' = K_0$. By the previous Lemma, we know that $N_G(H)$ is the collection of those elements of G which map L to L. In this collection, we have every element of $Aut(L/K_0)$ because of the extension hypothesis. Since L is normal, we have that no elements of L other than those of K_0 are fixed by every element of $N_G(H)$. But this means that $N_G(H)' = K_0$. Now finish with Theorem 3.0.20. □

Exercise 3.0.26. Let $K_0 \subseteq L \subseteq K$ with L/K_0 a PV extension and suppose that $K^\delta = K_0^\delta$. Show that any automorphism of L/K_0 extends to an automorphism of K/K_0.

Exercise 3.0.27. Let L/K_0 be a PV extension and let M/L be a differential field extension with $L^\delta = M^\delta$. Then for any $\sigma \in Aut(M/K_0)$, $\sigma(L) = L$.

Exercise 3.0.28. Show that the two previous lemmas (with the analogous definitions setup) are valid in the arbitrary first order context.

Some remarks on the above subsection: we have shown in a very general setting, that there is a Galois correspondence of sorts between closed subgroups and closed intermediate extensions. Such a theorem is satisfying but not particularly practical if one takes as the central aim of differential Galois theory understanding the solution sets of differential equations (for instance, we have learned nothing directly from the correspondence about what kind of intermediate differential field extensions there are - such intermediate field extensions correspond to certain kinds of differential subvarieties). This restriction is probably intrinsic; for instance, understanding when a differential equation is strongly minimal seems like a very difficult question. So, for general differential field extensions, this is probably a completely hopeless problem, but it is not in the case of PV extensions.

Finishing basic PV theory

Lemma 3.0.29. *Let K_0^δ be algebraically closed and let L/K be a PV extension. Suppose that we are given $z \in L$ and $A_i, B_i \subset L$ for $i \in I$ with the index set I arbitrary but A_i and B_i finite. Suppose that there is a differential isomorphism sending A_i to B_i for each i such that $\sigma(z) \neq z$.. Then there is a differential automorphism $\tau \in Aut(L/K_0)$ such that $\tau(z) \neq z$ and $\tau(A_i) = B_i$.*

[1]

Proof. Let σ be our differential isomorphism. Then $\sigma(u_i) = \sum_j k_{ij} u_i$ for some constants k_{ij} from $\mathcal{U}$ a saturated differentially closed field extending L. Now, any element of L is contained in $\text{Frac}(K[u_1, \ldots, u_n])$, so take two arbitrary elements $x, y \in L$. We can write each as a fraction in the $u = (u_1, ..., u_n)$ with coefficients in K_0. So, let $x = \frac{P(u)}{Q(u)}$, $y = \frac{R(u)}{S(u)}$. So, $y = \sigma(x)$ if and only if $S(u)P(\sigma(u)) = R(u)Q(\sigma(u))$. So, substituting the expression $\sigma(u_i) = \sum_j k_{ij} u_i$, we can see that the result is just a system of equations in $\bar{k}$. But now there is a solution to the system of equations corresponding to the condition that $\sigma(A_i) = B_i$, and there is some inequation expressing that $\sigma(z) \neq z$. So, now there is a solution in the base constant field K_0^δ. □

We obtain the following immediate corollary of the previous result.

Theorem 3.0.30. *Let K_0^δ be algebraically closed. Any PV extension of K_0 is normal.*

Lemma 3.0.31. *Let K/K_0 be finitely generated (as a δ-field extension) and let L be an intermediate differential field extension of K_0. Then L is finitely generated.*

The previous lemma can be proved in a rather elementary manner, but the following exercise gives hints of how to prove the result in manner of [18].

Exercise 3.0.32. Take $a_1, \ldots, a_n$ to be generators for K/K_0. If $\bar{a}$ is δ-transcendental over L, then we can see that $L = K_0$. So, assume the contrary, and take a characteristic set $A_1, \ldots, A_s$ for $\bar{a}$. Let Q be the collection of coefficients of the characteristic set of $\bar{a}$ over L. Certainly $K_{\langle} Q \rangle$ is a subfield of L, but we claim that $K_{\langle} Q \rangle = L$. Prove this by considering the following fields carefully:

- $K_0\langle Q \rangle$.
- $K_0\langle Q \rangle(R)$ where R is a set of parametric derivatives of the $\bar{a}$ (that is the derivatives of the elements of $\bar{a}$ which are not proper derivatives of leaders of any element of the characteristic set $A_1, \ldots, A_s$).
- $K_0(R)$.

[1][7] seems to state this lemma for arbitrary A_i and B_i, but I only see how the proof works in the case that A_i and B_i are finite. In fact, for extension of automorphisms, it only seems necessary to really use singletons for the A_i and B_i

Show that the latter two fields are identical by computing their index in K.

Show that the index of $K_0\langle Q\rangle$ in L is equal to the index of $K_0\langle Q\rangle(R)$ over $K_0(R)$ (which is one by what you just showed).

Theorem 3.0.33. *Let K_0 be a differential field of characteristic zero with K_0^D algebraically closed. Let K be a PV extension of K_0. Then any isomorphism between intermediate subfields can be extended to an automorphism of K/K_0.*

Proof. Apply Lemma 3.0.29 on the generating sets A_1 and B_1 using Lemma 3.0.31 on the generating sets to force them to be finite.[2] □

Remark 3.0.34. At a certain point, the previous several results fall back on the Noetherianity of the Zariski topology (used on the group of automorphisms of the PV-extension). One should note that an argument in this style seems very possible for generalization as long as one has a Noetherian topology lurking around (so, for instance, the Kolchin topology).

Let $K_0 \subset L \subset K$, and let K/K_0 be PV. Then K over L is also PV. By 3.0.30, we have that K/L is normal, and so we have shown that all intermediate subfields are closed.

Let $H \lhd G = Aut(K/K_0)$ a PV-extension. Let $L = H'$. Suppose that H is closed. Then any automorphism of L/K_0 extends to K. So, we have that $G/H = Aut(L/K_0)$. Assuming that L is closed and normal yields the same conclusion. We've shown that G is an algebraic group, and since K/L is a PV extension, it follows that $H = L'$ is an algebraic subgroup of G, and it remains to show that each such algebraic subgroup H is closed (in the sense that $H'' = H$). This fact would follow if we could show that H'' is always Zariski-dense in H. To that end, let f be some polynomial, in $\bar{k}$ which vanishes on H, but not on some element of H''. Take $K = K_0\langle u_1, \ldots, u_n\rangle$. The matrix

$$U = \begin{pmatrix} u_1 & \ldots & u_n \\ u_1' & \ddots & u_n' \\ u_1^{(n-1)} & \ldots & u_n^{(n-1)} \end{pmatrix}$$

is invertible over K, let $A = (a_{ij})$ be the inverse. Now, define

$$F(y_1, \ldots, y_n) = f(\sum_j a_{1j} y_1^{(j-1)}, \sum_j a_{1j} y_2^{(j-1)}, \ldots, \sum_j a_{1j} y_n^{(j-1)}, \sum_j a_{2j} y_1^{(j-1)}, \ldots, \sum_j a_{2j} y_1^{(j-1)}, \ldots).$$

(The previous expression is f evaluated on the entries of $A \cdot Y$, where Y is the matrix associated with the wronskian of $y_1, \ldots, y_n$.)

Now, consider $F(\sigma(u_1), \ldots, \sigma(u_n))$ with $\sigma \in H$.

[2] Kaplansky [7] does not seem to mention using a result like Lemma 3.0.31 in his proof of this Lemma (page 37), and he doesn't ever prove this result. Instead, what I think he intends is to take A_i to be singletons enumerating the domain and B_i to be the images of the singletons, and then apply the Lemma to these elements.

We have that F vanishes on $\sigma(\bar{u})$ for all $\sigma \in H$ but not for all $\sigma \in H''$. Consider the collection of all differential polynomials over K with this property, and take one, call it E, which has the minimal number of monomials appearing. We can assume that E is monic. Now, take the polynomial E_τ for $\tau \in H$ obtained by replacing the coefficients of E with their respective images under τ. There are two options:

1. $E - E_\tau$ is identically zero. Then every coefficient of E must have been in H', and the differential field H' is invariant under the action of H'', by definition. So, it must be that E_σ vanishes for all $\sigma \in H''$, contradicting our assumption.

2. $E - E_\tau$ is not identically zero. Then since $E - E_\tau$ (recall E is monic) has fewer monomials than E, we must have that $E - E_\tau$ vanishes at $\sigma(\bar{u})$ for all $\sigma \in H''$. But then we can find some $d \in K$ such that $E - d(E - E_\tau)$ has fewer monomials than E. But then note that $E - d(E - E_\tau)$ has the property that it vanishes identically on all of H but not all of H'', a contradiction to the minimality of E.

Thus, H is Zariski dense in H'', which implies that $H = H''$ whenever H is Zariski-closed. That is, H is closed in the Galois sense if and only if it is closed in the sense of the Zariski topology.

We have now proved:

Theorem 3.0.35. *Let K_0 be a differential field with K^δ algebraically closed, and K/K_0 a PV-extension. Then there is a bijective correspondence between intermediate differential fields and algebraic subgroups of $Aut(K/K_0)$. A closed subgroup H is normal if and only if the corresponding field extension L/K_0 is normal, and in that case, $G/H = Aut(L/K_0)$.*

Bibliography

[1] A. Buium. Intersections in jet spaces and a conjecture of s. lang. *Ann. of Math.*, (2) 136:557–567, 1992.

[2] Phyllis Cassidy. The classification of the semisimple differential algebraic groups and the linear semisimple differential algebraic lie algebras. *Journal of Algebra*, 121:169–238, 1989.

[3] James Freitag and Rahim Moosa. Finiteness theorems on hypersurfaces in partial differential-algebraic geometry. *Advances in Mathematics*, 314:726–755, 2017.

[4] James Freitag, Omar León Sánchez, and William Simmons. On linear dependence over complete differential algebraic varieties. *Communications in Algebra*, 44(6):2645–2669, 2016.

[5] O. Golubitsky, M. Kondratieva, and A. Ovchinnikov. On the generalized Ritt problem as a computational problem. *Journal of Mathematical Sciences*, 163(5):515–522, 2009.

[6] Matthew Harrison-Trainor, Jack Klys, and Rahim Moosa. Nonstandard methods for bounds in differential polynomial rings. *Journal of Algebra*, 360(0):71 – 86, 2012.

[7] Irving Kaplansky. *An introduction to differential algebra*. Actualités scientifiques et industrielles. Hermann, 1957.

[8] Ellis R. Kolchin. *Differential Algebra and Algebraic Groups*. Academic Press, New York, 1976.

[9] Ellis R. Kolchin. *Differential Algebraic Groups*. Academic Press, New York, 1984.

[10] Jerald J Kovacic. Differential schemes. In *In Differential Algebra and Related Topics*. Citeseer, 2002.

[11] David Marker. *Model Theory: an Introduction*. Springer, Graduate Texts in Mathematics, 217, Second Edition, 2002.

[12] David Marker, Margit Messmer, and Anand Pillay. *Model theory of fields*. A. K. Peters/CRC Press, 2005.

[13] Bruno Poizat. *A course in model theory*. Springer, New York, 2000.

[14] Joseph Fels Ritt. *Differential Algebra*. Dover Publications, New York, 1950.

[15] William Sit. The Ritt-Kolchin theory for differential polynomials. *Differential Algebra and Related Topics, Proceedings of the International Workshop, NJSU*, 2000.

[16] Dmitry Trushin. Global sections of structure sheaves of keigher rings. *Proceedings of the American Mathematical Society*, 140(2):505–511, 2012.

[17] William C Waterhouse. *Introduction to affine group schemes*, volume 66. Springer Science & Business Media, 2012.

[18] Israel Zuckerman. A new measure of a partial differential field extension. *Pacific Journal of Mathematics*, 15(1):357–371, 1965.