1 Notation

Throughout, we will be interested in the theory of differential fields with a finite set of derivations $\Delta = \{\delta_1, \ldots, \delta_m\}$. In particular we will investigate the Kolchin polynomial, which is a numerical polynomial which computes roughly how large of an extension of fields one obtains by adding successively higher order derivations of a fixed collection of elements to a base field k. Many proofs here are elaborations on proofs from [1] and [2].

We start with some definitions and notations for differential algebra. Given our set Δ of derivations, we let Θ be the commutative monoid of *derivative operators* generated by Δ. For $\theta = \delta_1^{a_1} \cdots \delta_m^{a_m} \in \Theta$, the order of θ is the quantity $\mathrm{ord}(\theta) = \sum_{i=1}^m a_i$. For a fixed $s \in \mathbb{N}$, we define $\Theta_{\leq s} = \{\theta \in \Theta \mid \mathrm{ord} \leq s\}$.

Given differential inderminates $\overline{y} = (y_1, \ldots, y_n)$, we define

$$\Theta\overline{y} = \{\theta y_i \mid \theta \in \Theta, 1 \leq i \leq n\}.$$

A similar definition can be given for $\Theta_{\leq s}\overline{y}$.

From now on, we fix a ranking $\preceq$ on $\Theta\overline{y}$, which is a linear order of order type ω so that

- $u \leq v \Rightarrow \theta u \leq \theta v$
- For any $\theta_1, \theta_2 \in \Theta$, and any $i, j \leq n$, $\mathrm{ord}(\theta_1) < \mathrm{ord}(\theta_2) \Rightarrow \theta_1 y_i < \theta_2 y_j$

For a fixed $u \in \Theta\overline{y}$, we define

$$(\Theta\overline{y})_{\preceq u} = \{v \in \Theta\overline{y} \mid v \preceq u\}.$$

We can define $(\Theta\overline{y})_{\prec u}$ analogously.

For a field k, we let $k\{\overline{y}\}$ be the free Δ-ring generated by $\overline{y}$, and is equal to the polynomial ring $k[\Theta\overline{y}]$. Given a polynomial $f \in k\{\overline{y}\}$, we call the highest rank term u showing up in f the *leader* of f, and we denote it u_f. We may express f as a polynomial in u_f with coefficients from $k[(\Theta\overline{y})_{\prec u_f}]$:

$$f = \sum_{i=0}^{d} a_i (u_f)^i.$$

We call a_i the *initial* of f, and denote it I_f. We define the *seperant* of f to be the formal partial derivative of f with respect to u_f, and denote it $S_f = \frac{\partial f}{\partial u_f}$. The degree of f, denoted $\deg(f)$, is the degree of f as a polynomial in u_f. We can assign to each polynomial f the ordered pair $\mathrm{rank}(f) = (u_f, \deg(f))$. When we comparing the ranks of differential polynomials, we order the ranks lexicograhically.

For f, g differential polynomials, we say f is reduced with respect to g if no proper derivative of u_g appears in f, and moreover if u_g appears in f then $\deg_{u_g}(f) < \deg_{u_g}(g)$. A set Λ of differential polynomials is said to be autoreduced if pairwise f is reduced with respect to to g for any distinct $f, g \in \Lambda$. One can check that autoreduced sets are always finite, and distinct elements of an autoreduced set have distinct leaders.

We can extend the ranking of polynomials to autoreduced sets by stipulating $\{f_1 < \cdots < f_r\} \leq \{g_1 < \cdots < g_s\}$ if either $s \leq r$ and $\mathrm{rank}(f_i) = \mathrm{rank}(g_i)$ for all $i \leq s$, or if there is an $i \leq \min\{s, r\}$ such that $\mathrm{rank}(f_j) = \mathrm{rank}(g_j)$ for all $j < i$, and $\mathrm{rank}(f_i) < \mathrm{rank}(g_i)$. For a prime Δ-ideal $\mathfrak{p} \subseteq k\{\overline{y}\}$, we say Λ is a characteristic set for $\mathfrak{p}$ if Λ a minimal autoreduced subset of $\mathfrak{p}$.

2 Numerical Polynomials

Definition 1. A *numerical polynomial* is a polynomial $f \in \mathbb{Q}[x]$ such that $f(t) \in \mathbb{Z}$ for all $t \in \mathbb{Z}$.

For $n \in \mathbb{N}$, we can define the so-called binomial coefficient polynomial $\binom{x}{n}$ as

$$\binom{x}{n} = \frac{x(x-1)\cdots(x-n+1)}{n!}.$$

For $n = 0$ we define $\binom{x}{n} = 1$. Notice that $\binom{x}{n}$ is numerical for all $n \in \mathbb{N}$, as evaluationg $\binom{x}{n}$ at a particular integer $x = a$ just corresponds to the actual binomial coefficient $\binom{a}{n}$, which is of course an integer. Moreover, sums, products and compositions of numerical polynomials are numerical. Hence, it follows that for any $a_0, \ldots, a_n \in \mathbb{Z}$, the polynomial $\sum_{i=0}^{n} a_i \binom{x+i}{i}$ is numerical. In fact, every numerical polynomial has this form.

Theorem 2. If $f(x)$ is a numerical polynomial of degree d, then there exist unique $a_0, \ldots, a_d \in \mathbb{Z}$ such that

$$f(x) = \sum_{i=0}^{d} a_i \binom{x+i}{i}.$$

Proof. As the degree of $\binom{x+i}{i}$ is i, there exist unique $a_0, \ldots, a_d \in \mathbb{Q}$ such that $f(x) = \sum_{i=0}^{d} a_i \binom{x+i}{i}$. Our goal is to show that in fact $a_0, \ldots, a_d$ must be integers. We may prove this by induction on d. The statement is clearly true when $f(x)$ is a constant. Now, suppose $d > 0$. Since f is numerical, so is the polynomial $f(x+1) - f(x)$. Also, the leading terms of $f(x+1)$ and $f(x)$ are the same, $f(x+1) - f(x)$ has smaller degree than f. We can compute

$$\begin{aligned} f(x+1) - f(x) &= \sum_{i=0}^{d} a_i \binom{x+1+i}{i} - a_i \binom{x+i}{i} \\ &= \sum_{i=0}^{d} a_i \binom{x+i}{i-1} \\ &= \sum_{i=0}^{d-1} a_{i+1} \binom{x+1+i}{i} \end{aligned}$$

Writing $f(x+1) - f(x)$ as a polynomial in $y = x + 1$, say $g(y)$, we have

$$g(y) = \sum_{i=0}^{d-1} a_{i+1} \binom{y+i}{i}.$$

By induction, we get that $a_1, \ldots, a_d$ are integers. We also get that a_0 is an integer as $a_0 = f(-1)$, and f is numerical. □

In the previous proof, the concept of the *forward difference operator* came up, which we can define by $\Delta[f](x) = f(x+1) - f(x)$. The inverse operator to this is summation: define the function $F(n) = \sum_{i=0}^{n-1} f(i)$. Then it is clear that $\Delta[F](n) = f(n)$. For a general function f, it is a hard problem to find a closed form for the function F in terms of f. However, in the case that f is a numerical polynomial, we are in luck.

Lemma 3. If f is a numerical polynomial, there exists a numerical polynomial F with $\Delta[F] = f$. Moreover, $\deg(F) \leq \deg(f) + 1$.

Proof. We note that Δ is linear. Hence, it suffices to prove the lemma in the case that $f(x) = \binom{x}{n}$ for some $n \in \mathbb{N}$. We can check that $\Delta[\binom{x}{n+1}] = \binom{x}{n}$, as this is just another form of the identity

$$\binom{n+1}{k+1} = \binom{n}{k} + \binom{n}{k+1}.$$

Clearly we have $\deg \binom{x}{n+1} = n + 1 \leq \deg \binom{x}{n} + 1$ □

We turn our attention back towards differential fields for a moment as motivation for the next section. Given a our set of commuting derivations $\Delta = \{\delta_1, \ldots, \delta_m\}$, we can generate the monoid of derivative operations

$$\Theta = \{\delta_1^{n_1} \cdots \delta_m^{n_m} \mid (n_1, \ldots, n_m) \in \mathbb{N}^m\}.$$

We can consider the divisibility relation on Θ as a monoid where θ_1 divides θ_2 if there exists some $\theta' \in \Theta$ such that $\theta_2 = \theta'\theta_1$. If we identify Θ with $\mathbb{N}^m$ in the natural way, the divisibility relation coincides with the product poset structure of $\mathbb{N}^m$. In this next section, we will analyze $\mathbb{N}^m$ with the product order, keeping in the back of our minds that the relation $(a_1, \ldots, a_m) \leq (b_1, \ldots, b_m)$ will later be taken to mean that we can get from $\delta_1^{a_1} \cdots \delta_m^{a_m}$ to $\delta_1^{b_1} \cdots \delta_m^{b_m}$ by applying more derivations.

For the next lemma, it will be useful to think of $\mathbb{N}^m$ not just as a poset, but as a lattice. As such, for $\overline{a} \in \mathbb{N}^m$, we will let $F(\overline{a})$ denote the principal filter above $\overline{a}$. That is, $F(\overline{a}) = \{\overline{x} \in \mathbb{N}^m \mid \overline{x} \geq \overline{a}\}$. Similarly, we will let $I(\overline{a})$ denote the principal ideal below $\overline{a}$.

Lemma 4. Fix $E \subseteq \mathbb{N}^m$. Let V_E be the set of points in $\mathbb{N}^m$ which are not greater than or equal to any element of E. Then exists a numerical polynomial ω_E such that for t sufficiently large,

$$\omega_E(t) = \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i \leq t \right\} \right|$$

Moreover, $\deg(\omega_E) \leq m$.

Proof. We can derive this from the Hilbert polynomial. Recall that if $S = \bigoplus_{i=0}^{\infty} S_i$ is a graded algebra over a field k which is finitely generated by elements of positive degree, then there exists a numerical polynomial $\mathcal{H}_S(x)$ such that for sufficiently large n,

$$\mathcal{H}_S(n) = \dim_k S_n.$$

For each $\overline{a} \in \mathbb{N}^m$, define the corresponding monomial $A_{\overline{a}} = x_1^{a_1} \cdots x_m^{a_m} \in k[x_1, \ldots, x_m]$. Now, define following homogeneous ideal of $R = k[x_1, \ldots, x_m]$ by

$$I = (A_{\overline{e}} \colon \overline{e} \in E).$$

Observe that for a general $\overline{a} \in \mathbb{N}^m$, we have $A_{\overline{a}} \in I$ if and only if $\overline{a} \notin V_E$. Consider the graded algebra

$$S = k[x_1, \ldots, x_m]/I.$$

We have $S_n = R_n/I_n$. Then $\dim_k(S_n) = \dim_k(R_n) - \dim_k(I_n)$. We have that $A_{\overline{a}}$ for $\overline{a}$ ranging over $\{(v_1, \ldots, v_m) \in \mathbb{N}^m, \sum_i v_i = n\}$ and ranging over $\{(v_1, \ldots, v_m) \notin V_E \colon \sum_i v_i = n\}$ are bases for R_n and I_n respectively. Hence,

$$\begin{aligned} \dim_k(S_n) &= \left| \left\{ (v_1, \ldots, v_m) \in N^m \colon \sum_i v_i = n \right\} \right| - \left| \left\{ (v_1, \ldots, v_m) \notin V_E \colon \sum_i v_i = n \right\} \right| \\ &= \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i = n \right\} \right| \end{aligned}$$

Then let $\mathcal{H}_S(x)$ be the Hilbert polynomial for S. There is some $n_0 \in \mathbb{N}$ so that for $n \geq n_0$, we have

$$\mathcal{H}_S(n) = \dim_k(S_n) = \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i = n \right\} \right|.$$

This is almost what we desire, except we want to take a summation of the first n values of this polynomial. Apply lemma 3 to $\mathcal{H}_S(x+1)$ to get a polnomial F so that

$$\Delta[F](x) = \mathcal{H}_S(x+1).$$

By adding a constant to F if necessary, we may assume that

$$F(n_0) = \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i \leq n_0 \right\} \right|$$

Now, it is an easy induction on $t \in \mathbb{N}$ to show that

$$F(n_0 + t) = F(n_0) + \sum_{i=1}^{t} \mathcal{H}_S(n_0 + i).$$

Hence,

$$\begin{aligned} F(n_0+t) &= \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i \leq n_0 \right\} \right| + \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon n_0 < \sum_i v_i \leq n_0 + t \right\} \right| \\ &= F(n_0+t) = \left| \left\{ (v_1, \ldots, v_m) \in V_E \colon \sum_i v_i \leq n_0 + t \right\} \right| \end{aligned}$$

This F is our desired ω_E

As for calculating $\deg(\omega_E)$, we have that $\deg(\mathcal{H}_S) \leq m-1$, as the graded algebra S has Krull dimension at most m, since it is the quotient of $k[x_1, \ldots, x_m]$. After applying lemma 3 to $\mathcal{H}_S(x+1)$, this raises the degree of F by at most 1, so $\deg(F) \leq m$. □

3 The Kolchin Polynomial

Now we have the tools to show the existence of the Kolchin polynomial. For this section, fix differential fields $k \subseteq K$.

Theorem 5. Let $a_1, \ldots, a_n \in K$. There exists a numerical polynomial $\omega_{\overline{a}/k}$ with the following properties:

(a) For t sufficiently large,
$$\omega_{\overline{a}/k}(t) = \operatorname*{trdeg}_k k((v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}}).$$

(b) $\deg(\omega_{\overline{a}/k}) \leq m$

(c) If $\mathfrak{p} = I(\overline{a}/k) \subseteq k\{y_1, \ldots, y_n\}$, if Λ is a characteristic set of $\mathfrak{p}$, and if E_i is the set of $\overline{e} \in \mathbb{N}^m$ such that $\delta_1^{e_1} \cdots \delta_m^{e_m} y_i$ is the leader of some element of Λ, then
$$\omega_{\overline{a}/k} = \sum_{i=1}^{n} \omega_{E_i}$$

Proof. Fix $f \in \Lambda$. Let $v_1, v_2, \ldots, v_r$ enumerate the elements of $\Theta\overline{y}$ of order strictly less than or equal to the leader of f. We may write $f = \tilde{f}(v_1, \ldots, v_r, u_f)$ for some $\tilde{f} \in k[z_1, \ldots, z_r, w]$. For $1 \leq i \leq r$ let $b_i = v_i(\overline{a})$. Then since $f(\overline{a}) = 0$, we have

$$\tilde{f}(b_1, \ldots, b_r, u_f(\overline{a})) = 0.$$

This witnesses that $u_f(\overline{a})$ is algebraic over the field $k((v(\overline{a}))_{v \preceq u_f})$. Now, for any $\theta \in \Theta$, we may apply the same argument to θf to show that $\theta u_f(\overline{a})$ is algebraic over the field $k((v(\overline{a}))_{v \preceq \theta u_f})$.

Let V denote the set of $v \in \Theta\overline{y}$ which are not derivatives of u_f, for any $f \in \Lambda$. Note: this should feel reminiscent of the V from lemma 4. Further along these lines, for $t \in \mathbb{N}$, let $V(t)$ denote the set of $v \in V$ with $\mathrm{ord}(v) \leq t$. It follows from the above that the field $k((v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}})$ is algebraic over $k((v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}})$. All that suffices to check is that $v(\overline{a})$ is algebraic over $k((v(\overline{a}))_{v \in \Theta_{\leq s}\overline{y}})$ for $v \in \Theta_{\leq t}\overline{y} \setminus V(t)$. But $\Theta_{\leq t}\overline{y} \setminus V(t)$ consists exactly of the derivatives of leaders of $f \in \Lambda$ of order at most t, which is exactly the v which we checked above. With this in mind, we have

$$\operatorname*{trdeg}_k k((v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}}) = \operatorname*{trdeg}_k k((v(\overline{a}))_{v \in V(t)}).$$

Hence, it will suffice to calculate the latter value. In fact, the family $(v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}}$ is algebraically independent over k. Suppose to the contrary that there were an algebraic dependence $P \in \mathfrak{p} \cap k[V]$. Note that since all $v \in V$ are not the derivatives of any u_f for any $f \in \Lambda$. Hence, any element of $k[V]$ is reduced with respect to Λ. Then $P \in \mathfrak{p}$ is impossible, as we can add P to Λ to form a autoreduced subset of $\mathfrak{p}$ which is smaller than Λ.

Since $(v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}}$ is algebraically independent, we can easily calculate

$$\operatorname*{trdeg}_k k((v(\overline{a}))_{v \in \Theta_{\leq t}\overline{y}}) = |V(t)|.$$

We can separate $V(t)$ into the sets as

$$V(t) = \bigcup_{i=1}^{n} V(t) \cap \Theta y_i.$$

Now we can readily apply lemma 4, as for each fixed y_i, the set Θy_i has the structure of $\mathbb{N}^m$. Moreover, $V \cap \Theta y_i$ is exactly the V from lemma 4 obtained from E_i. Then for sufficiently large t, we have

$$|V(t)| = \sum_{i=1}^{n} |V(t) \cap \Theta y_i| = \sum_{i=1}^{n} \omega_{E_i}(t).$$

Hence, we may take

$$\omega_{\overline{a}/k} = \sum_{i=1}^{n} \omega_{E_i}.$$

This proves (a) and (c). Moreover, each polynomial E_i has dimension bounded by m, so (b) is satisfied.

□

For a moment we will leave the realm of differential algebra to recall the definition of the dimension of a prime ideal, as well as a usefull lemma about it.

Definition 6. Let $\mathfrak{p}$ be a prime ideal of $k[\overline{y}]$. We define

$$\dim_k \mathfrak{p} = \operatorname*{trdeg}_k k(\overline{a}),$$

where $\overline{a}$ is a generic zero of $\mathfrak{p}$.

Proposition 7. If $\mathfrak{p}, \mathfrak{q} \subseteq k[\overline{y}]$ are two prime ideals with $\mathfrak{p} \subseteq \mathfrak{q}$ then $\dim_k \mathfrak{p} \geq \dim_k \mathfrak{q}$ with equality if and only if $\mathfrak{p} = \mathfrak{q}$.

Returning back to differential algebra, suppose $\mathfrak{p} \subseteq k\{\overline{y}\}$ is a prime Δ-ideal, and let $\overline{a}$ be a generic solution to $\mathfrak{p}$. We may define $\omega_{\mathfrak{p}}$ to be $\omega_{\overline{a}/k}$. We can reconsile this as follows. For $s \in \mathbb{N}$ let $\mathfrak{p}_s = \mathfrak{p} \cap k[\theta_{\leq s}\overline{y}]$. Then by definition, we have

$$\dim_k \mathfrak{p}_s = \operatorname*{trdeg}_k k((v(\overline{a}))_{v \in \theta_{\leq s}\overline{y}}).$$

Hence, for sufficiently large $s \in \mathbb{N}$, the Kolchin polynomial records

$$\omega_{\mathfrak{p}}(s) = \dim_k \mathfrak{p}_s.$$

Lemma 8. If $\mathfrak{p}, \mathfrak{q} \subseteq k\{\overline{y}\}$ are two Δ-ideals with $\mathfrak{p} \subseteq \mathfrak{q}$ and $\mathfrak{p} \neq \mathfrak{q}$ then $\omega_{\mathfrak{p}} > \omega_{\mathfrak{q}}$.

Proof. As $\mathfrak{p} \neq \mathfrak{q}$, we may find s_0 so that for all $s \geq s_0$ we have $\omega_{\mathfrak{p}}(s) = \dim_k \mathfrak{p}_s$, $\omega_{\mathfrak{q}}(s) = \dim_k \mathfrak{q}_s$, and $\mathfrak{p}_s \neq \mathfrak{q}_s$. Since $\mathfrak{p} \subseteq \mathfrak{q}$, we have $\mathfrak{p}_s \subseteq \mathfrak{q}_s$ for all s. Then by the previous proposition we have $\dim_k \mathfrak{p}_s > \dim_k \mathfrak{q}_s$ for all $s \geq s_0$. Hence, $\omega_{\mathfrak{p}}(s) > \omega_{\mathfrak{q}}(s)$ for all $s \geq s_0$. □

4 Forking

We move on to give an algebraic characterization of forking using the Kolchin polynomial. We will give the necessary definitions to state the characterization, and then build up to the theorem with some lemmas. Throughout this section we will be concerned with types over differential fields. Naturally to each type $p \in S_{\overline{y}}(k)$ we may associate a prime differential ideal $I_p = \{f \in k\{\overline{y}\} \mid f(\overline{y}) = 0 \in p\}$. For each $p \in S_{\overline{y}}(k)$ we will let Λ_p denote the characteristic set of I_p. Also, we let ω_p be the Kolchin polynomial associated to I_p, which is equivalently $\omega_{\overline{a}/k}$ for any $\overline{a}$ realizing p.

Suppose Λ_1 and Λ_2 are two autoreduced sets in $k\{\overline{y}\}$. We endow the relation $\Lambda_1 \sim \Lambda_2$ if Λ_1 and Λ_2 have the same set of leaders of their elements.

Definition 9. Let $\mathcal{M}$ be a model of a theory T in a language $\mathcal{L}$, and suppose $p \in S_{\overline{y}}(M)$ for some finite variable context x. We say a formula $\varphi(x, y)$ is *represented* in p if there is some $a \in M^y$ so that $\varphi(x, a) \in p$.

Now let $N \models T$ with $N \supseteq M$ and let $q \in S_{\overline{y}}(N)$ be a type extending p. We say q is an *heir* of p if every formula represented in q is represented in p.

Definition 10. Suppose $k \subseteq \ell$ are Δ-fields, and $p \in S_{\overline{y}}(k)$ and $q \in S_{\overline{y}}(\ell)$ with $p \subseteq q$. We say q *does not fork* over k if for all $M, N \models DCF_{0,m}$ with $k \subseteq M$, and $M \cup \ell \subseteq N$, there are $p_1 \in S_{\overline{y}}(M)$ and $q_1 \in S_{\overline{y}}(N)$ with $p \subseteq p_1$, $q \subseteq q_1$ and q_1 and heir of p_1.

Theorem 11. Let $k \subseteq \ell$ be Δ-fields. Suppose $p \in S_n(k)$ and $q \in S_n(\ell)$ with $q \supseteq p$. Then the following are equivalent:

(a) q forks over k

(b) $\Lambda_p \not\sim \Lambda_q$

(c) $\omega_q < \omega_p$.

We start by proving some lemmas that allow us to extend a type p to allow parameters from a larger, fields while retaining the same set of leaders of characteristic set Λ_p.

Lemma 12. For any $p \in S_{\overline{y}}(k)$, there is an extension $p' \supseteq p$ over k^{alg} with $\Lambda_p = \Lambda_{p'}$.

Proof. Let $\overline{a}$ be a generic realization of p, so that $I_p = I(\overline{a}/k)$, and let $\Lambda_p = \{f_1, \ldots, f_s\}$. Take $V = \Theta\overline{y} \setminus \{\theta u_f \mid \theta \in \Theta, f \in \Lambda\}$. As in the proof of lemma 4, we have $k\langle\overline{a}\rangle = k((v(\overline{a}))_{v\in\Theta\overline{y}})$ is algebraic over $k((v(\overline{a}))_{v\in V})$, and $((v(\overline{a}))_{v\in V})$ algebraically independent. Hence, $((v(\overline{a}))_{v\in V})$ is a transcendence basis for $k\langle\overline{a}\rangle$ over k. Then $(v(\overline{a})_{v\in V})$ remains a transcendence basis for $k^{\text{alg}}\langle\overline{a}\rangle$ over k^{alg}. Hence,

$$k^{\text{alg}}\langle\overline{a}\rangle = k^{\text{alg}}((v(\overline{a}))_{v\in V}, (\theta u_f(\overline{a}))_{\theta\in\Theta, f\in\Lambda})$$

Take $p' \in S_{\overline{y}}(k^{\text{alg}})$ so that $I_{p'} = I(\overline{a}/k^{\text{alg}})$. Let $V' = \{v \in V \mid v \preceq u_{f_s}$. Let $b_i = u_{f_i}(\overline{a})$, and let g_i^* be a minimal polynomial for b_i over $k^{\text{alg}}((v(\overline{a}))_{v\in V'}, (b_j)_{j<i})$. By clearing denominators if necessary, this gives rise to a polynomial $g_i \in k^{\text{alg}}[V', (u_{f_j})_{j\le i}]$ so that $g(\overline{a}) = 0$, and $u_{g_i} = u_{f_i}$. Moreover, since f_i induces another such polyomial for b_i, and g_i^* is minimal, we have that g_i^* divides this induced polynomial. Hence, every term that appears in g_i must be a factor of one that appears in f_i. Hence, the g_i form an autoreduced set. We claim $\Lambda = \{g_1, \ldots, g_s\}$ is a characteristic set for $I_{p'}$. Suppose not. Then there is some $1 \le i \le s+1$ and an h so that h is reduced with respect to each g_j for $j < i$, and either $i = s+1$ or the rank of h is less than that of g_i. In this case, h cannot have the same leader as any g_j. Moreover, since h is reduced with respect to each g_j that comes before it, it would follow that h lives in $k^{\text{alg}}[V]$. But this would contradict $(v(\overline{a}))_{v\in V'}$ being a transcendence basis for $k^{\text{alg}}\langle\overline{a}\rangle$ over k^{alg}. Now, since $u_{f_i} = u_{g_i}$, we have $\Lambda \sim \Lambda_p$. □

We will state and use the following facts from algebra without proof

Proposition 13. If R_0 is a subring of R so that R is a free R_0 module, and I is an ideal of R_0, then $(RI) \cap R_0 = I_0$.

Proposition 14. Let k be a field of characteristic 0, and let $k' \supseteq k$ be a field extension. Let $\mathfrak{p}$ be a prime ideal of a polynomial algebra $k[(x_j)_{j\in J}]$ and $\overline{a} = (a_j)_{j\in J}$ a generic zero of $\mathfrak{p}$. Then $k'\mathfrak{p}$ is prime if and only if for each $b \in k'$ that is algebraic over k, the minimal polynomial of b over k is irreducible over $k(\overline{a})$.

We also use the following theorem used as part of the axiomitization of $DCF_{0,m}$.

Theorem 15. For k a Δ-field, and $\Lambda \subseteq k\{\overline{y}\}$ a finite set, the following are equivalent:

(a) Λ is the characteristic set of a prime differential ideal of $k\{\overline{y}\}$

(b) Λ is coherent, and $(\Lambda) : H_\Lambda^\infty$ is a prime ideal containing no nonzero element that is reduced with respect to Λ.

If Λ is the characteristic set of a prime differential ideal $\mathfrak{p}$, then $\mathfrak{p} = [\Lambda] : H_\Lambda^\infty$.

Lemma 16. Let k be an algebraically closed Δ-field and let $p \in S_{\overline{y}}(k)$. Let $K \models DCF_{0,m}$ be a field extension of k. Then there is a $p' \in S_{\overline{y}}(K)$ such that $\Lambda_{p'} = \Lambda_p$.

Proof. Let $R_0 = k\{\overline{y}\}$ and $R = K\{\overline{y}\}$. Our goal will be to show that Λ_p is the characteristic set of a prime ideal in R. Then we may take p' so that $I_{p'} = R[\Lambda_p] : H_{\Lambda_p}^\infty$. By Theorem 15, it is equivalent to show that Λ_p is coherent (in R) and that $(R\Lambda_p) : H_{\Lambda_P}^\infty$ is a prime ideal containing no nonzero element which is reduced with respect to Λ.

Firstly, since I_p is a prime differential ideal in R_0, Theorem 15 tells us Λ_p is coherent in R_0. Passing to the larger ring R does not change the fact that Λ_p is coherent. Next we show $(R\Lambda_p) : H_{\Lambda_P}^\infty$ is prime. Proposition 13 gives us that

$$R((R_0\Lambda_p) : H_{\Lambda_p}^\infty) \cap R_0 = (R_0\Lambda_p) : H_{\Lambda_p}^\infty.$$

Now we can apply proposition 14 with $k' = K$ to get that $K((R_0\Lambda_p) : H^\infty_{\Lambda_p})$ is prime. The hypothesis holds trivially since k is algebraically closed. Note that $R = K\{\overline{y}\} = KR_0$. Therefore, for any ideal I of R_0 we have

$$KI = KR_0I = RI.$$

So, in fact we have $R((R_0\Lambda_p) : H^\infty_{\Lambda_p})$ is prime.

Next we claim $R((R_0\Lambda_p) : H^\infty_{\Lambda_p}) = (R\Lambda_p) : H^\infty_{\Lambda_p}$ so that $(R\Lambda_p) : H^\infty_{\Lambda_p}$ is prime. Clearly $(R\Lambda_p) : H^\infty_{\Lambda_p} \supseteq (R_0\Lambda_p) : H^\infty_{\Lambda_p}$. From this we get one direction of inculsion

$$(R\Lambda_p) : H^\infty_{\Lambda_p} \supseteq R((R_0\Lambda_p) : H^\infty_{\Lambda_p}).$$

Now we show the other direction. Suppose $g \in (R\Lambda_p) : H^\infty_{\Lambda_p}$. By definition, there is some $r \in \mathbb{N}$ so that $gH^r_{\Lambda_p} \in R\Lambda_p$. Clearly $R\Lambda_p \subseteq R((R_0\Lambda_p) : H^\infty_{\Lambda_p})$. Since we've shown the latter is prime, we must have either g or $H^r_{\Lambda_p}$ is in $R((R_0\Lambda_p) : H^\infty_{\Lambda_p})$. The latter is not the case because $H^r_{\Lambda_p}$ is nonzero and reduced with respect to each element of Λ_p. Hence, we must have $g \in R((R_0\Lambda_p) : H^\infty_{\Lambda_p})$. This proves the claim.

It remains to show that $(R\Lambda_p) : H^\infty_{\Lambda_p}$ contains no nonzero elements that are reduced with respect to each element of Λ_p. Suppose to the contrary that such an element g existed. We may write $g = \sum_{j=1}^r a_j A_j$ where $a_j \in K$, and each A_j is a monomial generated by $v \in \Theta\overline{y}$ of rank bounded by u_g, and reduced with respect to each element of Λ_p. Let $\Lambda_p = \{f_1, \ldots, f_s\}$. Since $g \in (R\Lambda_p) : H^\infty_{\Lambda_p}$, there is some $n \in \mathbb{N}$ so that $gH^n_{\Lambda_p} \in R\Lambda_p$. We may write $H_{\Lambda_p} = \sum_{i=1}^t b_i B_i$ where $b_i \in k$ and B_i is a monomial generated by $v \in \Theta\overline{y}$ of rank bounded by u_{f_s}. Now, $gH^n_{\Lambda_p} \in R\Lambda_p$ implies that we may write $gH^n_{\Lambda_p} = \sum_{i=1}^s P_i f_i$ for some $P_i \in R$. Finally, we may write each P_i as $\sum_{j=1}^{r_i} c_{i,j} C_{i,j}$, where $c_{i,j} \in K$ and $C_{i,j}$ is a monomial generated by $v \in \Theta\overline{y}$ bounded by u_{f_s}. We may treat elements of $\Theta\overline{y}$ as formal symbols, and define the sentence ϕ by

$$\exists (z_i)_{i\le r} \neq \overline{0}\ \exists (w_{i,j})_{i\le s, j\le r_s} \neq \overline{0}\ \left(\sum_{i=1}^t b_i B_i\right)^n \sum_{i=1}^r z_i A_i = \sum_{i=1}^s \left(\sum_{j=1}^{r_i} w_{i,j} C_{i,j}\right) f_i.$$

In reality, this equation of polynomials in $\Theta\overline{y}$ regarded as formal symbols is shorthand for a system of equations that equates the coefficients of these polynomials. Now, observe that the constants involved in φ are the b_i, and those that come from the f_i. These all come from k, so φ is a formula in $\mathcal{L}(k)$. Importantly, the tuples $(a_i)_{i\le r}$ and $(c_{i,j})_{i\le s, j\le r_s}$ witness that $K \models \phi$. Since this is a purely algebraic formula and k is algebraically closed, we also have $k \models \phi$. Hence, there are $(a'_i)_{i\le r} \in k$ not all 0 and $(c'_{i,j})_{i\le s, j\le r_s} \in k$ not all 0, satisfying the above. In other words, if we let

$$g' = \sum_{j=1}^r a'_i A_i \qquad P = \sum_{i=1}^s \left(\sum_{j=1}^{r_i} c'_{i,j} C_{i,j}\right) f_i.$$

Then we have $g' H^n_{\Lambda_p} = P \in R_0(\Lambda_p)$, hence $g' \in (R_0\Lambda_p) : H^\infty_{\Lambda_p}$. Moreover, since the monomials A_i are reduced with respect to each element of Λ_p, so is g'. This contradicts Λ_p being a characteristic set for the prime ideal I_p. This completes the proof. □

Lemma 17. Suppose $\Lambda_1 < \Lambda_2$ with $\Lambda_1 \not\sim \Lambda_2$. Then there exists some $f \in \Lambda_1$ that is reduced with respect to Λ_2

Proof. List $\Lambda_1 = \{f_1 < \cdots < f_s\}$ and $\Lambda_2 = \{g_1 < \cdots < g_r\}$. Since $\Lambda_1 < \Lambda_2$, there is a unique $i \le s$ such that $\operatorname{rank}(g_j) = \operatorname{rank}(f_j)$ for $j < i$, and either $i = r + 1$, or $\operatorname{rank}(f_i) < \operatorname{rank}(g_i)$. Observe that

f being reduced to g depends only on f and $\operatorname{rank}(g)$. Then since f_i is reduced with respect to each f_j for $j < i$, and $\operatorname{rank}(f_j) = \operatorname{rank}(g_j)$ for $j < i$, we get that f_i is reduced with respect to each g_j for $j < i$. Furthermore, for $j \geq i$, we have that $\operatorname{rank}(f) < \operatorname{rank}(g_i) \leq \operatorname{rank}(g_j)$. It follows that f is reduced with respect to g_j for $j \geq i$ as well. □

With one more lemma which we will state and not prove, we are ready to prove Theorem 11

Lemma 18. Suppose that for every $K \supseteq \ell$ with $K \models DCF_{0,m}$, there is a $p_1 \in S_{\overline{y}}(K)$ extending p such that every $q_1 \in S_{\overline{y}}(K)$ extending q represents a formula not represented by p_1. Then q forks over k.

We restate and prove the algebraic characterization of forking.

Theorem 11. Let $k \subseteq \ell$ be Δ-fields. Suppose $p \in S_n(k)$ and $q \in S_n(\ell)$ with $q \supseteq p$. Then the following are equivalent:

(a) q forks over k

(b) $\Lambda_p \not\sim \Lambda_q$

(c) $\omega_q < \omega_p$.

Proof. Fix p, q, k, ℓ as in the theorem.
(b) ⇒ (a) Suppose $\Lambda_q \not\sim \Lambda_p$. We will use lemma 18 to show q forks over k. Let $K \supseteq \ell$ with $K \models DCF_{0,m}$ be arbitrary. Consider the algebraic closure of k as being within K, proposition 12, there is a $\tilde{p} \in S_{\overline{y}}(k^{\text{alg}})$ so that $\Lambda_{\tilde{p}} \sim \Lambda_p$. Choose $\tilde{p}$ so that $\Lambda_{\tilde{p}}$ is of maximal rank. This ammounts to maximizing the degrees of the elements of $\Lambda_{\tilde{p}}$. Since necessarily $\Lambda_{\tilde{p}} \leq \Lambda_p$, this serves as an upper bound. Hence, this maximum is well defined. Now apply proposition 16 to $\tilde{p}$ to get a $p' \in S_{\overline{y}}(K)$ with $\Lambda_{p'} = \Lambda_{\tilde{p}}$.

Now, suppose $q' \in S_{\overline{y}}(K)$ is an extension of q. We wish to show q' is not an heir of p'. If $\Lambda_{q'} < \Lambda_{p'}$ and $\Lambda_{q'} \not\sim \Lambda_{p'}$ then $\Lambda_{q'}$ would contain some g that is reduced with respect to $\Lambda_{p'}$. Then $g(\overline{y}) = 0 \in q$ is a formula which cannot be represented in p'. Hence, the hypothesis of lemma 18 would be satisfied. Therefore, it suffices to show $\Lambda_{q'} < \Lambda_{p'}$ and $\Lambda_{q'} \not\sim \Lambda_{p'}$.

Let $\tilde{q} = q' \restriction k^{\text{alg}}$. Since $q' \supseteq q \supseteq p$, it follows that $\tilde{q} \supseteq p$. By the maximality of $\Lambda_{\tilde{p}}$ among the characteristic sets of types over k^{alg} extending p, we get $\Lambda_{\tilde{q}} \leq \Lambda_{\tilde{p}} = \Lambda_{p'}$. As $q' \supseteq \tilde{q}$, we also have $\Lambda_{q'} \leq \Lambda_{\tilde{q}}$. Thus, we have

$$\Lambda_{q'} \leq \Lambda_{p'}.$$

Finally, since $\Lambda_p \sim \Lambda_{p'}$ and $\Lambda_p \not\sim \Lambda_q \geq \Lambda_{q'}$, we get $\Lambda_p \not\sim \Lambda_q$.

(a) ⇒ (b) We prove the contrapositive. Suppose $\Lambda_q \sim \Lambda_p$. Let $M, N \models DCF_{0,m}$ with $k \subseteq M$ and $\ell \cup M \subseteq N$. Applying lemmas 12 and 16 to q, there is a $q' \in S_{\overline{y}}(N)$ extending q so that $\Lambda_{q'} \sim \Lambda_q$. Now, $q' \supseteq p$. Let $p' = q \restriction M \in S_{\overline{y}}(M)$. We wish to show q' is the heir of p'.

Borrowing the argument as in proposition 16, we have that $\Lambda_{p'}$ is a characteristic set in N. Therefore, $N[\Lambda_{p'} : H^{\infty}_{\Lambda_{p'}}$ is a prime Δ-ideal. Furthermore, $p' \subseteq q'$ so $N[\Lambda_{p'}] : H^{\infty}_{\Lambda_{p'}} \subseteq I_{q'}$. Moreover, we have $\Lambda_{q'} \leq \Lambda p' \leq \Lambda_p$ and $\Lambda_p \sim \Lambda_q \sim \Lambda_{q'}$. Hence, $\Lambda_{p'} \sim \Lambda_{q'}$. Now, since the Kolchin polynomial only depends on the leaders of the elements of an ideal's characteristic set, it follows that $\omega_{p'} = \omega_{q'}$. By lemma 8, it follows that $N[\Lambda_{p'}] : H^{\infty}_{\Lambda_{p'}} = I_{q'}$. Without loss of generality, we may assume

$$\Lambda_{q'} = \Lambda_{p'} = \{f_1, \ldots, f_s\}.$$

It remains to show that an arbitrary formula represented in q' is represented in p'. Suppose $\varphi(\overline{y}, \overline{a}) \in q'$. We want to find $\overline{b} \in M$ so that $\varphi(\overline{y}, \overline{b}) \in p'$. From $N[\Lambda_{p'}] : H^{\infty}_{\Lambda_{p'}} = I_{q'}$, we have that

$$q' = \left\{ \psi(\overline{y}) \in \mathcal{L}_{\overline{y}}(N) \mid DCF_{0,m} \cup \left\{ \bigwedge_{i=1}^{s} f_i(\overline{y}) = 0 \wedge g(\overline{y}) \neq 0 \mid g \text{ reduced wrt } \Lambda_{q'} \right\} \vdash \psi(\overline{y}) \right\}$$

Notice that the finite conjunction of formulas of the form $\bigwedge_{i=1}^{s} f_i(\overline{y}) = 0 \wedge g(\overline{y}) \neq 0$ can also be written in this form by taking the product of the various g's. Hence, by compactness, there is a single $g(\overline{y}, \overline{a})$ that is reduced with respect to $\Lambda_{q'}$ such that

$$DCF_{0,m} \vdash \left(\bigwedge_{i=1}^{s} f_i(\overline{y}) = 0 \wedge g(\overline{y}, \overline{z}) \neq 0 \right) \rightarrow \varphi(\overline{y}, \overline{z}).$$

But notice that $\bigwedge_{i=1}^{s} f_i(\overline{y}) = 0$ is in p' since $\Lambda_{p'} = \{f_1, \ldots, f_s\}$. Also, g is reduced with respect to $\Lambda_{p'}$ for the same reason. Hence, $g(\overline{y}, \overline{b}) \neq 0 \in p'$ for all $\overline{b}$. Thus, φ is represented in p'.

(c) $\Rightarrow$ (b) Suppose $\omega_q < \omega_p$. Since ω_p depends only on the leaders of Λ_p, if $\omega_p \neq \omega_q$ then we cannot have $\Lambda_p \sim \Lambda_q$.

(b) $\Rightarrow$ (c) Suppose $\Lambda_q \not\sim \Lambda_p$. Since we have $q \supseteq p$, we have $I_q \subseteq \ell\{\overline{y}\}$ contains $I_p \subseteq k\{\overline{y}\}$. It follows that $\omega_q \leq \omega_p$, since any algebraic relation bringing down the dimension of $(I_p)_s$ will also bring down the dimension of $(I_q)_s$. We wish to show strict inequality. It suffices to find some $v \in \Theta\overline{y}$ so that v is the derivative of some leader of a $g \in \Lambda_q$, but not the derivative of any leader in Λ_p. From lemma 17, there is some $g \in \Lambda_q$ that is reduced with respect to Λ_p. Then taking $v = u_g$ suffices. $\square$

References

[1] E.R. Kolchin. *Differential Algebra & Algebraic Groups.* Pure and Applied Mathematics. Academic Press, 1973.

[2] Tracey McGrail. The model theory of differential fields with finitely many commuting derivations. *Journal of Symbolic Logic*, 65(2):885–913, 2000.