Putting topologies on dp-minimal fields

Will Johnson

June 9, 2015

Conjecture 0.1. *If K is a field of dp-rank 1, then K is algebraically closed, real closed, or admits a valuation ring making K henselian with finite or algebraically closed residue field.*

1 Notation

We denote dp-rank of a definable or $\wedge$-definable set X by $\text{dp-rk}(X)$, and the dp-rank of $\text{tp}(a/S)$ by $\text{dp-rk}(a/S)$.

Fact 1.1. *Assuming finite dp-rank...*

1. *For X and Y non-empty,*
$$\text{dp-rk}(X \times Y) = \text{dp-rk}(X) + \text{dp-rk}(Y)$$
2. $\text{dp-rk}(X) > 0$ *if and only if X is infinite*
3. *If X is definable or type-definable over S, then*
$$\text{dp-rk}(X) = \max_{x \in X} \text{dp-rk}(x/S).$$
4. *If $b \in \text{acl}(aS)$, then* $\text{dp-rk}(b/S) \le \text{dp-rk}(a/S)$.
5. $\text{dp-rk}(ab/S) \le \text{dp-rk}(a/bS) + \text{dp-rk}(b/S)$.

We work in a monster model $\mathbb{C}$ of a dp-minimal field (i.e., a field of dp-rank 1). "Models" will refer to small elementary substructures $M \preceq \mathbb{C}$. We'll write concatenation multiplicatively and multiplication using $\cdot$.

We will assume that $\mathbb{C}$ is not strongly minimal. One already knows what strongly minimal fields look like, by Macintyre's theorem—they are algebraically closed.

If X and Y are subsets of $\mathbb{U}$, we'll use $X + Y$ and $X - Y$ to denote the usual things:

$$X + Y = \{x + y : x \in X,\ y \in Y\}$$

$$X - Y = \{x - y : x \in X,\ y \in Y\}$$

Note that if X and Y are definable, then so are $X+Y$ and $X-Y$, and if M is a model, then of course

$$X(M) - Y(M) = (X-Y)(M)$$

and so on.

If $\alpha \in \mathbb{C}$, we'll write $\alpha \cdot X$ to denote

$$\{\alpha \cdot x : x \in X\}.$$

We'll let $X -_\infty Y$ denote the set of $\delta \in \mathbb{C}$ such that there are infinitely many pairs $(x,y) \in X \times Y$ with $x - y = \delta$. A priori, this needn't be a definable set, even if X and Y are.

However...

2 Infinity is definable

Lemma 2.1. *Let $X \subset \mathbb{C}$ be definable. Then X is finite if and only if there is some $\alpha \in \mathbb{C}$ such that the map $f(x,y) = \alpha \cdot x + y$ is injective on $X \times X$.*

Proof. First suppose that X is finite. Then the set

$$S = \left\{ \frac{x_1 - x_2}{x_3 - x_4} : \vec{x} \in X^4,\ x_3 \neq x_4 \right\}$$

is finite. Since $\mathbb{C}$ is infinite, there is some $\alpha \in \mathbb{C} \setminus S$. Now suppose (x,y) and (x',y') are in $X \times X$, and

$$\alpha \cdot x + y = \alpha \cdot x' + y'.$$

Then

$$\alpha \cdot (x - x') = y' - y,$$

so either $(x,y) = (x',y')$ or $\alpha \in S$. By choice of α, we see $(x,y) = (x',y')$, and so f is injective on $X \times X$.

Conversely, suppose X is infinite. Then dp-rk$(X) = 1$, so dp-rk$(X \times X) = 2$, so $X \times X$ cannot embed into the dp-minimal field $\mathbb{C}$. $\square$

Corollary 2.2. *Th($\mathbb{C}$) eliminates $\exists^\infty$.*

Corollary 2.3. *If X and Y are definable, then so is $X -_\infty Y$.*

Corollary 2.4. *If M is a model, $X \subset \mathbb{C}^1$ is $\mathbb{C}$-definable, and $X \cap M$ is infinite, then $X \cap M$ contains $Y(M)$ for some infinite M-definable set Y. In other words, infinite externally definable sets contain infinite internally definable sets. (Note: we are not asserting that $Y \subset X$.)*

Proof. By honest definitions, there is some formula $\phi(x, y)$ such that for every finite subset A_0 of $X \cap M$, we can find $m \in M$ such that

$$A_0 \subset \phi(M; m) \subset X \cap M.$$

By elimination of $\exists^\infty$, there is some N_ϕ such that whenever $\phi(M; m)$ has size bigger than N_ϕ, it is infinite. Take A_0 of size bigger than N_ϕ. Then $\phi(M; m)$ is infinite, and contained in $X \cap M$. Take $Y = \phi(\mathbb{C}; m)$. □

3 Infinitesimals

Lemma 3.1. *If X and Y are infinite, then so is $X -_\infty Y$.*

Proof. Let S be some small set over which X and Y are defined. Take $(x_0, y_0) \in X \times Y$ of dp-rank 2 over S. Let $\delta = x_0 - y_0$. Note that

$$2 = \text{dp-rk}(x_0, y_0/S) = \text{dp-rk}(x_0, \delta/S) \leq \text{dp-rk}(x_0/S\delta) + \text{dp-rk}(\delta/S).$$

Since δ and x_0 live in the home sort, the dp-ranks on the right can be at most one. So they must equal one, and equality holds. It follows that $x_0 \notin \text{acl}(S\delta)$ and $\delta \notin \text{acl}(S)$.

If $\delta \notin X -_\infty Y$, then the $S\delta$-definable set of pairs $(x, y) \in X \times Y$ with $x - y = \delta$ would be finite or empty. As (x_0, y_0) is such a pair, we would have $x_0 \in \text{acl}(S\delta)$, a contradiction. So $\delta \in X -_\infty Y$.

Now $X -_\infty Y$ is S-definable. If it were finite, $\delta \in \text{acl}(S)$, which is false. So $X -_\infty Y$ is infinite. □

Corollary 3.2. *If X and Y are two infinite sets, then there is some translate $Y + a$ of Y such that $X \cap Y + a$ is infinite.*

Proof. This is what it means for $X -_\infty Y$ to be non-empty. □

Note that if X and Y are M-definable, then we can take $a \in M$.

Corollary 3.3. *If X and Y are definable infinite sets, there is a definable infinite set Z such that*

$$Z -_\infty Z \subset (X -_\infty X) \cap (Y -_\infty Y).$$

In particular, the family of sets of the form $X -_\infty X$, with X infinite and definable, is directed. *The same holds if we replace "definable" with "M-definable" for some model M.*

Proof. After translating Y, we may assume X and Y have infinite intersection, by the previous corollary. Let $Z = X \cap Y$. □

Definition 3.4. *Let M be a small model. Say that $\epsilon \in \mathbb{C}$ is* M-infinitesimal*, and write $\epsilon \approx_M 0$, if $\epsilon \in X -_\infty X$ for every infinite M-definable set X.*

Remark 3.5. *1. The set of M-infinitesimals is type-definable over M.*

2. 0 *is an M-infinitesimal*

3. *If $\Sigma(x)$ defines the M-infinitesimals, then for every finite subtype $\Sigma_0(x)$, there's an M-definable set X such that*

$$(x \in X -_\infty X) \implies \Sigma_0(x).$$

4. *There are infinitely many M-infinitesimals, because each set $X -_\infty X$ is infinite.*

5. *The set of M-infinitesimals is closed under multiplication by M, because for any $\alpha \in M^\times$, and any M-definable set X, one has*

$$\alpha \cdot (X -_\infty X) = (\alpha \cdot X) -_\infty (\alpha \cdot X).$$

Lemma 3.6. *Zero is the only M-infinitesimal element of M.*

Proof. By the last point of the remark, either all elements of $M^\times$ are M-infinitesimal, or none of them are.

We assumed that $\mathbb{C}$ was not strongly minimal. Take some M-definable set X which is neither finite nor cofinite. By Lemma 3.2, there is some $a \in M$ such that

$$D = X \cap ((\mathbb{C} \setminus X) + a)$$

is infinite. Then $D \cap (D + a)$ is empty, because if $x + a \in D + a$, then $x \in D \subset X$, so $x + a \notin (\mathbb{C} \setminus X) + a \supseteq D$.

It follows that $a \notin D -_\infty D$. As D is infinite, it follows that a is not an M-infinitesimal. As $a \in M$, it follows that all non-zero elements of M aren't M-infinitesimal. □

4 Sums of infinitesimals

The goal here is to show that the sum of two M-infinitesimals is M-infinitesimal.

Definition 4.1. *A $\mathbb{C}$-definable map $f : \mathbb{C}^1 \to \mathbb{C}^1$ is M-*small *if for every infinite M-definable set X, the intersection $X \cap f^{-1}(X)$ is infinite.*

So for instance, the map $x \mapsto x + a$ is M-small if and only if a is M-infinitesimal.

Lemma 4.2. *If $f = f_a$ is defined over some parameter a, and f_a is M-small, and* $\operatorname{tp}(a'/M')$ *is an heir of* $\operatorname{tp}(a/M)$ *for some $M' \succeq M$, then $f_{a'}$ is M'-small.*

Proof. Let X_b be some M'-definable set, for some tuple $b \in M'$. Suppose for the sake of contradiction that $X_b \cap f_{a'}^{-1}(X_b)$ is finite. Note that $\operatorname{tp}(b/Ma')$ is a coheir of $\operatorname{tp}(b/M)$, so it is finitely satisfiable in M. By elimination of $\exists^\infty$, we can therefore find $c \in M$ such that $X_C \cap f_{a'}^{-1}(X_c)$ is finite but X_c is infinite. Since $a' \equiv_c a$, it follows that $X_c \cap f_a^{-1}(X_c)$ is also finite, contradicting M-smallness of f_a. □

Lemma 4.3. *Suppose X is M-definable and infinite, f is an M-small map, and that $X(M) \cap f^{-1}(X(\mathbb{C})) = \emptyset$. Then we have a contradiction.*

Proof. Let a be a tuple over which f is defined, and write f as f_a. By induction, build a sequence $a_1, a_2, \ldots$, and $M_0, M_1, M_2, \ldots$ so that

- $M_0 = M$
- $\mathrm{tp}(a_i/M_{i-1})$ is an heir of $\mathrm{tp}(a/M)$.
- $M_i \succeq M_{i-1}$ and $a_i \in M_i$.

For each finite string $b \in \{0,1\}^{<\omega}$, let Y_b be the definable set

$$Y_b = \left\{ x \in X : \bigwedge_{i=1}^{|b|} (f_{a_i}(x))^{b_i} \right\}$$

where $\phi^1 = \phi$ and $\phi^0 = \neg\phi$. If all the Y_b are non-empty, then we have violated NIP.

Claim 4.4. *Each Y_b is infinite.*

Proof. We proceed by induction on the length of b. For the base case, length 0, $Y_b = X$, which is infinite by assumption.

Now suppose we know that Y_b is infinite. It suffices to show that Y_{b0} and Y_{b1} are infinite.

Let n be the length of b. Then Y_b is defined over M_n. Since Y_b is infinite, so is $Y_b(M_n)$. We claim that $Y_b(M_n) \subset Y_{b0}(\mathbb{C})$, so that Y_{b0} is infinite. The only thing to check is that if $x \in Y_b(M_n)$, then $f_{a_{n+1}}(x) \notin X$. Suppose otherwise. As $x \in \mathrm{dcl}(M_n)$ and $\mathrm{tp}(a_{n+1}/M_n)$ is an heir over M, it follows that $\mathrm{tp}(x/Ma_{n+1})$ is finitely satisfiable in M. Since $x \in Y_b \subset X$ and $f_{a_{n+1}}(x) \in X$, we can find $x' \in M$ such that $x' \in X$ and $f_{a_{n+1}}(x') \in X$. As $a_{n+1} \equiv_M a$, it follows that $f(x') = f_a(x') \in X$. So

$$x' \in X(M) \cap f^{-1}(X(\mathbb{C})) = \emptyset$$

a contradiction. So $Y_b(M_n) \subset Y_{b0}$, and Y_{b0} is infinite.

It remains to show that Y_{b1} is infinite. But by the previous lemma, we know that $f_{a_{n+1}}$ is certainly M_n-small. Since Y_b is infinite, so is $Y_b \cap f_{a_{n+1}}^{-1}(Y_b)$. So there are infinitely many $x \in Y_b$ such that $f_{a_{n+1}}(x) \in Y_b \subset X$. Each of these x is in Y_{b1}, so Y_{b1} is infinite. □

Now, because each Y_i is infinite, the formula

$$\phi(x, y) \iff f_x(y) \in X$$

has the independence property, a contradiction. □

Lemma 4.5. *Suppose f is M-small, and X is M-definable. Then there are only finitely many $x \in X(M)$ such that $f(x) \notin X$.*

Proof. If not, then by Corollary 2.4, we can find an infinite M-definable set Q such that

$$x \in Q(M) \implies x \in X \wedge f(x) \notin X$$

Let $X' = X \cap Q$. Then $X'(M) = Q(M)$ is infinite, and if $x \in X'(M)$, then $f(x) \notin X$, hence $f(x) \notin X'$. So

$$X'(M) \cap f^{-1}(X'(\mathbb{C})) = \emptyset,$$

contradicting Lemma 4.3. $\square$

Corollary 4.6. *If f is M-small, then for every M-definable set X there is a finite subset $S \subset M$ such that for $x \in M \setminus S$,*

$$x \in X \iff f(x) \in X$$

Proof. By the previous lemma, there are only finitely many $x \in M$ such that $x \in X$ fails to imply $f(x) \in X$. Replacing X with its complement, we get that there are only finitely many $x \in M$ such that $f(x) \in X$ fails to imply $x \in X$. $\square$

Corollary 4.7. *The set of definable bijections $\mathbb{C} \to \mathbb{C}$ which are M-small, is a group.*

Proof. The identity map is clearly M-small. Suppose f and g are M-small. We will show that $h := f \circ g^{-1}$ is M-small.

Let X be infinite and M-definable. By the previous corollary, for all but finitely many $x \in M$, we have the equivalences

$$g(x) \in X \iff x \in X \iff f(x) = h(g(x)) \in X$$

In particular, for infinitely many $x \in X(M)$, we have $g(x) \in X$ and $h(g(x)) \in X$, or equivalently,

$$g(x) \in X \cap h^{-1}(X).$$

By injectivity of g, it follows that $X \cap h^{-1}(X)$ is infinite. Since X was arbitrary, h is M-small. $\square$

Corollary 4.8. *The sum of two M-infinitesimals is M-infinitesimal. In fact, M-infinitesimals are a subgroup of the additive group.*

Proof. If ϵ_1 and ϵ_2 are M-infinitesimal, then the maps $x \mapsto x + \epsilon_i$ are M-small, so their composition is M-small, which means exactly that $\epsilon_1 + \epsilon_2$ is an infinitesimal.

We have already noted that 0 is an M-infinitesimal, and that M-infinitesimals are closed under multiplication by elements of M, including -1. $\square$

Corollary 4.9. *For any infinite definable set X, there is an infinite definable set Y such that*

$$(Y -_\infty Y) + (Y -_\infty Y) \subset X -_\infty X$$

Proof. Let M be a model over which X is defined. If $\Sigma(x)$ is the partial type over M asserting that x is an M-infinitesimal, we have just seen that $\Sigma(x), \Sigma(y) \vdash \Sigma(x+y)$. Since $x \in Y -_\infty Y$ is one of the formulas of $\Sigma(x)$, there must be a finite subtype $\Sigma_0(x)$ such that

$$\Sigma_0(x) \wedge \Sigma_0(y) \implies x + y \in X -_\infty X.$$

By part of Remark 3.5, there is some M-definable set Y such that

$$x \in (Y -_\infty Y) \implies \Sigma_0(x).$$

Then Y works. $\square$

5 The topology

Fix a model M. For the duration of this section, we will conflate X with $X(M)$.

Definition 5.1. *A* basic neighborhood *of $a \in M$ is a set of the form $(X -_\infty X) + a$ with X infinite and M-definable.*

The basic neighborhoods of a are filtered, infinite, and contain a, by Corollary 3.3 and Lemma 3.1.

Theorem 5.2. *Let M be a small model. Then there is a non-discrete Hausdorff topology on M such that for each $a \in M$, the "basic neighborhoods" of a are a neighborhood basis for the point a. Moreover, the additive group operations, as well as multiplication by fixed elements of M, are continuous.*

Proof. Say that a set $U \subset M$ (not necessarily definable) is open if for each $a \in U$, some "basic neighborhood" of a is contained in U.

The fact that "basic neighborhoods" are filtered implies that finite intersections of opens are open. It is clear that arbitrary unions of opens are open. So we certainly have a topology.

Claim 5.3. *If S is any subset of M and $a \in M$, then a is in the interior of S if and only if some "basic neighborhood" of a is contained in S.*

Proof. Let T denote the set of $a \in S$ such that some basic neighborhood of a is contained in T. We claim that T is open. Indeed, if $a \in T$, then there is some infinite M-definable set X such that $a + (X -_\infty X)$ is contained in S. By Lemma 4.9, we can find an infinite M-definable set Y such that

$$(Y -_\infty Y) + (Y -_\infty Y) \subset (X -_\infty X).$$

Then $a + (Y -_\infty Y)$ is contained in T—so some "basic neighborhood" of a is contained in T. As a was arbitrary, T is open.

Therefore T is contained in the interior S^{int} of S. Conversely, if $a \in S^{int}$, then by openness of S^{int} there is some "basic neighborhood" of a contained in S^{int}, hence in S. So $a \in T$. It follows that $T = S^{int}$. $\square$

Consequently, each "basic neighborhood" of a is an *actual neighborhood* of a, and the collection of basic neighborhoods of a is an actual neighborhood basis.

To see that addition is continuous, let $a, b \in M$ and $(a+b)+(X -_\infty X)$ be some basic neighborhood of $a+b$. By Lemma 4.9, we can find basic neighborhoods $a + (Y -_\infty Y)$ and $b + (Y -_\infty Y)$ such that

$$(a + (Y -_\infty Y)) + (b + (Y -_\infty Y)) \subset (a+b) + (X -_\infty X)$$

This shows that addition is continuous.

If $\alpha \in M$ is some fixed scalar, the map $x \mapsto \alpha \cdot x$ is continuous. This is trivial if $\alpha = 0$, and otherwise,

$$\alpha \cdot (a + ((\alpha^{-1} \cdot X) -_\infty (\alpha^{-1} \cdot X))) \subset \alpha \cdot a + (X -_\infty X)$$

so multiplication by α is continuous at a.

It follows that the group operations are continuous, so we have a topological group.

Because group operations and mulitplication by fixed scalars are continuous, Hausdorffness can be checked at 0 and 1. We noted earlier (Lemma 3.6) that 1 is not an M-infinitesimal. So there is some infinite M-definable set X such that $1 \notin X -_\infty X$. Let Y be an infinite set from Lemma 4.9. Then

$$(0 + (Y -_\infty Y)) \cap (1 + (Y -_\infty Y)) = \emptyset$$

so 0 and 1 can be separated by basic neighborhoods.

Finally, we check that the topology is not discrete, i.e., that $\{0\}$ is not an open set. If it were open, then by the claim, there would be a basic neighborhood of 0 containing nothing but 0. That is, there would be an infinite definable set X such that $X -_\infty X = \{0\}$. This contradicts Lemma 3.1. $\square$

Note that we have *not* shown that multiplication or division are continuous. I know how to show continuity of multiplication, which boils down to showing that products of infinitesimals are infinitesimal.

We have also not shown that there is a definable basis of opens; the basis is currently ind-definable. I think we could get a definable basis, if we knew that division were continuous away from 0. More on this later...

6 Germs of sets

The goal of this section is to show the following related results:

- Every infinitesimal type over M is M-definable.
- Every M-definable set has finite boundary.
- In a definable family of subsets of $\mathbb{C}$, there are only finitely many "germs at 0."

The third of these is the key to the others. We would like to mimic the proof that there are finitely many germs at 0 in the ordered case. To mimic the argument, it would seem like we need the following, which we will call Condition (*):

> There is an $\emptyset$-definable family X_a of infinite sets such that for any finite subset $S \subset \mathbb{C}^\times$, there is some $a \in \mathbb{C}$ such that $X_a -_\infty X_a$ is disjoint from S.

This would certainly be true if we had a definable basis, so failure of (*) is like a strong failure of the existence of a definable basis.

We will see that (*) failing has some peculiar consequences, which eventually yield a weak approximation to (*) that is still good enough–Lemma 6.8. (Later we will see that (*) always holds.)

Lemma 6.1. *Assume that (*) fails. If X is any infinite $\mathbb{C}$-definable set, and M is a small model, then $(X -_\infty X) \cap M \supsetneq \{0\}$.*

Proof. Suppose not. Then there is an infinite $\mathbb{C}$-definable set X such that $(X -_\infty X) \cap M = \{0\}$.

Write X as X_a. Now, for any $m_1, \ldots, m_n \in M^\times$,

$$\mathbb{C} \models \exists x : |X_x| = \infty \wedge \{m_1, \ldots, m_n\} \cap (X_x -_\infty X_x) = \emptyset$$

Hence

$$M \models \exists x : |X_x| = \infty \wedge \{m_1, \ldots, m_n\} \cap (X_x -_\infty X_x) = \emptyset$$

Then, as $m_1, \ldots, m_n$ were arbitrary,

$$M \models \forall y_1, \ldots, y_n \exists x : |X_x| = \infty \wedge \{y_1, \ldots, y_n\} \cap (X_x -_\infty X_x) = \emptyset$$

As this holds for all n, condition (*) holds, a contradiction. $\square$

Lemma 6.2. *Assume that (*) fails. If X is any infinite $\mathbb{C}$-definable set, and M is a small model, then $(X -_\infty X) \cap M$ is infinite.*

Proof. Let T denote the intersection. Suppose T is finite. By Hausdorffness of the topology on M, there is some basic neighborhood of 0 disjoint from the non-zero elements of T. So, there is some M-definable infinite set Y such that $(Y -_\infty Y) \cap M \cap T = \{0\}$. By Lemma 3.3, we can find an infinite $\mathbb{C}$-definable set Z such that

$$Z -_\infty Z \subset (X -_\infty X) \cap (Y -_\infty Y).$$

Now $(Z -_\infty Z) \cap M \subset (X -_\infty X) \cap M = T$, and $(Z -_\infty Z) \cap T = \{0\}$, so it follows that $(Z -_\infty Z) \cap M = \{0\}$. This contradicts Lemma 6.1. $\square$

Lemma 6.3. *Assume that (*) fails. Let $M \preceq N$ be an inclusion of models. Then the natural topology on M (coming from Theorem 5.2) is the same as the induced topology on M as a subset of N, with the natural topology on N.*

Note that this is false in the familiar settings, where M is often discrete as a subset of N!

Proof. Let U be an open subset of N. We first need to show that $M \cap U$ is open. Given $a \in M \cap U$, we need to find a basic M-neighborhood of a contained in U.

Translating everything, we may assume $a = 0$, for simplicity.

Since $0 \in U$, some basic N-neighborhood of 0 is in U. So, there is some infinite N-definable set X such that

$$X(N) -_\infty X(N) \subset U$$

By Lemma 4.9, there is some N-definable set Y such that

$$(Y(N) -_\infty Y(N)) + (Y(N) -_\infty Y(N)) \subset U.$$

By Lemma 6.2, $(Y -_\infty Y) \cap M$ is infinite.[1] By Corollary 2.4, there is some M-definable infinite set Q such that

$$Q(M) \subset (Y -_\infty Y) \cap M.$$

The set $(Y -_\infty Y)$ is closed under negation, so if $x, y \in Q(M)$, then x and $-y$ are in $(Y -_\infty Y)(N)$. Thus $x - y \in U$. It follows that

$$(Q -_\infty Q)(M) = Q(M) -_\infty Q(M) \subseteq Q(M) - Q(M) \subset U \cap M.$$

As Q is infinite, the set $(Q -_\infty Q)(M)$ is a basic M-neighborhood of 0. So 0 is in the M-interior of $U \cap M$.

This shows that every N-open set, intersected with M, is M-open.

It remains to show that all M-open sets arise this way. It is equivalent to show that all M-closed sets arise as intersections with M of N-closed sets. Let $C \subset M$ be M-closed. Let C' be the closure of M within N. We claim that $C = C' \cap M$. Otherwise, take $a \in C' \cap M \setminus C$. As $a \notin C$, and the complement of C is open, there is some basic M-neighborhood of a which is disjoint from C. So, there is an infinite M-definable set X such that

$$(a + (X -_\infty X))(M) \cap C = a + (X(M) -_\infty X(M)) \cap C = \emptyset$$

So the M-definable set $a + (X -_\infty X)$ does not intersect the small set C, as $C \subset M$. Therefore,

$$a + (X(N) -_\infty X(N)) \cap C = \emptyset$$

so some basic N-neighborhood of a also avoids C. Therefore a is not in the N-closure C' of C, contradicting the choice of a. $\square$

Corollary 6.4. *Assume (*) fails. If $M \preceq N$ is an inclusion of models, then M is topologically closed, in the natural topology on N.*

[1] Note that $(Y -_\infty Y) \cap N = (Y(N) -_\infty Y(N))$, so $(Y -_\infty Y) \cap M$ is the same thing as $(Y(N) -_\infty Y(N)) \cap M$.

Proof. By the preceding lemma, the natural topologies on M and N are induced from the natural topology on $\mathbb{C}$. So we may reduce to the case $N = \mathbb{C}$. Let $\overline{M}$ denote the topological closure of M in $\mathbb{C}$. The topology on $\mathbb{C}$ is $\mathrm{Aut}(\mathbb{C})$ invariant, so $\overline{M}$ is $\mathrm{Aut}(\mathbb{C}/M)$-invariant as a set. But it is also small: every element x of $\overline{M}$ is the ultralimit of some ultrafilter on M. By Hausdorffness of the topology on $\mathbb{C}$, ultralimits are unique. There is only a bounded number of ultrafilters on M, so $\overline{M}$ is also small. By automorphism invariance, it must consist entirely of elements algebraic over M, hence in M. So $\overline{M} = M$. $\square$

Using this, we can use compactness to bound the complexity of the formulas needed to isolate any point outside a model from the model.

Lemma 6.5. *Assuming (*) fails, there is a definable family X_a of infinite sets such that for any model $M \preceq \mathbb{C}$ and any $b \notin M$, there is some $a \in \mathbb{C}$ such that $b + (X_a -_\infty X_a)$ is disjoint from M.*

Proof. Let $T = Th(\mathbb{C})$. Let T^+ be the expansion where we add a new unary predicate P and a constant c, and assert the following:

- P is an elementary substructure of the universe
- c is not in P
- For each definable family X_* in the original language, there is no a such that $c+(X_a -_\infty X_a)$ is disjoint from P.

This theory is inconsistent, by Corollary 6.4. By compactness, we get a finite set of families $X_*^1, \ldots, X_*^n$ such that whenever $M \preceq N$ are models of T, and $c \in N \setminus M$, then there is some $a \in N$ and some $1 \leq i \leq n$ such that $c + (X_a^i -_\infty X_a^i)$ is disjoint from M. Adding extra parameters, we can replace the n different families with one family. $\square$

Corollary 6.6. *There is some definable family X_* such that for every finite set $S \subset \mathbb{C}$ and every $x \in \mathbb{C} \setminus \mathrm{acl}(S)$, there is an $a \in \mathbb{C}$ such that*

$$x + (X_a -_\infty X_a) \cap S = \emptyset$$

Proof. This condition is strictly weaker than (*), so we may assume (*) fails. Then let X_* be the family from Lemma 6.5. If $x \notin \mathrm{acl}(S)$, we can find a model M containing S, with $x \notin M$. Then Lemma 6.5 proves the existence of a such that $x + (X_a -_\infty X_a)$ is disjoint from M, so certainly disjoint from S. $\square$

Lemma 6.7. *There is some definable family Z_* such that whenever $x_1, \ldots, x_n$ are elements of the home sort, with $\text{dp-rk}(x_1, \ldots, x_n) = n$, then there is some a such that*

$$(Z_a -_\infty Z_a) \cap \{x_1, \ldots, x_n\} = \emptyset$$

Proof. Let Z_* be the family from Corollary 6.6. Let t be an element of the home sort such that $\text{dp-rk}(t, x_1, \ldots, x_n) = n + 1$. Then $t \notin \operatorname{acl}(t + x_1, t + x_2, \ldots, t + x_n)$ (or else the total dp-rank would be at most n), so there is some Z_a such that

$$t + (Z_a -_\infty Z_a) \cap \{t + x_1, \ldots, t + x_n\} = \emptyset$$

which is equivalent to the goal. □

Lemma 6.8. *There is some definable family N_* of basic $\mathbb{C}$-neighborhoods such that for any finite collection $X_1, \ldots, X_n$ of infinite definable sets, there is some $c \in \mathbb{C}$ such that each $N_i \setminus T_c$ is infinite.*

Proof. Let Z_* be the definable family from the previous lemma. We claim that $N_a := (Z_a -_\infty Z_a)$ works. Fix $X_1, \ldots, X_n$ definable. By elimination of $\exists^\infty$, there is some m such that for each i and any a, $X_i \setminus N_a$ has fewer than m or infinitely many elements.

The set $\prod_{i=1}^n X_i^m$ has dp-rank nm, so we can find some tuple c_{ij} from it, with dp-rank nm. By the previous lemma, we can find a such that N_a misses all the c_{ij}'s. Since we chose m elements from each X_i, missing m elements forces N_a to miss infinitely many elements. □

Finally, we can show that there are a small number of "germs at 0."

Definition 6.9. *Two $\mathbb{C}$-definable sets X and Y have the same* germ at 0 *if there is a basic $\mathbb{C}$-neighborhood of* 0 *whose intersection with X and Y are the same.*

Note that this is an equivalence relation, because of Corollary 3.3

Theorem 6.10. *Let X_a be a definable family of subsets of the home sort. Then there are only a bounded number of germs of infinity among the X_a's.*

Proof. Choose some small model M. If the conclusion of the theorem is false, then Morley-Erdos-Rado yields an M-indiscernible sequence $a_1, a_2, \ldots$ such that the X_{a_i} have pairwise-distinct germs at 0.

Let Y_i denote $X_{a_{2i}} \Delta X_{a_{2i+1}}$. So the Y_i are an M-indiscernible sequence of sets, and each Y_i intersects every basic $\mathbb{C}$-neighborhood of 0. Also, by NIP, the Y_i are k-inconsistent for some k.

Note that by Hausdorffness of the topology on $\mathbb{C}$, each Y_i intersects each basic $\mathbb{C}$-neighborhood *infinitely*, in fact.

Let N_* be the definable family of neighborhoods from Lemma 6.8.

Claim 6.11. *Given $n, m < \omega$, we can find c_j for $1 \leq j \leq m$ such that for each $1 \leq i \leq n$ and $1 \leq j \leq m$, the following set is infinite:*

$$Y_i \cap \bigwedge_{k \leq j} N_{c_k} \setminus \bigcup_{j < k \leq m} N_{c_k}$$

(In a sense, this means we're using the N_{c_k}'s to witness the order property, in parallel in each Y_i.)

Proof. We hold n fixed, and proceed by induction on m. For $m = 0$ we are just saying that each Y_i is infinite, which we know.

Suppose we have constructed $c_1, \ldots, c_{m-1}$. For $i \leq n$ and $j < m$, let B_{ij} denote $Y_i \cap \bigwedge_{k \leq j} N_{c_k} \setminus \bigcup_{j<k<m} N_{c_k}$. By induction, we can assume that each B_{ij} is infinite. Unraveling the definitions, we need to choose c_m so that the following conditions hold:

- For $j < m$, we need $B_{ij} \setminus N_{c_m}$ to be infinite.
- For $j = m$, we need $B_{i,m-1} \cap N_{c_m}$ to be infinite.

The second point will hold regardless of how we choose c_m, since $B_{i,m-1}$ is the intersection of Y_i with a neighborhood of 0, so further intersections with neighborhoods of 0 will continue to be infinite (as Y_i has 0 in its closure).

The first point is obtainable because the N_* came from Lemma 6.8. $\square$

Now, given the claim, it follows by compactness that we can find an infinite sequence $c_1, c_2, \ldots$ such that for each i, j, the set

$$Y_i \cap \bigwedge_{k \leq j} N_{c_k} \setminus \bigcup_{k>j} N_{c_k}$$

is non-empty. In fact, we can arrange that the c_k's are an indiscernible sequence. Then the formulas $N_{c_{2k}} \setminus N_{c_{2k+1}}$ are ℓ-inconsistent for some ℓ (by NIP), so we have an inp-pattern of depth 2, coming from the Y_i's and the $N_{c_{2k}} \setminus N_{c_{2k+1}}$'s. In an NIP theory, inp patterns can be converted to ict patterns, so we've contradicted dp-minimality. $\square$

Corollary 6.12. *The set of infinitesimal types over $\mathbb{C}$ is small/bounded.*

Proof. By the theorem, we can produce a small model M such that for every $\emptyset$-definable family of definable sets $X_\bullet$, and every $a \in \mathbb{C}$, there is some $b \in M$ such that $X_a \Delta X_b$ doesn't have 0 in its closure.

We claim that infinitesimal types over $\mathbb{C}$ are determined by their restrictions to M. Suppose p and q are distinct infinitesimal types over $\mathbb{C}$. Then there is some X_a such that p is in X_a and q is not. By choice of M, there is some $b \in M$ such that X_a and X_b have the same germ at 0. So there is some basic $\mathbb{C}$-neighborhood N such that $N \cap X_a = N \cap X_b$. As p and q both live in N, it follows that p is in X_b and q is not. So p and q have different restrictions to M. $\square$

Corollary 6.13. *If M is a small model, then every infinitesimal type over M is definable.*

Proof. Let p be an infinitesimal type over M. Every heir of p is an infinitesimal type, and so p must have a bounded number of heirs over $\mathbb{C}$. This forces p to be definable. $\square$

Corollary 6.14. *If p is a global infinitesimal type, then p is definable, and in fact,* $\operatorname{acl}(\emptyset)$*-definable.*

Proof. By the previous corollary, p is definable over $\mathbb{C}$. If the codes for the definition are not algebraic over $\emptyset$, then p has too many images under automorphisms of $\mathbb{C}$. $\square$

Lemma 6.15. *Let $M \preceq N$ be an inclusion of models, and suppose ϵ is N-infinitesimal. Then $N \mathop{\smile\!\!\!\!\!\!\,|}^{u}_{M} \epsilon$, i.e., $\mathrm{tp}(N/M\epsilon)$ is finitely satisfiable in M.*

Proof. If we let p be some global heir of $\mathrm{tp}(\epsilon/N)$, then p is M-definable, so p and hence $\mathrm{tp}(\epsilon/N)$ are heirs of $\mathrm{tp}(\epsilon/M)$. $\square$

Lemma 6.16. *Let X be a definable set. Let $a \notin \mathrm{acl}(\ulcorner X \urcorner)$. Then a is not in the boundary ∂X, i.e., there is a $\mathbb{C}$-definable neighborhood of a contained entirely in X or entirely in its complement.*

Proof. Replacing X with its complement, we may assume $a \in X$. By the assumption on algebraic closures, we can find a model M over which X is defined, but not a. Let N be an extension of M containing a. Assume for the sake of contradiction that no neighborhood of a is wholly within X. Then there is some N-infinitesimal ϵ such that $a + \epsilon \notin X$. By the previous lemma, $\mathrm{tp}(a/M\epsilon)$ is finitely satisfiable in M. By Lemma 4.5 applied to shift by ϵ, the set S of $m \in M$ such that $m \in X$ but $m + \epsilon \notin X$, is finite. The type of a over $M\epsilon$ asserts that $x \in X$, $x + \epsilon \notin X$, and $x \notin S$. This is not finitely satisfiable in M, a contradiction. $\square$

Theorem 6.17. *For any definable set X, the boundary ∂X of X is finite.*

Proof. It suffices to show that the frontier $X \setminus X^{int}$ is finite. This set is type-definable, as $a \in X \setminus X^{int}$ if $a \in X$ and for every definable family Y_*, there is no b such that $a + (Y_b -_{\infty} Y_b) \subset X$. We have also seen that this set is bounded, being contained in $\mathrm{acl}(\ulcorner X \urcorner)$. So it is finite. $\square$

7 Multiplication

Fix a model M.

Definition 7.1. *An element μ is* mulitiplicatively infinitesimal over M *if the map $x \mapsto \mu \cdot x$ is M-small.*

By Corollary 4.7, multiplicative infinitesimals form a subgroup of the multiplicative group.

Lemma 7.2. *If f is an M-small bijection and g is an M-definable bijection, then $g \circ f \circ g^{-1}$ is also M-small.*

Proof. Suppose X is M-definable and infinite. We need to show that there are infinitely many points in $X \cap g(f^{-1}(g^{-1}(X))$, or equivalently, that there are infinitely many points in $g^{-1}(X) \cap f^{-1}(g^{-1}(X))$. But $g^{-1}(X)$ is M-definable, so this follows by M-smallness of f. $\square$

Theorem 7.3. *1. If ϵ is an (additive) infinitesimal and μ is a multiplicative infinitesimal, then $\epsilon \cdot \mu$ is an additive infinitesimal.*

2. *If μ is a multiplicative infinitesimal, then $\mu - 1$ is an additive infinitesimal.*

3. *If ϵ is an additive infinitesimal, then $1 + \epsilon$ is a multiplicative infinitesimal.*

4. *The product of two additive infinitesimals is an additive infinitesimal.*

5. *The multiplication map is jointly continuous.*

Proof. 1. The maps $f(x) = x + \epsilon$ and $g(x) = \mu \cdot x$ are M-small. By Corollary 4.7, we can start composing them. In particular,

$$(g \circ f \circ g^{-1})(x) = \mu \cdot (\mu^{-1} \cdot x + \epsilon) = x + \mu \cdot \epsilon$$

is M-small. This means that $\mu \cdot \epsilon$ is an infinitesimal.

2. Let $f(x) = \mu \cdot x$, and $g(x) = x + 1$. By Lemma 7.2, the map $g^{-1} \circ f \circ g$ is M-small. This is the map

$$x \mapsto (x + 1) \cdot \mu - 1 = x \cdot \mu + \mu - 1$$

By Corollary 4.7, if we precompose this with f^{-1}, the result will still be an M-small map. The result is

$$x \mapsto \mu^{-1} \cdot x \cdot \mu + \mu - 1 = x + \mu - 1.$$

It follows that $\mu - 1$ is an (additive) infinitesimal.

3. Let ϵ be an (additive) infinitesimal over M, and $\mu = 1 + \epsilon$. Let X be an infinite M-definable set. We must show that there are infinitely many $x \in X$ such that $\mu \cdot x \in X$. By Theorem 6.17, $X(M) \setminus \partial X$ is infinite, so it suffices to show that if $a \in X(M) \setminus \partial X$, then $\mu \cdot a \in X$. As $a \in X(M) \subset M$, the element $a \cdot \epsilon$ is M-infinitesimal. As $a \notin \partial X$, there is some M-definable neighborhood N of a wholly contained in X. As $a \cdot \epsilon$ is an infinitesimal, $a + a \cdot \epsilon = \mu \cdot a$ is in N, hence in X. This shows that multiplication by μ is M-small, so μ is indeed a multiplicative infinitesimal.

4. If ϵ and δ are two additive infinitesimals, then by the previous bullet points, $1 + \epsilon$ is a multiplicative infinitesimal, and $(1 + \epsilon) \cdot \delta$ is an additive infinitesimal. So $\delta + \epsilon \cdot \delta$ is an infinitesimal. As infinitesimals are a group, we can remove the δ, and $\epsilon \cdot \delta$ is an infinitesimal.

5. Let a, b be elements of M, and N be an M-definable neighborhood of $a \cdot b$. If $\Sigma(x)$ is the partial type asserting that x is an M-infinitesimal, then

$$\Sigma(x - a) \wedge \Sigma(y - b) \implies \Sigma(x \cdot y - a \cdot b) \implies x \cdot y \in N$$

Indeed, if $x - a$ is an infinitesimal ϵ, and $y - b$ is an infinitesimal δ, then

$$x \cdot y - a \cdot b = (a + \epsilon) \cdot (b + \delta) - a \cdot b = a \cdot \delta + b \cdot \epsilon + \epsilon \cdot \delta.$$

We've just seen that $\epsilon \cdot \delta$ is an infinitesimal, and we already knew that infinitesimals were closed under multiplication by constants, so that $a \cdot \delta$ and $b \cdot \epsilon$ are infinitesimals. So the right hand side is an infinitesimal, and therefore $\Sigma(x \cdot y - a \cdot b)$ holds.

Now, by compactness, it follows that there's a finite subtype $\Sigma_0(x) \subset \Sigma(x)$ such that

$$\Sigma_0(x - a) \wedge \Sigma_0(y - b) \implies x \cdot y \in N$$

But $\Sigma_0(x-a)$ cuts out a definable neighborhood of a, and $\Sigma_0(y-b)$ cuts out a definable neighborhood of b, so we get continuity.

□

Corollary 7.4. *Condition (*) holds: there is a definable family N_* of neighborhoods of 0 such that for every finite set S of nonzero elements of $\mathbb{C}$, some N_a is disjoint from S.*

Proof. Fix some small model M. Let $\Sigma(x)$ be the big partial type over $\mathbb{C}$ asserting that x is infinitesimal over $\mathbb{C}$. If (*) fails, then by Lemma 6.1, $\Sigma(x) \wedge x \neq 0$ is finitely sastisfiable in M. So it can be completed to some global coheir of a type over M. Let p be this global coheir. Let $b \models p|M$, let M_b be a model containing M and b, and let $a \models p|M_b$. (So $(a, b) \models p \otimes p|M$, among other things.) Now p is a global infinitesimal type, so $p|M_b$ is infinitesimal—meaning that a is an M_b-infinitesimal. Since $1/b \in M_b$, it follows that a/b is M_b-infinitesimal, hence M-infinitesimal.

Next, note that $\mathrm{tp}(a/Mb)$ is finitely satisfiable in M, because p is a coheir over M. So $\mathrm{tp}(b/Ma)$ is an heir of $\mathrm{tp}(b/M)$. If M_a is a model containing M and a, we can move M_a over Ma so that $\mathrm{tp}(b/M_a)$ is an heir over M. Heirs of infinitesimal types are still infinitesimal, so b is an M_a infinitesimal. As before, it follows that b/a is an M_a-infinitesimal, hence an M-infinitesimal.

Now the product of a/b and b/a must be an M-infinitesimal. But it isn't. □

8 Division is continuous

Recall that $a \approx_M 0$ means that a is an M-infinitesimal. Extend this notation in the obvious way, so that $a \approx_M b$ means that $a - b \approx_M 0$.

Proposition 8.1. *Division is continuous on $\mathbb{C} \times \mathbb{C}^\times$.*

Proof. It suffices to show that the map $x \mapsto x^{-1}$ is continuous on $\mathbb{C}^\times$.

Claim 8.2. *If M is a small model, $a \in M$, and $\epsilon \approx 0$, then $1/(a + \epsilon) \approx 1/a$.*

Proof. Note that

$$\frac{1}{a} - \frac{1}{a+\epsilon} = \frac{1}{a}\left(1 - \frac{1}{1 + \frac{\epsilon}{a}}\right)$$

Multiplication by $1/a$ preserves infinitesimals, so we reduce to showing that $1 - 1/(1 + \epsilon)$ is infinitesimal when ϵ is. This follows from the second and third points of Theorem 7.3 and the fact that the multiplicative infinitesimals are a subgroup of $\mathbb{C}^\times$. □

Now, let $a \in \mathbb{C}$ be given. Let U be a definable neighborhood of $1/a$. Let M be a small model containing a and the definition of U. By the claim and compactness, there is some neighborhood V of a such that if $a + \epsilon \in V$, then $1/(a+\epsilon) \in U$. This shows continuity. $\square$

9 The topology is definable

Definition 9.1. *For M a small model, let E_M denote the set of M-infinitesimals. Let $\mathcal{O}_M$ denote the set of $x \in \mathbb{C}$ such that $x \cdot E_M \subseteq E_M$.*

Lemma 9.2. *The type-definable group E_M has no subgroups of bounded index.*

Proof. It is known that G^{00} exists in NIP theories, so E_M^{00} exists, and is M-definable. We need to show that E_M^{00} is all of E_M. Let $\epsilon \in E_M$ be an infinitesimal—we will show $\epsilon \in E_M^{00}$. If $\epsilon = 0$, this is easy to show, so we may assume $\epsilon \notin M$. Let $N \succeq M$ contain ϵ, and let ϵ' realize an heir of $\mathrm{tp}(\epsilon/M)$ to N. Thus $\epsilon \equiv_M \epsilon'$, so ϵ and ϵ' are in the same coset of E_M^{00}, and $\epsilon - \epsilon' \in E_M^{00}$. We claim that

$$\epsilon - \epsilon' \equiv_M \epsilon \tag{1}$$

If not, there is an M-definable set X containing ϵ but not $\epsilon - \epsilon'$. As $\mathrm{tp}(\epsilon'/N)$ is an heir of $\mathrm{tp}(\epsilon/M)$, the element ϵ' is an N-infinitesimal. It follows that $\epsilon \in \partial X$. But ∂X is finite and $\ulcorner X \urcorner$-definable, hence in M. So $\epsilon \in M$, a contradiction.

Therefore (1) holds, which ensures that ϵ and $\epsilon - \epsilon'$ are in the same coset of E_M^{00}, which must then by E_M^{00}. $\square$

Lemma 9.3. *Let G and H be type-definable subgroups of the additive group of $\mathbb{C}$. Then $G \cap H$ has bounded index in at least one of G and H.*

Proof. Suppose not, so $G/(G \cap H)$ and $H/(G \cap H)$ are both unbounded. Let M be a model over which H and G are defined. By Erdos-Rado, we can find an M-indiscernible sequence $a_1, a_2, \ldots$ in G such that $a_i - a_j \notin G \cap H$ for $i \neq j$, or equivalently, $a_i - a_j \notin H$. Similarly, we can find an M-indiscernible sequence $b_1, b_2, \ldots$ in H such that $b_i - b_j \notin G$ for $i \neq j$. The cosets $a_1 + H$ and $a_2 + H$ are disjoint, so by compactness, we can find an M-definable set H' containing H such that $a_1 + H' \cap a_2 + H' = \emptyset$. By indiscernibility, the sets $a_i + H'$ are pairwise disjoint. Similarly, we can find $G' \supset G$ such that the $b_i + G'$ are pairwise disjoint.

For each i, j, the intersection $a_i + H' \cap b_j + G'$ contains $a_i + b_j$, and hence is non-empty. Consequently, the $\{a_i + H'\}$ and $\{b_i + G'\}$ form an ict pattern of depth 2, contradicting dp-minimality. $\square$

Corollary 9.4. *For $a, b \in \mathbb{C}$, the sets $a \cdot E_M$ and $b \cdot E_M$ are comparable: one is a subset of the other.*

Proof. If a or b is zero, this is trivial. Otherwise, $a \cdot E_M$ and $b \cdot E_M$ are isomorphic to E_M, hence have no bounded-index quotients by Lemma 9.2. The result then follows by Lemma 9.3. $\square$

Recall that $\mathcal{O}_M$ denoted the set of $\alpha \in \mathbb{C}$ such that $\alpha \cdot E_M \subseteq E_M$.

Corollary 9.5. *$\mathcal{O}_M$ is a valuation ring in $\mathbb{C}$, and E_M is a proper ideal in $\mathcal{O}_M$.*

Proof. The set $\mathcal{O}_M$ is closed under addition and subtraction, because E_M is. It is clearly closed under multiplication and contains 1, so it is a ring.

To see that $\mathcal{O}_M$ is a valuation ring, we need to show that for any $a \in \mathbb{C}^\times$, either a or $1/a$ is in $\mathcal{O}_M$. That is, we need to show that $a \cdot E_M \subseteq E_M$ or $a^{-1} \cdot E_M \subseteq E_M$. Equivalently, we need to show that $a \cdot E_M$ and E_M are comparable, which follows by the previous lemma.

We have $E_M \subseteq \mathcal{O}_M$ because $E_M \cdot E_M \subseteq E_M$, by the penultimate part of Theorem 7.3. Given this, E_M is an ideal essentially by definition of $\mathcal{O}_M$. It is a proper ideal because $1 \notin E_M$. $\square$

Let Γ denote the corresponding value group and let $v : \mathbb{C}^\times \to \Gamma$ denote the corresponding valuation. Because E_M is contained in the maximal ideal of $\mathcal{O}_M$, elements of E_M have positive valuation. Also, if $x \in E_M$ and $v(y) \geq v(x)$, then $y \in E_M$. (Indeed, $y/x \in \mathcal{O}_M$ and $x \in E_M$.)

Theorem 9.6. *Products of non-infinitesimals are non-infinitesimal. E_M is the maximal ideal of $\mathcal{O}_M$.*

Proof. Suppose $a \cdot b \in E_M$ but $a, b \notin E_M$. Switching a and b, we may assume $v(a) \geq v(b)$. Then $v(a^2) \geq v(a \cdot b)$ so $a^2 \in E_M$.

As $a \notin E_M$, there is some M-definable neighborhood U of 0, with $a \notin U$. Let $\Sigma(x)$ be the partial type over M asserting that x is infinitesimal. Let ϵ be some non-zero M-infinitesimal. Note that

$$\Sigma(x^2) \implies \Sigma(\epsilon \cdot x)$$

To see this, suppose that x^2 is an infinitesimal. If $v(\epsilon) \geq v(x)$, then $v(\epsilon \cdot x) \geq v(x^2)$ so $\epsilon \cdot x$ is an infinitesimal. Otherwise, $v(\epsilon) < v(x)$, so x is an infinitesimal, and therefore so is $\epsilon \cdot x$.

Now, by compactness, it follows that there is some M-definable neighborhood V such that $x^2 \in V \implies \epsilon \cdot x \in U$. Because M is a model, there is some $e \in M$ such that

$$x^2 \in V \implies e \cdot x \in U. \tag{2}$$

Now, a^2 is an M-infinitesimal and $e^2 \in M$, so a^2/e^2 is also an M-infinitesimal, hence is in V. By (2), $e \cdot a/e = a$ is in U, a contradiction.

This shows that products of non-M-infinitesimals are non-M-infinitesimals. Now E_M is contained in the maximal ideal $\mathfrak{m}_M$ of $\mathcal{O}_M$. If the inclusion is strict, take $a \in \mathfrak{m}_M \setminus E_M$. Then $1/a \notin \mathcal{O}_M$, so there is some $c \in E_M$ such that $c/a \notin E_M$. Then c/a and a are non-infinitesimals whose product c is infinitesimal, a contradiction. So the maximal ideal of $\mathcal{O}_M$ is exactly E_M. $\square$

The fact that non-infinitesimals are closed under multiplication means concretely that if U and V are neighborhoods of 0, then $(\mathbb{C} \setminus U) \cdot (\mathbb{C} \setminus V)$ does not have 0 in its closure.

Corollary 9.7. *The topology is definable, i.e., has a definable basis. More precisely, for each model M there is a definable open set U containing 0 such that sets of the form $a \cdot U$ for $a \in M^\times$ are a neighborhood basis of 0, and therefore the family of sets of the form $a \cdot U + b$ is a basis for the topology on M.*

Proof. Let $\Sigma(x)$ be the partial type asserting that x is infinitesimal. Then $\Sigma(x) \cup \Sigma(1/x)$ is inconsistent, because 1 is not a product of two infinitesimals. By compactness, we can find a neighborhood U_0 of 0 such that $x \in U_0 \wedge x^{-1} \in U_0 \implies \perp$. In particular, if $x \in U_0$, then $1/x$ is not an infinitesimal, so $U_0(\mathbb{C}) \subseteq \mathcal{O}_M$. Let U be the interior of U_0—this is definable by Theorem 6.17. Of course U still has the property that it is contained in $\mathcal{O}_M$ and that 0 is in its interior.

We claim that $\{a \cdot U : a \in M^\times\}$ is a neighborhood basis of 0. Let $V = X -_\infty X$ be any M-definable neighborhood of 0. Take $\epsilon \approx_M 0$. Then

$$\begin{aligned} x \in U &\implies v(x) \geq 0 \implies v(\epsilon \cdot x) > 0 \\ &\implies \epsilon \cdot x \approx_M 0 \implies \epsilon \cdot x \in V. \end{aligned}$$

As M is a model, there is $a \in M$ approximating $\mathrm{tp}(\epsilon/M)$, such that $x \in U \implies a \cdot x \in V$, which means exactly that $a \cdot U \subseteq V$. As $a \cdot U$ is still open, and V was arbitrary, this is good enough to complete the proof. $\square$

Definition 9.8. *A definable set X is* bounded *if its image under the map $x \mapsto x^{-1}$ does not have 0 in its closure.*

From the above discussion, it follows that an M-definable set is bounded if and only if it is contained in $\mathcal{O}_M$. Consequently,

Remark 9.9. *If X and Y are bounded, so are $X \cdot Y$, $X + Y$, and $X - Y$.*

10 Another look at germs at 0

First we improve Theorem 6.10 using the definability of the topology.

Theorem 10.1. *Let X_a be a definable family of subsets of the home sort. Then there are only finitely many germs at infinity among the X_a's.*

Proof. Let $a \approx a'$ indicate that X_a and $X_{a'}$ have the same germ at 0. Because the topology is definable, this is a definable relation. By Theorem 6.10, the number of equivalence relations is bounded, hence finite. $\square$

Corollary 10.2. *Let Δ be a finite set of formulas. Under the restriction map from types over $\mathbb{C}$ to Δ-types over $\mathbb{C}$, the set of infinitesimal types over $\mathbb{C}$ has finite image.*

That is, there are only finitely many infinitesimal Δ-types. Equivalently, among the infinitesimal types (which are all definable), there are only finitely many Δ-definitions.

Proof. If $X \sim Y$ indicates that X and Y have the same germ at 0, then the boolean algebra of Δ-sets modulo $\sim$ is finite. The ultrafilters on this boolean algebra are the same thing as infinitesimal Δ-types. $\square$

Corollary 10.3. *Let X be a definable set with* 0 *in its closure. Then there is a definable finite-index subgroup $G \leq \mathbb{C}^\times$ such that X has the same germ at 0 as $g \cdot X$ for all $g \in G$.*

Proof. Look at the family of sets $g \cdot X$ as g ranges over $\mathbb{C}^\times$. If S is the finite set of germs at 0 of this family, then $\mathbb{C}^\times$ acts definably on S, because $Y \sim Z \implies g \cdot Y \sim g \cdot Z$. As S is finite, there is a finite index subgroup $G \leq \mathbb{C}^\times$ such that G acts trivially on S. I particular, G fixes the class of X. $\square$

Let $\mathbb{G}_m$ denote the multiplicative group. Because we are in an NIP setting, $\mathbb{G}_m^0$ and $\mathbb{G}_m^{00}$ exist.

Lemma 10.4. *Every infinitesimal type is multiplicatively stabilized by $\mathbb{G}_m^0$. That is, if $p(x)$ is a global infinitesimal type and $g \in \mathbb{G}_m^0$, then $p(x)$ is equivalent to $p(g \cdot x)$.*

Proof. If not, there is some $\mathbb{C}$-formula X which is in $p(g \cdot x)$ but not in $p(x)$. So $p(x) \vdash x \notin X \wedge g \cdot x \in X$. Therefore X and $g^{-1} \cdot X$ do not have the same germ at 0, because realizations of the infinitesimal type $p(x)$ are in one but not the other. This contradicts Corollary 10.3. $\square$

Lemma 10.5. *For each n, $\mathbb{C}^\times/(\mathbb{C}^\times)^n$ is finite.*

Proof. It suffices to show that this interpretable set is bounded. Since there are a bounded number of global infinitesimal types, it suffices to show that each coset C of $(\mathbb{C}^\times)^n$ contains an infinitesimal (i.e., the formula $x \in C$ is in some global infinitesimal type). Let M be a model over which C is defined, and let ϵ be an M-infinitesimal. Take $c \in C(M)$. Then $\epsilon^n \cdot c \approx_M 0$ and $\epsilon^n \cdot c \in C$. The unique heir of $\mathrm{tp}(\epsilon^n \cdot c/M)$ to $\mathbb{C}$ will be a global infinitesimal type in C. $\square$

Theorem 10.6. $\mathbb{G}_m^0 = \mathbb{G}_m^{00} = \bigcap_{n \in \mathbb{Z}_+} (\mathbb{C}^\times)^n$

Proof. For the first equality: suppose $g \in \mathbb{G}_m^0 \setminus \mathbb{G}_m^{00}$. Let $p(x)$ be some global infinitsimal type. Then $p(x)$ is equivalent to $p(g \cdot x)$ by Lemma 10.4, but these types are not in the same coset of $\mathbb{G}_m^{00}$, a contradiction.

By Lemma 10.5, the intersection $\bigcap_{n \in \mathbb{Z}_+} (\mathbb{C}^\times)^n$ has finite index in $\mathbb{C}^\times$, so it contains $\mathbb{G}_m^0$. But if G is any finite index subgroup of $\mathbb{G}_m$, then G contains every nth power, for n the index. So $G \supseteq \bigcap_{n \in \mathbb{Z}_+} (\mathbb{C}^\times)^n$. This shows that $\bigcap_{n \in \mathbb{Z}_+} (\mathbb{C}^\times)^n = \mathbb{G}_m^0$. $\square$

Proposition 10.7. *If $M \preceq \mathbb{C}$, $\epsilon \approx_M 0$, $\epsilon' \approx_M 0$, and $\epsilon/\epsilon' \approx_M 1$, then ϵ and ϵ' have the same type over M.*

Proof. Let $N \succeq M$ be a model containing an M-infinitesimal δ. Move N and δ so that $\mathrm{tp}(N/M\epsilon\epsilon')$ is finitely satisfiable in M. Then ϵ, ϵ', and $\epsilon/\epsilon' - 1$ are N-infinitesimals.

Let U be an M-definable open neighborhood of 0 as in Corollary 9.7, so that sets of the form $x \cdot U$ are a neighborhood basis of 0. Let $V = \delta \cdot U$.

Suppose for the sake of contradiction that ϵ and ϵ' have different types over M. Let X be an M-definable set containing ϵ' but not ϵ. Let G be the finite index subgroup of $\mathbb{G}_m$ ensured by Corollary 10.3. So, $g \cdot X$ has the same germ at 0 as X, for $g \in G$. We can take G to be defined over M.

Consider the N-definable sets

$$\begin{aligned} P^+ &= \{g \in \mathbb{G}_m \ : \ g \cdot S \subseteq S \text{ for } S \in \{V, V \cap X, V \setminus X\}\} \\ P &= \{g \cdot h^{-1} \ : \ g, h \in P^+\} \end{aligned}$$

Note that P^+ is a submonoid of $\mathbb{G}_m$ and P is a subgroup.

Claim 10.8. *P is an open set.*

Proof. Being a group, it suffices to show that P has non-empty interior. Because P has finite boundary, it suffices to show that P is infinite. Since G is finite index, G is infinite, so it suffices to show that $G(M) \subset P(\mathbb{C})$.

Suppose $g \in G(M)$; we will show $g \in P(\mathbb{C})$. Let n be the index of G in $\mathbb{G}_m$. Take some M-infinitesimal τ. Then τ^n and $g \cdot \tau^n$ are M-infinitesimals, so $\tau^n \cdot U \subset U$ and $g \cdot \tau^n \cdot U \subset U$. The type of τ over M is finitely satisfiable in M, so we can find some $t \in M$ such that

$$\begin{aligned} t^n \cdot U &\subset U \\ g \cdot t^n \cdot U &\subset U \end{aligned}$$

Multiplying by δ, we get the same statements for $V = \delta \cdot U$:

$$\begin{aligned} t^n \cdot V &\subset V \\ g \cdot t^n \cdot V &\subset V \end{aligned}$$

It remains to show that $g \cdot t^n$ and t^n are in P^+. Let a be $g \cdot t^n$ or t^n. Then $a \cdot V \subset V$ and $a \in G(M)$. Because $a \in G$, the M-definable sets X and $g \cdot X$ have the same germ at 0. So, if x is an M-infinitesimal, then $x \in X$ if and only if $a \cdot x \in X$. All elements of V are M-infinitesimals, so

$$x \in V \implies (x \in X \iff a \cdot x \in X)$$

This is exactly what it means for a to be in P^+.

So both $g \cdot t^n$ and t^n are in P^+, making g be an element of P. As g was an arbitrary element of $G(M)$, if follows that $G(M) \subset P$, so P is infinite, hence open. $\square$

Now P is an N-definable open subgroup of $\mathbb{G}_m$, so it contains an N-definable neighborhood of 1. As $\epsilon/\epsilon' \approx_N 1$, it follows that $\epsilon/\epsilon' \in P$. So we can write $\epsilon/\epsilon' = a/b$ for some $a, b \in P^+$.

As ϵ and ϵ' are N-infinitesimals, they are contained in the N-definable neighborhood V. So, by definition of P^+, we have the implications

$$\begin{aligned} \epsilon \in X &\implies b \cdot \epsilon \in X \\ \epsilon \notin X &\implies b \cdot \epsilon \notin X \\ \epsilon' \in X &\implies a \cdot \epsilon' \in X \\ \epsilon' \notin X &\implies a \cdot \epsilon' \notin X \end{aligned}$$

As $b \cdot \epsilon = a \cdot \epsilon'$, we get $\epsilon \in X \iff \epsilon' \in X$, contradicting the choice of X. □

This has the following interesting consequence, which might be helpful in working towards henselianity:

Theorem 10.9. *Let U be a definable open set, and f be a definable function on U which is differentiable with non-vanishing derivative. Then f is an open map.*

Proof. For $a \in U$ and V a neighborhood of a, we must show that $f(V)$ is a neighborhood of $f(a)$. Translating and rescaling, we may assume that $a = 0$, $f(0) = 0$, and $f'(0) = 1$. Let M be a model over which everything is defined. We must show that every M-infinitesimal is in $f(V)$. Let ϵ be an M-infinitesimal. The fact that $f'(0) = 1$ implies that

$$\frac{f(\epsilon) - f(0)}{\epsilon - 0} = \frac{f(\epsilon)}{\epsilon} \approx_M 1$$

In particular, $f(\epsilon)/\epsilon$ is in $\mathcal{O}_M$, so $f(\epsilon) \approx_M 0$, and then by Proposition 10.7, $f(\epsilon)$ has the same type as ϵ over M. As $f(\epsilon) \in f(V)$, we also have $\epsilon \in f(V)$. □

11 Finding higher-dimensional opens

Lemma 11.1. *Suppose S is a small set of parameters, a is an element non-algebraic over S, and ϵ is an infinitesimal over (some model containing) Sa. Then a is non-algebraic over $S\epsilon$.*

Proof. This follows by definability of infinitesimal types. Basically, if N is a model containing Sa and ϵ is N-infinitesimal, then $\mathrm{tp}(\epsilon/N)$ is $\mathrm{acl}(0)$-definable. If $a \in \mathrm{acl}(S\epsilon)$, there is an $S\epsilon$-definable finite set X containing a. By definability of $\mathrm{tp}(\epsilon/N)$, there is an $S\,\mathrm{acl}(\emptyset)$-definable set Y such that $X(N) = Y(N)$. Since X is finite, so is $Y(N)$. As Y is N-definable, Y itself must be finite. Then $a \in X(N) = Y(N)$ so $a \in \mathrm{acl}(S\,\mathrm{acl}(\emptyset))$. □

Theorem 11.2. *Let $X \subset \mathbb{C}^n$ be M-definable. Let $a = (a_1, \ldots, a_n)$ be an element of X with the property that for each i,*

$$a_i \notin \mathrm{acl}_M(\{a_1, \ldots, a_n\} \setminus a_i).$$

Then a is in the interior of X (using the product topology on $\mathbb{C}^n$).

Proof. Let N be a model containing the a_i's and $\ulcorner X \urcorner$. Let B be an M-definable open neighborhood of 0 whose rescalings are a neighborhood of 0. Let $e_1, e_2, \ldots, e_n$ be the basis vectors in $\mathbb{C}^n$.

Claim 11.3. *For $0 \leq k \leq n$, there is some $\epsilon \approx_N 0$ such that*

$$a + \sum_{j=1}^{k} \epsilon \cdot B \cdot e_j \subseteq X$$

Proof. We proceed by induction on k. For $k = 0$, this just says $a \in X$, which is given.

Now suppose $k > 0$, and, by induction, that $a + \sum_{j=1}^{k-1} \delta \cdot B \cdot e_j \subseteq X$ for some N-infinitesimal δ. For $z \in \mathbb{C}$, let

$$f(z) = (a_1, a_2, \ldots, a_{k-1}, z, a_k, \ldots, a_n)$$

Let Y be the set of z such that

$$f(z) + \sum_{j=1}^{k-1} \delta \cdot B \cdot e_j \subseteq X.$$

So $a_k \in Y$, and Y is defined over $a_{\neq k}\delta M$. By the previous Lemma, a_k is not algebraic over these parameters, so by Theorem 6.17, $a_k \notin \partial Y$. Hence a_k is in the interior of Y. Let ϵ be infinitesimal over everything so far. Then $a_k + B \cdot \epsilon \subset Y$, which exactly means that

$$a + \sum_{j=1}^{k-1} \delta \cdot B \cdot e_j + \epsilon \cdot B \cdot e_k \subseteq X.$$

Now ϵ is infinitesimal over δ, so $\epsilon \cdot B \subseteq \delta \cdot B$. So in fact,

$$a + \sum_{j=1}^{k-1} \epsilon \cdot B \cdot e_j + \epsilon \cdot B \cdot e_k \subseteq X.$$

completing the inductive proof of the claim. □

Now, taking $k = n$ in the claim, we see that a is in the interior of X. □

Corollary 11.4. *A definable set $X \subseteq \mathbb{C}^n$ has dp-rank n exactly if it has non-empty interior (in the product topology).*

Proof. If it has non-empty interior, it clearly has dp-rank n. Conversely, suppose X has dp-rank n. Let M be a model over which X is defined, and let $a = (a_1, \ldots, a_n)$ be an element of X with dp-rank n over M. From additivity of the dp-rank, no a_i can be in the M-algebraic closure of the others, so by the theorem, a is in the interior of X. □

We can also prove that dp-rank is definable in powers of the home sort. (This fails in imaginary sorts: consider the value group in p-adically closed fields.)

For $1 \leq i \leq n$, let $\pi_i : \mathbb{C}^n \to \mathbb{C}^{n-1}$ be the coordinate projection onto the coordinates other than the ith one. Say $Y \subset \mathbb{C}^n$ is π_i-finite if the map $Y \to \pi_i(Y)$ has finite fibers.

Lemma 11.5. *If $X \subset \mathbb{C}^n$ is definable, then either X has interior, or X can be written as a finite union $\bigcup_{j=1}^m Y_j$ of definable sets, each of which is π_i-finite for some i, possibly depending on j.*

Proof. Compactness and Theorem 11.2. □

Theorem 11.6. *Dp-rank is definable in families of subsets of $\mathbb{C}^n$.*

Proof. We proceed by induction on n. The base case $n = 1$ is Corollary 2.2. Suppose $n > 1$. The family of all definable subsets of $\mathbb{C}^n$ is ind-definable (a small union of definable families). It suffices to show that for $0 \leq k \leq n$, the subfamily of rank-k subsets is also ind-definable. For $k = n$, this follows by Corollary 11.4.

By the Lemma, if $X \subset \mathbb{C}^n$ does not have full rank, then X can be written as a finite union of π_i-finite sets for various i. Now, if Y is π_i-finite for some i, then $\text{dp-rk}(Y) = \text{dp-rk}(\pi_i(Y))$ because interalgebraic tuples have the same dp-rank. By induction, the family of π_i-finite Y with dp-rank k is ind-definable, for fixed i and k. If $\mathcal{F}_k$ denotes the family of definable subsets of $\mathbb{C}^n$ which have dp rank k and are π_i-finite for *some* i, then $\mathcal{F}_k$ is also ind-definable. A definable subset of $\mathbb{C}^n$ has dp-rank k if and only if it is a union of a set in $\mathcal{F}_k$ and finitely many sets in $\bigcup_{k' \leq k} \mathcal{F}_{k'}$. This is also an ind-definable family, completing the proof. □

Note that this also holds in a strongly minimal field. So we have proven that dp-rank is definable in dp-minimal fields.

By being more careful, we get a very crude sort of cell-decomposition.

Definition 11.7. *A definable set $X \subset \mathbb{C}^n$ is a* cell *if it has interior, or $n > 0$ and X is π_i-finite for some i and $\pi_i(X)$ is a cell.*

Proposition 11.8. *Every definable subset of $\mathbb{C}^n$ can be written as a disjoint union of finitely many cells.*

Proof. We proceed by induction on n. The base case $n = 0$ is clear. Suppose $X \subset \mathbb{C}^n$ is definable. The interior X^{int} of X is technically a cell, so it suffices to decompose the set $Z := X \setminus X^{int}$ into cells. The set Z has no interior, so by Lemma 11.5 it can be written as a finite union $\bigcup_{j=1}^m Y_j$ of definable sets, each of which is π_i-finite for some i. Shrinking the Y_j, we may arrange that the Y_j are disjoint. It remains to decompose an individual $Y = Y_j$ which is π_i-finite. By induction, we can decompose $\pi_i(Y)$ into cells. The preimage of each cell under $Y \to \pi_i(Y)$ will be a cell, and we are done. □

The following Lemma generalizes Theorem 6.17

Lemma 11.9. *If $X \subset \mathbb{C}^n$ is definable then the boundary ∂X of X has dp-rank less than n.*

Proof. The frontier $\overline{X} \setminus X$ has no interior, so it has dp-rank less than n. Similarly $X \setminus X^{int}$ has dp-rank less than n. The union of these two sets is ∂X. □

12 Something like Henselianity

Suppose $L/\mathbb{C}$ is a finite algebraic extension. Then there is a canonical topology on L coming from the identification of L with $\mathbb{C}^n$ (after choosing a basis) and then taking the product topology on $\mathbb{C}^n$. The choice of the basis doesn't matter, and the resulting topology on L makes addition, multiplication, and subtraction continuous. Our goal in this section is to show that this topology is a "V-topology," meaning that infintiesimals in L are the maximal ideal of a valuation ring.

Remark 12.1. *Fields of finite dp-rank are perfect. Indeed, if K has dp-rank n, then K^p has dp-rank n as well. If K/K^p had degree greater than 1, K would have dp-rank at least $2n > n$, a contradiction.*

The next result would tell us that $\mathcal{O}_M$ is henselian, if we knew that the Galois group of $\mathbb{C}$ was bounded (finitely many Galois extensions of each degree).

Theorem 12.2. *Let $L/\mathbb{C}$ be a finite extension, defined over $M \preceq \mathbb{C}$. (So $L = \mathbb{C}(\alpha)$ where α satisfies a monic irreducible polynomial $P(X)$ over M, of degree n.) Then $\mathcal{O}_M$ has a unique extension $\mathcal{O}$ to L. Moreover, the maximal ideal of $\mathcal{O}$ is exactly*

$$\left\{\sum_{i=0}^{n-1} a_i \cdot \alpha^i \ : \ a_i \in \mathfrak{m}_M\right\}$$

Furthermore, division is continuous in the product topology on L, and the product topology is a V-topology.

We will give the proof in characteristic $\neq 2$. The proof for characteristic 2 is analogous, except that the multiplicative group is replaced with the additive group, the squaring map is replaced with the Artin-Schreier map, the constants 1 and -1 are replaced with 0 and 1, respectively, and $1 + \mathfrak{m}_K$ is replaced with $\mathfrak{m}_K$.

Proof. Let K be some algebraic closure of $\mathbb{C}$ and $\mathcal{O}_K$ be some extension of $\mathcal{O}_M$ to K. Let $\alpha_1, \ldots, \alpha_n$ enumerate the roots of $P(X)$ in K. Consider the ring $K[X]/P(X)$. Identify L with the subring $\mathbb{C}[X]/P(X)$, and K and $\mathbb{C}$ and M with the subrings of $K[X]/P(X)$ in the obvious way.

For $x, y \in K[X]/P(X)$, let $x \approx y$ indicate that

$$x - y = \sum_{i=0}^{n-1} a_i \cdot X^i$$

with a_i in the maximal ideal $\mathfrak{m}_K$ of $\mathcal{O}_K$.

So $\approx$ extends $\approx_M$ on $\mathbb{C}$, and also, $\{x \in L : x \approx 0\}$ is the set of M-infinitesimals in L with respect to the product topology on L–i.e., the set of elements which are in every M-definable product-topology-neighborhood of 0.

Let $g : K[X]/P(X) \to K^n$ be the map of K-algebras sending X to $(\alpha_1, \ldots, \alpha_n)$. Then g is an isomorphism, because as a linear transformation, its matrix is a Vandermonde matrix coming from the α_i's, which are distinct.

Moreover, the α_i's are algebraic over M, so this Vandermonde matrix and its inverse have entries in $M^{alg} \subset K$. Because $M^\times \subseteq \mathcal{O}_M^\times$, the $\mathcal{O}_M$-valuation on M is trivial, which implies that the $\mathcal{O}_K$-valuation on M^{alg} is *also* trivial. Therefore the Vandermonde matrix and its inverse have entriesin $\mathcal{O}_K$, and in particular, both are elements of $GL_n(\mathcal{O}_K)$.

This has the following implication: for $x \in K[X]/P(X)$,

$$x \approx 0 \iff g(x) = (b_1, \ldots, b_n) \text{ where each } b_i \text{ is in } \mathfrak{m}_K$$

If $\pi_i : K^n \to K$ is the ith projection, then $\pi_i \circ g : K[X]/P(X) \to K$ is the map sending X to $\mathbb{C}$. Consequently, the compositions

$$L = \mathbb{C}[X]/P(X) \hookrightarrow K[X]/P(X) \xrightarrow{g} K^n \xrightarrow{\pi_i} K$$

are exactly the embeddings of L into K.

Because valued fields can be amalgamated, the extensions of $\mathcal{O}_M$ to L are exactly the pullbacks of $\mathcal{O}_K$ along these embeddings $L \hookrightarrow K$. This has the following consequence: for $x \in L$, the following four conditions are equivalent:

- x is in the maximal ideal of every extension of $\mathcal{O}_M$ to L
- For each embedding of L into K, x maps into $\mathfrak{m}_K$
- $g(x) = (b_1, \ldots, b_n)$ where each b_i is in $\mathfrak{m}_K$.
- $x \approx 0$

With all these preliminaries out of the way, we can begin proving things.

First of all, we check that division on L is continuous with respect to the product topology. Since the ring operations are continuous, one reduces to showing that the map $x \mapsto 1/x$ is continuous at $x = 1$. As in Proposition 8.1, this amounts to showing that for $x \in L$, $x \approx 1$ implies $x^{-1} \approx 1$. Now if $g(x) = (b_1, \ldots, b_n)$, then $g(x^{-1}) = (b_1^{-1}, \ldots, b_n^{-1})$, and we know

$$\begin{aligned} x \approx 1 &\iff \text{each } b_i \in 1 + \mathfrak{m}_K \\ x^{-1} \approx 1 &\iff \text{each } b_i^{-1} \in 1 + \mathfrak{m}_K \end{aligned}$$

But of course $1 + \mathfrak{m}_K$ is a subgroup of $K^\times$, so $x \approx 1 \iff x^{-1} \approx 1$. This shows that division is continuous.

Let A be the type-definable set of $x \in L$ such that $x \approx 1$. (We have just seen that A is a subgroup of $L^\times$.) Let $f : L \to L$ be the map $x \mapsto x^2$. The type-definable set A contains a $\mathbb{C}$-definable open neighborhood V of 1. Indeed, if ϵ is an M-infinitesimal, and U is a set as in Corollary 9.7, then $\epsilon \cdot U$ is an open neighborhood of 0 contained in the M-infinitesimals, so we can take

$$V = \left\{ 1 + \sum_{i=0}^{n-1} a_i \cdot X^i \ : \ a_i \in \epsilon \cdot U \right\}$$

The set V has dp-rank n because it has interior. The map f is finite-to-one, so $f(V)$ also has dp-rank n. By Corollary 11.4, it has interior. So the non-definable set $f(A)$, which contains $f(V)$, has interior. As $f(A)$ is a subgroup of the multiplicative group, $f(A)$ is open. Hence it contains some neighborhood of 1.

Consequently, the following is true: if N is any M-definable neighborhood of 1, then $f(N) \supseteq f(X)$ has 1 in its interior. So for every M-definable neighborhood N of 1, there is another definable neighborhood N' such that $f(N) \supseteq N'$. As M is a model, N' can be taken M-definable.

We prove that there are finitely many extensions of $\mathcal{O}_M$ to L. For this, it is safe to replace L with the Galois closure of L over M, which ensures that the various extensions of $\mathcal{O}_M$ to L are conjugate under $\mathrm{Gal}(L/M)$, and hence, they are pairwise incomparable. So by Lemma 12.3 below, we can find some $e \in L$ which is congruent to 1 with respect to some of these valuations, and 0 with respect to the others. Let $a = 2 \cdot e - 1$. Then a is congruent to 1 with respect to some of the valuations, and -1 with respect to the others.

So $a - 1$ is not in the maximal ideal of every extension of $\mathcal{O}_M$ to L, and neither is $a + 1$, but $a^2 - 1$ is. By the four equivalent conditions above, it follows that $a^2 \approx 1$ and $a \not\approx 1$ and $-a \not\approx 1$. So there is some M-definable neighborhood N of 1 in the product topology, which does not contain a or $-a$. By what we showed above, there is an M-definable neighborhood N' of 1, such that $f(N) \supseteq N'$. As $a^2 \approx 1$, $a^2 \in N'$. So there is some $b \in N$ such that $f(b) = a^2$, i.e., $b^2 = a^2$. But then $b = \pm a$, and both a and $-a$ are not in N by choice of N. So we have a contradiction.

So this establishes that $\mathcal{O}_M$ has a unique extension to any finite extension of L defined over M. Now we can drop the temporary Galois assumption.

Now let $\mathcal{O}$ be this unique extension, and $\mathfrak{m}$ be its maximal ideal. By the four equivalent conditions, we now know that $x \in \mathfrak{m}$ if and only if $x \approx 1$. This is the second claim of the theorem. We have already proven the third claim (division is continuous). The remaining claim, that the product topology is a V-topology, just means that the set of $x \in L$ such that $x \approx 1$ is the maximal ideal of a valuation ring, which we have just seen. $\square$

Lemma 12.3. *Let K be a field. Let $\mathcal{O}_1, \ldots, \mathcal{O}_n$ be pairwise incomparable valuation rings in K (so $\mathcal{O}_i \not\subseteq \mathcal{O}_j$ for $i \neq j$). Let $\mathfrak{m}_i$ be the corresponding maximal ideals. Then there exist elements $e_1, \ldots, e_n \in \bigcap_{i=1}^n \mathcal{O}_i$ such that $e_i \equiv 1 \mod \mathfrak{m}_i$, and $e_i \equiv 0 \mod \mathfrak{m}_j$ for $j \neq i$.*

Proof. Let $\nu_i : K \to \mathcal{O}_i/\mathfrak{m}_i \cup \{\infty\}$ denote the place corresponding to $\mathcal{O}_i$.

We proceed by induction on n, with $n = 2$ as the base case. For $n = 2$, incomparability yields elements $a \in \mathcal{O}_1 \setminus \mathcal{O}_2$ and $b \in \mathcal{O}_2 \setminus \mathcal{O}_1$. Then

$$v_1(a) \geq 0 > v_1(b)$$

$$v_2(b) \geq 0 > v_2(2)$$

So that $v_1(\tau) > 0$ and $v_2(\tau) < 0$ for $\tau = a/b$. Equivalently,

$$\nu_1(\tau) = 0 \text{ and } \nu_2(\tau) = \infty$$

Then

$$\nu_1\left(\frac{\tau}{\tau+1}\right) = \frac{0}{0+1} = 0$$
$$\nu_2\left(\frac{\tau}{\tau+1}\right) = \frac{\infty}{\infty+1} = 1$$

so we can take $e_2 = \tau/(\tau+1)$ and $e_1 = 1 - e_2$.

For the inductive step, suppose $n > 2$. By symmetry it suffices to produce e_1. By induction, we can find a and b such that

$$(\nu_1(a), \nu_2(a), \ldots, \nu_n(a)) = (1, 0, 0, \ldots, 0, 0, ?)$$

$$(\nu_1(b), \nu_2(b), \ldots, \nu_n(b)) = (1, 0, 0, \ldots, 0, ?, 0)$$

If $\nu_n(a) = \infty$, we can take $e_1 = 1/(a - 1 + a^{-1})$, because the rational map $1/(X - 1 + X^{-1})$ sends 0 and ∞ to 0, and 1 to 1. If $\nu_{n-1}(b) = \infty$, we can take $1/(b - 1 + b^{-1})$. Finally, if neither of these cases holds, then $e_1 = a \cdot b$ works. □

13 Henselianity

In this section, we show that $\mathcal{O}_M$ is henselian, among other things.

13.1 Generalities

First we recall some basic facts about valued fields.

Lemma 13.1. *Let L/K be a field extension and $\mathcal{O}$ be a valuation ring on K with maximal ideal $\mathfrak{m}$. Let $a_1, \ldots, a_n$ be elements of L. Then there is a valuation ring of L extending $\mathcal{O}$ and containing $a_1, \ldots, a_n$ if and only if there is no polynomial $P(X)$ with coefficients in $\mathfrak{m}$ such that $1 = P(a_1, \ldots, a_n)$.*

Proof. First suppose there is such a valuation ring. Then $P(a_1, \ldots, a_n)$ is a sum of terms with positive valuation, so it cannot equal 1. Conversely, suppose that 1 cannot be written as such a polynomial. This means that the extension of $\mathfrak{m}$ to the ring $\mathcal{O}[a_1, \ldots, a_n]$ is a proper ideal. Let $\mathfrak{n}$ be a maximal ideal containing this extension. By the usual machinery of extending places, there is some valuation ring on L dominating the localization of $\mathcal{O}[a_1, \ldots, a_n]$ at $\mathfrak{n}$. This will be a valuation ring extending $\mathcal{O}$, and containing $a_1, \ldots, a_n$. □

Lemma 13.2. *For n, m, there is some $d = d(n, m)$ such that the following holds. Let L/K be a degree m Galois extension of fields, and $\mathcal{O}$ be a valuation ring on K. Let $a_1, \ldots, a_n$ be elements of L. Then there is an extension of $\mathcal{O}$ to L containing $a_1, \ldots, a_m$ exactly if there is no polynomial $P(X)$ of degree less than $d(n, m)$, with coefficients in $\mathfrak{m}$, such that $1 = P(a_1, \ldots, a_n)$.*

Proof. Consider the theory asserting that L/K is a Galois extension of degree m, R is a valuation ring on L, no Galois conjugate of R contains $a_1, \ldots, a_n$, and 1 cannot be written as $P(a_1, \ldots, a_n)$ for any polynomial with coefficients from the maximal ideal of $R \cap K$. By the previous claim, this theory is inconsistent. Compactness yields a bound on the degree of the polynomial. $\square$

Lemma 13.3. *Let L/K be a finite extension of reasonably saturated perfect fields, and let $\mathcal{O}$ be a valuation ring on K. If the maximal ideal of $\mathcal{O}$ is type-definable, this is also true for any extension of $\mathcal{O}$ to L. Similarly for "definable" instead of "type-definable."*

Proof. Replacing L with its Galois closure over K, we may assume L/K is Galois. Let R be some extension of $\mathcal{O}$ to L. Because $\mathrm{Gal}(L/K)$ acts transitively on the extensions, there are no inclusions among the extensions, so R is maximal. Let $a_1, \ldots, a_n$ be some finite subset of R such that R is the only extension of $\mathcal{O}$ containing $a_1, \ldots, a_n$. For $x \in L$, the following are equivalent:

- x is in the maximal ideal of R
- $1/x$ is not in R
- No extension of $\mathcal{O}$ to L contains $a_1, \ldots, a_n, 1/x$.
- 1 can be written as a polynomial of degree less than d in $a_1, \ldots, a_n, 1/x$, with coefficients in the maximal ideal of $\mathcal{O}$.

This last condition is type-definable (or definable, respectively). $\square$

Any superring of a valuation ring is a valuation ring. So, given two valuation rings $\mathcal{O}_1$ and $\mathcal{O}_2$ of a field K, the ring $\mathcal{O}_3$ generated by $\mathcal{O}_1$ and $\mathcal{O}_2$ is a third valuation ring of K. It is the finest common coarsening of $\mathcal{O}_1$ and $\mathcal{O}_2$. We'll call $\mathcal{O}_3$ the *join* of $\mathcal{O}_1$ and $\mathcal{O}_2$.

Let $v_i, \Gamma_i, \mathfrak{m}_i$ be the associated data. As v_3 is a coarsening of v_i for $i = 1, 2$, there is some convex subgroup Δ_i of Γ_i such that v_3 is v_i composed with $\Gamma_i \to \Gamma_i/\Delta_i$.

Lemma 13.4. *An element $\gamma \in \Gamma_1$ is greater than Δ_1 exactly if $\gamma \geq v_1(x)$ for all $x \in K$ with $v_2(x) \leq 0$.*

Proof. Write $\gamma = v_1(c)$. First suppose $v_1(c) > \Delta_1$. This means exactly that $v_3(c) > 0$. Suppose $v_2(x) \leq 0$. Then $v_3(x) \leq 0$, so $v_3(c) > v_3(x)$, which implies $v_1(c) > v_1(x)$, and consequently $\gamma = v_1(c) \geq v_1(x)$.

Conversely, suppose $v_1(c)$ is less than some element of Δ_1. So $v_3(c) \leq 0$, and $c^{-1} \in \mathcal{O}_3$.

Let $D \subset \Gamma_1$ be the submonoid

$$\{v_1(x) \ : \ x \in K, \ v_2(x) \leq 0, \ v_1(x) \geq 0\}$$

Let Δ be the convex hull of $D \cup (-D)$. Then Δ is a convex subgroup of Γ_1. Let $\mathcal{O}_\Delta$ denote the valuation ring of the corresponding coarsening of $\mathcal{O}_1$.

We claim that $\mathcal{O}_\Delta$ contains $\mathcal{O}_3$. As a coarsening of $\mathcal{O}_1$, $\mathcal{O}_\Delta$ certainly contains $\mathcal{O}_1$, so we need to show $\mathcal{O}_2 \subset \mathcal{O}_\Delta$. Suppose $v_2(x) \geq 0$; we will show $x \in \mathcal{O}_\Delta$. If $v_1(x) \geq 0$ then $x \in \mathcal{O}_1 \subseteq \mathcal{O}_\Delta$. So suppose $v_1(x) < 0$. Then $v_2(x^{-1}) \leq 0$ and $v_1(x^{-1}) \geq 0$ so by definition of D, $v_1(x^{-1}) \in D$. So $v_1(x) \in \Delta$, which implies $x \in \mathcal{O}_\Delta^\times \subseteq \mathcal{O}_\Delta$.

So $c^{-1} \in \mathcal{O}_3 \subseteq \mathcal{O}_\Delta$. Therefore c is not in the maximal ideal of $\mathcal{O}_\Delta$, which means that $\gamma = v_1(c)$ is less than some element of Δ, hence less than some element of D. Consequently, there is some x such that

$$\gamma < v_1(x) \geq 0$$

$$v_2(x) \leq 0$$

So it is not true that $\gamma \geq v_1(x)$ for all x with $v_2(x) \leq 0$. $\square$

Lemma 13.5. *Suppose $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable. Suppose X_i is a subset of K for $i = 1, 2$, such that*

$$\mathfrak{m}_i \subseteq X_i \subseteq \mathcal{O}_i$$

For instance, X_i could be $\mathfrak{m}_i$ or $\mathcal{O}_i$. Then an element $a \in K$ is in $\mathfrak{m}_3$ if and only if $a \cdot X_2 \subseteq X_1$.

Proof. First suppose $a \in \mathfrak{m}_3$, i.e., $v_3(a) > 0$. If $x \in X_2 \subseteq \mathcal{O}_2$, then $v_2(x) \geq 0$, so $v_3(x) \geq 0$, so $v_3(a \cdot x) > 0$, so $v_1(a \cdot x) > 0$, so $a \cdot x \in \mathfrak{m}_1 \subset X_1$. This shows $a \cdot X_2 \subseteq X_1$.

Conversely, suppose $a \cdot X_2 \subseteq X_1$. By the incomparability assumption, we can find $b \in K$ such that $v_2(b) < 0$ and $v_1(b) \geq 0$. To show that $a \in \mathfrak{m}_3$, we need to show that $v_1(a) > \Delta_1$. By the previous lemma, it suffices to assume $v_2(c) \leq 0$ and show $v_1(a) \geq v_1(c)$.

As $v_2(b \cdot c) < v_2(c) \leq 0$, we see that $b^{-1} \cdot c^{-1} \in \mathfrak{m}_2 \subseteq X_2$. So $a \cdot b^{-1} \cdot c^{-1} \in X_1 \subseteq \mathcal{O}_1$ by assumption. This implies

$$v_1(a) - v_1(b) - v_1(c) \geq 0.$$

So $v_1(a) > v_1(b) + v_1(c) \geq v_1(c)$, which was what we wanted to show. $\square$

Corollary 13.6. *Suppose $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable and $\vee$-definable (equivalently, the maximal ideals are type-definable). Then the join is definable.*

Proof. We can find *definable* X_1 and X_2 as in the previous lemma, because disjoint type-definable sets are separated by definable sets. (So $\mathfrak{m}_i$ is separated from the complement of $\mathcal{O}_i$ by some definable set X_i.) The previous lemma then gives a definable condition for being in the maximal ideal of the join. $\square$

Recall that two non-trivial valuations are *independent* if they induce distinct topologies.

Lemma 13.7. *If $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable, then the join $\mathcal{O}_3$ is trivial if and only if $\mathcal{O}_1$ and $\mathcal{O}_2$ are independent. Moreover, $\mathcal{O}_1$ and $\mathcal{O}_2$ induce independent valuation rings on the residue field of $\mathcal{O}_3$.*

Proof. If the join is non-trivial, then $\mathcal{O}_i$ induces the same topology as $\mathcal{O}_3$, because non-trivial coarsenings don't change the topology. So $\mathcal{O}_1$ and $\mathcal{O}_2$ aren't independent.

Conversely, suppose that $\mathcal{O}_1$ and $\mathcal{O}_2$ aren't independent—they induce the same topology. Let $X_i = \mathcal{O}_i$. Then $\{\alpha \cdot X_i \ : \ \alpha \in K^\times\}$ is a neighborhood basis for the common topology.

Consequently, there is some α such that $\alpha \cdot X_2 \subseteq X_1$. This shows that there is a nonzero element in the maximal ideal of the join, so the join is non-trivial.

Let k_i denote the residue field of $\mathcal{O}_i$. The places $K \to k_1$ and $K \to k_2$ factor through $K \to k_3$. If $k_3 \to k_1$ and $k_3 \to k_2$ aren't independent, they would factor through their join $k_3 \to k'$ which would be nontrivial. Then $K \to k'$ contradicts the fact that $K \to k_3$ is the finest place which is a common coarsening of $K \to k_1$ and $K \to k_2$. □

13.2 Proving henselianity

Now we return to our setting of a monster dp-minimal field $\mathbb{C}$, not strongly minimal.

Lemma 13.8. *Let R be an M-definable non-trivial valuation ring on $\mathbb{C}$. Then $\mathcal{O}_M$ is a coarsening of R. In particular, both induce the same topology, which is just the standard topology on $\mathbb{C}$ we have been considering.*

Proof. Let $\mathfrak{m}_R$ denote the maximal ideal of R. We need to prove that $R \subseteq \mathcal{O}_M$, or equivalently, that $\mathfrak{m}_M \subseteq \mathfrak{m}_R$. That is, we need to show that every M-infinitesimal is in the maximal ideal of R.

Since R is non-trivial, the value group of R is infinite, which implies $\mathfrak{m}_R$ is infinite. By Theorem 6.17, it has interior. Since it is a subgroup of the additive group, it is actually open. So it contains an M-definable neighborhood of 0, which includes all infinitesimals. □

Lemma 13.9. *Let L be a finite extension of $\mathbb{C}$. Let R be a non-trivial valuation ring on L. Then R defines the product topology on L.*

Proof. Let M be a small model over which everything is defined. Then we have two definable V-topologies on L, both defined over M: one comes from R and one is the product topology. Let $\mathfrak{m}_1$ and $\mathfrak{m}_2$ be the M-infinitesimals with respect to these two topologies. As they are V-topologies, these are maximal ideals of some valuation rings on L. Also, $\mathfrak{m}_1$ and $\mathfrak{m}_2$ have the same restriction to K because the two topologies agree on K, by the previous lemma. By Theorem 12.2, $\mathfrak{m}_1$ and $\mathfrak{m}_2$ are the same. A ring topology is determined by its type-definable group of infinitesimals, so the two topologies are the same. □

In particular, two definable valuation rings on a finite extension of a dp-minimal field can never be independent, since they will always induce the same topology.

Theorem 13.10. *Any definable valuation R on a dp-minimal field $\mathbb{C}$ is henselian.*

Proof. If not, let $L/\mathbb{C}$ be a finite Galois extension such that R has multiple extensions to L. Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be two such extensions. They are definable by Lemma 13.3. From the transitive action of $\mathrm{Gal}(L/\mathbb{C})$ on the extensions, it is clear that $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable. Let $\mathcal{O}$ be their join; it is definable by Corollary 13.6. Let F be the residue field of $\mathcal{O}$. It is a finite extension of the residue field k of $\mathcal{O} \cap \mathbb{C}$. By Lemma 13.7, $\mathcal{O}_1$ and $\mathcal{O}_2$ induce independent valuations on F. This shows that F and k are infinite. As k is the image of a set of dp-rank 1 under a definable map, it is also dp-minimal. By Lemma 13.9 applied to F/k, we have a contradiction. □

Remark 13.11. *In fact, from the proof we get the following fact which we will use later: if R and R' are two definable valuations on a finite extension of $\mathbb{C}$, then R and R' are comparable. Otherwise, on the residue field of their join, one has two independent valuations, which cannot happen in finite extensions of dp-minimal fields.*

Theorem 13.12. *Let M be a small model. Then $\mathcal{O}_M$ is henselian.*

Proof. Suppose not. Let $L/\mathbb{C}$ be a finite Galois extension of $\mathbb{C}$ such that $\mathcal{O}_M$ has multiple extensions to L. Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be two such extensions. They are incomparable, so by Corollary 13.6, their join $\mathcal{O}_3$ is *definable.* It is also non-trivial. Indeed, suppose $M' \succeq M$ is a model over which L is defined. Then $\mathcal{O}_{M'}$ is a coarsening of $\mathcal{O}_M$ (because M'-infinitesimals are M-infinitesimals). By Theorem 12.2, there is a unique extension $\mathcal{O}_4$ of $\mathcal{O}_{M'}$ to L. Then $v_{M'}$ is $\pi \circ v_M$ for some projection $\pi : \Gamma \to \Gamma/\Delta$, and if we coarsen v_1 corresponding to $\mathcal{O}_1$ using the same Δ, we get a valuation on L extending $v_{M'}$ on $\mathbb{C}$, which must necessarily be v_4. So $\mathcal{O}_4$ is a coarsening of $\mathcal{O}_1$, and similarly of $\mathcal{O}_2$. This shows that the join of $\mathcal{O}_1$ and $\mathcal{O}_2$ is non-trivial, since it is finer than $\mathcal{O}_4$.

So there is a definable non-trivial valuation on $\mathbb{C}$. As M is an elementary substructure, there is also an M-definable valuation. It is henselian by Theorem 13.10. By Lemma 13.8, $\mathcal{O}_M$ is a coarsening of this henselian valuation, so $\mathcal{O}_M$ is henselian itself. $\square$

14 The canonical valuation

In this section, we will assume that $\mathbb{C}$ is a monster dp-minimal field. Recall that $\mathbb{C}$ is perfect by Remark 12.1.

Let $\mathcal{V}$ denote the set of valuation rings on $\mathbb{C}$. It is non-empty because of the trivial valuation. By remark 13.11, the members of $\mathcal{V}$ are totally ordered by inclusion (i.e., by coarseness). By Theorem 13.10, they are all henselian.

Suppose $L/\mathbb{C}$ is a finite extension of $\mathbb{C}$. If $\mathcal{O}$ is a definable valuation ring on $\mathbb{C}$, then $\mathcal{O}$ has a unique extension to L, by henselianity. We will denote this by $\mathcal{O}|L$. It is also definable, by Beth implicit definability. The map $\mathcal{O} \mapsto \mathcal{O}|L$ gives a strictly order-preserving bijection between definable valuation rings on $\mathbb{C}$ and L. (Strict order preservation is true on general grounds, but can be seen in this case by Remark 13.11 applied to L.)

Let $\mathcal{V}$ denote the set of all non-trivial definable valuation rings on $\mathbb{C}$, and let

$$\mathcal{O}_\infty = \bigcap_{\mathcal{O} \in \mathcal{V}} \mathcal{O}$$

denote the intersection.

Lemma 14.1. *$\mathcal{O}_\infty$ is a type-definable henselian valuation ring on $\mathbb{C}$. Moreover, the unique extension of $\mathcal{O}_\infty$ to a finite field extension L of $\mathbb{C}$, is*

$$\mathcal{O}_\infty|L := \bigcap_{\mathcal{O} \in \mathcal{V}} \mathcal{O}|L$$

which is also type-definable and henselian.

Proof. The intersection $\mathcal{O}_\infty$ is a valuation ring because $\mathcal{V}$ is totally ordered, so we are taking an intersection of a chain.

Type-definability of $\mathcal{O}_\infty$ follows because we are taking an intersection of a small number of definable families. (Every definable valuation ring is a member of a $\emptyset$-definable family, because whether or not a definable set is a valuation ring is expressible by a single formula.)

The same arguments show that $\mathcal{O}_\infty|L$ as defined above is a type-definable valuation ring on L. It is clearly a valuation ring extending $\mathcal{O}_\infty$.

Suppose $L/\mathbb{C}$ is a Galois extension. Each $\mathcal{O} \in \mathcal{V}$ is Henselian, so the extension $\mathcal{O}|L$ is $\mathrm{Gal}(L/\mathbb{C})$-invariant. Therefore so is the intersection $\mathcal{O}_\infty|L$. So $\mathcal{O}_\infty|L$ is a $\mathrm{Gal}(L/\mathbb{C})$-invariant extension of $\mathcal{O}_\infty$. As $\mathrm{Gal}(L/\mathbb{C})$ acts transitively on the extensions, it follows that $\mathcal{O}_\infty$ has a unique extension to L. As L was an arbitrary Galois extension, $\mathcal{O}_\infty$ is henselian. Therefore so is its extension to any finite field extension of $\mathbb{C}$. □

Theorem 14.2. *The residue field of $\mathcal{O}_\infty$ is algebraically closed, real closed, or finite. In the third case, $\mathcal{O}_\infty$ is definable, not just type-definable.*

To prove this, we will use some results from Jahnke and Koenigsmann's paper "Uniformly defining p-henselian valuations." First we review some facts about p-henselianity.

If K is a field, $K(p)$ denotes the compositum of all Galois extensions of K of p-power degree. $K(p) = K$ if and only if K has no cyclic Galois extensions of degree p. Say that K is "p-closed" in this case. The map $K \mapsto K(p)$ is a closure operation on subfields of some ambient algebraically closed field. (This boils down to the fact that if $K(\alpha)/K$ is p-cyclic and L/K is any extension whatsoever, then $L(\alpha)/L$ is trivial or p-cyclic.)

We will also start using the notation vK for the value group and Kv for the residue field of a valuation v on a field K.

Definition 14.3. *A field K satisfies enough$_p$ if it has all the pth roots of unity, and when $p = 2$, K has all the 4th roots of unity.*

Lemma 14.4. *Suppose K satisfies enough$_p$.*

1. *Any field extending K also satisfies enough$_p$.*
2. *If K is characteristic p, then K is p-closed if and only if the Artin-Schreier map is surjective. Otherwise, K is p-closed if and only if the pth-power map is surjective.*
3. *If v is a valuation on K, then the residue field Kv satisfies enough$_p$.*
4. *If v is equicharacteristic, and K is p-closed, so is Kv.*
5. *If K is not p-closed, neither is any finite extension of K.*

Proof. Everything is an easy exercise, but here are the proofs anyways:

1. Obvious.
2. Kummer theory.

3. If K has all the nth roots of unity, so does Kv.

4. In equicharacteristic p, we need to show the Artin-Schreier map on Kv is surjective, assuming it is surjective for K. This is proven in the obvious way: given $x \in Kv$, write x as the residue of $x' \in K$, let y' be an Artin-Schreier roots. Check that $v(y') = 0$, so the residue y of y' is an Artin-Schreier root of x. Do something similar with the pth power maps in the other case.

5. This is essetially the Artin-Schreier theorem. Because p-closure is a closure operation, if some finite extension L of K is p-closed, then $K(p) \subseteq L$ so $K(p)/K$ is finite. It follows that there is some field F, satisfying *enough*$_p$, such that some p-cyclic extension is p-closed. This extension must be $F(\sqrt[p]{a})$ for some $a \in F$, by Kummer theory. In characteristic prime to p, $F(\sqrt[p]{a})^\times$ must be p-divisible. The image of the norm map $F(\sqrt[p]{a})^\times \to F^\times$ must be p-divisible. The norm of $\sqrt[p]{a}$ is a (or $-a$ if $p = 2$), so a was already a pth power, a contradiction.

   In characteristic p, we get some field F which is not p-closed, such that $F(\alpha)$ is p-closed, where $\alpha^p - \alpha =: a \in F$, and $[F(\alpha) : F] = p$. The Artin-Schreier map on $F(\alpha)$ sends
$$\sum_{i=0}^{p-1} x_i \cdot \alpha^i \mapsto \sum_{i=0}^{p-1} x_i^p (a + \alpha)^i - x_i \alpha^i$$
   for $\vec{x} \in F^p$, because $\alpha^p = a + \alpha$. If we expand the right hand side in terms of the basis $\{\alpha^0, \alpha^1, \ldots, \alpha^{p-1}\}$, the coefficient of α^{p-1} is $x_{p-1}^p - x_{p-1}$, which must be in the image of the non-surjective Artin-Schreier map on F. So the Artin-Schreier map on $F(\alpha)$ is not surjective either, a contradiction.

□

Recall that a valuation v on a field K is *p-henselian* if it has a unique extension to $K(p)$. This is a weaker notion than henselianity.

The main theorem 3.1 of Jahnke and Koenigsmann's paper "Uniformly defining p-henselian valuations" implies[2] the following:

Fact 14.5. *Suppose K satisfies enough$_p$. Then there is a definable p-henselian valuation v on K such that at least one of the following is true:*

1. *v is the finest p-henselian valuation on K*

2. *The residue field Kv is p-closed.*

Note that v might be the trivial valuation.

In our case, this tells us the following:

[2]We have rigged *enough*$_2$ to ensure that Kv has the 4th roots of unity if it is characteristic zero, so Kv is not orderable, hence not "Euclidean."

Lemma 14.6. *Let K be a sufficiently saturated dp-minimal field and L be a finite extension satisfying enough$_p$. Then L has a definable henselian valuation v with finite or p-closed residue field. (v might be trivial.)*

Proof. Let v be the definable valuation on L from Jahnke and Koenigsmann. Its restriction to K is henselian by Theorem 13.10, so v is itself henselian. We will conflate v with its restriction to K.

Suppose the residue field vL is infinite. If vK is strongly minimal, then it is algebraically closed, and vL is an algebraic extension of vK, so vL is algebraically closed, hence p-closed.

Otherwise, Theorem 13.12 applied to vK shows that there is a non-trivial henselian valuation on vK. This induces a non-trivial henselian valuation on the finite extension vL of vK. This in turn induces a henselian valuation on L which is strictly finer than v. So v is not the finest henselian valuation on L, and certainly not the finest p-henselian valuation. Therefore vL is p-closed. □

Using this, we prove Theorem 14.2.

Proof. Let v_∞ be the valuation corresponding to $\mathcal{O}_\infty$. For v a definable valuation on $\mathbb{C}$, or v_∞, we'll abuse notation and identify v with its extension to any finite field extension L of $\mathbb{C}$. So we'll write Lv to denote the residue field of the extending valuation. It will be a finite extension of $\mathbb{C}v$.

If some definable valuation v on $\mathbb{C}$ has finite residue field, then v cannot be coarser than any other valuation, because finite fields don't admit non-trivial valuations. So $v = v_\infty$. Therefore v_∞ is definable and has finite residue field, and we're done.

So we may assume

$$\text{No definable valuation on } \mathbb{C} \text{ has finite residue field.} \tag{3}$$

We will show that $\mathbb{C}v_\infty$ is algebraically closed or real closed.

If v_∞ is equicharacteristic, let v_0 be the trivial valuation on $\mathbb{C}$. Otherwise, $1/p \notin \mathcal{O}_\infty$ for some $\mathcal{O}$, so some definable valuation ring avoids $1/p$. Let v_0 be the corresponding valuation. Then v_0 is mixed characteristic. So we have arranged that $\mathbb{C}v_0$ has the same characteristic as $\mathbb{C}v_\infty$.

Claim 14.7. *Suppose $L/\mathbb{C}$ is a finite extension satisfying enough$_p$. Then Lv_∞ is p-closed.*

Proof. By Lemma 14.4-3, Lv_0 satisfies *enough$_p$*. The residue field Lv_0 is infinite by (3). So it is an interpretable dp-minimal field. By Lemma 14.6, it has a definable valuation with finite or p-closed residue field. This induces a definable valuation v on L, as fine as v_0, such that Lv is finite or p-closed. By (3), Lv is p-closed. Also, Lv satisfies *enough$_p$* by Lemma 14.4-3.

By choice of v_∞, it is as fine as v. So, there is some valuation w on Lv_∞ whose residue field is Lv_∞. By choice of v_0, the valuation w will be equicharacteristic, so Lemma 14.4-4 applies. In particular, p-closedness of Lv implies p-closedness of Lv_∞. □

Suppose $\mathbb{C}v_\infty$ is not algebraically closed or real closed. It is perfect, because it is the residue field of a valuation on $\mathbb{C}v_0$, which is perfect by Remark 12.1.

Claim 14.8. *There is some finite Galois extension F of $\mathbb{C}v_\infty$ which satisfies $enough_p$ and is not p-closed for some prime p.*

Proof. Let F_0 be $\mathbb{C}v_\infty(\sqrt{-1})$ in characteristic not 2, and $\mathbb{C}v_\infty$ otherwise. By assumption, F_0 is perfect and not algebraically closed.

Let p be the least prime such that some extension of F_0 has order dividing p. Every polynomial in $F_0[X]$ of degree less than p splits, so F_0 has all the pth roots of unity. If $p = 2$, then F_0 also has the 4th roots of unity, by choice of F_0. So F_0 satisfies $enough_p$.

Let F/F_0 be a finite Galois extension with order dividing the prime p. Enlarging F a little, we may assume F is Galois over $\mathbb{C}v_\infty$. Taking a cylic subgroup of $\mathrm{Gal}(F/F_0)$ of degree p, we get an intermediate field K between F and F_0 such that $\mathrm{Gal}(F/K) \cong \mathbb{Z}/p$. Consequently, K is not p-closed. By Lemma 14.4 parts 1 and 4, the field F satisfies $enough_p$ and is not p-closed. □

Now we get a contradiction as follows: let $L/\mathbb{C}$ be a large enough finite extension of $\mathbb{C}$ such that L satisfies $enough_p$ and Lv_∞ contains the field F from Claim 14.8. The field Lv_∞ is a finite extension of F, so by Lemma 14.4-4, it is not p-closed, contradicting Claim 14.7. □

15 Defectlessness

Here is a well-known fact from valuation theory:

Fact 15.1. *Let $(L, w)/(K, v)$ be a finite extension of valued fields. Then*

$$[L : K] \geq |wL/vK| \cdot [Lw : Kv]$$

If K is henselian, then one in fact has

$$[L : K] = |wL/vK| \cdot [Lw : Kv] \cdot p^e$$

for some $e \geq 0$, where p is the residue characteristic (or 1 if the residue characteristic is zero). Either e or p^e is called the "defect," and K is said to be defectless if the defect for every finite extension L/K vanishes. (So, residue-characteristic zero henselian fields are defectless.)

Also, a valued field is *algebraically maximal* if it has no algebraic immediate extensions. Defectless fields are algebraically maximal, and finite extensions of defectless fields are defectless, so finite extensions of defectless fields are algebraically maximal.

The converse is also true, at least if you assume perfection. Suppose K is a perfect henselian field that is not defectless. Then some Galois extension L/K has defect. Let p be the residue characteristic. By the usual p-Sylow trick, we can find a sequence of intermediate fields

$$K < F_0 < F_1 < \cdots < F_n = L$$

such that F_0/K has degree prime to p and F_{i+1}/F_i is cyclic of degree p for each p. Then F_0/K is defectless, so one of the cyclic extensions F_{i+1}/F_i has defect. As F_{i+1}/F_i has degree p, this can only happen if it is an immediate extension.

Algebraically maximal fields are henselian (I guess because fields have maximal immediate extensions, which are maximally complete hence henselian, and the relative algebraic closure within this maximal immediate extension is the henselization, so henselization is always an immediate extension).

Remark 15.2. *If $K_1 \to K_2$ and $K_2 \to K_3$ are algebraically maximal places, then so is their composition $K_1 \to K_3$.*

Proof. If not, let $L_1 \to K_3$ be an immediate extension of $K_1 \to K_3$. The value groups are the same, so we can coarsen $L_1 \to K_3$ using the same convex subgroup, to get $L_1 \to L_2$ extending $K_1 \to K_2$. Algebraic maximality of $K_1 \to K_2$ ensures L_2 is strictly bigger than K_2 (since the value groups are the same). Then $L_2 \to K_3$ is an algebraic immediate extension of $K_2 \to K_3$, which is impossible. □

Theorem 15.3. *Let $\mathbb{C}$ be a reasonably saturated dp-minimal field. Then there is a not necessarily definable valuation v on $\mathbb{C}$ such that*

- *v is henselian and defectless*
- *The value group Γ satisfies $|\Gamma/n\Gamma| < \infty$ for all n.*
- *The residue field is a model of ACF_p or a characteristic zero local field.*
- *If $\mathbb{C}$ has characteristic p, then Γ is p-divisible.*
- *If v has mixed characteristic, then every element of $[-v(p), v(p)]$ is p-divisible.*

In the local field case, the elementary equivalence class of $\mathbb{C}$ in the language of pure rings is determined by that of Γ, because of Ax-Kochen-Ershov. This also happens in the other case, as we'll see in the next section.

Proof. The condition that $|\Gamma/n\Gamma| < \infty$ always follows directly from Lemma 10.5, so we will forget about it.

Let v_∞ be the canonical valuation from Theorem 14.2. If the (hyper-definable) residue field k of v_∞ is real closed or algebraically closed of characteristic zero, we can find a henselian valuation on k with residue field the real numbers or the complex numbers.

Next suppose that the residue field of v_∞ is finite and v_∞ is definable. We claim that $[-v(p), v(p)] \subset \Gamma$ is finite, i.e., there is bounded amounts of ramification. Otherwise, the maximal convex subgroup Δ of Γ avoiding $v(p)$, would be non-trivial. This group is type-definable, but it becomes definable in the Shelah expansion[3], because all cuts in the value group become definable. The Shelah expansion is still dp-minimal, by quantifier elimination in the Shelah expansion. In the Shelah expansion, we obtain an interpretable valued field

[3]The expansion by all externally definable sets.

with value group Δ and finite residue field. This contradicts a result of Kaplan-Scanlon-Wagner, that NIP valued fields cannot have finite residue field.

So there is bounded ramification. Let Δ now be the smallest convex subgroup of Γ containing $v(p)$. Coarsening v_∞ by Δ, we obtain a henselian valuation v whose residue field K is characteristic zero, and has a henselian valuation with finite residue field and value group $\mathbb{Z}$. Saturation of $\mathbb{C}$ implies that any countable chain of balls in $\mathbb{C}$ has non-empty intersection. This descends to K, so K is actually spherically complete and hence complete. So K is a complete discrete valuation ring with bounded ramification and finite residue field. It is therefore a local field of characteristic zero.

Finally, we have the confusing case, where v_∞'s residue field is algebraically closed of positive characteristic. In this case, we will take $v = v_\infty$. (Except when v_∞ is trivial, in which case the underlying field is a model of ACF_p, and can be expanded to a model of $ACVF_{p,p}$, in which all the conditions are clear. We will assume we are not in this case, so v_∞ is non-trivial.)

First consider the case of pure characteristic p. In this case, $\mathbb{C}$ is perfect, by Remark 12.1. By Kaplan-Scanlon-Wagner, p does not divide the degree of any finite extension of $\mathbb{C}$, so any henselian valuation on $\mathbb{C}$ will be defectless. The p-divisibility of the value group follows from Kaplan-Scanlon-Wagner.

The mixed characteristic case remains. In this case, unbounded ramification *must* occur. Otherwise, let τ be an element of minimal positive valuation, so τ generates the maximal ideal of $\mathcal{O}_\infty =: \mathcal{O}$. The pth power map on $k^\times$ is a bijection, so by the snake lemma applied to the pth power maps on the short exact sequence

$$1 \to (1 + \tau \cdot \mathcal{O})^\times \to \mathcal{O}^\times \to k^\times \to 1$$

we see that the cokernel of the pth power map on $1 + \tau \cdot \mathcal{O}$ maps into the cokernel of the pth power map on $\mathcal{O}^\times$, which must be finite by Lemma 10.5. But the image of the pth power map on $1 + \tau \cdot \mathcal{O}$ lands in the ideal generated by $p \cdot \tau$ and τ^p, which has infinite index in the ideal generated by τ, a contradiction.

Let v_0 be some definable valuation of mixed characteristic. So we can decompose the place $\mathbb{C} \to \mathbb{C}v_\infty$ as a composition $\mathbb{C} \to \mathbb{C}v_0 \to \mathbb{C}v_\infty$. The valuation v_0 is a coarsening by some convex subgroup Δ_0 of Γ, and v_0's valuation group is $\Gamma_0 := \Gamma/\Delta_0$. By the positive characteristic case, we know that the place $\mathbb{C}v_0 \to \mathbb{C}v_\infty$ is defectless, and its value group Δ_0 is p-divisible.

Because v_0 is definable, Γ_0 is interpretable. The unbounded ramification argument applies just as well to v_0. Let Δ_a be the largest convex subgroup of Γ_0 avoiding $v(p)$, and let Δ_b be the smallest convex subgroup of Γ_0 containing $v(p)$. In the Shelah expansion, these sets are definable, and Δ_a is now the value group of some interpretable valued field of characteristic p. By Kaplan-Scanlon-Wagner, Δ_a must be p-divisible. Now forget about the Shelah expansion.

There is a unique largest p-divisible convex subgroup G of Γ_0 and it is definable. We just saw that G contains Δ_a. By unbounded ramification, Δ_a is not definable itself, so by compactness, G must be strictly bigger. Since Δ_b is the smallest convex subgroup bigger than Δ_a, G must contain Δ_b.

Now, the convex subgroup of Γ generated by $v(p)$ sits in a short exact sequence with Δ_0 and Δ_b. Both groups are uniquely p-divisible: we just saw this for Δ_b, and for Δ_0, it followed from the positive characteristic case. So the convex subgroup in question is p-divisible.

It remains to show defectlessness. Let $L/\mathbb{C}$ be a finite extension. We will show that $L \to Lv_0$ and $Lv_0 \to Lv_\infty$ are both algebraically maximal. The case of $Lv_0 \to Lv_\infty$ follows from the positive characteristic case, i.e., the fact that Lv_0 has no algebraic extensions of degree divisible by p. Algebraic maximality of $L \to Lv_0$ remains.

Recall the groups Δ_a and Δ_b. Coarsening $L \to Lv_0$ according to these groups (or their convex closure in v_0L), we decompose the place $L \to Lv_0$ as a composition

$$L \to L_1 \to L_2 \to Lv_0$$

of henselian places, where L and L_1 have characteristic zero, L_2 and Lv_0 have positive characteristic, $L \to L_1$ has value group Γ_0/Δ_b, $L_1 \to L_2$ has value group Δ_b/Δ_a, and $L_2 \to Lv_0$ has value group Δ_a.

By the Shelah expansion trick L_2 is a perfect NIP field, so it has no algebraic extensions of degree divisible by p, and $L_2 \to Lv_0$ is automatically defectless and algebraically maximal.

Likewise, $L \to L_1$ is henselian of equicharacteristic zero, so it is defectless and algebraically maximal.

Finally, $L_1 \to L_2$ is algebraically maximal because it is spherically complete. Saturation of (L, v_0) ensures that any countable chain of balls in L has non-empty intersection. This descends to the place $L_1 \to L_2$. But since this place has archimedean value group, this implies actual spherical completeness. Spherically complete fields are algebraically maximal.

So $L \to L_1 \to L_2 \to Lv_0 \to Lv_\infty$ is algebraically maximal, completing the proof. $\square$

16 Ax-Kochen-Ershov type results

Theorem 16.1. *Let T_G be a complete theory of p-divisible ordered abelian groups. Let T be the theory of (henselian) defectless valued fields with value group satisfying T_G and with algebraically closed characteristic p residue field. Then T is complete after specifying the characteristic, and specifying the T_G-type of $v(p)$ in the mixed characteristic case.*

Proof. We make a few preliminary remarks on models of T. First of all, if M is a model of T, every finite algebraic extension of M has degree prime to p. Indeed, if L/M is such an extension, then L and M have the same residue field, because M's is algebraically closed. The "degree" of the value group extension is prime to p, because the value group is p-divisible.

As a consequence, models of M are perfect. Moreover, for any $m \in \mathbb{N}$, an element $x \in M$ has an mth root if and only if $\mathrm{val}(x)$ is divisible by m. This reduces easily to showing that elements of valuation zero have ℓth roots for all primes ℓ. If ℓ is not p, this follows from henselianity and the fact that $x^\ell - 1$ splits in the residue field. If ℓ is p, then any polynomial of the form $x^p - a$ is either irreducible or a has a pth root. In the former case, we contradict the fact that M has no algebraic extension of degree divisible by p.

Let L be the standard 1-sorted language of valued fields, expanded by all relations of the form $R(\mathrm{val}(x_1), \ldots, \mathrm{val}(x_n))$ where R is an n-ary formula in T_G, without parameters.

One can check that models of T exist (one can produce a valued field with the desired characteristic, value group, and residue field. Then pass to a maximal immediate extension. Finite extensions of spherically complete fields are spherically complete [because of "separating bases"], so they are algebraically maximal. Consequently, spherically complete = maximally complete fields are defectless). So it suffices to show that T has quantifier elimination in the language L.

It suffices to prove the following. Suppose M and N are models of T, N is $|M|^+$-saturated, and f is a maximal partial isomorphism from M to N (i.e., an isomorphism from a substructure of M to a substructure of N). Then f is total.

Let K be the domain of f. Then K is a subring by choice of the language. Even better...

- K is a field, because there is a unique way to extend the structure from K to $\mathrm{Frac}(K)$ compatible with $T_\forall$. (The new valuations are all linear combinations of the old ones, so their T_G-types are uniquely determined.)

- K is perfect, because both M and N contain the perfect closure of K and there's a unique way to extend the valuation to K^{perf}. The T_G-types of the new valuation are determined, as before.

- K is a henselian field, because both M and N contain the henselization of K so we could extend f otherwise.

- Any element $a \in K$ which has an mth root in M has one in K. If a does not have an mth root in K, then $x^m - a$ is irreducible (on general grounds). Whether or not a has an mth root is determined by whether $\mathrm{val}(a)$ is m-divisible. This is part of the quantifier free L-type of a, so $f(a)$ also has an mth root in N. The fields $K(\sqrt[m]{a})$ and $f(K)(\sqrt[m]{f(a)})$ are isomorphic, and there's a unique way to extend the L-structure from K to $K(a)$ (by henselianity of K), so we could extend f.

- The value group of K is pure in the value group of M, by the previous point. It is also pure in the value group of N, because this is determined by its T_G-type, and the map on the value groups is a partial elementary map.

- The residue field of K is separably closed. Otherwise, let $P(X)$ be a polynomial over K inducing a separable irreducible monic polynomial on the residue field. Because M is henselian and has algebraically closed residue field, $P(X)$ splits in M. Similarly, its pushforward under f splits in N. Now there is a unique way to extend the L-structure from K to the splitting field of $P(X)$, by henselianity of K. So we can extend f from K to the splitting field of $P(X)$, which is present in both M and N.

- K has all nth roots of unity, for n prime to p. This follows from henselianity and the fact that the residue field has all these roots.

- The Galois group of K is solvable. To see this, let L be the closure of K under radicals. This is a solvable extension of K, and the value group of L (with respect to the unique extension of the valuation from K to L) is now divisible. The residue field of L is an algebraic extension of a separably closed field, so it too is separably closed. Now L is perfect (because it is algebraic over the perfect field K), and every finite algebraic extension of L has degree a power of p (it must be a power of p times the degree of the residue field extension, which will be a power of p). So the Galois group of L is solvable, because it is a p-group. A solvable extension of a solvable extension is solvable.

- If $a \in M$ is algebraic over K, then $[K(a) : K]$ is a power of p. First of all, $\Gamma(K)$ is pure in $\Gamma(M)$, so there is no value group extension. The residue field extension can only involve a power of p, because K has separably closed residue field. The degree of $K(a)$ over K will be some power of p times the residue field degree.

- K is relatively algebraically closed in M. Otherwise, take $a \in M$ algebraic over K, but not in K. Within M^{alg}, let L be the Galois closure of $K(a)$. The extension $L \cap M/K$ doesn't extend the value group, so it has degree a power of p. On the other hand, $\mathrm{Gal}(M^{alg}/M)$ surjects onto $\mathrm{Gal}(L/L \cap M)$ (because M is perfect, among other things), and $\mathrm{Gal}(M^{alg}/M)$ has nothing to do with the prime p, so neither does $\mathrm{Gal}(L/L \cap M)$. This shows that the extension $L/L \cap M$ has degree prime to p.

  By Galois theory, $L \cap M$ is the fixed field of a Hall prime-to-p subgroup of $\mathrm{Gal}(L/K)$. By the same argument on the other side, $L \cap N$ is also a fixed field of a Hall prime-to-p subgroup. Because Hall subgroups are unique up to conjugation in solvable groups, $L \cap M$ and $L \cap N$ are actually isomorphic as fields over K. As usual, this contradicts the maximality of f, unless $K = L \cap M$. This implies that $[L : K] = [L : L \cap M]$ which is prime to p. So $[K(a) : K]$ is prime to p, contradicting the previous bullet point.

- K has algebraically closed residue field. It suffices to show that it has perfect residue field. We just saw that K is relatively algebraically closed in M, and every element of valuation zero in M has a pth root, so every element of K with valuation zero is a pth power as well. This shows that the residue field is perfect.

- K has the same residue field as M, because otherwise there is some $a \in M$ whose residue is transcendental over the residue field of K. There is a unique way to extend the valued field structure from K to $K(a)$ such that this holds, by basic facts about 1-types in ACVF (namely: if K is any subfield of a model of ACVF, there is a *unique* 1-type over K whose realizations have valuation zero and transcendental residue). By saturation of N, we can find an element realizing this type over $f(K)$. (Just take any ultralimit of elements of $f(K)$ having distinct residues.)

- K has the same value group as M. Otherwise, take $c \in M$ having a new valuation. This valuation is $\mathbb{Q}$-linearly independent from the value group of K, by purity of the value group. By saturation, we can find c' whose valuation realizes the same T_G-type over the value group of $f(K)$ that c does over the value group of K. Then as usual, $K(c)$ and $f(K)(c')$ are isomorphic.

- K is M. Otherwise, take $a \in M \setminus K$. Then $K(a)/K$ is a transcendental immediate extension. Every K^{alg}-definable ball containing a is actually K-definable (or else $K(a) \cap K^{alg}$ would be strictly bigger than K, contradicting relative algebraic closedness). Moreover, it is explicitly K-definable, in the sense that its center can be taken to be in K. (This can be done by averaging, using the fact that K has no extensions of degree divisible by p, which in turn comes from the fact that K is perfect, $\text{Gal}(M)$ surjects onto $\text{Gal}(K)$ by relative algebraic closedness, and $\text{Gal}(M)$ has nothing to do with p.) As a consequence, we can find a descending sequence of K-definable balls with centers in K, having no intersection in K^{alg}. This defines a unique ACVF-type over K^{alg}, and it is finitely satisfiable in K (hence realized in N), because it is the type of a pseudolimit of the sequence of the centers of the balls. Consequently, we can find a' in N such that $K(a')$ and $K(a)$ are isomorphic as valued fields. Because the extension is immediate, they are also isomorphic as L-structures.

$\square$