AN AHLFORS ISLANDS THEOREM
FOR NON-ARCHIMEDEAN MEROMORPHIC FUNCTIONS

ROBERT L. BENEDETTO

ABSTRACT. We present a p-adic and non-archimedean version of Ahlfors' Five
Islands Theorem for meromorphic functions, extending an earlier theorem of
the author for holomorphic functions. In the non-archimedean setting, the
theorem requires only four islands, with explicit constants. We present exam-
ples to show that the constants are sharp and that other hypotheses of the
theorem cannot be removed.

1. INTRODUCTION

In the 1930s, Ahlfors proposed his theory of covering surfaces [2] in complex
analysis as an analogue of Nevanlinna theory for domains, rather than for points.
One of the key theorems in the subject is the Five Islands Theorem:

Theorem (Ahlfors' Complex Five Islands Theorem). *Let $U_1, \ldots, U_5$ be simply con-
nected domains in the Riemann sphere with mutually disjoint closures. Then there
is a constant $h = h(U_1, \ldots, U_5) > 0$ with the following property: Let f be a complex
meromorphic function on the disk $|z| < 1$, and suppose that there is some $r \in (0, 1)$
with*

$$(1.1) \qquad\qquad S_{\mathbb{C}}(f, r) \geq h \cdot L_{\mathbb{C}}(f, r).$$

*Then there is a simply connected domain U contained in the disk $|z| < r$ such that
f is one-to-one on U and $f(U) = U_i$ for some $1 \leq i \leq 5$.*

Similar results hold for holomorphic functions, with only three islands $U_i \subseteq
\mathbb{C}$ required. Here, $S_{\mathbb{C}}(f, r)$ and $L_{\mathbb{C}}(f, r)$ (the *mean covering number* and *relative
boundary length*, respectively) are certain real quantities describing the image of f
on the open disk $|z| < r$. By the work of Dufresnoy [17], condition (1.1) may be
replaced by a condition of the form $f^{\#}(0) > \tilde{h}$, where $f^{\#}$ is the spherical derivative
of f, and $\tilde{h}$ is, like h, a constant which depends only on the domains $U_1, \ldots, U_5$.
In 1998, Bergweiler [7] proved the Five Islands Theorem without the theory of
covering surfaces by using a lemma of Zalcman [33], some Nevanlinna theory, and
quasiconformal perturbations. See [22], Chapters 5–6, for more details on the theory
of covering surfaces.

Initially, the Five Islands Theorem was used mainly in complex function theory.
Then, in 1968, Baker [3] applied it to complex dynamics, proving that the Julia set

of a complex entire function is the closure of the set of repelling periodic points. (The usual well-known proofs of repelling density for rational functions do not extend to entire functions.)

In this paper, we will consider non-archimedean fields. Recall that a non-archimedean field is a field K equipped with a non-trivial absolute value $|\cdot|$ satisfying the ultrametric triangle inequality $|x + y| \leq \max\{|x|, |y|\}$ for all $x, y \in K$. We set the following notation.

K	a complete, algebraically closed non-archimedean field with absolute value $\lvert \cdot \rvert$
$\mathcal{O}$	the ring of integers $\{x \in K : \lvert x \rvert \leq 1\}$ of K
k	the residue field of K

For example, K could be $\mathbb{C}_p$, the completion of an algebraic closure of $\mathbb{Q}_p$, the field of p-adic rationals. Alternately, K could be the completion of an algebraic closure of $\mathbb{F}((t))$, the field of formal Laurent series with coefficients in an arbitrary field $\mathbb{F}$. We refer the reader to [19, 30] for treatises on non-archimedean analysis.

Recall that the residue field k is defined to be $\mathcal{O}/\mathcal{M}$, where $\mathcal{M}$ is the maximal ideal $\{x \in K : |x| < 1\}$ of $\mathcal{O}$. If the original field K has positive characteristic, then $\operatorname{char} k = \operatorname{char} K$; however, if $\operatorname{char} K = 0$, then $\operatorname{char} k$ could be either zero or any prime $p > 0$. The equal characteristic case occurs for function fields; the different, or mixed, characteristic case occurs for $K = \mathbb{C}_p$.

There have been numerous studies in recent decades of non-archimedean versions of Nevanlinna theory. In 1971, Adams and Straus [1] proved some non-archimedean Nevanlinna-style results using methods much simpler than a full Nevanlinna theory. More recently, a number of authors have developed a broader non-archimedean Nevanlinna theory, including analogues of the First and Second Main Theorems; see [14] or [24] for expositions, and [12, 13, 15, 16, 25, 31] for some of the original papers.

At the same time, there has also been a growing interest in the dynamics of non-archimedean rational and entire functions; see, for example, [4, 5, 27, 28] and Section 7 of [32], or [10] for an expository introduction to the p-adic case. Although many of the fundamental results of complex dynamics have analogues in the non-archimedean setting, the question of non-archimedean repelling density remains open, even in the case of rational functions, at least in the case of positive residue characteristic. There have been some partial results: Hsia [23] has shown that the Julia set of a rational function is contained in the closure of *all* periodic points, and Bézivin [9] has proven repelling density if there is at least one repelling periodic point. However, as discussed in [6], there are serious obstacles to extending either result to prove repelling density completely.

Motivated by Baker's complex result on repelling density and the more recent developments in non-archimedean Nevanlinna theory, the author presented a non-archimedean version of Ahlfors' Islands Theorem for holomorphic functions in [6]. In that case, only two islands, rather than three, were required. However, an extra hypothesis was also needed, essentially stating that the analogue of $L_{\mathbb{C}}(f, r)$ is at some point larger than a constant which depends on the two islands. In this paper, we continue those investigations by presenting an analogue of the Islands Theorem for meromorphic functions in Theorem 7.2. We envisage that these non-archimedean islands theorems should be part of a non-archimedean theory of

Ahlfors' covering surfaces that is yet to be developed. We give an abbreviated statement of our result here.

Main Theorem (Non-archimedean Meromorphic Four Islands Theorem). *Let U_1, $U_2, U_3, U_4 \subseteq K \cup \{\infty\}$ be four disjoint open disks. Let ν_1 be a Berkovich point such that no residue class of ν_1 intersects more than two of U_1, U_2, U_3, U_4. Then there are real constants $C_1, C_2 \geq 0$ depending only on K and on U_1, U_2, U_3, U_4 with the following property.*

Let f be a meromorphic function on $D(0,1) = \{z \in K : |z| < 1\}$ such that $f^{\#}(0) > C_1$ and, for any point $\nu \in \mathcal{D}(0,1)$ in the open Berkovich disk such that $f_*(\nu) = \nu_1$, we have $\delta(f, \nu) \geq C_2$. Then there is an open disk $U \subseteq D(0,1)$ such that f is one-to-one on U and $f(U) = U_i$ for some $i = 1, 2, 3, 4$.

Here, $f^{\#}$ is a non-archimedean version of the spherical derivative (see Definition 3.2 and equation (4.3)), and $\delta(f, \nu)$ is a measure of distortion related to the non-archimedean relative boundary length (see equations (6.1) and (6.2)). The notation $\mathcal{D}(0,1)$, $\mathcal{P}^1(K)$, and $f_*(\nu)$ concerns the theory of Berkovich spaces and will be defined in Section 4. The full statement of the result, Theorem 7.2, includes precise formulas for the constants C_1 and C_2 (unlike Ahlfors' results, which did not explicitly describe h).

In Section 2, we will recall some basic facts about the non-archimedean projective line $\mathbb{P}^1(K)$. In Section 3 we will review some standard results about meromorphic functions on non-archimedean disks. Section 4 is a summary (without proofs) of the necessary elements of the Berkovich theory. Then, in Section 5, we will present a few relevant lemmas on the Berkovich projective line. In Section 6, we will study some particular functions from Berkovich spaces to $\mathbb{R}$, including the quantity $\delta(f, \nu)$. Section 7 is devoted to the statement and proof of the main theorem and a corollary. Finally, we will present some examples and address the sharpness of Theorem 7.2 in Section 8.

2. Non-archimedean disks and the projective line

For $a \in K$ and $r > 0$, we will denote by $D(a, r)$ and $\overline{D}(a, r)$ the open disk and closed disk (respectively) of radius r about a in K. If $r \in |K^{\times}|$, then $D(a, r) \subsetneq \overline{D}(a, r)$, whereas the two sets coincide if $r \notin |K^{\times}|$. All disks in K are both open and closed as topological sets, but we keep the labels "open disk" and "closed disk" because the two can behave differently under the action of holomorphic and meromorphic functions.

By ultrametricity, any point of a disk is a center; but because K is algebraically closed, the radius is well defined and is equal to the diameter. That is, $D(a, r) = D(b, s)$ if and only if $r = s$ and $b \in D(a, r)$; the analogous statement also holds for closed disks.

Let $\mathbb{P}^1(K)$ denote the projective line over K, with points represented in homogeneous coordinates by $[x, y]$, for $(x, y) \in K \times K \setminus \{(0, 0)\}$. We will usually identify $\mathbb{P}^1(K)$ with $K \cup \{\infty\}$ by taking $[x, y]$ to $z = x/y$, with $[1, 0]$ corresponding to $z = \infty$.

The metric on K induces a standard spherical metric on $\mathbb{P}^1(K)$, given by

$$\Delta(P_1, P_2) = \frac{|x_1 y_2 - x_2 y_1|}{\max\{|x_1|, |y_1|\} \max\{|x_2|, |y_2|\}},$$

where $P_i = [x_i, y_i]$. Clearly $0 \le \Delta(P_1, P_2) \le 1$. In affine coordinates,

$$\Delta(z_1, z_2) = \frac{|z_1 - z_2|}{\max\{1, |z_1|\} \max\{1, |z_2|\}}.$$

Note that for $z_1, z_2 \in \mathcal{O}$, we have $\Delta(z_1, z_2) = |z_1 - z_2|$. The topology on K induced by Δ is exactly the same as that induced by $| \cdot |$.

The group $\mathrm{PGL}(2, K)$ acts by linear fractional transformations on $\mathbb{P}^1(K)$. As on the Riemann sphere, given any six points $P_1, P_2, P_3, Q_1, Q_2, Q_3 \in \mathbb{P}^1(K)$, there is a unique $\eta \in \mathrm{PGL}(2, K)$ such that $\eta(P_i) = Q_i$ for all $i = 1, 2, 3$. Of course, this map η need not preserve distances. On the other hand, the subgroup $\mathrm{PGL}(2, \mathcal{O})$ of transformations $z \mapsto (az + b)/(cz + d)$ with $a, b, c, d \in \mathcal{O}$ and $|ad - bc| = 1$ is distance-preserving with respect to Δ. (See [6], Section 1, for example.)

A disk $V \subseteq \mathbb{P}^1(K)$ is defined to be either a disk in K in the usual sense or the complement in $\mathbb{P}^1(K)$ of a disk in K. Equivalently, V is a disk in $\mathbb{P}^1(K)$ if and only if it is the image under some $\eta \in \mathrm{PGL}(2, K)$ of a disk in K. (This is *not* the same as saying that V is a disk with respect to the spherical metric Δ.) We say V is open (respectively, closed) if it is either an open (respectively, closed) disk in K or the complement of a closed (respectively, open) disk in K.

3. Holomorphic and meromorphic functions on disks

We now summarize the relevant portions of the well-known theory of homomorphic and meromorphic functions on disks. Non-archimedean analysis on more general subsets of K has been described in a number of contexts; see, for example, [11, 19, 20, 21, 29], and the more accessible exposition in simpler cases in Chapter VI of [30]. For an abbreviated survey specific to disks, with most proofs included, see [6, Section 2].

Definition 3.1. Let $U \subseteq K$ be a disk.

a. Let $a \in U$ and let $g : U \to K$. We say g is *holomorphic* on U if we can write

$$g(z) = \sum_{i=0}^{\infty} c_i (z - a)^i \in K[[z - a]]$$

as a power series that converges to $g(z)$ for all $z \in U$.

b. Let $f : U \to \mathbb{P}^1(K)$. We say f is *meromorphic on* U if f is continuous on U, and if we can write f in homogeneous coordinates as

$$f(z) = [g(z), h(z)]$$

for all z in some dense subset of U, where g and h are holomorphic on U.

Thus, a holomorphic function is not just locally analytic but rigid analytic, in that its defining power series converges on the whole disk. Holomorphicity is well defined, in the sense that if $a, b \in U$, and if g can be written as a convergent power series centered at a, then g can also be written as a convergent power series centered at b. Naturally, any holomorphic function f is also meromorphic, by choosing $g = f$ and $h = 1$. Conversely, any meromorphic function which never takes on the value ∞ is in fact holomorphic.

Intuitively, a meromorphic function is simply the quotient of two holomorphic functions, as in complex analysis. The technical "dense subset" condition in Definition 3.1 is required only because the holomorphic functions g and h may have common zeros. As observed in [26], it may not be possible to choose g and h to

remove all common zeros if U is an open disk. Fortunately, this technicality will not affect us, because we will not be concerned with any specific representation g/h of a given meromorphic function f.

Derivatives of holomorphic and meromorphic functions are defined in the same way as in real and complex analysis, and they satisfy all the usual algebraic rules. In particular, if $f(z)$ is holomorphic or meromorphic on a disk U, then so is $f'(z)$.

If f is holomorphic on a disk U, and if $U' \subsetneq U$ is a smaller disk, then $f(U')$ is a disk. Moreover, $f(U')$ is open (respectively, closed) if and only if U' is open (respectively, closed); see, for example, [6, Lemma 2.2]. In addition, any two points in $f(U')$ have the same number of preimages in U', counting multiplicity.

Similarly, if f is meromorphic on a disk $U \subseteq K$, and if $U' \subsetneq U$ is a strictly smaller disk, then $f(U')$ is either all of $\mathbb{P}^1(K)$ or else a disk V in $\mathbb{P}^1(K)$; in the latter case, V is open (respectively, closed) if and only if U' is. In particular, in light of the Riemann mapping theorem, the appropriate analogue of a simply connected domain in the Riemann sphere is a disk in $\mathbb{P}^1(K)$.

Definition 3.2. Let $U \subseteq K$ be a disk, and let $f : U \to \mathbb{P}^1(K)$ be a meromorphic function. Let $a \in U$. If $f(a) \neq \infty$, the *spherical derivative of f at a* is

$$f^\#(a) = \frac{|f'(a)|}{\max\{1, |f(a)|^2\}}.$$

If $f(a) = \infty$, then we set $f^\#(a) = (1/f)^\#(a)$.

Note that $f^\#$ takes values in $[0, \infty) \subseteq \mathbb{R}$, not in K. The reader may verify that

$$f^\#(a) = \lim_{z \to a} \frac{\Delta(f(z), f(a))}{|z - a|}.$$

In particular, if $\eta \in \mathrm{PGL}(2, \mathcal{O})$ and f is meromorphic, then $(\eta \circ f)^\# = f^\#$, because η preserves Δ.

4. The Berkovich disk and projective line

Our discussion of meromorphic functions on $D(0, 1)$ will involve their action on larger spaces defined by Berkovich. We refer the reader to his papers, especially [8], for background on general Berkovich spaces and for proofs of most of their basic properties. For our purposes, the reader may find the exposition in [32] more useful, as it is specific to the case of disks and the projective line, which are all we need here. The same space for the projective line was independently discovered later by Rivera-Letelier [27, 28]; the set we will call $\mathcal{P}^1(K)$ is called $\mathbb{H} \cup \mathbb{P}^1(K)$ in his notation. Related seminorm constructions also appeared in earlier papers, such as [19], but Berkovich's more complete theory seems to be the proper language for a meromorphic islands theorem.

For $a \in K$ and $r \in |K^\times|$, the Berkovich disk $\overline{\mathcal{D}}(a, r)$ associated with $\overline{D}(a, r)$ is a topological space into which $\overline{D}(a, r)$ embeds naturally, but with extra points added in such a way that the resulting space is path-connected, compact, and Hausdorff. Most (or at least, the most useful) of these extra points correspond in a natural way to disks contained in $\overline{D}(a, r)$.

More specifically, let $\mathcal{A}$ be the ring of all holomorphic functions on $\overline{D}(a, r)$, and for any disk $\overline{D}(b, s) \subseteq \overline{D}(a, r)$, define the norm $\nu(b, s) = \| \cdot \|_{\nu(b,s)}$ on $\mathcal{A}$ by

(4.1) $$\|f\|_{\nu(b,s)} = \sup\{|f(x)| : x \in \overline{D}(b, s)\} = \max\{|c_i|s^i : i \geq 0\},$$

where $f(z) = \sum_{i=0}^{\infty} c_i(z-b)^i$. If $s \in |K^\times|$, then $\|f\|_{\nu(b,s)}$ can be considered the generic value of $|f(x)|$ on $\overline{D}(b,s)$, in the sense that most $x \in \overline{D}(b,s)$ (i.e., all but those in finitely many open subdisks $D(b',s)$) satisfy $|f(x)| = \|f\|_{\nu(b,s)}$. Note that $\|f \cdot g\|_{\nu(b,s)} = \|f\|_{\nu(b,s)} \cdot \|g\|_{\nu(b,s)}$.

Rigorously speaking, the Berkovich disk $\overline{\mathcal{D}}(a,r)$ is defined to be the set of all those multiplicative seminorms on $\mathcal{A}$ which are bounded with respect to $\nu(a,r)$. The natural topology is the Gel'fand topology, which is the weakest topology such that for every $f \in \mathcal{A}$, the map $\nu \mapsto \|f\|_\nu$ from $\overline{\mathcal{D}}(a,r)$ to $\mathbb{R}$ is continuous.

Berkovich classified the points of $\overline{\mathcal{D}}(a,r)$ into four types. *Type I* are those corresponding to points $x \in \overline{D}(a,r)$, with seminorm given by $\|f\|_x = |f(x)|$. The norms $\nu(b,s)$ described in equation (4.1) are either *type II* or *type III*, if s is in $|K^\times|$ or not, respectively. Last, there are also *type IV* points corresponding to infinite decreasing sequences of nested disks with empty intersection; however, such points can be safely ignored for our purposes. Indeed, removing the type IV points from $\overline{\mathcal{D}}(a,r)$ leaves a space which is still path-connected and Hausdorff, though no longer compact.

From the type II and III points, one can begin to see how $\overline{\mathcal{D}}(a,r)$ is path-connected. For example, if $x \in K$ and $r, R > 0$ are real numbers such that $0 < r < |x| < R$, we see the arrangement of the points 0 and x and the disks $\overline{D}(x,r)$ and $\overline{D}(0,R)$ in Figure 1(a). Bearing this picture in mind, we can see the path in Berkovich space from 0 to $\nu(x,r)$ as follows. Start at the type I point 0, which we can temporarily think of as a disk of radius $s = 0$. Increase the radius s through a path of type II and III points of the form $\nu(0,s)$ until we get to radius $s = R$. Then $\nu(0,R) = \nu(x,R)$ (since $\overline{D}(0,R) = \overline{D}(x,R)$), and so we may now consider disks $\overline{D}(x,s)$ centered at s. Decrease the radius s towards the new center x until we arrive at $\nu(x,r)$, corresponding to the disk $\overline{D}(x,r)$.

Intuitively, the space $\overline{\mathcal{D}}(a,r)$ looks like a tree branching out from the root point $\nu(a,r)$ with infinitely many branches at every type II point (which are dense in the tree), and with limbs ending at the type I and type IV points. See Figure 2(a), which shows a few of the (type II) points of branching, with a few of the branches at each such point; all of the endpoints at the bottom are type I points. The
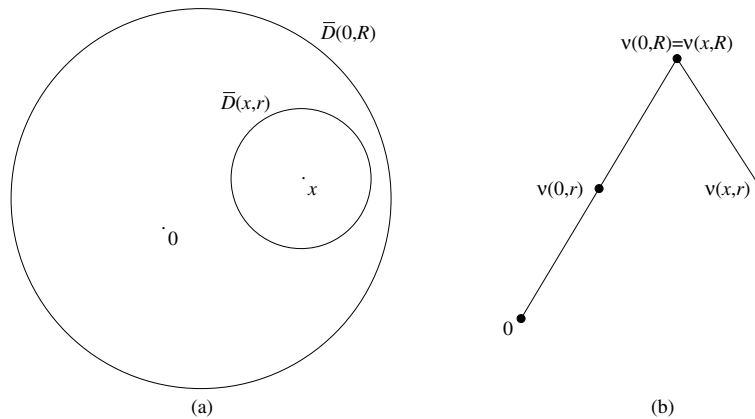

FIGURE 1. Path from 0 to $\nu(x,r)$ in $\overline{\mathcal{D}}(0,R)$.

infinitely many branches at a type II point $\nu(b,s)$ correspond to the infinitely many open subdisks $D(c,s)$ of $\overline{D}(b,s)$ of the same radius, as well as (if $s < r$) one more branch corresponding to increasing the radius. The type III points, meanwhile, are interior points with no branching. The topology is slightly weaker than the tree image might suggest at first; for example, the type I points (at the tips of the branches) are dense in the full space.

If f is meromorphic on $\overline{D}(a,r)$ and $\nu \in \overline{\mathcal{D}}(a,r)$, then we may define $\|f\|_\nu$ to be $\|g\|_\nu / \|h\|_\nu$, where $f = g/h$ for $g, h \in \mathcal{A}$. (Note that $\|f\|_\nu = \infty$ if and only if $\nu = b$ is type I and f has a pole at b.) As before, it is appropriate to think of $\|f\|_{\nu(a,r)}$ as the generic value of $|f(x)|$ for $x \in \overline{D}(a,r)$.

For an open disk $D(a,r)$, we can also associate a Berkovich space $\mathcal{D}(a,r)$ by taking the union (really, the direct limit) of sets $\overline{\mathcal{D}}(a,r_i)$, where $r_i \nearrow r$. The resulting space is still path-connected, Hausdorff, and locally compact, but it is no longer compact. Although $\mathcal{D}(0,1)$ will be one of our main objects of study, we will understand it by considering the subspaces $\overline{\mathcal{D}}(a,r)$ described above, for $a \in D(0,1)$ and $0 < r < 1$.

We may also define the Berkovich projective line $\mathcal{P}^1(K)$ by glueing two copies of $\mathcal{D}(0,r)$ (for some $r > 1$) as follows. A type I point x on one copy with $1/r < |x| < r$ is identified with $1/x$ on the other copy. Meanwhile, a type II or III point $\nu(b,s)$ with $1/r < |b| < r$ is identified with $\nu(1/b, s/|b|^2)$, since $\overline{D}(1/b, s/|b|^2)$ is the image of $\overline{D}(b,s)$ under $z \mapsto 1/z$. The type IV points are glued similarly.

Thus, $\mathcal{P}^1(K)$ looks like $\overline{\mathcal{D}}(0,1)$ with an extra copy of the open tree $\mathcal{D}(0,1)$ attached to the top (i.e., the ∞ end) of the point $\nu(0,1)$. The new top portion contains all points x of $\mathcal{P}^1(K)$ with $|x| > 1$, including ∞, as well as points $\nu(a,r)$ with $|a| > 1$ or $r > 1$. Like $\overline{\mathcal{D}}(0,1)$, the space $\mathcal{P}^1(K)$ is path-connected, Hausdorff, and compact. For a rough idea of the space, see Figure 2(b), which highlights the type I points $0, 1, \infty, \alpha$, for some $\alpha \in K$ with $0 < |\alpha| < 1$, as well as numerous type II points, including $\nu(0,1)$ and $\nu(0,|\alpha|)$. As is the case with $\overline{\mathcal{D}}(a,r)$, the space looks like a tree with infinite branching at a dense set of points along every edge.

Any disk in $\mathcal{P}^1(K)$ is associated with a unique point (of type II or III) of $\mathcal{P}^1(K)$. Indeed, any open or closed disk $D(a,r)$ or $\overline{D}(a,r)$ or its complement is associated with the point $\nu(a,r)$. Conversely, a type III point $\nu(a,r)$ is associated with exactly
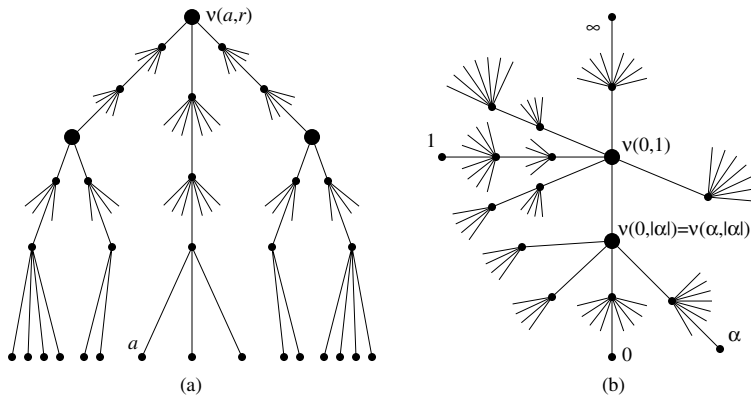

FIGURE 2. The Berkovich projective line.

two disks, namely $D(a, r) = \overline{D}(a, r)$ and its complement. Meanwhile, a type II point $\nu(a, r)$ is associated with infinitely many disks: every open disk $D(b, r)$ for $b \in \overline{D}(a, r)$, the open disk $\mathbb{P}^1(K) \setminus \overline{D}(a, r)$, and the (closed) complements of all these open disks. Note, however, that there is a one-to-one correspondence between type II and III points of $\mathcal{P}^1(K)$ and closed disks in K; the point $\nu(a, r)$ corresponds to $\overline{D}(a, r)$.

Viewed another way, a type II point $\nu(a, r)$ induces a partition of $\mathbb{P}^1(K)$ into the infinitely many open disks $D(b, r)$ contained in $\overline{D}(a, r)$, along with $\mathbb{P}^1(K) \setminus \overline{D}(a, r)$. Similarly, a type III point $\nu(a, r)$ partitions $\mathbb{P}^1(K)$ into two disks: $\overline{D}(a, r)$ and $\mathbb{P}^1(K) \setminus \overline{D}(a, r)$. In each case, we call the disks of the partition the *residue classes* of $\nu(a, r)$.

A point $\nu_1 \in \mathcal{P}^1(K)$ is type II or III if and only if $\mathcal{P}^1(K) \setminus \{\nu_1\}$ is not connected. In that case, the intersections of the various components of $\mathcal{P}^1(K) \setminus \{\nu_1\}$ with $\mathbb{P}^1(K)$ (i.e., the type I points) give precisely the residue classes of ν_1 described above.

A meromorphic function f on $\overline{D}(a, r)$ induces a continuous function $f_* : \overline{D}(a, r) \to \mathcal{P}^1(K)$. A fully rigorous description of f_* and its properties requires defining general Berkovich spaces as locally ringed spaces with patches given by general Berkovich affinoids. We refer the reader to [8] or to [32, Section 7] for the details; an equally rigorous description, specific to $\mathcal{P}^1(K)$ and in a different style, appears in [28, Section 4]. We will now describe f_* precisely, but we will skip the proofs.

Given f meromorphic on $\overline{D}(a, r)$ and $\nu \in \overline{\mathcal{D}}(a, r)$ (or given f meromorphic on $D(a, r)$ and $\nu \in \mathcal{D}(a, r)$), the image point $f_*(\nu) = \|\cdot\|_{f_*(\nu)}$ is the unique point (i.e., seminorm) in $\mathcal{P}^1(K)$ such that

(4.2) $$\|h\|_{f_*(\nu)} = \|f \circ h\|_\nu$$

for all $h \in K(z)$. (Recall that we allow $\|h\|_\nu$ to take on the value ∞ if ν is a type I point and h has a pole at ν.)

Unfortunately, equation (4.2) does not give much immediate insight into what f_* really looks like. Following [28], then, we present the following equivalent description.

First, if $\nu = x$ is a type I point, then $f_*(x)$ is simply $f(x)$.

Second, if $\nu = \nu(b, s)$ is a point of type II or III and f is holomorphic, we know from Section 3 that $f(D(b, s))$ is an open disk $D(f(b), \sigma)$. Then $f_*(\nu(b, s))$ is simply $\nu(f(b), \sigma)$. More generally, if f is meromorphic, it can be shown that there is a radius $s_0 < s$ such that for all s' with $s_0 < s' < s$, the image $f(V(s'))$ of the annulus $V(s') = \{x \in K : s' < |z - b| < s\}$ is itself an annulus of the form

$$\{y \in K : \sigma' < |z - \beta| < \sigma\} \quad \text{or} \quad \{y \in K : \sigma < |z - \beta| < \sigma'\},$$

where σ' varies with s' but β and σ are independent of s'. Then it turns out that

$$f_*(\nu(b, s)) = \nu(\beta, \sigma).$$

It follows quickly from equation (4.2) that f_* is the unique continuous extension of f to Berkovich space. In particular, if $f(\overline{D}(a, r)) = \overline{D}(b, s)$, then $f_*(\overline{\mathcal{D}}(a, r)) = \overline{\mathcal{D}}(b, s)$, with the analogous statement also holding for open disks or if the image is $\mathbb{P}^1(K)$. Moreover, if f is a non-constant meromorphic function, then f_* takes type I points to type I points, type II points to type II points, and so on.

Given f meromorphic on $\overline{D}(a, r)$, we can extend $f^{\#}$ from the type I points to all of $\overline{\mathcal{D}}(a, r)$ by setting

$$(4.3) \qquad f^{\#}(\nu) = \frac{\|f'\|_{\nu}}{\max\{1, \|f\|_{\nu}^2\}}.$$

Once again, it is easy to show that if $\eta \in \mathrm{PGL}(2, \mathcal{O})$, then $(\eta \circ f)^{\#}(\nu) = f^{\#}(\nu)$.

Definition 4.1. Given a point $\nu(a, \rho) \in \mathcal{P}^1(K)$ of type II or III, define

$$(4.4) \qquad r(\nu(a, \rho)) = \frac{\rho}{\max\{1, \rho^2, |a|^2\}}$$

to be the *spherical radius* of $\nu(a, \rho)$.

In fact, $r(\nu)$ is simply the diameter, with respect to the spherical metric Δ, of the smallest-diameter residue class of ν. For example, if $\overline{D}(a, \rho) \subseteq \overline{D}(0, 1)$, then $r(\nu(a, \rho)) = \rho$ is the usual radius of the associated disk. We also note that $r(\nu)$ is just $\exp[-\mathrm{dist}(\nu, \nu(0, 1))]$, where $\mathrm{dist}(\cdot, \cdot)$ is the metric on $\mathcal{P}^1(K) \setminus \mathbb{P}^1(K)$ which appears in [28, Section 3] as the metric on $\mathbb{H}$ and in [32, Section 2] as the "big model" metric. It is easy to verify that formula (4.4) is independent of the choice of a in $\overline{D}(a, \rho)$, and that $r \circ \eta_* = r$ for any $\eta \in \mathrm{PGL}(2, \mathcal{O})$.

The reader should be cautioned that r is *not* a continuous function with respect to the Gel'fand topology. However, the restriction of r to a path like the one described in Figure 1 is continuous.

5. Lemmas on the Berkovich projective line

The following lemmas concern the action of a meromorphic function on Berkovich points and will be needed to prove Theorem 7.2. They are implicit in the standard references, but we state them explicitly for the convenience of the reader.

Lemma 5.1. *Let f be meromorphic and non-constant on $D(0, 1)$, let $x \in D(0, 1)$, let $0 < r < 1$, and let $\nu_1 \in \mathcal{P}^1(K)$. Then there are only finitely many points $\nu \in \mathcal{D}(x, r)$ such that $f_*(\nu) = \nu_1$.*

Proof. This result follows easily from the machinery of [28, Section 4], but we provide a direct proof without extra machinery for the convenience of the reader.

Write $f = g/h$ for g and h holomorphic. If ν_1 is of type I, then by a change of coordinates, we may assume $\nu_1 = 0$. The conclusion is then well known by the finiteness of the Weierstrass degree of g on $D(x, r) \subsetneq D(0, 1)$; see, for example, [6, Lemma 2.2].

If ν_1 is of type II or III, then let $a = f(x)$. There must be a point $b \in \mathbb{P}^1(K)$ such that a and b lie in different residue classes of ν_1. By a change of coordinates, we may assume $a = 0$ and $b = \infty$, so that $\nu_1 = \nu(0, R)$ for some $R > 0$. By the previous paragraph, there are only finitely many zeros and poles of f in $D(x, r)$.

If $f_*(\nu(y, s)) = \nu_1$ for some $\overline{D}(y, s) \subseteq D(x, r)$, then because $f(\overline{D}(y, s))$ is either $\mathbb{P}^1(K)$ or a closed disk associated with ν_1, there must be a root of $f = 0$ or $f = \infty$ in $\overline{D}(y, s)$. Thus, any set of disjoint disks $\overline{D}(y_i, s_i)$ with $f_*(\nu(y_i, s_i)) = \nu_1$ must be finite, since there are only finitely many poles or zeros of f in $D(x, r)$, by the previous paragraph.

Meanwhile, if $y \in D(x, r)$ and $0 < s_1 < s_2 < r$ with $f_*(\nu(y, s_i)) = \nu_1$ for $i = 1, 2$, we claim that there must be a root y' of $f = 0$ or $f = \infty$ with $s_1 < |y' - y| \le s_2$. If not, then move y to 0 and write the holomorphic functions g, h as $g(z) = \sum_{j=0}^{\infty} a_j z^j$

and $h(z) = \sum_{j=0}^{\infty} b_j z^j$. Our assumption concerning the lack of roots implies that one term $a_m z^m$ of g and one term $b_n z^n$ of h are uniquely maximal in each sum for all $s_1 < |z| \le s_2$. Thus, for all such z,

$$|f(z) - cz^\ell| < |f(z)|,$$

where $c = a_m/b_n \in K$ and $\ell = m - n \in \mathbb{Z}$. Since $f_*(\nu(y, s_1)) = f_*(\nu(y, s_2)) = \nu(0, R)$, we must have $|cz^\ell| = R$ for all such z, which implies that $\ell = 0$ and $|c| = R$. In that case, however, $|f(z) - c| < R$ for all such z, meaning in particular that $|f(z) - c| < R$ for all $|z| = s_2$. Then $f_*(\nu(0, s_2)) \ne \nu(0, R)$, which is a contradiction and proves our claim.

Thus, any chain $\overline{D}(y_1, s_1) \supsetneq \overline{D}(y_2, s_2) \supsetneq \cdots$ of disks with $f_*(\nu(y_i, s_i)) = \nu_1$ must be finite, or else there would be infinitely many poles or zeros of f in $D(x, r)$. Together with the above fact that only finitely many disjoint disks $\overline{D}(y_i, s_i)$ can have $f_*(\nu(y_i, s_i)) = \nu_1$, the desired finiteness follows.

We will not need to consider the case that ν_1 is type IV in this paper, and we leave the proof of that case to the reader. $\square$

Lemma 5.2. *Let f be meromorphic on $D(0, 1)$, let $\nu \in \mathcal{P}^1(K)$, and let $x, y \in D(0, 1)$ such that $f(x)$ and $f(y)$ are in different residue classes of ν. Then $\nu \in f_*(\overline{\mathcal{D}}(x, |x - y|))$.*

Proof. The connected set $\overline{\mathcal{D}}(x, |x - y|)$ contains both x and y, but $f(x)$ and $f(y)$ lie in different components of $\mathcal{P}^1(K) \setminus \{\nu\}$. The result then follows by the continuity of f_*. $\square$

Lemma 5.3. *Let f be meromorphic on $D(0, 1)$, let $D(x, r) \subseteq D(0, 1)$, and let $\nu \in \mathcal{P}^1(K)$. If $\nu \in f_*(\mathcal{D}(x, r))$, then there is some $r' < r$ such that $\nu \in f_*(\overline{\mathcal{D}}(x, r'))$.*

Proof. As defined in Section 4, we may write $\mathcal{D}(x, r)$ as a union of closed Berkovich disks $\overline{\mathcal{D}}(x, r_i)$ for some sequence of radii $r_i \nearrow r$. The lemma then follows immediately. $\square$

6. Functions on Berkovich spaces

Theorem 7.2 and its proof will rely heavily on certain functions from $\mathcal{D}(0, 1)$ to $[0, \infty]$. We have already seen three such functions: the (non-continuous) spherical radius function $r(\nu)$ in Definition 4.1; and for any meromorphic function f on $D(0, 1)$, the (continuous) maps $\nu \mapsto \|f\|_\nu$ and $\nu \mapsto f^\#(\nu)$. (Note that r takes values in $[0, 1]$, $f^\#$ takes values in $[0, \infty)$, and $\|f\|_\nu$ attains the value ∞ only at (type I) poles of f.) Before introducing certain useful hybrids of these functions, we recall some convenient properties of the map $\nu \mapsto \|f\|_\nu$.

Given $f : D(0, 1) \to \mathbb{P}^1(K)$ meromorphic, if we fix a point $a \in D(0, 1)$, it is common to define the *valuation polygon* function $P_{f,a}(t)$ associated with f by

$$P(\log r) = P_{f,a}(\log r) = -\log \|f\|_{\nu(a, r)},$$

which is a continuous function from $(-\infty, 0)$ to $\mathbb{R}$. In fact:

- P is piecewise linear, with integer slopes, and
- the slope n at any point $t_0 = \log r_0$ (or just to the right of t_0 if P has a corner at t_0) is the number of poles minus the number of zeros of f in $\overline{D}(a, r_0)$.

Put another way, $\|f\|_{\nu(a,r)}$, viewed as a function of r, is continuous and piecewise monomial, with exponent at (or just to the right of) r_0 equal to the number of zeros minus the number of poles of f in $\overline{D}(a, r_0)$.

The proof is standard; see, for example, Section 3 of [29], Section VI.1.6 and that which follows it in [30], or Lemmas 4.4 and 4.5 of [6].

In light of the parenthetical comment in the second property above, for the rest of the paper we will say that the slope (resp., exponent) of a piecewise linear (resp., piecewise monomial) function at a corner t_0 is the slope (resp., exponent) just to the right of t_0.

Note that since f has only finitely many zeros and poles in $\overline{D}(a, R)$ for any given $0 < R < 1$, there are only finitely many radii $r \in (0, R]$ at which the exponent of $\|f\|_{\nu(a,r)}$ can change. Equivalently, for any $T \in (-\infty, 0)$, $P_{f,a}$ has corners at only finitely many points $t \in (-\infty, T]$.

Also note that if we fix $a \in D(0, 1)$ and let $f(z) = z - a$, then $r(\nu(a, \rho)) = \rho = \|f\|_{\nu(a,\rho)}$ for all $0 < \rho < 1$. Thus, even though $\nu \mapsto r(\nu)$ is not continuous on $\mathcal{D}(0, 1)$, it has a (continuous) valuation polygon associated with any given center a.

We will need three more functions specific to our setting. First, recall that in the complex case, Ahlfors' relative boundary length function $L_{\mathbb{C}}(r) = L_{\mathbb{C}}(f, 0, r)$ is the length (in the spherical metric) of the image of the circle $|z| = r$ under f. That is,

$$L_{\mathbb{C}}(r) = r \int_0^{2\pi} f^{\#}(re^{i\theta}) \, d\theta$$

in the complex setting. Viewing the integral as an average value, and recalling that in our setting, $f^{\#}$ generically has value $f^{\#}(\nu(a, r))$ on the "circle" $|z - a| = r$, it is natural to define, for a meromorphic function f on $D(0, 1)$,

$$(6.1) \qquad L(\nu) = L(f, \nu) = r(\nu) \cdot f^{\#}(\nu) = \frac{r(\nu) \cdot \|f'\|_\nu}{\max\{1, \|f\|_\nu^2\}},$$

which is a (non-continuous) function from $\mathcal{D}(0, 1)$ to $[0, \infty)$. For $\eta \in \mathrm{PGL}(2, \mathcal{O})$, note that $L(f, \nu) = L(\eta \circ f, \nu)$, because $f^{\#} = (\eta \circ f)^{\#}$.

Second, for f meromorphic on $D(0, 1)$ and $\nu \in \mathcal{D}(0, 1)$ of type II or III, define

$$(6.2) \qquad \delta(\nu) = \delta(f, \nu) = \frac{L(f, \nu)}{r(f_*(\nu))}.$$

(If $r(f_*(\nu)) = 0$ for some ν of type II or III, then f is constant, and we set $\delta = 0$ in this case.) If we write $f_*(\nu) = \nu(b, \rho)$, then $\delta(\nu) = r(\nu)\|f'\|_\nu/\rho$. In particular, if f is holomorphic on $D(a, r)$, and if we write $\nu = \nu(a, r)$ and $b = f(a)$, then $\delta(f, \nu) = r\|f'\|_\nu/\|f - b\|_\nu$; indeed, δ appeared in this form in [6].

The ratio δ is a measure of distortion, in the following sense. Intuitively, $L(f, \nu)$ measures the expected spherical radius of $f_*(\nu)$ based on the generic value of $f^{\#}$ at ν. Because $f^{\#}$ may be smaller than the size of the image would suggest (e.g., z^p has small derivative pz^{p-1} in residue characteristic p), the actual spherical radius $r(f_*(\nu))$ of the image may be larger than L. Thus, the smaller the distortion ratio δ, the further the actual spherical radius is from that predicted by the derivative.

Clearly $\delta(\eta \circ f, \nu) = \delta(f, \nu)$ for $\eta \in \mathrm{PGL}(2, \mathcal{O})$, by the invariance of r and L under η. The reader may verify that, in fact, the same equality holds for any $\eta \in \mathrm{PGL}(2, K)$.

We also leave it to the reader to verify that if char $k = 0$ and f is a non-constant meromorphic function, then $\delta(f, \nu)$ is identically equal to 1.

Third, given a point $\alpha \in K$ with $\alpha \neq 0, 1$ and a meromorphic function f on $D(0,1)$, define a real-valued function G on the type II and III points of $\mathcal{D}(0,1)$ by

$$(6.3) \qquad G(\nu) = G(f, \alpha, \nu) = \frac{(r(\nu) \cdot \|f'\|_\nu)^2}{\|f\|_\nu \|f - \alpha\|_\nu \|f - 1\|_\nu}.$$

Like L, the functions δ and G are usually not continuous. However, for a fixed center a, they have associated valuation polygons with properties analogous to the polygon $P_{f,a}$ associated with $\|f\|_\nu$. We will be interested mainly in the valuation polygon associated with G, which will involve the following integer-valued quantities. Given a point $b \in \mathbb{P}^1(K)$, then for any closed disk $\overline{D}(a, r) \subseteq D(0,1)$, define

$$(6.4) \qquad \begin{aligned} N_b(a, r) &= \text{number of roots of } f = b \text{ in } \overline{D}(a, r), \text{ and} \\ N_{\mathrm{ram}}(a, r) &= \text{number of ramification points of } f \text{ in } \overline{D}(a, r), \end{aligned}$$

where in each case we count points with multiplicity. Note that N_{ram} counts with multiplicity all points at which $f' = 0$; but it also counts ramification points at all multiple poles. That is, if $f = g/h$ with g and h holomorphic, then N_{ram} counts (with multiplicity) the zeros of $g'h - h'g$, less twice the number of common zeros of g and h.

Note also that because $r < 1$, any non-constant meromorphic function on $D(0,1)$ has only finitely many zeros on $\overline{D}(a, r)$, and therefore $N_b(a, r)$ is finite (cf. Lemma 5.1). Similarly, if f' is not identically zero (i.e., if f is not constant and not of the form $f(z) = g(z^p)$, where $p = \mathrm{char}\, K$), then $N_{\mathrm{ram}}(a, r)$ is also finite. We can now state the key properties of the valuation polygon associated with G.

Lemma 6.1. *Let f be a meromorphic function on $D(0,1)$ such that f' is not identically zero. Let $\alpha \in K \setminus \{0, 1\}$, and define G, N_b, and N_{ram} as in (6.3) and (6.4). Fix $a \in D(0,1)$, and define $\gamma : (-\infty, 0) \to \mathbb{R}$ by*

$$\gamma(\log r) = -\log G(\nu(a, r)).$$

Then γ is continuous and piecewise linear with slope at $\log r$ equal to

$$N_0(a, r) + N_\alpha(a, r) + N_1(a, r) + N_\infty(a, r) - 2N_{\mathrm{ram}}(a, r) - 2.$$

Proof. If we write $f = g/h$ with g, h holomorphic on $D(0,1)$, then for $\nu = \nu(a, r)$ with $\overline{D}(a, r) \subseteq D(0,1)$, we have

$$G(\nu(a, r)) = \frac{(r \cdot \|g'h - h'g\|_\nu)^2}{\|g\|_\nu \|g - \alpha h\|_\nu \|g - h\|_\nu \|h\|_\nu}.$$

The result then follows immediately from the properties of valuation polygons. $\square$

Lemma 6.2. *Let f be a meromorphic function on $D(0,1)$, and let $\alpha \in \overline{D}(0,1) \setminus D(1,1)$, with $\alpha \neq 0$. Then for any $\nu \in \mathcal{D}(0,1)$ of type II or III,*

$$\delta(f, \nu) \leq 1 \qquad \text{and} \qquad G(f, \alpha, \nu) \leq 1.$$

Moreover, if no residue class of $f_*(\nu)$ contains more than two of $0, 1, \alpha, \infty$, then

$$\delta(f, \nu)^2 = G(f, \alpha, \nu).$$

Proof. As the desired inequalities are trivial if f is constant, we may assume that f is not constant, and therefore that $f_*(\nu)$ is also type II or III. Write $\nu_1 = f_*(\nu)$.

For a holomorphic function g on $\overline{D}(0,1)$, it is easy to verify that $r(\nu)\|g'\|_\nu \leq \|g\|_\nu$ (see, for example, Lemma 4.2 of [6]), as the integer coefficients introduced in

differentiating all have absolute value at most one. Thus, writing $f = g/h$ for g, h holomorphic on $D(0,1)$, we see that

$$(6.5)$$
$$r(\nu)\|f'\|_\nu = \frac{r(\nu)\|g'h - h'g\|_\nu}{\|h\|_\nu^2} \leq \frac{\|g\|_\nu}{\|h\|_\nu} \cdot \max\left\{\frac{r(\nu)\|g'\|_\nu}{\|g\|_\nu}, \frac{r(\nu)\|h'\|_\nu}{\|h\|_\nu}\right\} \leq \|f\|_\nu$$

(using the easily verified fact that $\|\cdot\|_\nu$ is ultrametric), so that the same inequality holds for meromorphic functions as well.

To prove the bound on δ, choose $\eta \in \mathrm{PGL}(2, \mathcal{O})$ such that $\eta_*(\nu_1) = \nu(0, \rho)$, for some $\rho > 0$. Replacing f by $\eta \circ f$ (which does not affect δ), we have $\|f\|_\nu = \rho$. Meanwhile, we may write $\delta(\nu) = r(\nu)\|f'\|_\nu/\rho$, and therefore $\delta(\nu) = r(\nu)\|f'\|_\nu/\|f\|_\nu \leq 1$, by inequality (6.5).

For the bound on G, first suppose that $\nu_1 \in \mathcal{D}(0, |\alpha|)$. Then

$$\|f\|_\nu < \|f - \alpha\|_\nu = |\alpha|, \quad \text{and} \quad \|f - 1\|_\nu = 1,$$

so that by inequality (6.5),

$$G(\nu) = \frac{r(\nu)^2\|f'\|_\nu^2}{|\alpha|\|f\|_\nu} < \frac{r(\nu)^2\|f'\|_\nu^2}{\|f\|_\nu^2} \leq 1.$$

If ν_1 lies in either $\mathcal{D}(\alpha, |\alpha|)$, $\mathcal{D}(1, 1)$, or $\mathcal{P}^1(K) \setminus \overline{\mathcal{D}}(0, 1)$, the verification is similar. If ν_1 does not lie in any of these four Berkovich disks, then $|\alpha| \leq \|f\|_\nu = \|f - \alpha\|_\nu \leq 1$ and $\|f - 1\|_\nu = 1$, so that

$$G(\nu) = \frac{r(\nu)^2\|f'\|_\nu^2}{\|f\|_\nu^2} \leq 1.$$

For the final statement of the lemma, see Figure 3, which shows the only two possible ways the points in question could be arranged, according to the hypotheses. Figure 3(a) shows the case that $|\alpha| < 1$, so that $\nu_1 = \nu(0, \rho)$ for some $|\alpha| \leq \rho \leq 1$. Figure 3(b) shows the case that $|\alpha| = 1$, so that (by hypothesis) $|\alpha| = |\alpha - 1| = 1$, and we must have $\nu_1 = \nu(0, 1)$. In either case, we have

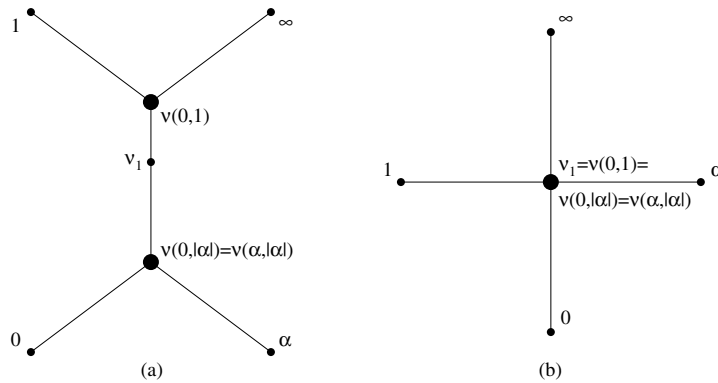$$|\alpha| \leq \|f\|_\nu = \|f - \alpha\|_\nu = r(\nu_1) \leq 1 = \|f - 1\|_\nu,$$

FIGURE 3. The two possible arrangements of $0, 1, \alpha, \infty$.

so that

$$\delta(\nu)^2 = \frac{r(\nu)^2 \|f'\|_\nu^2}{\|f\|_\nu^2} = G(\nu). \qquad \square$$

Lemma 6.3. *Let f be a meromorphic function on $D(0,1)$, and let $\alpha \in \overline{D}(0,1) \setminus \{0,1\}$. Fix $x \in D(0,1)$ and a radius $0 < R < 1$, and suppose that*

$$N_0(x,R) + N_\alpha(x,R) + N_1(x,R) + N_\infty(x,R) > 2N_{\mathrm{ram}}(x,R),$$

for N_b and N_{ram} as in (6.4). Then there is a point $y \in \overline{D}(x,R)$ such that for every $r \in (0,R]$,

$$N_0(y,r) + N_\alpha(y,r) + N_1(y,r) + N_\infty(y,r) > 2N_{\mathrm{ram}}(y,r).$$

Proof. Define $N_{\mathrm{tot}}(a,r) = N_0(a,r) + N_\alpha(a,r) + N_1(a,r) + N_\infty(a,r)$, and let $\{y_i\}_{i=1}^m$ be the finitely many roots of $f = 0, \alpha, 1, \infty$ in $\overline{D}(x,R)$.

If the desired conclusion fails, then for each $i = 1, \ldots, m$, there is some $r_i \in (0, R]$ such that $N_{\mathrm{tot}}(y_i, r_i) \leq 2N_{\mathrm{ram}}(y_i, r_i)$. If any two of the disks $\overline{D}(y_i, r_i)$ intersect, then one contains the other, and so we may discard the smaller one. We are left with a finite set $\{\overline{D}(y_i', r_i')\}_{i=1}^\ell$ of pairwise disjoint disks in $\overline{D}(x,R)$, each of which satisfies $N_{\mathrm{tot}}(y_i', r_i') \leq 2N_{\mathrm{ram}}(y_i', r_i')$, and which together contain all of the $\{y_i\}$. Thus,

$$N_{\mathrm{tot}}(x,R) = \sum_{i=1}^\ell N_{\mathrm{tot}}(y_i', r_i') \leq 2\sum_{i=1}^\ell N_{\mathrm{ram}}(y_i', r_i') \leq 2N_{\mathrm{ram}}(x,R),$$

contradicting the hypotheses and hence proving the lemma. $\qquad \square$

Note that by choosing r sufficiently small in the conclusion of Lemma 6.3, the point y must satisfy $f(y) \in \{0, 1, \alpha, \infty\}$. Furthermore, f is one-to-one on some small disk containing y.

7. The four islands theorem

We need the following specialized radius to define the key value μ, which will appear in Theorem 7.2.

Definition 7.1. Let $U_1, U_2, U_3, U_4 \subseteq \mathbb{P}^1(K)$ be four disjoint open disks in $\mathbb{P}^1(K)$, and for each $i = 1, 2, 3, 4$, let V_i be the smallest (closed) disk in $\mathbb{P}^1(K) \setminus U_i$ that contains all three of the other disks U_j. Choose a map $\eta_i \in \mathrm{PGL}(2, K)$ that sends V_i to $\overline{D}(0, r_i)$ and U_i to $\mathbb{P}^1(K) \setminus \overline{D}(0, R_i)$ for some radii $0 < r_i \leq R_i$, and let $s_i = r_i/R_i$. We define the *Ahlfors radius* of $\{U_1, U_2, U_3, U_4\}$ to be

$$s = \max\{s_1, s_2, s_3, s_4\}.$$

By "the smallest disk" V_i, we mean the intersection of all disks $V \subseteq \mathbb{P}^1(K) \setminus U_i$ that contain all three other disks U_j. We leave the verification of the following facts to the reader. First, this smallest disk V_i is indeed a closed disk. Second, a map η_i exists with the desired properties, and although neither η_i nor the resulting r_i and R_i are uniquely determined, the ratio s_i is well defined. (Indeed, $\log s_i$ is simply the modulus of the annulus $\mathbb{P}^1(K) \setminus (U_i \cup V_i)$, which, as is well known, is easily shown to be invariant under univalent meromorphic mappings.) Clearly $0 < s_i \leq 1$. Third, for any map $\eta \in \mathrm{PGL}(2, K)$, the Ahlfors radius of $\{\eta(U_1), \eta(U_2), \eta(U_3), \eta(U_4)\}$ is the same as that of $\{U_1, U_2, U_3, U_4\}$.

If we choose a point $a_i \in U_i$ for each $i = 1, 2, 3, 4$, then any three such points can be moved to $\{0, 1, \infty\}$ by some $\eta \in \mathrm{PGL}(2, K)$. The fourth then lands at some point $\alpha \in K \setminus \{0, 1\}$. Permuting $0, 1, \infty$ via combinations of $1/z$ and $1 - z$, we may assume that $\alpha \in \overline{D}(0, 1) \setminus D(1, 1)$ with $\alpha \neq 0$. (Note that this is the same condition that appeared in Lemma 6.2.) To say that the Ahlfors radius of the original disks is at most s is to say that the images of four disks under η are contained in $D(0, |\alpha|s)$, $D(\alpha, |\alpha|s)$, $D(1, s)$, and $\mathbb{P}^1(K) \setminus \overline{D}(0, 1/s)$.

We are now prepared to state our main result.

Theorem 7.2 (Non-archimedean Meromorphic Four Islands Theorem)*. Let K be a complete, algebraically closed non-archimedean field with residue field k, and let $p = \mathrm{char}\, k \geq 0$. If $p \geq 3$, define the real number*

$$E_p = \sum_{i=1}^{\infty} \frac{1}{p^i - 1}.$$

Let $U_1, U_2, U_3, U_4 \subseteq \mathbb{P}^1(K)$ be four pairwise disjoint open disks. Let $\nu_1 \in \mathcal{P}^1(K)$ such that no residue class of ν_1 intersects more than two of U_1, U_2, U_3, U_4.

Let s be the Ahlfors radius of $\{U_1, U_2, U_3, U_4\}$. Set

$$\mu = \begin{cases} 1 & \text{if } \mathrm{char}\, k = 0, \\ s^{1/2} & \text{if } \mathrm{char}\, k = 2, \\ s^{(\frac{1}{2} - \frac{1}{2p})} & \text{if } \mathrm{char}\, k = p \geq 3, \\ |p|^{-E_p} s^{1/2} & \text{if } \mathrm{char}\, k = p \geq 3 \text{ and } \mathrm{char}\, K = 0. \end{cases}$$

Let f be a meromorphic function on $D(0, 1)$ such that

(a). $f^{\#}(0) > 1/r(\nu_1)$*, and*

(b). *for any point $\nu \in \mathcal{D}(0, 1)$ such that $f_*(\nu) = \nu_1$, we have $\delta(f, \nu) \geq \mu$.*

Then there is an open disk $U \subseteq D(0, 1)$ such that f is one-to-one on U and $f(U) = U_i$ for some $i = 1, 2, 3, 4$.

Remarks. 1. Because the Ahlfors radius satisfies $0 < s \leq 1$, it follows that $0 < \mu \leq 1$.

2. In the $\mathrm{char}\, k = 0$ case, condition (b) becomes vacuous. Indeed, by the comments following equation (6.2), $\delta(\nu) = 1$ for all ν. In other words, only the condition on $f^{\#}$ is relevant in residue characteristic zero, and that condition depends only on the centers, but not the radii, of the four disks.

3. Outside the case $\mathrm{char}\, k = 0$, the statement of the theorem becomes stronger if μ is smaller, because more functions f will satisfy condition (b).

4. If $\mathrm{char}\, k = p \geq 3$ and $\mathrm{char}\, K = 0$, then two different values of μ are given above. Both are valid; but in practice, one should choose whichever is smaller, in light of Remark 7 above. The $s^{(1/2 - 1/(2p))}$ value is better for s close to 1, while the other value is better for s closer to 0.

5. The lower bound of $1/r(\nu_1)$ for $f^{\#}(0)$ in condition (a) is sharp, as we now observe. Choose $\lambda \in K$ with $|\lambda| \geq 1$. Let $f(z) = \lambda z$, let $\nu_1 = \nu(0, |\lambda|)$, and let each U_i be a disk of the form $D(a_i, \varepsilon)$, with $|a_i| = |\lambda|$ and $\varepsilon > 0$ as small as one wishes. Note that $f^{\#}(0) = C_1 = |\lambda|$, so that condition (a) just barely fails; however, all the other hypotheses of Theorem 7.2 hold. Nonetheless, the image $f(D(0, 1))$ fails to intersect any U_i, let alone map a subdomain onto one of them.

The sharpness of condition (b) is more subtle and will be considered in the examples of Section 8.

Theorem 7.2 will be a corollary of the following statement.

Theorem 7.3. *Let $0 < s \leq 1$, and let K, k, p, E_p, and μ be as in Theorem 7.2. Let $\alpha \in \overline{D}(0,1) \setminus D(1,1)$ with $\alpha \neq 0$, and set*

$$U_1 = D(0, |\alpha|s), \quad U_2 = D(\alpha, |\alpha|s), \quad U_3 = D(1, s), \quad U_4 = \mathbb{P}^1(K) \setminus \overline{D}(0, 1/s).$$

Let $\rho \in [|\alpha|, 1]$, and set $\nu_1 = \nu(0, \rho)$.
  Let f be a meromorphic function on $D(0,1)$. Suppose that

  (a). $f(0) \in \overline{D}(0, \rho) \setminus D(1,1)$,
  (b). $|f'(0)| > \rho$, *and*
  (c). $G(\nu) \geq \mu^2$ *for every* $\nu \in \mathcal{D}(0,1)$ *such that* $f_*(\nu) = \nu_1$, *where G is defined as in equation* (6.3).

Then there is an open disk $U \subseteq D(0,1)$ such that f is one-to-one on U and $f(U) = U_i$ for some $i = 1, 2, 3, 4$.

Proof. In our proof, we will proceed along the path in $\mathcal{D}(0,1)$ shown in Figure 1: we look at larger and larger disks about 0 until we come to a disk $\overline{D}(0, R)$ with certain nice properties. We choose an appropriate point $x \in \overline{D}(0, R)$, and then we look at smaller disks about x until we come to a disk $D(x, r)$ that maps one-to-one and onto one of the disks U_i. We break our proof into a number of steps.

(**i**). Using the equivalence of the two definitions of $\|f\|_\nu$ in (4.1), hypothesis (b) implies that there is some $\tilde{x} \in D(0,1)$ for which $|f(\tilde{x})| > \rho$. By hypothesis (a) and Lemma 5.2, there is some radius $0 < \tilde{r} < 1$ such that $\nu_1 \in f_*(\overline{\mathcal{D}}(0, \tilde{r}))$. By Lemma 5.1, we may assume that $\tilde{r}$ is the minimum such radius; hence, by hypothesis (a) and Lemma 5.2 again, $f(D(0, \tilde{r})) \subseteq \overline{D}(0, \rho) \setminus D(1,1)$.

(**ii**). Define the piecewise linear function $\gamma : (-\infty, 0) \to \mathbb{R}$ by

$$\gamma(\log r) = -\log G(\nu(0, r)),$$

as in Lemma 6.1. By step (i),

$$\|f\|_{\nu(0,\tilde{r})} \leq \rho, \quad \|f - \alpha\|_{\nu(0,\tilde{r})} \leq \rho, \quad \text{and} \quad \|f - 1\|_{\nu(0,\tilde{r})} = 1,$$

and, in addition, f is holomorphic on $D(0, \tilde{r})$. Hence f' is also holomorphic on $D(0, \tilde{r})$, and therefore $\|f'\|_{\nu(0,\tilde{r})} \geq |f'(0)|$; see equation (4.1). Applying these various bounds to equation (6.3), we obtain

$$(7.1) \qquad \gamma(\log \tilde{r}) \leq 2(\log \rho - \log |f'(0)| - \log \tilde{r}).$$

We claim that there is some $R \in [\tilde{r}, 1)$ such that

$$N_0(0, R) + N_\alpha(0, R) + N_1(0, R) + N_\infty(0, R) > 2N_{\text{ram}}(0, R).$$

Indeed, if there were no such R, then by Lemma 6.1, γ would have slope at most -2 on $[\log \tilde{r}, 0)$. By inequality (7.1), it would follow that

$$\gamma(\log r) \leq 2(\log \rho - \log |f'(0)| - \log r)$$

for all $r \in [\tilde{r}, 1)$. Since $|f'(0)| > \rho$, there would be some such r for which $G(\nu(0, r)) > 1$, contradicting Lemma 6.2. Thus, the desired radius R does exist, as claimed.

(**iii**). Since $R \geq \tilde{r}$, then by Step (i), $\nu_1 \in f_*(\overline{\mathcal{D}}(0, R))$. Moreover, by Lemma 6.3, we may choose $y \in \overline{D}(0, R)$ such that for every $r \in (0, R]$,

$$N_0(y, r) + N_\alpha(y, r) + N_1(y, r) + N_\infty(y, r) > 2N_{\text{ram}}(y, r).$$

We wish to find a (possibly different) point $x \in \overline{D}(0, R)$ and a (possibly smaller) radius R' so that the pair (x, R') satisfies the following four properties:

(1). $\nu_1 \in f_*(\overline{\mathcal{D}}(x, R'))$,

(2). $\nu_1 \notin f_*(\mathcal{D}(x, R'))$,

(3). $N_0(x, r) + N_\alpha(x, r) + N_1(x, r) + N_\infty(x, r) > 2N_{\mathrm{ram}}(x, r)$ for all $r \in (0, R']$, and

(4). $G(\nu(x, R')) \geq \mu^2$.

(**iv**). We will find x and R' by an inductive process. We begin with the point $y_0 = y$ and the radius $R_0 = R$ from step (iii). We have just observed that properties (1) and (3) already apply to the pair (y_0, R_0).

At stage $n \geq 1$ of the process, given a pair (y_{n-1}, R_{n-1}) for which properties (1) and (3) hold, define

$$\mathcal{R}'_{n-1} = \{r \in (0, R_{n-1}] : \nu_1 \in f_*(\overline{\mathcal{D}}(y_{n-1}, r))\},$$

which is non-empty because $R_{n-1} \in \mathcal{R}'_{n-1}$. Let $R'_{n-1} = \inf \mathcal{R}'_{n-1}$. It follows easily from the continuity of f_* that $R'_{n-1} > 0$. By definition of R'_{n-1}, the properties of (y_{n-1}, R_{n-1}), and Lemma 5.1, properties (1) and (3) apply to the pair (y_{n-1}, R'_{n-1}). In addition, property (2) applies to the pair, by Lemma 5.3. If $G(\nu(y_{n-1}, R'_{n-1})) \geq \mu^2$, then property (4) holds, and our process finishes by setting $(x, R') = (y_{n-1}, R'_{n-1})$. We may therefore assume that $G(\nu(y_{n-1}, R'_{n-1})) < \mu^2$.

By property (1), there is a disk $\overline{D}(z_n, R''_n) \subseteq \overline{D}(y_{n-1}, R'_{n-1})$ such that $f_*(\nu(z_n, R''_n)) = \nu_1$; hence $G(\nu(z_n, R''_n)) \geq \mu^2$, by hypothesis (c). Therefore, $\overline{D}(z_n, R''_n) \neq \overline{D}(y_{n-1}, R'_{n-1})$. In particular, $R''_n < R'_{n-1}$.

Define $\gamma_n(\log r) = -\log G(\nu(z_n, r))$, so that γ_n is continuous on the interval $[R''_n, R'_{n-1}]$, with $\gamma_n(R''_n) \leq -2\log\mu < \gamma_n(R'_{n-1})$. Therefore, γ_n has an absolute minimum at some radius $R_n \in [R''_n, R'_{n-1})$. Without loss, we may assume R_n is the largest radius in $[R''_n, R'_{n-1})$ for which γ_n attains its minimum. Thus, the piecewise linear function γ_n must have strictly positive slope at R_n; hence, by Lemma 6.1,

$$N_0(z_n, R_n) + N_\alpha(z_n, R_n) + N_1(z_n, R_n) + N_\infty(z_n, R_n) > 2N_{\mathrm{ram}}(z_n, R_n).$$

By Lemma 6.3, there is a point $y_n \in \overline{D}(z_n, R_n)$ such that for every $r \in (0, R_n]$,

$$N_0(y_n, r) + N_\alpha(y_n, r) + N_1(y_n, r) + N_\infty(y_n, r) > 2N_{\mathrm{ram}}(y_n, r).$$

In addition, $\nu_1 \in f_*(\overline{\mathcal{D}}(y_n, R_n))$, since $\overline{D}(z_n, R''_n) \subseteq \overline{D}(y_n, R_n)$. Thus, the pair (y_n, R_n) satisfies properties (1) and (3) above, so that our inductive process may repeat.

To show that the process must eventually stop, we claim, for every $n \geq 1$, that $y_{n-1} \notin \overline{D}(y_n, R'_n)$. Otherwise, because $y_n \in \overline{D}(y_{n-1}, R'_{n-1})$ and $R'_n \leq R_n < R'_{n-1}$, we would have $D(y_{n-1}, R'_{n-1}) = D(y_n, R'_{n-1})$. However, $\nu_1 \in f_*(\overline{\mathcal{D}}(y_n, R'_n)) \subseteq f_*(\mathcal{D}(y_n, R'_{n-1}))$, by condition (1) for (y_n, R'_n). At the same time, $\nu_1 \notin f_*(\mathcal{D}(y_{n-1}, R'_{n-1}))$, by condition (2) for (y_{n-1}, R'_{n-1}). This contradiction proves the claim.

In particular, the points y_n are all distinct. Meanwhile, $f(y_n) \in \{0, \alpha, 1, \infty\}$ for each y_n, by property (3) and the comments following Lemma 6.3. By Lemma 5.1, there are only finitely many such points in $\overline{D}(0, R)$; hence, the process must eventually stop.

(**v**). We now have a pair (x, R') satisfying properties (1)–(4) from step (iii) above. For the remainder of the proof, we proceed roughly as in Proposition 5.2 of [6].

As noted following Lemma 6.3, we have $f(x) \in \{0, \alpha, 1, \infty\}$. We may assume that $f(x) = 0$. Indeed, if $f(x) = \alpha$, then we can replace f by $\eta \circ f$, where $\eta(z) = (z - \alpha)/(1 - \alpha)$, which takes α to 0, fixes 1, ∞, and ν_1, and takes 0 to $\alpha/(\alpha - 1)$. It also maps each of the disks U_i into the disks of appropriate Ahlfors radius for these four centers, and it preserves properties (1)–(4) above. Similar arguments hold for $f(x) = \infty$ (with $\eta(z) = \alpha/z$) and for $f(x) = 1$ (with $\eta(z) = [\alpha(z - 1)]/[(\alpha - 1)z]$).

Since $\nu_1 \notin f_*(\mathcal{D}(x, R'))$, and since $f(x) = 0$, f cannot have poles in $D(x, R')$, by Lemma 5.2. Similarly, it cannot take on the value 1 in $D(x, R')$. Thus, f and f' are holomorphic on $D(x, R')$, and $N_1(x, r) = N_\infty(x, r) = 0$ for all $0 < r < R'$.

Let $R'' = \inf\{0 < r \le R' : \alpha \in f(\overline{D}(x, r))\}$ if this set is non-empty, or $R'' = R'$ if it is empty. For $R'' < r < R'$, we have $N_0(x, r) = N_\alpha(x, r)$, since any two points in the holomorphic image of a disk have the same number of preimages (see the comments preceding Definition 3.2). For any such r, then, condition (3) becomes

$$(7.2) \qquad N_0(x, r) + N_\alpha(x, r) + N_1(x, r) + N_\infty(x, r) - 2N_{\mathrm{ram}}(x, r) - 2 \ge 0,$$

because $N_0 + N_\alpha + N_1 + N_\infty = 2N_0$ is even.

Define the piecewise linear function γ_x by $\gamma_x(\log r) = -\log G(\nu(x, r))$, so that the left side of (7.2) is the slope of γ_x at $\log r$. By (7.2) and property (4), then,

$$\gamma_x(\log R'') \le \gamma_x(\log R') \le -2\log\mu.$$

Note that $f(D(x, R'')) \subseteq D(0, |\alpha|)$, by the minimality of R''. Hence, $\|f - \alpha\|_{\nu(x, R'')} = |\alpha|$ and $\|f - 1\|_{\nu(x, R'')} = 1$, and therefore

$$\gamma_x(\log R'') = \log|\alpha| + \log\|f\|_{\nu(x, R'')} - 2(\log R'' + \log\|f'\|_{\nu(x, R'')}),$$

so that

$$(7.3) \qquad \log|\alpha| + \log\|f\|_{\nu(x, R'')} + 2\log\mu \le 2(\log R'' + \log\|f'\|_{\nu(x, R'')}).$$

By the comments following equation (6.2), $\delta(f, \nu(x, r)) = r\|f'\|_{\nu(x, r)}/\|f\|_{\nu(x, r)}$ for all $0 < r \le R''$. In particular, inequality (7.3) may be rewritten as

$$(7.4) \qquad \log|\alpha| + 2\log\mu \le \log\|f\|_{\nu(x, R'')} + 2\log\delta(f, \nu(x, R'')).$$

Moreover, for all $r < R''$, property (3) is

$$(7.5) \qquad\qquad\qquad N_0(x, r) \ge 1 + 2N_{\mathrm{ram}}(x, r).$$

(**vi**). Let $\tilde{R}$ be the largest radius in $(0, R'']$ such that f is one-to-one on $D(x, \tilde{R})$. (Note that $\tilde{R} > 0$, by the comments following Lemma 6.3.) By inequality (7.5) and the fact that f is holomorphic on $D(x, R'')$ with $f(x) = 0$, it is easy to verify that

$$(7.6) \qquad N_0(x, r) \text{ is divisible by } p = \operatorname{char} k \text{ for all } r \in (\tilde{R}, R'')$$

and that

$$(7.7) \qquad \tilde{R}\|f'\|_{\nu(x, \tilde{R})} = \|f\|_{\nu(x, \tilde{R})}, \quad \text{or equivalently,} \quad \delta(f, \nu(x, \tilde{R})) = 1;$$

see, for example, Lemma 5.5 of [6].

We will now complete the proof by showing that f maps $U = D(x, \tilde{R})$ one-to-one and onto a disk containing $D(0, |\alpha|s)$. Since $f(x) = 0$ and f is holomorphic on U, we only need to show that $\|f\|_{\nu(x, \tilde{R})} \ge |\alpha|s$. We consider four separate cases, corresponding to the definition of μ in the statement of Theorem 7.2.

Case 1: $p = 0$. Since $f(x) = 0$, we know that $N_0(x, r) \geq 1$ for any $r > 0$; hence $\tilde{R} = R''$, lest (7.6) be violated. Inequality (7.4) with $s = \mu = 1$ and equation (7.7) then imply that $\|f\|_{\nu(x,\tilde{R})} \geq |\alpha|s$, and we are done.

Case 2: $p = 2$. Inequality (7.2) holds not just for $R'' < r < R'$ but also for $\tilde{R} < r < R''$, because $N_0(x, r)$ is divisible by 2 for all such r, by (7.6). Thus, γ_x does not decrease on $[\log \tilde{R}, \log R'']$, so that

$$\log |\alpha| + \log \|f\|_{\nu(x,\tilde{R})} + 2 \log \mu \leq 2(\log \tilde{R} + \log \|f'\|_{\nu(x,\tilde{R})}),$$

by inequality (7.3). By equation (7.7), then, $\|f\|_{\nu(x,\tilde{R})} \geq |\alpha|\mu^2 = |\alpha|s$, as desired.

Case 3: $p \geq 3$. For every integer $i \geq 0$, set $r_i = \sup\{r \in (0, R''] : N_0(x, r) < p^{i+1}\}$, and for every $i \geq 1$, define the piecewise linear function ψ_i by

$$\psi_i(\log r) = (p^i + 1) \log \|f\|_{\nu(x,r)} - 2p^i(\log r + \log \|f'\|_{\nu(x,r)})$$
$$= (1 - p^i) \log \|f\|_{\nu(x,r)} - 2p^i \log \delta(f, \nu(x, r)).$$

Note that the slope of ψ_i at $\log r$ is $(p^i + 1)N_0(x, r) - 2p^i(N_{\mathrm{ram}}(x, r) + 1)$, and that

$$\tilde{R} = r_0 \leq r_1 \leq \cdots \leq r_n = R''$$

where n is the smallest integer such that $N_0(x, R'') < p^{n+1}$.

For $r \in (r_{i-1}, R'')$, we have $N_0(x, r) \geq p^i$. Inequality (7.5) therefore gives us

$$2p^i N_{\mathrm{ram}}(x, r) + 2p^i \leq p^i N_0(x, r) + p^i \leq (p^i + 1)N_0(x, r),$$

which implies that ψ_i has non-negative slope everywhere on $(\log r_{i-1}, \log R'')$ and hence does not decrease on $[\log r_{i-1}, \log R'']$. Thus,
(7.8)
$$\log \|f\|_{\nu(x,r_{i-1})} - \log \|f\|_{\nu(x,r)} \geq \frac{2p^i}{p^i - 1} \left[\log \delta(f, \nu(x, r)) - \log \delta(f, \nu(x, r_{i-1})) \right]$$

for any $r \in [r_{i-1}, R'']$.

Using $i = 1$ and $r = R''$ in inequality (7.8), we obtain our desired result that

$$\log \|f\|_{\nu(x,\tilde{R})} \geq \log \|f\|_{\nu(x,R'')} + \frac{2p}{p - 1}[\delta(r, \nu(x, R'')) - \delta(r, \nu(x, \tilde{R}))]$$

$$\geq \frac{1}{p - 1} \left[p \log |\alpha| + 2p \log \mu - \log \|f\|_{\nu(x,R'')} \right]$$

$$\geq \log |\alpha| + \log s,$$

where the second inequality is due to (7.4) and (7.7), and the third is because $\mu = s^{(p-1)/(2p)}$ in this case and because $f(D(x, R'')) \subseteq D(0, |\alpha|)$.

Case 4: $p \geq 3$ **and** char $K = 0$. In this final case, preserving all the notation from Case 3, we use the functions ψ_i defined above for general i, rather than simply for $i = 1$. For fixed $i \geq 0$ and for $r \in (0, r_i]$, it is easy to verify that $\delta(f, \nu(x, r)) \geq |p|^i$, since f is holomorphic on $D(x, r_i)$ and $1 \leq N_0(x, r) < p^{i+1}$ for all $r < r_i$; see, for example, Lemma 4.2 of [6].

Substituting $r = r_i$ in inequality (7.8) and taking the sum over all $1 \leq i \leq n$,

$$\log \frac{\|f\|_{\nu(x,\tilde{R})}}{\|f\|_{\nu(x,R'')}} \geq \sum_{i=1}^{n} (2 + e_i) \left[\log \delta(f, \nu(x, r_i)) - \log \delta(f, \nu(x, r_{i-1})) \right]$$

$$= (2 + e_n) \log \delta(f, \nu(x, r_n)) + \sum_{i=1}^{n-1} (e_i - e_{i+1}) \log \delta(f, \nu(x, r_i))$$

$$\geq 2 \log \delta(f, \nu(x, R'')) + \left[n e_n + \sum_{i=1}^{n-1} (i e_i - i e_{i+1}) \right] \log |p|$$

$$\geq 2 \log \mu + (e_1 + e_2 + \cdots + e_n) \log |p| + \log |\alpha| - \log \|f\|_{\nu(x,R'')}$$

where $e_i = 2/(p^i - 1)$. Here, the equality is by equation (7.7) and because $r_0 = \tilde{R}$; the second inequality is because $e_i > e_{i+1}$ and $\delta(f, \nu(x, r_i)) \geq |p|^i$; and the third is by inequality (7.4). Since $|p| < 1$ and $e_1 + \cdots + e_n < 2E_p$, we get the desired inequality

$$\|f\|_{\nu(x,\tilde{R})} > |\alpha| |p|^{2E_p} \mu^2 = |\alpha| s. \qquad \square$$

Proof of Theorem 7.2. For each $i = 1, 2, 3, 4$, choose a point $a_i \in U_i$. Re-indexing if necessary, we may assume by hypothesis that neither a_1 nor a_2 lies in the same residue class of ν_1 as either a_3 or a_4. Similarly, we may assume that $f(0)$ does not lie in the same residue class as either a_3 or a_4.

There is a (unique) map $\eta \in \mathrm{PGL}(2, K)$ taking a_1 to 0, a_3 to 1, and a_4 to ∞. Let $\alpha = \eta(a_2)$. Each disk $\eta(U_i)$ is contained in one of the disks $D(0, |\alpha| s), \ldots, \mathbb{P}^1(K) \setminus \overline{D}(0, 1/s)$ listed in the statement of Theorem 7.3. It is straightforward to check that $\eta_*(\nu_1) = \nu(0, \rho)$ for some $\rho \in [|\alpha|, 1]$, and that the rest of the hypotheses of Theorem 7.3 are satisfied by $\eta \circ f$. (For example, hypothesis (c) follows from Lemma 6.2; and the lower bound on $|(\eta \circ f)'(0)|$ falls out of a brief computation using the given lower bound on $f^\#(0)$.) Hence, $\eta \circ f$ maps some disk one-to-one and onto one of the four disks $D(0, |\alpha| s)$, etc. Taking subdisks, one of the disks $\eta(U_i)$ is also a one-to-one image under $\eta \circ f$ (necessarily of a disk); Theorem 7.2 now follows by applying η^{-1}. $\qquad \square$

Corollary 7.4. *Let K, k, p, and E_p be as in Theorem 7.2. Let $U_1, U_2, U_3, U_4 \subseteq \mathbb{P}^1(K)$ be four pairwise disjoint open disks. Let $\nu_1 \in \mathcal{P}^1(K)$ such that no residue class of ν_1 intersects more than two of U_1, U_2, U_3, U_4. Define μ as in Theorem 7.2.*

Let f be a non-constant meromorphic function on K such that for any point $\nu \in \mathcal{P}^1(K) \setminus \{\infty\}$ for which $f_*(\nu) = \nu_1$, we have $\delta(f, \nu) \geq \mu$. Then there is an open disk $U \subseteq K$ such that f is one-to-one on U and $f(U) = U_i$ for some $i = 1, 2, 3, 4$.

Proof. Since f is non-constant and meromorphic on K, $f(K)$ omits at most one point of $\mathbb{P}^1(K)$; by Lemma 5.2, all type II and III points of $\mathcal{P}^1(K)$ are in the image of the Berkovich affine line $\mathcal{P}^1(K) \setminus \{\infty\}$. The hypotheses therefore imply that $\delta(f, \nu)$ is not identically zero (as would otherwise be possible for a function of the form $f(z) = h(z^p)$ in characteristic p, for example). Thus, $f^\#(x) > 0$ for some $x \in K$.

For $c \in K$ with $|c|$ sufficiently large, the map $g(z) = f(c(z - x))$ satisfies the hypotheses of Theorem 7.2 on $D(0, 1)$, and the corollary follows. $\qquad \square$

8. EXAMPLES

Theorem 7.2 differs from its complex counterpart in several noticeable ways. First and foremost, only four islands are required, as opposed to the five in the complex case. As observed in [6, Example 6], it would be impossible to reduce the number to three islands. On the other hand, in the case of positive residue characteristic, the non-archimedean theorem requires the extra condition that $\delta(f, \nu) \geq \mu$ for any ν mapping to ν_1. A weaker condition, that $\delta(f, \nu(0, r)) \geq \mu$ for *some* (large enough) r, is required for the holomorphic version [6]. In this section, we present examples to illustrate both that the lower bound of μ is essentially sharp, and that it is not enough to assume only one ν mapping to ν_1 satisfies the inequality, even if the number of islands is increased.

We address the sharpness μ in Example 8.1, but only in the case that char $k = p > 0 =$ char K. As in [6], we will show that the bound μ in Theorem 7.2 (really, Corollary 7.4, since we will consider meromorphic functions on K) is sharp if $p = 2$ and is almost sharp (except possibly for the coefficient $|p|^{-E_p}$) if $p \geq 3$. (If char $k = 0$, then condition (b) is vacuous and hence trivially sharp.) We conjecture that examples analogous to [6, Examples 3 and 4] would prove the sharpness of μ in the cases that char $K = p > 0$.

Example 8.1. Suppose char $K = 0$ and char $k = p \geq 2$. (For example, let $K = \mathbb{C}_p$.) Let E be an elliptic curve defined over K with identity point O, and let $n \geq 1$ be an integer. Assume that E has good ordinary reduction. (The conclusions we will reach also hold for multiplicative reduction, but that case is slightly more complicated because there are a number of different points ν which will map to the point ν_1 we will choose shortly.) Let E_1 be the set of points which map to $\overline{O}$ under reduction.

If we identify E_1 with the open unit disk, then by the characteristics of K and k, and because the corresponding formal group has height 1, there are p-torsion points $\{P_1, \ldots, P_{p-1}\}$ at distance $|p|^{1/(p-1)}$ from O in E_1. Moreover, there are no non-trivial torsion points closer than $|p|^{1/(p-1)}$ to O.

The multiplication-by-n map $[n] : E \to E$ has the property that $[n](-P) = -[n](P)$ for any point P on E. Meanwhile, the group $\{\pm 1\}$ acts on the curve E (with -1 taking P to $-P$) with quotient $\mathbb{P}^1$. It follows that there is a map $f_n : \mathbb{P}^1 \to \mathbb{P}^1$ for which

$$
\begin{array}{ccc}
E & \xrightarrow{\ [n]\ } & E \\
h \downarrow & & h \downarrow \\
\mathbb{P}^1 & \xrightarrow{\ f_n\ } & \mathbb{P}^1
\end{array}
$$

commutes, where h is the quotient map. (If E is written in the form $y^2 = $ (cubic in x), then h is just the x-coordinate map.) The function f_n is known as a Lattès map to dynamicists. It is a rational function of (geometric) degree n^2.

Let a_1, a_2, a_3, a_4 be the images under h of the 2-torsion points $E[2]$ of E. Note that $E[2]$ is the set of ramification points of h. For convenience, choose coordinates on $\mathbb{P}^1$ so that $h(O) = a_1 = 0$ and so that $h(E_1) = D(0, 1)$. Let $\nu_1 = \nu(0, 1)$; then because E has good reduction, we have $(f_n)_*(\nu) = \nu_1$ if and only if $\nu = \nu_1$. Furthermore, an examination of the map $[n]$ restricted to E_1 shows that $\delta(f_n, \nu_1) = |n|$. Let $\mu = |n|$.

Let $C_n = E[2n] \setminus E[2]$ be the set of $2n$-torsion points which are not 2-torsion, and let $B_n \subseteq \mathbb{P}^1(K)$ be the image of C_n under h. By considering the ramification of the map h, and knowing that f_n must have exactly $2 \deg f_n - 2 = 2n^2 - 2$ critical points, it is easy to check that B_n is the set of critical points of f_n, and that each point of B_n maps 2-to-1 to its image. One can also check that $f_n'(0) = n^2$.

Because h maps E_1 two-to-one onto $D(0,1)$, it is not difficult to show that $|h(P_i)| = |p|^{2/(p-1)}$ for all $i = 1, \ldots, p-1$. Similarly, the lack of non-trivial torsion points in E_1 at distance less than $|p|^{1/(p-1)}$ from O implies that

$$(8.1) \qquad (\{a_2, a_3, a_4\} \cup B_n) \cap D(0, |p|^{2/(p-1)}) = \emptyset.$$

In fact, if $p \nmid n$ and $p \geq 3$, then $(\{a_2, a_3, a_4\} \cup B_n) \cap D(0, 1) = \emptyset$.

Meanwhile, $B_n \cup \{a_1, a_2, a_3, a_4\}$ is precisely the preimage of $\{a_1, a_2, a_3, a_4\}$ under f_n. Thus, if we choose U_1, U_2, U_3, U_4 to be disjoint disks containing a_1, a_2, a_3, a_4, then the only disks U that could map one-to-one onto any U_i would have to contain exactly one of a_1, a_2, a_3, a_4 and cannot intersect B_n. By translating on E by 2-torsion points, it suffices to consider only one-to-one mappings of f_n from a disk U containing $a_1 = 0$ to $U_1 \subseteq D(0, 1)$. In fact, the preimage disk U we should consider is the largest disk about 0 which contains no points in $B_n \cup \{a_2, a_3, a_4\}$. If $p|n$ or $p = 2$, then by equation (8.1), this disk is $U = D(0, |p|^{2/(p-1)})$. Since $|f_n'(0)| = |n|^2$, the largest possible one-to-one image disk therefore has radius $s = |n|^2 \cdot |p|^{2/(p-1)}$. (See, for example, [6, Lemma 2.2].) On the other hand, if $p \nmid n$ and $p \geq 3$, the domain disk is $U = D(0, 1)$, and the image also has radius $s = 1$.

If $p = 2$, then there is one non-trivial 2-torsion point P_1 in E_1. Its image $h(P_1) = a_2$ is the point α of Section 7, with $|\alpha| = |p|^{2/(p-1)}$. Since no residue class of ν_1 contains more than two of a_1, a_2, a_3, a_4, the conditions of Corollary 7.4 hold for f_n. The image disk U_1 of the previous paragraph therefore has radius $|p|^{2/(p-1)}|n|^2 = |\alpha| \cdot \mu^2$; thus, the Ahlfors radius $s = \mu^2$ attains the bound $\mu = s^{1/2}$ in Theorem 7.2.

If $p \geq 3$, we have $|\alpha| = 1$, since none of $\{a_2, a_3, a_4\}$ lie in $D(0, 1)$. If $p \nmid n$, then $\mu = s = 1$, attaining the better (i.e., smaller) of the two values given for μ in this case in Theorem 7.2. Finally, if $p|n$, then the image disk has radius $s = \mu^2|p|^{2/(p-1)}$, which is the same order of magnitude (and only slightly worse in the coefficient) than the Theorem 7.2 bound of $s > \mu^2|p|^{2E_p}$.

Our final example will illustrate that having only one ν for which $f_*(\nu) = \nu_1$ with $\delta(f, \nu)$ bounded below by some fixed amount is not enough to guarantee an islands theorem, regardless of how many islands there are, how small they are, or how large the lower bound on $\delta(f, \nu)$ is. Like the previous example, Example 8.2 is only for the case that char $k > $ char $K = 0$. Of course, if char $k = 0$, then condition (b) is already vacuous, so there will be no counterexamples. On the other hand, if char $K = p > 0$, then we imagine that examples similar to the following one may be constructed.

Example 8.2. Assume that char $k = p > 0 = $ char K. For any integer $N \geq 0$ and any radius $0 < s \leq 1$, we select $N + 2$ islands as follows. Set $a_0 = 0 \in K$. For each $i = 1, \ldots, N$, choose $a_i \in K$ such that $|a_i| = 1$ and, for each $j \neq i$, $|a_i - a_j| = 1$. The first $N + 1$ islands will be the open disks $D(a_i, s)$, with $i = 0, \ldots, N$. The final island will be $\mathbb{P}^1(K) \setminus \overline{D}(0, 1/s)$, which is the open disk of spherical radius s centered at ∞. Let $\nu_1 = \nu(0, 1)$, which separates each island from every other.

Pick $n \geq 1$ large enough so that $|p^n| < s$, and choose $b \in K$ so that $0 < |b| < s$. Let $c = -b^{1+p^n}$. For each $i = 1, \ldots, N$, define

$$f_i(z) = \frac{(z - a_i)^{p^n} + c}{(z - a_i)^{p^n}} = 1 + \frac{c}{(z - a_i)^{p^n}}.$$

Then, define

$$f_0(z) = \frac{z^{1+p^n} + c}{z^{p^n}} = z + \frac{c}{z^{p^n}}, \quad \text{and} \quad f(z) = \prod_{i=0}^{N} f_i(z).$$

We consider f as a meromorphic function on K. The reader may check that $f_*(\nu_1) = \nu_1$ and $\delta(f, \nu(0,1)) = 1$, which by Lemma 6.2 is the maximum value δ can attain.

The only preimages of ∞ are $a_0, \ldots, a_N$, all of which are critical points. Thus, the island at ∞ has no one-to-one preimages. (If we were to use $\mathbb{P}^1(K)$ as our domain instead of K, then simply remove the final island (at ∞) from consideration.) In addition, if $|z - a_i| \geq 1$ for all $i = 0, \ldots, N$, then it is easy to see that $|f(z) - a_i| \geq 1$ also. In particular, any preimages of the islands $D(a_i, s)$ must lie in the disks $D(a_j, 1)$.

There are $1 + p^n$ preimages of 0 in $D(0,1)$, namely the roots of $f_0(z) = 0$, which are all of the form $\zeta^j b$, where ζ is a primitive $(1 + p^n)$-root of unity. In particular, the largest open disk containing exactly one such root has radius $|b|$. It is easy to compute that $|f'(\zeta^j b)| = 1$, and therefore the image of that largest open disk is a disk of radius $|b| < s$. Thus, the island at 0 is not a one-to-one image of a disk inside $D(0,1)$.

Similarly, for any fixed $i = 1, \ldots, N$, there are p^n preimages of 0 in $D(a_i, 1)$, namely the roots of $(z - a_i)^{p^n} = -c$. Those roots are of the form $x = a_i + \omega^j d$, where ω is a p^n-root of unity, and d is a p^n-root of $-c$. Any ω^j is distance $|p|^{1/(p-1)}$ from the nearest other ω^ℓ, so that the largest disk containing exactly one root of $(z - a_i)^{p^n} = -c$ has radius $|p|^{1/(p-1)}|c|^{1/p^n}$. We can compute that $|f'(x)| = |p|^n \cdot |c|^{-1/p^n}$, so that the largest possible one-to-one image disk has radius $|p|^{n+(1/(p-1))} < |p|^n < s$. Thus, there are no one-to-one preimages of the island at 0 anywhere in K.

For $1 \leq i \leq N$, the preimages of a_i in $D(0,1)$ must be points $x \in D(0,1)$ satisfying $|f_0(x)| = 1$. Since $|x| < 1$, this must mean $|x| = |c|^{1/p^n}$. Because of the pole at 0, the largest open disk about x that could conceivably map onto the island $D(a_i, s)$ must have radius at most $|c|^{1/p^n}$. (In fact, it will have slightly smaller radius than that, but the bound of $|c|^{1/p^n}$ will suffice for our purposes.) Some computation using the above value for $|x|$ shows that $|f'(x)| \leq \max\{1, |p|^n|c|^{-1/p^n}\}$, and therefore the largest possible one-to-one image disk has radius

$$\max\left\{|c|^{1/p^n}, |p|^n\right\} \leq \max\left\{|b|, |p|^n\right\} < s,$$

which fails to cover the island.

Before considering the final case, observe that

$$f(z) - z = z \cdot \left[-1 + \left(1 + \frac{c}{z^{1+p^n}}\right) \prod_{j=1}^{N} \left(1 + \frac{c}{(z - a_j)^{p^n}}\right)\right],$$

so that

$$|f(z) - z| \leq |z| \cdot \max\left\{\left|\frac{c}{z^{1+p^n}}\right|, \left|\frac{c}{(z-a_1)^{p^n}}\right|, \ldots, \left|\frac{c}{(z-a_N)^{p^n}}\right|\right\}.$$

Thus, if $|c|^{1/(1+p^n)} < |x - a_j| < 1$ for some $j = 1, \ldots, N$ and some $x \in K$, then $|f(x) - x| < |x - a_j|$. In particular, no such x can have image $f(x)$ in any of the $N + 2$ islands.

We are now ready to consider the final possibility, that there is a preimage of an island $D(a_i, s)$ in the disk $D(a_j, 1)$ for some $i, j \in \{1, \ldots, N\}$, where i and j may or may not be equal. In such a case, we have a point x in $D(x, a_j)$ with $f(x) = a_i$. This means that $|f_j(x)| = 1$, so that $|x - a_j| \geq |c|^{1/p^n}$. From this bound, it follows that $|f'(x)| \leq \max\{1, |p^n|/|x - a_j|\}$. Meanwhile, the largest possible disk mapping onto $D(a_i, s)$ cannot contain the pole at a_j, so that its radius must be at most $|x - a_j|$. Therefore, the largest possible one-to-one image disk about a_i has radius at most

$$|f'(x)| \cdot |x - a_j| \leq \max\{|x - a_j|, |p^n|\} \leq \max\{|c|^{1/(1+p^n)}, |p^n|\} = \max\{|b|, |p^n|\} < s,$$

where the second inequality is due to the previous paragraph. Thus, none of the islands has a one-to-one preimage anywhere in K, even though $f_*(\nu_1) = \nu_1$ with $\delta(f, \nu_1) = 1$.

We close by noting that an examination of the proof of Theorem 7.3 applied to Example 8.2 reveals why the hypothesis that *every* (or at least many) ν mapping to ν_1 must have $\delta(f, \nu) \geq \mu$. For that choice of f, we start from $R = 1$ and move inward, searching for a disk satisfying properties (1)–(4) of Step (iii) of the proof of Theorem 7.3. Properties (1), (3), and (4) already apply to $\overline{D}(0, R)$, but property (2) does not. The inductive process of Step (iv) would begin with $R_0 = 1$ and y_0 being one of the roots of $f = a_i$ in $D(a_i, 1)$ for some $i = 0, \ldots, N$. The minimal radius R_0' would be the smallest radius about y_0 for which $f(\overline{D}(y_0, R_0'))$ contained points outside $D(a_i, 1)$. That is, $R_0' = |y_0 - a_i|$, which we saw to be at most $|b|$. Even though the inequality of property (3) holds, the exponent of G is positive (in fact, equal to 1) for $r \in [R_0', R_0]$, so that as r shrinks from R_0 down to R_0', G also shrinks from 1 down to $R_0' \leq |b|$. Thus, although f is at last one-to-one on $D(y_0, R_0')$, the image is too small because there are poles too close to y_0, and the value of G (as well as δ, along with it) has shrunk too much. Then, without hypothesis (b) of the theorem, if we try to shrink to a disk $D(z_1, R_1'')$ according to the algorithm, we have no guarantee that G will increase, and therefore we have no guarantee that property (3) holds.

Acknowledgements

References

[1] W. Adams and E. Straus, Non-archimedean analytic functions taking the same values at the same points, *Illinois J. Math.* **15** (1971), 418–424. MR0277771 (43:3504)

[2] L. Ahlfors, Zur Theorie der Überlagerungsflächen, *Acta Math.* **65** (1935), 157–194. MR1555403

[3] I.N. Baker, Repulsive fixpoints of entire functions, *Math. Z.* **104** (1968), 252–256. MR0226009 (37:1599)

[4] R. Benedetto, *p*-adic dynamics and Sullivan's No Wandering Domains Theorem, *Compositio Math.* **122** (2000), 281–298. MR1781331 (2001f:37054)

[5] R. Benedetto, Components and periodic points in non-archimedean dynamics, *Proc. London Math. Soc.* (3) **84** (2002), 231–256. MR1863402 (2002k:11215)

[6] R. Benedetto, Non-archimedean holomorphic maps and the Ahlfors Islands Theorem, *Amer. J. Math.*, **125** (2003), 581–622. MR1981035 (2004a:30042)

[7] W. Bergweiler, A new proof of the Ahlfors five islands theorem, *J. Anal. Math.* **76** (1998), 337–347. MR1676971 (99m:30032)

[8] V. Berkovich, *Spectral theory and analytic geometry over non-Archimedean fields,* Amer. Math. Soc., Providence, 1990. MR1070709 (91k:32038)

[9] J.-P. Bézivin, Sur les points périodiques des applications rationnelles en analyse ultramétrique, *Acta Arith.* **100** (2001), 63–74. MR1864626 (2002j:37059)

[10] J.-P. Bézivin, Dynamique des fractions rationnelles *p*-adiques, monograph, 2005. Available online at `http://www.math.unicaen.fr/~bezivin/dealatex.pdf`

[11] S. Bosch, U. Güntzer, and R. Remmert, *Non-Archimedean Analysis: A Systematic Approach to Rigid Analytic Geometry,* Springer-Verlag, Berlin, 1984. MR746961 (86b:32031)

[12] A. Boutabaa, Théorie de Nevanlinna *p*-adique, *Manuscripta Math.* **67** (1990), 251–269. MR1046988 (91m:30039)

[13] A. Boutabaa and A. Escassut, Nevanlinna theory in characteristic *p* and applications, *Analysis and applications—ISAAC 2001 (Berlin)*, Kluwer, Dordrecht, 2003, 97–107. MR2022742 (2004k:30062)

[14] W. Cherry, A survey of Nevanlinna theory over non-Archimedean fields, *Bull. Hong Kong Math. Soc.* **1** (1997), 235–249. MR1605198 (99a:32047)

[15] W. Cherry and Z. Ye, Non-Archimedean Nevanlinna theory in several variables and the non-Archimedean Nevanlinna inverse problem, *Trans. Amer. Math. Soc.* **349** (1997), 5043–5071. MR1407485 (98c:11072)

[16] C. Corrales-Rodrigáñez, Nevanlinna theory on the *p*-adic plane, *Ann. Polon. Math.* **57** (1992), 135–147. MR1182179 (93h:30067)

[17] J. Dufresnoy, Sur les domaines couvertes par les valeurs d'une fonction méromorphe ou algébroïde, *Ann. Sci. École Norm. Sup. (3)* **58** (1941), 179–259. MR0012669 (7:56b)

[18] A. Escassut, Algèbres de Banach ultramétriques et algèbres de Krasner-Tate, *Astérisque* **10** (1973), 1–107. MR0364676 (51:930)

[19] A. Escassut, The ultrametric spectral theory, *Period. Math. Hungar.* **11** (1980), 7–60. MR571136 (81i:46098)

[20] A. Escassut, *Analytic Elements in p-adic Analysis,* World Scientific, Singapore, 1995. MR1370442 (97e:46106)

[21] D. Goss, A short introduction to rigid analytic spaces, *The arithmetic of function fields*, Ohio State Univ. Math. Res. Inst. Publ. **2**, Columbus, OH, 1991, 131–141. MR1196516 (93m:14020)

[22] W. Hayman, *Meromorphic Functions*, Oxford Univ. Press, London, 1964. MR0164038 (29:1337)

[23] L.-C. Hsia, Closure of periodic points over a non-archimedean field, *J. London Math. Soc.* (2) **62** (2000), 685–700. MR1794277 (2001j:11117)

[24] P.-C. Hu and C.-C. Yang, *Meromorphic Functions over Non-Archimedean Fields*, Mathematics and its Applications 522, Kluwer, Dordrecht, 2000. MR1794326 (2002a:11085)

[25] Hà Huy Khoái and My Vinh Quang, On *p*-adic Nevanlinna theory, *Complex Analysis, Joensuu 1987*, *Lecture Notes in Math.* **1351**, Springer, Berlin, 1988, 146–158. MR982080 (90e:11153)

[26] M. Lazard, Les zéros des fonctions analytiques sur un corps valué complet, *Inst. Hautes Études Sci. Publ. Math.* **14** (1962) 47–75. MR0152519 (27:2497)

[27] J. Rivera-Letelier, Dynamique des fonctions rationnelles sur des corps locaux, *Astérisque* **287** (2003), 147–230. MR2040006 (2005f:37100)

[28] J. Rivera-Letelier, Espace hyperbolique *p*-adique et dynamique de fonctions rationnelles, *Compositio Math.* **138** (2003), 199–231. MR2018827 (2004k:37090)

[29] P. Robba, Fonctions analytiques sur les corps valués ultramétriques complets, *Astérisque* **10** (1973), 109–218. MR0357841 (50:10307)

[30] A. Robert, *A Course in p-adic Analysis*, Springer-Verlag, New York, 2000. MR1760253 (2001g:11182)

[31] M. Ru, A note on *p*-adic Nevanlinna theory, *Proc. Amer. Math. Soc.* **129** (2001), 1263–1269. MR1712881 (2001h:11097)

[32] R. Rumely and M. Baker, Analysis and dynamics on the Berkovich projective line, preprint, 2004. Available online at `arxiv:math.NT/0407433`

[33] L. Zalcman, A heuristic principle in complex function theory, *Amer. Math. Monthly* **82** (1975), 813–817. MR0379852 (52:757)

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, AMHERST COLLEGE, AMHERST, MASSACHUSETTS 01002

E-mail address: `rlb@cs.amherst.edu`

URL: `http://www.cs.amherst.edu/~rlb`