THE JOURNAL OF SYMBOLIC LOGIC
Volume 86, Number 1, March 2021

THICKET DENSITY

SIDDHARTH BHASKAR

Abstract. We define a new type of "shatter function" for set systems that satisfies a Sauer–Shelah type dichotomy, but whose polynomial-growth case is governed by Shelah's two-rank instead of VC dimension. We identify the least exponent bounding the rate of growth of the shatter function, the quantity analogous to VC density, with Shelah's ω-rank.

§1. Introduction. The *shatter function* is a function from $\mathbb{N}$ to $\mathbb{N}$ that measures of the complexity of a set system. The shatter function of any set system satisfies the *Sauer–Shelah dichotomy*: it is either the binary exponential function $n \mapsto 2^n$, or is polynomially bounded. Whether or not the shatter function is polynomially bounded or exponential depends on whether a certain integer parameter, the *VC dimension*, is finite or infinite. In the finite case, the least exponent bounding the polynomial growth of the shatter function is a real number called the *VC density*.

VC density was discovered by Vapnik and Chervonenkis [13] and found important applications in probability theory, combinatorial geometry, and computational learning theory.[1] The relevance of VC density to theories without the independence property was pointed out by Laskowski [8], and subsequently developed by Aschenbrenner et al. [1, 2].

In the present paper, we associate a new function with any set system, which we call the *thicket shatter function*. It also satisfies the Sauer–Shelah dichotomy, but the quantity that distinguishes between polynomial and exponential growth is an instance of Shelah's local two-rank, and its rate of growth is an instance of Shelah's local ω-rank. In this context, we call these two quantities *thicket dimension* and *thicket density* to emphasize the analogy.

Seen from another angle, our work can be read as a way to calculate Shelah's local ω-rank using the asymptotic growth of certain finite combinatorial objects. Notably, this can be performed in *any* model of a theory, not just a saturated one.

Our work was foreshadowed by Tiuryn, whose Lemma 3.6 in [12] contains a special case of our Theorem 4.3 below, the Sauer–Shelah dichotomy for thicket shatter functions. It is remarkable that he was concerned with problems in dynamic logic—at best a distant relative of model theory, and even further from geometry and computational learning theory.

Received June 12, 2018.
2020 *Mathematics Subject Classification*. 03C45.
Key words and phrases. model theory, stability theory, VC-density, set systems.

[1] See Assouad [3] for an exposition of the fundamentals of VC theory.

0022-4812/21/8601-0005
DOI:10.1017/jsl.2020.55

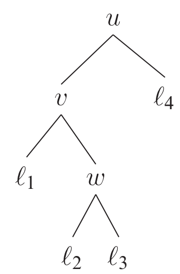

FIGURE 1. A binary tree with leaves $\ell_1,\ell_2,\ell_3,\ell_4$ and non-leaves u,v,w. Depending on the context, this could be ordered or unordered.

Organization of this paper. In §2, we discuss some basic facts about binary trees. In §3, we introduce thicket versions of the dimension and shatter function, and in §4, we prove the thicket version of the Sauer-Shelah dichotomy. In §5, we identify thicket density with a local model-theoretic rank, and in §6, we formulate a notion of degree, or multiplicity, for thicket density.

§2. Trees. Our fundamental objects of study are set systems and binary trees; the latter are the only sort of trees we will consider. Binary trees come in two varieties: *ordered*, and *unordered*, and this refers to whether we distinguish left from right children of non-leaves. Trees are normatively ordered; we shall say "unordered trees," when we mean it.

DEFINITION 2.1. A *tree* is either a single leaf or an ordered pair of subtrees, which we call *left* and *right*. An *unordered tree* is either a single leaf or an unordered pair of subtrees.

Notice that this definition allows for both finite trees and infinite trees of depth ω. In a set-theoretic account, a tree T would be defined as a nonempty prefix-closed subset of $2^{<\omega}$ such that for every $u \in 2^{<\omega}$, $u0 \in T \iff u1 \in T$. We prefer the "coinductive data type" definition presented here, which has the advantage of giving a more succinct definition of unordered trees. However, we will freely imagine a tree as a set of vertices, one of which is the *root*, some of which are *leaves*, equipped with a partial order defining the ancestor relation. We trust that this will cause no difficulty.

DEFINITION 2.2. For a tree T and vertices $u,v,w \in T$, we say

- $u \prec v$ in case $v \neq u$, but v is contained in the subtree with root u,
- $u \prec_L v$ if v is contained in the subtree whose root is the left child of u, and
- $u \prec_R v$ if v is contained in the subtree whose root is the right child of u.

For fixed $v \in T$, the set of vertices $P(v) = \{u : u \prec v\}$ is linearly ordered by $\prec$ and is partitioned by $P_L(v) = \{u : u \prec_L v\}$ and $P_R(v) = \{u : u \prec_R v\}$. For example in Figure 1, $P_L(\ell_3) = \{u\}$ and $P_R(\ell_3) = \{v,w\}$.

For vertices u, v in an unordered binary tree T, we cannot of course say that $u \prec_L v$ or $u \prec_R v$. What we can say is that v and w are contained in the same subtree of u, or different subtrees of u.

DEFINITION 2.3. For an unordered tree T, and vertices $u, v, w \in T$, we say

- $u \prec v$ in case $v \neq u$, but v is contained in the subtree with root u,
- $v \sim_u w$ in case $u \prec v$, $u \prec w$, and v and w lie in the same subtree of u, and
- $v \perp_u w$ in case $u \prec v$, $u \prec w$, and v and w lie in different subtrees of u.

Notice that, for any v and w, there is a unique u such that $v \perp_u w$. For example, if we interpret the tree in Figure 1 to be unordered, then $\ell_2 \sim_u \ell_3$, $\ell_2 \sim_v \ell_3$, but $\ell_2 \perp_w \ell_3$.

An important example of trees are the *finite, balanced* trees, which we call B_n:

DEFINITION 2.4. The tree B_0 is the single leaf. The tree B_{n+1} is the ordered pair of trees (B_n, B_n). Similarly, we define the unordered tree B_n° by letting B_0° be the single leaf, and B_{n+1}° be the unordered pair $\{B_n^\circ, B_n^\circ\}$.

DEFINITION 2.5. An *embedding* of the tree T_1 into the tree T_2 is an injection of the vertices of T_1 into the vertices of T_2 that preserves the $\prec_L$ and $\prec_R$ relations. An *embedding* of the unordered tree T_1 into the unordered tree T_2 is an injection of the vertices of T_1 into the vertices of T_2 that preserves the $\prec$, $\sim$, and $\perp$ relations. The *dimension* d of a tree (respectively, unordered tree) T is the largest n such that B_n (respectively, B_n°) can be embedded into T, or ∞ if there are arbitrarily large such n.

REMARK 2.6. For finite (ordered or unordered) trees T, dimension satisfies the following useful recursive identity. If T is a leaf, then $d = 0$. Otherwise, if d_1 and d_2 are the dimensions of its two subtrees, then

$$d = \begin{cases} \max\{d_1, d_2\} & \text{if } d_1 \neq d_2, \\ d_1 + 1 & \text{if } d_1 = d_2. \end{cases}$$

§3. Labeled trees and their solutions. A *set system* $(X, \mathcal{F}, \in)$ is a two-sorted structure with sorts X and $\mathcal{F}$, equipped with a single binary relation $\in \subseteq X \times \mathcal{F}$. Usually, we shall just write $(X, \mathcal{F})$, suppressing $\in$. A typical example of a set system is obtained by taking X to be some set, $\mathcal{F}$ to be a family of subsets of X, and $\in$ to be the containment relation. (But, in general, elements in $\mathcal{F}$ need not be extensional.) Given any set system $(X, \mathcal{F}, \in)$, its *dual* is the set system $(\mathcal{F}, X, \in^\star)$, where $F \in^\star x \iff x \in F$. Clearly, the dual of the dual of any set system is identical to the original.

DEFINITION 3.1. Let $(X, \mathcal{F})$ be a set system and T be a tree. An X*-labeling* of T is an assignment of elements of X to non-leaves of T. If T is an X-labeled tree with labeling $u \mapsto x_u$, and if $v \in T$ is a leaf, then we say $F \in \mathcal{F}$ *solves* v in case $(\forall u \prec v)(x_u \in F \iff u \prec_L v)$, or equivalently, $F \cap P(v) = P_L(v)$. In the special case that T has depth 0, i.e., is a single leaf v, the quantifier $\forall u \prec v$ is vacuous, so v has a solution in $\mathcal{F}$ iff $\mathcal{F}$ is nonempty (Figure 2).

If T is an unordered tree, an X-labeling is an assignment of vertices of T to elements of X. The solution of a single leaf is meaningless, but we can still speak of a solution to the whole tree.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

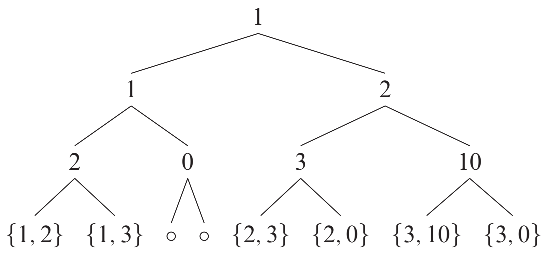

FIGURE 2. A $\mathbb{N}$-labeling of B_3. Leaves are labeled with solutions in $\binom{\mathbb{N}}{2}$, the family of all two-element subsets of $\mathbb{N}$, when they have one.

DEFINITION 3.2. If T is an X-labeled unordered tree and $L \subseteq T$ is the set of its leaves, a *solution* to T (in $\mathcal{F}$) is an assignment $L \to \mathcal{F}$, denoted $v \mapsto F_v$, satisfying

$$v_1 \sim_u v_2 \implies x_u \in F_{v_1} \leftrightarrow x_u \in F_{v_2},$$
$$v_1 \perp_u v_2 \implies x_u \in F_{v_1} \not\leftrightarrow x_u \in F_{v_2}.$$

In other words, any pair of sets labeling two leaves must disagree on whether they include the element of X labeling their most recent common ancestor, and agree on all other common ancestors.

DEFINITION 3.3. Say that a set system $(X, \mathcal{F})$ *admits* a tree T in case there is an X-labeling of T such that each leaf of T has a solution in $\mathcal{F}$. Similarly, $(X, \mathcal{F})$ *admits* an unordered tree T in case there exists an X-labeling of T which has a solution in $\mathcal{F}$. If a set system does not admit a tree or an unordered tree, then it *forbids* that tree.

REMARK 3.4. For a typical tree T, it is a much stronger statement to say that $(X, \mathcal{F})$ admits T, rather than $(X, \mathcal{F})$ admits the underlying unordered tree of T. Similarly, it is much stronger to say that $(X, \mathcal{F})$ forbids the underlying unordered tree of T, rather than $(X, \mathcal{F})$ forbids T. However, when $T = B_n$, there is no difference: admitting or forbidding B_n is equivalent to admitting or forbidding B_n°. This is because, if $(X, \mathcal{F})$ admits B_n°, it must admit *some* ordered tree which is obtained from B_n° by assigning "left" and "right" to each pair of siblings. But, each such tree is simply B_n.

REMARK 3.5. Suppose T is an unordered tree with subtrees T_1 and T_2. Then if $(X, \mathcal{F})$ admits T, there is an X-labeling of T with a solution in $\mathcal{F}$. Any solution can be partitioned into the leaves labeling T_1 and the leaves labeling T_2. If x is the label of the root of T, then the two parts of this solution disagree on x. Let $\mathcal{F}_x$ and $\mathcal{F}_{\bar{x}}$ be the elements of $\mathcal{F}$ containing and excluding x respectively. Then, either $(X, \mathcal{F}_x)$ admits T_1 and $(X, \mathcal{F}_{\bar{x}})$ admits T_2, or $(X, \mathcal{F}_x)$ admits T_2 and $(X, \mathcal{F}_{\bar{x}})$ admits T_1.

Conversely, suppose that for some $x \in X$, $(X, \mathcal{F}_x)$ admits T_1 and $(X, \mathcal{F}_{\bar{x}})$ admits T_2. Then $(X, \mathcal{F})$ admits T, by combining the X-labelings of T_1 and T_2 into an X-labeling of T, labeling the root by x. Similarly, if $(X, \mathcal{F}_x)$ admits T_2 and $(X, \mathcal{F}_{\bar{x}})$ admits T_1, then $(X, \mathcal{F})$ admits T.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

Definition 3.6. The *thicket dimension* of $(X,\mathcal{F})$, denoted $\dim(X,\mathcal{F})$, is the largest n such that B_n is admissible, or ∞ if there are arbitrarily large such n. If there are no such n, equivalently if $\mathcal{F}$ is empty, the dimension is – 1.

Remark 3.7. This is a well-known quantity that occurs in many different contexts. The thicket dimension of $(X,\mathcal{F})$ is equal to Shelah's rank $R(\mathtt{x}=\mathtt{x},\{\varphi\},2)$, where φ is the formula $\mathtt{x}\in\mathtt{F}$, relative to the theory $\mathrm{Th}(X,\mathcal{F})$ [7, 11]. It is also called *Littlestone dimension* in the context of computational learning theory [4]. Hodges calls it the *branching index* [7].

Definition 3.8. For a given set system $(X,\mathcal{F})$, let $\rho(n)$ be the maximum, as T varies over X-labelings of B_n, of the number of leaves of T with solutions in $\mathcal{F}$. The resulting function $\rho:\mathbb{N}\to\mathbb{N}$ is the *thicket shatter function* associated with $(X,\mathcal{F})$.

Remark 3.9. The thicket shatter function of any set system is bounded above by the binary exponential function $n\mapsto 2^n$.

Remark 3.10. To show that $\rho(n)\geq m$, it suffices to find an X-labeled tree T of depth *at most* n, such that at least m distinct leaves of T have solutions in $\mathcal{F}$. For then we could superimpose T on B_n, label all remaining non-leaves arbitrarily, and obtain an X-labeling of B_n at least m of whose leaves have solutions.

Dual Quantities. Given a set system $(X,\mathcal{F})$ and a tree T, an $\mathcal{F}$*-labeling* of T is an assignment of vertices of T to elements of $\mathcal{F}$. If T is an X-labeled tree with labeling $u\mapsto F_u$, and if $v\in T$ is a leaf, then we say $x\in X$ is a *solution* to v in case $(\forall u\prec v)(F_u\in x\iff u\prec_L v)$. This allows us to define corresponding "dual" versions of thicket dimension and the thicket shatter function, which are identical, respectively, to the thicket dimension and thicket shatter function of the dual set system $(X,\mathcal{F})^\star$.

Remark 3.11. VC dimension and dual VC dimension are each bounded above by a single exponent of the other [3]. On the other hand, thicket dimension and dual thicket dimension are each bounded above by a double exponent of the other, by the following argument.

Consider the *ladder dimension* of a set system $(X,\mathcal{F})$, which is the largest n such that there exists tuples $(x_1,\dots,x_n)$ from X and $(F_1,\dots,F_n)$ from $\mathcal{F}$, such that $x_i\in F_j\iff i<j$, or ∞ if there exist arbitrarily large such n. The ladder dimension of a set system is clearly equal to the dual ladder dimension, by reversing the roles of the x_i's and F_j's. On the other hand, the thicket and ladder dimensions of a set system are each bounded above by single exponent of the other (see [7], Lemma 6.7.9).

Therefore, the thicket dimension is bounded above by a single exponent of the ladder dimension, which is identical to the dual ladder dimension, and hence is bounded above by a double exponent of the dual thicket dimension. We do not know whether this double-exponent bound is tight.

§4. The Sauer–Shelah dichotomy. We now come to our first central result, that the Sauer–Shelah Lemma, relating the growth of the (usual) shatter function to VC dimension, holds verbatim in the thicket context.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

Definition 4.1. Define the function $\chi : \mathbb{N}^2 \to \mathbb{N}$ by $\chi(n,k) = \binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{k}$. Additionally, define $\chi(n,-1) = 0$, and $\chi(n,\infty) = 2^n$, for any $n \in \mathbb{N}$.

Remark 4.2. Since $\chi(n,k) = \sum_{0 \le i \le k} \binom{n}{i}$, it makes sense to identify $\chi(n,-1)$ with the empty sum, 0. Similarly, it makes sense to identify $\chi(n,\infty)$ with the formal sum $\binom{n}{0} + \binom{n}{1} + \binom{n}{2} + \cdots$. But as $\binom{n}{k} = 0$ when $k > n$, this sum is supported only by the finite part $\binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{n}$, which is 2^n.

Note that for $k < \omega$, $\chi(n,k) \in O(n^k)$.

Theorem 4.3. *For any set system $(X,\mathcal{F})$ and $k \in \{-1\} \cup \omega$,*

$$\dim(X,\mathcal{F}) = \infty \implies \forall n\, \rho(n) = \chi(n,\infty),$$
$$\dim(X,\mathcal{F}) \le k \implies \forall n\, \rho(n) \le \chi(n,k).$$

The first implication is immediate: if $\dim(X,\mathcal{F}) = \infty$, then $(X,\mathcal{F})$ admits B_n for each n, so $\rho(n) = 2^n$. If however $\dim(X,\mathcal{F}) \le k$, then $(X,\mathcal{F})$ forbids B_{k+1}. The conclusion follows by the second sentence of the next theorem.

Theorem 4.4. *For every finite unordered tree T of dimension d, there's a function $f(n) \in O(n^{d-1})$ such that if any set system $(X,\mathcal{F})$ forbids T, then $\rho(n) \le f(n)$. Moreover, when T is B_d, $\rho(n) \le \chi(n,d-1)$.*

Proof. Suppose that $(X,\mathcal{F})$ forbids T. We proceed by induction on the construction of T. If T is the single leaf B_0, then it has dimension 0. If $(X,\mathcal{F})$ forbids T, then $\mathcal{F}$ must be empty, so $\rho(n) = 0$, which is $O(n^{-1})$.

Otherwise, suppose that T has subtrees T_1 and T_2. Suppose their dimensions are d_1 and d_2, and suppose that $f_1(n) \in O(n^{d_1-1})$ and $f_2(n) \in O(n^{d_2-1})$ are given by induction. For $x \in X$, let $\mathcal{F}_x$ be the collection of those sets in $\mathcal{F}$ that include x, and let $\mathcal{F}_{\bar{x}}$ be the collection of those which exclude x. Let $\rho_x(n)$ and $\rho_{\bar{x}}(n)$ be the thicket shatter functions of $(X,\mathcal{F}_x)$ and $(X,\mathcal{F}_{\bar{x}})$ respectively. It is easy to see that for any $x \in X$, $\rho(n) \le \rho_x(n) + \rho_{\bar{x}}(n)$; simply take the X-labeled tree witnessing $\rho(n)$ and observe that the solutions to its leaves can be partitioned into those containing x and those not containing x.

Let $P_i(x)$ express that $(X,\mathcal{F}_x)$ admits T_i, and $Q_i(x)$ express that $(X,\mathcal{F}_{\bar{x}})$ admits T_i. Then by Remark 3.5,

$$(X,\mathcal{F}) \text{ admits } T \iff \exists x \in X \big((P_1(x) \wedge Q_2(x)) \vee (P_2(x) \wedge Q_1(x))\big).$$

Taking the negation, $(X,\mathcal{F})$ forbids T if and only if for all $x \in X$,

$$(\neg P_1(x) \wedge \neg P_2(x)) \vee (\neg P_1(x) \wedge \neg Q_1(x)) \vee$$
$$(\neg Q_2(x) \wedge \neg P_2(x)) \vee (\neg Q_2(x) \wedge \neg Q_1(x)).$$

By induction, this implies, for all $x \in X$,

$$(\rho_x \le f_1 \wedge \rho_x \le f_2) \vee (\rho_x \le f_1 \wedge \rho_{\bar{x}} \le f_1) \vee$$
$$(\rho_{\bar{x}} \le f_2 \wedge \rho_x \le f_2) \vee (\rho_{\bar{x}} \le f_2 \wedge \rho_{\bar{x}} \le f_1),$$

where, e.g., $\rho_x \le f_1$ abbreviates $\forall n \in \omega\, \rho_x(n) \le f_1(n)$. Label the four disjuncts (i)–(iv) respectively.

Suppose that for some $x \in X$, case (ii) holds. Then $\rho(n) \le \rho_x(n) + \rho_{\bar{x}}(n) \le 2f_1(n)$. The right-hand side is $O(n^{d_1-1})$, which is $O(n^{d-1})$, and we are done. Similar

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

reasoning applies if for some $x \in X$, case (iii) holds. If neither (ii) nor (iii) holds for any x, then

$$\forall x \in X\, \big((\rho_x \leq f_1 \wedge \rho_x \leq f_2) \vee (\rho_{\bar{x}} \leq f_2 \wedge \rho_{\bar{x}} \leq f_1)\big),$$

in which case,

$$\forall x \in X\, \big((\rho_x \leq f_1 \vee \rho_{\bar{x}} \leq f_1) \wedge (\rho_x \leq f_2 \vee \rho_{\bar{x}} \leq f_2)\big),$$

so

$$\forall x \in X\, (\rho_x \leq f_1 \vee \rho_{\bar{x}} \leq f_1) \wedge \forall x \in X\, (\rho_x \leq f_2 \vee \rho_{\bar{x}} \leq f_2).$$

We claim that for any function g, $(\forall x)(\rho_x \leq g \vee \rho_{\bar{x}} \leq g) \implies \rho \leq \int g$, where $\int g$ is defined by $(\int g)(n) = 1 + \sum_{k<n} g(k)$. If so ρ, would be bounded above by both $\int f_1$ and $\int f_2$, which are $O(n^{d_1})$ and $O(n^{d_2})$ respectively.[2] If $d_1 = d_2$, then $d = d_1 + 1$, and ρ would be bounded by a function in $O(n^{d_1})$, which is $O(n^{d-1})$. If $d_1 \neq d_2$, then suppose without loss of generality that $d_1 < d_2$. Then ρ would be bounded above by $\int f_1$, a function in $O(n^{d_1})$, which is again $O(n^{d-1})$, which completes the proof.

It remains to justify the claim. We show that $\rho(n) \leq (\int g)(n)$ by induction on n. If $n = 0$, $\rho(n) \leq 1 \leq (\int g)(n)$. Otherwise, consider the labeled tree w witnessing $\rho(n)$, and let r be the label of its root. By hypothesis either ρ_r or $\rho_{\bar{r}}$ is bounded above by g. Therefore, either the number of solutions to the left of the root or the number of solutions to the right must be bounded by $g(n-1)$. The number of solutions in the remaining half is bounded by $\rho(n-1)$, which by induction is at most $(\int g)(n-1)$. Therefore the total number of solutions is at most $g(n-1) + (\int g)(n-1) = (\int g)(n)$.

This concludes the proof of the first sentence. Finally, we must show that if T is B_d, $\rho(n)$ is bounded by $\chi(n, d-1)$. The base case $d = 0$ is the same as before. Otherwise, if $(X, \mathcal{F})$ forbids B_d, then for any $x \in X$, either $(X, \mathcal{F}_x)$ forbids B_{d-1} or $(X, \mathcal{F}_{\bar{x}})$ forbids B_{d-1}. Therefore, for any $x \in X$, $\rho_x \leq \chi(\cdot, d-1)$, or $\rho_{\bar{x}} \leq \chi(\cdot, d-1)$. Hence, by the above claim, $\rho \leq (\int \chi)(\cdot, d-1)$, which is $\chi(\cdot, d)$.[3] ⊣

The VC density of a set system is defined to be the least exponent bounding the growth rate of the shatter function. Analogously,

Definition 4.5. *Thicket density* $\text{dens}(X, \mathcal{F})$ is defined by

$$\inf\{c \in \mathbb{R} : \rho(n) \in O(n^c)\},$$

where $\rho(n)$ is the thicket shatter function of $(X, \mathcal{F})$. In case $\rho(n) \in O(n^c)$ for all $c \in \mathbb{R}$, $\text{dens}(X, \mathcal{F}) = -\infty$, and in case $\rho(n) \in O(n^c)$ for no $c \in \mathbb{R}$, $\text{dens}(X, \mathcal{F}) = \infty$. The *dual thicket density* is defined with the dual thicket shatter function $\rho^\star$ instead of ρ.

Remark 4.6. By Theorem 4.3, if $\dim(X, \mathcal{F}) = k < \infty$, then $\rho(n) \leq \chi(n, k)$, which is $O(n^k)$. Therefore, $\text{dens}(X, \mathcal{F}) \leq \dim(X, \mathcal{F})$ for any set system $(X, \mathcal{F})$, as $\mathbb{R} \cup \{-\infty, \infty\}$-valued quantities.

[2] The fact that $g \in O(n^p) \implies \int g \in O(n^{p+1})$ follows, for example, from elementary calculus, or from *Faulhaber's formula*, which expresses the sum $\sum_{k<n} k^p$ as a polynomial in n of degree $p+1$.

[3] This can be shown by induction on n and d. When both $n, d > 0$, $(\int \chi)(n, d-1) = (\int \chi)(n-1, d-1) + \chi(n-1, d-1)$, which is $\chi(n-1, d) + \chi(n-1, d-1)$ by induction. Observe that the latter quantity is $\chi(n, d)$, by Pascal's identity.

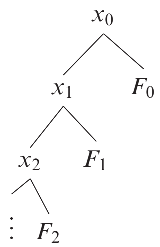

FIGURE 3. The unordered tree of Lemma 4.7.

LEMMA 4.7. *The thicket density of* $(X,\mathcal{F})$ *is* $-\infty$, 0, *or* ≥ 1, *depending on whether the cardinality of* $\mathcal{F}$ *modulo extensionality is zero, positive and finite, or infinite.*

PROOF. If $\mathcal{F}$ is empty, then $\rho(n) = 0$, and $\mathrm{dens}(X,\mathcal{F}) = -\infty$. If $\mathcal{F}$ has nonempty but has finitely many elements modulo extensionality, say at most B, then for all $n\rho(n) \leq B$. Therefore, $\mathrm{dens}(X,\mathcal{F}) = 0$. On the other hand, if $\mathcal{F}$ has infinitely many elements modulo extensionality, then we can extract a sequence $(x_i, F_i)_{i<\omega}$ such that for any $i < j$, F_i and F_j agree on x_h for $h < i$, but disagree on x_i. (Given $\mathcal{F}$, pick some $x \in X$ such that both $\mathcal{F}_x$ and $\mathcal{F}_{\bar{x}}$ are nonempty. One of them, call it $\mathcal{F}^\star$ must be infinite; pick some F that comes from the other, $\mathcal{F}^\dagger$. Replace $\mathcal{F}$ by $\mathcal{F}^\star$ and repeat.) Then the unordered tree pictured in Figure 3 has a solution, and, by truncating the tree at depth n for each n, we observe by Remark 3.10 that $\rho(n) \geq n+1$, hence $\mathrm{dens}(X,\mathcal{F}) \geq 1$. ⊣

An alternate formulation of density. Theorem 4.3 says roughly that, for set systems of finite thicket dimension, trees which are balanced cannot have very many realized leaves. Then it stands to reason that admissible trees, i.e., those with every leaf realized, must be far from balanced. This idea yields another way to define thicket density.

DEFINITION 4.8. Let $\sigma(n)$ be the minimum depth of any finite X-labeled tree, at least n of whose leaves have solutions.

If $(X,\mathcal{F})$ has infinite thicket density, then $\sigma(n) = \lfloor \log_2(n) \rfloor$, as witnessed by a balanced binary tree. If $(X,\mathcal{F})$ has zero density, then $\sigma(n)$ is undefined for arbitrarily large n. The interesting case is when $(X,\mathcal{F})$ has finite and positive density, in which case $\sigma(n)$ is bounded below by n^ε for some $\varepsilon > 0$. In fact, the rate of growth of σ is the inverse of the density:

LEMMA 4.9. *For any set system* $(X,\mathcal{F})$ *of finite positive density* δ, *if* $\varepsilon = \sup\{c \in \mathbb{R} : \sigma(n) \in \Omega(n^c)\}$, *then* $\varepsilon = \frac{1}{\delta}$.

PROOF. Fix n. Let T be any finite X-labeled binary tree, at least n of whose leaves have solutions, and let Δ be the depth of T. By Remark 3.10, $n \leq \rho(\Delta)$. Since T is arbitrary, by taking the T with the least possible Δ, we obtain $n \leq \rho(\sigma(n))$, for all $n \in \omega$. This implies that $\varepsilon \geq \frac{1}{\delta}$.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

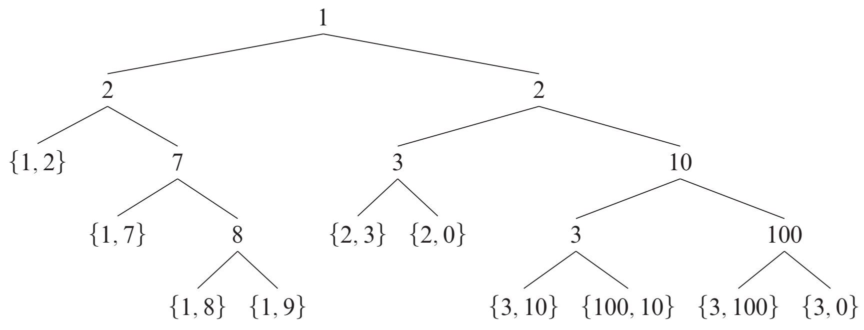

FIGURE 4. A $\mathbb{N}$-labeled tree with all leaves realized in $\binom{\mathbb{N}}{2}$. The depth of such trees grows like $\Omega(\sqrt{n})$, where n is the size.

On the other hand, fix Δ. Let T be any finite, balanced X-labeled binary tree of depth Δ, and let n be the number of leaves of T with solutions. Then $\Delta \geq \sigma(n)$, by definition of σ. Since T is arbitrary, by taking T with the greatest possible n, we obtain $\Delta \geq \sigma(\rho(\Delta))$, for all $\Delta \in \omega$. This implies that $\delta \leq \frac{1}{\varepsilon}$, which completes the proof (Figure 4). ⊣

§5. Rank. In this section, we establish an equivalence between thicket density and "Shelah's local ω-rank" $R(p,\{\varphi\},\omega)$, the local analogue of Morley rank.[4] From a model-theoretic perspective, the significance of this result is probably reversed, i.e., we identify $R(p,\{\varphi\},\omega)$ with the thicket density of a particular set system.

The arguments in the next two sections involve moving between lower bounds for ranks, the admissibility of certain trees, and the consistency of certain theories. They recall Shelah's proof of Unstable Formula Theorem (Theorem 2.2 of [11]). In this section and henceforth, we use `typewriter` script to distinguish syntactic variables, e.g., $\mathtt{x}$, from values, e.g., $x \in X$ or $a \in M$.

Concretely, given any partitioned first-order formula $\varphi(\mathtt{x},\mathtt{y})$, and for any model M of a complete theory T, define a set system $(M^{\mathtt{x}}, M^{\mathtt{y}}, \in)$ by $a \in b \iff M \models \varphi(a,b)$. Then, roughly speaking, we can calculate the rank $R(\mathtt{x}=\mathtt{x},\{\varphi\},\omega)$ of a finite φ-type modulo T using the thicket shatter function of the set system $(M^{\mathtt{x}}, M^{\mathtt{y}}, \in)$. Since the thicket shatter function depends only on T and not the particular choice of M, we can carry out this calculation in *any* model $M \models T$, not only a sufficiently saturated one.

In this section, fix a set system $(X,\mathcal{F})$, and a sufficiently saturated model M of $\mathrm{Th}(X,\mathcal{F},\in)$ (so that we can calculate ω-rank). Let M_X and $M_{\mathcal{F}}$ be the two sorts of M. Let φ be the formula $\mathtt{x} \in \mathtt{F}$. By a *$\varphi$-formula*, we mean a formula of the form $\varphi(a,\mathtt{F})$ or $\neg\varphi(a,\mathtt{F})$ for some $a \in M_X$. Whenever we assert that some sentence holds, we always mean relative to the model M.

[4]We do not know of a commonly accepted name for this rank. However, it is equivalent to $R(p,\{\varphi\},\infty)$, and the Cantor–Bendixson rank of the space of φ-types.

By a *finite φ-type*, we mean a conjunction of φ-formulas, including the *empty* conjunction $\top$. Two finite φ-types p and q are *contradictory* in case there exists some $a \in M_X$ such that one of $\{\varphi(a,\mathtt{F}), \neg\varphi(a,\mathtt{F})\}$ occurs as a conjunct in p, and the other occurs as a conjunct in q. In this case, we say that p and q *disagree on* $\varphi(a,\mathtt{F})$. By $p(\mathcal{F})$ we mean the subfamily of $\mathcal{F}$ satisfying the type p.

Definition 5.1. For any unordered tree T, let L be its set of leaves and N be its set of non-leaves. Let $\mathcal{L}_T$ be the signature $\{\in\}$ expanded by a new set of constant symbols $\{a_u : u \in N\}$ of sort X, and $\{b_v : v \in L\}$ of sort $\mathcal{F}$. Given in addition a finite φ-type p, define a first-order $\mathcal{L}_T$-theory Adm_p^T by:

1. $p(b_v)$ for any $v \in L$,
2. $\varphi(a_u,b_v) \not\leftrightarrow \varphi(a_u,b_w)$ if $v,w \in L$, $u \in N$, and $v \perp_u w$, and
3. $\varphi(a_u,b_v) \leftrightarrow \varphi(a_u,b_w)$ if $v,w \in L$, $u \in N$, and $v \sim_u w$.

Remark 5.2. The following hold of Adm_p^T:

- If $(X, p(\mathcal{F}))$ admits T, then $\mathrm{Th}(X,\mathcal{F})$ is consistent with Adm_p^T. If T is finite, then the converse holds as well.
- Since M is sufficiently saturated, the consistency of $\mathrm{Th}(X,\mathcal{F}) \cup Adm_p^T$ is equivalent the admissibility of T in $(M_X, p(M_{\mathcal{F}}))$. (Concretely: the existence of injections $u \mapsto a_u : N \to M_X$ and $v \mapsto b_v : L \to M_{\mathcal{F}}$ such that properties (1)–(3) of Definition 5.1 hold in M.)

Lemma 5.3. *Suppose that T is a finite-dimensional unordered tree and that there is an embedding of the unordered tree S into T.*[5] *Then if* $\mathrm{Th}(X,\mathcal{F})$ *is consistent with Adm_p^T, it is consistent with Adm_p^S.*

Proof. First observe that for every vertex $v \in T$, the subtree rooted at v contains some leaf. For otherwise (since each non-leaf has two children) T would have a complete infinite binary subtree, and not be finite-dimensional. Fix an embedding $\imath : S \to T$, and let $\jmath$ map leaves of S to leaves of T, such that for each leaf $\ell \in S$, $\jmath(\ell)$ is contained in the subtree rooted at $\imath(\ell)$. Then notice that for any non-leaf s and leaf ℓ in S, $s \prec \ell$ in S iff $\imath(s) \prec \jmath(\ell)$ in T. Moreover, if ℓ' is another leaf of S, then $\ell \perp_s \ell'$ in S iff $\jmath(\ell) \perp_{\imath(s)} \jmath(\ell')$ in T. Similarly, $\ell \sim_s \ell'$ in S iff $\jmath(\ell) \sim_{\imath(s)} \jmath(\ell')$ in T.

Suppose that T is admissible in $(M_X, p(M_{\mathcal{F}}))$, witnessed by the labeling $u \mapsto a_u$, $v \mapsto b_v$. For any non-leaf $s \in S$, label it by $a_{\imath(s)}$, and for any leaf $t \in S$, label it by $b_{\jmath(t)}$. Since the relations $\perp$ and $\sim$ are preserved by the map $\imath$ on non-leaves and $\jmath$ on leaves, properties (2) and (3) of Definition 5.1 carry over from T to S. Since $\jmath$ maps leaves to leaves, property (1) of Definition 5.1 carries over from T to S. Hence, S is admissible in in $(M_X, p(M_{\mathcal{F}}))$. ⊣

Next, we define the rank $R(p,\{\varphi\},\omega)$, for finite φ-types p. Here we follow the definition and notation of Pillay [10], who writes $R^{\varphi}_{\aleph_0}$, but we omit the cardinality $\aleph_0$, and just write R^{φ}.

[5] cf. Definition 2.5.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

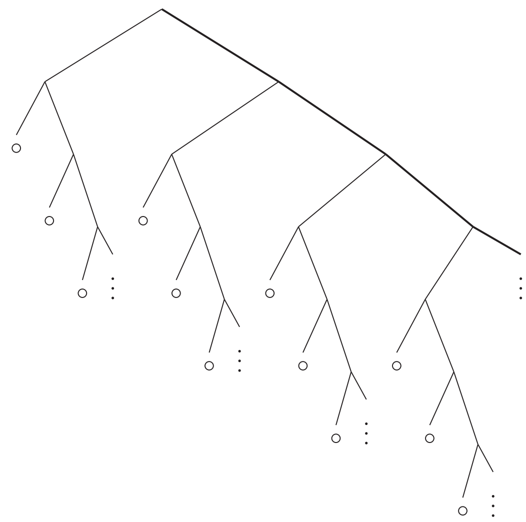

FIGURE 5. The two-branching tree from Definition 5.7. The *spine* is indicated by the thick edge. The vertices of any T_{k+1} can be partitioned into the vertices on the spine, plus countably many copies of T_k.

DEFINITION 5.4. For any finite φ-type $p(\mathtt{F})$, let

- $R^\varphi(p) \geq 0$ if p is consistent, i.e., there exists some $b \in M_{\mathcal{F}}$ such that $p(b)$.
- $R^\varphi(p) \geq \alpha + 1$ in case there is a pairwise contradictory family of finite φ-types $\{p_i : i < \omega\}$, such that for each i, $R^\varphi(p \wedge p_i) \geq \alpha$.
- For limit ordinal α, $R^\varphi(p) \geq \alpha$ just in case $R^\varphi(p) \geq \beta$ for all $\beta < \alpha$.

We say $R^\varphi(p) = \alpha$ in case $R^\varphi(p) \geq \alpha$ but $R^\varphi(p) \not\geq \alpha + 1$, $R^\varphi(p) = -\infty$ if $R^\varphi(p) \not\geq 0$, and $R^\varphi(p) = \infty$ in case $R^\varphi(p) \geq \alpha$ for all α.

REMARK 5.5. If $R^\varphi(p) \geq \omega$, then $R^\varphi(p) = \infty$. (Pillay [10] Exercise 6.53) Hence R^φ takes values in $\omega \cup \{-\infty, \infty\}$.

In fact, we can rearrange the quantifiers in the second clause to be a little bit stronger (Figure 5):

LEMMA 5.6. *Suppose that $\{p_i : i < \omega\}$ is a sequence of pairwise contradictory finite φ-types. Then there is an infinite set $S \subseteq \omega$ such that for any $r \in S$, there exists $a \in M_X$, such that for any $s > r$ in S, p_r and p_s disagree on $\varphi(a, \mathtt{F})$.*

PROOF. For any infinite set $S \subseteq \omega$, let m be its least element. Since p_m is finite, there must be some $a \in M_X$ and an infinite subset $S' \subseteq S \setminus \{m\}$ such that p_m disagrees with p_n on $\varphi(a, \mathtt{F})$ for any $n \in S'$. Let $S_0 = \omega$, and for each $i < \omega$, obtain S_{i+1} from S_i in this manner. Obtain m_i and a_i from S_i as above. Then $S_0 \supset S_1 \supset S_2 \supset \cdots$, $m_0 < m_1 < m_2 < \cdots$, and $m_i \in S_j \iff i \geq j$. Moreover, for every $i < j$, p_{m_i} and p_{m_j} disagree on $\varphi(a_i, \mathtt{F})$. Therefore, we may take $S = \{m_0, m_1, m_2, \ldots\}$. ⊣

DEFINITION 5.7. The 0-branching tree T_0 is a single leaf. For $k \geq 1$, the k-branching tree T_k is the unordered binary tree with subtrees T_k and T_{k-1}.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

REMARK 5.8. T_k has dimension k, and is isomorphic to the (unordered) tree $\{\sigma \in 2^{<\omega} : \sigma \text{ has at most } k \text{ many zeros}\}$. Therefore, the number of vertices in T_k at depth n is $\chi(n,k)$, since $\chi(n,k) = \sum_{i \le k} \binom{n}{i}$ is the number of binary strings of length n with at most k many zeros.

THEOREM 5.9. *$R^\varphi(p) \ge k$ if and only if* $\mathrm{Th}(X,\mathcal{F}) \cup Adm_p^{T_k}$ *is consistent.*

PROOF. We work in M, by Remark 5.2, and induct on k. If $k = 0$, then $R^\varphi(p) \ge k$ iff there exists some $b \in M_{\mathcal{F}}$ such that $p(b)$, iff the tree T_0 is admissible in $(M_X, p(M_{\mathcal{F}}))$, iff $\mathrm{Th}(X,\mathcal{F}) \cup Adm_p^{T_0}$ is consistent.

Otherwise, suppose that $R^\varphi(p) \ge k+1$, and let $\{p_i\}_{i<\omega}$ witness this. By Lemma 5.6, there is an infinite set $S \subseteq \omega$ such that for any $r \in S$, there is some $a \in M_X$, such that for any $s > r$ in S, p_r and p_s disagree on $\varphi(a,\mathtt{F})$. Let $\{q_i\}_{i<\omega}$ enumerate the set $\{p_i\}_{i\in S}$ in order. Let $a^{(i)}$ witness the φ-formula distinguishing q_i from q_j, for $j > i$.

Let N and L be the sets of non-leaves and leaves, respectively, of T_{k+1}. The tree T_{k+1} can be partitioned into a single infinite spine plus countably many copies of T_k.[6] Let $N = \bigcup_{i<\omega} N_i \cup N_s$, where N_i is the set of non-leaves of the i-th copy of T_k, and N_s are the vertices down the spine. Let $L = \bigcup_{i<\omega} L_i$, where L_i is the set of leaves of the i-th copy of T_k.

For each $i < \omega$, we have $R^\varphi(p \wedge q_i) \ge k$. Therefore, by induction, there are injections $u \mapsto a_u : N_i \to M_X$ and $v \mapsto b_v : L_i \to M_{\mathcal{F}}$ satisfying properties (1)–(3) for each $p \wedge q_i$. By taking the union over all i, we have a map $v \mapsto b_v : L \to M_{\mathcal{F}}$. Similarly, we can get a map $u \mapsto a_u : N \to M_X$ if we specify a_u for $u \in N_s$. But for $u \in N_s$ of distance i from the root, simply let $a_u = a^{(i)}$.

For any $v \in L$, b_v satisfies $p \wedge q_i$ for some i; in particular it satisfies p. Thus we have verified (1), and it remains to verify (2) and (3). Fix two leaves $v,w \in L$ and a common ancestor $u \in N$.

- If for some i, $v,w \in L_i$ and $u \in N_i$, then the relations $v \perp_u w$, $v \sim_u w$ are identical in the ambient tree T_{k+1} and the i-th copy of T_k, hence (2) and (3) are inherited from the given maps $N_i \to M_X$ and $L_i \to M_{\mathcal{F}}$.
- If for some i, $v,w \in L_i$ and $u \notin N_i$, then $v \sim_u w$, and u must be a vertex in N_s of distance at most i from the root. Therefore, a_u is $a^{(j)}$ for some $j \le i$. Since b_v and b_w both satisfy q_i, and q_i contains either $\varphi(a^{(j)},\mathtt{F})$ or its negation for each $j \le i$, b_v and b_w must agree on $\varphi(a_u,\mathtt{F})$.
- If for some $i < j$, $v \in L_i$ and $w \in L_j$, then u must be some vertex on the spine of distance at most i from the root; moreover, $q_i(b_v)$ and $q_j(b_w)$. Therefore, b_v and b_w disagree on $\varphi(a^{(i)},\mathtt{F})$, and agree on $\varphi(a^{(i')},\mathtt{F})$ for all $0 \le i' < i$. But, $a_u = a^{(i)}$ just in case $v \perp_u w$, and $a_u = a^{(i')}$ for some $0 \le i' < i$ just in case $v \sim_u w$. This concludes the forward direction.

In the other direction, suppose that we have injections $u \mapsto a_u : N \to M_X$ and $v \mapsto b_v : L \to M_{\mathcal{F}}$ satisfying (1)–(3). Let N_s, N_i, and L_i be as above, and let $a^{(i)}$ be the sequence of vertices along the spine. For $i < \omega$, define $p_i = \{\varphi^\star(a^{(i)},\mathtt{F})\} \cup \{\neg\varphi^\star(a^{(i')},\mathtt{F}) : i' < i\}$, where $\varphi^\star(a^{(i)},\mathtt{F})$ is either $\varphi(a^{(i)},\mathtt{F})$ or $\neg\varphi(a^{(i)},\mathtt{F})$, depending on which one the leaves in L_i satisfy (Figure 6).

[6] Under the identification of T_{k+1} with the set of binary strings i with at most $k+1$ zeros, the spine consists of $1^{<\omega}$, and the countably many copies of T_k are rooted at 1^i0, for $i < \omega$.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

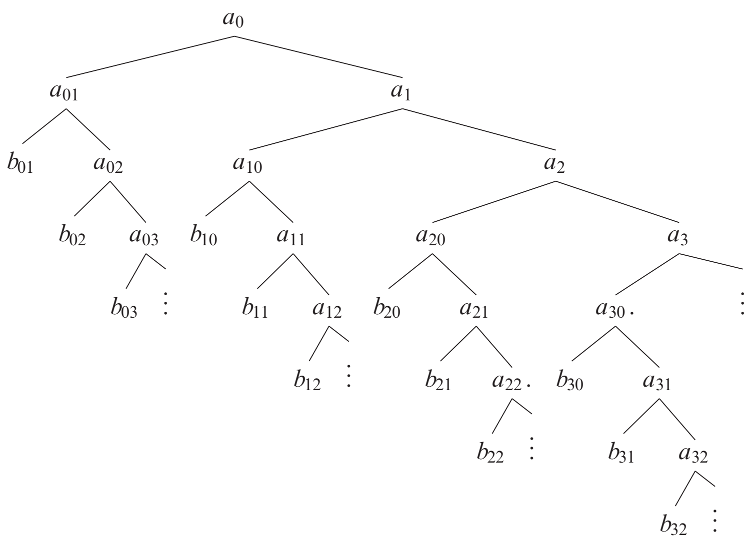

FIGURE 6. By Theorem 5.9, $R^{\varphi}(p) \geq 2$ is equivalent to the admissibility of the unordered tree T_2 in M, such that every b_{ij} satisfies p.

For each i, by restricting the maps $u \mapsto a_u$ and $v \mapsto b_v$ to N_i and L_i respectively, we get injections which inherit (2) and (3) from the original function. Towards (1), notice that for each $v \in L_i$, b_v satisfies $p \wedge p_i$, by definition of p_i. Hence, by induction, $R^{\varphi}(p \wedge p_i) \geq k$. Since for each $i < j$, p_i and p_j disagree on $\varphi(a^{(i)}, \mathtt{F})$, we have that $R^{\varphi}(p) \geq k+1$, which concludes the proof. ⊣

THEOREM 5.10. *The thicket density of* $(X, p(\mathcal{F}))$ *is identical to* $R^{\varphi}(p)$, *as* $\mathbb{R}_{\geq 0} \cup \{-\infty, \infty\}$*-valued quantities.*

PROOF. We show that both quantities are bounded above by each other.

For some $0 \leq k < \omega$, suppose that $R^{\varphi}(p) \geq k$. By Theorem 5.9, $\mathrm{Th}(X, \mathcal{F})$ is consistent with $Adm_p^{T_k}$. By Remark 5.2, for any finite subtree S of T_k, Adm_p^S is consistent, so $(X, p(\mathcal{F}))$ admits S. But if we obtain S by truncating T_k to depth n, then S has $O(n^k)$ leaves, thus ensuring $\mathrm{dens}(X, \mathcal{F}) \geq k$. Hence $R^{\varphi}(p) \leq \mathrm{dens}(X, p(\mathcal{F}))$.

Conversely, suppose that for some $0 \leq k < \omega$, $R^{\varphi}(p) \not\geq k$. By Theorem 5.9, $Adm_p^{T_k}$ is inconsistent with $\mathrm{Th}(X, \mathcal{F})$. By compactness, some finite fragment of $Adm_p^{T_k}$ is inconsistent with $\mathrm{Th}(X, \mathcal{F})$, and hence there exists a finite subtree S of T_k such that Adm_p^S is inconsistent with $\mathrm{Th}(X, \mathcal{F})$. By Remark 5.2, $(X, p(\mathcal{F}))$ forbids S. Since T_k has dimension k, S has dimension at most k. By Theorem 4.4, $\mathrm{dens}(X, p(\mathcal{F})) \leq k-1$. Hence $\mathrm{dens}(X, p(\mathcal{F})) \leq R^{\varphi}(p)$. ⊣

As an immediate corollary, we deduce that thicket density is integer-valued, in contrast to VC density. This fact was first proven by James Freitag and Dhruv Mubayi, using elementary combinatorics, and a very short elementary proof has

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

been found by Ross Berkowitz. (Both of these results are unpublished and were communicated to us in person.)

§6. Degree. Many model-theoretic ranks admit an associated notion of *degree* (sometimes called *multiplicity*); roughly, given a formula p, this is the maximum number of pairwise contradictory extensions of p which each have rank no less than p. For any fixed rank, this degree must be absolutely bounded by some cardinal κ, where the rank of a formula is at least $\alpha+1$ if there are at least κ many pairwise contradictory extensions of rank α. For example, we can define the degree D associated with the rank R^φ as follows.

DEFINITION 6.1. For any finite φ-type p, let $D(p)$ be the greatest integer $n<\omega$ such that there exist pairwise contradictory finite φ-types $p_1,\dots,p_n$, such that for each $1\le i\le n$, $R^\varphi(p\wedge p_i)=R^\varphi(p)$.[7]

In this section, we develop a notion of degree that seems natural and appropriate for thicket density. Like in §5, we fix a set system $(X,\mathcal{F})$ and a sufficiently saturated model M of $\mathrm{Th}(X,\mathcal{F})$. For a finite φ-type p to have *parameters from X* means every conjunct of p is of the form $\varphi(x,\mathtt{F})$ or $\neg\varphi(x,\mathtt{F})$ for some $x\in X$. For such a type p, let $\mathrm{dens}(p)$ abbreviate $\mathrm{dens}(X,p(\mathcal{F}))$, where $p(\mathcal{F})=\{F\in\mathcal{F}:p(F)\}$. More generally, for any formula $\psi(\mathtt{F})$, let $\mathrm{dens}(\psi)$ abbreviate $\mathrm{dens}(X,\psi(\mathcal{F}))$.[8]

DEFINITION 6.2. Given a vertex v in an X-labeled tree T, define the finite φ-type $p_v(\mathtt{F})$ to be the conjunction of literals $\varphi(x_u,\mathtt{F})$ for $u\prec_L v$, and $\neg\varphi(x_u,\mathtt{F})$ for $u\prec_R v$. Then p_v is a finite φ-type with parameters from X, and if v is a leaf, F solves v if and only if $p_v(F)$.

DEFINITION 6.3. Let p be a finite φ-type with parameters from X. An X-labeled tree T *factors* p in case $\mathrm{dens}(p\wedge p_v)=\mathrm{dens}(p)$, for each leaf v of T. Furthermore, T *irreducibly factors* p if, in addition, no proper extension T' of T factors p. We say that p is *irreducible* if it is irreducibly factored by a single leaf. We say that $(X,\mathcal{F})$ is irreducible in case the empty type is.

LEMMA 6.4. *For every finite φ-type p with parameters from X, if $0\le\mathrm{dens}(p)<\infty$, then p is irreducibly factored by some finite tree.*

PROOF. Fix the domain X. For each finite type $p\in S_\varphi(\mathcal{F})$, we build an X-labeled tree $\Pi(p)$ irreducibly factoring p via a nondeterministic construction.

- If p is irreducible, $\Pi(p)$ is a single leaf.
- Otherwise, let $x\in X$ satisfy $\mathrm{dens}(p\wedge\varphi(x,\mathtt{F}))=\mathrm{dens}(p\wedge\neg\varphi(x,\mathtt{F}))=\mathrm{dens}(p)$. Let $\Pi(p)$ have root x and left and right subtrees $\Pi(p\wedge\varphi(x,\mathtt{F}))$ and $\Pi(p\wedge\neg\varphi(x,\mathtt{F}))$ respectively.

We claim that this construction terminates after finitely many steps. For otherwise, there would be an infinite branch $x_0,x_1,x_2,\dots$ and an infinite sequence $p=p_0\subset p_1\subset p_2\subset\cdots$ such that for each $i<\omega$, either $p_{i+1}=p_i\wedge\varphi(x_i,\mathtt{F})$ or $p_{i+1}=p_i\wedge\neg\varphi(x_i,\mathtt{F})$,

[7] If there were arbitrarily large such n, the there must exist an infinite family p_i of pairwise contradictory φ-types, such that for each i, $R(p\wedge p_i)=R(p)$, contradicting Definition 5.4.

[8] We consider the density of one non-type, namely, $\mathrm{dens}(p\wedge(p_v\,\Delta\,p_w))$ in Lemma 6.5.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

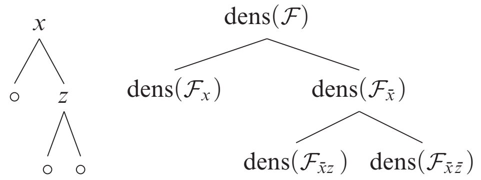

FIGURE 7. An illustration suggesting $\Pi(\top)$ (on the left), and the densities of the corresponding set systems (on the right). Each of these densities is the same, but the set systems labeling the leaves are irreducible, and cannot be "split" into subsystems of the same density by any element of X.

and $\text{dens}(p_i) = \text{dens}(p_i \wedge \varphi(x_i, \mathtt{F})) = \text{dens}(p_i \wedge \neg\varphi(x_i, \mathtt{F}))$. Let $p_i^\dagger$ be the extension of p_i that is not p_{i+1}. Then, for each $i < \omega \text{dens}(p_i^\dagger) = \text{dens}(p)$, and for each $i < j < \omega$, $p_i^\dagger$ and $p_j^\dagger$ agree on $\varphi(x_h, F)$ for $h < i$, and disagree on $\varphi(x_i, \mathtt{F})$.

Let $k = \text{dens}(p)$. By Theorems 5.9 and 5.10, for each $i < \omega$, we can find a consistent labeling of T_k by elements of M, such that each leaf satisfies $p_i^\dagger$. By stringing these trees along a spine labeled by the sequence $x_0, x_1, x_2, \ldots$, we get a consistent labeling of T_{k+1} by elements of M, such that each leaf satisfies p. But this implies that $\text{dens}(p) \geq k+1$, a contradiction (Figure 7). ⊣

We now show that even though there may not be a unique tree that irreducibly factors a set system, the number of vertices in any such tree is the same. Even stronger, the partition induced by such a tree is unique up to rearrangement by pieces of strictly smaller density. This is analogous to the situation for, e.g., Morley rank.

LEMMA 6.5. *Let p be a finite type with parameters from X. Suppose $\delta = \text{dens}(p)$ is finite and nonnegative, and suppose T_1 and T_2 are X-labeled trees that each irreducibly factor p. Then T_1 and T_2 have the same number of leaves. A fortiori, there is a bijection between the leaves of T_1 and the leaves of T_2, such that $\text{dens}(p \wedge (p_v \,\Delta\, p_w)) < \delta$ for any bijective pair (v, w).*

PROOF. For $i \in \{1, 2\}$, let L_i be the set of leaves of T_i. Since T_1 and T_2 each factor p, $\text{dens}(p \wedge p_v) = \text{dens}(p \wedge p_w) = \delta$ for each $v \in L_1$ and $w \in L_2$.

For any fixed $v \in L_1$, the types $\{p \wedge p_v \wedge p_w : w \in L_2\}$ partition $p \wedge p_v$. Hence $\text{dens}(p \wedge p_v) = \max\{\text{dens}(p \wedge p_v \wedge p_w) : w \in L_2\}$.[9] Hence, for some $w \in L_2$, $\text{dens}(p \wedge p_v \wedge p_w) = \delta$. However, such a w must be unique; if $\text{dens}(p \wedge p_v \wedge p_w) = \text{dens}(p \wedge p_v \wedge p_{w'}) = \delta$, then T_1 does not irreducibly factor p: the leaf v can be replaced by a non-leaf labeled by the least common ancestor of w and w' in T_2.

Reasoning symmetrically, for any $w \in L_2$, there is a unique $v \in L_1$ such that $\text{dens}(p \wedge p_v \wedge p_w) = \delta$. Hence, the relation $R(v, w) \iff \text{dens}(p \wedge p_v \wedge p_w) = \delta$ is the graph of a bijection, so $|L_1| = |L_2|$. Furthermore, for any bijective pair (v, w),

$$\text{dens}(p \wedge (p_v \,\Delta\, p_w)) = \max\{\text{dens}(p \wedge p_v \wedge p_{w'}), \text{dens}(p \wedge p_{v'} \wedge p_w) : v' \neq v, w' \neq w\},$$

[9] In general, $\text{dens}(p \vee q) = \max\{\text{dens}(p), \text{dens}(q)\}$. The thicket shatter function $\rho_{p \vee q}$ on is bounded below by both ρ_p and ρ_q, and bounded above by $\rho_p + \rho_q$.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

again, since the types $\{p \wedge p_v \wedge p_{w'}, p \wedge p_{v'} \wedge p_w\}_{v \neq v', w \neq w'}$ partition $p \wedge (p_v \,\Delta\, p_w)$. Hence, $\text{dens}(p \wedge (p_v \,\Delta\, p_w))$ must be strictly less than δ, since we maximize over finitely many densities all strictly less than δ. ⊣

Hence, we can now define a notion of degree for thicket density.

DEFINITION 6.6. For any finite φ-type p with parameters from X, the *thicket degree* $\deg(p)$ is the number of leaves of any tree irreducibly factoring p. The degree $\deg(X,\mathcal{F})$ is defined to be $\deg(\top)$, the degree of the empty conjunction.

REMARK 6.7. Given any tree T irreducibly factoring p, the types p_v are pairwise contradictory as v varies over the leaves of T, and moreover $R^{\varphi}(p \wedge p_v) = R^{\varphi}(p)$ for any leaf v. Hence $\deg(p) \leq D(p)$.

REMARK 6.8. By a finite $\varphi^{\star}$-type q with parameters from $\mathcal{F}$, we mean a finite conjunction of formulas of the form $\varphi(\mathrm{x},F)$ and $\neg\varphi(\mathrm{x},F)$, where F ranges over $\mathcal{F}$. We can define the *dual degree* , $\deg^{\star}(q)$, by switching the roles of X and $\mathcal{F}$ throughout. Concretely, the dual degree of q is the number of leaves of any $\mathcal{F}$-labeled binary tree T, such that $\text{dens}^{\star}(q \wedge q_v) = \text{dens}^{\star}(q)$ for any leaf v of T, and T is maximal with respect to this property. (The dual degree of q is simply the degree of q in the dual set system.)

REMARK 6.9. Even though thicket density is equivalent to the rank R, the thicket degree can differ from the degree D associated with the rank R. It suffices, and is easier, to exhibit a difference between the corresponding dual quantities. Let X be any infinite set, and let $\mathcal{F}$ be a partition of X with arbitrarily large finite sets, but no infinite set. Then $\text{dens}^{\star}(X,\mathcal{F}) \geq 1$, since X is infinite, but for any $F \in \mathcal{F}$, $\text{dens}^{\star}(F, \mathcal{F} \restriction F) = 0$, as F is finite (cf. Lemma 4.7). Hence, $(X,\mathcal{F})$ is (dually) irreducible: any $\mathcal{F}$-labeled tree that irreducibly factors $(X,\mathcal{F})$ must be a single leaf, and hence $\deg(X,\mathcal{F}) = 1$. On the other hand, using compactness, we can find an element $b \in M_{\mathcal{F}}$ that has infinitely many members and infinitely many non-members among elements of M_X. Therefore, $R^{\star}(\varphi(\mathrm{x},b))$ and $R^{\star}(\neg\varphi(\mathrm{x},b))$ are both at least 1, so $D^{\star}(\top) \geq 2$.

§7. Conclusion and future work. We have established a relationship between a measure of asymptotic growth of certain finite objects with the ordinal rank of an infinite object, a common phenomenon relating infinitary and finitary combinatorics. There are several lines of inquiry that our work raises. For example, from a technical standpoint, we do not know what other information about a set system is contained in its thicket shatter function. We suspect that other invariants (like the *leading coefficient*) might encode something interesting. Furthermore, there are many identities that the rank R^{φ} is known to satisfy, and it would be interesting to see if they might admit a purely combinatorial proof.

The similarity between the Sauer–Shelah Lemma and the present "thicket" version naturally raises the question of whether they both share a general setting. Chase and Freitag [5] answer this question positively by formulating a shatter function for *op-rank* of Guingona and Hill [6], which interpolates Shelah's two-rank with VC dimension, and establishing a dichotomy which interpolates both versions of the Sauer–Shelah Lemma. Most questions about the corresponding notion of density remain open, for example, what values it takes "between" the VC and thicket cases.

https://doi.org/10.1017/jsl.2020.55 Published online by Cambridge University Press

Finally, as we mentioned in §1, this work was originally inspired by reading Tiuryn's paper [12] separating deterministic from nondeterministic dynamic logic. Several branches of computer science, such as dynamic logic, descriptive complexity theory, and program schematology, are concerned with *programs that operate over first-order structures*. A fundamental invariant of a deterministic, sequential program its underlying *decision tree*, which encodes the sequence of tests made on the input data during execution of the program. Many algorithmic lower bounds are obtained by considering decision trees, for example, the $\Omega(n \log n)$ lower bound on the number of comparisons in any deterministic sorting program.

We can leverage the thicket shatter function, if it is bounded by a polynomial, to prove stronger lower bounds that what we would otherwise be able to show using binary decision trees. (In fact, Tiuryn does just this.) We suspect there are other applications in this vein. For example, Lynch and Blum pose the question of which first-order structures with domain the set of binary strings are *adequate*, which is roughly a polynomial-time version of Turing completeness [9]. We have some preliminary results which give a sufficient condition for inadequacy, which uses Theorem 4.3 in an essential way.

Acknowledgment. The present investigation started as a close reading of Tiuryn [12] suggested by the author's Ph.D. advisor Yiannis Moschovakis. It was greatly assisted by conversations with Anton Bobkov, Caroline Terry, Alex Kruckman, Jim Freitag, David Lippel, and Ross Berkowitz. In addition, Alex Kruckman, Larry Moss, and several anonymous referees provided feedback on various iterations of this manuscript, which substantially improved the presentation of the results. During the preparation of this paper the author was supported by Indiana University, Haverford College, and the Independent Research Fund Denmark.

REFERENCES

[1] M. Aschenbrenner, A. Dolich, D. Haskell, D. Macpherson and S. Starchenko, *Vapnik-Chervonenkis density in some theories without the independence property, II*. ***Notre Dame Journal of Formal Logic***, vol. 54 (2011), pp. 311–363.

[2] ———, *Vapnik-Chervonenkis density in some theories without the independence property, I*. ***Transactions of the American Mathematical Society***, vol. 368 (2015), pp. 5889–5949.

[3] P. Assouad, *Densite et dimension*. ***Annales de l'Institut Fourier***, vol. 33 (1983), no. 3, pp. 233–282.

[4] H. Chase and J. Freitag, *Model theory and machine learning*. ***The Bulletin of Symbolic Logic***, vol. 25 (2019), no. 3, pp. 319–332.

[5] ———, *Model theory and combinatorics of banned sequences*, 2020.

[6] V. Guingona and C. Hill, *On a common generalization of Shelah's* 2*-rank, dp-rank, and o-minimal dimension*. ***Annals of Pure and Applied Logic***, vol. 166 (2013), pp. 502–525.

[7] W. Hodges, ***A Shorter Model Theory***, Cambridge University Press, New York, 1997.

[8] M. C. Laskowski, *Vapnik-chervonenkis classes of definable sets*. ***Journal of the London Mathematical Society***, vol. s2–45 (1992), no. 2, pp. 377–384.

[9] N. A. Lynch and E. K. Blum, *Relative complexity of operations on numeric and bit-string algebras*. ***Mathematical Systems Theory***, vol. 13 (1980), pp. 187–207.

[10] A. Pillay, ***An Introduction to Stability Theory***, Clarendon Press, Oxford, 1983.

[11] S. Shelah, ***Classification Theory***, Elsevier Science Publishers, Amsterdam, 1978.

[12] J. Tiuryn, *A simplified proof of* $DDL < DL$. ***Information and Computation***, vol. 81 (1989), no. 1, pp. 1–12.

[13] V. N. Vapnik and A. Y. Ya Chervonenkis, *The uniform convergence of frequencies of the appearance of events to their probabilities*. ***Theory of Probability and Its Applications***, vol. 16 (1971), pp. 264–280.

DEPARTMENT OF COMPUTER SCINECE
HAVERFORD COLLEGE
HAVERFORD, PA 19041, USA
and
DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF COPENHAGEN
COPENHAGEN 2100, DENMARK
E-mail: sbhaskar@di.ku.dk