Building Uniformly Random Subtrees

Malwina Luczak *
Department of Mathematics
London School of Economics
Houghton Street
London WC2A 2AE, U.K.

Peter Winkler†
Bell Labs 2C-365
Lucent Technologies
700 Mountain Ave.
Murray Hill, NJ 07974 U.S.A.

February 22, 2004

Abstract

We prove the existence of, and describe, a (random) process which builds subtrees of a rooted d-branching tree one node at a time, in such a way that the subtree created at stage n is precisely a uniformly random subtree of size n. The union of these subtrees is a "uniformly random" infinite subtree, which we describe and generate in several ways. Generalization to generation of other combinatorial structures is also considered.

Keywords: random generation, random subtree, combinatorial structures, lattice animal, Catalan numbers

1 Introduction

Combinatorial structures of a given type (let us call them "organisms") are often equipped with a notion of size and a notion of containment. If we are interested in uniformly random organisms of size n, it is natural—and possibly useful—to try to *build* them one step at a time, in such a way that after step n what we have is an organism chosen uniformly at random from all organisms of size n. We call such a process a (uniform) *building scheme.*

A simple example is binary strings, where one string is said to be contained in another if it is an initial segment. In this case there is a unique building scheme which simply adds one random bit to the end. Notice that this process provides one way to define the notion of a "uniformly random" *infinite* binary string.

*Author's net address: malwina@planck.lse.ac.uk
†Author's net address: pw@lucent.com

On the other hand, suppose an organism of size n is a binary string containing exactly n 1's and n 0's, with the property that every initial segment contains at least as many 1's as 0's. These correspond to walks on $\mathbb{N}$ starting from and ending at the origin, often called "Dyck paths". One possible containment notion is to say that $x \subset y$ if x is obtained by deleting the last k 1's and the last k 0's of y, for some k. We will see later that in this case no building scheme exists; nor is it obvious how to extend the notion of a uniformly random organism to the infinite case. But a weaker containment notion does permit a building scheme and with it, a natural "uniform" distribution on infinite organisms.

The organisms with which we are concerned below are subtrees (containing the root) of an infinite d-branching tree $\mathbb{T}^d$. Alternatively they may be thought of as rooted binary plane trees (i.e., trees in which "left" and "right" children are distinguished).

Every node of the tree $\mathbb{T}^d$, including the root o, has exactly d children; thus $\mathbb{T}^d$ is regular of degree $d+1$ except that the root has degree only d. Henceforth an *n-subtree* is a subtree of $\mathbb{T}^d$ containing exactly n nodes including (unless $n = 0$) the root o. We freely confuse an n-subtree σ with its set of nodes, so that e.g. $|\sigma| = n$. We let $\mathbb{S}_n^d$ denote the set of n-subtrees in $\mathbb{T}^d$, and when d is fixed and understood, $s_n = |\mathbb{S}_n|$ will be the number of n-subtrees in $\mathbb{T}^d$. Finally, given any set A and a random variable $\mathbf{X}$, $\mathbf{X} \in_U A$ means that $\mathbf{X}$ is uniformly distributed over A.

Fig. 1 shows one stage of a building scheme for binary subtrees (that is, for $d = 2$). In the center of the figure are the five 3-subtrees; the white circles are nodes which are candidates to be added, each with the indicated probability. At left and right are the fourteen resulting 4-subtrees, each shown as part of a binary tree cut off at finite depth.

For each $d \geq 1$ we prove the existence of a building scheme, that is, a discrete-time random process $\{\mathbf{X}_n\}_{n\in\mathbb{N}}$ such that for all $n > 0$, $\mathbf{X}_n \in_U \mathbb{S}_n$ and $\mathbf{X}_n \subset \mathbf{X}_{n+1}$. This is a trivial statement for $d = 1$ but for $d = 2$ our subtrees correspond to the walks on $\mathbb{N}$ mentioned above, and also to triangulations of $(n+2)$-gons and other organisms counted by Catalan numbers, with appropriate natural containment notions.

Generation of such trees (at least for the case $d = 2$) has been considered, over several decades, in diverse publications in combinatorics and computer science. Methods which have appeared (implicitly or explicitly) include:

- Use of combinatorial bijection to reduce the problem to Łukasiewicz codes (also known as prefix codes, Polish prefix notation, Dyck paths, or gambler's ruin sequences). Drawing from one of the latter uniformly at random is equivalent to drawing a binary sequence of length $2n+1$ comprised of n 1's and n or $n+1$ 0's. (This is deducible from the "cyclic lemma" of Dvoretzky and Motzkin; see e.g. [10] or Chapters 1 and 3 of [9], or the "conjugacy principle" of Raney [27]; by means of a suitable rearrangement, all tree codes are uniformly generated in this way.) The latter task is realizable in linear time by the randomized sampling algorithm of [2].

- Use of the "recursive method" formalized in a much wider context by Flajolet *et al.* [13] on the basis of earlier studies of Nijenhuis & Wilf [23]. For the case at hand, the method provides a top-down generation procedure that is close to what is done in

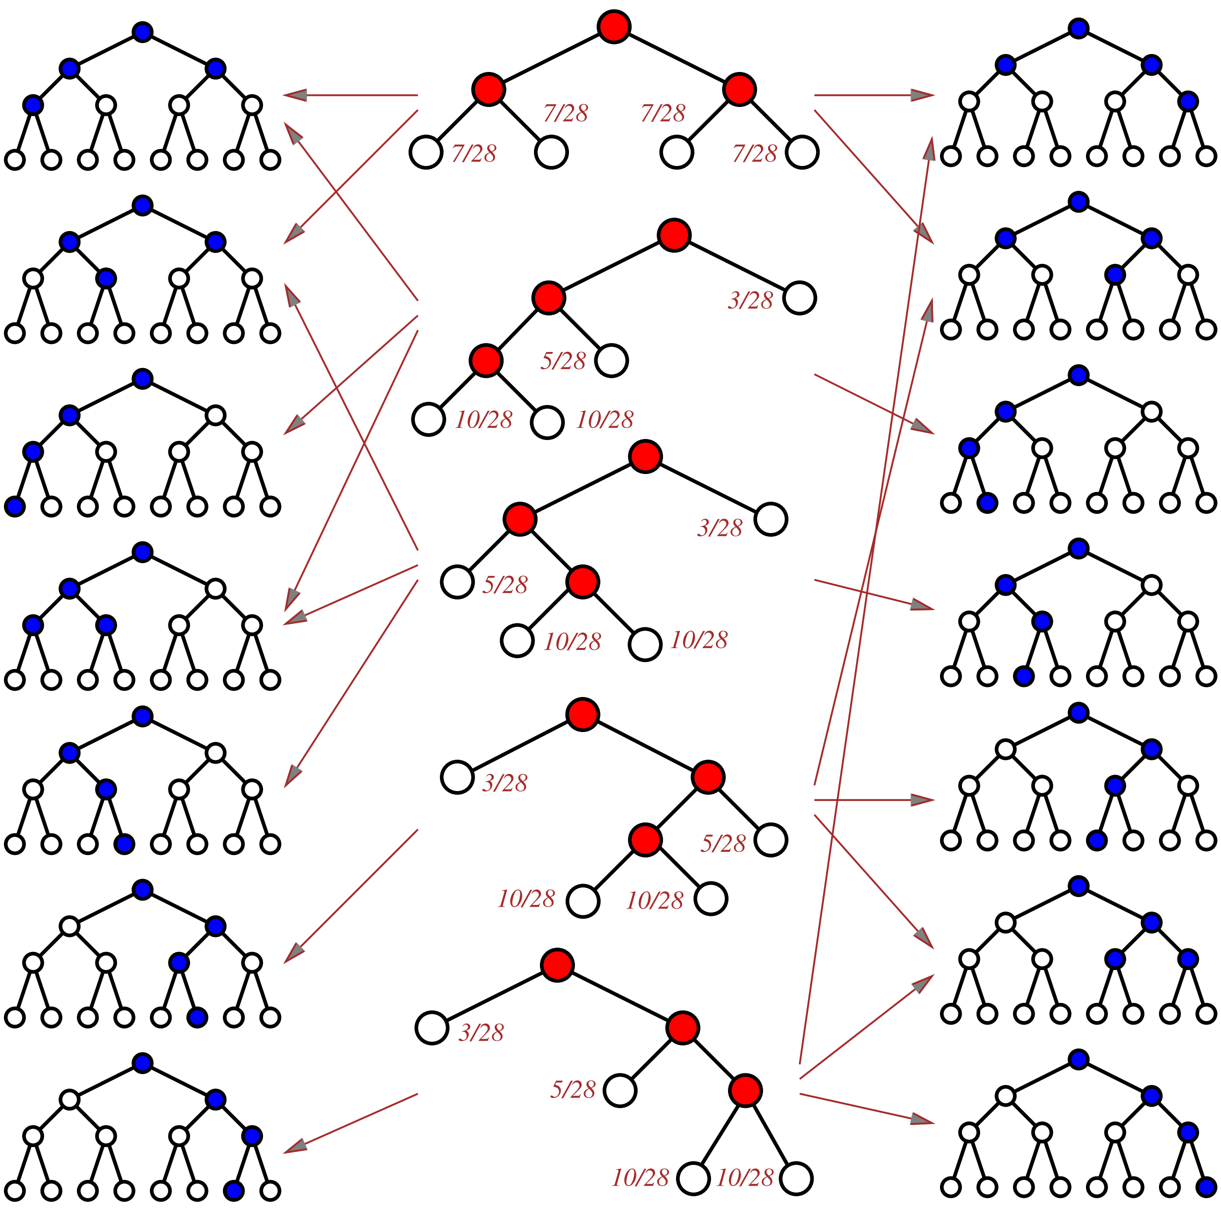

Figure 1: Part of a building scheme for binary subtrees

the present work, while being suboptimal relative to the algorithm based on bijective combinatorics.

- Development of a "generation tree" in the sense of Pinzani *et al.* [5, 4].
- Observation that a tree is a particular realization of a branching process that, for the sake of efficiency, can be taken as critical. This approach works well, especially if it is acceptable to only approximate the size of the resulting tree.

Here, we have an additional purpose: generating trees *incrementally*. Trees are grown by extending ends of branches (like trees in nature) and we require that at any given time n, the tree obtained must be uniformly random among trees of size n. (Note that adding leaves uniformly at random yields very different, non-uniform trees, which, for example, have expected height $\mathrm{O}(\log n)$ instead of the actual $\mathrm{O}(\sqrt{n})$.)

Considering in advance the $d = 2$ case, let $c_n = \frac{1}{n}\binom{2n-2}{n-1}$ denote the nth Catalan number; c_{n+1} counts the number of trees (in our sense) with n nodes. Thus, the left and right main branches of a random tree of size $n = k + \ell + 1$ have sizes k and ℓ with probability $\frac{c_{k+1}c_{\ell+1}}{c_{n+1}}$. The probability of any particular n-node tree is the product over all nodes of such splitting probabilities; of course, the product eventually telescopes to $1/c_{n+1}$.

To do an incremental construction, it suffices to compare the evolution of these probabilities as n changes to $n+1$. This is the essence of our argument, spelled out for the binary case in Sections 2 and 3 below; thanks to Lemma 3.1, the needed condition reduces to an immediate verification concerning Catalan numbers. The more complex situation when $d > 2$ is tackled in Section 4.

In Section 5 Galton-Watson trees are brought into the picture, for use in constructing the infinite trees of Section 6. In Section 7 we consider some other infinite objects, winding up with a short summary of conclusions in Section 8.

Before plunging into our building scheme we have just a few further remarks. Lyons, Pemantle and Peres' notion of "sized-biased" trees in [21] is related to our development. Our infinite "uniformly random" subtree appears in different guise in Janson [15]. In Pitman [24], plane trees and other objects are mapped to random walks in a manner similar to ours below. The flow property we establish to obtain our building schemes is equivalent to a well-known property in the theory of ranked posets, about which we will say more below.

2 n-subtrees and flows

The *boundary* of an n-subtree σ, denoted by $\partial\sigma$, is the set of nodes of $\mathbb{T}^d$ which do not belong to σ but which are adjacent to a node in σ. Since $|\partial\sigma| = (d-1)n+1$ for any n-subtree σ, it seems at first plausible that the process which at each step simply adds a uniformly random boundary point to the subtree should work. The difficulty is that n-subtrees have different numbers of nodes (namely, leaves) whose *removal* results in another subtree. For example, an n-subtree which consists of a path from o can only have been obtained in one way, so if $\mathbf{X}_{n-1}$ is a path we must assign relatively high probability to adding the next node at its end.

We define a bipartite graph $G_n := \langle \mathbb{S}_n, \mathbb{S}_{n+1}, E_n \rangle$ whose edge set E_n is the set of pairs $\{\sigma, \sigma'\}$ consisting of an n-subtree and an $(n+1)$-subtree which contains it. In order to have a building scheme we must have a *unit flow* $\phi: \; E_n \to \mathbb{R}_{\geq 0}$ such that for each $\sigma \in \mathbb{S}_n$,

$$\sum_{\sigma \in e \in E_n} \phi(e) = \frac{1}{s_n}$$

and for each $\sigma' \in \mathbb{S}_{n+1}$

$$\sum_{\sigma' \in e \in E_n} \phi(e) = \frac{1}{s_{n+1}} \; .$$

In the building scheme,

$$\phi(\sigma, \sigma') = \Pr\left(\mathbf{X}_n = \sigma \wedge \mathbf{X}_{n+1} = \sigma'\right) \; .$$

Thus, if we have built as far as $\mathbf{X}_n = \sigma$, we choose $\mathbf{X}_{n+1} = \sigma'$ where $\sigma \subset \sigma'$ with probability $s_n \phi(\sigma, \sigma')$.

Conversely, if such flows are defined for all n then they constitute a building scheme (though not necessarily an elegant one).

In general, the existence of flows such as this between levels of a ranked poset is equivalent to the so-called "LYM" or "BLYM" inequality[1], which says that the sum of the reciprocals of the ranks of the elements of an antichain cannot exceed 1. An equivalent formulation is that for any subset A of a level L_i, if B is the set of elements in L_{i+1} which lie above some element of A, then $|B|/|L_{i+1}| \geq |A|/|L_i|$. (The equivalence of the second formulation, sometimes called the "normalized matching property", seems to have been independently established by Harper [14] and Kleitman [16].) The LYM inequality has historically been used to prove that various finite posets, notably the Boolean algebra of subsets of a set, have the Sperner property: the largest antichain is the size of the largest level. The reader is referred e.g. to Aigner [1], Chapter VIII or Bollobás [7], Chapter 3 for further information.

Neither LYM formulation is of help to us here, as we cannot easily classify all antichains or all subsets of $\mathbb{S}_n$; we are reduced to actually constructing a flow. But it is perhaps worth noting that the existence of a building scheme does establish the Sperner property in the finite case. For example, it is trivial that subsets of a finite set have a building scheme, by adding one uniformly random new element at a time. But the probability that a particular subset appears when the scheme is implemented is the reciprocal of its rank. For subsets in an antichain these appearances are mutually exclusive events, the sum of whose probabilities is thus at most 1; Sperner's Theorem follows, and we have made no mention of binomial coefficients!

We now return to constructing flows for our subtrees. A subtree σ in $\mathbb{S}_n$ has d "branches" each consisting of a child of o and the subtree of σ spawned therefrom. Each branch may itself be thought of as a subtree (possibly empty) rooted at the aforementioned child; the sum of the sizes of the branches is thus $n-1$ as long as σ is not itself empty. The *profile* $\pi(\sigma) = \langle j_1, ..., j_d \rangle$ of σ is the list of sizes of the branches (ordered left to right) of σ. Profiles

[1] The letters stand for Bollobás, Lubell, Meshalkin and Yamamoto—all of whom worked independently.

are partially ordered coordinate-wise, so that $\sigma \subset \sigma'$ implies $\pi(\sigma) \le \pi(\sigma')$. We interpret a profile π as an *event* and write

$$\Pr_n(\pi) := \Pr_{\sigma \in_U \mathbb{S}_n} (\pi(\sigma) = \pi) \ .$$

If vertices of G_n representing subtrees of the same profile are identified, we obtain a weighted "profile graph" H_n whose vertices are profiles of n-subtrees and of $(n{+}1)$-subtrees, with π adjacent to π' if $\pi < \pi'$. The *weight* of a profile is its probability among uniform subtrees of the appropriate size. Given a flow ϕ as above, we set

$$\psi(\pi, \pi') = \sum_{\pi(\sigma)=\pi,\ \pi(\sigma')=\pi'} \phi(\sigma, \sigma')$$

to get a unit flow on H_n, with

$$\sum_{e \ni \pi} \psi(e) = \Pr_n(\pi)$$

for each n-profile π, and

$$\sum_{e \ni \pi'} \psi(e) = \Pr_{n+1}(\pi')$$

for each $(n{+}1)$-profile π'.

Suppose, conversely, that there is such a flow ψ on H_n for every n; then we can construct a building scheme recursively. The idea is that given $\mathbf{X}_n = \sigma$, we first decide *based solely on the profile of* σ which branch to extend. Then we extend that branch just as if it were a smaller subtree. This works because if $\sigma \in_U \mathbb{S}_n$ and has profile $\langle j_1, ..., j_d \rangle$, then for any i the ith branch of σ is itself uniformly random in $\mathbb{S}_{j_i}$. Specifically, if $\mathbf{X}_n = \sigma$ with $\pi(\sigma) = \pi$, we adopt the $(n{+}1)$-profile $\pi' > \pi$ with probability

$$\frac{\psi(\pi, \pi')}{\Pr_n(\pi)} \ .$$

The profile of the branch which is larger in π' than in π is then chosen by the same method, and so forth until an empty branch, which can only be extended in one way, is reached.

A building scheme constructed in this way will be called a *recursive* scheme. Note that the non-uniform scheme which appends a boundary point uniformly at random is, in fact, a recursive scheme; it is obtained by augmenting a branch of size k in an n-subtree with probability $((d-1)k+1)/((d-1)n+1)$. In the $d = 2$ case, the subtrees so obtained are uniformly random *binary search trees* (see e.g. [11]). Considering only *uniform* building schemes, however, we have shown:

Proposition 2.1. *The following are equivalent for any branching number $d \ge 1$:*

1. *there is a building scheme for subtrees of $\mathbb{T}^d$;*
2. *there is a recursive building scheme for subtrees of $\mathbb{T}^d$;*
3. *there is a flow ψ between profile graphs satisfying $\sum_{e \ni \pi} \psi(e) = \Pr_n(\pi)$ for each n-profile π, and $\sum_{e \ni \pi'} \psi(e) = \Pr_{n+1}(\pi')$ for each $(n{+}1)$-profile π'.*

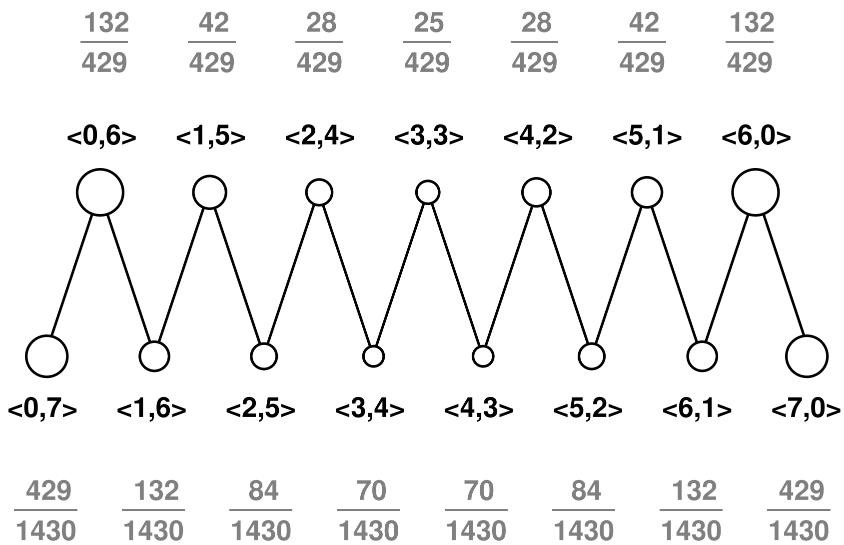

Figure 2: The weighted profile graph for binary subtrees of size 7 and 8

3 The binary case

In the $d = 2$ case there are n possible profiles for an n-subtree, so the graph H_n has n vertices in one part (which we imagine to be the "upper" vertices) and $n+1$ in the lower part. The edges themselves form a zig-zag path. Fig. 2 shows the graph H_7 with associated weights; the flow is not indicated.

The path structure of H_n means that the flow ψ, if it exists, will be unique; and it is apparent by trying to construct the flow from left to right that we require

$$\sum_{i=0}^{k} \Pr_{n+1}(\langle i, n-i\rangle) \leq \sum_{i=0}^{k} \Pr_{n}(\langle i, n-i-1\rangle) \leq \sum_{i=0}^{k+1} \Pr_{n+1}(\langle i, n-i\rangle) \tag{1}$$

for $0 \leq k \leq n-1$, in order for the flow to exist. Closed forms exist for these sums (see e.g. [29]) but we will not need them here.

Lemma 3.1. *In the binary case, if*

$$\Pr_{n+1}(\langle i+1, j\rangle) < \Pr_{n}(\langle i, j\rangle) > \Pr_{n+1}(\langle i, j+1\rangle)$$

for each i, j with $i + j = n-1$, then the flow ψ of Proposition 2.1 exists.

Proof. The first inequality of (1) follows directly, the second from

$$\sum_{i=0}^{k} \Pr_{n}(\langle i, n-i-1\rangle) = 1 - \sum_{i=k+1}^{n-1} \Pr_{n}(\langle i, n-i-1\rangle) > 1 - \sum_{i=k+2}^{n} \Pr_{n+1}(\langle i, n-i\rangle) = \sum_{i=0}^{k+1} \Pr_{n+1}(\langle i, n-i\rangle) .$$

□

As we have said, the number of binary n-subtrees is (as is easily checked) the $n+1^{\mathrm{st}}$ Catalan number, $\binom{2n}{n}/(n+1)$. Thus, the condition of Lemma 3.1 reduces to the observation that the ratios between successive Catalan numbers are increasing, or in subtree terms $s_{n+1}/s_n > s_n/s_{n-1}$, which is routine to check. We will see later a proof for a more general set of inequalities, needed when the branching number is arbitrary. For now we note that a building scheme does exist for subtrees of $\mathbb{T}^2$.

It is worth noting that in our binary recursive building scheme the relative probability of extending a branch of size $k > n/2$ is greater than the ratio $(k+1)/(n+1)$ which would produce uniform binary search trees. Hence our uniform subtrees are more lopsided and sparser than binary search trees, in fact much more; for example n-subtrees have height of order $\sqrt{n}$ ([12, 18]) whereas binary search trees of size n have height of order $\log n$ [11].

Before proceeding to higher branching number, it is perhaps worth taking a look at building schemes for other organisms counted by the Catalan numbers. The issue, of course, is that the containment notion may differ from case to case.

For example, the Catalan numbers also count the triangulations of an $(n+2)$-gon. Here it is natural to extend a triangulation by choosing one edge of the $(n+2)$-gon and building a new triangle on its exterior side, thus obtaining a triangulated $(n+3)$-gon. Clearly there are $n+2$ ways to do this as opposed to only $n+1$ ways to extend an n-subtree, so this containment notion is not the same; but it almost is. Distinguishing one edge of the $(n+2)$-gon provides a 1–1 correspondence to n-subtrees in which building on any polygon edge *except* the distinguished one corresponds to extending the subtree. (Building on the distinguished edge corresponds to adding a new parent to the subtree's root, then designating it as the new root.) Since the containment notion for the triangulations is thus weaker than for the subtrees, the building scheme works for polygon triangulations as well. Fig. 3 illustrates the containment relation for both between $n = 3$ and $n = 4$, the extra containments among the polygons shown by curved dashed lines.

Catalan numbers also count the ways to nest n pairs of parentheses, or equivalently, the number of binary sequences containing n 1's and n 0's all of whose initial segments contain at least as many 1's as 0's. One way to extend these is to insert a new 1 anywhere after the (formerly) last 1, and add a 0 at the end. This results in only one way to build any given sequence, and in such a case there cannot be a building scheme unless the number of structures increases by integer multiples (not the case here, where $s_2 = 2$ and $s_3 = 5$).

A more flexible way to extend is to turn any "0" into "100" or add "10" to the end. As we shall see (after Proposition 5.1) that turns out to be equivalent to our subtree containment.

A *rooted plane tree* is a rooted tree, with nodes of arbitrary degree, in which the children of any node are ordered (left to right). Thus, for example, the 4-node tree in which the root has two children and the left child has a child is different from the one in which the right child has a child. These are also counted by Catalan numbers (see e.g. [28], Chapter 5 and [11], esp. Section 6.3). In fact there is a natural correspondence between rooted plane trees on n nodes and binary sequences with n 1's and n 0's, conditioned as above; see Fig. 4 for an example. Turning any "0" into "100" or adding "10" to the end of the corresponding binary sequence is equivalent to adding a new child either to a childless node, or to the right of

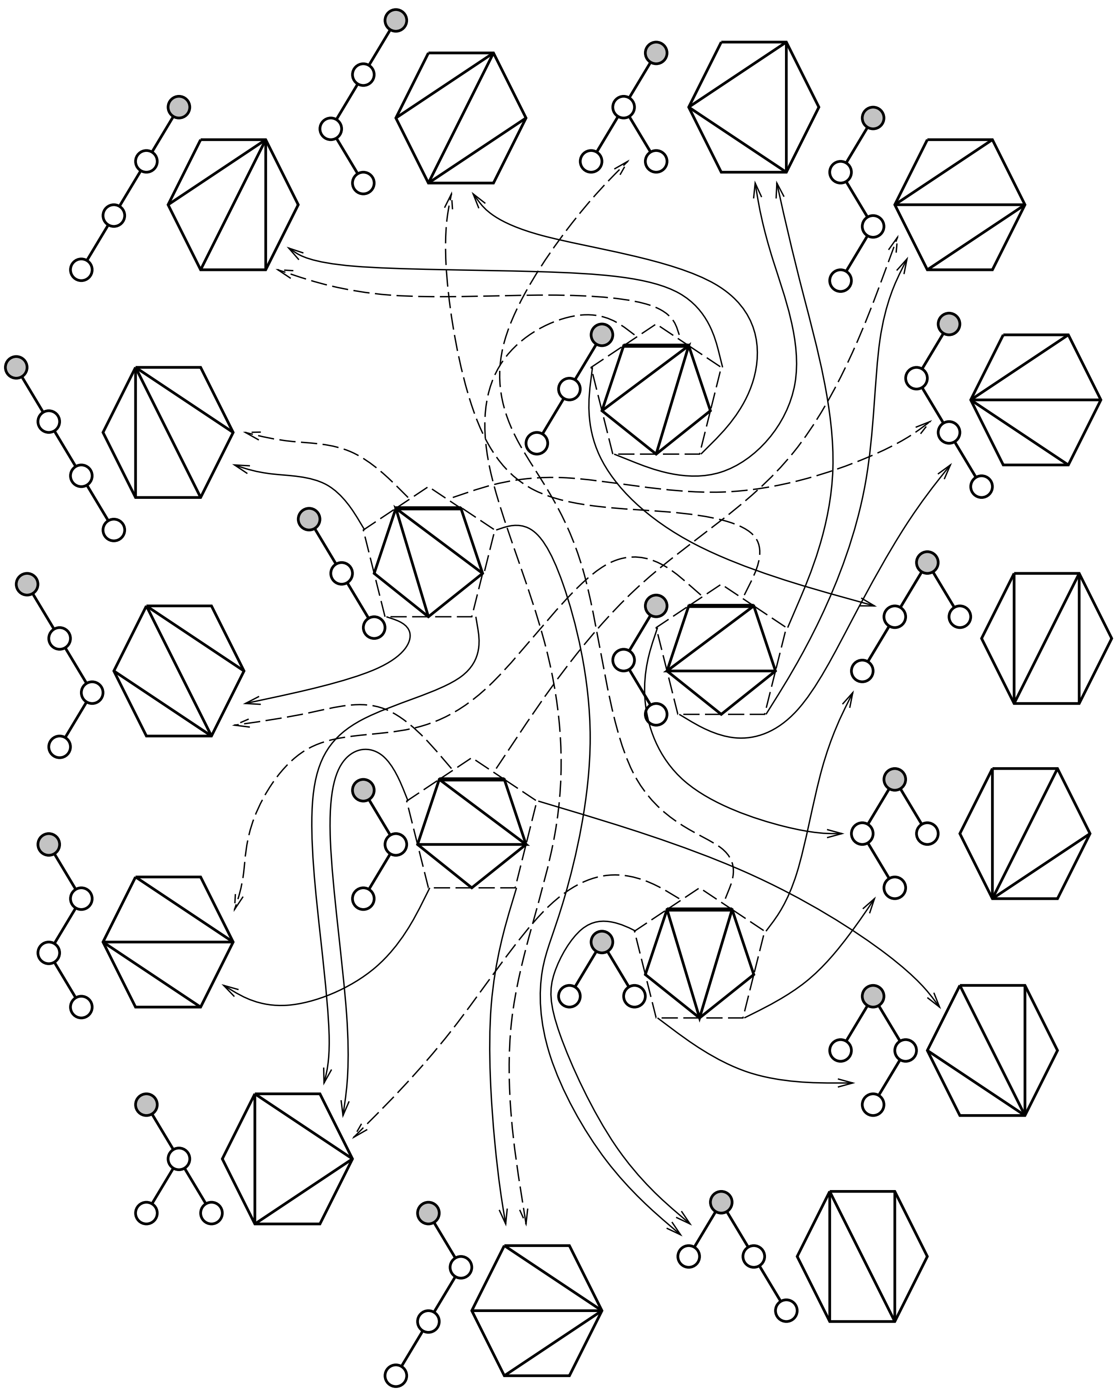

Figure 3: Containment among binary subtrees and polygonal triangulations, $n = 3, 4$

some existing child. Since the natural notion of containment permits this and more (adding new children to the left), our building scheme works for rooted plane trees.

Translated to binary sequences, the rooted plane tree containment amounts to allowing insertion of "10" anywhere, so we get a building scheme there too.

Many additional organisms counted by the Catalan numbers can be found in [30], Section 6.2 (esp. Exercise 6.19) and in Stanley's "Catalan Addendum," at http://www-math.mit.edu/∼rstan/ec/catadd.pdf, but we have not uncovered any additional interesting containment relations.

4 Higher branching number

We begin here the proof of our main result:

Theorem 4.1. *For each $d \geq 1$, there is a building scheme for the subtrees of $\mathbb{T}^d$.*

The profiles of n-subtrees of $\mathbb{T}^d$ for $d > 2$ constitute a discrete $(d-1)$-dimensional simplex with n profiles along each outer edge. The bipartite profile graph H_n connects two of these, corresponding to subtree sizes n and $n+1$, vertices in the former part having degree d.

If we fix the first $d-2$ entries of the profile at, say, $j_1, \ldots, j_{d-2}$, we get in each simplex a line of profiles, connected in H_n by a zig-zag path as in the binary case. We denote those lines by

$$\langle j_1, \ldots, j_{d-2}, *, * \rangle_n$$

and

$$\langle j_1, \ldots, j_{d-2}, *, * \rangle_{n+1}$$

respectively. Indeed, we will treat those lines as we did in the binary case, obtaining a unit flow between them. We will also create a flow between the line $\langle j_1, \ldots, j_{d-2}, *, * \rangle_n$ and the line $\langle j_1, \ldots, j_{i-1}, j_i{+}1, j_{i+1}, \ldots, j_{d-2}, *, * \rangle_{n+1}$, for each $i \leq d$–2; however, here the construction will turn out to be trivially easy.

Since we need a unit flow in total between the two simplices, we can only make use of fractions of the above flows; the values of these fractions will be determined by collapsing the lines into points and finding new flows on the quotient graph, and so forth. It would be nice if the quotient graph for branching number d were the same as the profile graph for $d-1$, but that is not the case, so our construction is iterative but not recursive.

Fig. 5 shows the $d = 3$, $n = 6$ case. Only the probabilities of the profiles are shown; the profiles themselves are indicated only by their position. The two simplices are superimposed, probabilities of 6-profiles in italics and probabilities of 7-profiles in boldface. Flows between horizontal lines are done first; then the horizontal lines are collapsed, resulting in the pair of vertical lines shown on the right side of the figure.

In the end, the flow from $\langle j_1, \ldots, j_d \rangle_n$ to $\langle j_1, \ldots, j_{i-1}, j_i+1, j_{i+1}, \ldots, j_d \rangle_{n+1}$ will be the product whose first term is

$$\psi(\langle j_1, \ldots, j_d \rangle, \langle j_1, \ldots, j_{i-1}, j_i+1, j_{i+1}, \ldots, j_d \rangle)$$

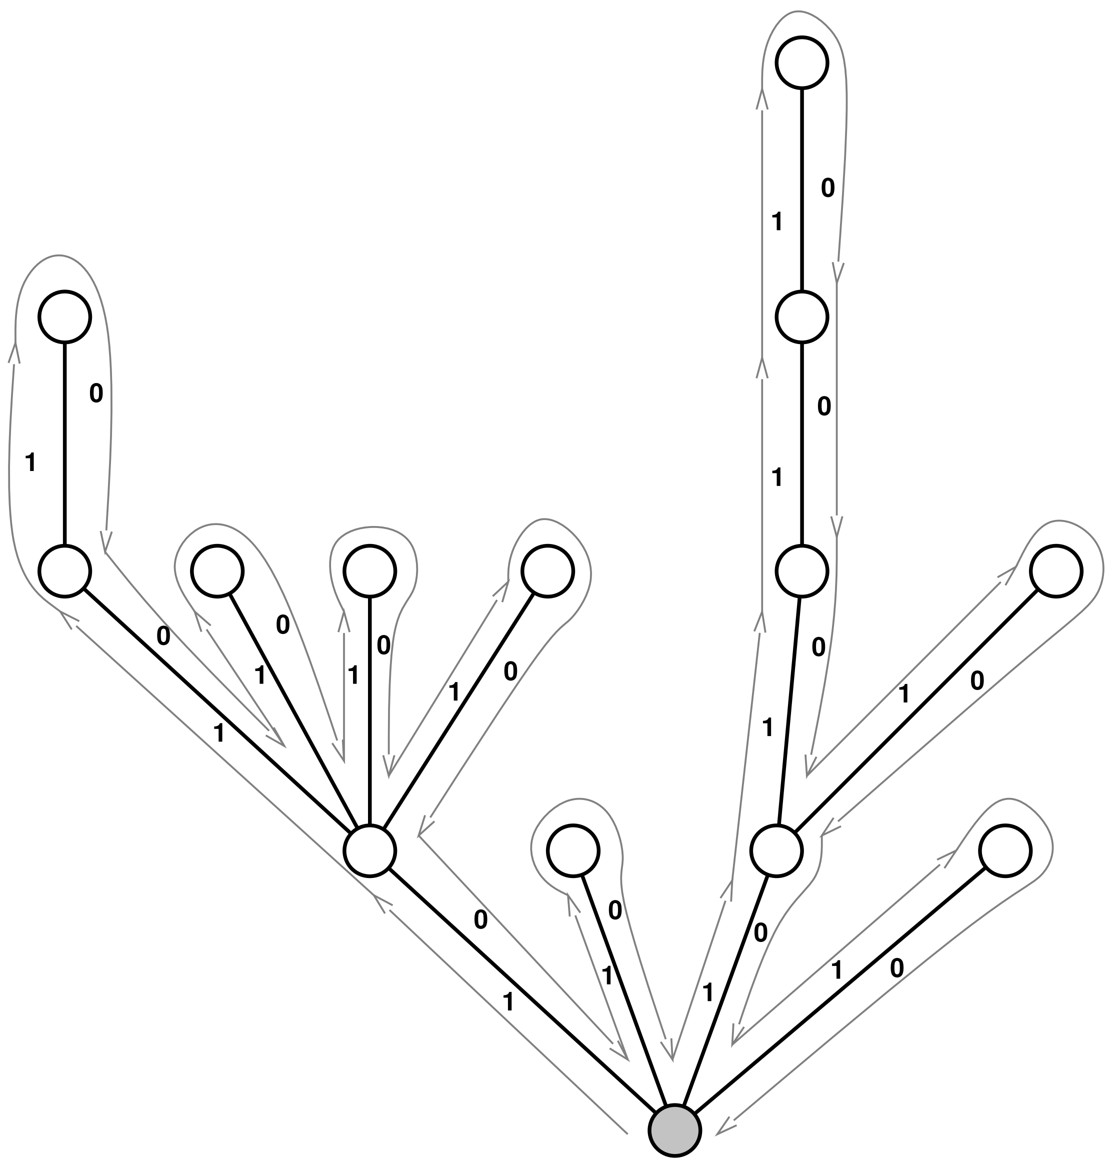

11100101010010111100010010

Figure 4: A rooted plane tree and its corresponding binary sequence

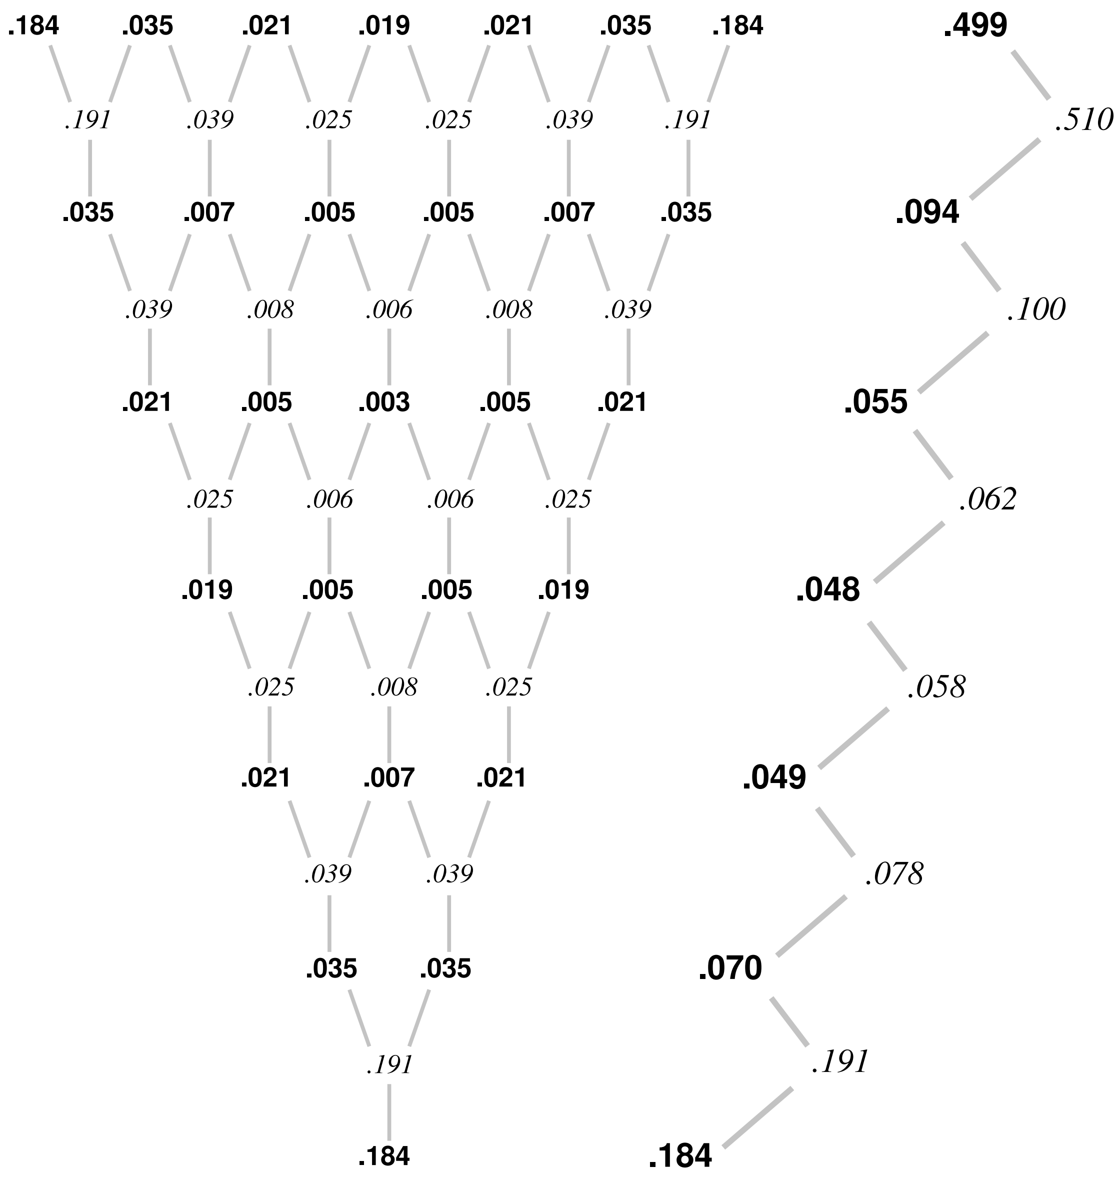

Figure 5: Profile probabilities before and after collapse, $d = 3$, sizes 6 and 7

which we can also write as

$$\psi(\langle j_1,\ldots,j_{d-1},*\rangle_n,\langle j_1,\ldots,j_{i-1},j_i{+}1,j_{i+1},\ldots,j_{d-1},*\rangle_{n+1})\ ,$$

and whose last ($(d-1)$st) term is

$$\psi(\langle j_1,*,\ldots,*\rangle_n,\langle j_1,*,\ldots,*\rangle_{n+1})\ .$$

We must now obtain unit flows at each level. We start with the easy case, where one line consist of points of the form

$$\langle j_1,\ldots,j_k,h,*,\ldots,*\rangle_n$$

for $h=0,1,\ldots,m$, and the other line's points are

$$\langle j_1,\ldots,j_{i-1},j_i{+}1,j_{i+1},\ldots,j_k,h,*,\ldots,*\rangle_{n+1}$$

also for $h=0,1,\ldots,m$, where $m=n-1-\sum_{\ell=1}^{k}j_\ell$. Here the graph is not a path but a parallel matching, so this approach is doomed unless the weights of the vertices at the ends of an edge are identical; but they are. To see this, let us introduce one more piece of notation: $s_n^{(j)}$ is the number of *sequences* of j subtrees whose sizes sum to n. Note that $s_n^{(d)}=s_{n+1}=s_{n+1}^{(1)}$ since a subtree consists of its sequence of d branches plus a root.

The weight of vertex $\langle j_1,\ldots,j_k,h,*,\ldots,*\rangle_n$ in the n-line is its relative probability, namely

$$\Pr_n\left(\langle j_1,\ldots,j_k,h,*,\ldots,*\rangle_n\,\middle|\,\langle j_1,\ldots,j_k,*,\ldots,*\rangle_n\right)$$

but this is just $s_h s_{m-h}^{(d-k-1)}/s_m^{(d-k)}$. The opposing profile weight

$$\Pr_{n+1}\left(\langle j_1,\ldots,j_{i-1},j_i{+}1,j_{i+1},\ldots,j_k,h,*,\ldots,*\rangle_{n+1}\,\middle|\,\langle j_1,\ldots,j_{i-1},j_i{+}1,j_{i+1},\ldots,j_k,*,\ldots,*\rangle_{n+1}\right)$$

yields the identical value.

Now we need to get a flow in the interesting case, where the n-line is as above but the $(n{+}1)$-line contains the points

$$\langle j_1,\ldots,j_k,h,*,\ldots,*\rangle_{n+1}$$

for $h=0,1,\ldots,m{+}1$. We will attempt to do this by verifying the condition of Lemma 3.1. Observe that

$$\Pr_{n+1}\left(\langle j_1,\ldots,j_k,h,*,\ldots,*\rangle_{n+1}\,\middle|\,\langle j_1,\ldots,j_k,*,\ldots,*\rangle_{n+1}\right)=\frac{s_h s_{m+1-h}^{(d-k-1)}}{s_{m+1}^{(d-k)}}$$

and

$$\Pr_{n+1}\left(\langle j_1,\ldots,j_k,h{+}1,*,\ldots,*\rangle_{n+1}\,\middle|\,\langle j_1,\ldots,j_k,*,\ldots,*\rangle_{n+1}\right)=\frac{s_{h+1} s_{m-h}^{(d-k-1)}}{s_{m+1}^{(d-k)}}\ .$$

Thus the condition we seek to verify is

$$\frac{s_h s^{(d-k-1)}_{m+1-h}}{s^{(d-k)}_{m+1}} < \frac{s_h s^{(d-k-1)}_{m-h}}{s^{(d-k)}_{m}} > \frac{s_{h+1} s^{(d-k-1)}_{m-h}}{s^{(d-k)}_{m+1}} .$$

The first inequality is equivalent to

$$\frac{s^{(d-k-1)}_{m+1-h}}{s^{(d-k-1)}_{m-h}} < \frac{s^{(d-k)}_{m+1}}{s^{(d-k)}_{m}}$$

and the second to

$$\frac{s_{h+1}}{s_h} < \frac{s^{(d-k)}_{m+1}}{s^{(d-k)}_{m}} .$$

If we can prove that

$$\frac{s_{n+1}}{s_n} < \frac{s^{(2)}_{n+1}}{s^{(2)}_{n}} < \frac{s^{(3)}_{n+1}}{s^{(3)}_{n}} < \cdots < \frac{s^{(d-1)}_{n+1}}{s^{(d-1)}_{n}} < \frac{s_{n+2}}{s_{n+1}}$$

for every n, then the two desired inequalities follow and the proof of Theorem 4.1 will be complete.

5 Galton-Watson processes and random walks on $\mathbb{Z}$

A natural way to construct random subtrees of unspecified size is via a Galton-Watson process with i.i.d. children. We fix the branching number d and a probability $p \in (0,1)$. The root o itself is "born" with probability p (otherwise our subtree will be empty) and each time a node is chosen, each of its d potential children is born independently with probability p. If the resulting subtree σ has size n then $\sigma \in_U \mathbb{S}_n$, since any subtree in $\mathbb{S}_n$ requires exactly n births and $(d-1)n+1$ failures and thus has the same probability

$$\Pr_n(\sigma) = p^n(1-p)^{(d-1)n+1} .$$

As is well known (see e.g. [3]), σ is finite with probability 1 if $p < 1/d$ (sub-critical) or $p = 1/d$ (critical). We will set $p = 1/d$ henceforth for definiteness.

At this time it is useful for us to imagine that births are considered in *depth-first* order, left to right, so that e.g. the entire left branch of σ is constructed before the next branch is begun. In this way each n-subtree σ corresponds to a unique birth/non-birth sequence $s_1(\sigma), s_2(\sigma), \ldots, s_{dn+1}(\sigma)$ which contains exactly n births and $(d-1)n+1$ non-births; moreover, up to any $k < dn+1$ the sequence $s_1(\sigma), \ldots, s_k(\sigma)$ must contain at least k/d births (otherwise the growth would have terminated by step k).

The sequence may be interpreted as a walk on $\mathbb{Z}$ beginning at 0, each step of which moves either $d-1$ steps to the right (corresponding to a birth) or 1 step to the left (a non-birth). The following propositions appear to have been known in some form for a century or more; see e.g. [27] or Chapter 11 of [20] for more information.

Proposition 5.1. *The correspondence birth = d–1 steps right, non-birth = 1 step left yields a bijection between n-subtrees and walks from 0 on $\mathbb{Z}$ with the following properties:*

1. *every step is either $d-1$ to the right or 1 to the left;*
2. *there are exactly n steps to the right; and*
3. *the final step brings the walk for the first time to the point -1.*

We thus have that $s_n = |\mathbb{S}_n|$ is the number of such walks; the probability that our Galton-Watson subtree has size n is $p_n = s_n p^n (1-p)^{(d-1)n+1}$. In the case of branching number 2, where steps to either side are length 1, these walks are well known to be counted by the Catalan numbers

$$s_n = \binom{2n}{n}/(n+1) ;$$

more generally,

$$s_n = \binom{dn}{n}/((d-1)n+1) .$$

This equality can be obtained using generating functions (see e.g. [17], Section 2.3.4.4, Problem 11, where the computation is left as a tricky exercise). The elementary approach below gives more information, and in particular yields directly the critical ratio inequalities needed for our flow construction.

We note here that extending an n-subtree to an n+1-subtree amounts to replacing a left step in the walk by a sequence consisting of a right step and then d left steps. If, as earlier, we leave off the last step and represent the walks corresponding to binary subtrees by binary strings with n 1's and n 0's, we must allow the addition of a "10" at the end as well as the conversion of a "0" to a "100".

The Galton-Watson construction can be generalized slightly to construct j subtrees instead of 1; we simply continue flipping the birth/non-birth coin until j trees are completed. Concatenating the sequences for those subtrees, we have

Proposition 5.2. *There is a 1-1 correspondence between sequences of j subtrees of total size n, and walks from 0 on $\mathbb{Z}$ with the following properties:*

1. *every step is either $d-1$ to the right or 1 to the left;*
2. *there are exactly n steps to the right; and*
3. *the final step brings the walk for the first time to the point $-j$.*

The number of such walks is thus $s_n^{(j)}$, and we are in a position to prove the following simple but critical result.

Lemma 5.3. *For every d, n and j, the ratio $s_n^{(j+1)}/s_n^{(j)}$ of the number of sequences of j+1 n-subtrees (of a d-branching tree) to the number of sequences of j n-subtrees is*

$$\frac{j+1}{j}\frac{dn+j}{(d-1)n+j+1} .$$

Proof. The idea is to define a bipartite graph with the $s_n^{(j)}$ sequences on the left, and the $s_n^{(j+1)}$ sequences on the right; we will arrange it so that the degree of every left node is $dn+j$, and the *mean* degree on the right is $(j/(j+1))((d-1)n+j+1)$.

To go from left to right, we insert a left step anywhere in the $s_n^{(j+1)}$ walk (including at the front) except at the end. Then we get a walk which ends after n right steps when it hits the point $-j-1$; exactly a $s_n^{(j+1)}$ walk. So the left-side vertices of our bipartite graph all have degree exactly $dn+j$.

To find the neighbors of a fixed $s_n^{(j+1)}$ walk we must remove a left step, but this left step cannot be from the final $((j+1)\text{th})$ tree since the resulting walk would then hit the point $-j$ twice. Hence the neighbors of such a walk are enumerated by the number of left steps in the first j trees. However, suppose we regard two $s_n^{(j+1)}$ walks as "equivalent" if we can get one from the other by rotating the $j+1$ subtrees. Then inside each class the average number of removable left steps is just $j/(j+1)$ times the total number of left steps. Since the latter is $((d-1)n+j+1)$, we have that the average degree of an $s_n^{(j+1)}$ walk is $j/(j+1)((d-1)n+j+1)$ as required.

Fig. 6 below shows the graph in the $d=3$, $j=1$, $n=2$ case. Each graph vertex is indicated by a solid oval containing the walk on $\mathbb{Z}^2$ and its corresponding subtree or pair of subtrees; equivalence classes are placed in dashed ovals. □

Corollary 5.4. *For every d and n, if $s_n^{(j)}$ is the number of sequences of j n-subtrees of a d-branching tree, then*

$$\frac{s_{n+1}}{s_n} < \frac{s_{n+1}^{(2)}}{s_n^{(2)}} < \frac{s_{n+1}^{(3)}}{s_n^{(3)}} < \cdots < \frac{s_{n+1}^{(d-1)}}{s_n^{(d-1)}} < \frac{s_{n+2}}{s_{n+1}} .$$

Proof. From Lemma 5.3 we see that $s_n^{(j+1)}/s_n^{(j)}$ is increasing in n, which implies each claimed inequality, once we remember that $s_n^{(d)} = s_{n+1}$. □

The proof of Theorem 4.1 follows. Lemma 5.3 also gives:

Corollary 5.5. *For every d, n and j,*

$$s_n^{(j)} = \frac{j}{n}\binom{dn+j-1}{n-1} .$$

Proof. By induction. □

6 Infinite subtrees

The subtree $\mathbf{X}_\infty := \bigcup_n \mathbf{X}_n$, which we think of as a *uniform random infinite subtree*, is not automatically an interesting object of study. The infinite version of the random binary search tree, for example, is just the complete binary tree (with probability 1); this reflects the fact

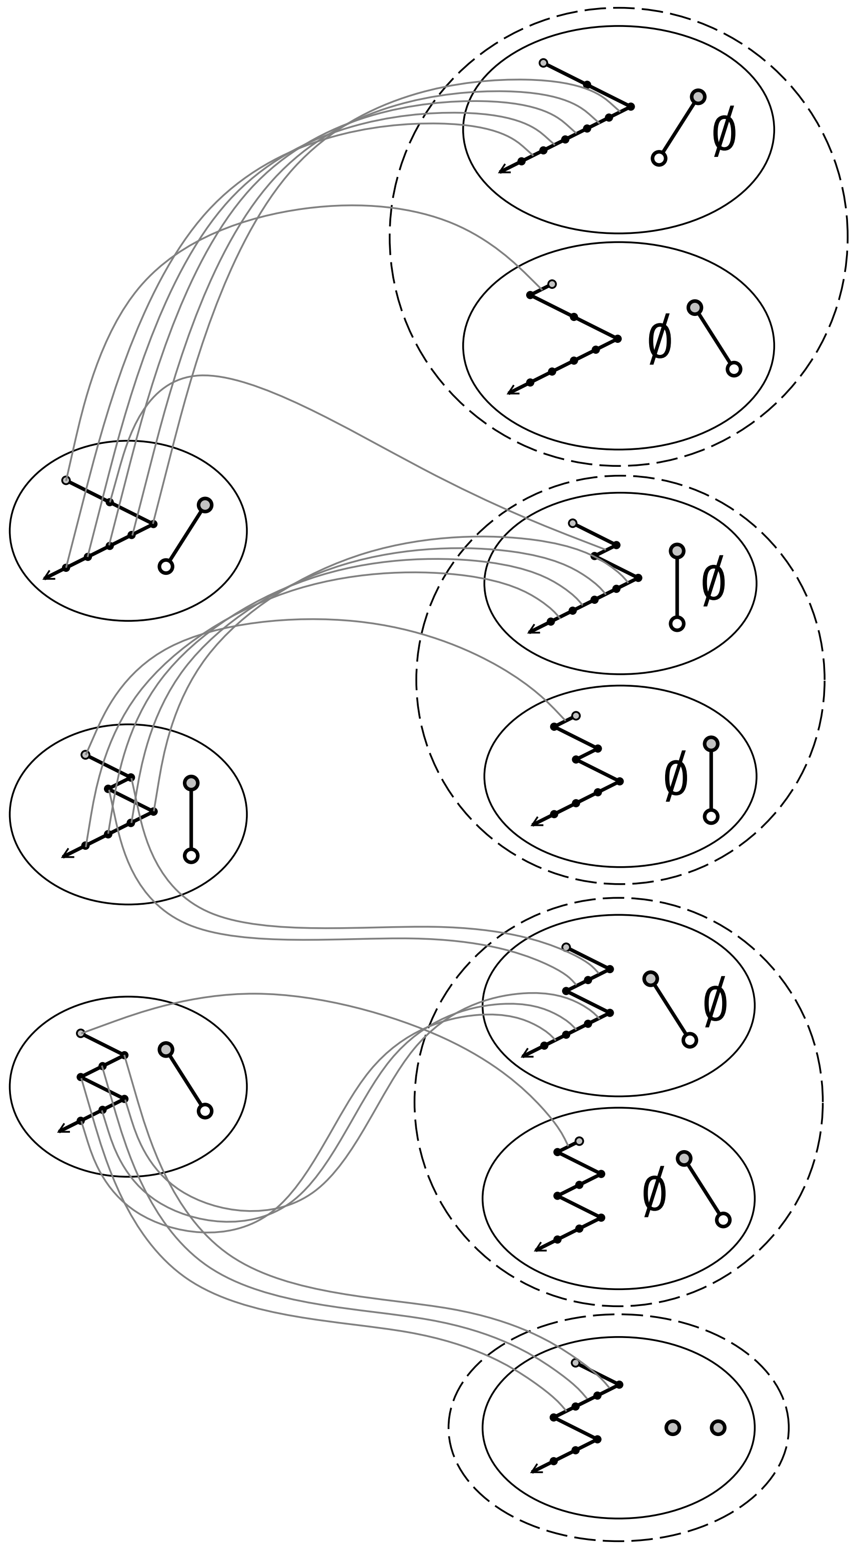

Figure 6: The graph in the proof of Lemma 5.3, for $d = 3$, $j = 1$ and $n = 2$

that for any height h, sufficiently large random search trees will contain all nodes of $\mathbb{T}^2$ at depth $\le h$.

In our case there is positive probability that even the leftmost child of the root will never show up; in fact we know from Lemma 5.3 that this probability is

$$\lim_{n\to\infty} \frac{s_n^{(d-1)}}{s_n^{(d)}} = \lim_{n\to\infty} \frac{d-1}{d} \frac{(d-1)n+d}{dn+d-1} = \left(\frac{d-1}{d}\right)^2 .$$

More generally, the limiting probability that an n-subtree has a *particular* k-subtree σ as its left branch is, by a similar calculation,

$$\left(\frac{d-1}{d}\right)^2 \alpha^k$$

where $\alpha = (d{-}1)^{d-1}/d^d$. However, in our critical Galton-Walton process for trees of branching number d, $\Pr(\sigma) = \alpha^k \cdot (d{-}1)/d$, and we know that these values sum to 1 over all σ because this process is finite with probability 1. It follows that the probability that the left branch of $\mathbf{X}_\infty$ is finite is $(d{-}1)/d$.

Since *some* branch of $\mathbf{X}_\infty$ is infinite, we conclude that the probability that more than one branch is infinite is zero. Applying the same argument at each level, we see that with probability 1, $\mathbf{X}_\infty$ contains just one infinite branch; and from every node on that branch, treated as a root, the $d-1$ branches which do not intersect the branch are independent critical Galton-Watson trees. This provides an easy way to describe $\mathbf{X}_\infty$, since the infinite branch itself can be chosen by repeated rolls of a d-sided die.

The construction of $\mathbf{X}_\infty$ using a simple Galton-Watson process is a bit tricky. We would like to run a critical process and identify it, as in the finite case, with a random walk on $\mathbb{Z}$ that starts off at 0 and at each step, goes to the right a distance $d{-}1$ with probability $1/d$ else to the left a distance 1, and then condition on never hitting -1 to ensure non-extinction of the critical process. Unfortunately, this causes two problems. The first one is that, with the depth-first order adopted earlier, the part of $\mathbf{X}_\infty$ to the right of the infinite branch may never be built. This can be dealt with quite easily, by building subtrees in breadth-first instead of depth-first order. We note that a 1-1 measure preserving correspondence between subtrees of size n and walks which hit -1 first at step $dn{+}1$ is still available.

The second issue is that the condition that the random walk never hits -1 has probability zero, and this is somewhat harder to deal with. To get around the difficulty, we shall use a random walk on $\mathbb{N}$ which behaves like a random walk on $\mathbb{Z}$ when the latter is conditioned on not dropping below 0. To do this we associate a weight $w(i) = i{+}1$ to each node i of $\mathbb{N}$, and step according to the following rule. From node i, we step to node $i{+}d{-}1$ with probability

$$\frac{\frac{w(i+d-1)}{d-1}}{\frac{w(i+d-1)}{d-1} + w(i{-}1)} = \frac{i{+}d}{d(i{+}1)} ,$$

and step to node $i-1$ with probability

$$\frac{w(i{-}1)}{\frac{w(i+d-1)}{d-1} + w(i{-}1)} = \frac{(d{-}1)i}{d(i{+}1)} .$$

In what follows we shall refer to this walk as W. For $d=2$, W is a node-weighted random walk [8]; for $d>2$, it can be regarded as an arc-weighted random walk on a digraph defined in an obvious way.

Theorem 6.1. *The weighted random walk W described above generates, via breadth-first labeling, the "uniformly random" infinite subtree $\mathbf{X}_\infty$.*

Proof. Let $\mathbf{Y}_\infty$ be the infinite subtree generated by W with breadth-first labeling, and let $\mathbb{T}^d_k$ be the subtree of $\mathbb{T}^d$ consisting of levels 0 through k. To prove Theorem 6.1, we show that for any $k \in \mathbb{N}$, $\mathbf{Y}_\infty \cap \mathbb{T}^d_k$ has the same distribution as $\mathbf{X}_\infty \cap \mathbb{T}^d_k$, and then it will follow from Carathéodory's extension theorem.

It follows from our construction that for each fixed subtree $S \subseteq \mathbb{T}^d_k$,

$$\mathbb{P}(\mathbf{X}_\infty \cap \mathbb{T}^d_k = S) = \lim_{n\to\infty} \mathbb{P}(\mathbf{X}_n \cap \mathbb{T}^d_k = S) \ .$$

Recall that each uniformly random $\mathbf{X}_n$ corresponds via breadth-first labeling to a unique walk W'_{nd+1} on $\mathbb{Z}$ of length $nd{+}1$, beginning at 0 and terminating upon first reaching -1, as in Proposition 5.1; the switch from depth- to breadth-first labeling has no effect in the finite case.

It is enough to prove that, given $\varepsilon > 0$, there exists an integer $n_\varepsilon \in \mathbb{N}$ such that for all $n \geq n_\varepsilon$,

$$\left|\mathbb{P}(\mathbf{X}_n \cap \mathbb{T}^d_k = S) - \mathbb{P}(\mathbf{Y}_\infty \cap \mathbb{T}^d_k = S)\right| < \varepsilon \ .$$

This reduces to showing that for every integer $m \in \mathbb{N}$ and every walk ω of length m the following holds. For all $\varepsilon > 0$, there exists $n_\varepsilon \in \mathbb{N}$ such that for all $n \geq n_\varepsilon$

$$|\mathbb{P}(W|_m = \omega) - \mathbb{P}(W'_{nd+1}|_m = \omega)| < \varepsilon \ . \tag{2}$$

Here $W|_m$ represents the initial segment of W of length m.

We prove (2) by induction on m. Fix a walk ω of length m and let $\omega^- = \omega|_{m-1}$, ending, say, at point j. Let us assume first that at its m^{th} step, ω moved to the left, to $j{-}1$. Then, as part of our infinite walk on $\mathbb{N}$,

$$\frac{\mathbb{P}(W|_m = \omega)}{\mathbb{P}(W|_{m-1} = \omega^-)} = \frac{(d{-}1)j}{d(j{+}1)} \ .$$

If W'_{nd+1} generates an n-subtree and begins with ω^-, then the rest of W'_{nd+1} is a walk which, when displaced distance j to the left, begins at 0 and ends $nd{+}1 - (m{-}1)$ steps later when it first hits $-j{-}1$. Such a walk contains r steps to the right, where $r(d{-}1) - (nd{+}1 - (m{-}1) - r) = -j{-}1$ and thus $r = n - (m{+}j{-}1)/d$. Recalling Proposition 5.2, the number of such walks is precisely the number $s_r^{(j+1)}$ of sequences of $j{+}1$ subtrees of size totaling r.

If W'_{nd+1} begins with ω, a similar argument computes the number of possibilities for the rest of W'_{nd+1} to be $s_r^{(j)}$.

It follows that

$$\frac{\mathbb{P}(W'_{nd+1}|_m = \omega)}{\mathbb{P}(W'_{nd+1}|_{m-1} = \omega^-)} = \frac{s_r^{(j)}}{s_r^{(j+1)}} = \frac{j}{j+1}\frac{(d-1)r+j+1}{dr+j} \to \frac{j}{j+1}\frac{(d-1)}{d}$$

so (2) for ω follows from (2) for ω^-. Moreover, (2) must then also hold for the walk ω' which consists of ω^- followed by a step to the right. $\square$

7 Measures on other infinite objects

Measures on various other infinite objects have recently attracted a fair amount of attention. For instance, consider self-avoiding walks on the lattice $\mathbb{Z}^d$ [22]. An n-step *self-avoiding walk* ω on $\mathbb{Z}^d$ is a sequence of lattice sites $(\omega(0), \dots, \omega(n))$ such that $\omega(i) \neq \omega(j)$ for all $i \neq j$. For $n \geq m$ let $P_{m,n}(\omega)$ be the fraction of n-step self-avoiding walks that extend a given m-step self-avoiding walk ω. Suppose $P_m(\omega) = \lim_{n\to\infty} P_{m,n}(\omega)$ exists for all m and all ω. Then the measures P_m will be consistent, that is $P_m(\omega) = \sum_{\rho>\omega} P_n(\rho)$, where the sum is over all n-step self-avoiding walks ρ extending ω. Hence one could define, via cylinder sets, the measure P_∞, that is, the uniformly random infinite self-avoiding walk.

The infinite self-avoiding walk is believed to exist in all dimensions, but this fact has only proved for $d > 4$ in the *spread-out model*, with steps (x, y) satisfying $0 \leq |x - y|_\infty \leq L$ for a suitable positive integer L; see Madras and Slade [22], Section 6.7. The proof relies on the convergence of the lace expansion. Also, Lawler [19] has constructed the infinite self-avoiding walk explicitly in high dimensions. In general dimensions what is available are merely some ratio limit theorems for the number of walks with specified end patterns, as in [22], Chapter 7.

More is known about a particular type of self-avoiding walk, the *bridge* on $\mathbb{Z}^d$. An n-step bridge is an n-step self-avoiding walk ω whose first components satisfy $\omega_1(0) < \omega_1(i) < \omega_1(n)$ for $1 \leq i \leq n$. For $n \geq m$ let $P^B_{m,n}(\omega)$ be the fraction of n-step bridges that extend an m-step self-avoiding walk ω. Note that this can only be non-zero if ω is a *half-space* walk, that is $\omega_1(0) < \omega_1(i)$ for all $i = 1, \dots, n$. Then a result in [22] Section 8.3 states that for any self-avoiding walk ω, the limit $P^B_m(\omega) = \lim_{n\to\infty} P^B_{m,n}(\omega)$ exists. This then guarantees existence and uniqueness of the probability measure P^B_∞ on infinite bridges. One proof of the result relies on renewal theory, but there is also an equivalent, constructive method of defining a measure on infinite bridges, somewhat in the spirit of our construction work.

Let b_n be the number of bridges of length n starting at the origin, and let c_n be the number of self-avoiding walks of length n starting at the origin. It is known—see [22] Sections 1.3 and 2.1 —that the limit

$$\mu = \lim_{n\to\infty} c_n^{1/n}$$

exists, and further that

$$\mu = \lim_{n\to\infty} b_n^{1/n} = \lim_{n\to\infty} \frac{b_{n+1}}{b_n}.$$

A bridge is said to be *irreducible* if one cannot decompose it into two non-trivial bridges [22]. Let λ_n denotes the number of n-step irreducible bridges; in [22], Chapter 4, Madras and Slade have shown using renewal theory that

$$\sum_{n=1}^{\infty} \frac{\lambda_n}{\mu^n} = 1.$$

This yields a probability distribution on irreducible bridges defined in an obvious way, and a process that builds a random infinite bridge one piece at a time. We select random pieces $\eta^{(1)}, \eta^{(2)}, \ldots$ from that distribution and define the infinite bridge ρ as a concatenation

$$\rho = \eta^{(1)} \cdot \eta^{(2)} \cdot \ldots ,$$

with corresponding probability measure Q_∞^B. Furthermore—see Section 8.3 of [22]—for every $m \geq 1$ and every m-step self-avoiding walk ω

$$Q_\infty^B\{\rho[0,m] = \omega\} = P_\infty^B\{\rho[0,m] = \omega\},$$

where $\rho[0,m]$ denotes the restriction of ρ to its first m steps. Thus indeed the "constructive" measure Q_∞^B is identical to the "existential" measure P_∞^B. It is an open problem whether an infinite bridge can be built one step at a time; and whether such a construction would yield the same probability measure.

For another example take lattice animals on a lattice L, where L can be the square lattice S, the hexagonal lattice H or the cubic lattice C. The link between connected subsets of these lattices and renewal sequences was explored in a paper by Rands and Welsh [25] —see also [26]. Let $A_n(L)$ be the set of equivalence classes under translation of the n-animals in the lattice L. For the square lattice S and the hexagonal lattice H one can define a unique bottom-left cell for any given animal A and a unique top-right cell; these definitions can be generalized to the cubic lattice. Hence it is easy to see that for each pair of positive integers m, n

$$A_n(L) \times A_m(L) \subseteq A_{m+n}(L), \quad L \in \{S, H, C\},$$

and hence

$$a_n(L)a_m(L) \leq a_{m+n}(L), \quad L \in \{S, H, C\},$$

where $a_n(L) = |A_n(L)|$. From that one can prove the existence of a constant $a(L)$ such that $a(L) = \lim_{n\to\infty}(a_n(L))^{1/n}$.

An n-animal A is *constructible* if $A \subseteq A_r(L) \times A_{n-r}(L)$ for some r, $0 < r < n$; otherwise A is inconstructible. Let $\Delta_n(L) = \{A \in A_n(L) : A \text{ is inconstructible }\}$ denote the set of inconstructible animals, and let $\delta_n(L) = |\Delta_n(L)|$. Then, using renewal techniques, Rands and Welsh show that

$$a_n(L) = \delta_n(L) + \sum_{r=1}^{n-1} \delta_r(L) a_{n-r}(L), \quad 1 \leq n < \infty.$$

Let $A(x) = \sum_{n=1}^{\infty} a_n(L)x^n$ and let $G(x) = \sum_{n=1}^{\infty} \delta_n(L)x^n$; standard calculations yield that

$$A(x) = \frac{1}{1 - G(x)}, \quad 0 \leq x < 1/a(L).$$

Furthermore, from the renewal theorem

$$G(1/a(L)) = \sum_{n=1}^{\infty} (a(L))^{-n} \delta_n(L) = \nu \leq 1.$$

[To conclude that ν is in fact equal to 1, one would have to prove that $A(x)$ is divergent at the critical point $x = 1/a(L)$; this is done in [6] but for a somewhat differently defined animal.] We can nevertheless construct an infinite lattice animal as a concatenation of inconstructible pieces, provided we endow each inconstructible piece of size n with a weight equal to $\nu^{-1}(a(L))^{-n}$. It is an open question whether the one-node-at-a-time technique that we used to build the infinite subtree applies here. If it does, then it would be interesting to know whether the resulting measure on infinite lattice animals coincides with the measure obtained by the piecewise construction described above.

Fig. 7 shows the edge-weights for a building scheme for lattice animals in the plane, of size up to 4. Symmetry classes are grouped for legibility. We leave it to the reader to try to extend the construction to infinity!

8 Conclusions

We have introduced the notion of a "building scheme", which constructs combinatorial objects incrementally, and in such a way that the object of each size n constructed by the process is a uniform random object of that size. Building schemes for subtrees of a d-branching tree (and some associated combinatorial structures) were constructed, and then used to make "uniformly random" infinite subtrees.

The general problem of devising a building scheme for a class of combinatorial objects, or even determining whether one exists, seems at this point to require *ad hoc* methods. We hope that our results show such schemes can provide insight, if not useful generation algorithms, in some combinatorial settings.

Acknowledgment

The authors are indebted to Svante Janson and to a couple of very thorough referees for pointing out some of the many connections to previous work. We suspect there are still many more to be found.

References

[1] M. Aigner, Combinatorial Theory, Springer-Verlag, Berlin, 1979.

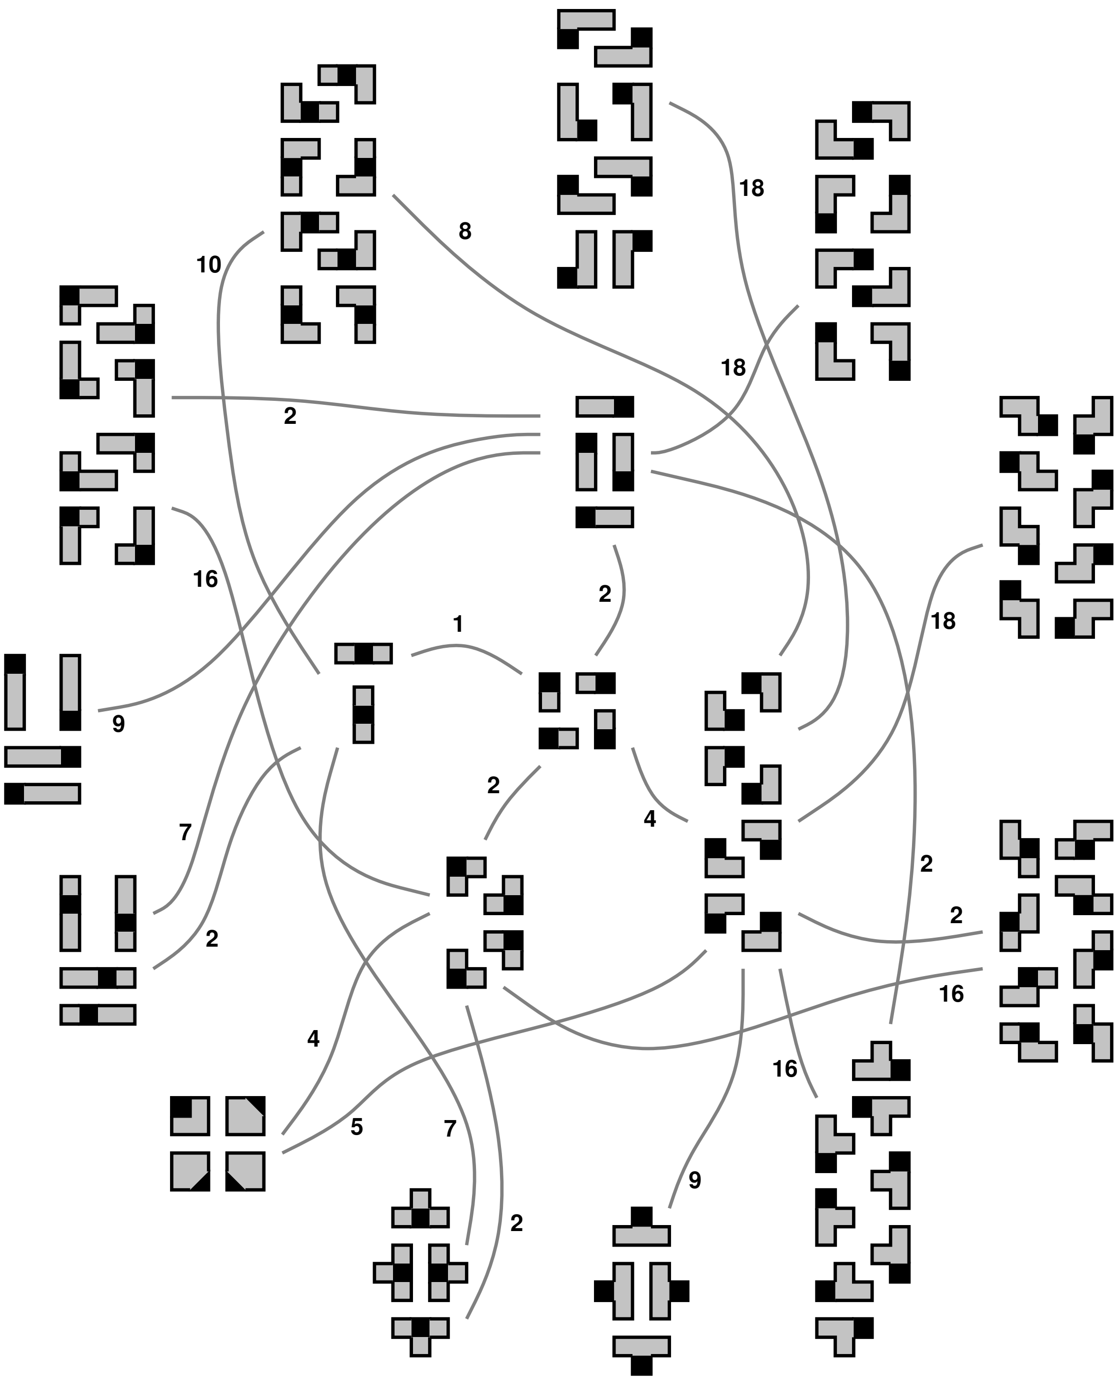

Figure 7: The start of a building scheme for lattice animals in the plane

[2] L. Alonso, J. L. Rémy and R. Schott, A linear-time algorithm for the generation of trees, Algorithmica 17(2) (1997), 162–182.

[3] K. B. Athreya and P. E. Ney, Branching Processes, Springer-Verlag, Berlin, 1972.

[4] C. Banderier, M. Bousquet-Mélou, A. Denise, P. Flajolet, D. Gardy and D. Gouyou-Beauchamps, Generating functions for generating trees, Discrete Math 246(1–3) (2002), 29–55.

[5] E. Barcucci, A. Del Lungo and E. Pergola, Random generation of trees and other combinatorial objects, Theor Comput Sci 218(2) (1999), 219–232.

[6] E. Barcucci, R. Pinzani, and R. Sprugnoli, The random generation of directed animals, Theor Comput Sci 127(2) (1994), 333–350.

[7] B. Bollobás, Combinatorics, Cambridge University Press, Cambridge, 1986.

[8] G. R. Brightwell and P. Winkler, Graph homomorphisms and phase transitions, J Combinat Theory B 77 (1999), 221–262.

[9] L. Comtet, Advanced Combinatorics, Reidel, Dordrecht, 1994.

[10] N. Dershowitz and S. Zaks, The cycle lemma and some applications, Europ J Combinatorics 11 (1990), 35–40.

[11] L. Devroye, Branching processes, in Probabilistic Methods for Algorithmic Discrete Mathematics, M. Habib, C. McDiarmid, J. Ramirez-Alfonsin and B. Reed (eds.), Springer-Verlag, Berlin, 1998.

[12] P. Flajolet and A.M. Odlyzko, The average height of binary trees and other simple trees, J Comput Syst Sci 25 (1982), 171–213.

[13] P. Flajolet, P. Zimmerman and B. van Cutsem, A calculus for the random generation of labelled combinatorial structures, Theor Comput Sci 132 (1994), 1–35.

[14] L. H. Harper, The morphology of partially ordered sets, J Combinat Theory A 17 (1974), 44–59.

[15] S. Janson, Ideals in a forest, one-way infinite binary trees and the contraction method, in Mathematics and Computer Science II, B. Chauvin, P. Flajolet, D. Gardy and A. Mokkadem (eds.), Trends in Mathematics 393–414, Birkhäuser, 2001.

[16] D. J. Kleitman, On an extremal property of antichains in partial orders. The LYM property and some of its implications and applications, in Combinatorics, M. Hall and J. van Lint (eds.), 77–90, Math. Centre Tracts 56, Amsterdam, 1974.

[17] D. E. Knuth, The Art of Computer Programming Vol. 1, Addison-Wesley, Reading MA, 1968.

[18] V.F. Kolchin, Random Mappings, Optimization Software Inc., New York, 1986 (translated from Slučajneye Otobraženija Nauka, Moscow, 1984).

[19] G. F. Lawler, The infinite self-avoiding walk in high dimensions, Ann Probab 17 (1989), 1367–1376.

[20] M. Lothaire, Combinatorics on Words, Encyclopedia of Mathematics and its Applications Vol. 17, Addison-Wesley, Reading, MA, 1983.

[21] R. Lyons, R. Pemantle and Y. Peres, Conceptual proofs of L log L criteria for mean behavior of branching processes, Ann Probab 23(3) (1995), 1125–1138.

[22] N. Madras and G. Slade, The Self-Avoiding Walk, Birkhäuser, 1993.

[23] A. Nijenhuis and H.S. Wilf, Combinatorial Algorithms, 2nd edition, Academic Press, New York, 1978.

[24] J. Pitman, Enumerations of trees and forests related to branching processes and random walks, Technical Report No. 482, University of California, 1997.

[25] B. M. I. Rands and D. J. A. Welsh, Animals, trees and renewal sequences, J Appl Mathematics 27 (1981), 1–17.

[26] B. M. I. Rands and D. J. A. Welsh, Animals, trees and renewal sequences: corrigendum, J Appl Mathematics 28 (1982), 107.

[27] G.N. Raney, Functional composition patterns and power series reversion, Trans AMS 94 (1960), 441–451.

[28] R. Sedgewick and P. Flajolet, Introduction to Analysis of Algorithms, Addison-Wesley, Reading MA, 1996.

[29] R. Sprugnoli, The generation of binary trees as a numerical problem, J ACM 39(2) (1992), 317–327.

[30] R.P. Stanley, Enumerative Combinatorics Vol. II, Cambridge University Press, Cambridge, UK, 1998.