PROBABILISTIC ALGORITHMS FOR TREES

Bruce E. Sagan
Michigan State University, East Lansing, MI 48824-1027

Yeong-Nan Yeh
Academica Sinica Nanking, Taipei, Taiwan, R.O.C.
(Submitted April 1987)

1. Introduction and Definitions

A *rooted tree*, τ, is a partially ordered set whose Hasse diagram is a tree (in the graph-theoretic sense of the term) having a unique minimal element called the *root*, see Figure 1a. If $|\tau| = n$, a *natural labeling* of τ is a bijection $T: \tau \to \{1, 2, \ldots, n\}$ such that $v < w$ in τ implies $T(v) < T(w)$. One such labeling is given in Figure 1b. In this case, we say T has *shape* τ. We let f_τ represent the number of natural labelings of τ.

The *hook* of a node $v \in \tau$ is

$$H_v = \{w \in \tau \mid w \geq v\}$$

with corresponding *hooklength* $h_v = |H_v|$. The hooklengths of our example tree are displayed in Figure 1c. The well-known hook formula [3] for the number of natural labelings states that

$$f_\tau = n! \Big/ \prod_{v \in \tau} h_v. \tag{1.1}$$

Thus, in our example $f_\tau = 7!/(7)(3)(2)(1)^4 = 120$.

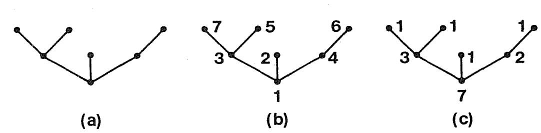

A tree, a labeling and the hooklengths

FIGURE 1

In Section 2 we will give a simple probabilistic proof of (1.1) inspired by an algorithm of Greene, Nijenhuis, and Wilf [1] for standard Young tableaux. The tree version has previously appeared in [5], but is included here for completeness. An algorithmic derivation of the hook-generating function for reverse tree partitions [which specializes to (1.1) as the variable approaches 1] can be found in [6].

A *Fibonacci tree* [9] is a finite lower-order ideal of the infinite poset in Figure 2a. The name derives from the easily proved fact that the number of Fibonacci trees with n nodes is the n^{th} Fibonacci number. For example, Figure 2b shows the five Fibonacci trees with four nodes. Let $\mathscr{F}_n$ be the set of all Fibonacci trees with n nodes, then

$$\sum_{\tau \in \mathscr{F}_n} f_\tau^2 = n! \qquad (1.2$$

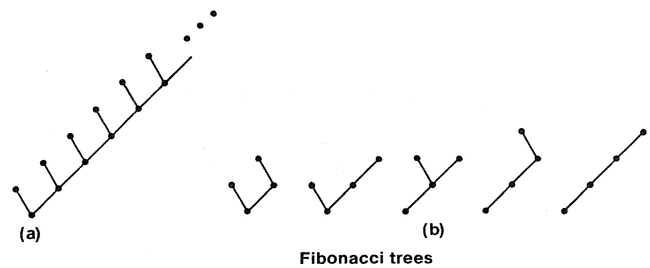

FIGURE 2

Formula (1.2) has a bijective proof due to Bender (reported in [9]). I Sections 3 and 4 below we will give two constructions that build a labeled tre $\tau \in \mathscr{F}_n$ with probability $f_\tau^2/n!$, thus proving (1.2) twice. The first algorith constructs the tree "from without" as done for tableaux in another paper o Greene et al. [2]. The second builds the tree "from within" and is based o work of Pittel [4].

2. Choosing a Labeling Uniformly

Let τ be a fixed shape with n nodes. The following algorithm can be use to choose a labeling of τ.

GNW1. Pick a node $v \in \tau$ uniformly at random, i.e., with probability $1/n$.

GNW2. If v is maximal (a leaf), then let $T(v) = n$ and return to GNW1 wit τ and n replaced by $\tau - \{v\}$ and $n - 1$, respectively (unless there are no node left, in which case the algorithm halts).

GNW3. If v is not maximal, then choose a different node $w \in H_v$ uniforml at random, i.e., with probability $1/(h_v - 1)$, and return to GNW2 with w in th role of v.

A sequence of nodes generated in the process of finding a vertex to b labeled (in this case by the loop between GNW2 and GNW3) is called a *trial*. A example of a typical trial is given in Figure 3.

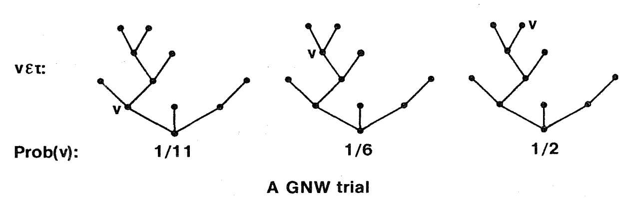

FIGURE 3

Theorem 1: If τ is a fixed rooted tree with n nodes, then GNW1-3 produce al labelings of τ uniformly at random. In fact, the probability of any give labeling is

$$\prod_{v \in \tau} h_v / n!$$

Proof: Let w be any maximal element of τ and let W be the set of vertices o the unique path from w to the root of τ (excluding w itself). Note that the are the only vertices whose hooklengths are changed if w is removed from τ during GNW2. Therefore, by induction, it suffices to show that the probability that w gets label n is

$$P(w) = (1/n) \prod_{v \in W} h_v/(h_v - 1)$$
$$= (1/n) \prod_{v \in W} \left(1 + \frac{1}{h_v - 1}\right).$$

But $1/n$ is the probability of choosing an initial node and each term in the expansion of the product corresponds to the probability of a unique trial ending in w. □

As an immediate corollary we have

Corollary 2: The number of labelings of a given tree τ with n nodes is

$$f_\tau = n! \Big/ \prod_{v \in \tau} h_v. \quad \square$$

3. Fibonacci Trees Grown from Without

It will be convenient to introduce coordinates for the infinite tree of Figure 2a. Let the nodes of the "spine" be $(i, 0)$ for $i = 0, 1, 2, \ldots$ while the leaves are denoted by $(i, 1)$ for the same range of i. Now, any Fibonacci tree can be specified by its coordinates as is done in Figure 4a.

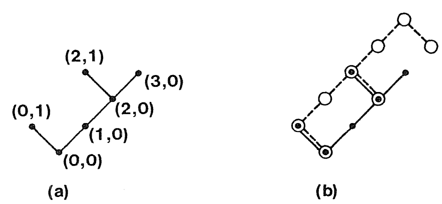

Coordinates and the associated tree

FIGURE 4

Given any vertex $v = (i, j)$, then v has *associate* $v' = (i, 1 - j)$. If τ is a Fibonacci tree with spine of length s, then the *associated tree* is

$$\tau' = \{v = (i, j) \mid v' \in \tau \text{ or } i = s + 1\};$$

see Figure 4b where the associated tree's nodes are the open circles. Note that τ' is "upside down" with root $r = (s + 1, 1)$.

Now suppose we wish to build a labeled Fibonacci tree, T. Assume that the first $m - 1$ vertices of T have already been constructed and given the labels $1, \ldots, m - 1$. Let τ be the current shape of T with associate τ' whose root is r. To add a node labeled m to T we proceed as follows:

WNG1. Choose a $v \in \tau' - \{r\}$ uniformly at random.

WNG2. If $v \notin \tau$, then add v to τ with label m and halt.

WNG3. If $v \in \tau$, say $v = (i, j)$, then return to WNG1 with τ' replaced by $\tau' - \{(i', j')\ i' \le i\}$.

Figure 5 presents an example of a trial generated by WGN1-3.

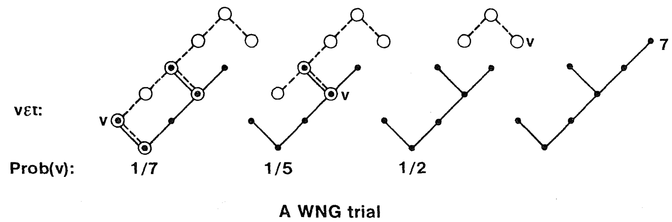

FIGURE 5

If this procedure is used iteratively for $m = 1, 2, \ldots, n$ to produce a labeled Fibonacci tree, then let $P(T)$ be the probability that labeling T is created. Thus, the total probability of producing a given shape τ is $P(\tau) = \Sigma P(T)$, where the sum is over all labelings T of τ.

Theorem 3: If τ is a Fibonacci shape with n nodes, then iteration of WNG1-3 produces all labelings of τ with total probability

$$P(\tau) = f_\tau^2/n!$$

Note: It is not true that WNG1-3 produces each labeling of τ with probability $P(T) = f_\tau/n!$.

Proof: Let τ have leaves $w_1, w_2, \ldots, w_k$ and define the subtrees $\tau_i = \tau - \{w_i\}$ for all i. Let $P(w_i|\tau_i)$ denote the probability that w_i gets labeled n after the algorithm constructs some labeling of τ_i. Hence, by the definitions above and induction,

$$P(\tau) = \sum_i P(\tau_i)P(w_i|\tau_i) = \sum_i (f_{\tau_i}^2/(n-1)!)P(w_i|\tau_i). \tag{3.1}$$

Let the w_i be arranged in order of increasing first coordinate, i.e.,

$$w_1 = (a_1, 1), \ldots, w_{k-1} = (a_{k-1}, 1), w_k = (a_k, j),$$

where $a_1 < \cdots < a_k$ and j may be 0 or 1. We need a couple of lemmas to help compute the quantities in (3.1).

Lemma 4: Let τ and the w_i be as above, then

$$f_\tau = \prod_{i=1}^{k-1} (n - a_i - i).$$

Proof: Using the hook formula (Corollary 2), we see that every term in the $n!$ is canceled by a hook of τ except those in the product above. □

Lemma 5: Let τ and the w_i be as above, then

$$P(w_i|\tau_i) = (1/n)\prod_{j=1}^{i-1}\left(1 + \frac{2}{n - a_j - j - 1}\right).$$

Proof: Initially we can pick any one of the n nodes in $\tau_i' - \{r\}$. Any trial ending at w_i can only pass through those w_j with $j < i$ and their associates w_j'. Landing on either of these two reduces the number of available nodes in $\tau_i' - \{r\}$ to $n - a_j - j - 1$, accounting for the second term of the binomial above. □

For notational convenience, let $b_i = n - a_i - i$. Hence, by Lemma 4,

$$f_\tau = b_1 b_2 \ldots b_{k-1}$$

and

$$f_{\tau_i} = (b_1 - 1) \ldots (b_{i-1} - 1) b_{i+1} \ldots b_{k-1}.$$

Also, from Lemma 5,

$$P(w_i | \tau_i) = (1/n)\left(1 + \frac{2}{b_1 - 1}\right) \ldots \left(1 + \frac{2}{b_{i-1} - 1}\right).$$

Thus,

$$f_{\tau_i}^2 P(w_i | \tau_i)/(n - 1)! = (1/n!)\left\{\prod_{1 \le j < i} (b_j^2 - 1)\right\}\left\{\prod_{i < j < k} b_j^2\right\}.$$

Plugging this expression into (3.1), we see that the sum of products telescopes (from the right-hand end) so that

$$P(t) = b_1^2 \ldots b_{k-1}^2/n! = f_\tau^2/n!$$

as desired. □

The obvious corollary is

Corollary 6: $\sum_{\tau \in \mathscr{F}_n} f_\tau^2 = n!$ □

We should also note that this algorithm has a "zone effect" similar to the original one for Young tableaux. Specifically, if $v = (a, 1)$ and $w = (b, 1)$ with $a_i < a$, $b < a_{i+1}$, then by Lemma 5 we have $P(v|\tau) = P(w|\tau)$. This observation will be useful in the next section.

4. Fibonacci Trees Grown from Within

Given $v \in \tau$, then v is a *singleton* if $v' \notin \tau$ and a *doubleton* otherwise. In Figure 6a, the singletons are (0, 0), (3, 0), (4, 0), and (6, 0), with the rest of the vertices being doubletons. If τ has a spine of length s, then the corresponding *extended tree* is

$$\tau'' = \tau \cup \{v' | v \in \tau \text{ is a singleton}\} \cup \{(s + 1, 0)\},$$

see Figure 6b. The elements of $\tau'' - \tau$ are organized into *zones*, which are maximal strings of vertices with consecutive first coordinates. Zones are numbered from the bottom up starting with zone 0, e.g., in Figure 6,

$$Z_0 = \{(0, 1)\},\ Z_1 = \{(3, 1), (4, 1)\},\ Z_2 = \{(6, 1), (7, 0)\}.$$

In the same way, the doubletons of τ are grouped into *bands* with band i directly below zone i. In our example, the bands are

$$B_0 = \emptyset,\ B_1 = \{(1, 0), (1, 1), (2, 0), (2, 1)\},\ B_2 = \{(5, 0), (5, 1)\}.$$

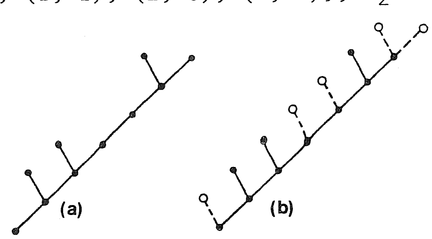

A tree and the extended tree

FIGURE 6

Finally, it will be convenient to have a total order on the vertices. If $v = (i, j)$ and $w = (x, y)$, then we will write $v \leq_t w$ if $i < x$ or $i = x$ and $j \leq y$.

Now, given a labeled Figonacci tree T of shape τ on $m - 1$ nodes, we find a node of $w \in \tau'' - \tau$ to label m by constructing a trial as follows. As usual, ":=" is the Pascal assignment symbol.

P1. Let $v := (0, 0)$ with probability 1. Let the set of predecessors of v be $P := \emptyset$.

P2. Set $P := P \cup \{v\}$.

P3. Pick w uniformly at random from among the set, D, of possible direct successors of $v = (i, j)$ defined by:

(a) if v is a doubleton, then $D = \{w \in \tau'' - P | w \geq_t v\}$.

(b) if v is a singleton, then let B be the band of largest index containing an element of P and let b be the maximum node of B (with respect to $\leq_t$). In this case

$$D = \{w \in \tau'' - P | w >_t b\} - \{w \text{ a singleton} | w \leq v\}.$$

If B does not exist, i.e., P consists only of singletons up to this point, then we take $b = (0, 0)$.

P4. If $w \in \tau'' - \tau$, then halt, else return to P2 with $w := v$.

Note that the trials generated by P1-4 do not necessarily respect the partial order in τ and the sequence of D's computed in P3 is not ordered by containment. For example, if a trial in the tree of Figure 6 has begun (0, 0), (4, 0), then the next node could be any one in τ'' except the two initial nodes and (3, 0). If the trial continues to (1, 1), then any nontrial vertex (i, j) with $i > 1$ is available for the next choice, including (3, 0). However, if the trial begins (0, 0), (1, 1), (4, 0), then the only possible successors are vertices (3, 1), (4, 1), (5, 0), (5, 1), (6, 0), (6, 1), and (7, 0).

Nevertheless, these rules do provide the desired distribution.

Theorem 7: If τ is a Fibonacci shape with n nodes, then iteration of P1-4 produces all labelings of τ with total probability

$$P(\tau) = f_\tau^2/n!$$

Proof: It suffices to show that Lemma 5 is still true when using P1-4. It will be convenient to reformulate the Lemma slightly for this setting. Let λ be a Fibonacci tree with $n - 1$ nodes and leaves $w_1, w_2, \ldots$ with first coordinates $a_1 < a_2 < \ldots$.

Lemma 8: With λ as above and $w \in \lambda'' - \lambda$ in the k^{th} zone, then the probability of terminating a P1-4 trial at w is

$$P(w) = (1/n) \prod_{w_j \in B_i,\ i \leq k} \left(1 + \frac{2}{n - a_j - j - 1}\right).$$

Proof: Induct on k. We will provide an explicit proof of the induction step, the anchor step being similar.

The trials t: $v_0 = (0, 0), v_1, \ldots, w$ are of two types, those that pass through an element of B_k and those that do not. The latter are in bijective probability preserving correspondence with trials $v_0, v_1, \ldots, w'$, where $w' \in Z_{k-1}$. In the former case, if $v_j \in B_k$ is the first such node then $v_0, \ldots, v_{j-1}, w'$ is a legal trial having the same probability as the initial segment of t. We will show below that the sum of the probabilities P of all possible final segments $v_j, v_{j+1}, \ldots, w$ is independent of both the particular node of B_k and the initial history of t. Thus, by induction, it suffices to demonstrate that

$$1 + P|B_k| = \prod_{w_j \in B_k} \left(1 + \frac{2}{n - a_j - j - 1}\right).$$

But the right side above telescopes to $(s + |B_k|)/s$, where s is the denominator corresponding to the largest leaf in B_k that has coordinates $(a, 1)$, say. It is easy to see that if we consider the subtree $\sigma = \{(i, j) \in \lambda | i > a\}$ then $s = |\sigma| + 1$. Hence, to finish the proof of the theorem, we need only show

Lemma 9: Let t, $v = v_j$, P, and σ be as above. Then P is independent of the set of nodes on t prior to v and of v itself (as long as $v \in B_k$). In fact, $P = 1/(|\sigma| + 1)$.

Proof: Let $\{v = u_1 \leq_t u_2 \leq_t \dots \leq_t u_m\}$ be the set of all possible vertices that could appear on t from v up to (but not including) w, i.e., the set of all elements above v that are either elements of B_k or singletons not previously on t. Because of these restrictions, the set of direct successors, $D(u_i)$, does not depend on the previous u_j chosen and, in fact, we have

$$\begin{aligned} D(u_i) &= \{u_j | j > i\} \cup \{v \in \sigma'' | v \text{ is not a singleton in } \sigma\} \\ &= D(u_{i-1}) - \{u_i\}. \end{aligned}$$

Thus,

$$|D(u_m)| = |\{v \in \sigma'' | v \text{ is not a singleton in } \sigma\}| = |\sigma| + 1$$

and

$$|D(u_i)| = |D(u_{i-1})| - 1.$$

Hence,

$$P = \frac{1}{|D(u_1)|}\left(1 + \frac{1}{|D(u_1)| - 1}\right) \cdots \left(1 + \frac{1}{|\sigma| + 1}\right) = \frac{1}{|\sigma| + 1}$$

as desired. □

Of course, Theorem 7 gives another proof of Corollary 6.

5. Remarks and Open Questions

Another point of similarity between Fibonacci trees and standard tableaux is the formula

$$\sum_{\tau \in \mathscr{F}_n} f_\tau = I_n, \tag{5.1}$$

where I_n is the number of involutions in the symmetric group S_n. The correspondence of Bender [9] mentioned in the introduction also proves (5.1). Is there a probabilistic way to demonstrate this, either for trees or tableaux?

A third family of posets that displays behavior similar to that of standard tableaux and rooted trees are the shifted standard tableaux [3]. The shifted analog of the hook formula (1.1) has been proved probabilistically by one of us [7]. It would be interesting to find an aleatory proof of the "sum of squares" equation in the shifted case (see [8] for the exact formula).

Finally, tableaux and shifted tableaux are intimately connected with representations of S_n. Ordinary tableaux give the degrees of ordinary irreducible representations (using matrices in GL_n), while their shifted cousins are related to projective ones (those using PGL_n, the projective linear group). In this setting, the analog of (1.2) expresses the fact that the sum of the squares of the irreducible degrees equals the order of the group. Can (1.2) itself be recast in this light? Specifically, is there a group of matrices G such that the degrees of the irreducible representations $\rho : S_n \to G$ are given by the f_τ?

References

1. C. Greene, A. Nijenhuis, & H. S. Wilf. "A Probabilistic Proof of a Formula for the Number of Young Tableaux of a Given Shape." *Adv. in Math 31* (1979):104-109.
2. C. Greene, A. Nijenhuis, & H. S. Wilf. "Another Probabilistic Method in the Theory of Young Tableaux." *J. Combin. Theory Ser. A* 37 (1984):127-135.
3. D. E. Knuth. *The Art of Computer Programming: Sorting and Searching*. Vol. 3. Reading, Mass.: Addison-Wesley, 1973.
4. B. Pittel. "On Growing a Random Young Tableau." *J. Combin. Theory Ser. A 41* (1986):278-285.
5. B. Sagan. "Partially Ordered Sets with Hooklenghts: An Algorithmic Approach." Ph. D. dissertation, M.I.T., 1979.
6. B. Sagan. "Enumeration of Partitions with Hooklengths." *European J. of Combin.* 3 (1982):85-94.
7. B. Sagan. "On Selecting a Random Shifted Young Tableau." *J. Algorithms 1* (1980):213-234.
8. B. Sagan. "Shifted Tableaux, Schur Q-Functions and a Conjecture of R. Stanley. *J. Combin. Theory Ser. A.* To appear.
9. R. P. Stanley. "The Fibonacci Lattice." *Fibonacci Quarterly 13* (1975): 215-232.

*****

Announcement

FOURTH INTERNATIONAL CONFERENCE ON FIBONACCI NUMBERS AND THEIR APPLICATIONS

Monday through Friday, July 30-August 3, 1990

Department of Mathematics and Computer Science Wake Forest University Winston-Salem, North Carolina 27109

International Committee: Horadam, A.F. (Australia), *Co-Chairman;* Philippou, A.N. (Cyprus), *Co-Chairman;* Ando, S. (Japan), Bergum, G. (U.S.A.), Johnson, M. (U.S.A.), Kiss, P. (Hungary), Filipponi, Piero (Italy), Campbell, Colin (Scotland), Turner, J.C. (New Zealand).

Local Committee: Fred T. Howard, *Co-Chairman;* Marcellus E. Waddill, *Co-Chairman;* Elmer K. Hayashi, Theresa Vaughan, Deborah Harrell.

CALL FOR PAPERS

The FOURTH INTERNATIONAL CONFERENCE ON FIBONACCI NUMBERS AND THEIR APPLICATIONS will take place at Wake Forest University, Winston-Salem, N.C., from July 30 to August 3, 1990. This Conference is sponsored jointly by the Fibonacci Association and Wake Forest University.

Papers on all branches of mathematics and science related to the Fibonacci numbers as well as recurrences and their generalizations are welcome. Abstracts are to be submitted by March 15, 1990, while manuscripts are due by May 1, 1990. Abstracts and manuscripts should be sent in duplicate following the guidelines for submission of articles found on the inside front cover of any recent issue of *The Fibonacci Quarterly* to:

Professor Gerald E. Bergum
The Fibonacci Quarterly
Department of Computer Science
South Dakota State University
P.O. Box 2201
Brookings, South Dakota 57007-0194