

**UNIFORM GENERATION OF
FORESTS OF RESTRICTED
HEIGHT**

M.D. Atkinson and J.-R. Sack

TR-230, OCTOBER 1993

School of Computer Science, Carleton University
Ottawa, Canada, K1S 5B6

Uniform generation of forests of restricted height

M.D. Atkinson

Department of Mathematical and Computational Sciences
North Haugh, St Andrews, Fife KY16 9SS, Scotland

J.-R. Sack⁺

School of Computer Science
Carleton University, Ottawa, Ontario, K1S 5B6, Canada

Keywords: Data structures; Design and analysis of algorithms; Trees and forests

Introduction

Tree structures occur throughout the theory of information storage. Typically each tree node contains some data information and some pointers to other tree nodes. In most cases the number k of pointer fields is fixed and the structure is referred to as a k -way tree. More formally, a k -way tree can be defined as being either empty (having no nodes) or as consisting of a root node together with an ordered sequence of k k -way trees. This formal definition suppresses the details about how much data is stored in a node and describes only the 'shape' of the k -way tree. The binary case $k=2$ is the most familiar. More generally we can define a k -way forest as being an ordered sequence of non-empty k -way trees. It is convenient to make a distinction between nodes which are roots of non-empty subtrees (internal nodes where data could be stored) and external nodes which define empty subtrees. The size of the forest is defined to be the number of internal nodes it contains.

It is very common, particularly in searching applications, to access the nodes of a forest by following pointers from one of the root nodes to another node in the forest. To measure how costly such searches are in the worst case the height of the forest is defined as the maximal number of nodes on any such search path. Often forests are organised so that their heights remain acceptably small in comparison to their size.

For developing and testing algorithms which manipulate forests it is useful to have methods for generating them in a uniform manner (so that each forest is generated with the same probability). Many methods have been given in the case of binary trees (2-way trees with one component) [1,2,3,7,8,9,10]. Some of these methods extend to k -way trees but there has been little previous work on uniformly generating trees of a fixed or bounded height.

⁺ Supported in part by the Natural Sciences and Engineering Research Council of Canada.

In this paper we consider k -way forests of fixed size n , height h , and number of components c . We give a generator for producing such forests uniformly at random thereby answering an open problem posed in [8]. Each random forest is generated in a linear number of operations in its size. The algorithm computes a table of values in a pre-processing phase. This table occupies $O(hn^2)$ locations and takes $O(hn^3)$ time to produce. Our method depends on a recurrence formula for the number of k -way forests of size n , height h , and number of components c . First we state the recurrence formula and then we describe our method.

Recurrences

In their study of the average height of a binary tree Brown and Shubert [4] (see also [5]) gave the following argument leading to a recurrence for the number of binary trees of a fixed height. Let $f(n,h)$ denote the number of binary trees on n nodes with height at most h . The quantities $f(n,h)$ have been investigated in [4,5] but no closed formula for them is known. Nevertheless they can be computed numerically by evaluating recurrence relations. It is clear that $f(n,0)=0$ if $n>0$ and $f(0,0)=1$. Moreover a non-empty binary tree of height at most h has left and right subtrees on j and $n-1-j$ nodes, for some j , which are of height at most $h-1$. Hence

$$f(n,h) = \sum_{j=0}^{n-1} f(j,h-1) * f(n-1-j,h-1)$$

Using this recurrence $f(n,h)$ can be computed in $O(n^3)$ steps and then the number of binary trees on n nodes of height exactly h can be computed as $f(n,h)-f(n,h-1)$. Once these quantities had been computed it would be possible, following the same type of argument as we give in the next section, to generate random binary trees of height h uniformly in time $O(n)$ per tree. This approach can be generalised to k -way trees but the analogous recurrence becomes a summation over all sequences (j_1, j_2, \dots, j_k) which sum to n and the overall time to compute the number of k -way trees rises to $O(n^{k+1})$. To obtain a manageable computation for general values of k we require a new approach. We regard a k -way tree as being formed in levels. Since the structure formed by all nodes below a certain level is, in general, a forest rather than a tree it is necessary to consider k -way forests. Although this requires somewhat more storage (because we have to keep track of how many components such forests have) the time requirement is then $O(n^3h)$ independent of k .

Let $t_k(n,h,c)$ denote the number of k -way forests with n internal nodes, height h , and c components.

- LEMMA 1
1. $t_k(n,h,c)=0$ if $c>n$ or $h>n$
 2. $t_k(n,0,c)=0$ unless $n=c=0$ when $t_k(0,0,0)=1$
 3. $t_k(n,h,c)=\sum_{s=0}^{kc} \binom{kc}{s} t_k(n-c,h-1,s)$ for $c,h \leq n$ and $h>0$

Proof. Parts 1 and 2 follow immediately from the definitions. For part 3 note that any k -way forest with n internal nodes, height h , and c components has c root nodes which have kc child nodes in all. Of these kc nodes some number s ($0 \leq s \leq kc$) will be internal nodes. The remainder of the forest is then determined by the forest with these s nodes as roots. The latter forest has $n-c$ internal nodes, height $h-1$, and s components so, once the s root nodes are chosen (and for a fixed s this can be done in $\binom{kc}{s}$ ways), there are $t_k(n-c,h-1,s)$ such forests. The result now follows.

Random Generation

We now discuss our method for producing k -way forests on n nodes uniformly at random. A k -way forest with n internal nodes, height h , and c components ($c \leq n$) is determined by

- (i) a choice (among the kc children of the root nodes of the forest components) of which s nodes are internal, and
- (ii) a choice of k -way forest with $n-c$ internal nodes, height $h-1$, and s components (rooted at the s nodes selected in (i))

We wish to select a forest so that every one of the $t_k(n,h,c)$ possibilities occurs with equal probability. The number of these forests which have precisely s internal nodes as children of the root is $\binom{kc}{s} t_k(n-c,h-1,s)$. Thus, to ensure that each of the forests occurs with equal probability we must choose the value s with probability

$$\frac{\binom{kc}{s} t_k(n-c,h-1,s)}{t_k(n,h,c)}$$

This can be done in the usual way: generate a random number in some range which is split into subranges of size proportional to these probabilities; select s according to the subrange in which the random number falls.

The next step is to select a subset of size s from among the kc children of the root nodes to be the internal children of the root nodes. A method for generating random combinations uniformly is given in [6, p. 137].

Finally we have to select, uniformly at random, a k -way forest with $n-c$ internal nodes, height $h-1$, and s components and root it at the set of s nodes chosen in the previous step (taken in left to right order). This is accomplished by recursive application of this method.

It is evident from the discussion that this procedure does generate k -way forests with n nodes, height h , and c components with equal probability. The forest so generated can be represented conveniently by an encoding given in [8] (primarily in the case of binary trees). In this encoding each forest is represented as a sequence of n 1's and $nk-n+c$ 0's. Essentially the sequence specifies which nodes on each level are internal. At the root level (the first level), all nodes are internal so the sequence begins with a segment of c 1's. At the second level there are kc nodes of which s , say, are internal; so the encoding now has a segment of length kc with precisely s 1's specifying which of the kc nodes are internal. At the third level there are ks nodes of which, say, t are internal; the coding will therefore have a segment of length ks with t 1's placed appropriately. It can be seen that the number u of 1's in each segment determines the length ku of the next segment. Since the first segment can be identified by its length $c = \text{number of 0's} - (k-1) * \text{number of 1's}$, the entire segmentation can be determined and so the encoding of the k -way forest uniquely determines the forest. Note however that not every binary sequence represents a forest; as pointed out in [8] there is a 'prefix' condition that has to be satisfied: the number of 1's is no less than the number of 0's in any prefix.

We now give a pseudo-code description of the algorithm in iterative terms which outputs the encoding of a random k -way forest with n internal nodes, height h , and c components:

GENERATE(n, h, c)

input: The number, c, of components of the k-way forest to be generated; its height, h; the number, n, of internal nodes

output: An encoding of the k-way forest generated uniformly at random

```

if c>n or h>n          then output "No k-way forest exists"
else if n=c=0 and h≠0 then output "No k-way forest exists"
    else if h=0 then encoding:= [ ]
        else
            begin
                encoding:=[1c];
                m:=n-c
                d:=c
                for j=1 to h do
                    begin
                        s:=select(m, h-j, d)
                        segment:=combination(k*d, s)
                        encoding:=append(encoding, segment)
                        m:=m - s
                        d:=s
                    end
                end
            end
        end
end

```

In this code the variables/functions have the following meanings:

encoding	The binary encoding of the k-way forest
j	The current level being generated
m	The number of nodes remaining to be generated
d	The number of nodes on the current level
select	A function which selects the integer s with probability proportional to $\binom{k*d}{s} t_k(m, h-j, s)$
combination	A function which returns a random combination of s integers from 1..k*d encoded as a bit vector
append	A function which concatenates its two arguments

Complexity Analysis

Assuming that the quantities $t_k(u, v, w)$ ($u \leq n, v \leq h, w \leq n$) are available the algorithm requires only $O(kn)$ operations. It constructs the k-way forest level by level. The construction time for a particular level is dominated by the time required to choose the number s of internal

nodes on the level (the `select` function) and the time required to obtain the s nodes themselves (the `combination` function). The `select` function is implemented by constructing a range which is split into $k*d$ subranges of length proportional to $\binom{k*d}{s} t_k(m, h-j, s)$. A random point in this range is generated and the value of s is selected according to the subrange containing the random point. The cost of these calculations is $O(k*d)$. However $k*d$ is the number of internal and external nodes on the next level. Since the total number of internal and external nodes is $nk+c$ the cost of all the `select` operations is $O(kn)$. Since the random combination generator of [6, p. 137] runs in time $O(k*d)$ a similar argument applies also to the total cost of the `combination` operations.

Before the random generation algorithm is used the values $t_k(u, v, w)$ must be calculated for $u \leq n$, $v \leq h$, $w \leq n$ (values $w > n$ may be required by the algorithm but we can exploit that $t_k(u, v, w) = 0$ for $w > v$). The recurrence formulae in Lemma 1 allow these n^2h quantities to be computed in $O(n^3h)$ steps. There are h levels of recursion, each step evaluates a sum of $kc+1$ quantities; each quantity is a product of a binomial and some $t_k(u, v, w)$. The values $t_k(u, v, w)$ and the binomials are known (after preprocessing) and can therefore be accessed in constant time. (the big O hides linear dependence on k since each $t_k(u, v, w)$ is a sum of $1+kw = O(kn)$ previously computed terms). Clearly the space requirement is $O(n^2h)$.

This discussion proves the following theorem.

THEOREM The algorithm `GENERATE` (n, h, c) is a uniform generator for k -way forests of height h with c components. Once the pre-processing stage is over (which has time requirement $O(n^3h)$ and space requirement $O(n^2h)$) it generates each k -way forest in time $O(n)$.

References

- [1] D.B. Arnold, M.R. Sleep: Uniform random number generation of n balanced parenthesis strings, *ACM Transactions on Programming Languages and Systems* 2 (1980), 122-128.
- [2] M.D. Atkinson and J.-R. Sack: Generating binary trees at random, *Inform. Process. Lett.* 41 (1992), 21-23.
- [3] M.D. Atkinson and J.-R. Sack: Uniform Generation of Combinatorial Objects in Parallel, to appear in *Journal of Parallel and Distributed Computing*, 1993.

- [4] G.B. Brown and B.O. Shubert: On random binary trees, *Math. Oper. Res.* 9 (1984), 43-65.
- [5] P. Flajolet, A. Odlyzko: The average height of binary and other simple trees, *J. Comput. and Syst. Sci.* 25 (1982), 171-213.
- [6] D.E. Knuth: *Semi-numerical Algorithms, The Art of Computer Programming*, vol 2, (2nd edition), Addison-Wesley, Reading, Massachusetts, 1981.
- [7] J.F. Korsh: Counting and randomly generating binary trees, *Inform. Process. Lett.* 45 (6) (1992), 291-294.
- [8] C.C. Lee, D.T. Lee, C.K. Wong: Generating binary trees of bounded height, *Acta Inform.* 23 (1986), 529-544.
- [9] H.W. Martin, B.J. Orr: A random binary tree generator, *Computing Trends in the 1990's*, ACM Seventeenth Computer Science Conference, Louisville, Kentucky (February 1989), ACM Press, 33-38.
- [10] R. Sprugnoli: The generation of binary trees as a numerical problem, *JACM* 39 (1992), 317-327.

**School of Computer Science, Carleton University
Recent Technical Reports**

- TR-210 Sorting Permutations with Networks of Stacks**
M.D. Atkinson, August 1992
- TR-211 Generating Triangulations at Random**
Peter Epstein and Jörg-Rüdiger Sack, August 1992
- TR-212 Algorithms for Asymptotically Optimal Contained Rectangles and Triangles**
Evangelos Kranakis and Emran Rafique, September 1992
- TR-213 Parallel Algorithms for Rectilinear Link Distance Problems**
Andrzej Lingas, Anil Maheshwari and Jörg-Rüdiger Sack, September 1992
- TR-214 Camera Placement in Integer Lattices**
Evangelos Kranakis and Michel Pocchiola, October 1992
- TR-215 Labeled Versus Unlabeled Distributed Cayley Networks**
Evangelos Kranakis and Danny Krizanc, November 1992
- TR-216 Scalable Parallel Geometric Algorithms for Coarse Grained Multicomputers**
Frank Dehne, Andreas Fabri and Andrew Rau-Chaplin, November 1992
- TR-217 Indexing on Spherical Surfaces Using Semi-Quadcodes**
Ekow J. Otoo and Hongwen Zhu, December 1992
- TR-218 A Time-Randomness Tradeoff for Selection in Parallel**
Danny Krizanc, February 1993
- TR-219 Three Algorithms for Selection on the Reconfigurable Mesh**
Dipak Pravin Doctor and Danny Krizanc, February 1993
- TR-220 On Multi-label Linear Interval Routing Schemes**
Evangelos Kranakis, Danny Krizanc, and S.S. Ravi, March 1993
- TR-221 Note on Systems of Polynomial Equations over Finite Fields**
Vincenzo Acciario, March 1993
- TR-222 Time-Message Trade-Offs for the Weak Unison Problem**
Amos Israeli, Evangelos Kranakis, Danny Krizanc and Nicola Santoro, March 1993
- TR-223 Anonymous Wireless Rings**
Krzysztof Diks, Evangelos Kranakis, Adam Malinowski, and Andrzej Pelc, April 1993
- TR-224 A consistent model for noisy channels permitting arbitrarily distributed substitutions, insertions and deletions**
B.J. Oommen and R.L. Kashyap, June 1993
- TR-225 Mixture Decomposition for Distributions from the Exponential Family Using a Generalized Method of Moments**
S.T. Sum and B.J. Oommen, June 1993
- TR-226 Switching Models for Non-Stationary Random Environments**
B. John Oommen and Hassan Masum, July 1993
- TR-227 The Probability of Generating Some Common Families of Finite Groups**
Vincenzo Acciario, September 1993
- TR-228 Power Roots of Polynomials over Arbitrary Fields**
Vincenzo Acciario, September 1993
- TR-229 Optimal Parallel Algorithms for Direct Dominance Problems**
Amitava Datta, Anil Maheshwari and Jörg-Rüdiger Sack, October 1993
- TR-230 Uniform Generation of Forests of Restricted Height**
M.D. Atkinson and J.-R. Sack