Balanced allocations for tree-like inputs

Andrei Z. Broder* Alan Frieze† Carsten Lund‡
Steven Phillips§ Nick Reingold¶

June 23, 1995

Abstract

We consider the following procedure for constructing a directed tree on n vertices: The underlying undirected tree is fixed in advance but the edges of the tree are presented in a random order (all orders are equally likely); each edge is oriented towards its endpoint that has the lower indegree at the time of its insertion. The question is what is $\mathbf{E}(\mathcal{M}(n))$, the expected maximum indegree? As we shall explain, this problem has connections with balanced allocations and with on-line load balancing.

Previous results by Azar, Naor, and Rom imply that if the insertion order is arbitrary, for any tree, $\mathcal{M}(n) = O(\log n)$ and that there are trees and insertion orders for which $\mathcal{M}(n) = \Omega(\log n)$. On the other hand, results by Azar, Broder, Karlin, and Upfal imply that if both the underlying tree and the insertion order are random, then $\mathbf{E}(\mathcal{M}(n)) = \Theta(\log \log n)$. Here we show an intermediate result: for any tree if the insertion order is random, then $\mathbf{E}(\mathcal{M}(n)) = O(\log n / \log \log n)$ and there are trees for which $\mathbf{E}(\mathcal{M}(n)) = \Omega(\log n / \log \log n)$.

Keywords: Analysis of algorithms, combinatorial problems, theory of computation, balanced allocations, online load balancing.

*Digital Systems Research Center, 130 Lytton Avenue, Palo Alto, CA 94301, USA. E-mail: broder@src.dec.com

†Department of Mathematics, Carnegie-Mellon University. A portion of this work was done while the author was visiting Digital SRC. Supported in part by NSF grants CCR9024935 and CCR9225008. E-mail: af1p@euler.math.cmu.edu

‡AT&T Bell Laboratories. E-mail: lund@research.att.com

§AT&T Bell Laboratories. E-mail: phillips@research.att.com

¶AT&T Bell Laboratories. E-mail: reingold@research.att.com

1 Introduction

We consider the following procedure for constructing a directed tree on n vertices: The underlying undirected tree is fixed in advance but the edges of the tree are presented in a random order (all orders are equally likely); each edge is oriented towards its endpoint that has the lower indegree at the time of its insertion. The question is what is $\mathbf{E}(\mathcal{M}(n))$, the expected maximum indegree? We show that for any tree if the insertion order is random $\mathbf{E}(\mathcal{M}(n)) = O(\log n/\log\log n)$ and there are trees for which $\mathbf{E}(\mathcal{M}(n)) = \Omega(\log n/\log\log n)$. The motivation for this problem comes from two directions: balanced allocations and on-line load balancing.

Suppose that we sequentially place n balls into n boxes by putting each ball into a randomly chosen box. This is known as the random allocation process and it has been extensively studied in the probability and statistics literature. (See e.g. [8, 9].) A standard result is that when the process has terminated, the fullest box has, with high probability (that is, $1 - o(1)$), $\ln n/\ln\ln n(1 + o(1))$ balls in it. Azar, Broder, Karlin, and Upfal [4] consider a variant whereby each ball comes with d possible destinations, chosen independently and uniformly at random. The ball is placed in the least full box among the d possible locations at the time of its insertion. Surprisingly when the process terminates the fullest box has only $\ln\ln n/\ln d + O(1)$ balls in it, an exponential difference. (This allocation method is called the *greedy* algorithm and the number of balls in a box is often called the *load* of the box.)

Given this gap, it is natural to consider an in-between model: assume that each ball has d possible destinations, fixed in advance, but the order in which the balls are inserted is random. If we restrict our attention to the case $d = 2$, we can view this process as the construction of a graph G with with n vertices and n edges as follows: the edges are presented in a random order (all orders are equally likely) and each edge is oriented towards its endpoint that has the lower indegree at the time of its insertion. Edges correspond to balls, the two endpoints of an edge are its two possible choices, and the orientation is done in the greedy fashion described above. The question becomes what is $\mathbf{E}(\mathcal{M}(n))$ the expected maximum indegree?

Clearly the maximum load can be as high as $n/2$ if all the choices for all the balls are the same two boxes. To avoid these trivial situations it is convenient to consider only the situation when in fact there is a way of putting the balls in boxes that is consistent with the available choices for each ball and that results in a maximum load equal to one. This means that the graph G has the property that the number of edges induced by any set of j vertices contains at most j edges. Furthermore if G is not connected the problem decomposes in an obvious way. So without loss of generality we can assume that G is connected which means G is actually a tree plus one extra edge. Finally it is easy to show that the addition of one edge changes $\mathcal{M}(n)$ by at most one, so we can assume that G is a tree T, which is our model. The proofs of [4] can be modified to show that if T is chosen uniformly at random then $\mathbf{E}(\mathcal{M}(n)) = \Theta(\log \log n)$.

The connection with on-line load balancing is as follows: We are given a set of n servers and a sequence of tasks. Each task comes with a list of servers on which it can be executed. The load balancing algorithm has to assign each task to a server on-line, with no information on future arrivals. The goal of the algorithm is to minimize the maximum load on any server. The quality of an on-line algorithm is measured by the *competitive ratio*: the ratio between the maximum load it achieves and the maximum load achieved by the optimal off-line algorithm that knows the whole sequence in advance. This load balancing problem models for example, communication in heterogeneous networks containing workstations, I/O devices, etc. Servers correspond to communication channels and tasks to requests for communication links between devices. A network controller must coordinate the channels so that no channel is too heavily loaded.

On-line load balancing has been studied extensively against worst-case adversaries [1, 2, 3, 5, 6, 7, 10]. What we are considering here are *permanent tasks*, that is, tasks that arrive but never depart. For this case, Azar, Naor and Rom [7] showed that the competitive ratio of the greedy algorithm is $O(\log n)$ and that no algorithm can do better. Their proof can be slightly modified to show that, in our terminology, for every tree and every order of insertion $\mathcal{M}(n) = O(\log n)$ and that there are trees and insertion orders for which $\mathcal{M}(n) = \Omega(\log n)$.

2 Upper bound

We consider the following procedure for constructing a tree on n vertices. The final tree T is fixed in advance, but each edge e is inserted at some time $t(e)$, where the t(e) are independent variables identically distributed uniformly over $[0, 1]$. Note that this implies that all insertion orders are equally likely.

Each edge e is oriented towards the endpoint that has the lower indegree at the time of e's insertion. The question is: for a given T, what is the expected maximum indegree?

Let $P(n,t,i)$ denote the maximum probability (over all trees) that a tree with n nodes contains an i-proof for the root at time t, where an i-proof for v at time t is a set of edges and their times that force the indegree on v to be at least i at time t.

Claim 1 *For any tree T*

$$P(n,t,i) \le \frac{nt^{i/2}}{i^{\alpha i}}, \tag{1}$$

where $\alpha = \ln 2 / \ln(4e^2)$.

Proof: The proof is by induction on i. Base case, $i = 1$:

$$P(n,t,1) \le \mathbf{Pr}(\exists \text{ one edge at time } t) \le nt \le nt^{1/2}.$$

The general case. Suppose that the root r has degree k and that its children are the vertices $1, \ldots, k$. We claim that

$$P(n,t,i) \le \sum_{1 \le j \le k} \int_0^t P(n-n_j, x, i-1) P(n_j, x, i-1)\, dx,$$

where n_j is the number of nodes in the tree hanging from j. The reason is that if at time x the edge $\{r,j\}$ completes an i-proof, there must exist prior to this time an $(i-1)$-proof for r that does not involve j, and an $(i-1)$-proof for j. These two events are independent since the trees that can contribute edges to these two proofs are disjoint, namely they are the two trees obtained from T by removing the edge $\{r,j\}$. By the induction hypothesis

$$\begin{aligned} P(n,t,i) &\le \sum_{1 \le j \le k} \int_0^t \frac{n_j(n-n_j)x^{i-1}}{(i-1)^{2\alpha(i-1)}}\, dx \\ &= \sum_{1 \le j \le k} \frac{n_j(n-n_j)t^i}{i(i-1)^{2\alpha(i-1)}} \le \frac{n^2 t^i}{i(i-1)^{2\alpha(i-1)}} \end{aligned}$$

Now either $nt^{i/2}/i^{\alpha i} \ge 1$ in which case (1) is trivial, or $nt^{i/2} \le i^{\alpha i}$ in which case

$$
\begin{aligned}
P(n,t,i) &\le \frac{nt^{i/2}}{i^{\alpha i}} \frac{nt^{i/2} i^{\alpha i}}{i(i-1)^{2\alpha(i-1)}} \le \frac{nt^{i/2}}{i^{\alpha i}} \frac{i^{2\alpha i}}{i(i-1)^{2\alpha(i-1)}} \\
&= \frac{nt^{i/2}}{i^{\alpha i}} \frac{(1+1/(i-1))^{2\alpha(i-1)}}{i^{1-2\alpha}} \le \frac{nt^{i/2}}{i^{\alpha i}} \frac{e^{2\alpha}}{i^{1-2\alpha}} \\
&\le \frac{nt^{i/2}}{i^{\alpha i}} \frac{e^{2\alpha}}{2^{1-2\alpha}} \le \frac{nt^{i/2}}{i^{\alpha i}}
\end{aligned}
$$

provided that

$$
\alpha \le \frac{\ln 2}{2(1+\ln 2)} = 0.204\ldots
$$

□

Corollary 2 *For any tree T on n vertices the expected maximum indegree is $O(\log n / \log\log n)$.*

Proof: Fix a tree T. From the claim above, for C a sufficiently large constant, $\mathbf{Pr}(\mathcal{M}(T) > C \ln n / \ln\ln n) < 1/n$. On the other hand $\mathcal{M}(n) \le n$. □

3 Lower Bound

For the lower bound, we assume that ties are broken randomly: when inserting an edge whose endpoints have the same indegree, the edge is oriented towards a randomly chosen endpoint.

Let $T_{k,d}$ be the k-ary tree of depth d, where d is the largest integer satisfying

$$
d\left(1 - \frac{1}{4d}\right)^k \le \frac{1}{4}. \tag{2}
$$

The *height* of a vertex in $T_{k,d}$ is the minimum distance to a leaf, so for example leaves have height 0 while the root has height d. For a node v in $T_{k,d}$, let the random variable $L(v,t)$ denote the load (indegree) of v at time t and let $\mathcal{E}(v,t)$ be the event that the edge to the parent of v has not yet been inserted at time t.

Claim 3 *Let v be a vertex with height at least i in $T_{k,d}$. Let $f(i)$ be the maximum over such v of $\mathbf{Pr}(L(v,i/d) < i \mid \mathcal{E}(v,i/d))$. Then for $0 \le i \le d$,*

$$
f(i) \le i\left(1 - \frac{1}{4d}\right)^k.
$$

Proof: The proof is by induction on i. The base case, $i = 0$, is clearly true. For the inductive step, assume the claim is true for $i - 1$ and let v have height at least i. If $L(v, i/d) < i$, then at least one of the following must be true:

1. $L\left(v, \frac{i-1}{d}\right) < i - 1$.

2. For each child w of v such that the edge $\{v, w\}$ is inserted during the interval $\left(\frac{i-1}{d}, \frac{i}{d}\right]$, either $L\left(w, \frac{i-1}{d}\right) < i-1$, or $L\left(w, \frac{i-1}{d}\right) = L\left(v, \frac{i-1}{d}\right)$ and the tie is broken in favor of w.

The first event occurs with probability at most $f(i-1)$. Secondly, consider a child w of v. The edge $\{v, w\}$ has probability $1/d$ of being inserted during the interval $\left(\frac{i-1}{d}, \frac{i}{d}\right]$. If it is inserted during this interval, then by the inductive hypothesis,

$$\mathbf{Pr}\left(L\left(w, \frac{i-1}{d}\right) < i - 1 \;\middle|\; \mathcal{E}\left(w, \frac{i-1}{d}\right)\right) \leq f(i-1).$$

Therefore the second event occurs with probability at most

$$\sum_{j=0}^{k} \binom{k}{j} \left(\frac{1}{d}\right)^j \left(1 - \frac{1}{d}\right)^{k-j} \left(f(i-1) + \frac{1}{2}\right)^j$$
$$= \left(\frac{1}{d}\left(f(i-1) + \frac{1}{2}\right) + 1 - \frac{1}{d}\right)^k \leq \left(1 - \frac{1}{4d}\right)^k$$

since $f(i-1) \leq 1/4$ by the choice of d and the events under consideration are independent. Thus

$$f(i) \leq f(i-1) + \left(1 - \frac{1}{4d}\right)^k \leq i\left(1 - \frac{1}{4d}\right)^k.$$

□

Corollary 4 *For any constant $c < 1$ and for large enough n, there is a tree T on at most n vertices for which the expected maximum indegree is at least $c \ln n / \ln\ln n$.*

Proof: Fix $c < 1$ and let k be sufficiently large. Let d be given by (2) above, so $d = (k/(4\ln k))(1 + o(1))$. The number of nodes in $T_{k,d}$ is $n =$

$1 + k + \cdots + k^d$, so $n = e^{(k/4)(1+o(1))}$ and therefore $k = 4 \ln n(1 + o(1))$ and $d = (\ln n / \ln \ln n)(1 + o(1))$.

For a fixed child of the root, the probability that it is not connected to the root and has indegree at least $d - 1$ at time $1 - 1/d$ is at least $3/(4d)$. Hence the probability that there exists such a child is at least

$$1 - \left(1 - \frac{3}{4d}\right)^k > 1 - \left(1 - \frac{1}{4d}\right)^k \geq 1 - \frac{1}{4d}.$$

The last inequality follows from (2) above. Therefore the expected maximum indegree of $T_{k,d}$ is at least $(d - 1)(1 - 1/(4d)) > d - 5/4$, which for large enough n is greater than $c \ln n / \ln \ln n$. $\square$

References

[1] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput competitive online routing. In *Proceedings of the 34th IEEE Conference on Foundations of Computer Science*, pages 32–40, 1993.

[2] B. Awerbuch, Y. Azar, S. Plotkin, and O. Waarts. Competitive routing of virtual circuits with unknown duration. In *Proceedings of 5th ACM/SIAM Symposium on Discrete Algorithms*, pages 321–327, 1994.

[3] Yossi Azar, James Aspnes, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line load balancing with applications to machine scheduling and virtual circuit routing. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 623–631, May 1993.

[4] Yossi Azar, Andrei Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 593–602, May 1994.

[5] Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. On-line load balancing. *Theoretical Computer Science*, 130:73–84, 1994.

[6] Yossi Azar, B. Kalyanasundaram, Serge Plotkin, K. Pruhs, and Orli Waarts. Online load balancing of temporary tasks. In *Workshop on Algorithms and Data Structures (WADS)*, Montreal, Canada, 1993.

[7] Yossi Azar, Joseph Naor, and Raphael Rom. The competitiveness of on-line assignments. In *Proceedings of the 3rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 203–210, 1992.

[8] Norman L. Johnson and Samuel Kotz. *Urn Models and Their Application.* John Wiley & Sons, 1977.

[9] Valentin F. Kolchin, Boris A. Sevastyanov, and Valdimir P. Chistyakov. *Random Allocations.* John Wiley & Sons, 1978.

[10] S. Phillips and J. Westbrook. Online load balancing and network flow. In *Proc. 25th ACM Symposium on Theory of Computing*, pages 402–411, 1993.