

OPTIMAL PLACEMENT OF IDENTICAL RESOURCES IN A DISTRIBUTED NETWORK[†]

Michael J. Fischer
University of Washington

Nancy D. Griffeth
Georgia Institute of Technology

Leo J. Guibas
Xerox Palo Alto Research Center

Nancy A. Lynch
Georgia Institute of Technology

ABSTRACT

The problem is considered of locating a number of identical resources at nodes of a tree so as to minimize the total expected cost of servicing a set of random requests for the resources. The cost of servicing a request is the tree distance from the requesting node to the node at which the resource satisfying the request is located. An algorithm for finding an optimal placement of t resources is presented which runs in time $O(t^2 \cdot |E|)$, where E is the edge set of the tree. In the special case of a complete binary tree with requests uniformly distributed over the n leaves, another algorithm works in time $O(t \log_2 n)$.

While optimal placements in general seem difficult to characterize, there is a very simple placement whose total cost differs from optimality by at most $|E|$. The expected cost of this placement on a complete binary tree is $C(n + \sqrt{tn})$.

1. INTRODUCTION

We consider the problem of locating some number t of identical resources at nodes of a distributed network in such a way as to minimize the expected "cost" of servicing a random set of t requests for those resources. Various different costs are likely to be important in different situations.

For example, suppose the resources are processors and requests are to establish a virtual connection with a processor which will be used for a very long period of time. In this case one might want to minimize the expected total of all the network distances (measured in some appropriate way) between requesting users and their assigned processors. Because of the long holding times, it is probably reasonable to expend considerable effort to find a good matching of requests to

resources. The total of the network distances provides a measure of the expected communication traffic introduced into the network by the computations.

For another example, suppose the resources are tickets to sporting events (or airline seats). A relevant cost is the expected waiting time until a buyer receives his ticket, or equivalently, the expected total waiting time for all buyers. In this situation, it is probably not reasonable to expend much effort in attaining a near-optimal matching, for the time to find the matching can easily exceed the eventual savings in locating a nearby ticket. Even so, the expected total distance in an optimal matching is significant as a lower bound for the expected total waiting time.

This paper investigates properties of optimal resource placements and provides fast algorithms for finding them. It also provides upper bounds for the costs of optimal placements and of certain nearly-optimal and easily-described placements. The network in all cases is assumed to be configured as a tree. Nevertheless, the upper bounds can often be applied to an arbitrary (connected) distributed network by constructing a spanning tree.

In this paper, we allow for the possibility that fractions of resources, and not only whole resources, might be located at some nodes of the network tree. This is reasonable if, for instance, the resources are large blocks of available data storage space. It would be perfectly permissible for a user to obtain parts of his needed storage from several different nodes. On the other hand, we assume that the requests are discrete -- each request is for one unit of resource.

In Section 2, we present our notation, definitions, and those results which apply to arbitrary trees and arbitrary probability distributions for arrivals of requests. §2.1 defines matchings and relates them to network flows. §2.2 defines the expected total cost of a placement in terms of expected total flow. Both the function describing the expected flow on each edge and the function describing the minimum possible expected total flow for any subtree are of a particularly simple form -- they are unbounded, convex, piecewise linear functions on the nonnegative reals, with

[†]This research was supported in part by the National Science Foundation under grants MCS77-02474, MCS77-15628, MCS78-01698, MCS80-03337, U.S. Army Research Office Contract Number DAAG29-79-C-0155, and the Office of Naval Research Contract Numbers N00014-79-C-0873 and N00014-80-C-0221.

all singularities at integers. This immediately implies (in §2.3) that there are always optimal resource placements consisting of whole numbers of resources located at each node; it is never necessary to place fractions of resources at any node in order to achieve an optimal placement. An algorithm using $O(t^2 \cdot \#edges)$ arithmetic operations is presented which always finds an optimal whole resource placement. This is the fastest algorithm we have which is completely general. Section 4 contains faster algorithms for special cases.

§2.4 considers what is required for a placement to optimize the expected flow over any particular edge in the tree -- i.e. to be locally optimal. Unfortunately, it is not possible in general to obtain a single placement which simultaneously optimizes the expected flow on all edges.

A fair whole placement is one which "almost optimizes" the flow on each edge -- that is, one in which the number of resources in each subtree is either the mean rounded up or down. It is always possible (and quite easy) to obtain a fair whole placement, and such placements are close to optimal.

In §2.5, we show that if there is a probability less than $\frac{1}{2}$ that any request will arrive in a particular subtree, then it is always bad to place even a very small fraction of a resource anywhere in that subtree.

In Section 3, we analyze the cost of an optimal placement for the case of a complete binary tree of n leaves but with an arbitrary probability distribution for request arrivals. An arbitrary fair whole resource placement has expected total cost at most $O(n + \sqrt{t} \sqrt{n})$. Note that in the very important case where t is roughly proportional to n (i.e. the number of resources in the network is proportional to the number of nodes), that the expected average cost per request is bounded by a constant, independent of the size of the network. This situation is very different from the case of centralizing the resources, where the average cost grows proportionately with the log of the number of nodes in the network.

2. NOTATION, DEFINITIONS AND GENERAL RESULTS

2.1. Trees, Matchings and Flows

A tree $T = (V, E)$ is an undirected acyclic graph, where V is the vertex (node) set and E is the edge set.

Let N denote the natural numbers, including 0, and let R^+ denote the nonnegative reals.

A set of requests is described by a function $r: V \rightarrow R^+$ such that $r(v)$ is the number of requests originating at node v . A placement of resources is described by a function $s: V \rightarrow R^+$

such that $s(v)$ is the number of resources placed at node v . We always assume $\text{total}(r) = \text{total}(s)$, where for any $g: V \rightarrow R^+$, $\text{total}(g) = \sum_{v \in V} g(v)$.

A matching is a function $m: V \times V \rightarrow R^+$. $m(u, v)$ gives the number of requests at node u which are satisfied by resources at node v . m is exact for r, s if for all $u, v \in V$, $\sum_{w \in V} m(u, w) = r(u)$ and $\sum_{w \in V} m(w, v) = s(v)$. Thus, an exact matching gives a complete correspondence between requests and resources. Clearly, an exact matching exists whenever $\text{total}(r) = \text{total}(s)$.

The cost of a matching is defined to be

$$\text{cost}(m) = \sum_{u, v \in V} m(u, v) \cdot d(u, v)$$

where $d(u, v)$ is the number of edges in the unique path from u to v in T . Let

$$\text{cost}(r, s) = \min \{ \text{cost}(m) \mid m \text{ is an exact matching for } r, s \}$$

A matching m is optimal for r, s if $\text{cost}(m) = \text{cost}(r, s)$.

It is convenient to relate a matching to a flow in a directed graph. Choose a node w , and direct the edges of the tree to point away from w . Corresponding to the usual way of drawing trees with the root at the top, we let $\text{high}(e)$ be the endpoint of edge e which is closest to w , and $\text{low}(e)$ be the endpoint farthest from w .

In the remainder of this paper, it will be convenient for w itself to have an edge entering it. Therefore, we add a new node v to serve as the root, and we let $\text{rootedge}(T) = (v, w)$, an edge from v to w . Resources will never be placed on v , and requests will never originate at v ; v and $\text{rootedge}(T)$ are just notational conveniences.

A flow is a function $f: E \rightarrow R$ which describes the "movement" of resources from their initial placement to corresponding requests. A positive value of $f(e)$ denotes a flow along the direction of the edge (i.e. towards the leaves) whereas a negative value of $f(e)$ denotes a flow towards the root. A flow is stable for r, s if for every $v \in V$,

$$r(v) + \sum_{e: \text{high}(e)=v} f(e) = s(v) + \sum_{e': \text{low}(e)=v} f(e')$$

Thus, at every node, the flow out equals the flow in.

The cost of a flow is defined to be

$$\text{cost}(f) = \sum_{e \in E} |f(e)|$$

Call a flow f optimal for r, s if it is stable

for r, s and has minimal cost over all such flows.

The following theorem formalizes the intuition the flows correspond to resources moving along edges and that in an optimal matching, resources do not move both directions along the same edge. We omit the straightforward but tedious proof.

Theorem 2.1.1. Let m be an optimal matching and f an optimal flow for r, s . Then $\text{cost}(m) = \text{cost}(f)$. Moreover, if either m or f has an integral range, then the other can be chosen so also.

The removal of any edge e of a tree T splits the tree into two disconnected components: $B(e)$, the nodes "below" e , and $A(e)$, the nodes "above" e . $B(e)$ and $A(e)$ are defined by:

$$B(e) = \{v \in V \mid v \text{ is in the subtree rooted by } \text{low}(e)\}$$

$$A(e) = V - B(e).$$

The flow along e in any flow for r, s depends only on the total number of requests and the total number of resources in $B(e)$. For any function $g: V \rightarrow R^+$ and $e \in E$, let

$$\text{total}(e, g) = \sum_{v \in B(e)} g(v).$$

If $\text{total}(e, r)$ exceeds $\text{total}(e, s)$, then there are more requests than resources in the subtree $B(e)$, so the excess must be satisfied by resources flowing into $B(e)$ via the edge e . On the other hand, if $\text{total}(e, s)$ exceeds $\text{total}(e, r)$, the excess resources in $B(e)$ are needed by requests in $A(e)$ (since $\text{total}(r) = \text{total}(s)$), so they must flow out of the subtree via e . By our convention that the sign of the flow denotes its direction, the directed flow in either case is given by

$$\text{dflow}_{r,s}(e) = \text{total}(e, r) - \text{total}(e, s).$$

Theorem 2.1.2. $\text{dflow}_{r,s}$ is the unique stable flow for r, s .

Proof. The argument sketched above shows that there is at most one stable flow for r, s . We leave to the reader to show that $\text{dflow}_{r,s}$ is stable for r, s . \square

Theorems 2.1.1 and 2.1.2 immediately yield:

$$\text{Theorem 2.1.3. } \text{cost}(r, s) = \text{cost}(\text{dflow}_{r,s}).$$

Thus, our original problem of studying optimal exact matchings reduces to the problem of studying a particular stable flow.

2.2. Expected Costs

In this section, we introduce probability

distributions for sets of requests and define costs of placements in terms of their expected costs given a randomly chosen set of requests.

If ϕ is a probability function on V and $t \in N$, then a random $r: V \rightarrow R^+$ can be chosen according to a probability distribution determined as follows: for each i in turn, $1 \leq i \leq t$, ϕ is used to select a vertex. Then $r(v)$ is the total number of times v is selected for each $v \in V$. We always assume $\phi(\text{root}) = 0$.

Fix ϕ, t as above, and let $s: V \rightarrow R^+$ with $\text{total}(s) = t$. Then $\text{expcost}(s)$ denotes the expected value of $\text{cost}(r, s)$, where r is chosen as described above, and $\text{minexpcost} = \min \{\text{expcost}(s) \mid s: V \rightarrow R^+ \text{ and } \text{total}(s) = t\}$.

The flow along an edge e depends only on $\text{total}(e, s)$, the number resources in the subtree below e , and not on their particular placement. Let $\text{expflow}_e(u)$ be the expected value of $|\text{dflow}_{r,s}(e)|$, where r is chosen randomly as described above, and s is any placement with $\text{total}(e, s) = u$.

The following two expansions are easy to see.

$$\text{Theorem 2.2.1. } \text{Expcost}(s) = \sum_{e \in E} \text{expflow}_e(\text{total}(e, s)).$$

$$\text{Theorem 2.2.2. } \text{Expflow}_e(u) = \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |u-i|, \text{ where } p = \sum_{v \in B(e)} \phi(v).$$

The next theorem shows that the expflow function has a simple form.

Theorem 2.2.3. For any fixed $\phi, t \geq 1$, and $e \in E$, expflow_e is an unbounded, convex, piecewise linear function from R^+ to R^+ , with all singularities occurring at integer values.

Proof. By Theorem 2.2.2, $\text{expflow}_e(u)$ is the sum of functions $g_i(u) = k_i \cdot |u-i|$, where k_i does not depend on $u, 0 \leq i \leq t$. Each g_i is a convex, piecewise linear function from R^+ to R^+ with a singularity at $u = i$, and at least one of the g_i 's is unbounded. Since addition preserves all four required properties, the result follows. \square

If $e \in E$, let E_e denote the set consisting of e , together with all edges below e in T , so $d \in E_e$ if $\text{low}(d) \in B(e)$. Define

$$\text{expcost}_e(s) = \sum_{d \in E_e} \text{expflow}_d(\text{total}(d, s))$$

and

$$\text{minexpcost}_e(u) = \min \{\text{expcost}_e(s) \mid \text{total}(e, s) = u\}.$$

Thus, $\text{expcost}_e(s)$ gives the portion of the expected cost of placement s due to flow in the subtree below (and including) e , and $\text{minexpcost}_e(u)$ gives the least such cost, subject to the constraint that exactly u resources be placed in the subtree.

From these definitions and Theorem 2.2.1, it should be clear that

$$\text{expcost}_e(s) = \text{expflow}_e(\text{total}(e,s)) + \sum_{i=1}^{\ell} \text{expcost}_{e_i}(s)$$

where e_1, \dots, e_{ℓ} are the immediate descendants of edge e . Optimizing over all s with $\text{total}(e,s) = u$, we get

Theorem 2.2.4. Let ϕ , t , and e be fixed. Let e_1, \dots, e_{ℓ} be the immediate descendant edges of e , and let $u \in R^+$. Then

$$\text{minexpcost}_e(u) = \text{expflow}_e(u) + \min \left\{ \sum_{i=1}^{\ell} \text{minexpcost}_{e_i}(u_i) \mid \sum u_i \leq u \right\}.$$

In order to obtain more information about the values of the minexpcost_e function, we first show that it has a simple form.

Theorem 2.2.5. For any fixed ϕ , $t \geq 1$, and e , minexpcost_e is an unbounded, convex, piecewise linear function from R^+ to R^+ , with all singularities occurring at integer values.

Proof. We use induction on edges in the tree, working from the leaves toward the root.

If e is a lowest edge, then there is only one edge, e , in E_e . Thus, $\text{minexpcost}_e = \text{expflow}_e$, which has the needed properties by Theorem 2.2.3.

Now assume the result holds for edges below e , and let e_1, \dots, e_{ℓ} denote the immediate descendant edges of e . Consider the expression for $\text{minexpcost}_e(u)$ given in Theorem 2.2.4. The first term has the needed properties by Theorem 2.2.3. It remains to show that the second term is convex, piecewise linear and has all singularities at integers.

Write f_i for minexpcost_{e_i} , $g(u)$ for $\min \{ \sum f_i(u_i) \mid \sum u_i \leq u \}$.

By the induction hypothesis, each f_i is unbounded, convex and piecewise linear with all singularities at integers. Hence, the derivative $f'_i(x)$ is defined at all non-singular points. We

define $f'_i(n)$ for n a singularity of f_i to be the limit of $f'_i(x)$ as x approaches n from above. By convexity and linearity of f_i , f'_i is a non-decreasing step function over the domain R^+ with steps occurring at integer points.

For each i , $1 \leq i \leq \ell$, let r_i be the smallest element of R^+ with $f'_i(r_i) \geq 0$. r_i exists since f_i is unbounded, and $r_i \in N$ by the properties of f'_i . We consider two cases:

Case 1. $u \geq \sum r_i$. Then $g(u) = \sum f_i(r_i)$, so g is constant for all such u .

Case 2. $u < \sum r_i$. Let $S = \{f'_i(x) \mid 1 \leq i \leq \ell, x \in [0, r_i]\}$. For $\sigma \in S$ and $1 \leq i \leq \ell$, define

$$x_i^{\sigma} = \text{glb} \{z \mid f'_i(z) \geq \sigma\}$$

$$y_i^{\sigma} = \text{lub} \{z \mid f'_i(z) \leq \sigma\}$$

By the properties of f'_i , we have $x_i^{\sigma}, y_i^{\sigma} \in N$, $x_i^{\sigma} \leq y_i^{\sigma} \leq r_i$, and $f'_i(z) = \sigma$ iff $z \in [x_i^{\sigma}, y_i^{\sigma})$. Hence, the intervals $[x_i^{\sigma}, y_i^{\sigma})$ for $\sigma \in S$ partition $[0, r_i)$, so the intervals $I_{\sigma} = [\sum_i x_i^{\sigma}, \sum_i y_i^{\sigma})$ for $\sigma \in S$ partition $[0, \sum r_i)$.

Choose $\sigma \in S$. We now show that g is linear over I_{σ} . Let $u \in I_{\sigma}$. Choose $y = (u_1, \dots, u_{\ell})$ such that $g(u) = \sum f_i(u_i)$ and $\sum u_i \leq u$. By the choice of u and using the facts that $f'_i(z) < \sigma$ iff $z < x_i^{\sigma}$ and $f'_i(z) > \sigma$ iff $z \geq y_i^{\sigma}$, it is straightforward to show that $x_i^{\sigma} \leq u_i \leq y_i^{\sigma}$ and $\sum u_i = u$. Hence, $f_i(u_i) = f_i(x_i^{\sigma}) + \sigma \cdot (u_i - x_i^{\sigma})$. Summing over all i , we get

$$\begin{aligned} g(u) &= \sum_{i=1}^{\ell} f_i(u_i) \\ &= \sum f_i(x_i^{\sigma}) + \sigma \cdot (\sum u_i - \sum x_i^{\sigma}) \\ &= \sum f_i(x_i^{\sigma}) + \sigma \cdot (u - \sum x_i^{\sigma}), \end{aligned}$$

a linear function of u as desired.

The convexity of g follows from the monotonicity of the f'_i . \square

Examination of the proof of Theorem 2.2.5 allows us to sharpen Theorem 2.2.4 by stating that the minimum cost can always be achieved by placing whole resources on all vertices.

Theorem 2.2.6. Let ϕ , t , e be fixed. Let e_1, \dots, e_k be the immediate descendant edges of e , and let $u \in \mathbb{N}$. Then $\text{minexpcost}_e(u) = \text{expflow}_e(u) + \min \left\{ \sum_{i=1}^k \text{minexpcost}_{e_i}(u_i) \mid \sum u_i \leq u \text{ and } u_i \in \mathbb{N} \right\}$.

2.3. Optimal Placement

We are interested in determining the "best" placement functions $s: V \rightarrow \mathbb{R}^+$ in the following sense. We say $s: V \rightarrow \mathbb{R}^+$ is optimal for ϕ , t provided $\text{total}(s) = t$ and $\text{expcost}(s) = \text{minexpcost}$. The first characterization result follows immediately from Theorem 2.2.6 and shows that there are optimal s which take on integral values only.

Theorem 2.3.1. For any probability function ϕ and any $t \in \mathbb{N}$, there exists $s: V \rightarrow \mathbb{N}$ which is optimal for ϕ , t .

Proof. Theorem 2.2.6 essentially provides an algorithm for producing such s . For any edge e of T with immediate descendants e_1, \dots, e_k , and any $u \in \mathbb{N}$, one determines values of s for all nodes below e by considering all possible decompositions u_1, \dots, u_k with $\sum u_i \leq u$ and $u_i \in \mathbb{N}$ for all i . For each such decomposition, one recursively determines values of s for all nodes below each e_i , and corresponding costs. The decomposition with the smallest total cost is chosen. \square

In order to analyze the cost of determining an optimal placement as above, we do not perform a straightforward recursive analysis of the algorithm described in the proof of Theorem 2.3.1. Rather, we take advantage of repeated work in various recursive calls. During the algorithms, one must calculate $\text{expflow}_e(u)$ for all $e \in E$ and all u , $0 \leq u \leq t$. The number of arithmetic operations involved in one calculation of $\text{expflow}_e(u)$ is $O(t)$ (if performed judiciously), independent of e and u . It is these costs which dominate the total count of arithmetic operations, so that an $O(t^2 \cdot |E|)$ analysis results. We summarize this discussion in the following theorem.

Theorem 2.3.2. There is an algorithm using $O(t^2 \cdot |E|)$ arithmetic operations which, for any tree $T = (E, V)$ probability function ϕ , and $t \in \mathbb{N}$, determines a placement of whole resources which is optimal for ϕ , t .

2.4. Optimizing Flow on Individual Edges

In this section, we show that the flow on each individual edge is optimized for a number equal to the median of an appropriate binomial distribution. We use this to bound the distance from optimal of two simple placements.

Let $n \in \mathbb{N} - \{0\}$, $0 \leq p \leq 1$. Let x be

a random variable whose value is the number of successes in n independent trials, each of whose probability of success is p . Define median(n, p) as the smallest $c \in \mathbb{R}^+$ such that $\Pr[x \leq c] \geq \frac{1}{2}$.

Theorem 2.4.1. Let ϕ be a probability function on V , $t \in \mathbb{N} - \{0\}$, $e \in E$. Then $\text{expflow}_e(u)$ is minimized at $u = \text{median}(t, p)$, where $p = \sum_{v \in B(e)} \phi(v)$.

Proof. Write f for expflow_e . By Theorem 2.2.3, f is minimized at an integer u which is the smallest $r \in \mathbb{N}$ with $f(r+1) - f(r)$ nonnegative. Now

$$\begin{aligned} f(r+1) - f(r) &= \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |r+1-i| \\ &\quad - \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |r-i| \end{aligned}$$

by Theorem 2.2.2,

$$\begin{aligned} &= \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} (|r+1-i| - |r-i|) \\ &= \sum_{i=0}^r \binom{t}{i} p^i (1-p)^{t-i} - \sum_{i=r+1}^t \binom{t}{i} p^i (1-p)^{t-i} \\ &= 2 \sum_{i=0}^r \binom{t}{i} p^i (1-p)^{t-i} - 1 \\ &= 2\Pr[x \leq r] - 1. \end{aligned}$$

Thus, u is the smallest $r \in \mathbb{N}$ with $2\Pr[x \leq r] - 1$ nonnegative, or $\Pr[x \leq r] \geq \frac{1}{2}$; that is, $u = \text{median}(t, p)$. \square

The following theorem follows immediately from Theorem 3.2 and Corollary 3.1 of Jogdeo and Samuels [1]. It is also implicit in some earlier work by Uhlmann [2,3].

Theorem 2.4.2 (Jogdeo and Samuels). Let $n \in \mathbb{N}$, $n \geq 1$, $0 \leq p \leq 1$. Then $\text{median}(n, p) \in \{\lfloor np \rfloor, \lceil np \rceil\}$.

Since np is the mean number of successes in n independent trials with success probability p , this theorem implies that the mean and median differ by less than one.

Thus, each edge individually has its flow optimized at a value which is either the mean rounded up or the mean rounded down. However, it is not always possible to achieve this local optimum consistently throughout the tree. In fact, in this very general setting, an optimal placement might require some subtree to contain a value other than the mean rounded up or down.

For any T , ϕ , t , and for each edge $e \in E$, let $p_e = \sum_{v \in B(e)} \phi(v)$. $s: V \rightarrow \mathbb{R}^+$ is an exact fair-share placement for T , ϕ , t if for each $e \in E$, $\text{total}(e,s) = tp_e$. $s': V \rightarrow \mathbb{N}$ is a fair whole placement for T , ϕ , t if for each $e \in E$, $\lfloor tp_e \rfloor \leq \text{total}(e,s') \leq \lceil tp_e \rceil$.

It is not difficult to see that for any T , ϕ , t , an exact fair-share placement and a fair whole placement exist. Let $s(v) = t \cdot \phi(v)$ for $v \in V$. Then s is an exact fair-share placement. Let $s'(v) = \lfloor t \cdot \phi(v) \rfloor$ if v is a leaf. Let $s'(v) = 0$ if v is the root. For v not a leaf or root, let e be the edge with $\text{low}(e) = v$, and let e_1, \dots, e_k be its immediate descendants. Let $s'(v) = \lfloor tp_e \rfloor - \sum_{i=1}^k \lfloor tp_{e_i} \rfloor$. An easy induction shows that s' is a fair whole placement for T , ϕ , t , and in fact, $\text{total}(e,s') = \lfloor tp_e \rfloor$ for every $e \in E$.

In the special case that ϕ is non-zero only on leaves, then there is a fair whole placement s'' which is non-zero only on leaves. It can be obtained by a simple recursive construction. Start at the root of T with t resources to distribute. Below any edge e in the tree, we will have either $\lfloor tp_e \rfloor$ or $\lceil tp_e \rceil$ to distribute. Assume e has descendants e_1, \dots, e_k . Distribute $\lfloor tp_{e_i} \rfloor$ or $\lceil tp_{e_i} \rceil$ to the i^{th} subtree. Since

$$\sum_{i=1}^k \lfloor tp_{e_i} \rfloor \leq \lfloor tp_e \rfloor \leq \lceil tp_e \rceil \leq \sum_{i=1}^k \lceil tp_{e_i} \rceil$$

the distribution is possible. We then proceed recursively on e_1, \dots, e_k .

Example 2.4.1. Let T be a 32-leaf complete binary tree (except for root edge (T)), $\phi(\ell) = \frac{15}{256}$ for each of the leftmost 16 leaves ℓ , $\frac{1}{256}$ for each of the rightmost 16 leaves, and 0 for all other nodes. Let $t = 16$. The placement s which has $s(\ell) = 1$ for each of the leftmost 16 leaves ℓ , and 0 elsewhere, has $\text{total}(e,s) = 16$ for e the left descendant of $\text{rootedge}(T)$. However, the mean number of requests in the subtree below e is $t \cdot \frac{15}{16} = 15$, so s is not a fair whole placement. Nevertheless, one can determine by exhaustive searching that s is optimal, and $\text{expcost}(s)$ is strictly smaller than $\text{expcost}(s')$ for any fair whole placement s' .

We note in general, however, that the optimal cost cannot be too much less than the cost of any exact fair-share placement or fair whole placement.


Theorem 2.4.3. Let s be an exact fair-share placement or a fair whole placement for T , ϕ , t . Then $\text{expcost}(s) \leq \text{minexpcost} + |E|$.

Proof. By the results of this section, expcost_e is minimized at some u , where $\lfloor tp_e \rfloor \leq u \leq \lceil tp_e \rceil$. Thus, $|\text{total}(e,s) - u| \leq 1$. But then $|\text{expcost}_e(\text{total}(e,s)) - \text{expcost}_e(u)| \leq 1$, by calculations similar to those in the proof of Theorem 2.4.1. Since no placement can do better than the optimal on each edge, s incurs at most an extra cost of 1 per edge. \square


For some choices of T , t and ϕ , of course, it is possible to consistently achieve the median on each edge -- for example, if the mean is an integer for every edge. But in general, we have no global optimality results. In §4 we obtain optimal results for a special case.

Example 2.4.2. If T is a complete binary tree with leaf vertices L , $|L| = 2^k$, $t = a \cdot 2^k$ for integer a , and ϕ the uniform distribution on the leaves, then the placement s such that $s(v) = a$ for each $v \in L$ and $s(v) = 0$ for $v \in V - L$ is optimal, because it achieves the integer mean on each edge.

Example 2.4.3. Let T be the tree depicted below, ϕ as indicated on the leaves and 0 elsewhere, and $t = 12$.



Then the means are indicated on each edge below, so all resources can be placed at the levels indicated by the circled numbers.



2.5. Nodes with Zero Placements

We conclude this section with a somewhat surprising characterization theorem. It says that if there is a probability less than $\frac{1}{2}$ of any request arriving in a subtree, then it is bad to place even a small fraction of a resource anywhere in that subtree.

Theorem 2.5.1. Let $e \in E$ satisfy $(1-p)^t > \frac{1}{2}$, where $p = \sum_{v \in B(e)} \phi(v)$. Let s be optimal for ϕ , t . Then $s(v) = 0$ for all $v \in B(e)$.

Proof. Assume not, and fix e, s exhibiting the contrary. Choose e' below e to be a lowest edge for which $x = s(\text{low}(e')) > 0$. Define a new placement s' , where $s'(\text{low}(e')) = 0$, $s'(\text{high}(e')) = s(\text{high}(e')) + x$, and $s'(v) = s(v)$ otherwise. We show that $\text{expcost}(s') < \text{expcost}(s)$, contradicting the optimality of s .

By Theorem 2.2.1, e' is the only edge for which s and s' have different expected absolute flows, so

$$\begin{aligned} & \text{expcost}(s) - \text{expcost}(s') \\ &= \text{expflow}_e(\text{total}(e', s)) - \text{expflow}_e(\text{total}(e', s')) \\ &= \text{expflow}_e(x) - \text{expflow}_e(0). \end{aligned}$$

Let $r = \sum_{v \in B(e')} \phi(v)$. Applying Theorem 2.2.2, the difference is

$$\begin{aligned} & \sum_{i=0}^t \binom{t}{i} r^i (1-r)^{t-i} (|x-i| - i) \\ & \geq x(1-r)^t + \sum_{i=1}^t \binom{t}{i} r^i (1-r)^{t-i} (-x) \\ & = x(2(1-r)^t - 1). \end{aligned}$$

Since $(1-r)^t \geq (1-p)^t > \frac{1}{2}$, the difference > 0 , giving the needed contradiction. \square

3. BOUNDS ON THE COSTS OF OPTIMAL PLACEMENTS

In this section, we assume that T is a complete (balanced) binary tree (except for rootedge(T)) with leaves L , and let $n = |L|$. Assume ϕ is arbitrary, except that as always $\phi(\text{root}) = 0$. We show that optimal placements are much better than centralized placements; in fact, their expected cost is linear in the number of leaves of the tree.

3.1. Cost of Exact Fair-Share Placements

Since we do not have a direct characterization of optimal placements, we instead bound the expected cost of an arbitrary exact fair-share placement. This upper bound, of course, provides an upper bound for optimal placements. Moreover, by Theorem 2.4.3, the costs cannot differ by more than $2n$.

Theorem 3.1.1. Let T, ϕ, L, n be as above. Let $t \in \mathbb{N}$, $t \geq 1$, and let $s: V \rightarrow \mathbb{R}^+$ be an exact fair-share placement with $\text{total}(s) = t$. Let $e \in E$ and $p_e = \sum_{v \in B(e)} \phi(v)$, $m = |L \cap B(e)|$. Then

$$\text{expcost}_e(s) \leq c\sqrt{m}tp_e$$

Here c is a fixed constant independent of the choice of T, ϕ, t, s , or e .

Before proving the theorem, we need some technical lemmas.

Lemma 3.1.2. For $t \in \mathbb{N}$, $t \geq 1$, $s \leq t-1$, $0 \leq p \leq 1$, it is the case that

$$\sum_{i=0}^s \binom{t}{i} p^i (1-p)^{t-i} (tp-i) = tp \binom{t-1}{s} p^s (1-p)^{t-s}.$$

Proof. $\binom{t}{i} p^i (1-p)^{t-i} (tp-i)$
 $= tp \binom{t}{i} p^i (1-p)^{t-i} - t \binom{t-1}{i-1} p^{i-1} (1-p)^{t-i}.$
 Since $\binom{t}{i} = \binom{t-1}{i} + \binom{t-1}{i-1}$, this expression is in turn equal to

$$\begin{aligned} & tp \binom{t-1}{i} p^i (1-p)^{t-i} + tp \binom{t-1}{i-1} p^{i-1} (1-p)^{t-i} \\ & \quad - t \binom{t-1}{i-1} p^{i-1} (1-p)^{t-i} \\ & = tp \binom{t-1}{i} p^i (1-p)^{t-i} - tp \binom{t-1}{i-1} p^{i-1} (1-p)^{t-(i-1)}. \end{aligned}$$

Thus $\sum_{i=0}^s \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$
 $= \binom{t}{0} (1-p)^t tp + tp \left[\sum_{i=1}^s \binom{t-1}{i} p^i (1-p)^{t-i} - \sum_{i=1}^s \binom{t-1}{i-1} p^{i-1} (1-p)^{t-(i-1)} \right]$
 $= tp(1-p)^t + tp \left[\sum_{i=1}^s \binom{t-1}{i} p^i (1-p)^{t-i} - \sum_{i=0}^{s-1} \binom{t-1}{i} p^i (1-p)^{t-i} \right]$
 $= tp(1-p)^t + tp \left[\binom{t-1}{s} p^s (1-p)^{t-s} - \binom{t-1}{0} p^0 (1-p)^t \right]$
 $= tp \binom{t-1}{s} p^s (1-p)^{t-s}. \quad \square$

Lemma 3.1.3. If $t \in \mathbb{N}$, $t \geq 1$, $0 \leq p < 1$, then

$$\sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |tp-i| = 2tp \binom{t-1}{\lfloor tp \rfloor} p^{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor}.$$

Proof.
$$\sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |tp-i|$$

$$= \sum_{i=0}^{\lfloor tp \rfloor} \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$

$$- \sum_{i=\lfloor tp \rfloor+1}^t \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$

$$= 2 \sum_{i=0}^{\lfloor tp \rfloor} \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$

$$- \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$

$$= 2tp \binom{t-1}{\lfloor tp \rfloor} p^{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor}$$

$$- \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$
by Lemma 3.1.2. But
$$\sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} (tp-i)$$

$$= \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} tp - \sum_{i=1}^t \binom{t}{i} p^i (1-p)^{t-i} i$$

$$= tp \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i}$$

$$- tp \sum_{i=1}^t \binom{t-1}{i-1} p^{i-1} (1-p)^{(t-1)-(i-1)} = tp - tp = 0.$$
□

Lemma 3.1.4. For $t \in \mathbb{N}$, $t \geq 1$, $0 \leq p < 1$, $tp \geq 1$, it is the case that

$$\binom{t-1}{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor} p^{\lfloor tp \rfloor} \leq (1 + \frac{1}{4t}) \cdot \frac{e}{\sqrt{2\pi \lfloor tp \rfloor}}$$

Proof. A version of Stirling's formula says that for all $n \in \mathbb{N}$, $n \geq 1$, it is the case that

$$\left(\frac{n}{e}\right)^n \sqrt{2\pi n} \leq n! \leq \left(\frac{n}{e}\right)^n \sqrt{2\pi n} \left(1 + \frac{1}{4n}\right).$$

$tp < t$ so $\lfloor tp \rfloor \leq t-1$ and $t - \lfloor tp \rfloor \geq 1$.

Then

$$\binom{t-1}{\lfloor tp \rfloor} = \frac{t - \lfloor tp \rfloor}{t} \cdot \binom{t}{\lfloor tp \rfloor}$$

$$\leq \left(\frac{t - \lfloor tp \rfloor}{\sqrt{t}}\right) \cdot \left(1 + \frac{1}{4t}\right)$$

$$\cdot \left(\frac{t^t}{\sqrt{2\pi \lfloor tp \rfloor} (t - \lfloor tp \rfloor) \cdot \lfloor tp \rfloor^{\lfloor tp \rfloor} (t - \lfloor tp \rfloor)^{t - \lfloor tp \rfloor}}\right).$$

Therefore,

$$\binom{t-1}{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor} p^{\lfloor tp \rfloor}$$

$$\leq \left(\frac{t - \lfloor tp \rfloor}{\sqrt{t}}\right) \cdot \left(1 + \frac{1}{4t}\right) \cdot \left(\frac{t-tp}{t - \lfloor tp \rfloor}\right)^{t-\lfloor tp \rfloor}$$

$$\cdot \left(\frac{tp}{\lfloor tp \rfloor}\right)^{\lfloor tp \rfloor} \cdot \left(\frac{1}{\sqrt{2\pi \lfloor tp \rfloor} (t - \lfloor tp \rfloor)}\right)$$

$$= \sqrt{\frac{t - \lfloor tp \rfloor}{2\pi t \lfloor tp \rfloor}} \cdot \left(1 + \frac{1}{4t}\right) \cdot \left(\frac{t-tp}{t - \lfloor tp \rfloor}\right)^{t-\lfloor tp \rfloor} \cdot \left(\frac{tp}{\lfloor tp \rfloor}\right)^{\lfloor tp \rfloor}.$$

But

$$\left(\frac{tp}{\lfloor tp \rfloor}\right)^{\lfloor tp \rfloor} = \left(1 + \frac{tp - \lfloor tp \rfloor}{\lfloor tp \rfloor}\right)^{\lfloor tp \rfloor} \leq \left(1 + \frac{1}{\lfloor tp \rfloor}\right)^{\lfloor tp \rfloor}$$

$$\leq e, \quad \frac{t - \lfloor tp \rfloor}{t} \leq 1 \quad \text{and} \quad \left(\frac{t-tp}{t - \lfloor tp \rfloor}\right)^{t-\lfloor tp \rfloor} \leq 1.$$

The lemma follows. □

Lemma 3.1.5. Let T , ϕ , t be as in Theorem

3.1.1. Let $e \in E$ and let $p = \sum_{v \in B(e)} \phi(v)$. Then $\text{expflow}_e(tp) < 6 \cdot \sqrt{tp}$.

Proof. If $p = 0$ or $p = 1$, then $\text{expflow}_e(tp) = 0$. Assume $0 < p < 1$. By Theorem 2.2.2 and Lemma 3.1.3,

$$\text{expflow}_e(tp) = \sum_{i=0}^t \binom{t}{i} p^i (1-p)^{t-i} |tp-i|$$

$$= 2tp \binom{t-1}{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor} p^{\lfloor tp \rfloor}$$

If $tp \geq 1$, Lemma 3.1.4 shows this is at most

$$2tp \cdot \left(1 + \frac{1}{4t}\right) \frac{e}{\sqrt{2\pi \lfloor tp \rfloor}}$$

$$\leq 2 \cdot tp \cdot \left(\frac{5}{4}\right) \cdot \frac{e}{\sqrt{2\pi}} \cdot \frac{\sqrt{\lfloor tp \rfloor}}{\lfloor tp \rfloor}$$

$$\leq 2 \cdot 2 \cdot \left(\frac{5}{4}\right) \cdot \frac{e}{\sqrt{2\pi}} \cdot \sqrt{tp}$$

$$< 6 \cdot \sqrt{tp}.$$

On the other hand, if $tp < 1$, then

$$2tp \cdot \binom{t-1}{\lfloor tp \rfloor} (1-p)^{t-\lfloor tp \rfloor} p^{\lfloor tp \rfloor} = 2tp \cdot (1-p)^t < 2\sqrt{tp}.$$

In any case,

$$\text{expflow}_e(tp) < 6 \cdot \sqrt{tp}. \quad \square$$

Proof of Theorem 3.1.1. Define a function

$G: \mathbb{N} \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ recursively:

$$G(0, u) = 6 \cdot \sqrt{u};$$

$$G(k, u) = \max \{G(k-1, u_1) + G(k-1, u_2) + 6 \cdot \sqrt{u} \mid u_1 + u_2 \leq u\}, \quad k > 1.$$

We first show by induction on m that $\text{expcost}_e(s) \leq G(\log_2 m, \text{tp}_e)$. Since s is an exact fair-share placement, $\text{total}(e, s) = \text{tp}_e$.

$m=1$: Then $\text{low}(e)$ is a leaf, so

$$\begin{aligned} \text{expcost}_e(s) &= \text{expflow}_e(\text{total}(e, s)) \\ &< 6 \cdot \sqrt{\text{tp}_e} \text{ by Lemma 3.1.5} \\ &= G(0, \text{tp}_e). \end{aligned}$$

$m > 1$: Then $m = 2^k$ since T is a complete binary tree. Let e_1, e_2 be the two immediate descendant edges of e . Then

$$\begin{aligned} \text{expcost}_e(s) &= \text{expflow}_e(\text{total}(e, s)) \\ &+ \sum_{i=1}^2 \text{expcost}_{e_i}(s). \end{aligned}$$

By induction,

$$\text{expcost}_{e_i}(s) \leq G(k-1, \text{tp}_{e_i}), \quad i=1, 2.$$

Again using Lemma 3.1.5, we have

$$\begin{aligned} \text{expcost}_e(s) &< 6 \cdot \sqrt{\text{tp}_e} + \sum_{i=1}^2 G(k-1, \text{tp}_{e_i}) \\ &\leq G(k, \text{tp}_e) = G(\log_2 m, \text{tp}_e) \end{aligned}$$

since $\text{tp}_{e_1} + \text{tp}_{e_2} \leq \text{tp}_e$.

We complete the proof of the theorem by showing that

$$G(k, u) \leq c \cdot 2^{k/2} \cdot \sqrt{u}$$

for some constant c .

First observe that for each $k \in \mathbb{N}$, $G(k, u)$, regarded as a function of u , is concave and monotonically increasing. This is clearly true for $k=0$. For $k > 0$, we know inductively that $G(k-1, u)$ is concave and monotonically increasing. Hence, $G(k-1, u_1) + G(k-1, u_2) \leq 2 \cdot G(k-1, (u_1 + u_2)/2) \leq 2 \cdot G(k-1, u/2)$ whenever $u_1 + u_2 \leq u$, so

$$(*) \quad G(k, u) = 2 \cdot G(k-1, u/2) + 6 \cdot \sqrt{u},$$

and $G(k, u)$ is concave and increasing.

We proceed to solve the recurrence equation

(*) First, substitute $a \cdot 2^k$ for u in (*) to

get

$$G(k, a \cdot 2^k) = 2 \cdot G(k-1, a \cdot 2^{k-1}) + 6 \cdot 2^{k/2} \cdot \sqrt{a}$$

Now, divide both sides by 2^k and express in terms of G' , where $G'(k, a) = G(k, a \cdot 2^k)/2^k$, to get

$$G'(k, a) = G'(k-1, a) + \frac{6 \cdot \sqrt{a}}{2^{k/2}}.$$

Also,

$$G'(0, a) = G(0, a) = 6 \cdot \sqrt{a}.$$

Hence,

$$G'(k, a) = \sum_{i=0}^k \frac{6 \cdot \sqrt{a}}{2^{i/2}} \leq 6 \cdot \sqrt{a} \sum_{i=0}^{\infty} 2^{-i/2}.$$

The series is convergent, so let $c = 6 \cdot \sum_{i=0}^{\infty} 2^{-i/2}$.

Then $G'(k, a) \leq c\sqrt{a}$. Substituting back, we get

$$G(k, u) = 2^k \cdot G'(k, u/2^k) \leq c \cdot 2^{k/2} \cdot \sqrt{u}$$

as desired.

We conclude that

$$\text{expcost}_e(s) \leq G(\log_2 m, \text{tp}_e) \leq c \cdot \sqrt{m \text{tp}_e}. \quad \square$$

Corollary 3.1.6. Let T, ϕ, L, n, t be as in Theorem 3.1.1. Let $s: V \rightarrow \mathbb{R}^+$ be an exact fair-share placement for T, ϕ, t , and let $s': V \rightarrow \mathbb{N}$ be a fair whole placement for T, ϕ, t . Then

$$\text{expcost}(s) \leq c \cdot \sqrt{nt}$$

and

$$\text{expcost}(s') \leq c \cdot \sqrt{nt} + 2n - 1.$$

Proof. The first bound comes from a direct application of Theorem 3.1.1 to $\text{rootedge}(T)$. The second bound follows from the first using Theorem 2.4.3 and that fact that an n -leaf binary tree has at most $2n-1$ edges (including $\text{rootedge}(T)$). \square

3.2. Cost of Centralized Placement

s is a centralized placement if $s(v) = 0$ everywhere except at $\text{low}(\text{rootedge}(T))$, and $s(\text{low}(\text{rootedge}(T))) = t$. For such an s ,

$$\text{expcost}(s) = t \log_2(n).$$

Note that the ratio of expected cost for a centralized placement to a fair whole placement is at least

$$\frac{t \log_2(n)}{c \sqrt{t} \sqrt{n} + 2n}$$

When t is small relative to n , then the centralized placement is superior, for reasons similar to those discussed in Section 2.5.

However, for $t = \Omega(n)$, the fair whole placement is better by a factor of $\Omega(\log_2 n)$, and for $t \gg n$, the ratio approaches \sqrt{t} , that is, the centralized placement is worse by a square.

Note also that for $t = \Omega(n)$, the fair whole placement has linear expected cost, so the expected cost per request is a constant, whereas the cost per request for a centralized placement is $\log_2 n$.

4. OPTIMAL PLACEMENTS FOR SPECIAL DISTRIBUTION

In this section, we give several characterization results and algorithms for optimal placements for a further restriction of the case considered in Section 3 in which we assume ϕ is the same on all leaves and zero elsewhere. Specifically, we assume the following for this section:

- (a) $T = (V, E)$ is a complete binary tree with leaves $L \subseteq V$, $n = |L|$, except that, as before, the root has a single emanating edge, $\text{rootedge}(T)$.
- (b) $\phi(v) = 1/n$ for all $v \in L$, and $\phi(v) = 0$ for all $v \in V - L$.
- (c) $t \in \mathbb{N}$, $t \geq 1$.

We define the level of a vertex to be its distance from the root. By our conventions regarding $\text{rootedge}(T)$, the leaves are at level $\log_2 n + 1$.

For the special case being considered in this section, certain of the relevant definitions can be generalized to reflect the symmetry in the tree. For example, if e_1 and e_2 are edges with $\text{high}(e_1)$ and $\text{high}(e_2)$ at the same level k , then $\text{expflow}_{e_1} = \text{expflow}_{e_2}$. Therefore, we write expflow_k in place of either. Similarly, we write minexpcost_k for minexpcost_e , where k is the level of $\text{high}(e)$.

4.1. A Bound on Levels with Nonzero Placements

Theorem 2.5.1 can be used to bound the level at which nonzero placement can occur in an optimal placement.

Theorem 4.1.1. Let $v \in V$ be at level $k \geq \lceil \log_2 t \rceil + 2$. Let s be optimal for ϕ , t . Then $s(v) = 0$.

Proof. By Theorem 2.5.1, it suffices to show that $(1 - \frac{1}{2^{\lceil \log_2 t + 1 \rceil}})^t > \frac{1}{2}$ for $t \in \mathbb{N}$. It

also suffices to consider t a power of 2. In this case, the inequality is just $(1 - \frac{1}{2t})^t > \frac{1}{2}$, which follows by an easy induction on $\log_2 t$. \square

Let $T_t = (V_t, E_t)$ denote the tree consisting of the vertices of T at levels from 0 to $\lceil \log_2 t \rceil + 1$ inclusive and the edges of T between them. The leaves L_t of T_t are the vertices at level $\lceil \log_2 t \rceil + 1$, and $n_t = |L_t|$. Let ϕ_t be defined on the vertices of T_t by $\phi_t(v) = 1/n_t$ for all $v \in L_t$ and $\phi_t(v) = 0$ for $v \in V_t - L_t$. If $s: V \rightarrow \mathbb{R}^+$ is such that $s(v) = 0$ for all $v \in V$ at levels below $\lceil \log_2 t \rceil + 1$, then $s_t: V_t \rightarrow \mathbb{R}^+$ can be defined from s by simply ignoring the missing vertices. Then it is easy to see that $\text{expcost}^{(T)}(s) = \text{expcost}^{(T_t)}(s_t) + t \cdot (\log_2 n - \lceil \log_2 t \rceil)$. (Superscripts distinguish the trees under consideration.) We then have the following.


Theorem 4.1.2. If $t > 1$, then $\text{minexpcost}^{(T)} = \text{minexpcost}^{(T_t)} + t \cdot (\log_2 n - \lceil \log_2 t \rceil)$.

Proof. (\geq) follows from Theorem 4.1.1 and the remarks.

(\leq) follows because any placement for T_t can be augmented to a placement for T by placing zeros on the additional nodes, thereby incurring the stated cost. \square

Thus, in the remainder of this section, we assume that the maximum level in T is at most $\lceil \log_2 t \rceil + 1$ (that is, $n < 2t$). The reader can then use Theorem 4.1.2 to infer corresponding results about cases where $n \geq 2t$.

Example 4.1.1. The case where $t = 1$ is somewhat peculiar. The reader can verify that for all T with maximum level number at least 2, the following pictures all represent optimal placements.



That is, in the first case,

$$s(v) = \begin{cases} 1 & \text{for } v \text{ the son of the root,} \\ 0 & \text{otherwise,} \end{cases}$$

while in the other two cases,

$$s(v) = \begin{cases} 1 & \text{for } v \text{ the left (resp. right) grandson of the root,} \\ 0 & \text{otherwise.} \end{cases}$$

This is the only value of t for which an optimal placement can have nonzero values below level $\lceil \log t \rceil + 1$.

Example 4.1.2. If $t = 2^k$ and $n > t$, then the placement with one resource on each of the t vertices at level $k+1$ is optimal: nothing is placed below this level, and an optimal placement for the whole tree results from an optimal placement within levels up to $\log t + 1$. But clearly putting one resource across all leaves of T_t is optimal, as seen in Example 2.4.2.

4.2. Algorithms for Finding Optimal Placements

In this section, we prove two sharper versions of Theorem 2.2.4 for the special case of this section, and use them as bases for two algorithms for finding optimal placements.

Theorem 4.2.1. Let $t \in \mathbb{N}$, $u \in \mathbb{R}^+$ and $k \in \mathbb{N}$, $k \leq \log_2 n - 1$. Then $\text{minexpcost}_k(u) = \text{expflow}_k(u) + 2 \cdot \min \{ \text{minexpcost}_{k+1}(u') \mid u' \in \{0, 1, \dots, \lfloor \frac{u}{2} \rfloor, \frac{u}{2} \} \}$.

Proof. \leq is clear. We show \geq . Write f for minexpcost_{k+1} . Consider any $u_1, u_2 \in \mathbb{R}^+$ with $u_1 + u_2 \leq u$ and $f(u_1) + f(u_2)$ minimal. We will produce $u' \in \{0, 1, \dots, \lfloor \frac{u}{2} \rfloor, \frac{u}{2} \}$ with $2f(u') \leq f(u_1) + f(u_2)$.

By Theorem 2.2.5, there exists $r \in \mathbb{N}$ such that r is the smallest element of \mathbb{R}^+ with $f(r) \leq f(s)$ for all $s \geq r$. We consider two cases.

Case 1. $u \geq 2r$.

Choose $u' = r$. $u' \in \{0, 1, \dots, \lfloor \frac{u}{2} \rfloor\}$ and $2f(u') \leq f(u_1) + f(u_2)$.

Case 2. $u < 2r$.

Then $u_1 + u_2 = u$. Choose $u' = \frac{u}{2}$. Then $2f(u') = 2f(\frac{u_1 + u_2}{2}) \leq f(u_1) + f(u_2)$ by convexity of f .

An appeal to Theorem 2.2.4 completes the proof.

Theorem 4.2.2. For any $t \in \mathbb{N}$, there exists s such that $s(v) = s(v')$ for all pairs of vertices v, v' at the same level, which is optimal for t .

Proof. s can be found by a recursive algorithm based on Theorem 4.2.1. \square

We proceed as before to analyze the cost of finding the placement of Theorem 4.2.2. During

During the algorithm, one must calculate $\text{expflow}_0(t), \text{expflow}_1(\frac{t}{2}), \dots, \text{expflow}_{\log(n)}(\frac{t}{n})$. In addition, one must calculate $\text{expflow}_1(i)$ for all $i \in \mathbb{N}$, $i \leq \frac{t}{2}$, $\text{expflow}_2(\frac{i}{2})$ for all $i \in \mathbb{N}$, $i \leq \frac{t}{2}$, $\text{expflow}_3(\frac{i}{4})$ for all $i \in \mathbb{N}$, $i \leq \frac{t}{2}, \dots, \text{expflow}_{\log(n)}(\frac{2i}{n})$ for all $i \in \mathbb{N}$, $i \leq \frac{t}{n}$.

This is a total of $O(t \log n)$ expflow computations. Since each such computation involves $O(t)$ arithmetic operations, we have the following theorem.

Theorem 4.2.3. There is an algorithm using $O(t^2 \log n)$ arithmetic operations which, for any tree T with n leaves satisfying the assumptions of this section and for any $t \in \mathbb{N}$, determines an s which is optimal for t such that $s(v) = s(v')$ for all pairs of vertices v, v' at the same level.

Note that the bound of Theorem 4.2.3 represents an improvement over the bound in Theorem 2.3.2 applied to the special case of this section; the placements produced by the two algorithms have somewhat different properties, however. The remainder of this subsection deals with integral placements, the situation considered in Theorem 2.3.2.

Theorem 4.2.4. Let $t, u, k \in \mathbb{N}$, $k \leq \log_2 n - 1$.

Then $\text{minexpcost}_k(u) = \text{expflow}_k(u) + \min \{ \text{minexpcost}_{k+1}(u_1) + \text{minexpcost}_{k+1}(u_2) \mid u_1, u_2 \in \mathbb{N}, u_1 + u_2 \leq u, \text{ and } |u_2 - u_1| \leq 1 \}$.

Proof. Again, we show \geq . Write f for minexpcost_{k+1} . By Theorem 4.2.1, there is a value $u' \in \{0, 1, \dots, \lfloor \frac{u}{2} \rfloor, \frac{u}{2} \}$ such that $\text{minexpcost}_k(u) = \text{expflow}_k(u) + 2f(u')$. If $u' \in \{0, \dots, \lfloor \frac{u}{2} \rfloor\}$, then we simply take $u_1 = u_2 = u'$. If $u' = \frac{u}{2} \notin \mathbb{N}$, then we take $u_1 = \lfloor \frac{u}{2} \rfloor$ and $u_2 = \lceil \frac{u}{2} \rceil$; in this case, piecewise linearity of f guarantees the required properties. \square

Theorem 4.2.5. For any $t \in \mathbb{N}$, there exists $s: V \rightarrow \mathbb{N}$ which is optimal for t . Moreover, s has the additional properties:

- if e_1 and e_2 are two edges with $\text{high}(e_1) = \text{high}(e_2)$, then $|\text{total}(e_1, s) - \text{total}(e_2, s)| \leq 1$,
- if e is any edge with $\text{high}(e)$ at level k , then $\text{total}(e, s) \leq \lceil \frac{t}{2^k} \rceil$.

Proof. Theorem 4.2.4 leads naturally to a recursive algorithm yielding a placement in $V \rightarrow N$ and obviously satisfying (a). To see that (b) is also satisfied, assume the contrary, and let e be a highest edge in the tree with

$$\text{total}(e,s) > \left\lceil \frac{t}{2^k} \right\rceil, \text{ where } \text{high}(e) \text{ is at level } k.$$

Since $\text{total}(e,s) \in N$, we have $\text{total}(e,s) \geq \left\lceil \frac{t}{2^k} \right\rceil + 1 \geq \frac{t}{2^k} + 1$. e is not rootedge(T), so let e' be

the edge $\neq e$ with $\text{high}(e) = \text{high}(e')$, and let e'' be the edge immediately above e in T . Then

$$\text{total}(e'',s) \leq \left\lceil \frac{t}{2^{k-1}} \right\rceil < \frac{t}{2^{k-1}} + 1.$$

Therefore,

$$\begin{aligned} \text{total}(e's) &\leq \text{total}(e'',s) - \text{total}(e,s) \\ &< \left(\frac{t}{2^{k-1}} + 1\right) - \left(\frac{t}{2^k} + 1\right) = \frac{t}{2^k}. \end{aligned}$$

But then

$$|\text{total}(e,s) - \text{total}(e',s)| > 1,$$

contradicting property (a). \square

Once again, we analyze the cost of determining an optimal placement with the properties of Theorem 4.2.5. One must calculate $\text{expflow}_0(t)$, and also $\text{expflow}_1(i)$ for all $i \in N$, $i \leq \lceil t/2 \rceil$, $\text{expflow}_2(i)$ for all $i \in N$, $i \leq \lceil \frac{\lceil t/2 \rceil}{2} \rceil = \lceil t/4 \rceil, \dots$, $\text{expflow}_{\log(n)}(i)$ for all $i \in N$, $i \leq \lceil t/n \rceil$. This is a total of $O(t)$ expflow computations, (since we are assuming that n is $O(t)$).

Theorem 4.2.6. There is an algorithm using $O(t^2)$ arithmetic operations, which for any tree T and for any $t \in N$, determines an $s: V \rightarrow N$ which is optimal for t , and which satisfies conditions (a) and (b) of Theorem 4.2.5.

4.3. A Fast Algorithm for Determining Optimal Placements

The results of Section 2.4 can be used to prune the algorithm's search space still further, leading to a much faster algorithm.

Theorem 4.3.1. Let $t, u, k \in N$,

$$k \leq \log(n) - 1, \quad \left\lfloor \frac{t}{2^k} \right\rfloor \leq u \leq \left\lceil \frac{t}{2^k} \right\rceil. \text{ Then}$$

$$\text{minexpcost}_k(u) = \text{expflow}_k(u)$$

$$+ \min\{\text{minexpcost}_{k+1}(u_1) + \text{minexpcost}_{k+1}(u_2)\}$$

$$u_1, u_2 \in \left\{ \left\lfloor \frac{t}{2^{k+1}} \right\rfloor, \left\lceil \frac{t}{2^{k+1}} \right\rceil \right\} \text{ and } u_1 + u_2 \leq u.$$

Proof. Again, we show \geq . Write f for minexpcost_{k+1} . By Theorem 4.2.4, there is a pair $u_1, u_2 \in N$ such that $u_1 + u_2 \leq u$, $|u_2 - u_1| \leq 1$ and $\text{minexpcost}_k(u) = \text{expflow}_k(u) + f(u_1) + f(u_2)$. Of all such pairs, choose one minimizing $u - (u_1 + u_2)$ and assume $u_1 \leq u_2$. We must check that $u_1, u_2 \in \left\{ \left\lfloor \frac{t}{2^{k+1}} \right\rfloor, \left\lceil \frac{t}{2^{k+1}} \right\rceil \right\}$.

$$\begin{aligned} \text{If } u_2 > \left\lceil \frac{t}{2^{k+1}} \right\rceil, \text{ then } u_2 &\geq \left\lceil \frac{t}{2^{k+1}} \right\rceil + 1, \text{ so} \\ u_1 &\geq \left\lfloor \frac{t}{2^{k+1}} \right\rfloor. \text{ Then } u \geq u_1 + u_2 \geq 2 \left\lceil \frac{t}{2^{k+1}} \right\rceil + 1 > \left\lceil \frac{t}{2^k} \right\rceil, \end{aligned}$$

a contradiction. Thus, u_2 (and therefore u_1) $\leq \left\lceil \frac{t}{2^{k+1}} \right\rceil$.

$$\begin{aligned} \text{If } u_1 < \left\lfloor \frac{t}{2^{k+1}} \right\rfloor, \text{ then } u_1 &\leq \left\lfloor \frac{t}{2^{k+1}} \right\rfloor - 1, \text{ so} \\ u_2 &\leq \left\lceil \frac{t}{2^{k+1}} \right\rceil. \text{ Then } u - (u_1 + u_2) \geq \left\lfloor \frac{t}{2^k} \right\rfloor \end{aligned}$$

$$- \left(\left\lfloor \frac{t}{2^{k+1}} \right\rfloor - 1 + \left\lceil \frac{t}{2^{k+1}} \right\rceil \right) \geq 1. \text{ Define a new decomposition}$$

w_1, w_2 by $w_1 = u_1 + 1, w_2 = u_2$. Then

$$\begin{aligned} w_1 + w_2 &\leq u \text{ and } |w_2 - w_1| \leq 1. \text{ Now, } \text{expflow}_{k+1}(w_1) \\ &\leq \text{expflow}_{k+1}(u_1) \text{ because } u_1 \leq \left\lfloor \frac{t}{2^{k+1}} \right\rfloor - 1, \text{ by} \end{aligned}$$

Theorems 2.2.3, 2.4.1 and 2.4.2. Thus, $f(w_1) \leq f(u_1)$, by Theorem 2.2.4. Hence, w_1, w_2 also satisfies the conditions used in choosing u_1, u_2 , but $u - (u_1 + u_2) > u - (w_1 + w_2)$, contradicting the minimality condition. \square

Theorem 4.3.2. For the special case of this section, there exists a fair whole placement s which is optimal for t .

Proof. Theorem 4.3.1 yields a recursive algorithm. \square

Note that a result similar to Theorem 4.3.2 does not hold in the general case -- recall Example 2.4.1.

The algorithm resulting from Theorem 4.3.2 is extremely fast. Namely, one must calculate


$$\text{expflow}_k(i) \text{ for } i = \left\lfloor \frac{t}{2^k} \right\rfloor, \left\lceil \frac{t}{2^k} \right\rceil, \text{ for each } k,$$

$1 \leq k \leq \log(n)$, for a total of only $O(\log n)$ expflow computations.

Theorem 4.3.3. There is an algorithm using $O(t \log n)$ arithmetic operations which for any tree T and for any $t \in \mathbb{N}$, determines a fair whole placement s which is optimal for t .

4.4. Example

Some of the optimal placements discovered by our algorithms are rather unexpected. For example, the following represents an optimal placement of 11 resources in a balanced binary tree with uniform probability distribution:



Acknowledgements

The authors thank Carl Spruill, Charles Blair and Mike Paterson for contributing their ideas and suggestions for some of the results in this paper.

References

1. K. Jogdeo and S.M. Samuels, "Monotone Convergence of Binomial Probabilities and a Generalization of Ramanujan's Equation," The Annals of Mathematical Statistics 39, 4 (1968), 1191-1195.
2. W. Uhlmann, "Ranggrößen als Schätzfunktionen," Metrika 7, (1963), 23-40.
3. W. Uhlmann, "Vergleich der hypergeometrischen mit der Binomial-Verteilung," Metrika 10, (1966), 145-158.