# Probabilistic Analysis of a Network Resource Allocation Algorithm* 

Nancy A. Lynch<br>Massachusetts Institute of Technology, Cambridge, Massachusetts

Nancy D. Griffeth

Georgia Institute of Technology, Atlanta, Georgia

Michael J. Fischer
Yale University, New Haven, Connecticut
AND
Leonidas J. Guibas
Stanford University, Stanford, California, and DEC/SRC
Palo Alto, California


#### Abstract

A distributed algorithm is presented, for allocating a large number of identical resources (such as airline tickets) to requests which can arrive anywhere in a distributed network. Resources, once allocated, are never returned. The algorithm searches sequentially, exhausting certain neighborhoods of the request origin before proceeding to search at greater distances. Choice of search direction is made nondeterministically. Analysis of expected response time is simplified by assuming that the search direction is chosen probabilistically, that messages require constant time, that the network is a tree with all leaves at the same distance from the root, and that requests and resources occur only at leaves. It is shown that the response time is approximated by the number of messages of one type that are sent during the execution of the algorithm, and that this number of messages is a nondecreasing function of the interarrival time for requests. Therefore, the worst case occurs when


[^0]requests come in so far apart that they are processed sequentially. The expected time for the sequential case of the algorithm is analyzed by standard techniques. This time is shown to be bounded by a constant, independent of the size of the network. It follows that the expected response time for the algorithm is bounded in the same way. © 1986 Academic Press, Inc.

## 1. Introduction

We consider the problem of allocating a number of identical resources to requests arriving at the sites of a distributed network. We assume that the network is configured as a tree. The nodes of the tree are processors and the edges are communication lines connecting the processors. Processes at a node may communicate only over the tree edges, with processes at other nodes. Resource allocation is managed by a collection of communicating resource allocation processes, one at each node. We will henceforth refer only to the node, identifying it with both the processor and the resource allocation process at the node.

From time to time, a request arrives at a node (potentially any node of the network) from the outside world. One of the resources should eventually be granted to the request, subject to the following conditions:

1. No resource is granted more than once. (Once granted, a resource is not returned. Thus, there is no legitimate reason to grant it more than once.)
2. At most one resource is granted to each request.
3. A node grants resources only to those requests which arrive at that node.
4. If the number of requests is no greater than the number of resources, then each request eventually receives a resource.
5. If the number of resources is no greater than the number of requests, then each resource is eventually granted to a request.

For convenience in describing allocation of specific resources to specific requests, we assume that each resource and request has a unique identifier.

The execution model for this distributed network is event-based. Two types of events may occur at a node: (1) a request may arrive from the outside world, and (2) a message may arrive from a neighbor in the tree. Each event triggers an indivisible step at the node. This step may include changing state, sending messages to other nodes, and granting resources to requests. (We ignore the time involved in this local processing when we measure the response time, considering only the communication time.) We assume that the communication lines are reliable, that is, each message is
delivered exactly once. However, we do not make any assumptions about he order of message arrivals.

There are many interesting approaches to solving this resource allocation problem. In a centralized approach, all resources are controlled by a single central node. When a request arrives at a possibly different node, a "buyer" is commissioned, who travels, via messages, to the central node to obtain a resource. The buyer then carries the resource back to the node where the request originated, so that the resource can be granted to the request at that location.

An alternative approach is to decentralize control of the resources, giving each node of the network control of some of them. In this approach, the buyers must search for the resources. An important choice to be made in designing an efficient search strategy is the choice between sending only one buyer to search for resources for each request and sending several buyers in parallel to search different parts of the tree. The former search strategy, which we call the sequential search strategy, avoids a number of problems arising from the parallel strategy, such as what to do about other buyers when one of them has found a resource. The next choice, if the sequential search strategy is used, is the choice of direction to search the tree. A good choice would involve guessing which nodes are most likely to have free resources when the buyer arrives at them.

Other strategies involve combining a decentralized search strategy with a dynamic resource redistribution strategy, letting resources search for requests (rather than vice versa), or giving nodes control of fractions of resources rather than whole resources.

One complexity measure which is useful for evaluating different strategies is the expected response time. This is a measure upon which any of the design choices could have a major impact. For example, the response time when using a centralized strategy must depend strongly on the network size. However, the decentralized strategies have the potential of depending on this size to a lesser extent.

In the first half of this paper, we present an algorithm for solving this resource allocation problem. Our algorithm is a decentralized solution in which each node controls some whole number of resources. A sequential search strategy is used, in which the direction to be searched is chosen nondeterministically. Certain neighborhoods of the node at which a request originates are exhausted before the search proceeds to more distant neighborhoods.

In order to gain some insight into the expected response time for our algorithm, we simulated its behavior, in some special cases. The nondeterministic choice of search direction was resolved by using a probabilistic choice, where the probabilities for the different directions depended on the initial placement of resources in those directions. We assumed an exponen-
tial distribution for time of arrival of requests, a uniform distribution for arrival location, and a normal probability distribution for message delivery time. We also assumed that all leaves of the network tree were at the same distance from the root, and that requests and resources occurred only at leaves. We first noted that expected response time was extremely good, with an upper bound that seemed to be independent of the size of the network. This was in marked contrast to a centralized algorithm. Next, we made a surprising observation: the expected response time appeared to be a nondecreasing function of the expected interarrival time for requests. If true, this observation would imply that the worst case for the algorithm was actually the case where requests come in so far apart that they are processed one at a time. This observation contradicted our preliminary intuitions about the algorithm: we had thought that the worst cases would arise when there was greatest competition among requests searching for resources.

Using these observations as hints, we were able to carry out a substantial amount of analysis of the algorithm's behavior, and this analysis comprises the second half of this paper. Namely, we prove an upper bound on the expected response time for a special case in which, among other restrictions, all leaves of the network tree are at the same distance from the root, and requests and resources occur only at leaves. First, we show that the response time can be bounded in terms of the number of messages of one type that are sent during the execution of the algorithm. Then we show that this number of messages is a nondecreasing function of the interarrival time for requests. Therefore, the worst case occurs when requests come in so far apart that they are processed sequentially. We analyze the expected time for the sequential case, showing it to be bounded by a constant, independent of the size of the network. It follows that the expected response time for the algorithm is also bounded by a constant.

Although the expected response time for our algorithm is very good, we do not claim that it is optimal. In fact, there are some simple changes that one would expect to yield improvements. Unfortunately, with these changes, the algorithm can no longer be analyzed using the same techniques; thus, we are not really certain that they are improvements at all.

There are several contributions in this paper. First, we think that the algorithm itself is interesting. Second, we have identified an interesting criterion for the performance of a distributed algorithm: that the performance be independent of the size of the network. Satisfying this criterion seems to require an appropriate, decentralized style of programming. Third, the analysis is decomposed in an interesting way: a sequential version is analyzed using traditional methods, and the performance of the concurrent algorithm is shown to be bounded in terms of the sequential algorithm. It is likely that this kind of decomposition will prove to be
useful for analysis of other distributed algorithms. For instance, a similar decomposition was used in the proof of correctness of a systolic stack (Guibas, and Liang, 1982).

The contents of the rest of the paper are as follows. Section 2 contains the algorithm, and Section 3 contains arguments for its correctness. Sections 4-6 contain the analysis of the algorithm. Section 4 proves the monotonicity result, which implies that the sequential case of the algorithm is worst. Section 5 analyzes the sequential case. Section 6 pulls together the results of Sections 4 and 5, thus giving a general upper bound. Finally, Section 7 describes some remaining questions.

## 2. The Algorithm

In this section, we present our algorithm. We begin with an informal decription, followed by a more formal presentation.

### 2.1. Informal Description

We assume that the network is a rooted tree. Our algorithm is a decentralized algorithm with a sequential searching strategy. Requests send buyers to search for resources. When a buyer finds a resource, it "captures" it. Each captured resource travels back to the origin of this buyer (or possibly some other buyer, if there is interference between the processing of concurrent requests), so that the grant can occur where the request originated.

When a request or buyer arrives at any node, any free resource at the node is captured. If there are no free resources there, a buyer is sent to a neighboring node, determined as follows. Each node keeps track of the latest estimate it knows, for the number of resources remaining in each of its subtrees. Each node sends a message informing its parent of each new request which has originated within the child's subtree. The estimate which a node keeps for the number of resources remaining in a subtree, is calculated from the initial placement of resources in that subtree, the number of requests which are known to have originated within that subtree, and the number of buyers which the node has already sent into that subtree. In order to decide on the direction in which to send a buyer, a node uses the following rules. First, it never sends a buyer out of its subtree if it estimates that its subtree still contains a resource. Second, it only sends a buyer downward to a child if it estimates that the child's subtree contains a resource. Third, if there is a choice of child to which to send the buyer, the node makes a nondeterministic choice. (Later, we will constrain this decision to a probabilistic choice using a particular random choice
function. This constraint will be important for the complexity analysis, but is not needed for the correctness of the algorithm.)

It is easy to see that any subtree which a node considers to contain no resources, actually contains no resources. Thus, no buyer is ever sent out of a subtree actually containing a resource. On the other hand, the perceived information about he availability of a resource in a child's subtree can be an overestimate, in case of interference among concurrent requests.

Example. Suppose that request $A$ enters at the node shown below, and its buyer travels upward until it reaches an ancestor that perceives the availability of a resource in one of its subtrees. Then the buyer travels downward toward that resource. Shortly before $A$ 's buyer reaches the resource, another request $B$ arrives at the node shown. Suppose $B$ 's buyer reaches the resource and captures it before $A$ 's buyer does. When (or before) $A$ 's buyer finally arrives at the resource's location, it will encounter the information that the resource is no longer there. Then $A$ 's buyer will be sent upward, backtracking in its search for a resource (see Fig. 1).

Although such interference can cause backtracking, the buyer will eventually find a resource if one exists. This is because no buyer ever leaves a subtree actually containing a resource. Several optimizations are incorporated into the algorithm, as follows:

1. Buyers, unlike requests, need not be uniquely identified. Instead, each node keeps track of the number of buyers received and sent and the net flow of buyers over each of its incident edges. Captured resources then travel in such a way as to negate net flow of buyers, and a buyer will eventually leave a subtree which does not contain a resource.
2. Buyers can travel "discontinuously." Assume node $v$ sends a buyer to a child node $w$, thinking that there is an available resource in w's subtree. Assume that, soon thereafter, $v$ receives a message from $w$, informing $v$ of an arrival of a new request in $w$ 's subtree, and implying that $v$ 's previous supposition of an available resource was false. Then $v$ knows that $w$ will


Figurf 1
eventually send some buyer back up to $v$, at which time $v$ should send the buyer in another direction. Since $v$ knows this will eventually occur, $v$ need not actually wait for the buyer to arrive from $w$; it can create a new buyer and send it in anticipation of the later return of the first buyer. Since the first buyer will not find any free resources in the subtree, this extra parallelism does no harm. In fact, with this optimization, it is no longer necessary for $w$ to return the buyer at all, since $v$ must ignore it when it returns to it in any case.
3. If each node knows how many resources were initially placed in each of its children's subtrees, then it is not even necessary for explicit buyers to be sent upward at all! All that is necessary is for nodes to send "ARRIVAL" messages upward to their parents, informing them of the arrival of new requests in their subtree. The parent is able to deduce the number of resources which the child would like to have sent down (i.e., the number of buyers emanating from the child's subtree), from the initial number of resources in the subtree, the number of arrivals in the subtree, and the number of buyers already sent down into the subtree. We will say more in a moment about how this deduction is made.

If information about newly-arrived requests (in the form of "ARRIVAL" messages) only flows upward in the tree, there is no way that a child can deduce that its parent would like it to send a resource upward. Thus, it is still necessary to send explicit buyers downward. Let us designate these explicit downward buyer messages as "BUYER" messages. Thus, the algorithm only uses two kinds of messages to search for resources: "ARRIVAL" messages flow upward to inform parents about new requests, and "BUYER" messages flow downward to inform a child that its parent would like the child to send up a resource.

The precise deduction which a parent can make about the number of buyers emanating from a child's subtree is as follows:

Let $a$ be the number of "ARRIVAL" messages which have been received by the parent from the child. Let $b$ be the number of "BUYER" messages which have been sent by the parent to the child. Let $p$ be the number of resources initially placed in the child's subtree. Then the number of buyers perceived as emanating from a child's subtree is max $(a+b-p, 0)$. This number is called the estimate of "virtual buyers" emanating from the subtree.

That is, if the total number of "ARRIVAL" and "BUYER" messages indicated above is no greater than the initial placement, no buyers are perceived as emanating from the subtree. On the other hand, if this total is greater, then the excess is perceived to be the number of buyers.

Analogously, the child node deduces an estimate of the number of "virtual buyers" it has sent out of its subtree, as follows. Let $a$ be the number
of "ARRIVAL" messages which the child has sent to its parent. Let $b$ be the number of "BUYER" messages which have been received by the child from its parent. Let $p$ be the number of resources initially placed in the child's subtree. Then the number of buyers the child perceives that it has sent out of its subtree (also called the estimate of "virtual buyers" sent out of the subtree $)=\max (a+b-p, 0)$. Because of message delays, the child and the parent may differ on their estimates of the number of virtual buyers.

In order to make the actual grants to specific requests, it seems necessary that each specific identified resource "travel" to a point of request origin, in order to get properly paired with a request. This travel requires a third kind of message to be sent around, namely, a specific "captured resource." The algorithm which sends resources around is particularly simple-resources are just sent in such a way as to negate the net flow of buyers. This part of the algorithm executes concurrently with, but has no effect on, the searching part.

### 2.2 Formal Description

In this subsection, we present a program implementing the algorithm described above. A sketch of a correctness proof is presented in the next section. Primarily, the proof consists of showing the correctness of the invariant assertions made at various points in the program. The reader may wish to examine the proof while reading the program.

We assume that the network is described by a rooted tree $T$. For uniformity, let the root of $T$ have an outgoing upward edge. (Messages sent along this edge will never be received by anyone.) We can then write a single program for all the nodes of $T$, including the root.

Let $V$ denote the set of vertices of $T$. Let RESOURCES $(v)$ denote the resources placed at vertex $v$, for each $v \in V$, and let $\operatorname{PLACE}(v)=$ $|\operatorname{RESOURCES}(v)|$ for all $v$. Let REQUESTS $(v)$ denote the requests arriving at $v$. We assume that all the sets RESOURCES $(v)$ and REQUESTS $(v)$ are finite. Let PARENT $(v)$, CHILDREN $(v)$, DESCENDANTS $(v)$, and NEIGHBORS $(v)$ denote the designated vertices and sets of vertices, for vertex $v$.

The kinds of messages used are "ARRIVAL," "BUYER," and messages corresponding to specific captured resources.

Program for Node $v, v \in V$. In the program for node $v$, we use RESOURCES as a shorthand for RESOURCES(v), and similarly for the other notation above.

It is convenient to think of the state of $v$ as consisting of "independent variables" and "dependent variables." The independent variables are just
the usual kind of variables, which can be read and assigned. The dependent variables are virtual variables whose values are defined in terms of the independent variables. These values can be read, but not modified. We can think of the reading of a dependent variable as shorthand for a read of several independent variables, together with a calculation of the function giving the dependency.

Independent Variables. REQUESTS, for the set of requests that originated at $v$,

ACTIVE, for the set of requests that originated at $v$, which are still unsatisfied,

FREE, for the set of resources in RESOURCES that have not yet been "captured" by requests,

GRANT, for a single captured resource on its way back to a request,
ARRIVALS $(w), w \in$ NEIGHBORS, for the number of "ARRIVAL" messages received from each of $v$ 's children, and sent to $v$ 's parent, respectively:

BUYERS $(w), w \in$ NEIGHBORS, for the number of "BUYER" messages sent to $v$ 's children and received from $v$ 's parent, respectively,

NETGRANTS $(w), w \in \operatorname{NEIGHBORS}$, for the net flow of captured resources out of $v$ to each of its neighbors,

NEXT, a temporary variable which can hold a vertex.
Initialization of Independent Variables. REQUESTS = ACTIVE $=\varnothing$, FREE $=$ RESOURCES, and all other variables are 0.

Dependent Variables and Their Dependencies. CAPTURED, for the set of resources in RESOURCES which have been captured;

Dependency:CAPTURED = RESOURCES - FREE.
SATISFIED, for the set of requests in REQUESTS which have been satisfied,

> Dependency:SATISFIED = REQUESTS - ACTIVE.

NETBUYERS $(w), w \in$ NEIGHBORS, for the net flow of buyers and virtual buyers into $v$ from each neighbor (Recall the definition of "virtual buyers" from the last subsection.),

Dependency: If $w=\operatorname{PARENT}$ then NETBUYERS $(w)=$ $\min (\operatorname{PLACE}(D E S C E N D A N T S) ~-~ A R R I V A L S(w)$, BUYERS $(w)$ ).
If $w \in$ CHILDREN, then NETBUYERS $(w)=-\min ($ PLACE
(DESCENDANTS $(w))$ - ARRIVALS $(w)$, BUYERS $(w)$ ).

These two equations can be understood as follows. Consider, for example, the first equation, for $w=$ PARENT. If PLACE(DESCENDANTS $) ~-~ A R R I V A L S ~(w)<\operatorname{BUYERS}(w)$, it means that the placement originally given for $v$ 's subtree is not adequate for handling the requests (arrivals) which have originated in $v$ 's subtree, together with the "BUYER" messages sent down from $w$. Therefore, all the resources in $v$ 's subtree are allocated to requests, either within or outside of $v$ 's subtree. Whether the net flow of buyers should be regarded as into or outward from $v$ 's subtree then depends solely on the sign of PLACE(DESCENDANTS ) - ARRIVALS $(w)$, without regard to the number of "BUYER" messages received from $w$. That is, if PLACE(DESCENDANTS $) \leqslant \operatorname{ARRIVALS}(w)$, then the sign is negative and the net flow of buyers is outward from $v$ 's subtree, while otherwise it is inward; in either case, its magnitude is equal to |PLACE(DESCENDANTS) ARRIVALS $(w) \mid$. On the other hand, if PLACE(DESCENDANTS) $\operatorname{ARRIVALS}(w) \geqslant \operatorname{BUYERS}(w)$, then the placement originally given for $v$ 's subtree is adequate for handling both the requests which have originated in $v$ 's subtree, together with the "BUYER" messages sent down from $w$. Therefore, the net flow of buyers is inward, and its amount is just equal to $\operatorname{BUYERS}(w)$, without regard to the other two values. The second equation is similar, with appropriate changes of sign.

Another way to understand the equations is as follows. Again, consider the first equation, for $w=\operatorname{PARENT}$. Then NETBUYERS $(w)=$ $\operatorname{BUYERS}(w)$ - VIRTBUYERS $(w)$, where the latter quantity is the number of virtual buyers which $v$ estimates it has sent to its parent. Using the expression which was derived in the preceding subsection for the number of virtual buyers, we see that NETBUYERS $(w)=\operatorname{BUYERS}(w)-$ $\max (\operatorname{ARRIVALS}(w)+\operatorname{BUYERS}(w)-\operatorname{PLACE}(D E S C E N D A N T S), 0)$. This is equal to $\min (\operatorname{PLACE}(D E S C E N D A N T S)-\operatorname{ARRIVALS}(w)$, BUYERS $(w)$ ), as needed. Again, the other calculation is similar.

The remaining dependent variables are:

NETBUYERS, for the total of all the NETBUYERS( $w$ ),
Dependency:NETBUYERS $=\sum_{w \in \text { NEIGHBORS }}$ NETBUYERS $(w)$.
NETFLOW, for the net flow of buyers into $v$,
Dependency:NETFLOW $=\mid$ REQUESTS $\mid+$ NETBUYERS.
NETGRANTS, for the net flow of grants out of $v$,

$$
\text { Dependency:NETGRANTS }=\sum_{w \in \text { NEIGHBORS }} \text { NETGRANTS }(w) .
$$

```
/* The following invariants hold at the beginning of any node step.
    NETFLOW = \CAPTURED|.
    ARRIVALS(PARENT) = |REQUESTS | + \Sigma w | CHILDAENARRIVALSS(w)].
    GRANT = 0.
    NETGRANTS = |CAPTURED| - |SATISFIED|. */
(Part 1)
If }M\mathrm{ is a request then
    [REQUESTS := REQUESTS U{M}
    ACTIVE := ACTIVE U {M}
    send "ARRIVAL" to PARFNT
    ARRIVALS(PARENT) := ARRIVALS(PARENT) + 1]
If }M=\mathrm{ "ARRIVAL" from w then
    [ARRIVALS(w) := ARRIVALS(w) + I
    send "ARRIVAL" to PARENT
    ARRIVALS(PARENT) := ARRIVALS(PARENT) + 1]
If M = "BUYER" then BUYERS(PARENT) := BUYERS(PARENT) + 1
1* Now slightly revised invariants hold:
    |CAPTURED} \leq NETFLOW \leq {CAPTURED| + 1.
    ARRIVALS(PARENT) = {REQUESTS | + \Sigma wecHILDRENARRIVALS(vi)]
    GRANT = 0.
    NETGRANTS = |CAPTURED| - |SATISFIED|.*/
(Part 2)
1* Next, if there is an excess request. service it. */
If NETFLOW = {CAPTURED| + 1 then
    if FREE }\not=
        then
            [choose s E FREE
            FREE := FREE - [s]
            GRANT := Ss]
        else
            [/* Send "BUYER" down into a subtree. "/
            S := {s G CHILDREN: PLACE(DESCENDANTS(s)) > ARRIVALS(s) + BUYERS(s)}
            /*S\not=\varnothing.*/
            choose NEXT E S
            send "BUYER" to NEXT
            BUYERS(NEXT) := BUYEERS(NEXT) + 1]
f* NETFLOW = |CAPTURED|.*/
(Part 3)
/* Process M if M is a captured resource message. */
If }M\mathrm{ is a resource from }W\mathrm{ then
    [NETGRANTS(w) := NETGRANTS(w) - 1
        GRANT := M]
/* Send a captured resource, if you have one, toward a request origin. */
If GRANT }\not=0\mathrm{ then
    /* NETGRANTS = |CAPTURED| - |SATISFIED| - 1. */
    [if ACTIVE }\ddagger\emptyset\mathrm{ then
        then
            [chooser E ACIIVE
            ACTIVE := ACTIVE - {r}
            output (r,GRANT)]
        e]se
            [choose w E NEIGHBORS with NETBUYERS(w) - NEIGRANTS(w) > 0
            send GRANT to w
            NETGRANTS(w):= NETGRANTS(w) + 1]
    GRANT := 0]
```

Figure 2

A code (see Fig. 2) is executed in response to the receipt of any message or request, $M$. The first part of the code does initial processing of messages, updating estimates, and sending any required "ARRIVAL" messages. After the first part of the code, there will be at most one excess request leit at the node, and if there is such a request left at the node, then the node is able to service that request, either locally or by sending a buyer into a subtree. In the second part of the code, the node decides where it can service such an excess request, and it does so. (In case a buyer is sent down into a subtree, the subtree is chosen nondeterministically. Later, we will refine the algorithm to use a probabilistic choice at this point.) Finally, in the third part of the code, the node processes an excess captured resource, if it happens to have one. (It cannot have more than one.) The node can have a captured resource either because $M$ was a captured resource message, or because (in the second part of the code) the node itself captured a local resource. The resource is granted to a local request if possible; otherwise, it is sent in such a way as to negate the net flow of buyers into the node along some edge. It is always possible to process such a captured resource in one of these ways.

## 3. Correctness of the Algorithm

Theorem 1. The given algorithm solves the resource allocation problem.
Proof. We claim that the node program given above implements the strategy described informally in the previous section. We do not give a proof of this correspondence here. Rather, we argue correctness of the key assertions of the program and give informal arguments for the rest of the proof of correctness of the algorithm.

The first portion of the algorithm, the initial processing of the first three kinds of messages, simply sends the appropriate "ARRIVAL" messages and records the proper changes to the various sets and counters.

For any of the three kinds of messages, node $v$ is finding out about a new request that needs to be processed. In some cases, $v$ will need to do more to help process the request. If the message is an "ARRIVAL," and node $v$ thinks that the corresponding request can be serviced in the sender's subtree, then $v$ has no further work to do. If the message is a request or an "ARRIVAL," and if node $v$ thinks that it is impossible to service that request in $v$ 's subtree, then the "ARRIVAL" message sent upward by $v$ will be counted by $v$ 's parent in its estimate of virtual buyers emanating from $v$ 's subtree. Thus, after sending this "ARRIVAL" message upward, $v$ will have no further work to do. Also, if the message is a "BUYER" and $v$ thinks that it is impossible to service the request in $v$ 's subtree, then $v$ need not do
anything more. However, if $v$ thinks that the new request can be serviced in its subtree, then it has some further work to do, in the second portion of the algorithm.
The second portion of the algorithm manages the disposition of any excess flow of requests into the node. We must first check that the number of excess requests after the initial processing of a single message can only be 0 or 1 . That is, we must verify that $\mid$ CAPTURED $\mid \leqslant$ NETFLOW $\leqslant$ $\mid$ CAPTURED $\mid+1$ between Parts 1 and 2 of the code.

A quick check of the cases shows that the only way this could fail to be true is if $M=$ "ARRIVAL" from a child $s$, and the result of processing $M$ causes $\operatorname{NETBUYERS}(s)$ to remain unchanged, while NETBUYERS(PARENT) decreases. In this case, we can deduce some relationships among the values of $v$ 's local variables at the beginning of the node step.

For every $t \in$ CHILDREN, it must be the case just before execution of Part 1 that

$$
\begin{aligned}
& \text { - NETBUYERS }(t)=\min (\operatorname{PLACE}(\operatorname{DESCENDANTS}(t))-\operatorname{ARRI}- \\
& \text { VALS }(t), \operatorname{BUYERS}(t)),
\end{aligned}
$$

so that
$-\operatorname{NETBUYERS}(t) \leqslant \operatorname{PLACE}(\operatorname{DESCENDANTS}(t))-\operatorname{ARRIVALS}(t)$. That is,
$\operatorname{NETBUYERS}(t) \geqslant \operatorname{ARRIVALS}(t)-\operatorname{PLACE}(\operatorname{DESCENDANTS}(t))$.
Since NETBUYERS( $s$ ) remains unchanged, then it must be the case that
$-\operatorname{NETBUYERS}(s) \quad \neq \quad \operatorname{PLACE}(\operatorname{DESCENDANTS}(s))-\operatorname{ARRI-}$
$\operatorname{VALS}(s)$.
(If they were equal, then PLACE(DESCENDANTS)(s))-ARRI$\operatorname{VALS}(s)=\min (\operatorname{PLACE}(\operatorname{DESCENDANTS}(s)) \quad-\operatorname{ARRIVALS}(s)$, $\operatorname{BUYERS}(s)$ ), and an increase to $\operatorname{ARRIVALS}(s)$ would cause a change to the minimum, thereby changing NETBUYERS $(s)$.)

Therefore, $\operatorname{NETBUYERS}(s) \neq \operatorname{ARRIVALS}(s)-\operatorname{PLACE}(D E S C E N-$ $\operatorname{DANTS}(s)$ ), and so

NETBUYERS $(s)>\operatorname{ARRIVALS}(s)-\operatorname{PLACE}(\operatorname{DESCENDANTS}(s))$.
Since NETBUYERS(PARENT) decreases, it means that NETBUYERS(PARENT) $=$ PLACE(DESCENDANTS) - ARRIVALS (PARENT).

Now consider NETFLOW $=\mid$ REQUESTS $\mid+$ NETBUYERS. The right side is equal to
$\mid$ REQUESTS $\mid$ + NETBUYERS $($ PARENT $)+\operatorname{NETBUYERS~}(s)$

$$
+\sum_{t \in \text { CHILDREN }, t \neq s} \text { NETBUYERS }(w) .
$$

By previous results, this is, in turn, strictly greater than

$$
\begin{aligned}
& \mid \text { REQUESTS } \mid+\operatorname{PLACE}(D E S C E N D A N T S) ~-~ A R R I V A L S ~ \\
& \text { (PARENT) }+\operatorname{ARRIVALS}(s)-\operatorname{PLACE}(D E S C E N D A N T S(s))+ \\
& \sum_{t \in \text { ChILDREN }, t \neq s} \operatorname{ARRIVALS}(t)-\operatorname{PLACE}(D E S C E N D A N T S \\
& (t)) .
\end{aligned}
$$

By the ARRIVAL invariant, this is equal to

## PLACE(DESCENDANTS) - PLACE(DESCENDANTS( $s$ ) ) - <br> $\sum_{t \in \text { Children }, t \neq s} \operatorname{PLACE}(D E S C E N D A N T S ~(t))=\operatorname{PLACE}(v)$.

Thus, NETFLOW $>\operatorname{PLACE}(v)$. However, the original invariant says that NETFLOW $=\mid$ CAPTURED $\mid$, and |CAPTURED $\mid$ is never permitted to be greater than $\operatorname{PLACE}(v)$, a contradiction.

We have thus shown that $\mid$ CAPTURED $\mid \leqslant$ NETFLOW $\leqslant \mid$ CAPTURED $\mid+1$ at the point where that claim is made. Thus, there is at most one excess request that requires disposition. In the case where there is an excess request, node $v$ must service that request in its subtree. There are two possibilities: either $v$ can service the request locally, or it cannot. If FREE $\neq \varnothing$, then a free resource is captured to service the excess request. If not, then a "BUYER" message must be sent down into some subtree. We must show that, in the event $\operatorname{FREE}=\varnothing$, it is possible to send such a "BUYER" message. That is, we must check that $S \neq \varnothing$ at the place where that claim is made.
Assume not. We will make some deductions about the values of the variables at the point where that claim is made. At that point, we know that $\operatorname{FREE}=\varnothing$, so that $\operatorname{PLACE}(v)=\mid$ CAPTURED $\mid$. We also know that

NETBUYERS(PARENT) $\leqslant$ PLACE(DESCENDANTS) ARRIVALS(PARENT),
by definition of NETBUYERS. Then
NETFLOW $=\mid$ REQUESTS $\mid+$ NETBUYERS(PARENT) + $\sum_{s \in \text { Children }}$ NETBUYERS $(s) \leqslant \mid$ REQUESTS $\mid+$ PLACE (DESCENDANTS) - ARRIVALS(PARENT) $+\sum_{s \in \text { Children }}$ NETBUYERS(s).

Because $S=\varnothing$, it follows that
$\operatorname{PLACE}(\operatorname{DESCENDANTS}(s))-\operatorname{ARRIVALS}(s) \leqslant \operatorname{BUYERS}(s)$ for each $s \in$ CHILDREN.
Therefore,

```
NETBUYERS(s) = - (PLACE(DESCENDANTS(s))}-\operatorname{ARRIVALS(s)})
```

Thus, the right-hand side of the next-to-last inequality is equal to
$\mid$ REQUESTS $\mid+\operatorname{PLACE}(D E S C E N D A N T S) ~-~ A R R I V A L S(P A R E N T) ~$
$-\sum_{s \in \text { Children }}(\operatorname{PLACE}(D E S C E N D A N T S(s))-\operatorname{ARRIVALS}(s))$.
This expression is, in turn, equal to
$\operatorname{PLACE}(D E S C E N D A N T S)-\sum_{s \in \text { Children }} \operatorname{PLACE}(\operatorname{DESCENDANTS}(s))$
$=\operatorname{PLACE}(v)=\mid$ CAPTURED $\mid$.
Thus, NETFLOW $\leqslant \mid$ CAPTURED $\mid$, a contradiction. Thus, we have shown that it is always possible to service an excess request.

Next, we must show that NETFLOW $=\mid$ CAPTURED $\mid$ between Parts 2 and 3 of the code. This means that after servicing any excess request, there is no remaining request to be serviced. Previous to Part 2, $\mid$ CAPTURED $\mid \leqslant$ NETFLOW $\leqslant \mid$ CAPTURED $\mid+1$. If NETFLOW was equal to |CAPTURED $\mid+1$, then the body of the conditional was executed. If the first case of the conditional held (i.e. the case for FREE $\neq \varnothing$ ), then $\mid$ CAPTURED $\mid$ is increased by 1 , so the invariant is restored. Otherwise, a "BUYER" message was sent to a child, $s$, for which $\operatorname{PLACE}(\operatorname{DESCENDANTS}(s))$ - $\operatorname{ARRIVALS}(s)>\operatorname{BUYERS}(s)$. This caused NETBUYERS $(s)$ to increase by 1, thereby increasing the value of NETFLOW and restoring the invariant.

The third portion of the algorithm manages the travel of captured resources back to requests. First, note that there can be only one captured resource assigned to GRANT at any node in a single step, since the two assignments to GRANT cannot both be executed during a single step. If the message is a captured resource, then no progress is done until the clause contains the second grant. Otherwise, this clause is skipped. We must argue that such a neighbor exists in this case.

Assume not. Then NETBUYERS $(s) \leqslant \operatorname{NETGRANTS}(s)$ for all $s \in$ NEIGHBORS. Now, NETFLOW $=\mid$ CAPTURED $\mid$, so that

$$
\begin{aligned}
\mid \text { CAPTURED } \mid & =\mid \text { REQUESTS } \mid+ \text { NETBUYERS } \\
& \leqslant \mid \text { REQUESTS } \mid+ \text { NETGRANTS } \\
& =\mid \text { REQUESTS }|+| \text { CAPTURED }|-| \text { SATISFIED } \mid-1 \\
& =\mid \text { ACTIVE }|+| \text { CAPTURED } \mid-1 .
\end{aligned}
$$

Therefore, $1 \leqslant \mid$ ACTIVE $\mid$, a contradiction.
Thus, we have checked that the key assertions hold and the code can be executed at all points. We have claimed (and tried to argue) that the algorithm follows the strategy of the preceding section, in setting up a flow of buyers from requests to resources. Eventually, the values of all the NETBUYERS $(w)$ variables will stabilize, and the values taken on by
corresponding NETBUYERS $(w)$ variables at either end of a single edge will be negations of each other. (We use the fact that there are only finitely many requests here. Eventually, no further requests will arrive, so no additional "ARRIVAL" messages will be sent. There is a bound on how many "BUYER" messages will be sent downward along any edge. Therefore, there are only finitely many total "ARRIVAL" and "BUYER" messages which get sent, so that eventually, they will all be delivered.) Similarly, all the REQUESTS variables will eventually stabilize.

Finally, we must consider the travel of captured resources to request origins. Define a new variable, $\operatorname{MESSAGES}(w)$, at node $v$, where $w \in \operatorname{NEIGHBORS}(v)$. Its value is defined to be the number of captured resource messages which have been sent from $w$ to $v$ but have not yet arrived. (Of course, neither $v$ nor $w$ actually "knows" the value of this variable.) For any time, $t$, after the NETBUYERS( $w$ ) and REQUEST values have stabilized, any node $v$, and any $w \in \operatorname{NEIGHBORS}(v)$, let $A(v, w, t)$ be the value of NETBUYERS $(w)-\operatorname{NETGRANTS}(w)+$ $\operatorname{MESSAGES}(w)$ at $v$ at time $t$. Note that $A(v, w, t)=-A(w, v, t)$ in all cases. Let $\operatorname{SUM}(t)$ denote $\sum_{v, w}|A(v, w, t)|$. We claim that any event which involves the receipt of a captured resource message does not change SUM $(t)$, while any event which involves the sending of a captured resource message decreases SUM $(t)$. Therefore, captured resource messages will not be sent forever: they will eventually subside, at which time they must have found a matching request.

First, consider an event involving the receipt of a captured resource, by $v$, from $w$. The only term in the sum which is affected is $A(v, w, t)$. The receipt of the messages causes $v$ 's values of $\operatorname{MESSAGES}(w)$ and NETGRANTS $(w)$ both to decrease by 1 , so that $A(v, w, t)$ is unchanged. Therefore, $\operatorname{SUM}(t)$ is unchanged. Second, consider an event involving the sending of a captured resource, by $v$, to $w$. The only terms in the sum which are affected are $A(v, w, t)$ and $A(w, v, t)$. At time $t$ just prior to the sending event, it must be that $v$ 's value of NETBUYERS $(w)$ NETGRANTS $(w)>0$, which implies that $A(v, w, t)>0$. The result of sending the message is to increase NETGRANTS $(w)$, which means that $A(v, w, t)$ gets decreased by 1 . Therefore, $|A(v, w, t)|$ gets decreased by 1 . Thus, also, $|A(w, v, t)|$ gets decreased by 1 , so that $\operatorname{SUM}(t)$ gets decreased by 2 .

## 4. Monotonicity Analysis

The rest of the paper is devoted to an analysis of the time requirements of the algorithm. Specifically, we measure the sum of the times between requests and their corresponding grants. For the purpose of carrying out
the analysis, certain restrictions will be made. These restrictions will be introduced as needed.

We begin with some basic definitions. Next, we introduce two restrictions which are needed throughout the analysis. Then we define and categorize the complexity measures of interest. We then prove a basic combinatorial result, and use it to prove the monotonicity of the number of "BUYER" messages as a function of interarrival time. Finally, we show that the expected running time of the algorithm is bounded by the expected time for the sequential case of the algorithm.

### 4.1. Definitions

Let $N$ denote the set of natural numbers, including 0 . Let $R^{+}$denote the set of nonnegative reals. If $f$ is a numerical function with domain $V$, then exend $f$ to subsets of $V$ by $f(W)=\sum_{v \in w} f(v)$.

Let $T$ be a rooted tree. We write vertices ${ }_{T}$, and leaves ${ }_{T}$ to denote the indicated sets of vertices of $T$. Let $\operatorname{root}_{T}$ denote the root. If $v \in$ vertices $_{T}$, we write $\operatorname{desc}_{T}(v)$ for the set of vertices of $T$ which are descendants of $v$ (including $v$ itself), parent ${ }_{T}(v)$ for $v$ 's parent in $T$, children ${ }_{T}(v)$ for $v$ 's children, and neighbors $T_{T}(v)$ for children $T_{T}(v) \cup$ parent $\left._{T}(v)\right\}$.

If $v$ is a vertex of $T$, let height $T_{T}(v)$ denote the maximum distance from $v$ to a leaf in its subtree. If $e$ is an edge in $T$, then define height ${ }_{T}(e)$ to be the same as height $T_{T}(v)$, where $v$ is $e$ 's upper endpoint. Let height ${ }_{T}$ denote height $_{T}\left(\operatorname{root}_{T}\right)$.

A placement for $T$ is a function $p$ : vertices $_{T} \rightarrow N$, representing the number of resources at each vertex. We write total $(p)$ for $p\left(\right.$ vertices $\left._{T}\right)$, the total number of resources in the entire tree. We say that $p$ is nonnull provided total $(p)>0$.

A weighted tree $T$ is an undirected, rooted tree with an associated probability density function, $\varphi_{T}$, on the leaves of $T$, such that $\varphi_{T}(v)>0$ for all leaves $v$. (This assumption is made for technical reasons, so that we can normalize probability functions without danger of dividing by 0 .) If $T$ is a weighted tree, $v \in$ internal $_{T}$, and $S$ is a nonempty subset of children $T_{T}(v)$, then let random $T_{T, S}$ denote the probability function which returns $s \in S$ with probability $\varphi_{T}\left(\operatorname{desc}_{T}(s)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(S)\right)$. Thus, random $T_{T, S}$ returns $s$ with probability proportional to the sum of the probability function values for the descendants of $s$.

### 4.2. Initial Restrictions

For the remainder of Section 4, we assume that the following two restrictions hold.

Restriction 1. T is a weighted tree, and the nondeterministic choice step in Part (2) of the algorithm uses a call to random $T_{T, S}$.

## Restriction 2. Delivery time for messages is always exactly 1.

Restriction 1 describes a particular method of choosing among alternative search directions. This method does not use all the information available during execution, but only the "static" probability distribution information available at the beginning of execution. One might expect a more adaptive choice method to work better; however, we do not know how to analyze such strategies.

Restriction 2 has the effect of restricting the executions under consideration; for example, all messages between any two nodes are pipelined-they arrive in the order in which they are sent. While we would like to understand the behavior of the algorithm in the presence of variable message delivery times, such analysis appears to be more difficult.

### 4.3. Cost Measures and Preliminary Results

A request pattern $r$ is a finite sequence of elements of vertices ${ }_{T} \times R^{+}$ whose second components are monotone nondecreasing. A request pattern represents the sequence of requests that occur, their locations and times. If $r$ is a request sequence, then length $(r)$ denotes its length.

A choice sequence $c$ for $v \in$ internal $_{T}$ is an infinite sequence of elements of children $_{T}(v)$, with infinitely many occurrences of each child. If $\mathscr{C}=\left\{c_{v}\right\}$ is a collection of choice sequences, one for each $v \in$ internal $_{T}$, then $\mathscr{C}$ can be used in place of probabilistic choices in an execution of the algorithm, as follows. Each internal node, $v$, makes choices among its children by choosing the first unused element of $c_{v}$ satisfying the inequality $\operatorname{PLACE}\left(\operatorname{desc}_{T}(s)\right)>\operatorname{ARRIVALS}(s)+\operatorname{BUYERS}(s)$. That is, $v$ chooses a child, $s$, for which $v$ thinks there are still remaining resources in $s$ 's subtree.

Let $p$ be a placement for $T$. Let $r$ be a request pattern, and $\mathscr{C}=\left\{c_{v}\right\}$ a collection of choice sequences, one for each $v \in$ internal $_{T}$. Then $\operatorname{cost}_{T, p}(r, \mathscr{C})$ is defined to be the total time from requests to corresponding grants, if requests arrive according to $r$ and $\mathscr{C}$ is used in place of probabilistic choices. (With suitable conventions for handling events which happen at the same time, the execution, and hence the cost, is uniquely defined for fixed $r$ and $\mathscr{C}$.)

The cost measure defined above can be broken up into two pieces, as follows. Let searchcost ${ }_{r, p}(r, \mathscr{C})$ be the total of the times from requests to corresponding captures of resources, if $r$ and $\mathscr{C}$ are used as above. Let returncost $_{T, p}(r, \mathscr{C})$ be the total of the times from captures to corresponding grants of resources.

Now we incorporate a probabilistic construction of $\mathscr{C}$ into the cost measure. If $r$ is a request pattern, let $\operatorname{cost}_{T, p}(r)$ denote the expected value of $\operatorname{cost}_{T, p}(r, \mathscr{C})$, where $\mathscr{C}$ is constructed using $\varphi_{T}$. (That is, for each $v \in$ internal $_{T}$, the sequence $c_{v}$ is constructed by successive choices from
among children $T_{T}(v)$, choosing $s$ with probability $\varphi_{T}\left(\operatorname{desc}_{T}(s)\right) /$ $\varphi_{T}\left(\operatorname{desc}_{T}(S)\right)$, where $S=\operatorname{children}_{T}(v)$. Among the sequences thereby generated are some for which it is not the case that each child occurs infinitely often. However, these sequences form a set of measure $O$, so that we can ignore them in calculating the expected cost measure.) We claim that $\operatorname{cost}_{T, p}(r)$ is exactly the expected total time from requests to grants, provided the algorithm is run in the normal way, using probabilistic choices. That is, the two strategies of constructing choice sequences independently of the algorithm and carrying out the probabilistic choices on-line give identical results.

Let $f$ denote an arbitrary probability density function whose domain consists of positive reals. Extend the domain of the function $\operatorname{cost}_{T, p}$ to the set of such functions by defining $\operatorname{cost}_{T, p}(f)$ to be the expected value of $\operatorname{cost}_{T, p}(r)$, where $r$ is of length total $(p)$, with its successive locations chosen independently using the distribution $\varphi_{T}$, and its successive interarrival times chosen using $f$. That is, at the time the algorithm begins, and at the time of each request, the probability that the next arrival occurs exactly $t$ units later is $f(t)$. We will be primarily interested in this cost, $\operatorname{cost}_{T, p}(f)$.

We define searchcost ${ }_{T, p}(r)$, searchcost ${ }_{T, p}(f)$, etc., analogously to our earlier definitions.

The following claim is true for all domains for which the definitions are valid.

Lemma 2. $\operatorname{cost}_{T, p}=\operatorname{searchcost}_{T, p}+$ returncost $_{T, p}$.

## Proof. Straightforward.

Next, we will relate the given cost measures to the total numbers of various kinds of messages sent during the execution of the algorithm. Note that during the execution of an algorithm, the estimates of "BUYER" and virtual buyer messages sent along an edge can be different at the two ends of the edge. However, after the entire execution of the algorithm is completed, the discrepancy disappears, so that the following definitions are unambiguous. Let $\operatorname{bnum}_{T, p}(r, \mathscr{C})$ denote the total number of "BUYER" messages sent on all edges during the execution of the algorithm on $r$ using $\mathscr{C}$. Let vbnum ${ }_{T, p}(r, \mathscr{C})$ denote the total number of virtual buyer messages sent on all edges during the execution of the algorithm on $r$ using $\mathscr{C}$. Let gnum $_{T, p}(r, \mathscr{C})$ denote the total number of captured resource messages sent on all edges during the execution of the algorithm on $r$ using $\mathscr{C}$. As before, define bnum $_{T, p}(r)$, bnum ${ }_{T, p}(f)$, etc.

Because of the fact that message delivery time is assumed to be exactly 1 , there are some relationships between the measures describing time costs and the measures describing numbers of messages. The following lemma
describes a set of relationships among the various measures. Note that all the statements are true over all possible domains of definition.

Lemma 3. (a) searchcost $t_{T, p} \leqslant$ bnum $_{T, p}+$ vbnum $_{T, p}$.
(b) returncost ${ }_{T, p}=$ gnum $_{T, p}$.
(c) $\operatorname{gnum}_{T, p} \leqslant$ bnum $_{T, p}+$ vbnum $_{T, p}$.
(d) $\operatorname{cost}_{T, p} \leqslant 2\left[\right.$ bnum $_{T, p}+$ vbnum $\left._{T, p}\right]$.

Proof. (a) This inequality is true because buyers continue to make progress up and down the edges of the tree; all time used by the algorithm is occupied by the transmission of appropriate buyer and virtual buyer messages. The reason that we have an inequality rather than an equation here is that buyers are permitted to travel "discontinuously," as described in Section 3.
(b) This equation is true because captured resources travel continuously via captured resource messages.
(c) We must show that each captured resource message always moves in such a way as to "negate" a buyer or virtual buyer message. This is a bit tricky to argue, because of the discrepancies between estimates at opposite ends of an edge. A captured resource only moves over an edge if the net flow of buyers into the node on that edge, as estimated at the near endpoint, is positive. By moving over that edge, the captured resource negates an incoming buyer or virtual buyer along that edge, as estimated at the near endpoint. Because of the assumption that all messages take exactly time 1 , by the time the captured resource reaches the far endpoint, the negated buyer or virtual buyer is also counted in the estimate of outgoing buyers and virtual buyers at the opposite endpoint. The arrival of the captured resource at the far endpoint can thus be regarded as negating an outgoing buyer or virtual buyer at the far endpoint as well.
(d) Straightforward by Lemma 2 and (a)-(c).

Now, we introduce an additional restriction, to remain in force for the remainder of Section 4.

Restriction 3. Thas all leaves at the same distance from the root, and $r$ and $p$ are nonzero only at leaves.

As usual, the following lemma is intended to hold for all valid domains of definition.

Lemma 4. bnum $_{T, p}=\operatorname{vbnum}_{T, p}$.

Proof. We sketch the argument for fixed $r$ and $\mathscr{C}$. For a particular edge $e$, let $a_{e}$ denote the number of request arrivals below $e, b_{e}$ denote the number of "BUYER" messages sent downward along $e$, and $p_{e}$ denote the number of resources placed below $e$, in the execution for $r$ and $\mathscr{C}$. Since all resources get matched to requests, we must have $a_{e}+b_{e} \geqslant p_{e}$, so that the number of virtual buyers sent over edge $e$ is exactly $\max \left(a_{e}+b_{e}-p_{e}, 0\right)=a_{e}+b_{e}-p_{e}$.

Now consider all the edges at any particular height $h$ in the tree. Since all resources and requests are at the leaves, and the branches are all of equal length, it is clear that

$$
\sum_{\operatorname{height}_{T}(e)=h} a_{e}=\operatorname{total}(p)=\sum_{\text {heightrfe }^{(e)=h}} p_{e} .
$$

Therefore,

$$
\sum_{\text {height }(e)=h} b_{e}=\sum_{\text {height }_{T}(e)=h}\left(a_{e}+b_{e}-p_{e}\right) .
$$

That is, the numbers of buyers and virtual buyers sent over edges of height $h$ are equal. Since this is true for all $h$, the result follows.

THEOREM 5. $\operatorname{cost}_{T, p} \leqslant 4\left(\right.$ bnum $\left._{T, p}\right)$
Proof. Immediate from Lemma 3 and Restriction 3.
Thus, in order to obtain an upper bound on $\operatorname{cost}_{T, p}(f)$, it suffices to prove a bound for bnum $_{T, p}(f)$.

### 4.4. A Combinatorial Result

This subsection contains a key combinatorial result which will be used in the subsequent analysis. We model the behavior of the algorithm at a single node $v$. The children of $v$ are modelled as a set of bins for resources. (Here, we do not concern ourselves about the tree structure beyond the children.) Let $c$ be a choice sequence for $v$. Each bin $s$ is initialized to contain a number $p(s)$ of resources.

The arrival of messages at $v$ is described by a script $S$. A script is a finite sequence of symbols, each of which is either a bin number $s$ or an " $X$." A bin number represents the arrival of an "ARRIVAL" message from the specified child. The symbol $X$ represents the arrival of a "BUYER" message from $v$ 's parent.

The processing of script $S$ on $c$ and $p$, is as follows. The elements of $S$ are processed sequentially. If $S(i)$ is:

```
s G bins, then
    if bin s is nonempty,
    then subtract 1 from the number of resources in s
    else if some bin is nonempty then
        [SELECT the first unused element of c describing a nonempty bin, t;
        subtract 1 from the number of resources in t]
x, then
    if some bin is nonempty then
        [SELECT the first unused element of c describing a nonempty bin, t;
        subtract 1 from the number of resources in t]
```

Define $\operatorname{SELECT}(S, c, p, i)$ to be the number of times bin $i$ is SELECTed during the course of processing $S$ on $c$ and $p$. (Note that a bin is only said to be SELECTed when the choice sequence is used to choose it, and not when it is explicitly chosen by the script.) Define choice $(S, c, p, j)$ to be equal to $k$ provided that when $S(j)$ is processed on $c$ and $p$, the $k$ th element of $c$ is used to select a bin. (If no element of $c$ is used, then choice $(S, c . p, j)$ is undefined.) It follows that $\operatorname{SELECT}(S, c, p, i)=$ $|\{j: c(\operatorname{choice}(S, c, p, j))=i\}|$,

For any script $S$, let binsequence $(S$ ) denote the subsequence of $S$ consisting of bin numbers. Script $S$ is said to dominate script $S^{\prime}$ provided that: (a) $T=T^{\prime}$, where $T=\operatorname{binsequence}(S)$ and $T^{\prime}=$ binsequence $\left(S^{\prime}\right)$, (b) the total number of $X$ 's in $S$ is at least as great as the total number of $X \mathrm{~s}$ in $S^{\prime}$, and (c) for each $i$, the number of $X$ 's in $S$ preceding $T(i)$ is at least as great as the number of $X^{\prime}$ s in $S^{\prime}$ preceding $T^{\prime}(i)$. The main result of this section is that, if $S$ dominates $S^{\prime}$, then $\operatorname{SELECT}(S, c, p, i) \geqslant \operatorname{SELECT}\left(S^{\prime}, c, p, i\right)$ for all $c, p$, and $i$.

We say that an interchange of two consecutive elements of a script $S$ is legal provided that the first element of the pair is an $X$. We say that a script $S^{\prime}$ is reachable from a script $S$ if $S$ can be transformed to $S^{\prime}$ by a series of legal interchanges. Note that $S$ dominates all scripts $S^{\prime}$ reachable from $S$; moreover, if $S$ dominates $S^{\prime}$, then $S^{\prime}$ can be augmented with some suffix of $X$ 's, to a script which is reachable from $S$.

Lemma 6. For any scripts $S$ and $S^{\prime}$ such that $S^{\prime}$ is reachable from $S$, and for any choice sequence $c$, placement function $p$ and bin $i$,

$$
\operatorname{SELECT}(S, c, p, i) \geqslant \operatorname{SELECT}\left(S^{\prime}, c, p, i\right)
$$

Proof. We prove this lemma by showing that if $S^{\prime}$ is reachable from $S$ by a single legal interchange, then the inequality holds. The lemma follows by induction on the number of legal interchanges.

Fix $S, c$, and $p$. Assume that $S^{\prime}$ is obtained from $S$ by interchanging $S(j)=X$ with $S(j+1)$. If $S(j+1)=X$, then $S=S^{\prime}$, so the result is obvious. So assume $S(j+1)=s \in$ bins. There are three cases.

Case 1. Bin $s$ is empty after processing $S(1) \cdots S(j-1)$ on $c$ and $p$. Then choices using $c$ are made for both $S$ and $S^{\prime}$ at both steps $j$ and $j+1$. Thus, choice $(S, c, p, j)=\operatorname{choice}\left(S^{\prime}, c, p, j\right)$ and choice $(S, c, p, j+1)=$ choice ( $S^{\prime}, c, p, j+1$ ). The number of resources remaining in each bin after step $j+1$ is the same for $S$ and $S^{\prime}$, and therefore processing continues identically for $S$ and $S^{\prime}$ after that point. Thus, $\operatorname{SELECT}(S, c, p, i)=$ $\operatorname{SELECT}\left(S^{\prime}, c, p, i\right)$.

Case 2. Bin $s$ contains more than one resource after processing $S(1) \cdots S(j-1)$ on $c$ and $p$, or else $c$ (choice $(S, c, p, j)$ ) is not bin $s$. (That is, the bin selected by the choice made at step $j$ is not bin $s$.) Then the effect of the pair of steps $j$ and $j+1$ is the same for both $S$ and $S^{\prime}$ : a resource is removed from bin $s$ and a resource is removed from bin $c$ (choice $(S, c, p, j)$ ). (When processing $S$, the choice from $c$ occurs first, while when processing $S^{\prime}$, the explicit removal from $s$ occurs first, but the net effect is the same.) Subsequent processing of the two scripts will be identical, and once again, $\operatorname{SELECT}(S, c, p, i)=\operatorname{SELECT}\left(S^{\prime}, c, p, i\right)$.

Case 3. Bin $s$ contains exactly one resource after processing $S(1) \cdots S(j-1)$, and $c($ choice $(S, c, p, j))=s$. (That is, the bin selected by the choice made at step $j$ is $s$.) In this case, the processing of $S$ uses choices from $c$ at both steps $j$ and $j+1$, because the choice of $s$ at step $j$ removes the last resource from bin $s$, and so a choice must also be made at step $j+1$. The processing of $S^{\prime}$ does not need a choice at step $j$, although it is forced to choose by the $X$ at step $j+1$. Thus, in both cases, step $j$ removes the last resource from bin $s$, while step $j+1$ makes a choice using $c$. Then choice $(S, c, p, j+1)=$ choice $\left(S^{\prime}, c, p, j+1\right)$; that is, the same entry in $c$ is used at step $j+1$ in both cases. The combined effect of steps $j$ and $j+1$ on the bins is the same for the two scripts. Subsequent processing is again identical, so $\operatorname{SELECT}(S, c, p, i)=\operatorname{SELECT}\left(S^{\prime}, c, p, i\right)$ for bin $i \neq s$, and $\operatorname{SELECT}(S, c, p, s)=\operatorname{SELECT}\left(S^{\prime}, c, p, s\right)+1>\operatorname{SELECT}\left(S^{\prime}, c, p, s\right)$.

We can now state the main result of this section.

Corollary 7. For any scripts $S$ and $S^{\prime \prime}$ such that $S$ dominates $S^{\prime}$, and for any choice sequence $c$, placement function $p$ and bin $i$,

$$
\operatorname{SELECT}(S, c, p, i) \geqslant \operatorname{SELECT}\left(S^{\prime}, c, p, i\right)
$$

Proof. Let $T$ be an augmentation of $S^{\prime}$ by a suffix of $X^{\prime}$ s, such that $T$ is reachable from $S$. Then Lemma 6 implies that $\operatorname{SELECT}(S, c, p, i) \geqslant$ $\operatorname{SELECT}(T, c, p, i)$. But the latter term is obviously at least as great as $\operatorname{SELECT}\left(S^{\prime}, c, p, i\right)$.

### 4.5. Expansions

In this subsection, we show that the number of "BUYER" messages sent is a monotone nondecreasing function of the interarrival time of the arrival distribution. We do this by comparing particular pairs of executions.

For $n \in N$, let $[n]$ denote $\{1, \ldots, n\}$. If $a \in R$ and $r=\left(v_{i}, t_{i}\right)_{i \in[k]}$, is a request pattern, then $a r$, the expansion of $r$ by $a$, is the request pattern $\left(v_{i}, a t_{i}\right)_{i \in[k]}$. That is, ar represents the request pattern in which the successive requests occur at the same locations as in $r$, but in which the times are expanded by the constant factor $a$.

We will compare executions for request pattern $r$ and request pattern ar, using the same choice sequence. We require a technical restriction, just to avoid the complications of having to consider multiple events occurring at the same node at the same time, in either execution. A request pattern $r$ is said to be a-isolated provided that no two requests occur in $r$ at the same time, and provided that the following holds: If $t_{1}$ and $t_{2}$ are two times at which requests arrive in $r$, where $t_{1} \neq t_{2}$, and if $k$ is an integer, then the following are true: (a) $t_{1}-t_{2} \neq 2 k$, and (b) $t_{1}-t_{2} \neq(2 / a) k$.

The next lemma is crucial to our analysis. Its truth was first observed empirically, and then proved analytically. It says that the number of "BUYER" messages sent during an execution cannot increase if the request pattern is expanded by a constant which is greater than or equal to 1 .

Lemma 8. If $a \geqslant 1$, and $r$ is a-isolated, then $\operatorname{bnum}_{T, p}(r, \mathscr{C}) \leqslant$ $\operatorname{bnum}_{T, p}(a r, \mathscr{C})$.

Proof. Fix $T, p$, and $\mathscr{C}$. Let $\operatorname{bsent}(r, e, t)$ denote the number of "BUYER" messages sent along edge $e$, in the execution for $r$ (using $T, p$, and $\mathscr{C}$ ), up to and including time $t$. Let $\operatorname{brec}(r, v, t)$ denote the number of "BUYER" messages received by vertex $v$, in the execution for $r$, up to and including time $t$. Let $\operatorname{arec}(r, e, t)$ denote the number of "ARRIVAL" messages received along edge $e$, in the execution for $r$, up to and including time $t$. We will show the following:

Claim. $\quad \operatorname{bsent}\left(r, e, t+\operatorname{height}_{T}(e)\right) \leqslant \operatorname{bsent}\left(a r, e, a t+\operatorname{height}_{T}(e)\right)$ for all $r$, $e$, and $t$.

This is a stronger claim than required for the lemma, since it shows an inequality not only for the total number of "BUYER" messages, but for the number along each edge, up to corresponding time.
$\operatorname{FACT} 1 . \quad \operatorname{arec}\left(r, e, t+\operatorname{height}_{T}(e)\right)=\operatorname{arec}\left(a r, e, a t+\operatorname{height}_{T}(e)\right)$.
This is so because the number of requests arriving in request sequence $r$ by time $t$ is the same as the number arriving in request sequence $a r$ by time $a t$, and messages just flow up the tree at a steady rate.

The rest of the proof proceeds by induction on height ${ }_{T}(e)$, starting with height $_{T}(e)=$ height $_{T}$, and working downward towards the leaves:

$$
\text { Base: }^{\text {height }}(e)=\text { height }_{T} .
$$

In this case, $e$ 's upper endpoint is root $_{T}$. The actions of root ${ }_{T}$ are completely determined by the "ARRIVAL" messages it receives, which are the same at corresponding times in the two executions, by Fact 1 . The claim follows.

Inductive step. height ${ }_{T}(e)<$ height $_{T}$. Let $v$ be the upper endpoint of edge $e$, and the lower endpoint of edge $e^{\prime}$.

FACT 2. $\quad \operatorname{brec}\left(r, v, t+\operatorname{height}_{T}(v)\right) \leqslant \operatorname{brec}\left(a r, v, a t+\operatorname{beight}_{T}(v)\right)$.
This is so because $\quad \operatorname{brec}\left(r, v, t+\right.$ height $\left._{T}(v)\right)=\operatorname{bsent}\left(r, e^{\prime}, t+\right.$ height $\left._{T}(v)-1\right)$ because all messages take exactly time $1,=\operatorname{bsent}\left(r, e^{\prime}\right.$, $\left.t-2+\operatorname{height}_{T}\left(e^{\prime}\right)\right), \quad \leqslant \operatorname{bsent}\left(a r, e^{\prime}, a(t-2)+\right.$ height $\left._{T}\left(e^{\prime}\right)\right) \quad$ by inductive hypothesis, $=\operatorname{bsent}\left(a r, e^{\prime}, a(t-2)+1+\operatorname{height}_{T}(v)\right)$, $=\operatorname{brec}(a r, v, a(t-2)+$ $\left.2+\operatorname{height}_{T}(v)\right) \leqslant \operatorname{brec}\left(a r, v, a t+\operatorname{height}_{T}(v)\right)$, since $a(t-2)+2 \leqslant a t$.

Now let us consider the situation from $v$ 's viewpoint. Node $v$ is comparing two executions, the first for $r$ and the second for ar. All $v$ sees is its incoming "ARRIVAL" and "BUYER" messages, and $v$ uses the same choice sequence in both cases. At corresponding times in the two executions (i.e., $t+$ height $_{T}(v)$ in the first execution vs at + height $_{T}(v)$ in the second execution), the same number of "ARRIVAL" messages have been received along each edge (by Fact 1), and an inequality holds for the number of "BUYER" messages which have been received (by FACT 2 ). We will show the needed inequality for the number of "BUYER" messages sent by $v$ along each edge, up to corresponding times.

Fix any time $t$. We compare the first execution up to time $t+$ height $_{T}(v)$ with the second execution up to time $a t+$ height $_{T}(v)$. We claim that this situation is modelled by the combinatorial problem presented in the preceding subsection. First, we represent $v$ 's inputs in each of the two executions by a script, i.e., a sequence of $X$ 's and "bins" the latter of which are identified with children of $v$. An $X$ models the arrival of a "BUYER," while $s \in$ bins model the receipt of an "ARRIVAL" message from $s$. (The fact that $r$ is a-isolated means that no two of $v$ 's inputs occur at the same time in the same execution, so a unique sequence can be obtained in each case.) Let $S$ and $S^{\prime}$ be the scripts for the first and second executions, respectively (up to the indicated times). The claims in the preceding paragraph imply that $S^{\prime}$ dominates $S$.

We claim that the processing described for the combinatorial problem
models the processing carried out at $v$ during execution of the algorithm. In particular, a SELECT of a bin $s$, if it occurs, models the sending of a "BUYER" message to $s$ and associated reduction of $v$ 's estimate of the number of resources remaining in $s$ 's subtree. With the given correspondence between the combinatorial problem and the executions, the conclusion of Corollary 7 translates immediately into the claim.

Lemma 9. If $a \geqslant 1$, and $r$ is $a$-isolated, then $\operatorname{bnum}_{T, p}(r) \leqslant \operatorname{bnum}_{T, p}(a r)$.
Proof. By Lemma 8, taking expectations.
Define bnum $_{T, p}(a, f)$ to be the expected value of bnum $T_{T, p}(a r)$, where $r$ is chosen according to $\varphi_{T}$ and $f$. The next theorem states that the expected number of "BUYER" messages is a monotone nondecreasing function of the interarrival time of the request distribution.

Theorem 10. (a) If $a \geqslant 1$, then $\operatorname{bnum}_{T, p}(f) \leqslant \operatorname{bnum}_{T, p}(a, f)$.
(b) If $0<a \leqslant b$, then bnum $_{T, p}(a, f) \leqslant \operatorname{bnum}_{T, p}(b, f)$.

Proof. (a) If a request sequence is chosen according to $\varphi_{\tau}$ and $f$, then with probability 1 , it will be $a$-isolated. The result then follows from Lemma 9 , by taking expectations over $r$.
(b) Let $g$ be the probability density function defined by $g(a t)=f(t)$. Since $b / a \geqslant 1$, we can apply Part (a) to $b / a$ and $g$, obtaining $\operatorname{bnum}_{T, p}(g) \leqslant \operatorname{bnum}_{T, p}(b / a, g)$. But $\quad \operatorname{bnum}_{T, p}(g)=\operatorname{bnum}_{T, p}(a, f) \quad$ and $\operatorname{bnum}_{T, p}(b / a, g)=\operatorname{bnum}_{T, p}(b, f)$, yielding the result.

### 4.6. Summary of Monotonicity Analysis

In this section, we have made the following restrictions, repeated here for convenience.

Restriction 1. $T$ is a weighted tree, and the nondeterministic choice step in Part (2) of the algorithm uses a call to random ${ }_{T, S}$.

Restriction 2. Delivery time for messages is always exactly 1.
Restriction 3. Thas all leaves at the same distance from the root, and $r$ and $p$ are nonzero only at leaves.

The major results we have proved in this section are that the expected response time is closely approximated by the expected number of "BUYER" messages (Theorem 5) and that the expected number of "BUYER" messages is a monotonic function of the interarrival time (Theorem 10). We can combine these two results, obtaining the following:

Theorem 11. $\operatorname{cost}_{T, p}(f) \leqslant 4 \lim _{a \rightarrow \infty} \operatorname{bnum}_{T, p}(a, f)$.
Proof. Consider what happens to the value of $\operatorname{bnum}_{T, p}(a, f)$ as $a$ increases. This value is monotonically nondecreasing, by Theorem 10. Also, it is bounded above, because no execution causes more "BUYER" messages to be sent on any edge than the number of resources initially placed below that edge. Therefore, the limit exists. The result follows immediately from Theorems 5 and 10.

That is, the expected cost of the algorithm for any probability function $f$ is bounded in terms of the limiting case of the algorithm, as the interarrival time approaches infinity. But note that as the interarrival time approaches infinity, the algorithm gravitates towards a purely sequential algorithm-one in which each request gets satisfied before the next one arrives. This kind of sequential algorithm is amenable to analysis of a more traditional kind, the subject of the next section of this paper. It seems quite surprising that the sequential case is the worst case. Our initial expectation was that cases where considerable interference between requests occurs would be the worst case. The monotonicity theorem indicates that that is not so. Of course, we have made a few assumptions in this section, most significantly the equal lengths of branches. It is quite likely that the sequential case will not be the worst case for an algorithm using more general tree topologies. The analysis in this more general case so far seems quite intractable, however.

## 5. Sequential Analysis

In this section, we analyze the sequential case of the algorithm. In the next section, we combine the results into an upper bound for the entire algorithm. Once again, we allow arbitrary weighted trees $T$, and allow $r$ and $p$ to be nonzero anywhere.

### 5.1. A Simplified Problem and Algorithm

The sequential case of the algorithm offers considerable simplification over the concurrent cases. There is no interference at all, since each request arrives after previous requests have been satisfied. This means that all the estimates of remaining resources are completely accurate. In fact, the result is equivalent to that of an algorithm in which all information is known globally.

The behavior of the algorithm in the sequential case can be modelled by repeated calls to the following sequential program, FIND. The program takes a weighted tree, a nonnull placement, and a vertex (the vertex at
which a request occurs) as input, and returns a vertex (the vertex at which the resource to be granted is located).

```
FIND(T,p,v)
Case
    p(v)>0 : return v
    p(desc
    p(v)=0 and p(desc
        [S:={w\inchildren (v):p{desc}\mp@subsup{C}{T}{}(w))>0
        return FIND(T,p,random
endcase
```

Thus, a request is satisfied, as before, in the smallest containing subtree which contains a resource; where there is a choice, the probability function is used.

Lemma 12. If $p$ is nonnull and $v \in$ vertices $_{T}$, then $\operatorname{FIND}(T, p, v)$ eventually halts and returns a vertex, $v$, with $p(v)>0$.

## Proof. Straightforward.

We next want to prove a lemma which will be useful in the later analysis. The content of the lemma is as follows. Let random ${ }_{T}$ denote the probability function which returns $s \in$ leaves $_{T}$ with probability $\varphi_{T}(s)$. Assume $T$ is a weighted tree and $p$ is a placement which is nonzero only at leaves. Consider the following two experiments:
(1) Call $\operatorname{FIND}\left(T, p, \operatorname{root}_{T}\right)$, and
(2) Call $\operatorname{FIND}\left(T, p\right.$, random $\left._{T}\right)$.

We claim that the "results" of these two experiments are the same. That is, for each $w \in \operatorname{vertices}_{T}$, the probability that experiment (1) returns $w$ is exactly the same as the probability that experiment (2) returns $w$. It will follow from the next lemma that this is so.

Some notation is helpful here. The result of FIND on a particular set of arguments can be expressed as a probability distribution of vertices. Let $\alpha_{T, p, v}$ denote the probability distribution of results for $\operatorname{FIND}(T, p, v)$. That is, $\operatorname{FIND}(T, p, v)$ returns $w$ with probability $\alpha_{T, p, v}(w)$.

Lemma 13. Let $T$ be a weighted tree, $p$ a nonnull placement for $T$. Assume that $p$ is nonzero only at leaves. Then the following are true:
(a) If $v \in$ internal $_{T}$, then $\alpha_{T, p, v}=\sum_{w \in \operatorname{children}_{T}(v)}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) /\right.\right.$ $\left.\left.\varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{T, p, w}\right]$.
(b) If $v \in$ vertices $_{T}$, then $\alpha_{T, p, v}=\sum_{w \in \operatorname{desc}_{T}(v)}\left[\left[\varphi_{T}(w) /\right.\right.$ $\left.\left.\varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{T, p, w}\right]$.

Proof. In the proof, we assume $T$ and $p$ are fixed and write $\alpha_{v}$ in place of $\alpha_{T, p, v}$, etc.
(a) We consider cases.

Case 1. $p\left(\operatorname{desc}_{r}(v)\right)=0$. Then since the algorithm immediately calls FIND on parent $T_{T}(v)$, we see that $\alpha_{v}=\alpha_{\text {parent }_{T(v)}}$. Similarly, for all children $w$ of $v$, we have $\alpha_{w}=\alpha_{\text {parent }_{T}(v)}$. Since $\sum_{w \in \operatorname{children} r(v)}\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) /\right.$ $\left.\varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right]=1$, the result follows.

Case 2. $p\left(\operatorname{desc}_{T}(v)\right)>0$. Since we are assuming that $v \in$ internal $_{T}$, we know that $p(v)=0$. Let $S=\left\{w \in \operatorname{children}_{T}(v): p\left(\operatorname{desc}_{T}(w) \gg 0\right\}\right.$. Then $S \neq \varnothing$. The third case in the algorithm holds, and we have that $\alpha_{v}=\sum_{w \in S}\left[\left[\varphi\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(S)\right)\right] \alpha_{w}\right]$. Now,

$$
\begin{aligned}
& \quad \sum_{w \in \operatorname{children}_{T}(v)}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{w}\right] \\
& \quad=\sum_{w \in S}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{w}\right] \\
& \quad+\sum_{w \in \operatorname{children}_{T(v)-S}}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{w}\right] .
\end{aligned}
$$

By the remark above, this sum is equal to

$$
\begin{aligned}
& \alpha_{v}\left[\varphi_{T}\left(\operatorname{desc}_{T}(S)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \\
& \quad+\sum_{w \in \operatorname{children} T(v)-S}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{w}\right] .
\end{aligned}
$$

If $w \in \operatorname{children}_{T}(v)-S$, we know that $p\left(\operatorname{desc}_{T}(w)\right)=0$, so that $\alpha_{w}=\alpha_{\text {parent }_{T}(w)}=\alpha_{v}$. So the sum above is equal to

$$
\begin{aligned}
\alpha_{v}\left[\varphi_{T}( \right. & \left.\left.\operatorname{desc}_{T}(S)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \\
& \quad+\alpha_{v} \sum_{w \in \operatorname{children}_{T}(v)-S}\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \\
= & \alpha_{v}\left[\varphi_{T}\left(\operatorname{desc}_{T}(S)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \\
& +\alpha_{v}\left[\left(\varphi_{T}\left(\operatorname{desc}_{T}(v)\right)-\varphi_{T}\left(\operatorname{desc}_{T}(S)\right)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right], \\
& =\alpha_{v} .
\end{aligned}
$$

(b) We proceed by induction on the height of $v$ :

Base: $v \in$ leaves $_{T}$.,

Then the only $w$ in $\operatorname{desc}_{T}(v)$ is $v$ itself, so the sum on the right is just $\left[\varphi_{T}(v) / \varphi_{T}(v)\right] \alpha_{v},=\alpha_{v}$, as needed.

Inductive step. $v \in$ internal $_{T}$. Then

$$
\begin{aligned}
\alpha_{v}= & \sum_{w \in \operatorname{children}_{T}(v)}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{w}\right] \quad \text { by part }(\mathrm{a}), \\
= & \sum_{w \in \operatorname{children}_{T}(v)}\left[\left[\varphi_{T}\left(\operatorname{desc}_{T}(w)\right) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right]\right. \\
& \left.\times \sum_{s \in \operatorname{dessc}_{T}(w)}\left[\left[\varphi_{T}(s) / \varphi_{T}\left(\operatorname{desc}_{T}(w)\right)\right] \alpha_{s}\right]\right], \quad \text { by inductive hypothesis, } \\
= & \sum_{w \in \operatorname{chidrren}_{T}(v)} \sum_{s \in \operatorname{desc}_{T}(w)}\left[\left[\varphi_{T}(s) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{s}\right], \\
= & \sum_{s \in \operatorname{desc}_{T}(v)}\left[\left[\varphi_{T}(s) / \varphi_{T}\left(\operatorname{desc}_{T}(v)\right)\right] \alpha_{s}\right]
\end{aligned}
$$

as needed.
Part (b) of this lemma, with $v=\operatorname{root}_{T}$, proves equivalence of the two experiments described prior to the lemma.

We can restrict attention to "request sequences" in place of request patterns, in the sequential case of the algorithm. Assume that $T$ is a weighted tree, and $p$ is a placement for $T$. A request sequence $r$ is a sequence of elements of vertices ${ }_{T}$, representing a sequence of request arrival locations. Similarly, a resource sequence $r$ is a sequence of elements of vertices $_{T}$, representing a sequence of resource locations. In either case, let length( $r$ ) denote the number of elements in a sequence. A resource sequence, $s$, is compatible with a placement $p$ provided that $\left|s^{-1}(v)\right| \leqslant p(v)$ for each $v \in$ vertices $_{T}$. (That is, the resource sequence grants at most the number of resources placed at each vertex.) If $r$ is a request sequence and $p$ is a placement with total $(p) \geqslant$ length $(r)$, then a matching of $r$ and $p$ is a pair $m=(r, s)$, where $s$ is a resource sequence compatible with $p$ and length $(r)=$ length $(s)$. A matching describes the successive locations of resources which are used to satisfy a sequence of requests.

Next, we define a probabilistic program which takes as inputs a request sequence $r$ and a placement $p$ with $\operatorname{total}(p) \geqslant \operatorname{length}(r)$, and returns a matching of $r$ and $p$. The procedure simply uses FIND repeatedly.

```
    MATCH(T,p,r)
For i = 1 to length(r) do
    [s(i) := FIND(T,p,r(i))
    p(s(i)) := p(s(i)) - 1]
```

Theorem 14. Let $r$ be a request sequence, $p$ a placement with total $(p) \geqslant$ length $(r)$. Then $\operatorname{MATCH}(T, p, r)$ will eventually halt and return $a$ matching of $r$ and $p$.

## Proof. Straightforward.

This algorithm is designed to behave exactly as the sequential case of our general algorithm.

### 5.2. Cost Measures

Let $\operatorname{dist}_{T}(u, v)$ denote the tree distance between $u$ and $v$. If $m=(r, s)$ is a matching, then seqcost $T_{T, p}(m)=\sum_{i} \operatorname{dist}_{T}(r(i), s(i))$. Thus, the "sequential cost" is just the sum of the tree distances between successive requests and their corresponding resources.

If $r$ is a request sequence with length $(r) \leqslant \operatorname{total}(p)$, then define seqcost $_{T, p}(r)$ to be the expected value of $\operatorname{seq}^{\operatorname{cost}}{ }_{T, p}(m)$, where $m$ is constructed using MATCH $(T, p, r)$. Let $\operatorname{seq}^{\operatorname{cost}}{ }_{T, p}$ denote the expected value of seqcost ${ }_{T, p}(r)$, where $r$ is of length total $(p)$, with its successive locations chosen independently according to $\varphi_{T}$. In the remainder of this section, we analyze seqcost ${ }_{T, p}$.

### 5.3. A Useful Recurrence

In this subsection, we present a solution to a system of recurrence equations. This solution will be useful in later subsections. Let $c \in R^{+}$. Define $G_{c}: N \times R^{+} \rightarrow R^{+}$by the equations:
$G_{c}(0, t)=0 \quad$ and
$G_{c}(k, t)=\max \left\{G_{c}\left(k-1, t_{1}\right)+G_{c}\left(k-1, t_{2}\right): t_{1}+t_{2} \leqslant t\right\}+c k \sqrt{t}$, for $k \geqslant 1$.
Lemma 15. For all $k \geqslant 0$, the following are true:
(a) The function mapping to $G_{c}(k, t)$ is concave downward and monotone nondecreasing.
(b) If $k \geqslant 1$ then $G_{c}(k, t)=2 G_{c}(k-1, t / 2)+c k \sqrt{t}$.

Proof. We proceed by induction on $k$. The base, $k=0$, is trivial. For the inductive step, let $k \geqslant 1$. If $t_{1}+t_{2} \leqslant t$, then $G_{c}\left(k-1, t_{1}\right)+$ $G_{c}\left(k-1, t_{2}\right) \leqslant 2 G_{c}\left(k-1,\left(t_{1}+t_{2}\right) / 2\right)$, since the inductive hypothesis states that the function mapping $t$ to $G_{c}(k-1, t)$ is concave. This is in turn $\leqslant 2 G_{c}(k-1, t / 2)$, by monotonicity. Therefore, $\quad G_{c}(k, t)=$ $2 G_{c}(k-1, t / 2)+c k \sqrt{t}$, showing (b). Since each term is concave and monotone, the sum is also, showing (a).

Theorem 16. $\quad G_{c}(k, t) \leqslant(3 \sqrt{2}+4) c \sqrt{t 2^{k}}$.

Proof. By Lemma 15, $G_{c}(k, t)=0$ if $k=0$, and $=2 G_{c}(k-1, t / 2)+$ $c k \sqrt{t}$ if $k \geqslant 1$. Expanding this recurrence, we see that

$$
\begin{aligned}
G_{c}(k, t) & =c\left[\sum_{i=0, \ldots, k-1} 2^{i}(k-i) \sqrt{t / 2^{i}}\right] \quad \text { for all } k \geqslant 0 \\
& =c \sqrt{t}\left[\sum_{i=0, \ldots, k-1}(k-i) \sqrt{2^{i}}\right]
\end{aligned}
$$

Let $x=1 / \sqrt{2}, n=2^{k}$. Then

$$
\begin{aligned}
G_{c}(k, t) & =(c \sqrt{t} \sqrt{n} / \sqrt{2})\left[\sum_{i=0, \ldots, k-1}(k-i) x^{k-i-1}\right], \\
& =(c \sqrt{t} \sqrt{n} / \sqrt{2})\left[\left(1+k x^{k+1}-(k+1) x^{k}\right) /(1-x)^{2}\right], \\
& =(c \sqrt{t} \sqrt{n}) /\left(\sqrt{2}(1-1 / \sqrt{2})^{2}\left[1+k x^{k+1}-(k+1) x^{k}\right],\right. \\
& =(c \sqrt{t} \sqrt{n})(3 \sqrt{2}+4)\left[1+k x^{k+1}-(k+1) x^{k}\right], \\
& =(c \sqrt{t} \sqrt{n})(3 \sqrt{2}+4)\left[1+(k x-k-1) x^{k}\right], \\
& \leqslant(c \sqrt{t} \sqrt{n})(3 \sqrt{2}+4) \quad \text { since } k x-k-1<0, \\
& =\left(c \sqrt{t 2^{k}}\right)(3 \sqrt{2}+4) .
\end{aligned}
$$

### 5.4. Recursive Analysis

Now, we require Restriction 3 and a new assumption, Restriction 4. These are to remain in force for the remainder of Section 5.

Restriction 3. Thas all leaves at the same distance from the root, and $r$ and $p$ are nonzero only at leaves.

Restriction 4. $\quad T$ is a complete binary tree.
If $T$ is a weighted tree, then a weighted subtree $T^{\prime}$ of $T$ consists of a subtree of $T$ together with a probability function $\varphi_{T^{\prime}}$ given by $\varphi_{T^{\prime}}(v)=\varphi_{T}(v) / \varphi_{T}\left(\right.$ vertices $\left._{T^{\prime}}\right)$. That is, the weights of the subtrees are just normalized restrictions of the weights of $T$. If $T$ is a weighted binary tree, let $l e f t_{T}$ and $\operatorname{right}_{T}$ denote the designated weighted subtrees of $T$.

If $T$ has height at least 1 , then let $T_{1}$ and $T_{2}$ denote left ${ }_{T}$ and right ${ }_{T}$, respectively. Let $p_{1}$ and $p_{2}$ denote $p \mid T_{1}$ and $p \mid T_{2}$, respectively. If $r$ is a request sequence, let $\operatorname{overflow}_{T, p}(r)$ denote $\| r^{-1}\left(\operatorname{vertices}_{r_{1}}\right) \mid-$ $p\left(\right.$ vertices $\left._{T_{1}}\right) \mid$, the difference between the number of requests that arrive in the left subtree and the number of resources placed there. Let overflow ${ }_{T, p}$ denote the expected value of overflow ${ }_{T, p}(r)$, where $r$ is a sequence of length $\operatorname{total}(p)$ chosen using $\varphi_{T}$.

The following is a key lemma which provides a recursive breakdown for the sequential cost. It says that the expected cost of matching breaks down into costs of matching within the two subtrees, plus a charge for the requests that overflow into the opposite subtree.

Lemma 17. Assume height $_{T} \geqslant 1$. Then $\operatorname{seq}^{\operatorname{cost}}{ }_{T, p} \leqslant \operatorname{seq}^{\operatorname{cost}}{ }_{T_{1}, p_{1}}+$ seqcost $_{T_{2}, p_{2}}+2$ height $_{T}$ overflow $_{T, p}$.

Proof. For any particular request sequence $r$, there is some particular number, overflow $_{T, p}(r)$, of requests that do not get satisfied within their own subtree, but rather overflow into the opposite subtree to find a resource. To be specific, assume that it is the left subtree from which any excess requests overflow. Let $r_{1}$ be the subsequence of $r$ consisting of requests arriving in $T_{1}$, truncated to length $p\left(T_{1}\right)$. Let $r_{2}$ be the subsequence of $r$ consisting of requests arriving in $T_{2}$. Recall that seqcost ${ }_{T, p}$ is the expectation of the search cost for enough requests to exhaust all resources present.

Before the time at which the left subtree overflows, the algorithm $\operatorname{MATCH}(T, p, r)$ runs exactly like $\operatorname{MATCH}\left(T_{1}, p_{1}, r_{1}\right)$ within the left subtree. Requests originating within $T_{1}$ become matched to exactly the same resources in both executions.

We now consider the right subtree. Requests which originate within the right subtree are handled in the algorithm $\operatorname{MATCH}(T, p, r)$ exactly as they are in the algorithm for $T_{2}$ and $p_{2}$. However, there are also overflow requests from the left subtree, which enter $T_{2}$ at its root rather than at its leaves. By Lemma 13, whenever such a request arrives, its probability of being matched to any particular resource is exactly the same as if the request had entered at a random leaf of $T_{2}$. All the requests remain independent, and these additional requests are just enough to exhaust the resources in $T_{2}$.

Now assume that the request sequences $r$ are chosen at random according to $\varphi_{T}$. They result in subsequences $r_{1}$ and $r_{2}$ which are chosen at random according to $\varphi_{T_{1}}$ and $\varphi_{T_{2}}$, respectively.

We claim that seqcost ${ }_{T, p}$, the expected cost taken over all $r$, is bounded by
seqcost $_{T_{1, p},}$, the expected cost taken over random sequences in the left subtree + seqcost $_{T_{2, p}}$, the expected cost taken over random sequences in the right subtree +2 height $_{T}$ overflow $_{T, p}$.

The third term allows for the expected overflow of requests, and assigns them the maximum cost, 2 height ${ }_{T}$.

Consider the first term. (The second term is analogous.) The first term allows for the expected cost incurred by an execution of MATCH on a ran-
dom sequence of requests within $T_{1}$. In case $T_{1}$ has its resources exhausted by requests originating within that subtree, this term measures exactly the expected cost for the matching of the first requests in $T_{1}$ to all the resources in $T_{1}$. This term ignores the cost incurred by any excess requests originating within $T_{1}$ which do not get matched within $T_{1}$. However, that is not a problem, since those costs are counted by the third term.

In case $T_{1}$ does not have its resources exhausted by requests originating within $T_{1}$, this term is actually greater than needed to measure the expected cost of matching all requests originating in $T_{1}$; in fact, it is enough to measure this cost of matching these requests, interspersed with enough other random requests (arriving at the leaves of $T_{1}$ ) to use up all the resources in $T_{1}$. We have already argued that these requests behave as if they were interspersed with other random requests, because requests arriving at the root match in the same way as if they arrived at random leaves. In this case, the first term does not account for the cost of matching those requests which enter $T_{1}$ at its root. However, that is not a problem since that cost is covered by the third term.

## 5.5. $d$-Fairness

We need to make another restriction on the algorithm, for the purpose of analysis. In particular, it is reasonable that the behavior of the algorithm should be best when the resources are distributed in the tree in some relationship with the probability distribution governing arrival of requests. (The paper by Fischer, Griffeth, Guibas, and Lynch, 1981, considers optimal placements of resources in a tree network.)

For $d \in R^{+}$, we say that a placement, $p$, is $d$-fair provided that the following is true. Let $u, v \in$ vertices $_{T}$, where $u \in \operatorname{desc}_{T}(v)$. Let $\varphi_{1}=\varphi_{T}\left(\operatorname{desc}_{T}(v)\right) \quad$ and $\quad \varphi_{2}=\varphi_{T}\left(\operatorname{desc}_{T}(u)\right)$. Then $\mid p\left(\operatorname{desc}_{T}(u)\right)-$ $\left(\varphi_{2} / \varphi_{1}\right) p\left(\operatorname{desc}_{T}(v)\right) \mid \leqslant d \sqrt{\left(\varphi_{2} / \varphi_{1}\right) p\left(\operatorname{desc}_{T}(v)\right)}$. That is, for each subtree, the number of resources placed in each of its subtrees is approximately proportional to the probability of arrivals in that subtree, and the difference is proportional to the "standard deviation." For any $T$, if $t$ and $d$ are sufficiently large, then techniques in (Fischer, Griffeth, Guibas, and Lynch, 1981) can be used to produce $d$-fair placements of $t$ resources in $T$.

From now on in Section 5, we assume the following:
Restriction 5. $p$ is d-fair (for some arbitrary but fixed d).
The next lemma says that restrictions of $d$-fair placements are also $d$-fair.
Lemma 18. Let $p$ be a d-fair placement for $T$. Let $T^{\prime \prime}$ be any subtree of $T$, and $p^{\prime}=p \mid$ vertices $_{T^{\prime}}$, the restriction of $p$ to vertices $_{T^{\prime}}$. Then $p^{\prime}$ is $d$-fair for $T^{\prime}$

Proof. Let $u, v \in \operatorname{vertices}_{T^{\prime}}$, with $u \in \operatorname{desc}_{T^{\prime}}(v)$. Let $\varphi_{1}=\varphi_{T}\left(\operatorname{desc}_{T}(v)\right)$, $\varphi_{2}=\varphi_{T}\left(\operatorname{desc}_{T}(u)\right), \varphi_{1}^{\prime}=\varphi_{T^{\prime}}\left(\operatorname{desc}_{T^{\prime}}(v)\right)$, and $\varphi_{2}^{\prime}=\varphi_{T^{\prime}}\left(\operatorname{desc}_{T}(u)\right)$. Then $\varphi_{1}^{\prime}=$ $\varphi_{1} / \varphi_{T}\left(\right.$ vertices $\left._{T^{\prime}}\right)$ and $\varphi_{2}^{\prime}=\varphi_{2} / \varphi_{T}\left(\right.$ vertices $\left._{T^{\prime}}\right)$, by definition. Therefore, $\varphi^{\prime}(2) / \varphi^{\prime}(1)=\varphi_{2} / \varphi_{1}$.
Note that $p^{\prime}\left(\operatorname{desc}_{T}(u)\right)=p\left(\operatorname{desc}_{T}(u)\right)$, and $p^{\prime}\left(\operatorname{desc}_{T}(v)\right)=p\left(\operatorname{desc}_{T}(v)\right)$. Thus,

$$
\begin{aligned}
&\left|p^{\prime}\left(\operatorname{desc}_{T^{\prime}}(u)\right)-\left(\varphi_{2}^{\prime} / \varphi_{1}^{\prime}\right) p^{\prime}\left(\operatorname{desc}_{T}(v)\right)\right| \\
&=\left|p\left(\operatorname{desc}_{T}(u)\right)-\left(\varphi_{2} / \varphi_{1}\right) p\left(\operatorname{desc}_{T}(v)\right)\right|, \\
& \leqslant d \sqrt{\left(\varphi_{2} / \varphi_{1}\right) p\left(\operatorname{desc}_{T}(v)\right)} \quad \text { since } p \text { is } d \text {-fair, } \\
&=d \sqrt{\left(\varphi_{2}^{\prime} / \varphi_{1}^{\prime}\right) p^{\prime}\left(\operatorname{desc}_{T^{\prime}}(v)\right)} \quad \text { as needed. }
\end{aligned}
$$

The final lemma of this subsection bounds the expected overflow for $d$ fair placements.

Lemma 19. overflow $T_{T, p} \leqslant(6+d) \sqrt{\varphi_{T}\left(\text { vertices }_{T_{1}}\right) \operatorname{total}(p)}$.
Proof. $\left\|r^{-1}\left(\operatorname{vertices}_{T_{1}}\right)\right\|-p\left(\right.$ vertices $\left._{T_{1}}\right)\|\leqslant\| r^{-1}\left(\operatorname{vertices}_{T_{1}}\right) \|-$ $\varphi_{T}\left(\operatorname{vertices}_{T_{1}}\right) \operatorname{total}(p)\left|+\left|\varphi_{T}\left(\operatorname{vertices}_{T_{1}}\right) \operatorname{total}(p)\right|-p\left(\operatorname{vertices}_{T_{1}}\right)\right|$.

The expected value of the first of these quantities, taken over $r$, is bounded by $6 \sqrt{\varphi_{T}\left(\text { vertices }_{T_{1}}\right) \operatorname{total}(p) \text {, using Lemma 3.1.5 of (Fischer, }}$ Griffeth, Guibas, and Lynch, 1981).

The second quantity is bounded by $d \sqrt{\varphi_{T}\left(\text { vertices }_{T_{1}}\right) \text { total }(p)}$, since $p$ is $d$-fair.

### 5.6. Sequential Analysis

Let $\operatorname{size}_{T}$ denote the number of vertices in $T$. We now give the main result of Section 5, that seqcost ${ }_{T, p}$ is $O\left(\sqrt{\operatorname{size}_{T} \operatorname{total}(p)}\right)$. This says, for example, that if $\operatorname{total}(p)$ is proportional to $\operatorname{size}_{T}$, then the total cost is proportional to total $(p)$. This implies that the average cost per request is just a constant, independent of the size of the network.

In order to prove this theorem, we require the followng restrictions, repeated here for convenience.

Restriction 3. Thas all leaves at the same distance from the root, and $r$ and $p$ are nonzero only at leaves.

Restriction 4. $\quad T$ is a complete binary tree.
Theorem 20. seqcost ${ }_{T, p}$ is $O \sqrt{\left.\operatorname{size}_{T} \operatorname{total}(p)\right)}$. (More specifically, seqcost $_{T, p} \leqslant(3 \sqrt{2}+4)((6+d) \sqrt{2}) \sqrt{\left.2^{\text {height }_{T}} \operatorname{total}(p) .\right)}$

Proof. By Theorem 16, it suffices to show that seqcost $T_{T, p} \leqslant$ $G_{c}\left(\right.$ height $\left._{T}, \operatorname{total}(p)\right)$, where $c=(6+d) \sqrt{2}$. We proceed by induction on height $_{T}$.

Base: height $_{r}=0$
Then $T$ has a single vertex, and seqcost $T_{T, p}=0$. The inequality is immediate.
Inductive step. height $_{T} \geqslant 1$. Then

$$
\begin{aligned}
& \operatorname{seq}^{\operatorname{cost}}{ }_{T, p} \\
& \leqslant \operatorname{seq}^{\operatorname{cost}_{T_{1}, p_{1}}}{ } \text { seqcost }_{T_{2}, p_{2}}+2 \text { height }_{T} \text { overflow }_{T, p^{\prime}} \\
& \text { by Lemma } 17 \text {, } \\
& \leqslant \operatorname{seq}^{\operatorname{cost}_{T_{1}, p_{1}}}+\operatorname{seq}^{\operatorname{cost}_{T_{2}, p_{2}}}+2 \operatorname{height}_{T}(6+d) \sqrt{\varphi_{T}\left(\operatorname{vertices}_{T_{1}}\right) \operatorname{total}(p)} \\
& \text { by Lemma } 19 .
\end{aligned}
$$

A similar inequality holds for $T_{2}$ in place of $T_{1}$ within the square root. Since at least one of $\varphi_{T}\left(\right.$ vertices $\left._{T_{1}}\right), \varphi_{T}\left(\right.$ vertices $\left._{T_{2}}\right)$ is no more than $\frac{1}{2}$, it follows that

$$
\begin{aligned}
& \operatorname{seq}^{\operatorname{cost}}{ }_{T, p} \leqslant \operatorname{seq}^{\operatorname{cost}}{ }_{T_{l}, p_{1}}+\operatorname{seq}^{\operatorname{cost}}{ }_{T_{2}, p_{2}}+2 \operatorname{height}_{T}(6+d) \sqrt{(1 / 2) \text { total }(p)}, \\
& =\operatorname{seq}^{\operatorname{cost}}{ }_{T_{1}, p_{1}}+\operatorname{seq}^{\operatorname{cost}}{ }_{T_{2}, p_{2}}+(6+d) \sqrt{2} \text { height }_{T} \sqrt{\operatorname{total}(p)} .
\end{aligned}
$$

By Lemma 18, we can apply the inductive hypothesis, which implies that the right-hand side of this inequality is at most equal to

$$
\begin{aligned}
& G_{c}\left(\text { height }_{T}-1, \operatorname{total}\left(p_{1}\right)\right)+G_{c}\left(\text { height }_{T}-1, \operatorname{total}\left(p_{2}\right)\right) \\
& \quad+(6+d) \sqrt{2} \operatorname{height}_{T} \sqrt{\operatorname{total}(p)}
\end{aligned}
$$

The definition of $G_{c}$ implies that this latter expression is at most equal to $G_{c}\left(\right.$ height $_{T}, \operatorname{total}(p)$ ), as needed.

## 6. The Final Analysis

In this section, we combine the monotonicity analysis and the sequential analysis, to obtain an upper bound for the expected cost for the algorithm.

### 6.1. Relationship Between the Costs

Now we require the following restrictions.
Restriction 1. $T$ is a weighted tree, and the nondeterministic choice step in Part (2) of the algorithm uses a call to random ${ }_{r, s}$.

Restriction 2. Delivery time for messages is always exactly 1.
With these restrictions, there is a close relationship between the costs of our general algorithm and the cost of the sequential algorithm MATCH.

Lemma 21. seqcost $_{T, p}=\lim _{a \rightarrow \infty}\left(\operatorname{bnum}_{T, p}(a, f)+\operatorname{vbnum}_{T, p}(a, f)\right)$.
Proof. There is an absolute upper bound on the time for our algorithm to satisfy a single request, in the absence of concurrent requests. Thus, as $a$ increases, the probability that there are any concurrent requests approaches 0 .

Therefore, the limiting case of the general algorithm behaves like MATCH. There is no backtracking, so the total search time just reduces to the sum of the distances from requests to the resources which satisfy them. This sum is just the total number of buyer and virtual buyer messages.

Now let us add one more restriction:
Restriction 3. Thas all leaves at the same distance from the root, and r and $p$ are nonzero only at leaves.
With this added restriction, we can prove a variant of the preceding lemma.

Lemma 22. seqcost $_{T, p}=2 \lim _{a \rightarrow \infty}\left(\operatorname{bnum}_{T, p}(a, f)\right)$.
Proof. By Lemmas 21 and 4.
This immediately implies the following bound on the cost of the general algorithm.

Lemma 23. $\operatorname{cost}_{T, p}(f) \leqslant 2 \operatorname{seq}_{\operatorname{cost}}^{T, p}$.
Proof. By Theorem 11 and Lemma 22.

### 6.2. The Main Theorem

Now we are ready to present the main result, the upper bound for the expected cost for the general algorithm. In order to apply the results of both the monotonicity analysis and the sequential analysis, we must assume the restrictions made for both cases. More specifically, we assume all of Restrictions 1-5:

Restriction 1. Tis a weighted tree, and the nondeterministic choice step in Part (2) of the algorithm uses a call to random ${ }_{T, S}$.

Restriction 2. Delivery time for messages is always exactly 1.

Restriction 3. Thas all leaves at the same distance from the root, and $r$ and $p$ are nonzero only at leaves.

Restriction 4. $T$ is a complete binary tree.
Restriction 5. $p$ is d-fair (for some arbitrary but fixed d).
Theorem 24. Let $f$ be a probability function. Then $\operatorname{cost}_{T, p}(f)$ is $O\left(\sqrt{\left.\operatorname{size}_{T} \text { total }(p)\right)} \quad\right.$ (More $\quad$ specifically, $\quad \operatorname{cost}_{T, p}(f) \leqslant 2(3 \sqrt{2}+4)$ $((6+d) \sqrt{2}) \sqrt{\left.2^{\text {height }_{T}} \operatorname{total}(p) .\right)}$

Proof. By Lemma 23 and Theorem 20.
In particular, provided that $\operatorname{total}(p)$ is proportional to $\operatorname{size}_{T}$, the expected average time taken by this algorithm to satisfy a single request is constant, independent of the size of the network.

Remark. It is possible to prove a variant of Theorem 24, for the case in which the placement $p$ is chosen at random using $\varphi_{T}$ (just as the request locations are chosen), rather than being $d$-fair. We sketch the ideas briefly.

First, we must extend the cost definitions to include expectations taken over placements of a particular length. Thus, we define $\operatorname{cost}_{T, t}(f)$ to be the expected value of $\operatorname{cost}_{T, p}(f)$ for $p$ with total $(p)=t$. Analogous definitions are made for seqcost $T_{T, t}$ and overflow ${ }_{T, 1}$. Lemma 23 then implies that $\operatorname{cost}_{T, t}(f) \leqslant 2 \operatorname{seq}^{(0 s t} t_{T, t}$. It is also easy to see that overflow ${ }_{T, t} \leqslant$ $d \sqrt{\varphi_{T}\left(\text { vertices }_{T_{1}}\right)} t$, for some constant $d$.

Next, we prove a consequence of Lemma 17 which says that seqcost $t_{T, t} \leqslant$ $\operatorname{Exp}_{\left(t_{1}, t_{2}\right){\text { with } t_{1}+t_{2}=t}}$ seqcost $\left._{T_{1}, t_{1}}+\operatorname{seq}^{\operatorname{cost}} T_{T_{2}, t_{2}}\right)+2$ height $_{T}$ overflow $_{T, 4}$. (Here, the expectation is taken over pairs which are obtained by using $\varphi_{T}$ to assign resources to $T_{1}$ or $T_{2}$.) This obviously implies that seqcost $T_{T, t} \leqslant$ $\max _{\left(t_{1}, t_{2}\right) \text { with } t_{1}+t_{2}=t}\left(\right.$ seqcost $_{T_{1}, t_{1}}+$ seqcost $\left._{T_{2}, t_{2}}\right)+2$ height $_{T}$ overflow $_{T, t}$.

Now we prove a variant of Theorem 20 which says that seqcost ${ }_{T, t}$ is $O\left(\sqrt{\text { size }_{T} t}\right)$. More specifically, we show that $\operatorname{seq}^{\operatorname{cost}} T_{T, t} \leqslant G_{c}$ (height $\left.{ }_{T}, t\right)$, where $c=d \sqrt{2}$. This is easily done by induction as before, using the new lemmas just described. Combining these results, we see that $\operatorname{cost}_{T, t}(f)$ is $O\left(\sqrt{\operatorname{size}_{T} t}\right)$.

## 7. Remaining Questions

There are several directions for remaining research. First, we would like to extend the analysis of the general algorithm. We would like to loosen our restrictions on tree shape, message delivery time, and locations for resources and requests. If we do this, is it possible to carry out an analysis
similar to the one in this paper? In particular, can the concurrent cases of the algorithm still be bounded in some way in terms of the sequential case?

We would also like to extend the analysis of the sequential algorithm MATCH. Here, we would like to loosen restrictions on tree shape and on locations for resources and requests. There are some apparent improvements in the algorithm, for example, adjusting the probabilities for the choice among children in response to knowledge of the number of resources remaining in each subtree. While this seems like an improvement, the resulting algorithm seems harder to analyze (since the recursive decomposition does not appear to work). Can any simple modifications be shown to be improvements?

We would like to compare the performance of this algorithm to that of alternative algorithms which solve the same problem. We have already observed that this algorithm performs much better than the centralized algorithm, which locates all resources at the root. How does it compare to algorithms which allow requests to search for resources in parallel rather than sequentially? What about algorithms which rebalance resources? Are there other interesting ideas for algorithms for this problem?

Finally, the general analysis strategy is quite attractive. Proving a monotonicity result which bounds the concurrent cases of an algorithm interms of the sequential case, and then analyzing the sequential case by traditional techniques, appears quite tractable. The use of this strategy for our algorithm appears to depend on many special properties of the algorithm and on restrictions on the execution. Is the strategy more generally useful? For what type of algorithms can it be used?

Received December 13, 1982; accepted September 19, 1983

## References

[^1]
[^0]:    * This work was supported in part by the Office of Naval Research Contract N00014-82-K-0154, by the Office of Army Research Contracts DAAG29-79-C-0155 and DAAG29-84K-0058, by the Army Institute Research in Management Information and Computer Systems Contract DAAK 70-79-D-0087, by the National Science Foundation Grants MCS-7924370, MCS-8116678, MCS-8200854, MCS-8200854, MCS-8306854, and DCR-8302391, and by the Defense Advanced Research Projects Agency (DARPA) Grant N00014-83-K-0125.

[^1]:    Fischer, M., Griffeth, N., Guibas, L., and Linch, N. (1981), Optimal placement of identical resources in a distributed network, in "Proc. 2nd International Conference on Distributed Computing"; Inform. and Control,
    Guibas, L. J., and Liang, F. M. (1982), Systolic stacks, queues, and counters, in "Proc. 2nd MIT VLSI Conference."

