# Electing a Leader in a Synchronous Ring 

GREG N. FREDERICKSON<br>Purdue University, West Lafayette, Indiana

AND<br>NANCY A. LYNCH<br>Massachusetts Institute of Technology, Cambridge, Massachusetts


#### Abstract

The problem of electing a leader in a synchronous ring of $n$ processors is considered. Both positive and negative results are obtained. On the one hand, if processor IDs are chosen from some countable set, then there is an algorithm that uses only $O(n)$ messages in the worst case. On the other hand, any algorithm that is restricted to use only comparisons of IDs requires $\Omega(n \log n)$ messages in the worst case. Alternatively, if the number of rounds is required to be bounded by some $t$ in the worst case, and IDs are chosen from any set having at least $f(n, t)$ elements, for a certain very fast-growing function $f$, then any algorithm requires $\Omega(n \log n)$ messages in the worst case.


Categories and Subject Descriptors: C.2.5 [Computer-Communications Networks]: Local Networksrings; F.1.1 [Computation by Abstract Devices]: Models of Computation-automata; unbounded-action devices. F.1.2 [Computation by Abstract Devices]: Modes of Computation-synchronous computation; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Theory, Validation
Additional Key Words and Phrases: Comparison algorithms, distributed algorithms, leader election, lower bounds, Ramsey's theorem, ring networks, synchronous computation

## 1. Introduction

Communication in a network can be performed in either a synchronous or an asynchronous mode. How does the choice of communication mode affect the computational resources required to solve a problem? We examine this question by considering the problem of electing a leader in a ring-shaped network. In this problem there are $n$ processors, which are identical except that each has its own unique identifier. At various points in time, one or more of the processors independently initiate their participation in an election to decide on a leader. The relevant resources for such a distributed computation are the total number of

[^0](c) 1987 ACM 0004-5411/87/0100-0098 \$00.75
messages used and the amount of time expended from the time that the first processor wakes up.

The problem of electing a leader efficiently has been studied by a number of researchers [1, 4-6, 8-11, 13]. The best previous deterministic algorithms have used $O(n \log n)$ messages for either bidirectional rings [4, 8,9] or unidirectional rings [6, 13]. These algorithms work for both the synchronous and asynchronous models, and use comparisons of IDs only. In addition, Burns has established a lower bound of $\Omega(n \log n)$ on the number of messages required if communication is asynchronous [4]. However, the proof in [4] does not extend to the case of synchronous communication. It is, therefore, quite natural to ask whether the $\Omega(n$ $\log n$ ) lower bound can be achieved in the synchronous case as well as in the asynchronous, or whether there are algorithms that somehow make use of the synchrony to limit the number of messages transmitted.

We obtain both positive and negative answers to our question of whether synchrony helps. On the one hand, we show that if processor IDs are chosen from some countable set (such as the integers), then there is an algorithm that uses only $O(n)$ messages in the worst case. The processors may initiate the algorithm at different rounds, and do not know the value of $n$. Our algorithm is thus an improvement on a probabilistic algorithm of [10] that uses $O(n)$ messages on average and assumes that the processors do know the value $n$. Unlike the earlier algorithms, our algorithm uses not only comparisons on IDs, but also the numerical value of the IDs to count rounds. However, the number of synchronous rounds used by our algorithm can be very large in the worst case. An algorithm similar to ours has been developed independently by Vitanyi [15].

On the other hand, we show that both the departure from the comparison model and the possibility of using a large number of rounds are necessary in order to obtain an algorithm of linear message complexity. More specifically, if the algorithm is restricted to use only comparisons of IDs, then we obtain an $\Omega(n \log n)$ lower bound for the number of messages required in the worst case. To achieve this bound we generate an assignment of IDs to processors that exhibits a large amount of "replication symmetry" around the ring. We give a relatively simple assignment of values if $n$ is a power of 2 , and a somewhat more involved assignment for general values of $n$. (More recently, a different assignment of IDs has been given in [2].)

Alternatively, if the number of rounds is required to be bounded by some $t$ in the worst case, then there is a (very fast-growing) function $f(n, t)$, which has the following very interesting property. If IDs are chosen from any set $T$ having at least $f(n, t)$ elements, then any $t$-bounded algorithm requires $\Omega(n \log n)$ messages in the worst case. In particular, if $t$ is a function of $n$, say $t(n)$, then any $t(n)$-bounded algorithm for a set $T$ with at least $f(n, t(n))$ elements exhibits the given lower bound on messages. We achieve this result by giving a transformation from any algorithm in what we call free form, over such a set $T$, to a comparison-based algorithm. The ideas for this transformation are derived from earlier work of Snir [14]. Both of our lower bound results hold even in the case that the number of processors in the ring is known to each processor, and all the processors are known to start at the same round.

## 2. The Algorithm

In this section we present an algorithm for electing a leader in a synchronous ring. The algorithm uses only $O(n)$ messages but may require a very large number of
rounds. The elected processor, and only this processor, eventually enters one of a set of distinguished "elected" states. The total number of messages used, including any messages that might be sent after the winner is elected, is $O(n)$. The algorithm presented is for a unidirectional ring, with communication assumed to be counterclockwise. Of course, essentially the same algorithm will work on a bidirectional ring. We assume that the unique ID of each processor is an integer. This assumption is reasonable if communication is implemented by transmitting packets of bits. In the description of the algorithm, we shall refer to the processor with ID $i$ as processor $i$.

The algorithm is initiated by individual processors deciding independently to wake up. The processors need not wake up at the same time, but no processor is allowed to wake after it has received a message from an awakened processor. When it wakes up, a processor (henceforth called a participating processor) spawns a message process, which moves around the ring, carrying the ID of the originating processor. The message process is charged one message for each edge that it traverses.

Our algorithm uses two ideas. The first is that message processes that originate at different processors are transmitted at different rates: The message process carrying processor ID $i$ travels at the rate of one message transmission every $2^{i}$ rounds. (Specifically, each processor delays for $2^{i}-1$ rounds before transmitting message process $i$.) Any slower message process that is overtaken by a faster message process is killed. Also, a message process carrying ID $i$ arriving at processor $j$ is killed if $j<i$ and processor $j$ has also spawned a message process. A message process that returns to its originator causes the originator to become elected.

Suppose that all participating processors were to wake up at the same round. The above strategy would then guarantee that the total number of messages is $O(n)$. To see this, consider the following. Let $i$ be the smallest ID of any participating processor. Message process $i$ traverses all edges, for a total cost of $n$. Consider any other message process, $j$. During message process $i$ 's circuit, either message process $i$ overtakes message process $j$, or else message process $j$ reaches processor $i$. In either case message process $j$ is killed by the time $i$ 's circuit is completed. Because of the different rates of travel, message process $j$ could travel at most distance $n /\left(2^{j-i}\right)$ during the time that message process $i$ travels distance $n$. Summing over all message processes, the total number of messages expended would be less than $2 n$.

However, this variable rate of transmission scheme is by itself not enough to realize $O(n)$ messages in the case in which not all participating processors wake up at the same time. The processors with smaller IDs could wake up correspondingly later and spawn message processes that would chase and ultimately overtake the slower message processes, but not before $\Omega(n)$ messages had been expended by each of $\Omega(n)$ message processes.

The second idea is to have a preliminary phase for each message process, before the variable rate phase begins. In this phase, all message processes travel at the same rate, one message transmission per round. When a processor decides it wants to participate, it spawns its message process and sends it off to its neighbor. The message process is transmitted around the ring until it encounters the next participating processor. At this point, the message process continues into the second phase, moving at its variable rate and acting as previously described.

Lemma 1. There is an algorithm that elects a leader in a synchronous ring of $n$ processors using fewer than $4 n$ messages and $O\left(n 2^{\prime}\right)$ time, where $i$ is the ID of the eventual winner.

Proof. We divide the messages into three categories, and bound each category separately. The categories are (1) the first-phase messages, (2) the second-phase messages sent before the eventual winner enters its second phase, and (3) the second-phase messages sent after the eventual winner enters its second phase.

First consider (1). Since exactly one message from the first phase will be transmitted along each edge, the total number of first-phase messages is exactly $n$. Next, consider (2). Every message process that is activated will enter its second phase within $n$ rounds of the time at which the first of the processors awakens. Thus, at most $n$ rounds need to be considered. Furthermore, message process $i$ sends no second phase messages during the rounds under consideration. Since the smallest ID that a winner can have is 0 , the smallest possible ID for a processor that is not an eventual winner is 1 . Thus the maximum number of second-phase messages for message process $j$ in these rounds is $n / 2^{j}$, for $j>0$. Summing, the total number of messages sent for all the message processes in these rounds is less than $n$.

Finally, consider (3). The argument is similar to the one used for the case in which all processors awaken at the same round. That is, message process $i$ makes a circuit for a total cost of $n$. Any other message process $j$ can send at most $n / 2^{j-i}$ phase-two messages during the time $i$ travels distance $n$. As before, the total number of messages used in (3) is less than $2 n$. Thus the total number of messages for all categories is less than $4 n$.

Because of the variable transmission rate, the number of rounds required is $O\left(n 2^{i}\right)$, where $i$ is the ID of the eventual winner.
The bound of $4 n$ messages for the algorithm above is reasonably tight. Consider the following example, where $f(n)=\log n-\log \log n$. Let processor 1 be at distance 1 from processor 0 , and let processor $k, k=2, \ldots, f(n)$, be at distance $k+$ $\mathrm{L}\left(2^{k-1}-2\right) n /\left(2^{k-1}-1\right) \mathrm{J}$ from processor 0 . Let processors 1 and 2 awaken at round 1 , and each processor $k, k=3, \ldots, f(n)$, awaken the round before it would be visited by the first-phase message from processor $k-1$. Similarly, let processor 0 awaken the round before it would be visited, which would be round $n-f(n)$. Message process $k, k=1, \ldots, f(n)$, will start its second phase at round $2+$ $\mathrm{L}\left(2^{k}-2\right) n /\left(2^{k}-1\right) \mathrm{J}$. It will be killed when it reaches processor 0 , or when it is overtaken near processor 0 , and thus will traverse at least $n /\left(2^{k}-1\right)-f(n)$ links before it is killed. There will be $n$ first-phase messages, at least $\Sigma_{k=1, \ldots, f(n)}$ $\left(n /\left(2^{k}-1\right)-f(n)\right)$ second-phase messages for message processes $k=1, \ldots, f(n)$, and $n-1$ second-phase messages for message process 0 . Of the second-phase messages for message process $k$, note that $\left[n /\left(2^{k}-1\right)-f(n)-2\right] / 2^{k}$ of them will fall under category (2), and the remainder under (3). For large $n$, the total is slightly more than $3.6 n$ messages in all.

It is possible to achieve a trade-off between the number of messages and the number of rounds by using powers of $c$ for any constant $c>1$, rather than powers of 2 . As before, there will be exactly $n$ messages in category (1). In category (2) there will be fewer than $\Sigma_{j=1 \ldots \ldots \infty} n / c^{j}=n /(c-1)$ messages, while in category (3) there will be fewer than $\Sigma_{j=0 \ldots \infty} n / c^{j}=n c /(c-1)$ messages. Thus, we obtain an algorithm that elects a leader in a synchronous ring of $n$ processors using fewer than $2 c n /(c-1)$ messages and at most $O\left(n c^{i}\right)$ rounds, where $i$ is the ID of the eventual winner. It is possible to retain the $2 c n /(c-1)$ message bound, while reducing the time to $O\left(n c^{i}\right)$, where $i$ is the minimum ID of all processors in the ring. The basic idea is to allow each processor to awaken and begin its algorithm (spawning its message process) as soon as it receives any message from its neighbor,
if it has not already awakened on its own. We thus obtain
Theorem 2. Let $c>1$. There is an algorithm that elects a leader in a synchronous ring of $n$ processors using fewer than $2 c n /(c-1)$ messages and $O\left(n c^{i}\right)$ time, where $i$ is the smallest ID of the processors in the ring.

Note that the algorithm works correctly in the case in which communication is purely asynchronous. It is only its complexity that depends on the synchrony. In the general asynchronous case, the algorithm is essentially the same as that of [5], and so exhibits a worst-case message behavior, which is $O\left(n^{2}\right)$.

## 3. Formal Model and Problem Statement

In this section we describe the formal model we use for our lower bounds. The contents of this section are summarized from [7], and the reader is referred to this paper for further details.
3.1 Algorithms. We use the following model for ring algorithms. Each processor is assumed to be identical to every other one, except for its own unique identifier, chosen from an ID space $X$, a totally ordered set. The processors all begin their identical election algorithms at the same time. Each processor behaves like an automaton as follows. Initially the state of the processor consists of its ID. At each round, the processor examines its state and decides whether to send a message to each of its neighbors, and what message to send. Then each processor receives any messages sent to it in that round. The processor uses its current state and these new messages to update its state. Certain of the states are designated as "elected" states.

It may be assumed, without loss of generality, that a ring algorithm is in a certain normal form. In this normal form, the state of each processor records exactly its initial ID and the history of messages received, and each message that is sent contains the entire state of the sending processor. We represent such history information by means of LISP $S$-expressions. The $S$-expressions that arise during computation are of a special type, which we call well formed. A well-formed $S$-expression over $X$ is either (1) an element of $X$, or (2) an expression of the form ( $s_{1}, s_{2}, s_{3}$ ), where $s_{2}$ is a well-formed $S$-expression over $X$, and each of $s_{1}$ and $s_{3}$ is either a well-formed $S$-expression over $X$ or the atom NIL. Let $\mathscr{S}(X)$ denote the set of well-formed $S$-expressions over $X$.

We refer to an algorithm in such a form as a free algorithm, and we restrict attention in this paper to algorithms that are free. An initial state of a processor will just be its ID. Each message will contain exactly the state of the sending processor. When a processor in state $s$ receives messages $s_{1}$ and $s_{2}$ from its counterclockwise and clockwise neighbors, respectively, its new state will be the $S$-expression ( $s_{1}, s, s_{2}$ ). If no message is received from a neighbor, the atom NIL is used in place of $s_{1}$ or $s_{2}$ as a placeholder. To complete the algorithm specification, we define a function that determines when messages are to be sent in either direction, and a designation of which states indicate that the processor has been elected. Thus, an algorithm over $X$ is a pair $(E, \mu)$, where $E \subseteq \mathscr{S}(X)$ is the set of elected states and $\mu$, a mapping from $\mathscr{S}(X) \times$ \{clockwise, counterclockwise\} to \{yes, no\}, is the message generation function. We assume that the set $E$ of elected states is "closed," so that once a processor has been elected, it will remain elected.
3.2 Executions. To facilitate discussion, we index the processors in the ring clockwise, as $0, \ldots, n-1$. (For convenience we are switching from the naming
convention which we used in Section 2. There, by "processor $i$ " we meant "the processor with ID $i$," whereas for the rest of the paper we mean "the processor with index $i$." We count indices modulo $n$. A ring of size $n$ over ID space $X$ is an $n$-tuple of elements of $X$, giving the initial IDs of the processors $0, \ldots, n-1$ in order. A configuration of size $n$ is an $n$-tuple of $S$-expressions in $\mathscr{S}(X)$, representing the states for the $n$ processors. A message vector of size $n$ is an $n$-tuple of ordered pairs of elements of $\mathscr{S}(X) \cup$ \{null\}. It represents the messages sent counterclockwise and clockwise by each of the $n$ processors.

An execution of an algorithm for ring $R$ of size $n$ is an infinite sequence of triples ( $C_{1}, M, C_{2}$ ), where $C_{1}$ and $C_{2}$ are configurations and $M$ is a message vector, all of size $n$. We require executions to satisfy several properties. First, the initial configuration must be $R$. Second, the second configuration in each triple must be the same as the first configuration in the next triple. Finally, each triple in an execution must describe correct message generation, as given by $\mu$, and correct state changes, as described earlier. An execution fragment is any finite prefix of an execution.

We now define our complexity measures. We measure the number of messages sent and the number of rounds taken only up to the point where a processor becomes elected. (This convention only serves to strengthen our lower bound.) For any execution $e$, let finishtime ( $e$ ) denote the number of the first round after which a processor has entered a state in $E$. Let messages( $e$ ) denote the number of messages sent during $e$, up to and including round finishtime $(e)$.
3.3 Election of a Leader. Let $X$ be an ID space with $|X| \geq n$. A ring algorithm over $X$ is said to elect a leader in rings of size $n$ provided that in each execution $e$ of the algorithm, for a ring $R$ of size $n$ over $X$, exactly one processor eventually enters a state in $E$.
3.4 Comparison Algorithms. We next define algorithms whose only operation with respect to processor IDs is to compare them. We say that two $S$ expressions, $s$ and $s^{\prime}$, over $X$ are order-equivalent provided that they are structurally equivalent as $S$-expressions, and, if two atoms from $s$ satisfy one of the order relations $<,=$, or $>$, then the corresponding atoms from $s^{\prime}$ satisfy the same relation. An algorithm is a comparison algorithm provided that, if $s$ and $s^{\prime}$ are order-equivalent well-formed $S$-expressions over $X$, then processors with states $s$ and $s^{\prime}$ transmit messages in the same direction or directions and have the same election status. That is, $\mu(s$, clockwise $)=\mu\left(s^{\prime}\right.$, clockwise), $\mu(s$, counterclockwise $)=$ $\mu\left(s^{\prime}\right.$, counterclockwise), and $s$ is in $E$ exactly if $s^{\prime}$ is in $E$.

## 4. Chains

In this section we describe the general theory needed for our lower bound proof for comparison algorithms. We introduce the concept of a chain, which describes information flow during an execution of a ring algorithm. The notion of a chain used in this paper is a substantial generalization of the notion of a chain used for a similar purpose in [7]. For comparison algorithms, we show that nonexistence of certain chains implies that certain processors in a ring remain indistinguishable.
4.1 Basic Definitions. A $k$-segment of a ring is a length $k$ sequence of consecutive processors in the ring, in clockwise order. Let $S$ and $T$ be two $k$ segments in a ring, with first processors $p$ and $q$, respectively, and last processors $p^{\prime}$ and $q^{\prime}$, respectively, and let $e$ be an execution (or execution fragment) of an algorithm in the ring. Then a clockwise chain in $e$ for $(S, T)$ is a length $k$
subsequence of the steps of $e, e_{i_{1}}, e_{i_{2}}, \ldots, e_{i_{k}}$ such that the following is true. In each step $e_{i j}$, a message is sent either by processor $p+j-2$ to processor $p+j-1$ or by processor $q+j-2$ to processor $q+j-1$. Thus, a clockwise chain for a pair of segments describes combined information flow clockwise in the two scgments from outside the two segments up to the last processors $p^{\prime}$ and $q^{\prime}$. A counterclockwise chain in $e$ for $(S, T)$ is defined analogously for information flow counterclockwise: in each step $e_{i j}$, a message is sent either by processor $p^{\prime}-j+2$ to processor $p^{\prime}-j+1$ or by processor $q^{\prime}-j+2$ to processor $q^{\prime}-j+1$.

Two length $k$ vectors of $X$-elements are said to be order-equivalent provided that the elements in corresponding positions satisfy the same ordering relations in the two vectors. That is, if the two vectors are $a$ and $b$, then $a_{i}$ and $a_{j}$ satisfy the same relation, $<,=$, or $>$, as $b_{i}$ and $b_{j}$. Two segments $S$ and $T$ are said to be orderequivalent in a particular ring $R$ provided that the sequences of initial IDs of the processors in the two segments are order-equivalent.

Let $e$ be an execution fragment. Then $\operatorname{maxc} w(e)$ is defined to be the maximum $k$ for which there are order-equivalent length $k$ segments $S$ and $T$ (possibly with $S$ $=T$ ), such that $e$ contains a clockwise chain for $(S, T)$. The quantity maxccw $(e)$ is defined analogously. Let $\operatorname{sum}(e)=\operatorname{maxcw}(e)+\operatorname{maxccw}(e)$.
4.2 Limitations on Chains. From the definitions of maxccw, maxcw, and sum, it follows that a length 0 execution $e$ has $\operatorname{maxcw}(e)=\operatorname{maxccw}(e)=\operatorname{sum}(e)=$ 0 . We show that chains cannot grow unreasonably quickly. The length of a longest chain can grow by at most 1 in any time step, and only if a message is sent in the appropriate direction.

Lemma 3. Let $e$ and $e^{\prime}$ be execution fragments for a ring $R$, such that $e^{\prime}$ consists of all but the last step of $e$. Then $(a) \operatorname{maxcw}(e) \leq \operatorname{maxcw}\left(e^{\prime}\right)+1$, with $\operatorname{maxcw}(e)=\operatorname{maxcw}\left(e^{\prime}\right)$ if no messages are sent clockwise at the last step of $e$, and (b) $\operatorname{maxccw}(e) \leq \operatorname{maxccw}\left(e^{\prime}\right)+1$, with maxccw $(e)=\operatorname{maxccw}\left(e^{\prime}\right)$ if no messages are sent counterclockwise at the last step of $e$.

Proof. We argue part (a). Part (b) is analogous. The second half of the claim is obvious. We argue the inequality $\max \operatorname{cw}(e) \leq \operatorname{maxcw}\left(e^{\prime}\right)+1$. We may assume that $\operatorname{maxcw}(e) \geq 1$, since otherwise the result is obvious.

Let $S$ and $T$ be order-equivalent segments of length $\operatorname{maxcw}(e)$ for which there is a clockwise chain in $e$. Let $S^{\prime}$ and $T^{\prime}$ be the segments of length maxcw $(e)-1$ consisting of all but the last processor in $S$ and $T$, respectively. Then $S^{\prime}$ and $T^{\prime}$ are order equivalent. Moreover, since only the last message in the chain could have been sent at the last step of $e$, it must be that $e^{\prime}$ contains a clockwise chain for ( $S^{\prime}, T^{\prime}$ ). Thus, $\operatorname{maxcw}\left(e^{\prime}\right) \geq \operatorname{maxcw}(e)-1$, as required.
4.3 Bisegments. We next introduce notation that allows us to describe at the same time a counterclockwise chain and a clockwise chain leading to the same processor. If $k_{1}$ and $k_{2}$ are positive integers, a ( $k_{1}, k_{2}$ )-bisegment is defined to be a pair of segments, the first of size $k_{1}$ and the second of size $k_{2}$, which overlap in a single processor. (The last processor of the first segment is the first of the second segment.) The processor that appears in both segments is called the center of the bisegment. The spanning segment of a bisegment is the segment obtained by concatenating the two segments in the bisegment and removing the duplicated center. Two bisegments are said to be order-equivalent in a particular ring provided that their spanning segments are order-equivalent. Two processors $p$ and $q$ are ( $k_{1}, k_{2}$ )-equivalent in a particular ring provided that their ( $k_{1}, k_{2}$ )-bisegments (i.e., the ( $k_{1}, k_{2}$ )-bisegments centered at $p$ and $q$ ) are order-equivalent.

Let $S=\left(S_{1}, S_{2}\right)$ and $T=\left(T_{1}, T_{2}\right)$ be two $\left(k_{1}, k_{2}\right)$-bisegments, and let $e$ be an execution or execution fragment. Then a clockwise chain in $e$ for $(S, T)$ is a clockwise chain in $e$ for ( $S_{1}, T_{1}$ ), and a counterclockwise chain for $(S, T)$ is a counterclockwise chain for ( $S_{2}, T_{2}$ ). A chain in $e$ for $(S, T)$ is either a clockwise chain or a counterclockwise chain for $(S, T)$.
4.4 Indistinguishability. In this subsection we show that, for comparison algorithms, the absence of long enough chains implies that certain processors must remain "indistinguishable." The absence of these chains then also implies that a correspondingly large number of messages will be sent in the next round.

Our notion of "indistinguishability" is defined as follows. If $S$ and $T$ are two ID sequences, each of length $k$, and $s$ and $t$ are two $S$-expressions, then $s$ is congruent to $t$ with respect to ( $S, T$ ) provided that $s$ and $t$ are structurally equivalent, and corresponding positions in $s$ and $t$ contain elements from corresponding positions of $S$ and $T$, respectively. If $S$ and $T$ are two segments of a particular ring, then $S$ and $t$ are congruent with respect to $(S, T)$ provided that $s$ and $t$ are congruent with respect to the corresponding sequences of IDs. Similarly, if $S$ and $T$ are two bisegments of a ring, we say that $s$ and $t$ are congruent with respect to $S$ and $T$ provided that they are congruent with respect to their spanning segments.

Lemma 4. Let e be an execution fragment of a comparison algorithm for ring $R$. Let $k_{1}$ and $k_{2}$ be positive integers. Let $p$ and $q$ be any pair of $\left(k_{1}, k_{2}\right)$-equivalent processors in $R$, and let $S$ and $T$ be their respective $\left(k_{1}, k_{2}\right)$-bisegments. If there are no chains in $e$ for $(S, T)$, then at the end of $e$, the states of $p$ and $q$ are congruent with respect to $(S, T)$.

Proof. The proof is by induction on the length of $e$.
Base. $|e|=0$. Neither $p$ nor $q$ has received any messages in $e$, so they will remain in states that are congruent with respect to $(S, T)$.

Inductive Step. $\quad|e|>0$. Assume as the induction hypothesis that the result holds for any execution fragment of length shorter than $|e|$ and any values of $k_{1}$ and $k_{2}$. Let $e^{\prime}$ denote $e$, except for its last step. Then by inductive hypothesis, $p$ and $q$ remain in states that are congruent with respect to $(S, T)$ up to the end of $e^{\prime}$. Consider what happens at the last step. Let $p^{\prime}$ and $q^{\prime}$ be the respective counterclockwise neighbors of $p$ and $q$, and $p^{\prime \prime}$ and $q^{\prime \prime}$ the respective clockwise neighbors.

Case 1. Both of the following hold: (a) Either $p^{\prime}$ and $q^{\prime}$ are in states that are congruent with respect to ( $S, T$ ) just after $e^{\prime}$, or else neither $p^{\prime}$ nor $q^{\prime}$ sends a message clockwise at the last step of $e$. (b) Either $p^{\prime \prime}$ and $q^{\prime \prime}$ are in states that are congruent with respect to ( $S, T$ ) just after $e^{\prime}$, or else neither $p^{\prime \prime}$ nor $q^{\prime \prime}$ sends a message counterclockwise at the last step of $e$.

In this case, it is easy to see that $p$ and $q$ remains in states that are congruent with respect to $(S, T)$, after $e$. For if $p^{\prime}$ and $q^{\prime}$ are in states that are congruent with respect to ( $S, T$ ) just after $e^{\prime}$, then, since the algorithm is a comparison algorithm, they both make the same decision about whether or not to send a message clockwise at the last step of $e$. If they both send a message, then the messages they send are just their respective states, which are congruent with respect to ( $S, T$ ). A similar argument applies to $p^{\prime \prime}$ and $q^{\prime \prime}$. It follows that $p$ and $q$ remain in states that are congruent with respect to ( $S, T$ ) after the last step of $e$.

Case 2. Processors $p^{\prime}$ and $q^{\prime}$ are in states that are not congruent with respect to ( $S, T$ ) just after $e^{\prime}$, and at least one of them sends a message clockwise at the last step of $e$.

If $k_{1}=1$ (i.e., if $p$ and $q$ are at the counterclockwise ends of their respective bisegments), then a clockwise chain for ( $S, T$ ) is produced by the message sent at the last step, a contradiction. So assume that $k_{1}>1$. Since $p$ and $q$ are ( $k_{1}, k_{2}$ )equivalent, it follows that $p^{\prime}$ and $q^{\prime}$ are ( $k_{1}-1, k_{2}+1$ )-equivalent. Let $S^{\prime}$ and $T^{\prime}$ denote their respective $\left(k_{1}-1, k_{2}+1\right)$-bisegments. $S^{\prime}$ and $T^{\prime}$ contain exactly the same processors as $S$ and $T$, respectively, but are centered at $p^{\prime}$ and $q^{\prime}$ rather than at $p$ and $q$. Since the states of $p^{\prime}$ and $q^{\prime}$ just after $e^{\prime}$ are not congruent with respect to ( $S, T$ ), they are also not congruent with respect to ( $S^{\prime}, T^{\prime}$ ). By the inductive hypothesis, there must be a chain in $e^{\prime}$ for ( $S^{\prime}, T^{\prime}$ ). If there is a counterclockwise chain in $e^{\prime}$ for ( $S^{\prime}, T^{\prime}$ ), then it is also a counterclockwise chain for ( $S, T$ ), so there is a counterclockwise chain in $e$ for $(S, T)$. On the other hand, if there is a clockwise chain in $e^{\prime}$ for ( $S^{\prime}, T^{\prime}$ ), then since at least one of $p^{\prime}$ and $q^{\prime}$ sends a message clockwise at the last step of $e$, we obtain a clockwise chain in $e$ for $(S, T)$. Either case is a contradiction.

Case 3. Processors $p^{\prime \prime}$ and $q^{\prime \prime}$ are in states that are not congruent with respect to ( $S, T$ ) just after $e^{\prime}$, and at least one of them sends a message counterclockwise at the last step of $e$.

The argument is analogous to the one for Case 2.
Thus, we have shown that absence of certain chains implies that certain processors must remain in congruent states. This lemma is actually stronger than we need for this paper, but this extra strength will probably be of use in handling other problems. In our subsequent analysis we use as an upper bound on $\operatorname{maxcw}(e)$ simply the number of distinct rounds in which messages are sent clockwise, and similarly for maxccw $(e)$. Thus, instead of the existence of a chain for ( $S, T$ ), we could have substituted the condition that either there are $k_{1}$ rounds in which messages are sent clockwise or there are $k_{2}$ rounds in which messages are sent counterclockwise. Reorganized in this way, our proof would be substantially the same as it is now (in fact, marginally simpler), but the revised lemma would give less information about the communication that must occur for congruence to be broken.

Two corollaries, which will be used in our lower bound proofs, follow from this lemma. The first one says that, when chains are short and there are lots of equivalent processors, any message that gets sent has many corresponding messages sent at the same time by other processors.

Corollary 5. Let $k$ be a positive integer. Assume ring $R$ is such that every $k$ segment has at least iorder-equivalent $k$-segments. Let e be any execution fragment of a comparison algorithm in $R, e^{\prime}$ be another fragment consisting of all but the last step of $e$, and assume that sum $\left(e^{\prime}\right)<k$. If some processor $p$ sends a message clockwise (or counterclockwise) at the last step of $e$, then there are at least $i$ processors that do the same.

Proof. Consider the case in which $p$ sends a message clockwise. The other case is analogous. Let $k_{1}=\operatorname{maxcw}\left(e^{\prime}\right)+1$ and $k_{2}=\operatorname{maxccw}\left(e^{\prime}\right)+1$. The $\left(k_{1}, k_{2}\right)$ bisegment for $p$ has at most $k$ elements, so that $p$ has at least $i\left(k_{1}, k_{2}\right)$-equivalent processors. Let $q$ be any one of these processors, and let $S$ and $T$ be the ( $k_{1}, k_{2}$ )bisegments centered at $p$ and $q$, respectively. Then there cannot be a chain in $e^{\prime}$
for $(S, T)$, by the definitions of maxcw and maxccw. But then Lemma 4 implies that $p$ and $q$ remain congruent with respect to ( $S, T$ ) at the end of $e^{\prime}$; since the algorithm is a comparison algorithm, $q$ also sends a message clockwise at the last step of $e$.

Lemma 4 also has the following consequence for comparison algorithms to elect a leader. This corollary says that long chains must be generated in order to elect a leader, if certain equivalent processors exist.

Corollary 6. Let $k$ be a positive integer. Let $R$ be a ring in which every $k$-segment $S$ has another order-equivalent $k$-segment $T$. Let e be any execution fragment of a comparison algorithm that elects a leader in $R$, such that a leader gets elected in $e$. Then $\operatorname{sum}(e) \geq k$.

Proof. Assume the opposite, that $\operatorname{sum}(e)=\operatorname{maxcw}(e)+\operatorname{maxccw}(e)<k$. Let $k_{1}=\operatorname{maxcw}(e)+1$ and $k_{2}=\operatorname{maxccw}(e)+1$. The $\left(k_{1}, k_{2}\right)$-bisegment for the processor $p$ that gets elected leader has at most $k$ elements, so that $p$ has a ( $k_{1}, k_{2}$ )equivalent processor $q \neq p$; let $S$ and $T$ be the ( $k_{1}, k_{2}$ )-bisegments centered at $p$ and $q$, respectively. Then there cannot be a chain in $e$ for $(S, T)$, by the definition of maxcw and maxcew. But then Lemma 4 implies that $p$ and $q$ remain congruent with respect to ( $S, T$ ); since the algorithm is a comparison algorithm, $p$ and $q$ cannot be distinguished as to leadership. This is a contradiction.

## 5. Lower Bound for Comparison Algorithms When n Is a Power of 2

In this section we restrict attention to algorithms that use comparisons only, and to rings in which the number of processors is a power of 2 . We present a lower bound of $n / 2(\log n+1)$ for the number of messages required for a comparison algorithm to elect a leader in this case. We handle the case of powers of 2 first because the assignment of IDs to processors that realizes the lower bound is simpler than for general values of $n$, and also because the constant of proportionality in the lower bound is larger than we have been able to achieve for general $n$.
5.1 Replication Symmetry. We first generate a labeling of the processors in a ring that has a large amount of replication symmetry. Let $\langle n\rangle$ denote $\{0, \ldots, n-1\}$. We assume that $n$ is a power of 2 , and let $X^{*}$ be the ID space consisting of the set $\langle n\rangle$, with the usual ordering.

For $j \in\langle n\rangle$, let reverse $(j)$ denote the integer whose binary representation is the reverse of the binary representation of $j$. We assign processor IDs so that processor $j$ has ID reverse $(j)$, for $j \in\langle n\rangle$. We call this pattern of IDs $Q_{n}$. We note that if a segment of $Q_{n}$ is of length at most $2^{i}$, then all ordering information about the IDs of processors in the segment is determined solely by the $i$ high-order bits.

Lemma 7. Let $S$ by any segment of $Q_{n}$ of length at most $2^{i}$, where $i<\log n$. Then there are at least $n / 2^{i}$ segments of $Q_{n}$ that are order-equivalent to $S$, including $S$ itself.

Proof. For each $i<\log n$, the processor IDs repeatedly cycle through the $2^{i}$ possible arrangements of $i$ high-order bits. Thus in a segment of length at most $2^{i}$, each ID differs from any other in its $i$ high-order bits. Any segment that is orderequivalent to $S$ will have its first processor at a distance that is any integral multiple of $2^{i}$ from the first processor in $S$. There are $n / 2^{i}$ such segments, including $S$ itself.
5.2 Lower Bound. We can now prove the lower bound for comparison algorithms when $n$ is a power of 2 . We make use of the following observation about comparison algorithms. Suppose $X$ and $X^{\prime}$ are arbitrary ID spaces, and $n$ is any integer. If $\mathscr{A}$ is a comparison algorithm over $X$ that elects a leader in a ring of size $n$ and uses at most $s$ messages, then there exists a comparison algorithm $\mathscr{A}^{\prime}$ over $X^{\prime}$ that elects a leader in a ring of size $n$ and uses at most $s$ messages. Thus a lower bound result over ID space $X^{*}$ ranslates directly into a lower bound result for any arbitrary ID space $X$.

Theorem 8. Assume $n$ is a power of 2 . Let $\mathscr{A}$ be a comparison algorithm over an arbitrary ID space $X$, which elects a leader in a synchronous ring of size $n$. Then there is an execution $e$ of $\mathscr{A}$ for which messages $(e) \geq(n / 2)(\log n+1)$.

Proof. It suffices to consider $X=X^{*}$. Let $e$ be the execution fragment on $Q_{n}$, which terminates just when the elected processor enters an "elected" state. By Lemma 7, every segment of length $n / 2$ has at least one other order-equivalent segment in $Q_{n}$. Thus by Corollary 6, execution $e$ must progress from having a sum of 0 to having a sum of at least $n / 2$.

Consider any step of $e$ at which the sum first stops being at most $k$, for any $k<2^{i}$. By Lemma 3, the sum increases by at most 2 at this step. Moreover, if no messages are sent clockwise (respectively, counterclockwise) at this step, then the sum increases by at most 1 .

Let $e^{\prime}$ be the prefix of $e$ preceding this step. Then $\operatorname{sum}\left(e^{\prime}\right)<2^{i}$. Lemma 7 implies that any segment of length $2^{\prime}$ has at least $n / 2^{\prime}$ order-equivalent segments in $Q_{n}$. Thus by Corollary 5, if any messages are sent clockwise at this step, then at least $n / 2^{i}$ messages are sent clockwise, and similarly for messages sent counterclockwise. Thus, if the sum increases by 1 at this step, at least $n / 2^{i}$ messages are sent, whereas if the sum increases by 2 at this step, then at least twice that number of messages are sent. It follows that the cost of increasing the sum from 0 to at least $n / 2$ can be apportioned as a cost of at least $n / 2^{i}$ for each increase from $k$ to $k+1$, where $k<2^{i}$.

We now total up the number of messages sent in e. Grouping increases by powers of 2 , we see that the number of messages sent must be at least

$$
\begin{aligned}
n+\sum_{1, \ldots, \log (n / 2)}\left(\frac{n}{2^{i}}\right)\left(2^{i}-2^{i-1}\right) & =n+\sum_{1, \ldots, \log (n / 2)} \frac{n}{2} \\
& =\left(\frac{n}{2}\right)(\log n+1)
\end{aligned}
$$

## 6. Lower Bound for Comparison Algorithms for General $n$

In the last section we generated an assignment of IDs to processors in the case in which $n$ was a power of 2 . The assignment possessed a large amount of replication symmetry, which allowed us to achieve the $\Omega(n \log n)$ lower bound. It does not appear possible to take our pattern $Q_{n}$ and then try to extend it in some way to accommodate extra processors. Such a strategy would introduce special treatment for the extra processors, which might change the behavior of the algorithm entirely, perhaps allowing some processor to become elected easily. Instead, we generate a pattern $P_{n}$ for any general value of $n$, such that a ring assigned IDs from $P_{n}$ possesses a large amount of replication symmetry. We then show that this replication symmetry causes the ring to require a large number of messages for election of a leader.
6.1 Hierarchical Organization of Processors. Fix a particular ring size $n \geq 1$. We generate a pattern $P_{n}$ of IDs, the elements of which are then assigned to processors 0 througn $n-1$, respectively. To achieve considerable replication symmetry, the construction of $P_{n}$ uses a hierarchical grouping of processors. The idea is that on any level of the hierarchy, two groups of processors should receive order-equivalent sequences of IDs. To have the construction work for general $n$, one type of group is not enough, so that at every level there will be two types of groups. We describe the grouping using a derivation tree of a context-free grammar. Later, we use the structure of the derivation tree to assign IDs to the $n$ leaves of the tree and thereby produce pattern $P_{n}$.

Define the context-free grammar $G$ as follows. The nonterminals, representing groups of processors, are $A_{i}$ and $B_{i}, 1 \leq i \leq d$, plus $B_{0}$. There is just one terminal symbol $p$, representing a processor. The start symbol is $B_{0}$. The productions are

$$
\begin{array}{ll}
B_{i} \rightarrow B_{i+1} A_{i+1} A_{i+1} B_{i+1} B_{i+1} A_{i+1} A_{i+1} B_{i+1} B_{i+1}, & \text { for } \quad 0 \leq i \leq d-1, \\
A_{i} \rightarrow A_{i+1} B_{i+1} B_{i+1} A_{i+1} A_{i+1} B_{i+1} B_{i+1} A_{i+1} A_{i+1}, & \text { for } \quad 1 \leq i \leq d-1, \\
B_{d} \rightarrow p^{\left(k_{d}\right)}, \quad \text { and } A_{d} \rightarrow p^{\left(a_{d}\right) .}
\end{array}
$$

The depth $d$ of the hierarchy is defined as $d=\left\llcorner\left(\log _{9} n\right) / 2 \mathrm{~J}\right.$. Note that in the last two productions, $B_{d}$ generates a string consisting of $b_{d} p$ symbols, and analogously for $A_{d}$. The quantities $a_{d}$ and $b_{d}$ will be defined later, in such a way as to guarantee that the length of the unique sentence generated by $G$ is $n$.

For each $i, 0 \leq i \leq d$, define the level $i$ sentential form of $G$ to be the unique string over $\left\{A_{i}, B_{i}\right\}$ derivable in $G$. There are exactly $9^{i}$ nonterminal symbols in the level $i$ sentential form. Moreover, for each $i$, the number of symbols $A_{i}$ is exactly one less than the number of symbols $B_{i}$.

Lemma 9. In the level $i$ sentential form of $G, 0 \leq i \leq d$, the number of symbols $A_{i}$ is $\left\lfloor 9^{i} / 2\right\rfloor$, and the number of symbols $B_{i}$ is $\Gamma 9^{i} / 21$.

Proof. By induction on $i$.
All $A_{i}$ nodes derive a terminal string of the same length; we call this length $a_{i}$. Similarly, all $B_{i}$ nodes derive a terminal string of the same length, which we call $b_{i}$. Let $c_{i}=\min \left(a_{i}, b_{i}\right)$, for all $i, 1 \leq i \leq d$.

We next describe how to select the values $a_{d}$ and $b_{d}$. They are chosen in such a way that the total length of the unique sentence derived in $G$ is exactly $n$, and so that $\left|b_{d}-a_{d}\right|$ is small. We use the following:

Lemma 10. Let $m, n \geq 0$ be integers. Then there are integers $a$ and $b$ such that $n=a m+b(m+1)$ and $|b-a| \leq m$.

Proof. Let $c=\lfloor n / m\rfloor, f=n-c m$, and $e=\lfloor(c-2 f+m) /(2 m+1)\rfloor$. We take $a=c-f-e(m+1)$ and $b=f+e m$. Then $a m+b(m+1)=(c-f-e(m+1)) m$ $+(f+e m)(m+1)=(c-f) m+f(m+1)=c m+f=n$. Also, $a-b+m=$ $c-f-e(m+1)-(f+e m)+m=c-2 f+m-(2 m+1) e=c-2 f+m-$ $(2 m+1)\lfloor(c-2 f+m) /(2 m+1)\rfloor$, which is between 0 and $2 m$. Thus $a-b$ is between $-m$ and $+m$, that is, $|a-b| \leq m$.

Let $m=\left\lfloor 9^{d} / 2\right\rfloor$. It is easy to see that $m$ is $\Theta\left(n^{1 / 2}\right)$ and, in particular, that $m \leq n^{1 / 2} / 2$. Using Lemma 10 , choose $a_{d}$ and $b_{d}$ to be integers such that $n=$ $a_{d} m+b_{d}(m+1)$ and $\left|b_{d}-a_{d}\right| \leq m$. We must show that $a_{d}$ and $b_{d}$ are nonnegative: if either of $a_{d}$ and $b_{d}$ is negative, then $\max \left(a_{d}, b_{d}\right) \leq m-1$, so $n=a_{d} m+$ $b_{d}(m+1) \leq 2\left(m^{2}\right) \leq n / 2$, a contradiction.

Lemma 11. The length of the unique sentence generated by $G$ is $n$.
Proof. By Lemma 9, there are exactly $\left[9^{d} / 2\right\rfloor=m$ symbols $A_{d}$ and exactly $\left\lceil 9^{d} / 2\right\rceil=m+1$ symbols $B_{d}$ in the level $d$ sentential form of $G$. Since $n=a_{d} m+$ $b_{d}(m+1)$, the result holds.

We have already noted that $m$ is $\Theta\left(n^{1 / 2}\right)$. Since $a_{d}$ is nonnegative, we have that $n \geq b_{d}(m+1)$. Using the lower bound on $m$, we see that $b_{d}$ is $O\left(n^{1 / 2}\right)$.

The final lemma of this subsection gives the exact value of the difference $c_{i}-c_{i+1}$, which we use in the analysis of the lower bound.

Lemma 12. The difference $c_{i}-c_{i+1}=4 \cdot 9^{d-(i+1)}\left(n-b_{d}\right) / m$, for $0 \leq i \leq d-1$.
Proof. Note that $c_{i}=\min \left(a_{i}, b_{i}\right)=\min \left(5 a_{i+1}+4 b_{i+1}, 4 a_{i+1}+5 b_{i+1}\right)=4 a_{i+1}+$ $4 b_{i+1}-\min \left(a_{i+1}, b_{i+1}\right)=4\left(a_{i+1}+b_{i+1}\right)-c_{i+1}$. Thus $c_{i}-c_{i+1}=4\left(a_{i+1}+b_{i+1}\right)$.

From the choice of $a_{d}$ and $b_{d}$, we have $a_{d}+b_{d}=\left(n-b_{d}\right) / m$. It follows that $a_{i+1}+b_{i+1}=9^{d-(i+1)}\left(n-b_{d}\right) / m$. Substituting into the expression for $c_{i}-c_{i+1}$ gives the desired result.
6.2 Labeling of Processors. Let $X$ be the ID space consisting of all strings of length $d+1$ whose elements are nonnegative integers, with the strings ordered lexicographically. $X$ is the ID space from which the pattern $P_{n}$ will be constructed.

We define $P_{n}$ by describing an assignment of IDs to $n$ processors, corresponding to the leaves of the derivation tree of $G$. In order to do this, we associate labels with the nodes of the derivation tree. The label of the root of the tree is the null string. If a node with a corresponding nonterminal $A_{i}$ or $B_{i}, 0 \leq i \leq d-1$, is labeled by the string $w$, then the labels of its nine children are, respectively, $w 0$, $w 1, w 2, w 3, w 8, w 7, w 6, w 5, w 4$. If a node with a corresponding nonterminal $A_{d}$ is labeled by the string $w$, then the labels of its $a_{d}$ children are, respectively, $w 0, w 1$, $\ldots, w\left(a_{d}-1\right)$. If a node with a corresponding nonterminal $B_{d}$ is labeled by the string $w$, then the labels of its $b_{d}$ children are, respectively, $w 0, w 1, \ldots, w\left(b_{d}-1\right)$. Processor IDs are generated by interpreting the labels of the leaves as elements of $X$, that is, as length $d+1$ strings of nonnegative integers, ordered lexicographically.

In the level $i$ sentential form of $G$, define an ordered pair of nonterminal symbols to be "of type $A>A$ " provided that it consists of the two symbols $A_{i} A_{i}$, and the label of the node of the first nonterminal is lexicographically greater than that of the second. We use analogous definitions for types $A<A, A>B, A<B, B>A$, $B<A, B>B$, and $B<B$. We now show that the level $i$ sentential form has equal numbers of consecutive pairs of nonterminals of the eight possible types.

Lemma 13. In the level $i$ sentential form of $G, 0 \leq i \leq d$, the number of occurrences of consecutive pairs of each of the eight types $A>A, A<A, A>B$, $A<B, B>A, B<A, B>B$, and $B<B$ is exactly $\left[9^{i} / 8 \mathrm{~J}\right.$.

Proof. It suffices to show that the numbers of occurrences of the eight types of pairs are equal, since the total number of pairs is exactly $9^{i}-1=8\left[9^{i} / 8 \mathrm{~J}\right.$. We proceed by induction on $i$. For the basis, $i=0$, the result is vacuously true. Assume that the result is true for $i$, and consider the level $i+1$ sentential form. There are two kinds of pairs of level $i+1$ nonterminals: those in which both elements are derived from the same level $i$ nonterminal node and those in which the two elements are derived from two different level $i$ nonterminal nodes. Each level $i$ nonterminal node generates a length 9 sequence of level $i+1$ nonterminals in which each of the eight types of pairs has exactly one occurrence. Therefore, there are equal numbers of the eight possible types among the pairs that are derived from
the same level $i$ nonterminal node. Also, each pair that is derived from two different level $i$ nonterminal nodes is of the same type as the corresponding pair of parent nodes; the inductive hypothesis implies that there are equal numbers of the eight possible types among these pairs, as well. The result follows.

In any level $i$ sentential form, note that the pair consisting of the last nonterminal node followed by the first nonterminal node, is of type $B>B$.

Having assigned the IDs in pattern $P_{n}$ to the processors of the ring, we state a lemma that describes the replication symmetry of the ring. This lemma is used in the next subsection to yield our lower bound for the number of messages required by a comparison algorithm to elect a leader.

Lemma 14. Consider a ring labeled with $P_{n}$. Let $1 \leq i \leq d$. Let $S$ be any segment of length at most $c_{i}+1$. Then there are at least $\left[9^{i} / 8 \mathrm{~J}\right.$ segments that are order-equivalent to $S$, including $S$ itself.

Proof. $S$ is contained in the subtrees of at most two nonterminal nodes at level $i$. These two are either two consecutive nonterminal nodes or the last and first nonterminals in the sentential form. Let $t$ be the type of this ordered pair of nonterminal nodes.

By Lemma 13, there are at least $\left[9^{i} / 8\right\rfloor$ instances of type $t$ consecutive pairs of nonterminal nodes in the level $i$ sentential form. Each of these instances of a pair of type $t$ contains a segment that is order-equivalent to $S$.

The following corollary is an immediate consequence of Lemma 14. It describes the replication symmetry that we obtain in our construction in a slightly different way. This new form is not used in this paper but may be of independent interest.

Corollary 15. There exists a constant c such that, for all positive integers $n$, there is a ring $R_{n}$ of $n$ processors having the following property. Let $S$ be any segment of $R_{n}$ of length $k$, where $k \geq \sqrt{n}$. Then there are at least $\left.L n / k c\right\rfloor$ segments in $R_{n}$ that are order-equivalent to $S$, including $S$ itself.
6.3 Lower Bound. In this section we state and prove our lower bound for the number of messages required by a comparison algorithm to elect a leader. We use the pattern $P_{n}$ constructed in the previous subsection and the two corollaries from Section 4.

Theorem 16. Let $\mathscr{A}$ be a comparison algorithm over an arbitrary ID space $X$ that elects a leader in a synchronous ring of size $n$. Then there is an execution $e$ of $\mathscr{A}$ for which messages $(e) \geq \Omega(n \log n)$.

Proof. Assume $n$ is fixed, and at least $9^{4}$. This ensures that the depth $d=\mathrm{L}\left(\log _{9} n\right) / 2 \mathrm{~J}$ is at least 2 . It suffices to consider the ID space $X$ consisting of length $d+1$ strings of nonnegative integers, ordered lexicographically. Assume the pattern $P_{n}$ is used to label the ring. Let $e$ be the execution fragment for the ring that terminates just when the elected processor enters an "elected" state. By Lemma 14, every segment of length $c_{2}+1$ has at least one other order-equivalent segment in the ring. (The lemma actually implies that there are at least nine others, but we do not require this fact here.) Thus, by Corollary 6, execution $e$ must progress from having a sum of 0 to having a sum of at least $c_{2}+1$.

Consider any step of $e$ at which the sum first stops being at most $k$, for any $k \leq c_{i}$. By Lemma 3, the sum increases by at most 2 at this step. Moreover, if no messages are sent clockwise (respectively, counterclockwise) at this step, then the sum increases by at most 1 .

Let $e^{\prime}$ be the prefix of $e$ preceding this step. Then $\operatorname{sum}\left(e^{\prime}\right)<c_{i}+1$. Lemma 14 implies that any segment of length $c_{i}+1$ has at least $\left[9^{i} / 8\right\rfloor$ order-equivalent segments in the ring. Thus, by Corollary 5 , if any messages are sent clockwise at this step, then at least $\left[9^{i} / 8\right\rfloor$ messages are sent clockwise, and similarly for messages sent counterclockwise. Thus, if the sum increases by 1 at this step, at least $\left[9^{i} / 8 \mathrm{~J}\right.$ messages are sent, whereas if the sum increases by 2 at this step, then at least twice that number of messages are sent. It follows that the cost of increasing the sum from 0 to at least $c_{2}+1$ can be apportioned as a cost of at least $\left[9^{i} / 8\right\rfloor$ for each increase from $k$ to $k+1$, where $k \leq c_{i}$.

We now total up the number of messages sent in $e$. Grouping increases according to level, we see that the number of messages sent must be at least

$$
\sum_{i=2, \ldots, d-1}\left\lfloor\frac{9^{i}}{8}\right\rfloor\left(c_{i}-c_{i+1}\right) .
$$

By Lemma 12, this quantity is equal to

$$
\begin{aligned}
& \sum_{i=2 \ldots, d-1}\left\lfloor\frac{9^{i}}{8}\right\rfloor\left(4 \cdot \frac{9^{d-(i+1)}\left(n-b_{d}\right)}{m}\right) \\
& \quad=4\left(\frac{n-b_{d}}{m}\right) \sum_{i=2, \ldots, d-1}\left\lfloor\frac{9^{i}}{8}\right\rfloor 9^{d-(i+1)} \\
& \quad \geq 4\left(\frac{n-b_{d}}{m}\right)\left[\sum_{i=2, \ldots, d-1}\left(\frac{9^{i}}{8}\right) 9^{d-(i+1)}-\sum_{i=2, \ldots, d-1} 9^{d-(i+1)}\right] .
\end{aligned}
$$

The first summation evaluates to $(d-2) 9^{d-1} / 8$, whereas the second is bounded above by $9^{d-2} / 8$. Thus, the message bound is at least

$$
4\left(\frac{n-b_{d}}{m}\right)\left[\frac{(d-2) 9^{d-1}}{8}-\frac{9^{d-2}}{8}\right]
$$

Since $m=\left\lfloor 9^{d} / 2\right\rfloor \leq 9^{d} / 2$, this is at least
$8\left(\frac{n-b_{d}}{9^{d}}\right)\left[\frac{(d-2) 9^{d-1}}{8}-\frac{9^{d-2}}{8}\right]=\left(n-b_{d}\right)\left[\frac{d-2}{9}-\frac{1}{81}\right]=\left(n-b_{d}\right)\left[\frac{d}{9}-O(1)\right]$.
Since $b_{d}$ is $O\left(n^{1 / 2}\right)$, the message bound is at least

$$
\begin{aligned}
& =\left(n-O\left(n^{1 / 2}\right)\right)\left[\frac{d}{9}-O(1)\right] \\
& =\left(n-O\left(n^{1 / 2}\right)\right)\left[\frac{(1 / 2) \log _{9} n}{9}-O(1)\right] \\
& =n \frac{(1 / 2) \log _{9} n}{9}-O(n) \\
& =\frac{n \log _{2} n}{18 \log _{2} 9}-O(n)
\end{aligned}
$$

## 7. Lower Bound for Time-Bounded Algorithms

In this section we prove our lower bound for time-bounded algorithms. We use the lower bound for comparison algorithms to do this. First, we show how to map from time-bounded algorithms to comparison algorithms. This result, presented in
the paracomputer model, is due to Snir [14], (Snir (personal communication) credits Yao [16] with inspiration for this result.) For completeness, we present a careful proof in our setting, even though a similar proof appears in [14]. We then infer the lower bound for time-bounded algorithms.
7.1 Definitions. In order to map from time-bounded to comparison algorithms, we require definitions describing the behavior of an algorithm within a bounded amount of time. We say that a free algorithm is a $t$-comparison algorithm provided that both of the following conditions hold:
(1) If $s$ and $s^{\prime}$ are order-equivalent $S$-expressions of parenthesis depth at most $t-1$, then $\mu^{\prime}(s$, clockwise $)=\mu^{\prime}\left(s^{\prime}\right.$, clockwise $)$ and $\mu^{\prime}(s$, counterclockwise $)=$ $\mu^{\prime}\left(s^{\prime}\right.$, countcrclockwise $)$.
(2) If $s$ and $s^{\prime}$ are order-equivalent $S$-expressions of depth at most $t$, and $a \in A$, then $s$ is in $E$ exactly if $s^{\prime}$ is in $E$.
During execution of a free algorithm, the $S$-expressions, which appear as states at the end of any round $t$, have depth exactly $t$. Thus, this definition says that the algorithm behaves as a comparison algorithm up to the end of the first $t$ rounds. We also add the qualifier "on inputs from $U$ " to this definition, provided that the appropriate conditions hold for the $S$-expressions that use atoms chosen from the set $U$.
7.2 Mapping a Time-Bounded Algorithm to a Comparison Aigorithm. In this subsection, we show how to convert a time-bounded algorithm to a comparison algorithm. The first step is to show that any free algorithm behaves as a comparison algorithm on a subset of its inputs. For the first lemma, we use a particular fastgrowing function $f(n, t)$. The precise definition of $f$ depends on Ramsey's Theorem and is implicit in the proof of the lemma.

Lemma 17. Fix n, t. Let $\mathscr{A}$ be any free algorithm over ID space $X$, where $X$ has at least $f(n, t)$ elements. Then there exists a subset $U$ of $X$, of size at least $n$, such that $\mathscr{A}$ is a $t$-comparison algorithm, on inputs from $U$.

Proof. Let $Y$ and $Z$ be two $n$-subsets of $X$, and let $Y=\left(y_{1}, y_{2}, \ldots, y_{n}\right)$ and $Z=\left(z_{1}, z_{2}, \ldots, z_{n}\right)$ be their representations in increasing order. Define $Y$ and $Z$ to be decision-equivalent if, for every $S$-expression of depth at most $t$ over $Y$, the corresponding $S$-expression over $Z$ (generated by substituting $z_{i}$ for $y_{i}, i=$ $1, \ldots, n$ ), gives rise to the same combination of choices: whether a message is sent counterclockwise, whether a message is sent clockwise, and whether or not the expression is in $E$. Decision-equivalence partitions the $n$-subsets of $X$ into finitely many equivalence classes. By Ramsey's theorem [3] there is a function $f(n, t)$ such that, if $X$ is of cardinality at least $f(n, t)$, then there is a subset $C$ of $X$ of cardinality $2 n-1$ such that all $n$-subsets of $C$ belong to the same equivalence class. Then take $U$ to be the set of the $n$ smallest elements of $C$.

That $U$ is the desired subset of $X$ is shown as follows. Consider two $m$-subsets $Y^{\prime}$ and $Z^{\prime}$ of $U$, where $m<n$. The sets $Y^{\prime}$ and $Z^{\prime}$ can be extended to sets $Y$ and $Z$, each of size $n$, by including the $n-m$ largest elements of $C$. Thus an $S$-expression over $Y^{\prime}$ (and thus over $Y$ ) will be decision-equivalent to the corresponding $S$-expression over $Z^{\prime}$ (and thus over $Z$ ).

The next lemma gives the mapping from free time-bounded algorithms to comparison algorithms.

Lemma 18. Fix $n$ and $t$. Let $\mathscr{A}$ be a free algorithm over ID space $X$ and alphabet $A$, where $X$ has at least $f(n, t)$ elements.

If $\mathscr{A}$ elects a leader in $t$ rounds, using at most $s$ messages in the worst case, then there exists a comparison algorithm $\mathscr{A}^{\prime}$, which elects a leader in $t$ rounds, using at most $s$ messages in the worst case.

Proof. By Lemma 17, there is a subset $U$ of $X$ of size at least $n$ such that $\mathscr{A}$ is a $t$-comparison algorithm on inputs of $U$. Consider any $S$-expression $L$, of depth less than $t$, with atoms in $X$. Define the value of the message decision function of $\mathscr{A}^{\prime}$ on this expression to be that of the message decision function of $\mathscr{A}$ on any $S$-expression $L^{\prime}$, with atoms from $U$, that is order-equivalent to $L$. Similarly, for any $S$-expression of depth at most $t$, with atoms in $X$, define membership in $E$ for $\mathscr{A}^{\prime}$ according to membership in $E$ for $\mathscr{A}$ of any order-equivalent $S$-expression with atoms in $U$. We define the message generation and decision functions so that $\mathscr{A}^{\prime}$ sends no messages after round $t-1$ and does not change any election status after round $t$.

Clearly algorithm $\mathscr{A}^{\prime}$ is a comparison algorithm. Since it simulates $\mathscr{A}$ on a sufficiently large subset of $X$, it can be seen to elect a leader in the same number of rounds, and with at most the same number of messages.

This lemma appears to be of much wider applicability than just to this work and Snir's. This result, or variants, should be very useful for the study of other orderinvariant problems on many different kinds of computation models. For example, see [12].
7.3 Lower Bound. Finally, we present our lower bound for time-bounded algorithms.

Theorem 19. Fix $n$ and $t$. Let $X$ be an arbitrary ID space with at least $f(n, t)$ elements. Let $\mathscr{A}$ be any algorithm over $X$ that elects a leader in a synchronous ring of size $n$, using at most time $t$. Then there is an execution, $e$, of $\mathscr{A}$ for which messages (e) is $\Omega(n \log n)$.

Proof. From Theorem 16, we know that there are constants $a$ and $b$ such that a comparison algorithm will have messages $(e) \geq a n \log n+b n$ for some execution $e$. Assume that there exists an algorithm $\mathscr{A}$ over $X$ that elects a leader in a synchronous ring of size $n$, using no more than time $t$ and fewer than $a n \log n+b n$ messages in the worst case. Then Lemma 18 implies that there exists a comparison algorithm that elects a leader in $t$ rounds and uses fewer than $a n \log n+b n$ messages in the worst case. This is a contradiction.

## 8. Remaining Questions

The general $\Omega(n \log n)$ bound that we have proved has a very small constant, $1 /\left(18 \log _{2} 9\right)$. In contrast, the best constant known for an upper bound is around $1.4[6,8]$. It remains to close this gap. For certain values of $n$, powers of 2 , we do have a narrower gap. It is possible that there are certain properties of the number $n$, for example, properties of its prime factorization, that affect the size of the constant. It would be interesting to understand these relationships.
acknowledgments. The authors thank Cynthia Dwork for pointing out the results of Snir, Maria Klawe for her encouragement in our attempts to obtain rings with replication symmetry, and Mark Tuttle for his comments on early versions of
the manuscript. Thanks also go to Mike Fischer and the referees for several suggestions on improving the presentation.

## REFERENCES

I. Angluin, D. Local and global properties in networks of processors. In Proceedings of the $12 t h$ Annual ACM Symposium on Theory of Computing (Los Angeles, Calif., Apr. 28-30). ACM, New York. 1980, pp. 82-93.
2. Attiya, C., Snir, M., and Warmuth, M. Computing on an anonymous ring. In Proceedings of the 4th Annual ACM Symposium on Principles of Distributed Computing (Minaki, Ontario, Canada, Aug. 5-7). ACM, New York, 1985, pp. 196-203.
3. Berge, C. Graphs and Hypergraphs. North-Holland, Amsterdam, 1973.
4. Burns, J. E. A formal model for message passing systems. Tech. Rep. TR-91. Indiana Univ., Bloomington, Ind., Sept. 1980.
5. Chang, E., and Roberts, R. An improved algorithm for decentralized extrema-finding in circular configurations of processes, Commun. ACM 22, 5 (May 1979), 281-283.
6. Dolev, D., Klawe, M., and Rodeh, M. An $O(n \log n)$ unidirectional distributed algorithm for extrema finding in a circle. J. Algorithms 3, 3 (Sept. 1982), 245-260.
7. Frederickson, G. N., and Lynch, N. The impact of synchronous communication on the problem of electing a leader in a ring. In Proceedings of the 16th Annual ACM Symposium on Theory of Computing (Washington, D.C., Apr. 30-May 2). ACM, New York, 1984, pp. 493-503.
8. Gallager, R. G., Humblet, P. A., and Spira, P. M. A distributed algorithm for minimumweight spanning trees. ACM Trans. Program. Lang. Syst. 5, 1 (Jan. 1983), 66-77.
9. Hirschberg, D. S., and Sinclair, J. B. Decentralized extrema-finding in circular configurations of processors. Commun. ACM 2311 (Nov. 1980), 627-628.
10. Ital, A., and Rodeh, M. Symmetry breaking in distributive networks. In Proceedings of 22nd Symposium on Foundations of Computer Science (Nashville, Tenn., Oct.), IEEE, New York, 1981, pp. 150-158.
11. LeLann, G. Distributed systems-Toward a formal approach. Information Processing, vol. 77. North Holland, Amsterdam, 1977. pp. 155-160.
12. Moran, S., Snir, M., and Manber, U. Applications of Ramsey's theorem to decision tree complexity. J. ACM 32, 4 (Oct. 1985), 938-949.
13. Peterson, G. L. An $O(n \log n)$ unidirectional algorithm for the circular extrema problem. $A C M$ Trans. Program. Lang. Syst. 4, 4 (1982), 758-762.
14. SNir, M. On parallel searching. Department of Computer Science, RR 83-21. Hebrew Univ. of Jerusalem, Jerusalem, Israel, June 1983.
15. Vitanyi, P. Distributed elections in an Archimedean Ring of Processors. In Proceedings of the 16th Annual ACM Symposium on Theory of Computing (Washington, D.C., Apr. 30-May 2). ACM, New York, 1984, pp. 542-547.
16. YaO, A. Should tables be sorted? J. ACM 28, 3 (July 1981), 615-628.

RECEIVED OCTOBER 1984; REVISED DECEMBER 1985; ACCEPTED JANUARY 1986


[^0]:    The work of G. N. Frederickson was supported by the National Science Foundation under grants MCS 82-01083 and DCR 83-20124. The work of N. Lynch was supported by NSF grants MCS 79-24370 and DCR 83-02391, U.S. Army Research Office contract DAAG29-84-K-0058, Office of Naval Research contract N00014-85-K-0168, and Advanced Research Projects Agency of the Department of Defense contracts N00014-75-C-0661 and N00014-83-K-0125.
    Authors' addresses: G. N. Frederickson, Department of Computer Sciences, Purdue University, West Lafayette, IN 47907; N. A. Lynch, Laboratory for Computer Science, MIT, Cambridge, MA 02139.
    Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

