A NEW LOWERBOUND TECHNIQUE FOR DISTRIBUTED EXTREMA FINDING ON RINGS OF PROCESSORS

Hans L. Bodlaender

RUU-CS-87-11
August 1987

Rijksuniversiteit Utrecht

Vakgroep Informatica

Budapestlaan 6 3584 CD Utrecht
Corr. adres: Postbus 80.012 3508 TA Utrecht
Telefoon 030-53 1454
The Netherlands

A NEW LOWERBOUND TECHNIQUE FOR DISTRIBUTED EXTREMA FINDING ON RINGS OF PROCESSORS

Hans L. Bodlaender

Technical Report RUU-CS-87-11
August 1987

Department of Computer Science
University of Utrecht
P.O. Box 80.012, 3508 TA Utrecht
The Netherlands

A NEW LOWERBOUND TECHNIQUE FOR DISTRIBUTED EXTREMA FINDING ON RINGS OF PROCESSORS*

Hans L. Bodlaender
Dept. of Computer Science, University of Utrecht
P.O. Box 80.012, 3508 TA Utrecht, the Netherlands

Abstract

With a new technique, using results from extremal graph theory, several new lowerbounds are derived for distributed extrema finding on rings of processors, where the ring size n is known in advance to the processors. A lowerbound of $\Omega(n \log n)$ is shown for the average number of messages, for unidirectional and bidirectional rings, for *any* ring size n, with the size of the index set I as small as cn, for any constant $c > 1$.

For unidirectional rings, the lowerbound for the average number of messages is improved to $\frac{1}{2} n \log n - O(n)$, requiring that n is a power of 2, and $|I| \geq n^2$. For bidirectional rings, we show that for all $\varepsilon > 0$, there are infinitely many n, such that the average number of messages sent on rings with fixed size n is at least $(\frac{1}{4} - \varepsilon) n H_n$, requiring that $|I| \geq n^2$.

1 Introduction

In this paper we consider the problem of finding a leader in an asynchronous ring of processors. Each processor is distinguished by a unique identification number, taken from some index set I. In this paper we assume that the size n of the ring is known in advance to the processors. There is no central controller. The problem is to design a distributed algorithm that "elects" a unique processor as leader (e.g. the highest numbered processor), using a minimum number of messages.

We assume that the processors work fully asynchronous and cannot use clocks or timeouts. Hence we can assume that the algorithm is message-driven: except for the first message upon initialization, a processor can only send messages as a result of the receipt of a message. We also assume that processors and the communication subsystem work error-free and that links work in a FIFO-manner.

There are basically two variants of the problem: the ring may be unidirectional (all messages go in one direction) or bidirectional (messages can go in both directions). For bidirectional algorithms, one has the variant where the ring has "a sense of direction", i.e. each processor has the same idea about "left" and "right", and the variant where

*A large part of this research was done, while the author was visiting the Laboratory of Computer Science of the Massachusetts Institute of Technology, with a grant from the Netherlands Organization for the Advancement of Pure Research (Z.W.O.).

processors do not have a sense of direction. We will assume the former case, which only strengthens the results.

Much work has been done to obtain good upper- and lowerbounds for the different variants of the problem.

Many bidirectional algorithms, using $O(n \log n)$ messages worst-case have been proposed [6,11,13,14,16,17,20]. Unidirectional algorithms, using $O(n \log n)$ algorithms can be found in [8,19]. In table 1, the best known upperbounds are summarized. None of these algorithm requires that processors know the ring size. (H_n is the n'th harmonic number, i.e. $H_n = \sum_{i=1}^{n} \frac{1}{i} \approx 0.69 n \log n$).

	average	worst-case
Unidirectional	nH_n [7]	$1.356n \log n + O(n)$ [8]
Bidirectional with sense of direction	$\frac{\sqrt{2}}{2} nH_n$ [4,10]	$1.356n \log n + O(n)$ [8]
Bidirectional without sense of direction	$\frac{\sqrt{2}}{2} nH_n$ [4,10]	$1.44n \log n + O(n)$ [16,17]

Table 1: Overview of upperbounds.

The first $\Omega(n \log n)$ lowerbound for the problem was obtained by Burns [6], for the worst-case number of messages on bidirectional rings without known ring size. Pachl, Korach and Rotem [18] obtained $\Omega(n \log n)$ lowerbounds for the average and worst-case number of messages on unidirectional and bidirectional rings without known ring size, and the worst-case number of messages on rings with known ring size. Similar lowerbounds, improving with a constant factor the results in [18], can be found in [2,3] and [12].

It has long been an open problem to determine the average number of messages on rings with a fixed ring size. Recently, Duris and Galil [9] obtained lowerbounds of $(\frac{1}{4} - \varepsilon) n \log n - O(n)$ for the average number of messages on unidirectional rings with fixed ring size, and $(\frac{1}{8} - \varepsilon) n \log n - O(n)$ for the average number of messages on bidirectional rings with fixed ring size. Their proof assumes that n is a power of 2, and requires that the size of the index set I is exponential in n.

In this paper we prove $\Omega(n \log n)$ lowerbounds for unidirectional and bidirectional rings with *any* fixed ring size n, where the index set I may be as small as cn, for any constant $c > 1$. For unidirectional rings we give an average case lowerbound of $\frac{1}{2} n \log n$ messages for rings with a fixed size n, with n a power of 2, and index set size $|I| \geq n^2$. For bidirectional rings, we show that for all $\varepsilon > 0$, there are infinitely many n, such that the average number of messages sent on rings with fixed size n is at least $(\frac{1}{4} - \varepsilon) n H_n$, for $|I| \geq n^2$.

Note that if $|I| - n$ is very small, then one can design algorithms which use less than $\Omega(n \log n)$ messages. For example, one can turn all processors with an identity, which is one of the $n-1$ smallest in I "inactive", and then run a variant of Petersons $1.44n \log n + O(n)$ unidirectional algorithm [19]. This gives an algorithm using $O(n \log(|I| - n))$ messages (worst-case). (This observation was made by Gerard Tel.)

This paper is organized as follows. In section 2 we give some definitions. Section 3 introduces all necessary definitions and results from extremal graph theory. Some new results are derived. In section 4 we give a simple lowerbound proof for the average number of messages on unidirectional rings with fixed size n, with n a power of 2, $|I| \geq cn$, c a constant > 1. The lowerbound is improved to $\frac{1}{2} n \log n - O(n)$ for $|I| \geq n^2$. In section 5 similar results (but with lower constants) are derived for bidirectional rings. In section 6 we prove an $\Omega(n \log n)$ lowerbound for the average number of messages on bidirectional (and

hence, also on unidirectional) rings with fixed ring size n, for arbitrary n, and $|I| \geq cn$.

2 Definitions

For an index set I, define $D(I)$ to be the set of finite, non-empty sequences of distinct elements of I. The concatenation of two strings $s = s_1 \cdots s_k$ and $t = t_1 \cdots t_l$ is denoted by $s \cdot t = s_1 \cdots s_k t_1 \cdots t_l$. The $l'th$ element of a string s is denoted by s_l. The length of a string $s = s_1 \cdots s_k$ is denoted by length(s)= k. The set of finite, non empty sequences of distinct elements of I with length k is denoted by $D_k(I) = \{s \in D(I) \mid \text{length}(s) = k\}$.

For the sake of analysis, we assume a (clockwise) numbering of the processors $1, 2, \ldots, n$. (n is the size of the ring; the numbering is not known to the processors). We say a ring is labeled with $s = s_1 \ldots s_n \in D_n(I)$, if for each i, $1 \leq i \leq n$, processor i has identity s_i.

Further we denote $X_k(I)$ to be the set of all sets of $\lfloor \frac{|I|}{k} \rfloor$ disjoint strings from $D_k(I)$, i.e. $X_k(I) = \{S \subseteq D_k(I) \mid |S| = \lfloor \frac{|I|}{k} \rfloor$ and $(\forall s, t \in S : s \neq t \Rightarrow \forall i, j \leq k : s_i \neq t_j))\}$.

For $k|n$, we say that a string $s \in D_n(I)$ is *derived* from $S \in X_k(I)$, if s is formed by concatenating $\frac{n}{k}$ different elements from S.

3 Definitions and results from extremal graph theory

In this section we review some results from extremal graph theory. The interested reader is referred to the book of Bollobás [5], for background, proofs, etc.

Define $\alpha(m, l)$ ($\bar{\alpha}(m, l)$) to be the maximum number of edges in a directed (undirected) graph with m vertices, that does not contain a cycle with length l, and let $\beta(m, l) = 1 - \frac{\alpha(m,l)}{m(m-1)}$.

Lemma 3.1
$\forall N, l, 3 \leq l \leq N : \alpha(N, l) \leq \bar{\alpha}(N, l) + \frac{1}{2}N(N-1)$.

Proof.
Let $G = (V, E)$ be a directed graph with $\bar{\alpha}(N, l) + \frac{1}{2}N(N-1) + 1$ edges. It follows that there are at least $\bar{\alpha}(N, l) + 1$ pairs of nodes v, w with $(v, w) \in E$ and $(w, v) \in E$. Hence G contains a cycle with length l. □

Theorem 3.2 [5]
Let $G = (V, E)$ be an undirected graph with $|V| = N$, $|E| = M$, and let $l \in \mathbb{N}^+$ be a positive natural number, such that $l \geq \lfloor \frac{1}{2}(N+3) \rfloor$, and $M > \binom{l-1}{2} + \binom{n-l+2}{2}$. Then G contains a cycle with length r, for every r, $3 \leq r \leq l$.

Corollary 3.3
Let $N > l \geq \frac{1}{2}N + 3$. Then

(i) $\bar{\alpha}(N, l) \leq \binom{l-1}{2} + \binom{N-l+2}{2}$.

(ii) $\alpha(N, l) \leq \binom{l-1}{2} + \binom{N-l+2}{2} + \frac{1}{2}N(N-1)$.

(iii) $\beta(N, l) \geq \frac{l}{N} - \frac{l^2}{N} \pm O(\frac{1}{N})$.

Proof.
(i), (ii) follow directly from theorem 3.2 and lemma 3.1.
(iii) can be derived as follows:

$$\begin{aligned}
\beta(N,l) &= 1 - \frac{\frac{1}{2}l(l-1)+\frac{1}{2}(N-l+2)(N-l+1)}{N(N-1)} - \frac{1}{2} \\
&= \frac{1}{2} - \frac{\frac{1}{2}N^2+\frac{1}{2}l^2-lN+3N-3l+2+\frac{1}{2}l^2-\frac{1}{2}l}{N(N-1)} \\
&= \frac{1}{2} - \frac{\frac{1}{2}N(N-1)+\frac{1}{2}N+l^2-lN+3N+3\frac{1}{2}l+2}{N(N-1)} \\
&= \frac{lN}{N(N-1)} - \frac{l^2}{N(N-1)} + O(\frac{1}{N}) \\
&= \frac{l}{N} - \frac{l^2}{N^2} + O(\frac{1}{N}).
\end{aligned}$$

□

Theorem 3.4 [5]
Let $G = (V, E)$ be an undirected graph with $|V| = N$, $|E| = M$. Let k be a natural number and let $M > 90kN^{1+1/k}$. Then G contains a cycle of length $2l$ for every integer l, $k \leq l \leq kn^{1/k}$.

Corollary 3.5
Let l be even; $4 \leq l \leq 4\sqrt{N}$. Then

(i) $\bar{\alpha}(N,l) \leq 180N\sqrt{N}$.

(ii) $\alpha(N,l) \leq 180N\sqrt{N} + \frac{1}{2}N(N-1)$.

(iii) $\beta(N,l) \geq \frac{1}{2} - \frac{180\sqrt{N}}{N-1}$.

Proof.
Use Theorem 3.4 with $k = 2$.

□

Next we derive some new results for graphs with two types of nodes. These results will be used in section 6.

Theorem 3.6
Let $G = (V, E)$ be an undirected graph, with $|V| = N$, $|E| = M$ and let $V = V_1 \cup V_2$, $V_1 \cup V_2 = \emptyset$, $|V_1| = N_1$, $|V_2| = N_2$.
Let l_1, l_2 be natural numbers, such that

$$\begin{aligned}
&l_1 \geq \tfrac{1}{2}N_1 + 3, \\
&l_2 \geq \tfrac{1}{2}N_2 + 3, \\
&\tfrac{1}{2}N(N-1) - M < \tfrac{1}{2}N_1(N_1-1) - \binom{l_1-1}{2} - \binom{N_1-l_1+2}{2}, \\
&\tfrac{1}{2}N(N-1) - M < \tfrac{1}{2}N_2(N_2-1) - \binom{l_2-1}{2} - \binom{N_2-l_2+2}{2}.
\end{aligned}$$

Let $3 \leq k_1 \leq l_1$, $3 \leq k_2 \leq l_2$. Suppose $\frac{1}{2}N(N-1) - M < \frac{1}{2}(k_1k_2)$.
Then G contains a cycle with exactly k_1 vertices from V_1 and k_2 vertices from V_2.

Proof.
First consider the subgraph of G, induced by V_1, $G[V_1]$. There are at most $\frac{1}{2}N(N-1)-M \leq \frac{1}{2}N_1(N_1-1) - \binom{l_1-1}{2} - \binom{N_1-l_1+2}{2} - 1$ unordered pairs (v,w), $v \neq w$, which correspond to a non-edge in G. It follows that $G[V_1]$ contains at least $\binom{l_1-1}{2} + \binom{N_1-l_1+2}{2} + 1$ edges, and hence, by theorem 3.2 it contains a cycle with k_1 vertices. With a similar argument one shows that $G[V_2]$ contains a cycle with k_2 vertices.

Now consider two fixed cycles, one in $G[V_1]$ with length k_1 and one in $G[V_2]$ with length k_2. Let the vertices in the cycle in V_1 be numbered $v_0, v_1, \ldots, v_{k_1-1}$; and the vertices in the cycle in V_2 be numbered $w_0, w_1, \ldots, w_{k_2-1}$.

Suppose now that G does not contain a cycle with exactly k_1 vertices from V_1 and k_2 vertices from V_2. Then for all i, j, $0 \leq i \leq k_1 - 1$, $0 \leq j \leq k_2 - 1 : (v_i, w_j) \in E \Rightarrow (v_{(i+1) \bmod k_1}, w_{(j+1) \bmod k_2}) \notin E$. If this is not the case, then one can construct the desired circuit, as in figure 3.

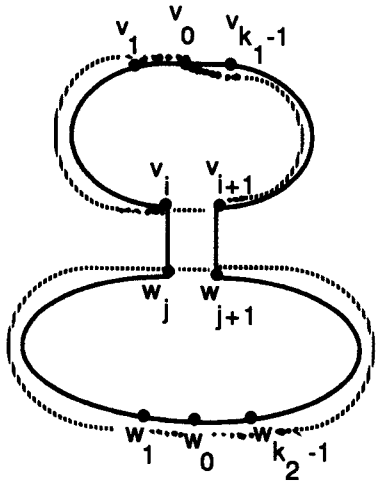

Figure 3.1.

It follows that there are at most $\frac{1}{2}(k_1k_2)$ pairs (v_i, w_j) with $(v_i, w_j) \in E$. Hence there are at least $\frac{1}{2}k_1k_2$ pairs (v_i, w_j) corresponding to a non-edge, so $\frac{1}{2}n(n-1) - |E| \geq \frac{1}{2}k_1k_2$. Contradiction. □

Corollary 3.7
Let $G = (V, E)$ be a directed graph, with $|V| = N$, $|E| = M$, and let $V = V_1 \cup V_2$, $V_1 \cap V_2 = \emptyset$, $|V_1| = N_1$, $|V_2| = N_2$.
Let l_1, l_2 be natural numbers, such that

$$l_1 \geq \tfrac{1}{2}N_1 + 3,$$
$$l_2 \geq \tfrac{1}{2}N_2 + 3,$$
$$N(N-1) - M < \tfrac{1}{2}N_1(N_1-1) - \binom{l_1-1}{2} - \binom{N_1-l_1+2}{2},$$
$$N(N-1) - M < \tfrac{1}{2}N_2(N_2-1) - \binom{l_2-1}{2} - \binom{N_2-l_2+2}{2}.$$

Let $3 \leq k_1 \leq l_1$, $3 \leq k_2 \leq l_2$. Suppose $N(N-1) - M < \frac{1}{2}(k_1k_2)$.
Then G contains a cycle with exactly k_1 vertices from V_1 and k_2 vertices from V_2.

Proof.
Similar to lemma 3.1. □

4 Lowerbounds for unidirectional rings with certain ring sizes

In this section we consider unidirectional rings, and assume that the ring size n is a power of 2. We prove a lowerbound of $(\frac{1}{c} - \frac{1}{c^2})n \log n - O(n)$ messages, for the average case on unidirectional rings, with index set I, with $|I| \geq cn$, $1 < c \leq 2$. For $|I| \geq n^2$, we prove a lowerbound of $\frac{1}{2}n \log n - O(n)$ messages, which improves a lowerbound of Duris and Galil [9].

For our analysis we first remark that as links operate in a FIFO-manner, the number of messages sent does not depend on the relative time it takes to send messages, in the unidirectional case. So we may as well assume that all processors start simultaneously at time 1, and each message takes unit time. As a consequence, it only depends on the id's of the $t-1$ processors, directly preceding a processor i, and its own id, whether or not processor i will send a message on time t. (This technique is very similar to techniques used in [18]).

Now consider some fixed ring size n and index set I. Let A be an asynchronous unidirectional leader finding algorithm for ring size n and index set I.

We may assume that after completion of the algorithm, every processor knows the identity of the leader. (Other variants differ in $O(n)$ messages, at most.)

Lemma 4.1
For all $r \in D_n(I)$, and $t < n-1$, there is at least one processor that sends a message at time t, on a ring labeled with r, when executing A, if $|I| \geq n+1$.

Proof.
Suppose not. Suppose processor i becomes the leading processor. At time t, processor $i-1$ (or n, if $i = 1$) cannot distinguish the case that processor i has identity r_i, or processor i has an identity, not in r. Contradiction. □

For all $k \leq \frac{1}{2}n$, we now define for each $S \in X_k(I)$ the following directed graph $G(S) = (S, E(S))$, by $E(S) = \{(s,t) \mid s,t \in S$; processor $2k$ will not send a message between time $k+1$ and $2k$, on a ring labeled with $r \in D_n(I)$, with $s = r_1 \cdots r_k$ and $t = r_{k+1} \cdots r_{2k}\}$.

Lemma 4.2
Let $k|n$, $k < \frac{1}{2}n$; let $S \in X_k(I)$. Then $G(S)$ does not contain a cycle with length $\frac{n}{k}$.

Proof.
Suppose $G(S)$ contains a cycle with length $\frac{n}{k}$, and let $s^1, \ldots, s^{\frac{n}{k}}$ be the successive nodes on this cycle. Let $r = s^1 \cdot s^2 \cdot \ldots \cdot s^{\frac{n}{k}}$. Now consider an execution of A on a ring labeled with r. (Note that $r \in D_n(I)$.)

It follows from lemma 4.1 that there is at least one processor that sends a message at time $2k$. So suppose processor $(ik+j)$, $0 \leq j \leq k-1$ sends a message at time $2k$. Then processor ik (or processor n, if $i = 0$), sends a message between time $k+1$ and

$2k$. It follows that $(s^{i-1}, s^i) = (r_{(i-2)k+1} \cdots r_{(i-1)k}, r_{(i-1)k+1} \cdots r_{ik}) \notin E(S)$. (Or, if $i = 1$, $(s^{\frac{n}{k}}, s^1) \notin E$). Contradiction. □

Theorem 4.3
Let $k|n$, $k < \frac{1}{2}n$. Then the number of messages, sent on a unidirectional ring with known ring size n between time $k+1$ and $2k$, averaged over all ring labelings $r \in D_n(I)$, is at least $\beta(\lfloor \frac{|I|}{k} \rfloor, \frac{n}{k}) \cdot n$.

Proof.
Consider some $S \in X_k(I)$. Since a non-edge in $G(S)$ corresponds to a message, sent by processor $2k$, between time $k+1$ and $2k$ the average number of messages sent by processor $2k$ between time $k+1$ and $2k$, over all rings, labeled with $r \in D_n(I)$, which are derived from S, is at least $\beta(\lfloor \frac{|I|}{k} \rfloor, \frac{n}{k})$. Note that each $r \in D_n(I)$ is derived from the same number of $S \in X_k(I)$. It follows that the average number of messages, sent by processor $2k$ between time $k+1$ and $2k$ is at least $\beta(\lfloor \frac{|I|}{k} \rfloor, \frac{n}{k})$. The result now follows by symmetry, because each processor can be taken as processor $2k$. □

We are now ready to prove the main results in this section.

Theorem 4.4
For all c, $1 < c < 2$, and all leader finding algorithms on unidirectional rings, where processors know the ring size n, the average number of messages sent on a ring of size n, n a power of 2, over all ring labelings $r \in D_n(I)$, with $|I| \geq cn$, is at least $(\frac{1}{c} - \frac{1}{c^2})n \log n - O(n)$.

Proof.
Denote the average number of messages, sent between time $2^i + 1$ and 2^{i+1}, over all $r \in D_n(I)$, by $av(2^i + 1, 2^{i+1})$. We now have the following lowerbound for the number of messages which must be estimated:

$$\begin{aligned}
\sum_{i=1}^{\log n-2} av(2^i+1, 2^{i+1}) &\geq \sum_{c=o}^{\log n-2} n \cdot \beta(\lfloor \frac{|I|}{2^i} \rfloor, \frac{n}{2^i}) \\
&\geq \sum_{i=1}^{\log n-2} n \left(\frac{\frac{n}{2^i}}{\lfloor \frac{|I|}{2^i} \rfloor} - \left(\frac{\frac{n}{2^i}}{\lfloor \frac{|I|}{2^i} \rfloor} \right)^2 - O(\frac{1}{n}) \right) \\
&\geq n \cdot \sum_{i=1}^{\log n-2} \left(\frac{n}{|I|} - (\frac{n}{|I|})^2 - O\left(\frac{\frac{n}{2^i}}{(\frac{|I|}{2^i})^2} \right) - O(\frac{1}{n}) \right) \\
&\geq n \left(\frac{1}{c} \log n - \frac{1}{c^2} \log n - \sum_{i=1}^{\log n-2} O(\frac{1}{2^i}) \right) - O(\log n) \\
&= (\frac{1}{c} - \frac{1}{c^2}) n \log n - O(n).
\end{aligned}$$

□

By taking a somewhat larger index set, one can improve the constant by a factor 2.

Theorem 4.5
For all leader finding algorithms on unidirectional rings, where processors know the ring size n, the average number of messages sent on a ring of size n, n a power of 2, over all ring labelings $r \in D_n(I)$, with $|I| \geq n^2$, is at least $\frac{1}{2}n \log n - O(n)$.

Proof.

$$\sum_{i=1}^{\log n-2} n\beta(\lfloor \frac{|I|}{2^i} \rfloor, \frac{n}{2^i}) \geq n \cdot \sum_{i=1}^{\log n-2} \left(\frac{1}{2} - \frac{180\sqrt{\frac{|I|}{2^i}}}{(\frac{|I|}{2^i}) - 1} \right) = \frac{1}{2} n \log n - O(n).$$

□

5 Lowerbounds for bidirectional rings for certain fixed ring sizes

In this section we consider bidirectional rings with fixed ring size n, with n a power of 2, or of the form $2 \cdot i!$.

The lowerbounds are of the type, where we average over all rings, labeled with strings $\in D_n(I)$, but where the delay times may be chosen in any manner, in order to obtain an as large as possible number of messages. All lowerbounds for the average number of messages for leaderfinding on asynchronous rings we know of, are of this type. Here we assume that all message delay times are equal, i.e. each message takes unit time. Further assume that when a processor receives two messages (from both neighbors) at the same moment, it handles the left one first. In this way we lose the implicit non-determinism, associated with asynchronous, bidirectional rings.

So we may assume that we have an asynchronous, message-driven algorithm, running on a synchronous ring. We again assume that all processors start to send at time 1. Note that it depends only on the id's of the processors with distance at most $t-1$ to processor i, whether or not processor i will send a message at time t or not.

Lemma 5.1
Let $|I| \geq n+1$. Then, for all $r \in D_n(I)$ and $t < \frac{1}{2}n$, there is at least one processor that sends a message at time t on a ring labeled r.

Proof.
Similar to lemma 4.1.

□

Now for all $k \leq \frac{1}{2}n$, $2|k$, $l \leq \frac{1}{2}k$ and each $S \in X_k(I)$ we define the following directed graph $H_l(S) = (S, E_l(S))$, by $E_l(S) = \{(s,t) \mid s,t \in S$; when a ring is labeled with $r \in D_n(I)$, with $s = r_1 \ldots r_k$, $t = r_{k+1} \ldots r_{2k}$, then none of the processors $\frac{1}{2}k+1, \frac{1}{2}k+2, \ldots, 1\frac{1}{2}k-1, 1\frac{1}{2}k$ sends a message at time $l\}$.

Lemma 5.2
Let $k|n$, $2|k$, $l \leq \frac{1}{2}k$, $k \leq \frac{1}{2}n$, $S \in X_k(I)$. Then $H_l(S)$ does not contain a simple directed cycle with length $\frac{n}{k}$.

Proof.
Suppose $H_l(S)$ contains a cycle with length $\frac{n}{k}$, say $s_1, \ldots, s_{\frac{n}{k}}$. Then on a ring labeled $s_1 \cdot s_2 \cdot \cdots \cdot s_{\frac{n}{k}}$ no processor sends a message at time l. Contradiction. □

Theorem 5.3
Let $k|n$, $2|k$, $l \leq \frac{1}{2}k$, $k \leq \frac{1}{2}n$. Then the average number of messages, sent at time l, over all rings, labeled with $r \in D_n(I)$, is at least $\frac{n}{k} \cdot \beta(\frac{|I|}{k}, \frac{n}{k})$.

Proof.
Consider some $S \in X_k(I)$. Since any non-edge in $H_l(S)$ corresponds to a message, sent at time l, by a processor in $\frac{1}{2}k+1 \ldots 1\frac{1}{2}k$, the average number of messages sent by processors $\frac{1}{2}k+1, \ldots, 1\frac{1}{2}k$ at time l is at least $\beta(\frac{|I|}{k}, \frac{n}{k})$. Again we argue that each $r \in D_n(I)$ is derived from the same number of $S \in X_k(I)$. It follows that the average number of messages, sent by processors $\frac{1}{2}k+1, \ldots, 1\frac{1}{2}k$ at time l, over all $r \in D_n(I)$ is at least $\beta(\frac{|I|}{k}, \frac{N}{k})$. The result now follows by symmetry, as every k successive processors can be taken as processors $\frac{1}{2}k+1, \ldots, 1\frac{1}{2}k$. □

Theorem 5.4
For all c, $1 < c < 2$, and all leader finding algorithms on bidirectional rings, where processors know the ring size n, the average number of messages sent on a ring of size n, n a power of 2, over all ring labelings $r \in D_n(I)$, with $|I| \geq cn$, is at least $\frac{1}{4}(\frac{1}{c} - \frac{1}{c^2})n \log n - O(n)$.

Proof.
It follows from theorem 5.3 that between times $\frac{1}{4}k+1$ and $\frac{1}{2}k$, at least $\beta(\frac{|I|}{k}, \frac{n}{k}) \cdot \frac{n}{4}$ messages are sent, on the average over all ring labelings $r \in D_n(I)$. Now the result follows, similar as in theorem 4.4. □

Similar as in theorem 4.5, one can improve the constant by taking $|I| \geq n^2$. In this way one obtains basically the same lowerbound as Duris and Galil [9], with the main difference that $|I|$ is here polynomial instead of exponential in n.

Theorem 5.5
For all leader finding algorithms on bidirectional rings, where processors know the ring size n, the average number of messages sent on a ring of size n, n a power of 2, over all ring labelings $r \in D_n(I)$, with $|I| \geq n^2$, is at least $\frac{1}{8}n \log n - O(n)$.

For n of the form $2m!$, we can obtain lowerbounds with (asymptotically) a better constant factor. Define for n even: $f(l,n) = \min\{k \geq 2l \mid 2|k \text{ and } k|n\}$.

Lemma 5.6
Let $n = 2(m!)$. Then $\sum_{l=1}^{\frac{1}{2}n} \frac{1}{f(l,n)} \geq \frac{1}{2}H_n - O(m)$.

Proof.
Write

$$\sum_{l=1}^{\frac{1}{2}n} \frac{1}{f(l,n)} = \sum_{i=1}^{m-1} \sum_{j=1}^{i} \sum_{l=j\cdot(\frac{\frac{1}{2}n}{(i+1)!})+1}^{(j+1)(\frac{\frac{1}{2}n}{(i+1)!})} \frac{1}{f(l,n)}$$

$$\geq \sum_{i=1}^{m-1} \sum_{j=1}^{i} (\frac{\frac{1}{2}n}{(i+1)!}) \cdot \frac{1}{(j+1)n(i+1)!}$$

$$= \sum_{i=1}^{m-1} \frac{1}{2}(H_i - 1)$$

$$\geq \sum_{i=1}^{m-1} \frac{1}{2} ln(i) - O(m)$$

$$= \tfrac{1}{2} ln((m-1)!) - O(m)$$

$$= \tfrac{1}{2} H_n - O(m).$$

($ln(x)$ denotes the logarithm of x to the base e.)
Note that $\sum_{l=1}^{\frac{1}{4}n} \frac{1}{f(l,n)} = \sum_{l=1}^{\frac{1}{2}n} \frac{1}{f(l,n)} \pm O(1)$. □

Theorem 5.7
For every $\varepsilon > 0$, there are infinitely many $n \in \mathbb{N}^+$, such that for all leader finding algorithms on bidirectional rings, where processors know the ring size n, the average number of messages sent on a ring with size n, over all ring labelings $r \in D_n(I)$, with $|I| \geq n^2$, is at least $(\frac{1}{4} - \varepsilon) n H_n$.

Proof.
It follows from theorem 5.3 that one can estimate the desired average by

$$\geq \sum_{l=1}^{\frac{1}{4}n} \frac{n}{f(l,n)} \cdot \beta\left(\frac{|I|}{f(l,n)}, \frac{n}{f(l,n)}\right)$$

$$\geq \sum_{l=1}^{\frac{1}{4}n} \frac{n}{f(l,n)} \cdot \left(\frac{1}{2} - \frac{180\sqrt{\frac{|I|}{f(l,n)}}}{\frac{|I|}{f(l,n)-1}}\right).$$

Note that

$$\sum_{l=1}^{\frac{1}{4}n} \frac{180\sqrt{\frac{|I|}{f(l,n)}}}{\frac{|I|}{f(l,n)-1}} = O(1).$$

When we take n of the form $2 \cdot (m!)$, then from lemma 5.6 and the observation that

$$\sum_{l=1}^{\frac{1}{4}n} \frac{1}{f(l,n)} = \sum_{l=1}^{\frac{1}{2}n} \frac{1}{f(l,n)} - O(1),$$

it follows that

$$\sum_{l=1}^{\frac{1}{4}n} \frac{n}{f(l,n)} \cdot \frac{1}{2} = \frac{1}{4} n H_n - O(m).$$

The result follows now easily by taking m large enough by given $\epsilon > 0$. □

Note that $\frac{1}{4} n H_n \approx 0.173 n \log n$.

6 Lowerbounds for arbitrary fixed ring sizes

In this section we prove an $\Omega(n \log n)$ lowerbound on the average number of messages on bidirectional rings with *any* fixed ring size n (so e.g. not only for n a power of 2). The same result for unidirectional rings follows directly as a corollary. Again $|I|$ may be as small as cn, for any constant $c > 1$.

We will assume that $|I| = cn$, with $1 < c < 2, c$ a constant. For larger index sets, the result follows easily from the result for smaller index sets.

Now suppose $n \in \mathbf{N}^+$ is given. We will use the following lemma, which can be found in [1].

Lemma 6.1
Let p, q be two positive integers, such that $(p, q) = 1$, i.e. p and q are relatively prime. Then, for all n, there exist integers r, s, such that $rp + sq = n$ and $|r - s| \leq \frac{(p+q)}{2}$.

We use this lemma to derive the following result.

Lemma 6.2
Let c^1 be a constant, with $0 < c^1 < 1$. Let $c^2 = \frac{(c^1)^2+1}{1-(c^1)^2}$. For all $k < c^1\sqrt{n}$, there are l_1, l_2, such that $l_1 k + l_2(k+1) = n$ and $(0 < l_1 \leq l_2 \leq c^2 l_1$ or $0 < l_2 \leq l_1 \leq c^2 l_2)$.

Proof.
From lemma 6.1 it follows that one can find l_1, l_2, such that $l_1 k + l_2(k+1) = n$, and $|l_1 - l_2| \leq \frac{(2k+1)}{2}$, hence $|l_1 - l_2| \leq k$. Also it follows that $\frac{n}{k} \geq l_1 + l_2 \geq \frac{n}{k+1} \geq \frac{1}{c^1}\sqrt{n} \geq \frac{1}{(c^1)^2} k$. If $l_1 \leq l_2$ then $l_1 - l_2 \leq (c^1)^2(l_1 + l_2)$, hence $l_2 \leq \left(\frac{(c^1)^2+1}{1-(c^1)^2}\right) l_1$. If $l_2 \leq l_1$, then similarly $l_1 \leq c^2 l_2$. □

We make the same assumptions on the message delays, etc. as in section 5. Let $K \leq \frac{1}{2}\lfloor c^1\sqrt{n} \rfloor$ be some fixed "time", and let $k = 2K$. Let l_1, l_2 as indicated in lemma 6.2 be given. Let $c^3 = \lfloor \frac{|I|}{n} \rfloor$. Note that $c^3 = O(c)$, i.e. is bounded by constants.

We now introduce the concept of *good string sets*. A good string set S is a collection of strings $\in D_k(I) \cup D_{k+1}(I)$, such that

- S contains exactly $c^3 \cdot l_1$ strings of length k, i.e. $\in D_k(I)$

- S contains exactly $c^3 \cdot l_2$ strings of length $k+1$, i.e. $\in D_{k+1}(I)$
- All strings in S are disjunct, i.e. $\forall s,t \in S, i \leq \text{length(s)}, j \leq \text{length(t)} : s \neq t \Rightarrow s_i \neq t_j$.

Let $Y_k(I)$ denote the set of all good string sets.

For a good string set S, define a graph $G_S = (S, E_S)$, with $E_S = \{(s,t) \mid$ On a ring, with a consecutive part labeled by $s \cdot t$, no processor with label in $s_{K+1} \ldots s_{\text{length}(s)} t_1 \ldots t_{\text{length}(t)-K}$ sends a message at time $K\}$.

Lemma 6.3
G_S does not contain a cycle with l_1 nodes, representing a string with length k_1 and l_2 nodes, representing a string with length k_2.

Proof.
Suppose $s_1, \ldots, s_{l_1+l_2}$ is such a cycle. Then $\text{length}(s_1 \cdot \ldots \cdot s_{l_1+l_2}) = l_1k_1 + l_2k_2 = n$, and no processor on a ring labeled with $s_1 \cdot \ldots \cdot s_{l_1+l_2}$ sends a message at time K. □

Lemma 6.4
Let $S \in Y_k(I)$. Let $c^4 = \frac{1}{2(1-c^3)^2(c^3)^2}$. Then $|S|(|S|-1) - |E(S)| \geq c^4|S|^2$.

Proof.
From lemma 6.3 and theorem 3.6 it follows that one of the following 3 cases must hold:

1. $|S|(|S|-1) - |E(S)| \geq \frac{1}{2}(c^3l_1)(c^3l_1 - 1) - \binom{l_1}{2} - \binom{c^3l_1 - l_1 + 2}{2}$.
2. $|S|(|S|-1) - |E(S)| \geq \frac{1}{2}(c^3l_2)(c^3l_2 - 1) - \binom{l_2}{2} - \binom{c^3l_2 - l_2 + 2}{2}$.
3. $|S|(|S|-1) - |E(S)| \geq \frac{1}{2}l_1l_2$.

Case 1. $|S|(|S|-1) - |E(S)| \geq \frac{1}{2}(c^3l_1)(c^3l_1 - 1) - \binom{l_1}{2} - \binom{c^3l_1 - l_1 + 2}{2}$

$= \frac{1}{2}((c^3)^2 - 1 - (c^3-1)^2)l_1^2 \pm O(l_1)$

$= c^3(l_1)^2 \pm O(l_1)$

$\geq c^3 \cdot (\frac{l_1+l_2}{(1+c^2)^2})^2 \geq c^4|S|^2$.

Case 2. Similar as case 1.

Case 3. $|S|(|S|-1) - |E(S)| \geq \frac{1}{2}l_1l_2 \geq \frac{1}{2}\left(\frac{l_1+l_2}{1+c^2}\right)^2 \geq c^4|S|^2$. □

For a good string set S, the set of ring labelings that can be derived from S is defined by

$N(S) = \{s^1 \cdot s^2 \cdot \ldots \cdot s^{l_1+l_2} \in D_n(I) \mid s^1, \ldots, s^{l_1+l_2} \in S$; exactly l_1 of the strings $s^1, \ldots, s^{l_1+l_2}$ have length k; exactly l_2 of the strings $s^1, \ldots, s^{l_1+l_2}$ have length $k+1\}$.

We denote the number of messages sent at time K on a ring, labeled with $r \in D_n(I)$ by $M(K,r)$. For $s \in N(S)$, we denote the i'th string from s, that is used to form s by s^i, i.e. $s = s^1 \cdot s^2 \cdot \ldots \cdot s^{l_1+l_2}$, with $s^1, s^2, \cdots, s^{l_1+l_2} \in S$.

Lemma 6.5
Let $S \in Y_k(I)$, and $s \in N(S)$.
Then $M(K,s) \geq \frac{1}{2}|\{i \in \{1,\cdots,l_1+l_2\} \mid (s^i, s^{i\oplus 1}) \notin E(S)\}|$, where $i \oplus 1 = i+1$, if $i < l_1+l_2$ and $(l_1+l_2) \oplus 1 = 1$.

Proof.
Let $i \in \{1,\cdots,l_1+l_2-1\}$ be given. (For $i = l_1+l_2$, the analysis is similar.) Suppose $(s^i, s^{i+1}) \notin E$. By definition, there is at least one processor with identity in $\{(s^i)_{K+1}, \cdots, (s^i)_{\text{length}(s^i)}, (s^{i+1})_1, \cdots, (s^{i+1})_{\text{length}(s^{i+1})-K}\}$, that sends a message at time K on a ring labeled s. Note that each message can be counted at most twice in this way. □

Denote $A_1 = \{(s,t) \mid s,t \in S, s \neq t, \text{length}(s) = \text{length}(t) = k\}$,
$A_2 = \{(s,t) \mid s,t \in s, s \neq t, \text{length}(s) = k, \text{length}(t) = k+1\}$,
$A_3 = \{(s,t) \mid s,t \in S, s \neq t, \text{length}(s) = k+1, \text{length}(t) = k\}$ and
$A_4 = \{(s,t) \mid s,t \in S, s \neq t, \text{length}(s) = \text{length}(t) = k+1\}$.
Denote $F(S) = \{(s,t) \mid s,t \in S, s \neq t, (s,t) \notin E(S)\}$.

Lemma 6.6

(i) $|\{r \in N(S) \mid (r^1, r^2) \in A_1\}| = O(|N(S)|)$.
(ii) $|\{r \in N(S) \mid (r^1, r^2) \in A_2\}| = O(|N(S)|)$.
(iii) $|\{r \in N(S) \mid (r^1, r^2) \in A_3\}| = O(|N(S)|)$.
(iv) $|\{r \in N(S) \mid (r^1, r^2) \in A_4\}| = O(|N(S)|)$.

Proof.
(i) Note that $|N(S)| = \binom{cl_1}{l_1} \cdot \binom{cl_2}{l_2} \cdot (l_1+l_2)!$, and $|\{r \in N(S) \mid \text{length}(r^1) = \text{length}(r^2) = k\}|$ $= 2\binom{cl_1}{2} \cdot \binom{cl_1-2}{l_1-2}\binom{cl_2}{l_2} \cdot (l_1+l_2-2)! = |N(S)| \cdot \frac{(l_1+l_2)(l_1+l_2-1)}{l_1(l_1-1)}$. By lemma 6.2 the result follows.
(ii) (iii)(iv) Similar. □

Lemma 6.7
There exists a constant c^5, such that for all $S \in Y_k(I)$: $|\{r \in N(S) \mid (r^1, r^2) \notin E(S)\}| \geq c^5|N(S)|$. ($c^5$ does not depend on k or n.)

Proof.
Observe that $|\{r \in N(S) \mid (r^1, r^2) \in A_1 \cap F(S)\}| = |\{r \in N(S) \mid (r^1, r^2) \in A_1\}| \cdot |A_1 \cap F(S)|/|A_1|$. (The argument here is that every pair (r^1, r^2), with $\text{length}(r^1) = k$ and $\text{length}(r^2) = k$ will appear as often as "start" of an r, derived from S.) The same observation is valid for A_2, A_3 and A_4. It now follows that

$$|\{r \in N(S) \mid (r^1, r^2) \notin E(S)\}| = \sum_{j=1}^{4} |\{r \in N(S) \mid (r^1, r^2) \in A_j \cap F(S)\}|$$

$$
\begin{aligned}
&= \sum_{j=1}^{4} \frac{|\{r \in N(S) \mid (r^1, r^2) \in A_j\}| \cdot |A_j \cap F(S)|}{|A_j|} \\
&= \sum_{j=1}^{4} \frac{O(|N(S)|)}{O(|S|^2)} \cdot |A_j \cap F(S)| \\
&= \frac{O(|N(S)|)}{O(|S|^2)} \cdot |(A_1 \cup A_2 \cup A_3 \cup A_4) \cap F(S)| \\
&= O(|N(S)|).
\end{aligned}
$$

(Use lemma 6.2, 6.4 and 6.6.) □

Note that, by symmetry, it also follows that for all $S \in Y_k(I)$, $i \in \{1, \cdots, l_1 + l_2\}$: $|\{r \in N(S) \mid (r^i, r^{i \oplus 1}) \notin E(S)\}| \geq c^5 |N(S)|$, where $i \oplus 1 = i + 1$ for $i \neq l_1 + l_2$, and $(l_1 + l_2) \oplus 1 = 1$.

Now we are ready to prove the main result of this section.

Theorem 6.8
For all $c > 1$, there exists a $C > 0$, such that for all $n \in \mathbf{N}^+$, and all leader finding algorithms on bidirectional (or unidirectional) rings where processors know the ring size n, the average number of messages sent, over all ring labelings $r \in D_n(I)$, with $|I| \geq cn$, is at least $Cn \log n$.

Proof.
Consider a good string set $S \in Y_k(I)$. The total number of messages sent at time K, over all $s \in D_n(I)$ derived from S, is

$$
\sum_{s \in N(S)} M(K, s) \geq \tfrac{1}{2} \sum_{s \in N(S)} |\{i \in \{1, \cdots, l_1 + l_2\} \mid (s^i, s^{i \oplus 1}) \notin E(S)\}|
$$

$$
= \tfrac{1}{2} \sum_{i=1}^{l_1 + l_2} |\{s \in N(S) \mid (s^i, s^{i \oplus 1}) \notin E(S)\}| \geq (l_1 + l_2) \cdot \tfrac{c^5}{2} \cdot |N(S)|.
$$

Hence, the average number of messages sent at time K, over all $S \in D_n(I)$, derived from S is $\Omega(l_1 + l_2) = \Omega(\frac{n}{K})$. As each $S \in D_n(I)$ is derived from the same number of good string sets, it follows that the average number of messages, sent at time K over all $s \in D_n(I)$ is $\Omega(\frac{n}{K})$.

As this is valid for all $K < \frac{1}{2} \lfloor c^1 \sqrt{n} \rfloor$, the theorem now follows. □

Acknowledgements

This work benefited very much from suggestions of and discussions with Zvi Galil and Manfred Warmuth.

[16] van Leeuwen, J., and R.B. Tan, *An improved upperbound for decentralized extrema-finding in bidirectional rings of processors*, Tech. Rep. RUU-CS-85-23, Dept. of Computer Science, Univ. of Utrecht, Utrecht, 1985. To appear in Distributed Computing.

[17] Moran, S., M. Shalom, and S. Zaks, *An algorithm for distributed leader finding in bidirectional rings without common sense of direction*, Tech. Rep. Technion, Haifa, 1985.

[18] Pachl, J., E. Korach, and D. Rotem, *Lowerbounds for distributed maximum-finding algorithms*, J. ACM 31 (1984) 905-918.

[19] Peterson, G.L., *An $O(n \log n)$ unidirectional algorithm for the circular extrema problem*, ACM Trans. Prog. Lang. & Syst. 4 (1982) 758-762.

[20] Santoro, N., E. Korach, and D. Rotem, *Decentralized extrema-finding in circular configurations of processors: an improved algorithm*, Congr. Numer. 34 (1982) 401-412.