

The Reconfigurable Ring of Processors: Fine-Grain Tree-Structured Computations

Arnold L. Rosenberg, *Fellow, IEEE*, Vittorio Scarano,
and Ramesh K. Sitaraman, *Member, IEEE*

Abstract—We study fine-grain computation on the Reconfigurable Ring of Processors (*RRP*), a parallel architecture whose processing elements (PEs) are interconnected via a multiline reconfigurable bus, each of whose lines has one-packet width and can be configured, independently of other lines, to establish an arbitrary PE-to-PE connection. We present a “cooperative” message-passing protocol that will, in the presence of suitable implementation technology, endow an *RRP* with message latency that is logarithmic in the number of PEs a message passes over in transit. Our study focuses on the computational consequences of such latency in such an architecture. Our main results prove that: 1) an N -PE *RRP* can execute a sweep up or down an N -leaf complete binary tree in time proportional to $\log N \log \log N$; 2) a broad range of N -PE architectures, including N -PE *RRPs*, require time proportional to $\log N \log \log N$ to perform such a sweep.

Index Terms—Communication protocols, dynamically reconfigurable parallel architectures, fine-grain parallel computing, rings of processors, tree-sweep algorithms.

1 INTRODUCTION

1.1 Overview

WE study fine-grain computation on the *Reconfigurable Ring of Processors (RRP)*, a parallel architecture whose processing elements (PEs) are interconnected via a multiline reconfigurable bus: Each line of the bus has one-packet width and can be configured, independently of other lines, to establish an arbitrary PE-to-PE connection.

Our study is inspired by a novel strategy for

- 1) designing the bus of an *RRP* and for
- 2) passing one-word messages along the lines of the bus in a way that yields message latency that is *logarithmic* in the number of PEs the message traverses—at least for (MOS) wafer-scale implementations of *RRPs*.

We are currently working on estimating the technological parameters that would enable our *Cooperative Message Transmission (COMET)* strategy to be realized. Our goal in the current paper, as in its companion paper [7], is to understand the computational consequences of a fine-grain logarithmic delay model for *RRPs*: Sufficiently good news would provide powerful motivation for paying the technological cost that would enable the model to be realized. Henceforth, we focus only on *RRPs* that have been implemented so as to achieve logarithmic message latency—perhaps, but not necessarily, via *COMET*. The interested reader should see [6] for another theoretical study of reconfigurable architectures with logarithmic message latency.

• A.L. Rosenberg and R.K. Sitaraman are with the Department of Computer Science, University of Massachusetts, Amherst, MA 01003.
E-mail: {rsnrbg, ramesh}@cs.umass.edu.

• V. Scarano is with the Dipartimento di Informatica ed Applicazioni “R.M. Capocelli,” Università di Salerno, 84081 Baronissi (SA), Italy.
E-mail: vitsca@dia.unisa.it.

Manuscript received 21 Mar. 1995.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number 105198.

In the remainder of this section, we describe, in turn, our main results, the detailed architecture of *RRPs*, and the *COMET* message-passing protocol. The subsequent sections of the paper describe our results and their proofs in detail.

1.2 Our Results

Our results are parameterized by the number of PEs in an *RRP* and the number of lines in its bus. We denote an N -PE ring whose bus has L lines an (N, L) -*RRP* when we want to make all parameters explicit, an N -*RRP* when we want to concentrate only on the number of PEs, and an *RRP* when the specific parameter values are not consequential.

In [7], we showed that *RRPs* can efficiently execute any *normal hypercube algorithm*,¹ a class which contains efficient algorithms for such diverse problems as sorting, matrix multiplication, and the Fast Fourier Transform. The upper-bound results of the current paper, which appear in Section 2, show that *RRPs* can perform, with close to optimal efficiency, any *leveled tree-structured algorithm (LTS algorithm)*, i.e., any algorithm that can be specified in terms of a fixed number of complete up- and/or down-sweeps on complete binary trees, performing a unit-time task each time a tree-node is encountered. Problems that can be solved by such algorithms include (segmented versions of) broadcast and accumulation (or, reduction), which are one-sweep algorithms, and parallel-prefix (or, scan), which is a two-sweep algorithm. When executed on an (N, L) -*RRP*, our algorithm performs each sweep in time bounded above by

$$\frac{\log^2 N}{\log L} + \log N \log \log L. \quad (1.1)$$

Thus, for values of L smaller than $N^{1/\log \log N}$, the time T_{SWEEP} for each sweep by the algorithm satisfies the inequality:

1. See, e.g., [4] for definitions.

$$T_{\text{Sweep}} \leq \frac{2 \log^2 N}{\log L},$$

while, for larger values of L , the time satisfies the following inequality, up to low-order terms:

$$T_{\text{SWEEP}} \leq 2 \log N \log \log N.$$

It follows that an N -RRP, endowed with sufficient communication bandwidth (as measured by the number of buslines), can execute LTS algorithms almost as fast as can an N -node hypercube of commensurate-power PEs: The slowdown incurred is only roughly $2 \log \log N$.

REMARK. Since we envision only chip- or wafer-scale RRP, the values of N will always be moderate, no more than, say, a few hundred. This small range of values for N has two consequences for our study. First, for such N , a bus with $N^{1/\log \log N}$ lines, as mandated when we discuss “larger values of L ,” is quite within the realm of feasibility. Second, the limited range of N renders asymptotic bounds (which begin to hold only “eventually”) inappropriate. This second fact has led us to use the evocative, but imprecise, phrase “low-order terms” when describing our bounds; of course, the analyses leading to the bounds are quite precise.

In Section 3, we show that the algorithms of Section 2 are within constant factors of optimal, in a variety of senses. Throughout this section, we use (one-to-all) broadcast as the prototypical one-sweep LTS algorithm, since it is—in a sense made precise in Section 4—the simplest such algorithm. In Section 3.1, we show that the time for performing a one-to-all broadcast on an (N, L) -RRP is bounded below by

$$T_{\text{Broadcast}} \geq \frac{1 \log^2 N}{4 \log L}.$$

Since this bound trivializes when L is as large as N^ϵ (because even communicating between PEs 0 and $N/2$ takes time proportional to $\log N$), we were motivated to find a nontrivial lower bound that holds for *all* RRP, no matter how many buslines they have. In Section 3.2, we prove such a bound: The time required for *any* N -RRP—no matter how many buslines it has—to perform a one-to-all broadcast is no smaller than

$$T_{\text{Broadcast}} \geq \frac{1}{15} \log N \log \log N.$$

This lower bound is surprisingly general: Not only does it apply to all N -RRPs, it applies also, with minor adaptation, to a much broader class of logarithmic-latency architectures. One can view this fact as suggesting that, at least for LTS algorithms, one could not gain appreciably—i.e., by more than a constant factor—if one replaced the ring-structured topology of RRP by a more highly connected topology, such as a mesh.

The extendibility of the lower bound in Section 3.2 both inspires and lays the technical groundwork for our very general lower bound in Section 4. We focus in that section on the following broad class of algorithms for *any parallel architecture that uses point-to-point communication*. We term

an algorithm *nontrivial* if it requires some one PE of the architecture to receive and/or send information—either directly or indirectly—to and/or from all other PEs. We prove that, to within constant factors, no nontrivial algorithm can be performed more efficiently than the (one-to-all) broadcast operation. As a consequence, any N -PE architecture of the sort discussed in Section 3.2 requires time proportional to $\log N \log \log N$ to execute any nontrivial algorithm. We stress that this lower bound is a fundamental limitation imposed by the logarithmic-latency model and is independent of specific architectural implementations.²

1.3 The Abstract RRP Architecture

1.3.1 A Static View

An (N, L) -RRP is a SIMD architecture that comprises N identical PEs,³ P_0, P_1, \dots, P_{N-1} , which we view as placed with equal spacing around a circle (see Fig. 1). There is a “bundle” of L lines, each having one-packet-width, passing outside the circle “over” the PEs and “through” the switches that provide the reconfiguration capability. Each PE of an (N, L) -RRP has L associated *communication sub-PEs* (CPEs) to help it manage the message traffic on its bus; specifically, CPE CP_{ij} controls switches on line j that allow PE P_i to participate (either as a source, a destination, or an intermediate node) in a dedicated point-to-point path along that line.

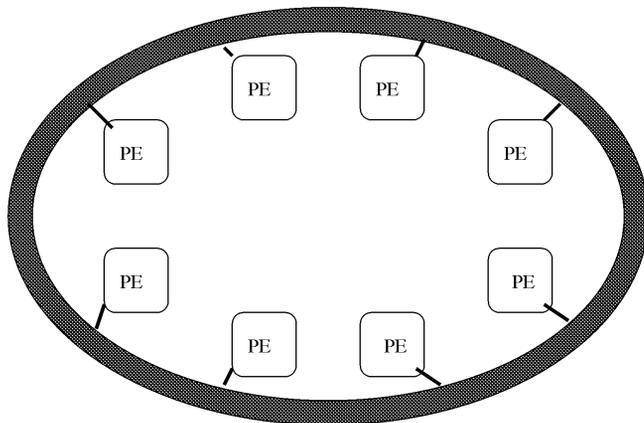


Fig. 1. An abstract view of an RRP.

We impose certain limitations on RRP, in order to minimize the technological resources required to implement them. First, as is implicit in the fact that buslines form *point-to-point* paths between pairs of communicating PEs, we assume that each message has exactly one sending PE and exactly one receiving PE; in particular, we do not support any “wired-or” or multireader bus capability. Second, we have our RRP observe a *single-port* communication regimen: In a single step, a PE can send at most one message and receive at most one (possibly, though not necessarily, involving distinct buslines). A word about these restrictions is in order. The constraint of single-port communication may well decrease the efficiency of RRP on a

2. It is instructive to see how our model and results compare with those of [6].

3. Throughout, when we discuss the PEs of an N -RRP, all PE-indices are computed modulo N .

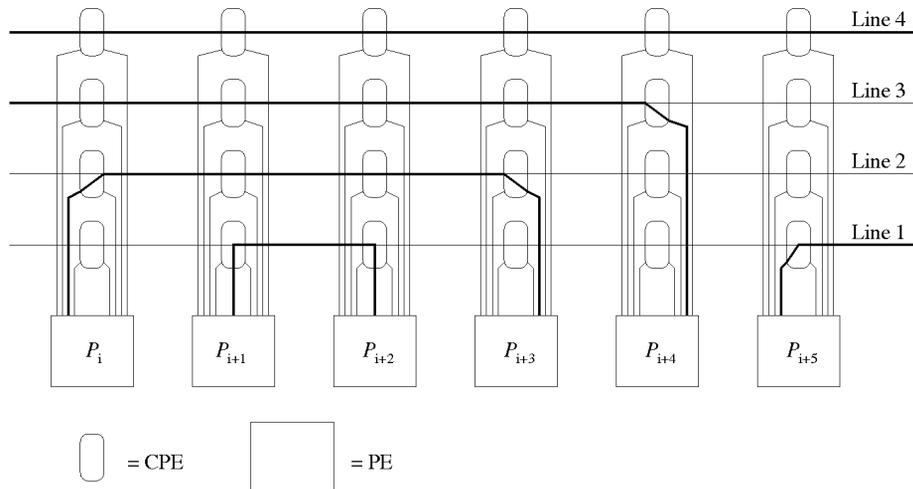


Fig. 2. A section of a bus of an $(N, 4)$ -RRP.

broad range of computational problems; we have yet to study the multiport version of our model. However, the absence of a multireader capability has less overall impact than one might initially expect. Clearly, a multireader bus capability would enable an N -RRP to perform one-to-all broadcasts in time proportional to $\log N$; however, this speedup would likely incur nonnegligible costs, as the multiple voltage drains would require either larger drivers, hence, a smaller RRP for a fixed amount of real estate, or slower CPEs, hence, a larger constant of proportionality on the time-function. Moreover, building on the techniques of Sections 3 and 4, one can show that a multireader bus capability would *not* accelerate operations such as accumulation, hence, would have only small impact on operations such as parallel-prefix.

The SIMD regimen observed by our RRP allows switch settings to be computed centrally and downloaded to the CPEs; hence, we shall not comment further on this aspect of RRP operation.

For reasons that will become clear when we describe the COMET message-passing protocol (in Section 1.4), we study only “fine-grain” computations by RRP, i.e., ones in which every inter-PE communication consists of a single one-packet message which is sent along a single busline.

Finally, to simplify algorithm specification and analysis, we assume henceforth that our RRP have numbers of PEs N and numbers of buslines L that are powers of 2; in particular, we shall write (when convenient) $N = 2^n$ and $L = 2^\ell$. These assumptions, which can be avoided by clerical modifications to our development, will affect only small additive terms in our bounds.

NOTATION. For any pair of PE-indices i and $j \neq i$, we denote by $B(i, j)$ the *block* of PEs $\{P_i, P_{i+1}, \dots, P_j\}$.

1.3.2 Path Formation and Communication

The first step in performing a communication between PEs P_i and P_j (in either direction) is to establish a dedicated path that connects these two PEs. This is achieved as follows: The SIMD controller appropriates a segment of some

busline which runs above either block $B(i, j)$ or block $B(j, i)$ and which is not currently used by any other communication. (Presumably, but not necessarily, if segments above both blocks are available, then the controller will choose the shorter one.) Say, for definiteness, that the available segment is a portion of busline k that runs above block $B(i, j)$. The controller then establishes the following connections, using the switches in the CPEs; see Fig. 2.

- PE P_i is connected to line k , via CPE $CP_{i,k}$;
- PE P_j is connected to line k , via CPE $CP_{j,k}$;
- each CPE $CP_{h,k}$ where $i \leq h < j$ is connected to CPE $CP_{h+1,k}$.

(If no busline is free above either block $B(i, j)$ or $B(j, i)$, then the message transmission must be deferred until a later time. The algorithms we present are carefully designed to avoid such an event.) After forming the dedicated path, the desired (one-packet) message is transmitted from P_i to P_j along line k .

MORE NOTATION. For any block of PEs $B(i, j)$, a communication between PEs P_i and P_j that is transmitted above block $B(i, j)$ —in either direction—is termed an $(i \leftrightarrow j)$ communication. Of course, every communication between P_i and P_j is either an $(i \leftrightarrow j)$ communication or a $(j \leftrightarrow i)$ communication.

1.3.3 Message Latency

We assume that the latency (or, duration) of a communication depends only on the distance the message travels (not the direction) and is *logarithmic* in this distance. Specifically, we assume that every $(i \leftrightarrow j)$ communication takes⁴

$$\lceil \log((j - i) \bmod N) \rceil$$

steps. In the next subsection, we sketch a possible implementation of RRP, which has the promise of implementing this delay model.

4. All logarithms are to the base 2 unless otherwise specified.

REMARK. We stress that the COMET protocol is intended to diminish only the *capacitive* delays of message transmission in MOS technologies; it does not affect transmission-line limitations. Therefore, our speedup scheme does not run into any conflicts with speed-of-light limitations [1], [8].

2 LTS ALGORITHMS FOR $RRPs$

This section is devoted to our upper-bound results. We begin, in Section 2.1, with an algorithm that allows an (N, L) - RRP to perform a single sweep of an LTS Algorithm in time

$$T(N; L) \leq \frac{\log^2 N}{\log L} + \log N \log \log L.$$

We then describe briefly, in Section 2.2, how to perform the operations of broadcast and accumulation using one-sweep LTS algorithms and how to compute the parallel-prefix using a two-sweep LTS algorithm; our descriptions are brief because these LTS algorithms are well known.

2.1 Generic Single-Sweep LTS Algorithms

In this subsection, we describe two intimately related parameterized families of one-LTS algorithms. Each algorithm $Down_{(N;L)}$ (resp., $Up_{(N;L)}$) simulates, on an (N, L) - RRP , a single *downward* (resp., *upward*) sweep on an N -leaf complete binary tree. For the sake of simplicity, we describe only algorithms that involve the entire RRP . However, because our algorithms actually work on a *path* of PEs, rather than on a ring, it should be clear that they are easily modified to compute *segmented* versions of the same operations.⁵ (This fact is essential for the correctness of our recursive simulations.) In fact, we exploit the wraparound structure of our rings only in our allowing *any* PE to play the role of the root of the tree in a simulated sweep.

2.1.1 Specification of Algorithms $Down_{(N;L)}$ and $Up_{(N;L)}$

Fix the parameters N and L , and focus on an (N, L) - RRP $\mathcal{R}_{N,L}$. Both our downward-sweep algorithm $Down_{(N;L)}$ and our upward-sweep algorithm $Up_{(N;L)}$ have $\mathcal{R}_{N,L}$ simulate the N -leaf complete binary tree \mathcal{T}_N via a recursive sequence of remappings of the nodes of \mathcal{T}_N onto the PEs of $\mathcal{R}_{N,L}$. These mappings and remappings are specified for $Down_{(N;L)}$ implicitly via the following recursive three-phase process.

Informal Specification of $Down_{(N;L)}$

Phase 1: *Sweeping down the top of \mathcal{T}_N .* $Down_{(N;L)}$ recursively uses the block of PEs $B(0, 2^k - 1)$ to execute the top $k + 1$ levels of \mathcal{T}_N . (We choose the parameter k later.) We assume that this execution ends with the 2^k level- k nodes of \mathcal{T}_N distributed, in left-to-right order, in the PEs of block $B(0, 2^k - 1)$.

Phase 2: *Remapping.* $Down_{(N;L)}$ remaps level k of \mathcal{T}_N by having each PE \mathcal{P}_i of block $B(0, 2^k - 1)$ send its current “state” to PE $\mathcal{P}_{i2^{n-k}}$.

(We assume here that fine-grain LTS algorithms produce

very small “states.” This is true for the motivating examples in the literature.)

Phase 3: *Sweeping down the bottom of \mathcal{T}_N .* $Down_{(N;L)}$ recursively uses all blocks of PES $B(i2^{n-k}, (i+1)2^{n-k} - 1)$, in parallel, to execute the bottom $n - k$ levels of \mathcal{T}_N . We assume that this execution ends with the leaves of \mathcal{T}_N (which are its level- n nodes) distributed, in left-to-right order, in the PEs of $\mathcal{R}_{N,L}$.

Choosing the parameter k . We have two goals that jointly determine our choice of the parameter k .

- 1) In order to minimize the communication overhead of our tree-sweep, we wish to be able to perform the remapping (Phase 2) at each level of the recursion via a single global communication. To accomplish this, we choose k so that $2^k \leq L + 1$. We reason that we can re-map up to L level- k nodes of \mathcal{T}_N via a single global communication, since $\mathcal{R}_{N,L}$ has precisely this many buslines. Since the leftmost level- k node does not move in the remapping (staying in PE \mathcal{P}_0), we arrive at the indicated inequality; since L is a power of 2, our inequality mandates making $k \leq \ell$.
- 2) Letting $T(M; L)$ denote the time that $Down_{(N;L)}$ takes to sweep down an M -leaf subtree of \mathcal{T}_N , we expect the recursive initiation (Phase 1) of the algorithm to take time $T(2^k; L)$ and the recursive extension (Phase 3) to take time $T(2^{n-k}; L)$. Heuristic arguments suggest that $T(M; L)$ is a convex function of M , whence the time for $Down_{(N;L)}$ (which is dominated by the sum $T(2^k; L) + T(2^{n-k}; L)$) will be minimized if we balance the times for Phases 2 and 3 by choosing k so that $2^k \approx 2^{n-k}$.

These two arguments lead us to choose the parameter k to be as large as possible, subject to the restriction that

$$2^k \leq \min\{L, \lceil \sqrt{N} \rceil\}.$$

We now present a detailed specification of Algorithm $Down_{(N;L)}$. Each recursive invocation of the algorithm requires two parameters: the index `root.index` of the PE of $\mathcal{R}_{N,L}$ that plays the role of the root of the current tree at that point in the recursion, and the number `num.leaves` of leaves of that tree. With no loss of generality, PE \mathcal{P}_0 of \mathcal{R}_N will play the role of the root of the initial tree \mathcal{T}_N ; of course, the initial number of leaves is N . The detailed specification appears in Fig. 3. We simplify the specification by assuming that `num.leaves` is always divisible by

$$\min\{L, \lceil \sqrt{\text{num.leaves}} \rceil\};$$

only small additive errors result from this assumption.

Algorithm $Up_{(N;L)}$ is just the dual of Algorithm $Down_{(N;L)}$, with Phases 1 and 3 temporally interchanged. We leave both the informal and formal specification of $Up_{(N;L)}$ to the reader.

2.1.2 Analysis of Algorithms $Down_{(N;L)}$ and $Up_{(N;L)}$

We now assess the time complexity of Algorithms $Down_{(N;L)}$ and $Up_{(N;L)}$. For simplicity, we analyze only Algorithm $Down_{(N;L)}(0, N)$. The reader can easily adapt the analysis to both upward and/or segmented sweeps.

5. Our *lower bounds* hold for arbitrary algorithms on $RRPs$, not just those that operate as ours do.

```

Algorithm  $\mathcal{D}own_{(N;L)}(\text{root.index}, \text{num.leaves})$ 
  Execute the task at  $\text{root.index}$ 
  while  $\text{num.leaves} > 1$  do
    begin
       $N' = \min\{L, \lceil \sqrt{\text{num.leaves}} \rceil\}$ 
      {Execute Phase 1: recursive initiation}
       $\mathcal{D}own_{(N;L)}(\text{root.index}, N')$ 
      {Execute Phase 2: remapping}
      for  $i = \text{root.index} + 1$  to  $\text{root.index} + N' - 1$  do in parallel
        Remap  $\mathcal{P}_i$  to  $\mathcal{P}_{iN'/N}$  via an  $(i \leftrightarrow iN'/N)$  communication
        {Execute Phase 3: recursive extension}
         $\mathcal{D}own_{(N;L)}(iN'/N, N'/N)$ 
      endfor
    end
  end
end Algorithm

```

Fig. 3. The broadcast algorithm $\mathcal{D}own_{(N;L)}$.

THEOREM 2.1. For $N \geq 4$ and $L \geq 2$, Algorithm $\mathcal{D}own_{(N;L)}(0, N)$ operates in time

$$T(N; L) \leq \log N \log \log L + \frac{\log^2 N}{\log L}. \quad (2.1)$$

PROOF. The following recurrence for $T(N; L)$ is clear from the description and specification of $\mathcal{D}own_{(N;L)}$.

$$T(N; L) \leq T(N'; N') + \log N + T\left(\frac{N}{N'}; L\right), \quad (2.2)$$

where, as before, $N' = \min\{L, \lceil \sqrt{N} \rceil\}$. The first term in (2.2) is the time for the Phase-1 recursive invocation of the algorithm, which sweeps down the first N' levels of \mathcal{T}_N ; the second term is the time for the message-transmissions that effect the Phase-2 remapping of the level- N' tree-nodes; the third term is the time for the Phase-3 recursive invocation of the algorithm, which sweeps down the remaining, bottom, levels of \mathcal{T}_N . Using the following easily verified initial condition for the recurrence,

$$T(4; 2) = 2,$$

we now bound the recurrence term by term.

We claim first that, for all $M \geq 4$,

$$T(M; M) \leq \log M \log \log M. \quad (2.3)$$

We proceed by induction on M . The basis is true, since $T(4; 4) \leq T(4; 2) = 2$. By invoking the inductive hypothesis and bound (2.2), we bound $T(M; M)$ for arbitrary $M > 4$ as follows.

$$\begin{aligned} T(M; M) &\leq 2T(\sqrt{M}; \sqrt{M}) + \log M \\ &\leq 2 \log(\sqrt{M}) \log \log(\sqrt{M}) + \log M \\ &= \log M \log \log M. \end{aligned}$$

We now use bounds (2.2), (2.3) to establish the claimed bound (2.1) on $T(N; L)$.

$$\begin{aligned} T(N; L) &\leq \log N' \log \log N' + \log N + T(N/N'; L) \\ &\leq (\log L \log \log L + \log N) \log_L N \\ &= \log N \log \log L + \log^2 N / \log L. \end{aligned}$$

The second step here uses bound (2.2) $\log_L N$ times to bound the $T(N/N'; L)$ term. \square

2.2 Specific LTS Algorithms

We describe here simple LTS algorithms for three fundamental operations: broadcasting, accumulation, and parallel-prefix [2], [3], [4].

One-to-all broadcasting. In the operation of *broadcasting*, one PE—with no loss of generality, \mathcal{P}_0 —sends a single-packet message M to all other PEs. Let the PEs of the architecture be mapped to the nodes of a complete binary tree in any way that places PE \mathcal{P}_0 at the root of the tree. Now, perform a sweep down the tree: as each PE-node receives the message M from its parent, it relays the message to both of its children. At the end of the downward sweep, each PE “knows” the message M .

In the sequel, we denote by $\mathcal{B}r_{(N;L)}$ the algorithm for broadcasting within an (N, L) -RRP, that is based on the downward sweep algorithm $\mathcal{D}own_{(N;L)}$.

Accumulation. The operation of *accumulation* (or, *reduction*) is defined for any binary associative operator \otimes . The \otimes -reduction of the vector $\langle x_0, x_1, \dots, x_{N-1} \rangle$ is the product $x_0 \otimes x_1 \otimes \dots \otimes x_{N-1}$. An architecture whose PEs are indexed from 0 to $N - 1$ in such a way that each PE \mathcal{P}_i initially “knows” value x_i can compute this reduction as follows. Assign the nodes of \mathcal{T}_N to the PEs in any way that assigns the leaves of \mathcal{T}_N to the PEs in their natural left-to-right order. Now, perform a sweep up the tree: Each PE computes the \otimes -product of the quantities it receives from its children and passes this product to its parent. At the end of this upward sweep, the PE that is assigned the root of \mathcal{T}_N “knows” the accumulated \otimes -product.

Parallel-prefix. The *parallel-prefix* (or, *scan*) operation is also defined for any binary associative operator \otimes . The \otimes -scan of the vector $\langle x_0, x_1, \dots, x_{N-1} \rangle$ is the vector $\langle y_0, y_1, \dots, y_{N-1} \rangle$, where each $y_i = x_0 \otimes x_1 \otimes \dots \otimes x_i$. An architecture whose PEs are indexed from 0 to $N - 1$ in such a way that each PE \mathcal{P}_i initially “knows” value x_i can compute this scan as follows. Assign the nodes of \mathcal{T}_N to the PEs in any way that assigns the leaves of \mathcal{T}_N to the PEs in their natural left-to-right order. Now, perform a sweep up the tree, followed by a sweep down the tree. During the upward sweep, the architecture performs an \otimes -reduction of the vector, but each PE retains (for the downward sweep) the value computed by its left child. During the downward sweep, each PE sends its retained value to its right child, which then computes the \otimes -product of its parent’s retained value by its retained value (in that order). At the end of the downward sweep, each PE \mathcal{P}_i “knows” the quantity y_i .

3 LOWER BOUNDS FOR LTS ALGORITHMS ON \mathcal{RRPs}

In this section, we prove lower bounds on $T^*(N; L) \stackrel{\text{def}}{=} \text{the fastest time in which an } (N, L)\text{-}\mathcal{RRP} \text{ can perform a one-to-all broadcast. Our formal bound takes the following form.}$

THEOREM 3.1. *For all integers $N > 4$ and $L \geq 7$,*

$$T^*(N; L) \geq \max \left\{ \frac{1}{4} \frac{\log^2 N}{\log L}, \frac{1}{15} \log N \log \log N \right\}.$$

The first bound in this maximization, which we prove in Section 3.1, says roughly that no $(N, L)\text{-}\mathcal{RRP}$ can perform a one-to-all broadcast in fewer than $(\log^2 N)/(\log L)$ steps. The second bound in this maximization, which we prove in Section 3.2, says, roughly, that no $N\text{-}\mathcal{RRP}$ —no matter how many buslines it has—can perform a one-to-all broadcast in fewer than $\log N \log \log N$ steps. Since broadcasting is intuitively the simplest one-sweep LTS algorithm (an intuition that is verified in Section 4), these lower bounds demonstrate that our one-sweep LTS algorithms in Section 2 (which operate within timebound (1.1)) cannot be sped up by more than a modest constant factor. We turn now to the proofs of our bounds.

3.1 A Lower Bound for Small L

The following lower bound on $T^*(N; L)$ establishes the optimality (to within a constant factor) of our one-pass LTS algorithms $\mathcal{Down}_{(N;L)}$ and $\mathcal{Up}_{(N;L)}$, for “small” values of L , i.e., values no greater than $N^{1/\log \log N}$.

THEOREM 3.2. *For all integers $N > 0$ and $L \geq 7$,*

$$T^*(N; L) \geq \frac{1}{4} \frac{\log^2 N}{\log L}. \quad (3.1)$$

PROOF. Say that the $(N, L)\text{-}\mathcal{RRP}$ \mathcal{R} has a PE \mathcal{P}_0 which wants to broadcast a one-word message M to all other PEs. We demonstrate that, no matter how \mathcal{R} disseminates M to its PEs (subject, of course, to the limitations of \mathcal{RRPs}), it can decrease only slowly the size of the largest remaining block of PEs that are “ignorant” of message M . To the end of verifying this, let us consider the execution of an arbitrary broadcast algorithm \mathcal{A} on \mathcal{R} .

At the beginning of the broadcast, only PE \mathcal{P}_0 “knows” message M , so the initial block of “ignorant” PEs is precisely the block $B_0 \stackrel{\text{def}}{=} B(1, N - 1)$.

Let T_0 be the first instant in the execution of Algorithm \mathcal{A} in which a message crosses either PE $\mathcal{P}_{\lfloor N/3 \rfloor}$ or PE $\mathcal{P}_{\lceil 2N/3 \rceil}$. By our delay model, $T_0 \geq \log N/3$. Because our \mathcal{RRP} has L buslines, at time T_0 , no more than $2L$ messages are “traveling” along the bus and directed to the block $B(\lfloor N/3 \rfloor, \lceil 2N/3 \rceil)$. Now, even if we assume that all of these $2L$ messages arrive at their destinations instantaneously at time T_0 (which can only speed up Algorithm \mathcal{A}), there must remain at

time T_0 a block of “ignorant” PEs B_1 of size

$$|B_1| \geq \left\lceil \frac{N}{3(2L+1)} \right\rceil.$$

Let us now apply the above reasoning to the block B_1 . Let T_1 be the first instant *after time T_0* in the execution of Algorithm \mathcal{A} in which a message crosses one of the two PEs located at $|B_1|/3$ positions to the right of the beginning of B_1 and $|B_1|/3$ positions to the left of the end of B_1 . (In other words, we imagine that we “reset the clock” after time T_0 and begin looking at the indicated PEs.) Since we allowed all messages that were in transit at time T_0 to get to their destinations instantaneously, two things are clear. First,

$$T_1 \geq \log \frac{1}{3} |B_1| \geq \log \frac{N}{9(2L+1)}.$$

This follows just from the fact that no messages were in transit at time T_0 (because we let all of them arrive instantaneously), so the message that determines time T_1 must have traversed at least one-third of block B_1 since time T_0 . Second, even if we assume that all of the $2L$ messages that could be in transit at time T_1 arrive at their destinations instantaneously, and that all of them are destined for PEs in the center of block B_1 , there must remain at time T_1 a block of “ignorant” PEs B_2 of size

$$|B_2| \geq \left\lceil \frac{N}{3^2(2L+1)^2} \right\rceil.$$

We continue this reasoning with block B_2 , which yields a block B_3 of “ignorant” PEs, then with block B_3 , and so on. Eventually, we infer the existence of a sequence of time instants, $T_0, T_0 + T_1, T_0 + T_1 + T_2 \dots$ in the execution of Algorithm \mathcal{A} , with associated blocks B_0, B_1, B_2, \dots of “ignorant” PEs, such that, for each i ,

$$T_i \geq \log \frac{1}{3} |B_i| \quad \text{and} \quad |B_i| \geq \frac{N}{3^i(2L+1)^i}.$$

Now, since Algorithm \mathcal{A} must eventually get the broadcast message to every PE, the total time taken by \mathcal{A} can be bounded as follows. Let

$$\gamma = \log_{3(2L+1)} N = \frac{\log N}{\log(3(2L+1))}.$$

Then, we have

$$\begin{aligned} T^*(N; L) &\geq \sum_{i=0}^{\gamma-1} T_i + 1 \\ &\geq \sum_{i=0}^{\gamma-1} \log \frac{N}{3^i[3(2L+1)]^i} + 1 \\ &= \gamma \log N - \gamma \log 3 - \left(\frac{\gamma}{2} \right) \log(3(2L+1)) + 1. \end{aligned}$$

If we now substitute for γ in this bound and simplify, we obtain the desired bound (3.1). Since our reasoning holds for any broadcast algorithm on an $(N, L)\text{-}\mathcal{RRP}$, the theorem follows. \square

By substituting L in the upper and lower bounds provided by Theorems 2.1 and 3.2, respectively, one establishes the following result.

COROLLARY 3.1. *If $L \leq N^{1/\log \log N}$, then Algorithm $\mathcal{B}r_{(N;L)}$ is within a constant factor of optimal on an (N, L) -RRP.*

We close this section by noting that the bound of Theorem 3.2 trivializes for large L (e.g., $L \geq N^{1/\log \log N}$) to the observation that $T^*(N; L)$ must be proportional to $\log N$. This bound tells us nothing about the complexity of the broadcast operation on RRP, because simply getting message M from PE P_0 to PE $P_{N/2}$ must take time proportional to $\log N$ (because of the subadditivity of logarithms). In the next subsection, we derive a lower bound on $T^*(N; L)$ that is nontrivial no matter how large L is, i.e., no matter how many buslines the RRP in question has. Moreover, this bound establishes the optimality (to within a constant factor) of $\mathcal{B}r_{(N;L)}$ (hence, of $\text{Down}_{(N;L)}$ and $\text{Up}_{(N;L)}$), even for large values of L .

3.2 A Lower Bound for All RRP

In this section, we prove that the time taken by an N -RRP \mathcal{R} to broadcast must be proportional to $\log N \log \log N$, no matter how many buslines \mathcal{R} has. As usual, we assume, with no loss of generality, that PE P_0 of \mathcal{R} is broadcasting a message M to all other PEs. The reader should note that we could rephrase the proof of this bound to hold for any one-sweep LTS algorithm. The rephrasing for downsweeps is little more than a change in terminology; the rephrasing for upsweeps is a bit more complicated, as one has to “run the proof backwards,” which requires a bit of reformulation.

THEOREM 3.3. *For all $N > 4$ and all L ,*

$$T^*(N; L) \geq \frac{1}{15} \log N \log \log N.$$

PROOF. We proceed by induction on N , using as a base all $N \leq 1,024$. Our lower bound is trivial for N in this range, since it asserts only that an RRP requires more than two steps to broadcast when $N > 4$. Let us, therefore, focus on an arbitrary fixed $N > 1,024$, and let us assume inductively that the theorem holds for all smaller N . Say that we have an optimal algorithm \mathcal{A} that broadcasts on an N -RRP \mathcal{R} within time $T^*(N; L)$. By Theorem 2.1, we know that \mathcal{A} operates within time $2 \log N \log \log N$. We can derive from \mathcal{A} a broadcast tree $BT(\mathcal{A})$ whose structure exposes how \mathcal{A} disseminates the broadcast message. The tree $BT(\mathcal{A})$ has node-set $\{0, 1, \dots, N-1\}$; $BT(\mathcal{A})$ has an edge from node i to node j precisely if, in Algorithm \mathcal{A} , PE P_j of \mathcal{R} receives the broadcast message for the first time via a direct communication from PE P_i . Note that node 0 is the root of $BT(\mathcal{A})$.

Henceforth, let us focus on the broadcast of a message M by Algorithm \mathcal{A} .

The intuitive flow of our proof is as follows. From the structure of $BT(\mathcal{A})$, we find in \mathcal{R} a block of adjacent PEs, called a *barrier*, that partitions a portion of the broadcast operation into three disjoint phases:

Phase 1. The root-PE P_0 supplies the PEs in the barrier with message M .

Phase 2. The PEs in the barrier supply a new set of “subroot” PEs with message M .

Phase 3. The new subroot-PEs broadcast message M to the PEs in their broadcast subtrees.

The sum of the times for these three phases is clearly a lower bound on the total time for Algorithm \mathcal{A} 's broadcast. The strategy of our proof is to show that there must exist a barrier for which each of these three operations takes a rather long time. We achieve this by showing that there must exist a barrier whose PEs must supply message M to “many” subroots, all of which are “far” from the barrier and each of which must broadcast message M to PEs of \mathcal{R} which belongs to a “big” subtree of $BT(\mathcal{A})$. One can view the rest of the proof as quantifying the quoted words in the preceding sentence.

We turn now to the details of the argument. Given a barrier B , we call any PE that receives the broadcast message—for the first time—directly from some PE of B a *subroot induced by B* . We say that a subroot P_i covers a PE P_j precisely when node j is in the subtree of $BT(\mathcal{A})$ rooted at node i ; this is equivalent to saying that P_j receives the broadcast message for the first time from P_i , either directly or indirectly. For any set of PEs P , we denote by $\text{COV}(P)$ the set of PEs $P \cup \{\text{PEs covered by PEs in } P\}$.

In order to define the barriers of interest, we first partition the (universal) block of PEs $B(0, N-1) = \{P_0, P_1, \dots, P_{N-1}\}$ into $\sqrt{N}-1$ subsets, each subset being the union of two blocks of PEs. For each

$$i \in \{1, 2, \dots, \sqrt{N}-1\},$$

the i th subset B_i is the set⁶

$$B_i \stackrel{\text{def}}{=} B \left(\left[-\frac{1}{2}(\sqrt{N}-1)i \right], \left[-\frac{1}{2}(\sqrt{N}-1)(i-1) \right] \right) \\ \cup B \left(\left[\frac{1}{2}(\sqrt{N}-1)(i-1) \right], \left[\frac{1}{2}(\sqrt{N})i \right] \right).$$

The k th barrier of \mathcal{R} comprises the PEs whose indices are in the set $B_k \stackrel{\text{def}}{=} \bigcup_{i=1}^k B_i$.

NOTATION. We denote the set of subroots induced by the k th barrier by R_k (see Fig. 4). Also, for notational clarity, we henceforth abbreviate the quantity $\frac{1}{15} \log N \log \log N$ by λ_N .

Our first lemma shows that \mathcal{R} must have a barrier which induces subroots that are located “far” from the barrier and that cover “many” PEs. Specifically, the lemma is our first step in quantifying the qualifiers we have been putting in quotes.

LEMMA 3.1. *There exists a $k \in \{1, 2, \dots, \lambda_N^2\}$ such that*

⁶ Recall that all arithmetic on the indices of PEs is modulo N , so that $P_{-i} = P_{N-i}$.

$$\left| \text{Cov}(R_k - B_{k+1}) \right| \geq \frac{N}{\lambda_N^2} - k\sqrt{N}. \quad (3.2)$$

PROOF. Let the parameter i range over the set $\{1, 2, \dots, \lambda_N^2\}$.

Note first that, for all i , the set of subroot PEs R_i covers at least $N - i\sqrt{N}$ PEs. Let us now consider the λ_N^2 sets of subroots $R_i \cap B_{i+1}$, for varying i (see Fig. 4). We claim that at least one of these sets covers fewer than $N(1 - 1/\lambda_N^2)$ PEs. Were this not the case, each of the sets $\text{Cov}(R_i \cap B_{i+1})$ would contain more than $N(1 - 1/\lambda_N^2)$ PEs. By the pigeon-hole principle, then, some PE would be covered by all λ_N^2 sets $R_i \cap B_{i+1}$. This, however, would imply that there is a “chain” of subroots, $r_1 \in R_1 \cap B_2, r_2 \in R_2 \cap B_3, \dots, r_{\lambda_N^2} \in R_{\lambda_N^2} \cap B_{\lambda_N^2+1}$,

such that each subroot r_j covers subroot r_{j+1} . Such a chain of subroot dependencies is equivalent to a path in $BT(\mathcal{A})$ of length $\geq \lambda_N^2$, hence would imply that Algorithm \mathcal{A} takes time $\geq \lambda_N^2$. This, however, contradicts our having chosen Algorithm \mathcal{A} from among those broadcast algorithms that operate within time proportional to λ_N . This contradiction proves the lemma. \square

Now, let k_0 be an index for which the set

$$\mathcal{R}'_{k_0} \stackrel{\text{def}}{=} R_{k_0} - B_{k_0+1}$$

has large coverage in the sense of (3.2). Our second lemma shows that there are “many” remote subroots in \mathcal{R}'_{k_0} , each of which covers a “big” subtree of PEs.

LEMMA 3.2. *There exists a nonempty set of subroots $\mathcal{R}''_{k_0} \subseteq \mathcal{R}'_{k_0}$ such that:*

- 1) *The number of PEs covered by the subroots in \mathcal{R}''_{k_0} exceeds*

$$\frac{1}{2} \left(\frac{N}{\lambda_N^2} - k_0\sqrt{N} \right).$$

- 2) *For each subroot $r \in \mathcal{R}''_{k_0}$, the number of PEs covered by r exceeds*

$$\frac{1}{4\lambda_N} \left| \mathcal{R}''_{k_0} \right| \left(\frac{N}{\lambda_N^2} - k_0\sqrt{N} \right).$$

PROOF. Let us index the elements of \mathcal{R}'_{k_0} in nondecreasing order of coverage, that is, so that

$$\left| \text{Cov}(\{r_1\}) \right| \geq \left| \text{Cov}(\{r_2\}) \right| \geq \dots \geq \left| \text{Cov} \left(\left\{ r_{\left| \mathcal{R}'_{k_0} \right|} \right\} \right) \right|.$$

For each $i \in \{1, 2, \dots, \left| \mathcal{R}'_{k_0} \right|\}$, let us denote by n_i and N_i , respectively, the cardinalities $n_i \stackrel{\text{def}}{=} \left| \text{Cov}(\{r_i\}) \right|$ and

$N_i \stackrel{\text{def}}{=} \left| \text{Cov}(\{r_1, r_2, \dots, r_i\}) \right|$. We now show that there is a nonempty set of subroots in \mathcal{R}'_{k_0} which comprises the desired set \mathcal{R}''_{k_0} .

Let k_1 be the largest index i such that $N_i \leq 2i\lambda_N n_i$. Note that, since $N_1 < 2\lambda_N N_1 = 2\lambda_N n_1$, we are sure that k_1 “exists,” i.e., that $k_1 \geq 1$. Let us abbreviate the quantity $\left| \mathcal{R}'_{k_0} \right|$ by S . Now, by definition of k_1 , we know that, for each subroot r_i , where $k_1 + 1 \leq i \leq \left| \mathcal{R}'_{k_0} \right|$,

$$n_i < \frac{N_i}{2i\lambda_N} < \frac{N_S}{2i\lambda_N}.$$

It follows that⁷

$$\sum_{j=k_1+1}^S n_j < \frac{N_S}{2\lambda_N} \sum_{j=1}^S \frac{1}{j} \leq \frac{N_S}{2\lambda_N} (1 + \ln S) \leq \frac{1}{2} N_S.$$

We now claim that the set $\{r_1, r_2, \dots, r_{k_1}\}$ can serve as the sought set of subroots \mathcal{R}''_{k_0} . To wit, Lemma 3.1 assures us that

$$\left| \mathcal{R}''_{k_0} \right| = N_{k_1} \geq \frac{1}{2} N_S \geq \frac{N}{2\lambda_N^2} - \frac{k_0\sqrt{N}}{2}.$$

Moreover, by definition, each subroot $r_i \in \mathcal{R}''_{k_0}$ covers at least

$$n_i \geq n_{k_1} \geq \frac{N_{k_1}}{2k_1\lambda_N} > \frac{1}{4k_1\lambda_N} \left(\frac{N}{\lambda_N^2} - k_0\sqrt{N} \right)$$

PEs. \mathcal{R}''_{k_0} thus yields the desired set of k_1 roots, which completes the proof. \square

We are now ready to prove Theorem 3.3. Let k_0 and k_1 be the integers produced in the proofs of Lemmas 3.1 and 3.2, respectively. We bound the time that the optimal algorithm \mathcal{A} takes for each of the three phases of the broadcast defined by the barriers we have selected.

The time for Phase 1. We claim that Algorithm \mathcal{A} (indeed, any algorithm) needs time

$$T_1 \geq T^* \left(\frac{k_1}{\log k_1 \log \log k_1}; L \right)$$

in order to generate the k_1 instances of message M that PEs within barrier $\mathcal{B}_{k_0} = \bigcup_{i=1}^{k_0} B_i$ must transmit (in Phase 2) to the k_1 “remote” subroot PEs of \mathcal{R}''_{k_0} .

In order to see this, say that, at the instant when the last copy of these k_1 instances of message M leaves barrier \mathcal{B}_{k_0} , there are p PEs in \mathcal{B}_{k_0} that “know” message M . Since only such knowledgeable PEs can transmit the message instances, some one of these p PEs must have transmitted at least k_1/p instances of M to the remote subroots. Furthermore, Algorithm \mathcal{A}

7. $\ln x$ denotes the natural logarithm of x .

must have taken time at least $T^*(p; L)$ to create p knowledgeable PEs. Thus, we have

$$T_1 \geq \min_p \max\{T^*(p; L), k_1/p\}. \quad (3.3)$$

Since $T^*(x; L) \geq \log x$ for all x (see the closing paragraph of Section 3.1), the value of p that minimizes maximization (3.3) can be determined to have the form $p_0 = k_1/\log k_1 + (\text{low-order terms})$. Therefore, since T^* is (clearly) monotonically increasing as a function of N , we have

$$T_1 \geq T^*\left(\frac{k_1}{\log k_1}; L\right) \geq T^*\left(\frac{k_1}{\log k_1 \log \log k_1}; L\right). \quad (3.4)$$

The time for Phase 2. The k_1 instances of message M that are sent from within barrier B_{k_0} to the “remote” subroot PEs must pass over block B_{k_0+1} in transit. The time T_2 necessary for the last message instance to make this trip can be no smaller than

$$T_2 \geq \frac{1}{2} \log N, \quad (3.5)$$

because of the distance traveled.

The time for Phase 3. Finally, we focus on the last subroot to receive the message—call it r_0 —that receives message M in Phase 2, and on the time it takes r_0 to relay message M to all the PEs in its subtree. By Lemma 3.2, r_0 covers no fewer than

$$\frac{N}{4k_1\lambda_N^3} - \frac{k_0\sqrt{N}}{4k_1\lambda_N} \geq \frac{N}{8k_1\lambda_N^3}$$

PEs; consequently, r_0 must take time at least

$$T_3 = T^*\left(\frac{N}{8k_1\lambda_N^3}; L\right) \quad (3.6)$$

to transmit message M to these PEs.

Adding bounds (3.4), (3.5), and (3.6), we find that the optimal time to perform a one-to-all broadcast in a N -RRP is no smaller than

$$\begin{aligned} T^*(N; L) &\geq T^*\left(\frac{k_1}{\log k_1 \log \log k_1}; L\right) + \frac{1}{2} \log N + T^*\left(\frac{N}{8k_1\lambda_N^3}; L\right) \\ &\geq \min_x \left\{ T^*\left(\frac{x}{\lambda_N}; L\right) + \frac{1}{2} \log N + T^*\left(\frac{N}{8x\lambda_N^3}; L\right) \right\}. \end{aligned} \quad (3.7)$$

Now, since $T^*(N; L) \leq 2 \log N \log \log N$, we know that the function T^* is subadditive; hence, the value of x that minimizes the final expression in (3.7) must satisfy

$$\frac{x}{\lambda_N} = \frac{N}{x\lambda_N^3},$$

so that

$$x = \frac{1}{2\sqrt{2}} \frac{\sqrt{N}}{\lambda_N}.$$

Thus, (3.7) yields the recurrence

$$T^*(N; L) \geq 2T^*\left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2}; L\right) + \frac{1}{2} \log N$$

which we now set out to solve.

First, using the inductive hypothesis, we obtain

$$\begin{aligned} T^*(N; L) &\geq 2T^*\left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2}; L\right) + \frac{1}{2} \log N \\ &\geq \frac{2}{15} \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right) \\ &\quad \log \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right) + \frac{1}{2} \log N \\ &\geq \frac{1}{15} \log N \log \left(\log \frac{225}{2\sqrt{2}} \sqrt{N} - \log(\log N \log \log N)^2 \right) \\ &\quad - \frac{2}{15} \log(\log N \log \log N)^2 \\ &\quad \log \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right) + \frac{1}{2} \log N. \end{aligned}$$

Now, for any $N > 1,024$, we have

$$\log \left(\frac{225}{2\sqrt{2}} \sqrt{N} \right) - \log(\log N \log \log N)^2 > \frac{5}{42} \log N;$$

therefore,

$$\begin{aligned} T^*(N; L) &\geq \frac{1}{15} \log N \left(\log \log N + \log \frac{5}{42} \right) \\ &\quad - \frac{2}{15} \log(\log N \log \log N)^2 \\ &\quad \log \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right) + \frac{1}{2} \log N \\ &\geq \frac{1}{15} \log N \log \log N + \left(\frac{1}{2} - \frac{1}{15} \log \frac{42}{5} \right) \log N \\ &\quad - \frac{2}{15} \log(\log N \log \log N)^2 \\ &\quad \log \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right). \end{aligned}$$

Since

$$\frac{1}{2} - \frac{1}{15} \log \frac{42}{5} > 0.29,$$

and since, for any $N > 1,024$,

$$0.29 \log N > \frac{2}{15} \log(\log N \log \log N)^2,$$

$$\log \log \left(\frac{225}{2\sqrt{2}} \frac{\sqrt{N}}{(\log N \log \log N)^2} \right),$$

we have finally shown that

$$T^*(N; L) \geq \frac{1}{15} \log N \log \log N,$$

as was claimed. \square

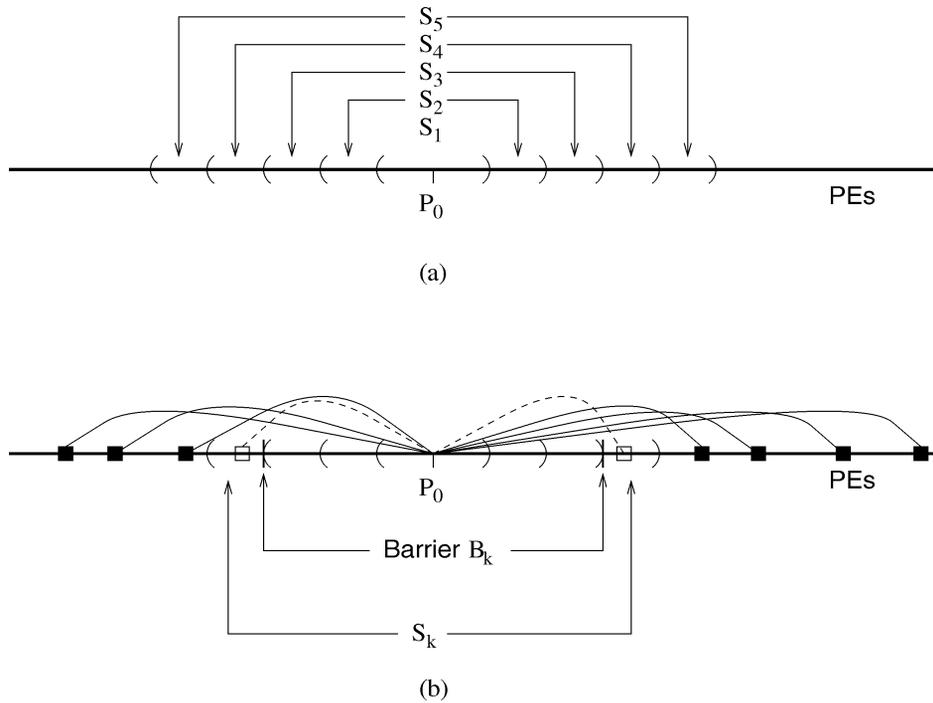


Fig. 4. Illustrating barriers (a) and the subroots induced by a barrier (b). In part (b), each open square denotes a root in S_k , and each filled square denotes a root in $R_k - S_k$.

4 NONTRIVIAL ALGORITHMS ON ARBITRARY ARCHITECTURES

This section is devoted to proving two rather surprising generalizations of Theorem 3.3. First, the lower bound of the theorem holds, not only for any single-sweep LTS algorithm, as shown in Section 3.2, but also for any algorithm that requires *nontrivial communication*. We say that an algorithm \mathcal{A} requires *nontrivial communication* (for short, is a *nontrivial algorithm*) if it requires some PE \mathcal{P} to receive/send information—directly or indirectly—from/to *all* other PEs during the course of the computation.

The second generalization of Theorem 3.3 is perhaps even more surprising. The lower bound of the theorem—even when generalized to all nontrivial algorithms—holds for a much broader class of parallel architectures than just *RRPs*. In fact, the bound holds for parallel machines that communicate via *arbitrary* point-to-point fixed interconnection networks, provided only that the machines have been implemented in a way that satisfies the following conditions (which are quite consistent with current technology).

- 1) The machines must be laid out in some fixed number of spatial dimensions. For simplicity, we assume two-dimensional layouts in what follows, but our argument can be adapted easily to three dimensions.
- 2) Each PE of a machine must occupy a “unit-square” and can send at most one message per time step. (Of course, this condition really just defines our units.)
- 3) The message latency of the machine must be *at least logarithmic* in the Manhattan (rectilinear) distance in the layout, between the sending and receiving PEs, and it must be unaffected by the direction of the communication.

Let \mathcal{M} be any N -PE parallel machine having point-to-point connections between PEs, which is implemented in such a way that these three conditions hold.

THEOREM 4.1. *The N -PE machine \mathcal{M} requires time proportional to $\log N \log \log N$ to execute any nontrivial algorithm \mathcal{A} .*

PROOF. An algorithm \mathcal{A} for machine \mathcal{M} can be nontrivial in two symmetric ways: Either there is at least one PE of \mathcal{M} whose final state is affected by the initial state of every other PE, or there is at least one PE of \mathcal{M} whose initial state affects the final state of every other PE. With no loss of generality, we concentrate on the former contingency, the latter following by “running the following argument backward.” Let us assume henceforth that PE \mathcal{P}_0 of \mathcal{M} is a “sink,” in the sense that its final state is affected by the initial state of every other PE. We prove the theorem in two steps. First, we establish that the time that \mathcal{M} takes to execute Algorithm \mathcal{A} can be no smaller than the time that \mathcal{M} takes to perform a one-to-all broadcast from PE \mathcal{P}_0 . Second, we bound from below the time that \mathcal{M} must take to perform this broadcast. The latter proof evolves in a manner similar to the proof of Theorem 3.3.

Let us attack the first portion of our proof by considering the *dependency tree* $DT(\mathcal{A})$ of Algorithm \mathcal{A} , that is constructed as follows: We assign PE \mathcal{P}_0 to be the root of $DT(\mathcal{A})$. We assign as the children of \mathcal{P}_0 those PEs—call them $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k$ —that communicated *directly* to \mathcal{P}_0 during the execution of \mathcal{A} . We assign as the children of \mathcal{P}_0 's child \mathcal{P}_1 those PEs that:

- communicated directly with \mathcal{P}_1 *before* its last communication with \mathcal{P}_0 *and*
- have not yet been assigned within the tree;

name these grandchildren of the root P_j , for $j = k + 1, k + 2, \dots$. We repeat the assignment process for PEs P_2, P_3, \dots , in order of the indices we are assigning during this process, until all PEs of \mathcal{M} are assigned to tree-nodes. The essential features of the assignment are that the children of PE P_i must have communicated directly with P_i before its last communication with its parent in the tree, and they must not yet have been assigned within the tree; we then index these new tree-nodes using the smallest as-yet unassigned indices. Now, it is obvious that $DT(\mathcal{A})$ is a spanning tree of the computation graph of Algorithm \mathcal{A} . Given that communication delay is not affected by the direction of the communication (by condition 3), we can conclude that the time taken by \mathcal{M} to execute Algorithm \mathcal{A} is no smaller than the time that \mathcal{M} would take to perform a one-to-all broadcast from P_0 , using $DT(\mathcal{A})$ as the broadcast tree (cf. the proof of Theorem 3.3).

We next bound from below the time taken to perform the broadcast. To this end, we formally define the two-dimensional analog of the blocks B_i of Theorem 3.3, for our proof follows the logical flow of the proof of that theorem.

Given any two-dimensional layout of machine \mathcal{M} , let us use the position of PE P_0 as a reference point, in order to partition the PEs of \mathcal{M} into two-dimensional blocks; we call these blocks *squares* as an aid to the intuition. We effect the partition as follows: square SQ_i comprises those PEs whose distance from P_0 in the tree $DT(\mathcal{A})$ is greater than $\frac{1}{2} N^{1/4} \sqrt{i-1}$ but less than $\frac{1}{2} N^{1/4} \sqrt{i}$; see Fig. 5.

Let us now define the k th barrier (of the layout) as the set of PEs $\bigcup_{i=1}^k SQ_i$. Finally, let us denote by R_k the set of subroots of the broadcast tree $DT(\mathcal{A})$ that learn the broadcast message directly from PEs in the k th barrier.

The following lemmas are the two-dimensional analogs of Lemmas 3.1 and 3.2 and are proved in much the same way.

Let T be the optimal time to broadcast using broadcast tree $DT(\mathcal{A})$.

LEMMA 4.1. *There exists a $k \in \{1, 2, \dots, T^2\}$, such that*

$$\left| \text{Cov}(R_k - SQ_{k+1}) \right| \geq \frac{N}{T^2} - k\sqrt{N}.$$

The interested reader can adapt the proof of Lemma 3.1 to prove Lemma 4.1. Let k_0 be an integer for which $R'_{k_0} \stackrel{\text{def}}{=} R_{k_0} - SQ_{k_0+1}$ is large as in Lemma 4.1.

LEMMA 4.2. *There exists a nonempty set of subroots $R''_{k_0} \subseteq R'_{k_0}$ such that:*

1) *The number of PEs covered by the set R''_{k_0} exceeds*

$$\frac{1}{2} \left(\frac{N}{T^2} - k_0\sqrt{N} \right).$$

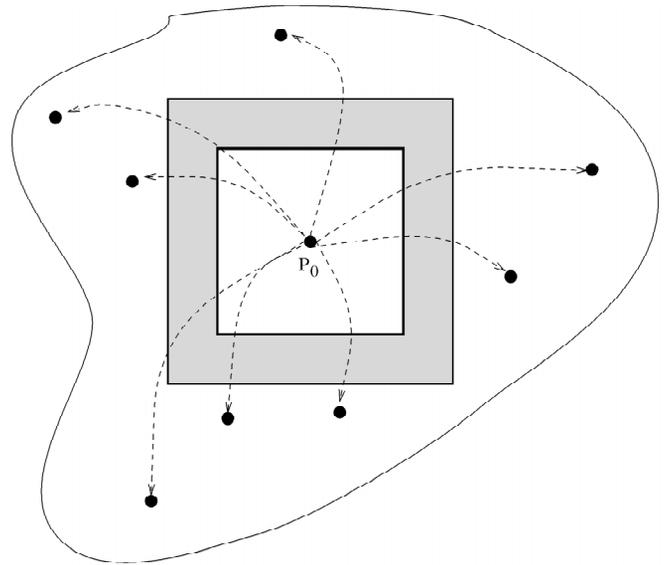


Fig. 5. Illustrating the squares and the barrier (shaded) in the two-dimensional layout of \mathcal{M} .

2) *For each subroot $r \in R''_{k_0}$, the number of PEs covered by r exceeds*

$$\frac{1}{8T|R''_{k_0}|} \left(\frac{N}{T^2} - k_0\sqrt{N} \right).$$

PROOF SKETCH. Define N_i, n_i , and S as in Lemma 3.2. Select x to be the largest i such that

$$n_i > \frac{N_i}{4iT}.$$

For each subroot r_i , where $x+1 \leq i \leq S$,

$$n_i < \frac{N_i}{4iT} < \frac{N_S}{4iT}.$$

Now,

$$\sum_{j=x+1}^S n_j \leq \sum_{j=1}^S n_j < \frac{N_S}{4T} \sum_{j=1}^S \frac{1}{j} \leq \frac{N_S}{4T} (\ln S + 1) \leq \frac{1}{2} N_S,$$

where the last inequality comes from the fact that for any nontrivial computation, $T \geq \frac{1}{2} \log N$. The rest of the proof is similar to Lemma 3.2 and is left to the reader. \square

The proof of the theorem now proceeds by induction on N and is similar to that of Theorem 3.3. \square

ACKNOWLEDGMENTS

It is a pleasure to acknowledge numerous helpful conversations with Marc Picquendar. A portion of this paper was presented at the Sixth IEEE Symposium on Parallel and Distributed Processing, Dallas, Texas, October 26-29, 1994. The research of Arnold L. Rosenberg and Vittorio Scarano was supported in part by U.S. National Science Foundation Grant CCR-92-21785; the research of Ramesh K. Sitaraman was supported in part by U.S. National Science Foundation Grant CCR-94-10077. A portion of this research was done

while Arnold L. Rosenberg was visiting the Department of Computer Science, The Technion, Haifa, Israel.

REFERENCES

- [1] G. Bilardi and F.P. Preparata, "Horizons of Parallel Computation," *J. Parallel and Distributed Computing*, vol. 27, pp. 172-182, 1996.
- [2] G.E. Blelloch, "Scans as Primitive Parallel Operations," *IEEE Trans. Computers*, vol. 38, pp. 1,526-1,538, 1989.
- [3] S.L. Johnsson, "Massively Parallel Computing: Data Distribution and Communication," *Parallel Architectures and Their Efficient Use: Proc. First Heinz-Nixdorf Symp.*, F. Meyer auf der Heide, B. Monien, and A.L. Rosenberg, eds., Paderborn, Germany, *Lecture Notes in Computer Science* 678, pp. 68-92. Berlin: Springer-Verlag, 1993.
- [4] F.T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*. San Mateo, Calif.: Morgan Kaufmann, 1992.
- [5] C. Mead and L. Conway, *Introduction to VLSI Systems*. Reading, Mass.: Addison-Wesley, 1980.
- [6] R. Miller, V.K. Prasanna-Kumar, D. Reisis, and Q.F. Stout, "Parallel Computations on Reconfigurable Meshes," *IEEE Trans. Computers*, vol. 42, pp. 678-692, 1993.
- [7] A.L. Rosenberg, V. Scarano, and R.K. Sitaraman, "The Reconfigurable Ring of Processors: Efficient Algorithms via Hypercube Simulation," *Parallel Processing Letters*, vol. 5, special issue on dynamically reconfigurable architectures, pp. 37-48, 1995.
- [8] P.M.B. Vitanyi, "A Modest Proposal for Communication Costs in Multicomputers," *Concurrent Computations: Algorithms, Architecture, and Technology*, S.K. Tewksbury, B.W. Dickinson, and S.C. Schwartz, eds., pp. 203-216. New York: Plenum Press, 1988.



Arnold L. Rosenberg received an AB in mathematics from Harvard College in 1962, and an AM and PhD in applied mathematics from Harvard University, in 1963 and 1966, respectively.

Dr. Rosenberg is Distinguished University Professor of Computer Science at the University of Massachusetts at Amherst, where he codirects the Theoretical Aspects of Parallel and Distributed Systems (TAPADS) Laboratory. Prior to joining UMass, he was a professor of computer science at Duke University from 1981 to 1986, and a research staff member at the IBM T.J. Watson Research Center from 1965 to 1981. He has held visiting positions at the Technion (Israel Institute of Technology), Yale University, and the University of Toronto.

Dr. Rosenberg's research focuses on theoretical aspects of parallel architectures and communication networks, with emphasis on developing algorithmic techniques for designing better networks and architectures and using them more efficiently. He is the author of more than 120 technical papers on these and other topics in theoretical computer science and discrete mathematics.

Dr. Rosenberg is a fellow of the ACM, a fellow of the IEEE, a Golden Core member of the IEEE Computer Society, and a member of SIAM. He is the editor-in-chief of *Theory of Computing Systems* (formerly, *Mathematical Systems Theory*).



Vittorio Scarano received the Laurea degree in computer science from the University of Salerno (Italy) in 1990 and the Dottorato di Ricerca (PhD) in applied mathematics and computer science from the University of Naples (Italy) in 1995.

Dr. Scarano is an assistant professor in the Dipartimento di Informatica ed Applicazioni "R.M. Capocelli" of the University of Salerno, Italy. From 1992 to 1994, he held visiting positions at the University of Massachusetts at Amherst and at Mount Holyoke College. In 1992, he visited

the University Eötvös Loránd in Budapest (Hungary).

Dr. Scarano's research is conducted in the area of theoretical aspects of parallel computation, mainly on structural properties of communication networks, parallel architectures and algorithms.



Ramesh K. Sitaraman received his BTech in electrical engineering from the Indian Institute of Technology, Madras. He obtained his PhD in computer science from Princeton University in 1993.

Dr. Sitaraman is currently an assistant professor of computer science at the University of Massachusetts at Amherst, where he is a member of the Theoretical Computer Science group, and codirects the Theoretical Aspects of Parallel and Distributed Systems (TAPADS) Laboratory.

Dr. Sitaraman's research focuses on the theoretical aspects of communicating and processing information in parallel and distributed architectures. His specific interests include interprocessor communication in parallel systems, routing in wireless networks, fault-tolerance, and algorithms for scheduling and load-balancing. He is a member of the IEEE and the ACM.