AD-A104 739 CARNEOIE-NELLON UNIV PITTSBURGH PA DEPT OF COMPUTER =-ETC F/B $9 / 2$ 1/0 COMPLEXITY: THE RED-BLUE PEBBLE GAME. (U) MAR BI J HONG, HT KUNS NODO14-76-C-0370

## UNCLASSIFIED CMU-CS-81-111

## DEPARTMENT of COMPUTER SCIENCE



## Carnegie-Mellon University

## 819 30 <br> n9?

## (V) COMPLEXITY: THE BED -BLUE PEBBLE GAME.



Hong, Jia-Wei and H. T. Mung
Department of Computer Science
Carnegic-Mellon University
Pittsburgh, Pennsylvania 15213


February 1981
[Last (elis) Maralta91]

(A preliminary version of this paper is to appear in the Proceedings of the 13th Annual AC.M Symposium on Theory of Computing, May 1981.)


The research was supported in part by the Office of Naval Research under Contracts N00014-76-C-0370 and N00014-80-C-0236, in part by the National Science Foundation under Grant MCS 78-236-76, and in part by the Defense Advanced Research Projects Agency under Contract F33615-78-C-1551 (monitored by the Air Force Office of Scientific Research). About the authors: Hong, Jia-wei is on leave from Peking Municipal Computing Center. Peking. China and is currently visiting the University of Rochester for the Spring of 1981. H.T. King is currently on leave from Carnegie-Mellon University at ESI.'s Advanced Processor Technology Group in San Jose. California. (ESL is a subsidiary of TRW.) Most of the research for this paper was carried out during the Fall of 1980 when both authors were at CMU.

$$
4=1
$$


#### Abstract

In this paper, the red-blue pebble game is proposed to model the input-output complexity of algorithms. Using the pebble game formulation, a number of lower bound results for the I/O, requirement are proven. For example, it is shown that to perform the $n$-point FFT (or the ordinary nan matrix multiplication algorithm) with a device of $O(S)$ memory, at least $\Omega\left(n \log n / \log S\right.$ ) (or $\Omega\left(n^{3} / \sqrt{S}\right)$, respectively) time is needed for the I/O. Similar results are obtained for algorithms for several other problems. All of the lower bounds presented are the best possible in the sense that they are achievable by certain decomposition schemes.


The results in this paper provide insight into the difficult task of balancing I/O and computation in specialpurpose system design. For example, for the $n$-point FFT, the I/O lower bound implies that an $S$-point device achieving a speed-up ratio $O(\log S)$ over the conventional $O(n \log n)$ implementation is all that one can hope for.


## 1. Int roduction

When a large computation is performed on a small device or memory, the computation must be decomposed into subcomputations. Executing subcomputations one at a time may require a substandial amount of I/O to store or retrieve intermediate results. Very often it is the I/O that dominates the speed of a computation. In fact, I/O is a typical botueneck for performance at all levels of a computer system. However, to the authors' knowledge the I/O problem was not previously modelled or sudied in any systematic or abstract manner. Similar problems were studied only in a few isolated instances [2,5]. This paper proposes a pebble game, called the red-blue pebble game, to model the problem, and presents a number of lower bound results for the I/O requirement. All the lower bounds presented can be shown to be the best possible, in the sense that they are achieved by certain decomposition schemes. The paper is organized according to the techniques used to derive these lower bounds.

In Section 2 we formally define the pebble game and point out its relation to the I/O problem. In Section 3 we show that lower bounds for I/O in the pebble game can be established by studying the so-called $S$-paritioning problem. This is the key result of the paper in the sense that it provides the basis for the derivation of all the lower bounds. In Section 4 we prove a lower bound for the FFT algorithm. Lower bounds in Section 5 are based on the information speed function, which measures how fast the number of vertices on which a given vertex "depends" can grow in a directed acyclic graph of a certain type. We demonstrate the dramatic difference between the I/O requirement for the odd-even transposition sorting network and that for the "snake-like" mesh graph. In contrast to the focus of Section 5. Section 6 studies independent computations for which there are very little information exchanges among vertices. There we obtain, for example, a lower bound for the ordinary matrix multiplication algorithm. In Section 7 we prove a general theorem on products of graphs. Using this theorem. one can determine the I/O required by a product of graphs, by examining only the individual graphs. A summary and concluding remarks are provided in Section 8.

Results of this paper impose upper bounds on the maximum possible speed-up obtainable with a specialpurpose hardware device when the bandwidth of the memory that supplies data to the device remains constant. For example, our lower bound on the I/O requirement for the $n$-point FFT (Corollary 4.1 ) implies that an $S$-point device can achieve a speed-up ratio of at most $O(\log S)$ over the conventional $O(n \log n)$ software implementation. Similarly, for matrix multiplication our result (Corollary 6.2) implies that a $\sqrt{S} \times \sqrt{S}$ device can achieve a speed-up ratio of at most $O(\sqrt{S})$.

## 2. The Red-Blue Pebble Game and Its Relation to the I/O Problem

As the usual pebble game (see, e.g., [4]), the red-blue pebble game is played on a directed acyclic graph ${ }^{1}$. At any point in the pebble game, some vertices of the graph will have red pebbles, some will have blue pebbles, some will have both red and blue pebbles and the remainder will have no pebbles at all. Following the notation of Pippenger [8], define a configuration as a pair of subsers of the vertices, one comprised of just the vertices having red pebbles, and the other just those having blue pebbles. Thus vertices beionging to the intersection of the two sets have both red and blue pebbles on them. The set of inputs (or outputs) of the graph is some designated set of vertices containing at least those vertices that have no predecessors (or successors, respectively). We assume that the set of inputs is disjoint from that of outputs. For all the examples discussed in the paper, only vertices that have no predecessors (or successors) are assumed to be inputs (or outputs, respectively), except in Section 7 where products of graphs are considered. The initial (or terminal) configuration is one in which only inputs (or outputs, respectively) have pebbles, and they are all blue pebbles. The rules of the red-blue pebble game are as follows.

R1. (Input) A red pebble may be placed on any vertex that has a blue pebble.
R2. (Output) A blue pebble may be placed on any vertex that has a red pebble.
R3. (Compute) If all the immediate predecessors of a vertex have red pebbles, a red pebble may be placed on that vertex.

R4. (Deiete) A pebble (red or blue) may be removed from any vertex.
A transition is an ordered pair of configurations, the second of which follows from the first according to one of the rules. A calculation is a sequence of configurations, each successive pair of which form a transition. A complete calculation is one that begins with the initial configuration and ends with the terminal configuration.

A graph on which the red-blue pebble game is played can model a computation performed on a two-level memory structure, consisting of say, a fast memory and a slow memory. Vertices represent operations and their results. An edge from one vertex to another indicates that the result of one operation is an operand of the other. An operation can be performed only if all the operands reside in the fast memory. Placing a red pebble using rule R3 corresponds to performing an operation and storing the result in the fast memory. Placing a blue pebble using rule R2 corresponds to storing a copy of a result (currently in the fast memory) into the slow memory, whereas placing a red pebble using R1 corresponds to retrieving a copy of a result (currently in the slow memory) into the fast memory. Removing a red or blue pebble using rule R4 means frecing a memory location in the fast or slow memory, respectively. The maximum allowable number of red

[^0]or blue pebbles on the graph at any point in the game corresponds to the number of words availabic for use in the fast or slow memory, respectively.

For the purpose of this paper, we assume that the fast memory can hold only S words, where S is a constant, while the slow memory is arbitrarily large. Thus when the pebble game is played on a graph, at most $S$ red pebbles, and any number of blue pebbles, can be on the graph at any dime. For any given graph, we are interested in the minimum I/O time Q , which is defined by
$\mathrm{Q}=$ the minimum number of transitions according to rule R 1 or R 2 required by any complete calculation.

For the FFT graph, it is not difficult to prove the following upper bound on Q by the decomposition scheme illustrated in Figure 2-1 ${ }^{2}$.

Theorem 2.1. For the $n$-point FFT graph,

$$
Q \cdot \log S=O(n \log n)
$$



Figure 2-1: (a) the 16-point FFT graph, and (b) decomposing the FFT graph, with $\mathrm{n}=16$ and $\mathrm{S}=4$.

However, for proving tight lower bounds on Q, we found that it was difficult to work with the red-blue pebble game directly. Instcad we study the $S$-partitioning problem, which is a "static" problem in the sense

[^1]that it does not apply rules on-the-fly as in a game. We show that lower bounds for the S-partitioning problem can be translated into lower bounds on $Q$ for the red-blue pebble game.

## 3. The S-Partitioning Problem and the Key Lemma

In this section we show that every complete calculation of the red-blue pebble game on a directed acyclic graph defines a partition of the graph. Let $G=(V, E)$ be a directed acyclic graph where $V$ and $E$ are the vertex and edge sets of $G$, respectively. A family of subsets of $V,\left\{V_{1}, V_{2}, \ldots, V_{h}\right\}$, is called an $S$-partition of $G$ for some positive integer $S$ if the following four properties hold.

Pl. The $V_{i}$ 's are disjoint and $U_{i=1}^{h} V_{i}=V$.
P2. For each $V_{i}, l \leq i \leq h$, there exists a dominator set $D_{i}$ for $V_{i}$ that contains at most $S$ vertices. (A dominator set for $V_{i}$ is defined to be a set of vertices in $V$ such that every path from an input of $G$ to a vertex in $V_{i}$ contains some vertex in the set.)

P3. For each $V_{i}, 1 \leq i \leq h$, the minimum set $M_{i}$ of $V_{i}$ has at most $S$ vertices. (The minimum set of $V_{i}$ is defined to be the set of those vertices in $V_{i}$ that do not have any sons belonging to $V_{i}$.)

P4. There is no cyclic dependence among vertex sets in $\left\{\mathrm{V}_{1}, \mathrm{~V}_{2}, \ldots, \mathrm{~V}_{\mathrm{h}}\right\}$. (A vertex set $\mathrm{V}_{\mathrm{i}}$ is said to depend on another vertex set $V_{j}$ if there is an edge in $E$ from a vertex in $V_{j}$ to a vertex in $V_{i}$.)

Theorem 3.1. Let $G=(V, E)$ be a directed acyclic graph. Any complete calculation of the redblue pebble game on $G$, using at most $S$ red pebbles, is associated with a $2 S$-partition of $G$ such that

$$
S \cdot h \geq q \geq S \cdot(h-1)
$$

where q is the I/O time required by the complete calculation and h is the number of vertex sets in the 2 S -partition.

Proof: Denote by $C$ any complete calculation. We can divide $C$ into a sequence of $h$ consecutive subcalculations, $C_{1}, C_{2}, \ldots . C_{h}$, for some $h$ such that in each $C_{i}, 1 \leq i \leq h-1$, there are exactly $S$ transitions using rule $R 1$ or $R 2$, and in $C_{h}$ there are no more than $S$ such transitions. For $i=1, \ldots, h$, define $V_{i}$ to be the largest vertex set in which each vertex satisfies the following three properties.
(i) During subcalculation $\mathrm{C}_{\mathrm{i}}$ it has a red pebble placed on it using rule R 1 or R 3 .
(ii) At the end of subcalculation $C_{i}$, it either has red pebbles, or blue pebbles that are placed on it during $\mathrm{C}_{\mathrm{i}}$, or has a son in $\mathrm{V}_{\mathrm{i}}$.
(iii) It does not belong to any $\mathrm{V}_{\mathrm{j}}$ with $\mathrm{j}<\mathrm{i}$.

We claim that the family $\left\{\mathrm{V}_{1}, \mathrm{~V}_{2}, \ldots, \mathrm{~V}_{\mathrm{h}}\right\}$ is a 2 S -partition of G . First we show that property Pl holds. By (iii) it follows immediately that the $\mathrm{V}_{\mathrm{i}}$ 's are disjoint. In the following we show that every vertex in $V$ belongs to some $V_{i}$. Suppose that a vertex. which is not an input. has a red or blue pebble on it at the end of some subcalculation $C_{i}$. Then there must exist a subcalculation $C_{j}, j \leq i$. during which the vertex has a red pebble placed on it using rule R3, and at the end of $C_{j}$ it either remains to have the red pebble or has a blue pebble that is placed on it during $C_{j}$. This implies
that the vertex belongs to $V_{k}$ for some $k \leq j$. Similarly one can show that if an input has a red pebble on it at the end of $C_{i}$, then it must belong to $V_{k}$ for some $k \leq i$. Because calculation $C$ is a complete calculation, all outputs have blue pebbles on them at the end of the last subcalculation $\mathrm{C}_{\mathrm{h}}$; thus they all belong to $\mathrm{U}_{\mathrm{i}=1}^{\mathrm{h}} \mathrm{V}_{\mathrm{i}}$. Consider now any immediate predecessor $u$ of an output $v$. Suppose that $v$ belongs to $\mathrm{V}_{\mathrm{i}}$. Then $v$ cannot have any pebble on it at the beginning of $\mathrm{C}_{\mathrm{i}}$ and thus must have a red pebble placed on it using R3 during $C_{i}$. This implies that we have one of the following two cases:

Case 1: Vertex $u$ has a red pebble on it at the end of subcalculation $\mathrm{C}_{\mathrm{i}-1}$. Then by reasons stated above, $u$ belongs to some $V_{j}, j \leq i-l$.

Case 2: Vertex $u$ has a red pebble placed on it using rule $R 1$ or $R 3$ during $C_{i}$. If $u$ does not belong to any $V_{j}$ with $j<i$, then because $u$ has a son $v$ in $V_{i}$, $u$ itself must belong to $V_{i}$.
We have shown that all the immediate predecessors of outputs belong to $U_{i=1}^{h} V_{i}$. Similarly, we can show that all the immediate predecessors of the immediate predecessors of outputs belong to $\mathrm{U}_{\mathrm{i}=1}^{\mathrm{h}} \mathrm{V}_{\mathrm{i}}$. Property Pl follows by induction. Note that both Case 1 and Case 2 above imply that if $V_{i}$ depends on $V_{j}$ then $j$ < $i$. Therefore there cannot be any cyclic dependence among $V_{i}$ 's, and thus property P4 hoids. For proving property P 2 for any $\mathrm{V}_{\mathrm{i}}, \mathrm{l} \leq \mathrm{i} \leq \mathrm{h}$, we consider two subsets of V , $V_{R}$ and $V_{B R}$, which are defined as follows.

- $\mathrm{V}_{\mathrm{R}}$ consists of those vertices that have red pebbles placed on them just before subcalculation $C_{i}$ begins.
- $\mathrm{V}_{\mathrm{BR}}$ consists of those vertices that have blue pebbles placed on them just before subcalculation $C_{i}$ begins and have red pebbles placed on them according to rule R1 during $C_{i}$.
It is easy to see that by property (i) in the definition of $V_{i}, V_{R} \cup V_{B R}$ forms a dominator set for $V_{i}$. Since there can be at most $S$ red pebbles on $G$ at any time, we have

$$
\left|V_{R}\right| \leq S
$$

The fact that at most $S$ transitions can use rule $R 1$ during $C_{i}$ implies that

$$
\left|V_{B R}\right| \leq s .
$$

Thus

$$
\left|V_{R} \cup V_{B R}\right| \leq\left|V_{R}\right|+\left|V_{B R}\right| \leq 2 S .
$$

We have shown that $\left\{\mathrm{V}_{1}, \mathrm{~V}_{2} \ldots, \mathrm{~V}_{\mathrm{h}}\right\}$ satisfies property P 2 . The proof of property P 3 is similar. By property (ii) in the definition of $\hat{V}_{i}$, we know that at the end of subcalculation $\mathrm{C}_{\mathrm{i}}$, every vertex in $M_{i}$, the minimum set of $V_{i}$, has red pebbles, or blue pebbles that are placed on it during $C_{i}$. Since there can be at most $S$ vertices having red pebbles placed on them at any time, and at most $S$ vertices having blue pebbles placed on them according to rule R2 during $C_{i}$, the minimum set $M_{i}$ can have at most $2 S$ vertices. We have shown that $\left\{V_{1}, V_{2}, \ldots, V_{h}\right\}$ is a $2 S$-partition of $G$. The theorem follows by noting that corresponding to each $\mathrm{V}_{\mathrm{i}}, 1 \leq \mathrm{i} \leq \mathrm{h}-1$, exactly S transitions using $R 1$ or $R 2$ are performed and to $V_{h}$, no more than $S$ such transitions are performed.

Let
$P(S)=$ the minimum number of vertex sets that any $S$-partition of $G$ must have.

We have, by Theorem 3.1, the key lemma of the paper.
Lemma 3.1. For any directed acyclic graph G , the minimum I/O time satisfies

$$
\mathrm{Q} \geq \mathrm{S} \cdot(\mathrm{P}(2 \mathrm{~S})-1)
$$

Using this lemma, lower bounds for P can be translated immediately into lower bounds for Q .

## 4. Lower Bounds for the FFT Computation

In this section we establish a lower bound on the minimum I/O time Q for the n -point FFT graph (see Figure 2-1(a)), by proving a lower bound on P .

Define an $S$-dominator partition of a graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$ to be a family of subsets of $\mathrm{V},\left\{\mathrm{V}_{1}, \mathrm{~V}_{2}, \ldots, \mathrm{~V}_{\mathrm{h}}\right\}$, satisfying properties $\mathrm{P} 1, \mathrm{P} 2$ and P 4 of an S -partition, but not necessarily property P 3 . Let
$P_{D}(S)=$ the minimum number of vertex sets that any $S$-dominator partition of $G$ must have.
Then clearly $P_{D}(S) \leq P(S)$, since any $S$-partition is also an $S$-dominator partition. The following theorem establishes a lower bound on $P_{D}(S)$, and thus a lower bound on $P(S)$.

Theorem 4.1. Suppose that $\mathrm{S} \geq 2$. The minimum number of vertex sets that any S -dominator partition of the $n$-point FFT graph must have satisfies

$$
P_{D}(S)=\Omega((n \log n) /(S \log S)) .
$$

Proof: Since there are a total of $\theta(n \log n)$ vertices in the $n$-point FFT graph, it suffices to prove that any vertex set $U$ that has a dominator set of size no more than $\mathrm{S}, \mathrm{S} \leq \mathrm{n}$, can have at most $S \log S+S$ vertices. We shall show this by induction on $n$. The assertion holds trivially for the case when $\mathrm{n}=2$. Assume now that it holds for the m -point FFT for any $\mathrm{m} .<\mathrm{n}$. We want to show that it holds for $\mathrm{m}=\mathrm{n}$. Consider the n -point FFT graph. We partition its vertex set into four disjoint sets $\mathrm{A}, \mathrm{B}, \mathrm{C}$ and D such that sets C and D equally partition the set of outputs. and sets A and $B$ equally separate the rest of the vertices. See Figure $4-1$ below. Let $d_{d} \cdot d_{B} . d_{C}$ or $d_{D}$ be the number of elements in the dominator set that belong to sets $A, B, C$ or $D$ respectively. Let $u_{A}, u_{B}$, $u_{C}$ or $u_{D}$ be the number of elements in the vertex set $U$ that belong to sets $A, B, C$ or $D$, respectively. It is easy to see that elements in set A that are also in the dominator set form a dominator set for set $\mathrm{U} \cap \mathrm{A}$. Thus by the inčucti, n hypothesis,

$$
\begin{equation*}
u_{A} \leq d_{A} \log d_{A}+d_{A} . \tag{1}
\end{equation*}
$$

Similarly, we have

$$
\begin{equation*}
u_{B} \leq d_{B} \log d_{B}+d_{B} \tag{2}
\end{equation*}
$$

Let $R_{A}$ or $R_{B}$ be the set of those horizontal paths from inputs to outputs on which there are verices in the dominator set that belong to $A$ or $B$, respectively. Then

$$
\begin{equation*}
\left|R_{A}\right| \leq d_{A} \text { and }\left|R_{B}\right| \leq d_{B} . \tag{3}
\end{equation*}
$$

For each vertex in $U \cap C$ or $U \cap D$. it either belongs to the dominator set or one of its immediate predecessors is on a horizontal path belenging to $R_{A}$ and the other on one belonging to $R_{B}$. Therefore

$$
u_{C} \leq d_{C}+\min \left(\left|R_{A}\right|,\left|R_{B}\right|\right) \text { and } u_{D} \leq d_{D}+\min \left(\left|R_{A}\right|,\left|R_{B}\right|\right),
$$



Figure 4-1: Partitioning the FFT graph for the induction proof.
from which we have

$$
\begin{equation*}
u_{C}+u_{D} \leq d_{C}+d_{D}+2 \min \left(\left\{R_{A}\left|,\left|R_{B}\right|\right) .\right.\right. \tag{4}
\end{equation*}
$$

By (1), (2), (3) and (4),

$$
u_{A}+u_{B}+u_{C}+u_{D} \leq\left[d_{A} \log d_{A}+d_{B} \log d_{B}+2 \min \left(d_{A} ; d_{B}\right)\right]+d_{A}+d_{B}+d_{C}+d_{D}
$$

Since $d_{A}+d_{B} \leq S-d_{C}-d_{D}$, we have

$$
u_{A}+u_{B}+u_{C}+u_{D} \leq\left(S-d_{C}-d_{D}\right) \log \left(S-d_{C}-d_{D}\right)+S \leq S \log S+S,
$$

which completes the induction proof.

By Lemma 3.1 we have the following lower bound result.
Corollary 4.1. For the n-point FFT graph,
$Q \cdot \log S=\Omega(n \log n)$.
Thus the I/O time for the $n$-point FFT when executed on a special-purpose device with $S$ words of memory is at least $\Omega(\mathrm{n} \log \mathrm{n} / \log \mathrm{S})$, implying that the maximum-possible speed-up ratio over the usual $O(\mathrm{n} \log \mathrm{n})$ implementation is at most $\mathrm{O}(\log S)$. This upper bound on the speed-up ratio holds no matter how fast the the device may be, since it is a consequence of the I/O consideration. The upper bound can be reduced only if the bandwidth of the memory that supplics data to the special-purpose device is increased. A systolic device that distributes $S$ words of memory in a linear processor array and achieves $\theta(\log S)$ speed-up for the $F F T$ is described by Kung [7].

## 5. Lower Bounds Based on Information Speed Functions

Many "rcgular" graphs $G=(\mathrm{V}, \mathrm{E})$ have the property that all inputs can reach all outputs through vertexdisjoint paths. In the proof of Theorem 4.1 we have alrady noted that the FFT graph has this propery. In the current section, this type of graph will be considered. The vertex-disjoint paths from inputs to outputs will be called lines, for simplicity. We say that the information speed function is $\Omega(\mathrm{F}(\mathrm{d})$ ) if for any two vertices $u$, von the same line that are $d$ apart there are at least $F(d)$ vertices in the graph satisfying the following two properties.

F1. At most one of these vertices can belong to a single line.
F2. Each of these vertices belongs to a path connecting $u$ and $v$.
The following theorem shows that lower bounds on Q can be obtained from lower bounds on F or upper bounds on $\mathrm{F}^{-1}$.

Theorem 5.1. For any graph where all inputs can reach all outputs through vertex-disjoint paths, if the information speed function is $\Omega\left(\mathrm{F}(\mathrm{d})\right.$ ) where F is monotonically increasing and $\mathrm{F}^{-1}$ exists, then

$$
\mathrm{Q} \cdot \mathrm{~F}^{-1}(\mathrm{~S})=\Omega(\mathrm{L}),
$$

where $L$ is ine total number of vertices on the vertex-disjoint paths or the lines.
Proof: As in the proof of Theorem 4.1, we will establish

$$
\mathrm{P}_{\mathrm{D}}(\mathrm{~S})=\Omega\left(\mathrm{L} /\left(\mathrm{S} \cdot \mathrm{~F}^{1}(\mathrm{~S})\right)\right)
$$

by showing that any vertex set $U$ in a $S$-dominator partition can have at most $O\left(S \cdot F^{1}(S)\right.$ ) vertices on the lines. Note that vertices in $U$ can be on at most $S$ lines, since the lines are vertex-disjoint and $U$ has a dominator set of size at most S . The theorem follows from the claim that on any line there can be at most $\mathrm{F}^{-1}(\mathrm{~S})+1$ vertices in U . Suppose that the claim is false for some line. Then on this line there are two vertices $u$ and $v$ in $U$ that are $F^{-1}(S)+1$ apart. Consequently, there are $\mathrm{F}\left(\mathrm{F}^{-1}(\mathrm{~S})+1\right)$ vertices satisfying properties F 1 and F 2 . If any of these vertices belongs to anocher vertex set $U^{\prime}$ in the $S$-dominator partition, then by property $F 2$ there will be a cyclic dependence among vertex sets in the S-dominator partition, violating property P 4 in Section 3 . Therefore all of these $F\left(F^{-1}(S)+1\right)$ vertices, which form a set of more than $S$ vertices, belong to $U$, and by property Fl they belong to distinct lines. This is a contradiction, since vertices in U can be on at most S lines.

Corollary 5.1. For the odd-even transposition sorting network (see, e.g., [6]) for sorting nelement runs,

$$
Q \cdot S=\Omega\left(n^{2}\right)
$$

for any $\mathrm{S}<\mathrm{n}$.
Proof: Consider the sub-network that includes only half of the inputs and outputs, as shown in Figure 5-1. It is easy to see that we can assume the sub-network has $n / 2$ lines with $\mathrm{L}=\theta\left(\mathrm{n}^{2}\right)$ and $F(d)=d / 2$ for $d \leq n$.

Corollary 5.2. For the $m \times n$ snake-like directed mesh as shown in Figure 5-2,

$$
Q=\Omega(m n),
$$



Figure 5-1: The odd-even transposition sorting network, where each " 0 " is a comparator.
for any $\mathrm{S}<\mathrm{m}$.
Proof: Consider as lines all the horizontal vertex-disjoint paths from inputs to outputs. It is easy to see that we can assume $\mathrm{F}(\mathrm{d})=\mathrm{m}$ for any $\mathrm{d} \geq 2$. Let U be any vertex set in an S -dominator partition of the graph. As in the proof of Theorem 5.1. we note that vertices in $U$ can be on at most $S$ lines, and that on any line chere can be at most two vertices in $U$. Therefore, U can have at most $O(S)$ vertices, and thus $P_{D}(S)($ or $P(S))=\Omega(m n / S)$. The corollary follows from Lemma 3.1.


Figure 5-2: The snake-like directed mesh.

Straightforward decomposition schemes will show that lower bounds in the above corollaries are best possible. We note that when S increases the I/O requirement Q for the odd-even transposition sorting network decreases at the rate of $\mathrm{l} / \mathrm{S}$, whereas that for the snake-like directed mesh remains unchanged essentially. We say that graphs like the latter are indecomposable.

## 6. Independent Evaluation of Multivariate Expressions

Given values for indeterminates $x_{1} \ldots \ldots x_{n^{\prime}}$, the problem is to evaluate multivariate polynomial expressions $y_{i}=y_{i}\left(x_{1}, \ldots, x_{n}\right), i=1,2, \ldots, m$. Assume that each $y_{i}$ is a sum of at least two terms and in each $y_{i}$, all the terms are distinct and have degrees $\leq \mathrm{D}$. An example of such a problem is matrix multiplication, where $\mathrm{D}=2$. An independent evaluation of $y_{i}{ }^{\prime} s$ is an algorithm or a directed acyclic graph with inputs $x_{i}$ 's and outputs $y_{i}$ 's satisfying the following properties.

El. In the evaluation of each $y_{\mathrm{i}}$, all (and only) those product terms which appear in the fully distributed expression of $y_{i}$ are computed first by multiplications, and then using these product terms $y_{i}$ is formed through a summation tree by additions or subtractions only. In particular, no multiplication can be performed after an addition or subtraction.

E2. Internal vertex sets of the summation trees for all the $y_{i}$ 's are disjoint from each other, that is, none of the internal vertices in one tree appears as an internal vertex in another. (Thus, evaluations of $y_{i}$ 's are independent from each other.)

Let $X$ be any set of $x_{i}^{\prime}$ 's or products in $x_{i}$ 's. For any output $y_{i}$, define $h\left(y_{i}, X\right)$ as the number of terms in $y_{i}$ that can be obtained from X directly or by multiplying elements in X . For any $\mathrm{Y} \subseteq\left\{y_{1}, \ldots, y_{m}\right\}$ we further define

$$
h(Y, X)=\Sigma_{y \in Y} h(y, X)
$$

For example, if $y_{1}=x_{1} x_{2}+x_{3}^{2} x_{1}, y_{2}=x_{1}^{2} x_{2}^{2}+x_{1} x_{3}^{4}, Y=\left\{y_{1}, y_{2}\right\}$, and $X=\left\{x_{1}, x_{2}^{2}, x_{3}^{2}\right\}$, then $h\left(y_{1}, X\right)=1$. $\mathrm{h}\left(\mathrm{y}_{2}, \mathrm{X}\right)=2$, and $\mathrm{h}(\mathrm{Y}, \mathrm{X})=3$. Define the $S$-combination number to be
$H(S)=\max \{\mathrm{h}(\mathrm{Y}, \mathrm{X})| | \mathrm{Y}|\leq \mathrm{S},|\mathrm{X}| \leq \mathrm{S}\}$.
We have the following result.
Theorem 6.1. Suppose that $H(S)=\Omega(S)$. Then for any independent evaluation of a multivariate expression of degree $\leq \mathrm{D}$,

$$
\mathrm{Q} \cdot \mathrm{D} \cdot \mathrm{H}(\mathrm{~S}) / \mathrm{S}=\Omega(|\mathrm{V}|),
$$

where $|\mathrm{V}|$ is the total number of vertices in the graph corresponding to the independent evaluation.

Proof: Let $\left\{\mathrm{V}_{1}, \mathrm{~V}_{2}, \ldots, \mathrm{~V}_{\mathrm{h}}\right\}$ be an S -partition of the graph associated with the independent evaluation. We shall prove the following.
(i) Each $\mathrm{V}_{\mathrm{i}}, 1 \leq \mathrm{i} \leq \mathrm{h}$, can have at most $\mathrm{H}(\mathrm{S})+2 \mathrm{~S}$ internal vertices. (An internal vertex is defined to be a vertex belonging to the internal vertex set of some summation tree.)
(ii) There are at least $|\mathrm{V}| /(2 \mathrm{D})$ internal vertices in the graph.

By property $P 3$ in the definition of $S$-partition. the minimum set of $V_{i}$ has at most $S$ vertices. This implies that $V_{i}$ can have nonempty intersections with internal vertex sets of at most $S$ summation trees, since by E2 each of such intersections has at least one distinct vertex in the minimum set. Thus, to bound the number of internal vertices that $V_{i}$ can have. we need only consider summation trees for $S y_{i}$ 's. By property P2 of S-partition, we note that $\mathrm{V}_{1}$ has a dominator
set $D_{i}$ of size no more than $S$. By the definition of $H(S)$, from $D_{i}$ one can form at most $H(S)$ terms appearing in the $S y_{i}$ 's. These terms. together with possible verices in $D_{i}$ that are already internal vertices, can generate at most $\mathrm{H}(\mathrm{S})+2 \mathrm{~S}$ internal vertices. We have shown (i). To prove (ii), let A be the total number of internal vertices in the graph corresponding to the independent evaluation. Then the total number of external vertices, or terms, in all the summation trees, is no greater than 2A. Each product term requires at most $D-1$ multiplications; thus the towl number of vertices $|\mathrm{V}|$ in the graph satisfies:

$$
|V| \leq 2 A(D-1)+A \leq 2 A D .
$$

This proves (ii). It follows from (i) and (ii) that

$$
h \geq(|V| / 2 D) /(H(S)+2 S)
$$

and by Lemma 3.1,

$$
\mathrm{Q}=\Omega(\mathrm{S} \cdot|\mathrm{~V}| /(\mathrm{D} \cdot(\mathrm{H}(2 \mathrm{~S})+2 \mathrm{~S}))) .
$$

The theorem follows from the assumption that $\mathrm{H}(\mathrm{S})=\Omega(\mathrm{S})$.
Corollary 6.1. For the ordinary matrix-vector multiplication algorithm for multiplying an mxn marrix with an $n$-vector,

$$
Q \cdot S=\Omega(m n)
$$

assuming that entries in the matrix can be generated on-the-fly and thus are not required to be input.
Proof: The corollary follows immediately by noting that $\mathrm{H}(\mathrm{S})=\Theta\left(\mathrm{S}^{2}\right)$ and $\mathrm{D}=1$.
Lemma 6.1. For matrix-matrix multiplication,

$$
H(S)=\theta\left(S^{3 / 2}\right) .
$$

Proof: We shall only prove $H(S)=O\left(S^{3 / 2}\right)$, since it is trivial to see $H(S)=\Omega\left(S^{3 / 2}\right)$. Consider the matrix multiplication, $A B=C$. Let $W$ be any set of entries in $A$ and $B$, with $|W| \leq S$. Partition A into two classes as follows. Class $A_{d}$ consists of all rows in $A$, each of which has at least $\sqrt{S}$ entries in $W$, and class $A_{d}^{\prime}$ consists of the rest of rows in A. Accordingly, matrix $C$ is partitioned into two classes, $A_{d} B$ and $A_{d}^{\prime} B$. Since $A_{d}$ can have at most $\sqrt{S}$ rows, and since in any row of $A_{d} B$ an entry in $B$ can appear at most once (and $B$ has no more than $S$ entries in $W$ ), the maximum number of terms in $A_{d} B$ that can be obtained by multiplying elements in $W$ is at most $S \cdot \sqrt{S}=S^{3 / 2}$. For entries in $A_{d}^{\prime} B$. each of them can be obtained by multiplying at most $\sqrt{S}$ elements in $W$, since each row in $A_{d}^{\prime}$ has at most $\sqrt{S}$ elements in $W$. Therefore, in any $S$ entries of $A_{d}^{\prime} B$, there are at most $S \cdot \sqrt{S}=S^{3 / 2}$ terms that can be obtained by multiplying elements in $W$.

By Theorem 6.1 and Lemma 6.1, we have the 10 wing result.
Corollary 6.2. For the ordinary matrix-matrix multiplication algorithm for multiplying mxk and kxn matrices,

$$
\mathrm{Q} \cdot \sqrt{\mathrm{~S}}=\mathrm{S}(\mathrm{mkn}) .
$$

## 7. Lower Bounds for Products of Graphs

As demonstrated in Sections 4 and 5 , one can establish lower bounds on $Q$ by proving upper bounds on the size of any vertex set that has a dominator set of size at most S . This is equivalent to proving lower bounds on
$D(n)=$ the minimum size of a dominator set for any vertex set having no less than $n$ vertices.
In this section we show that lower bounds on $D(n)$ for the product of two graphs can be obtained from lower bounds on $\mathrm{D}(\mathrm{n})$ for individual graphs. (See, for example, [3] for the definition of the product of two graphs.) Let $G_{1} \times G_{2}$ be the product of $G_{1}$ and $G_{2}$. A vertex $\left(v_{1}, v_{2}\right) \in G_{1} \times G_{2}$ is defined to be an input (or output) of $G_{1} \times G_{2}$ if $v_{1}$ is an input of $G_{1}$ or $v_{2}$ is an input of $G_{2}$, (or, respectively, $v_{1}$ is an output of $G_{1}$ and $v_{2}$ is an output of $G_{2}$.) Of course $D(n)$ depends on the graph on which it defines; we use $D_{1}(n), D_{2}(n)$ and $D(n)$ to distinguish the case when the graph is $G_{1}, G_{2}$ and $G$ respectively.

Lemma 7.1. If $f$ is a positive function such that $f(x) / x$ is non-increasing, $\Sigma a_{i} \geq T_{1} T_{2}$, and $0 \leq a_{i} \leq T_{2}$, then

$$
\left.\Sigma f\left(\mathrm{a}_{\mathrm{i}}\right) \geq \mathrm{T}_{2} \mathrm{fT}_{2}\right) .
$$

Proof:

$$
\Sigma f\left(\mathrm{a}_{\mathrm{i}}\right) \geq \sum \mathrm{a}_{\mathrm{i}} \mathrm{f}\left(\mathrm{~T}_{2}\right) / \mathrm{T}_{2} \geq \mathrm{T}_{1}\left(\mathrm{~T}_{2}\right) .
$$

Theorem 7.1. (The Production Theorem for Dominators)
If $D_{i}(n)=\Omega\left(d_{i}(n)\right)$ where $d_{i}, i=1,2$, is a positive, non-decreasing function such that $d_{i}(x) / x$ is non-increasing, then

$$
D\left(n_{1} n_{2}\right)=\Omega\left(\min \left\{n_{1} \cdot d_{2}\left(n_{2}\right), n_{2} \cdot d_{1}\left(n_{1}\right)\right\}\right) .
$$

Proof: Let $W$ be a subset in $V_{1} \times V_{2}$ of size $n_{1} n_{2}$. Define

$$
U_{2}=\text { the set of vertices } p_{2} \text { in } V_{2} \text { for which }\left|W \cap\left(V_{1} x\left\{p_{2}\right\}\right)\right| \geq n_{1}
$$

and

$$
U_{2}^{\prime}=V_{2}-U_{2}
$$

Clearly, we have $\left|U_{2}\right| \leq n_{2}$ giving

$$
\begin{equation*}
\left|W \cap\left(\left\{p_{1}\right\} x U_{2}\right)\right| \leq n_{2} \tag{5}
\end{equation*}
$$

and for $p \in U_{2}^{\prime}$,

$$
\begin{equation*}
\left|W \cap\left(V_{1} \times\left\{p_{2}\right\}\right)\right|<n_{1} . \tag{6}
\end{equation*}
$$

One of the following two cases must hold.

Case $1 .\left|W \cap\left(V_{1} \times U_{2}\right)\right| \geq n_{1} n_{2} / 2$.
Let $p_{1}$ be any vertex in $V_{1}$. Any dominator set for $W \cap\left(\left\{p_{1}\right\} x V_{2}\right)$ is of size at least $\mathrm{d}_{2}\left(\left|W \cap\left(\left\{p_{1}\right\} \times V_{2}\right)\right|\right)$. Thus the size of any dominator set for $W$ satisfies:

$$
D\left(n_{1} n_{2}\right) \geq \sum_{p_{1} \in V_{1}} d_{2}\left(\left|W \cap\left(\left\{p_{1}\right\} x V_{2}\right)\right|\right)
$$

Since $U_{2}$ is a subset of $V_{2}$ and $d_{2}$ is a non-decreasing function, we have

$$
D\left(n_{1} n_{2}\right) \geq \sum_{p_{1} \in V_{1}} d_{2}\left(\left|W \cap\left(\left\{p_{1}\right\} x U_{2}\right)\right|\right)
$$

By the definition of Case 1 ,

$$
\begin{equation*}
\Sigma_{p_{1} \in V_{1}}\left|W \cap\left(\left\{p_{1}\right\} \times U_{2}\right)\right| \geq n_{1} n_{2} / 2 \tag{7}
\end{equation*}
$$

By Lemma 7.1, it follows from (5) and (7) that

$$
\Sigma_{p_{1} \in V_{1}} d_{2}\left(\left|W \cap\left(\left\{p_{1}\right\} x U_{2}\right)\right|\right) \geq n_{1} \cdot d_{2}\left(n_{2}\right) / 2
$$

implying

$$
D\left(n_{1} n_{2}\right) \geq n_{1} \cdot d_{2}\left(n_{2}\right) / 2
$$

Case 2. $\left|\mathrm{W} \cap\left(\mathrm{V}_{1} \mathrm{xU}_{2}^{\prime}\right)\right|>\mathrm{n}_{1} \mathrm{n}_{2} / 2$.
Let $p_{2}$ be any vertex in $V_{2}$. Any dominator set for $W \cap\left(V_{1} \times\left\{p_{2}\right\}\right)$ is of size at least $d_{1}\left(\mid W \cap\left(V_{1} \times\left\{p_{2}\right\}\right)\right)$. Thus the size of any dominator set for $W$ satisfies:

$$
D\left(n_{1} n_{2}\right) \geq \Sigma_{p_{2} \in V_{2}} d_{1}\left(\left|W \cap\left(V_{1} \times\left\{p_{2}\right\}\right)\right| .\right.
$$

Since $U_{2}^{\prime}$ is a subset of $V_{2}$, we have

$$
D\left(n_{1} n_{2}\right) \geq \Sigma_{p_{2} \in \mathrm{U}_{2}^{\prime}} d_{1}\left(\mid W \cap\left(V_{1} \times\left\{p_{2}\right\} \mid\right) .\right.
$$

By the definition of Case 2 ,

$$
\begin{equation*}
\Sigma_{p_{2} \in U_{2}^{\prime}}\left|W \cap\left(V_{1} x\left\{p_{2}\right\}\right)\right| \geq n_{1} n_{2} / 2 \tag{8}
\end{equation*}
$$

By Lemma 7.1, it follows from (6) and (8) that
implying

$$
D\left(n_{1} n_{2}\right) \geq n_{2} \cdot d_{1}\left(n_{1}\right) / 2 .
$$

Let $L_{1}=\{V, E\}$ be a directed line where $V=\{1,2, \ldots, m\}$, and $E=\{(i, i+1) \mid i=1,2, \ldots, m-1\}$, with unique input " 1 " and output " $m$." We have $D_{L_{1}}(n)=1$ for any $n \leq m$. See Figure 7-1.
Let $L_{2}=L_{1} \times L_{1}$. Then

$$
D_{L_{2}}\left(n^{2}\right)=\Omega(\min \{l \cdot n, 1 \cdot n\}),
$$

giving
$D_{L_{2}}\left(n^{2}\right)=\theta(n)$.
Let $L_{3}=L_{2} L_{1}$. Then

$$
D_{L_{3}}\left(n^{3}\right)=\Omega\left(\min \left\{n \cdot n, n^{2} \cdot 1\right\}\right)
$$

giving

$$
D_{L_{3}}\left(n^{3}\right)=\theta\left(n^{2}\right) .
$$


$L_{1}$

$L_{1}$

$$
L_{2}=L_{1} \times L_{1}
$$

Figure 7-1: The product of two directed lines, where each " 0 " represents an input.
Let $L_{d}=L_{1} \times \ldots \times L_{1}$, that is, $L_{d}$ is the product of $d L_{1}$ 's. Then similarly,

$$
\begin{equation*}
D_{L_{d}}\left(n^{d}\right)=\theta\left(n^{d-1}\right) . \tag{9}
\end{equation*}
$$

Corollary 7.1. For the product $\mathrm{L}_{\mathrm{d}}$ with $\mathrm{d} \geq 2$,

$$
Q \cdot S^{1 /(d-1)}=\Omega\left(\mathrm{m}^{\mathrm{d}}\right)
$$

Proof: By (9), the maximum size of any vertex set that has a dominator set of size at most S is $O\left(S^{d /(d-1)}\right)$. Since there are a total of $m^{d}$ vertices in $L_{d}$, we have

$$
P(S)=\Omega\left(m^{d} / S^{d /(d-1)}\right)
$$

by which the Corollary follows from Lemma 3.1.

We have a similar product theorem for separators of a graph. For the special case $L_{d}$, bounds on the sizes of minimum separators have been established by A. L. Rosenberg [9].

## 8. Summary and Concluding Remarks

To compare I/O requirements for different algorithms, we propose the use of the following measure. The decomposability factor $\lambda(S)$ of an algorithm or graph $G=(\mathrm{V}, \mathrm{E})$ is defined to be the ratio between the sequential time of the algorithm, that is $|\mathrm{V}|$, and the minimum $\mathrm{I} / \mathrm{O}$ time Q when assuming S red pebbies are used. Thus,

$$
\mathrm{Q} \cdot \lambda(\mathrm{~S})=|\mathrm{V}| .
$$

For a given algorithm, $|\mathrm{V}|$ is fixed. We see that the larger the $\lambda(\mathrm{S})$ is, the less the I/O is required. A summary of results of this paper on specific algorithms or graphs, expressed in terms of bounds on $\lambda(\mathrm{S})$, is as follows:

Algorithms or Graphs
Matrix-vector multiplication (ordinary algorithm)
Odd-even transposition sorting network
Matrix-matrix multiplication (ordinary algorithm)
$L_{\mathrm{d}^{\prime}}(\mathrm{d} \geq 2)$
FFT
Snake-like directed mesh
$\lambda(\mathbf{S})$
$\theta(\mathbf{S})$
$\theta(\mathrm{S})$
$\theta(\sqrt{5})$
$\theta\left(S^{1 /(d-1)}\right)$
$\theta(\log S)$
$\theta(1)$

It is also possible to establish upper bounds on $\lambda(\mathrm{S})$ for a class of algorithms for solving a given problem. For example, it has been shown recently that for any sorting algorithm based on the decision tree model, $\lambda(S)=O(\log S)[10]$.

The problem of establishing bounds on $\lambda(\mathrm{S})$ is closely related to several other graph partitioning problems. We intend to work on some of these partitioning problems in the future, and show how they are related to the 1/O complexity problem addressed in this paper.

## Acknowledgments

The authors would like to thank S. W. Song for his comments on this paper.

## References

[1] Cook, S.A. and Sethi, R.
Storage Requirements for Deterministic Polynomial Time Recognizable Languages.
J. Comp. and Sys. Sci. 13:25-37, 1976.
[2] Floyd, R.W.
Permuting Information in Idealized Two-Level Storage.
In Miller, R.E. and Thatcher, J.W. (editor), Complexity of Computer Computations, pages 105-109. Plenum Press, New York, 1972:
[3] Harary, F.
Graph Theory.
Addison-Wesley, Reading, Massachusetts, 1969.
[4] Hopcroft, J.E., Paul, W. and Valiant, L.G.
On Time Versus Space.
Journal of the ACM 24:332-337, 1977.
[5] Kernighan, B.W. and Lin, S.
An Effective Heuristic Procedure for Partitioning Graphs.
Bell Systems Technical Journal 49:291-308, February, 1970.
[6] Knuth, D. E.
The Art of Computer Programming. Volume 3: Sorting and Searching. Addison-Wesley, Reading, Massachusetts, 1973.
[7] Kung, H.T.
Special-Purpose Devices for Signal and Image Processing: An Opportunity in VLSI.
In Proceedings of the SPIE, Vol. 241, Real-Time Signal Processing III, pages 76-84. The Society of Photo-Optical Instrumentation Engineers, July, 1980.
[8] Pippenger, N.
Pebbling.
In Proceedings of the Fifth IBM Symposium on Mathematical Foundations of Computer Science. Academic \& Scientific Programs, IBM Japan, May, 1980.
[9] Rosenberg, A. L.
Private communication.
[10] Song, S.W.
I/O Complexily and Design of Special-Purpose Hardware for Sorting.
VLSI Document V075, Carnegie-Mellon University, Department of Computer Science, February, 1981.


| REPORT DOCU!.. EHTATION PAGE | READ INSTRUCTIONS BEFORE CO:FiETMG FDR:G |
| :---: | :---: |
| 1. REFOAT MUMAEA CMU-CS-81-111 | 3. RECIPIENT:S CATALOE AUMBEA $1237$ |
| 4. TITLE (and Subitila) <br> I/O COMPLEXITY: THE RED-BLUE PEBBLE GAME | 5. TYGE OF REPORT A PERIOO COVEAED <br> Interi:n |
|  | 6. performing org. report mumber |
| 7. authonfes HONG, JIA-WEI <br> AND H <br> T. <br> KUNG | 2. COMtRACT OR GAAMT nJmben(o) N00014-76-C-0370 |
| 9. PERFORMING ORGANIZATICN NAME AND ADORESS Carnegie-Mellon University Computer Science Department Pittsburgh, PA 15213 | 10. PROJRAMELEMEGT.FEOIECT. TASK AREA \& WORK UNIT RUMBERS |
| II. CONTROLLINCOFFICE NAME AND ADORESS Office of Naval Research Arlington, VA 22217 | 12. REPORT OATE |
|  | February 1981 |
|  | 13. NUMAER OF PAGES 19 |
|  | 15. SECUR1TY CLASS. (ot mio roport) |
|  | UNCLASSIFIED |
|  | ise. OECLASSIFICATION. SOWNGRADINE |

16. DISTRIBUTION STATEMENT (of thia Report)

Approved for public release; distribution unlimited.
17. DISTRIEUTION STATEMENT (aI the abstract ontorod in Block 20, $1 /$ dilforonc fran Rapart)
18. SUPPLEMENTARY NOTES

20. ABSTMACT (Contlnue en reverse dide if neceetery and laundly ay bleck mamber)



[^0]:    ${ }^{\text {T The red-blue pebble game discussed in this paper is not related in any way to the black-and-white pebble game introduced by Cook }}$ and Sethi [1].

[^1]:    ${ }^{2}$ All logarithms used in this paper are to base 2.

