

CR05: Data Aware Algorithms – Year 2020

Part 1: Introduction to data movement and pebble-game models

Loris Marchal

September 9, 2020

1 Algorithm Design and Data Movement

Let us consider here a very simple problem, the matrix product, to show how the design of the algorithm can influence the memory usage and the amount of data movement¹. We consider two square $n \times n$ matrices A and B , and we compute their product $C = AB$.

Algorithm 1: SIMPLE-MATRIX-MULTIPLY(n, C, A, B)

```
for  $i = 0 \rightarrow n - 1$  do
  for  $j = 0 \rightarrow n - 1$  do
     $C_{i,j} = 0$ 
    for  $k = 0 \rightarrow n - 1$  do
       $C_{i,j} = C_{i,j} + A_{i,k}B_{k,j}$ 
```

We consider that this algorithm is executed on a simple computer, consisting of a processing unit with a fast memory of size M . In addition to this limited memory, a large but slow storage space is available. In the following, we assume that this space is the disk and has unlimited available storage space.² Our objective is to minimize the data movement between memory and disk, also known as the volume of I/O (input/output), that is the number of A , B and C elements that are loaded from the disk to the memory, or written back from the memory to the disk. We assume that the memory is limited, and cannot store more than half a matrix, i.e., $M < n^2/2$.

In Algorithm 1, all B elements are accessed during one iteration of the outer loop. Thus, as the memory cannot hold more than one half of B , at least $n^2/2$ elements must be read. For the n iterations, this leads to $n^3/2$ read operations. This is huge as it is the same order of magnitude of the number of computations (n^3).

Fortunately, it is possible to improve the I/O behavior of the matrix product by changing the algorithm. We set $b = \sqrt{M/3}$ and assume that n is a multiple of b . We consider the blocked version of the matrix product, with block size b , as detailed in Algorithm 2. In this algorithm $C_{i,j}^b$ denotes the block of size b at position (i, j) (all elements $C_{k,l}$ such that $ib \leq k \leq (i+1)b - 1$ and $jb \leq l \leq (j+1)b - 1$).

Each iteration of the inner loop of the blocked algorithm must access 3 blocks of size b^2 . Thanks to the choice of b , this fits in the memory, and thus, each of these $3b^2$ elements are read and written at most once. This leads to at most $2M$ data movements. Since there are $(n/b)^3$ iteration of the inner loop, the volume of I/O of the blocked algorithm is $O((n/b)^3 \times 2M) = O(n^3/\sqrt{M})$.

¹A large part of this section is adapted from [9].

²Note that this study detailed for the case “main memory vs. disk” may well apply to other pairs of storage such as “fast small cache vs. large slower memory”.

Algorithm 2: BLOCKED-MATRIX-MULTIPLY(n, C, A, B)

```

 $b \leftarrow \sqrt{M/3}$ 
for  $i = 0, \rightarrow n/b - 1$  do
  for  $j = 0, \rightarrow n/b - 1$  do
    for  $k = 0, \rightarrow n/b - 1$  do
      Simple-Matrix-Multiply( $n, C_{i,j}^b, A_{i,k}^b, B_{k,j}^b$ )
  
```

2 (Black) pebble game for memory minimization

2.1 Definition

We present here the first theoretical model that was proposed to study the space complexity of programs. This model, based on a pebble game, was originally used to study register allocation. Registers are the first level of storage, the fastest one, but also a scarce resource. When allocating registers to instructions, it is thus crucial to use them with caution and not to waste them. The objective is thus to find the minimum amount of registers that is necessary for the correct execution of the program. Minimizing the number of registers is similar to minimizing the amount of memory. For the sake of consistency, we present the pebble game as a problem of memory size rather than register number minimization. This does not change the proposed model nor the results, but allows us to present all results of this document with the same model and formalism.

The pebble-game was introduced by Sethi [6] to study the space complexity of “**straight-line**” programs, that is, programs whose control flow does not depend on the input data. A straight-line program is modeled as a directed acyclic graph (DAG): a vertex represents an instruction, and an arc between two vertices $i \rightarrow j$ means that the results of the vertex i is used for the computation of j .




Figure 1: Graph corresponding to the computation of the expression $7 + (5 - z) \times (1 + x) - (1 + x - t) / (2 + z) + v$

When processing a vertex, all its inputs (as well as the result) must be loaded in memory, and the goal is to execute the program using the smallest amount of memory. Memory slots are modeled as pebbles, and executing the program is equivalent to playing a game on the graph with the following rules:

- (PG1) A pebble may be removed from a vertex at any time.
- (PG2) A pebble may be placed on a source node at any time.
- (PG3) If all predecessors of an unpebbled vertex v are pebbled, a pebble may be placed on v .

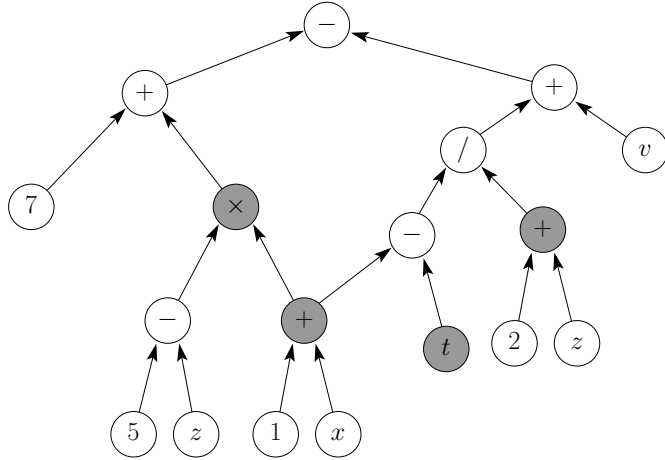


Figure 2: Playing the (black) pebble game on the graph of Figure 1. Dark nodes are the ones currently pebbled, meaning that four values are now in memory: $(5 - z) \times (1 + x)$, $1 + x$, t and $2 + z$.

The goal of the game is to put a pebble on each output vertex at some point of the computation, and to minimize the total number of pebbles needed to reach this goal. In this game, pebbling a node corresponds to loading an input in memory (rule PG2) or computing a particular vertex (rule PG3). From a winning strategy of this game, it is thus straightforward to build a solution to the original memory allocation problem.

2.2 Variants and hardness

Note that the game does not ensure that each vertex will be pebbled only once. Actually, in some specific graphs, it may be beneficial to pebble several times a vertex, that is, to compute several times the same values, to save a pebble needed to store its value. A variation of the game, named the *Progressive Pebble Game*, forbids any recomputation, and thus models the objective of minimizing the amount of memory without any increase in the computational complexity. In this latter model, the problem of determining whether a directed acyclic graph can be processed with a given number of pebbles has been shown NP-hard by Sethi [6]. The more general problem allowing recomputation is even more difficult, as it has later been proved PSPACE-complete by Gilbert, Lengauer and Tarjan [2]. Another variant of the game slightly changes rule PG3 and allows to *shift* a pebble to an unpebbled vertex if all its predecessors are pebbled. Van Emde Boas [10] shows that it can at most decrease the number of pebbles required to pebble the graph by one, but in the worst case the saving is obtained at the price of squaring the number of moves needed in the game.

Another variant, called the black-white pebble game has been proposed to model non deterministic execution [3, 4], where putting a white pebble on a node corresponds to guessing its value; a guessed vertex has to be actually computed later to check the value of the guess.

2.3 Special case: tree-shaped graphs

A simpler class of programs consists in trees rather than general graphs: for example, arithmetic expressions are usually described by in-trees (rooted directed trees with all edges oriented towards the root), unlike the one of Figure 1 which uses a common sub-expression twice.

2.3.1 Complete binary trees

We first consider **Complete binary trees** and establish the following result for the variant when shifting pebbles is not allowed:




Figure 3: Complete binary tree of depth $k = 4$ and pebbles when the last leaf (27) is pebbled. This figure as well as the followings used to illustrate the pebble games are borrowed from [5].

Theorem 1. *Any pebbling strategy (with or without recomputation) for the complete balanced binary tree $T(k)$ of depth k uses at least $k + 2$ pebbles and $2^{k+1} - 1$ steps. There is a pebbling strategy that reaches both bounds.*

Proof. For the first lower bound, for a given strategy, consider the last time when it pebbles a leaf, called v . We assume that no vertex on the path P from v to the root is already pebbled (otherwise pebbling v at this time is unnecessary and can be avoided). Since v is the last pebbled leaf, for any other leaf v' , there should be a pebble on the path from v' to some node on the path P from v to the root. Minimizing this number is obtained by pebbling all the descendants not in P of nodes of P , which needs k pebbles. Together with the pebble on v , it makes $k + 1$ pebbles. When the parent of v is pebbled, $k + 2$ pebbles are required

The second lower bound simply derives from the fact that each vertex of the tree should be pebble once.

Finally, the following postorder pebbling strategy reaches both bounds:

1. Pebble the left subtree, leave a pebble on its root
2. Pebble the right subtree, leave a pebble on its root
3. Shift the pebble of one child of the root to the root, remove the other pebble

A tree of depth 0 (single node) can be pebbled with one pebble. Hence, a tree of depth k can be pebbled with $k + 2$ pebbles using this strategy. Furthermore, it only pebbles once each vertex. \square

2.3.2 General trees

We now get back to the variant without shifting pebble. Even in this case, the problem gets much simpler than for general graph: Sethi and Ullman [7] designed an optimal scheme which relies on the following theorem.³

³When pebble shifting is allowed, the minimum number of pebbles needed to pebble a tree is the Strahler number (as outlined in [1]), which is a measure of a tree's branching complexity that appears in natural science, such as in the analysis of streams in a hydrographical basin [8] or in biological trees such as animal respiratory and circulatory systems https://en.wikipedia.org/wiki/Strahler_number.

Theorem 2. *An optimal solution for the problem of pebbling an in-tree with the minimum number of pebbles using rules PG1 – PG3 is obtained by a depth-first traversal which orders subtrees by non-increasing values of $P(i)$, where the peak $P(v)$ of the subtree rooted at v is recursively defined by:*

$$P(v) = \begin{cases} 1 & \text{if } v \text{ is a leaf} \\ \max(k + 1, \max_{i=1 \dots k} P(c_i) + i - 1) & \text{where } c_1, \dots, c_k \text{ are the children of } v \\ & \text{such that } P(c_1) \geq P(c_2) \geq \dots \geq P(c_k) \end{cases}$$

The first step to prove this result is to show that depth-first traversals are dominant, i.e., that there exists an optimal pebbling scheme which follows a depth-first traversal. Pebbling a tree using a depth-first traversal adds a constraint on the way a tree is pebbled: consider a tree T and any vertex v whose children are c_1, \dots, c_k , and assume w.l.o.g that the first pebble that is put in the subtree rooted at v is in the subtree rooted at its leftmost child c_1 . Then, in a depth-first traversal, the following pebbles must be put in the subtree rooted at c_1 , until c_1 itself holds a pebble (and all pebbles underneath may be removed). Then we are allowed to start pebbling other subtrees rooted at c_2, \dots, c_k . These traversals are also called *postorder*, as the root of a subtree is pebbled right after its last child.

To prove that depth-first traversals are dominant, we notice that whenever some pebbles have been put on a subtree rooted at some vertex v , it is always beneficial to completely pebble this subtree (until its root v , which uses a single vertex) before starting pebbling other subtrees.

The second step to prove Theorem 2 is to justify the order in which the subtrees of a given vertex must be processed to give a minimal number of pebbles. This result follows from the observation that after having pebbles $i - 1$ subtrees, $i - 1$ pebbles should be kept on their respective roots while the i^{th} subtree is being pebbled. When the root is pebbled, $k + 1$ pebbles are needed, which is why it also appear in the maximum.

2.4 Space-Time Tradeoffs

2.4.1 Example with the FFT graph




Figure 4: FFT graph with 8 input/output vertices (depth $k = 3$)

The FFT graph (see Figure 4) contains many complete binary trees, hence it inherits some of the previous results: we need at least $k + 1$ pebbles to pebble the FFT graph of depth k . A

pebbling scheme reaching this number of pebbles is obtained by pebbling each output one after the other. A simple optimization is possible: when two vertices of the third level are pebbled, two outputs (on the fourth level) can be pebbled at once. However, this results in a large number of re-computations: all level-3 and 4 vertices are pebbled once, each level-2 vertex is pebbled twice, and each input (level-1) vertices are pebbled four times. Using a larger number of pebbles allows to reduce the number of steps. For example, with $2n = 2^{k+1}$ pebbles, each level can be pebble one after the other, which allows to reach the minimum number of steps.

2.4.2 Grigoriev's Lower-Bound Method

Definition 1 ($w(u,v)$ -flow). *Let f be a function from \mathcal{A}^n to \mathcal{A}^m . f has a $w(u,v)$ -flow if for all partition (X_0, X_1) of the n inputs and all partition (Y_0, Y_1) of the m outputs, such that $|X_1| = u$ and $|Y_1| = v$, there exists a subfunction h of f obtained by assigning all variables in X_0 to fixed values and discarding output values in Y_0 , such that h has at least $|A|^{w(u,v)}$ distinct points in the image of its domain f .*

This definition is illustrated on Figure 5. The idea is to characterize a significant information flow from X_1 to Y_1 , even if inputs not in X_1 are fixed and outputs not in Y_1 are discarded.

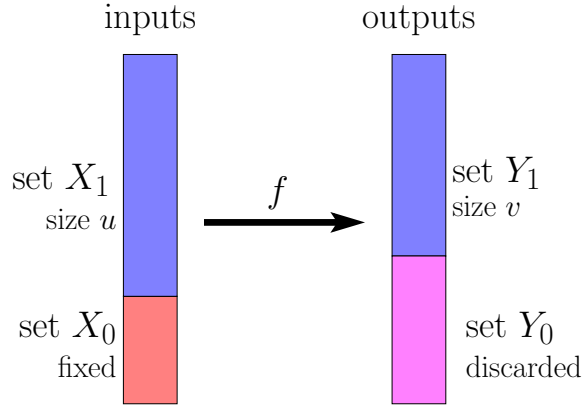


Figure 5: Definition of a flow

A more general property is defined below.

Definition 2 ((α, n, m, p) -independent function). *A function $f : \mathcal{A}^n \mapsto \mathcal{A}^m$ is an (α, n, m, p) -independent function for $\alpha \geq 1$ and $p \leq m$ if it has a $w(u,v)$ -flow with $w(u,v) \geq v/\alpha - 1$ for all u, v such that $n - u + v \leq p$.*

The matrix multiplication will be used as a running example in this section. We consider any standard algorithm for multiplying two matrices of size $N \times N$, that is, we exclude Strassen algorithm as well as its variants and concentrate on algorithms performing N^3 multiplications.

Lemma 1. *The matrix multiplications of two matrices of size $N \times N$ is $(1, 2N^2, N^2, N)$ -independent.*

Proof. Note that $n = N^2$ and $p = N$ so that $n - u + v \leq p$ with $u = |X_1|$ and $v = |Y_1|$ and is indeed $|X_0| + |Y_1| \leq N$. We consider here the worst case (equality).

We consider the matrix product $C = AB$, a set of inputs X_0 (elements of A and B) and a set of outputs Y_1 (elements of C) such that $|X_0| + |Y_1| = N$. The outputs of Y_1 lie in at most $|Y_1|$ columns of C (denoted by C_1), while the inputs of X_0 lies in at most $|X_0|$ columns of A (some elements may even be in B). Thus, at least $N - |X_0|$ columns of A contains only elements in X_1 (elements that are allowed to vary), we call these columns A_1 (note that $|A_1| \geq N - |X_0| \geq |Y_1| \geq |C_1|$). We fix the entries in B in $\{0, 1\}$ so that AB is a permutation of the columns of A : some of these elements are in X_0 and should be fixed anyway, the others are in X_1 and are also allowed to be

fixed in the subfunction h . We design the permutation matrix B so that (some of) the columns of A_1 are sent to columns of C_1 . Hence, when varying the inputs of X_1 over R , the outputs of Y_1 can take up to $R^{|Y_1|}$ values, which corresponds to a $|Y_1|$ -flow, larger than the $(|Y_1| - 1)$ -flow requested by the theorem. \square

The following theorem establishes a lower bound on the number of steps based on the existence of the flow. The proof uses the fact that during an interval where b outputs are pebbled, the information coming from the variables corresponding to the initial position of the S pebbles at the beginning of the interval is limited. Hence, if f has a large flow, a lot of inputs must be read during such an interval, which gives a lower bound on the total number of steps.

Theorem 3. *Let $f : \mathcal{A}^n \mapsto \mathcal{A}^m$ be a function with a $w(u, v)$ -flow. Any pebbling of any DAG computing f that takes T steps and S space respects:*

$$T \geq \lfloor m/b \rfloor (n - d)$$

for any $b \leq m$ and d the largest integer such that $w(d, b) \leq S$.

Proof. We consider such a function f and a pebbling scheme that computes it with S pebbles and T steps. Let $b \leq m$ an integer.

We divide the pebbling schemes into phases, such that exactly b outputs are pebbled in each phase (except possibly the last phase in which less outputs may be pebbled). We consider such a regular phase. Let Y_1 be the set of outputs pebbled in this phase ($|Y_1| = b$), x_0 (resp. x_1) the number of inputs pebbled inside the phase (resp. outside the phase). Since f has a $w(u, v)$ -flow, there exists an assignment of the x_0 inputs such that the outputs in Y_1 have at least $|\mathcal{A}|^{w(n-x_0, b)}$ different values. If $w(n-x_0, b) > S$, there are more possible outputs in Y_1 than the possible values computed at the beginning of the phase, represented by the (at most) S pebbles located in the graph in the beginning of the phase. As the other x_0 inputs are supposed fixed in the flow, this contradicts the assumption on f . Hence x_1 cannot be larger than d (by definition of d). Thus, the number of inputs pebbled during the phase satisfies $x_0 \geq n - d$.

Since there are $\lfloor m/b \rfloor$ regular phases, the number of times the inputs are pebbled is at least $\lfloor m/b \rfloor (n - d)$ and so is the total number of steps in the pebbling scheme. \square

For independant function, choosing b appropriately gives the following corollary:

Corollary 1. *Let $f : \mathcal{A}^n \mapsto \mathcal{A}^m$ be (α, n, m, p) -independant. For every pebbling of every DAG computing f using S pebbles and T steps, we have*

$$\lceil \alpha(S + 1) \rceil T \geq mp/4$$

Proof. By definition f has a $w(u, v)$ -flow satisfying $w(u, v) > v/\alpha - 1$ for all u, v with $n - u + v \leq p$. We choose $b = v = \lceil \alpha(S + 1) \rceil$, and thus have $v/\alpha - 1 \geq S$. We define d as in the previous theorem (larger integer with $w(d, b) \leq S$). We first assume that $(n - d) + b \leq p$. By definition of the flow, we then have $w(d, b) > S$, which contradicts the definition of d . Hence, we have $(n - d) + b > p$, that is $n - d > p - \lceil \alpha(S + 1) \rceil$.

We notice that $\lfloor m/x \rfloor \geq (m - x + 1)/x$. The bound of the previous theorem now translates to: translates to:

$$T \geq \frac{(m - \lceil \alpha(S + 1) \rceil + 1)}{\lceil \alpha(S + 1) \rceil} (p - \lceil \alpha(S + 1) \rceil)$$

We consider two cases:

- Case 1: $\lceil \alpha(S + 1) \rceil \leq p/2$. Then $p - \lceil \alpha(S + 1) \rceil \geq p/2$ and $m - \lceil \alpha(S + 1) \rceil \geq m/2$ as $m \geq p$, which leads to the desired bound.
- Case 2: $\lceil \alpha(S + 1) \rceil > p/2$, then $\lceil \alpha(S + 1) \rceil T \geq mp/2$ since $T > m$ (all outputs need to be pebbled).

\square

We now get back to the multiplication of two $N \times N$ matrices. We proved earlier that it was $(1, 2N^2, N^2, N)$ -independent. This leads to the following bound:

Theorem 4. *Every pebbling strategy for any straight-line program computing the multiplication of two $N \times N$ matrices uses a space S and time T respecting the following inequality:*

$$(S + 1)T \geq N^3/4$$

Note that the standard algorithm that computes each $C_{i,j}$ values one after the other by computing the inner (dot) product of two vectors reaches this bound (at least asymptotically): each dot product can be performed with 4 pebbles in time $4N$, and there are N^2 of them to compute, hence $ST = 16N^3$ for this algorithm.

Material used for this course and further documentation

Refer to chapter 10 (“Space-Time Tradeoffs”) of the book “Models of Computation” by J. Savage [5] for more information on the black pebble game.

References

- [1] Philippe Flajolet, Jean-Claude Raoult, and Jean Vuillemin. The number of registers required for evaluating arithmetic expressions. *Theoretical Computer Science*, 9(1):99–125, 1979.
- [2] John R. Gilbert, Thomas Lengauer, and Robert Endre Tarjan. The pebbling problem is complete in polynomial space. *SIAM J. Comput.*, 9(3):513–524, 1980.
- [3] Thomas Lengauer. Black-white pebbles and graph separation. *Acta Informatica*, 16(4):465–475, 1981.
- [4] Friedhelm Meyer auf der Heide. A comparison of two variations of a pebble game on graphs. *Theoretical Computer Science*, 13(3):315–322, 1981.
- [5] John E. Savage. *Models of Computation: Exploring the Power of Computing*. Addison-Wesley, 1998.
- [6] Ravi Sethi. Complete register allocation problems. In *Proceedings of the 5th Annual ACM Symposium on Theory of Computing (STOC'73)*, pages 182–195, New York, NY, USA, 1973. ACM Press.
- [7] Ravi Sethi and J.D. Ullman. The generation of optimal code for arithmetic expressions. *J. ACM*, 17(4):715–728, 1970.
- [8] Arthur N Strahler. Hypsometric (area-altitude) analysis of erosional topography. *Geological Society of America Bulletin*, 63(11):1117–1142, 1952.
- [9] Sivan Toledo. A survey of out-of-core algorithms in numerical linear algebra. In *External Memory Algorithms and Visualization*, pages 161–180. American Mathematical Society Press, 1999.
- [10] Peter van Emde Boas and Jan van Leeuwen. Move rules and trade-offs in the pebble game. In *Theoretical Computer Science 4th GI Conference*, pages 101–112. Springer, 1979.