

Efficient Scheduling of Tiled Iteration Spaces onto a Fixed Size Parallel Architecture

Maria Athanasaki, Evangelos Koukis and Nectarios Koziris

National Technical University of Athens,
School of Electrical and Computer Engineering,
Computing Systems Laboratory
e-mail: {maria, vkoukis, nkoziris}@cslab.ntua.gr

Abstract. In this paper we propose several alternative methods for the compile time scheduling of Tiled Iteration Spaces onto a cluster of SMP nodes with a fixed number of nodes. We elaborate on the distribution of computations among processors, provided that we have chosen either a non-overlapping communication mode, which involves successive computation and communication steps, or an overlapping communication mode, which supposes a pipelined, concurrent execution of communication and computations. In order to utilize the available processors as efficiently as possible, we can either adopt a cyclic assignment schedule, or assign neighboring tiles to the same CPU, or adapt the size and shape of tiles to the underlying architecture size. We theoretically and experimentally compare four different schedules, so as to design one which achieves the minimum total execution time, depending on the cluster configuration, the internal characteristics of the underlying architecture and the iteration space size and shape. Experimental results in a small Linux cluster of SMP nodes, using the GM library over the Myrinet high performance interconnect, illustrate the merits and drawbacks of each approach.

1 Introduction

One of the most intriguing areas in the field of parallel computing is the efficient mapping of loop iterations onto a parallel system, taking into account its architectural internal characteristics. In order to achieve maximum acceleration of the final program, one of the key issues to be considered is minimization of the communication overhead. Among other solutions given in the literature, researchers have dealt with this problem by applying the supernode or tiling transformation. Supernode partitioning of the iteration space was proposed by Irigoin and Triolet in [13]. They introduced the initial model of loop tiling and gave conditions for a tiling transformation to be valid. Later, Ramanujam and Sadayappan in [19] showed the equivalence between the problem of finding a set of extreme vectors for a given set of dependence vectors and the problem of finding a tiling transformation H that produces valid, deadlock-free tiles. The problem of determining the optimal shape was surveyed, and more accurate conditions were also given by others, as in [22], [5], [11], [12].

Nevertheless, all above approaches do not investigate an ideal execution scheme in order to reduce the overall tiled space completion time. Hodzic and Shang [10] proposed

a method to correlate optimal tile size and shape, based on overall completion time reduction. Their approach considers a straightforward execution scheme, where each processor executes all tiles along a specific dimension, by interleaving computation and communication phases. In [9] we proposed an alternative method for the problem of scheduling the tiles to single CPU nodes. Each atomic tile execution involves a communication and a computation phase and this is repeatedly done for all time planes. We compacted this sequence of communication and computation phases, by overlapping them for the different processors. The proposed method acts like enhancing the performance of a processor’s datapath with pipelining [18], because a processor computes its tile at k time step and concurrently receives data from all neighbors to use them at $k + 1$ time step and sends data produced at $k - 1$ time step.

In [4] we extended the method proposed in [9] for executing Tiled Iteration Spaces in SMP nodes (Symmetric MultiProcessors). We grouped together neighboring tiles along a hyperplane. Hyperplane-grouped tiles are concurrently executed by the CPUs of the same SMP node. In this way, we eliminate the need for tile synchronization and communication between intranode CPUs. As far as scheduling of groups is concerned, we take advantage of the overlapping execution scheme of [9] in order to “hide” each group communication volume within the respective computation volume. Under the above execution scheme, the iteration space involves the overlapped execution of communication and computation phases between successive groups of tiles. We thus avoid most of the communication overhead by allowing for actual computation to communication overlapping.

However, the proposed schedule assumes the availability of an unlimited number of SMP nodes. In [3] Andronikos et al. have proposed an assignment scheme onto a fixed number of nodes, however the complexity of evaluating which tiles should be assigned to which node is too high. In [6], [7] Boulet et al. and Calland et al. have theoretically proven the optimality of a cyclic assignment of 2-dimensional tiles onto a fixed number of single CPU nodes. On the other hand, Manjikian and Abdelrahman have presented in [16] an alternative method for scheduling Tiled Iteration Spaces onto a fixed number of SMP nodes, without taking into account that there is no need for communication among CPUs of the same SMP node, since the data required are located in the node’s shared memory.

In this paper, we propose four different methods for scheduling tiled iteration spaces onto an existing clustered system with a fixed number of SMP nodes: the cyclic assignment schedule, the mirror assignment schedule, the cluster assignment schedule and the retiling schedule. Firstly, we adapt the method proposed in [4] for a cluster of SMPs with a fixed number of nodes. We discuss the approaches of [6], [7], [16] and generalize them for n -dimensional Spaces, taking into account the particularity of immediate exchange of data among CPUs of the same SMP node. In addition, we apply to all four schedules, two alternative execution schemes, the overlapping [9] and the non-overlapping [10] communication scheme and we discuss the merits and drawbacks of each combined approach.

The rest of this paper is organized as follows: In Section 2 we provide the mathematical background and terminology used throughout the paper and we briefly revise concepts, such as grouping transformation, described in our previous work. In Section

3 we adapt the theory proposed in [4] for a fixed number of SMP nodes, using four different mapping methods. In Section 4 we use some exemplary Iteration Spaces, so as to experimentally delve into the advantages of each schedule. We deduce that our experimental results strongly confirm our theory. Finally, in Section 5 we summarize our conclusions.

2 Algorithmic Model - Grouping Transformation

Our proposed method can be applied to any code segment which can be transformed into a Tiled Iteration Space. However, without lack of generality, in this paper our model consists of perfectly nested FOR-loops with uniform data dependencies, as in [4],[9].

Throughout this paper, the following notation is used: N is the set of natural numbers, n is the number of nested FOR-loops of the algorithm. $J^n \subset Z^n$ is the set of loop indices: $J^n = \{\mathbf{j}(j_1, \dots, j_n) | j_i \in Z \wedge l_i \leq j_i \leq u_i, 1 \leq i \leq n\}$. Each point in this n -dimensional integer space is a distinct instantiation of the loop body.

In a Supernode or Tiling Transformation, the Iteration space J^n is partitioned into identical n -dimensional parallelepiped areas (tiles or supernodes) formed by n independent families of parallel hyperplanes. Tiling transformation is defined by the n -dimensional square matrix H . Each row vector of H is perpendicular to one family of hyperplanes forming the tiles. Dually, tiling transformation can be defined by n linearly independent vectors, which are the sides of the tiles. Similar to matrix H , matrix P contains the side-vectors of a tile as column vectors. It holds $P = H^{-1}$.

Formally, tiling transformation is defined as follows:

$$r : Z^n \longrightarrow Z^{2n}, r(\mathbf{j}) = \begin{bmatrix} [H\mathbf{j}] \\ \mathbf{j} - H^{-1}[H\mathbf{j}] \end{bmatrix},$$

where $[H\mathbf{j}]$ identifies the coordinates of the tile that index point $\mathbf{j}(j_1, j_2, \dots, j_n)$ is mapped to and $\mathbf{j} - H^{-1}[H\mathbf{j}]$ gives the coordinates of \mathbf{j} within that tile relative to the tile origin. The resulting Tile Space J^S is defined as follows: $J^S = \{\mathbf{j}^S | \mathbf{j}^S = [H\mathbf{j}], \mathbf{j} \in J^n\}$. It can be also written as $J^S = \{\mathbf{j}^S(j_1^S, \dots, j_n^S) | j_i^S \in Z \wedge l_i^S \leq j_i^S \leq u_i^S, 1 \leq i \leq n\}$, where l_i^S, u_i^S can be directly computed from the functions $l_1, \dots, l_n, u_1, \dots, u_n$ and the tiling matrix H , as described in [1], [8].

In the rest of this paper we shall consider that the non-overlapping and overlapping execution schemes, extensively discussed in [9] (sections 3,4), [20] and the concept of grouping, introduced in [4] (section 4) are known.

For example, let us consider an n -dimensional rectangular Tile Space J^S , whose bounds are defined as follows: $0 \leq j_i^S < u_i^S, i = 1, \dots, n$ and $u_1^S \geq u_i^S, i = 2, \dots, n$. It is grouped according to the matrices

$$P^G = \begin{bmatrix} 1 & -m_2 & \dots & -m_n \\ 0 & m_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & m_n \end{bmatrix}, H^G = (P^G)^{-1} = \begin{bmatrix} 1 & 1 & \dots & 1 \\ 0 & \frac{1}{m_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{1}{m_n} \end{bmatrix} \quad (1)$$

Thus, a tile \mathbf{j}^S belongs to group $\mathbf{j}^G = (\sum_{i=1}^n j_i^S, \lfloor \frac{j_2^S}{m_2} \rfloor, \dots, \lfloor \frac{j_n^S}{m_n} \rfloor)^T$. Following the overlapping execution scheme, if there are as many SMP nodes as required, it will be executed in the SMP node (j_2^G, \dots, j_n^G) during the time step $t = \sum_{i=1}^n j_i^G = \sum_{i=1}^n j_i^S + \sum_{i=2}^n \lfloor \frac{j_i^S}{m_i} \rfloor$ (according to the scheduling vector $\mathbf{II}^G = (1, 1, \dots, 1)$). Thus, the number of steps required for the completion of the algorithm will be:

$$P_{\text{unlimited-overlap}} = 1 + \sum_{i=1}^n (u_i^S - 1) + \sum_{i=2}^n \lfloor \frac{u_i^S - 1}{m_i} \rfloor \Rightarrow$$

$$P_{\text{unlimited-overlap}} = \sum_{i=1}^n u_i^S + \sum_{i=2}^n \lceil \frac{u_i^S}{m_i} \rceil - 2n + 2 \quad (2)$$

Similarly, if we follow the non-overlapping execution scheme, then group $\mathbf{j}^G = (\sum_{i=1}^n j_i^S, \lfloor \frac{j_2^S}{m_2} \rfloor, \dots, \lfloor \frac{j_n^S}{m_n} \rfloor)^T$ will be executed during the time step $t = j_1^G = \sum_{i=1}^n j_i^S$ (according to the scheduling vector $\mathbf{II}^G = (1, 0, \dots, 0)$). Thus, the number of steps required for the completion of the algorithm will be:

$$P_{\text{unlimited-non-overlap}} = 1 + \sum_{i=1}^n (u_i^S - 1) \Rightarrow$$

$$P_{\text{unlimited-non-overlap}} = \sum_{i=1}^n u_i^S - n + 1 \quad (3)$$

3 Scheduling onto a Fixed Number of SMPs

3.1 Cyclic Assignment to SMPs

In [6], [7] the optimality of the cyclic assignment of 2-dimensional tiles onto a fixed number of processors was theoretically proven. However, the calculations in [6], [7] did not take into account the communication overhead involved. Generalizing this approach for n -dimensional tiles and for clusters of SMP nodes, we consider that the available SMP nodes form a virtual $(n-1)$ -dimensional mesh of $p_2 \times \dots \times p_n = p$ SMP nodes. We cyclically assign the groups to the SMP nodes. That is, we assign group \mathbf{j}^G to the SMP node $(j_2^G \% p_2, \dots, j_n^G \% p_n)$, as indicated in Fig. 1. Therefore, each SMP node will execute $\lceil \frac{u_2^S}{m_2 p_2} \rceil \times \dots \times \lceil \frac{u_n^S}{m_n p_n} \rceil = \lceil \frac{u_2^G}{p_2} \rceil \times \dots \times \lceil \frac{u_n^G}{p_n} \rceil$ rows of groups (where $u_i^G = \lceil \frac{u_i^S}{m_i} \rceil$, $i = 2, \dots, n$).

If the rows of groups assigned to an SMP node, are executed in lexicographic order, the row (x, j_2^G, \dots, j_n^G) will be executed in the SMP node $(j_2^G \% p_2, \dots, j_n^G \% p_n)$ after $\sum_{i=2}^n \left[\lceil \frac{j_i^G}{p_i} \rceil \prod_{k=i+1}^n \lceil \frac{u_k^G}{p_k} \rceil \right]$ rows, imposing a latency of $u_1^S \sum_{i=2}^n \left[\lceil \frac{j_i^G}{p_i} \rceil \prod_{k=i+1}^n \lceil \frac{u_k^G}{p_k} \rceil \right]$ time steps. In addition, as deduced from Fig. 1, the location of a group, relatively to the corresponding chunk origin, is $(j_1^{G'}, j_2^G \% p_2, \dots, j_n^G \% p_n)$, where $j_1^{G'} = j_1^S + \sum_{i=2}^n j_i^S \% m_i p_i$.

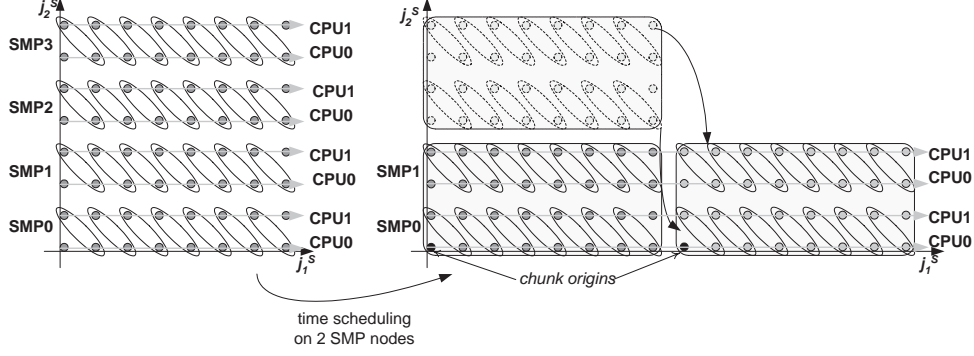


Fig. 1. Cyclic assignment to SMPs

Therefore, if the underlying architecture allows for concurrent execution of computations and communication, following the overlapping execution scheme, group j^G will be computed during the time step $t(j^G) = j_1^{G'} + \sum_{i=2}^n j_i^G \% p_i + u_1^S \sum_{i=2}^n \left[\lfloor \frac{j_i^G}{p_i} \rfloor \prod_{k=i+1}^n \lceil \frac{u_k^S}{p_k} \rceil \right]$.

Thus, the number of steps required for the completion of the algorithm will be $P_{cyclic-overlap} = \max t(j^G) - \min t(j^G) + 1 \Rightarrow$

$$P_{cyclic-overlap} = \sum_{i=2}^n \left[(u_i^S - 1) \% m_i p_i + (\lceil \frac{u_i^S}{m_i} \rceil - 1) \% p_i \right] + u_1^S \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \quad (4)$$

The first term of the right-hand part in formula (4) represents the time required for filling the pipeline, while the second term corresponds to the time each processor is busy executing calculations.

If we should do with a conventional communication architecture as node interconnect (i.e. without NIC support for relieving the CPU from the communication burden), following the non-overlapping execution scheme, group j^G will be computed during the time step $t(j^G) = j_1^{G'} + u_1^S \sum_{i=2}^n \left[\lfloor \frac{j_i^G}{p_i} \rfloor \prod_{k=i+1}^n \lceil \frac{u_k^S}{p_k} \rceil \right]$. Thus, the number of steps required for the completion of the algorithm will be $P_{cyclic-non-overlap} = \max t(j^G) - \min t(j^G) + 1 \Rightarrow$

$$P_{cyclic-non-overlap} = \sum_{i=2}^n [(u_i^S - 1) \% m_i p_i] + u_1^S \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \quad (5)$$

3.2 Mirror Assignment to SMPs

Let us consider another schedule, if we assign the tiles to SMP nodes as indicated in Fig. 2. That is, we assign group j^G to the SMP node

$$\left(\begin{array}{c} j_2^G \% p_2 \text{ if } \text{even}(j_2^G / p_2) \\ (p_2 - 1) - j_2^G \% p_2 \text{ if } \text{odd}(j_2^G / p_2) \end{array} \right), \dots, \left(\begin{array}{c} j_n^G \% p_n \text{ if } \text{even}(j_n^G / p_n) \\ (p_n - 1) - j_n^G \% p_n \text{ if } \text{odd}(j_n^G / p_n) \end{array} \right)$$




Fig. 2. Mirror assignment to SMPs

This schedule has the advantage that there is no need for data transfer along the boundaries of chunks of tiles, thus less time is wasted for communication.

Then, like the cyclic assignment schedule, if the chunks of groups are executed in lexicographic order, the chunk containing row (x, j_2^G, \dots, j_n^G) will be executed after $\sum_{i=2}^n \left[\lfloor \frac{j_i^G}{p_i} \rfloor \prod_{k=i+1}^n \lceil \frac{u_k^G}{p_k} \rceil \right]$ chunks. The latency imposed by each of the previous chunks, when combining the mirror assignment schedule with the overlapping execution scheme, is greater than the respective one when applying the cyclic assignment schedule. It, thus, equals to $u_1^S + \sum_{i=2}^n [(m_i + 1)p_i] - 2n + 2$, as the computation of a whole chunk should be finished before the computation of the next chunk starts. In addition, as deduced from Fig. 2, the position of a group, relatively to the corresponding chunk origin, is $(j_1^{G'}, j_2^G \% p_2, \dots, j_n^G \% p_n)$, where $j_1^{G'} = j_1^S + \sum_{i=2}^n j_i^S \% m_i p_i$. Therefore, group j^G will be computed during the time step

$t(j^G) = j_1^{G'} + \sum_{i=2}^n j_i^G \% p_i + \left[u_1^S + \sum_{i=2}^n [(m_i + 1)p_i] - 2n + 2 \right] \sum_{i=2}^n \left[\lfloor \frac{j_i^G}{p_i} \rfloor \prod_{k=i+1}^n \lceil \frac{u_k^G}{p_k} \rceil \right]$. Thus,

the number of steps required for the completion of the algorithm will be $P_{mirror-overlap} = \max t(j^G) - \min t(j^G) + 1 \Rightarrow$

$$\begin{aligned}
 P_{mirror-overlap} = & \sum_{i=2}^n \left[(u_i^S - 1) \% m_i p_i + (\lceil \frac{u_i^S}{m_i} \rceil - 1) \% p_i \right] - \sum_{i=2}^n [(m_i + 1)p_i] + 2n - 2 + \\
 & + \left[u_1^S + \sum_{i=2}^n [(m_i + 1)p_i] - 2n + 2 \right] \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil
 \end{aligned} \tag{6}$$

If there is no shortage of processors ($u_i^S \leq m_i p_i, \forall i = 2, \dots, n$), the proposed schedules are equivalent. Otherwise, it can be easily deduced from (4),(6) that $P_{cyclic-overlap} < P_{mirror-overlap}$. Their difference is due to the fact that, following the mirror assignment schedule, every time the computation of a chunk finishes and the computation of the next one starts, there are some idle time steps for some of the processors, as indicated in Fig. 2. The cyclic schedule is thus preferable to the mirror one.

Similarly, following the non-overlapping execution scheme, group j^G will be computed during the time step $t(j^G) = j_1^{G'} + (u_1^S + \sum_{i=2}^n m_i p_i - n + 1) \sum_{i=2}^n \left[\lfloor \frac{j_i^G}{p_i} \rfloor \prod_{k=i+1}^n \lceil \frac{u_k^G}{p_k} \rceil \right]$. Thus, the number of steps required for the completion of the algorithm will be $P_{mirror-non-overlap} = \max t(j^G) - \min t(j^G) + 1 \Rightarrow$

$$P_{mirror-non-overlap} = \sum_{i=2}^n [(u_i^S - 1) \% m_i p_i] - \sum_{i=2}^n m_i p_i + n - 1 + \left[u_1^S + \sum_{i=2}^n m_i p_i - n + 1 \right] \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \quad (7)$$

It can be deduced from (5),(7) that $P_{cyclic-non-overlap} \leq P_{mirror-non-overlap}$. However, since the communication overhead is not hidden under the computation time, this schedule may sometimes result in a shorter total execution time, due to better exploitation of the available bandwidth. In particular, if there are only two SMP nodes along a dimension, no SMP node should both send and receive data along that dimension. Thus, the communication overhead will be halved.

3.3 Cluster Assignment to SMPs




Fig. 3. Cluster assignment to SMPs

Alternatively, following the approach of [16], generalizing it for n -dimensional spaces and taking into account that there is no need for communication among processors of the same SMP node, we may assign neighboring rows of tiles to the same CPU, as

indicated in Fig. 3. In order to achieve this schedule, we cluster together neighboring tiles $(j_1^S, j_2^S, \dots, j_n^S)$, mapping them to a supertile or “TILE” labeled as $(j_1^S, \lfloor \frac{j_2^S}{\lceil \frac{u_2^S}{m_2 p_2} \rceil} \rfloor, \dots, \lfloor \frac{j_n^S}{\lceil \frac{u_n^S}{m_n p_n} \rceil} \rfloor)$. Thus, the corresponding “GROUP” will be $j^G = (j_1 + \sum_{i=2}^n \lfloor \frac{j_i^S}{\lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor, \lfloor \frac{j_2^S}{\lceil \frac{u_2^S}{m_2 p_2} \rceil} \rfloor, \dots, \lfloor \frac{j_n^S}{\lceil \frac{u_n^S}{m_n p_n} \rceil} \rfloor)$ and, following the overlapping execution scheme, it will be executed during the time “STEP” $t(j^S) = j_1 + \sum_{i=2}^n \lfloor \frac{j_i^S}{\lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor + \sum_{i=2}^n \lfloor \frac{j_i^S}{m_i \lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor$. As a “TILE” consists of $\prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil$ tiles, a “STEP” will be equivalent to $\prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil$ time steps (excluding the DMA initialization and synchronization time). Thus, the total number of steps required for the completion of the algorithm will be

$$P_{cluster-overlap} = \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil (\max t(j^S) - \min t(j^S) + 1) \Rightarrow$$

$$P_{cluster-overlap} = \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \left(u_1^S - 2n + 2 + \sum_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil + \sum_{i=2}^n \lfloor \frac{u_i^S}{m_i \lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor \right) \quad (8)$$

Lemma 1. *It holds that $P_{cyclic-overlap} \leq P_{cluster-overlap}$.*

Proof: *Omitted due to lack of space.* ◻

Thus, this schedule results to a worse number of execution steps than the previous one. Their difference is due to the fact that, in this schedule, the filling of the pipeline is slower. In case that $u_1^S \gg u_i^S$ ($i = 2, \dots, n$), the time each processor is busy, outflanks the pipeline filing time and it holds that $P_{cyclic-overlap} \simeq P_{cluster-overlap}$. However, the previous mathematical lemma has not taken into consideration the time required for the initialization of messages and for synchronization. Since the cluster assignment schedule requires less messages to be sent and less synchronization, in some cases it may be practically proven more efficient, as we will show in §4.

Similarly, following the non-overlapping execution scheme, tile $(j_1^S, j_2^S, \dots, j_n^S)$, corresponding to “GROUP” $j^G = (j_1 + \sum_{i=2}^n \lfloor \frac{j_i^S}{\lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor, \lfloor \frac{j_2^S}{\lceil \frac{u_2^S}{m_2 p_2} \rceil} \rfloor, \dots, \lfloor \frac{j_n^S}{\lceil \frac{u_n^S}{m_n p_n} \rceil} \rfloor)$ is executed during the time “STEP” $t(j^S) = j_1 + \sum_{i=2}^n \lfloor \frac{j_i^S}{\lceil \frac{u_i^S}{m_i p_i} \rceil} \rfloor$. A computation “subSTEP” is equivalent to $\prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil$ computation substeps, but a communication “subSTEP” is equivalent to less than $\prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil$ communication substeps. In particular, if the communication load is equal along all communication dimensions (as resulted by the method proposed in [22]), the amount of data to be transferred, as indicated in Fig. 4, is $\prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \sum_{i=2}^n \frac{1}{(n-1) \lceil \frac{u_i^S}{m_i p_i} \rceil} \leq \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil$ times the communication load of a tile. Thus, the total number of steps required for the completion of the algorithm will be




Fig. 4. Clustering communication

$$P_{cluster-non-overlap} = C (\max t(j^S) - \min t(j^S) + 1) \quad (\text{where } 1 \leq C \leq \prod_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil) \Rightarrow$$

$$P_{cluster-non-overlap} = C \left(u_1^S - n + 1 + \sum_{i=2}^n \lceil \frac{u_i^S}{m_i p_i} \rceil \right) \quad (9)$$

In conclusion, comparing to the cyclic assignment schedule, this method has the drawback of slower pipeline filling. However, it results to less communication overhead, which significantly reduces the total execution time, especially when the non-overlapping execution scheme is applied.

3.4 Retiling




Fig. 5. Retiling

A more efficient schedule can be obtained, if we adapt the size of tiles to the available number of SMPs (Fig. 5). That is, we retile the initial Iteration Space, so as to get $u_i^{S'} = m_i p_i$, ($i = 2, \dots, n$) and $u_1^{S'} = u_1^S \prod_{i=2}^n \frac{u_i^S}{m_i p_i}$. Then, the size of a “new” tile will be

equal to the size of an “old” tile and, consequently, a “new” computation step will be equivalent to an “old” computation step. Following the overlapping execution scheme, the number of time steps required for the completion of the algorithm, according to the formula (2), will be $P_{retile-overlap} = \sum_{i=1}^n u_i^{S'} + \sum_{i=2}^n \lceil \frac{u_i^{S'}}{m_i} \rceil - 2n + 2 \Rightarrow$

$$P_{retile-overlap} = \sum_{i=2}^n [(m_i + 1) p_i] - 2n + 2 + u_1^S \prod_{i=2}^n \frac{u_i^S}{m_i p_i} \quad (10)$$

Using the non-overlapping execution scheme, the number of time steps required for the completion of the algorithm, according to the formula (3), will be

$$P_{retile-non-overlap} = \sum_{i=1}^n u_i^{S'} - n + 1 \Rightarrow$$

$$P_{retile-non-overlap} = \sum_{i=2}^n m_i p_i - n + 1 + u_1^S \prod_{i=2}^n \frac{u_i^S}{m_i p_i} \quad (11)$$

From (5),(11), we can deduce that $P_{retile-non-overlap} \leq P_{cyclic-non-overlap}$. In addition, a “new” computation substep is equivalent to an “old” computation substep, but a “new” communication substep is equivalent to less than an “old” communication substep. In particular, as in the cluster assignment schedule, if the communication load is equal along all communication dimensions, the amount of data to be transferred is $\sum_{i=2}^n \frac{1}{(n-1) \frac{u_i^S}{m_i p_i}} \leq 1$ times the communication load of an “old” tile.

In conclusion, in every case this schedule is preferable to previously proposed ones, assuming that there are no factors constraining the tile shape, such as false sharing, or cache locality [14], [15], [21]. It can fully exploit the computational power of all the SMP nodes and it achieves a perfect load balance, without imposing any additional complexity to the initial schedule. But if, apart from parallel scheduling, there are other factors constraining the tile size and shape, this schedule will be proven to be inefficient, since it totally reorganizes the execution order of iterations.

4 Experimental Results

4.1 Experimental Platform

In order to evaluate the proposed methods, we use a Linux SMP cluster with 2 identical nodes. Each node has 1GB of RAM and 2 Pentium III @ 1266 MHz CPUs. The cluster nodes communicate through a Myrinet high performance interconnect, using the GM low level message passing system.

In order to utilize the available processors in each SMP node as efficiently as possible, our implementation uses one multi-threaded process per SMP, with the number of threads equal to the number of CPUs. Multithreading support is based on the LinuxThreads library. Threads executing on the same SMP communicate using shared memory, eliminating the need for message passing. For the data exchange between processes executing on different SMPs, Myricom’s GM version 1.6.3 is used [17]. GM is a

low-level message passing library for Myrinet. It comprises a library used by userspace programs, an OS driver (in our case, a Linux kernel module) and a Myrinet Control Program (MCP), which is executed on the LANai, the embedded RISC microprocessor on the Myrinet NIC. The GM driver is used during the execution of a userspace process to open and close *ports* and to allocate and free memory suitable for DMA transfers. A port is a communication endpoint, used as the interface between a userspace process and the NIC. Having opened a port, a process can communicate directly with the NIC without the need for system calls, bypassing the operating system. Thus, all data exchange is performed directly to and from userspace buffers.

To provide flow control between the host and the NIC, sending and receiving messages is regulated by *tokens*. Initially, a process possesses a finite number of send and receive tokens. To be able to receive a message, the process must provide GM with a buffer in DMAable memory, relinquishing a receive token. When a message is received, the DMA engine on the Myrinet NIC places it directly into the userspace buffer. The process polls for new messages and retrieves the receive token when a message arrives. The same applies to sending messages: The process relinquishes a send token by requesting the transmission of a message from a userspace buffer, then retrieves it when the send operation completes and an appropriate send completion callback function is executed by GM. As the data exchange between the host memory and the NIC is undertaken by the DMA engine on the NIC, without involving the CPU, overlapping of communication with computation is possible.

4.2 Experimental Data

We performed several series of experiments in order to evaluate and compare the practical speedups obtained using each one of the four alternative schedules, combined with both the alternative execution schemes. Our test application code was the following:

```
for (i=1; i<=X; i++)
  for(j=1; j<=Y; j++)
    for(k=1; k<=Z; k++)
      A[i][j][k]=func(A[i-1][j][k],
                    A[i][j-1][k],A[i][j][k-1]);
```

where A is an array of $X \times Y \times Z$ floats and $X, Y \ll Z$. Without lack of generality, we consider, as a tile, a rectangle with ij , ik and jk sides. The dimension k is the largest one, so all tiles along the k -axis are mapped onto the same processor, as proposed in [2], [9]. Each tile has i , j dimensions equal to x and k dimension equal to z . Thus, there are $\frac{X}{x}$ tiles along dimension i , $\frac{Y}{x}$ tiles along dimension j and $\frac{Z}{z}$ tiles along dimension k .

After implementing all four schedules in combination with both execution schemes, as described by the pseudo-code of Table 1, we measured the performance of all schedules and compared it with their theoretically expected performance. For various tile sizes, we conducted a series of experiments for each combination of schedule and execution scheme, varying the iteration space size. In Figs 6-8 we have plotted our experimental results along with the respective theoretical curves. As a measure of performance,

Table 1. Execution Schemes Implementation

Non Overlapping Execution Scheme	Overlapping Execution Scheme
<pre> gm_provide_receive_buffer() do poll the GM event queue process the event until data received </pre>	<pre> Pre-computation Part of Communication If on first tile Execute a non-overlapping receive gm_provide_receive_buffer() for tile (t1 + 1, t2, t3) gm_send_with_callback() for tile (t1 - 1, t2, t3) </pre>
<pre> gm_send_with_callback() do poll the GM event queue process the event until data sent Barrier for Threads in SMP </pre>	<pre> Post-computation Part of Communication do poll the GM event queue process the event until send & receive completed Barrier for Threads in SMP If on last tile Execute a non-overlapping send </pre>

we have used the ratio of the speedup obtained to the best possible speedup. That is, we have depicted the ratio of the speedup obtained to the number of processors used. Thus, the closer a plot is to 1, the more efficient a schedule is. As can be seen in Figs 6-8, the practical completion times of our experiments differ to our theoretical predictions by at most 3%. For the overlapping communication schedules, this can be attributed to both the DMA engine on the Myrinet NIC and the CPU trying to access data in memory simultaneously.

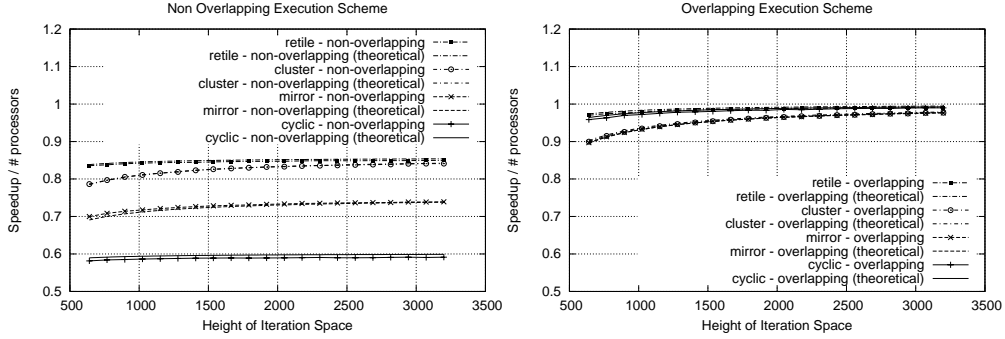


Fig. 6. Experimental Data: Tile Size $32 \times 32 \times 32$

One can easily deduce that in almost all cases, the retiling schedule achieves the best performance, both theoretically and experimentally. This result was expected, since the retiling schedule absolutely adjusts tiles to the existing configuration of a cluster. However, in our experiments we have eliminated the effect of cache miss penalties by using small iteration space widths. If our iteration space dimensions which are not assigned to the same processor were too long, the retiling schedule could have destroyed the data locality achieved by optimally selected small tiles.

Note also that the cluster assignment schedule using tile size x is equivalent to the retiling schedule using tile size $4x$. This was expected, considering that by construction

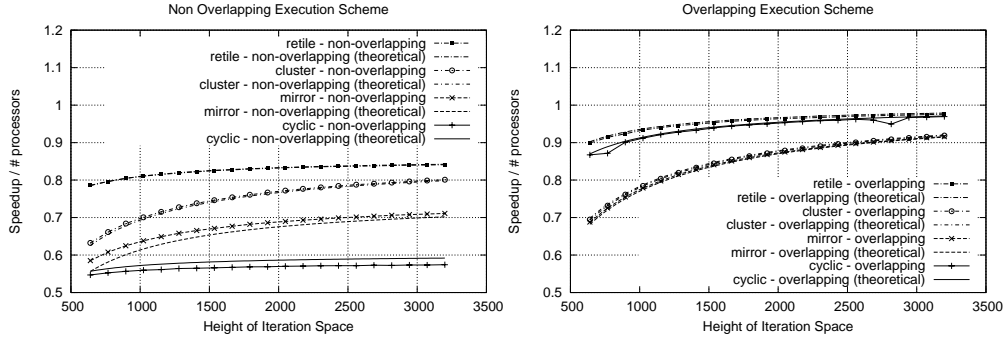


Fig. 7. Experimental Data: Tile Size $128 \times 32 \times 32$

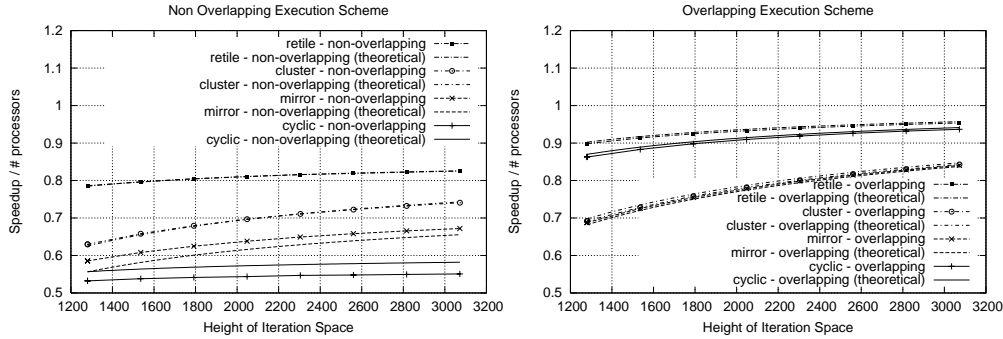


Fig. 8. Experimental Data: Tile Size $256 \times 32 \times 32$

the iterations executed and the data sent in these two cases are the same. What differs is the execution order of iterations but here we have eliminated the cache misses overhead, in order to test our schedules' optimality and not data locality.

When following the non-overlapping execution scheme, the difference among the performance of the four schedules is mainly due to the volume of the data to be transferred. As depicted in Fig. 9, the mirror assignment schedule involves double the communication of retiling and cluster assignment schedule, while the cyclic assignment schedule involves 6 times the same communication volume.

When following the overlapping execution scheme, since the communication volume is hidden under computation, their difference is due to the time steps that each SMP has to stall waiting for the required data to arrive. The number of these time steps is the same for both the retiling and the cyclic assignment schedule. However, using the cluster or the mirror assignment schedule, results in a multiple number of idle time steps, as depicted in Figs 1, 2.

In addition, note that all schedules achieve better performance for long Iteration Spaces. This is due to the fact that when the mapping dimension of the Iteration Space is comparatively short, the time required for the last processor to start computing after the first data have arrived, is not negligible in comparison to the total execution time.




Fig. 9. Communication among SMPs

5 Conclusions

In this paper, we presented and experimentally compared four different methods for scheduling Tiled Iteration Spaces onto a cluster with a fixed number of SMPs. We concluded that the most efficient schedule is in most cases obtained when we adapt the size and shape of tiles to the size of the underlying architecture (retiling schedule). However, in case it is not possible, or it is not desired, since tiles are already optimally selected considering data locality [14], [15], [21], we propose either a cyclic assignment schedule, or clustering together neighboring tiles and handling them as a super-tile. The cyclic approach is preferable when the communication and computation substeps can be overlapped. In the opposite case, we propose the cluster assignment schedule, which considerably reduces the volume of data to be transferred.

Acknowledgement

Maria Athanasaki is partially supported by a research student scholarship, awarded by the A.S. Onassis public benefit foundation.

References

1. C. Ancourt and F. Irigoin. Scanning Polyhedra with DO Loops. In *Proceedings of the Third ACM SIGPLAN Symposium on Principles & Practice of Parallel Programming (PPoPP)*, pages 39–50, Williamsburg, VA, Apr 1991.
2. T. Andronikos, N. Koziris, G. Papakonstantinou, and P. Tsanakas. Optimal Scheduling for UET/UET-UCT Generalized N-Dimensional Grid Task Graphs. *Journal of Parallel and Distributed Computing*, 57(2):140–165, May 1999.
3. T. Andronikos, N. Koziris, G. Papakonstantinou, and P. Tsanakas. Optimal Scheduling for UET-UCT Grids Into Fixed Number of Processors. In *Proceedings of 8th Euromicro Workshop on Parallel and Distributed Processing (PDP2000)*, IEEE Press, pages 237–243, Rhodes, Greece, Jan 2000.
4. M. Athanasaki, A. Sotiropoulos, G. Tsoukalas, and N. Koziris. Pipelined Scheduling of Tiled Nested Loops onto Clusters of SMPs using Memory Mapped Network Interfaces. In *Proceedings of the ACM Supercomputing 2002 (SC2002)*, Baltimore, Nov 2002.
5. P. Boulet, A. Darté, T. Risset, and Y. Robert. (Pen)-ultimate tiling? *INTEGRATION, The VLSI Journal*, 17:33–51, 1994.

6. P. Boulet, J. Dongarra, Y. Robert, and F. Vivien. Tiling for Heterogeneous Computing Platforms. Technical Report UT-CS-97-373, University of Tennessee, Knoxville, 1997.
7. P. Y. Calland, J. Dongarra, and Y. Robert. Tiling with Limited Resources. In L. Thiele, J. Fortes, K. Vissers, V. Taylor, T. Noll, and J. Teich, editors, *Application Specific Systems, Architectures, and Processors, ASAP'97, IEEE Computer Society Press*, pages 229–238, Jul 1997. Extended version available on the web at <http://www.ens-lyon.fr/~yrobert>.
8. G. Goumas, M. Athanasaki, and N. Koziris. Automatic Code Generation for Executing Tiled Nested Loops Onto Parallel Architectures. In *Proceedings of the ACM Symposium on Applied Computing (SAC 2002)*, pages 876–881, Madrid, Spain, Mar 2002.
9. G. Goumas, A. Sotiropoulos, and N. Koziris. Minimizing Completion Time for Loop Tiling with Computation and Communication Overlapping. In *Proceedings of IEEE Int'l Parallel and Distributed Processing Symposium (IPDPS'01)*, San Francisco, Apr 2001.
10. E. Hodzic and W. Shang. On Supernode Transformation with Minimized Total Running Time. *IEEE Trans. on Parallel and Distributed Systems*, 9(5):417–428, May 1998.
11. E. Hodzic and W. Shang. On Time Optimal Supernode Shape. *IEEE Trans. on Parallel and Distributed Systems*, 13(12):1220–1233, Dec 2002.
12. K Hogstedt, L. Carter, and J. Ferrante. On the Parallel Execution Time of Tiled Loops. *IEEE Trans. on Parallel and Distributed Systems*, 14(3):307–321, Mar 2003.
13. F. Irigoien and R. Triolet. Supernode Partitioning. In *Proceedings of the 15th Ann. ACM SIGACT-SIGPLAN Symp. Principles of Programming Languages*, pages 319–329, San Diego, California, Jan 1988.
14. M. Kandemir, J. Ramanujam, and A. Choudary. Improving Cache Locality by a Combination of Loop and Data Transformations. *IEEE Trans. on Parallel and Distributed Systems*, 48(2):159–167, Feb 1999.
15. M. Lam, E. Rothberg, and M. Wolf. The Cache Performance and Optimizations of Blocked algorithms. In *Second International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 63–74, Santa Clara, California, Apr 1991.
16. N. Manjikian and T. S. Abdelrahman. Exploiting Wavefront Parallelism on Large-Scale Shared-Memory Multiprocessors. *IEEE Trans. on Parallel and Distributed Systems*, 12(3):259–271, Mar 2001.
17. Myricom. GM: A Message-Passing System for Myrinet Networks, 2002. <http://www.myri.com/scs/GM/doc/html>.
18. D. Patterson and J. Hennessy. *Computer Organization & Design. The Hardware/Software Interface*, pages 364–367. Morgan Kaufmann Publishers, San Francisco, CA, 1994.
19. J. Ramanujam and P. Sadayappan. Tiling Multidimensional Iteration Spaces for Multi-computers. *Journal of Parallel and Distributed Computing*, 16:108–120, 1992.
20. A. Sotiropoulos, G. Tsoukalas, and N. Koziris. Enhancing the Performance of Tiled Loop Execution onto Clusters using Memory Mapped Network Interfaces and Pipelined Schedules. In *Proceedings of the 2002 Workshop on Communication Architecture for Clusters (CAC'02), Int'l Parallel and Distributed Processing Symposium (IPDPS'02)*, Fort Lauderdale, Florida, Apr 2002.
21. M. Wolf and M. Lam. A Data Locality Optimizing Algorithm. In *ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI)*, Toronto, Ontario, Jun 1991.
22. J. Xue. Communication-Minimal Tiling of Uniform Dependence Loops. *Journal of Parallel and Distributed Computing*, 42(1):42–59, 1997.