# A Data Locality Optimizing Algorithm 

Michael E. Wolf and Monica S. Lam<br>Computer Systems Laboratory<br>Stanford University, CA 94305


#### Abstract

This paper proposes an algorithm that improves the locality of a loop nest by transforming the code via interchange, reversal, skewing and tiling. The loop transformation algorithm is based on two concepts: a mathematical formulation of reuse and locality, and a loop transformation theory that unifies the various transforms as unimodular matrix transformations.

The algorithm has been implemented in the SUIF (Stanford University Intermediate Format) compiler, and is successful in optimizing codes such as matrix multiplication, successive over-relaxation (SOR), LU decomposition without pivoting, and Givens QR factorization. Performance evaluation indicates that locality optimization is especially crucial for scaling up the performance of parallel code.


## 1 Introduction

As processor speed continues to increase faster than memory speed, optimizations to use the memory hierarchy efficiently become ever more important. Blocking [9] or tiling [18] is a well-known technique that improves the data locality of numerical algorithms $[1,6,7,12,13]$. Tiling can be used for different levels of memory hierarchy such as physical memory, caches and registers; multi-level tiling can be used to achieve locality in multiple levels of the memory hierarchy simultaneously.

To illustrate the importance of tiling, consider the example of matrix multiplication:

```
for }\mp@subsup{I}{1}{}:=1\mathrm{ to }
    for }\mp@subsup{I}{2}{}:=1\mathrm{ to }
        for }\mp@subsup{I}{3}{}:=1\mathrm{ to }
```

[^0]$$
\mathrm{C}\left[I_{1}, I_{3}\right]+=\mathrm{A}\left[I_{1}, I_{2}\right] \quad * \mathrm{~B}\left[I_{2}, I_{3}\right] ;
$$

In this code, although the same row of $C$ and $B$ are reused in the next iteration of the middle and outer loop, respectively, the large volume of data used in the intervening iterations may replace the data from the register file or the cache before it can be reused. Tiling reorders the execution sequence such that iterations from loops of the outer dimensions are executed before completing all the iterations of the inner loop. The tiled matrix multiplication is

```
for }I\mp@subsup{I}{2}{}:=1\mathrm{ to }n\mathrm{ by }
    for III :=1 to n by s
        for }\mp@subsup{I}{1}{}:=1\mathrm{ to }
            for I}\mp@subsup{I}{2}{}:=I\mp@subsup{I}{2}{}\mathrm{ to min(II 2 +s-1,n)
            for }\mp@subsup{I}{3}{}:=I\mp@subsup{I}{3}{}\mathrm{ to min(II_ +s-1, n)
                C}[\mp@subsup{I}{1}{},\mp@subsup{I}{3}{}]+=\textrm{A}[\mp@subsup{I}{1}{},\mp@subsup{I}{2}{}]*\textrm{B}[\mp@subsup{I}{2}{},\mp@subsup{I}{3}{}]
```

Tiling reduces the number of intervening iterations and thus data fetched between data reuses. This allows reused data to still be in the cache or register file, and hence reduces memory accesses. The tile size $s$ can be chosen to allow the maximum reuse for a specific level of memory hierarchy.

The improvement obtained from tiling can be far greater than from traditional compiler optimizations. Figure 1 shows the performance of $500 \times 500$ matrix multiplication on an SGI 4D/380 machine. The SGI 4D/380 has eight MIPS/R3000 processors running at 33 Mhz . Each processor has a 64 KB direct-mapped first-level cache and a 256 KB direct-mapped second-level cache. We ran four different experiments: without tiling, tiling to reuse data in caches, tiling to reuse data in registers [5], and tiling for both register and caches. For cache tiling, the data are copied into consecutive locations to avoid cache interference [12].

Tiling improves the performance on a single processor by a factor of 2.75 . The effect of tiling on multiple processors is even more significant since it not only reduces the average data access latency but also the required memory bandwidth. Without cache tiling, contention over the


Figure 1: Performance of $500 \times 500$ double precision matrix multiplication on the SGI 4D/380. Cache tiles are 64 $\times 64$ iterations and register tiles are $4 \times 2$.
memory bus limits the speedup to about 4.5 times. Cache tiling permits speedups of over seven for eight processors, achieving an impressive speed of 64 MFLOPS when combined with register tiling.

### 1.1 The Problem

The problem addressed in this paper is the use of loop transformations such as interchange, skewing and reversal to improve the locality of a loop nest. Matrix multiplication is a particularly simple example because it is both legal and advantageous to tile the entire nest. In general, it is not always possible to tile the entire loop nest. Some loop nests may not be tilable. Sometimes it is necessary to apply transformations such as interchange, skewing and reversal to produce a set of loops that are both tilable and advantageous to tile.

For example, consider the example of an abstraction of hyperbolic PDE code in Figure 2(a). Suppose the array in this example is larger than the memory hierarchy level of interest; the entire array must be fetched anew for each iteration of the outermost loop. Due to dependences, the loops must first be skewed before they can be tiled. This is also equivalent to finding non-rectangular tiles. Figure 2 contains the entire derivation of the tiled code, which we will use to illustrate our locality algorithm in the rest of the paper.

There are two major representations used in loop transformations: distance vectors and direction vectors [2, 17]. Loops whose dependences can be summarized by distance vectors are special in that it is advantageous, possible and easy to tile all loops [10, 15]. General loop nests, whose dependences are represented by direction vectors, may not be tilable in their entirety. The data locality problem addressed in this paper is to find the best combination of loop interchanges, skewing, reversal and tiling that maximizes the data locality within loop nests, subject to the constraints of direction and distance vectors.

Research has been performed on both the legality and the desirability of loop transformations with respect to data locality. Early research on optimizing loops with direction vectors concentrated on the legality of pairwise transformations, such as when it is legal to interchange a pair of loops. However, in general, it is necessary to apply a series of primitive transformations to achieve goals such as parallelism and data locality. This has led to work on combinations of primitive transforms. For example, Wolfe [18] shows how to determine when a loop nest can be tiled; two-dimensional tiling can be achieved via a pair of transformations known as "strip-mine and interchange" [14] or "unroll and jam" [5]. Wolfe also shows that skewing can make a pair of loops tilable. Banerjee discusses general unimodular transforms for two-deep loop nests [4]. A technique used in practice to handle general $n$-dimensional loop nests is to determine a priori the sequence of loop transforms to attempt. This technique is inadequate because certain transformations, such as loop skewing, may not improve code, but may enable other optimizations that do so. Which of these to perform will depend on which other optimizations will be enabled: the desirablity of a transformation cannot be evaluated locally. Furthermore, the correct ordering of optimizations are highly program dependent.

On the desirability of tiling, previous work concentrated on how to determine the cache performance and tune the loop parameters for a given loop nest. Porterfield gives an algorithm for estimating the hit rate of a fully-associative LRU (least recently used replacement policy) cache of a given size [14]. Gannon et al. uses reference windows to determine the minimum memory locations necessary to maximize reuse in a loop nest [8]. These evaluation functions are useful for comparing the locality performance after applying transformations, but do not suggest the transformations to apply when a series of transformations may first need to be applied before tiling becomes feasible and useful.

If we were to use these evaluation functions to find the suitable transformations, we would need to search the transformation space exhaustively. The previously proposed method of enumerating the all possible combinations of legal transformations is expensive and not even possible if there are infinitely many combinations, as is the case when we include skewing.

## (a): Extract dependence information

```
for \(I_{1}:=0\) to 5 do
    for \(I_{2}:=0\) to 6 do
        \(\mathrm{A}\left[I_{2}+1\right]:=1 / 3 *\left(\mathrm{~A}[k]+\mathrm{A}\left[I_{2}+1\right]+\mathrm{A}\left[I_{2}+2\right]\right) ;\)
            \(D=\{(0,1),(1,0),(1,-1)\}\).
```

(b): Extract locality information


Uniformly generated set $=\left\{\mathrm{A}\left[I_{2}\right], \mathrm{A}\left[I_{2}+1\right], \mathrm{A}\left[I_{2}+2\right]\right\}$.

| reuse category | reuse vector space | potential reuse factor |
| :---: | :---: | :---: |
| self-temporal | $\operatorname{span}\{(1,0)\}$ | $s$ |
| self-spatial | $\operatorname{span}\{(1,0),(0,1)\}$ | $l$ |
| group | $\operatorname{span}\{(1,0),(0,1)\}$ | 3 |

Loops carrying reuse $=\left\{I_{1}, I_{2}\right\}$.
(c): Search transformation space

| localized space | transformation | sources of locality | accesses per iteration |
| :---: | :---: | :---: | :---: |
| $\operatorname{span}\{(0,1)\}$ | $T=\left[\begin{array}{cc}1 & 0 \\ 0 & 1\end{array}\right]$ | self-spatial, group | $1 / l$ |
| $\operatorname{span}\{(1,0)\}$ | not possible | - | - |
| $\operatorname{span}\{(1,0),(0,1)\}$ | $T=\left[\begin{array}{cc}1 & 0 \\ 1 & 1\end{array}\right]$ | self-temporal, self-spatial, group | $1 /(l s)$ |

The best legal choice is to tile both $I_{1}$ and $I_{2}$.
(d): Skew to make inner loop nest fully permutable

```
for I}\mp@subsup{I}{1}{\prime}:=0\mathrm{ to }5\mathrm{ do
    for }\mp@subsup{I}{2}{\prime}:=\mp@subsup{I}{1}{\prime}\mathrm{ to }6+\mp@subsup{I}{1}{\prime}\mathrm{ do
    A}[\mp@subsup{I}{2}{\prime}-\mp@subsup{I}{1}{\prime}+1]:=1/3* (A[\mp@subsup{I}{2}{\prime}-\mp@subsup{I}{1}{\prime}
            + A [I2
            T=[\begin{array}{ll}{1}&{0}\\{1}&{1}\end{array}]
            D'=TD={(0,1),(1, 1),(1,0)}
```


(e): Final code

```
for }I\mp@subsup{I}{2}{\prime}:=0\mathrm{ to }11\mathrm{ by 2 do
    for }\mp@subsup{I}{1}{\prime}:=0\mathrm{ to 5 do
        for I2
            A[I}\mp@subsup{I}{2}{\prime}-\mp@subsup{I}{1}{\prime}+1]:=1/3 * (A[\mp@subsup{l}{2}{\prime}-\mp@subsup{I}{1}{\prime}
            + A [I2 - II +1] + A[II - I' +2]);
```



Figure 2: Example of the locality optimization algorithm on a hyperbolic PDE-style loop nest.

### 1.2 An Overview

This paper focuses on maximizing data locality at the cache level. Although the basic principles in memory hierarchy optimization are similar for all levels, each level has slightly different characteristics, requiring slightly different considerations. Caches usually have small set associativity, so cache data conflicts can cause desired data to be replaced. We have found that the performance of tiled code fluctuates dramatically with the size of the data matrix, due to cache interference [12]. We show that this effect can be mitigated by copying reused data to consecutive locations before the computation, or choosing the tile size according to the matrix size. Both of these optimizations can be performed after code transformation, and thus cache interference need not be considered at code transformation time.
Another major difference between caches and registers is their capacity. To fit all the data used in a tile into the faster level of memory hierarchy, transformations that increase the dimensionality of the tile may reduce the tile size. However, as we will show, the reduction of memory accesses can be a factor of $s^{d}$ when $d$-dimensional tiles of lengths $s$ are used. Thus for typical cache sizes and loop depths found in practice, increasing the dimensionality of the tile will reduce the total number of memory access, even though the length of a tile side may be smaller for larger dimensional tiles. Thus the choice of tile size can be postponed until after the application of optimizations to increase the dimensionality of locality.

The problem addressed in this paper is the choice of loop transforms to increase data locality. We describe a locality optimization algorithm that applies a combination of loop interchanges, skewing, reversal and tiling to improve the data locality of loop nests. The analysis is applicable to array references whose indices are affine functions of the loop indices. The transformations are applicable to loops with not just distance dependence vectors, but also direction vectors as well. Our locality optimization is based on two results: a new transformation theory and a mathematical formulation of data locality.

Our loop transformation theory unifies common loop transforms including interchange, skewing, reversal and their combinations, as unimodular matrix transforms. This unification reduces the legality of all compound transformations to satisfying the same simple constraints. In this way, a loop transformation problem can be formulated as solving for the transformation matrix that maximizes an objective function, subjected to a set of constraints. This matrix model has previously been applied only to distance vectors. We have extended the framework to handle direction vectors as well.

The second result is a formulation of an objective function for data locality. We introduce the concepts of a reuse vector space to capture the potential of data locality optimization for a given loop nest. This formulation collapses
the space of transformed code into equivalence classes, hence allowing pruning of the search for the best transformation.

Our locality algorithm uses the evaluation function and the legality constraints to reduce the search space of the transforms. Unfortunately, finding the optimal transformation still requires a complex algorithm that is exponential in the loop nest depth. We have devised a heuristic algorithm that works well for common cases found in practice.

We have implemented the algorithm in the SUIF (Stanford University Intermediate Format) compiler. Our algorithm applies unimodular and tiling transforms to loop nests, handles non-rectangular loop bounds, and generates non-uniform tiles to handle non-perfectly loop nests. It is successful in tiling numerical algorithms such as matrix multiplication, successive over-relaxation (SOR), LU decomposition without pivoting, and Givens QR factorization. For simplicity, we assume here that all loops are perfectly nested. That is, all computation is nested in the innermost loop.

In Section 2, we discuss our dependence representation and the basics of unimodular transformations. How tiling takes advantage of reuse in an algorithm is discussed in Section 3. In Section 4, we describe how to identify and evaluate reuse in a loop nest. We use those results to formulate an algorithm to improve locality. Finally, we present some experimental data on tiling for a cache.

## 2 A Loop Transformation Theory

While individual loop transformations are well understood, ad hoc techniques have typically been used in combining them to achieve a particular goal. Our loop transformation theory offers a foundation for deriving compound transformations efficiently [16]. We have previously shown the use of this theory to maximizing the degree of parallelism in a loop nest; we will demonstrate its applicability to data locality in this paper.

### 2.1 The Iteration Space

In this model, a loop nest of depth $n$ corresponds to a finite convex polyhedron of iteration space $\mathcal{Z}^{n}$, bounded by the loop bounds. Each iteration in the loop corresponds to a node in the polyhedron, and is identified by its index vector $\vec{p}=\left(p_{1}, p_{2}, \ldots, p_{n}\right) ; p_{i}$ is the loop index of the $i$ loop in the nest, counting from the outermost to innermost loop. The iterations are therefore executed in lexicographic order of their index vectors. That is, if $\overrightarrow{p_{2}}$ is lexicographically greater than $\overrightarrow{p_{1}}$, written $\overrightarrow{p_{2}} \succ \overrightarrow{p_{1}}$, iteration $\overrightarrow{p_{2}}$ executes after iteration $\overrightarrow{p_{1}}$.

Our dependence representation is a generalization of distance and direction vectors. A dependence vector in an $n$-nested loop is denoted by a vector $\vec{d}=\left(d_{1}, d_{2}, \ldots, d_{n}\right)$. Each component $d_{i}$ is a possibly infinite range of integers,
represented by $\left[d_{i}^{\min }, d_{i}^{\text {max }}\right]$, where

$$
d_{i}^{\min } \in \mathcal{Z} \cup\{-\infty\}, d_{i}^{\max } \in \mathcal{Z} \cup\{\infty\} \text { and } d_{i}^{\min } \leq d_{i}^{\max }
$$

A single dependence vector therefore represents a set of distance vectors, called its distance vector set:

$$
\mathcal{E}(\vec{d})=\left\{\left(e_{1}, \ldots, e_{n}\right) \mid e_{i} \in \mathcal{Z} \text { and } d_{i}^{\min } \leq e_{i} \leq d_{i}^{\max }\right\}
$$

Each of the distance vector defines a set of edges on pairs of nodes in the iteration space. Iteration $\overrightarrow{p_{2}}$ depends on iteration $\overrightarrow{p_{1}}$, and thus must execute after $\overrightarrow{p_{1}}$, if for some distance vector $\vec{e}, \overrightarrow{p_{2}}=\overrightarrow{p_{1}}+\vec{e}$. By definition, since $\overrightarrow{p_{2}} \succ$ $\overrightarrow{p_{1}}, \vec{e}$ must therefore be lexicographically greater than $\overrightarrow{0}$, or simply, lexicographically positive.

The dependence vector $\vec{d}$ is also a distance vector if each of its components is a degenerate range consisting of a singleton value, that is, $d_{i}^{\min }=d_{i}^{\max }$. For short, we simply denote such a range with the value itself. There are three common ranges found in practice: $[1, \infty]$ denoted by ' + ', $[-\infty,-1]$ denoted by ' - ', and $[-\infty, \infty]$ denoted by ' $\pm$ '. They correspond to the previously defined directions of ' $<$ ', ' $>$ ', and ' $*$ ', respectively [17].

We have extended the definition of vector operations to allow for ranges in each of the component. In this way, we can manipulate a combination of distances and directions simply as vectors. This is needed to support the matrix transform model, discussed below.

Our model differs from the previously used model in that all our dependence vectors are represented as lexicographically positive vectors. In particular, consider a strictly sequential pair of loops such as the one below:

$$
\begin{aligned}
& \text { for } I_{1}:=0 \text { to } n \text { do } \\
& \text { for } I_{2}:=0 \text { to } n \text { do } \\
& b:=g(b) ;
\end{aligned}
$$

The dependence of this program would previously be represented as (' $*$ ',' $*$ '). In our model, we represent them as a pair of lexicographically positive vectors, $\left(0,{ }^{\prime}+\prime\right),\left({ }^{\prime}+\prime, ' \pm\right.$ ' $)$. The requirement that all dependences are lexicographically greatly simplifies the legality tests for loop transformations. The dependence vectors define a partial order on the nodes in the iteration space, and any topological ordering on the graph is a legal execution order, as all dependences in the loop are satisfied.

### 2.2 Unimodular Loop Transformations

With dependences represented as vectors in the iteration space, loop transformations such as interchange, skewing and reversal, can be represented as matrix transformations.

Let us illustrate the concept with the simple example of an interchange on a loop with distances. A loop interchange transformation maps iteration ( $p_{1}, p_{2}$ ) to iteration ( $p_{2,1}$ ). In matrix notation, we can write this as

$$
\left[\begin{array}{ll}
0 & \mathbf{1} \\
1 & 0
\end{array}\right]\left[\begin{array}{l}
p_{1} \\
p_{2}
\end{array}\right]=\left[\begin{array}{l}
p_{2} \\
p_{1}
\end{array}\right]
$$

The elementary permutation matrix $\left[\begin{array}{ll}0 & 1 \\ 1 & 0\end{array}\right]$ thus performs the loop interchange transformation on the iteration space.

Since a matrix transformation $T$ is a linear transformation on the iteration space, $T \vec{p}_{2}-T \vec{p}_{1}=T\left(\vec{p}_{2}-\vec{p}_{1}\right)$. Therefore, if $\vec{d}$ is a distance vector in the original iteration space, then $T \vec{d}$ is a distance vector in the transformed iteration space. Thus in loop interchange, the dependence vector $\left(d_{1}, d_{2}\right)$ is mapped into

$$
\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right]\left[\begin{array}{l}
d_{1} \\
d_{2}
\end{array}\right]=\left[\begin{array}{l}
d_{2} \\
d_{1}
\end{array}\right]
$$

in the transformed space. Therefore if the transformed dependence vector remains lexicographically positive, the interchange is legal.

The loop reversal and skewing transform can similarly be represented as matrices [3, 4, 16]. (An example of skewing is shown in Figure 2(d).) These matrices are unimodular matrices, that is, they are square matrices with integral components and a determinant of one or negative one. Because of these properties, the product of two unimodular matrices is unimodular, and the inverse of a unimodular matrix is unimodular, so that combinations of unimodular loop transformations and inverses of unimodular loop transformations are also unimodular loop transformations. Under this formulation, there is a simple legality test for all transforms.

Theorem 2.1. Let $D$ be the set of distance vectors of $a$ loop nest. A unimodular transformation $T$ is legal if and only if $\forall \vec{d} \in D: T \vec{d} \succ \overrightarrow{0}$.

The elegance of this theory helps reduce the complexity of the implementation. Once the dependences are extracted, the derivation of the compound transform simply consists of matrix and vector operations. After the transformation is determined, a straightforward algorithm applies the transformation to the loop bounds and derives the final code.

## 3 The Localized Vector Space

It is important to distinguish between reuse and locality. We say that a data item is reused if the same data is used in multiple iterations in a loop nest. Thus reuse is a measure that is inherent in the computation and not dependent on the particular way the loops are written. This reuse may not lead to saving a memory access if intervening iterations flush the data out of the cache between uses of the data.

For example, reference $\mathrm{A}\left[I_{2}\right]$ in Figure 3 touches different data within the innermost loop, but reuses the same elements across the outer loop. More precisely, the same data $A\left[I_{2}\right]$ is used in iterations $\left(I_{1}, I_{2}\right), 1 \leq I_{1} \leq n$. There is reuse, but the reuse is separated by accesses to $n-1$

```
for }\mp@subsup{I}{1}{}:=1\mathrm{ to }n\mathrm{ do
    for }\mp@subsup{I}{2}{}:=1\mathrm{ to }n\mathrm{ do
        f(A[I I], A[I I ]);
```

Figure 3: A simple example.
other data. When $n$ is large, the data is removed from the cache before it can be reused, and there is no locality. Therefore, a reuse does not guarantee locality.

Specifically, if the innermost loop contains a large number of iterations and touches a large number of data, only the reuse within the innermost loop can be exploited. We can apply a unimodular transformation to improve the amount of data reused in the innermost loop. However, as shown in this example, reuse can sometime occur along multiple dimensions of the iteration space. To exploit multi-dimensional reuse, unimodular transformations must be coupled with tiling.

### 3.1 Tiling

In general, tiling transforms an $n$-deep loop nest into a $2 n$-deep loop nest where the inner $n$ loops execute a compiler-determined number of iterations. Figure 4 shows the code after tiling the example in Figure 2(a), using a tile size of $2 \times 2$. The two innermost loops execute the iterations within each tile, represented as $2 \times 2$ squares in the figure. The two outer loops, represented by the two axes in the figure, execute the 12 tiles. As the outer loop nests of tiled code controls the execution of the tiles, we will refer to them as the controlling loops. When we say tiling, we refer to the partitioning of the iteration space into rectangular blocks. Non-rectangular blocks are obtained by first applying unimodular transformations to the iteration space and then applying tiling.

Like all transformations, it is not always possible to tile. Loops $I_{i}$ through $I_{j}$ in a loop nest can be tiled if they are fully permutable [11, 16]. Loops $I_{i}$ through $I_{j}$ in a loop nest are fully permutable if and only if all dependence vectors are lexicographically positive and for each dependence vector, either $\left(d_{1}, \ldots, d_{i-1}\right)$ is lexicographically positive, or the $i$ th through $j$ th components of $\vec{d}$ are all non-negative. For example, the components of dependences in Figure 2(b) are all non-negative, and the two loops are therefore fully permutable and tilable. Full permutability is also very useful for improving parallelism [16], so parallelism and locality are often compatible goals.

After tiling, both groups of loops, the loops within a tile and the loops controlling the tiles, remain fully permutable. For example, in Figure 4, loops $I I_{1}^{\prime}$ and $I I_{2}^{\prime}$ can be interchanged, and so can $I_{1}^{\prime}$ and $I_{2}^{\prime}$. By interchanging $I_{1}^{\prime}$ and $I_{2}^{\prime}$, the loops $I I_{2}^{\prime}$ and $I_{2}^{\prime}$ can be trivially coalesced to produce the code in Figure 2(e). This transformation has previously been known as "strip-mine and interchange".
"Unroll and jam" is yet another equivalent form to tiling. Since all loops within a tile are fully permutable, any loop in an $n$-dimensional tile can be chosen to be the coalesced loop.

### 3.2 Tiling for Locality

When we tile two innermost loops, we execute only a finite number of iterations in the innermost loop before executing iterations from the next loop. The tiled code for the example in Figure 3 is:

$$
\begin{aligned}
& \text { for } I I_{2}:=1 \text { to } n \text { by } s \text { do } \\
& \text { for } I_{1}:=1 \text { to } n \text { do } \\
& \text { for } I_{2}:=I I_{2} \text { to } \max \left(n, I I_{2}+s-1\right) \text { do } \\
& f\left(\mathrm{~A}\left[I_{1}\right], A\left[I_{2}\right]\right) \text {; }
\end{aligned}
$$

We choose the tile size such that the data used within the tile can be held within the cache. In this example, as long as $s$ is smaller than the cache size (in words), $\mathrm{A}\left[I_{2}\right]$ will still be present in the cache when it is reused. Thus tiling increases the number of dimensions in which reuse can be exploited. We call the iterations that can exploit reuse the localized iteration space. In the example above, reuse is exploited for loops $I_{1}$ and $I_{2}$ only, so the localized iteration space includes only those loops.

The depth of loop nesting and the number of variables accessed within a loop body are small compared to typical cache sizes. Therefore we should always be able to choose suitable tile sizes such that the reused data can be stored in a cache. Since we do not need the tile size to determine if reuse is possible, we abstract away the loop bounds of the localized iteration space, and characterize the localized iteration space as a localized vector space. Thus we say that tiling the loop nest of Figure 3 results in a localized vector space of $\operatorname{span}\{(0,1),(1,0)\}$.

In general, if $n$ is the first loop with a large bound, counting from innermost to outermost, then reuse occurring within the inner $n$ loops can be exploited. Therefore the localized vector space of a tiled loop is simply that of the innermost tile, whether the boundary between the controlling loops and the loops within the tile be coalesced or not.

## 4 Evaluating Reuse and Locality

Since unimodular transformations and tiling can modify the localized vector space, knowing where there is reuse in the iteration space can help guide the search for the transformation that delivers the best locality. Also, to choose between alternate transformations that exploit different reuses in a loop nest, we need a metric to quantify locality for a specific localized iteration space.

```
for }H\mp@subsup{I}{1}{\prime}:=0\mathrm{ to }5\mathrm{ by 2 do
    for }I\mp@subsup{I}{2}{\prime}:=0\mathrm{ to 11 by 2 do
        for II := III to min( I', +1, 5) do
        for I2 := max}(I\mp@subsup{I}{1}{\prime},I\mp@subsup{I}{2}{\prime}) to min(6+I, , II_ ' +1) do
            A[IL}]:=1/3*(\textrm{A}[\mp@subsup{I}{2}{\prime}-1]+\textrm{A}[\mp@subsup{I}{2}{\prime}]+\textrm{A}[\mp@subsup{I}{2}{\prime}+1])
```



Figure 4: Iteration space and dependences of tiled code from Figure 2(a).

### 4.1 Types of Reuse

Reuse occurs when a reference within a loop accesses the same data location in different iterations. We call this self-temporal reuse. Likewise, if a reference accesses data on the same cache line in different iterations, it is said to possess self-spatial reuse. Furthermore, different references may access the same locations. We say that there is group-temporal reuse if the references refer to the same location, and group-spatial reuse if they refer to the same cache line. Examples of each type of reuse are given below.

Let us first consider reuse within a single reference. In Figure 3, the reference $\mathrm{A}\left[I_{1}\right]$ has self-temporal reuse in the innermost loop because it accesses the same element for all iterations ( $I_{1}, I_{2}$ ), where $1 \leq I_{2} \leq n$. Similarly, A [ $I_{2}$ ] has temporal reuse in the outermost loop. In both cases, the same data is reused $n$ times; that is, the memory accesses are reduced to $1 / n$th of the original if the reuse is exploited.

Besides self-temporal reuse, A [ $\left.I_{2}\right]$ also has self-spatial reuse in the innermost loop, since each cache line is reused $l$ times, where $l$ is the cache line size. Likewise, $\mathrm{A}\left[I_{1}\right]$ has self-spatial locality in the outermost loop. Altogether, each reference has temporal reuse of a factor of $n$, and an additional spatial reuse of a factor of $l$. Therefore, in either case each cache line can be reused $n l$ times if the localized space encompasses all the reuse.

By definition, temporal reuse is a subset of spatial reuse; reusing the same location is trivially reusing the same cache line. That is, loops carrying temporal reuse also carry spatial reuse. While the same data can be reused arbitrarily many times depending on the program, the additional factor of improvement from spatial reuse is limited to a factor of $l$, where $l$ is the cache line size.

We now consider group reuse, reuse among different references. Trivially, identical references within the same loop nest will result in the same cache behavior as if there had just been one reference. Now consider, for example, the references $\mathrm{A}\left[I_{2}\right], \mathrm{A}\left[I_{2}+1\right]$ and $\mathrm{A}\left[I_{2}+2\right]$ in Figure 2. In addition to any self reuse these references might have, it is easy to see that they have a factor of three reuse in
the $I_{2}$ loop.
In contrast, consider the example in Figure 3. In this 2dimensional iteration space, the iterations that use the same data between the two groups are the $k$ th column and the $k$ th row. As it is, only iterations near the diagonal of the space can exploit locality. Furthermore, no unimodular or tiling transformation can place uses of the same data close to each other. Thus, multiple references to the same array do not necessarily result in significant locality.

The difference between the $\mathrm{A}\left[I_{2}\right], \mathrm{A}\left[I_{2}+1\right], \mathrm{A}\left[I_{2}+2\right]$ references and the $A\left[I_{1}\right], A\left[I_{2}\right]$ references is that the former set of references has similar array index functions, differing only in the constant term. Such references are known as uniformly generated references. The concept of uniformly generated references is also used by Gannon et al. [8] in estimating their reference windows.

Definition 4.1 Let $n$ be the depth of a loop nest, and $d$ be the dimensions of an array A. Two references $\mathrm{A}[\vec{f}(\vec{\imath})]$ and $\mathrm{A}[\vec{g}(\vec{\imath})]$, where $\vec{f}$ and $\vec{g}$ are indexing functions $Z^{n} \rightarrow Z^{d}$, are called uniformly generated if

$$
\vec{f}(\vec{\imath})=H \vec{\imath}+\vec{c}_{f} \text { and } \vec{g}(\vec{\imath})=H \vec{\imath}+\vec{c}_{g}
$$

where $H$ is a linear transformation and $\vec{c}_{f}$ and $\overrightarrow{\boldsymbol{c}}_{g}$ are constant vectors.

Since little exploitable reuse exists between nonuniformly generated references, we partition references in a loop nest into equivalence classes of references that operate on the same array and have the same $H$. We call these equivalence classes uniformly generated sets. In the degenerate case where a uniformly generated set consists of only one element, we have only self reuse for that reference. In the example in Figure 2, the indexing functions of the references $\mathrm{A}\left[I_{2}\right], \mathrm{A}\left[I_{2}+1\right]$ and $\mathrm{A}\left[I_{2}+2\right]$ can be written as

$$
\begin{aligned}
& {\left[\begin{array}{ll}
0 & 1
\end{array}\right]\left[\begin{array}{l}
I_{1} \\
I_{2}
\end{array}\right]+\left[\begin{array}{l}
0
\end{array}\right]} \\
& {\left[\begin{array}{ll}
0 & 1
\end{array}\right]\left[\begin{array}{l}
I_{1} \\
I_{2}
\end{array}\right]+\left[\begin{array}{l}
1
\end{array}\right]}
\end{aligned}
$$

$$
\left[\begin{array}{ll}
0 & 1
\end{array}\right]\left[\begin{array}{l}
I_{1} \\
I_{2}
\end{array}\right]+[2]
$$

respectively. These references belong to a single uniformly generated set with an $H$ of $\left[\begin{array}{ll}0 & 1\end{array}\right]$.

### 4.2 Quantifying Reuse and Locality

So far, we have discussed intuitively where reuse takes place. We now show how to identify and quantify reuse within an iteration space. We use vector spaces to represent the directions in which reuse is found; these are the directions we wish to include in the localized space. We evaluate the locality available within a loop nest using the metric of memory accesses per iteration of the innermost loop. A reference with no locality will result in one access per iteration.

### 4.2.1 Self-Temporal

Consider the self-temporal reuse for a reference $\mathrm{A}[\mathrm{H} \vec{\imath}+$ $\vec{c}]$. Iterations $\vec{\imath}_{1}$ and $\vec{\imath}_{2}$ reference the same data element whenever $H \vec{\imath}_{1}+\vec{c}=H \vec{\imath}_{2}+\vec{c}$, that is, when $H\left(\vec{\imath}_{1}-\vec{\imath}_{2}\right)=\overrightarrow{0}$. We say that there is reuse in direction $\vec{r}$ when $H \vec{r}=\overrightarrow{0}$. That is, reuse is exploited if $\vec{r}$ is included in the localized vector space. The solution to this equation is $\mathrm{ker} H$, a vector space in $\mathcal{R}^{n}$. We call this the self-temporal reuse vector space, $R_{S T}$. Thus,

$$
R_{S T}=\operatorname{ker} H .
$$

For example, the reference $\mathrm{C}\left[I_{1}, I_{3}\right]$ in the matrix multiplication in Section 1 produces:

$$
\begin{gathered}
H=\left[\begin{array}{lll}
1 & 0 & 0 \\
0 & 0 & 1
\end{array}\right] \\
R_{S T}=\operatorname{ker} H=\operatorname{span}\{(0,1,0)\} .
\end{gathered}
$$

Informally, we say there is self-temporal reuse in loop $I_{2}$. Since loop $I_{2}$ has $n$ iterations, there are $n$ reuses of $C$ in this loop nest. Similar analysis shows that $\mathrm{A}\left[I_{1}, I_{2}\right]$ has self-temporal reuse in the $I_{3}$ direction and $\mathrm{B}\left[I_{2}, I_{3}\right]$ has reuse in the $I_{1}$ direction.
In this example nest, every reuse vector space is onedimensional. In general, the reuse vector space can have zero or more dimensions. If the dimensionality is zero, then $R_{S T}=\emptyset$ and there is no self-temporal reuse. An example of a two-dimensional reuse vector space is the reference $A\left[I_{1}\right]$ within a three-deep nest with loops ( $I_{1}, I_{2}, I_{3}$ ). In general, if the number of iterations executed along each dimension of reuse is $s$, then each element is reused $s^{\operatorname{dim}\left(R_{s t}\right)}$ times.

As discussed in Section 3, a reuse is exploited only if it occurs within the localized vector space. Thus, a reference has self-temporal locality if its self-temporal reuse space $R_{S T}$ and the localized space $L$ in the code have a nonnull intersection. The dimensionality of $R_{S T} \cap L$ indicates
the quantity of self-temporal reuse utilized: the number of memory accesses is, simply,

$$
1 / s^{\operatorname{dim}\left(R_{S T} \cap L\right)},
$$

where $s$ is the number of iterations in each dimension.
Consider again the matrix multiplication code. The only localized direction in the untiled code is the innermost loop, which is span $\{(0,0,1)\}$. It coincides exactly with $R_{S T}\left(\mathrm{~A}\left[I_{1}, I_{2}\right]\right.$ ), resulting in locality for that reference. Similarly, the empty intersection with the reuse vector space of references $\mathrm{B}\left[I_{2}, I_{3}\right]$ and $\mathrm{C}\left[I_{1}, I_{3}\right]$ indicates no temporal locality for these references. In contrast, the localized vector space of the tiled matrix multiplication spans all the three loop directions. Trivially, there is a non-empty intersection with each of the reuse spaces, so self-temporal reuses are exploited for all references.

All references within a single uniformly generated set have the same $H$, and thus the same self-temporal reuse vector space. Therefore the derivation of the reuse vector spaces and their intersections with the localized space need only be performed once per uniformly generated set, not once per reference. The total number of memory accesses is the sum over that of each reference. For typical values of tile sizes chosen for caches, the number of memory accesses will be dominated by the one with the smallest values of $\operatorname{dim}\left(R_{S T} \cap L\right)$. Thus the goal of loop transforms is to maximize the dimensionality of reuse; the exact value of the tile size $s$ does not affect the choice of the transformation.

### 4.2.2 Self-Spatial

Without loss of generality, we assume that data is stored in row major order. Spatial reuse can occur only if accesses are made to the same row. Furthermore, the difference in the row index expression has to be within the cache line size. For a reference such as A $\left[I_{1}, I_{2}\right]$, the memory accesses in the $I_{2}$ loop will be reduced by a factor of $l$, where $l$ is the line size. However, there is no reuse for $\mathrm{A}\left[I_{1}, 10 * I_{2}\right]$ if the cache line is less than 10 words. For any stride $k: 1 \leq k \leq l$, the potential reuse factor is $l / k$.

All but the row index must be identical for a reference $\mathrm{A}[H \vec{\imath}+\vec{c}]$ to possess self-spatial reuse. We let $H_{S}$ be $H$ with all elements of the last row replaced by 0 . The self-spatial reuse vector space is simply

$$
R_{S S}=\operatorname{ker} H_{S} .
$$

As discussed above, temporal locality is a special case of spatial locality. This is reflected in our mathematical formulation since

$$
\operatorname{ker} H \subset \operatorname{ker} H_{S}
$$

If $R_{S S} \cap L=R_{S T} \cap L$, the reuse occurs on the same word, the spatial reuse is also temporal, and there is no additional
gain due to the prefetching of cache lines. If $R_{S S} \cap L \neq$ $R_{S T} \cap L$, however, different elements in the same row are reused. If the transformed array reference is of stride one, all the data in the cache line is also reused, resulting in a total of $1 /\left(l^{\mathrm{dim}\left(R_{S T} \cap L\right)}\right)$ memory accesses per iteration. If the stride is $k<l$, the number of accesses per iteration is $k /\left(l s^{\operatorname{dim}\left(R_{S T} \cap L\right)}\right)$. As an example, the locality for the original and fully tiled matrix multiplication is tabulated in Table 1.

### 4.2.3 Group-Temporal

To illustrate the analysis on group-temporal reuse, we use the code for a single SOR relaxation step:

$$
\begin{aligned}
& \text { for } I_{1}:=1 \text { to } n \text { do } \\
& \text { for } I_{2}:=1 \text { to } n \text { do } \\
& \mathrm{A}\left[I_{1},\right. \\
& \left.I_{2}\right]:=0.2 \star\left(\mathrm{~A}\left[I_{1}, I_{2}\right]+\mathrm{A}\left[I_{1}+1, I_{2}\right]\right. \\
& \\
& \\
& \\
& \\
& \\
& \\
& \\
& \left.+\mathrm{A}\left[I_{1}-1, I_{1}, I_{2}-1\right]\right)+\mathrm{A}\left[I_{1}, I_{2}+1\right]
\end{aligned}
$$

There are five distinct references in this innermost loop, all to the same array. The reuse between these references, can potentially reduce the number of accesses by a factor of 5 . If all temporal and spatial reuses are exploited, the total number of memory accesses per iteration is reduced from 5 to $1 / l$, where $l$ is the cache line size. With the way the nest is currently written, the localized space consists of only the $I_{2}$ direction. The reference $A\left[I_{1}, I_{2}-1\right]$ uses the same data as $\mathrm{A}\left[I_{1}, I_{2}\right]$ from the previous iteration, and $\mathrm{A}\left[I_{1}, I_{2}+1\right]$ from the second previous iteration. However, $\mathrm{A}\left[I_{1}-1, I_{2}\right]$ and $\mathrm{A}\left[I_{1}+1, I_{2}\right]$ must necessarily each access a different set of data. The number of accesses per iteration is thus $3 / l$. We now show how to mathematically calculate and factor in this group-temporal reuse.

As discussed above, group-temporal reuse need only be calculated for elements within the same uniformly generated set. Two distinct references $\mathrm{A}\left[H \vec{\imath}+\vec{c}_{1}\right]$ and $\mathrm{A}\left[H \vec{\imath}+\vec{c}_{2}\right]$ have group temporal reuse within a localized space $L$ if and only if

$$
\exists \vec{r} \in L: H \vec{r}=\vec{c}_{1}-\vec{c}_{2} .
$$

To determine whether such an $\vec{r}$ exists, we solve the system of equations $H \vec{r}=\vec{c}_{1}-\vec{c}_{2}$ to get a particular solution $\vec{r}_{p}$, if one exists. The general solution is $\operatorname{ker} H+\vec{r}_{p}$, and so there exists a $\vec{r}$ satisfying the above equation if and only if

$$
\left(\operatorname{span}\left\{\vec{r}_{p}\right\}+\operatorname{ker} H\right) \cap L \neq \operatorname{ker} H \cap L .
$$

In the SOR example above, $H$ is the identity matrix and ker $H=\emptyset$. Thus there is group reuse between two references in $L$ if and only if $\overrightarrow{r_{p}}=\vec{c}_{1}-\vec{c}_{2} \in L$. When the inner loop is $I_{2}$, there is reuse between $A\left[I_{1}, I_{2}-1\right]$ and $\mathrm{A}\left[I_{1}, I_{2}\right]$, since

$$
(0,-1)-(0,0) \in \operatorname{span}\{(0,1)\}
$$

Reuse does not exist between $\mathrm{A}\left[I_{1}, I_{2}-1\right]$ and $\mathrm{A}\left[I_{1}+\right.$ $1, I_{2}$ ], since

$$
(0,-1)-(1,0) \notin \operatorname{span}\{(0,1)\} .
$$

In fact, the references fall into three equivalence classes:

$$
\begin{gathered}
\left\{\mathrm{A}\left[I_{1}, I_{2}-1\right], \mathrm{A}\left[I_{1}, I_{2}\right], \mathrm{A}\left[I_{1}, I_{2}+1\right]\right\} \\
\left\{\mathrm{A}\left[I_{1}+1, I_{2}\right]\right\} \\
\left\{\mathrm{A}\left[I_{1}, I_{2}\right]\right\}
\end{gathered}
$$

Reuse exists and only exists between all pairs of references within each class. Thus, the effective number of memory references is simply the number of equivalence classes. In this case, there are three references instead of five, as expected.

In general, the vector space in which there is any grouptemporal reuse $R_{G T}$ for a uniformly generated set with references $\left\{\mathrm{A}\left[H \vec{\imath}+\overrightarrow{c_{1}}\right], \ldots, \mathrm{A}\left[H \vec{\imath}+\overrightarrow{c_{g}}\right]\right\}$ is defined to be

$$
R_{G T}=\operatorname{span}\left\{\vec{r}_{2}, \ldots, \vec{r}_{g}\right\}+\operatorname{ker} H
$$

where for $k=2, \ldots, g, \vec{r}_{k}$ is a particular solution of

$$
H \vec{r}=\vec{c}_{1}-\vec{c}_{k}
$$

For a particular localized space $L$, we get an additional benefit due to group reuse if and only if $R_{G T} \cap L \neq$ $R_{S T} \cap L$. To determine the benefit, we must partition the references into equivalence classes as shown above. Denoting the number of equivalence classes by $g_{T}$, the number of memory references for the uniformly generated set per iteration is $g_{T}$, instead of $g$.

### 4.2.4 Group-Spatial

Finally, there may also be group-spatial reuse. For example, in the following loop nest

$$
\begin{aligned}
& \text { for } I_{1}:=1 \text { to } n \text { do } \\
& f\left(A\left[I_{1}, 0\right], A\left[I_{1}, 1\right]\right)
\end{aligned}
$$

the two references refer to the same cache line in each iteration.

Following a similar analysis as above, the group-spatial vector space $R_{G S}$ for a uniformly generated set with references $\left\{\mathrm{A}\left[H \vec{\imath}+\overrightarrow{c_{1}}\right], \ldots, \mathrm{A}\left[H \vec{\imath}+\overrightarrow{c_{g}}\right]\right\}$ is defined to be

$$
R_{G S}=\operatorname{span}\left\{\vec{r}_{2}, \ldots, \vec{r}_{g}\right\}+\operatorname{ker} H_{S}
$$

where $H_{S}$ is $H$ with all elements of the last row replaced by 0 and for $k=2, \ldots, n, \vec{r}_{k}$ is a particular solution of

$$
H \vec{r}=\vec{c}_{S, 1}-\vec{c}_{S, k}
$$

where $\vec{c}_{S, i}$ denotes $\vec{c}_{i}$ with the last row set to 0 . The relationships between the various reuse vector spaces are therefore

$$
R_{S T} \subset R_{S S}, R_{G T} \subset R_{G S}
$$

| Reference | Reuse |  |  | Untiled $\left(L=\operatorname{span}\left\{\vec{e}_{3}\right\}\right)$ |  | Tiled $\left(L=\operatorname{span}\left\{\vec{e}_{1}, \vec{e}_{2}, \vec{e}_{3}\right\}\right)$ |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $R_{S T}$ | $R_{S S}$ | $R_{S T} \cap L$ | $R_{S S} \cap L$ | $\operatorname{cost}$ | $R_{S T} \cap L$ | $R_{S S} \cap L$ | cost |
| $\mathrm{A}\left[I_{1}, I_{2}\right]$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $\operatorname{span}\left\{\vec{e}_{2}, \vec{e}_{3}\right\}$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $1 / s$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $\operatorname{span}\left\{\vec{e}_{2}, \vec{e}_{3}\right\}$ | $1 /(l s)$ |
| $\mathrm{B}\left[I_{2}, I_{3}\right]$ | $\operatorname{span}\left\{\vec{e}_{1}\right\}$ | $\operatorname{span}\left\{\vec{e}_{1}, \vec{e}_{3}\right\}$ | $\emptyset$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $1 / l$ | $\operatorname{span}\left\{\vec{e}_{1}\right\}$ | $\operatorname{span}\left\{\vec{e}_{1}, \vec{e}_{3}\right\}$ | $1 /(l s)$ |
| $\mathrm{C}\left[I_{1}, I_{3}\right]$ | $\operatorname{span}\left\{\vec{e}_{2}\right\}$ | $\operatorname{span}\left\{\vec{e}_{2}, \vec{e}_{3}\right\}$ | $\emptyset$ | $\operatorname{span}\left\{\vec{e}_{3}\right\}$ | $1 / l$ | $\operatorname{span}\left\{\vec{e}_{2}\right\}$ | $\operatorname{span}\left\{\vec{e}_{2}, \vec{e}_{3}\right\}$ | $1 /(l s)$ |

Table 1: Self-temporal and self-spatial locality in tiled and untiled matrix multiplication. We use the vectors $\vec{e}_{1}, \vec{e}_{2}$ and $\vec{e}_{3}$ to represent $(1,0,0),(0,1,0)$ and $(0,0,1)$ respectively.

Two references with index functions $H \vec{\imath}+\overrightarrow{c_{i}}$ and $H \vec{\imath}+\overrightarrow{c_{j}}$ belong to the same group-spatial equivalence class if and only if

$$
\exists \vec{r} \in L: H_{S} \vec{r}=\vec{c}_{S, i}-\vec{c}_{S, j}
$$

We denote the number of equivalence sets by $g_{S}$. Thus, $g_{S} \leq g_{T} \leq g$, where $g$ is the number of elements in the uniformly generated set. Using a probabilistic argument, the effective number of accesses per iteration for stride one accesses, with a line size $l$, is

$$
g_{S}+\left(g_{T}-g_{S}\right) / l
$$

### 4.2.5 Combining Reuses

The union of the group-spatial reuse vector spaces of each uniformly generated set captures the entire space in which there is reuse within a loop nest. If we can find a transformation that can generate a localized space that encompasses these spaces, then all the reuses will be exploited. Unfortunately, due to dependence constraints, this may not be possible. In that case, the locality optimization problem is to find the transform that delivers the fewest memory accesses per iteration.

From the discussion above, the memory accesses per iteration for a particular transformed nest can be calculated as follows. The total number of memory accesses is the sum of the accesses for each uniformly generated set. The general formula for the number of accesses per iteration for a uniformly generated set, given a localized space $L$ and line size $l$, is

$$
\frac{g_{S}+\left(g_{T}-g_{S}\right) / l}{l^{e} s^{\operatorname{dim}\left(R_{S S} n L\right)}}
$$

where

$$
e= \begin{cases}0 & R_{S T} \cap L=R_{S S} \cap L \\ 1 & \text { otherwise } .\end{cases}
$$

We now study the full example in Figure 2(a) to illustrate the reuse and locality analysis. The reuse vector spaces of the code are summarized in Figure 2(b). For the uniformly generated set in the example,

$$
H=\left[\begin{array}{ll}
0 & 1
\end{array}\right], H_{S}=\left[\begin{array}{ll}
0 & 0
\end{array}\right]
$$

The reuse vector spaces are

$$
\begin{aligned}
R_{S T} & =\operatorname{ker} H=\operatorname{span}\{(1,0)\} \\
R_{S S} & =\operatorname{ker} H_{S}=\operatorname{span}\{(0,1),(1,0)\} \\
R_{G T} & =\operatorname{span}\{(0,1)\}+\operatorname{ker} H=\operatorname{span}\{(0,1),(1,0)\} \\
R_{G S} & =\operatorname{span}\{(0,1)\}+\operatorname{ker} H_{S}=\operatorname{span}\{(0,1),(1,0)\}
\end{aligned}
$$

When the localized space $L$ is span $\{(0,1)\}$,

$$
\begin{aligned}
& R_{S T} \cap L=\emptyset \\
& R_{S S} \cap L=\operatorname{span}\{(0,1)\} \neq R_{S T} \\
& R_{G T} \cap L=\operatorname{span}\{(0,1)\} \\
& R_{G S} \cap L=\operatorname{span}\{(0,1)\}=R_{G T}
\end{aligned}
$$

Since $g_{T}=1$, the total number of memory accesses per iteration is $1 / l$. Similar derivation shows the overall number of accesses per iteration for the localized space of $\operatorname{span}\{(1,0),(0,1)\}$ to be $1 /(l s)$.
The data locality optimization problem can now be formulated as follows:

## Definition 4.2 For a given iteration space with

## 1. a set of dependence vectors, and

## 2. uniformly generated reference sets

the data locality optimization problem is to find the unimodular andlor tiling transform, subject to data dependences, that minimizes the number of memory accesses per iteration.

## 5 An Algorithm

Our analysis of locality shows that differently transformed loops can differ in locality only if they have different localized vector spaces. That is, all transformations that generate code with the same localized vector space can be put in an equivalence class and need not be examined individually. The only feature of interest is the innermost tile; the outer loops can be executed in any legal order or orientation. Similarly, reordering or skewing the tiled loops themselves does not affect the vector space and thus need not be considered.

From the reuse analysis, we can identify a subspace that is desirable to make into the innermost tile. The question of whether there exists a unimodular transformation that is legal and creates such an innermost subspace is a difficult one. An existing algorithm that attempts to find such a transform is exponential in the number of loops [15]. The general question of finding a legal transformation that minimizes the number of memory accesses as determined by the intersection of the localized and reused vector spaces is even harder.

Although the problem is theoretically difficult, loop nests found in practice are generally simple. Using characteristics of programs as a guide, we simplify this problem by (1) reducing the set of equivalent classes, and (2) using a heuristic algorithm for finding transforms.

### 5.1 Loops Carrying Reuse

Although reuse vector spaces can theoretically be spanned by arbitrary vectors, in practice they are typically spanned by a subset of the elementary basis of the iteration space, that is, by a a subset of the loop axes. In the code below, the array A has self-temporal reuse in the vector space spanned by $(0,0,1)$; we say that the loop $I_{3}$ carries reuse.

$$
\begin{aligned}
& \text { for } I_{1} \\
& \text { for } I_{2} \\
& \quad \text { for } I_{3} \\
& \quad \mathrm{~A}\left[I_{1}, I_{2}\right]:=\mathrm{B}\left[I_{1}+I_{2}, I_{3}\right] ;
\end{aligned}
$$

On the other hand, the B array has a self-temporal reuse in the $(1,-1,0)$ direction of the iteration space, which does not correspond to either loop $I_{1}$ or $I_{2}$ but rather a combination of them. This latter situation is not as common.

Instead of using an arbitrary reuse vector space directly to guide the transformation process, we use its smallest enclosing space spanned by the elementary vectors of the iteration space. For example, the reuse vector space of $\mathrm{B}\left[I_{1}+I_{2}, I_{3}\right]$ is spanned by $(1,0,0)$ and $(0,1,0)$, the directions of loops $I_{1}$ and $I_{2}$ respectively. If we succeed in making the innermost tile include both of these directions, we will exploit the self-temporal reuse of the reference. Informally, this procedure reduces the arbitrary vector spaces to a set of loops that carry reuse.

With this simplification, a loop in the source program either carries reuse or it does not. We partition all transformations into equivalence classes according to the set of source loop directions included in the localized vector space. For example, both loops in Figure 2(a) carry reuse. Transformations are classified depending on whether the transformed innermost tile contains only the first loop, the second or both.

### 5.2 An Algorithm to Improve Locality

We first prune the search by finding those loop directions that need not or cannot be included in the localized vector
space. These loops include those that carry no reuse and can be placed outermost legally, and those that carry reuse but must be placed outermost due to legality reasons. For example, if a three-deep loop has dependences

$$
\left\{\left(1,{ }^{\prime} \pm^{\prime},{ }^{\prime} \pm ’\right),\left(0,1, ‘ \pm^{\prime}\right),\left(0,0,{ }^{\prime}+\prime\right)\right\}
$$

and carry locality in all three loops, there is no need to try all the different transformations when clearly only one loop ordering is legal.

On the remaining loops $\mathcal{I}$, we examine every subset $I$ of $\mathcal{I}$ that contains at least some reuse. For each set $I$, we try to find a legal transformation such that the loops in $I$ are tiled innermost. Among all the legal transformations, we select the one with the minimal memory accesses per iteration. Because the power set of $\mathcal{I}$ is explored, this algorithm is exponential in the number of loops in $\mathcal{I}$. However, $\mathcal{I}$, a subset of the original loops, is typically small in practice.

There are two steps in finding a transformation that makes loops $I$ innermost. We first attempt to order the outer loops-the loops in $\mathcal{I}$ but not in $I$. Any transformation applied to the loops in $\mathcal{I}-I$ that result in these loops being outermost and no dependences being violated by these loops is sufficient [16]. If that step succeeds, then we attempt to tile the $I$ loops innermost, which means finding a transformation that turns these loops into a fully permutable loop nest, given the outer nest. Solving these problems exactly is still exponential in the number of dependences. We have developed a heuristic compound transformation, known as the SRP transform, which is useful in both steps of the transformation algorithm.

The SRP transformation attempts to make a set of loops fully permutable by applying combinations of permutation, skewing and reversal [16]. If it cannot place all loops in a single fully permutable nest, it simply finds the outermost nest, and returns all remaining loops and dependences left to be made lexicographically positive. The algorithm is based upon the observations in Theorem 5.1 and Corollary 5.2.

Theorem 5.1 Let $N=\left\{I_{1}, \ldots, I_{n}\right\}$ be a loop nest with lexicographically positive dependences $\vec{d} \in D$, and $D^{i}=$ $\left\{\vec{d} \in D \mid\left(d_{1}, \ldots, d_{i-1}\right) \nsucc \overrightarrow{0}\right\}$. Loop $I_{j}$ can be made into a fully permutable nest with loop $I_{i}$, where $i<j$, via reversal andlor skewing, if

$$
\begin{aligned}
& \forall \vec{d} \in D^{i}:\left(d_{j}^{\min } \neq-\infty \wedge\left(d_{j}^{\min }<0 \rightarrow d_{i}^{\min }>0\right)\right), \text { or } \\
& \forall \vec{d} \in D^{i}:\left(d_{j}^{\max } \neq \infty \wedge\left(d_{j}^{\max }>0 \rightarrow d_{i}^{\min }>0\right)\right) .
\end{aligned}
$$

Proof: All dependence vectors for which $\left(d_{1}, \ldots, d_{i-1}\right) \succ \overrightarrow{0}$ do not prevent loops $I_{i}$ and $I_{j}$ from being fully permutable and can be ignored. If

$$
\forall \vec{d} \in D^{i}:\left(d_{j}^{\min } \neq-\infty \wedge\left(d_{j}^{\min }<0 \rightarrow d_{i}^{\min }>0\right)\right)
$$

then we can skew loop $I_{j}$ by a factor of $f$ with respect to loop $I_{i}$ where

$$
f \geq \max _{\left\{\vec{d} \mid \vec{d} \in D^{\bullet} \wedge d_{i}^{\min } \neq 0\right\}}\left\lceil-d_{j}^{\min } / d_{i}^{\min }\right\rceil
$$

to make loop $I_{j}$ fully permutable with loop $I_{i}$. If instead the condition

$$
\forall \vec{d} \in D^{i}:\left(d_{j}^{\max } \neq \infty \wedge\left(d_{j}^{\max }>0 \rightarrow d_{i}^{\min }>0\right)\right)
$$

holds, then we can reverse loop $I_{j}$ and proceed as above.

Corollary 5.2 If loop $I_{k}$ has dependences such that $\exists \vec{d} \in$ $D: d_{k}^{\min }=-\infty$ and $\exists \vec{d} \in D: d_{k}^{\max }=\infty$ then the outermost fully permutable nest consists only of a combination of loops not including $I_{k}$.

The SRP algorithm takes as input the loops $N$ that have not been placed outside this nest, and the set of dependences $D$ that have not been satisfied by loops outside this loop nest. It first removes from $N$ those serializing loops as defined by the $I_{k} \mathrm{~s}$ of Corollary 5.2. It then uses an iterative step to build up the fully permutable loop nest $F$. In each step, it tries to find a loop from the remaining loops in $N$ that can be made fully permutable with $F$ via possibly multiple applications of Theorem 5.1. If it succeeds in finding such a loop, it permutes the loop to be next outermost in the fully permutable nest, adding the loop to $F$ and removing it from $N$. Then it repeats, searching through the remaining loops in $N$ for another loop to place in $F$. This algorithm is known as SRP because the unimodular transformation it performs can be expressed as the product of a skew transformation (S), a reversal transformation (R) and a permutation transformation ( P ).

We use SRP in both steps of finding a transformation that makes a loop nest $I$ the innermost tile. We first apply SRP iteratively to those loops not in $I$. Each step through the SRP algorithm attempts to find the next outermost fully permutable loop nest, returning all remaining loops that cannot be made fully permutable and returning the unsatisfied dependences. We repeatedly call SRP on the remaining loops until (1) SRP fails to find a single loop to place outermost, in which case the algorithm fails to find a legal ordering for this target innermost tile $I$, or (2) there are no remaining loops, in which case the algorithm succeeds. If this step succeeds, we then call SRP with loops $I$, and all remaining dependences to be satisfied. In this case, we succeed only if SRP makes all the loops fully permutable.

Let us illustrate SRP algorithm using the example in Figure 2. Suppose we are trying to tile loops $I_{1}$ and $I_{2}$. First an outer loop must be chosen. $I_{1}$ can be the outer loop, because its dependence components are all nonnegative. Now loop $I_{2}$ has a dependence component that is negative, but it can be made non-negative by skewing
with respect to $I_{1}$ (Figure 2D). Loop $I_{2}$ is now placed in the same fully permutable nest as $I_{1}$; the loop nest is tilable (Figure 2(e)).

In general, SRP can be applied to loop nests of arbitrarily depth where the dependences can include distances and directions. In the important special case where all the dependences in the loop nest to be ordered are lexicographically positive distance vectors, the algorithm can place all the loops into a single fully permutable loop nest. The algorithm is $O\left(n^{2} d\right)$, where $n$ is the loop nest depth and $d$ is the number of dependence vectors. With the extension of a simple 2D time-cone solver[16], it becomes $O\left(n^{3} d\right)$ but can find a transformation that makes any two loops fully permutable, and therefore tilable, if some such transformation exists.

If there is little reuse, or if data dependences constrain the legal ordering possibilities, the algorithm is fast since $\mathcal{I}$ is small. The algorithm is only slow when there are many carrying reuse loops and few dependences. This algorithm can be further improved by ordering the search through the power set of $\mathcal{I}$ using a branch and bound approach.

## 6 Tiling Experiments

We have implemented the algorithm described in this paper in our SUIF compiler. The compiler currently generates working tiled code for one and multiple processors of the SGI 4D/380. However, the scalar optimizer in our compiler has not yet been completed, and the numbers obtained with our generated code would not reflect the true effect of tiling when the code is optimized. Fortunately, our SUIF compiler also includes a $\mathbf{C}$ backend; we can use the compiler to generate restructured $C$ code, which can then be compiled by a commercial scalar optimizing compiler.

The numbers reported below are generated as follows. We used the compiler to generate tiled $C$ code for a single processor. We then performed, by hand, optimizations such as register allocation of array elements[5], moving loop-invariant address calculation code out of the innermost loop, and unrolling the innermost loop. We then compiled the code using the SGI's optimizing compiler. To run the code on multiple processors, we adopt the model of executing the tiles in a DO-ACROSS manner[16]. This code has the same structure as the sequential code. We needed to add only a few lines to the sequential code to create multiple threads, and initialize, check and increment a small number of counters within the outer loop.

### 6.1 LU Decomposition

The original code for LU-decomposition is:

$$
\begin{aligned}
& \text { for } I_{1}:=1 \text { to } n \text { do } \\
& \text { for } I_{2}:=I_{1}+1 \text { to } n \text { do }
\end{aligned}
$$



Figure 5: Performance of $500 \times 500$ double precision LU factorization without pivoting on the SGI 4D/380. No register tiling was performed.

$$
\begin{aligned}
& \mathrm{A}\left[I_{2}, I_{1}\right] /=\mathrm{A}\left[I_{1}, I_{1}\right] ; \\
& \text { for } I_{3}:=I_{1}+1 \text { to } n \text { do } \\
& \mathrm{A}\left[I_{2}, I_{3}\right]-=\mathrm{A}\left[I_{2}, I_{1}\right] * \mathrm{~A}\left[I_{1}, I_{3}\right] ;
\end{aligned}
$$

For comparison, we measured the performance of the original untiled LU code on one, four and eight processors. For this code, we allocated the element $\mathrm{A}\left[I_{2}, I_{1}\right]$ to a register in the duration of the innermost loop. We parallelized the middle loop, and ran it self-scheduled, with 15 iterations per task. (The results for larger and smaller granularities are virtually identical.) Figure 5 shows that we only get a speedup of approximately 2 on eight processors, even though there is plenty of available parallelism. This is because the memory bandwidth of the machine is not sufficient to support eight processors that are each reusing so little of the data in their respective caches.

This LU loop nest has a reuse vector space that spans all three loops, so our compiler tiles all the loops. The LU code is not perfectly nested, since the division is outside of the innermost loop. The generated code is thus more complicated:

$$
\begin{aligned}
& \text { for } I I_{2}=1 \text { to } n \text { by } s \text { do } \\
& \text { for } I I_{3}=1 \text { to } n \text { by } s \text { do } \\
& \text { for } I_{1}=1 \text { to } n \text { do } \\
& \text { for } I_{2}=\max \left(I_{1}+1, I I_{2}\right) \text { to } \\
& \min \left(n, I I_{2}+s-1\right) \text { do } \\
& \text { if } I I_{3} \leq I_{1}+1 \leq I I_{3}+s-1 \text { then } \\
& \mathrm{A}\left[I_{2}, I_{1}\right] /=\mathrm{A}\left[I_{1}, I_{1}\right] ; \\
& \text { for } I_{3}=\max \left(I_{1}+1, I I_{3}\right) \text { to } \\
& \min \left(n, I I_{3}+s-1\right) \text { do } \\
& \mathrm{A}\left[I_{2}, I_{3}\right]-=\mathrm{A}\left[I_{2}, I_{1}\right] * \mathrm{~A}\left[I_{1}, I_{3}\right] ;
\end{aligned}
$$

As with matrix multiplication, while tiling for cache reuse is important for one processor, it is crucial for multiple processors. Tiling improves the uniprocessor execution by about $20 \%$; more significantly, the speedup is much closer to linear in the number of processors when tiling is used. We ran experiments with tile sizes $32 \times 32$ and $64 \times 64$. The results coincide with our analysis that the marginal performance gain due to increases in the tile size is low. In this example, a $32 \times 32$ tile, which holds $1 / 4$ the data of a $64 \times 64$ tile, performed virtually as well as the larger on a single processor. It actually outperformed the larger tile in the eight processor experiment, mostly due to load balancing.

### 6.2 SOR

We also ran experiments with the SOR code in Figure 7(a), which is a two-dimensional equivalent of the example in Figure 2(a). Figure 6 shows the performance results for three versions of this nest, where $t$ is 30 and $n+1$, the size of the matrix, is 500 . None of the original loops is parallelizable. The first verson, labeled "DOALL", is transformed via wavefronting so that the middle loop is a DOALL loop [16] (Figure 7(b)). This transformation unfortunately destroys the original locality within the code. Thus, performance is abysmal on a single processor, and speedup for multiple processors is equally abysmal even though again there is plenty of available parallelism in the $I_{2}^{\prime}$ loop.


Figure 6: Behavior of 30 iterations of a $500 \times 500$ double precision SOR step on the SGI 4D/380. The tile sizes are $64 \times 64$ iterations. No register tiling was performed.

The code for the second version, labeled " 2 D tile", is shown in Figure 7(c). This version tiles only the inner-

```
(a): SOR nest
for \(I_{1}:=1\) to \(t\) do
    for \(I_{2}:=1\) to \(n-1\) do
        for \(I_{3}:=1\) to \(n-1\) do
        \(\mathrm{A}\left[I_{2}, I_{3}\right]:=0.2^{*}\left(\mathrm{~A}\left[I_{2}, I_{3}\right]+\mathrm{A}\left[I_{2}+1, I_{3}\right]+\mathrm{A}\left[I_{2}-1, I_{3}\right]+\mathrm{A}\left[I_{2}, I_{3}+1\right]+\mathrm{A}\left[I_{2}, I_{3}-1\right]\right) ;\)
```

(b): "DOALL" SOR nest

```
for \(I_{1}^{\prime}:=5\) to \(2 N-2+3 t\) do
    doall \(I_{2}^{\prime}:=\max \left(I_{1}\left(I_{1}^{\prime}-t-2 N+2\right) / 2\right)\) to \(\min \left(t,\left(I_{1}^{\prime}-3\right) / 2\right)\) do
        for \(I_{3}^{\prime}:=\max \left(1+I_{2}^{\prime}, I_{1}^{\prime}-2 I_{2}^{\prime}-n+1\right)\) to \(\min \left(n-1+I_{2}^{\prime}, I_{1}^{\prime}-1-2 I_{2}^{\prime}\right)\) do
        \(\mathrm{A}\left[I_{3}^{\prime}-I_{2}^{\prime}, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}\right]:=0.2^{*}\left(\mathrm{~A}\left[I_{3}^{\prime}-I_{2}^{\prime}, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}\right]+\mathrm{A}\left[I_{3}^{\prime}-I_{2}^{\prime}+1, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}\right]\right.\)
            \(\left.+\mathrm{A}\left[I_{3}^{\prime}-I_{2}^{\prime}-1, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}\right]+\mathrm{A}\left[I_{3}^{\prime}-I_{2}^{\prime}, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}+1\right]+\mathrm{A}\left[I_{3}^{\prime}-I_{2}^{\prime}, I_{1}^{\prime}-2 I_{2}^{\prime}-I_{3}^{\prime}-1\right]\right) ;\)
```

(c): "2-D Tile" SOR nest

```
for }\mp@subsup{I}{1}{\prime}:=1\mathrm{ to t do
    for }I\mp@subsup{I}{3}{\prime}:=1\mathrm{ to }n-1\mathrm{ by }s\mathrm{ do
        for }\mp@subsup{I}{2}{\prime}:=1\mathrm{ to }n-1\mathrm{ do
            for I}\mp@subsup{I}{3}{\prime}:=I\mp@subsup{I}{3}{\prime}\mathrm{ to min}(n-1,I\mp@subsup{I}{3}{}+s-1) d
```


(d): "3-D Tile" SOR nest

```
for \(I I_{2}^{\prime}:=2\) to \(n-1+t\) by \(s\) do
    for \(I I_{3}^{\prime}:=2\) to \(n-1+t\) by \(s\) do
        for \(I_{1}^{\prime}:=1\) to \(t\) do
            for \(I_{2}^{\prime}:=I I_{2}^{\prime}\) to \(\min \left(n-1+t, I I_{2}^{\prime}+s-1\right)\) do
            for \(I_{3}^{\prime}:=I I_{3}^{\prime}\) to \(\min \left(n-1+t, I I_{3}^{\prime}+s-1\right)\) do
                \(\mathrm{A}\left[I_{2}^{\prime}, I_{3}^{\prime}\right]:=0.2^{*}\left(\mathrm{~A}\left[Z_{2}^{\prime}-I_{1}^{\prime}\right]\left[I_{3}^{\prime}-I_{1}^{\prime}\right]+\mathrm{A}\left[I_{2}^{\prime}-I_{1}^{\prime}+1, I_{3}^{\prime}-I_{1}^{\prime}\right]\right.\)
                    \(\left.+\mathrm{A}\left[I_{2}^{\prime}-I_{1}^{\prime}-1, I_{3}^{\prime}-I_{1}^{\prime}\right]+\mathrm{A}\left[I_{2}^{\prime}-I_{1}^{\prime}, I_{3}^{\prime}-I_{1}^{\prime}+1\right]+\mathrm{A}\left[I_{2}^{\prime}-I_{1}^{\prime}, I_{3}^{\prime}-I_{1}^{\prime}-1\right]\right) ;\)
```

Figure 7: Code for the different SOR versions.
most two loops. The uniprocessor performance is significantly better than even the eight processor wavefronted performance. Thus although wavefronting may increase parallelism, it may reduce locality so much that it is better to use only one processor. However, the speedup of the 2D tile method is still limited because self-temporal reuse is not being exploited.

To exploit all the available reuse, all three loops must be included in the innermost tile. The SRP algorithm first skews $I_{2}$ and $I_{3}$ with respect to $I_{1}$ to make tiling legal, and then tiles to produce the code in Figure 7(d). This "3D tile" version of the loop nest is best for a single processor, and also has the best speedup in the multiprocessor version.

## 7 Conclusions

In this paper, we propose a complete approach to the problem of improving the cache performance for loop nests. The approach is to first transform the code via interchange, reversal, skewing and tiling, then determine the tile size, taking into account data conflicts due to the set associativity of the cache [12]. The loop transformation algorithm is based on two concepts: a mathematical formulation of
reuse and locality, and a matrix-based loop transformation theory.

While previous work on evaluating locality estimates the number of memory accesses directly for a given transformed code, we break the evaluation into three parts. We use a reuse vector space to capture the inherent reuse within a loop nest; we use a localized vector space to capture a compound transform's potential to exploit locality; finally we evaluate the locality of a transformed code by intersecting the reuse vector space with the localized vector space.

There are four reuse vector spaces: self-temporal, selfspatial, group-temporal, and group-spatial. These reuse vector spaces need to be calculated only once for a given loop nest. We show that while unimodular transformations can alter the orientation of the localized vector space, tiling can increase the dimensionality of this space.

The reuse and localized vector spaces can be used to prune the search for the best compound transformation, and not just for evaluating the locality of a given code. First, all transforms with identical localized vector space are equivalent with respect to locality. In addition, transforms with different localized vector spaces may also be
equivalent if the intersection between the localized and the reuse vector spaces is identical. A loop transformation algorithm need only to compare between transformations that give different localized and reuse vector space intersections.

Unlike the stepwise transformation approach used in existing compilers, our loop transformer solves for the compound transformation directly. This is made possible by our theory that unifies loop interchanges, skews and reversals as unimodular matrix transformations on dependence vectors with either direction or distance components. The algorithm extracts the dependence vectors, determines the best compound transform using locality objectives to prune the search, then transforms the loops and their loop bounds once and for all. This theory makes the implementation of the algorithm simple and straightforward.

## References

[1] W. Abu-Sufah. Improving the Performance of Virtual Memory Computers. PhD thesis, University of Illinois at Urbana-Champaign, Nov 1978.
[2] U. Banerjee. Data dependence in ordinary programs. Technical Report 76-837, University of Illinios at Urbana-Champaign, Nov 1976.
[3] U. Banerjee. Dependence Analysis for Supercomputing. Kluwer Academic, 1988.
[4] U. Banerjee. Unimodular transformations of double loops. In 3rd Workshop on Languages and Compilers for Parallel Computing, Aug 1990.
[5] D. Callahan, S. Carr, and K. Kennedy. Improving register allocation for subscripted variables. In Proceedings of the ACM SIGPLAN '90 Conference on Programming Language Design and Implementation, June 1990.
[6] J. Dongarra, J. Du Croz, S. Hammarling, and I. Duff. A set of level 3 basic linear algebra subprograms. ACM Transactions on Mathematical Software, pages 1-17, March 1990.
[7] K. Gallivan, W. Jalby, U. Meier, and A. Sameh. The impact of hierarchical memory systems on linear algebra algorithm design. Technical report, University of Illinios, 1987.
[8] D. Gannon, W. Jalby, and K. Gallivan. Strategies for cache and local memory management by global program transformation. Journal of Parallel and Distributed Computing, 5:587-616, 1988.
[9] G. H. Golub and C. F. Van Loan. Matrix Computations. Johns Hopkins University Press, 1989.
[10] F. Irigoin and R. Triolet. Computing dependence direction vectors and dependence cones. Technical Report E94, Centre D'Automatique et Informatique, 1988.
[11] F. Irigoin and R. Triolet. Supernode partitioning. In Proc. 15th Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, January 1988.
[12] M. S. Lam, E. E. Rothberg, and M. E. Wolf. The cache performance and optimizations of blocked algorithms. In Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems, April 1991.
[13] A. C. McKeller and E. G. Coffman. The organization of matrices and matrix operations in a paged multiprogramming environment. CACM, 12(3):153-165, 1969.
[14] A. Porterfield. Software Methods for Improvement of Cache Performance on Supercomputer Applications. PhD thesis, Rice University, May 1989.
[15] R. Schreiber and J. Dongarra. Automatic blocking of nested loops. 1990.
[16] M. E. Wolf and M. S. Lam. A loop transformation theory and an algorithm to maximize parallelism. IEEE Transactions on Parallel and Distributed Systems, July 1991.
[17] M. J. Wolfe. Techniques for improving the inherent parallelism in programs. Technical Report UIUCDCS-R-78-929, University of Illinois, 1978.
[18] M. J. Wolfe. More iteration space tiling. In Supercomputing '89, Nov 1989.


[^0]:    This research was supported in part by DARPA contract N00014-87-K. 0828.

    Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
    © 1991 ACM 0-89791-428-7/91/0005/0030...\$1.50
    Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation. Toronto, Ontario, Canada, June 26-28, 1991

