

Efficient synthesis of out-of-core algorithms using a nonlinear optimization solver[☆]

Sandhya Krishnan^a, Sriram Krishnamoorthy^{a,*}, Gerald Baumgartner^b, Chi-Chung Lam^a,
J. Ramanujam^c, P. Sadayappan^a, Venkatesh Choppella^{a, d}

^aDepartment of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA

^bDepartment of Computer Science, Louisiana State University, Baton Rouge, LA 70803, USA

^cDepartment of Electrical and Computer Engineering, Louisiana State University, Baton Rouge, LA 70803, USA

^dIndian Institute of Information Technology and Management—Kerala, Technopark, Thiruvananthapuram, Kerala 695 581, India

Received 8 August 2004; received in revised form 1 February 2005; accepted 6 June 2005

Abstract

We address the problem of efficient out-of-core code generation for a special class of imperfectly nested loops encoding tensor contractions arising in quantum chemistry computations. These loops operate on arrays too large to fit in physical memory. The problem involves determining optimal tiling of loops and placement of disk I/O statements. This entails a search in an explosively large parameter space. We formulate the problem as a nonlinear optimization problem and use a discrete constraint solver to generate optimized out-of-core code. The solution generated using the discrete constraint solver consistently outperforms other approaches by up to a factor of four. Measurements on sequential and parallel versions of the generated code demonstrate the effectiveness of the approach.

© 2005 Published by Elsevier Inc.

Keywords: Data locality optimization; Out-of-core algorithms; Program transformation; Compiler optimization; Discrete constrained search; Tensor contractions

1. Introduction

Many scientific and engineering applications need to operate on data sets that are too large to fit in the physical memory of the machine. We focus on the domain of electronic structure calculations in quantum chemistry [17,36,40]. These calculations require contractions (generalized matrix multiplications) of multi-dimensional tensors that are often larger than available physical memory. In such situations, it is necessary to develop *out-of-core* algorithms that explicitly

orchestrate the movement of blocks of data between main memory and secondary disk storage. We are developing a program synthesis tool called the Tensor Contraction Engine (TCE) [4,5] to facilitate the development of parallel programs for this domain, by automatically transforming high-level tensor contraction expressions into efficient parallel programs.

In this paper, we address the following problem that arises in the context of the TCE system. We are given an imperfectly nested loop structure containing a collection of tensor contraction computations expressed in an “abstract” form, that is, without concern for whether the arrays can fit within available physical memory. The problem consists of generating a “concrete” form of the code by suitably tiling the loops and inserting the necessary disk I/O statements so as to minimize the total cost of disk I/O. In the case of code generation for a parallel system, the problem also involves distributing the workload among processors and inserting

[☆] Supported in part by the National Science Foundation through Grants CHE-0121676, CHE-0121706, CCF-0073800 and EIA-9986052.

* Corresponding author.

E-mail addresses: krishnas@cse.ohio-state.edu (S. Krishnan), krishnsr@cse.ohio-state.edu (S. Krishnamoorthy), gb@csc.lsu.edu (G. Baumgartner), clam@cse.ohio-state.edu (C.-C. Lam), jxr@ece.lsu.edu (J. Ramanujam), saday@cse.ohio-state.edu (P. Sadayappan), choppell@iiitmk.ac.in (V. Choppella).

the required communication. The search space of possible placements of the disk I/O statements and possible combinations of tile sizes is explosively large. We formulate the problem as a non-linear optimization problem and use a general-purpose discrete constraint solver to generate optimized out-of-core code.

The paper is organized as follows: in Section 2, we explain the computational context for which the data locality optimization approach is developed. In Section 3, we review related work in the area. Section 4 describes the Discrete Constrained Search (DCS) solver [9,51–53] and outlines the steps used to convert the abstract code specification into concrete code. Section 5 illustrates the code-generation process using a representative example. Our experimental results in Section 6 demonstrate that the DCS-based approach to out-of-core code generation is efficient and effective.

2. The computational context

The optimization presented in this paper has been developed in the context of the TCE [5,13], a domain-specific compiler for ab initio quantum chemistry calculations. The TCE takes as input a high-level specification of a computation expressed as a set of tensor contraction expressions and transforms it into efficient parallel code. Several compile-time optimizations are incorporated into the TCE: algebraic transformations to minimize operation counts [33,34], loop fusion to reduce memory requirements [30–32], space–time trade-off optimization [11], communication minimization [12], and data locality optimization [13,14] of memory-to-cache traffic.

A tensor contraction expression comprises a collection of multi-dimensional summations of the product of several input arrays. As an example, consider the following contraction, used often in quantum chemistry calculations to transform a set of two-electron integrals from an atomic orbital (AO) basis to a molecular orbital (MO) basis:

$$B(a, b, c, d) = \sum_{p,q,r,s} C1(s, d) \times C2(r, c) \times C3(q, b) \times C4(p, a) \times A(p, q, r, s).$$

This contraction is referred to as a four-index transform. Here, $A(p, q, r, s)$ is a four-dimensional input array initially stored on disk, and $B(a, b, c, d)$ is the transformed output array to be placed on disk at the end of the computation. The arrays $C1$ – $C4$ are called the transformation matrices. In practice, these four arrays are identical; we identify them by different names in order to be able to distinguish them in the text.

The indices p, q, r , and s have the same range N , denoting the total number of orbitals, which is equal to $O + V$. O denotes the number of occupied orbitals and V denotes the

number of unoccupied (virtual) orbitals. Likewise, the index ranges for a, b, c , and d are the same, and equal to V . Typical values for O range from 10 to 300; the number of virtual orbitals V is usually between 50 and 1000.

The calculation of B is done in the following four steps to reduce the number of floating point operations from $O(V^4N^4)$ in the initial formula (8 nested loops, for p, q, r, s, a, b, c , and d) to $O(VN^4)$

$$B(a, b, c, d) = \sum_s C1(s, d) \times \left(\sum_r C2(r, c) \times \left(\sum_q C3(q, b) \times \left(\sum_p C4(p, a) \times A(p, q, r, s) \right) \right) \right).$$

This operation-minimization transformation results in the creation of three intermediate arrays

$$\begin{aligned} T1(a, q, r, s) &= \sum_p C4(p, a) \times A(p, q, r, s), \\ T2(a, b, r, s) &= \sum_q C3(q, b) \times T1(a, q, r, s), \\ T3(a, b, c, s) &= \sum_r C2(r, c) \times T2(a, b, r, s). \end{aligned}$$

Assuming that the available memory less than V^4 (which for $V = 800$ and double precision arrays is about $3TB$), none of $A, T1, T2, T3$, and B can fit entirely in memory. Therefore, the intermediates $T1, T2$, and $T3$ need to be written to disk on production, and read from disk before consumption in the next step. Since none of these arrays can be fully stored in memory, it may not be possible to perform all multiplication operations by reading each element of the input arrays from the disk only once. This could result in the amount of disk I/O volume being much larger than the total volume of the data on disk.

For illustration purposes, we focus on the following contraction (a two-index transform):

$$B(m, n) = \sum_{i,j} C1(m, i) \times C2(n, j) \times A(i, j).$$

The operation minimal form of the two-index transform and the corresponding intermediate array are as follows:

$$B(m, n) = \sum_i C1(m, i) \times \left(\sum_j C2(n, j) \times A(i, j) \right),$$

$$T(n, i) = \sum_j C2(n, j) \times A(i, j).$$

```

double T(V,N)
T(*,*) = 0
B(*,*) = 0
FOR i = 1, N
  FOR n = 1, V
    FOR j = 1, N
      T(n,i) += C2(n,j) * A(i,j)
    END FOR j,n,i
  END FOR i
FOR i = 1, N
  FOR n = 1, V
    FOR m = 1, V
      B(m,n) += C1(m,i) * T(n,i)
    END FOR m,n,i
  END FOR i
(a)

```

```

double T(V,N)
T(*,*) = 0
B(*,*) = 0
FOR i, n, j
  T(n,i) += C2(n,j) * A(i,j)
END FOR j,n,i
FOR i, n, m
  B(m,n) += C1(m,i) * T(n,i)
END FOR m,n,i
(b)

```

```

double T
B(*,*) = 0
FOR i, n
  T = 0
  FOR j
    T += C2(n,j) * A(i,j)
  END FOR j
FOR m
  B(m,n) += C1(m,i) * T
END FOR m
END FOR n,i
(c)

```

Fig. 1. Example of the use of loop fusion to reduce memory requirements. Loops i and n are fused to reduce T to a scalar: (a) unfused code, (b) compact notation and (c) fused code.

Fig. 1 shows the computation of the array B and illustrates how memory requirements for the computation of B may be reduced using loop fusion. Such an “abstract form” of the computation cannot be directly compiled and executed because it does not take into account the available physical memory (if the arrays fit within the virtual memory, execution of the compiled code is possible, but will exhibit extremely poor performance due to excessive overhead of paging to move virtual memory pages between disk and physical memory). The transformation of abstract forms into concrete forms that can be efficiently executed is addressed in Section 4. Concrete forms have explicit disk I/O statements to move data between disk-resident arrays and their in-memory counterparts. Fig. 1(a) shows the abstract form of the computation before loop fusion. The computation consists of two loop nests: a first loop that produces the intermediate $T(1 : V, 1 : N)$, and a second loop that uses T to produce the result $B(1 : V, 1 : V)$.

In Fig. 1(b), each loop nest is abbreviated as a single “For” loop with a sequence of indices. Fig. 1(c) illustrates the result of loop fusion. Note that all loops in each of the two loop nests in Fig. 1(a) are fully permutable and there are no fusion-preventing dependences between the loops. Hence, the common loops i and n , shown underlined, can be fused. After loop fusion, the storage requirements for T can be reduced because there is no longer a need for an explicit dimension of T corresponding to any loop indices that are fused between the producer of T and the consumer of T —storage elements can be reused over sequential iterations of fused loops. In this example, T can be contracted to a scalar as shown in Fig. 1(c). Although the total number of arithmetic operations remains unchanged, the significant reduction in size of the intermediate array T implies that it may be completely stored in memory, without the need for any disk I/O for it. In contrast, if $V \times N$ is larger than the available memory, the unfused version would result in T being written out to disk after it is produced in the first loop, and then read in from disk for the second loop.

Given a fused *abstract* form of the computation, in the form of an imperfectly nested loop structure (i.e., a nested loop structure in which at least one loop other than the inner-most contains more than one statement, as in Fig. 1(c)) the out-of-core code-generation process requires consideration of a number of issues. Each loop in the imperfectly nested loop structure is split into a tiling and an intra-tile loop. Given a tiled loop structure, there are a number of different candidate positions for placing disk I/O statements. Thus, a search space consisting of two dimensions, placement of disk read/write statements and the tile sizes, needs to be explored. A disk I/O statement transfers blocks of data between the disk resident array and its in-memory counterpart. The size of the in-memory buffer is a function of the tile sizes and placement of the corresponding disk I/O statement. The task of the out-of-core code-generation algorithm is to tile the loops, determine the optimal placements of disk I/O statements, and tile sizes that minimize the disk I/O cost while satisfying the memory limit constraints.

3. Related work

We have addressed various issues arising in the synthesis context described above focusing primarily on minimizing memory-to-cache data movement [13,14]. In Cociorva et al. [14], we developed an integrated approach to fusion and tiling transformations for a restricted class of loops arising in the context of our program synthesis system; these assumed restrictions were subsequently removed [13]. We developed a tile size search procedure to estimate the total capacity miss cost for each of a large number of combinations of tile sizes for the various loops of an imperfectly nested loop set. After finding the best combination of tile sizes, we made adjustments to address spatial locality considerations—by adjusting the tile sizes for any loop indexing the fastest varying dimension of any array to be larger than the cache line size.

Krishnan [28] extended this approach to the disk-memory hierarchy using a greedy approach to disk read/write placement. For each set of tile sizes, Krishnan's algorithm places read/write statements immediately inside those loops at which the memory limit is just exceeded. In Krishnan et al. [29], we describe an algorithm to determine effective tile sizes. This algorithm explores the tile size search space using the set of candidate fusion structures with disk I/O placements as input. The search space was divided into feasible and infeasible solution spaces and their boundary was shown to contain the optimal solution. We developed an algorithm to locate the boundary efficiently and used steepest ascent hill-climbing to determine an efficient solution for the tile sizes.

There has been some work in the area of software techniques for optimizing disk I/O. These include parallel file systems, compile time [6–8,21–23,43,46] and runtime libraries and optimizations [10,49]. Several compiler techniques for optimizing out-of-core programs in high-performance Fortran are discussed in [6,7]. Bordawekar et al. [8] develop a scheduling strategy to eliminate additional I/O arising from communication among processors. Solutions to choreographing disk I/O with computation are presented by Paleczny et al. [46]; they organize computations into groups that operate efficiently on data accessed in chunks. Compiler-directed pre-fetching is discussed by Mowry et al. [43], which is orthogonal to compiler transformations discussed in this paper. ViC* (Virtual C*) [16] is a preprocessor that transforms out-of-core C* programs into in-core programs with appropriate calls to the ViC* I/O library. Kandemir et al. [21–23] developed file layout and loop transformations for reducing I/O. None of these techniques address model-driven automatic tile size selection for optimizing I/O and all of them deal only with perfectly nested loops.

Considerable research on loop transformations for locality in nested loops has been reported in the literature [15,37,38,41,54]. Nevertheless, a performance-model driven approach to the integrated use of loop fusion and loop tiling for enhancing locality in imperfectly nested loops has not been addressed in these works. Loop tiling for enhancing data locality has been studied extensively [3,15,24–26,47,48,54,55], and analytical models of the impact of tiling on locality in perfectly nested loops have been developed [20,35,42]. Mitchell et al. [42] provide analytical models for multi-level tiling of matrix–matrix multiplication. Ahmed et al. [1,2] have developed a framework that embeds an arbitrary collection of loops into an equivalent perfectly nested loop that can be tiled; this allows a cleaner treatment of imperfectly nested loops. Lim et al. [39] developed a framework based on affine partitioning and blocking to reduce synchronization and improve data locality. Specific issues of locality enhancement, I/O optimization and automatic tile size selection have not been addressed in the works that can handle imperfectly nested loops [1,2,39,48].

4. Proposed approach

We use the DCS solver to compute the best placement of disk I/O statements that would minimize the disk access cost while satisfying the memory limit constraints. DCS [9,51–53] is a software package for determining the constrained global minima (CGM) in the discrete variable space of a single-objective, discrete, constrained non-linear programming problem (NLP). A web interface to the DCS solver is available [50]. It uses AMPL, *A Modeling Language for Mathematical Programming* [19], as the problem input format. Due to the limitations in AMPL in modeling arbitrary discrete variables, their current implementation can only solve problems with continuous variables by discretizing them.

The out-of-core code-generation process translates the abstract code into concrete code by loop tiling and placement of disk I/O statements. We fully explore the search space of disk I/O placements and tile sizes by formulating the search as a non-linear constrained minimization problem, where the objective function is the disk I/O cost. The solution to be determined is constrained by the memory limit and minimum I/O block size for efficient disk I/O. We input the formulated non-linear problem to the DCS system, which determines the optimal combination of placement of disk I/O statements and tile sizes.

We continue with the two-index transform example introduced in Section 2 for transforming atomic orbitals into molecular orbitals. Fig. 2(a) shows an abstract code for the two-index transform. We assume that the arrays involved are too large to fit into the physical memory of the machine. The arrays involved in the loop structure fall into the following three categories: input arrays that initially reside on disk (A , $C1$ and $C2$), intermediate arrays produced and consumed within the computation and not required on completion (T), and output arrays that must finally be written to disk (B). Fig. 2(b) shows the parse tree corresponding to the abstract code in Fig. 2(a). To simplify the tree representation, each sequence of perfectly nested loops is represented by a single-node labeled with the corresponding sequence of loop indices.

The input to the out-of-core code-generation algorithm consists of the abstract code, the loop ranges and the memory limit of the machine. The algorithm consists of the following three steps:

- (1) *Loop tiling*: We split each loop into a tiling loop and an intra tile loop and propagate the intra tile loops down to the leaves. For example, as shown in Fig. 3, loop i is split into tiling loop iT and intra-tile loop iI . Fig. 3(b) shows the parse tree for the tiled abstract code in Fig. 3(a).
- (2) *Candidate placements*: For each array, we enumerate all feasible placements of disk read/write statements. Any placement surrounded by a loop index that is not involved in the I/O statement is ignored, as this I/O




Fig. 2. Example of abstract code and corresponding parse tree for 2-index transform: (a) abstract code for the 2-index transform and (b) parse tree for the 2-index transform.

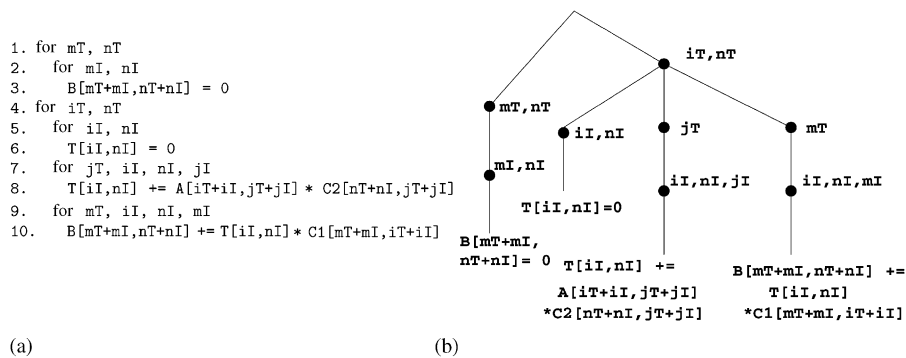


Fig. 3. Example of abstract code and corresponding parse tree for the tiled version of the 2-index transform: (a) abstract code for the tiled 2-index transform and (b) parse tree for the tiled 2-index transform.

statement can be moved out of the loop reducing the I/O cost.

- (3) *DCS Input construction*: Given the enumeration from step 2, we construct non-linear equations for the objective function and constraints and provide them as input to the DCS solver. The DCS solver outputs the disk read/write placement for each array and the tile sizes that minimize the disk I/O cost and satisfy the memory limit constraint.

4.1. Candidate placements

Given a tiled imperfectly nested loop structure (Fig. 3), we consider various possible placements of reads for input arrays, reads and writes for intermediates, and writes (and reads, if required) for output arrays. In enumerating the candidate placements, there are some constraints that must be satisfied.

- (1) *Input array constraints*: The read statement for an input array may only be placed to be executed before the statement where it is consumed. For example, in Fig. 3(a), the read for input array A can be placed anywhere before line 7.

- (2) *Output array constraints*: The write for an output array may only be placed after the statement where it is produced. For example, in Fig. 3(a), the write for output array B can be placed anywhere after line 9.

- (3) *Intermediate array constraints*: For intermediate arrays, we have two cases to consider: the array is either kept in memory or written to disk. If the array is kept in memory, there will be no disk I/O statements inserted for the array. On the other hand, if it is written to disk, there is a constraint imposed on its disk read/write placement. For example, in Fig. 3(a), intermediate T is produced in statement 7 and consumed in statement 9. If we consider these statements in the parse tree in Fig. 3(b), the lowest common ancestor for both the statements is loop nT . The write statement for the production and read statement for the consumption must be inside this nT loop.

The approach to enumerating the placements for input, output and intermediate arrays is sketched below; details may be found in [28].

- (1) *Input arrays*: Each loop index surrounding the consumption of an input array is considered as a candidate position for placing the read. At any candidate position, if


```

Input Arrays: (Read Placements)
A: iI, nT
C2: iI, jT
C1: iI, nT

Output Arrays: (Write Placements)
B:
  Write Placement: iI, mT
  Read Required : Yes, Yes

Intermediates: (Write & Read Placements)
T: In Memory

```

(a)

```

for mT, nT
  for mI, nI
    B[mI,nI] = 0
  Write BDisk[mT:mT+Tm-1,nT:nT+Tn-1]
for iT, nT
  for iI, nI
    T[iI,nI] = 0
  for jT
    C2[1:Tn,1:Tj] =
      Read C2Disk[nT:nT+Tn-1,jT:jT+Tn-1]
    A[1:Ti,1:Tj] =
      Read ADisk[iT:iT+Ti-1,jT:jT+Tj-1]
    for iI, nI, jI
      T[iI,nI] += C2[nI,jI] * A[iI,jI]
    for mT
      B[1:Tm,1:Tn] =
        Read BDisk[mT:mT+Tm-1,nT:nT+Tn-1]
      C1[1:Tm,1:Tj] =
        Read C1Disk[mT:mT+Tm-1,jT:jT+Tj-1]
      for iI, nI, mI
        B[mI,nI] += T[iI,nI] * C1[mI,iI]
      Write BDisk[mT:mT+Tm-1,nT:nT+Tn-1]

```

(b)

Fig. 4. Candidate I/O placements and final concrete code. $N_m = N_n = 35000$, $N_i = N_j = 40000$, memory limit = 1 GB, double precision arrays: (a) candidate I/O placements and (b) final concrete code for 2-index transform.

there exists a redundant loop immediately surrounding it, then we ignore that position and move further up. A redundant loop for a read statement is one that does not index the array being read. We also ignore those read placements that cause the in-memory version of the input array to be a scalar or a vector. This is because the resulting concrete code will involve in-memory matrix–matrix products using level-3 BLAS kernels [18], and scalar and vector operands will result in poor performance. Consider the abstract tiled code in Fig. 3. All loops surrounding statement 7 are candidate positions for placing the read for array A . Loops jI and nI are ignored so that the in-memory version of array A is at least two-dimensional. Loop jT is not considered because the surrounding loop nT is redundant for array A . Another important check that needs to be made is that the in-memory version of the array fits in memory. For every candidate position, we compute the memory cost of the corresponding *local buffer* assuming a tile size of one. If the buffer does not fit in memory, we do not move further up.

- (2) *Output arrays*: The algorithm for enumerating write placements for an output array is exactly the same as that for input arrays, except that if any redundant loop surrounds the write statement, we need to insert a corresponding read for the array before the production. This is required as we will be re-accessing the disk array for every iteration of the redundant loop. For example, consider statement 9 in Fig. 3(a), where the output array B is produced. If the write for array B is placed just after loop mT , an extra read will be required as the write will be surrounded by the redundant loop iT .
- (3) *Intermediate arrays*: If an intermediate array is written to disk, the algorithm for enumerating the disk read/write statements is exactly the same as for input/output arrays,

except that the constraint specified earlier for intermediate arrays must be satisfied.

Fig. 4(a) shows the candidate read and write placements computed for each array in the code shown in Fig. 3(a). Fig. 4(b) shows the final concrete code for the two-index transform using the candidate read and write placements shown in Fig. 3(a). Note that in Fig. 4(b), for a loop index x , the index of the tiling loop is denoted xT and the index of the intra-tile loop is denoted xI ; in addition, the tile size for this loop index is denoted Tx .

4.2. DCS input construction

If all possible combinations of disk I/O placements, shown in Fig. 4(a), are considered for all the arrays, a very large number of cases will have to be evaluated. Our approach to avoid explicit evaluation for each combination of I/O placements is to encode the placement into the formulation of a nonlinear optimization problem that is input to the DCS system, as explained below. DCS attempts to minimize an objective function subject to equality and inequality constraints. The input to DCS consists of *input parameters*, *variables*, *objective function*, and a set of *constraints*.

4.2.1. Input parameters

The input parameters for our problem are the memory limit of the machine and the ranges N_i, N_j, \dots of the loop indices i, j, \dots

4.2.2. Variables

The variables in our case include tile sizes T_i, T_j, \dots for loops i, j, \dots where each tile size variable has a lower bound of 1 and an upper bound of the full loop range. In addition to tile size variables, placement variables, $\lambda_i, i = 0, 1, 2, \dots$, are introduced for arrays having more than one candidate

placement. The placement variables corresponding to an array encode all the possible placements for the disk I/O for the array. The values chosen for these variables in the solution from the solver uniquely determine the disk I/O placement for that array.

4.2.3. Objective function

The objective function for our problem is the disk I/O cost. The disk I/O cost for an I/O statement is the product of the size of the array being read/written and the ranges of any redundant loops surrounding the statement. Consider the two possible read placements for input array A shown in Fig. 4(a). For the first read placement above loop iI , the disk I/O cost will be:

$$D1_A = (Nn/Tn) \times Size_A,$$

where the total size of array A is multiplied by the range of the redundant loop nT . The disk I/O cost for the second read placement (above loop nT), is $D2_A = Size_A$. Since there are two possible placements for A , $\lceil \log_2(2) \rceil = 1$ placement variable λ_0 is introduced as follows to express the disk I/O cost:

$$(\lambda_0 \times D1_A) + ((1 - \lambda_0) \times D2_A).$$

If $\lambda_0 = 1$, the first placement is selected, else if $\lambda_0 = 0$, the second one is selected. As explained later, the placement-encoding variables are constrained to have a value of 0 or 1.

4.2.4. Constraints

The total space utilized for in-memory buffers is constrained to be within the memory limit. A static memory cost model is used, in which all the in-memory buffers are allocated memory at compile time. The total memory cost is the sum of the memory usage for all the individual in-memory buffers. The memory cost for an in-memory buffer is the product of the ranges of its indices. The memory cost expression for array A can be constructed, along the same lines as the disk cost expression, as follows. For the read placement above loop iI , the in-memory buffer for input array A will be $A[iI, jI]$, which makes the memory cost $M1_A = Ti \times Tj$. On the other hand, for the read placement above loop nT , the in-memory buffer is $A[iI, j]$, thus making the memory cost $M2_A = Ti \times Nj$. The memory limit constraint using placement variable λ_0 is

$$(\lambda_0 \times M1_A) + ((1 - \lambda_0) \times M2_A) \leq MemoryLimit.$$

The placement variables are constrained to take values 0 or 1 as follows:

$$\lambda_i \times (1 - \lambda_i) = 0, \quad i = 0, 1, 2, \dots$$

We also introduce constraints on the minimum size of the in-memory version of an array. The arrays are stored in a blocked fashion on disk. The block sizes of the arrays are equal to the size of their in-memory versions, determined by the out-of-core code generation algorithm. A block is the

basic unit of I/O and is chosen to be large enough to make the disk seek time negligible compared to the block transfer time. Krishnamoorthy et al. [27] observed that the incremental improvement obtained in the ratio of transfer time to seek time became negligible, and approached the performance of sequential I/O, beyond a certain block size. The in-memory version of the array, and hence the block size, is constrained to be larger than this block size. For the system on which the experiments were conducted, and whose configuration is described in Section 6, the block size for reads must be at least 2 MB, while that for writes must be at least 1 MB.

In this manner, we can construct disk cost, memory cost and other constraint expressions for all arrays. Using these expressions, we build the input to DCS using the AMPL format [19]. DCS minimizes the objective function, that is, the disk I/O cost expression, while satisfying the memory limit, placement variable and buffer size constraints. DCS outputs values for the placement variables and tile sizes, thus providing the parameters for the concrete code.

The code generated for a multi-processor system uses the Global Arrays (GA) and Disk Resident Arrays (DRA) libraries [44,45]. GA provides a shared-memory programming model while encouraging locality of access. DRA extends the shared-memory model to secondary storage. GA/DRA provide an array abstraction in which the portion of data to be accessed is specified as a section of the array. In the generated code, the reads and writes from the disk are performed by the read and write routines in DRA. The in-memory computation is performed using kernel matrix multiplication libraries in GA. The I/O operations and the in-memory computations are collective operations.

5. Illustration

In this section, we illustrate the process of code generation using the 4-index transform. Consider the abstract code

```

1.  for a, q, r, s
2.    T1[a,q,r,s] = 0
3.  for a, p, q, r, s
4.    T1[a,q,r,s] += C4[p,a] * A[p,q,r,s]
5.  for a, b, c, d
6.    B[a,b,c,d] = 0
7.  for a, b
8.    for c, s
9.      T3[c,s] = 0
10.  for r, s
11.    T2 = 0
12.    for q
13.      T2 += C3[q,b] * T1[a,q,r,s]
14.    for c
15.      T3[c,s] += C2[r,c] * T2
16.    for c, d, s
17.      B[a,b,c,d] += C1[s,d] * T3[c,s]
```

Fig. 5. Abstract code for the 4-index transform example.

```

1. for aT, qT, rT, sT, aI, qI, rI, sI
2.   T1[aT+aI,qT+qI,rT+rI,sT+sI] = 0
3. for aT, pT, qI, rT, sT, aI, pI, qI, rI, sI
4.   T1[aT+aI,qT+qI,rT+rI,sT+sI] += C4[pT+pI,aT+aI]*A[pT+pI,qT+qI,rT+rI,sT+sI]
5. for aT, bI, cT, dT, aI, bI, cI, dI
6.   B[aT+aI,bT+bI,cT+cI,dT+dI] = 0
7. for aT, bT
8.   for cT, sT, aI, bI, cI, sI
9.     T3[cT+cI,sT+sI,aI,bI] = 0
10.  for rT, sT
11.    for aI, bI, rI, sI
12.      T2[aI,bI,rI,sI] = 0
13.    for qT, aI, bI, rI, sI, qI
14.      T2[aI,bI,rI,sI] += C3[qT+qI,bT+bI] * T1[aT+aI,qT+qI,rT+rI,sT+sI]
15.    for cT, aI, bI, rI, sI, cI
16.      T3[cT+cI,sT+sI,aI,bI] += C2[rT+rI,cT+cI] * T2[aI,bI,rI,sI]
17.  for cT, dT, sT, aI, bI, cI, dI, sI
18.    B[aT+aI,bT+bI,cT+cI,dT+dI] += C1[sT+sI,dT+dI] * T3[cT+cI,sT+sI,aI,bI]

```

Fig. 6. Abstract code for the tiled 4-index transform.

Table 1
Candidate disk I/O placements and placement variables for the arrays in the 4-index transform

Array	Possible placements	Placement variables
T1	In Memory + $\{pI, aI, sT, rT, pT\} \times \{bI, aI, qT, sT, bT\}$	λ_0 — λ_4
C4	$\{qT, pT, aT\}$	λ_5 — λ_6
A	$\{aI, sT, rT, qT\}$	λ_7 — λ_8
T2	In Memory + $\{rI, bI\} \times \{rI, bI, cT\}$ + $\{qT\} \times \{rI, bI\}$	λ_9 — λ_{12}
C3	$\{aI, rT, aT\}$	λ_{13} — λ_{14}
T3	In Memory + $\{rI, bI, aI, cI\} \times \{cI, bI, aI, dT, cT\}$ + $\{rT\} \times \{cI, bI, aI, dT\}$	λ_{15} — λ_{19}
C2	$\{aI, sT, aT\}$	λ_{20} — λ_{21}
B	$\{cI, bI, sT, dT, cT, bT\}$	λ_{22} — λ_{24}
C1	$\{aI, sT, aT\}$	λ_{25} — λ_{26}

for the 4-index transform in Fig. 5. The 4-index transform involves four contractions, requiring three intermediate arrays. These are T1, T2, and T3 in the abstract code shown. Loops are fused so that two of the intermediates are significantly reduced in size, leaving only T1 to be a four-dimensional array. This fused abstract code was tiled and the intra-tile loops were moved to be innermost in the loop structure. The tiled abstract code for the 4-index transform is shown in Fig. 6.

Then, the possible placements for disk I/O statements are enumerated for all the arrays. The enumeration starts at the first tiling loop or the first redundant intra-tile loop and proceeds up the fusion graph. The list of candidate I/O placements is shown in Table 1. An I/O placement for an array denoted by a loop index specifies that the disk read (write) for that array is inserted above (below) the loop corresponding to that index, surrounding the use of the array. The loops surrounding the update of an array are considered for write placements, and those surrounding the read-only use of an array are considered for read placements. The read and write placements noted correspond to these loops. The read (write) placements are shown for the input (output) arrays. For intermediate arrays, the I/O cost for production

and consumption of the array are enumerated. This is shown by the cross-product of the write and read placements. The input and output arrays are disk-resident. The intermediate arrays can potentially be in memory, which is also enumerated as a possible I/O placement. The placement variables allocated to each array are also shown in Table 1. Consider the I/O placements enumerated for array T2. Array T2 could be in memory, represented by the first “placement” possibility. If T2 is disk-resident, there are three possible write and read placements, forming nine placement pairs. Each placement pair uniquely determines the read and write placement for that array. The placement pair $qT \times cT$ results in the read and write at the same node in the fusion tree, and is equivalent to T2 residing in memory. Hence it is discarded, leaving eight placements pairs, as shown.

The I/O and memory costs are the sum total of I/O and memory costs for all the arrays. Each possible placement of I/O statement for an array has a potentially different contribution to the I/O and memory cost. The contributions to the I/O and memory cost by the array T1 for some of the possible placements are shown in Table 2. Every element

Table 2
Contributions to disk I/O cost and memory cost by the array T1 in the 4-index transform

Placement variables					I/O placement		I/O cost (*Size _{T1} *8)		Memory cost
λ_0	λ_1	λ_2	λ_3	λ_4	Produce	Consume	Read	Write	
1	1	1	1	1	In Mem	In Mem	0	0	$Na*Nq*Nr*Ns*8$
1	1	1	1	0	pI	bI	Np/Tp	Np/Tp	$Tq*Tr*Ts*8$
1	1	1	0	1	pI	aI	Np/Tp	Np/Tp	$Ta*Tq*Tr*Ts*8$
⋮									
0	1	0	1	0	pT	bI	1	Nb/Tb	$Ta*Nq*Nr*Ns*8$
⋮									
0	0	1	0	1	-	-	-	-	$2*MemSize$
⋮									
0	0	0	0	0	-	-	-	-	$2*MemSize$

in a disk-resident array needs to be accessed from disk at least once. The I/O cost shown is the factor of redundant I/O over the minimum access of $Size_{T1} * 8$, where $Size_{T1}$ is the size of array T1 in the tiled code, and the size of each element is 8. When the placement variables have value 11111, T1 is in memory and has no I/O cost. T1 has 25 possible placements and any values of placement variables that are beyond 00110 do not correspond to any legal I/O placement. These values for the placement variables are pruned away by specifying the memory cost to be higher than the available memory. The table also shows that the read and write costs can potentially be different, as in the case of placement variables being equal to 01010. When determining the overall I/O costs, the read and write costs can be weighted by the average read and write times. Sequential access (read/write) time was found to be a good approximation of the actual access time, once the block size is larger than a threshold.

The tile sizes are limited to the valid range by the constraints

$$1 \leq Ta, Tb, Tc, Td \leq 190,$$

$$1 \leq Tp, Tq, Tr, Ts \leq 180.$$

The constraints on the placement variables to limit their values to either 0 or 1 is given by

$$\forall i \in \{0, \dots, 26\} \lambda_i * (1 - \lambda_i) = 0.$$

The I/O sizes are constrained to be large enough for efficient I/O on the target system. The I/O sizes are just the size of the in-memory buffers and hence can be computed from the memory costs. The read and write constraints, respectively, for the array T1 are given by

$$8 * \sum_{i,j,k,l,m=0}^1 Placement(i, j, k, l, m)$$

$$* MemCost(i, j, k, l, m) \geq readbufsize,$$

Table 3
Parameters used in the construction of the optimization problem for the 4-index transform example

Mem. limit (GB)	Min. read size (MB)	Min. write size (MB)	Read time (ns/byte)	Write time (ns/byte)
2	2	1	16	20

Table 4
Tile sizes for the 4-index transform example

Ta	Tb	Tc	Td	Tp	Tq	Tr	Ts
48	95	95	190	90	60	180	180

$$8 * \sum_{i,j,k,l,m=0}^1 Placement(i, j, k, l, m)$$

$$* MemCost(i, j, k, l, m) \geq writebufsize,$$

where

$$Placement(i, j, k, l, m) = \begin{cases} 1 & \text{if } (i = \lambda_0 \wedge j = \lambda_1 \wedge k \\ & = \lambda_2 \wedge l = \lambda_3 \wedge m = \lambda_4), \\ 0 & \text{otherwise.} \end{cases}$$

The parameters to complete the construction of the optimization problem are shown in Table 3. We determined the parameters for the Itanium-2 cluster at the Ohio Supercomputer Center, discussed in Section 6. The parameters correspond to a single-processor execution using local disks. The loop bounds Na, Nb, Nc , and Nd are set to 190 and Np, Nq, Nr , and Ns are set to 180. The array sizes ($Size_A, Size_{T1}, \dots$) are also specified.

The optimization problem thus constructed is solved using the non-linear optimization solver. The result is interpreted to obtain the concrete code shown in Fig. 7. The tile sizes determined are shown in Table 4.

```

1.  for aT, qT, rT, sT
2.    for aI, qI, rI, sI
3.      T1[aI,qI,rI,sI] = 0
4.    Write T1Disk[aT:aT+47,qT:qT+59,rT:rT+179,sT:sT+179]
5.  for aT, pT
6.    C4[1:90,1:48] = Read C4Disk[pT:pT+89,aT:aT+47]
7.    for qT, rT
8.      A[1:90,1:60,1:180,1:Ns] = Read ADisk[pT:pT+89,qT:qT+59,rT:rT+179,1:Ns]
9.    for sT
10.     T1[1:48,1:60,1:180,1:180] =
11.       Read T1Disk[aT:aT+47,qT:qT+59,rT:rT+179,sT:sT+179]
12.     for aI, pI, qI, rI, sI
13.       T1[aI,qI,rI,sI] += C4[pI,aI] * A[pI,qI,rI,sT+sI]
14.     Write T1Disk[aT:aT+47,qT:qT+59,rT:rT+179,sT:sT+179]
15.   C1[1:Ns,1:Nd] = Read C1Disk[1:Ns,1:Nd]
16. for aT, bT
17.   for cT, sT, aI, bI, cI, sI
18.     T3[cT+cI,sT+sI,aI,bI] = 0
19.   C3[1:Nq,1:95] = Read C3Disk[1:Nq,bT:bT+94]
20.   for rT
21.     C2[1:180,1:Nc] = Read C2Disk[rT:rT+179,1:Nc]
22.   for sT
23.     for aI, bI, rI, sI
24.       T2[aI,bI,rI,sI] = 0
25.     for qT
26.       T1[1:48,1:60,1:180,1:180] =
27.         Read T1Disk[aT:aT+47,qT:qT+59,rT:rT+179,sT:sT+179]
28.       for aI, bI, rI, sI, qI
29.         T2[aI,bI,rI,sI] += C3[qT+qI,bI] * T1[aI,qI,rI,sI]
30.       for cT, aI, bI, rI, sI, cI
31.         T3[cT+cI,sT+sI,aI,bI] += C2[rI,cT+cI] * T2[aI,bI,rI,sI]
32.     for cT, dT
33.       for aI, bI, cI, dI
34.         B[aI,bI,cI,dI] = 0
35.       for sT, aI, bI, cI, dI, sI
36.         B[aI,bI,cI,dI] += C1[sT+sI,dT+dI] * T3[cT+cI,sT+sI,aI,bI]
37.       Write BDisk[aT:aT+47,bT:bT+94,cT:cT+94,dT:dT+189]

```

Fig. 7. Concrete code for the 4-index transform example. $N_p = N_q = N_r = N_s = 190$, $N_a = N_b = N_c = N_d = 180$. Memory limit = 2 GB.

6. Experimental results

We evaluated the developed approach, referred to as the DCS approach, by comparing it with two alternatives. The first, referred to as the equi-tile-size approach, is used in state-of-the-art quantum chemistry codes. In this approach, equal tile sizes are chosen for all loop indices. The tile sizes are made large enough to fully utilize the available memory. The placement of I/O statements is determined in a greedy fashion. For a given set of tile sizes, the I/O statements are placed at that position in the parse tree in which the total size of the data accessed in that subtree, rooted at that position, just fits in the available memory.

The second approach, referred to as the uniform sampling approach, was developed for locality optimization of the disk-memory hierarchy [13,28]. A greedy approach to disk I/O placement is used, where for each set of tile sizes, the algorithm places read/write statements immediately inside those loops at which the memory limit is exceeded. The tile size search space is sampled uniformly in a logarithmic fashion along each dimension. This sampled search space is then explored using a brute force approach.

Performance was evaluated on an Itanium-2 Cluster at the Ohio Supercomputer Center. Each node in the cluster is a dual Itanium-2 900 MHz system running Linux 2.4.18. Each node has 4 GB of memory and an 80 GB SCSI hard disk. The

generated code was compiled using the Intel Itanium Fortran Compiler for Linux (efc version 7.1). For code-generation purposes, the physical memory available to the computation is specified as half the available memory to minimize any paging effects.

The performance of the concrete code generated by the three approaches was evaluated for three computations. The first is a three-contraction computation, whose abstract codes are shown in Figs. 5 (four-index transform), 8 (CCD kernel), and 9 (three-contraction). It is a synthetic computation, in which the outputs of two tensor contractions are contracted again. The loop structure of the computation is shown in Fig. 9. This is a prevalent subtree in tensor contraction operator trees, though it does not occur independently in many codes. The problem size is varied by increasing the range of the loop index a , with a corresponding increase in size of the T array. The results are shown in Fig. 10. The size of the T array is shown on the x -axis, with the predicted and measured disk I/O cost shown along the y -axis. It can be observed that the DCS approach is consistently better than the uniform sampling approach, which in turn is superior to using equal tile sizes. The experimentally measured performance matches prediction very closely.

The second computation used for testing is the four-index transform, which was introduced earlier. The imperfectly nested loop structure shown in Fig. 5 was used. The number

```

1.  for a, i
2.    for b, j
3.      t_2[b,j] = 0
4.    for c, k
5.      t_1 = ((2.0 * v_ovvo[k,a,c,i]) + (-1.0 * v_ovov[k,a,i,c]))
6.      for b, j
7.        t_2[b,j] += (t[b,c,j,k] * t_1)
8.    for b, j
9.      t_3 = 0
10.     for c, k
11.       t_3 += (t[c,b,j,k] * v_ovvo[k,a,c,i])
12.     i0[a,b,i,j] = ((t_2[b,j] + (-1.0 * t_3)) + v_vvoo[a,b,i,j])

```

Fig. 8. Abstract code for the kernel from the CCD equation.

```

1.  for a, b, i, j
2.    X[a,b,i,j] = 0
3.  for b, d
4.    for j, l
5.      t_2[j,l] = 0
6.    for c, e
7.      t_1 = 0
8.      for f, k
9.        t_1 += (D[b,f,e,k] * N[d,c,f,k])
10.     for j, l
11.       t_2[j,l] += (S[c,j,e,l] * t_1)
12.   for a, i, j, l
13.     X[a,b,i,j] += (T[a,d,i,l] * t_2[j,l])

```

Fig. 9. Abstract code for the three-contraction example.

of virtual orbitals (V) was varied, to vary the problem size. Fig. 11 shows the predicted and measured disk I/O costs for the various problem sizes considered, shown in terms of the size of array A , the dominant input array affecting the memory requirement. For this example too, the DCS approach is superior to uniform sampling, which is better than the equi-tile-size approach. For the equi-tile-size approach, the experimentally measured data is a bit lower than prediction; we believe that this is a consequence of caching of produced-consumed files in physical memory.

The third computation considered is a sub-computation from the Coupled Cluster Doubles (CCD) equation [17,36,40] for ab initio electronic structure modeling. The computation is given by the loop structure shown in Fig. 8. The predicted and measured disk I/O cost for various virtual orbital ranges was evaluated. The results are shown in Fig. 12. For this computation, the DCS approach is again superior to the other two. However, surprisingly, we find that the uniform sampling approach is not consistently better than the equi-tile-size approach. We believe that this is due to sharp peaks and troughs in the disk-IO-cost as a function of tile size, causing uniform search to miss many local optima.

Overall, the graphs show that the predicted disk I/O closely matches the measured cost. The equal-tile-size approach generally performs worse than the other two approaches, while the DCS approach performs consistently better than the other two approaches. It is up to four times

better than the equal-tile-size approach and up to two times better than the uniform sampling approach.

For the parallel context, we use the Global Arrays (GA) [45] and Disk-Resident Arrays (DRA) [44] framework. The GA model provides an abstraction of global shared multi-dimensional arrays, transparently implemented on systems with physically distributed memory. The DRA model extends the global shared abstraction to disk-resident multi-dimensional arrays, permitting an arbitrary multi-dimensional segment of a DRA to be moved into a memory-resident GA. The aggregate memory available on all the processors was used as the memory available for the computation, and tiling was done as in the sequential case.

The measured disk I/O times for the parallel code generated for the three-contraction example and the four-index transform are shown in Figs. 13 and 14 respectively. The DCS approach performs significantly better than the other approaches for the different processor counts considered. In particular, considerable improvement in performance can be observed for smaller processor counts. This shows the improved resource-utilization by the DCS approach, potentially enabling larger problems to be computed on a given machine more efficiently.

6.1. Computational complexity of the three evaluated approaches

The three approaches represent varying degrees of complexity in the code generation process, in terms of the determination of I/O placements and tile sizes. With the equal-tiles approach, the tile sizes and the I/O placements are determined in a straightforward manner based solely on the available memory. This approach vastly simplifies code generation, while not necessarily resulting in quality code. The choice of equal tile sizes for all the arrays ignores various aspects of the program such as the reuse characteristics of the different arrays, thus resulting in sub-optimal code.

The uniform sampling approach takes a greedy approach to placement of I/O statements, and searches the sample space of possible tile sizes to choose the set of tile sizes that minimize the I/O cost. The greedy I/O placement strategy,

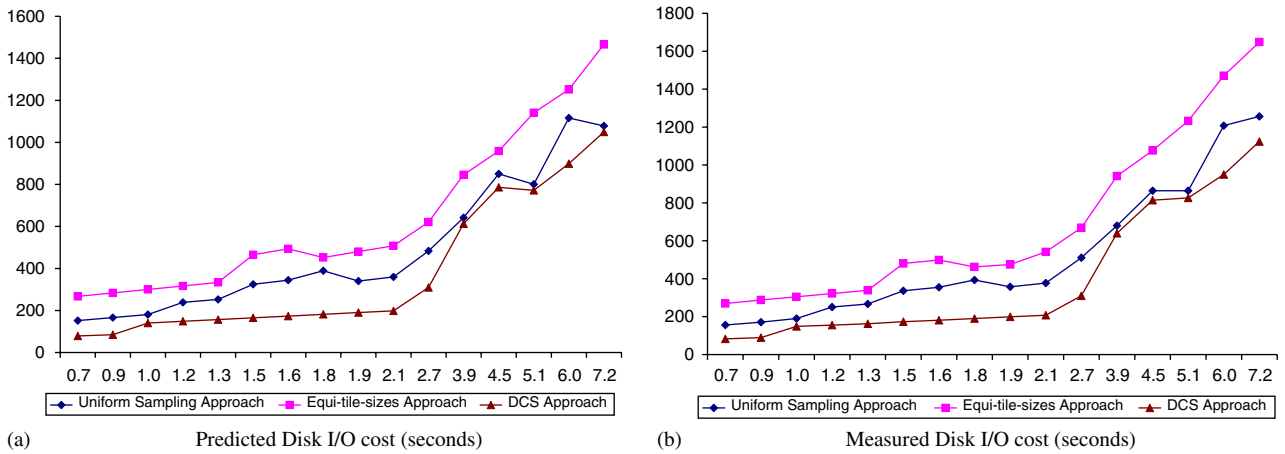


Fig. 10. The (a) predicted and (b) measured disk I/O cost, in seconds, for the three-contraction example. The size of T array is shown along x -axis, in gigabytes.

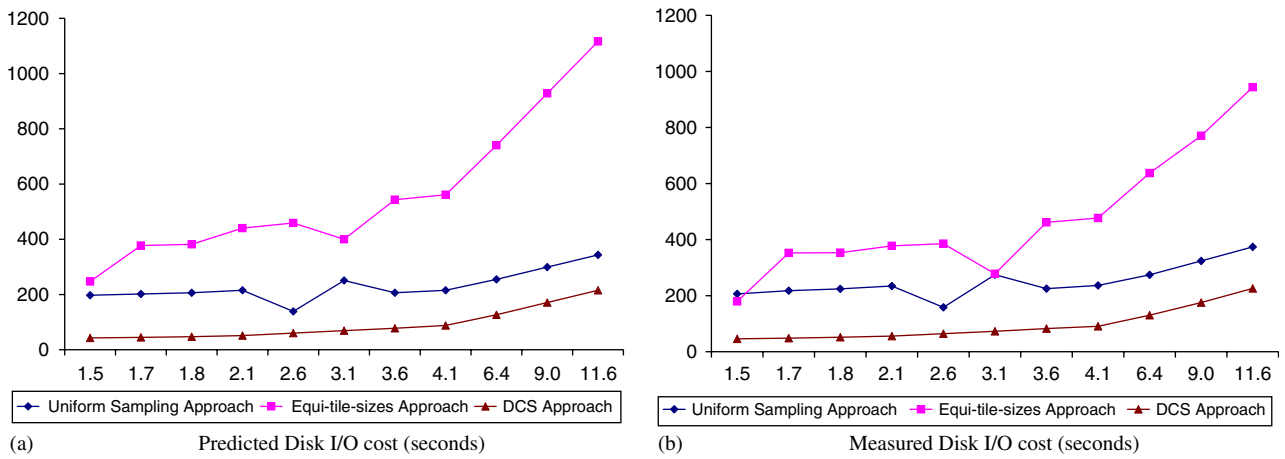


Fig. 11. The (a) predicted and (b) measured disk I/O cost, in seconds, for the four-index transform. The size of A array is shown along x -axis, in gigabytes.




Fig. 12. The (a) predicted and (b) measured disk I/O cost, in seconds, for the computation from the CCD equation. The virtual orbital ranges (V) are shown along the x -axis.

while simplifying code-generation complexity, is not always the best strategy. It also decouples the two phases of the code-generation problem, thus potentially affecting the run-

ning time. Also, the sampling nature of the tile size search does not always produce optimal code, while significantly affecting code-generation time. The dimensionality of the




Fig. 13. The measured disk I/O cost, in seconds, for the generated parallel code for the three-contraction example. The number of processors is shown along the x-axis.

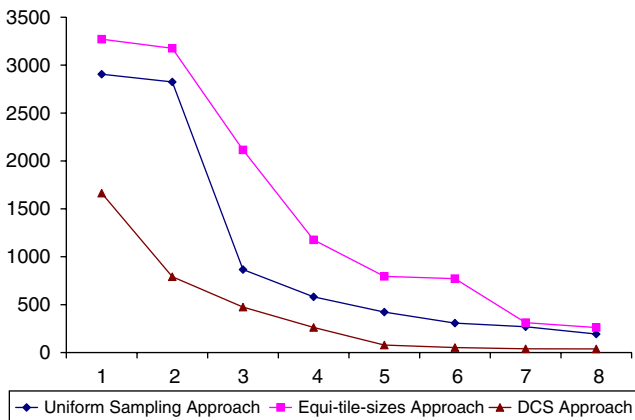


Fig. 14. The measured disk I/O cost, in seconds, for the generated parallel code for the four-index transform. The number of processors is shown along the x-axis.

search space is linearly proportional to the number of loop indices.

The DCS approach produces a composite cost function combining the effects of I/O placements and tile sizes on the disk I/O cost. The search over this composite space uses heuristics established in the solver. The encoding of the I/O placements is such that the search space to be explored for the I/O placements increases logarithmically with the number of loop indices. Hence, the complexity of the sample space to be explored is still linear in the number of loop indices, while generally generating a more globally optimal solution.

7. Conclusion

We have described an approach to the synthesis of out-of-core algorithms for a class of imperfectly nested loops. The approach was developed for the implementation in a

component of a program synthesis system targeted at the quantum chemistry domain. The determination of optimal placements of disk I/O statements and choice of tile sizes requires search in a very large search space. By formulating it as a non-linear constrained optimization problem, and use of a general-purpose constrained optimization solver, code was generated that outperforms other approaches by up to a factor of four.

Acknowledgements

We sincerely thank Prof. Benjamin Wah and Yi Xin Chen of the University of Illinois for their significant help with using the Discrete Constrained Search (DCS) Solver. We express our appreciation to the reviewers of the paper for their suggestions. We are grateful to the Ohio Supercomputer Center (OSC) for the use of their computing facilities.

References

- [1] N. Ahmed, N. Mateev, K. Pingali, Synthesizing transformations for locality enhancement of imperfectly-nested loops, in: Proceedings of the ACM International Conference on Supercomputing, 2000, pp. 141–152.
- [2] N. Ahmed, N. Mateev, K. Pingali, Synthesizing transformations for locality enhancement of imperfectly-nested loops, *Internat. J. Parallel Programming* 29 (5) (2001) 493–544.
- [3] J. Anderson, S. Amarasinghe, M. Lam, Data and computation transformations for multiprocessors, in: Proceedings of the Fifth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, July 1995, pp. 166–178.
- [4] G. Baumgartner, A. Auer, D. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Hirata, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, A. Sibiryakov, Synthesis of high-performance parallel programs for a class of ab initio quantum chemistry models, *Proc. IEEE* 93 (2) (2005) 276–292.
- [5] G. Baumgartner, D. Bernholdt, D. Cociorva, R. Harrison, S. Hirata, C. Lam, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, A high-level approach to synthesis of high-performance codes for quantum chemistry, in: Proceedings of the Supercomputing 2002, November 2002.
- [6] R. Bordawekar, Techniques for compiling I/O intensive parallel programs, Ph.D. Thesis, Department of Electrical and Computer Engineering, Syracuse University, April 1996.
- [7] R. Bordawekar, A. Choudhary, K. Kennedy, C. Koelbel, M. Paleczny, A model and compilation strategy for out-of-core data-parallel programs, in: Proceedings of the Fifth ACM Symposium on Principles and Practice of Parallel Programming, 1995, pp. 1–10.
- [8] R. Bordawekar, A. Choudhary, J. Ramanujam, Automatic optimization of communication in out-of-core stencil codes, in: Proceedings of the 10th ACM International Conference on Supercomputing, 1996, pp. 366–373.
- [9] Y. Chen, Optimal anytime search for constrained nonlinear programming, Master's Thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, IL, May 2001.
- [10] Y. Chen, M. Winslett, Y. Cho, S. Kuo, Automatic parallel I/O performance optimization in Panda, in: Proceedings of the 10th Annual ACM Symposium on Parallel Algorithms and Architectures, 1998, pp. 108–118.
- [11] D. Cociorva, G. Baumgartner, C. Lam, P. Sadayappan, J. Ramanujam, M. Nooijen, D. Bernholdt, R. Harrison, Space–time trade-off

- optimization for a class of electronic structure calculations, in: Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation, 2002, pp. 177–186.
- [12] D. Cociorva, X. Gao, S. Krishnan, G. Baumgartner, C. Lam, P. Sadayappan, J. Ramanujam, Global communication optimization for tensor contraction expressions under memory constraints, in: Proceedings of the 17th International Parallel and Distributed Processing Symposium (IPDPS), 2003.
- [13] D. Cociorva, J. Wilkins, G. Baumgartner, P. Sadayappan, J. Ramanujam, M. Nooijen, D. Bernholdt, R. Harrison, Towards automatic synthesis of high-performance codes for electronic structure calculations: data locality optimization, in: Proceedings of the International Conference on High Performance Computing, vol. 2228, Springer, Berlin, 2001, pp. 237–248.
- [14] D. Cociorva, J. Wilkins, C. Lam, G. Baumgartner, P. Sadayappan, J. Ramanujam, Loop optimization for a class of memory-constrained computations, in: Proceedings of the 15th ACM International Conference on Supercomputing (ICS'01), 2001, pp. 500–509.
- [15] S. Coleman, K. McKinley, Tile size selection using cache organization and data layout, in: Proceedings of the SIGPLAN '95 Conference on Programming Languages Design and Implementation, 1995, pp. 279–290.
- [16] T. Cormen, A. Colvin, ViC*: a preprocessor for virtual-memory C*, Technical Report PCS-TR94-243, Dartmouth College, November 1994.
- [17] T. Crawford, H. Schaefer III, An introduction to coupled cluster theory for computational chemists, in: K. Lipkowitz, D. Boyd (Eds.), Reviews in Computational Chemistry, vol. 14, Wiley, New York, 2000, pp. 33–136.
- [18] J. Dongarra, J. Du Croz, I. Duff, S. Hammarling, A set of level-3 basic linear algebra subprograms, ACM Trans. Math. Software 16 (1) (1990) 1–17.
- [19] R. Fourer, D. Gay, B. Kernighan, AMPL: A Modeling Language for Mathematical Programming, 2002.
- [20] S. Ghosh, M. Martonosi, S. Malik, Cache miss equations: a compiler framework for analyzing and tuning memory behavior, ACM Trans. Programming Language Systems 21 (4) (1999) 703–746.
- [21] M. Kandemir, A. Choudhary, J. Ramanujam, An I/O conscious tiling strategy for disk-resident data sets, J. Supercomput. 21 (3) (2002) 257–284.
- [22] M. Kandemir, A. Choudhary, J. Ramanujam, R. Bordawekar, Compilation techniques for out-of-core parallel computations, Parallel Comput. 24 (3–4) (June 1998) 597–628.
- [23] M. Kandemir, A. Choudhary, J. Ramanujam, M. Kandaswamy, A unified framework for optimizing locality, parallelism, and communication in out-of-core computations, IEEE Trans. Parallel Distributed Systems 11 (7) (2000) 648–668.
- [24] I. Kodukula, N. Ahmed, K. Pingali, Data-centric multi-level blocking, in: Proceedings of the SIGPLAN Conference Programming Language Design and Implementation, 1997, pp. 346–357.
- [25] I. Kodukula, K. Pingali, Data-centric transformations for locality enhancement, Internat. J. Parallel Programming 29 (3) (2001) 319–364.
- [26] I. Kodukula, K. Pingali, R. Cox, D. Maydan, An experimental evaluation of tiling and shackling for memory hierarchy management, in: Proceedings of the ACM International Conference on Supercomputing (ICS 99), 1999, pp. 482–491.
- [27] S. Krishnamoorthy, G. Baumgartner, D. Cociorva, C. Lam, P. Sadayappan, On efficient out-of-core matrix transposition, Technical Report OSU-CIRSC-9/03-T52, The Ohio State University, Columbus, OH, September 2003.
- [28] S. Krishnan, Data locality optimization for synthesis of out-of-core programs, Master's Thesis, The Ohio State University, Columbus, OH, September 2003.
- [29] S. Krishnan, S. Krishnamoorthy, G. Baumgartner, D. Cociorva, C. Lam, P. Sadayappan, J. Ramanujam, D. Bernholdt, V. Choppella, Data locality optimization for synthesis of efficient out-of-core algorithms, in: Proceedings of the International Conference on High Performance Computing, 2003.
- [30] C. Lam, Performance optimization of a class of loops implementing multi-dimensional integrals, Ph.D. Thesis, The Ohio State University, Columbus, OH, August 1999.
- [31] C. Lam, D. Cociorva, G. Baumgartner, P. Sadayappan, Memory-optimal evaluation of expression trees involving large objects, in: Proceedings of the International Conference on High Performance Computing, Springer, Berlin, 1999, pp. 103–110.
- [32] C. Lam, D. Cociorva, G. Baumgartner, P. Sadayappan, Optimization of memory-usage and communication requirements for a class of loops implementing multi-dimensional integrals, in: Proceedings of the 12th Workshop on Languages and Compilers for Parallel Computing (LCPC), Springer, Berlin, 1999, pp. 350–364.
- [33] C. Lam, P. Sadayappan, R. Wenger, On optimizing a class of multi-dimensional loops with reductions for parallel execution, Parallel Processing Lett. 7 (2) (1997) 157–168.
- [34] C. Lam, P. Sadayappan, R. Wenger, Optimization of a class of multi-dimensional integrals on parallel machines, in: Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing, 1997.
- [35] M. Lam, E. Rothberg, M. Wolf, The cache performance and optimizations of blocked algorithms, in: Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, 1991, pp. 63–74.
- [36] T. Lee, G. Scuseria, Achieving chemical accuracy with coupled cluster theory, in: S. Langhoff (Ed.), Quantum Mechanical Electronic Structure Calculations with Chemical Accuracy, 1997, pp. 47–109.
- [37] W. Li, Compiling for NUMA parallel machines, Ph.D. Thesis, Cornell University, August 1993.
- [38] W. Li, Compiler cache optimizations for banded matrix problems, in: Proceedings of the International Conference on Supercomputing, 1995, pp. 21–30.
- [39] A. Lim, S. Liao, M. Lam, Blocking and array contraction across arbitrarily nested loops using affine partitioning, in: Proceedings of the Eighth ACM SIGPLAN Symposium on Principles and Practices of Parallel Programming, ACM Press, New York, 2001, pp. 103–112.
- [40] J. Martin, Benchmark studies on small molecules, in: P. Schleyer, P. Schreiner, N. Allinger, T. Clark, J. Gasteiger, P. Kollman, H. Schaefer III (Eds.), Encyclopedia of Computational Chemistry, vol. 1, Wiley, New York, 1998, pp. 115–128.
- [41] K. McKinley, S. Carr, C. Tseng, Improving data locality with loop transformations, ACM TOPLAS 18 (4) (July 1996) 424–453.
- [42] N. Mitchell, K. Hogstedt, L. Carter, J. Ferrante, Quantifying the multi-level nature of tiling interactions, Internat. J. Parallel Programming 26 (6) (June 1998) 641–670.
- [43] T. Mowry, A. Demke, O. Krieger, Automatic compiler-inserted I/O prefetching for out-of-core applications, in: Proceedings of the Second Symposium on Operating Systems Design and Implementations, 1996, pp. 3–17.
- [44] J. Nieplocha, I. Foster, Disk resident arrays: an array-oriented I/O library for out-of-core computations, in: Proceedings of the Sixth Symposium on the Frontiers of Massively Parallel Computation, 1996, pp. 196–204.
- [45] J. Nieplocha, R. Harrison, R. Littlefield, Global arrays: a nonuniform memory access programming model for high-performance computers, J. Supercomput. 10 (1996) 197–220.
- [46] M. Paleczny, K. Kennedy, C. Koelbel, Compiler support for out-of-core arrays on parallel machines, Technical Report 94509-S, Rice University, Houston, TX, December 1994.
- [47] G. Rivera, C. Tseng, Eliminating conflict misses for high performance architectures, in: Proceedings of the International Conference on Supercomputing, 1998, pp. 353–360.
- [48] Y. Song, Z. Li, New tiling techniques to improve cache temporal locality, in: Proceedings of the ACM SIGPLAN Conference

on Programming Language Design and Implementation, 1999, pp. 215–228.

- [49] R. Thakur, R. Bordawekar, A. Choudhary, R. Ponnusamy, T. Singh, PASSION runtime library for parallel I/O, in: Proceedings of the Scalable Parallel Libraries Conference, 1994, pp. 119–128.
- [50] B. Wah, Y. Chen, Web Interface for Discrete Constrained Search Solver. <http://manip.crhc.uiuc.edu/csa>
- [51] B. Wah, Y. Chen, Constrained genetic algorithms and their applications in nonlinear constrained optimization, in: Proceedings of the 11th IEEE International Conference on Tools with Artificial Intelligence, 2000, pp. 286–293.
- [52] B. Wah, Y. Chen, Optimal anytime constrained simulated annealing for constrained global optimization, in: Proceedings of the ACM Symposium on Principles and Practice of Constraint Programming, Springer, Berlin, 2000, pp. 425–439.
- [53] B. Wah, Y. Chen, Hybrid constrained simulated annealing and genetic algorithms for nonlinear constrained optimization, in: Proceedings of the IEEE Congress on Evolutionary Computation, 2001, p. 2001.
- [54] M. Wolf, M. Lam, A data locality algorithm, in: Proceedings of the ACM SIGPLAN Conference Programming Language Design and Implementation, 1991, pp. 30–44.
- [55] M. Wolf, D. Maydan, D. Chen, Combining loop transformations considering caches and scheduling, *Internat. J. Parallel Programming* 26 (4) (1998) 479–503.




Sandhya Krishnan received the B.E. degree in Computer Engineering from the Bombay University, India in March 2001, and graduated with a Masters in Computer and Information Science from The Ohio State University, Columbus, Ohio in September 2003. She continued to work as a Systems Developer on the same project till May 2004. Her research interests include high-performance computing and development of optimization algorithms.



Sriram Krishnamoorthy was born in Chennai, India in 1981. He received his B.E. degree from the College of Engineering, Guindy, Anna University in 2002. He has been pursuing the Ph.D. degree at the Ohio State University, under the supervision of Prof. P. Sadayappan, since September 2002. His research interests include high-performance computing, out-of-core algorithms and optimizations for scientific computing.



Gerald Baumgartner received the Dipl. Ing. degree from the University of Linz, Austria, and M.S. and Ph.D. degrees from Purdue University, all in Computer Science. He began his academic career at The Ohio State University in 1997. Since 2004, he is at the Department of Computer Science at Louisiana State University. His research interest includes compiler optimizations, the design and implementation of domain-specific and object-oriented languages, desktop grids, and development and testing tools for object-oriented and embedded systems programming.



Chi-Chung Lam graduated from The Ohio State University with a Ph.D. degree in Computer and Information Science in 1999. The title of his dissertation was “Performance optimization of a class of loops implementing multi-dimensional integrals”.



J. Ramanujam received the B. Tech. degree in Electrical Engineering from the Indian Institute of Technology, Madras, India in 1983, and his M.S. and Ph.D. degrees in Computer Science from The Ohio State University in 1987 and 1990 respectively. He is currently the John E. & Beatrice L. Ritter Distinguished Professor in the Department of Electrical and Computer Engineering at Louisiana State University. His research interests are in compilers for high-performance computer systems, embedded systems, software optimizations for low-power

software optimizations for low-power computing, high-level hardware synthesis, parallel architectures and algorithms. He has published over 120 papers in refereed journals and conferences in these areas in addition to several book chapters and a book. He received the National Science Foundation’s Young Investigator Award in 1994. In addition, he has received the best paper award at the 2003 International Conference on High-Performance Computing (HiPC 2003) for his work with others on compiler optimizations for quantum chemistry computations.



P. Sadayappan received the B. Tech. degree from the Indian Institute of Technology, Madras, India, and an M.S. and Ph.D. from the State University of New York at Stony Brook, all in Electrical Engineering. He is currently a Professor in the Department of Computer Science and Engineering at The Ohio State University. His research interests include Compile/Runtime Optimization and Scheduling and Resource Management for Parallel/Distributed Systems.



Venkatesh Choppella is on the faculty at the Indian Institute of Information Technology and Management—Kerala. He holds a B. Tech. from the Indian Institute of Technology, Kanpur and a Ph.D. from the Indiana University, both in Computer Science. He has held research and engineering positions at Xerox Corporation, Hewlett-Packard Company, Indiana University and Oak Ridge National Laboratory. His interests are in programming languages, compilers for domain specific languages, automated deduction, and software engineering.