

Automated Operation Minimization of Tensor Contraction Expressions in Electronic Structure Calculations

Albert Hartono¹, Alexander Sibiryakov¹, Marcel Nooijen³, Gerald Baumgartner⁴,
David E. Bernholdt⁶, So Hirata⁷, Chi-Chung Lam¹, Russell M. Pitzer²,
J. Ramanujam⁵, and P. Sadayappan¹

¹ Dept. of Computer Science and Engineering

² Dept. of Chemistry, The Ohio State University,
Columbus, OH, 43210 USA

³ Dept. of Chemistry, University of Waterloo,
Waterloo, Ontario N2L BG1, Canada

⁴ Dept. of Computer Science

⁵ Dept. of Electrical and Computer Engineering,
Louisiana State University, Baton Rouge, LA 70803 USA

⁶ Computer Sci. & Math. Div., Oak Ridge National Laboratory,
Oak Ridge, TN 37831 USA

⁷ Quantum Theory Project, University of Florida,
Gainesville, FL 32611 USA

Abstract. Complex tensor contraction expressions arise in accurate electronic structure models in quantum chemistry, such as the Coupled Cluster method. Transformations using algebraic properties of commutativity and associativity can be used to significantly decrease the number of arithmetic operations required for evaluation of these expressions, but the optimization problem is NP-hard. Operation minimization is an important optimization step for the Tensor Contraction Engine, a tool being developed for the automatic transformation of high-level tensor contraction expressions into efficient programs. In this paper, we develop an effective heuristic approach to the operation minimization problem, and demonstrate its effectiveness on tensor contraction expressions for coupled cluster equations.

1 Introduction

Currently, manual development of accurate quantum chemistry models is very tedious and takes an expert several months to years to develop and debug. The Tensor Contraction Engine (TCE) [2, 1] is a tool that is being developed to reduce the development time to hours/days, by having the chemist specify the computation in a high-level form, from which an efficient parallel program is automatically synthesized. This should enable the rapid synthesis of high-performance implementations of sophisticated ab-initio quantum chemistry models, including models that are too tedious for manual development by quantum chemists. An important first step in the synthesis process of the TCE

is that of algebraic manipulation of the input tensor contraction expressions, to find an equivalent form with minimized operation count.

We illustrate the operation minimization problem using simple examples. Consider the following tensor contraction expression involving three tensors t , f and s , with a and c representing virtual orbital indices with range V , and i and j representing occupied orbital indices with range O . Computed as a single nested loop computation, the number of multiply-accumulate operations needed would be O^2V^2 .

$$(1) \quad r_i^a += \sum_{c,k} t_i^c f_c^k s_k^a, \quad (\text{cost } O^2V^2)$$

However, by performing a two-step computation with an intermediate I , it is possible to compute the result using $2OV^2$ operations:

$$(2) \quad I_c^a = \sum_k f_c^k s_k^a, \quad (\text{cost } OV^2)$$

$$(3) \quad r_i^a += \sum_c t_i^c I_c^a, \quad (\text{cost } OV^2)$$

Another possibility using O^2V computations, which is more efficient when $V > O$ (as is usually the case), is shown below:

$$(4) \quad I_i^k = \sum_c t_i^c f_c^k, \quad (\text{cost } O^2V)$$

$$(5) \quad r_i^a += \sum_k I_i^k s_k^a, \quad (\text{cost } O^2V)$$

The above example illustrates the problem of single-term optimization, also called strength reduction: find the best sequence of two-tensor contractions to achieve a multi-tensor contraction. Different orders of contraction can result in very different operation costs; for the above example, if the ratio of V/O were 10, there is an order of magnitude difference in the number of arithmetic operations for the two choices.

For more complex expressions with several tensors to be contracted, the number of possible ways of forming intermediates is exponential in the number of tensors. The single-term optimization problem is a generalization of the well known matrix-chain multiplication problem, but while the latter has a simple polynomial time dynamic programming solution, the former problem has been shown to be NP-complete [8].

Let us next consider an expression with two terms:

$$(6) \quad r_{ij}^{ab} += \sum_{c,d} t_i^c s_j^d v_{cd}^{ab} + \sum_{c,d} u_{ij}^{cd} v_{cd}^{ab}, \quad (\text{cost } 2O^2V^4)$$

If each term were individually optimized via strength reduction, we would have:

$$(7) \quad I_{id}^{ab} = \sum_c t_i^c v_{cd}^{ab}, \quad (\text{cost } OV^4)$$

$$(8) \quad r_{ij}^{ab} += \sum_d s_j^d I_{id}^{ab} + \sum_{c,d} u_{ij}^{cd} v_{cd}^{ab}, \quad (\text{cost } O^2V^3 + O^2V^4)$$

However, a better approach to reducing the overall operation cost would be as follows:

$$(9) \quad I_{ij}^{cd} = t_i^c s_j^d + u_{ij}^{cd}, \quad (\text{cost } O^2V^2)$$

$$(10) \quad r_{ij}^{ab} += \sum_{c,d} I_{ij}^{cd} v_{cd}^{ab}, \quad (\text{cost } O^2V^4)$$

Thus it can be seen that single-term optimization (strength reduction) is not an optimal strategy. We have to look at the expression in the global context to determine the optimal evaluation. Evaluation of the binary terms in an expression constitute an

intrinsic cost to evaluating the tensor product. Little can be done (except for instances of factorization of the form $AC + BC \rightarrow (A + B)C$) to reduce the cost of the binary terms. However, the cost of evaluating terms that are ternary or higher can be greatly reduced by combining them with the binary terms that have to be evaluated anyway. In the present example, the additional cost of evaluating the ternary term is reduced from $O^2V^4 + O^2V^3$ to O^2V^2 . The expensive O^2V^4 multiplication that would be counted in single term optimization disappears as the multiplication has to be done anyway in the binary term.

The goal of operation minimization is to find an optimal or near-optimal factorization of the input tensor contraction expression to evaluate the ternary+ terms, given the presence of the binary terms that carry an intrinsic basic cost. In this paper, we develop algorithms for operation minimization. The solution presented in this paper is an extension of the techniques first reported in [11]. The conceptually simplest approach is to use an exhaustive search algorithm that is guaranteed to determine the optimal factorization. However, its runtime grows exponentially with the number of terms in the tensor contraction expression, making it infeasible for use on the more complex coupled cluster equations. We then develop heuristic search strategies for the operation minimization problem. The best algorithm is found to be a random-descent heuristic, which is then used to explore the generated solutions for a range of values of V/O . The results validate the effectiveness of the use of an automated approach to generating operation-minimal factorizations for large tensor contraction expressions.

2 Related Work

Compilers use common subexpression elimination to reduce the number of arithmetic operations [4]. However they do not consider algebraic properties such as associativity and distributivity. Computer algebra systems typically contain factorization algorithms, e.g., for finding roots of polynomials [3]. Similarly, an algorithm based on factor graphs can be used to factor functions of many variables into products of local functions [6]. However, the emphasis of these approaches is mainly on symbolic manipulation instead of on minimizing operation counts based on index range information. Winograd [13] addressed the general problem of evaluating multiple expressions that share common variables using the minimum number of arithmetic operations. Miller [9] suggested several analytical and numerical techniques for reducing the operation count in computational electromagnetic applications.

The work presented in this paper builds upon methods developed in a recent thesis by Sibiriyakov [11]. The problem of strength reduction for arbitrary tensor contraction expressions was addressed in [8, 7]. We are not aware of other work that has addressed in a general manner the operation minimization problem that we consider in this paper. In developing efficient implementations of electronic structure methods such as the coupled cluster methods [10, 12], quantum chemists have used domain heuristics for strength reduction and factorization for specific kinds of tensor contraction expressions, but have not developed approaches to solve the operation minimization problem for arbitrary tensor contraction expressions.

3 Operation Minimization Algorithms

In this section, we outline several algorithms for the operation minimization problem. Due to space limitations, pseudo-code and details are omitted, but may be found in [5].

3.1 Exhaustive Search

We first describe an exhaustive search algorithm that systematically evaluates all possible factorizations of the input tensor contraction expression to determine the form with lowest operation count. This search is implemented recursively using memoization, which is equivalent to a dynamic programming approach implemented in a top-down manner.

Considering a particular tensor as a factor, exhaustive search enumerates all possible factorizations, which grows exponentially. If a factor appears in n terms of the tensor contraction expression, the number of possible factorizations with respect to that factor is $2^n - n$. For instance, all possible factorizations of an expression $AB + AC + AD$ are $AB + AC + AD$, $A(B + C) + AD$, $A(B + D) + AC$, $A(C + D) + AB$, and $A(B + C + D)$.

As in standard dynamic programming, a storage table is maintained with solutions for subexpressions; hence, we need not re-evaluate subexpressions that have been previously considered. Matching two equivalent expressions requires generating canonical forms of both expressions. If a canonical form of a subexpression is found as a key in the storage table, the corresponding entry value, which is the optimal solution of the subexpression, is fetched and replicated. The indices of replica may differ from the original subexpression; thus, renaming indices of the replica is required.

The exhaustive search algorithm is guaranteed to find the operation-minimal factorization of the input expression, but since its time complexity grows exponentially with the number of terms, it may be impractical to use in optimizing expressions with a large number of terms. We therefore also implemented several heuristic search strategies for operation minimization.

3.2 Time-Limited Exhaustive Search with Tier-Based Partitioning

By imposing a time limit, we can avoid an indefinitely long search time that often occurs in exhaustive search. Each time exhaustive search exceeds the specified time limit, we suspend the search and store the result of the partially executed exhaustive search. Afterward the original expression is divided into two smaller groups each with half the original group's terms. These two subsets of terms can be individually factorized using the same time-limited exhaustive search and the partially factored terms can then be recombined. The cost of the combined expression is compared with the result of the timed exhaustive search that was previously interrupted. The result with the minimum operation count is returned. The splitting process is continued till each group of terms is successfully factorized within the time limit.

Prior to each splitting, it is essential to sort the expression terms in a decreasing order of term cost (as determined by single-term optimization), allowing higher-ordered terms to be placed and optimized in the same group after the splitting.

In the worst case, a splitting can occur whenever a time-limited exhaustive search is applied. We can view these recursive splittings as a binary tree in which each node represents one exhaustive search. The maximum number of nodes in such a binary tree will be $\sum_{i=0}^{\log_2 N} \binom{N}{2^i} \approx 2N$, where N indicates the number of terms in the original expression. Therefore, the time complexity of time-limited exhaustive search is $O(TN)$, where T is the given time limit. For generating experimental results in this paper, the time limit used for exhaustive search was set to ten minutes. In addition, very similar results were also obtained with a shorter time limit of ten seconds, demonstrating the efficacy of the algorithm in finding a reasonable solution quickly.

We evaluated another approach to partitioning based on tier groups. Two terms are placed into the same tier group if they have the same number of tensors. The terms in the input expression are first partitioned into tier groups. Optimizing and combining tier groups is done incrementally. Suppose we have groups at tier 2, tier 3, and tier 4. We first optimize tier 2 with timed exhaustive search. The optimized tier 2 terms are then grouped together with the unoptimized tier 3 terms and then optimized with timed exhaustive search. Then this result is grouped with the unoptimized tier 4 terms and then optimized again with timed exhaustive search.

3.3 Direct Descent Search

Direct descent is a greedy algorithm that chooses the best local factorization at every step. All pairs of terms that can be combined using the distributive property of multiplication over addition are considered, and the transformation that provides the greatest reduction in operation count is chosen. For a particular factor, if the number of factorizable terms in the expression is n the total number of possible two-term factorizations is $n(n-1)/2$. At every step, all possible two-term factorizations are evaluated and the best one is chosen; this process is repeated until no more factorization is possible. Based on the number of factorizations considered at each step, the runtime complexity of direct descent obviously grows polynomially with the degree of terms in the input expression.

3.4 Random Descent Search

Random descent search is a modified version of direct descent search that attempts to avoid some local minima by making random choices for two-term factorization at the initial steps. These random moves are then followed by direct descent moves. Through experimentation, it was found best to make the number of random moves to be one fourth of the total number of terms in the input expressions. Using too many random factorizations at the initial steps was found to give poorer results; too few random factorizations at the initial steps did not contribute a significant improvement. In order to further minimize the operation count, we first execute a direct descent search and store its factorization result; after that, one hundred attempts of random descent are repeated. The best result from these one hundred tries is compared with the result of direct descent that was initially executed. The result with the minimum operation count is returned.

4 Experimental Results

We have implemented the algorithms for searching a formula with minimum number of operations as described previously. They were tested on complex tensor contraction expressions that appear in the “coupled cluster” family of quantum chemical methods. We used the coupled cluster equations including just single and double excitations (CCSD) and also single, double, and triple excitations (CCSDT) as representatives of the many different computational chemistry methods based on tensor contraction expressions. These methods involve coupled equations which determine the single excitation amplitudes (referred to here the “T1” equation), double excitation amplitudes (T2), and in the case of CCSDT, triple excitations (T3).

Table 1 shows the number of terms in each of the equations, along with the number of arithmetic operations of evaluating the equation. The number of arithmetic operations depends upon O and V , which vary depending on the molecule and desired quality of the simulation, but a typical range is $1 \leq V/O \leq 100$. To provide concrete comparisons, we set O to 10 and V to 100, giving the numerical values in Table 1. This is representative of calculations of modest size that could be done on a workstation and will be used throughout this paper unless otherwise noted.

In order to focus on the effects of the optimization algorithms, we eliminate the binary terms from the input equations and consider only the ternary and higher terms, as described in the introduction. The optimal cost of the binary terms of each equation can be seen in the last column of Table 1.

Table 2 illustrates the results obtained by optimizing the five equations with the algorithms described previously. Exhaustive search results are shown only for the T1

Table 1. Characteristics of the fully unfactorized input equations used in this experiment

| Equation | Number of terms | Operation count (no optimization) | Operation count (with single term optimization only) | Operation count of optimal binary terms |
|----------|-----------------|-----------------------------------|--|---|
| CCSD T1 | 14 | $1.78 \cdot 10^{10}$ | $3.11 \cdot 10^8$ | $2.24 \cdot 10^8$ |
| CCSD T2 | 31 | $4.90 \cdot 10^{13}$ | $3.58 \cdot 10^{10}$ | $2.27 \cdot 10^{10}$ |
| CCSDT T1 | 15 | $2.08 \cdot 10^{10}$ | $2.31 \cdot 10^9$ | $2.22 \cdot 10^9$ |
| CCSDT T2 | 37 | $6.13 \cdot 10^{13}$ | $3.00 \cdot 10^{11}$ | $2.45 \cdot 10^{11}$ |
| CCSDT T3 | 47 | $9.34 \cdot 10^{16}$ | $3.26 \cdot 10^{13}$ | $2.26 \cdot 10^{13}$ |

Table 2. The operation counts of expressions optimized by different algorithms

| Equation | Ternary+ operation count | | | | |
|----------|--------------------------|-------------------|---|----------------------|----------------------|
| | Single Term | Exhaustive | Timed Exhaustive with Tier-Based Partitioning | Direct Descent | Random Descent |
| CCSD T1 | $8.65 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ |
| CCSD T2 | $1.31 \cdot 10^{10}$ | – | $5.18 \cdot 10^9$ | $5.33 \cdot 10^9$ | $5.14 \cdot 10^9$ |
| CCSDT T1 | $8.65 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ | $4.73 \cdot 10^7$ |
| CCSDT T2 | $5.52 \cdot 10^{10}$ | – | $5.57 \cdot 10^9$ | $5.57 \cdot 10^9$ | $5.37 \cdot 10^9$ |
| CCSDT T3 | $9.94 \cdot 10^{12}$ | – | $9.80 \cdot 10^{11}$ | $9.17 \cdot 10^{11}$ | $9.17 \cdot 10^{11}$ |

Table 3. Ternary+ operation count for CCSD T2

| Optimized for V/O | Actual V/O | | | | | |
|----------------------|---|-------------------|-------------------|-------------------|-------------------|----------------------|
| | 1 | 2 | 3 | 5 | 10 | 100 |
| 2 | $7.41 \cdot 10^6$ | $4.52 \cdot 10^7$ | $1.46 \cdot 10^8$ | $7.20 \cdot 10^8$ | $7.47 \cdot 10^9$ | $4.41 \cdot 10^{13}$ |
| 5 | $1.05 \cdot 10^7$ | $5.88 \cdot 10^7$ | $1.73 \cdot 10^8$ | $7.08 \cdot 10^8$ | $5.14 \cdot 10^9$ | $4.67 \cdot 10^{12}$ |
| 10 | $1.07 \cdot 10^7$ | $5.95 \cdot 10^7$ | $1.74 \cdot 10^8$ | $7.13 \cdot 10^8$ | $5.15 \cdot 10^9$ | $4.67 \cdot 10^{12}$ |
| Opt. V/O | Leading terms of cost function in symbolic form | | | | | |
| 2 | $2O^3V^3 + 2O^4V^2 + 4OV^4 + 10O^2V^3 + 10O^3V^2 + 6O^4V + \dots$ | | | | | |
| 5 | $4O^3V^3 + 4O^4V^2 + 6O^2V^3 + 8O^3V^2 + 6O^4V + \dots$ | | | | | |
| 10 | $4O^3V^3 + 4O^4V^2 + 6O^2V^3 + 8O^3V^2 + 6O^4V + \dots$ | | | | | |

Table 4. Ternary+ operation count for CCSDT T2

| Optimized for V/O | Actual V/O | | | | | |
|----------------------|---|-------------------|-------------------|-------------------|-------------------|----------------------|
| | 1 | 2 | 3 | 5 | 10 | 100 |
| 2 | $7.85 \cdot 10^6$ | $4.77 \cdot 10^7$ | $1.54 \cdot 10^8$ | $7.52 \cdot 10^8$ | $7.70 \cdot 10^9$ | $4.43 \cdot 10^{13}$ |
| 5 | $9.88 \cdot 10^6$ | $5.89 \cdot 10^7$ | $1.78 \cdot 10^8$ | $7.55 \cdot 10^8$ | $5.63 \cdot 10^9$ | $5.28 \cdot 10^{12}$ |
| 10 | $1.09 \cdot 10^7$ | $6.13 \cdot 10^7$ | $1.80 \cdot 10^8$ | $7.40 \cdot 10^8$ | $5.37 \cdot 10^9$ | $4.88 \cdot 10^{12}$ |
| Opt. V/O | Leading terms of cost function in symbolic form | | | | | |
| 2 | $2O^3V^3 + 2O^4V^2 + 4OV^4 + 12O^2V^3 + 12O^3V^2 + 6O^4V + \dots$ | | | | | |
| 5 | $4O^3V^3 + 2O^4V^2 + 12O^2V^3 + 16O^3V^2 + 6O^4V + \dots$ | | | | | |
| 10 | $4O^3V^3 + 4O^4V^2 + 8O^2V^3 + 10O^3V^2 + 6O^4V + \dots$ | | | | | |

equations because they are the only ones small enough for this approach to be feasible. All heuristic algorithms find the optimal factorization in small cases (i.e., the T1 equations), and in the other cases produce very similar results, which have less than one percent differences. The direct descent results illustrate its tendency to get stuck in local minima and not find an optimal factorization. Random descent sometimes offers improvement and consistently performs the best of all of the algorithms described.

To examine the behavior of the random descent search in more detail, we examined the effect of varying the V/O ratio. Changing this ratio will change the actual costs of each term and may even change the optimal factorization. Tables 3 and 4 illustrates these effects. Operation counts are shown for V/O ratios ranging from 1 to 10, for factorizations of the CCSD T2 and CCSDT T2 equations that have been optimized explicitly for V/O ratios of 2, 5, and 10. This is representative of how the results of this optimization are likely to be used in practice: code will be automatically generated for selected values of V/O spanning the range of interest, and at runtime, the best available version will be selected based on the actual V and O for the molecule under study. The results clearly illustrate the value of tailoring the factorization to the V/O of interest. For example, using the factorization that is optimal for $V/O = 2$ for a calculation in which $V/O = 100$ yields an operation count that is an order of magnitude larger than in the more optimal case using the algorithm optimized for $V/O = 10$. Conversely, using an algorithm optimized for the large $V/O = 10$ ratio is clearly not optimal if V/O is

small, e.g. 1. In Table 3 the symbolic operation count is given for the CCSD T2 equation for the various factorization solutions. For the larger V/O ratios of 5 and 10, no terms are used that scale as V^4 or OV^4 , as are present (and underlined) in the factorization scheme optimized for $V/O = 2$. Instead the algorithm prefers to use more terms that scale as O^3V^3 (four instead of two), and there is an interesting tradeoff between terms that formally scale as N^5 vs. N^6 , where N indicates O or V .

Similar conclusions can be drawn from Table 4. Comparing the optimized algorithms for $V/O = 2$ and $V/O = 10$ for various values of actual V/O ratios we clearly see that for small V/O values, the $V/O = 10$ solution is about 30% more costly than the $V/O = 2$ solution, while at the other end of the spectrum the $V/O = 10$ solution is about an order of magnitude more efficient. Similar tradeoffs as in case of the CCSD T2 equations are at work to determine the best overall factorization scheme, depending on the actual V and O values.

The present computer optimized factorization can be contrasted with current (hand-written) implementations of coupled cluster methods. In traditional implementations, factorization is considered only at a symbolic level, trying to reduce the V exponent first, then the O exponent, then the factor in front of the cost term; typically, little attention is paid to terms beyond the highest order (in the sum of the O and V exponents). This approach doesn't fully consider the ratio of V/O , and the possibility that terms considered lower order might result in an operation count comparable to higher order terms with a different balance of O and V exponents. The equation parts of Tables 3 and 4 illustrates this idea, with the symbolic costs of the different factorizations found for different V/O ratios. Comparing, for example, the costs of the CCSD T2 equation factored for $V/O = 2$ and $V/O = 5$, we observe that N^6 terms have *higher* coefficients in $V/O = 5$ than in $V/O = 2$, while the OV^4 term has been entirely eliminated from $V/O = 5$. The larger V space means that it is more cost-effective to evaluate more N^6 terms with a lower V exponent than the N^5 term OV^4 . A similar cross-over occurs in the CCSDT T2 equations between $V/O = 2$ and $V/O = 5$. Moreover, changes in term coefficients take place in the CCSDT T2 equations where $V/O = 5$ and $V/O = 10$. In the $V/O = 10$ equation the coefficient of the O^4V^2 term is higher than that in the $V/O = 5$ equation, while it is vice versa for the O^2V^3 term. This is an important result because it runs counter to the intuition (and accepted practice) of most quantum chemists. Let us note that it is not only the ratio V/O which determines the optimal factorization, but also their individual sizes as there is a trade-off between overall N^5 and N^6 terms in achieving optimal performance.

Interestingly, the random descent algorithm not necessarily finds the optimal factorization. This can be seen from Table 4 where the algorithm optimized for $V/O = 10$ performs better for the actual $V/O = 5$ ratio than the algorithm that was explicitly optimized for this case. This shows that there is still some room for improvement.

Table 5. Percentage of ternary+ operation count

| V/O | 1 | 2 | 3 | 5 | 10 | 100 |
|-------------------|-------|-------|-------|-------|-------|------|
| %ternary+ (CCSD) | 51.72 | 42.50 | 36.98 | 30.47 | 18.47 | 2.26 |
| %ternary+ (CCSDT) | 18.94 | 13.57 | 10.45 | 7.08 | 3.87 | 0.33 |

To put these results, obtained for the ternary and higher terms only, in proper perspective, Table 5 shows the percentage of the total computational cost due to the ternary and higher terms as a function of the V/O ratio. It is seen that the ternary+ cost is a sizable fraction of the overall calculation for lower V/O ratios, but it can rapidly become a rather small fraction of the overall cost if V/O is large, in particular for CCSDT. In such cases single term optimization might suffice. However it must be noted that these percentages do not tell the entire story. On modern hierarchical memory systems, the binary terms can typically be implemented with significantly greater efficiency than the ternary+ terms, so that a simple operation count underestimates their importance to the overall computation. Further, the present optimization scheme for factorization is quite generally applicable, and the efficiency gains can be expected to be very relevant for computational schemes that are not dominated by the binary terms. The factorization of more complicated, yet efficient theoretical models in quantum chemistry will be explored in our future work.

5 Conclusions

In this paper we presented heuristic and exhaustive search algorithms for operation minimization of complex tensor expressions occurring for example in quantum chemistry. It has been demonstrated that optimal factorization depends on the precise sizes of the index ranges of the tensors involved, and therefore very different computer implementations will be optimal for different size problems. We found that the random descent algorithm works best, although this search algorithm can be expensive for complicated cases. The time-limited exhaustive search with tier-based partitioning algorithm is often a cost-effective alternative; in addition, we believe this can provide a suitable starting point that can be subject to further optimizations.

Acknowledgments. This work has been supported in part by the U.S. National Science Foundation, the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory (ORNL), and by a Discovery grant from the Natural Sciences and Engineering Research Council of Canada. ORNL is managed by UT-Battelle, LLC for the US Dept. of Energy under contract DE-AC-05-00OR22725.

References

1. 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, and A. Sibiryakov. Synthesis of high-performance parallel programs for a class of ab initio quantum chemistry models. *Proceedings of the IEEE*, 93(2):276–292, February 2005.
2. G. Baumgartner, D.E. Bernholdt, D. Cociorva, R. Harrison, S. Hirata, C. Lam, M. Nooijen, R. Pitzer, J. Ramanujam, and P. Sadayappan. A high-level approach to synthesis of high-performance codes for quantum chemistry. In *Proc. of Supercomputing 2002*, November 2002.
3. B. Buchberger, G. Collins, and R. Loos, editors. *Computer Algebra: Symbolic and Algebraic Computation*. Springer-Verlag, New York, 1983.

4. C.N. Fischer and R.J. LeBlanc Jr. *Crafting a Compiler*. Benjamin/Cummings, 1991.
5. A. Hartono, A. Sibiryakov, M. Nooijen, G. Baumgartner, D. Bernholdt, S. Hirata, C. Lam, R. Pitzer, J. Ramanujam, and P. Sadayappan. Automated operation minimization of tensor contraction expressions in electronic structure calculations. Technical Report OSU-CISRC-2/05-TR10, Computer Science and Engineering Department, The Ohio State University, 2005.
6. F. Kschischang, B. Frey, and H. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, February 2001.
7. C. Lam. *Performance Optimization of a Class of Loops Implementing Multi-Dimensional Integrals*. PhD thesis, The Ohio State University, Columbus, OH, August 1999.
8. C. Lam, P. Sadayappan, and R. Wenger. On optimizing a class of multi-dimensional loops with reductions for parallel execution. *Parallel Processing Letters*, 7(2):157–168, 1997.
9. E.K. Miller. Solving bigger problems by decreasing the operation count and increasing computation bandwidth. *Proceedings of the IEEE*, 79(10):1493–1504, October 1991.
10. G.E. Scuseria, C.L. Janssen, and H.F. Schaefer III. An efficient reformulation of the closed-shell coupled cluster single and double excitation (CCSD) equations. *The Journal of Chemical Physics*, 89(12):7382–7387, 1988.
11. A. Sibiryakov. Operation Optimization of Tensor Contraction Expression. Master's thesis, The Ohio State University, Columbus, OH, August 2004.
12. J.F. Stanton, J. Gauss, J.D. Watts, and R.J. Bartlett. A direct product decomposition approach for symmetry exploitation in many-body methods. I. Energy calculations. *The Journal of Chemical Physics*, 94(6):4334–4345, 1991.
13. S. Winograd. *Arithmetic Complexity of Computations*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1980.