INTERNATIONAL JOURNAL OF QUANTUM CHEMISTRY, VOL. 40, 809–827 (1991)

Exact and Approximate Solutions to the One-Center McMurchie–Davidson Tree-Search Problem

BENNY G. JOHNSON, PETER M.W. GILL, AND JOHN A. POPLE
Department of Chemistry, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

Abstract

We have attempted to optimize the cost (the total number of floating-point operations required) of using the McMurchie–Davidson R_{NLMj} recurrence relation. Rigorous solutions of the tree-search problem inherent in the cost minimization are given for total angular momentum $L \leq 7$. For $L \geq 8$, the rigorous search algorithm is prohibitively expensive, and we propose an approximate algorithm that generates highly optimized trees. Cost comparisons demonstrate that the present scheme is consistently superior to two others currently in use.

1. Introduction

Many popular methods [1–7] for computation of two-electron repulsion integrals (ERIs) over Gaussian basis functions [8] are based upon the use of recurrence relations to construct the needed ERIs recursively from easily computed quantities involving the incomplete gamma function. The familiar scheme [1] due to McMurchie and Davidson (MD) involves certain one-center integrals related by a two-term recurrence relation, and some of the more recent algorithms for computing ERIs and their derivatives, such as the PRISM method of Gill and Pople [7], also make use of these one-center integrals and the associated recurrence relation.

Optimization of the use of the recurrence relation has been addressed by Gill, Head–Gordon, and Pople [4], as well as by Saunders [9], but prior to these, the topic had received little attention because the cost generally did not represent a major portion of the total cost of ERI evaluation by MD-based methods. However, in the PRISM algorithm of Gill and Pople, the cost of the MD recursion is significant in certain cases. For example, computing an entire class of $(d, sp|d, sp)$ with degree of contraction equal to two using the PRISM CCTTT pathway [7] requires 57,766 floating-point operations (FLOPs), of which 28,189 or 48.8% correspond to the MD work. It is therefore of critical importance to optimize this transformation step.

In this paper, we consider the problem of optimizing the cost, in terms of the number of FLOPs (adds, subtracts, multiplies, and divides), of producing the one-center integrals by examining the tree-search problem that is inherent in the application of the MD recurrence relation. We describe search algorithms that were

 CCC 0020-7608/91/060809-19$04.00

used to determine the minimum cost of transformation for maximum total angular momentum $L \leq 7$. For $L \geq 8$, the search time required by these algorithms for rigorous elucidation of the optimal cost is prohibitive, and we propose and discuss a new algorithm for finding efficient transformations, or "trees," to the needed one-center integrals. Finally, we compare our method with other proposed methods for generating efficient trees and demonstrate that our method produces the most FLOP-efficient transformations found so far.

2. Notation and Motivation

We introduce the integrals $[\mathbf{r}]^{(m)}$, defined [4] by

$$[\mathbf{r}]^{(m)} = 2^{(m+1)}\pi^{5/2}(\zeta\eta)^{-1}(\zeta + \eta)^{-1/2}\int_0^1 (\theta t)^{r+2m}H_{r_x}(R_x\theta t)H_{r_y}(R_y\theta t)H_{r_z}(R_z\theta t) \times \exp[-R^2\theta^2 t^2]\,dt, \tag{1}$$

where the H_{r_i} is the usual Hermite polynomial of degree r_i; ζ, η, and θ are given in terms of Gaussian exponents; and R_x, R_y, and R_z are the components of a position vector, with $R^2 = R_x^2 + R_y^2 + R_z^2$. The integer vector $\mathbf{r} = (r_x, r_y, r_z)$ gives the components of angular momentum of the integral in each Cartesian direction. The total angular momentum of the integral is defined as $r = r_x + r_y + r_z$. In the paper by MD [1], the integral $[\mathbf{r}]^{(m)}$ was designated by $R_{r_x r_y r_z m}$.

In computing ERIs or ERI derivatives having a total angular momentum L, it is in general necessary to compute the full set of $[\mathbf{r}]^{(0)}$ with $0 \leq r \leq L$. This is accomplished by first forming the integrals $[\mathbf{0}]^{(m)}$, $0 \leq m \leq L$, which involve the incomplete gamma function, and transforming them to $[\mathbf{r}]^{(0)}$ by the familiar MD identity [1, 4]:

$$[\mathbf{r}]^{(m)} = R_i[\mathbf{r} - \mathbf{1}_i]^{(m+1)} - (r_i - 1)[\mathbf{r} - \mathbf{2}_i]^{(m+1)}, \tag{2}$$

where i represents a Cartesian variable (x, y, or z), and $\mathbf{1}_i$ and $\mathbf{2}_i$ denote, respectively, one and two multiples of the unit vector in the i-direction. We shall refer to (2) as the MDRR (McMurchie–Davidson Recurrence Relation). The integrals $[\mathbf{r} - \mathbf{1}_i]^{(m+1)}$ and $[\mathbf{r} - \mathbf{2}_i]^{(m+1)}$ on the right-hand side of the MDRR will be referred to, respectively, as a parent and grandparent of $[\mathbf{r}]^{(m)}$.

It will be useful to introduce the following shorthand notations for denoting integrals and sets of integrals:

1. We use $[j\,k\,l]^{(m)}$ to indicate the single integral $[\mathbf{r}]^{(m)}$, where $\mathbf{r} = (j, k, l)$.
2. We use $\{j\,k\,l\}^{(m)}$ to indicate the set of six (or fewer) $[\mathbf{r}]^{(m)}$ whose components are generated by all permutations of j, k, l.

The application of the MDRR to produce an integral of higher angular momentum from one or two integrals of lower angular momentum involves either one

FLOP ($r_i = 1$), two FLOPS ($r_i = 2$), or three FLOPS ($r_i \geq 3$).* Note that, since the MDRR can be used to increase angular momentum in any Cartesian direction, any $[\mathbf{r}]^{(m)}$ is expressible in terms of $[\mathbf{r}]^{(m+1)}$ of lower angular momentum in up to three different ways, with each way possibly incurring a different cost. It is therefore evident that there are many different pathways possible to the $[\mathbf{r}]^{(0)}$ from the $[\mathbf{0}]^{(m)}$ and that finding the most economical pathway involves a tree-search problem.

It would be desirable, given a value of total angular momentum L, to find the tree for that angular momentum value, or L-tree, requiring the minimum number of FLOPS to produce the $[\mathbf{r}]^{(0)}$ from the $[\mathbf{0}]^{(m)}$. L-trees that are optimal in this sense are not difficult to find for $L \leq 4$; however, when $L \geq 5$, rigorous elucidation of the minimum-FLOPS L-tree is nontrivial, and for $L \geq 8$, it remains an unsolved problem.

3. Rigorously Optimal L-Trees ($L \leq 7$)

We proceed to derive rigorous solutions to the tree-search problems for $L \leq 7$. The derivation of these trees serves also to illustrate our notation and to suggest a motivation for the algorithms proposed later. Although our aim is to produce trees that begin with $[\mathbf{0}]^{(m)}$ integrals and lead to $[\mathbf{r}]^{(0)}$ integrals, using the MDRR to *increase* angular momentum, the discussion is facilitated if we consider the tree in reverse; that is, we begin with the set of desired $[\mathbf{r}]^{(0)}$ and *reduce* the angular momentum of these, producing additional integrals that must be reduced and proceeding in this way until all have been reduced to $[\mathbf{0}]^{(m)}$. Thus, we will speak of the *reduction cost* (1, 2, or 3 FLOPS) incurred by applying the MDRR to form a particular integral from its parent and grandparent.

The 2-Tree

We omit consideration of the 1-tree as it is trivial, and begin our discussion with the 2-tree. Figure 1 is a McMurchie–Davidson 2-tree, showing all of the equations required to compute the $[\mathbf{r}]^{(0)}$ from the $[\mathbf{0}]^{(m)}$. Figure 2 is a tabular representation of the same tree. Each row corresponds to a distinct vector $\mathbf{r}$, and each column, to an m value. The entry in the table for each integral $[\mathbf{r}]^{(m)}$ in the tree consists of a letter (x, y, or z), indicating the component of angular momentum reduced when applying the MDRR, and a parenthesized number (1, 2, or 3) indicating the associated reduction cost. We assume that the $[\mathbf{0}]^{(m)}$ have been pre-

*The stated costs are for applying the MDRR to the *uncontracted* integrals. Gill and Pople have derived analogs of the MDRR for contracted and half-contracted integrals [7], which are also implemented in the PRISM method. For the new recurrence relations, a constant is added to each of the individual costs given above, and the optimal tree is a function not only of the total angular momentum but also of the angular momenta of the individual basis functions. However, trees optimized for uncontracted integrals perform well for contracted integrals, and so, to avoid a proliferation of special cases, the appropriate tree for uncontracted integrals is always used. Therefore, in this paper, when we speak of trees, FLOP counts, etc., we are specifically referring to the case of uncontracted integrals.

$$[0\ 0\ 1]^{(0)} = R_z[0\ 0\ 0]^{(1)}$$
$$[0\ 1\ 0]^{(0)} = R_y[0\ 0\ 0]^{(1)}$$
$$[1\ 0\ 0]^{(0)} = R_x[0\ 0\ 0]^{(1)}$$

$$[0\ 0\ 1]^{(1)} = R_z[0\ 0\ 0]^{(2)}$$
$$[0\ 1\ 0]^{(1)} = R_y[0\ 0\ 0]^{(2)}$$
$$[1\ 0\ 0]^{(1)} = R_z[0\ 0\ 0]^{(2)}$$

$$[0\ 0\ 2]^{(0)} = R_z[0\ 0\ 1]^{(1)} - [0\ 0\ 0]^{(1)}$$
$$[0\ 1\ 1]^{(0)} = R_y[0\ 0\ 1]^{(1)}$$
$$[0\ 2\ 0]^{(0)} = R_y[0\ 1\ 0]^{(1)} - [0\ 0\ 0]^{(1)}$$
$$[1\ 0\ 1]^{(0)} = R_z[1\ 0\ 0]^{(1)}$$
$$[1\ 1\ 0]^{(0)} = R_x[0\ 1\ 0]^{(1)}$$
$$[2\ 0\ 0]^{(0)} = R_x[1\ 0\ 0]^{(1)} - [0\ 0\ 0]^{(1)}$$

Figure 1. A McMurchie–Davidson 2-tree (cost = 15 FLOPS).

[**r**]	m = 0	m = 1	m = 2
[0 0 0]	*	*	*
[0 0 1]	z(1)	z(1)	
[0 1 0]	y(1)	y(1)	
[1 0 0]	x(1)	x(1)	
[0 0 2]	z(2)		
[0 1 1]	y(1)		
[0 2 0]	y(2)		
[1 0 1]	z(1)		
[1 1 0]	x(1)		
[2 0 0]	x(2)		

Figure 2. Tabular representation of the McMurchie–Davidson 2-tree in Figure 1.

computed and that, therefore, there is no cost associated with the first row of the tree. The sum of all numbers in parentheses is the cost in FLOPs of the 2-tree. It follows that each element of the table in Figure 2 corresponds to an equation in Figure 1. Henceforth, we will avoid the more cumbersome equation listing and only the tabular form of trees will be given.

We are now ready to discuss the 2-tree shown. First of all, note that *any* integral belonging to a set $\{0\,0\,j\}^{(m)}$, $j \geq 1$, can be reduced in only one way (namely, by reducing the "j"), producing a parent in $\{0\,0\,j-1\}^{(m+1)}$ and grandparent in $\{0\,0\,j-2\}^{(m+1)}$ (if $j > 1$) which are of the same form. Hence, we are led to our first rule:

Rule 1: *All integrals $\{0\,0\,j\}^{(m)}$, $0 \leq j \leq L - m$, $0 \leq m \leq L$, are present in any L-tree.*

From now on we shall refer to these integrals as "axial" integrals because their angular momentum vectors lie along one of the Cartesian axes. We note that 12 of the 15 integrals in the 2-tree are axial.

The only nonaxial integrals are $\{0\,1\,1\}^{(0)}$. Each of these may be reduced in either of two ways (each costing one FLOP) and yields an axial parent that, by the previous conclusion, must already be part of the tree; hence, the reductions introduce no new integrals to the tree. Therefore, reducing either component will give the same contribution to the total cost, and our choices in Figure 1 are arbitrary. In that all $[\mathbf{r}]^{(0)}$ have been reduced to $[\mathbf{0}]^{(m)}$, the tree is now complete, and summing the individual reduction costs yields a total of 15 FLOPs. It is also evident from the preceding discussion that *any* 2-tree must cost exactly 15 FLOPs.

The 3-Tree

Figure 3 depicts a McMurchie–Davidson 3-tree. By Rule 1, all axial integrals appear in the tree. Furthermore, it is clear that the reductions of the $\{0\,1\,1\}^{(0)}$ discussed in the first example again have no consequence in determining the overall cost of the tree; these reductions are carried out as for the 2-tree.

The integrals $\{0\,1\,2\}^{(0)}$ are the first for which the reduction procedure is nontrivial. These key reductions are given in bold face in Figure 3. Reduction of the component of unit angular momentum costs one FLOP and involves an axial parent, whereas reduction of a component of two units costs two FLOPs and involves a parent in the set $\{0\,1\,1\}^{(1)}$ (which must itself be subsequently reduced) and an axial grandparent. It is clear that reduction of the unit component is preferred for two reasons: first, it invokes the least expensive special case of the MDRR, and, second, it introduces no new integrals (which, themselves, would have to be reduced) to the tree. This argument applies, in general, to the reduction of integrals in $\{0\,1\,j\}^{(m)}$, $j \geq 2$, since the parent involved in reducing the "1" is an axial integral and will therefore be in the tree. As an indication of this preferred manner of reduction, we refer to such integrals as "axial child" integrals and are led to our second rule:

Rule 2: *Axial child integrals are to be reduced in the component of unit angular momentum.*

It should be apparent that any reductions chosen for the remaining integral $[1\,1\,1]^{(0)}$ and its parent (which is a member of $\{0\,1\,1\}^{(1)}$), have the same effect on

	m			
[**r**]	0	1	2	3
[0 0 0]	*	*	*	*
[0 0 1]	z(1)	z(1)	z(1)	
[0 1 0]	y(1)	y(1)	y(1)	
[1 0 0]	x(1)	x(1)	x(1)	
[0 0 2]	z(2)	z(2)		
[0 1 1]	y(1)	y(1)		
[0 2 0]	y(2)	y(2)		
[1 0 1]	z(1)			
[1 1 0]	x(1)			
[2 0 0]	x(2)	x(2)		
[0 0 3]	z(3)			
[0 1 2]	**y**(1)			
[0 2 1]	**z**(1)			
[0 3 0]	y(3)			
[1 0 2]	**x**(1)			
[1 1 1]	x(1)			
[1 2 0]	**x**(1)			
[2 0 1]	**z**(1)			
[2 1 0]	**y**(1)			
[3 0 0]	x(3)			

Figure 3. Tabular representation of a McMurchie–Davidson 3-tree (cost = 41 FLOPS).

the overall cost; here, we arbitrarily reduce the x-component of $[1\,1\,1]^{(0)}$ and the y-component of its parent.

It is worthwhile to observe that, since the tree contains the minimum number of integrals possible and each is formed in the manner costing the fewest FLOPs, the cost of 41 FLOPs for the whole tree is the minimum possible for a 3-tree.

An important feature of the 3-tree is that it contains two "holes," i.e., Figure 3 gives a complete path from the $[\mathbf{0}]^{(m)}$ to the $[\mathbf{r}]^{(0)}$, but the integrals $[1\,0\,1]^{(1)}$ and $[1\,1\,0]^{(1)}$ are not needed as intermediates. It is true, in general, that not all the intermediate integrals are needed, as has been previously noted [4, 9], and eliminating unnecessary integrals must be a primary focus in developing algorithms to produce efficient trees.

The 4-Tree

For the 4-tree shown in Figure 4, all integrals can be reduced as in previous examples except for the six integrals $\{0\,2\,2\}^{(0)}$ and $\{1\,1\,2\}^{(0)}$. Applying the MDRR to

$[\mathbf{r}]$	m = 0	1	2	3	4
[0 0 0]	*	*	*	*	*
[0 0 1]	z(1)	z(1)	z(1)	z(1)	
[0 1 0]	y(1)	y(1)	y(1)	y(1)	
[1 0 0]	x(1)	x(1)	x(1)	x(1)	
[0 0 2]	z(2)	z(2)	z(2)		
[0 1 1]	y(1)	y(1)			
[0 2 0]	y(2)	y(2)	y(2)		
[1 0 1]	z(1)				
[1 1 0]	x(1)				
[2 0 0]	x(2)	x(2)	x(2)		
[0 0 3]	z(3)	z(3)			
[0 1 2]	y(1)	y(1)			
[0 2 1]	z(1)				
[0 3 0]	y(3)	y(3)			
[1 0 2]	x(1)				
[1 1 1]	x(1)				
[1 2 0]	x(1)	x(1)			
[2 0 1]	z(1)	z(1)			
[2 1 0]	y(1)				
[3 0 0]	x(3)	x(3)			
[0 0 4]	z(3)				
[0 1 3]	y(1)				
[0 2 2]	**y**(2)				
[0 3 1]	z(1)				
[0 4 0]	y(3)				
[1 0 3]	x(1)				
[1 1 2]	**x**(1)				
[1 2 1]	**z**(1)				
[1 3 0]	x(1)				
[2 0 2]	**z**(2)				
[2 1 1]	**y**(1)				
[2 2 0]	**x**(2)				
[3 0 1]	z(1)				
[3 1 0]	y(1)				
[4 0 0]	x(3)				

Figure 4. Tabular representation of a McMurchie–Davidson 4-tree (cost = 86 FLOPS).

any of these will yield a parent that is not needed to form axial and axial child integrals. We now observe that for $\{1\,1\,2\}^{(0)}$, reduction of a component having unit angular momentum involves a parent that is an axial child integral; hence, such a reduction choice can ultimately add no more than two FLOPS to the cost of the tree. Since the only other reduction choice (reducing the "2") costs two FLOPS, it can never be a more economical option. This argument applies to all integrals in the sets $\{1\,1\,j\}^{(m)}, j \geq 2$, and it is evident that an optimal tree must exist in which all such integrals that are present in the tree are reduced in a component of unit angular momentum. For this reason, we adopt our third rule ($\{1\,1\,j\}$ rule):

Rule 3: *Integrals in the sets $\{1\,1\,j\}^{(m)}$, $j \geq 2$ are never to be reduced in the component of angular momentum indicated by "j."*

It is important to note that this rule does not dictate which unit component *is* to be reduced—only that the "j" should *not* be reduced. Selection of the appropriate component for reduction depends upon the context of the reductions to be performed. We observe that in this case, without loss of generality, we may reduce an arbitrary unit component of one of the $\{1\,1\,2\}^{(0)}$. However, upon doing so, it is evident that one of the $\{0\,2\,2\}^{(0)}$ can now be formed in a manner costing the fewest FLOPS, while introducing no additional parents and grandparents to the tree (which would also have to be reduced). Indeed, we may formulate a general rule (common parent rule):

Rule 4: *If the parent and grandparent associated with a minimum-FLOPS reduction of an integral are present in a tree, the integral is to be reduced to the associated parent and grandparent integrals.*

In fact, each of the $\{1\,1\,2\}^{(0)}$ may be reduced in such a way that the reduction of a member of $\{0\,2\,2\}^{(0)}$ follows from the common parent rule. Because of the potential for common parents, we say that $\{1\,1\,2\}^{(0)}$ and $\{0\,2\,2\}^{(0)}$ are "coupled" sets of integrals. For this tree, we reduce $[1\,1\,2]^{(0)}$ by applying the MDRR to the x-component. It then follows from the common parent rule that the y-component of $[0\,2\,2]^{(0)}$ is to be reduced. These key reductions and others are given in bold in Figure 4. The total cost is 86 FLOPS, and we note that by the same argument as for the 3-tree it follows that 86 FLOPS is the minimum cost for a 4-tree.

Before proceeding, we observe that it is a general consequence of the axial child and $\{1\,1\,j\}$ rules (Rules 2 and 3) that no integrals $[1\,1\,1]^{(m)}$, $\{0\,1\,1\}^{(m)}$, $m > 0$, are present in any tree in which these rules are observed, because they can arise only from reductions that violate the rules (excepting, of course, one of the $\{0\,1\,1\}^{(1)}$ that must be present from the reduction of $[1\,1\,1]^{(0)}$. Therefore, for any tree, the first 10 rows are completely determined, and, henceforth, we consider only integrals with angular momentum greater than or equal to three.

The 5-Tree

Figure 5 shows the portion of a 5-tree consisting of the integrals with total angular momentum greater than or equal to four. The axial and axial child integrals may be immediately reduced, and the possibilities for reduction of integrals $\{1\,1\,2\}^{(0)}$ and $\{1\,1\,3\}^{(0)}$ are limited by the $\{1\,1\,j\}$ rule. The reductions of the integrals $\{0\,2\,2\}^{(0)}$, $\{1\,2\,2\}^{(0)}$, and $\{0\,2\,3\}^{(0)}$ are not determined by any previous rules. However, by an argument similar to that establishing the $\{1\,1\,j\}$ rule, we conclude that reduction of a "2" for integrals in the sets $\{0\,2\,j\}^{(m)}$, $j \geq 3$, will never prevent us from obtaining an optimal tree. Hence, we propose the $\{0\,2\,j\}$ rule:

Rule 5: *Integrals in the sets $\{0\,2\,j\}^{(m)}$, $j \geq 3$ are to be reduced in the component of angular momentum indicated by "2."*

This new rule governs the reduction of $\{0\,2\,3\}^{(0)}$ and allows for subsequent reduction of the $\{1\,1\,3\}^{(0)}$ by the common parent rule.

However, for the nine integrals $\{0\,2\,2\}^{(0)}$, $\{1\,1\,2\}^{(0)}$, and $\{1\,2\,2\}^{(0)}$ that remain unreduced, there are no integrals already in the tree due to previous reductions that could be used as parents. Simply proceeding by reducing the minimum com-

[r]	m = 0	m = 1
[0 0 4]	z(3)	z(3)
[0 1 3]	y(1)	y(1)
[0 2 2]	y(2)	
[0 3 1]	z(1)	z(1)
[0 4 0]	y(3)	y(3)
[1 0 3]	x(1)	x(1)
[1 1 2]	x(1)	
[1 2 1]	z(1)	**z**(1)
[1 3 0]	x(1)	x(1)
[2 0 2]	z(2)	z(2)
[2 1 1]	y(1)	
[2 2 0]	x(2)	**x**(2)
[3 0 1]	z(1)	z(1)
[3 1 0]	y(1)	
[4 0 0]	x(3)	x(3)
[0 0 5]	z(3)	
[0 1 4]	y(1)	
[0 2 3]	y(2)	
[0 3 2]	z(2)	
[0 4 1]	z(1)	
[0 5 0]	y(3)	
[1 0 4]	x(1)	
[1 1 3]	x(1)	
[1 2 2]	**z**(2)	
[1 3 1]	z(1)	
[1 4 0]	x(1)	
[2 0 3]	x(2)	
[2 1 2]	**y**(1)	
[2 2 1]	**z**(1)	
[2 3 0]	x(2)	
[3 0 2]	z(2)	
[3 1 1]	y(1)	
[3 2 0]	y(2)	
[4 0 1]	z(1)	
[4 1 0]	y(1)	
[5 0 0]	x(3)	

Figure 5. Tabular representation of a partial McMurchie–Davidson 5-tree, showing key reductions which lead to the minimum cost of 160 FLOPS.

ponent of angular momentum for each of these (thereby minimizing the individual reduction costs) cannot give a total cost less than 161 FLOPS, which turns out to be one FLOP greater than the optimal number. In fact, *for* L $\geq$ *5, it is no longer possible to produce optimal trees by any scheme that reduces all needed integrals by applying the MDRR in a direction costing the fewest FLOPS* and a more elaborate method is necessary.

We now turn our attention to the problem of determining the optimal reductions for the nine $[\mathbf{r}]^{(0)}$ integrals. Although Rules 1–5 do not dictate a definite reduction procedure for these, they have afforded a simplification of the problem that is considerable, as will be seen.

The possible parents and grandparents of the remaining $[\mathbf{r}]^{(0)}$ are as follows:

$$\{0\,2\,2\}^{(0)} \longrightarrow \{0\,1\,2\}^{(1)} + \{0\,0\,2\}^{(1)}$$

$$\{1\,1\,2\}^{(0)} \longrightarrow \{0\,1\,2\}^{(1)}$$

$$\{1\,2\,2\}^{(0)} \longrightarrow \{0\,2\,2\}^{(1)}, \{1\,1\,2\}^{(1)} + \{1\,0\,2\}^{(1)}.$$

Note that none of the $[\mathbf{r}]^{(0)}$ integrals already reduced by Rules 1–5 have parents or grandparents in any of the above sets (except for those with axial parents and/or grandparents), nor is there a possibility for the $[\mathbf{r}]^{(1)}$ integrals resulting from the reduction of the nine remaining $[\mathbf{r}]^{(0)}$ to have parents (except for axials) in common with any of the other needed $[\mathbf{r}]^{(1)}$. We see, therefore, that the rules have "decoupled" the problem of reducing the nine integrals from the reduction of the rest of the integrals in the tree. This decoupled problem can easily be solved by examining all possible reduction choices for the nine integrals by a simple computer program. Upon doing so, the optimal 5-tree cost is found to be 160 FLOPS, as mentioned earlier.

In Figure 5, the reductions in bold allow a tree costing 160 FLOPS. Note that a nonminimum component is reduced for $[1\,2\,2]^{(0)}$, but that this particular reduction produces a parent $[1\,2\,1]^{(1)}$ that is reducible in one FLOP, versus the two FLOPS required to reduce the parent $[0\,2\,2]^{(1)}$ associated with the minimum-FLOPS reduction, and involves a grandparent $[1\,2\,0]^{(1)}$ that is also used as a parent for $[2\,2\,0]^{(0)}$. Most importantly, when $[1\,2\,1]^{(1)}$ and $[2\,2\,0]^{(1)}$ are reduced as shown, a common parent results, yielding a net savings of one FLOP.

After solving the 5-tree problem, it is useful to consider the more general problem of how to optimally reduce any given subset of $[\mathbf{r}]^{(L-5)}$, $3 \le r \le 5$, to axial integrals (which are guaranteed to be present in any L-tree by Rule 1). Since the optimal procedure for reduction of all integrals in a tree with the same m value does not depend upon the manner in which the integrals for all lower m values were reduced, the solution of this problem would allow the remainder of any partial L-tree ($L > 5$) in which the $[\mathbf{r}]^{(0)}$ have all been reduced to $[\mathbf{r}]^{(L-5)}$ to be solved optimally. Again, our rules allow us to decompose the problem into two separate ones: reduction of needed integrals in the sets $\{0\,2\,2\}^{(L-5)}$, $\{1\,1\,2\}^{(L-5)}$, and $\{1\,2\,2\}^{(L-5)}$, for which Rules 1–5 are not immediately applicable, and the other $[\mathbf{r}]^{(L-5)}$, for which Rules 1–5 fully determine the reductions.

For the set $\{0\,2\,2\}^{(L-5)} \cup \{1\,1\,2\}^{(L-5)} \cup \{1\,2\,2\}^{(L-5)}$, there are 2^9 different subsets, one of which is the complete set considered in the solution of the 5-tree. Each of the 2^9 related problems is quite modest in complexity and trivially solvable by computer. Furthermore, each of these problems needs to be solved only once, with the results stored for future reference when considering L-trees with $L > 5$. Once these optimal reduction schemes are obtained, any such L-tree may be completed efficiently from the $[\mathbf{r}]^{(L-5)}$ by the algorithm given in Scheme 1.

SCHEME 1. Algorithm for optimal reduction of nonaxial $[\mathbf{r}]^{(L-5)}$, $3 \le r \le 5$.

Reduce $\{0\,1\,j\}^{(L-5)}$, $j \ge 2$, by axial child rule.
Reduce $\{0\,2\,3\}^{(L-5)}$ by $\{0\,2\,j\}$ rule, and resulting $\{0\,1\,3\}^{(L-4)}$ by axial child rule.
Reduce $\{1\,1\,3\}^{(L-5)}$ by $\{1\,1\,j\}$ rule, observing the common parent rule. If for any of these a common parent (from reduction of the appropriate $\{0\,2\,3\}^{(L-5)}$) is unavailable, reduce either unit component, and reduce the subsequent parent in $\{0\,1\,3\}^{(L-4)}$ by the axial child rule.
Obtain optimal reductions for needed $\{0\,2\,2\}^{(L-5)}$, $\{1\,1\,2\}^{(L-5)}$, and $\{1\,2\,2\}^{(L-5)}$, from list of pregenerated solutions.

The 6-Tree

In the spirit of the approach for the 5-tree, we seek an efficient method for optimally reducing a given set of $[\mathbf{r}]^{(L-6)}$, $3 \leq r \leq 6$, to axial integrals. Any axial child integrals and $\{0\,2\,j\}^{(L-6)}$, $j \geq 3$, may, of course, immediately be reduced. Note that any $\{1\,1\,4\}^{(L-6)}$ that are not reducible by the common parent rule may be reduced in either unit component, as the only integrals that could have a parent in common with these are the $\{0\,2\,4\}^{(L-6)}$. This leaves for consideration only the integrals shown in Figure 6, given with their possible nonaxial parents and grandparents.

All these reductions are coupled in that each set of $[\mathbf{r}]^{(L-6)}$ listed has at least one set of possible parents or grandparents in common with those of at least one other $[\mathbf{r}]^{(L-6)}$ set. Since there can be up to 22 of these $[\mathbf{r}]^{(L-6)}$, it is inefficient to determine the optimal reductions by examining all possibilities (up to $2^{12}3^{10}$) for this coupled set. However, we observe that there exist two coupled subsets of the $[\mathbf{r}]^{(L-6)}$ that are nearly disjoint, indicated by the groupings in Figure 6, and only the $\{1\,2\,3\}^{(L-6)}$ are members of both the coupled subsets. This indicates that, for each possible reduction of the needed $\{1\,2\,3\}^{(L-6)}$ (up to $3^6 = 729$), the larger problem of simultaneously considering the reductions of up to 16 remaining integrals is decomposed into two smaller problems, namely, the reduction of $\{0\,2\,2\}^{(L-6)}$, $\{1\,1\,2\}^{(L-6)}$, $\{1\,2\,2\}^{(L-6)}$, and $[2\,2\,2]^{(L-6)}$ and, separately, the reduction of $\{1\,1\,3\}^{(L-6)}$ and $\{0\,3\,3\}^{(L-6)}$.

For efficiency in solving any given set of $[\mathbf{r}]^{(L-6)}$, the solutions for all possible special cases of these two subproblems may be precomputed and then merely referenced, as was done for certain sets of $[\mathbf{r}]^{(L-5)}$. However, in this case, the possible presence of integrals in the sets $\{1\,1\,2\}^{(L-5)}$ and $\{1\,2\,2\}^{(L-5)}$ (from reduction of

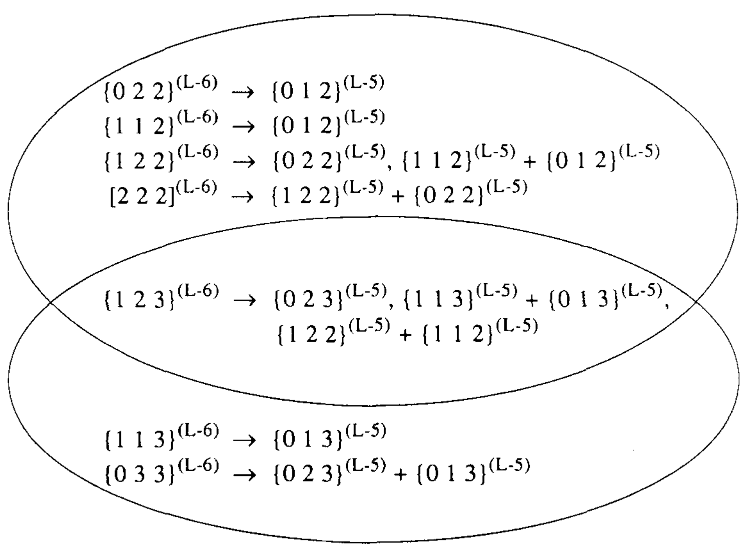

Figure 6. Coupled sets of $[\mathbf{r}]^{(L-6)}$ and their nonaxial parents and grandparents.

$\{1\,2\,3\}^{(L-6)}$) must be considered in the first subproblem, as they are potential parents or grandparents for some of the integrals to be reduced. Similarly, in the second subproblem, we must consider the possibility that integrals in $\{0\,1\,3\}^{(L-5)}$ and $\{0\,2\,3\}^{(L-5)}$ might be present and, hence, could be used as parents and grandparents. Also, integrals in $\{1\,1\,3\}^{(L-5)}$ must be considered, as they could potentially have parents in common with the other relevant $[\mathbf{r}]^{(L-5)}$. These extra considerations merely increase the number of different special cases of the two subproblems whose solutions are precomputed and do not affect the efficiency of the resulting tree.

After these "dictionaries" of optimal reductions have been compiled, the $[\mathbf{r}]^{(L-6)}$ reductions may be solved efficiently by the algorithm given in Scheme 2. Since all possible reductions are examined *only* for the needed $\{1\,2\,3\}^{(L-6)}$, the maximum number of different combinations is only 3^6, as opposed to $2^{12}3^{10}$.

In particular, this algorithm may be used to determine an optimal 6-tree by setting $L = 6$ and requiring the full set of nonaxial $[\mathbf{r}]^{(0)}$, $3 \leq r \leq 6$ (except for $[1\,1\,1]^{(0)}$, which will not be considered because of the $\{1\,1\,j\}$ rule). The cost of the tree is obtained as the cost of the reductions performed by the algorithm, plus two FLOPs for reducing $[1\,1\,1]^{(0)}$ and its parent, plus the cost of reduction for all axials, which can easily be shown for an L-tree to be $3/2\ (3L^2 - 3L + 2)$. This gives the minimum 6-tree cost as 268 FLOPs.

The 7-Tree

After reducing as many integrals as possible with Rules 1–5 for a given set of $[\mathbf{r}]^{(L-7)}$, $3 \leq r \leq 7$, we are still faced with over 40 $[\mathbf{r}]^{(L-7)}$ remaining; the reductions of which are, in general, highly coupled so as to preclude an efficient method similar to that of the 6-tree for determining their optimal reduction. However, since a highly efficient method is available for solving the reduction problem for any given set of $[\mathbf{r}]^{(L-6)}$, the approach in finding the optimal 7-tree cost was to determine a set of partial 7-trees, consisting of all the $[\mathbf{r}]^{(0)}$ and certain of the $[\mathbf{r}]^{(1)}$, $3 \leq r \leq 6$, which are viable (in the sense that every $[\mathbf{r}]^{(0)}$, $4 \leq$

SCHEME 2. Algorithm for optimal reduction of nonaxial $[\mathbf{r}]^{(L-6)}$, $3 \leq r \leq 6$.

Reduce $\{0\,1\,j\}^{(L-6)}$, $j \geq 2$, by axial child rule.
Reduce $\{0\,2\,j\}^{(L-6)}$, $j \geq 3$, by $\{0\,2\,j\}$ rule, and resulting $\{0\,1\,j\}^{(L-5)}$ by axial child rule.
Reduce $\{1\,1\,4\}^{(L-6)}$ by $\{1\,1\,j\}$ rule, observing the common parent rule. If for any of these a common parent (from reduction of the appropriate $\{0\,2\,4\}^{(L-6)}$) is unavailable, reduce either unit component, and reduce the subsequent parent in $\{0\,1\,4\}^{(L-5)}$ by the axial child rule.
For each possible reduction of needed $\{1\,2\,3\}^{(L-6)}$
  Obtain optimal reductions for needed $\{0\,2\,2\}^{(L-6)}$, $\{1\,1\,2\}^{(L-6)}$, $\{1\,2\,2\}^{(L-6)}$, and $[2\,2\,2]^{(L-6)}$ in presence of needed $\{1\,1\,2\}^{(L-5)}$ and $\{1\,2\,2\}^{(L-5)}$ from list of pregenerated solutions.
  Obtain optimal reductions for needed $\{1\,1\,3\}^{(L-6)}$ and $\{0\,3\,3\}^{(L-6)}$ in presence of needed $\{0\,1\,3\}^{(L-5)}$, $\{0\,2\,3\}^{(L-5)}$, and $\{1\,1\,3\}^{(L-5)}$ from list of pregenerated solutions.
  Obtain cost of optimally reducing resulting $[\mathbf{r}]^{(L-5)}$ by algorithm in Scheme 1.
Next $\{1\,2\,3\}^{(L-6)}$ reduction
Select the $\{1\,2\,3\}^{(L-6)}$ reductions, along with the associated optimal reductions for other integrals, which yield the minimum sum of reduction costs for the $[\mathbf{r}]^{(L-6)}$ and associated $[\mathbf{r}]^{(L-5)}$.

$r \leq 7$, can be formed from the $[\mathbf{r}]^{(1)}$), such that at least one of the partial trees will yield an optimal tree if all $[\mathbf{r}]^{(1)}$ and subsequent intermediate integrals are reduced optimally. An optimal 7-tree was then found by completing each partial tree by the algorithm in Scheme 2 and noting the minimum cost obtained.

Because of the vast number of reduction choices available for the $[\mathbf{r}]^{(0)}$ in a 7-tree, it is evident that any feasible procedure for optimizing the 7-tree must consider only a minute fraction of all viable $[\mathbf{r}]^{(1)}$ sets. Fortunately, simplifications are possible that reduce the partial trees that must be examined to a manageable number.

First, note that by Rules 3–5, applied to $\{0\,2\,j\}^{(0)}$ and $\{1\,1\,j\}^{(0)}$, $3 \leq j \leq 5$, all $\{0\,1\,j\}^{(1)}$, $3 \leq j \leq 5$ will be in any such $[\mathbf{r}]^{(1)}$ set; hence, we need to consider these and the $[\mathbf{r}]^{(0)}$ associated with them no further. The problem of completing the $[\mathbf{r}]^{(1)}$ sets is simplified considerably if we consider each different angular momentum value ($r = 3, \ldots, 6$) *separately,* by generating for each angular momentum all sets of $[\mathbf{r}]^{(1)}$ that are "parent-viable," meaning that each $m = 0$ integral with angular momentum $r + 1$ has at least one parent in the set. These are easily generated by examining all reduction choices for the associated $[\mathbf{r}]^{(0)}$ and collecting all distinct sets of parents that result. For example, there are 18 $[\mathbf{r}]^{(0)}$ with angular momentum seven that remain after initial application of Rules 1–5 for which there are $2^6 3^{12}$ ($>10^7$) reduction choices, but only 45,605 distinct parent-viable sets.

After initially applying Rules 1–5 to the full set of $[\mathbf{r}]^{(0)}$ and relevant intermediate integrals generated, we note that it is not a restriction to arbitrarily choose the reduction of the integral having all three components equal with highest total angular momentum ($[2\,2\,2]^{(0)}$). By predetermining that the x-component be reduced, we require that $[1\,2\,2]^{(1)}$ be present in the parent-viable sets for angular momentum five and that $[0\,2\,2]^{(1)}$ be present in the sets for angular momentum four. This reduces the number of five-sets roughly by one-third and halves the number of four-sets. Less restrictively, the possibilities can be slightly limited further for angular momentum less than or equal to five by recognizing that some of the integrals must be used as grandparents as well as parents. For example, a reduction of $[2\,2\,3]^{(0)}$ demands that one of the three integrals $[0\,2\,3]^{(1)}$, $[2\,0\,3]^{(1)}$, and $[2\,2\,1]^{(1)}$ be present in any 7-tree. Therefore, any set of $[\mathbf{r}]^{(1)}$ with angular momentum five that does not contain at least one of these three need not be considered, even if it is parent-viable. Such "grandparent restrictions" allow approximately 10% of parent-viable sets for angular momentum less than or equal to five to be discarded.

Once all relevant parent-viable sets have been obtained, all sets of $[\mathbf{r}]^{(1)}$ differing by one unit of angular momentum are paired such that each $[\mathbf{r}]^{(0)}$ of appropriate angular momentum can be formed by the MDRR from integrals in the paired sets. As for the individual parent-viable sets, in certain cases it is possible to eliminate a viable pairing because of other considerations, hence, reducing the number of 7-trees that need ultimately be considered. Particularly, since any set of $[\mathbf{r}]^{(1)}$ with angular momentum three will contain only axial child integrals, which reduce immediately to axials by Rule 2 (and, hence, are not coupled with any other integrals other than trivially), it is evident that we need pair each parent-viable set with angular momentum four with only *one* set with angular momentum three, if we choose the pairing such that cost of completely reducing the $[\mathbf{r}]^{(0)}$ with angular momentum five (using the integrals in the paired sets as intermediates) is minimized.

We may reduce the number of viable pairings considered for sets of integrals with angular momentum five and six by noting that $[\mathbf{r}]^{(1)}$ integrals with angular momentum six are used only as parents in a 7-tree (never as grandparents). Therefore, any pairing may be discarded that contains any integrals with angular momentum six that are not necessary to form all $[\mathbf{r}]^{(0)}$ with angular momentum seven in the cheapest manner allowed by the available $[\mathbf{r}]^{(1)}$, as the unnecessary integrals can only add to the cost of the tree. This eliminates approximately 90% of viable pairings from further consideration, a substantial reduction.

Since reduction of an integral by the MDRR involves only integrals differing by one unit of angular momentum, no more intermediate pairings are necessary, and we may proceed directly to construct $[\mathbf{r}]^{(1)}$ sets, implied by the pairings already obtained, which are viable for constructing all $[\mathbf{r}]^{(0)}$. Note that any such set of $[\mathbf{r}]^{(1)}$ that contains an unnecessary integral may be discarded, as was done when pairing set of integrals of angular momentum five and six. This ultimately results in approximately 7.7×10^5 partial trees, each of which is completed optimally by the algorithm in Scheme 2 to yield a 7-tree. Since all restrictions we have introduced eliminate only partial trees that are nonviable, contain unnecessary integrals, or are trivially equivalent to another partial tree considered, the trees obtained having the minimum cost, which is 418 FLOPs, are optimal.

When the total angular momentum is greater than seven, the time required to rigorously determine the minimum L-tree cost by algorithms such as these is prohibitive. We now turn our attention to the development of an algorithm that generate trees for higher L values, which, although not proven optimal, are believed to be the most efficient yet known.

4. Near-Optimal L-Trees ($L \geq 8$)

A motivation for constructing an algorithm to produce near-optimal trees for high angular momentum values can be drawn from the fact that it is quite often the case that an optimal L-tree "contains" an optimal $(L - 1)$-tree, in that it is only necessary to remove some integrals from an optimal L-tree to reveal an optimal $(L - 1)$-tree. Furthermore, the optimal L-tree (for $L > 1$) is not unique—there are actually many distinct L-trees that have the minimum possible cost. For example, out of all the 7-trees examined in the exhaustive search described above, 298 were optimal. This suggests that if an efficient 8-tree is desired, an excellent starting point ("skeletal tree") is one of the optimal 7-trees. We require that all the integrals in the 7-tree be present in the 8-tree, and then complete the 8-tree by adding the $[\mathbf{r}]^{(0)}$ with $r = 8$ and selecting a scheme for performing all necessary reductions. This procedure may then be continued inductively, using the best $(L - 1)$-trees as skeletal trees for constructing L-trees.

It now remains only to specify a method for carrying out the remaining reductions that are necessary to complete the tree. Of course, Rules 1–5 may be used to determine as many reductions as possible, and once all integrals have been reduced to a set of $[\mathbf{r}]^{(L-6)}$, we may proceed according to the algorithm in Scheme 2. For the integrals that remain, there are three factors that must be taken into account when deciding their reductions:

1. The immediate cost (1, 2, or 3 FLOPS) of applying the MDRR to reduce an integral.
2. The cost of reducing any new integrals introduced to the tree by a reduction.
3. The effect of a current reduction on other, later reductions (such as by introducing integrals to the tree that would allow some subsequent reductions to follow by the common parent rule).

As a means of succinctly considering all these effects when faced with ambiguous reduction choices for a particular integral, we introduce the notion of a cost index for each nonzero component of an integral, which serves as an indication of the contribution of reduction of the component to the total cost of the tree. It is designed such that reduction of components with lower indices should be preferred to reduction of components with higher indices when striving to produce trees of minimum total cost.

Component cost indices were taken as the sum of three contributions, one for each of the three considerations above:

1. The cost of applying the MDRR to reduce the component.
2. The minimum over all possible reduction combinations of any new parent and/or grandparent introduced by the reduction, of the MDRR costs of reducing these plus the minimum cost of reducing any new integrals that would be added to the tree by the particular reduction choice of the new parent and/or grandparent.
3. The sum of the minimum cost index of all other needed integrals with the same value of m, which would result *if* the particular component under consideration were reduced. The cost index for an integral is defined as the minimum over all nonzero components of the sum of the first two contributions to the component cost index given above.

In defining the component cost index, we place more emphasis on the third consideration, since as L increases, the proportion of reductions that cost three FLOPS tends asymptotically to one, and, hence, minimization of the number of integrals in the tree through common parents becomes more and more important.

At this point, a simple example is useful to illustrate the computation of a component cost index. We consider again the nine integrals $\{0\,2\,2\}^{(0)}$, $\{1\,1\,2\}^{(0)}$, and $\{1\,2\,2\}^{(0)}$ discussed in the rigorous solution for $L = 5$ and compute the cost index of the y-component of $[1\,2\,2]^{(0)}$. Although the notion of a component cost index was not used in solving the 5-tree problem, this particular case demonstrates the index without the lengthiness involved in its computation for integrals in trees of higher angular momentum. The first contribution is two FLOPS, and since the reduction would introduce $[1\,1\,2]^{(1)}$ and $[1\,0\,2]^{(1)}$ to the tree, the second contribution is seen to be three FLOPS (one each for the reduction of the two $[\mathbf{r}]^{(1)}$, plus one for the reduction of the member of $\{0\,1\,2\}^{(2)}$ that must appear from the reduction of $[1\,1\,2]^{(1)}$). It now remains only to sum the cost indices of the other integrals. It is obvious that the integral cost indices of $[1\,1\,2]^{(0)}$ and $[2\,0\,2]^{(0)}$, which could be reduced by the common parent rule if the y-component of $[1\,2\,2]^{(0)}$ is reduced, are, respectively, one and two FLOPS. On the other hand, the indices for remaining members of $\{1\,1\,2\}^{(0)}$ and $\{0\,2\,2\}^{(0)}$, whose reductions are not coupled to the current one, are seen to be two and three FLOPS, respectively. This demonstrates the general effect of lowering of cost indices through the potential for common parents. Similarly, it is easily shown that the integral cost index of $[2\,1\,2]^{(0)}$ is three FLOPS and that of $[2\,2\,1]^{(0)}$ is four FLOPS, with the former being less, again, due to the potential for a common parent with $[1\,2\,2]^{(0)}$. This gives the third contribution as 20 FLOPS, and, hence, the component cost index is $2 + 3 + 20 = 25$ FLOPS. It is interesting to note that the cost index of the x-component of $[1\,2\,2]^{(0)}$, given the same set of needed $[\mathbf{r}]^{(0)}$, is 27 FLOPS; hence, the component cost index would indicate that a nonminimum component of angular momentum be reduced for $[1\,2\,2]^{(0)}$. As previously demonstrated, it is necessary that at least one of the $\{1\,2\,2\}^{(0)}$ be reduced in a nonminimum component in order to obtain an optimal 5-tree.

After the cost index has been used to guide the reduction of a set of $[\mathbf{r}]^{(m)}$ to a set of $[\mathbf{r}]^{(m+1)}$, we may sieve certain of the $[\mathbf{r}]^{(m+1)}$ from the tree, which will only increase the total cost, as was done when exhaustively searching the 7-tree. An integral $[\mathbf{r}]^{(m+1)}$ is considered unnecessary if its removal does not cause the minimum cost of computing the needed $[\mathbf{r}]^{(m)}$ from the remaining $[\mathbf{r}]^{(m+1)}$ to increase. Such integrals, if allowed to remain in the tree, can do nothing but introduce integrals by their reduction that may themselves be unnecessary and make spurious contributions to the cost indices of other integrals and their components. Once a complete tree is obtained, its minimum possible cost, given the intermediate $[\mathbf{r}]^{(m)}$, is trivially obtainable, as it is simply the sum over all individual integrals of the minimum cost of formation of each integral *from the integrals available in the tree.*

We are now ready to give an algorithm for finding efficient trees, which is outlined in Scheme 3. The algorithm was used to find trees for $8 \leq L \leq 16$. All best $(L - 1)$-trees were used as skeletal trees in searching for L-trees, starting with the 298 known optimal 7-trees as a basis for finding 8-trees.

5. Results and Discussion

Table I summarizes the results of the present study in terms of the best costs obtained for trees with $L \leq 16$. (Copies of our optimized trees, if desired, are available from the authors upon request.) Also listed are the corresponding costs for two other methods for generating trees: the sieve method of Gill, Head-Gordon, and Pople [4], and that which we term the "full" method, in which *all*

SCHEME 3. Algorithm for producing efficient McMurchie–Davidson L-trees.

Produce a skeletal L-tree from an efficient $(L - 1)$-tree.
For $m = 0, \ldots, L - 7$
Reduce as many $[\mathbf{r}]^{(m)}$ as possible by Rules 1–5.
For each of $[\mathbf{r}]^{(m)}$ that are needed but have not yet been reduced, select first by increasing maximum component of angular momentum, then by dictionary order.
Compute the cost index for each nonzero component of angular momentum.
Reduce the component of angular momentum having the minimum cost index. If the minimum cost index is not unique, apply the following criteria, in order, to determine which component with minimum index is to be reduced.

1. Select the component whose reduction introduces the fewest new integrals to the tree.
2. Select the maximum component of angular momentum.
3. If there are exactly two components with minimum cost index, select the component according to the cyclic scheme

$$(x \text{ and } y) \Rightarrow x, \qquad (y \text{ and } z) \Rightarrow y, \qquad (z \text{ and } x) \Rightarrow z.$$

4. If there are three components with identical cost indices, select the x-component.

Reduce as many $[\mathbf{r}]^{(m)}$ as possible by the common parent rule.
Next $[\mathbf{r}]^{(m)}$
Remove any $[\mathbf{r}]^{(m+1)}$ from the tree whose removal does not increase the minimum cost of forming the needed $[\mathbf{r}]^{(m)}$.
Next m
Reduce the needed $[\mathbf{r}]^{(L-6)}$ integrals optimally by the algorithm in Scheme 2.
Compute the minimum cost of transforming $[\mathbf{0}]^{(m)} \rightarrow [\mathbf{r}]^{(0)}$ using the intermediate $[\mathbf{r}]^{(m)}$.

TABLE I. FLOP-costs of McMurchie–Davidson L-trees obtained by various methods.

L	Present	Sieve[a]	Full[b]
1	3*	3*	3*
2	15*	15*	15*
3	41*	41*	43
4	86*	86*	95
5	160*	161	180
6	268*	272	312
7	418*	428	507
8	622	648	783
9	890	936	1161
10	1233	1302	1665
11	1668	1776	2322
12	2219	2358	3162
13	2866	3060	4218
14	3638	3924	5526
15	4554	4944	7125
16	5633	6135	9057

[a]The sieve method of Gill, Head-Gordon, and Pople [4].
[b]*All* $[\mathbf{r}]^{(m)}$, $0 \le r \le L - m$, $0 \le m \le L$ are formed as cheaply as possible.
*FLOP-cost is rigorously the smallest possible.

intermediate $[\mathbf{r}]^{(m)}$ are formed. An asterisk beside a cost indicates that the cost is optimal.

It should be noted that the L-tree costs by the present algorithms are always equal to, or smaller than, those by other algorithms. The failure of the other methods to perform as well as the present one can be attributed to violation of one or more of the three considerations given in the last section for selecting reductions. It is obvious that, since *all* integrals are generated in the "full" method, it will become inferior to the others as soon as the possibility of unnecessary integrals arises. As was previously noted, this happens for $L \geq 3$, and it is seen that the "full" method is the most expensive for $L \geq 3$.

Conversely, the sieve method [4] concentrates solely on removing unnecessary integrals. This is done by systematically eliminating from the complete tree any integral $[\mathbf{r}]^{(m)}$, $m > 0$, which is redundant in that every integral which could be formed from $[\mathbf{r}]^{(m)}$ could also be formed from some other integral in the tree. This process is repeated until no further integrals can be removed. Although many unnecessary integrals are removed by this procedure, no account is explicitly taken of reduction costs in determining *which* of the removable integrals should actually be eliminated. Hence, one would expect that the sieve method would produce L-trees that are far more economical than are the corresponding "full" trees (in that approximately one-third of the integrals in the "full" tree may be removed), but are not as efficient as trees produced by the present algorithm. Comparison of the costs in Table I shows that this is indeed the case.

Furthermore, improvements over the present scheme may be obtainable for high values of L by regarding the minimization of the number of integrals in the tree as the foremost concern, since, as previously noted, in the limit of large L, all integrals have the same reduction cost.

Finally, we note that, although total tree cost in FLOPS is the proper quantity to be optimized in most cases, this is not always the case. Certain computers are capable of performing more than one FLOP simultaneously. This suggests that, for these machines, the trees should be reoptimized using costs for the special cases of the MDRR of one ($r_i \leq 2$) or two ($r_i \geq 3$) when an add and a multiply may be "chained" together or of one everywhere when an add and two multiplies may be chained. Of course, the second case is simply equivalent to minimizing the number of integrals in the tree. To give a particular example, a 5-tree exists that contains fewer integrals than do known 160-FLOP 5-trees, yet it costs more than 160 FLOPS; hence, a tree that is a solution of one of the three optimization problems is not necessarily a solution of the other two. However, approaches similar to the ones used here in optimizing the total number of FLOPS should be useful in performing the new optimizations.

Acknowledgments

We thank Vic Saunders, whose discussion of the tree-search problem at the Methodology of the Evaluation of Integrals in LCAO Calculations workshop at the Argonne National Laboratory stimulated us to publish our work on this topic. This work was supported by the National Science Foundation under grant CHEM-8918623.

Bibliography

[1] L. E. McMurchie and E. R. Davidson, J. Comput. Phys. **26**, 218 (1978).
[2] S. Obara and A. Saika, J. Chem. Phys. **84**, 3963 (1986).
[3] M. Head-Gordon and J. A. Pople, J. Chem. Phys. **89**, 5777 (1988).
[4] P. M.W. Gill, M. Head-Gordon, and J. A. Pople, Int. J. Quantum Chem. **S23**, 269 (1989).
[5] P. M.W. Gill, M. Head-Gordon, and J. A. Pople, J. Phys. Chem. **94**, 5564 (1990).
[6] T. P. Hamilton and H. F. Schaefer III, Chem. Phys. **150**, 163 (1991).
[7] P. M.W. Gill and J. A. Pople, Int. J. Quantum Chem. **40**, 753 (1991).
[8] S. F. Boys, Proc. R. Soc. Lond. **A200**, 542 (1950).
[9] V. Saunders, The optimization of the recursion for the McMurchie–Davidson R(i, j, k) integrals. Presented at the Methodology of the Evaluation of Integrals in LCAO Calculations workshop, Argonne National Laboratory, August (1990).

Received November 20, 1990
Accepted for publication December 27, 1990