

WIS-CS-75-243

COMPUTER SCIENCES DEPARTMENT
University of Wisconsin
1210 West Dayton Street
Madison, Wisconsin 53706

March 13, 1975

ON PARSING AND COMPILING ARITHMETIC EXPRESSIONS
IN PARALLEL COMPUTATIONAL ENVIRONMENTS

by

Charles N. Fischer

Technical Report #243

MARCH 1975

ABSTRACT

The problem of parsing and compiling arithmetic expressions in parallel computational environments is considered. It is seen that the concept of Operator Precedence can be generalized to allow encodings of one or more arithmetic expressions to be transformed directly into encodings of their corresponding derivation trees. The algorithm which performs this transformation is compact, efficient (linear in both time and space), and highly concurrent. Further, it can be extended to compile arithmetic expressions directly into object code (in the form of quadruples). The extension preserves the compactness, efficiency (linearity) and highly concurrent nature of the original algorithm.

1. Introduction

With the advent of parallel computers such as the Illiac IV and CDC Star-100, the design and analysis of algorithms which utilize parallelism has become of great interest. Rather surprisingly, however, the problem of designing compilers to run on such parallel computers has received relatively little attention ([1],[3],[4]). In this paper we consider the problem of parsing and compiling arithmetic infix expressions in parallel environments. Such expressions are of interest in that they are an integral part of virtually all programming languages; in fact in some languages (such as APL) they represent the only non-trivial structure present. We suggest it is entirely plausible that on suitable parallel computers all the arithmetic expressions occurring in a program segment, or perhaps an entire program, might be parsed (or compiled) concurrently in an efficient manner.

First we introduce Arithmetic Infix Grammars (AIG's) as a means of describing the class of arithmetic infix expressions. Next it is seen that a number of arithmetic expressions can be parsed efficiently and with a high degree of concurrency. In particular, encodings of such expressions can be transformed directly into encodings of their corresponding derivation trees. This enables us to avoid the overhead, inherent in serial parsing techniques, of repeatedly finding and replacing single occurrences of various productions. Finally, this parsing technique is extended to allow arithmetic expressions to be compiled directly into object code (in the form of quadruples). This extension allows us to by-pass conventional parsing entirely and yet is still both efficient and highly concurrent.

2. Arithmetic Infix Grammars

The structure of an arithmetic infix expression is determined by the properties of its operators. Usually the set of operators

used is partitioned into a number of classes. Each operator class is distinguished by its priority (often called precedence) and its associativity (either right or left).

Thus a given Arithmetic Infix Language will be characterized by n pairwise disjoint operator classes OP_1, \dots, OP_n . All unary operators will be placed in OP_n . Operators in OP_{i+1} will have a higher priority than those in OP_i, OP_{i-1}, \dots , etc. The highest priority operators (the unary ones) are applied first, then the next highest, etc. Further, each operator class, OP_i , has an associativity, $ASSOC[i]$. -1 denotes left associativity; 1 denotes right associativity. By definition, $ASSOC[n] = 1$ (right). Finally, we shall assume a class $OP_0 \equiv \{\#\}$ is added. $\#$ is a null operator used as an endmarker and to separate adjacent arithmetic expressions. We assume $ASSOC[0] = -1$.

For example, an Arithmetic Infix Language having the operators $+, -, *, /, **, SQRT$ and the usual semantics would be characterized as follows:

$$OP_0 = \{\#\}, OP_1 = \{+,-\}, OP_2 = \{*,/\}, OP_3 = \{**\}, OP_4 = \{SQRT\}$$

$$ASSOC[0] = -1, ASSOC[1] = -1, ASSOC[2] = -1, ASSOC[3] = 1, ASSOC[4] = 1$$

We may now define an Arithmetic Infix Grammar which generates the Arithmetic Infix Language characterized by OP_0, \dots, OP_n and $ASSOC[0], \dots, ASSOC[n]$ as $G = (V, V_T, P, S)$.

$$V_T = \{ID, (,)\} \cup OP_0 \cup \dots \cup OP_n$$

$$V = V_T \cup \{S, N_0, \dots, N_{n+1}\}$$

P is the union of the following sets of productions:

- (1) $\{S \rightarrow \#N_i \# \mid 0 \leq i \leq n+2\}$
- (2) $\{N_n \rightarrow \theta N_i \mid \theta \in OP_n \text{ and } n \leq i \leq n+2\}$
- (3) $\{N_{n+1} \rightarrow (N_i) \mid 1 \leq i \leq n+2\}$
- (4) For $0 \leq j \leq n-1$:

If ASSOC[j] = -1

then $\{N_j \rightarrow N_i \circ N_k\} \in OP_j$ and $j \leq i \leq n+2$ and $j < k \leq n+2$
else $\{N_j \rightarrow N_i \circ N_k\} \notin OP_j$ and $j < i \leq n+2$ and $j \leq k \leq n+2$

Note that, by definition, $N_{n+2} \equiv ID$.

Theorem 1. If G is an AIG then G is unambiguous.

Proof: It is easy to establish that G is LR(1). ■

3. Parsing Arithmetic Infix Expressions

In general a parser serves two functions. It verifies that an input string is, in fact, in L(G) (that is, derivable from S) and, given that it is, it produces a parse of it.

If G is an AIG, then, as we shall see, it is especially simple to test whether an input string is in L(G). First let $U \subseteq V_T$ for some AIG. Then we can define Follow(U) = $\{bev_T | S \xrightarrow{+} \dots ab\dots and acU\}$. Follow(U) is the set of terminal symbols that may legitimately follow some symbol in U.

Let UN (the set of unary operators) = OP_n and BIN (the set of binary operators) = $OP_0 \cup \dots \cup OP_{n-1}$. It can easily be verified that the Follow sets listed below are correct.

Set: UN BIN $\{()\}$ $\{()\}$ {ID}
Follow: W W W W' W'

where $W \equiv UN \cup \{(\cdot, ID)\}$ and $W' \equiv BIN \cup \{()\}$.

Next, let acV_T and xev_T^* . Then COUNT(x,a) is a count of the number of occurrences of a in x. For example,

COUNT(aabaa,a) = 4 and COUNT(aabaa,b) = 1.

Finally, let xev_T^+ for G some AIG. Then x is Well Formed (with respect to G) iff:

(1) x is of the form #x'#

(2) $x = x_1x_2 \Rightarrow COUNT(x_1,() - COUNT(x_1,)) \geq 0$ and $x = x_3\#x_4 \Rightarrow COUNT(x_3\#,() - COUNT(x_3\#,)) = 0$ for $x_1, x_2, x_3, x_4 \in V_T^*$.

(3) $x = \dots ab\dots \Rightarrow beFollow(\{a\})$.

Note that condition (1) requires correct endmarkers and that condition (2) requires proper parenthesis structure. The following establishes the importance of Well Formedness:

Theorem 2. Let G be an AIG. Then $x \in L(G)$ iff x is Well Formed.

Proof: Clearly $x \in L(G) \Rightarrow x$ is Well Formed. Assume x is Well Formed and write x as $\#x_1\#x_2\#\dots\#x_m\#$ where # does not occur in x_1, \dots, x_m . Write x_1 as $\dots(x')\dots$ where x' contains no "(" or ")"; otherwise, let $x' = x_1$. Using Follow sets, x' must be of the form $(UN^*Z)(BIN UN^*Z)^*$ where $Z = \{ID, N_{n+1}\}$. Unary operators in x' can be reduced to obtain x'' which is of the form $Y(BIN Y)^*$ where $Y = \{ID, N_{n+1}, N_n\}$. A simple induction on the number of operators in x'' will establish that $N_i \xrightarrow{+} x''$ for $1 \leq i \leq n+2$ (recall $ID \equiv N_{n+2}$). Now if we had $x' = x_1$ then we have established that $N_i \xrightarrow{+} x_1$; otherwise we have reduced $x_1 = \dots(x')\dots$ to $\dots(N_i)\dots$. But $\dots(N_i)\dots$ can be reduced to $\dots N_{n+1}\dots$ and the above process is then repeated. Thus finally x_1 is reduced to N_i for $1 \leq i \leq n+2$. x_2, \dots, x_m can be reduced in like manner. Thus x can be reduced to $\#N_i\#N_j\#\dots\#N_q\#$ for $1 \leq i, j, \dots, q \leq n+2$. It is then easy to verify that $S \xrightarrow{+} \#N_i\#N_j\#\dots\#N_q\#$. ■

Note that Well Formedness can easily be tested in linear time. In fact conditions (1) and (2) require but a single scan of the input while condition (3) can be tested in a highly concurrent manner. We may now attack our main objective - that of parsing arithmetic infix expressions.

The order of application of operators in an expression depends

Theorem 3. Let some operator θ_i occur in a subtree rooted by the operator θ_j in the reduced derivation tree of some arithmetic infix expression. Then $PREC[i] > PREC[j]$.

Proof: Consider the corresponding (non-reduced) derivation tree. Since this tree is generated by some AIG, it must be the case that:

- (1) The parenthesis level of θ_i is \geq than that of θ_j .
- (2) If the parenthesis levels of θ_i and θ_j are equal then

$$CLASS(\theta_i) \geq CLASS(\theta_j)$$

(3) If $CLASS(\theta_i) = CLASS(\theta_j)$ then

- (a) If $ASSOC[CLASS(\theta_i)] = -1$ then θ_i is to the left of θ_j , that is, $i < j$
- (b) If $ASSOC[CLASS(\theta_i)] = 1$ then θ_i is to the right of θ_j , that is, $i > j$.

But these conditions ensure that (by the definition of $PREC$) that $PREC[i] > PREC[j]$. ■

The reduced derivation tree of an expression contains almost exactly the same structure as the corresponding non-reduced derivation tree (only non-terminals, endmarkers and parentheses are lost). Thus we will consider the problem of parsing an expression to be that of determining its reduced derivation tree. The following functions are key to such a determination.

Let Q be a vector of unique integers. Then define:

$$\begin{aligned} SR(i,Q) &= \{j \mid i < j < |Q| \text{ and} \\ &\quad \neg \exists k (i < k \leq j \text{ and } Q[k] < Q[i])\} \\ SL(i,Q) &= \{j \mid j < i \text{ and} \\ &\quad \neg \exists k (j \leq k < i \text{ and } Q[k] < Q[i])\} \end{aligned}$$

SR scans right from i , getting the largest set of consecutive integers having Q values $> Q[i]$. SL does the same, scanning left. We then define

$$L(i,Q) = \begin{cases} \text{If } SL(i,Q) = \emptyset \text{ then } 0 \\ \text{else } j \in SL(i,Q) \text{ st } \forall k \in SL(i,Q) \ Q[j] \leq Q[k] \end{cases}$$

$$R(i,Q) = \begin{cases} \text{If } SR(i,Q) = \emptyset \text{ then } 0 \\ \text{else } j \in SR(i,Q) \text{ st } \forall k \in SR(i,Q) \ Q[j] \leq Q[k] \end{cases}$$

L (R) simply finds that index in SL (SR) that has the smallest Q value, or returns 0. The import of these functions is established in the following.

Theorem 4. Let $INPUT \in L(G)$ for G some AIG and let $PREC(OPS)$ be as above.

Then (1) $SR(i,PREC(OPS))$ ($SL(i,PREC(OPS))$) = the indices of all the operators in the right (left) subtree of the i -th operator.
 (2) $R(i,PREC(OPS))$ ($L(i,PREC(OPS))$) = the index of the operator that roots the right (left) subtree of the i -th operator if it exists, otherwise 0.

Proof: (1) Consider SR (SL is analogous). All operators in θ_i 's right subtree are to the immediate right of θ_i and by Thm. 3 their indices will be included in $SR(i,PREC(OPS))$. Let θ_j ($j > i$) be the first operator in the reduced derivation tree that is to the right of θ_i but not in θ_i 's right subtree. It must be that $PREC[i] > PREC[j]$ (if θ_i were not a descendent of θ_j then θ_i and θ_j would have a common ancestor lying between them). Thus $SR(i,PREC(OPS))$ is exactly the set of indices of the operators in θ_i 's right subtree.

(2) If $SR(i,PREC(OPS)) = \emptyset$ then no operators are in θ_i 's right subtree and $R(i,PREC(OPS)) = 0$ as required. Otherwise, the root of θ_i 's right subtree is (by Thm. 3) that operator in the subtree with the minimum $PREC$ value. But this is just $R(i,PREC(OPS))$. A similar argument holds for $L(i,PREC(OPS))$. ■

Observe that the L and R functions are just what we need to determine the reduced derivation tree. $R(i,PREC(OPS))$, for

example, either points to the root of θ_j 's right subtree or tells us (by returning 0) that the first ID to θ_j 's right is the correct right subtree.

We are now ready to present an AIG parser.

Algorithm 5. (An AIG parser)

Input: One or more arithmetic infix expressions, separated by #'s, stored in INPUT.
Output: Encodings of the reduced derivation trees of the arithmetic infix expressions stored in EXPR_ROOTS, LEFT_SUBTREE, RIGHT_SUBTREE.

[1] If INPUT does not satisfy conditions (1), (2) and (3) of the Well Formedness requirement then stop and signal an error in INPUT.

[2] Calculate PREC(OPS).

[3] Let #_INDEX = the indices of all but the rightmost # in OPS (That is OPS[#_INDEX] = #, #, ..., #)

Let UNARY_INDEX = the indices of all unary operators in OPS

LEFT_SUBTREE = L(1, PREC(OPS)), ..., L(|OPS|-1, PREC(OPS))

LEFT_SUBTREE[UNARY_INDEX] = -1

RIGHT_SUBTREE = R(1, PREC(OPS)), ..., R(|OPS|-1, PREC(OPS))

EXPR_ROOTS = RIGHT_SUBTREE[#_INDEX]

The encodings of the reduced derivation trees are quite straightforward. Recall that INPUT is of the form $\#E_1\#E_2\#\dots\#E_r\#$ where E_1, \dots, E_r are individual arithmetic expressions. In particular, E_i is delimited by the i -th and $i+1$ st #'s. By construction the PREC value of the i -th # is $<$ the PREC values of all the operators in E_i and is $>$ the PREC value of the $i+1$ st #. This means that if the R value of the i -th # is j then if $j > 0$ then θ_j is the root of E_i 's reduced derivation tree and if $j = 0$ then E_i is a single ID to the immediate right of the i -th #.

To find the roots of the reduced derivation trees of E_1, \dots, E_r

we then simply look up EXPR_ROOTS[1], ..., EXPR_ROOTS[r]. Further, if θ_j is an operator \dagger # then if RIGHT_SUBTREE[j] \dagger 0 then it points to the operator that is the root of θ_j 's right subtree.

If RIGHT_SUBTREE[j] = 0 then θ_j 's right subtree is the first ID to the right of θ_j . So too, if θ_j is binary then LEFT_SUBTREE[j] either points to the operator that is the root of θ_j 's left subtree or indicates that θ_j 's left subtree is the first ID to θ_j 's left. LEFT_SUBTREE[j] is = -1 if θ_j is unary.

Reconsidering the example used earlier, we had

INPUT = #ID*(ID+ID**ID) - SQRT ID#SQRT((ID+ID)/ID)#

PREC(OPS) = -1,46,141,196,19,102,-7,104,255,158,-11.

If we apply Algorithm 5 we get:

LEFT_SUBTREE = 0,0,0,0,2,-1,5,-1,0,9

RIGHT_SUBTREE = 5,3,4,0,6,0,8,10,0,0

EXPR_ROOTS = 5,8

The reader may easily verify that the reduced derivation trees of the two expressions are correctly encoded.

Theorem 6. Let G be some AIG. Then

(1) If INPUT \in L(G) then Alg 5 produces a correct encoding of the reduced derivation tree(s) of the arithmetic expression(s) in INPUT. Further, if the L and R functions are evaluated in $O(|INPUT|)$ time and space then Alg 5 will execute in $O(|INPUT|)$ time and space.

(2) If INPUT \notin L(G) then Alg 5 will signal on error and stop.

Proof: (1) Correctness: Follows from Thm 2 and Thm 4.

Linearity: As noted earlier steps [1] and [2] can be done in linear time and space and by assumption so can step [3]. Follows from Thm 2.

At this point we reemphasize the fact that Alg 5 transforms an encoding of the input expressions (the PREC vector) directly into an encoding of the corresponding reduced derivation trees (the EXPR_ROOTS, LEFT_SUBTREE, and RIGHT_SUBTREE vectors). The role of the AIG is purely descriptive. Thus the usual overhead of finding and replacing occurrences of productions is avoided.

4. Calculating the L and R Functions

We now consider the problem of calculating L and R. For convenience let $\bar{L}(Q)$ denote $L(1, Q), \dots, L(|Q|-1, Q)$ and similarly for $\bar{R}(Q)$. Clearly $\bar{L}(Q)$ and $\bar{R}(Q)$ are closely related, in fact $\bar{L}(Q) = 0, (|Q| - \bar{R}REV(Q[|Q|-1, \dots, 1])) \text{ Mod } |Q|$ where $\bar{V}REV$ denotes V reversed.

Note that the Mod function is used solely to insure that the integers in $\bar{L}(Q)$ are in the range $|Q|-1$ to 0.

Thus we need consider only methods of calculating $\bar{R}(Q)$. One method is the following*.

Algorithm 7

Input: A vector Q of unique integers
Output: $\bar{R}(Q)$ stored in RESULT

[1] Stack $(-\infty, \infty)$ onto an empty push down stack where doublets are of the form

(STK_POINTER, STK_Q)
 POINTER + 1; Q[|Q|] + $-\infty$
 Do while POINTER \leq |Q| :

If STK_Q(top stack element) \leq Q[POINTER]
 then stack (POINTER, Q[POINTER])
 POINTER + POINTER + 1
 PREV + 0

*this algorithm was originally suggested by John Hopcroft

else RESULT[STK_POINTER(top stack element)] + PREV
 PREV + STK_POINTER(top stack element)
 pop off the top doublet on the stack

Theorem 8. Given input Q Alg 7

- (a) Computes $\bar{R}(Q)$ correctly,
- (b) In time and space bounded by $O(|Q|)$.

Proof: (a) When POINTER = r for $1 \leq r < |Q|$, $(r, Q[r])$ is stacked (after perhaps popping a number of stack elements).

$(r, Q[r])$ remains on the stack until POINTER = j where $Q[j] < Q[r]$. That is, j is the first index not in $SR(r, Q)$. Thus $r+1, \dots, j-1$ are the elements of $SR(r, Q)$. Each has already been considered, including that value k ($r+1 \leq k \leq j-1$) having the minimum Q value (if $SR(r, Q) \neq \emptyset$). By definition $k = R(r, Q)$. Further, it will be stacked immediately above $(r, Q[r])$ since $r+1, \dots, k-1$ have Q values larger than $Q[k]$. Thus when $(r, Q[r])$ is popped, $PREV = k = R(r, Q)$. Also, if $j = r+1$, then $SR(r, Q) = \emptyset$ and $PREV = 0 = R(r, Q)$.

- (b) During each iteration of step [2] an element of vector Q is pushed or popped.

Further, each element of Q is pushed and popped at most once.

It may appear that the above algorithm is too "serial" to utilize parallel environments efficiently. However this need not be the case. Recall that in general the input to our parser will be of the form $\#E_1\#E_2\#\dots\#E_m\#$. For a given E_r ($1 \leq r \leq m$) all its L and R values will of necessity lie within E_r . Further, the R value of all but the rightmost # points to the root of the expression to its immediate right. The L values of #'s are not used and need not be computed. Thus the input may be divided into a number of independent segments $\#E_1\#\#E_2\#\#E_m\#$. Alg 7 may then be run concurrently on each of these segments. In fact this approach is similar to that taken by Zozel, et al [4] in performing a conventional Operator Precedence parse in parallel.

We now consider an algorithm which is "more parallel" in structure. Let a REPEAT b = the vector b, b, \dots, b of length a . Further, let B be a boolean vector and V, W, V' be arbitrary vectors all of the same length. Then V, W MASK $B = V'$ where for $1 \leq r \leq |V|$ $V'[r] = \text{if } B[r] = 1$ then $W[r]$ else $V[r]$. Thus (1,2,3,4), (5,6,7,8) MASK 0,1,1,0 = 1,6,7,4. (Note that REPEAT and MASK are actual STAR-100 instructions.)

Algorithm 9.

Input: Q a vector of unique integers and $LIMIT < |Q|$

Output: A vector RESULT and a bit vector HAVERESULT.

[1] $BIG \leftarrow \text{any integer} > \text{Max}(\text{Abs}(Q))$

$R_RANGE \leftarrow 1, \dots, |Q|-1$

$QI \leftarrow (R_RANGE + |Q| * Q[R_RANGE]) \text{ CONCAT}$

$(LIMIT \text{ REPEAT } (|Q| * -BIG))$

$BIGVECTOR \leftarrow MINVAL + (|Q|-1) \text{ REPEAT } (|Q| * BIG)$

$HAVERESULT \leftarrow (|Q|-1) \text{ REPEAT } 0$

[2] For POS from 1 to LIMIT do:

$HAVERESULT \leftarrow HAVERESULT \vee (QI[POS+R_RANGE]$

$< QI[R_RANGE])$

$COMPARE_VAL \leftarrow QI[POS+R_RANGE], BIGVECTOR$

MASK HAVERESULT

$MINVAL \leftarrow \text{Min}(MINVAL, COMPARE_VAL)$

[3] $RESULT \leftarrow MINVAL \text{ mod } |Q|$

Theorem 10. Given Q and $LIMIT$ as inputs to Alg 9,

(a) If $HAVERESULT[r] = 1$ then $RESULT[r] = R(r, Q)$.

(b) The algorithm is bounded in space by $O(|Q|)$ and in time by $O(LIMIT * |Q|)$.

Proof: (a) For each index r ($1 \leq r < |Q|$) we encode r and $Q[r]$ into $QI[r]$ in such a way that $QI[r] < QI[j]$ iff $Q[r] < Q[j]$. $MINVAL$ and $BIGVECTOR$ are initialized to an encoding of a very large Q value and a 0 index. In step [2], $HAVERESULT[r]$

is set to 1 the first time that $Q[r+POS] < Q[r]$. Thus $SR(r, Q) = \{r+1, \dots, r+POS-1\}$. Up to this point the minimum of $|Q| * BIG, QI[r+1], \dots, QI[r+POS-1]$ has been accumulated in $MINVAL[r]$. After $HAVERESULT[r]$ is set to 1, $MINVAL[r]$ is unchanged (because $MINVAL[r] < BIGVECTOR[r] = |Q| * BIG$). Thus if $SR(r, Q) = \emptyset$ $MINVAL[r] = |Q| * BIG$ and when this is decoded $RESULT[r] = 0$. Otherwise $MINVAL[r] = \text{min}(QI[r+1], \dots, QI[r+POS-1])$ (by definition) $QI[R(r, Q)]$. When this is decoded, $RESULT[r] = R(r, Q)$.

(b) Clearly each of the vectors used is $O(|Q|)$ in length. Thus each iteration of step [2] takes $O(|Q|)$ and a $O(LIMIT * |Q|)$ time bound follows. ■

We are still faced with the problem of choosing $LIMIT$ so that if the value of $R(r, Q)$ is needed $HAVERESULT[r] = 1$. One way (quite obviously) is to set $LIMIT = |Q|$. Does this mean Alg 9's time bound is really $O(|Q|^2)$? In practice not. Recall that the input in general will be of the form $\#E_1 \#E_2 \# \dots \#E_m \#$. If the number of operators in each E_r is $< k$ then we may set $LIMIT = k$ (since the R values of the operators cannot extend beyond the #'s).

Since in practice individual expressions are fairly short, we can set $LIMIT$ to some reasonably small fixed value (say 20) and a linear algorithm is obtained. Those very rare cases for which $LIMIT$ is too small can be handled by Alg 7 and linearity is still maintained.

In conclusion then, we may utilize either Alg 7 or Alg 9 to compute the L and R functions. In either case the result is a compact, efficient and highly concurrent parser for arithmetic infix expressions.

5. Generating Object Code for AIG's

In compilation, parsing is not an end unto itself. Rather, it is closely connected with another phase of compilation - code generation. Since the actual code generated by a compiler depends on the particular computer for which it is targeted, we shall deal

(3) Once a result is stored in a temporary, it is never changed since each temporary is assigned to exactly once.

Note that condition (3) points up a major weakness of our current method of allocating temporaries. Since each temporary is assigned to but once, we need as many temporaries as there are operators to evaluate an expression. Clearly a method of allocation which "reuses" temporaries is needed. One solution is as follows.

Assume that for each operator $\theta_j \neq \#$ we maintain a field RESULT_TEMP. If RESULT_TEMP[j] = k, then θ_j 's quadruple stores its result into T_k . For EXP_ROOT[i] (the root of the i-th expression in INPUT) we set RESULT_TEMP[EXP_ROOT[i]] = EXP_ROOT[i]. The values of RESULT_TEMP for the other operators in the i-th expression are obtained by iterating the following:

If RESULT_TEMP is known for θ_j then if θ_j 's right subtree is rooted by θ_k then RESULT_TEMP[k] = RESULT_TEMP[j]. Further, if θ_j is binary and θ_j 's left subtree is rooted by θ_p then if θ_j 's right subtree is an ID then RESULT_TEMP[p] = RESULT_TEMP[j] else RESULT_TEMP[p] = p.

Note that the value of RESULT_TEMP is passed, if possible, to the right son of an operator. If the right son is an ID then it is passed to the left son (if it is not an ID). By using this allocation scheme in the above example we obtain:

RESULT_TEMP = ?,2,2,2,5,5,?,8,8,8.

The corresponding quadruples are:

- L1: (*,ID₃,ID₄,T₂) L2: (+,ID₆,ID₇,T₈)
- L1+1: (+,ID₂,T₂,T₂) L2+1: (/ ,T₈,ID₈,T₈)
- L1+2: (*,ID₁,T₂,T₂) L2+2: (SQRT, ,T₈,T₈)
- L1+3: (SQRT, ,ID₅,T₅)
- L1+4: (-,T₂,T₅,T₅)

Note that 2 (rather than 5) and 1 (rather than 3) distinct temporaries are used. In general, it is easy to verify, that the number of different temporaries used is equal to the number of operators having only ID's as operands. While this is not always

the minimum number of temporaries required, it is usually a good approximation.

We are now ready to consider an algorithm which both parses and generates quadruples for arithmetic infix expressions. Let B be a boolean vector and V an arbitrary vector and assume $|B| = |V|$. Then B COMPRESS V is a vector composed of those elements of V whose corresponding B elements are 1. Thus 1,0,1,1 COMPRESS 5,6,7,8 = 5,7,8. (Note that COMPRESS is an actual STAR-100 instruction).

Algorithm 11. (An AIG parser and code generator)

Input: (a) m arithmetic infix expressions, separated by #'s, stored in INPUT

(b) FIRST_ADR, a vector of length m equal to the starting addresses of the m sets of quadruples which will be generated.

Output: Encodings of quadruples to evaluate the m expressions of INPUT stored in QUAD_ADR, RESULT_TEMP, LEFT_SUBTREE and RIGHT_SUBTREE.

[1] Use Alg 5 to parse INPUT

[2] (Calculate QUAD_ADR and RESULT_TEMP as follows)

RIGHT_SUBTREE_SIZE = |SR(1,PREC(OPS))|, ..., |SR(|OPS|-1,PREC(OPS))|

QUAD_ADR[EXPR_ROOTS] = FIRST_ADR +

RIGHT_SUBTREE_SIZE[#_INDEX] - 1

RESULT_TEMP = 1, ..., |OPS| - 1

OP_VECT = EXPR_ROOTS

Do while |OP_VECT| > 0

RIGHT_SONS = RIGHT_SUBTREE[OP_VECT]

LEFT_SONS = LEFT_SUBTREE[OP_VECT]

QUAD_ADR[(RIGHT_SONS > 0) COMPRESS RIGHT_SONS] =

QUAD_ADR[(RIGHT_SONS > 0) COMPRESS OP_VECT] - 1

```

QUAD_ADR[(LEFT_SONS > 0) COMPRESS LEFT_SONS] =
  QUAD_ADR[(LEFT_SONS > 0) COMPRESS OP_VECT] -
  RIGHT_SUBTREE_SIZE[(LEFT_SONS > 0) COMPRESS
    OP_VECT] - 1
LARGER_SONS = Max(LEFT_SONS, RIGHT_SONS)
RESULT_TEMP[(LARGER_SONS > 0) COMPRESS LARGER_SONS] =
  RESULT_TEMP[(LARGER_SONS > 0) COMPRESS OP_VECT]
BOTH_SONS = LEFT_SONS CONCAT RIGHT_SONS
OP_VECT = (BOTH_SONS > 0) COMPRESS BOTH_SONS

```

The encoding of the quadruples is very similar to that used in Alg 5. Associated with each operator $\theta_j \neq \#$ is a quadruple. Its location is $QUAD_ADR[j]$. If $RIGHT_SUBTREE[j] = 0$ then the quadruple's right argument is the first ID to the right of θ_j . If $RIGHT_SUBTREE[j] = k \neq 0$ then the right argument is T_p where $p = RESULT_TEMP[k]$. Similarly, if $LEFT_SUBTREE[j] = 0$ then the quadruple's left argument is the first ID to θ_j 's left and if $LEFT_SUBTREE[j] = k > 0$ then the argument is T_p where again $p = RESULT_TEMP[k]$. Further, if $LEFT_SUBTREE[j] = -1$ then θ_j is unary and thus has no left argument. The quadruple stores its result in $RESULT_TEMP[j]$.

Theorem 12. Let G be some AIG. Then

(1) If $INPUT \in L(G)$ and INPUT contains m arithmetic infix expressions then

(a) Alg 11 will produce encodings of m sets of quadruples. The i-th set of quadruples ($1 \leq i \leq m$) will correctly evaluate the i-th arithmetic expression and will start at $FIRST_ADR[i]$.

(b) If the L and R functions and $RIGHT_SUBTREE_SIZE$ are evaluated in $O(|INPUT|)$ time and space then Alg 11 will execute in $O(|INPUT|)$ time and space.

(2) If $INPUT \notin L(G)$ then Alg 11 will signal an error and stop.

Proof: (1a) By Thm 6(1), Alg 5 parses INPUT correctly and by Thm 4(1) $RIGHT_SUBTREE_SIZE$ contains correct values. We first establish the following 3 facts:

(i) Let $\theta_j \neq \#$ root a subtree containing k operators.

Then the k quadruples associated with these operators have addresses $QUAD_ADR[j] + 1 - k, \dots, QUAD_ADR[j]$. (This can be established via a simple induction on the height of θ_j 's subtree)

(ii) Let θ_j be as in (i). Then if θ_q is in the subtree rooted by θ_j then $RESULT_TEMP[q] = RESULT_TEMP[j]$ or p where θ_p is in θ_j 's subtree. (Again a simple induction on the height of θ_j 's subtree establishes this)

(iii) Let θ_r be any operator $\neq \#$. Then $RESULT_TEMP[r] = s$ where θ_s roots a subtree containing θ_r . (Initially, if $EXPR_ROOTS[i] = t$ for $1 \leq i \leq m$ then $RESULT_TEMP[t] = t$. Further, if θ_u is a son of θ_v then $RESULT_TEMP[u] = u$ or $RESULT_TEMP[v] = RESULT_TEMP[v]$).

Let θ_j be as in (i) and assume $QUAD_ADR[j] = r$. Then we will establish that if the quadruples at $r-k+1, \dots, r$ are executed then the value of the subtree rooted by θ_j will be computed correctly and stored in $RESULT_TEMP[j]$.

Consider first the case in which θ_j 's subtree is of height 1. Then all of θ_j 's arguments are ID's and the single quadruple at r correctly computes the expression. By induction, we now assume the above to be established for subtrees of height $\leq m$ and consider subtrees of height $m+1$. If θ_j has only one son, θ_p , that is an operator (the other son is an ID or θ_j is unary), it must be of height m and it contains $k-1$ operators. But then quadruples $r-k+1, \dots, r-1$ correctly compute θ_p 's subtree and store the result in $RESULT_QUAD[p]$. θ_j 's quadruple at r immediately uses the result (and perhaps an ID) to correctly compute θ_j 's subtree. If θ_j has two sons θ_p and θ_q which are operators then the height of each subtree is $\leq m$. Further, if θ_j 's subtree has t operators, then θ_p 's subtree has $k-t-1$. This means $QUAD_ADR[q] = r-1$ and $QUAD_ADR[p] = r-1-t$.

Thus executing quadruples $r - k + 1$ to $r - 1 - t$ correctly evaluates θ_p 's subtree and stores the result in $RESULT_TEMP[p]$ and executing quadruples $r - t$ to $r - 1$ correctly evaluates θ_q 's subtree and stores the result in $RESULT_TEMP[q]$. Now by construction $RESULT_TEMP[q] = RESULT_TEMP[j]$ and $RESULT_TEMP[p] = p \neq RESULT_TEMP[q]$ (by (iii)). Further, θ_p does not occur in θ_q 's subtree. Therefore by (ii), no operator in θ_q 's subtree has $RESULT_TEMP = p$. Thus if quadruples $r - k + 1$ to $r - 1$ are executed, $RESULT_TEMP[p]$ and $RESULT_TEMP[q]$ contain the values of θ_j 's left and right subtrees. Thus executing quadruples $r - k + 1$ to r correctly evaluates θ_j 's subtree.

Letting θ_j = the root of the i -th expression yields the desired result.

(1b) By Thm 6(1) step [1] is linear. By assumption the calculation of $RIGHT_SUBTREE_SIZE$ also is. All vectors used are bounded by $|OPS| < |INPUT|$. Further each iteration of the while loop is proportional to $|OP_VECT|$. But each operator $\neq \#$ in OPS appears in OP_VECT exactly once.

(2) Follows from Thm 6(2). ■

Note that $RIGHT_SUBTREE_SIZE$ can easily be calculated when $RIGHT_SUBTREE$ is. In fact, in Alg 7 when $RIGHT_SUBTREE[i] = RESULT[i]$ $= R(i,Q)$ is stored, $RIGHT_SUBTREE_SIZE[i] = |SR(i,Q)| = POINTER - i - 1$. In like manner, when Alg 9 sets $HAVERESULT[i] = 1$, $|SR(i,Q)| = POS - 1$. Thus the results of Thms 8 and 10 immediately apply to the calculation of $RIGHT_SUBTREE_SIZE$ as well as $RIGHT_SUBTREE$. Finally, it is important to observe that the effect of the while loop in Alg 11 is to examine the roots of all the individual expressions in INPUT, then all their sons, etc. This means that the while loop, as well as the rest of the algorithm, can be executed in a highly concurrent manner.

6. Conclusion

As we have seen, it is possible to develop compact, efficient and highly concurrent algorithms to parse and generate code for arithmetic expressions. The question of how effective such algorithms would be in actual parallel environments remains unanswered. Clearly, careful implementations and tests are called for. However preliminary tests are favorable.

APL may be used to model parallel vector computers such as the STAR-100. This is because in APL it is much faster to perform an operation on a vector than it is to perform the same operation iteratively on its components.

A slight variant of Alg 11 (using Alg 9 to compute the L and R functions) was coded in APL\360 as was a conventional serial SLR(1) parser and code generator. In all cases in which INPUT contained more than a minimum number of operators (5 to 7), Alg. 11 proved faster. As the number of operators in INPUT was increased, a limiting ratio of approximately 4 to 1 in execution times was quickly reached.

These tests, of course, are by no means conclusive but they do suggest that the above algorithms could be most competitive with serial techniques on suitable parallel computers.

References

1. Fischer, C. "On Parsing Context Free Languages in Parallel Environments". Ph.D. Thesis, Cornell University, Ithaca, N.Y., 1975.
2. Floyd, R. "Syntactic Analysis and Operator Precedence". JACM, 10:3, pp. 316-333, 1963.
3. Lincoln, N. "Parallel Programming Techniques for Compilers". SIGPLAN Notices, Vol. 5, No. 10, 1970.
4. Zozel, M. "A Parallel Approach to Compilation". ACM Symposium on the Principles of Programming Languages, pp. 59-70, 1973.