

Department of Applied Analysis and
Computer Science

Research Report CS-74-07

May 1974

Local vs Global Strategies for
Adaptive Quadrature

by

Michael A. Malcolm
R. Bruce Simpson

This work was partially supported by the National Research
Council of Canada.

ABSTRACT

A study has been made of two alternative subinterval selection strategies for adaptive quadrature. The more commonly used is termed a local acceptance criterion (as in SQUANK or CADRE), and the other is termed a global acceptance criterion (as in AIND). An efficient programming technique for the global alternative is given, and an analytic model for predicting the subinterval distributions selected by classes of adaptive algorithms is developed. To test the predictions of the model, the well known algorithm, SQUANK was reprogrammed to make it operate under the global criterion. Experiments with these routines were carried out comparing their actual performance with the predicted performance and with the performance of CADRE and AIND, and some conclusions are drawn.

Keywords: adaptive quadrature, Simpson's rule, SQUANK, AIND, CADRE.

Introduction

In the past five years there have been a number of significant developments in adaptive quadrature methods for automatic numerical integration. Following a lull after McKeeman's original paper in 1963 [6], several algorithms have been published [1], [4], [7], [8] which differ in their local quadrature rules, strategies for subdividing the interval of integration, and criteria for termination. Some of these algorithms have special features for detecting round-off error [4], [8], for detecting singular integrand behavior [1], or for returning an indication of the validity of results [1], [9]. Compared with one another empirically, these subroutines are found to require considerably different numbers of functions evaluations for the same integrands [7], [8], [9]. In general, there appears to have been little analysis of the performance of these algorithms with the notable exception of [11].

In this paper we attempt to understand some of these differences, focussing on the effects of the interval subdivision strategy and termination criterion.

1. Interval subdivision strategies

In order to focus on the interval subdivision process, we shall regard an adaptive quadrature scheme as an algorithm for processing a sequence of subintervals $\{I_n\}$. The main components of an adaptive quadrature algorithm are:

i) a local quadrature procedure for evaluating

$$Q(I_n) \approx \int_{I_n} f(x)dx$$

ii) a method for calculating $E(I_n)$, an estimate for

$$|Q(I_n) - \int_{I_n} f(x)dx|$$

iii) criteria for deciding which subinterval in the sequence $\{I_n\}$ to subdivide at each stage, and for deciding when to terminate.

Using a user-specified absolute error tolerance, ϵ , the objective is to produce a number, res , for which

$$(1.1) \quad |res - \int_a^b f(x)dx| \leq \epsilon.$$

We shall be particularly interested in adaptive Simpson's rules with components i) and ii) as used in SQUANK [4]. Using $S^{(m)}(I_n)$ to denote Simpson's rule applied m times to I_n , these can be expressed as

$$Q(I_n) = S^{(2)}(I_n) \text{ and } E(I_n) = |S^{(1)}(I_n) - S^{(2)}(I_n)|/15.$$

In this section, we will concentrate on component iii).

The intervals of the sequence I_n fall into three categories: I_n has already been subdivided and discarded ($n \in \underline{Dis}$), or has been accepted by the algorithm ($n \in \underline{Acc}$), or is pending further examination ($n \in \underline{Pen}$). A program implementing an adaptive quadrature algorithm must

store the data associated with $\{I_n | n \in \text{Pen}\}$ using some data structure, and can accumulate the contributions to res from I_n for $n \in \text{Acc}$ as soon as the acceptance is recognized. A discussion of the relationships between subdivision strategies and data structures for storing the Pen set has been published by Rice [10].

A common interval acceptance criterion, using $\ell(I_n)$ to denote the length of I_n , is

$$(1.2) \quad E(I_n) \leq \ell(I_n)\epsilon / (b-a).$$

An interval I_n is accepted if and only if it satisfies (1.2) (e.g., CADRE [1] and SQUANK [4]). Algorithms using (1.2) generally use a stack to store Pen and start with $I_1 = [a,b] \in \text{Pen}$; at each stage, I_n (the top of the stack) either satisfies (1.2) and is accepted, or is bisected and its right and left halves (in that order) are placed on the stack.

Criterion (1.2) has the features that:

- i) The decision is based entirely on information available from I_n .
- ii) If the error estimate $E(I_n)$ is in fact a bound for the error in $Q(I_n)$ for every $n \in \text{Acc}$ when the algorithm terminates, then the user's tolerance is guaranteed to be met (assuming exact arithmetic).

On the other hand, the local error criterion (1.2) decreases linearly with the interval length, and hence is most stringent as a tolerance, in regions where the adaptive process is working at subdivision the hardest.

Clearly feature ii) is a crucial one; however, i) may be relaxed. One alternative, then, is to retain as pending, all the intervals in $\{I_n\}$ which haven't been discarded, so that

$$(1.3) \quad \bigcup_{m \in \text{Pen}} I_m = \{x | a \leq x \leq b\}.$$

Then we can use

$$(1.4) \quad \sum_{m \in \text{Pen}} E(I_m) \leq \epsilon$$

as an acceptance criterion for the entire pending set. If at some stage of the process, (1.4) is not satisfied, we can determine $M \in \text{Pen}$ such that

$$(1.5) \quad E(I_M) \geq E(I_m), \text{ for all } m \in \text{Pen},$$

and bisect I_M , adding its left and right halves to the pending set and rechecking (1.4). We shall refer to criterion (1.4) as a global acceptance criterion; (1.4) and (1.5) form the strategy of the subroutine SQUAGE discussed in Section 3. The subroutine AIND [8], [9] uses this strategy except for a slight modification of (1.4) to account for an additional user-specified relative tolerance.

The intention of the global strategy (1.4) and (1.5) is to select subintervals so that the local errors are roughly equal in magnitude, rather than scaled by the length of the subintervals. Thus, algorithms using global strategies should work no harder on subintervals where the integrand is difficult to integrate, than on subintervals where it is

easy. Hence, a global criterion has the potential both for reducing the number of subintervals used and for generating a large list of data associated with the pending set. In this sense, it appears to be a space-time trade off. In the next section, we shall outline an analytic predictive technique for helping to quantify this trade off.

2. Prediction of subinterval distributions

In this section, we develop a technique for predicting the subinterval distributions resulting from the strategies of Section 1. We then examine a simple example to illustrate the expected reduction in the number of subintervals resulting from the global acceptance criterion, (1.4) instead of the local one (1.2). We conclude with a discussion of the difference in performance between these strategies when applied to functions with endpoint singularities.

Our discussion is based on the specific quadrature rules and estimates employed in SQUANK [4], and a modification of SQUANK (discussed in Section 3) to use the global acceptance criterion. However, it will be apparent that the technique has more general applicability. As mentioned in the preceding section, SQUANK uses

$$(2.1) \quad Q(I_n) = S^{(2)}(I_n),$$

and

$$(2.2) \quad E(I_n) = |S^{(1)}(I_n) - S^{(2)}(I_n)|/15.$$

This estimate is based on the first term of the asymptotic expansion of the error in a power series in h [5]. In fact

$$(2.3) \quad E(I_n) = \lambda(I_n)^5 |f^{iv}(c)|/(4^4 \cdot 180), \quad c \in I_n.$$

Moreover, $15E(I_n)$ is known to be a (sharp) bound on the error if $f^{iv}(x)$ has constant sign for $x \in I_n$ [12].

SQUANK employs the local interval acceptance criterion

$$(2.4) \quad E(I_n) \leq \lambda(I_n)\epsilon/(b-a)$$

((1.2) of Section 1). Suppose that for a given integrand, $f(x)$, SQUANK selects intervals with endpoints x_i :

$$a = x_0 < x_1 < x_2 < \dots < x_{L-1} < x_L = b.$$

Since each subinterval is acceptable under (2.4), we can conclude that for $i = 0, 1, \dots, L-1$,

$$(2.5) \quad (x_{i+1} - x_i)^5 |f^{iv}(c_i)|/(4^4 \cdot 180) = \theta_i (x_{i+1} - x_i)\epsilon/(b-a),$$

where $x_i < c_i < x_{i+1}$ and $0 \leq \theta_i < 1$. Assuming that $f^{(iv)}(c_i) \neq 0$, this can be written

$$(2.6) \quad x_{i+1} = x_i + h_i 4(180)^{1/4} / |f^{iv}(c_i)|^{1/4}; \quad x_0 = a,$$

for $h_i = (\theta_i \epsilon / (b-a))^{1/4}$, $i = 0, \dots, L-1$. We can regard (2.6) as a rather crude single-step method for solving the initial-value problem

$$(2.7) \quad dx(s)/ds = 4(180)^{1/4} / |f^{iv}(x(s))|^{1/4}; \quad x(0) = a,$$

using step sizes h_i . The variability of h_i due to the unknown fraction θ_i complicates the connection between (2.7) and (2.6). We shall assume for $i = 0, \dots, L-1$,

- (i) that $1/16 \leq \theta_i \leq 1$ on the hypothesis that if $\theta_i < 1/16$, a parent interval of $[x_i, x_{i+1}]$ would have been accepted, and
- (ii) that θ_i can be replaced by the mean of a uniform distribution over its range, i.e. $17/32$.

Hence, we replace h_i by

$$(2.8) \quad h = (17/32)^{1/4} (\epsilon/(b-a))^{1/4}.$$

If $x(s)$ is the solution of (2.7), then we might expect the approximation

$$(2.9) \quad x_i \approx x(ih), \quad i = 0, 1, \dots, L,$$

to improve as ϵ approaches zero, subject to the validity of assumptions (i) and (ii) above. In particular, defining s_i^* to be the parameter value for which $x(s_i^*) = b$, we will use

$$(2.10) \quad L_{\text{pred}}(f; a, b, \epsilon) = \lfloor s^*/h \rfloor = \lfloor (32(b-a)/17)^{1/4} s^*/\epsilon^{1/4} \rfloor$$

as a prediction for the number of subintervals selected by SQUANK. The predicted number of function evaluations required by SQUANK is given by $4 L_{\text{pred}}(f; a, b, \epsilon) + 1$.

For example, consider $f(x) = (1+c-x)^{-1}$ for positive parameter c and $a = 0$, $b = 1$. In this case, (2.7) is

$$(2.11) \quad \frac{dx(s)}{ds} = 4(15/2)^{1/4}(1+c-x(s))^{5/4}; \quad x(0) = 0,$$

with solution

$$(2.12) \quad (1+c-x(s))^{-1/4} - (1+c)^{-1/4} = (15/2)^{1/4}s.$$

Substituting $x(s^*) = b = 1$ gives

$$s_i^* = (2/15)^{1/4}(c^{-1/4} - (1+c)^{-1/4}),$$

and

$$(2.13) \quad L_{\text{pred}}(f; 0, 1, \epsilon) = \lfloor (64/255)^{1/4}(c^{-1/4} - (1+c)^{-1/4})\epsilon^{-1/4} \rfloor.$$

Experiments with this example are discussed in Section 4.

We now turn to predicting the subinterval distributions produced by the global acceptance criterion of Section 1. Consider the modification of SQUANK to be an algorithm using (2.1) and (2.2) and which accepts the entire pending set of subintervals if

$$(2.14) \quad \sum_{m \in \text{Pen}} E(I_m) \leq \epsilon$$

(using (1.4) and (1.5) of Section 1). Suppose for a given integrand $f(x)$, this algorithm selects subintervals with endpoints y_i :

$$a = y_0 < y_1 < \dots < y_{G-1} < y_G = b.$$

Then, from (2.3), we can see that the error estimate for $[y_i, y_{i+1}]$ can be expressed as

$$(y_{i+1} - y_i)^5 |f^{iv}(d_i)| / (4^5 \cdot 45), \quad y_i < d_i < y_{i+1},$$

and we can define $\theta_i \geq 0$ ($i = 0, \dots, G-1$) by

$$(2.15) \quad (y_{i+1} - y_i)^5 |f^{iv}(d_i)| / (4^5 \cdot 45) = \theta_i \epsilon.$$

Using (2.14), we can conclude that

$$(2.16) \quad \sum_{i=0}^{G-1} \theta_i \leq 1.$$

As previously, if $f^{iv}(d_i) \neq 0$, we can rewrite (2.15) as

$$(2.17) \quad y_{i+1} = y_i + h_i 4(45)^{1/5} / |f^{iv}(d_i)|^{1/5}; \quad y_0 = a,$$

where $h_i = (\theta_i \epsilon)^{1/5}$, ($i = 0, 1, \dots, G-1$). Equation (2.17) can be interpreted as a numerical method for solving

$$(2.18) \quad dy(t)/dt = 4(45)^{1/5} / |f^{iv}(y(t))|^{1/5}; \quad y(0) = a,$$

using variable step sizes h_i . We shall again make the crude assumption that

- (iii) θ_i can be replaced by some constant (average) value and, from (2.16), we select $1/G$ for this constant.

Hence, we replace h_i by

$$(2.19) \quad h = (\epsilon/G)^{1/5}.$$

If $y(t)$ is the solution of (2.18) then, we might expect the approximation

$$y_i \approx y(ih), \quad i = 0, 1, \dots, G,$$

to improve as ϵ approaches zero, and be subject to assumption particular, if we define t^* to be the parameter value for which $y(t^*) = b$,

then we can determine the predicted number of subintervals $G_{\text{pred}}(f; a, b, \epsilon)$ using

$$t_i^* = G_{\text{pred}} h = G_{\text{pred}} \epsilon^{1/5} / G_{\text{pred}}^{1/5}.$$

This gives

$$(2.20) \quad G_{\text{pred}}(f; a, b, \epsilon) = \lfloor t_i^{*5/4} / \epsilon^{1/4} \rfloor.$$

(Note that we have chosen the notation "L" and "Lpred" to indicate local strategy; "G" and "Gpred" indicates global strategy.)

Returning to the example $f(x) = (1+c-x)^{-1}$, $a = 0$, $b = 1$, we have, for (2.18)

$$dy(t)/dt = 4(15/8)^{1/5}(1+c-y(t)); y(0) = 0,$$

with solution

$$(2.21) \quad \log_e((1+c)/(1+c-y(t))) = 4(15/8)^{1/5}t$$

and

$$t_i^* = (8/15)^{1/5} \log_e((1+c)/c) / 4.$$

Using (2.20), we predict for the number of subintervals

$$(2.22) \quad G_{\text{pred}}(f; 0, 1, \epsilon) = \lfloor 2^{-7/4} 15^{-1/4} (\log_e((1+c)/c))^{5/4} / \epsilon^{1/4} \rfloor.$$

Comparing this with (2.13), we see that the local acceptance criterion yields a predicted number of subintervals which is $O(c^{-1/4})$ as c approaches zero, whereas the global one yields a predicted number of subintervals which is $O(|\log c|^{5/4})$.

These techniques predict a sharp difference between the results of using the local and global acceptance criteria on integrands with endpoint singularities. If we consider $f(x) = (1-x)^\alpha$; $a = 0$; $b = 1$, we find for the local acceptance criterion (2.4) that a finite number of subintervals is predicted only if $\alpha > 0$. That is, using

$$p_4(\alpha) = |\alpha(\alpha-1)(\alpha-2)(\alpha-3)|,$$

$$(2.23) \quad L_{\text{pred}}(f;0,1,\epsilon) = \lfloor (p_4(\alpha)/180\epsilon)^{1/4}/\alpha \rfloor.$$

However, the global strategy predicts a finite number of subintervals for $\alpha > -1$, i.e.,

$$(2.24) \quad G_{\text{pred}}(f;0,1,\epsilon) = \lfloor (5/4(1+\alpha))^{5/4} (p_4(\alpha)/45\epsilon)^{1/4} \rfloor.$$

Experiments confirming this difference are discussed in Section 4. The ability of adaptive schemes using the local acceptance criterion to handle $f(x) = (1-x)^\alpha$ for $\alpha > 0$ has also been observed by Rice [4], in which a convergence theorem for an adaptive trapezoidal rule is given.

3. Programming the global strategy

The global acceptance strategy (1.4) and (1.5) has been implemented in a program written in ANSI Standard Fortran. The technique used to store information associated with the pending set of intervals will be discussed in this section. The Fortran program is presented at the end of this section.

As discussed in Section 1, the global strategy requires that all intervals be kept in the pending set until the acceptance criterion

$$(3.1) \quad \sum_{m \in \underline{\text{Pen}}} E(I_m) \leq \epsilon$$

is satisfied. At each stage of the process prior to satisfying condition (3.1), the program must determine an $M \in \underline{\text{Pen}}$ such that

$$(3.2) \quad E(I_M) \geq E(I_m), \text{ for all } m \in \underline{\text{Pen}}.$$

This could be accomplished by searching the entire list of subintervals in $\underline{\text{Pen}}$ at each stage. However, this would require $n-1$ floating-point comparisons at each stage, where n is the number of subintervals currently in the list. If the list is kept sorted in decreasing order of the error estimates, then no comparisons are required for determining M , but more work is required for inserting new subintervals into the list. A bubble sort at each stage would require up to n floating-point comparisons for each new subinterval added to the list. (This is the technique used by AIND [8], [9].) If a binary search is used to determine where to insert the new subintervals, then only $\lceil \log_2 n \rceil$ comparisons are needed for each new subinterval, but up to n items in the list will need to be moved to different memory locations. Using N to denote the value of n when (3.1) is satisfied, we have a total of

$$\sum_{n=1}^{N-1} (n-1) = N^2/2 - 3N/2 + 1$$

comparisons required for the first of the above methods. The other techniques require similar (i.e. $O(N^2)$) amounts to work. (An estimate for N was given in (2.20).) For integrands for which N becomes large (say greater than 10^3), these comparisons (or memory fetches and stores) may require a major portion of the computation time.

Fortunately, a completely ordered list is not necessary for determining M in (3.2). It is sufficient to store the subintervals in a partially-ordered binary tree, sometimes called a "heap". Each node of the binary tree (illustrated in Figure 3.1) contains the information for one of the subintervals. The nodes are partially ordered in the sense that the error estimate E for each node is no less than that for each of its subnodes: i.e.,

$$E(I_{t_i}) \geq E(I_{t_{2i}}), \text{ and}$$

$$E(I_{t_i}) \geq E(I_{t_{2i+1}}), \quad i = 1, 2, \dots$$

Hence, at each stage M is simply t_1 , the root node (see Fig.3.1).

The interval subdivision process consists of removing and bisecting the root node subinterval and inserting the data for the two resulting subintervals into nodes of the tree. It is important that the new data be inserted into the tree so as to keep it balanced, i.e. the only vacant nodes occur at the lowest level. This is achieved by inserting the data for the newly created left-hand subinterval in the (just vacated) root node; and the data for the other new subinterval is placed in the left-most vacant node in the lowest level of the tree. (In Fig.3.1, this would




Fig.3.1 A balanced binary tree for $n = 12$

be t_{13} .) After each insertion, the new node must be repositioned in the tree (if necessary) to maintain partial ordering of the tree. Each new node may be moved up to $\lfloor \log_2 n \rfloor$ levels in the tree during this repositioning. The new root node may require up to $2\lfloor \log_2 n \rfloor$ comparisons as it moves down the tree; and the new node inserted at the lowest level in the tree may require at most $\lfloor \log_2 n \rfloor$ comparisons. This gives at most

$$\begin{aligned} & \sum_{i=2}^{N-1} (2\lfloor \log_2 i \rfloor + \lfloor \log_2 (i+1) \rfloor) \\ & \leq 2 \log_2(N-1)! + \log_2(N!/2) \\ & \leq 3 N \log_2 N \end{aligned}$$

total comparisons. The use of a partially-ordered binary tree has reduced the number of comparisons and memory references from $O(N^2)$ to $O(N \cdot \log N)$. Using this technique, the global strategy requires little computational overhead for very large values of N .

The idea of a partially-ordered binary tree was used in a well-known sorting algorithm due to Floyd [2]. (See Knuth [3] for a discussion of Floyd's tree sort algorithm.) It is a useful data structure for implementing other algorithms of numerical mathematics, e.g. Jacobi's method and Gauss-Southwell relaxation.

The binary tree is implemented in the Fortran subroutine SQUAGE (Simpson's quadrature used adaptively global error) using the integer array T. Nodes are inserted in the tree by the subroutine INSERT. New root nodes are moved down the tree (if necessary) for the subroutine BUBLDN; other new nodes are positioned in the tree by the subroutine BUBLUP. Each node in the tree contains the following information about its subinterval:

X1 left end point of the subinterval
L length of the subinterval
F1, F2, F3, F4, F5, F6 function values at equally-spaced
 abscissas (left to right) throughout the interval
ERREST the (scaled) error estimate for the subinterval

As presently programmed, the array space reserved for the binary tree and subinterval information is 9000 memory cells. This is considerably more storage than that required by subroutines using local termination criteria. On some machines this could be a serious disadvantage of the global strategies. However, many computing systems are now equipped with vast amounts of either real or "virtual" memory. It is worth noting that the subroutine SQUAGE together with the function subprogram which evaluates the integrand would normally occupy only a few pages of virtual storage. Hence, during the quadrature computation the working set is relatively small, thus satisfying the hypotheses of many virtual memory page swapping algorithms. On machines of this sort, the extra memory space required by the arrays in SQUAGE would not be a significant disadvantage. Higher-order adaptive methods using global strategies, such as AIND, require much less memory space than SQUAGE for the storage of RAM.

We are not presenting SQUAGE as a serious competitor of the better quadrature subroutines in the literature. It is given here only as a programming example, and was written to compute the experimental results given in the next section. SQUAGE lacks many important attributes of serious quadrature subroutines such as noise detection and a garbage collection mechanism for the case of binary-tree overflow.

```
REAL FUNCTION SQUAGE(A, B, EPS, ERR, NO, FUN)
INTEGER NO
REAL A, B, EPS, ERR, FUN
```

```
C
C SQUAGE -- SIMPSON'S QUADRATURE USED ADAPTIVELY WITH A BINARY
C TREE OF ERROR ESTIMATES
```

```
C 23 FEBRUARY 1974
C MICHAEL MALCOLM
C BRUCE SIMPSON
```

```
C SQUAGE - FIFTH-ORDER ESTIMATE OF INTEGRAL OF FUN(X) FOR
C A.LT.X.LT.B
C EPS - REQUESTED ACCURACY (ABSOLUTE)
C ERR - ESTIMATE OF ACHIEVED ERROR
C NO - NUMBER OF FUNTION EVALUATIONS USED IN QUADRATURE
C FUN - REAL FUNCTION OF ONE REAL VARIABLE TO BE INTEGRATED
```

```
C
C DECLARE BINARY TREE.
```

```
C
C REAL X1(1000),L(1000),F1(1000),F2(1000),F3(1000),
C 2 F4(1000),F5(1000),ERREST(1000)
C INTEGER T(1000)
```

```
C
C X1 - LEFT END POINT OF INTERVAL
C L - LENGTH OF INTERVAL
C F1-F5 - FUNCTION VALUES AT POINTS DELIMITING THE QUARTERSECTIONS OF
C THE INTERVAL
C ERREST - SCALED ERROR ESTIMATE FOR INTERVAL
C T - USED FOR INDEXING THE OTHER ARRAYS SUCH THAT THEY FORM A
C BINARY TREE
```

```
C
C OTHER VARIABLES.
```

```
C
C REAL L2, SEPS, SUM, ERR5
```

```
C
C L2 - HALF THE LENGTH OF THE ROOT INTERVAL
C SEPS - EPS SCALED BY 180. FOR ERROR TEST
C SUM - USED FOR ACCUMULATING TOTAL ESTIMATE OF INTEGRAL
C ERR5 - USED FOR ACCUMULATING TOTAL FIFTH-ORDER CORRECTION TERM
```

```
C
C INTEGER N, J, T1
```

```
C
C N - THE CURRENT NUMBER OF INTERVALS
```

```
C
C REAL PX1, PF1, PF3, PF5
```

```
C
C PX1-PF5 USED TO SATISFY SILLY ANSI FORTRAN RULES ABOUT PARAMETER
C PASSING.
```

```
C
C COMMON /SQUA/ X1, L, F1, F2, F3, F4, F5, ERREST, ERR5, T, N
C SEPS = 180.*EPS
C N = 1
C T(1) = 1
```

```
ERR5 = 0.
CALL INSERT(1,A,B-A,FUN(A),FUN(0.5*(A+B)),FUN(B), FUN)
GC TO 101
100 IF (ERR5.LF.SEPS) GC TO 200
101 T1 = T(1)
ERR5 = ERR5 - DABS(ERREST(T1))
L2 = 0.5*L(T1)
PX1 = X1(T1)
PF1 = F1(T1)
PF3 = F2(T1)
PF5 = F3(T1)
N = N+1
T(N) = N
CALL INSERT(N, PX1, L2, PF1, PF3, PF5, FUN)
CALL BUBLUP
T1 = T(1)
PX1 = X1(T1) + L2
PF1 = F3(T1)
PF3 = F4(T1)
PF5 = F5(T1)
CALL INSERT(1, PX1, L2, PF1, PF3, PF5, FUN)
CALL BUBLDN
IF (N.LT.1000) GO TO 100
C
C NO CONVERGENCE, WRITE ERROR MESSAGE
C
WRITE(6,150)
150 FORMAT(26H MEMORY OVERFLOW IN SQUAGE)
C
C THE ERROR TEST IS NOW SATISFIED, AND THE ESTIMATE CAN NOW BE SUMMED.
C
200 SUM = 0.
ERR5 = 0.
DO 300 J=1,N
SUM = SUM + L(J)*(F1(J) + 4.*F2(J) + 2.*F3(J) + 4.*F4(J)
2 + F5(J))
ERR5 = ERR5 + ERREST(J)
300 CONTINUE
SQUAGE = SUM/12. + ERR5/180.
EPR = DABS(ERR5/180.)
no = 4*n + 1
RETURN
END
SUBROUTINE INSERT(I, PX1, PL, PF1, PF3, PF5, FUN)
INTEGER I
REAL PX1, PL, PF1, PF3, PF5, FUN
C
C THIS SUBROUTINE SUBLIVIDES THE INTERVAL PASSED AS AN ARGUMENT TO
C OBTAIN AN ERROR ESTIMATE, AND THEN INSERTS IT AS THE I-TH NODE IN
C THE BINAPY TREE. TWO FUNCTION EVALUATIONS ARE PERFORMED DURING
C EACH CALL OF THIS SUBROUTINE.
C
REAL X1(1000),L(1000),F1(1000),F2(1000),F3(1000),
2 F4(1000),F5(1000),ERREST(1000),ERR5
INTEGER T(1000), N, TI
```

```
COMMON /SQUA/ X1, L, F1, F2, F3, F4, F5, ERREST, ERR5, T, N
TI = T(I)
X1(TI) = PX1
L(TI) = PL
F1(TI) = PF1
F2(TI) = FUN(PX1 + 0.25*PL)
F3(TI) = PF3
F4(TI) = FUN(PX1 + 0.75*PL)
F5(TI) = PF5
ERREST(TI) = -PL*(PF1 - 4.*F2(TI) + 6.*PF3 - 4.*F4(TI)
2          + PF5)
ERR5 = ERR5 + DABS(ERREST(TI))
RETURN
ENL
SUBROUTINE BUBLUP
```

C
C THE ITEM IN NODE N OF THE TREE IS BUBBLED UP THE TREE TO
C MAINTAIN PARTIAL ORDERING OF THE TREE.
C

```
REAL X1(1000),L(1000),F1(1000),F2(1000),F3(1000),
2 F4(1000),F5(1000),ERREST(1000),ERR5
INTEGER T(1000), N
COMMON /SQUA/ X1, L, F1, F2, F3, F4, F5, ERREST, ERR5, T, N
INTEGGER I, J, ITEMP, TI, TJ
I = N
10 J = I/2
IF (J.EQ.1) GO TO 100
TI = T(I)
TJ = T(J)
IF (DABS(ERREST(TI)).LE.DABS(ERREST(TJ))) GO TO 100
ITEMP = T(I)
T(I) = T(J)
T(J) = ITEMP
I = J
GO TO 10
100 RETURN
END
SUBROUTINE BUBLDN
```

C
C THE ITEM IN THE ROOT NODE IS BUBBLED DOWN THE TREE TO
C MAINTAIN PARTIAL ORDERING OF THE TREE.
C

```
REAL X1(1000),L(1000),F1(1000),F2(1000),F3(1000),
2 F4(1000),F5(1000),ERREST(1000),ERR5
INTEGER T(1000), N, TI, TJ, TJ1
COMMON /SQUA/ X1, L, F1, F2, F3, F4, F5, ERREST, ERR5, T, N
INTEGGER I, J, ITEMP
I = 1
10 J = 2*I
IF (J.GT.N) GO TO 100
IF (J.EQ.N) GO TO 20
TJ = T(J)
TJ1 = T(J+1)
IF (DABS(ERREST(TJ)).GE.DABS(ERREST(TJ1))) GO TO 20
J = J+1
```

```
20  TI = T(I)
    TJ = T(J)
    IF (DABS(ERREST(TI)).GE.DABS(ERREST(TJ))) GC TO 100
    ITEMP = T(I)
    T(I) = T(J)
    T(J) = ITEMP
    I = J
    GC TO 10
100  RETURN
    END
```

4. Experimental computations

In this section, we give the results of three computational experiments which test the model given in Section 2. We also compare the performance of SQUANK and SQUAGE with two related subroutines: CADRE and AIND. These computations were all performed on the University of Waterloo's Honeywell 6050 computer.

The first experiment involves the integrand $f(x) = (1+c-x)^{-1}$ for a fixed error tolerance and various positive values of c . The second studies the same integrand for $c = .01$ while varying the absolute error tolerance ϵ . The third experiment uses the integrand $f(x) = (1-x)^\alpha$, as α varies. All integrations are over the interval $0 \leq x \leq 1$.

Both CADRE and AIND use higher-order local quadrature procedures (on each subinterval) than SQUANK or SQUAGE. CADRE [1] uses a dynamic extrapolation process for its local quadrature procedure; when it senses that the extrapolation may not be working, it stops to make an adaptive subdivision of the interval. In this sense, it is based on a high (variable) order local quadrature rule. It uses a local interval acceptance criterion. The local quadrature rule for AIND is the 21-point Kronrod formula obtained by adding 11 points to the 10-point Gaussian formula [8]. Hence, AIND uses a very high-order local quadrature rule. It is the only published algorithm known to the authors which uses a global interval acceptance criterion.

For the study of $f(x) = (1+c-x)^{-1}$, an absolute error tolerance of $\epsilon = 10^{-6}$ was used. CADRE and AIND have both absolute and relative error tolerances appearing in their calling sequences, and they attempt

to meet the least stringent of the two requirements. Hence, for purposes of comparison, we set the relative tolerances to zero for these tests. The predicted number of function evaluations for SQUANK was developed in Section 2 and is

$$(4.1) \quad 4 L_{\text{pred}}(f;0,1,\epsilon)+1 = 4L(64/255)^{1/4}(c^{-1/4} - (1+c)^{-1/4})\epsilon^{-1/4} J+1 .$$

For SQUAGE it is

$$(4.2) \quad 4 G_{\text{pred}}(f;0,1,\epsilon)+1 = 4L2^{-7/4}15^{-1/4}(\log_e(1+c^{-1}))^{5/4}\epsilon^{-1/4} J+1 .$$

Plots of $-\log_{10} c$ versus $\log_{10} N$ are given in Figures 4.1 and 4.2 to show the comparison between predicted and actual numbers of function evaluations for SQUANK and SQUAGE. (here, N denotes the number of function evaluations.)

In these Figures, the predicted growth rates of N with c appear to be borne out; however, the actual numbers of function evaluations tend to be greater than predicted by a factor of about 1.3 for SQUANK, and 1.1 for SQUAGE. This probably represents a minor failing of assumptions (i) and (iii) of Section 2. In Figure 4.3, we show the actual numbers of function evaluations used as they vary with c for all four subroutines mentioned. In Table 4.1, we tabulate the actual errors made. The performances of CADRE and AIND for that integrand are also listed in Table 4.1, for comparison.

Turning now to the second experiment, we fix the integrand at $f(x) = (1.01-x)^{-1}$ and look at the effects of varying ϵ . Table 4.2 lists the numbers of function evaluations predicted and used by SQUAGE




FIGURE 4.1




FIGURE 4.2




FIGURE 4.3

LOG ₁₀ (c)	SQUANK		SQUAGE		CADRE		AIND	
	A	B	A	B	A	B	A	B
- .5	-1.5 E-8	45	-1.5 E-8	37	4.2 E-9	48	4.8 E-9	21
-1.0	-3.0 E-8	93	-6.0 E-8	65	7.8 E-9	80	8.8 E-8	63
-1.5	0	157	0	105	9.9 E-9	96	1.8 E-8	147
-2.0	6.0 E-8	241	0	145	1.8 E-8	128	4.7 E-7	189
-2.5	-6.0 E-8	361	-1.2 E-7	189	1.7 E-8	160	9.0 E-8	273
-3.0	0	513	-6.0 E-8	241	2.1 E-8	176	-9.9 E-9	357
-3.5	-1.2 E-7	733	-2.4 E-7	289	3.0 E-8	208	3.9 E-7	399
-4.0	0	957	-3.6 E-7	341	3.0 E-8	240	9.8 E-8	483
-4.5	0	1153	2.4 E-7	405	3.3 E-8	256	3.2 E-8	567
-5.0	0	1381	3.6 E-7	457	4.2 E-8	288	3.4 E-7	609

Column A - error

Column B - number of function evaluations used

A Comparison of Performance on $\int_0^1 (1+c-x)^{-1} dx$ for $\epsilon = 10^{-6}$

Table 1

Log ₁₀ ε	SQUANK			SQUAGE			CADRE		AIND	
	A	B	C	A	B	C	A	B	A	B
-3	- 4.1	37	34	- 3.5	29	22	- 4.9	56	- 4.3	147
-4	- 5.7	77	61	- 4.7	49	40	- 5.9	80	- 4.3	147
-5	- 6.8	137	108	- 6.4	85	72	- 7.1	120	- 6.3	189
-6	- 8.3	241	193	- 7.8	145	129	- 7.7	128	- 6.3	189
-7	- 9.8	425	344	- 9.3	253	229	- 8.7	176	- 9.1	231
-8	-11.3	757	612	-10.7	453	408	-10.0	240	- 9.1	231
-9	-12.8	1325	1089	-12.3	805	726	-10.9	256	- 9.1	231
-10	-14.3	2373	1938	-13.8	1429	1292	-12.1	384	-12.6	273
-11	-15.8	4249	3446	-15.3	2537	2298	-12.9	480	-12.6	273
-12	-16.5	7585	6128	-15.8	4001	4087	-13.5	512	-12.6	273

Column A - $\log_{10}|\text{error}|$

Column B - Number of function evaluations used

Column C - Number of function evaluations predicted

Table 4.2 A Comparison of Performance on $\int_0^1 (1.01-x)^{-1} dx$

and SQUANK. Again, the predictions for SQUANK and SQUAGE are low by factors of about 1.3 and 1.1, respectively. However, the actual numbers used seem to grow like $\epsilon^{-1/4}$, as predicted. The performances of CADRE and AIND are also listed in Table 4.2. The much slower growth rate of N with ϵ for CADRE and AIND demonstrates an advantage of using high-order local quadrature procedures.

It should be noted from Tables 4.1 and 4.2 that SQUAGE shows a tendency to meet the prescribed error tolerance more closely (from below) than SQUANK.

Our third experiment involves integrands $f(x) = (1-x)^\alpha$ for various values of α . The model in Section 2 predicts that the subinterval selection process using the local acceptance criterion (1.2) does not terminate if $\alpha < 0$. Although SQUANK uses (1.2), it has a limited storage space for the stack used to store the subintervals in Pen. Thus, after a maximum number of bisections (when the stack is full), the current subinterval is accepted without reference to (1.2). In the version of SQUANK given in [4], the maximum number of bisections (maximum stack length) is set at 30. However, for this experiment, the maximum was reset to 500 to avoid the effect of default acceptance of subintervals not accounted for in the model of Section 2. (For these integrands, the default acceptance of subintervals is benign in the sense that the error tolerance is always met using the unchanged version of SQUANK.)

Using $p_4(\alpha) = |\alpha(\alpha-1)(\alpha-2)(\alpha-3)|$, of Section 2, the predicted numbers of function evaluations as given in Section 2 are, for SQUANK

$$(4.3) \quad 4 L_{\text{pred}}(f;0,1,\alpha)+1 = 4L(p_4(\alpha)/180\epsilon)^{1/4}/\alpha_{J+1}, \quad \text{if } \alpha \geq 0,$$

and for SQUAGE,

$$(4.4) \quad 4 G_{\text{pred}}(f;0,1,\epsilon) = 4L(5/4(1+\alpha))^{5/4}(p_4(\alpha)/45\epsilon)^{1/4}_{J+1}, \quad \text{if } \alpha > -1.$$

Predictions (4.3) and (4.4) were tested using error tolerance $\epsilon = 10^{-6}$. The results for SQUANK are given in Figure 4.4; those for SQUAGE are given in Figure 4.5.

In addition to these three computations, a general comparison between SQUANK and SQUAGE was made by testing them using a quadrature testing program written by Professor P. Keast of the University of Toronto (Department of Mathematics). Keast's program uses 131 test functions divided into seven groups, each group exhibiting a specific behavior (e.g., oscillation, end-point singularity, etc.). The quadrature subroutine is tested on each of these functions at 10 different absolute tolerance levels: $\epsilon = 10^{-n}$, $n = 1(1)10$. The observations made from the results of Keast's tests on SQUANK and SQUAGE are similar for all groups except groups 5 and 6. With the exception of groups 5 and 6, SQUAGE always took fewer function evaluations, typically by a factor of 2 or 3 at the smaller values of ϵ . SQUAGE also tended to come closer to the prescribed tolerance, in general. Group 5 is composed of functions with singularities, like $f(x) = |x-c|^\alpha$ for c in the range of integration, and $-1 < \alpha < 0$. For this group, SQUANK's




FIGURE 4.4




FIGURE 4.5

performance is markedly inferior to SQUAGE's, as our previous discussion would lead us to expect. Group 6 is comprised of various step functions. For these, SQUANK generally took about 20 to 40 function evaluations and gave up, apparently interpreting the discontinuities as noise in the function evaluations. Since SQUAGE doesn't attempt to recognize roundoff error, or other noise, it continued on and met the tolerance for about half of the functions in Group 6. Of course, for this group, SQUAGE generally used substantially more function evaluations.

It is perhaps worth noting that Group 5 represents an important class of functions for which SQUAGE is successful but SQUANK fails, using vastly more function evaluations. Group 6 is not as important since neither of the subroutines is designed for integrands having discontinuities. Using Group 6 for testing these particular quadrature subroutines is, as Lyness has remarked, a little like building a truck and then testing it by driving it over a river.

There were no examples in which both subroutines met the error tolerance with SQUANK requiring fewer function evaluations. However, simple examples can be constructed in which both are successful, but SQUANK requires fewer subintervals; so no simple statement concerning relative numbers of function evaluations seems proveable.

Keast's tests for SQUANK took .575 hours of processor time whereas the tests for SQUAGE took .205 hours. Thus, the number of function evaluations appears to be a reasonable measure of complexity when comparing the two subroutines.

5. Conclusions

The results of the preceding section suggest that the analytical model of Section 2 works well enough to be a useful tool for studying adaptive quadrature algorithms. This model indicates that the number of subintervals selected may be regarded as the product of two factors: one that summarizes the interaction between the algorithm and the integrand; the other depends only on the absolute error tolerance.

For both SQUANK and SQUAGE this second factor is $\epsilon^{-1/4}$, and it is clearly derived from the use of Simpson's rule as a local quadrature procedure. In this sense, then, we can call both SQUAGE and SQUANK fourth-order algorithms, and note that the interval acceptance criterion has no influence on the order.

The connection between the order of the local quadrature rule and the global performance of an adaptive quadrature method has also been observed by Rice [11] in the context of the rate of convergence to zero of the quadrature error. These results seem to be consistent with his, when the error estimate used is actually an error bound for each subinterval.

While it may not influence the order of an adaptive quadrature algorithm, the subinterval selection strategy appears to influence its performance in at least three ways i) by affecting the number of function evaluations required for an integral; ii) by changing the domain of integrands which can be handled by the algorithm; iii) by affecting the closeness with which the user's tolerance is achieved. Each of these

ways has been observed both in the experimental computations and the general testing using Keast's program. The first is clearly evident in the reduced number of function evaluations used by SQUAGE as compared with SQUANK when applied to $f(x) = (1+c-x)^{-1}$, and in the 2.05/5.75 ratio of running times for the general tests on these two routines. The second influence is demonstrated by the difference in success between SQUAGE and SQUANK when applied to $(1-x)^\alpha$ for $-1 < \alpha < 0$. This effect has also been observed by Rice and given a more rigorous treatment in [10] and [11]. The third influence mentioned above of the subinterval selection process can be seen in Tables 4.1 and 4.2, and was also evident in the general testing.

The advantages of global strategies over local strategies were seen to be in some sense a space-time tradeoff. The reduced numbers of required function evaluations, and apparent increase in domain of applicability using the global strategy comes at the expense of additional memory requirements. However, the memory requirements of SQUAGE are not unreasonable for many applications; and when a higher-order local quadrature rule is used, such as in AIND, the additional memory requirements are very small.

An adaptation of Floyd's tree sort algorithm has proven to be an effective way of programming the global strategy when a large number of subintervals are anticipated. Using this technique, the additional overhead involved in the global strategy is negligible.

6. References

- [1] deBoor, C., "CADRE: An algorithm for numerical quadrature", in Mathematical Software (ed., J. Rice), pp.417-449, New York: Academic Press, 1971.
- [2] Floyd, R.W., "Treesort 3", Algorithm 245, Comm. ACM, Vol.7, No.12, Dec.1964, 701.
- [3] Knuth, D.E., "Sorting and searching", The Art of Computer Programming, Vol.3, Reading Mass: Addison Wesley.
- [4] Lyness, J.N. "SQUANK (Simpson quadrature used adaptively - noise killed)", Algorithm 379, Comm. ACM, Vol.13, No.1, April 1970, 260-262.
- [5] Lyness, J.N., "Notes on the adaptive Simpson quadrature routine", J. ACM, Vol.16, 1969, 483-495.
- [6] McKeeman, W.M., "Adaptive numerical integration by Simpson's rule", Algorithm 145, Comm. ACM, Vol.5, 1962, 604.
- [7] Patterson, T.N.L., "Algorithm for automatic numerical integration over a finite interval", Comm. ACM, Vol.16, No.11, Nov. 1973, 694-699.
- [8] Piessens, R., "An algorithm for automatic integration", Angewandte Informatik, Sept. 1973, 399-401.
- [9] Piessens, R., "A quadrature routine with round-off error guard", Report TW 17, March 1974, Appl. Math. and Programming Div., Katholieke Universiteit Leuven.
- [10] Rice, J., "An educational adaptive quadrature algorithm", SIGNUM Newsletter, Vol.8, No.2, April 1973, 27-41.
- [11] Rice, J., "A metalgorithm for adaptive quadrature", to appear in J. ACM.
- [12] Rowland, J.H. and Y.L. Varol, "Exit criteria for Simpson's compound rule", Math. Comp., Vol.26, No.119, July 1972, 699-704.