

NBER WORKING PAPER SERIES

ROSEPACK Document No. 2:

AUTOMATING STEM-AND-LEAF DISPLAYS

David C. Hoaglin*
Stanley S. Wasserman**

Working Paper No. 109

COMPUTER RESEARCH CENTER FOR ECONOMICS AND MANAGEMENT SCIENCE
National Bureau of Economic Research, Inc.
575 Technology Square
Cambridge, Massachusetts 02139

November 1975

Preliminary

NBER working papers are distributed informally and in limited numbers.

This report has not undergone the review accorded official NBER publications; in particular, it has not yet been submitted for approval by the Board of Directors.

*NBER Computer Research Center and Harvard University. Research supported in part by National Science Foundation Grants DCR 75-08802 to the National Bureau of Economic Research and SOC 72-05257 to Harvard University.

**NBER Computer Research Center and Harvard University. Research supported in part by National Science Foundation Grant DCR 70-03456 A04 to the National Bureau of Economic Research.

Abstract

The stem-and-leaf display is a natural semi-graphic technique to include in statistical computing systems. This paper discusses the choices involved in implementing both automated and flexible versions of the display, develops an algorithm for the automated version, examines various implementation considerations, and presents a set of semi-portable FORTRAN subroutines for producing stem-and-leaf displays.

Table of Contents

1. Introduction	1
2. Stem-and-Leaf Displays	2
Exhibit 1	3
3. How Many Lines?	5
4. What Scaling?	6
5. A Scaling Algorithm	8
6. Stems, Units, and Leaves	10
7. Some Refinements	11
8. Examples	12
Exhibit 2	13
Exhibit 3	14
Exhibit 4	15
Exhibit 5	16
Exhibit 6	17
9. Acknowledgments	18
References	19
Appendix: A FORTRAN Implementation	A1
A1. Organization of the Subroutines	A1
A2. Numerical Aspects	A4
A3. Calculating the Depths	A5
A4. Installation As a Command	A6
A5. Error Checking	A6
A6. FORTRAN Listings	A7

1. Introduction

The stem-and-leaf display has so many uses in everyday data analysis that it should be a natural component of any modern statistical computing system. In implementing it, however, one must make many decisions. For example, how many lines of output should the display occupy? How should it be scaled to give a pleasant and effective appearance? How many different leaf digits should fall on one line -- ten, five, or two? Some of the answers doubtless involve personal taste, so the system must have enough flexibility to produce the display the user really wants. Why, then, should one consider automation, which would seem to deprive the user of that control? There is actually a strong justification for trying to produce automatically a reasonable stem-and-leaf display for a given batch of data. If the system is interactive, the user will usually need to ask for only minor changes in order to produce the finished display. More extensive trial would be easy, but it would generally be unnecessary. In a non-interactive system the much longer delay in receiving the next output places a premium on getting close to the "right" stem-and-leaf display on the first try so that at most one further iteration is necessary. In what follows we discuss several basic rules which lead "automatically" to a reasonable stem-and-leaf display. An appendix presents a set of "semi-portable" FORTRAN subroutines and discusses some of the details of implementation.

2. Stem-and-Leaf Displays

For starting to look at a batch or sample of data the stem-and-leaf display, developed by John W. Tukey [7,8], provides a flexible and effective technique. The basic idea is to let the most significant digits of the data values themselves do most of the work of sorting the batch into numerical order and displaying it. In the simplest form one chooses a suitable pair of adjacent digits in the data, splits each data value between these two digits, allocates a separate line in the display for each possible string of leading digits (the stem), and writes down the first trailing digit (the leaf) of each data value on the line corresponding to its leading digits. (The name "stem-and-leaf" comes about by analogy to espaliered trees or shrubs, which are trained so that their trunks grow vertically against a wall and their branches grow horizontally along it.) An example readily shows how the process works.

Frohliger and Kane [2] report the pH values for 26 samples of precipitation collected at a location in Allegheny County, Pennsylvania, from December 1973 to June 1974. In chronological order the data values are

4.57, 5.62, 4.12, 5.29, 4.64, 4.31, 4.30, 4.39, 4.45,
5.67, 4.39, 4.52, 4.26, 4.26, 4.40, 5.78, 4.73, 4.56,
5.08, 4.41, 4.12, 5.51, 4.82, 4.63, 4.29, 4.60.

In this case it is reasonable to split between the second and third digits; for example, 4.57 yields 4.5|7 . The necessary lines (17 in all) run from 4.1 to 5.7, and writing down the leaves in chronological order gives the raw display on the left in Exhibit 1. In the finished display the decimal points have been dropped in favor of a reminder that all data values are in units of .01, an occasional asterisk indicates that the leaves are one-digit, and a column of depths (which we will shortly define) has been added to the left of the stems. In overall appearance the display is similar to a histogram with an interval width of .1; the leaves add numerical detail.

Exhibit 1

Stem-and-Leaf Displays for Precipitation pH Data

<u>raw</u>			<u>finished</u>			
			(unit = .01)			
4.1		22	2	41*		22
4.2		669	5	42		669
4.3		1099	9	43		1099
4.4		501	12	44*		501
4.5		726	(3)	45		726
4.6		430	11	46		430
4.7		3	8	47*		3
4.8		2	7	48		2
4.9				49		
5.0		8	6	50*		8
5.1				51		
5.2		9	5	52		9
5.3				53*		
5.4				54		
5.5		1	4	55		1
5.6		27	3	56*		27
5.7		8	1	57		8

A data value can be assigned a rank by counting in from each end of the batch. The depth of the data value is the smaller of these two ranks. Since a number of summary values (such as the median and the quartiles or the hinges) can easily be defined in terms of their depths, it is helpful to present a set of depths with the display. Except for one middle line, the number in the depth column is the maximum depth associated with data values on that line. Thus the depth of 4.29 is 5. The "middle line" (absent when the batch size is even and the median falls between lines) includes the median, and the depth column shows in parentheses the number of leaves (in the example, 3) on this line. If the display has been prepared by hand, adding the count on the middle line and the depths on the two adjacent lines provides a simple check that no data values have been omitted. In the example, $12 + 3 + 11 = 26$.

Either display in Exhibit 1 reveals quite a lot about the behavior of the precipitation pH data: a rather flat distribution of values from 4.1 to 4.7 with scattered values trailing off above that to 5.3 and a clump of four values from 5.51 to 5.78. It is worth remarking that Frohlinger and Kane give the range and average of their data but do not comment on the distribution of the sample, which hardly lends itself to such a simple summary and may suggest a mixture of two populations.

3. How Many Lines?

Experience suggests that an effective choice of the number of lines involves the number of data values in the batch as well as the range to be covered. In view of the similarities between stem-and-leaf displays and histograms, we should be able to calculate the maximum number of lines according to a rule given by Dixon and Kronmal [1]:

$$L = [10 \times \log_{10} n] ,$$

where n is the number of data values and $[x]$ is the largest integer not exceeding x . (Interestingly, Dixon and Kronmal based their histogram rule on a suggestion by Tukey!) This seems to give very reasonable values of L over the range $20 < n < 300$, where almost all applications fall. Values of n smaller than 20 may need special treatment. These cases are also more likely to arise when comparing several batches in parallel stem-and-leaf displays, a situation we would want to handle differently anyway. Batches of 300 or so are usually cumbersome in a stem-and-leaf display, but the rule should still cope with them reasonably well. There is nothing sacred about the constant factor 10 in the definition of L , and further experience may lead to a different value. As an alternative rule, Velleman [9] has suggested $L = [2\sqrt{n}]$.

4. What Scaling?

Using L as a rough limit on the number of lines in the display, we must now determine the interval of values corresponding to each line. The simple way to do this (implicit in the earliest form of stem-and-leaf display, as in Exhibit 1) is to arrange that the interval width be a power of 10. This we can easily accomplish by dividing R , the range of the batch, by L and rounding the quotient up (if necessary) to the nearest power of ten. A segment of the stems for such a display might look like this:

```
0|  
1|  
2|  
3|
```

and each line would receive leaves 0 through 9. It soon became clear that the display was sometimes too crowded, having too many leaves per line.

Tukey's response was to split the lines and repeat each stem:

```
0*|  
0·|  
1*|  
1·|  
2*|  
2·|
```

putting leaves 0 through 4 on the * line and 5 through 9 on the · line.

In such a display the interval width is 5 times a power of 10.

This change made a great improvement, but there were still some cases in which the result was too crowded even in the split-stem form and too straggly in the original form at the next lower power of 10. To cure these troubles, Tukey added the third form, five lines per stem:

0*	
t	
f	
s	
0.	

with leaves 0 and 1 on the * line, 2 and 3 on the t line, 4 and 5 on the f line, 6 and 7 on s line, and 8 and 9 on the . line. (As a reminder in starting to place leaves it is convenient that the three lettered lines contain leaves whose words begin with that letter.) Here the interval width is 2 times a power of 10.

5. A Scaling Algorithm

Viewing the three forms together, all we need to do is divide R by L and round the quotient up to 1, 2, or 5 times a power of 10. This rule of thumb is useful when one is preparing a stem-and-leaf display by hand, but we must formulate it as a specific algorithm for a computer. This is straightforward if we adapt the algorithm of Dixon and Kronmal [1].

Their algorithm, intended for scaling axes in graphs and histograms, uses a set of "round" numbers p_1, \dots, p_m such that $1 \leq p_1 < p_2 < \dots < p_m < 10$. For scaling stem-and-leaf displays we need only $p_1 = 1, p_2 = 2, p_3 = 5$. To fix notation, we let

- A = smallest value in the batch,
- B = largest value in the batch, and
- S = scale factor (the interval width)

and recall that

- L = number of lines in the display, and
- R = B - A.

Thus we want $L \times S \geq R$ with S as small as possible and $S = p_i \times 10^k$ for some i and k . Like Dixon and Kronmal, we begin by taking $t = \lceil \log_{10}(R/L) \rceil$. If $(R/L)/10^t \leq p_3 = 5$, we let $k = t$ and find the smallest p_i which is at least $(R/L)/10^k$. Otherwise, $k = t+1$ and the desired p_i is $p_1 = 1$.

Next we calculate the number of stems actually required for the display and determine whether S must be increased. We easily find that the number of lines actually used is $\text{sign}(B) \times \lceil |B/S| \rceil - \text{sign}(A) \times \lceil |A/S| \rceil + 1$ (plus 1 if

needed for the -0 stem). This can exceed L , and if it does, we replace S with the next larger "round" number, increasing k if necessary. This completes the scaling.

6. Stems, Units, and Leaves

In producing the display we must calculate the stems and then separate each data value into a starting part and a leaf. To do this, we recall that all entries in a stem-and-leaf display may be regarded as integer multiples of a power of 10, referred to as the unit. If we denote this unit by U , we can very easily calculate it from S : When $S = 2 \times 10^k$ or $S = 5 \times 10^k$, $U = 10^k$; and when $S = 10^k$, $U = 10^{k-1}$.

In calculating the stems (or more precisely, the labels on the lines) we must take into account the special role of zero. We implicitly number the lines of the display in steps of S/U from $\text{sign}(A) \times [|A/S|] \times (S/U)$ to $\text{sign}(B) \times [|B/S|] \times (S/U)$, including -0 if it occurs. We drop the low-order digit of each such number, and when S is 2 or 5 times a power of ten, that digit indicates which of the special labels such as t and \cdot to use for an intermediate line.

Finally we are ready to separate each data value into a starting part and a leaf and collect the leaves on their proper stems. We will generally cut each data value to a multiple of U by simply discarding low-order digits. (This makes it easy to match an entry in the display with its original data value whenever that value requires further attention.) In that integer multiple, then, the last digit is the leaf, and the other digits are the starting part. For example, with $U = .01$ we would cut 2.213 to 2.21 and separate it into a starting part of 22 and a leaf of 1.

7. Some Refinements

In a sense this discussion has removed the stem-and-leaf display from its natural context, exploratory data analysis. Since resistant analyses (which are little affected by changes in a small fraction of the data) are an important part of the exploratory mode, it is quite unwise for an automated stem-and-leaf display to depend so heavily on the extreme values, A and B. Instead, we should begin by setting aside any serious outliers and basing A and B on those values which remain.

Since the batch will already be sorted (or will be sorted as the first step in producing a stem-and-leaf display), it is easy to isolate outliers. One useful rule of thumb [7] finds the upper and lower hinges (approximate quartiles) and their difference dH and sets aside any data values further than $1.5 \times dH$ from the nearer hinge. In printing the display we can add the lines "hi" and "lo" beyond the set of stems and list those outlying values.

Another important refinement lets us handle parallel displays (with one common set of stems) for comparing several batches. Here the tentative rule (subject to further experience) is to regard the individual batches as combined into a single batch for the purposes of scaling and forming stems. This evidently provides a set of stems which covers the combined range, and using the total number of data values to determine the number of lines will accommodate moderate shifts from batch to batch. (Large shifts may lead us to line up the batches by subtracting at least a rough adjustment from each.)

8. Examples

A few examples should clarify the suggestions for automating stem-and-leaf displays. We present four which are representative of a broad range of possibilities.

In the first example (Exhibit 2) all steps are straightforward. We can easily round R/L up to give $S=5$ without using logarithms, and the line requirement (10) does not exceed L (14). The display seems quite satisfactory.

If we do not check for possible outliers, the new feature in the second example (Exhibit 3) is the need to increase S from .02 to .05 in order to keep the line requirement within L . The resulting display may seem rather bunched on the stems 23^* and $23\cdot$. A closer look uncovers four possible stray values at the low end and leads to a somewhat more reasonable display (Exhibit 4). Since the four low values do not stray far, we might also use 15 lines and avoid separating the low values from the rest of the batch.

Exhibit 5 shows how to handle $+0$ and -0 when the batch contains both positive and negative values. In this case no data values need to be set aside, and it is a simple task to construct the display.

The fourth and last example (Exhibit 6) shows both parallel displays and an obvious outlier. After we set aside the one straying value, the display is reasonably effective, but the straggling behavior of the third batch has clearly taken a toll.

All in all, the present "automated" approach usually seems to produce the stem-and-leaf displays one might prefer after some trial and error.

Exhibit 2

Data (hardness of aluminum die castings [6, p. 42])

53.0, 70.2, 84.3, 55.3, 78.5, 63.5, 71.4, 53.4,
82.5, 67.3, 69.5, 73.0, 55.7, 85.8, 95.4, 51.1,
74.4, 54.1, 77.8, 52.4, 69.1, 53.5, 64.3, 82.7,
55.7, 70.5, 87.5, 50.7, 72.3, 59.5

Calculations

$$n = 30, L = [10 \times \log_{10} 30] = [14.77] = 14$$

$$A = 50.7, B = 95.4, R = B - A = 44.7$$

$$R/L = 44.7/14 = 3.19; S = 5, U = 1$$

$$\text{lines required} = [95.4/5] - [50.7/5] + 1 = 10$$

Display

(UNIT = 0.1000E+01)

7	5	I	0123334
11	5.	I	5559
13	6	I	34
3	6.	I	799
14	7	I	001234
8	7.	I	78
6	8	I	224
3	8.	I	57
1	9	I	
1	9.	I	5

Exhibit 3

Data

2.346, 2.334, 2.365, 2.417, 2.399, 2.354,
2.339, 2.368, 2.257, 2.358, 2.326, 2.334,
2.313, 2.298, 2.371, 2.197, 2.378, 2.395,
2.335, 2.207, 2.398, 2.211, 2.273, 2.213,
2.330, 2.359, 2.468, 2.352, 2.463, 2.398

Calculations

$$n = 30, L = [10 \times \log_{10} 30] = 14$$

$$A = 2.197, B = 2.468; R = 0.271$$

$$R/L = 0.271/14 = 0.0194; S = .02, U = .01$$

$$\text{lines required} = [2.468/.02] - [2.197/.02] + 1 = 15$$

Since this exceeds L, S must be increased to .05.

Display (unit = .01)

1	21.	9
4	22*	011
7	22.	579
15	23*	12333334
15	23.	555566779999
3	24*	1
2	24.	66

Exhibit 4

Data (in Exhibit 3)

Calculations

hinges: 2.313 and 2.378, dH = .065

cutoff values: 2.2155 and 2.4755

Set aside 2,197, 2.207, 2.211, and 2.213 to get A = 2.257 and B = 2.468 with n = 26 and L = 14.

R/L = 0.211/14 = .015; S = .02, U = .01

Lines required = 12

Display

4 LO I 2.1970 2.2070 2.2110 2.2130

(UNIT = 0.1000E-01)

5 F I 5
6 S I 7
7 22. I 9
8 23. I 1
14 T I 233333
5 F I 45555
11 S I 6677
7 23. I 9999
3 24. I 1
2 T I
2 F I
2 S I 66

Exhibit 5

Data

-1.4, 1.2, 0.5, -0.3, -0.8, 0.4, -1.3, 0.5,
0.7, -0.2, 0.1, 2.3, 0.1, -0.8, -2.6, 0.7
-0.9, -0.4, 0.2, 1.3

Calculations

$$n = 20, L = [10 \times \log_{10} 20] = 13$$

$$A = -2.6, B = 2.3; R = 4.9$$

$$R/L = 4.9/13 = .38; S = .5, U = .1$$

$$\text{lines required} = [2.3/.5] + [2.6/.5] + 2 = 11$$

Display

(UNIT = 0.1000E+00)

1	-2.	I	0
1	-2	I	
1	-1.	I	
3	-1	I	43
6	-0.	I	988
9	-0	I	432
4	0	I	1124
7	0.	I	5577
3	1	I	23
1	1.	I	
1	2	I	3

Exhibit 6

Data (counties, including independent cities, by region in U.S. states)

Northeast: 8, 16, 14, 10, 21, 62, 67, 5, 14

North Central: 102, 92, 99, 105, 83, 87, 115, 93, 53, 88, 67, 72

South: 67, 75, 3, 67, 159, 120, 64, 24, 82, 100, 77, 46, 95, 254, 130, 55

West: 29, 14, 58, 63, 5, 44, 56, 17, 32, 36, 29, 39, 33

Calculations (scale as one batch)

Hinges: 24 and 88, dH = 64

cutoff values: -72 and 184

Set aside 254 (Texas) to get A = 3 and B = 159 with n = 49 and L = 16

R/L = 156/16; S = 10, U = 1

Lines required = 16

Display (unit = 1 county)

0*	85		3	5
1	6404			47
2	1		4	99
3*				2693
4			6	4
5		3	5	86
6*	27	7	774	3
7		2	57	
8		378	2	
9*		293	5	
10		25	0	
11		5		
12*			0	
13			0	
14				
15*			9	
hi			254	

9. Acknowledgments

The present implementation of stem-and-leaf displays has evolved from earlier one-line-per-stem versions programmed by Michael D. Godfrey for the instructional computing package SNAP/IEDA on the IBM 7094 and refined by Hale F. Trotter when SNAP/IEDA was converted to the IBM 360. We are also indebted to Paul W. Holland, Paul F. Velleman, and Roy E. Welsch for articulating the needs of a variety of users, to Neil E. Kaden and Virginia Klema for valuable discussions of semi-portability, and to Stephen C. Peters for assistance in testing and debugging.

References

- [1] Dixon, W.J., and Kronmal, R.A., 1965, "The Choice of Origin and Scale for Graphs," Journal of the Association for Computing Machinery 12, 259-261.
- [2] Frohlinger, J.O., and Kane, R., 1975, "Precipitation: Its Acid Nature," Science 189, 455-457.
- [3] Hoaglin, D.C., and Welsch, R.E., 1974, "MIT-SNAP, An Interactive Data Analysis System," Sloan School of Management, Massachusetts Institute of Technology.
- [4] Ryder, B.G., 1974, "The PFORT Verifier," Software -- Practice and Experience 4, 359-377.
- [5] Ryder, B.G., and Hall, A.D., 1975, "The PFORT Verifier," Computing Science Technical Report #12, Bell Laboratories, Murray Hill, New Jersey.
- [6] Shewhart, W.A., 1931, Economic Control of Quality of Manufactured Product, D. Van Nostrand, Inc., Princeton, New Jersey.
- [7] Tukey, J.W., 1970, Exploratory Data Analysis (Limited Preliminary Edition), Volume 1, Addison-Wesley, Reading, Massachusetts.
- [8] Tukey, J.W., 1972, "Some Graphic and Semigraphic Displays," Statistical Papers in Honor of George W. Snedecor (T.A. Bancroft, editor), Iowa State University Press, Ames, Iowa.
- [9] Velleman, P.F., 1975, "Interactive Computing for Exploratory Data Analysis I: Display Algorithms," American Statistical Association: Proceedings of the Statistical Computing Section, 1975, American Statistical Association, Washington, D.C.

APPENDIX

A FORTRAN Implementation

To illustrate various aspects of automating stem-and-leaf displays, we have developed a set of "semi-portable" FORTRAN subroutines. ("Semi-portable" implies that it should be possible to compile the subroutines under many different versions of FORTRAN, but that a few machine-dependent details remain. We have used the PFORT Verifier [4,5] to check adherence to PFORT, a large, carefully defined, portable subset of American National Standard FORTRAN. Only one departure from PFORT remains: for clarity we have retained subscript expressions involving more than one integer variable, as in "N-I+1".)

In addition to presenting the FORTRAN listings for the subroutines, this appendix briefly describes the subroutines and their roles in producing the stem-and-leaf display, discusses machine-dependent numerical aspects of the algorithm, explains the calculations of the depths printed with the display, reports on how stem-and-leaf has been installed as a command in a statistical computing system, and reviews the points at which error checking is desirable.

A1. Organization of the Subroutines

In this implementation the process of producing a stem-and-leaf display for a single batch of data has been modularized into four components (SLDSPY, SLSCAL, SLLEAF, and SLPRNT) and two utility routines (SLSORT and IFLOOR). This particular arrangement provides the flexibility for pre-chosen scaling and for parallel displays using a common set of stems.

SLDSPY (sequenced SLAB) is the driver routine. It takes as input the batch of data and the necessary scratch storage, and it controls the

succeeding steps in producing the display. After calling SLSORT to sort the data, it checks for possible outliers and withholds them from the display. (Outliers are identified according to a simple rule of thumb based on the "hinges" (approximate quartiles) of the batch: set aside any data value further than 1 step beyond the hinge, where a step is 1.5 times the difference between the hinges.) SLDSPY then calls SLSCAL to scale the display, calls SLLEAF to lay out the stems and calculate the leaves, prints a heading and the low outliers, calls SLPRNT to print the display, and finally prints the high outliers. The stem-and-leaf display is communicated from SLLEAF to SLPRNT in four arrays: ISTEMS and LABELS together determine the starting part or stem to be printed on each line, LEAVES contains the single-digit leaves, and ILFCNT gives the number of leaves on each line of the display. Using this structure, one could display several batches side-by-side on the same page by having a common set of stems (in ISTEMS and LABELS) and one set of leaves (as in LEAVES and ILFCNT) for each batch. In practice it usually suffices to use a common scaling for the batches and print the displays one after another, letting the user cut and paste to achieve the desired effect.

SLSCAL (sequenced SLAD) uses the Dixon/Kronmal scaling algorithm to set the interval of values corresponding to all lines in the display at 1, 2, or 5 times a power of 10.

SLLEAF (sequenced SLAF) handles the calculations involved in setting up the starting parts (in ISTEMS and LABELS), converting each data value into a leaf, and placing that leaf on the proper line. The statements from line SLAF1210 to line SLAF1540 may need some explanation. Basically, the statements involving stems implement the implicit numbering of lines in steps of S/U from $\text{sign}(A) \times [|A/S|] \times (S/U)$ to $\text{sign}(B) \times [|B/S|] \times (S/U)$, as described in the section "Stems, Units, and Leaves". The built-in

FORTTRAN function INT(x) has the same effect as $\text{sign}(x) \times [|x|]$, and the variable IS has the value S/U. The statement "KS1 = KSF / 10" discards the low-order digit of KSF and puts the digits of the stem in KS1. That low-order digit is recovered and saved in LAB. If the lowest data value in the display is negative, each negative stem is shifted down by 1 to allow for the stem "-0", which is represented by the value -1 in ISTEMS. When the number of lines per stem (equal to 10/IS) is 2 or 5, this representation of negative stems actually introduces a jump in the implicit numbering scheme, and the test for KS = -10 at line SLAF1530 resets KS so that the next iteration will continue with the stem "+0". The role of the factor "ONE + EPS" in calculating stems and leaves is discussed in the next section.

SLPRNT (sequenced SLAH) prints the numerical unit used in the display and then prints the display, accompanied at the left by the column of depths. More details on these depths are given in a later section. The parameter LINWID (passed from SLDSPY) provides flexibility in producing output for various devices (including interactive terminals and line printers) with different line widths. The value of LINWID is the number of spaces in the output line; if it is less than 24, there will be no room for any leaves, and a value greater than 123 may produce incorrect output (because FORMAT statements provide for at most 100 leaves per line).

SLSORT (sequenced SLAJ) is a conveniently available and fairly straightforward sorting routine.

IFLOOR (sequenced SLAL) is simply the greatest integer or "floor" function required by SLSCAL. (If desired, it could readily be placed in-line in SLSCAL.)

A2. Numerical Aspects

In implementing a semi-graphic technique one might expect to have no difficulty with numerical details. As stem-and-leaf algorithms illustrate, matters are not quite so simple, but fortunately the numerical considerations are few and straightforward. In a naive algorithm the most likely indication of difficulty is the appearance of a data value as an incorrect leaf, possibly on the wrong line. For example, a data value of 6.2 might appear with a leaf of 1 in the computer-produced display. The reason is that most computers represent numbers in a "binary" form (the common bases are 2 and 16) with a fixed number of digits. Since .2, for example, has a non-terminating binary expansion, it must be rounded to fit into the fixed-length computer word, and the result is a number slightly less than .2. The scaling calculations will not necessarily correct for this, and the leaf may appear incorrectly as 1.

Since the floating-point representation error takes the form of a relative error and causes difficulty in calculating leaves only when the error reduces the magnitude of the data value, we use the factor "ONE + EPS" in SLSCAL and SLLEAF to compensate for it. The machine-dependent constant EPS is a small positive number chosen to allow for representation error and the effects of succeeding calculation. For single-precision arithmetic on the IBM 360 or 370 the value $\text{EPS} = 10^{-6}$ is a conservative choice which allows for an error of 1 in the next-to-last of the 6 hexadecimal digits in the floating-point fraction ($16^{-5} = 2^{-20} \approx 10^{-6}$). For a double-precision implementation on the same computers the corresponding value would be $\text{EPS} = 2.5 \times 10^{-16}$.

Another numerical problem can arise when all data values agree in their first several digits. In this case the digits which determine the leaves may be affected by representation and roundoff errors, and a jumbled stem-and-leaf display may result. This is entirely a consequence of the finite-precision floating-point representation, and the best solution is for the user to drop the leading digits common to all data values before entering the data into the computer (otherwise the damage may already have been done). A conservative test for this problem is (in the notation of Section 5) to ask whether $S/\max(|A|, |B|) < 10 \times \text{EPS}$. If so, it is desirable to give the user a warning.

A3. Calculating the Depths

The information necessary to produce the column of depths is readily available by summing the entries of the array ILFCNT, and the calculations are handled in lines SLAH0920 to SLAH0940 in SLPRNT. If we let k_i denote the cumulative count of data values up through line i and n_i denote the number of values on line i (that is, $n_i = \text{ILFCNT}(i)$), then the depth to be printed for line i is the smaller of $k_{i-1} + n_i$ and $n - k_{i-1}$. The exceptional "middle" line is defined as having $|(n - k_{i-1} - n_i) - k_{i-1}| < n_i$ (that is, the count above the line and the count below the line differ by less than n_i -- it is straightforward to show that this happens on at most one line), and the number printed is n_i . To avoid using additional FORMAT statements, it is convenient to print the "depth" value on the "middle" line without parentheses and also to print the value in the depth column when the line has no leaves.

A4. Installation As a Command

The subroutines we have developed are the basis for the STEM command in the instructional package MIT-SNAP [3]. In designing such a command it is important to offer a high degree of automation (so that beginning students need not grapple with a long list of options) as well as considerable flexibility (so that more experienced users can produce the displays they want). Thus in MIT-SNAP the simplest request for a stem-and-leaf display produces the automated version we have described, including lists of outlying values at each end of the batch. Optional parameters enable the user to control

- the maximum number of lines in the display,
- the basic unit (power of ten) for the data,
- the number of lines per stem (only 1,2, and 5 are permitted),
- the cut-off values at each end (data values outside these are listed separately),
- the forced inclusion of all data values in the display, and
- the use of a common scaling for several batches.

Velleman [9] has described some alternative choices and options as used in LEDA.

A5. Error Checking

In preparing a program for general use, one should check for input which would cause errors or lead to garbage as output. An "automated" stem-and-leaf-display routine can avoid most problems by leaving few decisions to the user, but a few checks remain to be made:

- Is the number of data values too small ($n < 4$)?
- Is the working storage (ILFCNT, ISTEMS, LABELS, and LEAVES) large enough ($M \geq MSTEMS$)?

- Is the output line too narrow (LINWID < 33) or too wide (LINWID > 123)?
- Are the largest and smallest data values to be displayed different (XHI \neq XLO)?
- Does the data have so many significant digits that the leaves are likely to be affected by roundoff error?

The parameter IERR, returned by the driver routine SLDSPY, indicates which (if any) of these difficulties has occurred.

Anyone who uses the component subroutines to get a display different from the automated one is expected to assume the responsibility of checking for errors. In MIT-SNAP the STEM command provides much greater flexibility and checks its optional parameters for validity.

A6. FORTRAN Listings

SLAR0430
 SLAR0440
 SLAR0450
 SLAR0460
 SLAR0470
 SLAR0480
 SLAR0490
 SLAR0500
 SLAR0510
 SLAR0520
 SLAR0530
 SLAR0540
 SLAR0550
 SLAR0560
 SLAR0570
 SLAR0580
 SLAR0590
 SLAR0600
 SLAR0610
 SLAR0620
 SLAR0630
 SLAR0640
 SLAR0650
 SLAR0660
 SLAR0670
 SLAR0680
 SLAR0690
 SLAR0700
 SLAR0710
 SLAR0720
 SLAR0730
 SLAR0740
 SLAR0750
 SLAR0760
 SLAR0770
 SLAR0780
 SLAR0790
 SLAR0800
 SLAR0810
 SLAR0820
 SLAR0830
 SLAR0840

ON OUTPUT:

IERR CONTAINS AN INTEGER BETWEEN 0 AND 5 AND IS USED FOR
 ERROR CHECKING. THE FOLLOWING TABLE LISTS POSSIBLE
 VALUES FOR IERR AND THE CORRESPONDING ERRORS.

IERR	ERROR
0	NONE. DISPLAY SUCCESSFUL.
1	N TOO SMALL.
2	N TOO SMALL.
3	LIMIT TOO SMALL OR TOO LARGE.
4	ZERO RANGE FOR DATA IN DISPLAY.
5	LEAF DIGITS TOO NEAR MACHINE ROUND-OFF LEVEL.

TEMPORARY STORAGE:

ILEFNT IS AN ARRAY OF LENGTH M. IT MUST BE DIMENSIONED
 IN THE CALLING PROGRAM. ILEFNT IS USED TO STORE
 THE NUMBER OF LEAVES ON EACH LINE OF THE DISPLAY.

ISTEMS IS AN ARRAY OF LENGTH M. IT MUST BE DIMENSIONED
 IN THE CALLING PROGRAM. ISTEMS IS THE ARRAY OF
 STARTING PARTS OR STEMS FOR THE DISPLAY.

LABELS IS AN ARRAY OF LENGTH M. IT MUST BE DIMENSIONED
 IN THE CALLING PROGRAM. LABELS IS THE ARRAY OF
 LABELS (0, 2, 4, 5, 6, OR 8) FOR THE DISPLAY.

LEAVES IS AN ARRAY OF LENGTH M. IT MUST BE DIMENSIONED
 IN THE CALLING PROGRAM. LEAVES IS THE ARRAY OF
 LEAVES FOR THE DISPLAY.

*****APPLICATION AND USAGE RESTRICTIONS:
 SLDSPY CALLS A SEQUENCE OF SUBROUTINES. THE FOLLOWING IS A BRIEF
 DESCRIPTION OF EACH ROUTINE PRESENTED IN THE ORDER CALLED
 BY SLDSPY.
 SLSORT SORTS THE DATA IN INCREASING ORDER.
 SLSCAL DETERMINES THE SCALE FACTOR AND UNIT FOR THE DISPLAY.
 SLEAF DETERMINES THE STEMS AND LEAVES.
 SLPRNT HANDLES DISPLAY PRINTING.

SLAB0850
 SLAB0860
 SLAB0870
 SLAB0880
 SLAB0890
 SLAB0900
 SLAB0910
 SLAB0920
 SLAB0930
 SLAB0940
 SLAB0950
 SLAB0960
 SLAB0970
 SLAB0980
 SLAB0990
 SLAB1000
 SLAB1010
 SLAB1020
 SLAB1030
 SLAB1040
 SLAB1050
 SLAB1060
 SLAB1070
 SLAB1080
 SLAB1090
 SLAB1100
 SLAB1110
 SLAB1120
 SLAB1130
 SLAB1140
 SLAB1150
 SLAB1160
 SLAB1170
 SLAB1180
 SLAB1190
 SLAB1200
 SLAB1210
 SLAB1220
 SLAB1230
 SLAB1240
 SLAB1250
 SLAB1260

SLOSPY CALCULATES THE 'LO' AND 'HI' DATA VALUES, PRINTS THESE
 VALUES, AND FINDS THE MAXIMUM NUMBER OF LINES IN THE DISPLAY.

***ALGORITHM NOTES:

LINES IN THE SUBROUTINE CONTAINING 'CDP' IN THE FIRST THREE
 COLUMNS WILL HELP THE USER IN CONVERTING THE ROUTINE FROM SINGLE
 TO DOUBLE PRECISION. 'REAL' DEFINITION CARDS MUST BE REMOVED.

IN ADDITION, ALL CONSTANTS INITIALIZED IN THE SUBROUTINE
 SHOULD BE CHANGED FROM 'E' TO 'D' FOR DOUBLE-PRECISION
 CONVERSION.

ROUTINE CONTAINS THE DATA SET REFERENCE NUMBER FOR THE WRITE
 STATEMENTS. IOUNT IS INITIALLY SET TO '6' IN LINE
 SLAB1470, BUT MAY BE CHANGED BY THE USER TO AN INTEGER BETWEEN
 1 AND 99 INCLUSIVE (THE UPPER LIMIT IS SPECIFIED BY THE
 INSTALLATION AND MAY BE LESS THAN 99).

****REFERENCES:

(1) HOAGLIN, D.C. AND MASSERMAN, S.S., 'AUTOMATING
 STEM-AND-LEAF DISPLAYS'. NBER WORKING PAPER SERIES. WORKING
 PAPER NO. 109. NOVEMBER 1975.

(2) TUKEY, J.W., 'SOME GRAPHIC AND SEMI-GRAPHIC DISPLAYS',
 STATISTICAL PAPERS IN HONOR OF GEORGE W. SNEDECOR (I.A.A.
 RANCROFT, EDITOR). AMES, IOWA: IOWA STATE UNIVERSITY
 PRESS, 1972.

****HISTORY:

WRITTEN BY STANLEY MASSERMAN AND DAVID HOAGLIN (NBER COMPUTER
 RESEARCH CENTER) 10 DECEMBER 1974.

DATE LAST MODIFIED: 4 DECEMBER 1975.

****GENERAL:

QUESTIONS AND COMMENTS SHOULD BE DIRECTED TO:
 SUPPORT STAFF MANAGER
 COMPUTER RESEARCH CENTER FOR ECONOMICS AND MANAGEMENT SCIENCE
 NATIONAL BUREAU OF ECONOMIC RESEARCH
 575 TECHNOLOGY SQUARE, NINTH FLOOR
 CAMBRIDGE, MASS. 02139


```

C          CALL SLSCAL (XHI, XLO, MSTEMS, IERR, IS, NSTEMS, SCALE, UNIT)
C
C          :::::::::: CHECK FOR ABNORMAL EXIT FROM SLSCAL ::::::::::
C          IF ( IERR .EQ. 5 ) RETURN
C
C          IF ( M .GE. NSTEMS ) GO TO 80
C          IERR = ?
C          RETURN
C
C          80 CALL SLLEAF (X, ILOW, IHI, N, NN, NSTEMS, IS, SCALE, UNIT,
C          * ILFCNT, ISTEMS, LABELS, LEAVES)
C
C          WRITE (IOUNIT,500) N
C
C          :::::::::: DETERMINE NUMVAL, THE NUMBER OF HIGH OR LOW VALUES
C          ALLOWED PER LINE ::::::::::
C          NUMVAL = (LINWID - 23) / 10
C
C          :::::::::: PRINT LOW VALUES ::::::::::
C          K = ILOW - 1
C          IF ( K .LT. 1 ) GO TO 100
C          J2 = 0
C          IF ( J2 .EQ. K ) GO TO 100
C          J1 = J2 + 1
C          J2 = J2 + MINO (NUMVAL, (K - J1 + 1))
C          IF ( J1 .EQ. 1 )
C          *   WRITE (IOUNIT,600) K, (X (I), I = J1, J2)
C          IF ( J1 .NE. 1 )
C          *   WRITE (IOUNIT,610) (X (I), I = J1, J2)
C          GO TO 90
C
C          100 WRITE (IOUNIT,700)
C          CALL SLPRNT (ILFCNT, ILOW, IOUNIT, ISTEMS, LABELS, LEAVES,
C          * LINWID, N, NN, NSTEMS, UNIT)
C
C          :::::::::: PRINT HIGH VALUES ::::::::::
C          K = IHI + 1
C          J = N - IHI
C          IF ( K .GT. N ) GO TO 120

```

```

SLAB2110
SLAB2120
SLAB2130
SLAB2140
SLAB2150
SLAB2160
SLAB2170
SLAB2180
SLAB2190
SLAB2200
SLAB2210
SLAB2220
SLAB2230
SLAB2240
SLAB2250
SLAB2260
SLAB2270
SLAB2280
SLAB2290
SLAB2300
SLAB2310
SLAB2320
SLAB2330
SLAB2340
SLAB2350
SLAB2360
SLAB2370
SLAB2380
SLAB2390
SLAB2400
SLAB2410
SLAB2420
SLAB2430
SLAB2440
SLAB2450
SLAB2460
SLAB2470
SLAB2480
SLAB2490
SLAB2500
SLAB2510
SLAB2520

```

```

SLAB2530
SLAB2540
SLAB2550
SLAB2560
SLAB2570
SLAB2580
SLAB2590
SLAB2600
SLAB2610
SLAB2620
SLAB2630
SLAB2640
SLAB2650
SLAB2660
SLAB2670
SLAB2680
SLAB2690
SLAB2700
SLAB2710
SLAB2720

      J2 = K - 1
      IF ( J2 .EQ. N ) GO TO 120
      J1 = J2 + 1
      J2 = J2 + MINO ( NUMVAL, ( N - J1 + 1 ) )
      IF ( J1 .EQ. K )
      *   WRITE ( IOUNIT, 601 ) J, ( X ( I ), I = J1, J2 )
      *   IF ( J1 .NE. K )
      *     WRITE ( IOUNIT, 610 ) ( X ( I ), I = J1, J2 )
      GO TO 110

C 120 RETURN
C
500  FORMAT ( 1X // 10X, 27HSTEM-AND-LEAF DISPLAY, N = , I6 /// )
600  FORMAT ( 1X, I5, 13X, 5HLO I , 10F10.4 )
601  FORMAT ( 1X // 1X, I5, 13X, 5HHI I , 10F10.4 )
610  FORMAT ( 24X, 10F10.4 )
700  FORMAT ( 1H )
C
C  :::::::::::  LAST LINE OF SLDSPY  :::::::::::
      END

```

```

SUBROUTINE SLSLAL (XHI, XLO, MSTEMS, IERR, IS, NSTEMS, SCALE,
* UNIT)
C *****PARAMETERS:
C INTEGER MSTEMS, IERR, IS, NSTEMS
C REAL XHI, XLO, SCALE, UNIT
C *****LOCAL VARIABLES:
C INTEGER I, IA, IR, K
C REAL FPS, FIVE, ONE, RLOG2, RLOG5, T, TEN, TWO, Y, ZERO
C *****FUNCTIONS:
C INTEGER IFLOOR, INT
C REAL ABS, ALOG10, AMAX1, FLOAT
C *****:
C *****PURPOSE:
C THIS SUBROUTINE CALCULATES THE 'DEFAULT' SCALE FACTOR AND
C UNIT, THE EXACT NUMBER OF LINES, AND THE NUMBER OF LINES PER
C STEM IN A STEM-AND-LEAF DISPLAY.
C *****PARAMETER DESCRIPTION:
C ON INPUT:
C XHI CONTAINS THE LARGEST DATA VALUE USED IN THE DISPLAY.
C XLO CONTAINS THE SMALLEST DATA VALUE USED IN THE DISPLAY.
C MSTEMS CONTAINS THE MAXIMUM NUMBER OF LINES IN THE DISPLAY.
C THIS VALUE IS CALCULATED IN THE SUBROUTINE SLDSPY.
C ON OUTPUT:
C IERR IS USED BY THE SUBROUTINE SLDSPY FOR ERROR CHECKING.
C IF THE LEAF DIGITS ARE TOO CLOSE TO THE ROUND-OFF
C LEVEL OF THE MACHINE, IERR IS SET TO 5 AND CONTROL IS
C RETURNED TO THE CALLING SUBROUTINE.
C IS, WHOSE INTEGER VALUE IS SCALE/UNIT, CONTAINS EITHER 10, 5,
C OR 2. THE NUMBER OF LINES PER STEM IS (10 / IS).
C NSTEMS CONTAINS THE EXACT NUMBER OF LINES.

```

```

SLAD0010
SLAD0020
SLAD0030
SLAD0040
SLAD0050
SLAD0060
SLAD0070
SLAD0080
SLAD0090
SLAD0100
SLAD0110
SLAD0120
SLAD0130
SLAD0140
SLAD0150
SLAD0160
SLAD0170
SLAD0180
SLAD0190
SLAD0200
SLAD0210
SLAD0220
SLAD0230
SLAD0240
SLAD0250
SLAD0260
SLAD0270
SLAD0280
SLAD0290
SLAD0300
SLAD0310
SLAD0320
SLAD0330
SLAD0340
SLAD0350
SLAD0360
SLAD0370
SLAD0380
SLAD0390
SLAD0400
SLAD0410
SLAD0420

```

C SCALE CONTAINS THE SCALE FACTOR, OR NUMERICAL VALUE OF ONE
C LINE, IN THE DISPLAY.
C
C UNIT CONTAINS THE APPROPRIATE POWER OF 10 FOR THE DISPLAY.
C
C *****APPLICATION AND USAGE RESTRICTIONS:
C SLSAL IS CALLED FROM THE SUBROUTINE SLDSPY. THE CALCULATIONS
C PERFORMED HERE ARE NECESSARY IN THE ACTUAL COMPUTING OF THE STEMS
C AND LEAVES FOR A STEM-AND-LEAF DISPLAY.
C
C THE ALGORITHM USED IN FINDING THE SCALE FACTOR IS FOUND IN (1).
C
C XHI SHOULD BE DISTINCT FROM XLO, AS THE SUBROUTINE COMPUTES THE
C BASE 10 LOGARITHM OF THEIR DIFFERENCE.
C
C *****ALGORITHM NOTES:
C LINES IN THE SUBROUTINE CONTAINING 'CDP' IN THE FIRST THREE
C COLUMNS WILL HELP THE USER IN CONVERTING THE ROUTINE FROM SINGLE
C TO DOUBLE-PRECISION. 'REAL' DEFINITION CARDS MUST BE REMOVED.
C
C IN ADDITION, ALL CONSTANTS INITIALIZED IN THE SUBROUTINE
C SHOULD BE CHANGED FROM 'E' TO 'D' FOR DOUBLE-PRECISION
C CONVERSION.
C
C FPS IS A MACHINE-DEPENDENT TOLERANCE WHICH ALLOWS FOR ERRORS
C INHERENT IN REPRESENTING DECIMAL NUMBERS IN 'BINARY' BASES.
C (IT ALLOWS ROUGHLY THE LAST FOUR BITS OF THE FLOATING POINT
C FRACTION TO BE AFFECTED BY 'REPRESENTATION ERROR'.)
C
C FOR SUCCESSFUL CALCULATION OF THE LEAVES,
C (SCALE / MAX (ABS (XHI), ABS (XLO)) SHOULD BE GREATER
C THAN 10 * FPS.
C
C *****REFERENCES:
C (1) DIXON, W.J. AND KRONMAL, R.A., 'THE CHOICE OF ORIGIN AND
C SCALE FOR GRAPHS', JOURNAL OF THE ACM, VOL. 12 (1965),
C PAGE 259.
C
C
C
C DOUBLE PRECISION DY, FPS, FIVE, ONE, RLOG2, RLOG5,
C * SCALE, T, TEN, TWO, UNIT, XHI, XLO, Y, ZERO

SLAD0430
SLAD0440
SLAD0450
SLAD0460
SLAD0470
SLAD0480
SLAD0490
SLAD0500
SLAD0510
SLAD0520
SLAD0530
SLAD0540
SLAD0550
SLAD0560
SLAD0570
SLAD0580
SLAD0590
SLAD0600
SLAD0610
SLAD0620
SLAD0630
SLAD0640
SLAD0650
SLAD0660
SLAD0670
SLAD0680
SLAD0690
SLAD0700
SLAD0710
SLAD0720
SLAD0730
SLAD0740
SLAD0750
SLAD0760
SLAD0770
SLAD0780
SLAD0790
SLAD0800
SLAD0810
SLAD0820
SLAD0830
SLAD0840

```

COP DOUBLE PRECISION ABS, ALOG10, AMAX1, FLOAT
C
COP ABS (DY) = DARS (DY)
COP ALOG10 (DY) = DLOG10 (DY)
COP AMAX1 (DY) = DMAX1 (DY)
COP INT (DY) = DJINT (DY)
COP FLOAT (I) = DFLOAT (I)
C
C ::::::: CONSTANT INITIALIZATION :::::::
TEN = 10.0E0
FIVE = 5.0E0
TWO = 2.0E0
ONE = 1.0E0
ZERO = 0.0E0
C ::::::: SET EPS, MACHINE-DEPENDENT TOLFRANCE :::::::
EPS = 1.0E-6
C
C RLOG2 = ALOG10 (TWO)
C RLOG5 = ALOG10 (FIVE)
C
C ::::::: THE FOLLOWING LINES DETERMINE K, THE CORRECT EXPONENT
C OF 10 FOR UNIT. :::::::
T = ALOG10 ((XHI - XLO) / FLOAT (MSTEPS))
C ::::::: FLOOR SIMULATES FLOOR FUNCTION :::::::
K = JFLOOR (T)
Y = T - FLOAT (K)
C
C :::::::
C
C NEXT CALCULATE SCALE, UNIT, AND IS, THE CORRECT
C MULTIPLE OF UNIT TO DETERMINE THE VALUE OF EACH LINE
C IN THE DISPLAY.
C
C :::::::
C
C IF ( Y .LE. RLOG5 ) GO TO 10
IS = 10
GO TO 20
10 IS = 5
IF ( Y .LE. RLOG2 ) IS = 2
IF ( Y .GT. ZERO ) GO TO 20

```

SLAD0850
SLAD0860
SLAD0870
SLAD0880
SLAD0890
SLAD0900
SLAD0910
SLAD0920
SLAD0930
SLAD0940
SLAD0950
SLAD0960
SLAD0970
SLAD0980
SLAD0990
SLAD1000
SLAD1010
SLAD1020
SLAD1030
SLAD1040
SLAD1050
SLAD1060
SLAD1070
SLAD1080
SLAD1090
SLAD1100
SLAD1110
SLAD1120
SLAD1130
SLAD1140
SLAD1150
SLAD1160
SLAD1170
SLAD1180
SLAD1190
SLAD1200
SLAD1210
SLAD1220
SLAD1230
SLAD1240
SLAD1250
SLAD1260

SLADI690
SLADI700
SLADI710
SLADI720
SLADI730

* IERR = 5

RETURN

***** LAST LINE OF SLSCAL *****

END

C

C

SLAF0430
 SLAF0440
 SLAF0450
 SLAF0460
 SLAF0470
 SLAF0480
 SLAF0490
 SLAF0500
 SLAF0510
 SLAF0520
 SLAF0530
 SLAF0540
 SLAF0550
 SLAF0560
 SLAF0570
 SLAF0580
 SLAF0590
 SLAF0600
 SLAF0610
 SLAF0620
 SLAF0630
 SLAF0640
 SLAF0650
 SLAF0660
 SLAF0670
 SLAF0680
 SLAF0690
 SLAF0700
 SLAF0710
 SLAF0720
 SLAF0730
 SLAF0740
 SLAF0750
 SLAF0760
 SLAF0770
 SLAF0780
 SLAF0790
 SLAF0800
 SLAF0810
 SLAF0820
 SLAF0830
 SLAF0840

THE NUMBER OF LINES PER STEM IS 10/IS. IS IS COMPUTED
 IN SLSCAL.

SCALE CONTAINS THE SCALE FACTOR, OR NUMERICAL VALUE OF EACH
 LINE. SCALE IS ALSO COMPUTED IN SLSCAL.

UNIT CONTAINS THE APPROPRIATE POWER OF 10 FOR THE DISPLAY.
 UNIT IS COMPUTED IN SLSCAL.

ON OUTPUT:

ILFCNT IS AN ARRAY OF LENGTH NSTEMS. IT CONTAINS THE NUMBER
 OF LEAVES ON EACH LINE OF THE DISPLAY. ILFCNT IS USED
 BY THE SUBROUTINE SLPRNT IN DETERMINING HOW THE LEAVES
 ARE TO BE PLACED ON THE STEMS.

ISTEMS IS AN ARRAY OF LENGTH NSTEMS, CONTAINING THE 'STEM'
 PART OF EACH LINE IN THE DISPLAY.

LABELS IS ALSO AN ARRAY OF LENGTH NSTEMS. STORED IN
 LABELS ARE THE VALUES 0, 2, 4, 5, 6, OR 8 WHICH ARE
 USED BY SLPRNT TO DETERMINE THE 'TAG' (T, F, OR S) TO
 BE PLACED ON EACH LINE.

LLEAF IS AN ARRAY OF LENGTH NN CONTAINING THE 'LEAF'
 PART OF EACH DATA VALUE IN THE DISPLAY.

***APPLICATION AND USAGE RESTRICTIONS:
 SLLEAF IS CALLED BY THE SUBROUTINE SLUSPY. IT USES THE BATCH OF
 DATA SORTED BY SLSORT AND THE OUTPUT OF SLSCAL (NSTEMS, IS,
 SCALE, AND UNIT) TO PRODUCE THE STEMS, LABELS, AND
 LEAVES FOR THE DISPLAY. THESE ARRAYS ARE PASSED TO SLPRNT FOR
 PRINTING THE STEM-AND-LEAF DISPLAY.

***ALGORITHM NOTES:
 LINES IN THE SUBROUTINE CONTAINING 'CDP' IN THE FIRST THREE
 COLUMNS WILL HELP THE USER IN CONVERTING THE ROUTINE FROM SINGLE
 TO DOUBLE-PRECISION. 'REAL' DEFINITION CARDS MUST BE REMOVED.

IN ADDITION, ALL CONSTANTS INITIALIZED IN THE SUBROUTINE
 SHOULD BE CHANGED FROM 'E' TO 'D' FOR DOUBLE-PRECISION
 CONVERSION.

SLAF1270
 SLAF1280
 SLAF1290
 SLAF1300
 SLAF1310
 SLAF1320
 SLAF1330
 SLAF1340
 SLAF1350
 SLAF1360
 SLAF1370
 SLAF1380
 SLAF1390
 SLAF1400
 SLAF1410
 SLAF1420
 SLAF1430
 SLAF1440
 SLAF1450
 SLAF1460
 SLAF1470
 SLAF1480
 SLAF1490
 SLAF1500
 SLAF1510
 SLAF1520
 SLAF1530
 SLAF1540
 SLAF1550
 SLAF1560
 SLAF1570
 SLAF1580
 SLAF1590
 SLAF1600
 SLAF1610
 SLAF1620
 SLAF1630
 SLAF1640
 SLAF1650
 SLAF1660
 SLAF1670
 SLAF1680

```

C      LAB = IARS (KSF - KS1 * 10)
C      ::::::: SUBTRACT 1 TO ALLOW FOR -0 STEM :::::::
C      IF ( X (ILOW) .GE. ZERO ) GO TO 20
C      KS1 = KS1 - 1
C      KS = KS - 10
    
```

```

C      20 ISTEMS (1) = KS1
C      LEAVES (1) = KLT
C      LABELS (1) = LAB
C      ILECNT (1) = 1
C      :::::::
    
```

THE STEMS AND THE LABELS ARE DETERMINED FOR THE REMAINING ENTRIES IN THE ARRAYS ISTEMS AND LABELS, BY INCREMENTING THE FIRST STEM AND LABEL.

```

C      :::::::
C      DO 30 I = 2, NSTEMS
C      KS = KS + IS
C      ISTEMS (I) = KS / 10
C      LABELS (I) = IARS (KS - ISTEMS (I) * 10)
C      ILECNT (I) = 0
C      ::::::: ADJUST KS TO ALIGN LABELS OF POSITIVE AND
C      NEGATIVE STEMS :::::::
C      IF ( KS .EQ. (-10) ) KS = -IS
    
```

```

C      30 CONTINUE
C      :::::::
C      THE NEXT SECTION CALCULATES THE REMAINING LEAVES AND
C      PLACES THEM ON THE CORRECT LINE OF THE DISPLAY BY
C      ADJUSTING ILECNT.
    
```

```

C      :::::::
C      NI = ILOW + 1
C      DO 60 I = NI, NIH
C      IF ( X (I) .NE. ZERO ) GO TO 40
C      KS = 0
    
```



```

SLAH0850
SLAH0860
SLAH0870
SLAH0880
SLAH0890
SLAH0900
SLAH0910
SLAH0920
SLAH0930
SLAH0940
SLAH0950
SLAH0960
SLAH0970
SLAH0980
SLAH0990
SLAH1000
SLAH1010
SLAH1020
SLAH1030
SLAH1040
SLAH1050
SLAH1060
SLAH1070
SLAH1080
SLAH1090
SLAH1100
SLAH1110
SLAH1120
SLAH1130
SLAH1140
SLAH1150
SLAH1160
SLAH1170
SLAH1180
SLAH1190
SLAH1200
SLAH1210
SLAH1220
SLAH1230
SLAH1240
SLAH1250
SLAH1260

C
DO 50 I = 1, NSTEMS
C      : : : : : ISTEM IS THE CURRENT STEM : : : : :
C      ISTEM = ISTEMS (I)
C      : : : : : KS SPECIFIES THE CURRENT LINE IN THE CYCLE : : : : :
C      KS = LABELS (I)
C      : : : : : ID IS THE DEPTH OF THE CURRENT LINE : : : : :
C      ID = MINO (IDEPH + ILFCNT (I), N - IDEPTH)
C      IF (IARS (N - (IDEPH + ILFCNT (I)) - IDEPTH) .LT. ILFCNT (I))
C      *      ID = ILFCNT (I)
C      IF ( ILFCNT (I) .EQ. 0 ) GO TO 20
C      KL1 = KL2 + 1
C      KL2 = KL2 + ILFCNT (I)
C      : : : : : CHECK THAT LEAF COUNT IS NOT GT LINE WIDTH : : : : :
C      KLEND = KL2
C      IF ( ILFCNT (I) .GT. LEWID ) KL2 = KL1 + LEWID - 1
C      IF ( ISTEM .EQ. (-1) ) GO TO 10
C      IF ( ISTEM .LT. 0 ) ISTEM = ISTEM + 1
C      : : : : : PRINTING OF LINES WITH LEAVES ON THEM : : : : :
C
C      IF ( KS .EQ. 0 ) WRITE (IUNIT,501) ID, ISTEM,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 5 .OR. KS .EQ. 8 ) WRITE (IUNIT,505) ID,
C      ISTEM, (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 2 ) WRITE (IUNIT,502) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 4 ) WRITE (IUNIT,503) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 6 ) WRITE (IUNIT,504) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      GO TO 40
C      : : : : : SPECIAL ARRANGEMENTS FOR THE -0 STEM : : : : :
C
C      IF ( KS .EQ. 0 ) WRITE (IUNIT,510) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 5 .OR. KS .EQ. 8 ) WRITE (IUNIT,515) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 2 ) WRITE (IUNIT,512) ID,
C      (LEAVES (J), J = KL1, KL2)
C      *
C      IF ( KS .EQ. 4 ) WRITE (IUNIT,513) ID,

```


SLAH1690
SLAH1700
SLAH1710
SLAH1720
SLAH1730
SLAH1740
SLAH1750
SLAHL760
SLAH1770
SLAH1780
SLAH1790

```
C 510 FORMAT (1X, I5, 12X, 6H-0, I, 100I1)
C 512 FORMAT (1X, I5, 12X, 6H T I, 100I1)
C 513 FORMAT (1X, I5, 12X, 6H F I, 100I1)
C 514 FORMAT (1X, I5, 12X, 6H S I, 100I1)
C 515 FORMAT (1X, I5, 12X, 6H-0, I, 100I1)
C 600 FORMAT (24X, 100I1)
C
C :::::::::: LAST LINE OF SLPRNT ::::::::::
C END
```

```

C      SUBROUTINE SLSORT (V1, N)
C      ***PARAMETERS:
C      INTEGER N
C      REAL V1 (N)
C      ***LOCAL VARIABLES:
C      INTEGER I, J, K, L, M
C      REAL X
C      ***FUNCTIONS:
C      NONE
C      ::::::::::::::::::::::::::::::::::::::::::::
C      ***PURPOSE:
C      THE SUBROUTINE PERFORMS A SHELL SORT. THIS ALGORITHM IS
C      TAKEN FROM (1).
C      ***PARAMETER DESCRIPTION:
C      ON INPUT:
C      V1 CONTAINS THE ELEMENTS OF THE BATCH TO BE SORTED. AS
C      THIS ALGORITHM SORTS DESTRUCTIVELY, THE ORIGINAL ORDER
C      OF THE ELEMENTS IS LOST.
C      N MUST BE SET TO THE NUMBER OF ELEMENTS IN V1.
C      ON OUTPUT:
C      V1 CONTAINS THE SORTED ELEMENTS OF THE BATCH.
C      ***APPLICATION AND USAGE RESTRICTIONS:
C      SLSORT IS CALLED FROM THE SUBROUTINE SLDSPY. THE SORTING
C      PERFORMED HERE IS NECESSARY IN DETERMINING THE STEM-AND-LEAF
C      DISPLAY.
C      ***ALGORITHM NOTES:
C      LINES IN THE SUBROUTINE CONTAINING 'CDP' IN THE FIRST THREE
C      COLUMNS WILL HELP THE USER IN CONVERTING THE ROUTINE FROM SINGLE
C      TO DOUBLE-PRECISION. 'REAL' DEFINITION CARDS MUST BE REMOVED.
C      ***REFERENCES:
C      (1) HOOTHROYD, J., ALGORITHM 201, COMMUNICATIONS OF THE ACM,

```

```

SLAJ0010
SLAJ0020
SLAJ0030
SLAJ0040
SLAJ0050
SLAJ0060
SLAJ0070
SLAJ0080
SLAJ0090
SLAJ0100
SLAJ0110
SLAJ0120
SLAJ0130
SLAJ0140
SLAJ0150
SLAJ0160
SLAJ0170
SLAJ0180
SLAJ0190
SLAJ0200
SLAJ0210
SLAJ0220
SLAJ0230
SLAJ0240
SLAJ0250
SLAJ0260
SLAJ0270
SLAJ0280
SLAJ0290
SLAJ0300
SLAJ0310
SLAJ0320
SLAJ0330
SLAJ0340
SLAJ0350
SLAJ0360
SLAJ0370
SLAJ0380
SLAJ0390
SLAJ0400
SLAJ0410
SLAJ0420

```

SLAJ0430
SLAJ0440
SLAJ0450
SLAJ0460
SLAJ0470
SLAJ0480
SLAJ0490
SLAJ0500
SLAJ0510
SLAJ0520
SLAJ0530
SLAJ0540
SLAJ0550
SLAJ0560
SLAJ0570
SLAJ0580
SLAJ0590
SLAJ0600
SLAJ0610
SLAJ0620
SLAJ0630
SLAJ0640
SLAJ0650
SLAJ0660
SLAJ0670
SLAJ0680
SLAJ0690
SLAJ0700

VOLUME 6 (1963), PAGE 445.

.....

C
C
C
C
C
C

DOUBLE PRECISION V1 (N), X

I = 1
1 I = I + 1
IF (I .LE. N) GO TO 1
M = I - 1
2 M = M / 2
IF (M .EQ. 0) RETURN
K = N - M

C

DO 4 J = 1, K
L = J
3 IF (L .LT. 1) GO TO 4
IF (V1 (L + M) .GE. V1 (L)) GO TO 4
X = V1 (L + M)
V1 (L + M) = V1 (L)
V1 (L) = X
L = L - M
GO TO 3

4 CONTINUE

C

GO TO 2
..... LAST LINE OF SLSORT
END

C

```

C      FUNCTION IFLOOR (Y)
C      ***PARAMETER:
C      REAL Y
C      ***LOCAL VARIABLES:
C      INTEGER I
C      REAL ZERO
C      ***FUNCTIONS:
C      INTEGER INT
C      REAL FLOAT
C      ::::::::::::::::::::::::::::::::::::
C      ***PURPOSE:
C      THIS FUNCTION SIMULATES THE FLOOR FUNCTION.
C      ***PARAMETER DESCRIPTION:
C      ON INPUT:
C      Y CONTAINS A REAL NUMBER. IFLOOR RETURNS THE LARGEST INTEGER
C      LESS THAN OR EQUAL TO Y.
C      ***APPLICATION AND USAGE RESTRICTIONS:
C      NONE
C      ***ALGORITHM NOTES:
C      LINES IN THE FUNCTION CONTAINING 'CDP' IN THE FIRST THREE
C      COLUMNS WILL HELP THE USER IN CONVERTING THE FUNCTION FROM SINGLE
C      TO DOUBLE-PRECISION. 'REAL' DEFINITION CARDS MUST BE REMOVED.
C      IN ADDITION, ALL CONSTANTS INITIALIZED IN THE SUBROUTINE
C      SHOULD BE CHANGED FROM 'F' TO 'D' FOR DOUBLE-PRECISION
C      CONVERSION.
C      ::::::::::::::::::::::::::::::::::::
C      CDP DOUBLE PRECISION DY, Y, ZERO
C      CDP DOUBLE PRECISION FLOAT
C      CDP INT (DY) = IDINT (DY)
C      CDP FLOAT (I) = DFLOAT (I)
C
SLAL0010
SLAL0020
SLAL0030
SLAL0040
SLAL0050
SLAL0060
SLAL0070
SLAL0080
SLAL0090
SLAL0100
SLAL0110
SLAL0120
SLAL0130
SLAL0140
SLAL0150
SLAL0160
SLAL0170
SLAL0180
SLAL0190
SLAL0200
SLAL0210
SLAL0220
SLAL0230
SLAL0240
SLAL0250
SLAL0260
SLAL0270
SLAL0280
SLAL0290
SLAL0300
SLAL0310
SLAL0320
SLAL0330
SLAL0340
SLAL0350
SLAL0360
SLAL0370
SLAL0380
SLAL0390
SLAL0400
SLAL0410
SLAL0420

```

SLAL0430
SLAL0440
SLAL0450
SLAL0460
SLAL0470
SLAL0480
SLAL0490
SLAL0500
SLAL0510

```
C ::::::::::: CONSTANT INITIALIZATION :::::::::::
ZERO = 0.0E0

C IFLOOR = INT (Y)
IF ( Y .LT. ZERO .AND. Y .NE. FLOAT (IFLOOR) )
*   IFLOOR = IFLOOR - 1
RETURN
C ::::::::::: LAST LINE OF FUNCTION IFLOOR :::::::::::
END
```