

OMNICALC 2 FOR THE TIMEX 2068

The 'tom2' and 'tom64000' programs are exactly equivalent to the Spectrum versions except that:

1. Interface 1 / Microdrive commands are not available.
2. The amount of space available for models is smaller because of the switching code which the Timex uses. This effectively reduces the maximum size model by about 800 cells.

LIST OF CONTENTS

Introduction

A Worked Example

Loading the Program from Tape

Using Full Size Printers

Putting Information In

Positions & Equations

List of Functions

Errors and Unexpected Results

Hints on Using the Mathematical Functions

The Once/Repeat Options in the Calculate Function

*MICROSPHERE COMPUTER SERVICES LTD 1984
72 Rosebery Road, London N10 2LA. Tel.01-883 9411*

OMNICALC 2 **USER GUIDE**

Introduction

OMNICALC 2 is a spreadsheet analysis program designed to make full use of the Sinclair Spectrum. It gives you access to a grid of numbers which can be up to 99 columns across or 250 rows deep. All the mathematical functions of the Spectrum can be used, making it easy to build up sophisticated models for financial, engineering and many other applications. You can even set up conditional expressions, so such things as tax bands, differential interest rates and variable overheads can be accommodated easily.

The program has been designed to be simple to use. All input is prompted. The keyboard is programmed to check each character as you put it in, giving you a different sound to normal if it detects an error. Other errors are handled by displaying meaningful error messages.

You can set up headings (both for rows and columns). Whenever any data in a row or column is being displayed, the relevant heading is also displayed, so you know precisely where you are at all times.

In most applications, many values are derived in fundamentally the same way. For example all values in a "VAT" column may simply be 15% of values in a sales column. OMNICALC 2 recognises this. It incorporates extensive repeat facilities, so rather than laboriously keying in many similar equations, you can simply put in a general expression and repeat this all the way down the column.

Full totalling and sub-totalling facilities are present, and a special function exists to enable you to put in the months of the year quickly. All these features will help you to minimise the effort it takes to build up a model, giving you more time to experiment with it.

You can choose to display numbers either with no decimal places or to two decimal places (although the numbers are held to a greater accuracy to avoid the build-up of rounding errors). You can print out your results, both to the ZX-printer and to a full-size printer connected with a suitable interface (and OMNICALC 2 can adjust to the increased line length on full-size printers, to produce professional reports).

OMNICALC 2 has many other useful features. The ability to link models through an independent workarea; to solve iteratively; to use microdrives and modems through the RS232 port on the Sinclair Interface 1; and to draw histograms of your results and then print them out, make OMNICALC 2 an exceedingly powerful tool.

The remainder of this manual should tell you all you need to know to realise its full potential. The next section gives you a worked example. You are advised to sit down with your Spectrum and work through it. It will not take very long and will give you a good understanding of how to use the program.

OMNICALC 2 USER GUIDE

A Worked Example

The easiest way to illustrate how to use the program is to go through a simple example. Imagine that you are a manufacturer, and you expect sales of 1000 in January and thereafter you expect sales to grow by between 1% and 2% each month until the end of the year. Imagine also that your net costs are 90% of sales value, and that you are interested in seeing your profit in each month and your cumulative profit at the end of each month.

In this example you will need 4 columns (sales; costs; profit and cumulative profit). But let us include a couple of extra columns for any constants we need (1 column for the value of the constant and 1 for some text to describe it. It is a good idea to keep your constants together and separate from the rest of the matrix). You will also need 12 rows (one for each month). You might also like to include a total for sales and costs, making 13 lines in all.

You can now switch on your Spectrum and load the program. (Follow the instructions in the section called "Loading the program from tape", at this stage loading program "tom2"). After the load the first thing that you will see will be the prompt "Load/Init?" As you want to start from scratch you press **I**. The prompt "no. of columns?" will then appear, and you will see a flashing L at the bottom of the screen, which means that the program will expect you to press, ENTER when you have completed your input. Enter 6 and press ENTER. You will then see the prompt "no. of rows?" You will only need 13 lines but the program displays 15 lines at a time, and this is the minimum number that it will accept, so enter 15 (and press ENTER).

The program will then display the grid. You will see columns 1 -3 and rows a-o. The grid cursor can be seen flashing in box a1. At the bottom you will notice that the prompt "function?" has appeared.

When the "function?" prompt is being displayed, pressing **5**, **6**, **7** or **8** will move the cursor around the screen. (The keys have arrows on them to show which direction the cursor will move in). If you hold the shift key down at the same time, then the whole screen will move to show you the next line or column in the direction the arrows point.

Always remember that if you press something by mistake, pressing the EDIT key (shifted 1) will get you out of trouble.

Let us start by putting in the heading information. Line the cursor up in column 1 and move it up as far as it will go. This will take it up above the horizontal line that contains the column numbers. Now press **I** (for Input). In reply to the prompt "text?" type in SALES. When you press ENTER you will see that SALES now appears in the box. Note that all text is put in capitals (the program will not let you put in small letters) so that you never get confused between text and the row numbers (which are always small letters). Now move the cursor to the right into column 2, press **I** and type in COST. Similarly move the cursor to the right again, press **I** once more, and

type in PROFIT. Finally move the screen so that column 4 appears (use shifted 8), put CUM in the top box in column 4 and then move the cursor down to the second box in the column and type in PROFIT. The column headings are now set up. Whenever any data in the column is displayed, the associated column heading will also be shown on the screen so that you know precisely what the column refers to.

For the row headings, move the cursor down to row a, and then across to the far left of the screen. Press **I** and when the prompt "text?" appears, put in JAN.

You now want to put in FEB in row b. MAR in row c etc. OMNICALC 2 has a special feature to allow you to put in the months of the year quickly. Simply press **Y**. Nothing will appear to happen, because the results of this function are not shown on the screen until the screen is rebuilt. But, if you look at the top left hand corner of the screen you will see that a Y has appeared. (The function being executed or just finished always appears in this position so that you can see what you are doing, in case you press a function by mistake).

If you now cause the screen to be rebuilt (press **G** for Goto and then key in a1 in response to the "position" prompt) you will see the other months of the year appear.

You are now ready to start putting in the other mathematical parts of the model. Position the cursor over box a1 and press **I**. This time the prompt "Equation/Value/Blank/Text?" will appear, as the computer has worked out that this is a position in the body of the matrix, and all these possibilities exist. You want to put in a one-off value (the sales in January will be 1000, remember) so press **V**, and in reply to the "value?" prompt type 1000.

Now you expect the growth to be between 1% and 2% each month, so presumably you will want to try various values between 1% and 2%. The easiest way to accomplish this is to put the growth in a box somewhere and then to refer to this as a constant in subsequent equations. (The alternative is to incorporate it into the equations directly, but this would involve changing equations whenever you wanted to change the growth, rather than changing a single value in a box). Move the cursor until it is resting over box a6, press **I** for input, then "**V**" and then key in **1** (assuming that you want to start off at 1%. It would be nice to be able to put a label near this constant so that you do not forget what it is, so move the cursor to the left to box a5 and press **I**. If you reply **T** to the next prompt, the prompt "text?" will appear. You can now fill this box with text. Type in GROWTH

Now move the cursor to box b1. This corresponds to the sales in February. The sales in February will be GROWTH% higher than the sales in January and so you want to put an equation in to reflect this. Press **I**, reply **E** to the "Equation/Value/Blank/Text?" prompt and then type in:

$$a1*(1 + ka6/100)$$

In this equation the 100 is there to turn the growth (a6) from a percentage into a decimal. ka6 means the constant at a6. (The k in front of the grid reference denotes that it is a constant), a1 means the value in box a1, but the computer has stored this as

"the value in the box immediately above". In other words it has worked out a relative location, because you have not declared it as a constant.

After typing in the equation when you press ENTER, the computer will do the calculation and the value 1010 will appear in box b1.

If you now look at box c1, the equation you would write to derive it would be

$$b1*(1 + ka6/100)$$

This looks remarkably like the equation you wrote for box b1. In fact it is the same equation. Because a1 was taken to mean "the value in the box immediately above", if you apply the same equation to c1, the box immediately above is now b1. Notice that because you used k in front of a6, this term will not change, so it will still be taken as meaning "the constant at a6".

Rather than writing a fresh equation for box c1, you can simply repeat the equation for b1. In fact all the remaining values in the column are derived from the same equation. They are all

$$\text{The value in the box above} * (1 + ka6/100)$$

You can therefore repeat this equation all the way down the column. To do this, leave the cursor over box b1 (so the program knows which equation you want to repeat), press **R** (for Repeat), and in reply to the "position?" prompt type in I1 (i.e. L1). (You want the equation repeated down the column to the line corresponding to December, which is line I). Nothing will happen immediately as the computer will not do the calculation until you tell it to by using function C, but it will have checked to make sure that you are not asking it to do the impossible (ie that any relative grid references in the equation correspond to absolute grid references in all new boxes governed by the equation).

If you now press **C**, reply 0 to the "Once/Repeat?" prompt and then reply **R** to the "Row/Column?" prompt, you will see values appear all the way down the column.

Now position the cursor over box c1 and press **D** (for Decode) and you will see the equation $b1 * (1 + ka6/100)$ appear at the bottom of the screen telling you exactly how c1 is calculated and proving that the equation really was repeated.

Now move the cursor to box a2 (in the COST column). The costs are worked out as 0.9 multiplied by the sales, so press **I**, then **E**, and in reply to the "equation?" prompt, type in $0.9 * a1$. The remaining entries in this column take exactly the same relative form so use function **R** to repeat the equation down the column to position I2 (i.e. L2).

Moving on to box a3 in the PROFIT column, the profit is worked out as the sales less the cost, so the equation is simply $a1 - a2$. This relationship holds for the whole column, so input the equation into box a3 and then use function **R** to repeat the equation down to box I3 (i.e. L3).

The fourth column is the CUMULATIVE PROFIT column. The cumulative profit to the end of January is simply the profit made in January, so against the box a4 you can simply enter the equation "a3" (i.e. the value in a4 should equal whatever is in a3). For the second and subsequent months the cumulative profit is the previous cumulative profit plus any new profit made in the month. Against box b4, then you can put in the equation "b3 + a4", and then repeat this equation down to box l4 (i.e. L4).

All that remains is to put in some totals. You may like to leave a line between the months and the totals, so rather than putting the totals in on line m you might use line n (the totalling functions will ignore the blank line in between).

Position the cursor over box n1 and press T (for Total). You wish to total all entries in the column, so in reply to the "Row/Column?" prompt type in C. You want a total, so in reply to the "Total/Sub-tot?" prompt, type in T. The totals will not be done straight away (but only during function C), but the value zero will appear in the box to show that it has been accepted. You can either repeat this procedure for boxes n2 and n3 (there is no point in doing it for n4), or simply use function R (which will also cause totals and sub-totals to be repeated) to repeat the total to position n3.

All the information for the model has now been entered, so you can now calculate. Press function C, then reply 0 to the "Once/Repeat?" prompt, and then R in reply to the "Row/Column?" prompt. The values will be calculated and you should get a total profit figure of 1268.

All figures are currently being displayed in integer mode, with every figure rounded down to the next whole number. You can change the mode in which the figures are displayed to show 2 decimal places by using function M. This will not take effect immediately, but if you press G (for Goto) and then reply a1 to the "position?" prompt, when the screen is rebuilt you will see the more accurate figures. (The Goto function is very useful to move about amongst larger models which may have up to 99 columns or 250 rows).

If this was a real model, you might want to save it (function S) or print it (function P) at this point. How to do this is explained in the List of Functions section.

Let us now go one stage further and produce a histogram from the information in the model.

Firstly we must tell the program which data we wish to plot. We do this by putting it into the "workarea". Press "W" in reply to the "function?" prompt and then "E" in reply to the "Extract/Insert?" prompt. When the "position?" prompt appears, type in:

$$a2 - 12/a3 - 13 \text{ (i.e. } A2 - L2/A3 - L3)$$

and press ENTER

This loads 24 values into the workarea. The first 12 are the costs and the next 12 are the profits.

We can now proceed to actually plot the histogram, so press "GRAPHICS" (shifted 9). The first prompt you will see will be "no. of groups?". In this case we have two groups - costs and profits - so type in 2 and press ENTER. The prompt "no. in group?" will then appear. There are 12 values in each group, so key in 12 and press ENTER.

After this two prompts will appear. The first one is "stack Y/N?" and asks whether you want all values plotted in separate columns, or whether you want them stacked (i.e. each column to contain one value from each group stacked on top of each other). The second prompt is "gap Y/N?" and asks whether you want the columns to be separated out for increased clarity. This first time around reply "N" and "N" to these prompts.

You should now have a histogram on the screen, consisting of 24 columns, alternating between solid columns and stippled columns. The solid columns refer to the first group, which in this case is costs, and the stippled columns to the second group, the profits. (If there was a third group, this would be shown as outlined). On the left of the screen are numbers showing the scale of the histogram.

If you now press EDIT (shifted 1) to escape from the histogram and then plot the equivalent stacked histogram (press GRAPHICS, then reply 2 to the "no. in groups?" prompt; 12 to the "no. in group?" prompt; Y to the "stacked Y/N?" prompt and N to the "gap Y/N?" prompt) you will notice several differences.

There are now only 12 columns, but each column has two parts - a solid base with a stippled area on top. As before the solid group represents the first group plotted i.e. the costs and the stippled area represents the second group, i.e. the profit. Notice that the height of each part of a column represents one value, and the total height of each column represents the sum of one cost and one profit value, which in this case equates to sales (costs + profit = sales).

You might at this stage like to experiment and produce the histograms again, only this time replying Y to the "gap Y/N?" prompt.

On all the histograms you will notice that apart from the columns themselves and the scaling numbers, there is nothing else on the screen, except a flashing cursor.

If you wanted to produce a hard-copy of the histogram, you could move the cursor around the screen, adding text to produce a very professional, easy-to-read diagram.

This was a very simple model. You can try experimenting with different growths by changing the values in box a6 and then recalculating. If you consider what you have done, however, you will see that you have succeeded in calculating some 50 values, and presenting this in a clear form all in the space of a few minutes. Furthermore, changes in the base information (the growth and the expected sales in January) can be fed in very quickly and you can then produce a new set of results within seconds.

OMNICALC 2 **USER GUIDE**

Loading the Program From Tape

The OMNICALC 2 tape has 2 programs on it. These are:

"tom2" and "tom64000"

The two programs are identical in function, the only difference between them being that "tom2" occupies the very top of the Spectrum memory, whereas "om64000" leaves the bytes between 64000 and 65535 free in case you require complex printer interface software that will not fit into the 256 byte LPRINT buffer (but this means of course that "tom2" can handle slightly larger models than "tom64000").

To load one of these programs from tape proceed as follows.

1. Set the volume control of your cassette recorder to somewhere near the middle. Rewind the tape to the beginning.
2. Connect the EAR socket of the cassette recorder to the EAR socket of the computer. Make sure that there is nothing in the MIC socket of the computer.
3. Type in either LOAD "tom2" or LOAD "tom64000" depending on which version of the program you want. (The load is obtained simply by pressing the J key, the quotes are obtained by pressing the SYMBOL SHIFT and the P key at the same time, and you must press ENTER at the end of the command.
4. If the program fails to load, try again at a different volume level.

Note that if you are loading code either above address 64000 or in the LPRINT buffer to control a printer interface, you should load such code before loading "tom2" or "tom64000".

When the program has loaded, the prompt "Load/lnit?" will appear on the screen.

Using Full Size Printers

OMNICALC 2 is designed to support full size printers. It allows you to tell it how many columns your printer can accept, and will then format print-outs to this width.

Most Centronics-type printer interfaces consist of some hardware and some driving software. If your interface needs printer driver software, then you have a choice as to where you put it. You can either slot it into the LPRINT buffer (if it will fit) and use program "tom2", or you can load it anywhere from address 64000 upwards and use program "tom64000", (remembering that using "tom64000" gives you slightly less space for your spreadsheet). Most interface suppliers can supply software that fits in the LPRINT buffer but this may not include routines to enable you to take a hard-copy of the screen. In general, if your driver software has text-printing and screen-copying routines that can both fit in the LPRINT buffer, or if the text-printing routine does fit and you do not require a printed copy of your histograms, then use "tom2". Regardless of which program you use, remember that you should load the driver software before you load the program.

Kempston produce a Centronics-type interface that is hardware only. Unfortunately there are two versions of this interface. To get hard-copies of your histograms from the most common version you should initialise the interface by keying in the COPY command relevant to your printer, and only then load "tom2". If this does not work then you must have the earlier version and Kempston will be able to advise you.

If your printer supports graphics, and if you possess suitable screen-copy software, you can produce hard-copies of histograms produced in the program. To do this you have to tell the program where your hard-copy routine starts, by using option G of function X. (The Kempston hardware interface case has already been discussed. For other interfaces, the address you need to tell the program is the address in the RAND USR xxxxx statement that you would use to get a screen copy from BASIC).

Putting Information In

All input to the program is done in response to prompts. Prompts can be split into two types:- those requiring just one character to be put in, and those requiring a string of characters.

If just one character is required, then no cursor will appear at the bottom of the screen, and when the character has been pressed, there is no need to press ENTER.

There are a number of situations where you are asked to make a choice from several alternatives and in these cases when the prompt is displayed, you will see a series of keywords. To make your choice, you should press the key corresponding to the first letter of the desired keyword.

If a string of characters is required, then a cursor (a flashing "L") will appear at the bottom of your screen. As you put in information, this will be displayed at the bottom. When you have finished putting in the information you must press ENTER before the computer will pay any attention to it. You are not allowed to put in nothing, ie. there must be at least one character before you press ENTER. Depending on the prompt you are responding to, the program will limit the characters you are permitted to put in (eg. if it expects a number, it will only allow you to use the keys 0 - 9). You can, however, always move the cursor about within the input using the shifted 5 and shifted 8 keys, and delete characters using the shifted 0 key (ie. the normal Spectrum editing features are supported).

Input at the bottom of the screen can be up to 3 lines long. You will normally only need more than one line if you are putting in a complex equation. If you put in a character that would take the input over 3 lines, then the input will be totally rejected and the error message TOO LONG will be displayed. Other errors may be handled in two ways. If the input has format or syntax errors, then the input will be redisplayed with a marker showing where the error was detected, and you should correct the error and reinput. If, on the other hand, the input is syntactically correct, but is illogical, the input will be rejected and an appropriate error message displayed.

In general, you can terminate input and cancel everything you have done so far in the function you are using by pressing the EDIT key (shifted 1). The only exceptions to this are when you first enter the program, and at certain points during subsequent

loads/initialisations where the computer is not holding sufficient information to be able to do any further work.

Positions & Equations

The information held by the program is laid out in the form of a grid with numbers representing the columns, and letters the lines. Any position in the grid can be defined by a grid reference eg. a1 the first line 1st column position, c8 the third line 8th column position etc. Note that after all the single letters have been used up, then double letters are used, ie. line 27 is represented by the letters aa, line 28 by ab (all the way up to line 250 being represented by the letters ip if your grid is this deep).

Whenever the program prompts for a position you should supply the grid reference in this form.

Grid references are also used in equations. For example, if you want the value at grid reference a2 to be twice the value at grid reference a1, you would put an equation in to derive a2 as "2 * a1". The program will thereafter always multiply a1 by 2 whenever it is working out what a2 should be.

One of the features of OMNICALC 2 is that this equation is relative ie. rather than storing it in the simple form 2 * a1 it will have taken it to mean 2 * the value in the box immediately to the left. Thus if you want b2 to equal 2 * b1, you can use exactly the same equation. In fact in most models you will set up with OMNICALC 2 you will find that there are relatively few basic equations, and that these equations are used to derive whole rows or whole columns.

There are occasions when you may not want a grid reference to be relative, but rather you want the value at that location to act as a constant in equations. In this case you should put a k in front of the relevant grid reference when you use it in equations. Whenever this equation is then used, it will always pick up the value at that specific reference.

At any time the screen will only show a small portion of the grid - a so called "window" on the grid. This window comprises 3 columns and 15 lines. In addition, the header information relevant to each column and each line will be displayed, so that you can see precisely what each value represents.

List of functions

Unshifted 5, 6, 7 and 8 - moves the cursor around the screen, in the direction shown on the arrows on these keys.

Shifted 5, 6, 7 and 8 - moves the "window" on the data that is currently being displayed (e.g. shifts the whole screen one column to the left, etc.)

Graphics (Shifted 9) - plots a histogram. Before using this function you should ensure that the data to be plotted is put into the workarea (see function W), and that if a hard-copy of the histogram is required, the Graphics address is set up correctly. (See function X).

Values in the range -3000 million to +3000 million can be plotted. Values outside this range will result in the "calc. error" error message being produced. Scaling is done automatically, with very large numbers being shown in thousands (in which case the figures 000 will appear in the top left hand corner of the histogram), or in millions (in which case the letter M will appear in the top left hand corner).

Two formats of histograms are allowed. The histogram can either have one column for every value, or can be "stacked" so that each column contains one value from each group (see below) stacked on top of each other.

The first prompt you will see will be "no. of groups?" This refers to the number of types of item you are plotting. Thus, if you have 12 cost figures and 12 profit figures, there are 2 groups - profit and cost. You are allowed up to 3 groups. On the resultant histogram, values in the first group will show up as solid, values in the second group as stippled, and values in the third group as outlined. Note that the groups should be present in the workarea one after the other, so that for example the first twelve items in the workarea should be costs, the second twelve profits.

The next prompt will be "no. in group?" and you should type in the number of values in each group (so if you have 12 costs and 12 profits the no. in group is 12).

If you have more than one group the prompt "Stack Y/N?" will appear, and you should reply Y or N depending on the desired format.

The program then works out if there are sufficient spare columns to allow you a gap between sets of information. If there is, it will display the prompt "gap Y/N?" and you should reply Y or N accordingly. Note that if there is more than one group and the histogram is not in stacked format, the program will only put gaps after there has been an occurrence of each group. (So that you will see one cost column, one profit column then one gap rather than a cost column, a gap, a profit column, a gap etc.)

The histogram will then be displayed. This is centred on the screen. You can then annotate the histogram as required, by moving the cursor around the screen (by the shifted 5, 6,7 & 8 keys) and typing in the titles you require. You cannot delete characters (you should overwrite with a space instead). To enable you to put in a key to the columns, the characters <=, <>, and >= (reached by symbol shift Q, W & E) have been programmed to produce a representation of the format used in plotting the groups. Capital letters are available by pressing CAPS shift as well as the desired letter.

To produce a hard-copy of the histogram you should press GRAPHICS (shifted 9). Some hard-copy routines (such as the one controlling the ZX printer) do not copy the bottom two lines of the screen, but as the histogram never occupies the bottom 3 lines, there will always be at least one blank line below the histogram for annotations that are to be printed.

To return to the main program you should press EDIT

A - Amend Area - inserts/deletes rows and columns. If you are deleting, the cursor must be over an item in the row or column you want to delete before you enter the function. If you are inserting, the column or row over which the cursor is resting will be duplicated.

You will first see the prompt "Insert/Delete?" and you should reply I or D depending on what you want to do. After this the prompt "Row/Column?" will appear and you should reply R or C accordingly. The program will then check to ensure that you are not reducing or increasing the size of the model beyond acceptable limits (the program requires that there always be at least 15 rows and 3 columns, and never more than 250 rows or 99 columns), and that, in the case of an insertion, there is physically enough space in memory for the new line or column.

It is important to realise that the insertion or deletion will change the effect of the equations in the model. For example, if box a3 is calculated by the equation "a2 * 2" and you delete column 2, the new column 2 will be the old column 3 and if you decode box a2, it will now read "a1 * 2" (ie. it still means the box immediately to the left multiplied by 2). If you go one stage further and delete column 1, the new column 1 will be the old column 2, but now there is a box immediately to the left and the equation makes no sense. You will not only get a calculation error (code 3) should you try to calculate, but also when you try to decode the equation, the grid references will make no sense.

The important thing to remember is that insertions and deletions do not change the equations at all, so that if the equation was relative (ie. if it contains grid references that are not defined as constants), the equation would still be relative, but would now refer to a different box, or at worst a non-existent box.

As has been mentioned, insertions duplicate a column or row. This means that your new row or column can be either of two. This is deliberate, as the program itself cannot decide whether you want the new row or column to be calculated in exactly the same way as the original, or whether you want a blank row or column before the original or after the original. Unless you specifically want the new row or column to be a duplicate of, or at least calculated in the same way as the original, you are recommended to do the insertion and then to put blanks in one or other of the rows or columns to avoid confusion.

C - Calculate -forces a recalculation of all variable values in the spreadsheet, i.e. evaluates all the equations, sub-totals and totals).

Two major options are available. If you reply 0 to the "Once/Repeat?" prompt, all the values in the spreadsheet will be calculated once. If you reply R, then the spreadsheet will be evaluated over and over again, until the value in the last box (ie. the bottom right hand corner box) becomes zero. See the section entitled "The Once/Repeat Options in the Calculate Function" for more details.

You can also decide whether calculation should proceed row by row or column by column, by your reply to the "Row/Column?" prompt. In most cases it does not matter which of these options you choose, but there are cases where they give different results. For example, in a typical cash-flow analysis, the closing balance from one

column becomes the opening balance in the next. If you evaluate such a model row by row, you will be transferring the closing balance values to the opening balances before you have had a chance to calculate the closing balance. Thus, in this case, you should use the column option.

The calculation can be stopped at any time by pressing the BREAK key (but it may take a few seconds to take effect).

D - Decode - used to tell you how the value at the cursor position is derived. This will either be an equation or the codes BL, VL, TX, RT, RS, CT or CS standing for blank, value, text, row total, row sub-total, column total and column sub-total respectively.

G - Goto -used to move the "window" on the data, so that the top left hand corner is the line/column reference you supply in response to the following "position?" prompt. If the position you enter does not exist, then the input will be redisplayed with an error marker showing where the error was detected. If the position exists, but is so far to the right or to the bottom of the grid that a full screen would not be given if this was the left hand corner position, the left hand corner will be deduced, but a screen will be produced which includes the position supplied. In all cases the cursor will be positioned over the position supplied.

I - Input - used to enter information at the location where the cursor is resting. If this is in a heading position (above the black horizontal line with the column reference numbers, or to the left of the vertical band containing the column reference letters), then text will be expected and the prompt "text?" will be displayed. In text mode all letters (which will automatically be treated as capitals), all numbers, all punctuation marks and all arithmetic operators which occupy only one character (+, -, * etc.) are accepted. The graphics mode is not available. Only the first seven characters of any input are accepted (spaces count as a character).

If the cursor is resting over the body of the data, then the prompt "Equation/Value/Blank/Text?" will appear. If the value in the box is dependent on values held at other locations, and can therefore change if these other values change, then the Equation option should be selected. If the value is a once-off value which will not change, then the Value option should be used. The Blank option merely blanks out whatever is in the box, so that nothing is displayed. The Text option is used to put text into the matrix. If the Blank option is selected, the box will be blanked out and the function will end. If one of the other options is used, then the prompts "equation?", "value?" or "text?" will be displayed, depending on the option chosen.

For values and equations a mathematical expression should now be entered. This expression may include position references, arithmetic operators, numbers and most arithmetic functions. (The exceptions being those concerned with string handling). In short, the program will only allow you to put in symbols which could be used in forming a valid mathematical expression. The RND function is supported as are the logical operators (AND, OR and NOT).

Upon entry, the program will validate the expression, checking that any position references are meaningful and that a valid expression has been formed. If an error is detected during these checks, the input will be redisplayed with an error marker and

the cursor positioned after the offending character (or characters). If the expression is valid, then its result will be calculated and displayed. If you have indicated that it is an equation, then the equation will be stored in memory and used in subsequent calculations.

If the text option is chosen, then up to 7 characters of text may be entered. To compress this text into as small a space as possible, there are certain restrictions as to which characters you can use. Letters will automatically be treated as capitals, the '\' character will be treated as the pound sign, and ASCII codes 123 - 127 will be interpreted as ASCII codes 91 - 95.

K - copy - copies a single value, equation, sub-total, total or piece of text from the box where the cursor is resting to a box specified by you in response to the "position?" prompt. Unlike the R function, the new box can be anywhere in the model. Note that the new box is not evaluated straight away, and if it now refers to an equation, or is a sub-total or total, you will have to use the calculate function to see its correct value.

L - Load - this function is also used if you want to Initialise (ie. start from scratch) as both have the effect of overwriting existing information.

The prompt "Load/Init?" will be displayed. If you reply I, the existing spreadsheet will be erased (but not the workarea) and you will no longer be able to escape from the function. The prompt "no. of columns?" will appear. You should reply with the number of columns of data that you expect to use (which must be in the range of 3-99). The prompt "no. of rows?" will then appear, and you should key in the number of rows you require in the model. (The permissible range is 15-250).

The program will then determine if there is sufficient space to accommodate a spreadsheet of this size. (There is space for at least 4000 cells, ie. the number of rows multiplied by the number of columns can be at least 4000).

If there is space, the top left part of the grid will be displayed. Otherwise you will be asked for the "number of columns?" again.

If you choose the Load option by pressing L in response to the "Load/Init?" prompt, you will first see the prompt "Data/Work?" If you want to load spreadsheet data, reply D. If you want to load previously saved values into the workarea, then reply W. You will then be asked for the device you wish to load the information from by the "channel?" prompt. You should reply with a value in the range 0-10, with 0 meaning tape, 1-8 meaning the microdrive on that number, 9 meaning the RS232 port and 10 meaning the ZX-net. Finally you will be asked for the name of the file by the "name?" prompt. You must specify a name, and this can be bordered by quotation marks (in accordance with the normal Spectrum convention) if you wish, otherwise the program will automatically supply them. (If you have forgotten the name of your file you can enter a dummy name, start the tape and, when the program displays the file name from the tape, break into the program through the load again. In the case of a microdrive, the CAT option in function X lets you look at the files on your cartridge.)

When you press ENTER after keying in the name, the program will start the load. If it is successful, then the prompt "verify Y/N?" will be displayed. This verify after the

load is provided so that you can double check that no errors occurred during the load. If you reply Y then you should rewind the tape and set it playing again (if using tape files), or arrange for the save to be repeated on the Spectrum that is sending the file (if you are using the RS232 port or the net.)

Note that the program checks to see if you are using Interface 1 on your Spectrum. If it is not present, it knows that you can only be using tape files, and so does not give you the "channel?" prompt.

When you first load the program, this function is called automatically. However, at this stage, the program has no spreadsheet to work on, and it will assume that the first thing you want to load is some spreadsheet data, and so will not display the "Data/Work?" prompt. (At certain other times the program may also detect that you do not have a spreadsheet in memory, and again will skip the "Data/Work?" prompt.)

You can interrupt the LOAD at any time by pressing the BREAK key.

M - Mode - used to switch between integer and 2-decimal place display modes. When the computer is first switched on, integer mode will be selected. At any time you can switch to 2-decimal place mode by pressing M. Any existing data on the screen will not be changed, but anything new will be displayed in the new format (or if you move the window, or use the Calculate or Goto functions, the data that appears will be in the new format).

If you are already in 2-decimal mode, you can switch back to integer mode by pressing M again. Note that the largest number that can be displayed without over-running into the previous column is 9999999, in integer mode and 9999, in 2-decimal place mode. If the numbers might exceed these limits, then leave a blank column to the left of any suspect column to avoid over-running other figures.

P - Print - used to print either any rectangle from within the model, or all data within the model, depending on your reply to the "Sector/All?" prompt.

If you choose the Sector option, you will first be asked for the "no. of columns?" you wish to print. This refers to the number of columns of data rather than the number of characters that will fit on one line of your printer. You will then be asked for the "no. of rows?" and you should enter the number of rows of data that you want to print. Finally you will be asked for a "position?", and you should reply with the grid reference of the cell that is to be the top left-hand corner of the rectangle to be printed.

If you select the All option, then you will be asked for the "no. of columns?" (Again this refers to data columns and not characters per line of your printer). For this option this is all the information you need to supply for the printer to print out the entire model, (printing as many columns as you specified on the first pass, and then starting again with the next set of columns etc).

For both options, column headings will be printed. The line heading will also be printed and so the following values represent the maximum number of columns that you can use.

ZX printer	3
80 column printer	9
132 column printer	15

You need not always use these maximum values. Even if you have a 132 column printer, you can still print 1 column at a time if you so choose.

Both print options can be stopped at any time simply by pressing BREAK.

Q- Quick - The slowest part of the program is converting between the floating point binary numbers used by the computer and decimal numbers. If you have a model where a lot of data has to be changed every time it is run, then you do not necessarily need to see what is already there to be able to put in the new values correctly.

In quick mode, existing data is not displayed (although the headings are, so that you can ensure that you are in the right place). This means that shifting the window is virtually instantaneous, making entry of the new values quicker. To enter quick mode, you simply use function Q, and to return to normal mode you press Q again. (Note that all data will be displayed following a Calculate or Goto; it is only shifting the window that is speeded up.

Note also that if at any time you appear to have lost some information, it is probably because you are in quick mode).

R - Repeat - used to repeat the contents (and method of derivation) of the box where the cursor is resting, to another position either to the right on the same row, or down in the same column. All intervening positions also take on the same contents (and method of derivation).

Note that if you are repeating an equation, it would normally be relative, so that if the initial equation was $a1 * 2$ and you are repeating it downwards, the value in the next box will be calculated as $b1 * 2$, the next $c1 * 2$ etc. Using this function does not cause automatic calculation of any boxes affected. To actually see what effect the repetition has had, you will have to use the Calculate function.

S - Save - used to save the data of the model, or an extract from the model that has been put into the workarea. You will first be asked whether you are saving the entire spreadsheet data, or just the workarea by the "Data/Work?" prompt, and you should reply D or W accordingly.

You will then be asked for the "name?" of the file you wish to save. This name can be up to 10 characters long, and all letters will be treated as capitals. You can place quotation marks around the name if you wish otherwise the program will supply them automatically.

If you are saving to tape, a prompt telling you to start the tape will be displayed, and you can press any key to start the save.

When the save is completed, the prompt "verify Y/N?" will appear. For saves to tape you would normally wish to verify, and so would press Y (rewinding the cassette and setting it up to play rather than record).

If you do verify, and it is successful, you will be returned to the model. Otherwise you can try verifying again, or repeating the save procedure.

T Total/Sub-total - used to set up totals and subtotals either down a column or across a line. You will first see the prompt "Row/Column?" and you should reply R or C accordingly. Next the prompt "Total/Sub-total?" will appear and you should reply T or S depending on the form of totalling you require.

The rules for how the four permutations work are as follows.

Column Total - accumulates all values above it in this column with the exception of other Column Totals or Column Sub-totals. (It will however add in any Row Sub-totals or Row Totals that it meets).

Row Total - accumulates all values to the left of it in this row with the exception of previous Row Sub-totals and Row Totals. (It will add in any column sub-totals or totals that it meets).

Column Sub-total - works upwards accumulating any values in the column until it reaches either a previous column sub-total or the first line. It ignores any column totals, but will include row subtotals and totals.

Row Sub-totals - works to the left accumulating any values in the row until it either reaches another row sub-total or the first column. It ignores any row totals, but will include column sub-totals and totals.

These rules are so designed that you can use the totals several times in a row or column to give cumulative totals, without affecting sub-totalling or subsequent totalling.

Note that totals and sub-totals are not evaluated immediately, but only during the Calculate function.

W -Workarea - transfer values between the spreadsheet and the workarea.

If you want to pull out information from the spreadsheet into the workarea, reply E to the "Extract/Insert?" prompt. If you want to move information from the workarea back into the spreadsheet, reply I.

The "position?" prompt will then appear. This is asking you which values you want to pull out (option E), or where in the spreadsheet you want to put the values that are in the workarea (option I). You can reply by individual grid references, or groups of grid references. If you have a group of references, then the positions referring to the start and end of the group must be separated by a minus sign. You can pull out several grid references (or groups) providing that you separate them by the "/" character. The workarea can accommodate up to 60 references.

As an example, imagine that you want to pull out the first 10 entries in column 1 (ie. a1 to j1); the top entry in column 2, and the first five entries in column 3 (a3 to e3). In replying to the position prompt, you would type in:

```
a1 -j1/a2/a3 - e3
```

The first 16 positions in the workarea will now contain the values from these 16 positions.

Note that this function assumes that you are dealing with values. If you pull out any grid references containing text, or which are derived by equations, then subsequently use the I option in function W to copy the workarea to the spreadsheet, any equations will be lost and text will be translated into a (odd-looking) value.

The workarea has many uses.

Firstly, it is used in the plotting of histograms (The values you wish to plot must first have been pulled into the workarea).

Secondly, it can be used in moving values around in a spreadsheet. (For example, if you want to transpose 2 columns, you can pull information from the first of the columns into the workarea, delete that column in the spreadsheet, so that the second column now occupies the correct position, then insert a column and put the stored information from the workarea into that column).

Thirdly, as the workarea can be saved and loaded independently, it can be used to communicate between spreadsheets. For example, closing balances from the end of one year can be pulled out from one model to be the opening balance in the following year's model, or summary figures from many divisions can be extracted and used in producing a balance sheet for a corporation.

X – Extended commands - this function allows you to carry out certain microdrive/Interface 1 commands in the Spectrum, while still in the program. It is also used to set up the "graphics address", which is the address of your screen-copy routine and is used in producing hard-copies of histograms.

In reply to the "Open 3/Drop 3/Cat/Erase/Baud/G-ad?" prompt you should reply 0, D, C, E, B or G according to what you want to do.

The Open 3 option assigns an RS232 "b" channel to the printer. This is used when your printer is connected to the Spectrum via the RS232 port, and ensures that any print-out is directed to the printer. Note that it is a "b" channel, (as graphics information from a screen-copy routine can also be required to go to the printer) and so a line feed will not be sent automatically.

The Drop 3 option closes the channel created by the Open 3 option, so that print-out is now sent to the ZX-printer. It is only really needed if you have two RS232 devices, only one of which is a printer, and now want to use the other one (or if you used the Open 3 option by mistake, in which case you should select this option, then press

BREAK and reply N to the "retry Y/N?" prompt. You must press BREAK, because otherwise the computer would wait for ever for an acknowledgement from a non-existent RS232 device).

The Cat option is used to display a list of all the files on a microdrive. You will be asked to stipulate which microdrive by the "channel?" prompt, and should reply with the microdrive number (1 - 8). The catalogue will then be displayed. If there are many files on the cartridge, you may see the prompt "scroll?" at the bottom of the screen, in which case when you have finished looking at this group of files you should press Y, which will display the rest. (If you reply N to the scroll prompt the program may assume that an error has occurred, and you will have to reply N to the "retry Y/N?" prompt to get back to the main program).

When all the files on the cartridge have been displayed, the program will display the prompt "ready Y/N?" and when you press Y the list of files will be cleared from the screen and the spreadsheet will reappear.

The Erase option is used to erase a file from a microdrive cartridge. You will first have to tell the program the location of the file you wish to erase, by giving the relevant microdrive number in reply to the "channel?" prompt. You will then be asked for the "name?" of the file, and should reply accordingly. The program will then attempt to erase the file. Note that if you enter a wrong file name, you will get no indication that the file does not exist and that also, for some reason internal to the Interface 1, the erase does not always work. You should always, therefore, use the Cat option after the erase to ensure that the file is no longer there. You should also note that the program only allows capital letters in files. This is so that you cannot accidentally erase programs or files set up by other programs within OMNICALC 2.

The Baud option allows you to set the baud rate for the RS232 interface from within the program. You will be prompted to put in the baud rate by the "rate?" prompt. You should reply with the highest baud rate that your peripheral can accept.

The G-ad option is used to set up the address of the routine that you are using to produce hard-copies of histograms. When the program is first loaded, this is set up to point to the COPY routine in the Spectrum ROM, and unless you are using the ZX or Alphacom printer, will need to be changed. This is discussed in more detail in the section called "Using full size printers".

To change the graphics address, you should key in the address you want, when the "address?" prompt appears. Please take great care to enter the correct address, as getting the wrong address may cause the program to crash.

Y - Year - provides a fast method of feeding in the months of the year. The cursor should be positioned over a box containing the first three letters of a month of the year (eg.JAN, FEB etc). When function Y is pressed, the remaining months of the year will be set up in boxes to the right (if the box is a column heading).

If it is a line heading, the remaining months will be set up downwards.

If the initial box contains some other information (eg. JAN 83) then the other information will be copied (so that the next box reads FEB 83 etc.)

Errors and Unexpected Results

Error messages

The error messages have been designed to be reasonably self-explanatory. Fuller details follow:

> 60 (functions W
& Graphics) -

If you are using function W, the program has detected that you have asked for more than 60 references in reply to the "position?" prompt, and the cursor will be resting where the program detected this condition.

If it occurs during the GRAPHICS function it means that the "no. of groups?" multiplied by the "no. in group?" has exceeded 60.

CALC. ERROR – (functions C, I, K, R & GRAPHICS) - If this occurs during the GRAPHICS function it means that you have tried to plot a histogram from values outside the acceptable range (which is plus or minus 3000 million). It may be that this is because you have inadvertently put some text into the workarea which is being interpreted strangely.

For the other functions a grid reference at which the error occurred is given and this is followed by the Spectrum error code (see appendix B of the Spectrum manual).

From our experience the following four sources of error are the most common:

- 1) You have been using function A to delete columns and lines, and now have some equations that refer to columns or rows which do not exist. (The error code in this case will be 3.)
- 2) You are trying to take the square root of (or exponentiate) a negative number.
- 3) If the error occurs during functions K or R and you get the error code 3 it is because the equation you are using contains relative grid references and there are no corresponding grid references in the destination position.
- 4) You are trying to divide by zero (or by a blank). For example, if column 3 is derived as column 1 divided by column 2, you must have an entry in each cell in column 2 to make the division possible. (You can get around this problem without putting values in each cell by adding an insignificant amount to the divisor; eg. instead of $a1/a2$ use $a1/(a2+0.00000001)$. This will have a negligible effect where $a2$ is non-zero.)

CANNOT EXTEND - (function Y) - The box over which the cursor is resting does not contain a recognised month of the year (or it contains DEC which cannot be extended anyway) or the dimensions of the grid do not allow any extension to take place.

EDIT PRESSED - (any function) - The edit key was pressed during input. Nothing you have done within this function has affected the data.

INV. REPEAT POS - (function R) - You cannot repeat to this position, either because both the column and line references of this position are different to the cursor location, or because this position is not below or to the right of the cursor location.

I/O ERROR/BREAK - (functions S, L, & X) - Either you pressed the BREAK key, or an error has occurred. If you are using tape then this is a tape loading error. If you are using any of the Interface 1 facilities (including microdrives), then you will normally get a more detailed error message on the line below, which you can look up in the Interface 1 manual.

If you are using microdrives, did not press the BREAK key, and no other message was produced, it means that either you are trying to load a model into too small a space (Probably because you set it up using the original OMNICALC or created it in program "tom2" and are now using "tom64000"), or because there is not enough room to set up a microdrive channel.

This latter case is very unlikely to occur without prior warning (NO SPACE messages appearing in Functions A or I perhaps). If it does occur, and you now wish to load data on another model, you can simply use the I option in function L to set up a much smaller spreadsheet, thereby making more space and then load in your new model. In other instances you will have to try to make more space. (The first thing to do is to try INSERTING columns until the "NO SPACE" error occurs, then to delete any columns you managed to create. This has the effect of forcing the program to remove any redundant equations, and so may give you a bit more space.)

NO INTERFACE - (function X) - You are trying to use microdrives or one of the facilities of Interface 1 without the interface being connected.

NOT NUMBERS - (functions A, D, K, R & T) - The cursor is resting over a header field, so that these functions which are concerned with the body of the grid cannot be used. (To blank out a value in a header field, you should use function I and key in spaces).

NOT HEADING - (function Y) - The cursor is not resting over a header field.

NO SPACE - (functions A & I) - There is no more space either to insert a new column or row, or for another equation. The program has already checked to make sure all equations you have entered are actually being used (and has deleted any that are not).

The only way to get around this error is by making more space, either by deleting redundant columns or rows, or blanking out cells that hold no useful information. If

you still cannot make enough space, you may have to consider splitting your spreadsheet into two and use the workarea to pass values between them.

Note that when you first set up your spreadsheet, enough space was left for at least 50 average equations (which, in the majority of cases, is sufficient). However, if you have been adding columns or rows, or if you set up the model on "tom2" and are now using "tom64000", you may have eaten away into this free space.

OUT OF RANGE - (functions A, L, P,S, X & during the initial entry into the program) - Where you are setting up the dimensions of the model, the number of columns must be in the range 3 - 99, and the number of rows in the range 15 - 250. If this error occurs during function A, it means that you are attempting an insertion or deletion which would take the dimensions outside these ranges.

If you are replying to the "channel?" prompt, your answer must be in the range 0-10 during saves and loads, and 1 - 8 for the microdrive only options in function X.

In function P, the program limits the maximum number of columns you can print at once to 19 (which is equivalent to 160 characters per line). Remember too, that if you are using the Sector option in function P, you must key in a number in the range 1 - 250 in reply to the "no. of rows?" prompt.

For the graphics address and Baud options within function X, the program only checks to ensure that your input is in the range 0 - 65535, and it is up to you to make sure that you get the information right.

TOO BIG – REINPUT – (During initialisation) - There is not enough room in memory to accommodate the number of columns and number of lines you requested (and still leave some room for equations).

TOO LONG – (Function I) - While you were putting in information at the bottom of the screen you exceeded 3 lines.

Unexpected results

If you get results that are unexpected, you should first ensure that you have calculated (used function C, in other words), as it is only during this function that totals, sub-totals and equations which you have copied or repeated are actually evaluated. You should then check that you are calculating in the correct sequence, so that boxes whose values are used in deriving other boxes are evaluated first (remember that you can set the calculation to go row by row or column by column). If you have inserted or deleted rows and columns, it might also be that certain of your equations now refer to the wrong boxes and you should use function D to test for this.

OMNICALC 2

USER GUIDE

Hints on Using the Mathematical Functions

OMNICALC 2 supports all the mathematical functions of the TS2068. This makes it ideal for engineering applications as well as financial projections.

Two facilities are worthy of special note. Firstly the RND function, which produces a pseudo-random number between 0 and 1, and secondly the logical mathematical operators AND, OR and NOT.

The RND function is an important simulation device. Looking at the worked example, it assumed that sales increased smoothly. This is very unlikely to happen - there may be an underlying upward trend, but sales in each month may show peaks and troughs.

If sales are expected to be random within plus or minus 10% of the expected sales, we could reflect this quite simply by an equation such as: Random sales are:- Expected sales * (.9 + 2 * RND).

More sophisticated equations can be used to model statistical distributions, or to reflect relationships between months sales. (If, for example, you have a good sales month in January, it may well be that some of February's sales have come early).

The advantage of incorporating RND in such equations is that you will get different random values every time you calculate, so apart from pressing C you do not have to put in much effort. Done several times, this may reveal potential problem situations that you might not have considered.

The logical operators AND, OR and NOT can be used to put conditions into the model. The costs in the worked example are, to say the least, simplistic. A more realistic situation might be that there are fixed non-production costs of 300, a material cost of 0.3 per unit sales, and that factory labour and overhead costs are a minimum of 300, no matter how many are being made, but that everything over 1000 units of sales value produced will incur an excess factory cost of 0.4 per unit.

We can then write a mathematical expression for the costs as:

$$(300 + 0.3 * SALES + 300) + 0.4 * (SALES - 1000) \text{ AND } (SALES > 1000)$$

The right hand part of this expression means that a cost of 0.4 per unit is incurred on sales in excess of 1000 if (for AND read if) sales exceed 1000.

Differential tax bands etc. can be put in by similar expressions.

The Once/Repeat Options in the Calculate Function

Instead of just calculating everything in the model just once, there is an option in function C to allow you to have everything recalculated over and over again until a condition (that you can stipulate) is met.

The option is selected by replying R to the "Once/Repeat?" prompt that appears after you press function C. If you do select it then after doing the calculations, the value in the last (ie. the bottom right hand corner) box will be inspected by the program. If it is zero, then calculations will terminate. If, on the other hand, it is non-zero, the calculations will be repeated.

Uses for this may not be immediately obvious, and so some examples are given below.

Disaster simulations - If you are using the RND function to provide random numbers to feed into your model, you may want to know if certain conditions that you would classify as disasters will ever occur. To do this, you could simply set up the model to perform repeated calculations, but to stop if a disaster condition occurred, and then start it running. You could then go away and return some time later. If a disaster condition has occurred, the circumstances under which it would happen can be inspected. If not, you can use the BREAK key to stop the calculations, happy in the knowledge that nothing is likely to go wrong.

Range assessments - Again if you are using the RND function in your simulations, you might be interested in knowing how often the results of your calculations fell within certain ranges - to give you a distribution of likely outcomes. Here you might set up the calculation to be performed say 100 times, and then have expressions in several boxes which act as counters.

(Say that you were using boxes aa1, aa2 and aa3 as counters and you wanted to increment aa1 if a profit figure in box z6 was in the range 0 - 100, or increment aa2 if it was in the range 100 - 200, or increment aa3 if it was above 200. You could set up a counter to do this in box aa1 by the expression "aa1 + (1 AND z6 < 100)". For box aa2 you could use "aa2 + (1 AND Z6 >= 100) - (1 AND z6 > 200)", and for box aa3 you could use "aa3 + (1 AND z6 >= 200)". At the end of the 100 calculations you could inspect these three boxes to see the relative frequencies of each profit level, which would give you a clear idea of the average expected profit and tell you about what the chances of getting more or less than this were).

Where you want to perform a certain number of repetitions, you are recommended to use the expression:

Last-box - 1 + (No. of repetitions AND Last-box = 0). (You should substitute the grid reference of your last box, and the figure for the number of repetitions you want in this expression).

The advantage of this formulation is that it will automatically set the number of repetitions to the number you want every time you run the model (so that having set it up once, you can forget it).

Numerical analysis - In numerical analysis you are normally trying to find a value (or values) that offer the 'best' solution to a complicated series of expressions. Best may mean highest or lowest or closest to some predetermined value. Such things as the calculation of Internal Rate of Returns are specialised examples of numerical analysis.

The conditions under which you stop repeating the calculations vary from case to case. If you are trying to find the highest or lowest, you normally stop when the difference between this calculation and the one immediately before is so small as to be insignificant. If you are trying to get to some predetermined value, you normally stop when you are either at the value, or so close to it to make no difference in practical terms.

To give an example, let us assume that you are trying to find a value which makes the value in box z6 zero, and that practically speaking, you will treat plus or minus one-tenth as being the same as zero. In your last box you would set up the expression:

$$1 \text{ AND ABS } z6 > 0.1$$

which will take value 1 if z6 is not in the range -0.1 to +0.1, and the value 0 when it is in this range.