Chapter 2– Finding Square Roots by Iteration

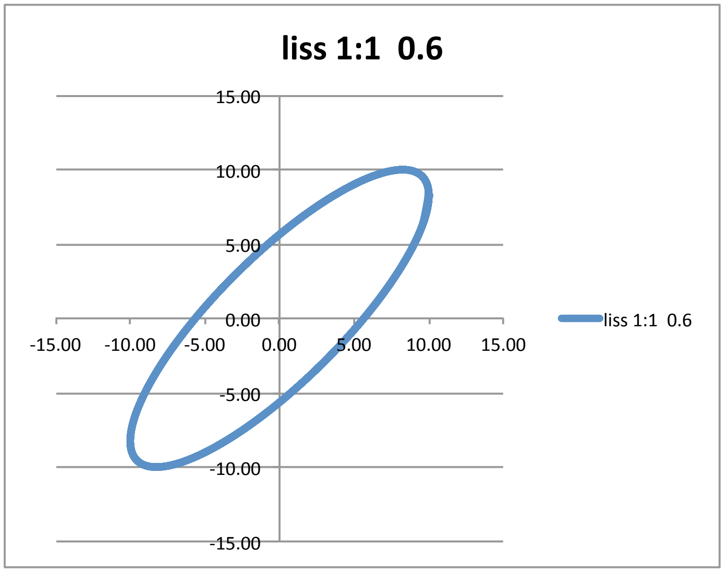

This chapter introduces the idea of finding a result by iteration. Iteration means simply that we will find a result by repeating a series of operations:

1. choose a starting value
2. use the starting value to calculate a better value
3. go back to step 2, but use the latest "better value" for the next calculation
4. continue repeating the calculation until the result is good enough

This seems a bit vague, since I've said nothing about what the calculation may involve, nor how to judge whether the result is "good enough." Finding a square root by iteration gives a perfect example: the calculation amounts to dividing two numbers, and then finding an average. And "good enough" is really just a question of how many decimal places are required.

An Aside on Estimating Square Roots

It is useful to be able to look at a number and mentally estimate the square root. For a number like 56, the reasoning is straightforward. 56 is somewhere between 49, or 7^2, and 64, or 8^2. So a reasonable estimate is that the square root is about 7.5. (The actual square root of 56 is ~7.483, so our estimate is pretty good.) If the number were 52 – closer to 49, we might dial the estimate back to 7.2. This is all easy if the number is in the range from 1 to 100, since the squares of integers from 1 to 10 are learned in elementary school.

Table 2-1 Squares of integers

N	1	2	3	4	5	6	7	8	9	10
N^2	1	4	9	16	25	36	49	64	81	100

What about numbers bigger than 100? Estimate the square root of 73,495, for instance. Here's where scientific notation comes in handy. If we write 73,495 in scientific notation as 7.3495×10^4, then we can estimate the square root of the two parts of the number. The square root of 7.3495 must be between 2 and 3, probably closer to 3 since 7.3 is closer to 9 than 4, maybe around 2.7. The square root of 10^4 is exactly 10^2. So the estimated square root is 2.7×10^2, or simply 270. (The actual square root is ~271.1.)

If we wanted the square root of 7,349, scientific notation gives us 7.349×10^3, and we run into a problem with finding the square root of 10^3, since 3 is an odd number. In this case, we can express the number as 73.49×10^2, and now we can guess the square root as somewhere around 8.5×10^1, or simply 85.

Here is the process, laid out step by step:

7,349	original number
7.349×10^3	express in scientific notation
73.49×10^2	if needed, adjust the decimal point so that: • Exponent is even, i.e., 2 instead of 3 • The coefficient (73.49) is kept in the range from 1 to 100
8.5×10^1	estimate square root for 73, and divide exponent by 2
85	convert back from scientific to decimal format (actual ~85.73)

Finding a Square Root by Iteration

Now let's use Excel to harness a numerical method for finding square roots. Yes, it is true that Excel has a square root function – but the real lesson here is to illustrate an iterative numerical algorithm that will home in on the square root. The basic routine is simple enough:

Take a number N, and estimate its square root S.

1. Divide N by S, call the result D
2. Find the average of S and D, i.e. (S+D)/2
3. Now use this average value as a new value for S, and go back to step 1

Stay in the 1,2,3 loop until the new value for S is the same as the previous value.

If the initial choice of S is smaller than the true square root, then D will be larger, and vice versa. The average of S and D will be closer to the true square root than either S or D, so we use this average value as S for the next iteration of the loop. As the loop is repeated, the average values converge very quickly on the square root. This notion of convergence – our iterative approximations approaching the true value – is an important and much-studied facet of computation. This square root algorithm is said to converge quickly; a few repetitions of the loop suffice to get an accurate result. Other examples later in the book converge much more slowly – and we will always be on the lookout for ways to speed the convergence.

Normally when one speaks of iterative loops, programmers start thinking of constructs like: *for (i=0; i<10; i=i+1).* That's the usual way to program a loop in many programming languages. With Excel, we can do the same thing by setting up formulas in a spreadsheet, and repeating those formulas through a range of cells, in such a way that each successive iteration (usually a line in the spreadsheet) operates on the results of the previous calculation. This sounds more complicated than it is; Excel makes it very easy.

Excel Setup

Here is the initial setup. The cells that will hold numbers have been formatted to display 12 places after the decimal point.

<Formatting cells refresher: Select the cells with your mouse. You can select one or many, an entire column, an entire row, etc. Then click the Format tab, and click Cells… with your mouse. (Alternately, right click after selecting the cells, and then choose Format Cells from the popup.) This will bring up a box with a tab bar at the top; select *Number* in the bar, and then select *Number* under the *Category* column. This causes the box to change its options, and you should now be able to select the number of Decimal places to display, whether to use commas to separate thousands, etc.>

Cell B2 holds the number whose square root we wish to find, here set to 67. (We don't have to enter all the 0s; just enter 67. Formatting the cell to display 12 decimal places automatically displays 12 zeros after the decimal point.) Cell C2 holds the Excel formula for the square root of the contents of cell B2: =SQRT(B2)

NOTE: you can manually type B2, or just click in cell B2, and it will be entered at the cursor. You still have to type the parentheses and function name though.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	=SQRT(B2)
3		S	D
4	initial guess->	8.000000000000	

Figure 2-1 Excel setup for finding square root

As soon as the formula is entered in C2, by pressing return, the cell displays not the formula, but the result of evaluating the formula, i.e., the square root of the number in cell B2, namely 8.1853… This cell is useful for comparison to our numerical computations, just so we know we're on the right track.

The S and D are just labels for the data columns. We enter an estimate for the square root in cell B4, i.e., 8. And cell C4 holds the formula for calculating D, i.e., N/S, or in excel terms, = $B$2/B4.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	=$B$2/B4

Figure 2-2 Add formula to calculate D

Note the use of an absolute cell reference for N, $B$2, and a relative cell reference for S, B4, i.e., the cell just to the left.

Next we add a formula in cell B5 to calculate the next S, which will be the average of the previous S and D values, (S + D)/2, or in the spreadsheet
= (B4 + C4)/2

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		=(B4 + C4)/2	

Figure 2-3 Find average of S and D

Now we could enter another formula in cell C5 – we want to divide the contents of cell B2 by the contents of the cell to the left of C5, i.e., B5. But this is exactly what the formula in C4 did; it divided $B$2 by the contents of the cell on its left, i.e., B4. We can just copy that formula from C4. Select cell C4 by clicking in it, and copy it by dragging the selection box down one cell.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		8.187500000000	

Figure 2-4 Select cell C4

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		8.187500000000	

Figure 2-5 Drag down to repeat formula

As soon as you release the mouse button, the formula will be copied, and cell C5 will display the new D value.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		8.187500000000	8.183206106870

Figure 2-6 Formula copied & evaluated

Now we've done all the real work. At this point, B5 contains the formula for averaging the S and D values in the cells one row above. C5 contains the formula for dividing the latest S value into cell B2 (here containing 67). We can copy both those formulas at once, and replicate them in cells below. Select both cells B5 and C5, grab the square at the bottom right of cell C5, and drag down several rows.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		8.187500000000	8.183206106870
6			
7			
8			
9			
10			

Figure 2-7 Copy formulas to several rows

When the mouse button is released, the formulas are replicated, and the cells are all updated with the next several iterations of calculations. Each row represents one iteration of the loop.

	A	B	C
1		Number	actual sqrt
2		67.000000000000	8.185352771872
3		S	D
4	initial guess->	8.000000000000	8.375000000000
5		8.187500000000	8.183206106870
6		8.185353053435	8.185352490310
7		8.185352771872	8.185352771872
8		8.185352771872	8.185352771872
9		8.185352771872	8.185352771872
10		8.185352771872	8.185352771872

Figure 2-8 Row by row iteration

As you can see, by row 7 we have the square root worked out to 12 decimal places. This particular algorithm converges especially quickly. Row 6 has the square root correct to 6 decimal places – 8.185353 (with rounding). The next iteration – row 7 – adds another 6 decimal places of accuracy.

Debug Tip – If a cell or cells in your spreadsheet behaves oddly –

- Displays a 0 when you were expecting something else
- Displays #VALUE! or #REF!

- it is usually an issue with a cell reference. Double check the cell references in the misbehaving cell. Something is pointing your formula to a blank cell (excel will evaluate as 0), or a cell that contains text instead of a number, or maybe a cell that no longer exists, because of a recent deletion of a row or column.

Things to try

Perform these experiments by changing the Number (cell B2) and/or the initial guess (cell B4). Excel will happily recalculate everything when either of these cells is changed.

1. What if the initial estimate is way off? How many additional iterations (here 1 iteration is one row in the spreadsheet) are needed if the estimate is 2 instead of 8?

2. Try the algorithm with a number that is a perfect square, like 81. Enter an initial guess that is intentionally off, like 6 or 7, just to exercise the algorithm. This makes it easy to watch the algorithm work its way to the correct value, as 0s and 9s propagate through the fractional part of the S and D values.

3. Try the algorithm with some really large numbers – millions. You can use the scientific notation method to get a good initial estimate of the square root, or just make a wild guess and let the algorithm do the work.

4. What happens if you use a negative number for the initial estimate?

5. What happens if you try to find the square root of a negative number, i.e., -67. (Spoiler alert: terrible things happen. The algorithm blows up pretty badly – but it is worth taking a look, just to watch how it goes wrong. This is the spreadsheet equivalent of a train wreck.)

An Aside on Roots as special cases of exponentials - x^n

The traditional mathematical symbol for square root of 67, for example, is $\sqrt{67}$. We have trouble typing this, which is why Excel uses a more keyboard friendly format: SQRT(67). An alternate mathematical way of writing the same thing is $67^{1/2}$. If this is new to you, remember that:

$67^{1/2} \times 67^{1/2} = 67^1 \equiv 67$ (add the exponents when multiplying). So $67^{1/2}$ meets the definition of a square root – *a number that produces a specified quantity when multiplied by itself*. We will use this notation later in the book. Note that we can have other fractional exponents:

$67^{1/3}$ is the cube root of 67, commonly written as $\sqrt[3]{67}$.

$67^{0.01}$ is 67 to the 0.01 power, which could be written as $\sqrt[100]{67}$; however no one quite knows what to call the 100th root – square root and cube root are the most commonly used root names. So for most fractional roots/exponents, it is more common to use the exponential format, and to speak of 67 to the power of 0.01.

Caution: don't make the mistake of thinking that $67^{3/8}$ means $67^3/67^8$. When dividing, the exponents are subtracted, so $67^3/67^8$ is 67^{-5}, or alternately $1/67^5$. So just what is $67^{3/8}$? If forced to find this with only pencil and paper, we could do it with the skills we have already.

$67^{1/2} \cong 8.185$ (the square root of 67)

$67^{1/4} \cong 2.861$ (the square root of 8.185)

$67^{1/8} \cong 1.691$ (the square root of 2.861)

$67^{3/8} = 67^{1/4} * 67^{1/8} \cong 2.861 * 1.691 \cong 4.839$

Seen in this light, square roots are a special case of exponentials x^n, where n=1/2, and cube roots are the special case where n=1/3. These are useful special cases, because of the tie-in to geometry – areas and volumes.

Another special case of historic importance is the case where x=10. It turns out that any positive real number R can be represented as 10^N, if we can find the correct N. Around 1617 Henry Briggs began the calculations to produce a table of logarithms (also known as common logarithms, base 10 logarithms, or decimal logarithms) in which he tabulated the logs (the exponent N) corresponding to integers from 1 to 1000, and he calculated these log entries to 14 digits – by hand. From this table, you could look up the log of 675 for example, log (675) = 2.829303773, and you'd know that $10^{2.829303773}$ is 675. Fascinating hobby, eh?

Such a table would look like this (leaving out ~985 entries, but here's enough to get the general idea):

N	log10
1	0.00000000000000
2	0.30102999566398
3	0.47712125471966
4	0.60205999132796
5	0.69897000433602
6	0.77815125038364
7	0.84509804001426
8	0.90308998699194
9	0.95424250943933
10	1.00000000000000
…	…
996	2.99825933842370
997	2.99869515831166
998	2.99913054128737
999	2.99956548822598
1000	3.00000000000000

Figure 2-9 base 10 log table

You can see in the table that 10^0=1, 10^1=10, and 10^3 = 1000, all as expected. The table can be used in either direction, that is:

- log 3 = 0.47712125471966
- $10^{.77815125038364}$ = 6

Probably you are totally unimpressed, especially if you started high school in 1970 or later. Texas Instruments invented the pocket calculator in 1967; they began selling a version in 1972. Before that time, calculations had to be carried out on expensive noisy bulky mechanical calculators, or – before 1900, largely by hand. Logarithms were used to simplify calculations. For a frivolous example, suppose one wishes to calculate 2 x 3.

$10^{\log3}$ = 3 log3 = 0.47712125471966

$10^{\log 2} = 2$ log2 = 0.30102999566398

$2 \times 3 = 10^{\log 2} \times 10^{\log 3} = 10^{\log 2 + \log 3} = 10^{0.77815125038364}$

And $10^{0.77815125038364}$ is just 6, as you can see from the table. That may not seem like a big simplification, but consider a more challenging calculation, like multiplying 247.6 x 8976; then the calculation consists of:

1. look up the logs of the multiplicands
2. add them together
3. look up the antilog of the sum (antilog means use the table backwards)

For example:

log 247.6 =	2.393750640	find by table lookup
log 8976 =	3.953082844	find by table lookup
sum of logs =	6.346833484	just add the logs
antilog 6.346…=	2222457.6	find by table lookup (backwards)

The great advantage for one calculating by hand is that the multiplication of large numbers is avoided. In its place there is a much easier addition of two large numbers, and three things to look up in tables.

This still seems horribly tedious now that everyone over 12 years old carries around a powerful computer nearly all the time (a device which can also be used to make phone calls, take photographs, etc.), but in the 1960s, high school math, chemistry, and physics textbooks generally had log tables in the back, as an aid to calculation. The tables in textbooks were not generally so precise – maybe only 10 digit accuracy instead of 14. Full-fledged nerds might be caught carrying around CRC tables – a book containing more accurate log tables, plus hundreds of pages of other mathematical tables. Despite the tedium, this sort of calculation saw humankind through the early 20th century. (CRC stood for Chemical Rubber Company, the original name of the CRC Press, which is now up to the 104th edition of its flagship CRC Handbook of Chemistry and Physics – which at 1580 pages, is probably too big for even the nerdiest to carry around. Well – E-reader versions are available.)

The French mathematician Urbain Le Verrier had no better means of computing when he calculated the position of a then-undiscovered planet – Neptune – in the year 1846. He was able to calculate its position from the unexplained discrepancies in the orbit of the known planet Uranus. Unable to interest French astronomers in his prediction – probably they regarded him as a crackpot mathematician – he finally mailed his prediction to a German colleague, Johann Galle at the observatory in Berlin. The observatory received the letter in the afternoon, and discovered Neptune just after midnight that same night, within 1° of the location predicted by Le Verrier. It is difficult to believe today that all the number-crunching underlying this discovery could have been done without a computer; at that time log tables were state-of-the-art.

After all this, you may be wondering if calculators use the algorithm we demonstrated in Excel to calculate square roots. It depends – if you have a very old or very simple calculator, one that is limited to addition, subtraction, multiplication, division, and square roots, it may actually use the algorithm that we used. If it is a more robust calculator – in particular if it can calculate x^y , then it probably uses a more general algorithm to do that, and that same algorithm will handle the special case where y = ½.

Is there any reason to use logarithms today? Well yes – sometimes one encounters exponential equations, for instance:

$327 = 17^x$ With the skills learned in algebra 1, this simple equation is embarrassingly difficult. However, if you take the log of both sides, it quickly transforms into something easily solved.

$\log(327) = \log(17^x)$

$= x \bullet \log(17)$ Remember that $\log(17^2) = \log(17 \bullet 17) = \log(17) + \log(17)$, or just $2 \bullet \log(17)$.

In general, $\log(17^x)$ is $x \bullet \log(17)$

$x = \log(327) / \log(17) = 2.515 / 1.230 = 2.044$

Bottom line: logarithms are your friend! They transform awkward exponential equations to a form easily handled. We'll run into more exponential functions in later chapters, and they will bedevil you one way or another forever (they lie at the heart of interest rate calculations, should you take out a loan for college, or a car, or house…).