Gerald's Column
by Gerald Fitton

My column is a little shorter than usual this month. I am pleased to report that this is not because the number of letters you've sent me have reduced, indeed, if anything there has been more than usual. In the past, my monthly correspondence has consisted of many letters on a few major themes from which I have been able to choose one or two which I think will be of general interest in this column. Receiving many letters on the same topic has its advantages for me; I can think it through once and then send similar answers to many people. This month the topics have varied widely and I've had to spend more time than usual on generating individual replies.

Anyway, what I decided to do was to look at some of my earlier correspondence and select a couple of items for this column which were previously omitted for want of space rather than select current items which might be of limited appeal.

Impression as a Graphics Library

Roger King sent me this idea. Examples of the files to which I refer are on the Archive monthly disc in the directory Library.

Essentially, what Roger has done is to use his Impression Publisher (it also works in Style) as a database of Draw and Sprite graphics. He has created a new master page which contains twelve graphics frames each 50mm by 50mm. Above each graphic frame is a small text box 50mm by 12mm containing the title of the graphic. The graphics are dropped into the graphics frames and a suitable text inserted in the title frame. Choose the text for these title frames with care since they are really an 'index' to the the library of graphics.

There are many advantages to using this system for storing graphics rather than the usual Acorn directory structure but the one I have found most useful is the ease with which I can search through the database. Since you will have chosen the titles for your graphics carefully you will be able to search for that text through your library document. Place the mouse pointer in any small text frame on the page, tap <F4> and enter the search data in the 'Find what:' dialogue box. Select 'Whole document' and click on 'Find next'. Style and Publisher will look for the search string in the document. Once you have found your graphic you can Save it using the 'File – Save graphic' menu.

Precision

It seems years ago that I received a letter from Nik Kelly on this subject. In his letter Nik pointed out that, using a variety of spreadsheets and calculators, he had tried calculating the standard deviation of the following six numbers: 1.797, 1.797, 1.798, 1.799, 1.799, 1.799. He was surprised to find that the answers ranged from 9.8319 e–4 to 9.8336 e–4. He concluded that the are accurate to only three significant figures.

Since I received Nik's letter I have had many more complaining about the precision of certain spreadsheet calculations. Now I know that you may not be interested in standard deviations but this article (like most of my articles) has a wider application than the

calculation of standard deviations. Indeed, I use it only as an example.

Nik concludes his letter with:

"We must limit reported standard deviation results to 3 significant figures, as recommended by HLR's draft 'Precision & Accuracy' SOP. We need to challenge calculators and spreadsheets for arithmetic precision and accuracy, and we must also check the integrity of their statistics formulae. To establish benchmarks, we need sets of data calculated to much longer precision than those shown above. How do we know which of them were accurate?"

I don't wish to anticipate the detail of my answer but, using Fireworkz and PipeDream 4 the 'right' way the 'accurate' answer I get is: 0.000983192080250172. I believe that my answer is accurate to much more than the three significant figures quoted by Nik.

A Digression

It was a long time ago when I was first confronted by this problem of apparent lack of precision when carrying out particular types of calculation. In fact it was first demonstrated to me during a 'Numerical Methods' lecture given as part of my Maths Degree Course at Imperial College.

Essentially, the reason for the lack of precision is this. If your calculating equipment is limited to a fixed number of significant figures then small differences between relatively large numbers can not be calculated to the same accuracy as the original data.

Let me give you an example which doesn't need a calculator. Let us suppose that you have two numbers which are close to 500 000, a and b. Suppose a rounds off to 499 999 and b rounds off to 499 998. Both a and b are accurate to six significant figures. However, the value of (a − b) is accurate to only one significant figure. We have lost precision equivalent to five significant figures.

Given the explanation above I'm sure that you can devise for yourself an example for which a and b are accurate to twelve significant figures but, after subtraction, (a − b) is accurate to only three figures.

In the case of the SD calculation above the original data differs only in the fourth significant figure even though most of the calculators used have at least twelve digit precision. Calculating SDs involves finding the differences between squares of numbers; in finding differences between squares it is all too easy to choose a method in which a significant amount of precision is lost!

The example chosen by Nik of calculating a small standard deviation is a classic in the sense that, even when I was being taught, it was the calculation of a SD which my lecturer chose to demonstrate the phenomenon! He said at the time, that the specific example he chose was at least 50 years old because it was the one his lecturer chose when he was taught over 40 years previously! My lecturer was approaching retirement and I got my Maths degree well before the days of computers and hand calculators.

The method we used at college in those days was to use a mechanical calculator (such as a

Facit) where you turned a handle to multiply! Such mechanical calculators were capable of 12 digit precision. The standard deviation example we were given by the lecturer to calculate failed at the third significant figure in a similar way to the way Nik's failed when we used the 'standard' inaccurate method.

My lecturer added that loss of precision when calculating differences was well documented way back in 1665!

The Fireworkz std(list) method

If you want to calculate the standard deviation of a list of numbers then, in many spreadsheets, you can use a function similar to the Fireworkz built in std(list) function. Fireworkz gives the result: 0.00098319208141798 to the formula:

std(1.797,1.797,1.798,1.799,1.799,1.799).

I believe that the first eight significant figures are correct and that the ninth is wrong.

Gerald's Method

The 'correct' method of calculating SD values which are small when compared with the mean of the numbers is that which I show in the file [StdDev] which you will find on the monthly disc. It is the method explained to me those many years ago by my Numerical Methods Lecturer.

The first thing that you must appreciate is that, if you reduce every value in the set of data by a fixed amount then the true value of the SD of the modified data is equal to the SD of the original data. For example, the SD of the three numbers 1001,1002,1003 is exactly the same as the SD of the numbers 1,2,3. The original numbers have each been reduced by 1000.

The recommended method of calculating the SD of a set of data is to first calculate the mean. Then you subtract the mean from each number. Finally, you calculate the SD of the reduced data. The value for the SD given by this method is: 0.000983192080250172. I believe that this is about as accurate as you will get with double precision arithmetic.

When I reduced each of Nik's values by 1.8 (which is close to the mean but not exactly equal to the mean) before using the Std(list) function I got the answer: 0.000983192080250173. Compare this answer with that given in the previous paragraph when the (exact) mean is deducted from the data and you will see that it differs by 1 in the fifteenth significant figure. I assure you that the number with the 3 at the end is slightly less accurate than the 2 given by subtracting the exact mean.

The Bi-directional Parallel Port

If I have to choose a single topic from this month's correspondence then this must be it. I shall not relate all the many different ways in which this topic has arisen. I'll restrict myself to just one.

It would seem that a few of the well known printer manufacturers believe that everybody must have a computer with a bi-directional printer port. [By the way, my hyphen is added deliberately for emphasis; it is usually omitted].

For very many years now all parallel port cards for the IBM type of personal computer have been bi-directional. This means that the port will accept data flow both out of and into the computer. As we all know PCs become redundant much more quickly than Archimedes do and, as a consequence, the number of PCs with uni-directional parallel ports is negligible. On the other hand the parallel port of many of the Archimedes machines up to and including the A540 are uni-directional and these machines still exist and, I have no doubt, will continue to do so for many years to come.

The 'old' standard for a parallel printer port is one called 'Centronics' (after the company which first developed it). This standard is a uni-directional standard; data flow is in one direction from the computer to the printer.

When you connect a printer built to the Centronics standard to a more modern bi-directional parallel port it still works. The Risc PC (amongst others) has a bi-directional parallel port. Hence, if you connect a printer built to the Centronics standard to a Risc PC it will work.

Now here is the bad news. Until recently all printers conformed to the Centronics standard and all of them worked with all Archimedes computers. It seems that some printers now appearing on the market are no longer built to the Centronics (backward compatible) standard but to a 'Bi-tronic' standard. Of course, since nearly all PCs still in use have a bi-directional parallel port, problems in that market are practically non existent. The owner of a PC can go to the High Street store and buy a 'Bi-tronic' printer with impunity!

If you buy a printer built to the 'Bi-tronic' standard (from your local supermarket!) then it will not work with those earlier Archimedes computers such as the A540 which have a uni-directional parallel port.

I had intended to include a list in this article of all those computers which are known to be 'Centronics incompatible' but my list would not be sufficiently comprehensive to be useful. I think that my advice must be that you buy your computer through an Archimedes dealer rather than at the High Street store since they are more likely to be aware of this problem and they won't sell you a problem.

Marketing

This month I shall briefly mention this topic in the hope that this will solicit opinion from you and allow me to expand on the theme in a future article.

Regular readers of this column will know that I teach "Anything to do with Maths". One of my subjects is Marketing. In case you are wondering "How is Marketing linked to Maths?" the link is via Statistics. Quoting the text book, "Statistical data and analysis is used to plan and control marketing activities".

What even regular readers might not know is that in addition to being able to teach Marketing I had much practical experience when I worked in industry. Indeed, I built up a small company from nothing (just me and no ready designed products) to a turnover of £0.5M in six years – then the owner sold it! I owned 5% of the company so I got my 5% cut of the £0.75M that it was sold for and a three year contract with the new owners. At the end of the three years I was made redundant at a ten minute meeting during which it was suggested that if I bought the company car with part of my redundancy pay then it would avoid the necessity of being given a lift home! Enough of these reminiscences.

When it comes to pricing a new product the classic teaching is that, in the initial stage, you can 'cream' the market at a high price. Later, you want to pick up more sales and so, as time goes by, you gradually reduce the price.

Although not in the text book for the course I teach I suggest to my students that this classic scenario of a high initial price which is gradually reduced to maintain sales applies only in rare circumstances. I quote counter examples where it is better to start at a low price so that the market is established quickly, and then, when the market is established and everybody wants the product, the price can be increased.

So that you can more easily understand what I mean let us take the Acorn upgrade from RISC OS 2 to 3 as an example. By offering RISC OS 3 to all users of RISC OS 2 at a low initial price for a well defined limited period, the change over from RISC OS 2 to 3 was rapid. Because many Archimedes users changed almost immediately it wasn't long before software appeared which would not run on a RISC OS 2 machine. Because this new software was so desirable, those who had not upgraded decided to do so albeit at the higher price.

Let me take another example. Computer Concepts brought in an upgrade to Impression II, namely Style and Publisher. They marketed this upgrade by offering a low price upgrade from Impression II and Impression Junior but only for a well defined limited period of time. The result was that nearly all those who were regular users of the earlier version upgraded within the time limit. This established Style and Publisher in the market and allowed Computer Concepts to drop Impression II from the range of products they sold. More subtly, they were able to reduce the technical backup to Impression II to virtually zero without any complaint from their customers. To my way of thinking, this is an example of good Marketing practice by CC.

From my correspondence I have no doubt that you know of many examples of companies who have made a substantial upgrade to their product for which they can charge. Some have introduced the upgrade (or new product) at a high initial price to existing users only to reduce it later. Let me know whether you think that this was good marketing.

I do not want to anticipate my conclusions but I will quote Bob Ardler who regrets that he didn't heed (or was it "read"?) my advice to upgrade from Impression II to Style whilst the price was cheap.

Finally

Please write to me at the Abacus Training address given on the inside back cover of this Archive magazine rather than via the Archive office.

It does help me if you send me a disc containing a worked example of anything even slightly complex. At the very least it helps avoid any ambiguity about what you mean. In many cases, if you include some typical data (as a disc file) which you need processed, then it saves me a great deal of time since I don't have to create dummies which might not be representative of your problem.