

# Capacity Calculations:

## Handle with Care

This paper discusses avoiding calculation results that are more precise than is justified by precision of corresponding measurement input data.



### **About the Author**

Neil J. Gunther, M.Sc., Ph.D., is an internationally known computer performance and IT researcher who founded Performance Dynamics in 1994. Dr. Gunther was awarded Best Technical Paper at CMG'96 and received the prestigious A.A. Michelson Award at CMG'08. In 2009 he was elected Senior Member of both ACM and IEEE. His latest thinking can be read on his blog at [perfdynamics.blogspot.com](http://perfdynamics.blogspot.com)

## 1. Introduction

With the availability of electronic calculators, spreadsheets, and other tools for easily performing otherwise tedious calculations, it is easy to overlook results that are more precise than is justified by the precision of the corresponding input data. The concept of precision is very important and can impact results that arise from measurements in surprising ways.

To appreciate the difference between counting and measurement, recall the year 2000 Presidential election. Prior to that election, voting in the United States was assumed to be a straightforward counting process; count the votes to determine which candidate has the largest number. In the State of Florida, however, a new problem arose—scrutinizing the small difference between two very large numbers. Ballot officials needed to accurately count a few hundred votes out of some 6,000,000 ballots cast. The political pressure for greater accuracy soon led to deeper question: What is a vote? But this is really a question of measurement—not counting. Everyone failed to realize that counting is exact but measurement is not. Strange new terms like “hanging,” “dimpled” and “pregnant” “chads” were really lame attempts to convert measurements into counts. A better approach would have been to assign a real number to those votes. For example, hanging = 0.75, dimpled = 0.50, pregnant = 0.25 and add those real numbers to get the total count. Of course, the total would have been a real number: “George Bush wins by 103.64 votes!” Apparently, the notion that a vote has to be an integer was so psychologically ingrained that this approach remained untried.

This kind of dilemma arises out of a failure to recognize that there are two classes of number:

1. Exact numbers (those we identify with the positive integers or cardinals in mathematics)
2. Measured numbers (those we identify with the reals in mathematics)

The final zero in the number 50.0 is not there for cosmetic purposes. It tells us something very important about the level of precision used in the measurement.

Exact numbers are numbers that are exact by definition, e.g., 3600 seconds corresponds to one hour. There’s no question about it because it’s true by definition. Mathematically, exact numbers are associated with the integers. When you ask for seating in a restaurant, the number of people you provide is an exact number—an integer. That’s because people come in integral multiples, not fractions. Measured numbers, on the other hand, are quantities that are estimated by including significant digits (“sigfigs”) without the benefit of any natural integer multiplier. Mathematically, measured values are associated with the real numbers. For example, the real number,  $\pi$ , refers to the measurement of a circular circumference using the diameter as the “yard-stick”. As the Greeks discovered to their dismay, the circumference is not an exact multiple of the diameter. It cannot be expressed as an exact number and therefore it is given the Greek meta-name,  $\pi$ , instead.

Averaging<sup>1</sup> is another example of a process closely associated with measurements and estimates. Even though people occur naturally as integer-valued quantities, the average family in the United States is reckoned to have 2.37 children. This estimate is real-valued and is expressed here to 3 significant digits or 3 sigfigs. Time-based averages occur throughout computer performance analysis and capacity planning because most of the metrics we use are sampled and then averaged over some prescribed measurement interval.

Let’s define some other terms associated with measurement, more precisely.

## 1.1 Significant Digits

A significant digit is one which is actually measured. The number of significant digits in a measurement depends on the type of the measuring device. No matter what the measuring device, there will always be some uncertainty in the measurement. Both the device and the observer add their own uncertainty to the measurement. As I've already explained, this point reached world-wide significance during the confusion surrounding the Florida vote count (See more in Section 6) in the 2000 Presidential elections.

## 1.2 Accuracy

In everyday parlance, we tend to use the words accuracy and precision synonymously but in science and engineering they are clearly distinguished. Accuracy refers to how close a measurement is to the expected value. Using an archery target analogy where an arrow represents a measurement and the bulls-eye represents the expected (or accepted) value, accuracy corresponds to the distance between the arrows and the bulls-eye.

Mathematically, it's the maximum error we introduce because we truncate the digits. By convention, this is taken to be one half of the value of the least significant digit present.

## 1.3 Precision

Using the archery analogy, precision is the distance between each arrow, irrespective of where they lie on the target with respect to the bulls-eye. The grouping of arrows could be tightly clustered but a long way from the bulls-eye.

Mathematically, it's the number of digits available to represent the mantissa<sup>2</sup>. Exact numbers (or integers) have infinite precision. But beware! It is possible to have high precision with poor accuracy.

### 1.3.1 Calculating Pi

In 1853 William Shanks published a calculation of  $\pi$  to 607 decimal places. Twenty years later, he published a result that extended this precision to 707 decimal places. This was the most precise numerical definition of  $\pi$  for its time and adorned many classroom walls. In 1949 a computer was used to calculate  $\pi$ , and it was discovered that William Shanks's result was in error starting at a point near the 500th decimal place all the way to the 707th decimal place. Nowadays, with the benefit of a true value for  $\pi$  to 100,000 decimal places, we can say that William Shanks's techniques generated a precise result, but the value he obtained was not accurate.

## 2. Count by Zero

Here are the rules for assigning significance to a digit.

### Algorithm I

Always scan Left to Right  
Is there an explicit decimal point?

YES: Locate the first non-zero digit  
Count it and ALL digits (including zeros) to its right

NO: Insert a decimal point on the end  
Locate the last non-zero digit prior to the decimal point  
Count it and ALL digits to its left  
Ignore all zeros trailing that digit

By “count” we mean to include that digit in the count of significant digits. That’s it in a nutshell.

Let’s take the number 50.0 as an example. It does have an explicit decimal point. Therefore, we scan it from left to right. The first non-zero digit is ‘5’. This is the first significant digit. We continue to count all digits thereafter, including zeros. There are two zeros. So, there is a total of three significant digits.

### 2.1 Examples

Number	Type	SigFigs	Remarks
50	Integer	1	Implicit decimal point
.00005	Real	1	
65000	Integer	2	Implicit decimal point
0.0067	Real	2	
2500	Integer	2	Implicit decimal point
0.070	Real	2	
0.00341	Real	3	
0.157	Real	3	
28.0	Real	3	
40300	Integer	3	Implicit decimal point
$6.02 \times 10^{23}$	Real	3	Scientific notation
30.07	Real	4	
200300	Integer	4	Implicit decimal point
6.751	Real	4	
1.0040	Real	5	
50.0005	Real		

Table 1: Significant digits for a variety of integers and real numbers.

What about the number 5060? It does not have an explicit decimal point. Therefore, we insert a decimal point and then locate the last non-zero digit prior to the decimal point. That would be the '6' digit. We now count that digit all digits to its left-including zeros. Thus there are 3 significant digits. Table 1 shows more examples of SigFigs output that you should check for yourself.

### 3. Precision vs. Scale

According to Algorithm I in Section 2, the numbers 11000 and 0.011 have the same number of significant digits. How can that be? Surely more effort was put into measuring the decimal places of the second number.

To understand why this is an illusion, imagine you are getting your doctor's prescription filled at a pharmacy and she is required to measure out 100 milliliters of a liquid to make up your prescription. A milliliter is one one-thousandth of a liter so, 0.100 liter is the same as 100 milliliters.

The pharmacist has a choice of measurement vessels. For example, she could measure 10 (deciliter) divisions in a one liter graduated cylinder or she could just fill a 100 milliliter graduated cylinder to the top graduation.

In either case:

- The amount of reading effort is about the same in each case.
- The real difference is the size of the measuring device, not the amount measured.

In other words, the quantities: 11000 and 0.011, distinguish scale, not precision. The precision is the same, 2 sigfigs. It makes no difference to the precision whether we write 11000 microliters or 0.011 liters.

### 4. Rounding Rules

Consider the number 4.246 expressed to three sigfigs. Because there are four digits in the number, we need to drop the '6'. When I was in school, we were taught to "round up" the '4' because "6 is bigger than 5". The reported number becomes 4.25, correct to three sigfigs. Many of you will be familiar with this rounding convention.

But what if the number was 4.245? I was also taught to round up the '4' when the next digit is '5' or greater. Many of you will also be familiar with that convention. It turns out that this rule has been updated in recent times because the old rule introduces an inherent bias<sup>3</sup> in the rounding process.

The "new" rule requires that we look at digits beyond the '5' as well as examine whether the preceding digit is odd or even. In this case, there are no digits beyond the '5' and the digit preceding it is even. The new rule states that we should simply drop the '5' and leave the '4' alone. The reported number is therefore 4.24, correct to three sigfigs; not 4.25, as you might have been anticipating on the basis of the old rule.

By making the following notational definitions, we can encapsulate these rules in the form of an algorithm.

1. Denote by the  $X$  the value of the last digit to be reported.
2. Denote by  $p(X)$ , the position of  $X$ .
3. Denote by  $Y$  the value of the digit at  $p(X) + 1$ .
4. Denote by  $Z$  the value of the digit at  $p(X) + 2$ .

We can write the new rounding rules in the form of Algorithm II:

### Algorithm II

```

Examine Y:
If Y < 5: Goto (h) ... (a)
If Y > 5: X = X + 1, Goto (h) ... (b)
If Y == 5: Examine Z ... (c)
    If Z >= 1: Y = Y + 1, apply (a) or (b) ... (d)
    If Z == NULL or A string of zeros: ... (e)
        Examine the parity of X (New Rule)
        If X == odd:
            X = X + 1 ... (f)
        Else X = X ... (g)
        Goto (h)
Drop Y and all trailing digits ... (h)

```

The “old” rules are (a)-(d) and (h). The previous example followed “New” rule (g).

## 4.1 Examples

The following Table 2 shows more examples of applying the new rounding rules.

Number	Rounded	SigFigs	Remarks
1.2345	1.234	4	New rule (e-f).
1.2355	1.236	4	Same as old rule.
1.23550	1.236	4	New rule (e-f).
1.23552	1.236	4	New rule (d)
0.000343	0.0003	1	Same as old rule.
4589	4590	3	Same as old rule.

Table 2: The new rounding rules.

Tools like EXCEL get this wrong. Setting the Cell Format to General:  $=\text{ROUND}(4.245, 2) \neq 4.25$ . Note that the second argument in ROUND indicates the number of places after the decimal point, rather than the number of sigfigs.

## 4.2 The Golden Rule

When a calculation involves measurements with different numbers of significant digits, the result should have the same number of significant digits as the least of those among the measurements.

## 4.3 Sum Rule

A sum or difference can never be more precise than the least precise number in the calculation. So, before adding or subtracting measured quantities, round them to the same degree of precision as the least precise number in the group to be summed.

## 4.4 Example

Sum the following numbers: 2.95, 32.7, and 1.414. The first two numbers have the least precise values viz, 3 sigfigs. Setting the digits of each number in their respective columns produces:

$$\begin{array}{r} 2.95 \\ 32.7 \\ + 1.414 \\ \hline \end{array}$$

Next, the fractional digits (those following the decimal point) are rounded to one decimal place. The sum then reads:

$$\begin{array}{r} 3.0 \\ 32.7 \\ + 1.4 \\ \hline 37.1 \end{array}$$

The result is 37.1, correct to 3 sigfigs.

## 4.5 Product Rules

When two numbers are multiplied, the result often has several more digits than either of the original factors. Division also frequently produces more digits in the quotient than the original data possessed, if the division is continued to several decimal places. Results such as these appear to have more significant digits than the original measurements from which they came, giving the false impression of greater accuracy than is justified. To correct this situation, the following rule is used.

### 4.5.1 Equal Sigfigs

In order to multiply or divide two measured quantities having an equal number of significant digits, round the answer to the same number of significant digits as are shown in one of the original numbers.

### 4.5.2 Unequal Sigfigs

If one of the original factors has more significant digits than the other, round the more accurate number to one more significant digit than appears in the less accurate number. The extra digit protects the answer from the effects of multiple rounding.

### 4.5.3 Final Rounding

After performing the multiplication or division, round the result to the same number of sigfigs as appear in the less accurate of the original factors.

## 4.6 Example

Calculate:  $2.95 * 0.90462$ . There are 3 significant digits in the least precise number (the first factor). But, applying rule(4.5.2), we retain 4 significant digits of the second factor.

$$\begin{aligned} &\Rightarrow 2.95 * 0.9046 \\ &= 2.66857 \\ &\Rightarrow 2.67 \text{ (rounded up)} \end{aligned}$$

Result: 2.67 is the correct result to 3 significant digits. The  $\Rightarrow$  symbol should be read as “becomes” to distinguish it from the “=” sign since that step involves a non-mathematical transformation with regard to precision.

## 5. Expressing Errors

There are three common and acceptable ways to display errors.

1. Absolute error: Half the smallest sigfig
2. Relative error:  $\text{Error} = [(\text{measured} - \text{expected}) / \text{expected}]$
3. Error bars: See Fig. 1.

Consider, for example, the first entry in Table (1): 50. There is only 1 sigfig for the least precise input value. That sigfig is in the 10's column. Therefore, half that sigfig (i.e.,  $10/2 = 5$ ) can be assigned as the absolute error i.e.,  $50 \pm 5$  would be an appropriate way to express the error for that quantity. This corresponds to  $5/50$  or  $\pm 10\%$ .

Conversely, if we had measured a value of 60 when we were expecting 50, the relative error would be:

$$\begin{aligned} \text{Error} &= \frac{\text{measured} - \text{expected}}{\text{expected}} \\ &= \frac{60 - 50}{50} \\ &= \frac{10}{50} \\ &= 20\% \end{aligned}$$



Since the error may differ for each data point, the (vertical) error bars in Fig. 1, should have different heights that reflect that error. Fig. 1 used the EXCEL "Y Error Bars" tab and it automatically computes the standard error. But this is better than nothing.

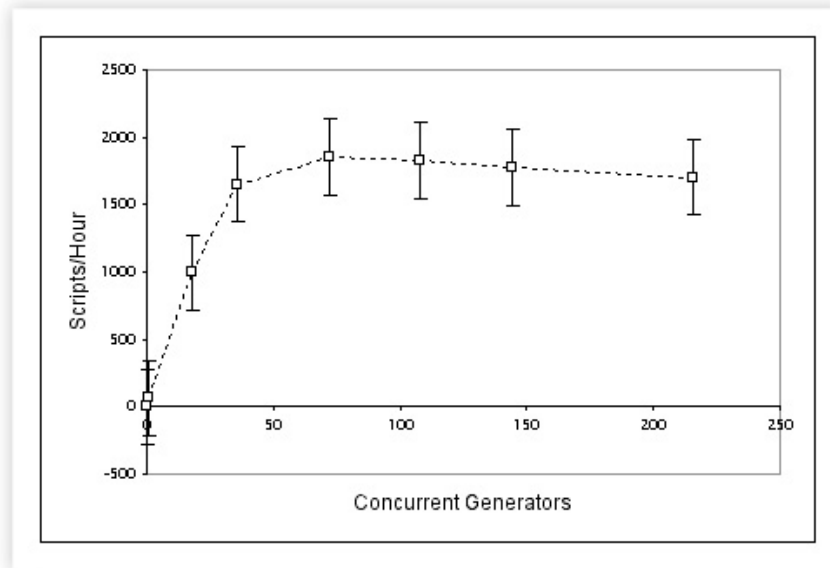


Figure 1: Error bars on a throughput plot.

## 6. The Florida Vote — A Slight Return

Because the ballot measuring equipment (both manual and mechanical) was not capable of producing the necessary kind of precision, the whole problem eventually degenerated into a question of, What is a vote!? In actual fact, as I mentioned in my opening remarks, both the Voting officials and the American public were being exposed to a problem in Experimental Physics: How to accurately measure the small difference between two large numbers. Voting officials were trying to discern a few hundred votes out of some 6,000,000 cast. The gross count in Florida had more or less the same number of votes in favor of each candidate, and so it became a question of precision on the order of 1 part in 6 million.

If the same techniques of experimental physics that are used to measure subatomic interactions could have been applied to the Florida vote count, it could have been determined with a precision equivalent to 1 vote in 6 billion. In other words, quantum measurement techniques could easily have determined whether Bush or Gore won to within a single vote ... even if the whole world had voted!

## 7. Downloads

Perl scripts for the significance and rounding algorithms are available for downloading as:

1. **Algorithm I**
2. **Algorithm II**

## Footnotes

<sup>1</sup> The process of averaging means that information is lost. In fact, the process of simply adding numbers together loses information.

<sup>2</sup> The part of the number after the decimal point.

<sup>3</sup> Using the old rule, you would have rounded down if the next digit was either of (1, 2, 3, or 4) but rounded up if the next digit was either of (5, 6, 7, 8, or 9). Over a large number of rounding samples, you would tend to round down 4-9ths of the time but round up 5-9ths of the time! By selecting out the '5', we are left with rounding up if the next digit is one of (6, 7, 8, or 9) i.e., 4-9ths of the time. In the case of '5' exactly, we only round up only half time based on whether or not the preceding digit is odd. The overall effect is to make the rounding process balanced.

# TeamQuest Corporation

[www.teamquest.com](http://www.teamquest.com)

Follow the TeamQuest Community at:

## Americas

[info@teamquest.com](mailto:info@teamquest.com)

+1 641.357.2700

+1 800.551.8326

## Europe, Middle East and Africa

[emea@teamquest.com](mailto:emea@teamquest.com)

Sweden

+46 (0)31 80 95 00

United Kingdom

+44 (0)1865 338031

Germany

+49 (0)69 6 77 33 466

## Asia Pacific

[asiapacific@teamquest.com](mailto:asiapacific@teamquest.com)

+852 3579-4200

## Copyright ©2010 TeamQuest Corporation All Rights Reserved

TeamQuest and the TeamQuest logo are registered trademarks in the US, EU, and elsewhere. All other trademarks and service marks are the property of their respective owners. No use of a third-party mark is to be construed to mean such mark's owner endorses TeamQuest products or services.

The names, places and/or events used in this publication are purely fictitious and are not intended to correspond to any real individual, group, company or event. Any similarity or likeness to any real individual, company or event is purely coincidental and unintentional.

NO WARRANTIES OF ANY NATURE ARE EXTENDED BY THE DOCUMENT. Any product and related material disclosed herein are only furnished pursuant and subject to the terms and conditions of a license agreement. The only warranties made, remedies given, and liability accepted by TeamQuest, if any, with respect to the products described in this document are set forth in such license agreement. TeamQuest cannot accept any financial or other responsibility that may be the result of your use of the information in this document or software material, including direct, indirect, special, or consequential damages.

You should be very careful to ensure that the use of this information and/or software material complies with the laws, rules, and regulations of the jurisdictions with respect to which it is used.

The information contained herein is subject to change without notice. Revisions may be issued to advise of such changes and/or additions. U.S. Government Rights. All documents, product and related material provided to the U.S. Government are provided and delivered subject to the commercial license rights and restrictions described in the governing license agreement. All rights not expressly granted therein are reserved.