electronics COOLING

2022 GUIDE TO PRACTICAL STATISTICS FOR THERMAL MANAGEMENT

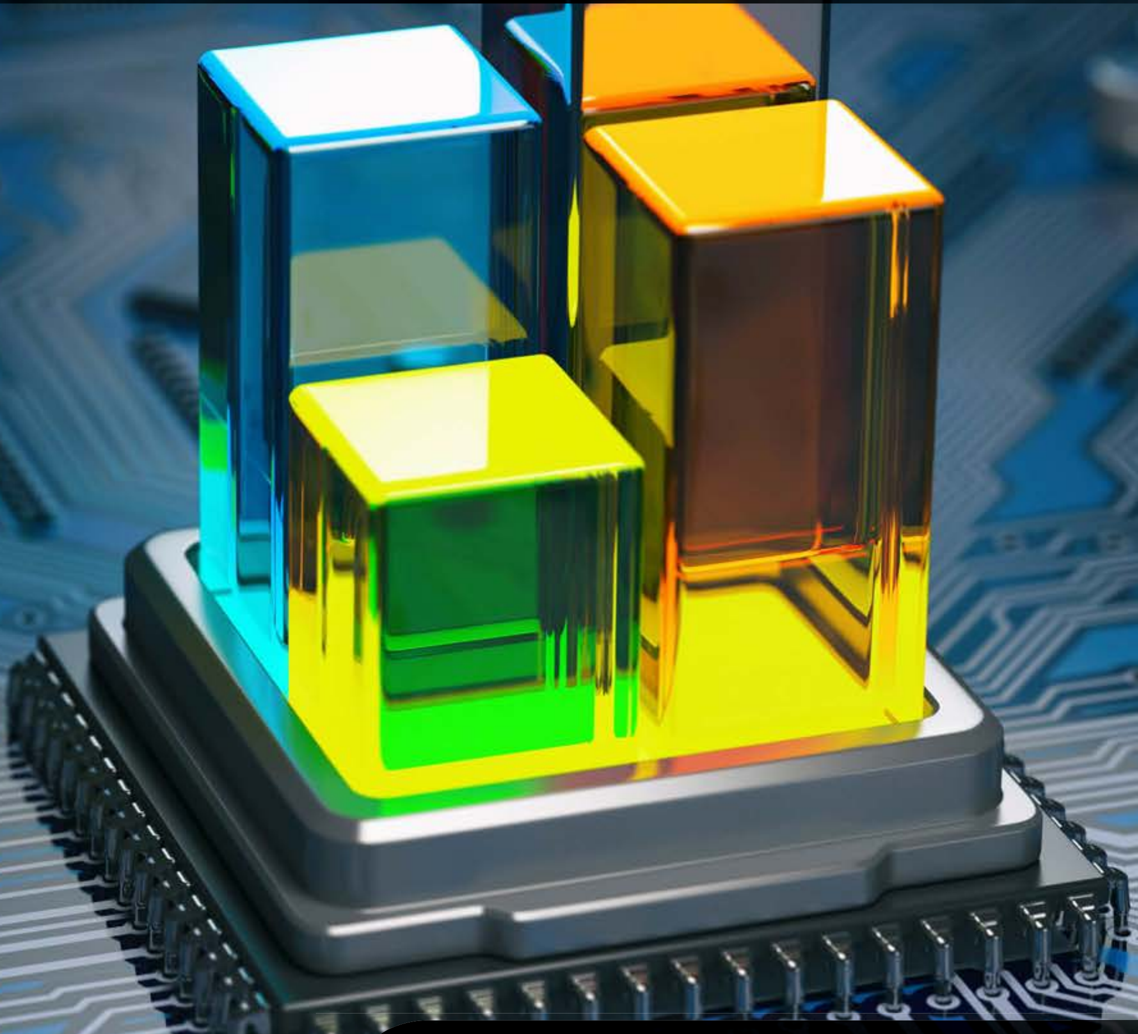

FEATURED IN THIS EDITION

ALPHA | Your partner for thermal solutions

Online Heatsink Customization

Online Heatsink Customization

1: Specify Customization | 2: Review Customization | 3: Input Contact Information | 4: Submit Request

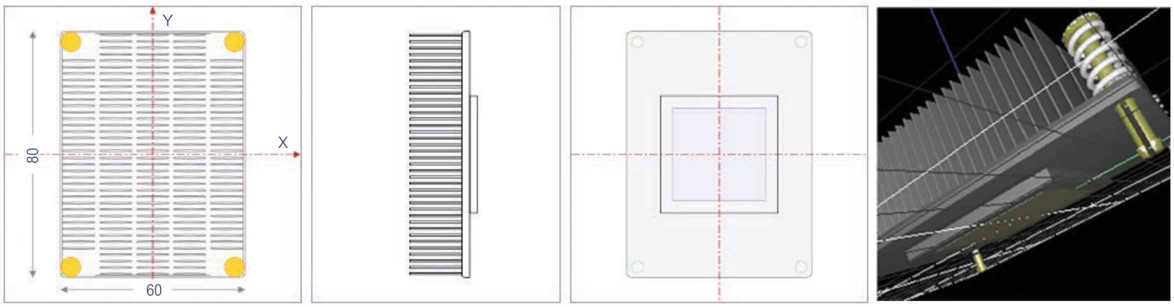

Customization Options for UB8060

Estimated Price (US$)

Quantity	Unit price
1	164.57
5	53.15
10	37.89
50	14.95
100	13.08
200	Quote

Finalize

Customization Summary [Units:mm]

Change Height / Thickness	Fin Height	17	Base Thickness	5
Add TIM	Type	T-PCM585-30S	Size	30 x 30
Add Pedestal	Size	30 sq.	Height	1.5

New Features

- 3D view is updated immediately, providing instant design visualization
- Improved customization user interface
- Fin removal can be specified
- 2D/3D drawings and quotes are available in 0-2 business days
- Proto lead time is 1-2 weeks with No tooling/NRE fees

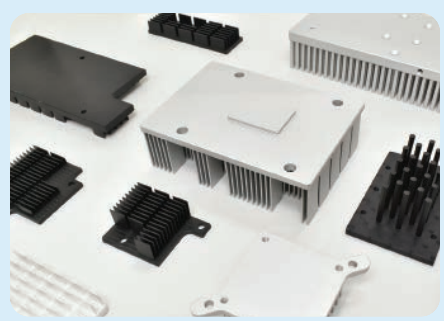

ALPHA Co., Ltd.
Head Office
www.micforg.co.jp
256-1 Ueda, Numazu City, Japan 410-0316
Tel: +81-55-966-0789 Fax: +81-55-966-9192
Email: alpha@micforg.co.jp

ALPHA NOVATECH, INC.
USA Subsidiary
www.alphanovatech.com
473 Sapena Ct. #12, Santa Clara, CA 95054 USA
Tel: +1-408-567-8082 Fax: +1-408-567-8053
Email: sales@alphanovatech.com

www.alphanovatech.com

TABLE OF CONTENTS

EDITORIAL

Ross Wilcoxon
Associate Technical Editor
Technical Fellow, Mechanical Engineer, Collins Aerospace

I Think That This Is Probably About Statistics

Let's get things straight from the outset - I am not a statistician. However, over the past couple decades I have often fielded requests from co-workers to analyze test data, recommend test approaches for design of experiments, etc. This has provided me with opportunities to learn, forget, and relearn methods for analyzing test data to identify trends and determine, for example, if a trend is statistically significant. In many cases, a big part of that begins with identifying useful ways to plot the data – but formal statistical tests eventually find their way into the analysis.

A few years ago I decided that it would be somewhat fun and educational (for me at least) to put together a short course that teaches a few of the statistical lessons that I have learned during my career. I quickly addressed the most important aspect of this: coming up with what I thought was a good title of "I think that this course is probably about statistics". My primary criteria to qualify as a 'good' title was something that both described my viewpoint and also amused me. My description of statistics is that it is a set of tools that apply the laws of probability to deal with the uncertainty that is inherent to any data. After putting that course together, I realized that it could provide suitable material for a series of columns in Electronics Cooling Magazine.

My goals with these columns have been to be mostly correct and always useful; I assume that classically trained statisticians might occasionally cringe at the ad hoc way I occasionally deal with some of the concepts though. I am not always as disciplined as I should be in defining, for example, what terms are related to an entire population or if they describe the characteristics of a discrete set of measurements. I trust that my readers will gamely stick with me as I eventually (hopefully) reach my objective of providing a better understanding of fundamental concepts of statistics. In addition, whenever possible I have attempted to provide the reader with methods for conducting their own statistical analysis using spreadsheet tools such as Microsoft Excel.

Over time, I plan to add more columns to this document to describe other statistical topics that I find interesting and useful. So be on the lookout for updates with additional topics.

–Ross Wilcoxon

PRODUCT MATRIX

Continuing steep reductions in the prices of consumer electronics products have made them more affordable in developing and developed economies. Helped by innovations in sensor technology in the form of MEMS components and augmented by communications technology, recent years have witnessed yet another explosion of connected devices such as wearables, fashion-tech wearables, smart clothing, smart jewelry, fitness gadgets, virtual assistants, home security appliances, drones, virtual reality (VR) / augmented reality (AR) headsets, fashion-tech, and even more.

Not all the foregoing categories will need thermal management solutions, and many are well established product lines that are in their maturity and phase out stages of the life cycle. However, their newer versions and respins will inevitably use cost-effective thermal management solutions.

In the design of newer generation consumer electronics products that do use thermal management solutions, the choices are equally challenging for vendors and designers alike – offer the best solution for the problem that is reliable, manufacturable and cost-effective. In the context of increasing expectation and decreasing price tags, achieving this feat is no easy task.

The following tables list typical products, albeit partially, in the consumer electronics and IoT market place (the names were arbitrarily chosen and Electronics Cooling® is not recommending them in any way). Thermal management product component types are also listed.

PRODUCT MATRIX OF CONSUMER DEVICES (A PARTIAL LIST)		THERMAL MANAGEMENT COMPONENTS														
PRODUCT TYPE	PRODUCT EXAMPLES	Electro-Active Polymers	Fans	Gap Pads	Heat Pipes	Heat Sinks	Heat Spreaders	Higher Thermal Conductivity Plastics	Lightweight Liquid Cooling Systems	Liquid Cooling Systems	Phase Change Materials	Resistive Heating Elements	Solders	Thermal Gaskets	Thermal Insulators	Thermal Interface Materials
Mobile Phones	Apple iPhone, Samsung S7, LG, Motorola, Xiaomi			X	X		X	X			X		X	X		X
Digital Cameras	Canon, Nikon, Sony, Olympus , Hasselblad			X			X	X					X			X
Wearables (arm/wrist wear)	Fitbit, Pebble, Apple Watch, Xiaomi, Garmin, Samsung Gear, Withings			X			X	X							X	X
Smart/Comfort Clothing	AiQ Smart Clothing, Sensoria	X							X			X				
Virtual Reality (VR) / Augmented Reality (AR) Headsets	Facebook Oculus VR, Microsoft Hololens, Sensic's OSVR, HTC Vive			X			X	X							X	X
Virtual Assistants	Amazon Echo, Amazon Alexa Device, Google Home		X	X	X		X	X					X			X
Drones	GoPro, Parrot, DJI		X	X	X	X	X	X					X			X
Television & Entertainment Hardware	4K Ultra High-Definition (4K UHD), OLED TV/ Displays, Apple TV, Google Chrome, Harmon Kardon, TiVO		X	X	X	X	X	X					X	X		X
Gaming Desktops / Laptops	Acer, Velocity Micro, Lenovo, SkyTech		X	X	X	X	X	X		X	X		X			X
IOT GATEWAYS, NODES, MOTES, TAGS, ETC. (A PARTIAL LIST)		THERMAL MANAGEMENT COMPONENTS														
Mobile Phones	Apple iPhone, Samsung, LG, Motorola, Xiaomi			X	X		X	X			X		X	X		X
Smart Home IoT Gateways	Amazon Echo, Google OnHub, Motorola, Libelium Meshlium, Rigado IoT Gateways, Logitech Harmony		X	X	X		X	X		X	X		X			X
Edge Node: Occupancy Sensors	Libelium Plug & Sense, Lutron			X	X		X	X			X		X			X
Edge Node: Home alarm systems	Comcast, Libelium Plug & Sense		X	X	X		X	X		X	X		X			X
Edge Node: Power monitoring	Verdigris' Einstein system		X	X	X		X	X		X	X		X			X
Edge Node: Smart lighting	Philips Hue Bridge, Lutron		X	X	X		X	X		X	X		X			X

THERMAL MANAGEMENT MANUFACTURERS		TYPE OF PRODUCT/SERVICE						
COMPANY	CONTACT INFORMATION - URL	Couplers	Design Services	Fans	Heat Sinks & Devices	Liquid Cooling & Thermo Electric	Software	Thermal Interface Materials
AAVID THERMALLOY	www.aavid.com/				X	X		X
Alpha Novatech, Inc.	www.alphanovatech.com		X		X			
Amulaire	www.amulaire.com/			X				
Adaptive Corporation-A TriMech Company	www.adaptivecorp.com						X	
Ansys, Inc.	www.ansys.com/products/fluids		X					
Autodesk	www.autodesk.com/products/cfd/overview		X					
BOYD Corporation	www.boydcorp.com		X	X	X	X		X
Cadence	www.cadence.com						X	
Celsia	www.celsiainc.com		X		X	X	X	
CEJN USA	www.cejn.us					X		
CPC Worldwide	www.cpcworldwide.com	X						
Cradle North America Inc.	www.cradle-cfd.com/		X					
Delta Electronics (Americas) Ltd.	www.delta-fan.com/			X				
DUPONT	www.dupont.com							X
Element Six Technologies	www.e6.com							X
Ellsworth Adhesives	www.ellsworth.com							X
Ferrotec NORD	www.ferrotec-nord.com/					X		
FLIR Commercial Systems, Inc.	www.flir.com/applications/professional-tools/		X					
Fujipoly America Corp	www.fujipoly.com/usa	X	X					X
Future Facilities	www.futurefacilities.com		X				X	
Henkel	www.henkel.com							X
Hexagon	www.hexagonmi.com/mscsoftware		X				X	
Indium Corporation	www.indium.com							
International Manufacturing Services	www.ims-resistors.com				X			
JARO Thermal (NRC Electronics)	www.jarothermal.com/			X	X	X		
Knight-Orion	www.orionfans.com/			X				
Laird Thermal Systems	www.lairdthermal.com							X
Leader Tech	www.leadertechinc.com							X
LMB Fans	www.lmbfans.com			X				

THERMAL MANAGEMENT MANUFACTURERS		TYPE OF PRODUCT/SERVICE						
COMPANY	CONTACT INFORMATION - URL	Couplers	Design Services	Fans	Heat Sinks & Devices	Liquid Cooling & Thermo Electric	Software	Thermal Interface Materials
Malico Inc.	www.malico.com				X	X		
Master Bond Inc.	www.masterbond.com							X
Mentor Graphics	www.mentor.com/products/mechanical/		X					
Mersen	www.mersen.us				X	X		
Nanoramic Laboratories	www.nanoramic.com							X
NeoGraf Solutions	https://neograf.com/				X			X
Panasonic	https://industrial.panasonic.com/ww/products/thermal-solutions							X
Polymer Science	www.polymerscience.com							X
Rogers Corporation - Advanced Connectivity Solutions	www.rogerscorp.com/advanced-electronics-solutions				X			X
Rosenberg USA, Inc.	www.rosenbergusa.com			X				
Sanyo Denki America	www.sanyodenki.us			X				
Shin-Etsu MicroSi	www.microsi.com/product-category/packaging/							X
Sager Electronics	www.sager.com			X	X	X		X
Schlegel Electronic Materials	www.schlegelemi.com							X
Shiu Li Technology Co., LTD	www.shiuli.com.tw							X
Siemens	www.plm.automation.siemens.com/global/en		X				X	
SIKA	www.sika.net/en/products/flow-measuring-instruments.html						X	
Staubli Corporation	www.staubli.com					X		
Techsil	www.techsil.co.uk/applications/thermal-management							X
T-Global Technology	www.tglobal.com/en							X
ThermoElectric Cooling America	www.thermoelectric.com/					X		
TTI, Inc.	www.tti.com			X	X	X		X
TRAN-TEC	www.tran-tec.com				X			
Universal Sciences	www.universal-science.com				X			X

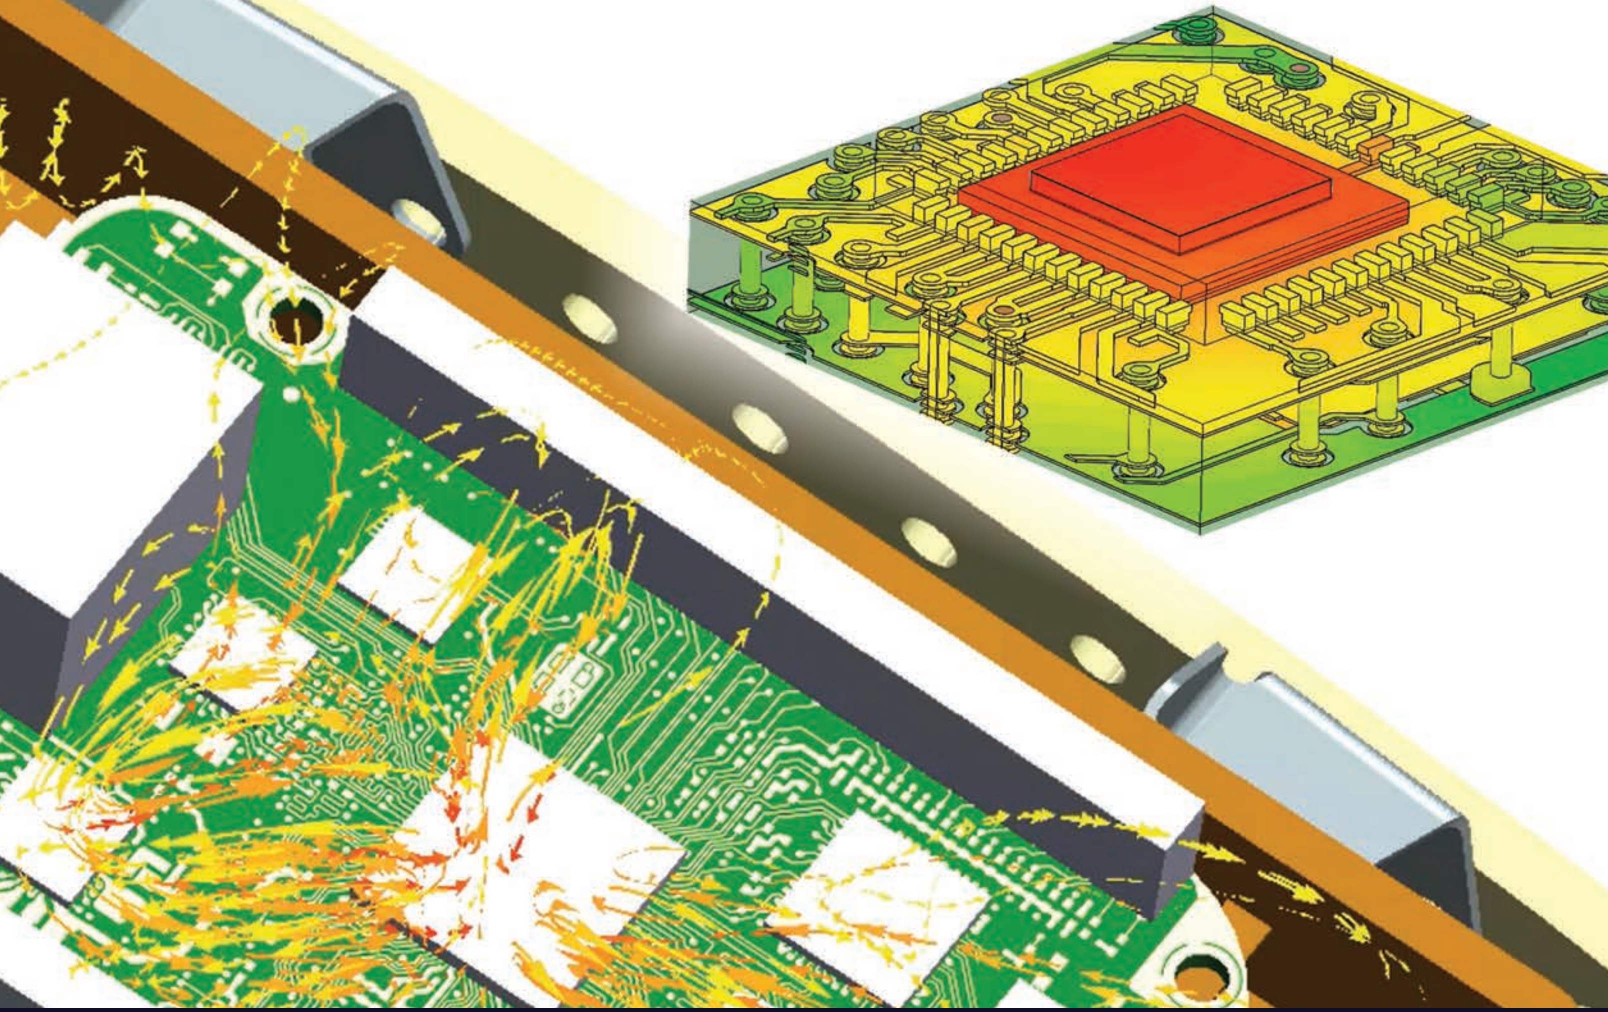

STATISTICS CORNER: PROBABILITY

STATISTICS CORNER: PROBABILITY

INTRODUCTION

I've decided to introduce a new column to Electronics Cooling Magazine – and it's not just because I have run out of good ideas for "Thermal Facts and Fairy Tales" columns. For 2022, I will publish a series of columns in which I try to provide the readers with some insight into the field of statistics and a few tools for effective use of statistical methods. After a couple decades in industry, I have observed that a number of experienced engineers can be intimidated by the topic of statistics – these columns will attempt to reduce the level of intimidation. While I have been interested in statistics for a few decades now, I don't claim to be an expert. I will do my best in these columns to get things as right as I can as well as to make things useful and practical.

Statistical analysis is needed because data always have some degree of uncertainty; a value that we determine from a single measurement, or even set of measurements, is not necessarily going to tell us exactly what value we will determine with additional measurements. Statistical analysis uses the mathematics of probability to create tools that we can use to deal with that uncertainty. This column discusses some aspects of probability concepts to set the basis for how the mathematics of probability can be applied to address uncertainty in statistical analysis.

Any discussion of statistical analysis must include a discussion on probability. Since the entire field of probability and statistical analysis began with gamblers attempting to improve their chances of winning, it seems appropriate that this discussion on probability begins with a game of chance: namely, throwing dice. To begin, I assume that we have an infinite amount of time and patience that allows us to make a lot of throws, the dice that are not loaded (on any given throw they are equally likely to fall with any of its sides up), and we are not playing Dungeons and Dragons, so our dice only have six sides. In other words, I will use an Excel spreadsheet to simulate throwing dice, I trust that the random function is in fact fairly random and that I can calculate the result of throwing a die with the equation "=ROUNDDOWN(RAND()*6,0)+1".

Figure (1) shows what fraction of 30,000 throws of 1-6 dice, as calculated using a simple Excel spreadsheet, had a total value of 1 to 36. For a single die, we would expect that the values 1 through 6 would each occur approximately 1/6th of the time – which is about what reasonably close to what was found in the calculations. As the number of dice included in the throws increases from 1 to 6, the distributions change from a flat line to a triangle to an increasingly 'bell shaped curve'.

Figure 1. Probability Distributions for 30,000 Simulated Dice Throws

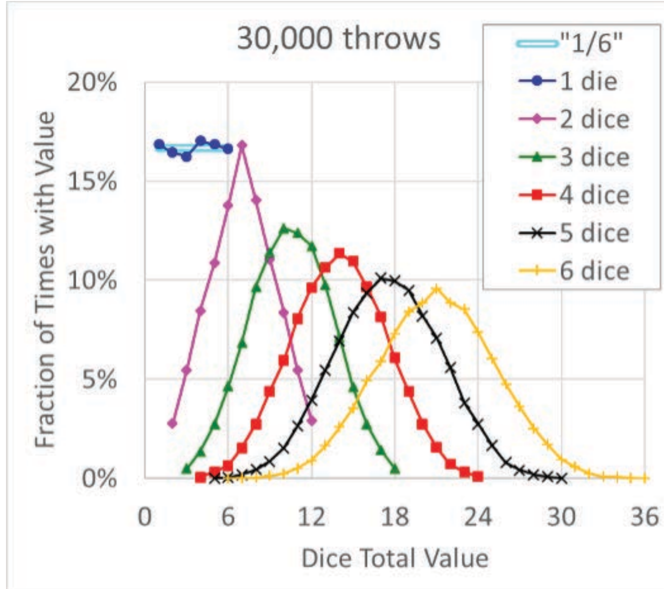

Figure (2) shows the same data but plots the cumulative distribution that show what portion of the throws had a total value that was equal to or less than a value of between 1 and 36. One of the fundamental tenets of probability theory is that the probability of the sum of all possible outcomes is equal to one, which is both logical and illustrated in the figure. In these cumulative distributions, the plots transition from a straight line to a 'tilted S shaped curve' as the number of dice increases from 1 to 6.

Figure 2. Cumulative Distributions for 30,000 Dice Throws

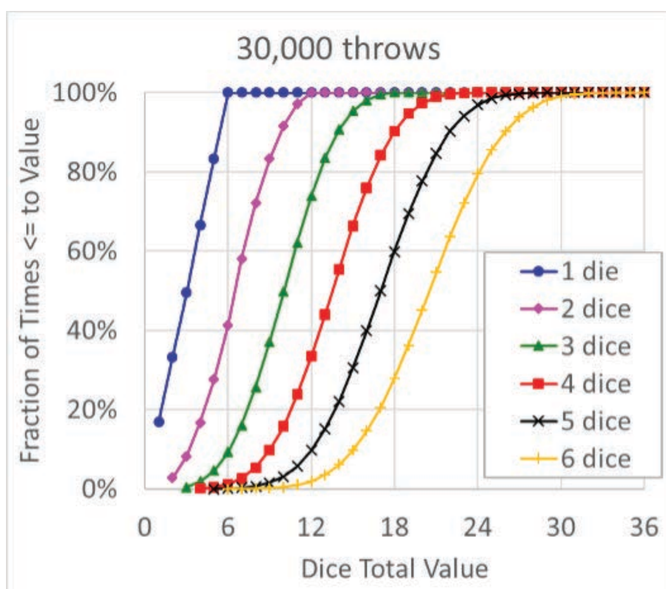

Readers with some (any?) background in statistics likely can see where this is going – the 'bell' and 'tilted S' shaped curves start to look like the Normal distribution that is widely used in statistical analysis. The discussion on that topic will be in the next column in this series.

A question that may be asked is "What would happen if we had to use actual dice and we didn't have the time needed to throw them 30,000 times?". Again, we can simulate that, with results shown for the probability and cumulative distributions for sets of only 30 throws of 1 through 6 dice in *Figure (3)* and *Figure (4)*. *Figure (3)* is best described as an incoherent mess: for the '1 die' data, two values fell exactly on their expected theoretical value of 16.7% while two other values were ~60% higher or lower than that. Data for more than one die do not appear to be much better behaved.

1/6th, so in four rolls his chances should be 4 * 1/6 = 2/3. Since that value is larger than 50% and he was playing even odds (the loser pays the same amount regardless of who it is), he had concluded that it was, on average a winning bet. But when he extended the game to two dice and gave himself 24 attempts to roll a double 6, which by his reasoning should have had the same probability (24 * 1/6 * 1/6 = 2/3), he began to lose money. He asked the mathematician Blaise Pascal to help him understand why his luck had changed.

Figure 3. Probability Distributions for 30 Simulated Dice Throws

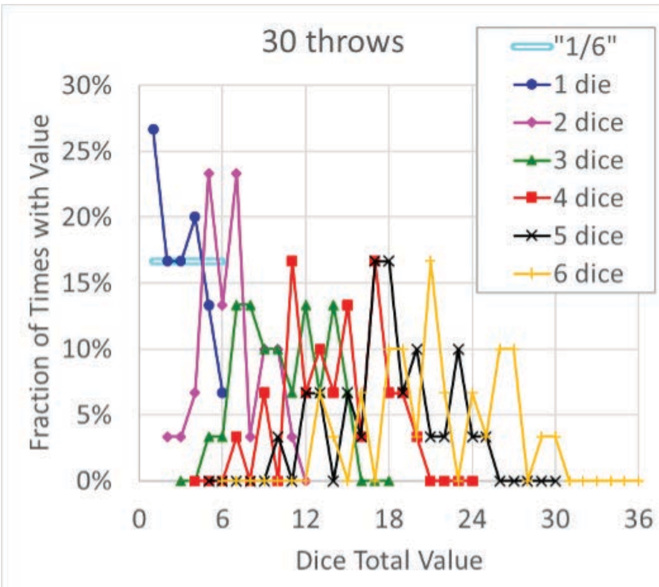

Figure 4. Cumulative Distributions for 30 Dice Throws

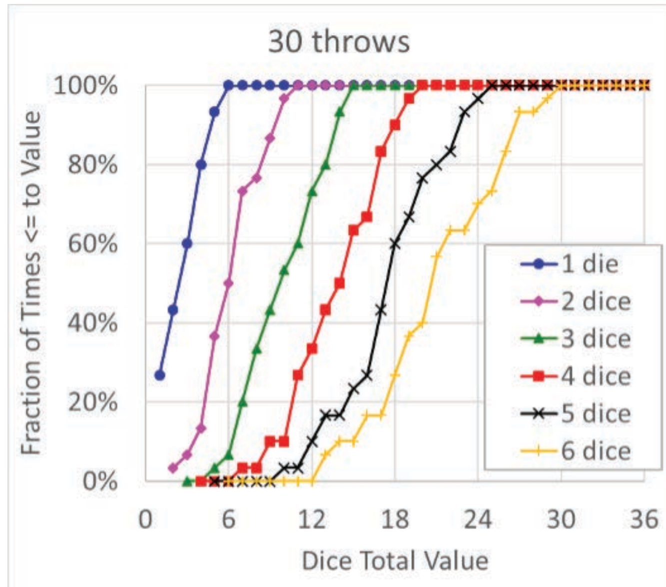

While the cumulative data for 30 throws, *Figure (4)*, shows considerably more jitter than their counterparts for 30,000 throws (*Figure 2)*, the cumulative distributions appear to be much less random than the raw distribution data (*Figure 3*). A comparison of *Figure (3)* and the curves that show the same data in *Figure (4)* illustrates why some data, such as from reliability testing, is often plotted in terms of a cumulative distribution rather than probability.

If a situation is governed by known physics, it can be relatively straightforward to estimate probabilities of a single event. In the case of rolling a single, non-loaded, six-sided die, it should seem obvious that there is a 1/6th chance of any of the six possible outcomes occurring. However, probability calculations can start to become less intuitive when we begin to consider combinations of multiple events. For example, consider the classic question that is considered to have been the beginning of mathematical analysis of probability – the likelihood of rolling a specific value within a specific number of attempts [1]. De Mere, a gambler in the 1600's, tended to win more often than not when he bet that he would roll a 6 within four attempts. His reasoning for why he would win was that the chances of rolling a 6 in one roll was

When calculating probabilities of multiple events, two things that should be kept in mind are that the calculated probability of any outcome must never exceed 100% and that it is often useful to think in terms of an event not happening. In de Mere's case, one simply has to consider the first point to recognize that his equation was incorrect. If the chances of rolling a 6 in four attempts is 2/3, then that equation states that the probability of rolling a 6 in eight attempts will be 133% (4/3). Clearly, this is not possible. To correctly determine the probability of rolling a 6 in four attempts, one can consider the probability of ***not*** rolling a 6 in one attempt and multiply that times itself four times. The probability of not rolling a 6 is (5/6 = 83.3%), so the probability of not rolling a 6 in four attempts is $(5/6)^4$ = 48.2%. Since the probability of not rolling a 6 in those four attempts plus the probability of rolling a 6 in the same attempts must equal 100%, the probability of rolling a 6 in four attempts is 100%-48.2% = 51.8%. This probability is greater than 50%, so with even odds it makes sense that de Mere was coming out ahead. On the other hand, when using the same approach the probability of rolling double 6's in 24 attempts can be calculated as $1 – (35/36)^{24}$ = 49.1%, which is less than 50% and therefore not a good bet at even odds.

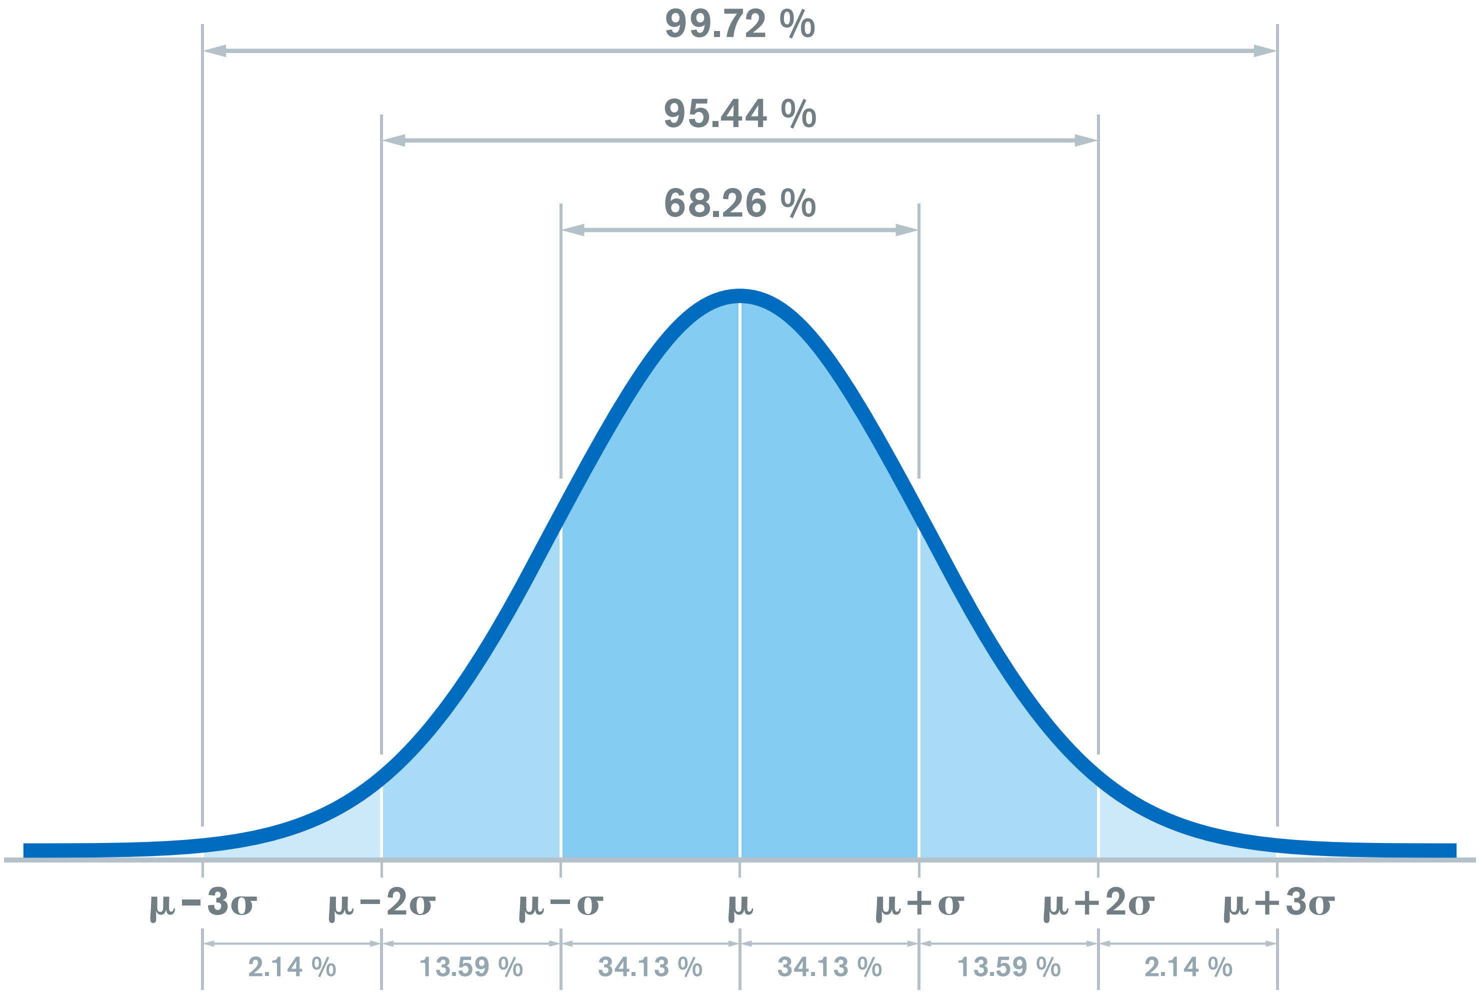

SUMMARY

Probability theory is fascinating and, even the most cursory overview of it, encompasses far more than can be addressed in this short article. This is particularly true if one considers the topic of conditional probability [2], in which the probability of an event depends on another probabilistic event. The example described in this article illustrates that, in a reasonably well-behaved population of data, the effects of measurement variability tend to wash out and lead us to familiar-looking distributions. But it may require a lot of samples from that population to get there. If we only look at a small portion of the population, the distribution won't necessarily appear as a nice bell shape.

Future articles in this series will discuss some of the statistical approaches used to extract useful information and understand the distribution characteristics of data sets that are smaller than 30,000 that led to the smooth curves shown in *Figure (1)* and *(2)*. Topics will include different parameters used to characterize a data set, what confidence we have regarding the uncertainty of those parameters, different models for distributions and how to use them, how to determine if one set of data is different from another, how many samples do we need for a given test, etc.

References

1. https://introductorystats.wordpress.com/2010/11/12/one-gambling-problem-that-launched-modern-probability-theory/, accessed August 24, 2019
2. https://www.probabilitycourse.com/chapter1/1_4_0_conditional_probability.php

STATISTICS CORNER: NORMAL DISTRIBUTION

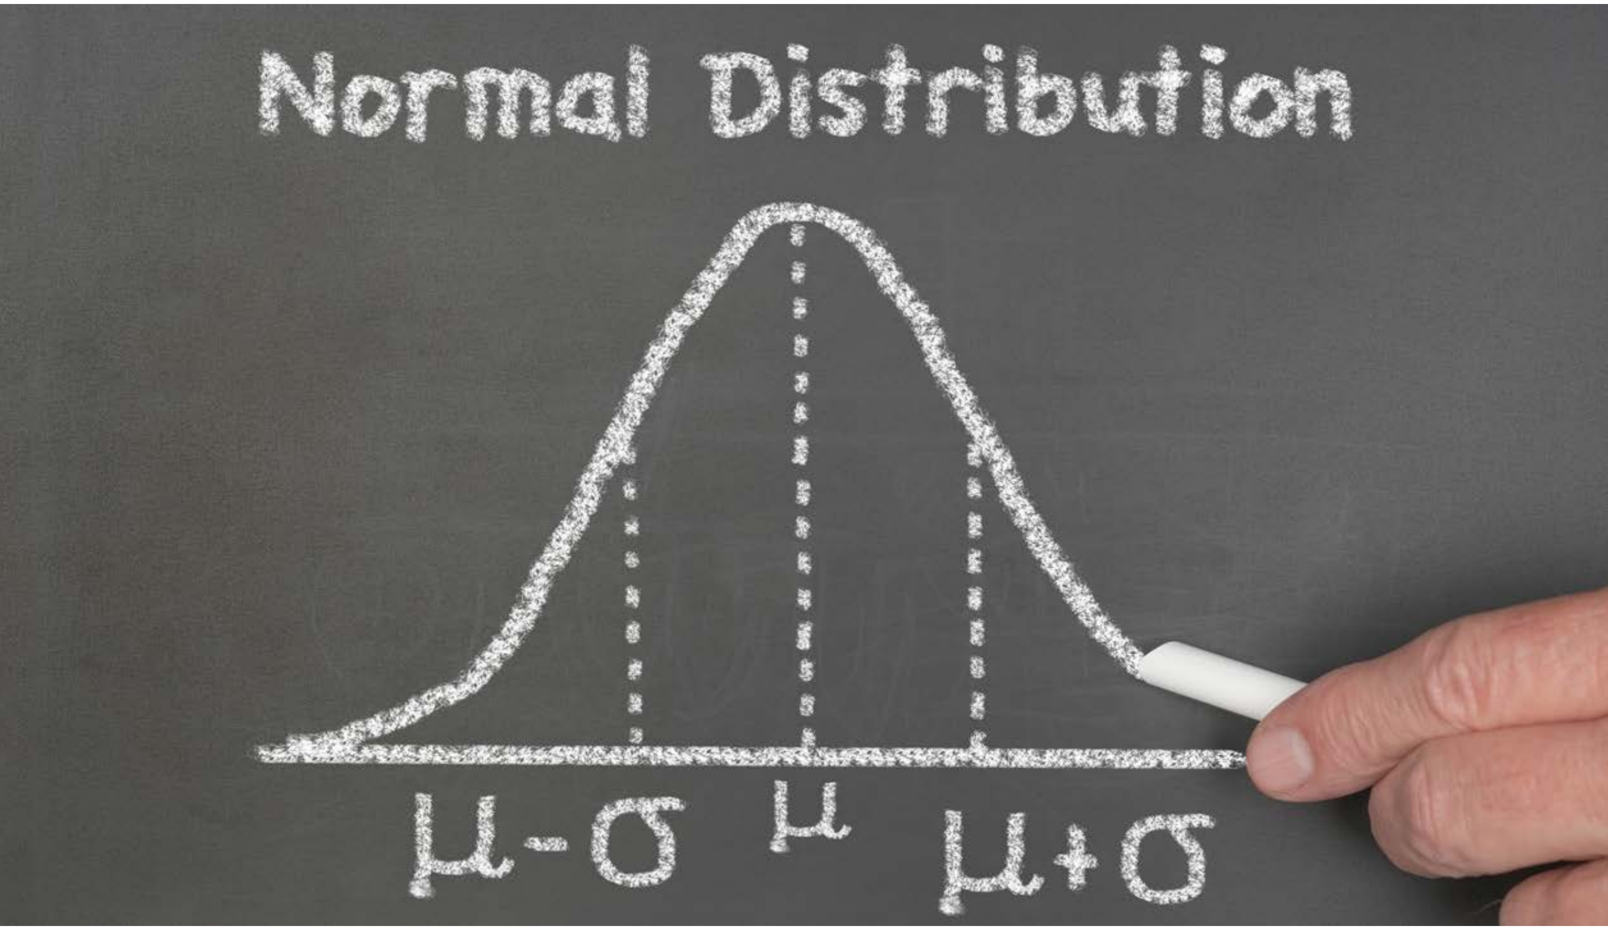

STATISTICS CORNER: NORMAL DISTRIBUTION

The previous column in this series [1] discussed statistical probability and showed that the plot of the probability of a given value occurring within a population can look like a hill in which there is a peak in the middle that tapers off to increasingly smaller slopes on each side. The example in that column referred to the scores produced by shaking a number of dice and the 'hilly' plot was described with the more common term of 'bell shaped curve'. The purpose of establishing a probability distribution to describe a population of data with uncertainty is that it provides a mathematical framework for dealing with that uncertainty - if a reasonable mathematical model of the distribution curve can be defined. As we will see in subsequent columns, many different models for probability distributions exist; the selection of the correct model depends on characteristics of the population and what data are available for analysis.

This column focuses on the probability model that is most widely used and most recognized: the Gaussian, or normal, distribution. The normal distribution, as written in terms of a probability density function is shown in *Equation {1}:*

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \qquad \{1\}$$

Equation {1} allows one to estimate the probability of a given value, x, occurring in a population that is defined with the two parameters μ and σ. While the normal distribution equation itself may not be familiar to everyone, the terms μ and σ should be recognizable to anyone who deals with data: μ is the mean (or average) value and σ is the standard deviation. If the mean and standard deviation of a population are known, the probability that a randomly selected member of that population will have a value of x can be calculated with *Equation {1}*. Or, if you are like me and rely on spreadsheets to do most of your calculations, you can use the function @norm.dist(x, μ, σ, false)[1].

The mean[2] and standard deviations can be calculated for a set of N samples, x_1, x_2, …x_n, by:

$$\mu_s = \frac{\sum_{i=1}^{N} x_i}{N} \qquad \{2\}$$

$$\sigma_s = \sqrt{\frac{\sum_{i=1}^{N}(x_i-\mu_0)^2}{N}} \qquad \{3\}$$

Equations {2} and *{3}* include the subscript 's' for the mean and standard deviation as a reminder that the values determined from a sample (the set of data drawn from a population) are not exactly equal to those for the entire population (all possible elements). Traditionally, the mean and statistical deviation of a sample are written as $\overline{x}$ and s, respectively. A future column in this series will discuss how to use the values of μ_s and σ_s, also known as $\overline{x}$ and s, to determine a range in which we can be confident that the population mean, μ, actually lies. Note that the mean and standard deviation of a data set can be calculated in a spreadsheet with the functions @average('data') and @stdev('data') respectively, where 'data' refers to the cells that contain the data.

Our familiarity with the term 'average' can lead to its occasional misuse, which can be avoided if we keep its relationship to the normal distribution in mind. The version of the classic illustration of a misuse of the term 'average' is the example of 9 people, who all have a net worth of $500,000, are sitting in a bar. In walks the founder of a "multinational conglomerate technology company that focuses on e-commerce, cloud computing, digital streaming, and artificial intelligence" who has a net worth of $140B. Using *Equation {2}*, the average net worth of the individuals in the bar suddenly increases to $14B, which may be mathematically correct but not physically relevant. The primary basis for this discrepancy is the fact that the population of the 10 individuals in the bar is not representative of a normal distribution. In cases like this, the median may be a more appropriate parameter for reporting a typical value. The median is the middle value in a ranked list of the sample set such that an equal number of values are greater than and less than it (the spreadsheet function for calculating median is @median('data'). When the mean and median of a data set are substantially different, such as in the aforementioned example of people in the bar, one should suspect that the sample set is not normally distributed.

A fundamental strength and justification for utilizing the normal distribution is the Central Limit Theorem, which shows that when multiple, independent random variables are added, the result tends towards a normal distribution – even if the variables themselves are not normally distributed (see for example reference [2]). Reference [1] discussed the results of throwing dice, beginning with the assumption that a single die is 'fair' such that the probability of it showing any particular value is equal to the inverse of the number of sides on the die (i.e., 1/6th for a six sided die). The distribution of scores that result from throwing that one die is certainly not normal; it would be a

1 The 'false' in this equation specifies that the probability distribution (a bell-shaped curve that goes to zero as x goes to infinity) is calculated. If 'true' is used instead, the function returns the cumulative distribution (an S shaped curve that goes to 1 as x goes to infinity)

2 *Equation {2}* defines the arithmetic mean, which is the same as the arithmetic average. The more generic terms mean and average are often used interchangeably but can refer to different definitions.

straight line with equal probability of 16.7% for each of the six values. With two die, the distribution had a triangular shape and, as the number of die increased, the distribution looked more and more like a bell shaped curve.

Figure (1) shows simulated results for throwing 6 six-sided die 100 and 5000 times. The larger number of throws leads to a more well-behaved distribution of results (red bars appear to be more 'bell shaped'). However, both data sets produce very similar normal distributions with mean values of ~21 and standard deviations of 4.2. This illustrates the power of the normal distribution and the results of the central limit theorem. Even when we have results from a small data set that in of itself does not appear to be normally distributed (have a 'clean' bell shaped curve), if the data were drawn from a normal distribution there is a good chance that they can be used to accurately estimate the fundamental characteristics of that population.

The normal distribution provides a straightforward method for using measurable characteristics of a data set (the mean and standard deviation of the sample) to estimate the probabilities of future measurements falling within a prescribed range of values (related to the properties of the entire population).

This is incredibly useful in that it allows us to perform tasks such as:

• Determining how many samples must be measured to have confidence that a population has been adequately characterized.

• Comparing different data sets and decide with they are from the same population or not (e.g. to determine if differences in their mean values are 'statistically significant').

• Deciding whether a value that seems to be an outlier is likely to be from the population that we are evaluating or if it is due to a factor such as a measurement error.

• Defining how much confidence we should have in a curve fit we generate from a data set.

These types of practical tools are all topics that will be discussed in future columns, now that these foundational topics of probability and distributions have been covered.

Figure 1. Simulated results for throwing 6 six-sided die

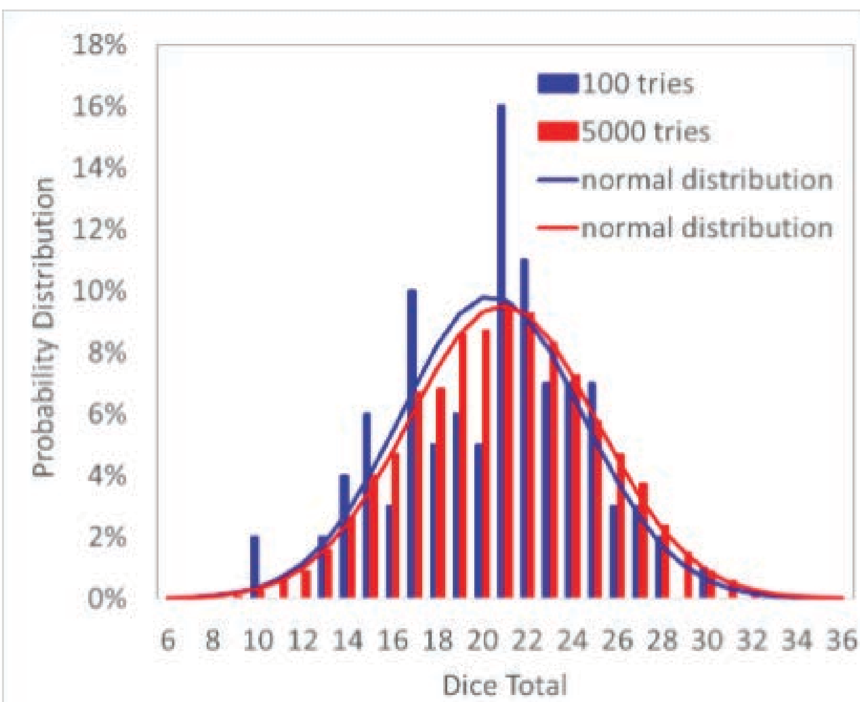

References

1. Ross Wilcoxon, "Statistics Corner – Probability", Electronics Cooling Magazine, Spring 2020, pp. 16-18
http://sphweb.bumc.bu.edu/otlt/MPH-Modules/BS/BS704_Probability/BS704_Probability12.html

STATISTICS CORNER: CONFIDENCE INTERVALS

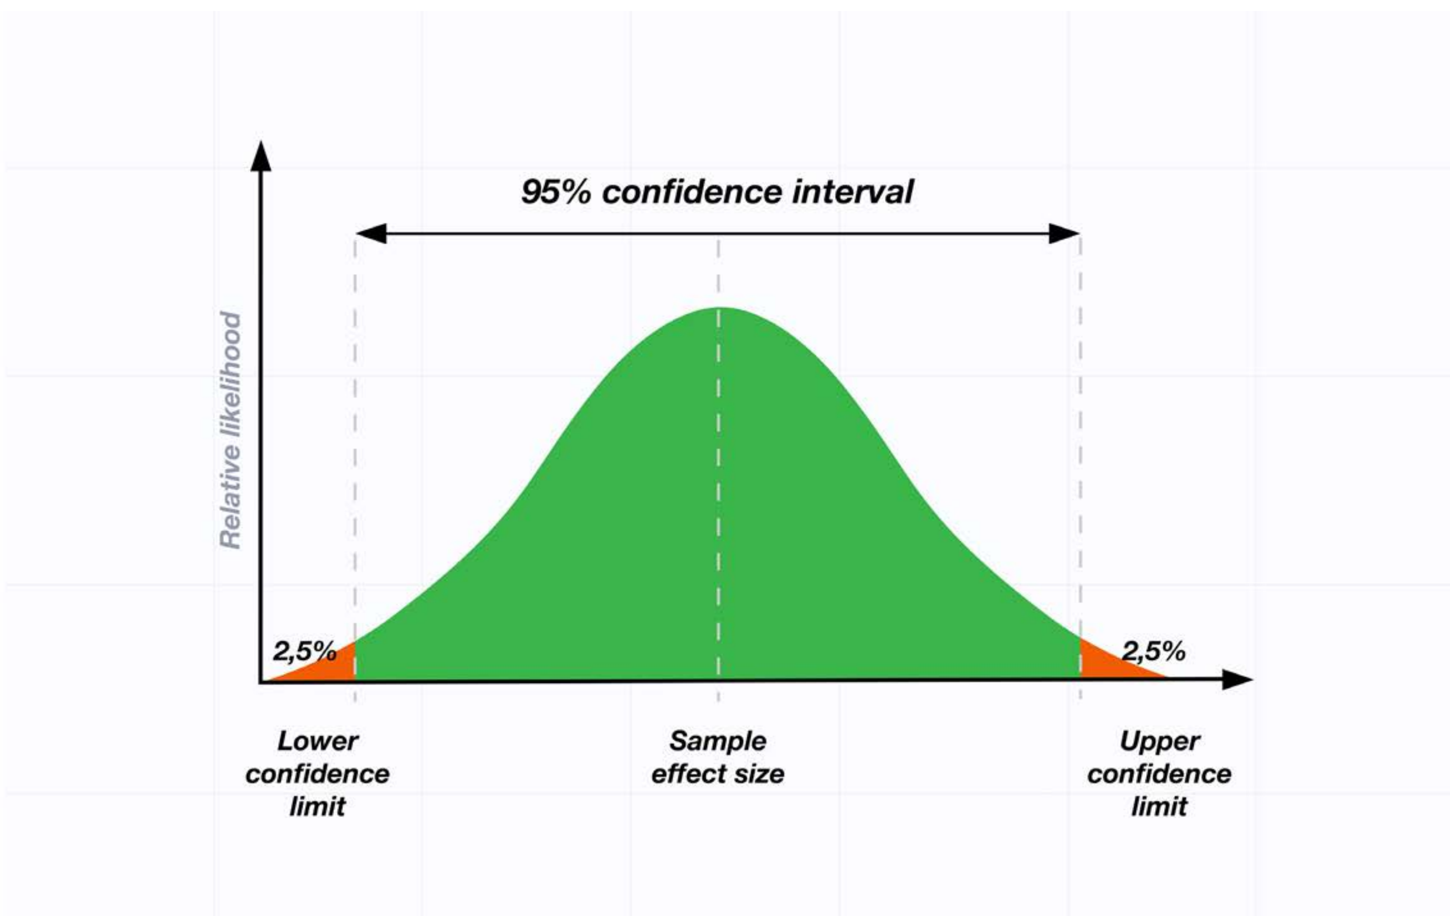

STATISTICS CORNER: CONFIDENCE INTERVALS

BASIS OF A VARIETY OF STATISTICAL ANALYSES

The previous articles in this series [1, 2] described how the mean and standard deviations of a set of data are calculated and how they can be used to estimate the characteristics of a population using the normal distribution. Since measured data typically represents only a subset of an entire population, one should recognize that the estimated mean and standard deviation values determined from a sample set are not likely to be exactly equal to the true values of the entire population. In other words, when we calculate the mean, i.e., average, of a data set, we should recognize that there is actually a range of values in which the true population mean lies – and the size of that range depends on the confidence level that we assign to our estimate. This article describes a process that can be used to determine that range as a function of the number of data points and confidence level. More generally, this outlines the overall process for linking probability distributions with confidence levels that is the basis of a variety of statistical analyses.

Figure (1) shows an example of a normal distribution in which the mean value is 8 and the standard deviation is 0.75. Two lines are included in this plot: the probability distribution corresponds to the probability of any specific value occurring within the population while the cumulative distribution indicates what portion of the population is less than or equal to a given value[3]. The cumulative distribution curve is equal to the area under the curve of the probability distribution.

Figure 1. Normal distribution for mean of 8 and standard deviation of 0.75

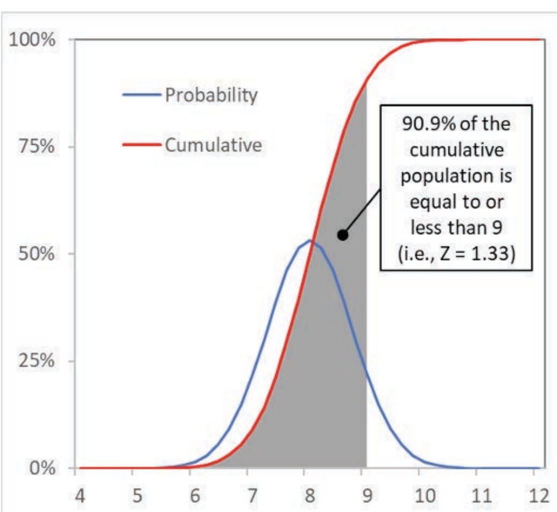

Z-value is defined as the difference of a value from the mean, normalized by the standard deviation. For example, with the mean of 8 and standard deviation of 0.75, a measurement of 9 would have a Z-value of (9-8)/0.75 = 1.33. The cumulative normal distribution for this value is 90.9%[4]; in other words, 90.9% of the population with a normal distribution would have a Z-value that is less than or equal to 1.33. This is illustrated in *Figure (1)* as the area under the cumulative probability curve for values of 9 and below.

As mentioned previously, we don't generally know the true mean and standard deviation of an entire population because we only make measurements of a subset of it. The larger the subset (the more samples we use in making our estimates), the more accurate our estimates of the true values should be. When we use a small sample size in an analysis, we can account for the additional uncertainty due to sample size by 'flattening' the bell curve of the standard normal distribution and increasing the size of the 'tails' (the areas under the curve that are far from the midpoint) using the Student t-distribution.

The t-distribution looks very similar to a normal distribution, but its specific shape depends on the number of degrees of freedom. The degrees of freedom (DoF) for a sample set corresponds to the number of independent values used in the analysis. In this case, we can consider a data set of n data point to be comprised of n-1 independent values; the difference between the sample size and DoF is due to the use of the data points to estimate the population mean [3]. Thus, one data point is not independent and the other data points can be used to assess to assess the uncertainty. *Figure (2)* shows t-distributions for 2, 4 and 10 DoF and compares them to the normal distribution, again for a mean of 8 and standard deviation of 0.75. As the DoF increases, the t-distribution converges to the normal distribution; above 30 DoF, the t-distribution is virtually identical to the normal distribution.

---

[3] To use Excel to calculate the values of these curves, use the function =norm.dist(x, mean, stdev, dist), where x is the x-axis value, mean and stdev are the mean and standard deviation of the population respectively, and 'dist' is FALSE for the probability distribution and TRUE for the cumulative distribution.

[4] This value can be calculated with Excel in two different ways (at least). The simple approach is the function =norm.dist(9,8,0.75,TRUE). Another approach is to use the Z-value with a standard normal distribution, which has a mean of 0 and standard deviation of 1. Thus, the function would be =norm.dist(1.33,0,1,TRUE). Both of these functions will return the same value of 0.9088.

Figure 2. Normal and t-distributions for mean of 80 and standard deviation of 0.75

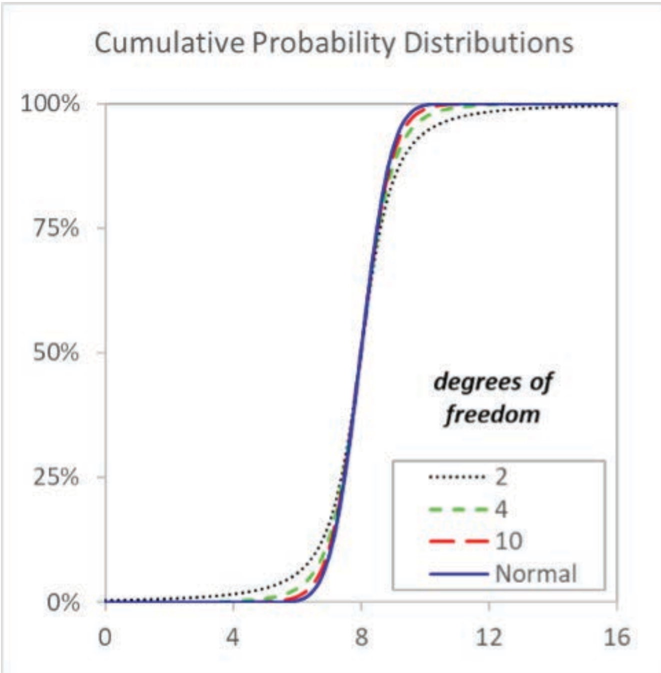

THE CENTRAL LIMIT THEOREM

With this baseline information on distributions established, we can now describe how they are used to define what range of values the true mean of a population is based on a small number of measurements. The Central Limit Theorem is a critical component in establishing this. The Central Limit Theorem states that a population of terms $(X-\mu)/(\sigma/n^{1/2})$ tends towards being normally distributed, where X is a value in the population, μ is the mean, σ is the standard deviation and n is the number of samples used to estimate the population characteristics. This applies not only to the data, but also to the distribution that we use to assess our confidence in the mean that is calculated from a sample set.

For example, assume that 50 measurements of a heat sink show that its thermal resistance is 8°C/W with a standard deviation of 0.75°C/W. How confident can we be that the actual population mean is somewhere between 7.9 and 8.1°C/W? Using the Central Limit Theorem, we calculate a test statistic as $(7.9-8)/(0.75/50^{1/2}) = -0.943$. We can look this value up on a standardized normal distribution table, which shows that 17.3% of a normal population that has a mean of 0 and standard deviation of 1 will have a value of -0.943 or less[5]. Since the normal distribution is symmetric, we will also find that 17.3% of the population will have value of 0.943 or greater. Thus, 34.6% of the normal distribution is either less than 7.9 or greater than 8.1 and there is a confidence band of ~65% that the true mean is between 7.9 and 8.1.

Typically, we are more interested in conducting the reverse analysis – namely, what range of values corresponds to a prescribed confidence band. Also, we may not have the luxury to have a sufficient number of measurements (more than 30) to justify using a standard normal distribution, rather than a t-distribution, in our calculations. The steps for determining the range of mean values correspond to a specified confidence band are shown in *Table (1)*. This table includes example calculations for testing on ten heat sinks that again showed a mean thermal resistance of 8°C/W with a standard deviation of 0.75°C/W. The goal of the analysis is to determine the range of values that we can be 90% confident that the true mean lies within.

Figure (3) shows the 90% confidence intervals calculated for the t- and normal distributions for data with the same mean and standard deviations, but different sample sizes. As the sample size increases, the lines for the confidence intervals come converge. At small sample sizes, the confidence bands that are calculated using the more appropriate t-distribution are much wider than those calculated using the standard normal distribution. In general, it may be questionable that an extremely small sample size of 2-3 samples is necessarily representative of a population

Table 1. Steps for determining the range of mean values corresponding to a specified confidence band

#	Step	Example (with Excel Functions)
1	Measure n samples and calculate the nominal mean, μ, and standard deviation, σ	n = 10, μ = 8°C/W, σ = 0.75°C/W
2	Define confidence interval, C.I. that defines the range of means, leading to the size of the tails outside the C.I., $\alpha/2$	C.I. = 90%; two tails (one on each side of the distribution) => area in each tail is (1-0.9)/2 = $\alpha/2$ = 0.05
3	Calculate the inverse t-distribution for the tail size and degrees of freedom, $t_{\alpha/2}$	$t_{\alpha/2}$ = t.inv((1-α/2),(10-1)) = t.inv(0.95,9) = 1.833 *t.inv inputs: probability of being outside the tail (1-α/2), DoF (sample size – 1)*
4	Calculate the distance from the lower limit of the confidence interval to the mean, Δ	$\Delta = t_{\alpha/2} * \sigma / n^{1/2}$ =1.833 * 0.75 / sqrt(10) = 0.43
5	Calculate the range that corresponds to the confidence interval = $\mu \pm \Delta$	Confidence interval: 8 - 0.43 to 8 + 0.43 C.I. = 7.57 to 8.43°C/W

5 Or we can use the Excel function =norm.dist(-0.943,0,1,true) if that seems easier…

– and that small of a sample size also leads to an uncomfortably large confidence interval. But when dealing with typical sample populations with 5-20 measurements, the t-distribution provides a reasonable approach for estimating the confidence interval and can be evaluated for whether it is likely from a normal distribution (a discussion for a future article). When the sample size is greater than 30, the normal distribution can be used to assess the confidence interval. However, since the t-distribution converges to the normal distribution at large sample sizes, one can continue to use the t-distribution even with larger populations. So in general, if these equations are incorporated into a tool such as Excel, it is appropriate to use the t-distribution even for very large sample sizes.

SUMMARY

In summary, the goal of the first three articles in this series has been to provide a sufficient background to allow readers to better understand future articles aimed at providing practical statistical analysis approaches. Hopefully, the articles did not achieve a 'worst of both worlds' status in which they included more theory than engineers might want to see and less theory than statisticians would expect. Regardless of whether they achieved that or not, with a basic statistical foundation in place we can now move on in future articles to describe tools and analysis methods for solving the types of statistical problems that engineers may encounter.

Figure 3. Confidence bands calculated for different sample sizes of heat sink resistance

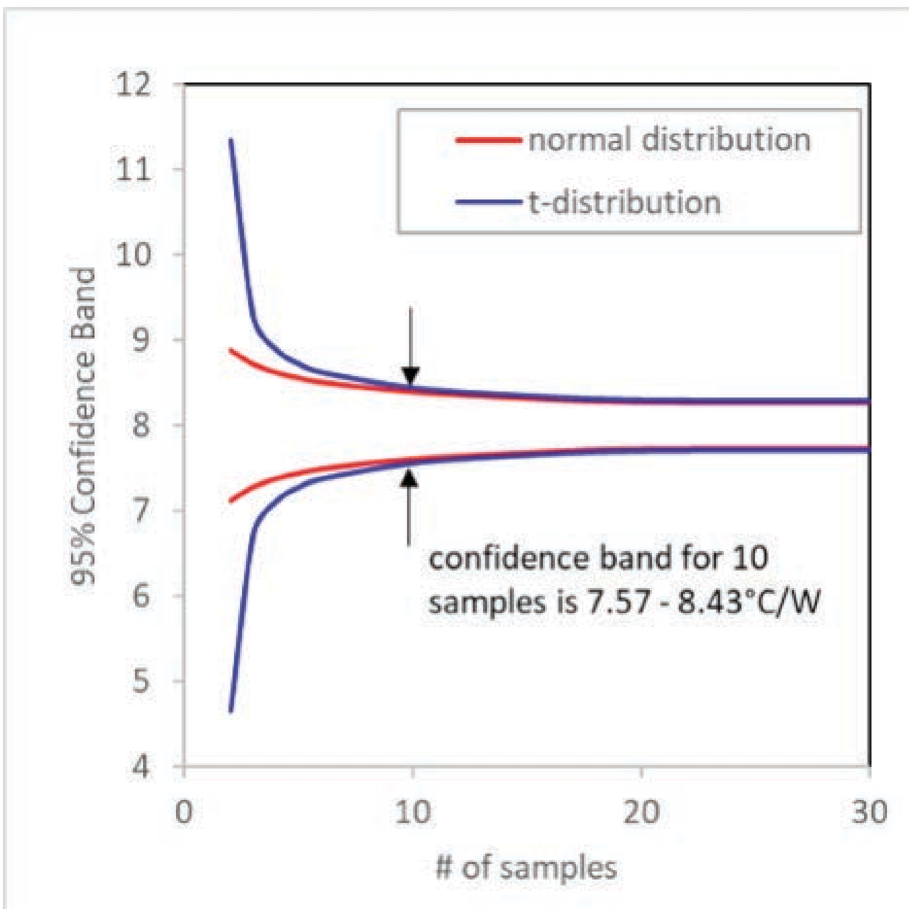

References

1. Ross Wilcoxon, "Statistics Corner—Probability", Electronics Cooling Magazine, Spring 2020
2. Ross Wilcoxon, "Normal Distribution", Electronics Cooling Magazine, Summer 2020
3. https://blog.minitab.com/blog/statistics-and-quality-data-analysis/what-are-degrees-of-freedom-in-statistics


Call for Authors and Contributors!
Want to be a part of the next issue of Electronics Cooling? Have an article or blog post you'd like to write for Electronics-Cooling.com?
Let us know at
editor@electronics-cooling.com
electronics COOLING
www.Electronics-Cooling.com

STATISTICS CORNER: COMPARING POPULATIONS

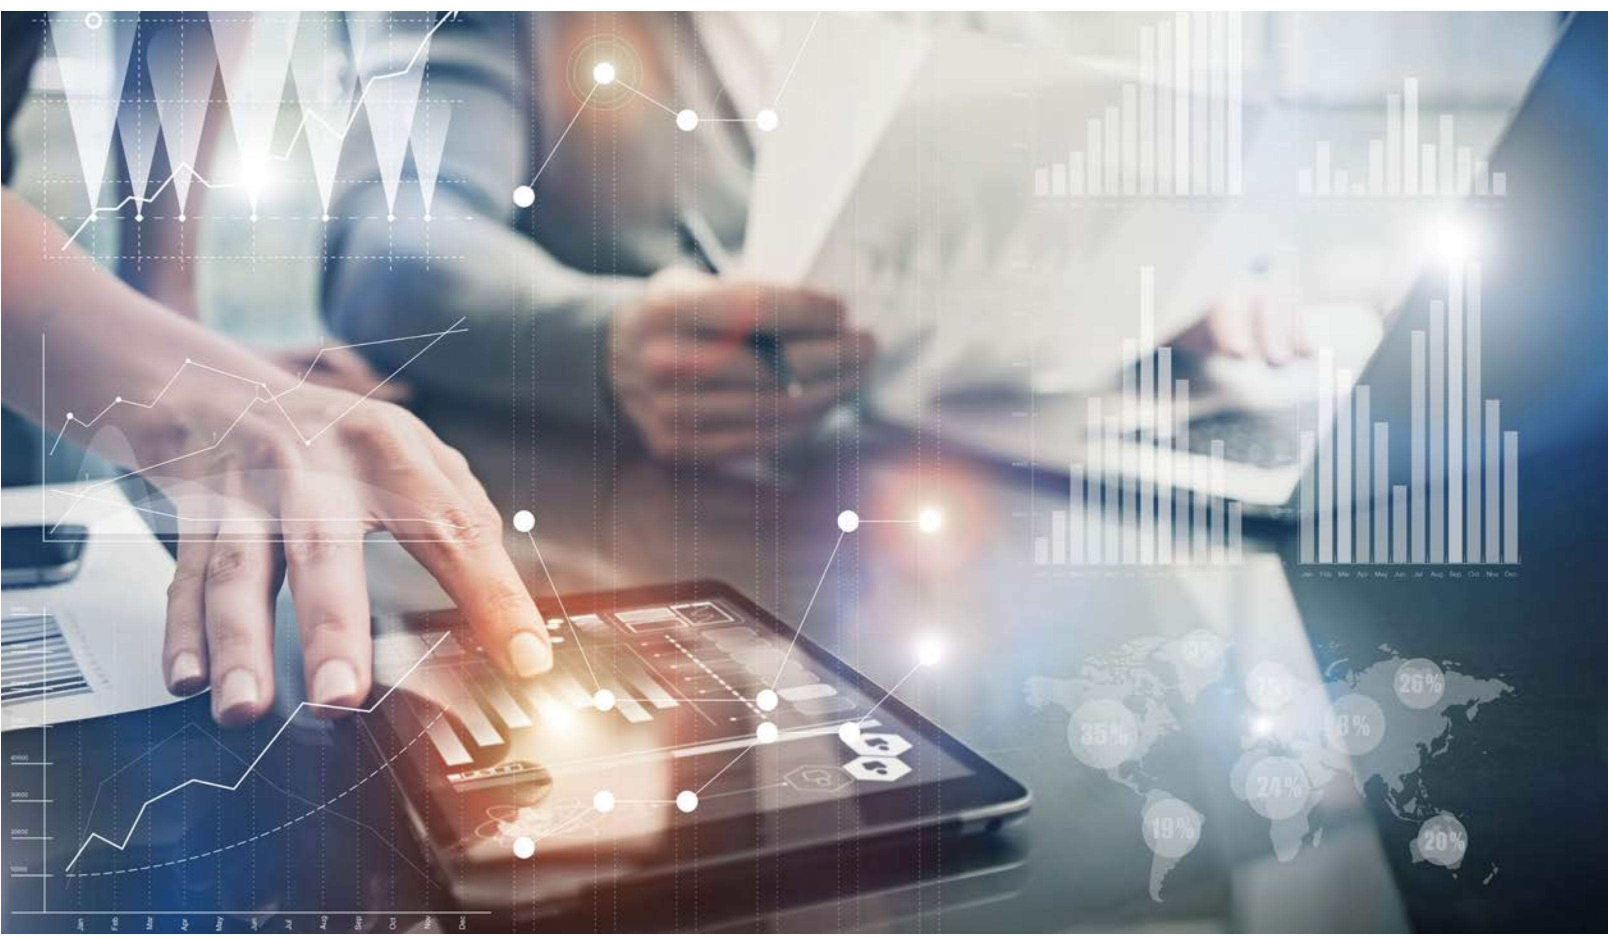

STATISTICS CORNER: COMPARING POPULATIONS

INTRODUCTION

Previous articles in this series described the normal distribution and how it is used to relate probability and confidence levels [1, 2]. A practical application of the use of the confidence interval was to describe how to determine the range of values in which the true mean of a population falls within, based on the mean and standard deviation calculated from a set of samples drawn from the population. This range depends on the confidence level that is selected for the analysis and the number of samples used to estimate the population. The fewer samples or the higher the confidence level that is selected, the wider the range in which the true mean value may lie.

COMPARING POPULATIONS

This article discusses how the confidence band concept is extended to compare different sample sets to determine, within a set confidence level, whether the two sample sets have the same mean. This provides a statistically established method to determine if two data sets are different, thereby demonstrating whether a treatment, design change, etc., lead to an improvement.

The following example illustrates the overall procedure for comparing two populations to determine whether they are different. Somewhat rudimentary testing[6] was used to evaluate a natural convection heat sink in the three different orientations illustrated in *Figure (1)*. In the Horizontal orientation, the base of the heat sink was aligned with the direction of gravity and the fins were perpendicular to gravity. In the Vertical orientation, the heat sink base was again aligned with gravity, but the heat sink was rotated 90° such that the fins were parallel to gravity. In the Flat orientation, the heat sink faced upwards with the base perpendicular to the direction of gravity. *Table (1)* shows individual thermal resistances calculated from different tests and calculated values for the number of data points, average, and standard deviation for each heat sink orientation. In this testing, power was dissipated from a heater attached to the back of the heat sink. The average heat sink temperature was determined with four thermocouples attached to the heat sink and the thermal resistances were calculated by dividing the temperature difference (average heat sink temperature minus ambient air temperature) by the heater power dissipation. Tests were repeated over a number of days with different heater powers and with the orientations relative to gravity randomized. Testing did not control for the effects of radiation, minor room drafts, heat losses from the back of the heat sink, etc. Therefore, the data shown in *Table (1)* are primarily useful for comparing the effects of orientation, and do not represent a controlled investigation to determine the precise heat sink resistances. While testing was conducted over a broad range of power dissipations, the data in *Table (1)* are limited to rest data for power dissipation in the range of 20-25W.

Figure 1. Heat sink orientations for natural convection testing

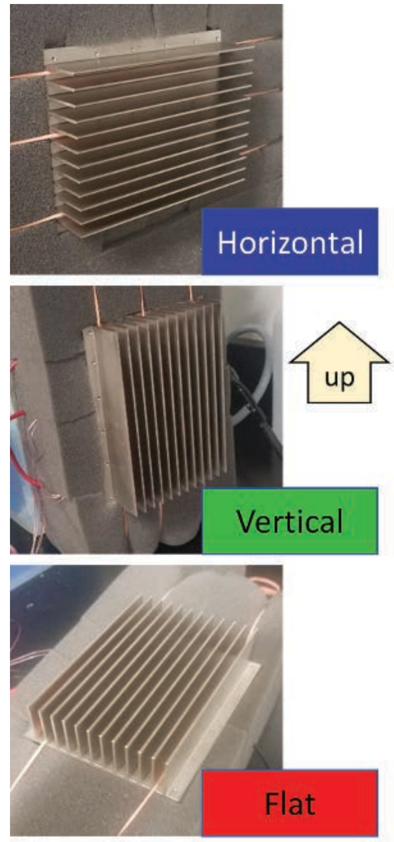

Table 1. Heat sink thermal resistance test data

	Horizontal	Vertical	Flat
Calculated Thermal Resistance Values (K/W)	3.75	1.82	1.93
	3.80	1.83	1.96
	3.82	1.85	1.97
	3.82	1.85	2.00
	3.94	1.90	2.01
		1.91	2.01
		1.92	2.05
		2.13*	2.08
			2.08
Sample Size	5	8	9
Average	3.82	1.90	2.01
St. Deviation	0.0676	0.1003	0.0541

[6] Note the rubber bands and packing foam used in the test setup.

Based on the average values of the thermal resistances, it appears that the Horizontal orientation has a significantly higher thermal resistance (of 3.82 K/W) than the Vertical orientation (resistance of 1.88 K/W) and the Flat orientation (2.01 K/W). The questions that we will attempt to answer in this analysis are 1) what is our confidence that the Horizontal thermal resistance is actually different from the other orientations and 2) is there a statistically significant difference in the thermal resistances of the Vertical and Flat heat sink orientations?

Using the procedure described in [1], we can use the calculated mean, standard deviation and sample size to determine the range in which the true mean lies, for a given confidence level, of each set of measurements. For example, with the horizontal data, in which 5 measurements produced an average of 3.82 and standard deviation of 0.676, we can be 95% confident that the true mean falls between 3.74 and 3.91[7]. The parameters used to calculate the confidence intervals and the ranges for the 95% confidence interval for the mean thermal resistance for each of the three heat sink orientations are shown in *Table (2)*.

Table 2. Terms used to determine 95% confidence intervals for thermal resistance of each orientation

		Horizontal	Vertical	Flat
Distribution Statistics	# samples, n	5	8	9
	mean, μ	3.824	1.901	2.010
	st dev, σ	0.06763	0.1003	0.05414
Parameters for 95% Confidence	$t_{0.025}$	2.776	2.365	2.306
	delta, Δ	0.0840	0.0839	0.0416
Range in Which Mean Falls	min	3.74	1.82	1.97
	max	3.91	1.98	2.05

The results in *Table (2)* indicate that we can be more than 95% confident that the thermal resistance of the heat sink in the Horizontal orientation is higher than in the other two orientations, since there is no overlap between its range and the others (3.74 is greater than both 1.98 and 2.05). However, with this approach and using a 95% confidence level, indicates the Vertical and Flat thermal resistances are not statistically different since there is a slight overlap in their ranges (1.97 falls between 1.82 and 1.98).

This analysis can be improved by recognizing that the size of the range for a given confidence level is due to a combination of the standard deviations and the number of samples of each test set. Even if the means of two populations are different, their standard deviations may be similar enough that we can 'pool' the data and increase the effective sample size by accounting for the number of samples in each data set.

We can use the F-distribution to determine whether we can pool data to increase the effective sample size. Without going into the detail that it deserves, the F-distribution is a statistical distribution that can tell us the probability that the variances, i.e., standard deviations, of two populations are different – in a similar manner that the t-distribution indicates the probability that the means of two populations are different. We can use the Excel function @f.test(array1, array2) to determine whether the standard deviations of two populations are different, and therefore that data can be pooled to determine the confidence range for the mean. The arrays in that function are the two sets of measurements under consideration and the function returns the probability that the standard deviations are different.

In comparing the Vertical and Flat data, the F-test function returns a value of 10.5%. Since this is greater than 5%, we cannot conclude (at a 95% confidence level) that the variances (standard deviations) of the two populations are different. Therefore, we can pool the data in applying the t-test that determines whether the means of the two population are different. The pooled standard deviation, σ_p, can be calculated using the estimated standard deviations (σ) and sample sizes (n) of each sample set, as shown in *Equation {1}* [3].

$$\sigma_p^2 = \frac{(n_1-1)\sigma_1^2+(n_2-1)\sigma_2^2}{n_1+n_2-2} \qquad \{1\}$$

The pooled standard deviation of the Vertical and Flat data is calculated to be 0.006262. This is then be used to calculate the test statistic, T, for the pooled data, which is determined with *Equation {2}*:

$$T = \frac{|\mu_1-\mu_2|}{\sqrt{\sigma_p^2\left({}^1/_{n_1}+{}^1/_{n_2}\right)}} \qquad \{2\}$$

For the pooled Vertical and Flat data, this is calculated to be 2.8259. This value is then compared to the t-value that is calculated for the selected confidence level and the total degrees of freedom, which is the sum of the samples from the two data sets minus 2 (since one degree of freedom was 'consumed' in calculating each of the two means). For a 95% confidence level and a two tailed distribution, the input parameter for the t-test is $\alpha/2$ = (1-0.95)/2 = 0.025. In Excel, a value of $\alpha/2$ that is less the 0.5 will return a negative number, so the t-value for

7 Calculation method: n = 5, μ = 3.824, σ = 0.06763, d.f. = n-1 = 4. 95% confidence for a two tailed distribution => α /2 = (1-0/95)/2 = 0.025; t-statistic in Excel is $t_{0.025}$ = abs(t.inv(α/2, d.f.) = abs(t.inv(0.025, 4) = 2.7764
Δ = $t_{0.025}$* σ /n1/2 = 2.7764*0.06763/$5^{1/2}$ = 0.084, mean range = μ ± Δ = 3.824 - 0.084 to 3.824 + 0.084 = 3.74 to 3.91

the pooled sample can be calculated with =-t.inv(0.025, (8+9-2)) = 2.13144[8]. Since this value is less than the test statistic of 2.8259, we can be 95% confident that average thermal resistance for the Vertical heat sink is different from the Flat heat sink. If we repeat the analysis with a confidence level of 98.72%, the t-statistic is equal to the T-value of 2.8259. If one prefers to type in fewer equations, a simpler approach to reach the same conclusion is to use the function =1 – t.test(array1, array2, 2, 2), where the two arrays are the two sets of data for the Vertical and Flat orientations. In this function, the first 2 indicates that the distribution is two-tailed, and the second 2 indicates that the two data sets are homoscedastic (with the same variance, as determined by the F-test). This function, with the two sets of data returns a value of 0.9872; since this value is larger than the 95% we can therefore conclude with a 95% confidence level that the thermal resistances for the Flat and Vertical orientations are different.

If the F-test finds that the variance of the data sets ae unequal, the Smith-Satterthwaite procedure [3] can be used to estimate an effective number of degrees of freedom in the pooled data (n_{pooled}) using *Equation {3}*

$$n_{pooled} = rounddown \left\langle \frac{\left[\sigma_1^2/n_1+\sigma_2^2/n_2\right]^2}{\frac{\left[\sigma_1^2/n_1\right]^2}{n_1-1}+\frac{\left[\sigma_2^2/n_2\right]^2}{n_2-1}} \right\rangle \quad \{3\}$$

where the function rounddown<> rounds the result down to the integer value.

SUMMARY

The t-test provides a method for estimating a range that we can be confident that a population mean falls within, based on a limited sample size. This article described how that can be extended to compare two data sets to determine whether the populations have different mean values. The F-test provides a similar method for determining whether the variances of two populations are the same in order to justify whether to data can pooled together to increase the effective sample size and thereby increase the confidence of conclusions.

AUTHOR NOTES

1. Readers with a reasonable background in statistics may have noticed that I have been quite careless in my treatment of the mean and standard deviation of a population. A rigorous approach would be much more careful to differentiate between actual quantities relevant for the population (the true mean and standard deviation) compared to the estimated values based on the subset of the population that is sampled. The goal of these article is to provide reasonably simple tools for drawing statistical conclusions without diving too far into the statistical details. I hope that my attempts to minimize confusion that can be generated by additional notations and parameters do not actually increase confusion due to a lack of sufficient context.

2. I slightly modified the data set in order to better demonstrate the impact of pooling data sets. The highest thermal resistance value for the Vertical orientation (2.13, which is indicated with an asterisk) was actually measured to be 1.93 in testing. It is left to the interested reader to calculate how the use of the correct value affects the overall confidence level.

References
1. Ross Wilcoxon, “Statistics Corner – Normal Distribution”, Electronics Cooling Magazine, Summer 2020, pp. 10-11
2. Ross Wilcoxon, “Statistics Corner – Confidence Intervals”, Electronics Cooling Magazine, Spring 2021, pp. 10-12
3. J. S. Milton and Jesse Arnold, Introduction to Probability and Statistics: Principles and Applications for Engineering and the Computing Sciences, McGraw-Hill, 1986, pp. 297-308

[8] Alternatively, Excel will also return the correct positive value if we use 1- α/2: =t.inv(0.975,15) = 2.13144.

STATISTICS CORNER: REGRESSION ANALYSIS

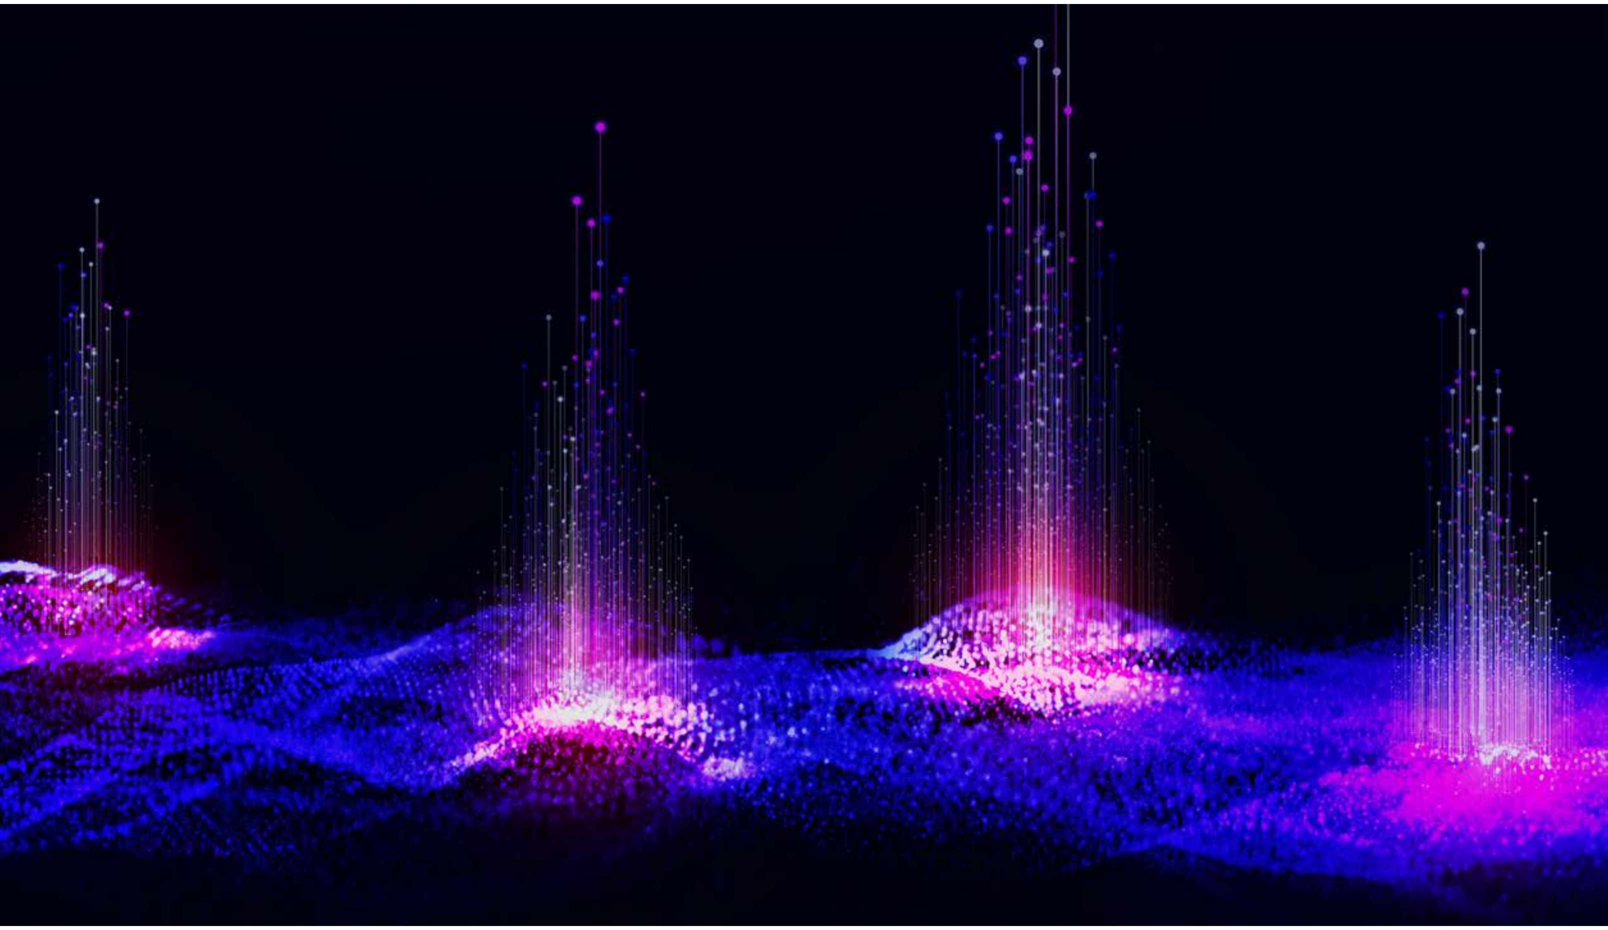

STATISTICS CORNER: REGRESSION ANALYSIS

Throughout their careers, engineers and scientists are all likely to encounter and utilize the results of regression analysis, which is "a set of processes for estimating the relationships between a dependent variable and one or more independent variables" [1]. In other words, regression analysis uses a set of data to estimate a relationship between the independent 'predictor(s)' and a 'response' or 'output' parameter. In its simplest form, a response, y, may be linearly related to a single predictor, x, in a relationship of y = mx + b. Regression analysis provides a method for estimating values of the constants m (the slope) and b (the intercept).

Regression analysis can be accomplished with different approaches that could include, at least theoretically, a piece of wood, drywall screws, rubber bands and a welding rod[9], as shown in *Figure (1)*. In this regression analysis, screws were put into x-y locations of a graph drawn on the wood and rubber bands between the welding rod and the screws hold the welding rod in an equilibrium position that allow the slope and the intercept to be determined.

Figure 1. Regression analysis done the hard way

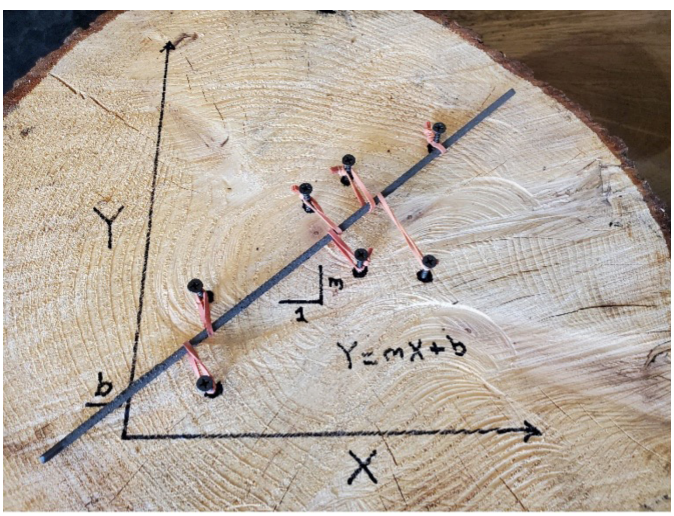

A slightly easier and certainly more accurate approach for conducting a regression analysis is the use of Least Squares. In this approach, rather than the location of the welding rod that leads to a balance in the forces generated by the rubber bands in *Figure (1)*, the 'welding rod' corresponds to the straight line that produces the smallest value of E, where E is the sum of squares of the distance in the y-direction between the line and each data point. The equations for calculating the coefficients for a least-squares estimate for linear regression with a single predictor, i.e., y = mx + b, are shown in *Equations {1}* and *{2}* [2]:

$$m = \frac{S_{xy}}{S_{xx}} = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sum(x_i - \bar{x})^2} \qquad \{1\}$$

$$b = \bar{y} - m\bar{x} \qquad \{2\}$$

where x_i and y_i are the x-y values for the ith data point and $\bar{x}$ and $\bar{y}$ are the mean values of the x and y data respectively. The coefficient of determination, R^2, is another important parameter in regression analysis. This term describes how well the regression analysis describes the data: an R^2 of 1 indicates a perfect fit while a value of 0 indicates that the regression analysis does not predict the response from the input data. R^2 is calculated using *Equation {3}*:

$$R^2 = \frac{S_{yy} - SSE}{S_{yy}} = \frac{\sum(y_i - \bar{y})^2 - \sum(y_i - b - mx_i)^2}{\sum(y_i - \bar{y})^2} \qquad \{3\}$$

where SSE is known as the sum of squares error.

Equations {1-3} are implemented in any software that does regression analysis. For example, several methods can be used in Microsoft Excel to determine regression coefficients. Methods that this author has used are summarized in *Figure (2)*. *Figure (3)* shows an example of an Excel regression analysis, using Option 1 as described in *Figure (2)*, for the x-y values that were used in the demonstration illustrated in *Figure (1)*.

For only one independent variable:
Option 1. Create an x-y chart of the data being analyzed, right click on the data the chart, select "Add Trendline…", check boxes for "Display Equation on Chart" and "Display R-squared on chart"

Option 2. Enter the functions"=slope(y-values, x-values)", "=intercept(y-values,x-values)", and/or "=rsq(y-values, x-values)", where 'x-values' and 'y-values' are cells that contain the x and y values of the data set being analyzed.

For one or more independent variables (multiple x's):
Option 3. Highlight a suitable range of cells, type in the function "=linest(y-values, x-values, true, true)", and instead of hitting 'Enter' hit Control-Shift-Enter, because this is an array formula. Relevant statistics are generated in the array (the correct size of the array depends on how many sets of x-values are selected) (note, the configuration of the output parameters does not correspond with the input configuration, so it is recommended that before using this function for the first time, they generate a dummy set of data with known coefficients so that they can know exactly where the important output values are in the generated array).

9 In other words, random stuff that I had laying around my house on a weekend.

Option 4. Add the Data Analysis Add-in, go to the 'Data' tab, select 'Data Analysis' to open a pick list of data analysis tools, select 'Regression' and define inputs in the dialogue box that is displayed.

Option 5: Guess coefficients for each independent variable and put them into a range of cells. Calculate the value of y using these coefficients for each set of x-values and sum the error for each data point, i.e., the square of the difference between measured y and calculated y. Then use the Excel Solver Add-in to minimize that sum by varying coefficients. Depending on how good the initial guessed coefficient values are and the nature of the modeled regression curve, this approach may converge to the correct values or may spiral off to 'infinity and beyond'.

Figure 2. Methods for Regression Analysis in Excel

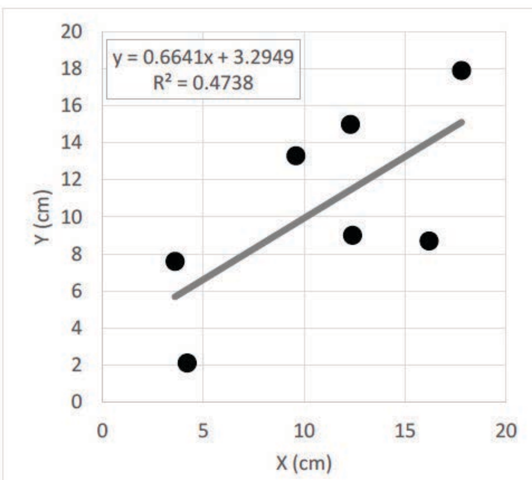

Figure 3. Excel-based regression analysis for same data as Figure 1

A previous column in this series described how probability distribution concepts could be used to a confidence interval for a limited set of data. When measurements are used to determine an average value, we can determine what range of values the actual average of the falls within a range to a given confidence level [3]. The confidence interval depends on the variance of the measurements (standard deviation) and the number of measurements made. The t-distribution was used in the calculation of the range.

In the same manner that we estimate a mean value within a confidence interval, confidence intervals also apply to the coefficients (slope and intercept) determined through regression analysis. These intervals are determined with *Equations {4}* and *{5}* [2]:

Confidence band on the slope: $m \pm t_{\alpha/2} \frac{S}{\sqrt{S_{xx}}}$ {4}

Confidence band on the intercept:

$$b \pm t_{\alpha/2} \frac{S\sqrt{\sum x^2}}{\sqrt{nS_{xx}}} \quad \{5\}$$

Where $t_{\alpha/2}$ is the t-distribution corresponding to the confidence level and degrees of freedom, n is the number of data points, $\sum x^2$ is the sum of all x values,

$$S_{xx} = \sum (x_i - \bar{x})^2, \text{ and } S = SSE/(n-2) = \frac{\sum (y_i - b - mx_i)^2}{n-2}.$$

Another confidence interval of interest is the value of y that is predicted by the regression analysis for any x-value. This confidence interval accounts for the combined effects of the confidence bands associated with the slope and intercept and is shown in Equation {6}.

Confidence band on y-value:

$$y = mx + b \pm t_{\alpha/2} S \sqrt{A + \frac{1}{n} + \frac{(x-\bar{x})^2}{S_{xx}}} \quad \{6\}$$

Where A = 0 if we are estimating the confidence band on the average y value for the population tested and A = 1 for an individual item.

Data from the heat sink assessment discussed in [4] will illustrate how these equations are used to determine confidence intervals of regression coefficients. A flat plate heat sink was tested in still air under a range of orientations relative to gravity. Results for ~20W dissipation values for a range of angles are shown in Figure 4, which includes regression analysis results with R^2 of ~86%. While this R^2 value is reasonable, the suitability of the fit is probably somewhat questionable: the values at the low and high range of measurements are above the fit while those in the middle are below. This is often an indication that the regression analysis may not be capturing the fundamental physics that influence the results.

Figure 4. Test data for natural convection heat sink at different orientations

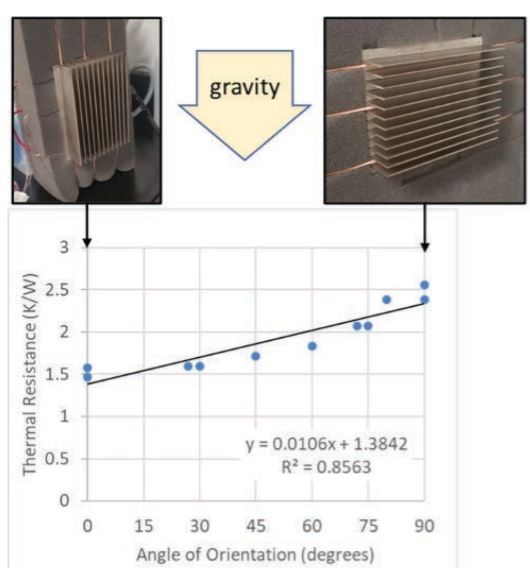

When assessing the results in terms of the physics that cause the heat sink thermal resistance to change with its orientation relative to gravity, it seems reasonable that the buoyant flow that drives natural convection will depend on the cosine of the orientation angle, rather than the angle itself. *Figure (5)* shows the resulting correlation between the thermal resistance as a function of the cosine of its angle relative to gravity. This appears to improve the fit substantially; the R^2 increases from 86% to 95%. Given this improvement in the fit, the subsequent analysis assumes that the cosine of the angle, rather than the angle itself, is the correct independent variable for regression analysis.

Figure 5. Natural convection heat sink resistance vs. cosine of orientation

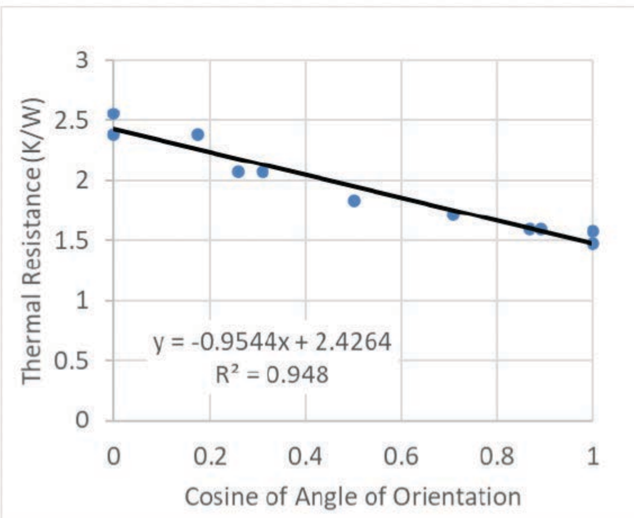

Table (1) shows the eleven data points used to generate the previous plots while the values of the parameters used in, or resulting from, the regression analysis are shown in *Table (2)* along with brief descriptions of how they are calculated.

Table 1. Measured data

θ (deg)	cos (θ)	R_{th} (K/W)
0	1.000	1.575
45	0.707	1.711
60	0.500	1.835
30	0.866	1.591
27	0.891	1.594
75	0.259	2.069
90	0.000	2.380
0	1.000	1.472
72	0.309	2.074
80	0.174	2.383
90	0.000	2.561

Table 2. Calculated parameters for the regression analysis confidence interval

Parameter	Value	Equation
n	11	number of data points
μ_x	0.519	average of all x values
μ_y	1.931	average of all y values
S_{xx}	1.527	sum of each value of $(x-\mu_x)^2$
S_{yy}	1.467	sum of each value of $(y-\mu_y)^2$
S_{xy}	-1.457	sum of each value of $(x-\mu_x)^*(y-\mu_y)$
Σx^2	4.487	sum of each value of x^2
m	-0.954	$= S_{xy} / S_{xx}$
b	2.426	$= \mu_y - m^* \mu_x$
SSE	0.0764	sum of each value of $(y - b - mx)^2$
R^2	94.8%	$= (S_{yy} - SSE)/ S_{yy}$
S	0.0921	= SSE/(n-2)

The confidence intervals for the regression coefficients depend on what confidence level is defined. For example, for a 95% confidence level, the t-statistic would be calculated for a probability of 0.975 (1-(1-0.95)/2) and 9 degrees of freedom (sample size of 11 minus 2) as 2.262. The confidence bands for the coefficients are then:

Slope confidence band:

$$\pm t_{\alpha/2}\frac{S}{\sqrt{S_{xx}}} = 2.262 * \frac{0.0921}{sqrt(1.527)} = 0.169$$

Intercept confidence band:

$$\pm t_{\alpha/2}\frac{S\sqrt{\Sigma x^2}}{\sqrt{nS_{xx}}} = 2.262 * \frac{0.0921 * sqrt(4.487)}{sqrt(11 * 1.527)} = 0.110$$

Since the nominal slope and intercept are -0.954 and 2.426, respectively, we can be 95% confident that the slope is between -1.123 and -0.786 (i.e., -0.954±0.169) and the intercept is between 2.317 and 2.536. Using *Equation {6}*, we can determine the confidence bands for the population and individual measurements, which are plotted in *Figure (6)*.

Figure 6. Confidence bands for regression analysis of heat sink

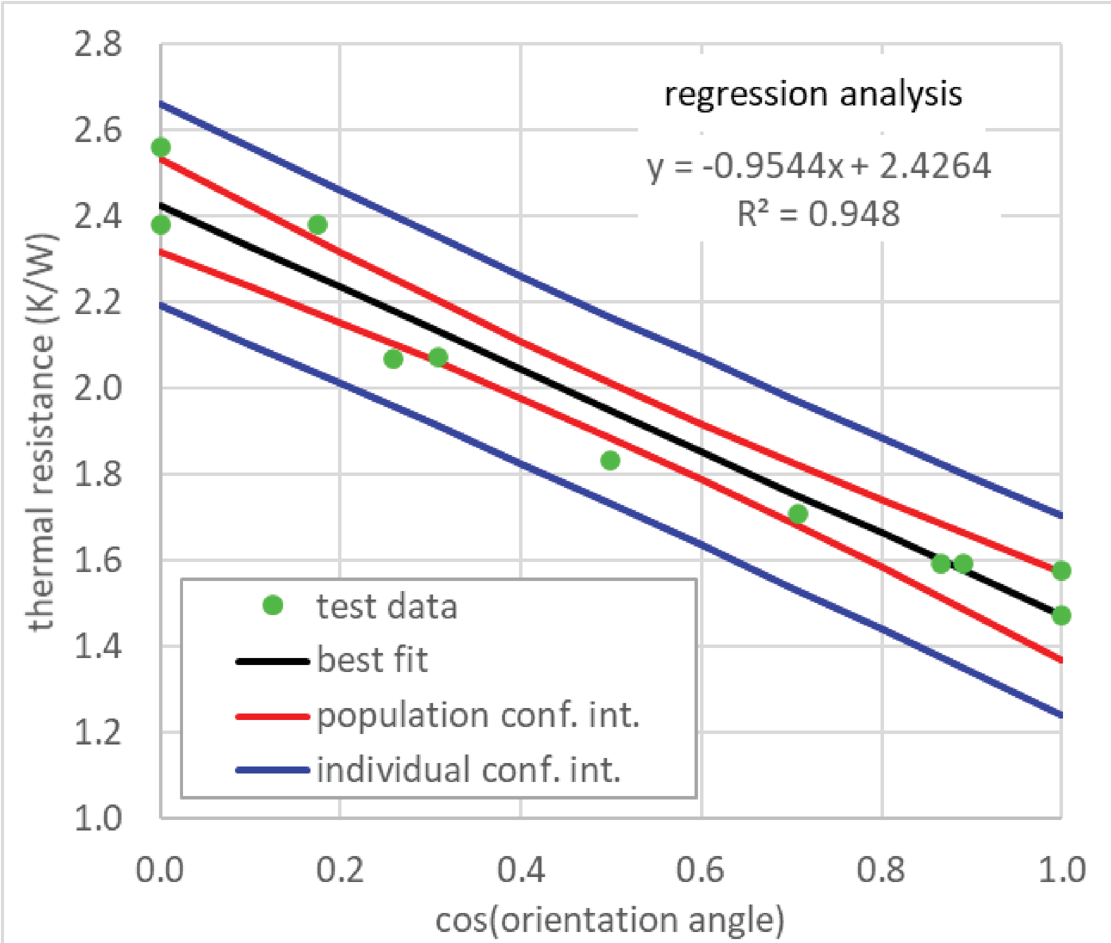

SUMMARY

Conducting a regression analysis can be a relatively straightforward process. Tools are widely available, or the basic equations can easily be implemented into a spreadsheet, to determine a curve fit between independent and dependent variables. One needs to keep in mind, however, that these tools will provide a curve fit, regardless of whether the correct variables have been input to them. As in this case, recognizing the physics of the situation led to a change in the independent variable so that a better fit was obtained. Also, this article described how to calculate confidence bands for the coefficients resulting from a regression analysis, since one must recognize that those values are merely estimates.

References

1. https://en.wikipedia.org/wiki/Regression_analysis (accessed 8/7/21)
2. J. S. Milton and Jesse Arnold, Introduction to Probability and Statistics: Principles and Applications for Engineering and the Computing Sciences, McGraw-Hill, 1986, pp. 297-308
3. Ross Wilcoxon, “Statistics Corner – Confidence Intervals”, Electronics Cooling Magazine, Spring 2021
4. Ross Wilcoxon, “Statistics Corner – Comparing Populations”, Electronics Cooling Magazine, Summer 2021

STATISTICS CORNER: WEIBULL DISTRIBUTION

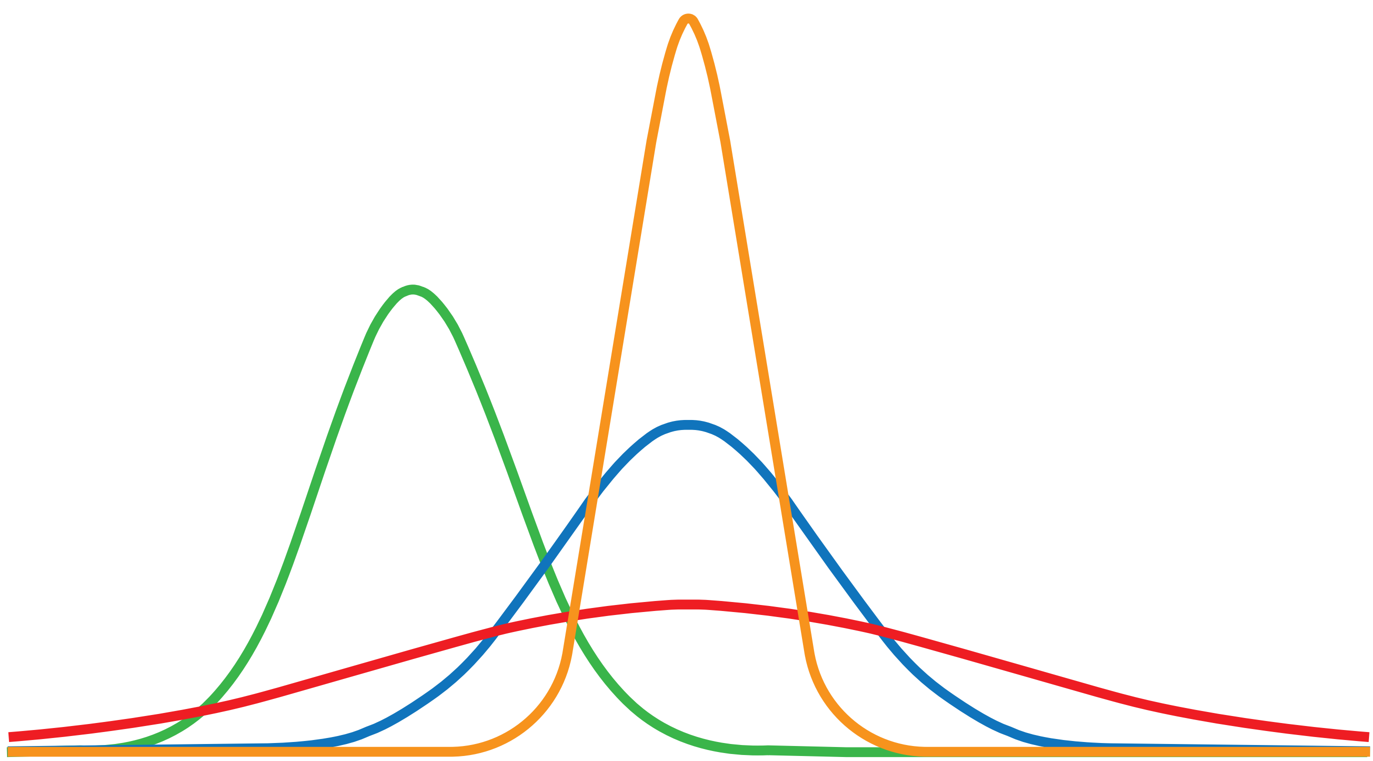

STATISTICS CORNER: WEIBULL DISTRIBUTION

A little over a pandemic ago, the first article in this series on statistical analysis mentioned that a fundamental aspect of statistics is that one assumes a mathematical model that describes the distribution of a data set and then uses that model to estimate the probability that a given value or set of values will occur [1]. This allows us, for example, to estimate whether two sets of data are from the same underlying population or if they are statistically different. The statistical analyses discussed thus far have primarily assumed that a population has a normal distribution. However, there certainly are other distributions that can, and should, be used for different types of data. *Table (1)* lists a few common statistical distributions with brief descriptions and examples of how they are applied.

Table 1. Common statistical distributions (adapted from Table 2.2 of Ref. [2])

Distribution	Field of Application	Example Applications
Normal	Physical properties	Material properties, component values (such as resistance or capacitance), dimensions, etc.
Log Normal	Life phenomena of components; asymmetric distributions with large differences in observed values	Performance measurements for a population of systems such as cars, light bulbs, etc.
Weibull	Same as log normal and cases such as wear out, where failure rate changes	Life of components such as bearings, gears, electronic components, corroding materials, strength distribution of materials, survival probability in thermal cycling and shock, etc.
Exponential	System (multi-component) failures, including random failures	Non-wear out failures, life to failure of machines with many possible component failure modes
Binomial	Define correct sample size for assessing a large population with a known defect probability	Inspection for process / quality control. Sampled population large enough that parameters don't change when sampled
Hypergeometric	Same as Binomial, but for smaller populations	Probability of finding X defects out of a finite population of Y parts
Poisson	Situations where the number of occurrences can be measured, but not the number of times they don't occur	Long lines at the store, accidents, breakdowns, etc.

Lognormal and Weibull distributions are often applied to analyze reliability data for situations such as the wear out of solder joints that have been subjected to multiple thermal cycles. As a reminder, the formula for the probability of a given value of x in a normal distribution is shown in *Equation {1a}*. The mean, μ, and standard deviation, σ, for a population x can be easily calculated. Therefore, the probability distribution, f(x), of the normal distribution can be calculated directly with *Equation {1a}*. The cumulative distribution, F(x), which is the area under the probability distribution curve to the left of a value x, is found by integrating the probability distribution, as shown in *Equation {1b}*.

$$f_{normal}(x) = \frac{1}{\sigma\sqrt{2\pi}} exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] \quad \{1a\}$$

$$F(x) = \int_{-\infty}^{x} f(u)du \quad \{1b\}$$

The lognormal distribution is also calculated using *Equations {1}*, except that the values of x that are used to calculate μ, σ, and f are all the natural logs of the population values.

The Weibull distribution is defined such that its cumulative distribution[10] is calculated as shown in *Equation {2}*.

$$F_{Weibull}(x) = 1 - exp\left[-\left(\frac{x}{\theta}\right)^{\beta}\right] \quad \{2\}$$

As with the probability distribution of the normal distribution, the cumulative Weibull distribution can be directly calculated for any value of x once the two terms that characterize the distribution are known. Instead of the mean and standard deviation used in the normal distribution, the Weibull distribution uses the characteristic life (also known as the scale parameter), θ, and the shape parameter, β. Physically, these two terms are similar to their normal distribution counterparts: the characteristic life is analogous to the mean, but instead of indicating the 50% failure point (for failure data), θ corresponds to the point at which 63.2% (1-1/*e*) of the population would fail. Because of this, the characteristic life is often referred to as N63. The shape parameter, β, also known as the Weibull slope, is analogous to the inverse of the standard deviation. The larger the shape factor, the smaller the spread in the data.

[10] Interested readers can find the Weibull probability distribution function from any number of sources. Since it is somewhat less intuitive than the cumulative distribution and not really relevant to the point being made in this article, I'm not including it here.

While the Weibull coefficients (θ and β) are physically analogous to the normal distribution coefficients (μ and σ), there are no formulas for directly calculating the Weibull terms as there are with the normal distribution terms. The Weibull coefficients can be calculated using a different approach that is ultimately at least part of the reason that the Weibull distribution has been widely used for analyzing reliability data.

The approach for determining the Weibull coefficients begins by first rearranging *Equation {2}*, as shown below.

$$\frac{1}{1-F(x)} = exp\left[\left(\frac{x}{\theta}\right)^{\beta}\right]$$

Taking the natural log of both sides of the equation twice gives us:

$$ln\left(ln\frac{1}{[1-F(x)]}\right) = \beta * \ln\left[\frac{x}{\theta}\right]$$

This can be written as:

$$ln\left(ln\frac{1}{[1-F(x)]}\right) = \beta \ln(x) - \beta \ln(\theta)$$

This produces a linear equation of the form *Y=mX+C*, where the terms for Y, X, and C are:

$$Y = ln\left(ln\frac{1}{[1-F(x)]}\right), X = \ln(x), C = -m\ln(\theta)$$

Once the regression analysis has determined the linear coefficients m and C, they can be used to determine the Weibull coefficients. The shape factor is equal to the slope of the regression analysis (β = m) and the characteristic life is determined by rearranging the above equation for C as θ = exp(-C/m).

One easy way to generate a set of fatigue data is to count the number of times paper clips can be bent from 0° to 90°, as shown in *Figure (1)*, before breaking. Data for the number of bends needed to break fourteen paper clips are shown in the first column of *Table (2)*. As discussed later, this table also includes processed data used to calculate the Weibull coefficients. As shown in the bottom of the table, the measurements had an average of 17.5 bends with a standard deviation of 6.048.

Figure 1. Paper clip fatigue test configurations

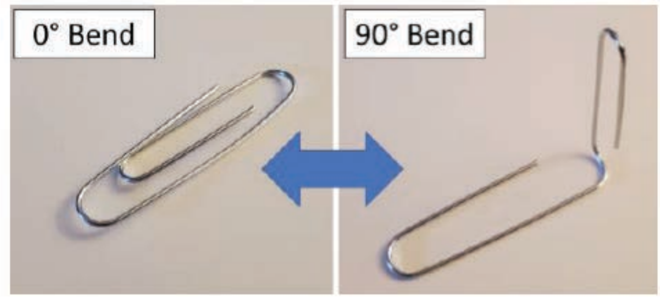

One possible way to define the cumulative failure distribution, F(x), would be to divide the rank of the failure by the total number of samples. For example, if 10 components are tested, the first failure would have F(x) = 0.1, the second one would be 0.2, etc. This approach, however 'pushes' F to higher values; for example, only the first failure would be categorized to the lowest 10% while the last two failures would be categorized to the highest 10%. A better approach for calculating the values of F(x) is to use the median ranks, which are typically calculated with *Equation {3}* [2]

$$F(x) = \frac{i-0.3}{n+0.4} \qquad \{3\}$$

In this equation, i is the rank of the failure and n is the total number of samples. For example, if 10 samples are tested, the first failure would have F(x) = (1-0.3)/(10+0.4) = 0.0673, the second failure would have F(x) = (2-0.3)/10.4 = 0.163, etc.

Since the calculation of F(x) requires that the order of failures be determined (first, second, etc.), the Excel @ rank() function can be used to determine the order of failures[11]. Because of the way that this function deals with ties (each tie has the same rank), a small amount of 'noise' was added to the data to prevent any ties when calculating the ranks used for F(x). This noise was generated using the random function, which generates a random number between 0 and 1, by adding the term 'rand()/100' to the measured data, x, to create x*. The x* value was only used to determine the rank, which was then implemented in *Equation {3}* to calculate F(x).

[11] The syntax to use this function to determine the rank of a value of X within a set of data Y, use =rank(X,Y,1) where 1 specifies ascending order so that the first failure is 1, the second failure 2, etc.

Table 2. Paper clip failures: raw and processed data

Data	Noise Added		Linearized Terms	
x	x*	F(x*)	$ln\left[ln\left(\frac{1}{1-F}\right)\right]$	ln(x)
26	26.006	0.9514	1.107	3.258
14	14.007	0.3958	-0.685	2.639
15	15.005	0.4653	-0.468	2.708
22	22.009	0.6736	0.113	3.091
12	12.003	0.2569	-1.214	2.485
19	19.000	0.6042	-0.076	2.944
17	17.000	0.5347	-0.268	2.833
13	13.008	0.3264	-0.929	2.565
10	10.003	0.0486	-2.999	2.303
11	11.002	0.1181	-2.074	2.398
11	11.006	0.1875	-1.572	2.398
26	26.002	0.8819	0.759	3.258
24	24.000	0.7431	0.307	3.178
25	25.004	0.8125	3.219	0.515
17.5	**<= average =>**		**2.806**	
6.048	**<= stan, dev. =>**		**0.351**	

Once values of x and F(x) were determined, the linearized data of Y = ln(ln(1/(1-F) and X =ln(x) were calculated. Regression coefficients were determined using Excel functions m=slope(Y,X) and C=intercept(Y,X), as discussed in [3]. *Table (3)* shows the values determined for m and C as well as the Weibull coefficients determined from them.

Table 3. Regression and resulting Weibull coefficients

Regression slope, m	3.088
Regression intercept, C	-9.198
Shape Factor, β	**3.088**
Characteristic Life, θ	**19.66**

Notes
- β= regression slope, m
- θ= exp(-C/m)

Figure (2) plots the paper clip data along with corresponding fits of the data using many of the distributions listed in Table 1. The Excel equations used to generate these curves are shown in *Table (4)*.

Figure 2. Paper clip data fit to different statistical distributions

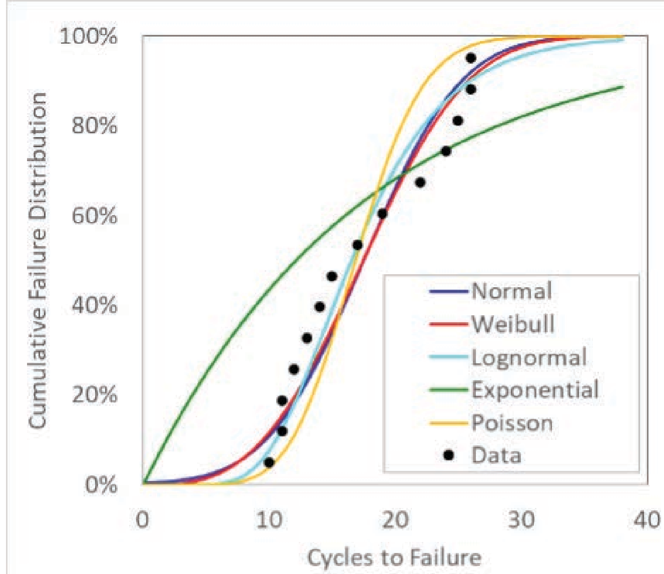

Table 4. Excel functions used to calculate statistical distributions in Figure 2.

Distribution	Excel Equation Used
Normal	=NORM.DIST(N, 17.5, 6.048, TRUE)
Weibull	=1-EXP(-((N/19.66)^3.088))
Lognormal	=LOGNORM.DIST(N, 2.806, 0.351, TRUE)
Exponential	=EXPON.DIST(N, 1/17.5, TRUE)
Poisson	=POISSON.DIST(N, 17.5, TRUE)

Notes
- Numerical values used in equations are shown in bold in Table 2 & Table 3.
- N is the number of cycles for which the CDF is found (for this plot, values of N = 2, 4, ... 38 were used
- For Exponential distribution, a value of 24 gave a somewhat better fit than the population average

An interesting observation from *Figure (2)* is that, except for the Exponential distribution, all of the distributions provided a reasonably good fit to the measured data. The primary differences between the different distributions are in the tails and the knees (where F(x) ~0-10% and ~90-100%). One may then ask then, if the normal distribution, which is used almost everywhere, produces similar results to the Weibull distribution, why is Weibull primarily used for reliability analysis?

One reason is that Weibull distributions, as well as the other non-normal distributions shown, have a physically accurate limit. In those distributions, by definition, F(0) is equal to zero[12].

[12] The 3-parameter Weibull distribution includes an additional value, x0, which is offset value that forces the function to have a value of F(x_0) = 0. The term x0, which corresponds to a failure-free life of the component, is physically useful but mathematically somewhat problematic since it prevents the Weibull distribution from being directly linearized.

In contrast, a normal distribution of failure data includes some portion of the population supposedly failing at negative cycles. This is typically a small value - for the data in *Table (2)*, the normal distribution calculates that 0.318% of the paper clips would fail before the first bend.

The much more important reason why the Weibull distribution has been used for assessing certain types of data, such as that from reliability testing, is that it does not rely on complete knowledge of the entire test population. If an entire set of samples is tested until all fail, the normal distribution – or better yet the lognormal distribution, which passes through F(0) = 0, will likely describe the data as well as the Weibull. However, due to constraints in time, budget, testing availability, etc., testing may stop before all samples have failed. If one characterizes the population only using that fraction of parts that have failed, the calculated average life will be much smaller since those components with longer lives would not be included in the calculations.

For example, if the paper clip testing had been stopped after a maximum of 15 bends, only half of the samples would have registered a failure. The average life of those parts would be 12.3 cycles, rather than the 17.5 that was found with the entire population.

Figure (3) plots the same distributions as were shown in *Figure (2)*, but with truncated data that only included those failures that occurred within 15 cycles. It is not surprising that those distributions that depend on the mean values, which were lower in the truncated data, are pushed to the left and provide a less accurate description of the population of the entire data set. Because the Weibull distribution uses regression analysis of the failures as part of the entire test population, it continues to agree with the actual data.

Figure 3. Failure distributions calculated only using failures up to 15 cycles

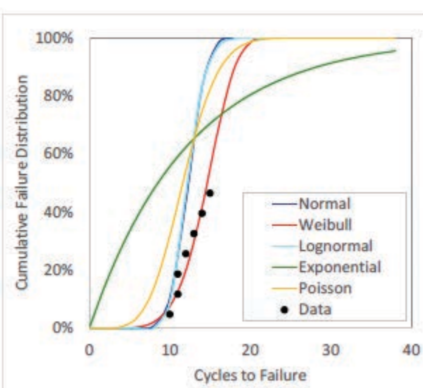

Readers who are familiar with Weibull plots will notice that *Figure (2)* and *Figure (3)* use an unusual format for showing Weibull fits. Traditionally, those plots use linearized axes so that the data are shown relative to a straight line. *Figure (4)* shows the data and the various distributions as plotted in this more conventional format.

Figure 4. Paper clip data from Figure 2, plotted on linearized axes

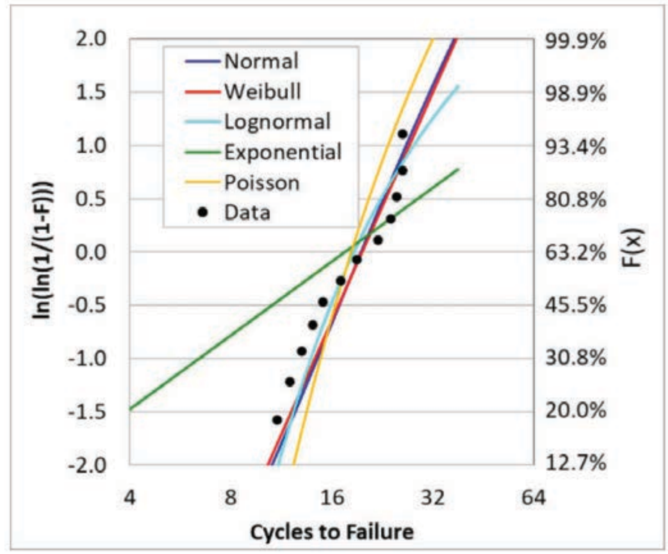

CONCLUSIONS

This article briefly discussed statistical distributions other than the normal distribution and specifically focused on the Weibull distribution. The analysis presented here showed that most, but not all, of the distributions provided a reasonably good fit when data for the entire population were available. However, the Weibull distribution was also able to directly determine an accurate representation of reliability data, even when testing had been stopped before all samples were tested.

Different distribution models are used for different types of analyses. It is important that, when using these different distributions, one has some understanding of why a given distribution is typically used and to recognize its limits.

References

1. Ross Wilcoxon, "Statistics Corner – Probability", Electronics Cooling Magazine, Spring 2020
2. J. S. Milton and Jesse Arnold, Introduction to Probability and Statistics: Principles and Applications for Engineering and the Computing Sciences, McGraw-Hill, 1986, pp. 297-308
3. Ross Wilcoxon, "Statistics Corner – Regression Analysis", Electronics Cooling Magazine, Fall 2021


Break the same old pattern.
Problem First. Product Last.
Strategy | Content | Data | Technology
LECTRIX®
Digital Marketing for the B2B Electronics Industry
1.484.688.0300 | info@lectrixgroup.com
www.lectrixgroup.com

STATISTICS CORNER: MODIFYING SAMPLE SIZE

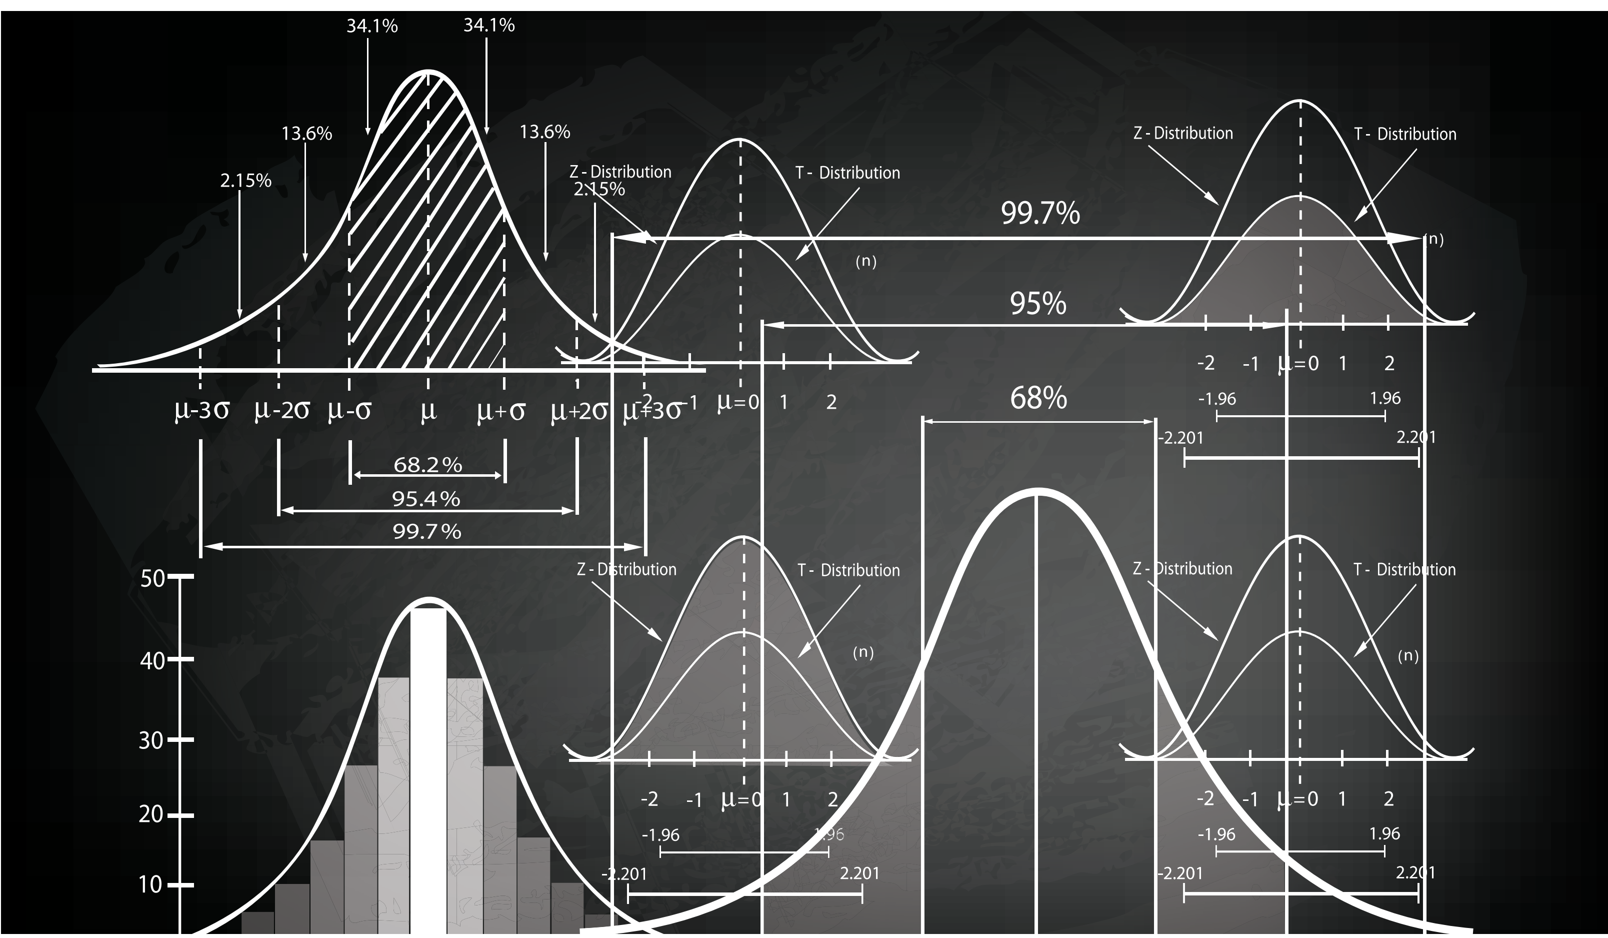

STATISTICS CORNER: MODIFYING SAMPLE SIZE

ABSTRACT

Reliability verification often requires that a specific number of components be tested to a predetermined level of testing to demonstrate that none of the samples fail. This article describes a statistical approach for justifying the use of fewer samples by subjecting them to a more severe level of testing.

BACKGROUND

Reliability verification often includes accelerated testing in which a set of components must survive a prescribed set of severe environmental conditions. For example, components used in many avionics systems are commonly required to demonstrate solder joint integrity for 500 thermal cycles of -55 to +125°C [1].

One question that may arise in reliability testing is how many components must pass the test to demonstrate reliability. The number of components needed in a test is related to both the reliability that is to be demonstrated and the confidence level required for the results. Bayes (Success Run Theorem) formula[13] provides a relationship between these three parameters [2] and can be written as shown in *Equation {1}*.

$$R = (1 - C)^{(1/(n+1)} \qquad \{1\}$$

In this equation, R is the reliability, C is the confidence level and n = number of samples that are tested. Both the reliability and the confidence level have values between 0 and 1. *Figure (1)* plots the confidence level resulting from different samples sizes for select values of reliability, as a function of the number of samples.

Equation {1} can be rearranged to show the number of samples required for a given reliability and confidence level, as shown in *Equation {2}*.

$$n = \frac{ln(1-C)}{ln(R)} - 1 \qquad \{2\}$$

Figure 1. Reliability and confidence level for different sample sizes

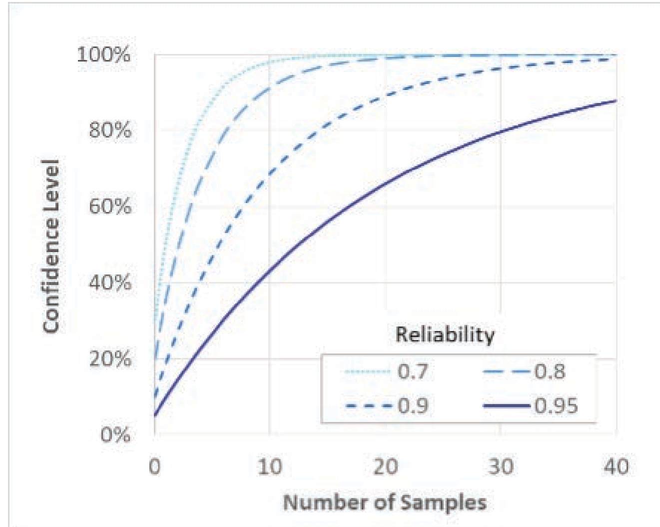

To demonstrate, for example, a reliability of 90% with a confidence level of 80%, a test would need to include at least 15 components[14].

ANALYSIS

Established test procedures often define how many components must be included in a given test, presumably based on an assessment of the reliability and confidence level required. In some cases, there may be a desire to test a different number of components than what has been prescribed for a test. This may occur if the components are extremely expensive or if the evaluation approach, which determines whether a component has failed, is very expensive or time consuming.

Reducing the number of samples in a reliability test reduces the confidence level that a given test demonstrates a particular reliability value. However, this can be offset by testing components more severely (to demonstrate a higher reliability) and thereby justify the use of a smaller number of test samples. This article describes a statistical analysis that provides a basis for adjusting the severity of testing to reduce the number of test components required to demonstrate the same reliability and confidence level as a larger population. This analysis has been described by many others, such as in [2].

The analysis is based on the assumption that the failure distribution of the test articles follows a Weibull distribution. As described previously, the Weibull distribution provides a relationship between the number of test cycles and the population's failure rate. The Weibull distribution

[13] This equation is based on setting the cumulative distribution function for a binomial distribution to be equal to 1 minus the confidence level.

[14] *Equation {2}* finds the number of samples to be n = 14.27, which is rounded up to 15.

is often used to analyze failure data and can be defined as shown in *Equation {3}*.

$$F(N) = 1 - e^{-\left(\frac{N}{\theta}\right)^{\beta}} \qquad \{3\}$$

For this discussion, it is assumed that component testing is evaluated using thermal cycling, but the same approach can be used for any type of accelerated testing. In *Equation {3}*, F(N) is the cumulative failure distribution of the population and N is the number of thermal cycles to which a population has been subjected. The failure characteristics of the population can then be described with the parameters θ, which is known as the characteristic life, and the shape factor β. *Figure (2)* shows examples of Weibull distributions for different values of shape factors and characteristic lives.

Figure 2. Weibull distributions

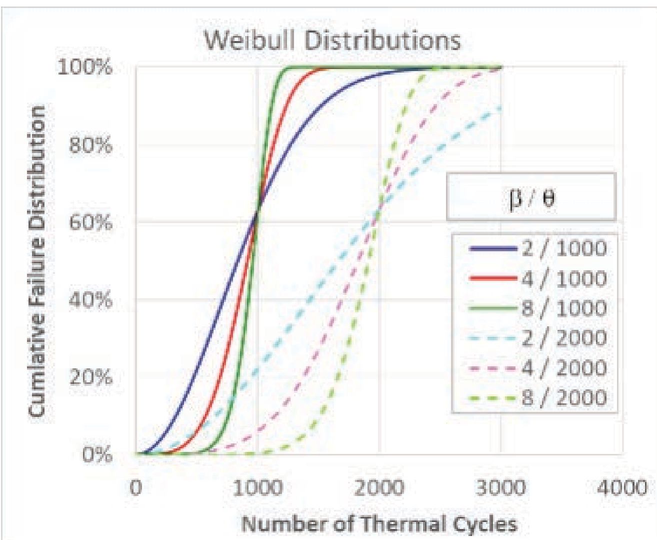

Regardless of the shape factor, all distributions with the same characteristic life cross through that value at 63.2%, hence the characteristic life also being known as N63. The cumulative failure distribution begins at zero and as the number of cycles increases asymptotically approaches 100%. The transition from F = 0 to F = 1 is sharper when the shape factor is larger. In terms of the more familiar Normal distribution, the characteristic life is analogous to the median of the population while the shape factor is analogous to the inverse of the standard deviation.

The failure distribution is related to the reliability by:

$$R = 1 - F. \qquad \{4\}$$

In other words, if 5% of a population has failed it has a reliability of 95%.

Equations {1}, {3}, and *{4}* can be combined to determine the number of thermal cycles required to demonstrate the same confidence level for a given reliability level for the different sample sizes. Combining *Equations {1}* and *{3}* relates the Weibull coefficients to the reliability:

$$R = 1 - F, \text{ so } R = e^{-\left(\frac{N}{\theta}\right)^{\beta}} \qquad \{5\}$$

Combining this with *Equation (4)* produces a relationship between the confidence level and the Weibull coefficients, leading to *Equation (6)* after taking the natural log of both sides.

$$ln(1 - C) = -(n + 1)\left(\frac{N}{\theta}\right)^{\beta} \qquad \{6\}$$

If components from the same population with a given characteristic life and shape factor are used in two tests with different sample sizes (n), *Equation {6}* can be written as:

$$-\theta^{\beta} ln(1 - C) = constant = (n_1 + 1)N_1^{\beta} = (n_2 + 1)N_2^{\beta} \qquad \{7\}$$

In this equation, n_1 and n_2 are the sample sizes in tests 1 and 2 respectively while N_1 and N_2 are the numbers of thermal cycles used in tests 1 and 2 respectively. The last two terms in *Equation {7}* can be rearranged to produce:

$$N_2 = N_1 * \left(\frac{n_1+1}{n_2+1}\right)^{1/\beta} \qquad \{8\}$$

Equation {8} provides a method for calculating how many additional thermal cycles must be conducted for a test population with a reduced size. For example, consider a test that mandates that a population of 32 samples must survive for 500 thermal cycles. However, only 20 samples are available for testing. Assuming that the population will exhibit a failure distribution with a Weibull shape factor of 5, the number of thermal cycles required would be:

$$N_2 = N_1 * \left(\frac{n_1 + 1}{n_2 + 1}\right)^{1/\beta}$$

$$= 500 * \left(\frac{32 + 1}{20 + 1}\right)^{1/5}$$

$$= 500 * \left(\frac{33}{21}\right)^{0.2}$$

$$= 500 * 1.095$$

$$= 547.3 = 548\ cycles$$

The required number of cycles is rounded up to the nearest whole number.

The critical assumption in determining an equivalent number of thermal cycles for a test is the value used for the Weibull shape factor. The shape factor depends on the type of component being tested, the type of test, and the consistency of the manufacturing build. One study, for example, has found that thermal cycle testing of surface mount components on assemblies built in a production environment had shape factors in the range of 3-6 [3]. This range seems reasonable based on this author's experience, but it can vary depending on the consistency of the manufacturing processes as well as the type of component.

For reference, *Figure (3)* plots equation {8} for a number of values of shape factors and assuming a standard test population size of 32 components. For example, consider a standard test that requires that 32 components survive 100 cycles. If the component is assumed to have a shape factor of 3, and is tested with only 10 samples, the test would need be conducted for ~1.45x (145) cycles to demonstrate the same confidence level.

Figure 3. Equivalent test cycles for reduced sample size

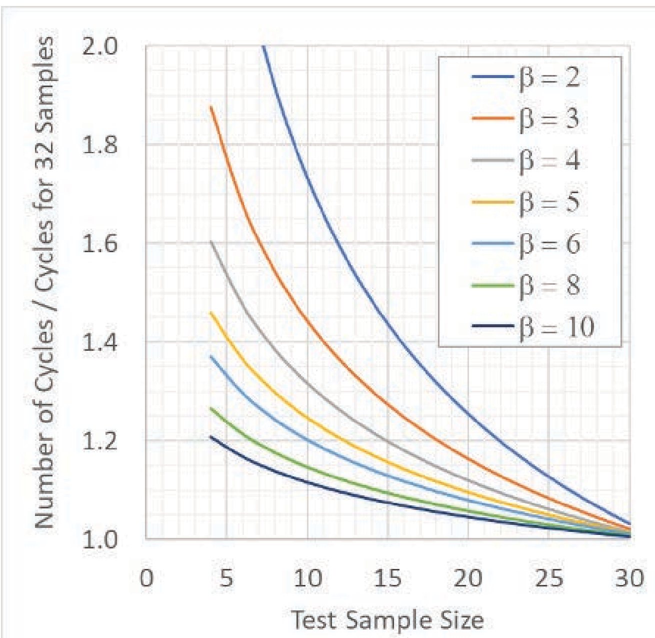

One should be cautious in applying this approach to justify substantially reducing the number of samples by testing more severely. For example, with this approach one could theoretically reduce the sample size for a test population with shape factor of 3 from 32 to 3 by doubling the test duration. This is dangerous due to the risk of initiating a new failure mechanism that would not otherwise appear in a shorter accelerated test. In addition, the uncertainty of the shape factor for a given population is magnified when used with larger ratio of test cycles as shown in *Equation (8)*. As a rule of thumb, one should generally not apply *Equation (8)* to justify a smaller sample size that requires increasing the test duration by more than ~40%.

SUMMARY

This article described an approach for determining how many additional thermal cycles are required to meet a given reliability level if testing is conducted with fewer than the standard number of test samples. The approach assumes that the failure distribution of the components follows a Weibull distribution of cumulative failure as a function of test cycles. It also assumes that the minimum expected Weibull shape factor for the component population s can be estimated with a reasonable degree of confidence.

References

1. Hillman, et al., "The Impact of Reflowing a Pbfree Solder Alloy using a Tin/Lead Solder Alloy Reflow Profile on Solder Joint Integrity", Int. Conf. on Lead-free Soldering, CMAP 2005, Toronto
2. Lipson and Sheth, "Statistical Design and Analysis of Engineering Experiments", McGraw-Hill, 1973
3. Liu and Lewis, "A Comparison of Statistical Distribution Functions to Predict BGA Attach Reliability", IEEE Transactions on Components and Packaging Technologies, Vol. 31, Issue. 3, 2008

Index of ADVERTISERS

Alpha Novatech, Inc.
473 Sapena Ct. #12,
Santa Clara, CA 95054

t: 408.567.8082
e: sales@alphanovatech.com
w: www.alphanovatech.com
page: 2

Electronics Cooling
1000 Germantown Pike,
Plymouth Meeting, PA 19462

t: 484.688.0300
e: info@electronics-cooling.com
w: www.electronics-cooling.com
page: 21

Leader Tech Inc.
12420 Race Track Rd
Tampa, Florida 33626

t: 866.832.4364
e: sales@leadertechinc.com
w: www.leadertechinc.com
page: 13

LECTRIX
1000 Germantown Pike,
Plymouth Meeting, PA 19462

t: 484.688.0300
e: info@lectrixgroup.com
w: www.lectrixgroup.com
page: 36

SIEMENS Digital Industries Software
8005 SW Boeckman Road
Wilsonville, OR 97070

t: 800.592.2210
e: www.plm.automation.siemens.com/global/en/contact-us.html
w: https://www.plm.automation.siemens.com/global/en/
page: 8

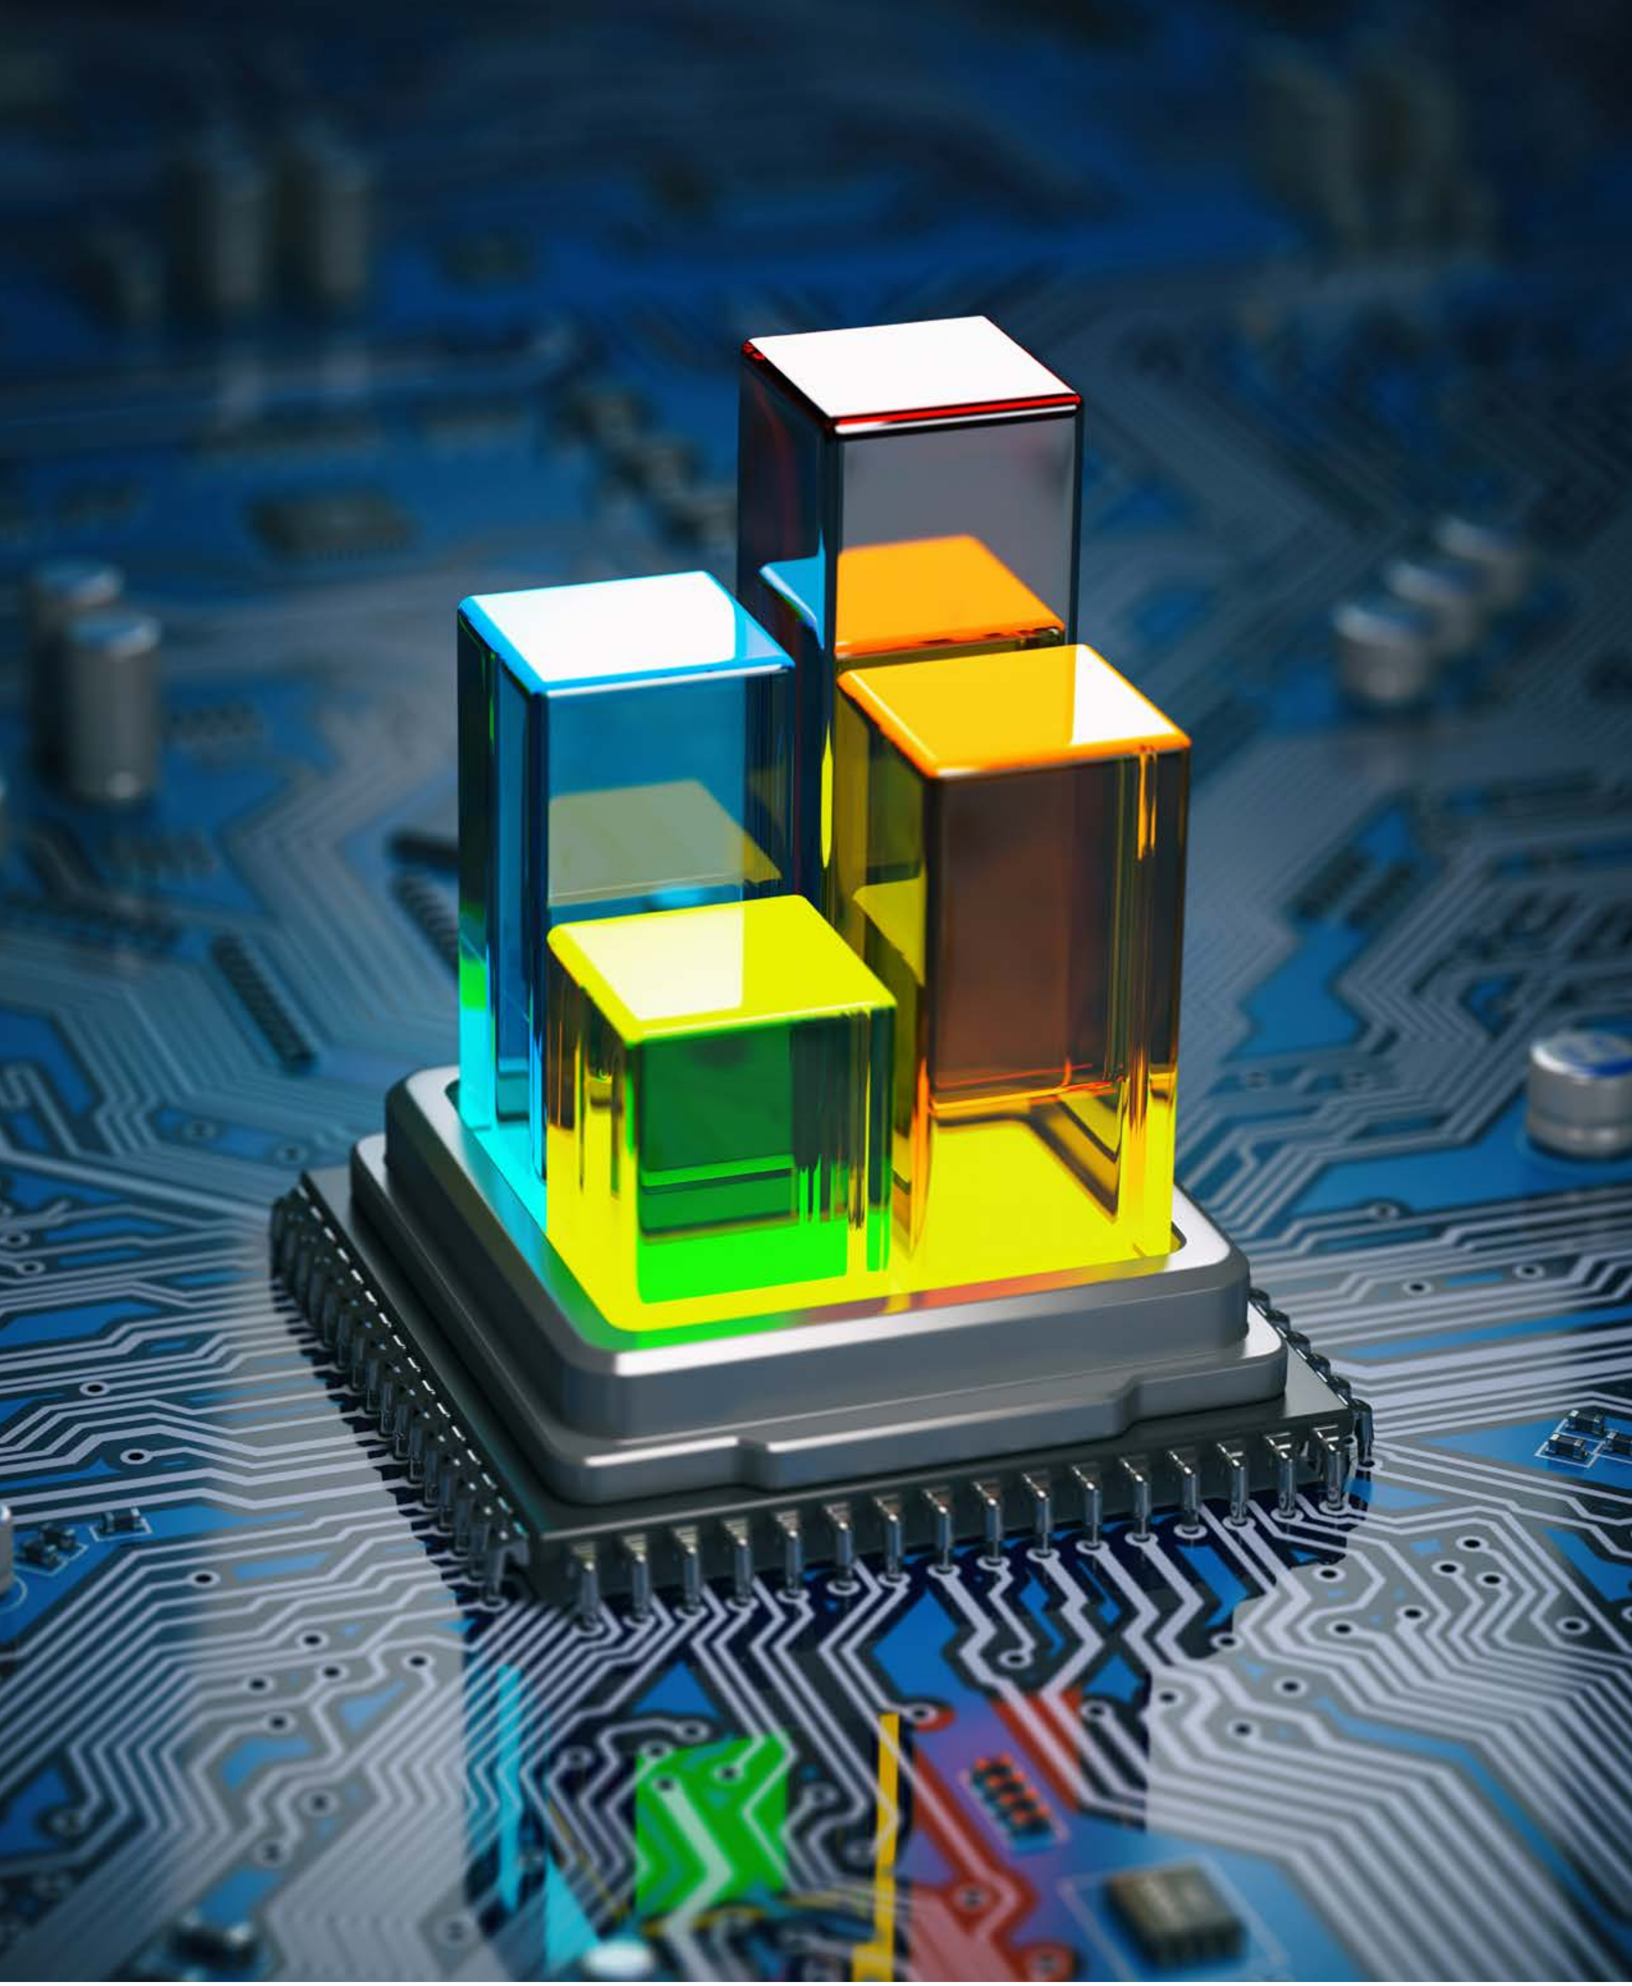

www.electronics-cooling.com