The Exponential Distribution: the Good, the Bad and the Ugly. A Practical Guide to its Implementation.

Kenneth E. Murphy; ARINC, Albuquerque
Charles M. Carter; ARINC, Albuquerque
Steven O. Brown; ARINC, Albuquerque

Key Words: Misuse, Exponential distribution, Redundant systems, Simulation, Memory-less property

SUMMARY AND CONCLUSIONS

Our paper discusses the widespread applicability and mathematical tractability of the exponential distribution, as well as the profound implications of Drenick's Theorem. We demonstrate the risks of the exponential assumption when used to represent redundant subsystems and when used within simulations. We caution our readers about the use of the exponential distribution to represent component repairs and the failure of many to appreciate the implications of the exponential distribution's memory-less property. In summary, this paper provides insight into the good aspects of using the exponential distribution but more importantly, into the ubiquitous misuse of the most commonly implemented reliability distribution.

1.0 INTRODUCTION

In the world of rapidly changing technology, we often find ourselves in situations where we need to quickly develop meaningful estimates of a system's reliability, availability, and maintainability (RAM) characteristics. The pressure from competitive world markets often forces reliability engineers to expedite their analyses by assuming a system's components fail and repair exponentially. Because this approach is widespread within the reliability community, we surmise that few stop to consider the appropriateness of this assumption. The exponential assumption is mathematically tractable, but with easy to use reliability simulators running on giga-hertz processors, it begs the question whether it should be used so blindly. Furthermore, it does not take much in the way of system complexity before the traditional exponential redundancy equations begin to become less than ideal RAM estimators.

This paper provides a practical look at some of the Good, the Bad and the Ugly aspects of performing analysis using the exponential distribution. We used the generic reliability simulator Raptor(Ref.1) to explore the limitations of the exponential assumption for several scenarios, thus bringing into question the need for implementing such assumptions. Examples of the Good aspects of the exponential distribution include: its widespread applicability, its mathematical tractability, and the tendency for "complex" repairable systems to tend towards being well represented by the exponential distribution. Examples of the Bad and the Ugly attributes of the exponential distribution include: its use in representing redundant systems, its use in representing repair data, and the lack of a thorough understanding of the implications of the exponential distribution's memory-less property.

Acronyms

MTTF	Mean Time To Failure
RAM	Reliability, Availability and Maintainability

2.0 THE GOOD

2.1 *Widespread Applicability*

The exponential distribution (also known as the negative exponential distribution) has a density function represented by Equation 1, and is probably more recognizable by its reliability function shown in Equation 2.

$$f(t) = \lambda e^{-\lambda t} \tag{1}$$

$$R(t) = e^{-\lambda t} \tag{2}$$

The Greek letter lambda represents a system's failure rate and the variable t represents the particular time of interest. The exponential distribution has a constant hazard rate that lends itself to representing the frequency of failures of many electronic components that have survived their infant mortality phase. Given the technology trend of the late twentieth-century towards a heavy reliance on electronic components, even in traditional mechanical systems such as cars, it is not surprising that the exponential distribution represents a significant number of systems' failure frequencies.

Kececioglu(Ref. 2) implies that the use of the exponential distribution to model electron tubes, resistors and capacitors is appropriate. In addition, Abernathy(Ref. 3) cites its use to represent failures due to lightning strikes on transformers, foreign object damage in aircraft engines and woodpecker attacks on power poles. Furthermore, our experience in testing a wide array of mature systems such as radars, aircraft and spacecraft electronics, satellite constellations, communications equipment and computer networks has indicated (at least empirically) that the exponential distribution fits the failure of various components quite well. This widespread applicability of the exponential distribution for numerous diverse systems makes it arguably the most versatile reliability distribution.

2.2 *Mathematical Tractability*

Another good aspects of the exponential distribution are the unique mathematical properties it possesses. First, the failure distribution is completely described by a single parameter (known or estimated). This parameter is the Mean Time To Failure (MTTF) and is often denoted with the Greek letter θ. Humans have a relatively easy time estimating the central tendency of data (i.e., its mean value) based on experience or intuition. Few humans, however, are able to adequately estimate the degree of spread of the very same data. Most reliability distributions additionally require estimations for their variance, shape or scale, parameters that indicate to some extent the degree of spread of the data. The Pearson VI distribution, for example, requires the knowledge of two different shape parameters and a scale parameter before it can be used for RAM analysis. Simply put, the exponential distribution is relatively easy to define with a minimum amount of estimation.

A second mathematical uniqueness associated with the exponential distribution is the fact that the failure rate is completely defined by knowing the mean life, and vice versa. The relationship between the two is shown in Equation 3.

$$\lambda = \frac{1}{\theta} \tag{3}$$

For the exponential distribution, the failure rate is constant; thus, Equation 3 is always true. For other distributions, the failure rate is not constant. For example, components governed by a wear-out function typically have low failure rates initially that tend to increase with time.

The last unique mathematical property of the exponential distribution is the fact that its reliability function is defined by a simplistic transcendental function (recall Equation 2). Thus, when attempting to determine the reliability of a series of independent components, one need merely multiply together their reliability functions, which implies the addition of their exponents. For example, the expression shown as Equation 4 reminds the reader of the algebraic law that allows the easy manipulation of exponentials that are multiplied together.

$$e^{x} \cdot e^{y} \cdot e^{z} = e^{(x+y+z)} \tag{4}$$

More importantly, the exponential distribution's representation by a simple transcendental function is convenient when integrating. The reader is reminded of the integration rule (Equation 5) for an exponential function with a linear exponent (i.e., exponential distribution's density and reliability functions).

$$\int e^{cx} dx = \frac{1}{c} e^{cx} + constant \tag{5}$$

Since most of the mathematical manipulations that are performed in answering reliability questions involve the integration of a distribution's density or reliability function, integrating the exponential function is quite easy. This point becomes more apparent when considering the difficultly in attempting to integrate the normal distribution's density function as shown in Equation 6.

$$\int \frac{1}{\sigma\sqrt{2\pi}} e^{\left[\frac{-(x-\mu)^2}{2\sigma^2}\right]} dx \tag{6}$$

What many in the reliability community often do not appreciate is that the easy mathematical manipulations that are valid only for the exponential distribution are the result of the mathematical tractability propeties that it possesses.

2.3 *Drenick's Theorem*

Drenick(Ref. 4) published a paper in which he proved that, under certain constraints, systems which are composed of a "large" quantity of non-exponentially distributed subcomponents tend themselves toward being exponentially distributed. This profound proof allows reliability practitioners to disregard the failure distributions of the pieces of the system since it is known that the overall system will fail exponentially. Given that most systems are composed of a large number of subcomponents, it would seem that Drenick's Theorem is a reliability analysis godsend. The usefulness of this theorem, however, lies in the applicability of the proof's constraints.

Kececioglu(Ref. 5) delineated the constraints of Drenick's Theorem quite well as follow:

1. The subcomponents are in series.
2. The subcomponents fail independently.
3. A failed subcomponent is replaced immediately.
4. Identical replacement subcomponents are used.

If the four conditions above are met, then as the number of subcomponents and the time of operation tends toward infinity, system failures tend towards being exponentially distributed regardless of the nature of the subcomponents' failure distributions.

The proof of this theorem, which will not be reviewed in this paper, uses some rather sophisticated mathematics and we believe that a demonstration of the theorem would benefit the reader more. Thus, we used the Raptor(Ref. 1) reliability simulator to model the essence of Drenick's Theorem. We modeled the tires of a motorcycle, a dump truck and a semi-tractor trailer truck. We chose to model tires since it is well known that tires do not fail exponentially but in fact fail according to a wear-out function. For our simulations, the tires failed normally with a mean of 40,000 miles and a standard deviation of 10,000 miles. The tires failed independently and were instantaneously replaced (or repaired) with an identical spare. With all four of Drenick's constraints met, our simulation should indicate the validity of his theorem.

We simulated each type of vehicle for 200,000 and 500,000 miles, collected the failure times and then used curve-fitting software to determine how well the failure data tended towards being exponentially distributed. We fit the data to a Weibull distribution and given Drenick's Theorem, the Weibull β-values should tend towards one. Recall that a Weibull distribution with a beta parameter equal to one is by definition an exponential distribution. As the results of Table 1 indicate, the vehicle with more components and a larger simulation time does indeed tend toward being exponentially distributed.

		Number of Tires		
		2	10	18
Miles	200,000	1.55	0.91	0.95
	500,000	1.45	0.98	1.01

Table 1. Weibull Beta Values

A tangent discussion of Drenick's Theorem is that if the system has already reached steady state, then the system will tend toward being exponentially distributed even faster. What Drenick is saying is that if a system's components have varying degrees of life already exhausted when you begin to apply his theorem, then the tendency for the system to behave as an exponential system occurs more quickly. The Raptor software has a feature that allows for randomization of the number of miles already consumed for any tire before the start of a simulation. We conducted such a simulation for the dump truck for 200,000 miles, and the results are shown in Table 2. The simulation results conformed well to Drenick's Theorem, as is indicated by the β-value more quickly tending towards one by merely randomizing the life exhausted on the dump truck's tires.

200,000 Miles for 10 Tires

	β
All New Tires (Repeated from Tabel 1)	0.91
Random Life Exhausted Tires	0.97

Table 2. Weibull Beta Values

Drenick's Theorem is a powerful tool to aid in the reliability analysis of a system when used correctly. The requirement that the number of subcomponents must tend towards infinity is a difficult constraint to meet. Determining if a system contains enough subcomponents for Drenick's Theorem to be valid can only realistically be answered by using a simulator. Furthermore, the constraint that requires the time of operation to approach infinity may be unrealistic, since many systems would be taken out of service before this time is reached. For example, is it likely that a dump truck would operate for 500,000 miles? When used correctly, Drenick's Theorem is a very powerful aspect of the exponential distribution.

3.0 THE BAD

3.1 *Redundancy Approximation*

A system comprised of exponential (and non-exponential, according to Drenick) components in series is itself exponentially distributed. A system comprised of exponential components in any redundancy configuration can be shown to be unquestionably non-exponentially distributed. Yet, it is very common for reliability engineers to ignore this fact and assume redundant subsystems are exponentially distributed. This assumption can ultimately lead to serious errors in their analysis.

Let us review how the assumption of exponentiality for redundant systems is typically applied by stepping through the flawed logic. First, the reader should recall that a system of identical non-repairable redundant *k-out-of-n* subcomponents, which themselves are exponentially distributed, has a MTTF (Murphy per Ref. 6) as described by Equation 7.

$$\theta = MTTF_{k/n} = \frac{\sum_{i=k}^{n} \frac{1}{i}}{\lambda} \qquad (7)$$

Some reliability practitioners now assume that the mean value can be inverted and the results used in Equation 2, because they believe Equation 3 is valid. This process is described by Equations 8 and 9.

$$\lambda_{k/n\,Flawed} = \lambda_F = \frac{1}{MTTF_{k/n}} \qquad (8)$$

$$R(t)_{k/nFlawed} = R(t)_{k/nF} = e^{-\lambda_F t} \qquad (9)$$

It seems quite a shame that this error is committed so often, especially since there exists a true reliability value for identical non-repairable *k-out-of-n* exponential components, which is shown as Equation 10 (Murphy per Ref. 6).

$$R(t)_{k/nTrue} = R(t)_{k/nT} = \sum_{i=k}^{n} \binom{n}{i} e^{-\lambda t i}\left[1 - e^{-\lambda t}\right]^{n-i} \qquad (10)$$

To demonstrate the impact of assuming redundant systems are exponentially distributed, we calculated the reliability of several redundant systems using both Equations 9 and 10. We then calculated the percent error induced by the exponential assumption for redundant systems. We conducted the analysis using non-repairable exponential components with a mean life of 100 hours. We calculated the true and flawed reliability at the particular time of 25 hours.

		n				
		1	2	3	4	5
k	1		11.0	11.8	11.1	10.3
	2			15.3	17.6	16.9
	3				17.1	21.4
	4					17.3
	5					

Table 3. Percent Errors of Truth Minus Flawed R(25)

As the results of Table 3 indicate, rather large errors in the reliability at a particular time can occur by using Equations 8 and 9. Admittedly, we chose this particular example which resulted in large errors to explain our point; times other than 25 hours could yield very small errors and thus the flawed method could be a close approximation to reality. Nonetheless, given that Equation 10 exists, it is difficult to understand why the flawed method is still used by some reliability practitioners.

3.2 *Relatively Long Simulation Times*

Often, reliability practitioners complain that reliability simulators do not quickly produce expected results when the exponential distribution is used to represent failure times. This crankiness could be fueled by two misunderstandings. The first is a lack of appreciation of simulation theory and the other is a failure to thoroughly understand the impact of an exponentially distributed component with a large mean life. Proper simulation theory is a vast subject too large to discuss in this paper, but we would recommend our readers to review Murphy and Carter's(Ref. 7) paper on the appropriate length and number of trials a simulation requires to generate good results.

Many reliability engineers do not realize that the exponential distribution has a relatively large variance compared to its mean value. In fact, the variance of the exponential distribution is the square of its mean life. Thus, an exponential system that fails every 100 hours has a variance of 10,000 hours^2. An exponential system with a mean of 1,000 hours has a variance of one million hours^2. These large variances make it difficult for short simulations to fully mimic a component that fails exponentially.

A normally distributed component with a mean life of 1,000 hours and variance of one million hours^2 is unlikely to be produced by any quality-oriented manufacturer, but the same values for an exponential component are quite reasonable.

The explosive growth of the variance for exponential components with realistic mean values causes simulators to require very long run times to accurately represent them. Grosh(Ref. 8) simply explains the nature of the exponential distribution as follows: "although there is a preponderance of small values there is a *heavy tail* that corresponds to a nontrivial number of very large values". For short simulations, the exponential distribution can yield poor results compared to other reliability distributions.

To further explain this point we conducted a Raptor simulation using a single component that fails exponentially with a mean of 100 hours. We captured the times to failure of the component and then used a curve-fitter to see how long it would take the simulation to mimic the component failing exponentially. The results are shown in Table 4.

	Number of Failures				
	10	25	100	250	500
β	1.46	1.06	1.00	1.09	1.01
η	172.41	143.20	104.83	102.82	102.50

Table 4. Simulating a Known Exponential Component

The reader should appreciate that the component was defined within the simulation to fail exponentially but it took quite a few failures to actually represent this fact. Many reliability practitioners often think that 25 or so failures is sufficient based on their misunderstanding of the Central Limit Theorem, but as our results indicate this is just not the case. Our example needed several hundred failures before the Weibull β-value approached the value of one and its η-value approached 100, as it should if the component fails exponentially.

Simulations that contain exponential components will require relatively long run times; that is just the nature of the beast. This is often perceived as a bad aspect of the exponential distribution, but if more properly identified as a natural constraint, it becomes less of an issue as computer processors continue to improve.

4.0 THE UGLY

4.1 *As a Repair Distribution*

We find it odd that many reliability engineers choose to represent the repair data of a component as exponentially distributed. In our experience, we have never seen repair data come even close to fitting the exponential distribution. We are not saying that repair data should not be represented by the exponential distribution, but, given our experience, it would not be our first (or second, third, etc.) choice.

We suspect that the pervasive use of Markovian analysis techniques in the last century has lead to the widespread use of the exponential repair assumption. Without the assumption of exponentiality, the Markovian analysis techniques would be mathematically intensive. However, with the number of easy to use reliability simulators available today, the need for Markovian analysis techniques, and thus the implementation of the exponential repair assumption, is questionable. Sometimes your favorite analysis techniques have to be retired for more modern and accurate analysis methods.

Our experience has indicated, instead, that the lognormal distribution is an excellent choice for representing repair data. Reliability engineers should seriously consider not using the exponential distribution for repair data unless a good fit, as indicated by a curve-fitter, justifies its use.

4.2 *Failure to Appreciate the Memory-less Property*

The exponential distribution is the only continuous distribution that possesses the memory-less property. Simply stated, the memory-less property asserts that the amount of time already exhausted for an exponential system has no bearing whatsoever on the likelihood of the system failing in the future.

For example, suppose you purchase a 100-year-old house that contains exponentially distributed circuit breaker fuses. These fuses have been protecting the house the entire 100 years and have never failed. What should you say to a friend who gives you a box of brand new fuses (that are identical to the original fuses) as a house-warming gift, implying that you should replace the antique fuses? A good reliability engineer would explain that there is no need to replace non-failed exponential components based on the memory-less property, and ask the friend for a more useful house-warming gift. However, many people would not be comfortable with the decision to leave 100-year-old fuses in the house, believing that the old fuses could not be as good as identical brand new fuses. Either the fuses are not governed by the exponential distribution, in which case it may be wise to replace the fuses, or the person is unfamiliar with the memory-less property.

Another example might help to further explain the memory-less property to those readers still not appreciating its importance. Let us assume that you purchased six glasses three years ago. Somehow, one of the glasses is placed so far back into the cupboard that you never have the opportunity to use that particular glass. The other glasses are used on a daily basis. Assume that the only failure mode for a glass is the act of dropping it from a sufficient height to cause breakage. You decide to host a large party for your reliability friends and you realize that you will need all the glasses that you possess for the raucous party. Scouring the cupboard, you find all of your glasses, including the misplaced glass. If the glasses are exponentially distributed, is the glass with no usage time over the past three years less likely to fail than the other glasses that have been used daily? If you understand the memory-less property of the exponential distribution, then you understand that all six glasses have an equal chance of failing during the party. It is completely irrelevant that some of the glassed have three years of life exhausted while one glass has no life exhausted.

If the reader claims a system is exponentially distributed, then it obeys the memory-less property. If the reader does not believe that any system could truly obey the memory-less property, that is, some amount of wear on the product occurs over time, then the reader should never use the exponential distribution. A product is exponentially distributed if and only if it obeys the memory-less property.

The memory-less property of the exponential distribution really should be considered a Good aspect of the distribution but we consider it an Ugly aspect since few truly understand or appreciate the property's real meaning. They often assume the exponential distribution is valid for their systems, although they refuse to accept the memory-less property for the same systems. Such a use of the exponential distribution is illogical.

5.0 SYNOPSIS

The exponential distribution has great potential to ease reliability analysis when used properly. Unfortunately, the exponential distribution is overly used to the point of abuse. This paper attempted to explain the Good aspects of the exponential distribution, which are commonly recognized throughout the reliability community. More importantly, this paper attempted to bring to light the Bad and the Ugly attributes of the exponential distribution, which are realized when it is misapplied. The ultimate goal of this paper was to increase the reader's sophistication with respect to the exponential distribution and to encourage more practitioners to preach against its misuse.

REFERENCES

1. RAPTOR Software, Version 6.0, ARINC Incorporated, 2001.

2. Dimitri Kececioglu, *Reliability Engineering Handbook Vol I*, PTR Prentice Hall, New Jersey, 1991, pp 261.

3. Robert B. Abernethy, *The New Weibull Handbook*, Reliability Analysis Center, New York, 1996, pp 8-12.

4. R. F. Drenick, *The Failure Law of Complex Equipment*, The Journal of the Society for Industrial Applications of Mathematics, December 1960, Vol 8, No. 4, pp 680-689.

5. Dimitri Kececioglu, *Reliability Engineering Handbook Vol II*, PTR Prentice Hall, New Jersey, 1991, pp 341-349.

6. Kenneth E. Murphy, *Reliability Theory Workbook*, USAF Printing, New Mexico, 1994, pp 2.2-2.

7. Kenneth E. Murphy and Charles M. Carter, *How long Should I Simulate, and for How Many Trials? A Practical Guide to Reliability Simulations.* Annual Reliability and Maintainability Symposium, 2001 Proceedings, pp. 207-212.

8. Doris L. Grosh, *A Primer of Reliability Theory*, John Wiley & Sons, New York, 1989, pp 43.

BIOGRAPHY

Kenneth E. Murphy
ARINC
2309 Renard Place SE, Suite 200
Albuquerque, NM, 87106, USA

kmurphy@arinc.com

Kenneth E. Murphy graduated from the University of Colorado with a Bachelor of Science degree in Aerospace Engineering in 1989. In 1990, he graduated from the University of Alabama with a Masters of Science in Systems Engineering and a minor in Reliability Engineering. Ken was a reliability engineer for the United States Air Force for eleven years where he introduced his visions of the way reliability analysis and testing should be conducted. Ken developed the theory of a quick but robust reliability simulator in 1992 and three years later his idea became a reality when the free reliability software tool called RAPTOR was distributed to the world. In 1996 Ken became the chief reliability engineer for the Air Force's operational testing agency where he continued to lead the development of RAPTOR versions 3.0 and 4.0. Ken is currently a principal reliability engineer with the ARINC Corporation where he is building a reliability skunk-works shop responsible for the development of new versions of RAPTOR, reliability analysis consulting, and RAPTOR training.

Charles M. Carter
ARINC
2309 Renard Place SE, Suite 200
Albuquerque, NM, 87106, USA

ccarter@arinc.com

Charles M. Carter has a B.S. in Aeronautical and Astronautical Engineering from the University of Illinois and a M.S. in Systems Engineering from the Air Force Institute of Technology. As an officer in the USAF, Chuck was assigned to Vandenberg AFB, CA, where he coordinated ground processing of Titan II and Titan IV payloads and served on the launch crew. He was also assigned to Kirtland AFB, NM, where he worked as a reliability engineer evaluating space and electronics systems. He later worked for SAIC at Johnson Space Center as a payload safety engineer, performing safety analysis on Space Shuttle and International Space Station payloads. Chuck is currently a principal engineer with ARINC and is responsible for development of the RAPTOR simulation engine. He is the original author of the RAPTOR simulation engine and developed the algorithms RAPTOR uses to track all system state changes.

Steven O. Brown
ARINC
2309 Renard Place SE, Suite 200
Albuquerque, NM, 87106, USA

sbrown@arinc.com

Steven O. Brown graduated from Memphis State University with a B.S. degree in electrical engineering in 1986. In 1992, he graduated from the Air Force Institute of Technology with a M.S. degree in Operations Research specializing in simulation. As a commissioned officer in the U. S. Air Force for 14 years, Steve performed a wide range of duties from technical to managerial. In 1992 while serving as a reliability engineer for the Air Force's Test and Evaluation Center, Steve developed the idea of building a generic reliability model which could be used to model the availability of ground station-satellite systems. Steve was the original simulation engine designer for generic satellite availability. In 1997, after a three-year detour to Florida, Steve returned to the RAPTOR team and assumed the role of lead designer of the graphical user interface. In 2001, he separated from the Air Force to work RAPTOR full time. Steve is currently a principal reliability analyst with the ARINC Corporation where he is primarily responsible for RAPTOR's graphical user interface, but also develops a wide variety of other simulation tools.