# Improving Availability via Staggered Systems
## Part 2: Mitigating Redundant Failures via System Staggering

**Dr. Bruce D. Holenstein**  >>  **President & CEO**  >>  Gravic, Inc.

**Dr. Bill Highleyman**  >>  **Managing Editor**  >>  Availability Digest

**Paul J. Holenstein**  >>  **Executive Vice President**  >>  Gravic, Inc.

I n Part 1 of this series, we introduced the concept of Mean Time To Failure, MTTF, to replace the commonly used measure, MTBF. MTBF is the Mean Time Between Failures and is constant. It is memoryless. If MTBF is 1,000 hours and 500 hours have passed, MTBF is still 1,000 hours. The probability that the system will fail in the next 1,000 hours is still the same.

### Mean Time To Failure (MTTF)

However, this is not the real world. As time passes, the probability of failure grows nearer. This is what MTTF measures. As time goes on, the mean time to failure gets shorter in most real-life scenarios. This is illustrated in Figure 1 and Figure 2. In these figures, a typical failure probability distribution for a system, $p_f(t)$, is shown. Its value at time $t_1$ is $p_1$. The probability that the system will fail in some small interval, $\Delta t$, around time $t_1$ is $p_1\Delta t$. The MTTF is the average of these probabilities starting from some particular time. As shown in Figure 1 starting at time zero, the MTTF is

$$MTTF = \sum_{i=0}^{\infty} t_i p_i \Delta t \qquad [1]$$

However, if we wait a time T as shown in Figure 2, the MTTF is now

$$MTTF = \frac{\sum_{i=1}^{\infty} (t_i - T)p_i \Delta t}{\sum_{i=1}^{\infty} p_i \Delta t} = \frac{\sum_{i=1}^{\infty} t_i p_i \Delta t}{\sum_{i=1}^{\infty} p_i \Delta t} - T \qquad [2]$$

Comparing Equations [1] and [2], it is clear that MTTF has become smaller as time goes on.
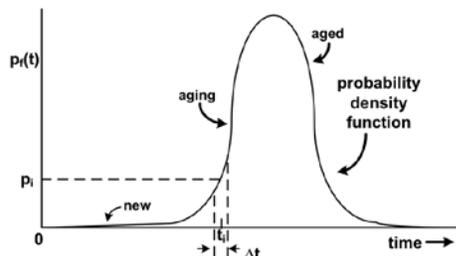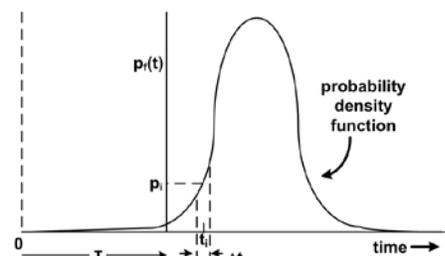


Figure 1: MTTF at Time 0



Figure 2: MTTF at Time T

### System Staggering

We also introduced in Part 1 the concept of staggering system starting times in a redundant system to improve reliability. If the systems are started at the same time, their peak failure probabilities are aligned and the probability that they will fail at the same time is high. However, if they are started at different times (staggered starting times), the peak failure probability of one system likely will be aligned with a much smaller probability of failure of the other system. Therefore, the probability that both systems will fail at the same time is significantly less.

In this Part 2, we analyze the improvement in availability that can be obtained via staggered starts.

### MTTF for Staggered Starts

The staggered redundant system we will be analyzing comprises System 1 and System 2. Let s be the amount of stagger time. That is, System 2 has been started after System 1 by a time of s.

The failure probability distribution for System 1 is $p_{f1}(t)$. The failure probability distribution for System 2 is $p_{f2}(t-s)$ since it was started at a time s after System 1.

Should System 1 fail, it will take an average time of MTR (mean time to repair) to return it to operation. The probability that system 1 will still be failed at time $t$ is

probability that System 1 will still be failed at time $t = \int_{t-MTR}^{t} p_{f1}(t')dt'$ \qquad [3]

The probability that System 2 will fail by time T while System 1 is still down is

probablility that System 1 and System 2 will fail by time $T = \int\limits_{0}^{T}\left[p_{f2}(t-s)\int\limits_{t-MTR}^{t}p_{f1}(t')dt'\right]dt$  (4)

The MTTF of the staggered system is

$$MTTF(s) = \int\limits_{0}^{\infty}t\left[p_{f2}(t-s)\int\limits_{t-MTR}^{t}p_{f1}(t')dt'\right]dt$$  (5)

For small MTR, Equation (5) can be simplified since $p_{fi}(t)$ will be almost constant over the time MTR. Thus,

Probability of a dual-system failure as MTR → 0 =

$$\lim_{MTR\to0}\int\limits_{0}^{T}\left[p_{2}(t-s)\int\limits_{t-MTR}^{t}p_{1}(t')dt'\right]dt = MTR\int\limits_{0}^{T}p_{2}(t-s)p_{1}(t)dt$$  (6)

MTTF is a function of the stagger time, *s*. To determine the optimum value for s (that is, the value that will yield the maximum MTTF value), differentiate MTTF(s) with respect to s and set the result to zero:

$$\frac{d}{ds}MTTF(s) = 0$$  (7)

Solving for s will enable one to pick the optimal value of the stagger time to reduce the mean time to failure.

## Software Failures

Software has different failure modes than hardware. Software bugs lurk in areas that seldom get executed and typically go undetected until the erroneous code is executed. When this happens, the system typically fails.

As an example, consider a software counter overflow. The program includes a transaction counter that is incremented on every transaction. However, if the transaction counter should reach its limit, the next transaction will cause a counter overflow that, in our example, will cause the system to crash.

The probability of failure of the system, $p_f(t)$, is shown in Figure 3. $p_f(t)$ is normally zero (the system is operational). However, when the counter overflows, $p_f(t)$ jumps to one (the system is down). The system will remain down until it is rebooted, at which time the system is again operational and $p_f(t)$ returns to zero. The availability of the system is shown in Figure 4.
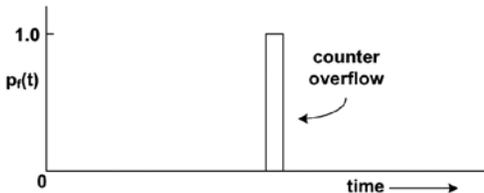


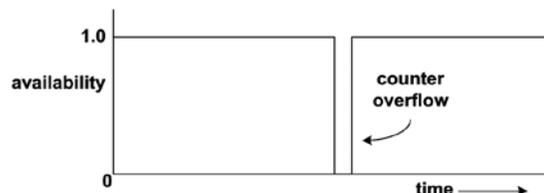Figure 3: Software Counter Overflow Failure



Figure 4: Availability Due to a Software Counter Overflow Failure

In a redundant system, if the two systems are started simultaneously or are not sufficiently staggered, the software fault will occur at the same time in both systems, as shown in Figure 5 and Figure 6. When both systems fail, the capacity of the redundant systems falls to zero.
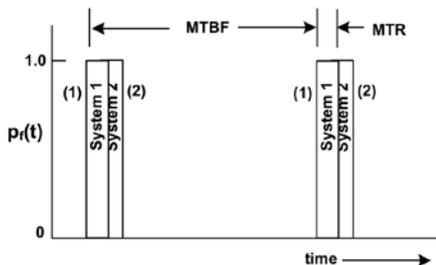


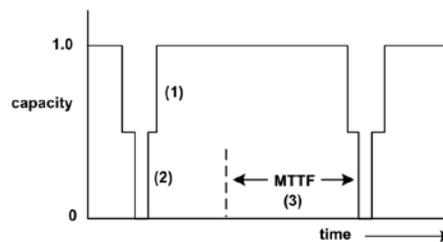Figure 5: Insufficiently Staggered Software Systems



Figure 6: Availability of Insufficiently Staggered Software Systems

However, if the two systems are sufficiently staggered, as shown in Figure 7 and Figure 8, one system will always be operational; and the redundant system capacity will not fall below 0.5.
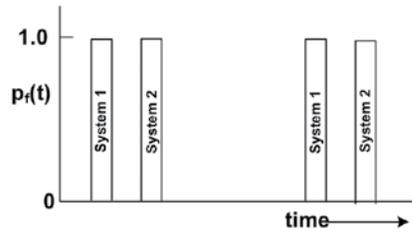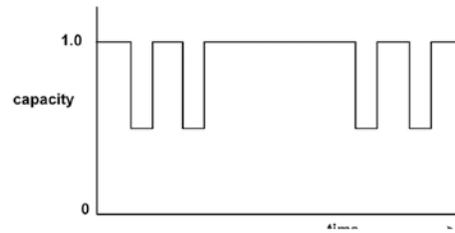
Figure 7: Staggered Software Systems



Figure 8: Availability of Staggered Software Systems

## Correlation

A primary goal in choosing the proper stagger time s is to minimize the correlation between the failure probability distributions of System 1 and System 2. If the two distributions are highly correlated, then the peak of one distribution likely aligns with the peak of the other distribution. The probability of a dual system failure is higher than if the peaks of each distribution are not aligned.

The correlation between the two failure probability distributions can be measured by the correlation coefficient. Let there be two data sets, $\{x_1...x_n\}$ and $\{y_1...y_n\}$. The degree to which these two data sets are correlated is given by the correlation coefficient, r:

$$r = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}}$$

[8]

where $\bar{x}$ and $\bar{y}$ are the means of the data sets.

The meaning of the correlation coefficient can be seen more clearly graphically, as shown in Figure 9 and Figure 10. Figure 9 shows two highly correlated distribution functions, x(t) and y(t). Note that at each point in time, $(x_i - \bar{x})$ and $(y_i - \bar{y})$ are either both positive or negative. Therefore, the correlation coefficient, $r$, is positive since the summation in the numerator, $(x_i - \bar{x})(y_i - \bar{y})$, is always positive.

Similarly, Figure 10 shows two distribution functions that are poorly correlated. At many points in time when $(x_i - \bar{x})$ is positive, $(y_i - \bar{y})$ is negative and vice versa. Therefore, the numerator in r can approach zero.

Note that for a pair of perfectly correlated functions, $(x_i - \bar{x}) = k(y_i - \bar{y})$, and r is equal to 1. For a pair of functions that are exactly anti-correlated, $(x_i - \bar{x}) = -k(y_i - \bar{y})$, and r is equal to -1. Therefore, the correlation coefficient, r, ranges from +1 for perfectly correlated functions to -1 for perfectly anti-correlated functions. If r = 0, there is no correlation between the functions.
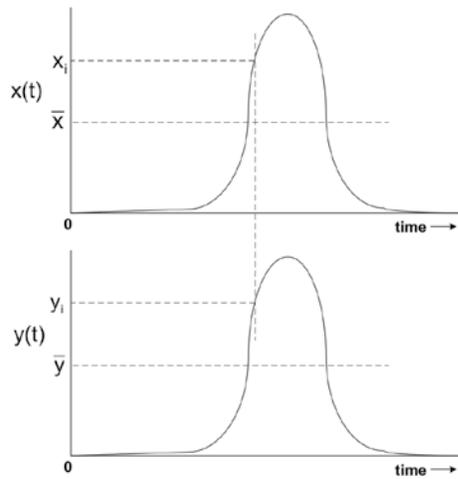

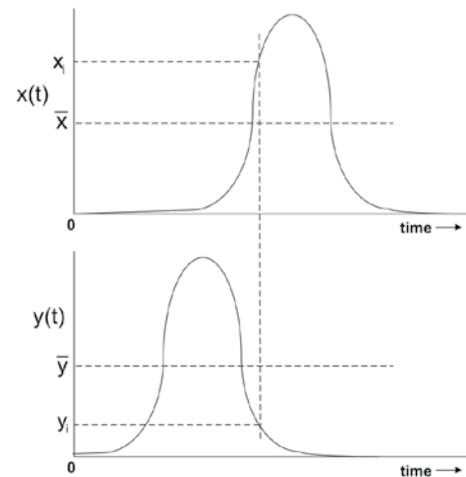
Figure 9: Highly Correlated Distributions



Figure 10: Poorly Correlated Distributions

As previously discussed, Equation 7 yields the optimal value for the stagger time s as it takes into consideration MTR. However, a close approximation that is useful in practice is to simply use the stagger value that minimizes the correlation between the failure probability distribution for System 1 and System 2. As stagger time is varied, the correlation coefficient between the two probability distributions will look something like that shown in Figure 11.

Assuming that the failure probability distributions for System 1 and System 2 are nearly identical, the correlation coefficient, r, will be +1.0 when they are started simultaneously (s = 0). As their starting times are staggered, the correlation coefficient will decrease. Ultimately, it will increase as the starting times become once again coordinated. The stagger time at which the correlation coefficient reaches its minimum value is the stagger time that will lead to the maximum redundant system reliability.
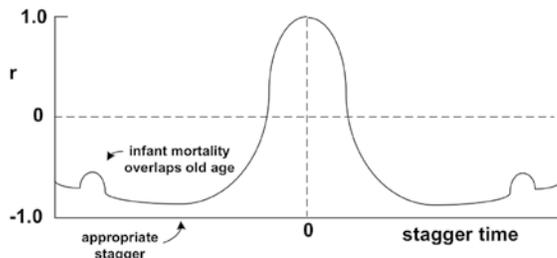
Figure 11: Correlation Coefficient *r* as a Function of Stagger Time *s*

## Methods for Improving the Availability of Redundant Systems

The availability of a redundant system can be improved by incorporating a Failure Analysis Engine that continually monitors the current MTTF of the system and a Failure Prevention Engine that performs one or more appropriate actions to increase the availability of the system should its MTTF fall below an acceptable value as determined by the Failure Analysis Engine.[1]

### Failure Analysis Engine

Typical functions for a Failure Analysis Engine include:
1. Determine the probability distribution of failures for each subsystem in the redundant system.
2. Determine the stagger time that leads to the lowest MTTF.
3. With the distributions of the two systems appropriately staggered, determine the probability distribution of failures for the redundant system by multiplying the probability of failure at each instant of time for the two independent systems.
4. Determine the MTTF of the redundant system via the summation methods described in the Section entitled "Mean Time to Failure."
5. Periodically recalculate the MTTF from the current time.
6. If the MTTF falls below a critical threshold, issue a warning message; and inform the Failure Prevention Engine.
7. When the Failure Prevention Engine completes whatever actions it is going to take, return to Step 1.

### Failure Prevention Engine

When the Failure Prevention Engine is notified by the Failure Analysis Engine that the system MTTF has fallen below an acceptable threshold, it will take actions to improve the availability of the system by lengthening its MTTF. Typical actions include:
1. Restart a node if it is about to fail due to a software problem.
2. Replace a critical hardware component if hardware failures happen after some period of time (such as a solid-state disk drive).
3. Dispatch a repairman if manual intervention is required.
4. Replace a system in its entirety if it is nearing end-of-life.
5. Restart the nodes according to a staggering schedule if both nodes are apt to fail simultaneously from a software or hardware problem.

[1] Note that the functions of the described Engines may be accomplished by operator inspection and intervention until a product becomes available that automates and makes them transparent.

6. Reroute users to the node with the lowest probability of failure (i.e., the longest MTTF).
7. Add another node to the cluster or network of nodes such that the overall probability of failure of the cluster is reduced.
8. Move to an active/active configuration (in which both nodes are processing transactions) or to a sizzling-hot-takeover configuration (an active/active system in which only one node is processing transactions) to lower the recovery time (Recovery Time Objective, or RTO) and the amount of data loss (the Recovery Point Objective, or RPO).

## Summary

The prior art for calculating estimated availability from any point in time is flawed because it is based on memoryless random variables. The calculation of the average time to the next failure is always the same regardless of how long a system has been in service. Rather, an alternative term, Mean Time To Failure (MTTF), provides an estimate of the expected time of failure that decreases as time goes on.

The reliability of a redundant system is optimized by minimizing the probability that both systems will fail simultaneously. If they both have the same failure probability distribution, then when one system is most likely to fail, so is the other system.

By staggering the system starting times so that their probability distributions are not aligned, the time that the two systems are most likely to fail are different. When one system is most likely to fail, the probability that the other system will fail is significantly reduced. Therefore, the probability of a dual system failure is reduced. Redundant system reliability can be greatly enhanced by staggering the starting times of the two systems. This strategy applies both to hardware failures and to software failures.

A method for estimating the remaining availability of a system is to use a Failure Analysis Engine to calculate the system's MTTF from the current time based on the known failure probability distribution of the system. This can be accomplished by sampling the failure probability distribution and summing the probabilities of failure at each of a series of small time intervals. If the MTTF falls below a critical threshold, actions can be taken via a Failure Prevention Engine to mitigate the potential failure based on the expected cause of the failure.

*Dr. Bruce D. Holenstein, President and CEO. Dr. Holenstein leads all aspects of Gravic, Inc. as President and CEO. He started company operations with his brother, Paul, in 1980, and is presently leading the company through the changes needed to accommodate significant future growth. His technical fields of expertise include algorithms, mathematical modeling, availability architectures, data replication, pattern recognition systems, process control and turnkey software development. Dr. Holenstein is a well-known author of articles and books on high availability systems. He received his BSEE from Bucknell University and his Ph.D. in Astronomy and Astrophysics from the University of Pennsylvania.*

*Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him a popular seminar speaker and a sought-after technical writer. Dr. Highleyman is a past chairman of ITUG, the former HP NonStop Users' Group, the holder of numerous U.S. patents, the author of Performance Analysis of Transaction Processing Systems, and the co-author of the three-volume series, Breaking the Availability Barrier.*

*Paul J. Holenstein is Executive Vice President, Gravic, Inc. He has direct responsibility for the Gravic, Inc. Shadowbase Products Group and is a Senior Fellow at Gravic Labs, the company's intellectual property group. He has previously held various positions in technology consulting companies, from software engineer through technical management to business development, beginning his career as a Tandem (HPE NonStop) developer in 1980. His technical areas of expertise include high availability designs and architectures, data replication technologies, heterogeneous application and data integration, and communications and performance analysis. Mr. Holenstein holds many patents in the field of data replication and synchronization, writes extensively on high and continuous availability topics, and co-authored Breaking the Availability Barrier, a three-volume book series. He received his BSCE from Bucknell University, a MSCS from Villanova University, and is an HPE Master Accredited Systems Engineer (MASE). To contact the author, please email: SBProductManagement@gravic.com . Hewlett Packard Enterprise directly sells and supports HPE Shadowbase Solutions (www.ShadowbaseSoftware.com); please contact your local HPE account team.*