Fault Tolerance

1 Introduction

Fault tolerance is the ability of a system to continue functioning in spite of malfunctions or faults. As a notion, it is tightly coupled with the concept of reliability, the lack of defects and the availability of a system.

Reliability R(t) of a system at time t is the probability that the system operates without failure in the interval [0,t), given that the system was performing correctly at time 0. As a probability function, its values lie in the [0, 1] interval. This measure is suitable for applications in which even a momentary disruption can prove costly. One example is computers that control physical processes such as an aircraft, for which failure would result in catastrophe. Mathematically, if a system starts functioning at time 0 and a failure occurs at time T, the system's reliability function may be expressed as:

$$R(t) = P(T > t \mid OK@t = 0) \quad (1)$$

Another metric we can use to measure a system's fault tolerance is the *Mean Time Between Failures*, or *MTBF*. This parameter can be derived through the observation of the behaviour of a system during its operational lifetime. The simplest model that includes fault tolerance is one in which the system switches between only two states: one in which it is fully operational and another one in which the system is completely failed. Transitions between the two states happen upon the occurrence of a failure, or after the system has undergone a repair, as depicted in Figure 1.

We can apply this model to simple systems, such as a lightbulb - which can be either on or burned out, a wire in a circuit - which could be either connected or has a break in it. To some extent, we can apply it also to more complex systems, such as a car or a web server, but we will have to define what "operational" and "failed" mean in this context. For example, "operational" for the webserver could mean that it is completely available for requests from its clients and "failed" could mean that it is either completely unresponsive due to a failure or undergoing maintenance work.

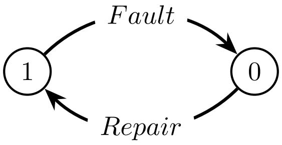

Figure 1: Simple state transition graph for a system with failure and repair

If we draw a timeline of the system described above, we will have intervals in which the system is operational that alternate with the intervals in which the system is down for repair, as in Figure 2. Starting a time 0, the system is operational until it encounters a failure. This first Time-To-Fail interval is noted TTF_1. Conversely, the following interval represents the time in which the system has been down and awaiting repair, so we can name it the first Time-To-Repair, or TTR_1. This succession of intervals continues throughout the entire lifetime of the system.

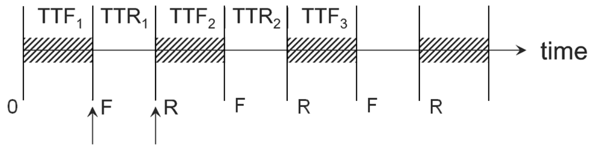

Figure 2: The lifetime of a system with consecutive functioning and repair episodes

Measuring these intervals and averaging their values over a long observational period yields two important metrics: the *Mean Time to Failure*, or *MTTF* which is an average of all Time-To-Fail (TTF) intervals, and the *Mean Time to Repair*, or *MTTR*, which is the average of all Time-To-Repair intervals.

$$MTTF = \sum_i \frac{TTF_i}{n} \qquad\qquad MTTR = \sum_i \frac{TTR_i}{n} \tag{2}$$

Using these two notions, we can define the *Mean Time Between Failures, MTBF* as the average expected time between two failures for a repairable system:

$$MTBF = MTTF + MTTR \tag{3}$$

There are not many systems which are designed to operate continuously without interruption and without any kind of maintenance. In most cases, we are interested not only in the reliability of the system, but also in the number of failures and the time needed to recover the system from a failed state. For a webserver we are interested in maximizing its uptime, meaning the fraction of time that the system is in the operational state. This parameter is expressed by *Availability.*

Availability A(t) of a system at time t is the probability that the system is functioning correctly at the instant of time t. A(t) is also referred as point availability, or instantaneous availability. This measure is appropriate for applications in which continuous performance is not vital but where it would be expensive to have the system down for a significant amount of time. An airline reservation system needs to be highly available, because downtime can put off customers and lose sales; however, an occasional (very) short-duration failure can be well tolerated.

Often it is necessary to determine the *Interval*, or *Mission Availability.* It is defined by

$$A(T) = \frac{1}{T}\int_0^T A(t)dt \tag{4}$$

A(T) is the value of the point availability averaged over some interval of time *T*. This interval might be the life-time of a system or the time to accomplish some particular task.

Finally, it is often found that after some initial transient effect, the point availability assumes a time-independent value. In this case, we are referring to *Steady-state availability*, also known as Long-term Availability denoted by *A(∞)*, which is defined as:

$$A(\infty) = \lim_{T\to\infty} A(T) = \lim_{T\to\infty}\left(\frac{1}{T}\int_0^T A(t)dt\right) \tag{5}$$

A(∞) can be interpreted as the probability that the system will be up at some random point in time, and is meaningful only in systems that include repair of faulty components. If a system cannot be repaired, the point availability *A(t)* is equal to the systems reliability, i.e. the probability that the system has not failed between 0 and t. Thus, as *T* goes to infinity, the steady-state availability of a non-repairable system goes to zero

$$A(\infty) = 0 \tag{6}$$

The long-term availability $A(\infty)$, or more simply written A, can be calculated from MTTF, MTBF, and MTTR as follows:

$$A = \frac{MTTF}{MTBF} = \frac{MTTF}{MTTF + MTTR} \tag{7}$$

It is possible for a low-reliability system to have high availability: consider a system that fails every hour on the average but comes back up after only a second. Such a system has an MTBF of just 1 hour and, consequently, a low reliability; however, its availability is high: A = 3599/3600 = 0.99972.

Steady-state availability is often specified in terms of downtime per year. Table 1 shows examples for some of the values for availability and the corresponding downtime.

Availability(%)	Downtime per year	Downtime per month	Downtime per week
90% ("one nine")	36.5 days	72 hours	16.8 hours
99% ("two nines")	3.65 days	7.2 hours	1.68 hours
99.9% ("three nines")	8.76 hours	43.2 minutes	10.1 minutes
99.99% ("four nines")	52.56 minutes	4.32 minutes	1.01 minutes
99.999% ("five nines")	5.26 minutes	25.9 seconds	6.05 seconds
99.9999% ("six nines")	31.5 seconds	2.59 seconds	0.605 seconds

Table 1: Availability and the corresponding downtime per year.

Availability is typically used as a measure for systems where short interruptions can be tolerated. Networked systems, such as telephone switching and web servers, fall into this category. A customer of a telephone system expects to complete a call without interruptions. However, a downtime of three minutes a year is considered acceptable. Surveys show that web users lose patience when web sites take longer than eight seconds to show results. This means that such web sites should be available all the time and should respond quickly even when a large number of clients concurrently access them. Another example is the electrical power control system. Customers expect power to be available 24 hours a day, every day, in any weather condition. In some cases, a prolonged power failure may lead to health hazards, due to the loss of services such as water pumps, heating, light, or medical attention. Industries may suffer substantial financial loss.

2 Probability Theory Basics

Probability is the branch of mathematics which studies the possible outcomes of given events together with their relative likelihoods and distributions. In common language, the word "probability" is used to mean the chance that a particular event will occur expressed on a linear scale from 0 (impossibility) to 1 (certainty).

The first axiom of probability theory states that the value of probability of an event A lies between 0 and 1:

$$0 \leq P(A) \leq 1 \tag{8}$$

$\overline{A}$ denotes the event *not A*. For example, if A stands for it rains, $\overline{A}$ stands for it does not rain. The second axiom of probability theory says that the probability of an event A is equal to 1 minus the probability of the event A:

$$P(\overline{A}) = 1 - P(A) \tag{9}$$

Suppose that one event, A is dependent on another event, B. Then $P(A \mid B)$ denotes the conditional probability of event A, given event B. The fourth rule of probability theory states that the probability $P(A \cdot B)$ that both A and B will occur is equal to the probability that B occurs times the conditional probability $P(A \mid B)$:

$$P(A \cdot B) = P(A \mid B) \cdot P(B) \tag{10}$$

If P(B) is greater than zero, then equation (10) can be written as

$$P(A \mid B) = \frac{P(A \cdot B)}{P(B)} \tag{11}$$

An important condition that we will often assume is that two events are mutually independent. For events A and B to be independent, the probability *P(A)* does not depend on whether B has already occurred or not, and vice versa.

Thus, $P(A \mid B) = P(A)$. So, for independent events, the rule (10) reduces to

$$P(A \cdot B) = P(A) \cdot P(B) \tag{12}$$

This is the definition of independence, that the probability of two events both occurring is the product of the probabilities of each event occurring. Situations also arise when the events are mutually exclusive. That is, if A occurs, B cannot, and vice versa. As such, we can write the following

$$P(A \cdot B) = P(B \cdot A) = 0 \tag{13}$$

This is the definition of mutually exclusiveness, that the probability of two events both occurring is zero.

Let us now consider the situation when either A, or B, or both event may occur. The probability $P(A+B)$ is given by

$$P(A + B) = P(A) + P(B) - P(A \cdot B) \tag{14}$$

Combining (13) with (14), we get the following expression for mutually exclusive events

$$P(A + B) = P(A) + P(B) \tag{15}$$

As an example, consider a system consisting of three identical components A, B and C, each having a reliability R. Let us compute the probability of exactly one out of three components failing, assuming that the failures of the individual components are independent. By rule (9), the probability that a single component fails is $1 - R$. Then, by rule (12), the probability that a single component fails and the other two remain operational is $(1 - R)R^2$. Since the probabilities of any of the three components to fail are the same, then the overall probability of one component failing and other two not is $3(1 - R)R^2$. The three probabilities are added by applying rule (15), because the events are mutually exclusive.

3 Hardware Fault Tolerance

Hardware fault tolerance is the most mature area in the general field of fault tolerant computing. Many hardware fault-tolerance techniques have been developed and used in practice in critical applications ranging from telephone exchanges to space missions. In the past, the main obstacle to a wide use of hardware fault tolerance has been the cost of the extra hardware required. With the continued reduction in the cost of hardware, this is no longer a significant drawback, and the use of hardware fault-tolerance techniques is expected to increase. However, other constraints, notably on power consumption, may continue to restrict the use of massive redundancy in many applications.

The single most important parameter used in the reliability analysis of hardware systems is the component failure rate, which is the rate at which an individual component suffers faults. This is the expected number of failures per unit time that a currently good component will suffer in a given future time interval.

Failure rate λ is the expected number of failures per unit time. For example, if a processor fails, on average, once every 1000 hours, then it has a failure rate 1/1000 failures/hour. Often failure rate data is available at component level, but not for the entire system. This is because several professional organizations collect and publish failure rate estimates for frequently used components (diodes, switches, gates, flip-flops, etc.). At the same time the design of a new system may involve new configurations of such standard components.

The failure rate depends on the current age of the component, any voltage or physical shocks that it suffers, the ambient temperature, and the technology. The dependence on age is usually captured by what is known as the *bathtub curve* (see Figure 3). Failure rate changes as a function of time. For hardware, a typical evolution of failure rate over a systems life-time is characterized by the phases of *infant mortality* (I), *useful life* (II) and *wear-out* (III). Failure rate is high at first due to frequent failures in weak components with manufacturing defects overlooked during manufacturers testing (poor soldering, leaking capacitor, etc.), then stabilizes after a certain time and then increases as electronic or mechanical components of the system physically wear out.

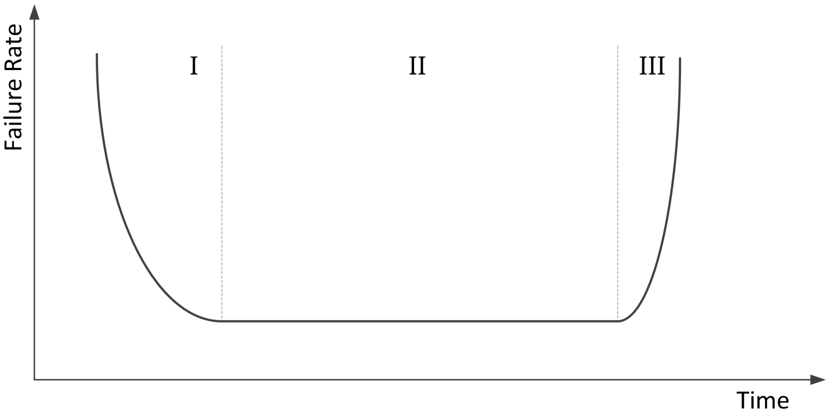

Figure 3: The ”bathtub curve” of failure intensity versus time

When components are very young, their failure rate is quite high. This is due to the chance that some components with manufacturing defects slipped through manufacturing quality control and were released. As time goes on, these components are weeded out, and the component spends the bulk of its life showing a fairly constant failure rate. As it becomes very old, aging effects start to take over, and the failure rate rises again.

The impact of the other factors can be expressed through the following empirical failure rate formula:

$$\lambda = \pi_L \pi_Q (C_1 \pi_T \pi_V + C_2 \pi_E) \tag{16}$$

where the notations are as follows:

- λ - Failure rate of component.
- π_L - Learning factor, associated with how mature the technology is.
- π_Q - Quality factor, representing manufacturing process quality control (ranging from 0.25 to 20.00).
- π_T - Temperature factor, with values ranging from 0.1 to 1000. It is proportional to $e^{\frac{E_a}{kT}}$, where E_a is the activation energy in electron-volts associated with the technology, k is the Boltzmann constant (0.8625 104 eV/K), and T is the temperature in Kelvin.
- π_V - Voltage stress factor for CMOS devices; can range from 1 to 10, depending on the supply voltage and the temperature; does not apply to other technologies (where it is set to 1).
- π_E - Environment shock factor; ranges from very low (about 0.4), when the component is in an air-conditioned office environment, to very high (13.0) when it is in a harsh environment.
- C_1, C_2 - Complexity factors; functions of the number of gates on the chip and the number of pins in the package.

Further details can be found in MIL-HDBK-217E, which is a handbook produced by the U.S. Department of Defense.

Devices operating in space, which is replete with charged particles and can subject devices to severe temperature swings, can thus be expected to fail much more often than their counterparts in air-conditioned offices, so too can computers in automobiles (which suffer high temperatures and vibration) and industrial applications.

Software failure rate usually decreases as a function of time. A possible curve is shown in Figure 4. The three phases of evolution are: test/debug (I), useful life (II) and obsolescence (III).

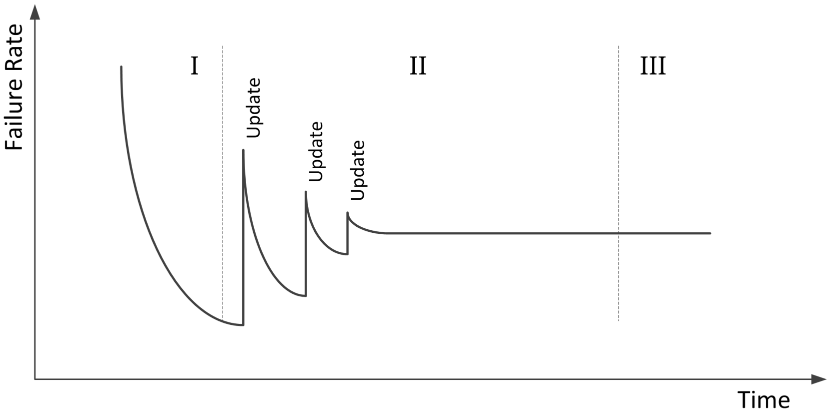

Figure 4: The failure intensity curve for software versus time

Software failure rate during useful life depends on the following factors:

1. software process used to develop the design and code
2. complexity of software,
3. size of software,
4. experience of the development team,
5. percentage of code reused from a previous stable project,
6. rigor and depth of testing at test/debug (I) phase.

There are two major differences between hardware and software curves. One difference is that, in the useful-life phase, software normally experiences an increase in failure rate each time a feature upgrade is made. Since the functionality is enhanced by an upgrade, the complexity of software is likely to be increased, increasing the probability of faults. After the increase in failure rate due to an upgrade, the failure rate levels off gradually, partly because of the bugs found and fixed after the upgrades. The second difference is that, in the last phase, software does not have an increasing failure rate as hardware does. In this phase, the software is approaching obsolescence and there is no motivation for more upgrades or changes.

4 Failure Rate, Reliability, and Mean Time to Failure

In this section, we consider a single component of a more complex system, show how reliability and Mean Time Between Failures (MTBF) can be derived from the basic notion of failure rate. Consider a component that is operational at time t = 0 and remains operational until it is hit by a failure.

Suppose now that all failures are permanent and irreparable. Let T denote the lifetime of the component (the time until it fails), and let f(t) and F(t) denote the probability density function of T and the cumulative distribution function of T, respectively. These functions are defined for $t \geq 0$ only (because the lifetime cannot be negative) and are related through

$$f(t) = \frac{dF(t)}{dt} \qquad F(t) = \int_0^t f(\tau)d\tau \tag{17}$$

f(t) represents (but is not equal to) the momentary probability of failure at time t. To be exact, for a very small $\Delta t, f(t)\Delta t \approx Prob(t \leq T \leq t + \Delta t)$. Being a density function, f(t) must satisfy

$$f(t) \geq 0, \forall t \geq 0 \;\; and \;\; \int_0^\infty f(t)dt = 1 \tag{18}$$

F(t) is the probability that the component will fail at or before time t

$$F(t) = Prob(t \leq T) \tag{19}$$

R(t), the reliability of a component (the probability that it will survive at least until time t), is given by

$$R(t) = Prob(T > t) = 1 - F(t) \tag{20}$$

f(t) represents the probability that a new component will fail at time t in the future. A more meaningful quantity is the probability that a good component of current age t will fail in the next instant of length Δt. This is a conditional probability, since we know that the component survived at least until time t. This conditional probability is represented by the failure rate (also called the hazard rate) of a component at time t, denoted by $\lambda(t)$, which can be calculated as follows:

$$\lambda(t) = \frac{f(t)}{1 - F(t)} \tag{21}$$

Since $f(t) = \frac{dF(t)}{dt} = -\frac{dR(t)}{dt}$, we can rewrite the expression above

$$\lambda(t) = -\frac{dR(t)}{dt}\frac{1}{R(t)} \tag{22}$$

Certain types of components suffer no aging and have a failure rate that is constant over time, $\lambda(t) = \lambda$. In this case,

$$\frac{dR(t)}{dt} = -\lambda R(t) \tag{23}$$

The solution to (23), assuming R(0) = 1, is

$$R(t) = e^{-\lambda t} \tag{24}$$

Therefore, a constant failure rate implies that the lifetime T of the component has an exponential distribution, with a parameter that is equal to the constant failure rate λ.

This law is known as the *exponential failure law*. The plot of reliability as a function of time is shown in Figure 5.

The exponential failure law is very valuable for the analysis of reliability of components and systems in hardware. However, it can only be used in cases when the assumption that the failure rate is constant is adequate.

To summarize the definitions we have derived above

$$f(t) = \lambda e^{-\lambda t} \qquad F(t) = 1 - e^{-\lambda t} \qquad R(t) = e^{-\lambda t} \qquad for\; t \geq 0 \tag{25}$$

The expected value of a random variable, intuitively, is the long-run average value of repetitions of the experiment it represents. For example, the expected value in rolling a six-sided die is 3.5 because, roughly

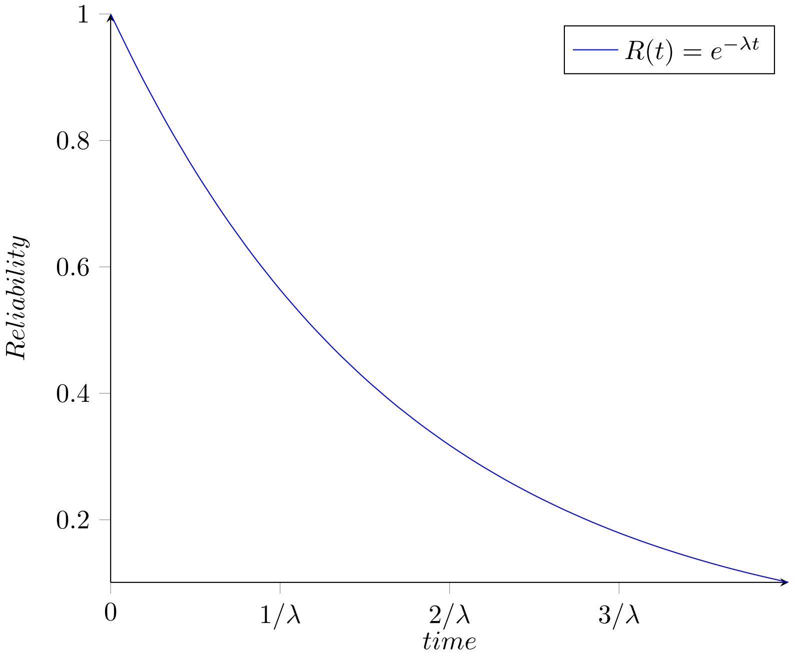

Figure 5: Reliability curve

speaking, the average of all the numbers that come up in an extremely large number of rolls is very nearly always quite close to three and a half.

Suppose random variable X can take value x_1 with probability p_1, value x_2 with probability p_2, and so on, up to value x_k with probability p_k. Then the expectation of this random variable X is defined as

$$E[X] = p_1x_1 + p_2x_2 + p_3x_3 + .. + p_kx_k \tag{26}$$

Since all probabilities p_i add up to one ($p_1 + p_2 + .. + p_k = 1$), the expected value can be viewed as the weighted average, with p_i being the weights:

$$E[X] = \frac{p_1x_1 + p_2x_2 + p_3x_3 + .. + p_kx_k}{1} = \frac{p_1x_1 + p_2x_2 + p_3x_3 + .. + p_kx_k}{p_1 + p_2 + .. + p_k} \tag{27}$$

For example, let X represent the outcome of a roll of a fair six-sided die. More specifically, X will be the number of pips showing on the top face of the die after the toss. The possible values for X are 1, 2, 3, 4, 5, and 6, all equally likely (each having the probability of $\frac{1}{6}$. The expectation of X is

$$E[X] = 1 \cdot \frac{1}{6} + 2 \cdot \frac{1}{6} + 3 \cdot \frac{1}{6} + 4 \cdot \frac{1}{6} + 5 \cdot \frac{1}{6} + 6 \cdot \frac{1}{6} = 3.5 \tag{28}$$

If one rolls the die n times and computes the average (arithmetic mean) of the results, then as n grows, the average will almost surely converge to the expected value, a fact known as the strong law of large numbers. One example sequence of ten rolls of the die is 2, 3, 1, 2, 5, 6, 2, 2, 2, 6, which has the average of 3.1, with the distance of 0.4 from the expected value of 3.5. The convergence is relatively slow: the probability that the average falls within the range 3.5±0.1 is 21.6% for ten rolls, 46.1% for a hundred rolls and 93.7% for a thousand rolls.

If the probability distribution of X admits a probability density function f(x), then the expected value can be computed as

$$E[X] = \int_{-\infty}^{\infty} xf(x)dx \tag{29}$$

The MTBF of an irreparable component equal to its expected lifetime, E[T]. As the random variable is time, which is always greater than zero, we can rewrite (29) as

$$MTBF = E[T] = \int_0^\infty tf(t)dt \tag{30}$$

Substituting $f(t) = -\frac{dR(t)}{dt}$ we get

$$MTBF = -\int_0^\infty t\frac{dR(t)}{dt}dt = -tR(t)\mid_0^\infty + \int_0^\infty R(t)dt \tag{31}$$

The value of $-tR(t)$ is equal to 0 at $t = 0$ and also to zero at $t \to \infty$, as the reliability of every system asymptotically drops to zero given a long enough time, ($R(\infty) = 0$). Thus, we can write

$$MTBF = \int_0^\infty R(t)dt \tag{32}$$

Given an exponential reliability function with a constant failure rate λ, we can rewrite (32) as

$$MTBF = \int_0^\infty R(t)dt = \int_0^\infty e^{-\lambda t}dt = \frac{1}{\lambda} \tag{33}$$

Although a constant failure rate is used in most calculations of reliability (mainly owing to the simplified derivations), there are cases for which this simplifying assumption is inappropriate, especially during the infant mortality and wearout" phases of a components life (Figure 3). In such cases, the Weibull distribution is often used. This distribution has two parameters, λ and β, and has the following density function of the lifetime T of a component:

$$f(t) = \lambda\beta t^{\beta-1}e^{-\lambda t^\beta} \tag{34}$$

The corresponding failure rate in this case will not be constant and it will be dependant on time

$$\lambda(t) = \lambda\beta t^{\beta-1} \tag{35}$$

This failure rate is an increasing function of time for $\beta > 1$, is constant for $\beta = 1$, and is a decreasing function of time for $\beta < 1$. This makes it very flexible, and especially appropriate for the wear-out and infant mortality phases.

The component reliability for a Weibull distribution is

$$R(t) = e^{-\lambda t^\beta} \tag{36}$$

The MTBF associated with this reliability is

$$MTBF = \int_0^\infty R(t)dt = \frac{\Gamma(\beta^{-1})}{\beta\lambda^{\beta-1}} \tag{37}$$

where $\Gamma(x) = \int_0^\infty y^{x-1}e^{-y}dy$ is the gamma function, which is a generalization for real numbers of the factorial function and satisfies the following

- $\Gamma(x) = (x-1)\Gamma(x-1)$ for $x > 1$
- $\Gamma(0) = \Gamma(1) = 1$
- $\Gamma(n) = \Gamma(n-1)!$ for an integer n, n = 1, 2, ..

Note that the Weibull distribution includes as a special case ($\beta = 1$) the exponential distribution with a constant failure rate λ.

5 Dependability model types

In this section we consider two common dependability models: reliability block diagrams and Markov processes. Reliability block diagrams belong to a class of combinatorial models, which assume that the failures of the individual components are mutually independent. Markov processes belong to a class of stochastic processes which take the dependencies between the component failures into account, making the analysis of more complex scenarios possible.

5.1 Reliability Block Diagrams

Combinatorial reliability models include reliability block diagrams, fault trees, success trees and reliability graphs. In this section we will consider the oldest and most common reliability model: reliability block diagrams. A reliability block diagram presents an abstract view of the system. The components are represented as blocks. The interconnections among the blocks show the operational dependency between the components. Blocks are connected in series if all of them are necessary for the system to be operational. Blocks are connected in parallel if only one of them is sufficient for the system to operate correctly. A diagram for a two-component serial system is shown in Figure 6 a). Figure 6(b) shows a diagram of a two-component parallel system. Models of more complex systems may be built by combining the serial and parallel reliability models.

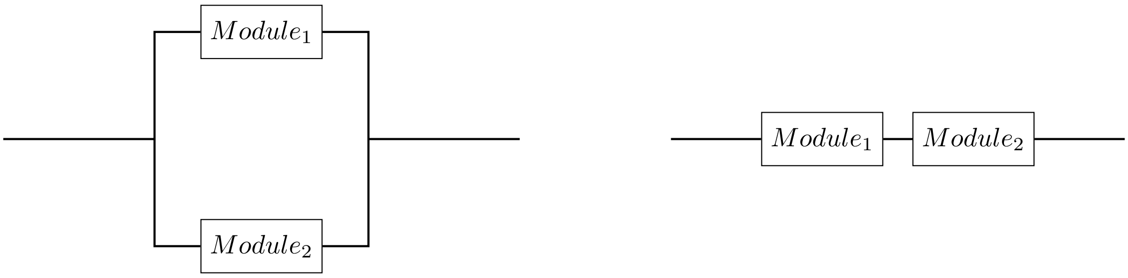

Figure 6: Reliability diagrams for a parallel (a) and a series (b) system

As an example, consider a system consisting of two duplicated processors and a memory. The reliability block diagram for this system is shown in Figure 7. The processors are connected in parallel, since only one of them is sufficient for the system to be operational. The memory is connected in series, since its failure would cause the system failure.

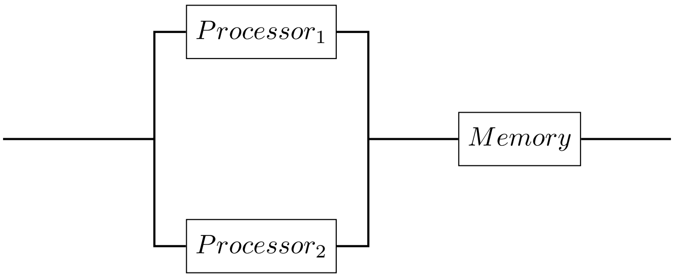

Figure 7: Reliability diagram for a three-component system

Reliability block diagrams are a popular model, because they are easy to understand and to use for modeling systems with redundancy. They are also easy to evaluate using analytical methods. However, reliability block diagrams, as well as other combinatorial reliability models, have a number of serious limitations.

First, reliability block diagrams assume that the system components are limited to the operational and failed states and that the system configuration does not change during the mission. Hence, they cannot model standby components, repair, as well as complex fault detection and recovery mechanisms. Second, the failures of the individual components are assumed to be independent. Therefore, the case when the sequence of component failures affects system reliability cannot be adequately represented.

5.2 Canonical Structures

In this section, we consider some canonical structures, out of which more complex structures can be constructed. We start with the basic series and parallel structures, continue with non-series/parallel ones, and then describe some of the many resilient structures that incorporate redundant components (next referred to as modules). In the next sub-sections, we will use the following notations:

- $R_i = p_i$, the reliability of block i, meaning the probability that functional block i is working properly
- $Q_i = q_i = 1 - p_i$, the probability that functional block i is defective
- R, the reliability of the whole system (i.e. the probability that the whole system is functioning properly)
- $Q = 1 - R$, the probability that the whole system is defective

5.2.1 Series Structures

A series system is defined as a set of N modules connected together so that the failure of any one module causes the entire system to fail. Note that the diagram in Figure 9 is a reliability diagram and not always an electrical one; the output of the first module is not necessarily connected to the input of the second module.

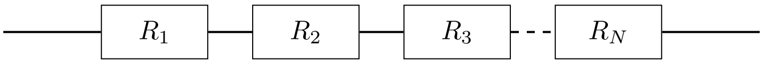

Figure 9: Reliability diagram for a series system

For such a system to function properly, all its units must function properly. Assuming that the modules in Figure 9 fail independently of each other, the reliability of the entire series system is the product of the reliabilities of its N modules.

Denoting with $R_s(t)$ the reliability of the whole system we can write the following,

$$R_s = P(1 \wedge 2 \wedge 3 \wedge .. \wedge N) = p_1 \cdot p_2 \cdot p_3 \cdot \cdot ... \cdot p_N \tag{38}$$

If we denote by $R_i(t)$ the reliability of module i, we can rewrite the equation,

$$R_s = \prod_{i=1}^{N} R_i \tag{39}$$

Also,

$$Q_s = 1 - R_s = 1 - \prod_{i=1}^{N} (1 - Q_i) \tag{40}$$

$$\prod_{i=1}^{N}(1-Q_i) = 1-(Q_1+Q_2+..+Q_N)+(Q_1Q_2+Q_1Q_3+..+Q_{N-1}Q_N)-.. \tag{41}$$

Usually, in order for the whole system to have a high reliability, each block needs to have a high reliability $R_i \geq 0.9$, which means that Q_i is very small, so we can neglect factors that contain a product of at least two Q_i factors. Therefore, we can rewrite (41) as,

$$\prod_{i=1}^{N}(1-Q_i) \approx 1-(Q_1+Q_2+..+Q_N) = 1-\sum_{i=1}^{N} Q_i \tag{42}$$

If we input this into Equation (40), we get

$$Q_s = 1-(1-\prod_{i=1}^{N}(1-Q_i)) = \sum_{i=1}^{N} Q_i \tag{43}$$

If module i has a constant failure rate, denoted by λ_i, then, according to Equation (24), $R_i(t) = e^{-\lambda_i t}$, and consequently

$$R_s(t) = \prod_{i=1}^{N} e^{\lambda_i t} = e^{-\sum_{i=1}^{N}\lambda_i t} = e^{-\lambda_s t} \tag{44}$$

From Equation (40) we see that the series system also follows an exponential repartition and has a constant failure rate equal to λ_s (the sum of the individual failure rates). Using the relation derived in Equation (33), the MTBF of the series system is therefore

$$MTBF_S = \frac{1}{\lambda_s} = \frac{1}{\sum_{i=1}^{N}\lambda_i} = \frac{1}{\sum_{i=1}^{N}\frac{1}{MTBF_i}} \tag{45}$$

This means that

$$MTBF_S < MTBF_i, \forall i = \overline{1,N} \tag{46}$$

The failure rate of a series system increases with the number of units that are linked in series.

$$\lambda_S = \sum_{i=1}^{N}\lambda_i \tag{47}$$

For identical systems, with the same failure rate, $\lambda_i = \lambda$, we can simplify the equation above to:

$$\lambda_S = N\lambda \tag{48}$$

Example Consider the series structure in Figure 10. The four modules in this diagram represent the instruction decode unit (R_{ID}), execution unit (R_{EU}), data cache (R_{DC}), and instruction cache (R_{IC}) in a microprocessor. All four units must be fault-free for the microprocessor to function, although the way they are physically connected does not resemble a series system.

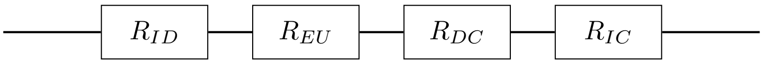

Figure 10: Reliability diagram for a series system

Let's assume the modules have the following constant reliabilities: $R_{ID} = 0.9$, $R_{EU} = 0.95$, $R_{DC} = 0.99$, $R_{IC} = 0.89$. Then, the total reliability of the microprocessor is

$$R_S = R_{ID}\cdot R_{EU}\cdot R_{DC}\cdot R_{IC} = 0.9\cdot 0.95\cdot 0.99\cdot 0.89 \approx 0.75 \tag{49}$$

As a general rule, the reliability of a series structure is lower than the reliability of its individual components. Probabilistically, this can be explained by the fact there are more states in which two modules can fail when working together than individually. It can be noted that, for the processor to have a 99.9% reliability (which is a common figure for today's PCs), the reliability of each of the four subsystems needs to be at least $R = \sqrt[4]{0.999} \approx 0.9998$ If we increase the number of components that are linked in series even further, the overall reliability will decrease asymptotically towards zero.

For example, if we link together an ever increasing number of systems with reliability $R = 0.9$, we will get the following decrease in overall reliability:

- 2 systems: $R_S = 0.9^2 = 81\%$
- 3 systems: $R_S = 0.9^3 = 72.9\%$
- 4 systems: $R_S = 0.9^4 = 65.61\%$
- 5 systems: $R_S = 0.9^5 = 59.05\%$
- 6 systems: $R_S = 0.9^6 = 53.14\%$

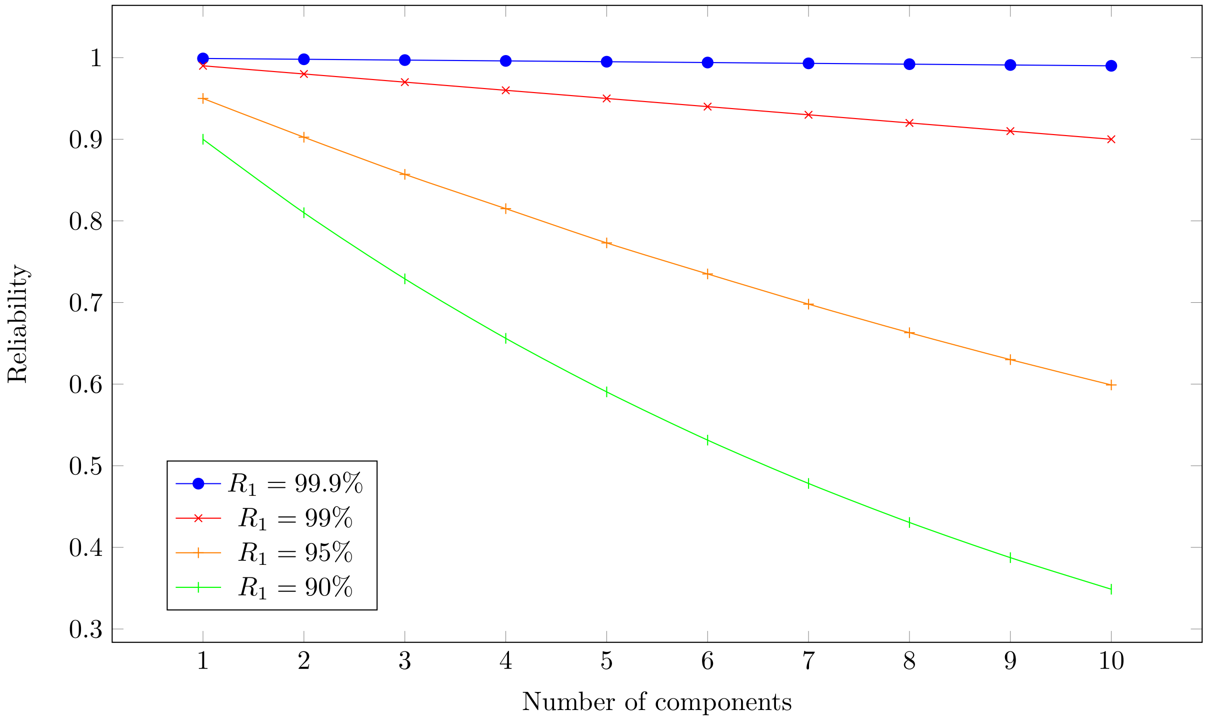

Figure 11: Reliability of a series system with increasing number of components: single component reliability from 90% to 99.9%

5.2.2 Parallel Structures

A parallel system is defined as a set of N modules connected together so that it requires the failure of all the modules for the system to fail, as in Figure 12.

To get to a reliability formula for the parallel structure, we will have to first consider the probability that the whole system will malfunction (Q_p). This will happen when all the blocks malfunction, so block 1, block 2 through block N are all defective. We can express that by,

$$Q_P = P(\overline{1} \wedge \overline{2} \wedge .. \wedge \overline{N}) = \prod_{i=1}^{N} Q_i \tag{50}$$

where all blocks are independent and Q_i is the probability that block i is faulty.

We can therefore express the reliability of a parallel structure of N modules by

$$R_P = 1 - Q_P = 1 - \prod_{i=1}^{N} Q_i = 1 - \prod_{i=1}^{N}(1 - R_i) \tag{51}$$

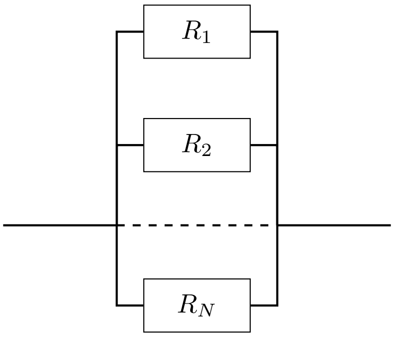

Figure 12: Reliability diagram for a parallel system

If every module has a constant failure rate λ_i, then we can write

$$R_P = 1 - \prod_{i=1}^{N}(1 - e^{-\lambda_i t}) = \sum_{i=1}^{N} e^{-\lambda_i t} - \sum_{i=1,j=1,i\neq j}^{N} e^{-(\lambda_i+\lambda_j)t} + .. + (-1)^{N+1}\prod_{k=1}^{N} e^{-\lambda_k t} \tag{52}$$

To calculate the MTBF, we follow the rule derived in Equation (32):

$$MTBF_P = \int_0^\infty R_P(t)dt = \tag{53}$$

$$= \sum_{i=1}^{N}\int_0^\infty e^{-\lambda_i t}dt - \sum_{i=1,j=1,i\neq j}^{N}\int_0^\infty e^{-(\lambda_i+\lambda_j)t}dt + .. + (-1)^{N+1}\int_0^\infty \left(\prod_{k=1}^{N} e^{-\lambda_k t}\right)dt = \tag{54}$$

$$= \sum_{i=1}^{N}\int_0^\infty e^{-\lambda_i t}dt - \sum_{i=1,j=1,i\neq j}^{N}\int_0^\infty e^{-(\lambda_i+\lambda_j)t}dt + .. + (-1)^{N+1}\int_0^\infty \left(e^{-\sum_{k=1}^{N}\lambda_k t}\right)dt \tag{55}$$

We can simplify Equation (53) by integration:

$$MTBF_P = \sum_{i=1}^{N}\frac{1}{\lambda_i} - \sum_{i=1,j=1,i\neq j}^{N}\frac{1}{\lambda_i+\lambda_j} + .. + (-1)^{N+1}\frac{1}{\sum_{k=1}^{N}\lambda_k} \tag{56}$$

If all systems have the same failure rate $\lambda_i = \lambda_j = .. = \lambda_N = \lambda$, we can rewrite Equation (54):

$$MTBF_P = \frac{N}{\lambda} - \frac{N}{2\lambda} + .. + (-1)^{N+1}\frac{1}{N\lambda} = \frac{1}{\lambda}\sum_{i=1}^{N}(-1)^{N+1}\frac{C_N^i}{i} \tag{57}$$

We can easily substitute the sum in Equation (55) with the partial sum of the harmonic series:

$$\sum_{i=1}^{N}(-1)^{N+1}\frac{C_N^i}{i} = \sum_{i=1}^{N}\frac{1}{i} \tag{58}$$

Therefore, we can simplify Equation (55) and write the MTBF of a parallel system with N identical components as:

$$MTBF_P = \frac{1}{\lambda}\sum_{i=1}^{N}\frac{1}{i} \tag{59}$$

Note that the harmonic series is divergent, so a parallel system does not have a constant failure rate. The failure rate decreases with the increase of the systems that are linked in parallel. We can derive the global failure rate of a system with N identical modules with failure rate λ that are connected in parallel using the result in Equation (59):

$$\lambda_P = \frac{\lambda}{\sum_{i=1}^{N} \frac{1}{i}} \tag{60}$$

Example A system consists of two components in parallel, as in Figure 13. What is the total reliability, MTBF and the failure rate of the system?

We can derive the reliability formula from the general form

$$R_P = 1 - \prod_{i=1}^{2}(1 - R_i) = 1 - (1 - R_1)(1 - R_2) = R_1 + R_2 - R_1 R_2 \tag{61}$$

Presuming $R_1 = e^{-\lambda_1 t}$ and $R_2 = e^{-\lambda_2 t}$, the MTBF of the system can be expressed as

$$MTBF_P = \int_0^\infty R_P(t)dt = \int_0^\infty e^{-\lambda_1 t}dt + \int_0^\infty e^{-\lambda_2 t} - \int_0^\infty e^{-(\lambda_1+\lambda_2)t}dt \tag{62}$$

$$MTBF_P = \frac{1}{\lambda_1} + \frac{1}{\lambda_2} - \frac{1}{\lambda_1 + \lambda_2} \tag{63}$$

If both modules are identical, meaning $R_1 = R_2 = R$, then we can simplify the reliability formula to

$$R_P = 2R - R^2 \tag{64}$$

The MTBF will then be equal to

$$MTBF_P = \frac{3}{2\lambda} \tag{65}$$

and the failure rate of the parallel module will be equal to

$$\lambda_P = \frac{2}{3}\lambda \tag{66}$$

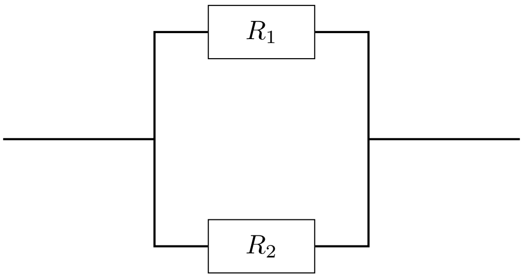

Figure 13: Reliability diagrams for a parallel system with two components

It is worth noting that, as individual reliability functions are $0 \leq R < 1$, R_P will always be greater than R, which means that the reliability of the parallel system will always be greater than the reliability of its individual components.

$$R_P = 2R - R^2 > R, \forall R \in [0, 1) \tag{67}$$

For example, if the reliability of the individual components is $R_1 = R_2 = 0.9$, then, the total reliability of the system is $R_P = 0.99$. If we increase the number of systems in parallel, as in Figure 14, the overall system reliability will also increase:

- 2 systems: $R_P = 1 - (1 - 0.9)^2 = 99\%$

- 3 systems: $R_P = 1 - (1 - 0.9)^3 = 99.9\%$
- 4 systems: $R_P = 1 - (1 - 0.9)^4 = 99.99\%$
- 5 systems: $R_P = 1 - (1 - 0.9)^5 = 99.999\%$
- 6 systems: $R_P = 1 - (1 - 0.9)^6 = 99.9999\%$

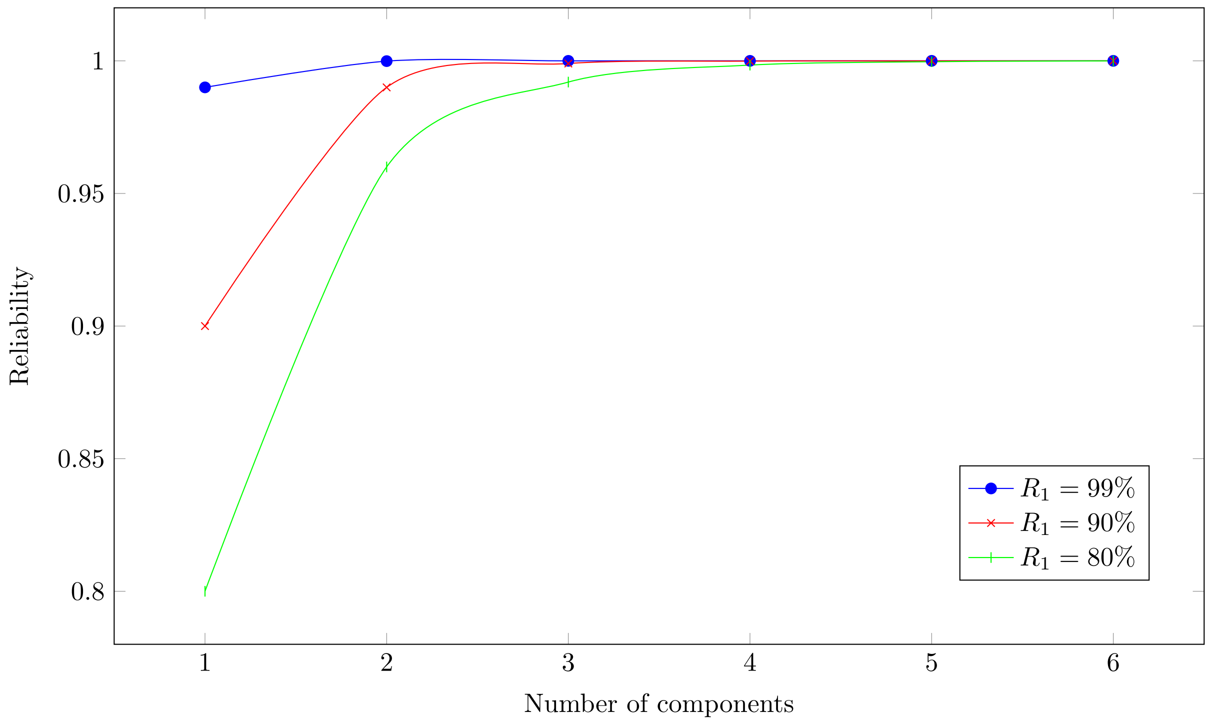

Figure 14: Reliability of a parallel system with increasing number of components: single component reliability of 80%, 90% and 99%

5.2.3 Combination of Series and Parallel

While many smaller systems can be accurately represented by either a simple series or parallel configuration, there may be larger systems that involve both series and parallel configurations in the overall system. Such systems can be analysed by calculating the reliabilities for the individual series and parallel sections and then combining them in the appropriate manner. Such a methodology is illustrated in the following example.

Example Calculating the Reliability for a Combination of Series and Parallel.

Consider a system with three components. Units 1 and 2 are connected in series and Unit 3 is connected in parallel with the first two, as shown in Figure 15.

What is the reliability of the system if $R_1 = 99.5\%$, $R_2 = 98.7\%$ and $R_3 = 97.3\%$ at 100 hours?

First, the reliability of the parallel segment consisting of Units 1 and 2 is calculated: $R_{1,2} = 1 - (1 - R_1)(1 - R_2) = 1 - (1 - 0.995)(1 - 0.987) = 0.999935$

The reliability of the overall system is then calculated by treating Units 1 and 2 as one unit with a reliability of 99.9935% connected in series with Unit 3. Therefore: $R_{1,2,3} = R_{1,2}R_3 = 0.97294$

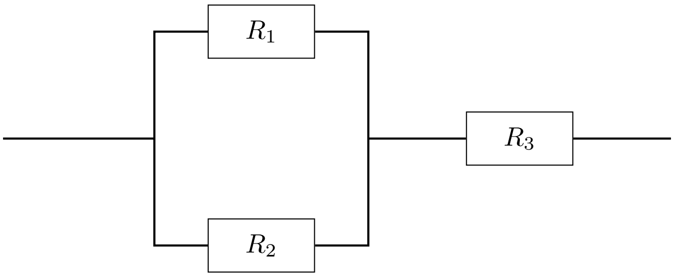

Figure 15: Reliability diagram for a simple series-parallel system