Rel: Estimating Digital System Reliability

Qualitatively, the *reliability* of a digital system is the likelihood that it works correctly when you need it. Marketing and sales people like to say that the systems they sell have "high reliability," meaning that the systems are "pretty likely to keep working." However, savvy customers ask questions that require more concrete answers, such as, "If I buy 100 of these systems, how many will fail in a year?" To provide these answers, digital design engineers are often responsible for calculating the reliability of the systems they design, and in any case they should be aware of the factors that affect reliability.

reliability

Quantitatively, reliability is expressed as a mathematical function of time:

$$R(t) = \text{Probability that the system still works correctly at time } t$$

Reliability is a real number between 0 and 1; that is, at any time $0 \le R(t) \le 1$. We assume that $R(t)$ is a monotonically decreasing function; that is, failures are permanent and we do not consider the effects of repair. Figure Rel-1 is a typical reliability function.

The foregoing definition of reliability assumes that you know the mathematical definition of probability. If you don't, reliability and the equivalent probability are easiest to define in terms of an experiment. Suppose that we were to build and operate N identical copies of the system in question. Let $W_N(t)$ denote the number of them that would still be working at time t. Then,

$$R(t) = \lim_{N \to \infty} W_N(t) / N$$

That is, if we build *lots* of systems, $R(t)$ is the fraction of them that are still working at time t. When we talk about the reliability of a single system, we are simply using our experience with a large population to estimate our chances with a single unit.

It would be very expensive if the only way to to compute $R(t)$ was by experiment—to build and monitor N copies of the system. Worse, for any t, we wouldn't know the value of $R(t)$ until time t had elapsed in real time. Thus, to answer the customer's question posed earlier, we'd have to build a bunch of systems and wait a year; by then, our potential customer would have purchased something else.

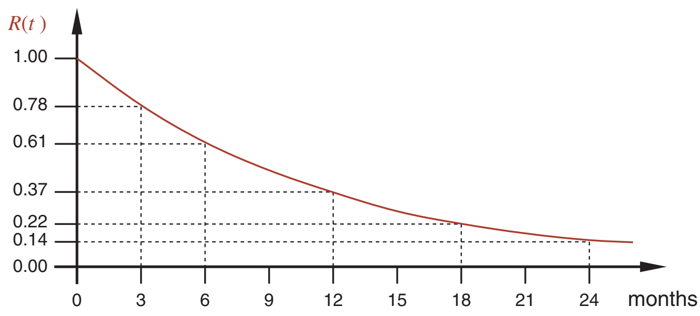

Figure Rel-1 Typical reliability function for a system.

Instead, we can estimate the reliability of a system by combining reliability information for its individual components, using a simple mathematical model. The reliability of mature components (e.g., 74FCT CMOS chips) may be known and published based on actual experimental evidence, while the reliability of new components (e.g., a Sexium microprocessor) may be estimated or extrapolated from experience with something similar. In any case, a component's reliability is typically described by a single number, the "failure rate" described next.

Rel.1 Failure Rates

failure rate

The *failure rate* is the number of failures that occur per unit time in a component or system. In mathematical formulas, failure rate is usually denoted by the Greek letter λ. Since failures occur infrequently in electronic equipment, the failure rate is measured or estimated using many identical copies of a component or system. For example, if we operate 10,000 microprocessor chips for 1,000 hours, and eight of them fail, we would say that the failure rate is

λ

$$\lambda = \frac{8 \text{ failures}}{10^4 \text{ chips} \cdot 10^3 \text{ hours}} = (8 \cdot 10^{-7} \text{ failures/hour})/\text{chip}$$

That is, the failure rate of a single chip is $8 \cdot 10^{-7}$ failures/hour.

The actual process of estimating the reliability of a batch of chips is not nearly as simple as we've just portrayed it; for more information, see the References. However, we can use individual component failure rates, derived by whatever means, in a straightforward mathematical model to predict overall system reliability, as we'll show later in this section.

Since the failure rates for typical electronic components are so small, there are several scaled units that are commonly used for expressing them: percent failures per 10^3 hours, failures per 10^6 hours, and failures per 10^9 hours. The last unit is called a *FIT*:

FIT

$$1 \text{ FIT} = 1 \text{ failure}/(10^9 \text{ hours})$$

In our earlier example, we would say that $\lambda_{\text{microprocessor}} = 800$ FITs.

The failure rate of a typical electronic component is a function of time. As shown in Figure Rel-2, a typical component has a high failure rate during its early life, during which most manufacturing defects make themselves visible; failures during this period are called *infant mortality*. Infant mortality is why manufacturers of high-quality equipment perform *burn-in*—operating the equipment for 8 to 168 hours before shipping it to customers. With burn-in, most infant mortality occurs at the factory rather than at the customer's premises. Even without a thorough burn-in period, the seemingly stingy 90-day warranty offered by most electronic equipment manufacturers does in fact cover most of the failures that occur in the first few *years* of the equipment's operation (but if it

infant mortality

burn-in

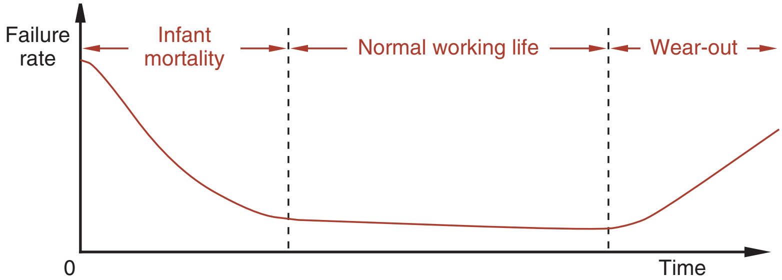

Figure Rel-2 The "bathtub curve" for electronic-component failure rates.

fails on the 91st day, that's tough!). This is quite different from the situation with an automobile or other piece of mechanical equipment, where "wear and tear" increases the failure rate as a function of time.

Once an electronic component has successfully passed the burn-in phase, its failure rate can be expected to be pretty much constant. Depending on the component, there may be *wear-out* mechanisms that occur late in the component's life, increasing its failure rate. In the old days, vacuum tubes often wore out after a few thousand hours because their filaments deteriorated from thermal stress. Nowadays, most electronic equipment reaches obsolescence before its solid-state components start to experience wear-out failures. For example, even though it's been over 25 years since the widespread use of EPROMs began, many of which were guaranteed to store data for only 10 years, we haven't seen a rash of equipment failures caused by their bits leaking away. (Do you know anyone with a 10-year-old PC or VCR?)

wear-out

Thus, in practice, the infant-mortality and wear-out phases of electronic-component lifetime are ignored, and reliability is calculated on the assumption that failure rate is constant during the normal working life of electronic equipment. This assumption, which says that a failure is equally likely at any time in a component's working life, allows us to use a simplified mathematical model to predict system reliability, as we'll show later in this section.

There *are* some other factors that can affect component failure rates, including temperature, humidity, shock, vibration, and power cycling. The most significant of these for ICs is temperature. Many IC failure mechanisms involve chemical reactions between the chip and some kind of contaminant, and these are accelerated by higher temperatures. Likewise, electrically overstressing a transistor, which heats it up too much and eventually destroys it, is worse if the device temperature is high to begin with. Both theoretical and empirical evidence support the following widely used rule of thumb:

- An IC's failure rate roughly doubles for every 10°C rise in temperature.

This rule is true to a greater or lesser degree for most other electronic parts.

Note that the temperature of interest in the foregoing rule is the *internal* temperature of the IC, not the ambient temperature of the surrounding air. A power-hogging component in a system without forced-air cooling may have an internal temperature as much as 40–50°C higher than ambient. A well-placed fan may reduce the temperature rise to 10–20°C, reducing the component's failure rate by perhaps a factor of 10.

Rel.2 Reliability and MTBF

For components with a constant failure rate λ, it can be shown that reliability is an exponential function of time:

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

The reliability curve in Figure Rel-1 is such a function; it happens to use the value $\lambda = 1$ failure/ year.

Another measure of the reliability of a component or system is the *mean time between failures (MTBF)*, the average time that it takes for a component to fail. For components with a constant failure rate λ, it can be shown that MTBF is simply the reciprocal of λ:

mean time between failures (MTBF)

$$\text{MTBF} = 1/\lambda$$

Rel.3 System Reliability

Suppose that we build a system with m components, each with a different failure rate, $\lambda_1, \lambda_2, \ldots, \lambda_m$. Let us assume that for the system to operate properly, *all* of its components must operate properly. Basic probability theory says that the system reliability is then given by the formula

$$\begin{aligned} R_{\text{sys}}(t) &= R_1(t) \cdot R_2(t) \cdot \cdots \cdot R_m(t) \\ &= e^{-\lambda_1 t} \cdot e^{-\lambda_{2(1)} t} \cdot \cdots \cdot e^{-\lambda_m t} \\ &= e^{-(\lambda_1 + \lambda_2 + \cdots + \lambda_m)t} \\ &= e^{-\lambda_{\text{sys}} t} \end{aligned}$$

where

$$\lambda_{\text{sys}} = \lambda_1 + \lambda_2 + \cdots + \lambda_m$$

Thus, system reliability is also an exponential function, using a composite failure rate λ_{sys} that is the sum of the individual component failure rates.

The constant-failure-rate assumption makes it very easy to determine the reliability of a system—simply add the failure rates of the individual components to get the system failure rate. Individual component failure rates can be obtained from manufacturers, reliability handbooks, or company standards.

For example, a portion of one company's standard is listed in Table Rel-1. Since failure rates are just estimates, this company simplifies the designer's job by considering only broad categories of components, rather than giving "exact" failure rates for each component. Other companies use more detailed lists, and

Component	Failure Rate (FITs)
SSI IC	90
MSI IC	160
LSI IC	250
VLSI microprocessor	500
Resistor	10
Decoupling capacitor	15
Connector (per pin)	10
Printed-circuit board	1000

Table Rel-1
Typical component failure rates at 55°C.

some CAD systems maintain a component failure-rate database, so that a circuit's composite failure rate can be calculated automatically from its parts list.

Suppose that we were to build a single-board system with a VLSI microprocessor, 16 memory and other LSI ICs, 2 SSI ICs, 4 MSI ICs, 10 resistors, 24 decoupling capacitors, and a 50-pin connector. Using the numbers in Table Rel-1, we can calculate the composite failure rate,

$$\begin{aligned}\lambda_{\text{sys}} &= 1000 + 16 \cdot 250 + 2 \cdot 90 + 4 \cdot 160 + 10 \cdot 10 + 24 \cdot 15 + 50 \cdot 10 + 500 \text{ FITs} \\ &= 7280 \text{ failures} / 10^9 \text{ hours}\end{aligned}$$

The MTBF of our single-board system is $1/\lambda$, or about 15.6 years. Yes, small circuits really can be that reliable. Of course, if we include a typical power supply, with a failure rate of over 10,000 FITs, the MTBF is more than halved.