JOURNAL OF PROPULSION AND POWER
Vol. 14, No. 5, September–October 1998

Quality Issues in Propulsion

John P. McCarty*
McCarty Group, Inc., Huntsville, Alabama 35806
and
Garry M. Lyles†
NASA Marshall Space Flight Center, Huntsville, Alabama 35812

In this paper, a high-quality propulsion system is one that has high reliability. Thus, quality is a high probability within tolerance performance or operation. Because failures are out of tolerance performance, the probability that failures will occur is the difference between high- and low-quality systems. Failures can be described at three levels: 1) the system failure (the detectable end of a failure), 2) the failure mode (the failure process), and 3) the failure cause (the start). Failure causes can be evaluated and classified by type. The results of typing flight history failures show that most failures are in unrecognized modes and result from human error or noise, that is, failures are when engineers learn how things really work. Although the study is based on U.S. launch vehicles, a sampling of failures from other countries indicates that the finding has broad application. The parameters of the design of a propulsion system are not single valued, but have dispersions associated with the manufacturing of parts. Many tests are needed to find failures, if the dispersions are large relative to tolerances, that could contribute to the large number of failures in unrecognized modes.

Nomenclature

C = occurrence of failure cause
g = performance function
K = number of system failures
k = system failure index
L = number of failure causes
l = failure cause index
n = number of
P_f = probability of failure
$P(\ )$ = probability of ( )
x = predictor variable associated with a failure cause
α = average expected probability of failure of new failure causes introduced when a change is made
β = rate of reduction proportionality constant
δ = integration constant
ζ = expected probability of occurrence of a failure cause
$\xi(\ )$ = expected value of ( )
ρ = expected probability of failure
Ψ = function

Subscripts
c = cause
e = element
k = system failure index
l = failure cause index
m = number of predictor variables associated with a failure cause
t = total launch vehicle data set

Introduction

THIS paper will focus on a different aspect of propulsion engineering—propulsion system quality—and will use the term "quality" differently than normal aerospace usage. This different usage is of interest because of its importance to successful propulsion systems. From our perspective a high-quality propulsion system is one that has high reliability. That is, a high-quality propulsion system is based on a design that has a high probability that the performance or operation of the system will result in the mission terminating within the defined tolerance limits. There have been numerous examples of designs that contribute to low-quality propulsion systems. This paper will attempt to explore some of the reasons that such examples are more numerous in the community than is desirable.

Received Sept. 2, 1997; revision received April 15, 1998; accepted for publication April 29, 1998.

*President. E-mail: mccarty.mg@worldnet.att.net. Fellow AIAA.
†Manager, Advanced Space Transportation Office.

Failures

Performance or operation of a propulsion system that results in the mission terminating outside of the defined tolerance limits is a failure. Each failure can be described at three levels: system failure, failure mode, and failure cause. The system failure is the detectable end that describes the individual performance or output deficiency that prevents successful termination of the mission. Therefore, system failure represents the termination of the failure propagation events in a state that precludes transition of the mission to a final success. Each of the individual propulsion system failures has associated with it one or more failure modes, one of which caused the observed system failure. The failure mode identifies the individual failure mechanism that causes a deviation from the design expectation of a part or assembly feature to lead to a system failure. The failure mode traces the sequence of events between the failure cause and the system failure. Each mode has associated with it one or more failure causes, one of which precipitated the failure. Failure cause identifies the deviation in the feature from the design expectation that initiates the failure mode. The deviation that initiates the failure mode is often unrecognized until the system failure occurs.

Space Shuttle Main Engine Fuel Preburner Burnthrough

The fuel preburner (FPB) of the Space Shuttle Main Engine (SSME) provides the turbine drive gases for the high-pressure fuel turbopump (HPFTP). The FPB consists of two propellant manifolds, a central igniter, an injector, and a short combustion section that is welded into the hot gas manifold. The hot gas manifold also includes the oxidizer preburner, the turbine exhaust gas flow path, and the main combustion chamber injec-

tor. The FPB injector is a coaxial element type with a fuel annulus surrounding an oxidizer core. Combustion takes place below the injector in three compartments formed by baffles. The outer structural shell of the combustor is cooled by hydrogen flowing between a concentric cylindrical liner and the outer shell.

The failure of Engine 0006 in 1980 was a result of a hole burned through the outer structural wall.[1] Until the loss of turbine drive gas through the hole resulted in an external fire, the effects of failure propagation on combustion or other performance were not measurable, although the increased wall heat transfer resulted in the structural failure and the external hole. Postfailure inspection revealed that the individual liquid oxygen (LOX) element posts were not concentric with the fuel annuli, causing a fuel restriction on the outboard side of the outer row elements. Further inspection showed that the lack of concentricity was caused by a deformity, an outward bowing of the injector faceplate halfway between the center and the outer row of elements. This feature deviation, which was unique to this preburner, was not observed until the burnthrough. Periodic inspections for all preburners to verify outer row element concentricity were instituted as a result of this failure.

Value of Testing

In the previous example, the feature deviation that was the failure cause was the geometrical shape of the injector faceplate that was unique to that FPB. It is not always the case that the deviation is a geometrical parameter unique to a specific part. In many cases the deviation is an error in the analytical representation of the behavior of the design, a deviation that is generic to all parts. Since failure causes result from deviations that are unique to a specific part or errors in analysis that are generic to all parts, it is observed that failures are when engineers learn how things actually work.

In the early development of the SSME HPFTP, subsynchronous whirl rotordynamic instability at about one-half shaft speed was observed once the test speeds of the pump exceeded two times the first critical speed. The potential for this instability had been recognized and analysis accomplished, based on analytical predictions of the rotor forces, which showed that the rotor system was stable. Being a true instability, the whirl was self-initiating. After initiation, the amplitude increased rapidly, and within a few cycles, because of the bending of the relatively flexible shaft, the internal clearances closed and rubbing occurred at many locations. At this point the system stiffness increased significantly, limiting further increases in amplitude and raising the first critical speed and whirl frequency. Bearing loads increased significantly in the limit cycle condition, and were three times higher on the turbine end bearings than on the pump end, leading to a larger number of turbine end-bearing failures.

Twenty-two potential drivers were identified. A vigorous analysis and test program eventually concluded that two factors were the most significant: The hydrodynamic cross-coupling forces of the pump interstage seals, and the low natural frequency of the rotating assembly. A series of design changes were made over a period of time to decrease the cross-coupling forces and provide damping at the seals, and to stiffen the shaft and bearing supports to increase the critical speed. The changes eventually increased the whirl inception speed to above the operating region.

As higher speed operation resulted from the design changes that increased the whirl inception speed, it became evident that the turbine end bearings were overheating and that the mechanism was not related to whirl. Turbopump instrumentation indicated that the hydrogen coolant flow to the bearing was low and was allowing hot gas to backflow into the bearing from the upstream face of the first turbine rotor. Detailed reexamination of the bearing coolant flow analysis in conjunction with the physical evidence led to the conclusion that the flow path had been incorrectly modeled and that a free vortex existed in the path at the base of the pump shaft. A paddle wheel was added that changed the free vortex into a forced vortex and reduced the pressure loss by a factor of almost 50.

Characteristics of Failure Causes

A propulsion failure implies output or operation that leads to a state that is out of performance or functional limits. This failure could be caused by an element failure, by an input variable or an environmental variable that is outside expected limits, or it could be caused by a description error in the input/output transfer function such that the element did not give the expected output when supplied with an input.

Typically, components have multiple failure modes, e.g., a valve could fail open, fail closed, fail to track commands, etc. In general, each component failure mode has multiple causes that could originate from structural, mechanical, chemical, or other sources. Each of these failure causes has associated with it a probability density function and related cumulative density function. The following list indicates some of the characteristics that provide the variables that control the density functions of each cause.

1) Design and analysis: Environment characterization (loads and distribution functions), property characterization (material design values and distribution functions), process characterization (dimensional tolerances), and modeling errors (inappropriate assumptions, math errors, communication errors).

2) Fabrication: Fabrication variability, tool wear, setup errors, and inspection errors.

3) Manufacturing assembly: Tolerance accumulations, assembly process variations, assembly errors, and inspection errors.

4) Launch preparation and checkout: Assembly process variation, assembly errors, and test and inspection errors.

5) Environment: Current atmospheric conditions and history, and performance of related systems.

Any typical launch-propulsion system involves hundreds of components; thousands of parts; hundreds of thousands to millions of process steps, part features, and failure causes. Therefore there is a very large number of failure causes associated with each system. The probability of failure is the union of the probabilities of occurrence of all of the failure causes, C_{kl}

$$P_f = \cup_{k=1}^{K} \cup_{l=1}^{L_k} P(C_{kl}) \tag{1}$$

where K, a countable large number, is the number of system failures, and L_k is the number of failure causes that can result in the Kth system failure.[2] In general, each failure cause, C_{kl}, has a different performance function that determines the probability of occurrence:

$$g_{kl}(\boldsymbol{x}) = g_{kl}(x_1, x_2, \ldots, x_m), \quad k = 1, 2, \ldots, K, \quad l = 1, 2, \ldots, L_k \tag{2}$$

where the m parameters of the vector $\boldsymbol{x}$ can, in general, assume random values. The m parameters can be design factors, environment factors, manufacturing factors, assembly factors, etc. For each performance function, a failure region and a nonfailure region can be defined, such that the occurrence of any failure cause[3] is defined as

$$C_{kl} = [g_{kl}(\boldsymbol{x}) \leq 0] \tag{3}$$

Once the processing is complete at launch or test commit, the failure cause performance functions for that sample are fixed. Consequently, the production and preparation processing can be viewed as drawing a discrete experimental realization of the extremely numerous C_{kl} from sampling populations. It is the cumulative effect of the C_{kl} samples that determines the probability of failure.

For any test or flight, the causes of failure can be sorted from the largest to the smallest probability of occurrence on the basis of the expected value, $\xi(C_{kl})$, and the number of causes falling in the interval $d\zeta$ about each value ζ of the following equation can be counted:

$$\zeta = \xi(C_{ij}) \tag{4}$$

The number of causes, n_c, at each ζ can be graphed against ζ. Although the shape of the graph is not evident, some characteristics can be described. ζ must be between zero and one. All values of n_c are finite and to obtain an overall probability of failure of less than 1, the integral of the $\zeta - n_c$ curve must be less than 1. In general, the right-hand tail of the curve, which represents the more likely causes, is expected to determine the probability of failure during a test or launch. If one or a few of the causes have a high probability of occurrence, then it is expected that the more frequent occurrence of the cause(s) would lead to correction and elimination, leading in turn to a lower probability of failure controlled by the more numerous, but less likely causes. This improvement in failure rate because of element evolution results from design changes. However, any design change, because it is a change, introduces new failure causes.

HPFTP Turbine End Thermal Shield Nut

The turbine end bearings and support structure of the HPFTP is protected from FPB hot gases by a gas barrier and a thermal shield. A single nut, in line with the pump shaft, retains the thermal shield. As a result of thermal strain, cracks developed in the retainer and nut. The parts were redesigned to relieve the thermal strain and installed in an HPFTP on Engine 2013. On test 901-364 in 1982, a failure occurred. Posttest investigation showed that an HPFTP turbine end-bearing failure had occurred, leading to extensive damage to the engine. The failure cause was finally isolated; it was a hot gas leak path through the thermal shield and gas barrier that had been created by the redesign. Hot gas leakage through this path mixed with the bearing coolant, increasing the coolant temperature, so that inadequate cooling was provided for the bearing.

Failure Cause Types

The history of 11 U.S. launch vehicle families, numbering 1,757 launches through the end of 1992, was investigated. Of these launches, 310 were failures. These failures were evaluated to differentiate the failures into two groups: Design mode failures (DM) and unrecognized mode failures (UM). DMs are those that occur in a mode of behavior that the system was or should have been explicitly designed to prevent. Unrecognized mode failures are those that occur in a mode of behavior that the system was not explicitly designed to prevent. Because DM failures are present in any design, random failures are expected and specific requirements and criteria must be established to control their frequency of occurrence. Objective data and information along with appropriate models can be used to evaluate the probability of occurrence for DM type failures, and tests can be performed to verify that the infrequency criteria are met. Failure then is caused by expected, but infrequent, random events. Unrecognized modes, on the other hand, are not subject to explicit evaluation. Causes of UM failures include ignorance and oversight. ''Ignorance'' includes cases where the specific system failure or cause was not recognized by the profession or was known but analysis methods and physical data were not available to adequately describe the failure or cause. ''Oversight'' includes cases where the failure was known to the profession but was not accounted for by the designers or producers of the specific element.

Results of Typing Flight History Failures

Nine broad categories for failure causes were developed to facilitate consistent classification of each failure: 1) Well understood and modeled failure because of extreme value of demand or capability; 2) designer did not allow for basic mode of behavior well understood by existing technology; 3) failure in a basic mode of behavior that was recognized but poorly understood by existing technology; 4) failure in a basic mode of behavior that was not recognized; 5) failure because of construction (manufacture, assembly, or launch operations) not providing a prescribed feature; 6) failure because of construction variability in a nonprescribed feature; 7) failure because of misuse beyond prescribed feature limits; 8) failure because of misuse beyond a nonprescribed feature limit; and 9) failure because of an independent event (deliberate mistake, sabotage). The information on each failure was evaluated and it was classified relative to the nine failure cause categories or designated as unknown when the available information was not sufficient to permit assignment to a category. Each failure classified in the nine categories was then assigned a location in a two-dimensional array as DM or UM, and noise, human error, or design feature equaling a limit state according to the following classification scheme: Design feature = limit state, 1 (DM); human error, 2, 7 (DM), 5, 8 (UM); and noise, 3 (DM), 4, 6, 9 (UM). Those failures designated as unknown were assigned to the ''unk-unk'' location of the array. The results are shown in Fig. 1. Note that no failures were assigned to categories 7 or 9.

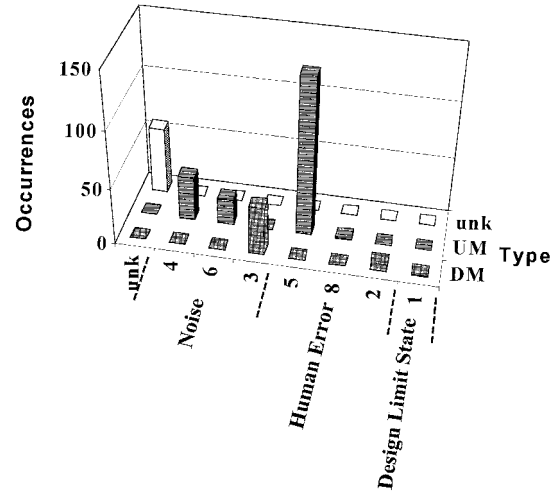

Fig. 1 Failure type.

As can be seen relatively few failures are design mode. More prevalent are the unrecognized mode failures associated with human errors or unrecognized variations in critical parameters, i.e., noise. These data tend to support the hypothesis of Brown.[4] Brown shows that the failure rates predicted by the objective information used in structural reliability analysis in civil construction projects are too small by a factor of 10 or more. Brown notes that this set of objective parameters excludes many actual failure causes: Factors such as numerical mistakes, incorrect idealization, omissions, construction errors, and the political and financial climate. Because unrecognized modes are not subject to explicit evaluation by engineering analysis, historical data are necessary to address failure rate expectations, under the assumption that similar industrial practices are used, so that the population characteristics remain unchanged. Although the present study is based on 11 U.S. launch vehicle families, evaluation of a sample of failures from other countries indicates that the finding has broad application.

Context

As noted in the previous text, it is expected that the more frequent occurrence of the more likely cause(s) will lead to correction and elimination, leading in turn to a lower probability of failure controlled by the more numerous, but less likely causes. This process of learning and improvement is

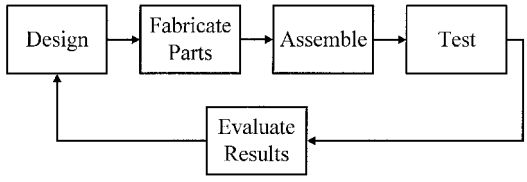

Fig. 2 Design process.

shown schematically in Fig. 2. A design is prepared and analyzed for its expected behavior relative to performance, fabrication, assembly, and test. Although the ability of models to estimate the expected behavior is improving, it is still necessary to test to confirm that the expected results, including failure causes with low probabilities of occurrence, were achieved.

Dispersions, Tolerances, and Tests

The engineering analysis and models used to develop the estimates of expected behavior are becoming quite sophisticated in their ability to represent complex phenomena in increasing levels of detail. The parameter values either used as input or predicted as output can be represented to many significant digits. Despite the precision of these parameter values that come from the models, design parameters are not single valued relative to part-to-part variation, test-to-test variation, or even time-to-time variation within the same test or flight. Parameter values are dispersed about a central value. Some of the characteristics that influence this parameter dispersion were illustrated earlier. The characteristics of the dispersion may allow modeling by a normal distribution or may require a more complex representation. Once the dispersion characteristics are understood, limits of the dispersion range may be defined to include an arbitrarily large fraction of the expected population. It then becomes possible to understand the influence of the dispersion range on the quality of the design.

Significant to this understanding is the relationship of the dispersion characteristics and the tolerance limits, and the influence that the relationship has on the ability to efficiently achieve a high-quality design. Tolerance limits define the range over which a parameter may vary and the system is still expected to perform within performance and without failure. Figure 3 illustrates two different examples of the relationship of dispersions and tolerance range. Also shown in the figure is the parameter value at which a failure would occur, and which is shown as a single value for simplicity.

Figure 3a shows a case in which the dispersion range is as large as the tolerance range. In any design, failure points are estimated with models and then the estimate is either confirmed or re-estimated based upon subsequent test or flight operating experience. The question arises as to how many tests or flights are necessary to achieve a degree of confidence that the failure point for the parameter is outside the dispersion range. Figure 4 illustrates the answer for a situation that is reasonably modeled by a binomial probability distribution. As can be seen, the answer is much more than is desirable for an efficient project. Multiply this single parameter illustration by the number of parameters in a propulsion system to obtain an appreciation of the effect in a typical system. Figure 3b shows a case in which the dispersion range is small relative to the tolerance range. For this case it now becomes possible to select a value of the parameter at which to test the element so that one test or flight operating experience can confirm that, for the parameter, the failure point is outside the dispersion range. Clearly the case illustrated in Fig. 3b is more desirable.

Mathematics of Changes and Failures

When the desired outcome of a test or flight, as noted by postevent evaluation, is not achieved it is necessary to prepare a design change to improve the situation and repeat the cycle illustrated in the previous text. The design change may be a

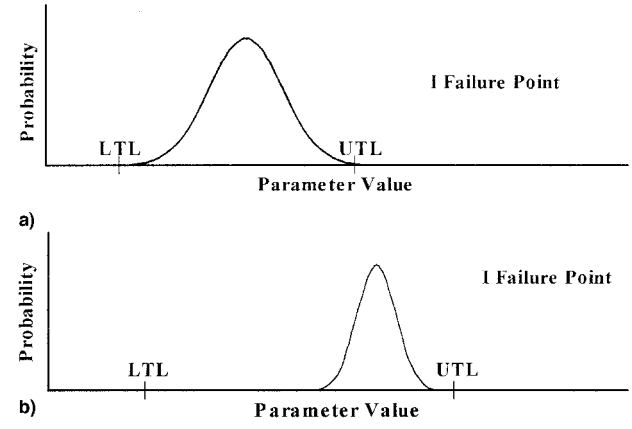

Fig. 3 Parameter value range relationship: a) large and b) small dispersion relative to tolerance.

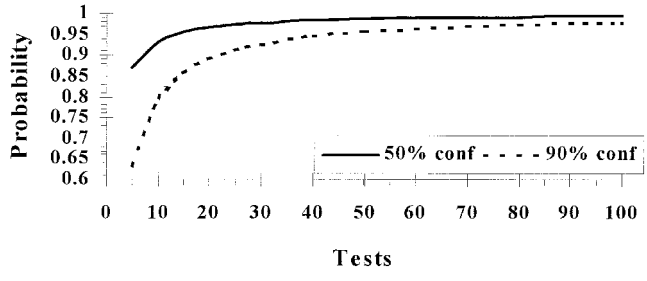

Fig. 4 Demonstration test requirements.

new part configuration, a new processing procedure, or a new set of flight rules. However, the price of any design change, whether made to eliminate or reduce the probability of occurrence of a failure cause, to accommodate vendor or part obsolescence, or to achieve performance improvements required or desired by mission requirements, because it is a change, is the introduction of new failure causes. As noted earlier, in fact, some of the new failure causes could have higher probabilities of occurrence than the eliminated failure cause or could lead to more severe system failures than were present prior to the change.

Therefore, for any element, the rate of change of the probability of failure has two trends. One trend is a decrease in the probability of failure that is expected to be correlated to the size of the probability of failure, because the higher probability cause is more likely to occur and be corrected. The other trend reflects the fact that in any change, new failure causes are introduced that provide a limit, even if small, below which the probability of failure cannot be reduced. Consequently, the rate of change of probability of failure relative to flight exposures is found to be

$$\frac{\partial \rho}{\partial n_e} = -\beta_e(\rho - \alpha_e) \tag{5}$$

As flight history proceeds and design changes are made to improve performance or reduce the probability of occurrence of a failure cause, a change could well increase the probability of failure. This would introduce a deviation to the expected probability of failure until the cause is diagnosed and a subsequent change is introduced, as observed in the empirical data. The initial probability of failure of each element should reflect the state of knowledge in the industry at the time the element was designed and introduced into service, whether that introduction occurred early or late in the overall flight history. Therefore, the initial probability of failure for an element should be higher for an element introduced early than for one

introduced later. Allowing for this undefined function of the overall flight history, Eq. (5) can be integrated to give

$$\rho = \alpha_e \left(1 + \frac{\delta_e}{\alpha_e} \Psi_t e^{-\beta_e n_e} \right) \tag{6}$$

where Ψ_t is the undefined function of the overall flight history. The initial value of the probability of failure when the element is introduced is at $n_e = 0$, which gives

$$\rho_{n_e=0} = \alpha_e \left[1 + \frac{\delta_e}{\alpha_e} \Psi_t \right] \tag{7}$$

and the asymptotic probability of failure at $n_e = \infty$ is

$$\rho_{n_e=\infty} = \alpha_e \tag{8}$$

The asymptotic factor, α_e, measures the lower limit on probability of failure because of the expected probability of failure of new failure causes introduced as a result of changes. The historical improvement factor with learning parameters δ_e and β_e describes the expected failure rate reduction path as the element evolves from first use. Clearly, a low asymptotic failure rate is dependent on the ability to detect failure causes and introduce corrective design changes that contain only a low probability of occurrence causes. The rate of reaching the asymptotic region is also dependent on the ability to detect failure causes.

Conclusions

This paper examined a different aspect of propulsion engineering—propulsion system quality—and characterized a high-quality propulsion system as one that has high reliability. A high-quality propulsion system has a high probability that the performance and operation of the system will result in the mission terminating within the defined tolerance limits. Because performance or operation of propulsion systems that results in the mission terminating outside of the defined tolerance limits is a failure, it is important to understand failures and their causes. Failure causes can be evaluated and classified by type, and the results of typing flight history failures show that most failures are in unrecognized modes and result from human error or noise, from which it is concluded that failures are when engineers learn how things really work. The expected more frequent occurrence of the higher likelihood failure cause(s) leads to correction and elimination, and to a lower probability of failure controlled by the more numerous lower frequency of occurrence causes. When the desired outcome of a test or flight is not achieved, a design change is prepared to improve the situation. The price of any design change, because it is a change, is the introduction of new failure causes. These new failure causes provide a limit, even if small, below which the probability of failure cannot be reduced. Numerous research opportunities to improve the quality of rocket propulsion systems are illustrated throughout this paper.

References

[1]Biggs, R. E., "Space Shuttle Main Engine, The First Ten Years," American Astronautical Society National Conference and Annual Meeting, Los Angeles, CA, Nov. 1989.

[2]McCarty, J. P., "A Critical Function Technique for Modeling Launch Vehicle Reliability," Ph.D. Dissertation, Dept. of Industrial and Systems Engineering, Univ. of Alabama, Huntsville, AL, 1996.

[3]Hasofer, A. M., and Lind, N. C., "Exact and Invariant Second Moment Code Format," *Journal of Engineering Mechanics Division*, Vol. 100, 1974, pp. 111–121.

[4]Brown, C. B., "A Fuzzy Safety Measure," *Journal of the Engineering Mechanics Division*, Vol. 105, 1979, pp. 855–872.