

Launch Vehicle Abort Analysis for Failures Leading to Loss of Control

John M. Hanson^{*}

NASA Marshall Space Flight Center, Huntsville, Alabama, 35812

Ashley D. Hill[†]

Dynamic Concepts, Inc., Huntsville, Alabama, 35806

Bernard B. Beard[‡]

ARES Corporation, Huntsville, Alabama, 35805

Launch vehicle ascent is a time of high risk for an onboard crew. There is a large fraction of possible failures for which time is of the essence and a successful abort is possible if the detection and action happens quickly enough. This paper focuses on abort determination based on data already available from the Guidance, Navigation, and Control system. This work is the result of failure analysis efforts performed during the Ares I launch vehicle development program. The two primary areas of focus are the derivation of abort triggers to ensure that abort occurs as quickly as possible when needed, but that false aborts are avoided, and evaluation of success in aborting off the failing launch vehicle.

^{*}Lead for Flight Mechanics Analysis and Design, Flight Mechanics and Analysis Division, Mail Code EV40, Marshall Space Flight Center, Huntsville, AL 35812, <john.m.hanson@nasa.gov>, AIAA Senior Member.

[†]Senior Engineer, Dynamic Concepts, Inc., 6700 Odyssey Drive, Ste. 202, Huntsville, AL 35806, <ahill@dynamic-concepts.com>.

[‡]Senior Consultant, ARES Corporation Tennessee Valley Office, 200 Sparkman Drive, Huntsville, AL 35805, <bbbeard@arescorporation.com>.

Nomenclature

$F(x)$ = cumulative distribution function for a generalized Pareto distribution

H_M = total window of interest (altitude, speed, etc.)

h = correlation height (or speed, etc.) for the data

P_M = desired false abort probability

p^* = probability setting for calculating the x^* trigger value

x = independent variable in a generalized Pareto distribution

x^* = desired setting for a trigger as calculated by a generalized Pareto distribution

μ = threshold setting for a generalized Pareto distribution

ξ = shape parameter for a generalized Pareto distribution

σ = scale parameter for a generalized Pareto distribution

Introduction

THERE are currently a number of potential humans-to-orbit developments being pursued. Several of these efforts involve putting a crew on a launch vehicle that was not specifically designed with crew launch in mind. Whether or not that is the case, it is desirable that the crew have abort capability to quickly escape a failing vehicle. The purpose of this paper is to provide knowledge gained during the development of the Ares I abort warning system for launch vehicle failures that lead to loss of control. The two primary areas of focus are the design of loss of control abort triggers that provide rapid warning in failure situations (without causing false aborts), and analysis of rapidly developing loss of control failure situations, to examine whether the crew is able to depart the vehicle prior to reaching vehicle failure limits or limits beyond which abort is not feasible.

Historical Launch Vehicle Failures and Crewed Launch Vehicle Abort Systems

A number of historic launch vehicle failures were ultimately a failure to control flight. Some examples include the first Pegasus XL, the first Delta III, the Conestoga launch vehicle, and the first Ariane 5 launch [1]–[3]. Other types of launch vehicle failures (such as engine failures) may lead to loss of control as one consequence [2]. Cases where the vehicle is under control but will not reach orbit do not require immediate abort (in a crewed flight), so the

crew could wait and abort later or wait until the engine shuts down due to the low performance condition rather than aborting immediately. Although a first flight of a new vehicle is typically the riskiest, things can and do go wrong on later flights. Examples of problems in ascent flight during crewed missions include the Challenger disaster, Apollo 12 being struck by lightning, Apollo 13 experiencing pogo problems (that nearly exceeded structural limits), and two Soyuz failures that made abort necessary [4].

Historical launch vehicles that carried humans have always provided for some means of abort or escape. Project Mercury had a rocket design to pull the capsule away from the launch vehicle [4]. It was jettisoned after the first 3 min and 50 s of flight. It used an attitude rate limit as the loss of control parameter for automatic abort. In Project Gemini, ejection seats were provided for the first 15,000 ft, and the de-orbit motors were used to separate the spacecraft from the rocket for later aborts (after engine shutdown). The Titan's nitrogen tetroxide and hydrazine propellants were considered to be much less of an explosive hazard than the kerosene and liquid oxygen used on the Mercury Redstone and Atlas [4]. Attitude rates were again monitored, but the abort modes were manual.

Project Apollo had a launch escape tower for pulling the capsule away, similar to the Mercury approach. On the Saturn V, the Emergency Detection System (EDS) provided for automatic abort if either excessive angular rates or loss of thrust on two first stage engines during boost were detected [4]. A constant angular rate limit for pitch and yaw followed by a switch to a different constant rate provided the trigger limits along with a higher rate limit for roll (mission success was significantly less sensitive to roll) [5]. The constant angular rate limit protected for certain rapid flight control loss situations such as actuator hardover [6]. First stage hardover was the limiting case (failed most rapidly), but the same rate limit was used in second stage flight. Due to limited computing capability, only a limited set of failure cases were examined in setting the limits, and limits were set sufficiently high that false aborts were believed to be unlikely.

On the Space Shuttle, potential loss of control situations due to failure did not allow for successful abort (the first four Shuttle flights had ejection seats for early ascent, but these were removed for subsequent flights). The Shuttle had to maintain control in any abort situation and continue to fly, or else the crew and spacecraft would be lost. This approach is inherently less safe than allowing for all the situations where the vehicle is no longer able to fly and the crew has a means of departing the failing vehicle.

The Soyuz spacecraft has a rocket for abort, which has been used twice as mentioned above: once for a pad failure in 1983 and earlier for a stage separation failure in 1975. The original automated system included detection

of loss of control, premature booster separation, combustion chamber loss of pressure, velocity decrease, and loss of thrust [4]. China's Shenzhou spacecraft also features a rocket escape system for abort capability in the atmosphere [4]. SpaceX's Dragon spacecraft includes a pusher-type rocket for aborting from a failing launch vehicle, which is not jettisoned, unlike the others already discussed that were jettisoned during ascent.

A number of future crewed launch vehicle concepts include the ability to depart a failing vehicle, so some kind of EDS-like system will be used for situations requiring immediate abort. At the same time, analysis capabilities have vastly increased since the Apollo days. Today, a much better capability exists (due to computing capacity) to design the abort limits to stay out of any no-failure, false-positive regions and also to ensure positive abort triggering as soon as possible when it is needed. The capability is also much better for evaluating the success of these triggers in allowing the crew to depart in a timely manner (with, for example, numerous Monte Carlo simulations). Likewise, onboard computing capabilities have vastly increased, so that the onboard logic can consider parameters that change during flight.

Ascent Trajectories and Aborts

Vertical lift-off trajectories (for all orbital launch vehicles) typically have a critical initial portion where the launch tower must be cleared. Then the vehicle pitches over to head downrange, with most of its speed relative to the rotating Earth still vertical. It picks up speed and builds dynamic pressure and drag as it heads downrange and is still in a relatively dense portion of the atmosphere. A rocket-propelled abort system is primarily valuable early in flight, for a number of reasons: the crew must be pulled (or pushed) away from a failing vehicle if it fails on or near the launch pad; pulling or pushing is also needed when drag is high; the spacecraft must be pulled or pushed high enough to allow it to successfully land or splash down; and explosions of launch vehicles are most dangerous to the crew when in the atmosphere. Abort at the pad requires a substantial total impulse to take the crew far from the failing vehicle. Loss of control failures tend to be particularly time-critical during the region of high dynamic pressure, since structural loads build up rapidly with increasing aerodynamic angles (as the vehicle starts to lose control). Abort at high drag conditions requires overcoming the drag and dynamic pressure effects to get the crew away from the (possibly still accelerating) vehicle with sufficient net relative acceleration. After this, drag and dynamic pressure fall off while speed continues to increase, as the atmospheric density continues its reduction. At some point, there is usually a staging event where booster motors or a first stage are jettisoned.

The rocket continues to increase its speed until reaching orbital insertion conditions, with potentially another staging event during later ascent. If the design of the escape rocket is such that it also functions as propulsion for the spacecraft later, it may be kept on board, but otherwise it is typically jettisoned since there is a significant payload to orbit penalty for keeping it longer than necessary. The highest vehicle accelerations tend to occur late in flight for each stage, and if there are loss of control failures, these times can be critical since the vehicle has the least inertia and so will have the highest angular accelerations. If the crew escape system is jettisoned during ascent to enhance payload capability, then the engines will have to be shut down before the crew can depart. Issues can arise in all parts of flight, so all regions of flight should be investigated.

Failures that Lead to Loss of Control

There are a large number of possible failures that might cause loss of control. For example, thrust vector control (TVC) problems may be caused by hardware issues with the actuators or power supplies, or software issues with improper commands being sent to the actuators. These problems could manifest themselves by actuators failing to hardover (maximum angle), or failing in place, or failing to the null position. Engine nozzle damage may lead to performance losses, control issues, or both. A recontact between the upper stage engine nozzle and the lower stage during stage separation might damage the nozzle, or some joint failure or nozzle burnthrough might cause nozzle damage. In many cases, the number of engines will influence whether a particular failure leads to the need for abort. A solid booster could suffer a case or joint burnthrough, where part of the exhaust escapes through a growing hole. Failure in the navigation system or an insufficiently tested software issue could lead to making bad flight control commands. A sizable fraction of other potential failures (e.g., catastrophic engine failure) that are not failures of flight control systems might also lead to loss of control, and possibly a sufficiently fast reaction might save the crew.

Goals and Metrics in Aborting from a Failing Launch Vehicle

Some failures can occur so rapidly that there will be insufficient time for the crew to escape no matter what abort determination is used, and other failures are so benign that warning is not needed. This paper focuses on failures where timely warning might save the crew. Some of these require immediate warning and abort; others allow for a bit more time.

Because the types of failures that can occur are so wide ranging, it is not possible to ensure that all of them are detected if the objective is to understand where every failure originated. However, the Guidance, Navigation, and

Control (GN&C) impacts of these various failures lead to the ability to predict the need for abort based on calculations using GN&C data. For the most part, this means that data that already exist on board can be used to recommend the need for abort. For example, if navigated attitude rates go beyond those ever expected without failure, then assuming a failure must have occurred allows for an abort recommendation without understanding the source of the failure. There will be plenty of time after the flight to diagnose the failure.

Abort trigger settings designed for rapid loss of control situations must inherently be a compromise. The trigger settings must be sufficiently high so that false aborts (or false positives) do not occur (when there is no failure but the attitude rates or angles go outside the abort setting), but at the same time, a higher trigger setting delays crew departure and increases the chance that the crew will not survive a failed situation. One may consider the following equation for the portion of flights that result in loss of crew:

$$\text{Loss of crew percentage} = \% \text{ where failure requires abort and triggers work} \times (1 - \text{abort effectiveness}) + \% \text{ where no failure occurred and there is a false positive} \times (1 - \text{abort effectiveness}) + \% \text{ where failure requires abort and there is a false negative}$$

A false negative means that an abort was required and the sensors have failed or otherwise did not detect the need for abort. Abort effectiveness is the fraction of abort cases where the crew survives. Subtracting it from 1 gives the fraction of cases where the crew is lost during abort for any reason, including not successfully escaping the vehicle, suffering system failures some time during the abort and landing process, and unsuccessful recovery after abort landing. Abort is considered to be an inherently risky operation, so that it is always desirable not to abort unless it is necessary (besides the fact that aborting when it is not necessary ruins the mission). A false negative requires that there is a failure occurring which requires abort. This is two failures deep: there is a failure that requires abort which this sensor should detect, and the sensor fails to detect it. Assuming that success is more likely than failure (usually much more likely for particular failure modes), the false negative likelihood is reduced by that factor. With GN&C attitude rates and attitude errors, the rates tend to get worse fast when losing control, and the main issue becomes getting the crew off in time. So the compromise is between false positive and abort effectiveness (higher trigger setting = lower false positive = reduced abort effectiveness) rather than between false positive and false negative. If the vehicle is losing control, the rates and errors get worse with time and even if the trigger is set higher it will be

passed, so it is not a false negative. But if it is higher, a structural or attitude rate limit might be passed and the crew does not survive (so abort effectiveness is decreased). This is much more likely than a false negative since GN&C sensors are redundant and highly reliable. Also, a false negative could be associated with sensor failure, meaning that a failure from some cause other than GN&C sensing occurred and abort is needed, but the sensor data are bad so the need is not detected. But if the GN&C sensors have failed prior to some other failure, the vehicle will have lost control anyway and abort is needed (again mitigating any false negative chance; abort would have occurred when the GN&C sensors failed). So the main compromise for loss of control triggers is between false positives and abort effectiveness:

$$\text{Loss of crew percentage} \cong \% \text{ where failure requires abort and triggers work} \times (1 - \text{abort effectiveness}) + \% \\ \text{where no failure occurred and there is a false positive} \times (1 - \text{abort effectiveness})$$

Since abort is risky, it is very important that false positives be kept very low but that the trigger settings are no higher than needed, in order to maximize abort effectiveness. Determining appropriate settings and measuring the effectiveness in getting the crew off in time is the focus of the rest of this paper.

A further metric, used later in this paper, is success in abort triggering, when there is a failure that would typically require abort. This metric is related to the ability to get the crew off prior to some demise criteria, and does not include whether the abort systems subsequently perform as needed for survival. This percentage is given by (where a failure occurs in all cases):

$$\text{Success percentage in abort triggering} = \% \text{ where the vehicle will meet its demise and the triggering enables the} \\ \text{crew to depart prior to demise} + \% \text{ where the vehicle does not meet its demise (whether or not abort was} \\ \text{recommended by the triggering)} - \% \text{ where the vehicle does not meet its demise but the abort triggering} \\ \text{recommended abort anyway.}$$

The last part of the equation does not necessarily result in loss of crew, but forces the crew into a risky situation that ends the mission unnecessarily.

The Ares/Orion System

The Constellation Program included the Ares I rocket with the Orion spacecraft mounted on top. A Launch Abort System (LAS) was to provide abort capability during First Stage operation and until the Upper Stage was successfully operating. In an abort, the LAS pulls the Crew Module (CM) away from the rest of the launch vehicle and the CM parachutes to safety. After LAS jettison early in Upper Stage flight, aborts called for shutting the Upper Stage engine down, followed by a normal separation where the Orion CM and Service Module separate from the Upper Stage. For most of the rest of ascent, in an abort situation, the CM again comes down on parachutes. If abort occurs late enough in ascent, Orion might have aborted to orbit.

If Ares was in the process of failing, and was able to detect the need for abort, then Ares would recommend abort to Orion. Before LAS jettison, if the abort situation was time-critical, Orion's computers would immediately initiate the abort sequence. After LAS jettison, if the abort situation was time-critical, Ares would notify Orion of the need to abort and would also shut the Upper Stage engine down. The Orion crew could also initiate abort at any time. In failure situations that were not time-critical, the crew would be involved in the decision. In general, loss of control situations happen quickly enough that the computers would make the decision.

Ares I Analysis for Failures Resulting in Loss of Control

During the Ares I design effort, high-fidelity simulation was used to model the launch vehicle dynamics as it ascended to orbit. Simulating the ascent in Monte Carlo simulation with all vehicle and environmental dispersions included, for various launch times and vehicle models, provided knowledge of the statistics of vehicle attitude rates and attitude angles when there are no failures [7] [8]. Then, designing abort trigger settings to avoid false aborts (false positives) involves setting the trigger levels above the levels that are seen without failures. Care must also be taken to ensure that some events such as strong wind gust, that may not have been in the simulation, do not lead to abort. After designing the triggers, additional Monte Carlo simulations where a particular type of failure occurs at random times on every flight (again with vehicle and environmental dispersions included) allows for evaluation of abort effectiveness. For each sample trajectory, the failure occurs, the triggers determine when the abort happens, and the situation is evaluated at the actual time of Orion departure from the failing Ares to determine whether any conditions were exceeded that allow for Orion to safely depart. The rest of the paper develops the trigger design

methods, provides some lessons learned concerning trigger design during the Ares I effort, and describes the abort effectiveness evaluation.

Loss of Control Abort Triggers and Selection of Trigger Levels

Trigger Possibilities

Navigation and flight control data are available for the purpose of deciding whether an abort is necessary. In particular, the information that flight control uses in order to control to the guidance commands are available at high rates and are appropriate for this purpose. Attitude rates, attitude rate errors (difference between the guidance command and the actual rate), and attitude errors (between the guidance commands and the actuals) are all parameters that could be used to determine the need for an abort. Each of these values may be further separated into body pitch, yaw, and roll channels.

An additional trigger possibility for detecting loss of control would be if a measurement of engine actuator angles or a measurement of TVC health is available with high reliability. The actuator angle could be compared against the expected actuator angle from the command and the modeled actuator dynamics. Trigger settings for the angle error could be set in the same manner as for other triggers, as discussed below. Whether or not actuator error or TVC health becomes part of the baseline triggers would depend on how much they contribute to abort success, and on how much additional cost and complexity would be associated with the sensors. Since the events being protected for are not that likely, the sensors would have to be highly reliable, assuming that failure of the sensors would lead to a false abort. These actuator triggers will not be investigated further in this paper.

Generating the Abort Trigger Settings

The abort triggers should be set at a level higher than any seen in no-failure dispersed simulation, with a setting sufficiently high that a false abort recommendation is very unlikely. Early on, the triggers for Ares were set according to a number of standard deviations that assumed the data are Gaussian, targeting for a desired false abort rate. Trigger settings were done at a number of different points along the trajectory, calculating the setting using all the dispersion results to generate the data. Dispersion Monte Carlo runs included winter and summer trajectories, both ends of the launch window, and sluggish and sporty vehicle models. The Monte Carlo simulations included light to moderate wind gusts, so a gust increment was added to the calculated abort trigger to ensure that abort would not be recommended simply due to a severe wind gust. These gust increments were calculated using flexible vehicle

simulations with tuned gust frequencies at maximum dynamic pressure, where the tuning was done to challenge the vehicle the most. The gust effect was scaled by dynamic pressure to apply to different flight phases.

Assuming a Gaussian (and setting the abort trigger to a 10^{-6} level) has led to experiencing false aborts in Monte Carlo runs where the random seeds are changed and there is no other change in the modeling (with 2000 samples). So clearly the trigger setting was not sufficiently high to ensure no false aborts. Most of the Monte Carlo data fail the Anderson-Darling [9] test for normality. It makes sense that these values do not behave as Gaussians, since they are highly nonlinear, closed-loop flight control responses to the various perturbations that are happening in the trajectory.

As a result of the issue with non-Gaussian behavior, a new approach was adopted, since the need here is to characterize the tail of the distribution, not the mean and standard deviation. Extreme value theory [10] deals with the tails of distributions and has been applied in many fields, including hydrology, meteorology, insurance, reliability, gerontology, economics, and evolutionary biology. There are two popular approaches: “block maxima” that uses a “generalized extreme value distribution” and “peak over threshold” that uses a “generalized Pareto distribution (GPD).” The idea is that once the distribution is known, a low false abort probability may be set, and the location on the distribution corresponding to that false abort probability determined. The approach taken here uses a peak-over-threshold method and fits a GPD to the data. Reference 10 provides some development of the GPD approach and justification for its use.

In the peak-over-threshold method, a threshold is chosen that is above most of the data (for example at the 95% level). The parameter that is calculated is the difference between the measured value and the threshold setting. The approach is to find σ and ξ to best fit the experimental cumulative distribution function (CDF) of the data that are above the threshold, given by

$$F(x) = 1 - [1 + \xi((x - \mu)/\sigma)]^{-1/\xi}, \quad (1)$$

where x is the experimental sample and $x - \mu$ is the difference between the value and the threshold for cases where x exceeds the threshold. Then to find the value x^* associated with a certain low target probability p^* (e.g. 10^{-6}), use

$$x^* = \mu + (\sigma/\xi) [(p^*)^{-\xi} - 1] . \quad (2)$$

The free statistics package “R” [The R Project for Statistical Computing, <http://www.r-project.org/> (retrieved 3 Feb. 2011)] has a function “fitgpd” for the purpose of determining σ and ξ . Threshold selection involves a tradeoff between bias and uncertainty. Too low a threshold includes many points “not in the tail” that bias the fit. Too high a threshold increases the uncertainty because too few points are used in the fit. Studies for Ares I indicated that using the top 5% of Monte Carlo data is adequate (e.g., take the top 100 out of 2000 or top 1000 out of 20,000 and perform the prescribed procedure). Much of the time, the fit works very well. An example is in Fig. 1, a graph of the experimental distribution, the best-fit GPD, and GPD with zero shape parameter, for a case where the Gaussian fit did not work well. Even if ξ is set to zero and σ only is fit, the fit works very well. The data, best-fit GPD, and GPD with zero shape parameter all fall nearly on top of each other in the figure. On the other hand, sometimes there are one or two data points that lie well beyond the rest. In this case, the fit does not work as well, and goodness of fit tests show that the fit is poor. Also, due to the random and nonlinear nature of the data, using a fitted ξ leads to an abort trigger with significant spikes (see Fig. 2). In Fig. 2, the trigger setting for a 10^{-6} probability of false abort is graphed at each altitude, using a GPD fit. ξ is the shape parameter, and changes the shape of the distribution significantly when there are outliers. For this reason, ξ was set to zero for defining the distributions used to generate the abort triggers. Use of $\xi = 0$ actually yields a Gumbel distribution for the extreme values. This choice is not as accurate a fit of the individual tail data points in each case, but avoids the large variation in shape that causes the spikes. The Gumbel distribution has been used for extreme value calculations of flood heights, extreme wind velocities, and runway roughness.

Notice how many altitude data points there are in the trigger settings calculated in Fig. 2. This was done because flight control parameters are continuous functions of altitude (or speed, etc.). However, using a GPD fit at every altitude point would set the trigger too low by a factor that depends on the correlation of successive altitude slices (i.e., the false abort probability might be 10^{-6} at each small altitude slice, but add up to much more than that over all altitudes). So the correlation for each parameter must be estimated. The general approach to this calculation (as taken in the Ares effort) is to calculate p^* by using

$$p^* \sim P_M(h/H_M) , \quad (3)$$

where P_M may be 1×10^{-6} . If 2000 ft is the correlation height and 200,000 ft is the total altitude window, then p^* will be the desired false abort probability divided by 100. That is, in order to get a total false abort probability of 1×10^{-6} , if the total window is divided into 100 sections, the probability in each section must be $1/100$ of the total. To determine the correlation height, calculate the correlation coefficient of the value of interest (say, pitch rate) measured at one altitude versus another altitude. When the correlation drops below some value, say 0.5 (picking a higher value will be more conservative), that determines the correlation height.

Even with the GPD approach, portions of flight will be problematic due to the nature of the data (see Fig. 3). The top curve is the GPD trigger setting for a 1×10^{-6} probability. The Monte Carlo data are shown, with many of the curves lying on top of each other (2000 samples are included). The data are obviously non-Gaussian. For example, near 120,000 ft there is a single outlier at about a Gaussian 6 sigma level (the 3 sigma level is at the top of the mass of data and the outlier is at least twice as high), with a Gaussian probability of about 1×10^{-9} , with only 2000 samples taken. The distinct outliers in one case exceed the GPD fit. So some additional margin must be added in these cases. Figure 4 shows the development of final trigger settings for Ares I analysis. The maximum of the Monte Carlo data is shown as the lowest curve. The heavy solid curve is what the GPD approach generates. For most of flight, this curve is well above the maximum Monte Carlo data, but there are several issues. First, there are regions where the candidate trigger setting is very close to the maximum value seen in simulation. Even worse, some cases resulted in the candidate trigger level being less than the worst-case value from the simulation. Cases like this can result when the data are ill behaved, such as having a single outlier from the rest of the closely bunched data. Related to that is the fact that, when the data are fairly benign (e.g., in Fig. 4, at higher altitudes, the maximum attitude rate seen is ≈ 0.5 deg/s, not a large value), the trigger for a low false abort rate would be set to a low value when that low value does not in fact represent even a slight hazard. Some margin on top of the maximum value is always desirable. A final issue is that, when a value in the data occurred at a particular altitude, what if in flight that value occurred at a slightly different altitude? An abort would be undesirable. For example, the drop in the maximum data curve just above 60,000 ft could be slightly later, and would thus exceed the heavy solid trigger curve.

With these needs in mind, two adjustments were made to the triggers. First, on the vertical axis, a minimum separation was enforced (based on engineering judgment) between the trigger setting and the maximum simulated value (generating the Trig Adj 1 curve). Second, on the horizontal axis, the trigger was not allowed to go below the value required by the vertical axis adjustment for some range on the x -axis (Trig Adj 2, again based on engineering

judgment; the data used for calculation is not the original trigger setting but what the Trig Adj 1 limit would be without the original trigger). Finally, the top-most curve shows the final trigger setting, resulting from adding the wind gust effect as described earlier.

Pitch and yaw channels may be combined to form a single root sum square (RSS) trigger for faster response. This method certainly applies for axially symmetric vehicles and may be applicable to other vehicles as well. It may not be as useful in regions of flight where pitch and yaw responses are clearly not the same (e.g., a yaw maneuver just after lift-off to avoid the tower or a pitch maneuver just after staging). Figure 5 shows pitch and yaw axes with hypothetical pitch and yaw trigger settings, and an RSS setting. The drawing demonstrates that a combination of pitch and yaw error or rate will be detected faster (in many cases) with an RSS trigger as opposed to separate pitch/yaw triggers. Any motion that is in a direction other than pure pitch or pure yaw will leave the RSS region faster than it will cross the individual pitch and yaw trigger settings. The same GPD approach with trigger adjustments generates the trigger settings.

Similarly, attitude rate may be combined with attitude error, or attitude rate error may be combined with attitude error in an approach dubbed “phase plane” triggers. Figure 6 shows an example that demonstrates how this trigger method can lead to faster detection of a problem. The small circles show Monte Carlo no-failure dispersed data. Hypothetical individual attitude error and attitude rate triggers are shown, along with an ellipse that combines them. When there is a very rapid attitude rate buildup, such as would occur with an actuator hardover (at least for a single-engine launch vehicle), the phase plane trigger does not offer an advantage unless the starting condition before failure has significant attitude error. If the starting condition includes significant attitude error, then motion primarily in the rate direction will cross the ellipse boundary faster than the single parameter boundary. When the rates and attitude errors build up over time, such as might happen with an actuator that has failed in place (so at the time it failed, the vehicle was in trimmed flight) or a multi-engine vehicle with an actuator hardover, the phase plane trigger can be substantially faster.

The data used to generate the phase plane triggers must be well behaved (somewhat like a bell-shaped curve of probability density) in order to generate a reasonable statistical ellipse. (Procedures for generating the appropriate ellipse will be discussed shortly.) Figure 7 shows how the data may be reasonable for this purpose in some flight phases and not in others. The figure shows a locus of rate versus attitude error points, for dispersed flight with no failures, for three different times in flight. The data at 24 s and at 100 s are not good for designing a probability

ellipse, whereas the data at 55 s are. An ellipse could be drawn around the 24-s data, but determining the statistical likelihood of exceeding the ellipse setting would be problematic. The data at 24 s are skewed in part because of varying vehicle pitch-over profiles during early ascent. The data at 100 s are well behaved, but there is not much to be gained by using an extremely eccentric ellipse as compared to the individual triggers.

A similar approach to that used for the individual triggers may be used to design the phase plane triggers by converting the two-dimensional phase plane parameters into a normalized radius (see Fig. 8). The procedure involves first normalizing each parameter by its standard deviation. The second step is to generate a radius (square root of the sum of squares). Third, generate the trigger setting for this one-dimensional radius using the GPD method as before. Next, denormalize to get the individual axes back. The center of the phase plane ellipse is at the mean of the two sets of data. Finally, apply the trigger adjustments (as in Fig. 4) to the individual axes (see Fig. 9). In Fig. 9, each individual phase plane parameter (a) and (b) is adjusted as in Fig. 4 and the result is applied to the axis of the phase plane ellipse (c). Use of the phase plane trigger data (for Ares I) was well behaved in the high dynamic pressure region, and time is critical there. On the other hand, in other regions of flight, the data were not behaved as well (Fig. 7) and the regions were more benign for abort timing, so the phase plane triggers were not used. Of the three samples in Fig. 7, case (b) only supports use of phase plane triggers (fitting an ellipse to the data).

The methods for designing the various abort triggers have now been defined. The general philosophy used was that the desire is to get the crew off (or trigger abort) as quickly as possible once it is known that some kind of failure is making the vehicle lose control. Of course, the results are only as good as the input models and uncertainties used in the Monte Carlo simulation. If any of the models describing the vehicle in the simulation are incorrect, the trigger settings might be inappropriate and could cause false aborts in actual flight, or delays in aborting beyond what is necessary, or both. The uncertainty in knowing whether the models truly represent the vehicle is one reason for adding the engineering judgment increments mentioned earlier. Certainly test flights should be used to validate the model correctness (including such items as noise in the data) to the extent possible prior to using these trigger settings to determine abort. There may be portions of flight where the trigger settings would still allow quite a bit of time before vehicle demise (e.g., when dynamic pressure is very low and expected flight does not require much maneuvering). In these cases, it may be desirable to increase the trigger settings in order to allow for temporary issues (from mismodeling or for whatever reason) to clear up.

Some Lessons Learned from Ares I for Abort Trigger Design

This section lists some lessons learned from Ares I concerning abort trigger design, which should be revisited for future applications in order to determine whether they apply to a new vehicle. It was seen that roll attitude error does not require abort; the vehicle can generally proceed to orbit successfully if the roll angle is not successfully controlled. Similarly, significant roll rates can be sustained without requiring abort. Separation and jettison events are not very sensitive to roll rates. The pitch and yaw flight control should be able to maintain control to the guidance commands at fairly substantial roll rates. These statements are generally true for axially symmetric launch vehicles; as mentioned earlier, this lesson was known to the Saturn abort trigger designers. A reasonable value for a roll rate abort trigger might be somewhat below a value for which abort is no longer feasible; i.e., the abort and successful recovery of the crew have their own roll rate limits. With a single engine, a separate reaction control system or other method might be used for roll control, allowing for separate consideration of the roll channel. In multiple engine cases where an engine failure leads to use of only one engine, some roll rate may build up. A roll rate trigger may be unnecessary if there are no credible failure modes identified where roll rate can build up to unacceptable values while pitch and yaw are still controlled.

Attitude rates directly from the navigation, and rate errors and attitude errors from the flight control processing, examined before going through the flight control filters, are better than the flight control filtered data. The filtered data are smoother (lower peaks), but the quickness of the response with the unfiltered data more than makes up for the increase in abort trigger setting necessary with the data taken before the filter. False abort probability was kept at the same level in both cases, so for the unfiltered abort triggers, the trigger level was set higher to compensate for the higher peaks in the data. Most of the sensor noise has already been removed by the navigation system prior to output, so the data before the flight control filter is not very noisy; it mainly has higher peaks that are real, and effects are seen more quickly. A small magnitude of noise is seen in the data, but the real peaks are much larger. This was the conclusion in Ares I analysis, and resulted in increased abort success (getting the crew off prior to demise) relative to flight control filtered data. The improvement in detection time varies depending on the particular failure dynamics of each case and the characteristics of the particular filter. For any new vehicle, raw navigation data should be compared to data with filters tailored for the abort application, so that the best result determines the triggers used.

The trigger values will be a nonlinear function of flight condition (altitude, speed, or some other independent variable), since the dispersed values of each parameter being measured vary during flight, and setting the triggers to

the worst case would unnecessarily delay abort at other conditions. Trigger settings for the lift-off region (prior to tower clearance) should be designed separately from the rest of first stage ascent, since this region contains several unique qualities, including flight control response to stabilize the vehicle once the nozzles clear the launch platform, a nearby tower that is an additional demise condition, and a potential steering maneuver to enhance tower clearance in the presence of ground winds. Since this is a short time period, time or altitude may be used as the independent variable.

For first stage ascent, time is not typically a good choice of independent variable since it does not vary with vehicle dispersions. Use of altitude should provide better trigger behavior than time, since a slower vehicle reaches the same altitude later, so maximum dynamic pressure and guidance commands based on altitude would come later as well. Typically, open-loop guidance is based on altitude or speed rather than on time. The trigger could be based on the same parameter that the guidance uses. Speed may not be monotonic late in first stage ascent, which would favor using altitude.

Since staging times vary, and since typically closed-loop guidance commands a steering maneuver once the upper stage flight begins, it is best to base the parameters after staging on an independent variable that is a delta since staging. Delta altitude since staging is a good choice since it is monotonic whereas speed might not be. Use of a delta value makes the dispersed data used to generate the abort triggers more consistent.

As orbital velocity is approached, use of inertial speed as the independent parameter makes sense since the target shutdown condition is based on inertial velocity. Since trigger timing is critical near the end of flight when the vehicle is lightest and has the most forward center of mass (failures that cause loss of control will yield the fastest rising attitude rates), it is important to focus on conditions of interest near orbital injection, and inertial velocity does that. A failure that leads to significant attitude rates, such as an actuator hardover near orbit insertion, may cause problems for spacecraft separation. Even though the abort may be triggered rapidly, attitude rates may be large by the time the engine has shut down due to the substantial impulse imparted by the engine during the shutdown process. Spacecraft that keep the launch abort system on for all of ascent will be able to start the abort process before the engine shutdown is complete, but will still have larger attitude rates near the time of orbit injection as compared to earlier times.

If integral attitude control is used, even though it is designed not to be saturated, it is advisable to have logic that increases the attitude error abort trigger setting should the integral control become saturated. This is because the

saturation generally does not lead to the need to abort, but causes a larger than normal attitude error that typically takes some time to be removed. Similar possibilities should be investigated for other flight control methods.

Generally, when not using the phase plane triggers, the most effective triggers are attitude rate errors, followed by attitude rates. Attitude errors catch the fewest failures, but they are sometimes the first triggers to be exceeded for cases that slowly drift from the desired attitude.

Evaluation of Abort Effectiveness

Suppose that the abort triggers have been determined using the procedures described above. Then failures may be simulated and the ability of the system to abort prior to vehicle demise may be determined. If the demise criteria are exceeded prior to departure, it is assumed the crew would be lost. The lost cases all contribute to reducing the abort effectiveness. One obvious demise criterion would be if a vehicle structural load indicator is exceeded that indicates the structure is failing. Some additional criteria may be added if there is further time delay before this failure becomes catastrophic to the crew. Note that a simple indicator such as dynamic pressure times angle of attack (to represent the bending load) is not accurate when there are loss of control failures, since this limit incorrectly assumes the vehicle is in trimmed flight. A more sophisticated indicator that includes the effects of the actuator angles and attitude accelerations should be used. This will generally result in a smaller load than would exist in trimmed flight (at the same angle of attack), thereby allowing more time for abort. Of course, if abort success was determined to be adequate using a simple indicator such as dynamic pressure times angle of attack, the more detailed indicator might not be needed. For abort analyses, since the vehicle does not need to survive and continue flying, all margins may be used up before the limit is exceeded.

Another possible demise criterion would be if a structural load indicator is exceeded that means the spacecraft will be unable to abort (e.g., due to structural reaction to the release of bending moment that occurs just after abort). Again, a dynamic pressure times angle of attack indicator is only valid for trimmed flight, so a higher fidelity indicator is desirable. Another problem with spacecraft departure that could be included as a demise condition is recontact during the abort separation. There may also be limits to the abort stability and maneuvering capability that could be implemented as attitude rate limits, attitude error limits, or others.

Some other potential demise criteria are possible. Certainly the launch tower must be avoided. If altitude rate becomes negative in parts of flight where it should be positive, that might be considered a demise condition. For this and the attitude error limit, potentially abort would still be possible, but there are certainly values of attitude error

and trajectory error for which it is very desirable to have the crew gone. For example, a negative altitude rate means the launch escape system will accelerate the crew towards the ground, significantly reducing the time available for preparation for landing including parachute deployment. A related but alternative criteria might be the time available before a certain altitude is reached during the descent, which must be sufficient to prepare for landing. Another possible demise indicator is if crew physiological limits are exceeded. Additional demise criteria are possible depending on the application.

Failure modes that could be modeled include actuator failure. One potential actuator failure is TVC fail to hardover, using the modeled TVC dynamics (or alternate failure dynamics), assuming either a hardware or software failure. This could be one or both (e.g., pitch and yaw) actuators. One could also have a TVC fail in place from a hardware failure, which means the starting condition after the failure would be trimmed flight. Finally, there could be a TVC fail to null from a hardware failure. There could be engine nozzle failures of various kinds that might lead to loss of a percentage of thrust and an additional side thrust. If the nozzle fails forward of the actuator attach points, it would lead directly to loss of control. For solid boosters, another failure mode is some kind of case or joint burn-through that grows with time. This would lead to loss of some thrust and a side force as well (and potentially structural failure later if the hole gets big enough) [11]. A case burst (rapid explosion) is less likely to be a failure that can be analyzed using the methods discussed here. There could be a reaction control system failure (full or partial, on or off). Since main engines are generally providing at least pitch and yaw control and sometimes three-axis control, these failures may not yield abort situations. Stage separation recontact is another possibility, modeled by, for example, an impulse to the nozzle transferring impulse to one or more actuators, damage to the engine nozzle, and loss of thrust with added side thrust. The result could be rapid loss of control (if the actuator is physically damaged or if the force exceeds actuator capabilities), inability to achieve orbit due to performance issues, or just a decreased likelihood of reaching orbit, depending on the severity [12]. Other failures could be modeled, or it could be determined that these modeled failures more or less capture the range of things that can happen. If there are multiple engines, some of these failure types may not have as severe consequences as when there is only one engine providing the flight control. On the other hand, sometimes a slow divergence in control at high dynamic pressure is more challenging for abort triggering and timing than a rapid divergence (which shows up quickly as a rate violation before the angles get large).

Statistical results may be obtained for each type of failure by running a Monte Carlo simulation where the failure occurs at random times. For failures late in flight, picking a time based on some parameter (such as guidance-calculated time to go) that runs all the way to main engine shutdown is important, since attitude rates at the time of crew departure will be highest near orbital insertion when the vehicle has the least mass and the most forward center of mass. It may be advisable to run a Monte Carlo simulation for the lift-off region of flight separately in order to capture good statistical results for failures in this region. Each Monte Carlo simulation should have all the regular parameters (winds, engine performance, other vehicle and environmental uncertainties) varying randomly so that good statistics may be obtained for what happens when there is a failure. It should be noted that this procedure will probably not capture a case where the failure happens at the time of maximum dynamic pressure times total angle of attack (considered a key load indicator), since that would require the failure to randomly occur on the trajectory with the worst load at the worst time.

Once the failure occurs, the simulation will model the dynamic effects. When the first trigger is passed, logic in the code determines the time delay until the crew departs. This time includes time to confirm that there is indeed a problem, decide to command the abort (assumed an automated decision for loss of control situations), and for the abort motor to ramp up in thrust. When the abort motor is no longer present, the engine shutdown begins after a delay and the shutdown must be complete before the crew can depart. Then the question to be answered is whether any of the demise criteria are exceeded before the crew leaves. If the abort motor is not jettisoned, successful separation clearance would be the criterion of most obvious interest, followed by successful attitude control.

Success in abort triggering was defined previously. It includes escaping the vehicle prior to demise (contributing to abort effectiveness), as well as cases where the vehicle did not experience demise in this particular failure scenario. The vast majority of these latter cases should be ones where the failure was sufficiently benign that an abort was not needed and no abort triggers were exceeded. It is important to capture the cases where there was no vehicle demise, to understand how challenging each failure type is, and to capture an understanding of the appropriateness of the trigger settings. Cases where abort did occur but the vehicle did not experience demise should be investigated, should any of these cases occur, since these cases subtract from the success of the abort triggering. In at least some of these, for example, actuator fail in place late in flight, stage separation recontact with minor damage, or actuator failure in place just prior to staging, continuation of the flight may be possible and the abort triggers should be adjusted. These cases are not, strictly speaking, false positives, since a failure did occur that

normally requires abort. However, for many failure types, there will come a time prior to orbit insertion (or staging) where the failure is sufficiently benign in its effects that the mission could proceed. A further potential step could be to simulate the relative motion after crew departure, the crew recovery dynamics, and the failing vehicle dynamics after departure in order to investigate the safety of the crew as it gets away from the vehicle, but that is beyond the scope of this paper.

Some sample methods for understanding results of Monte Carlo failure simulations are shown in Figs. 10–13. In each of these simulations, a single failure type was modeled, with random time of failure, during a particular phase of flight (e.g., lift-off region, first stage ascent, upper stage prior to launch abort system jettison, and upper stage after launch abort system jettison). Figure 10 shows an example of resulting success fraction for the time period between clearing the tower and 30 seconds into flight, with the different shades corresponding to different types of failure and the different groups corresponding to different combinations of abort triggers. The first three failures are different types of actuator (thrust vector control) failure, followed by two different models for a solid motor burn through that grows with time, and a nozzle joint failure.

Figure 11 is a typical graph of time available for abort, for a particular failure type. Zero on the y-axis corresponds to the time of crew departure from the launch vehicle, after the trigger settings are exceeded and after the time delay for data latencies and for the launch abort system rocket to come up to sufficient thrust. The time is how much time there is after zero prior to vehicle demise as defined by the various demise criteria. The shades of gray and black correspond to different reasons for demise. Negative times should be investigated since they correspond to loss of crew. Figure 12 shows the value of a particular load indicator from the various Monte Carlo samples, evaluated at each vehicle location at the time of abort. Zero is above the nose of the vehicle and the aft end is at about station 4000. This load indicator is normalized so that a value >1 means the limits are exceeded.

Finally, Fig. 13 shows the value of attitude rate for failures that occur during upper stage flight after the launch abort system is jettisoned, for an actuator hardover failure. In these cases, the main engine is shut down, and the issue is whether the spacecraft can be successfully separated at the rates that result. Rates are not surprisingly highest late in flight when the vehicle mass and inertia are lowest. The different shades of gray are for use of different abort triggers. Each case of modified abort triggering must be evaluated in a separate Monte Carlo simulation since the dynamics of the vehicle during engine shutdown following the trigger exceedance are modeled. Use of a different trigger means the engine would be shut down at a different time, resulting in a different rate after

shutdown. In the case of Ares I, two sets of triggers were evaluated: a baseline set that included attitude error, attitude rate, and attitude rate error, and an augmented set that added actuator position error to the baseline set. The two failure scenarios considered—actuator hardover and actuator fail in place—resulted in four Monte Carlo runs.

What should be done if the abort success percentages are not satisfactory? There may be cases where, if the spacecraft waits until it is sure a failure is happening, it will not be able to perform an abort. It could be argued that the right thing to do is to abort prior to reaching any spacecraft abort limits. Rather than setting abort triggers based on spacecraft abort limits, this paper set the triggers based on knowing that there is a failure situation. The problem with determining the need to abort based on spacecraft abort limitations is that the false abort probability increases and aborts could occur when there is not a problem. Another problem with setting triggers based on spacecraft abort capability limitations is that all triggers would be near the limits of spacecraft capability, putting stress on the abort capability. In regions where there is margin between knowing that a failure occurred and these limits, it would mean staying on a failing vehicle longer than necessary, with the possibility that catastrophic launch vehicle events could occur during the delay. The approach in this paper assumes that the driving need is to be sure something is going wrong and not to wait any longer than necessary once failure is known to be occurring. If the abort success percentages are not satisfactory in a particular region of flight, there are many potential ways to address the problem, including compromise between abort success and false abort probability (which would reduce the trigger setting in a particular region), investigating additional abort triggers or failure sensors, working on the abort time delay, and changing certain aspects of the launch vehicle or spacecraft design to be more robust.

Conclusions

This paper describes how to generate abort triggers for failures that lead to loss of control in such a manner as to avoid false aborts while at the same time maximizing the time available for the crew to depart. This process is quite a bit more complicated than simply choosing a value higher than the no-failure simulated results. The triggers use navigated attitude rates, along with attitude errors and attitude rate errors from the flight control. Some manifestations of these triggers combine them into a single, more effective trigger. An extreme value distribution was used to generate abort trigger settings at the desired probability level, with care taken to account for correlation between adjacent regions of flight. Then, recognizing difficulties with very non-Gaussian behavior and the need to protect for movement of features to slightly different regions of flight, engineering judgment increments were added so that at least a minimum separation between the trigger settings and all no-failure dispersed data was achieved.

Finally a wind gust increment was added so that abort would not be commanded due to a severe gust. During the development of these abort triggers, a number of lessons were learned about which triggers are more effective than others; these lessons are included in the paper.

The paper also describes how to use the trigger settings to generate abort success measures, which could lead to additional refinement of the trigger settings. Failures that may lead to loss of flight control are modeled in simulation. These failures are simulated in the presence of flight dispersions in order to evaluate the success of getting the crew off prior to vehicle demise. Demise may be defined in different ways, for example, striking the launch tower, exceeding an attitude rate limit, or exceeding vehicle structural load indicator limits. The triggers are used along with an appropriate time delay to determine whether the crew had time to abort prior to reaching the defined limits. Use of the approaches developed here should provide for a higher probability of abort success as compared to more simple approaches that were used in the past when less computing power was available for either analysis or onboard computation.

References

- [1] Harland, D. M., and Lorenz, R. D., *Space Systems Failures: Disasters and Rescues of Satellites, Rockets and Space Probes*, Springer-Praxis Publishing, Chichester, UK, May 2005. Pegasus XL p. 102, Delta III p.151, Ariane 5 p. 146.
- [2] Isakowitz, S. J., Hopkins, J. P., Jr., and Hopkins, J. B., *International Reference Guide to Space Launch Systems*, 3rd ed., American Institute of Aeronautics and Astronautics, 1999. Some failure listings and descriptions Ariane p. 21 and 37, Atlas p. 67, Delta p. 111, Proton p. 293, Soyuz pp. 376–377, Pegasus p. 271 and 279.
- [3] Hill, A. D., “An Investigation of the Flight Vibrations Experienced During the METEOR Mission of the Conestoga Launch Vehicle,” Master’s Thesis, The University of Alabama in Huntsville, Huntsville, AL, 1996, available at the UAH Library and from University Films International, <http://disexpress.umi.com/dxweb>, (retrieved October 10, 2011).
- [4] Shayler, D., *Space Rescue: Ensuring the Safety of Manned Spaceflight*, Springer-Praxis Publishing, Chichester, UK, Nov. 2008, Chapters 5 and 6.

- [5] “Apollo-Saturn Emergency Detection System Description, AS-500 Series Vehicles,” MSFC No. III-5-509-7, IBM No. 66-966-0022, prepared under contract NAS 8-14000 by International Business Machines Corporation, Huntsville, AL, July 1, 1967, available from the NASA Technical Reports Server at http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19770076380_1977076380.pdf, (retrieved October 10, 2011).
- [6] Hyle, C. T., Foggatt, C. E., and Weber, B. D., “Apollo Experience Report – Abort Planning,” NASA TN D-6847, June 1972.
- [7] Hanson, J. M., and Hall, C. E., “Learning about Ares I from Monte Carlo Simulation,” AIAA-2008-6622, AIAA Guidance, Navigation, and Control Conference, Honolulu, HI, August 2008.
- [8] Hanson, J. M., and Beard, B. B., “Applying Monte Carlo Simulation to Launch Vehicle Design and Requirements Analysis,” NASA/TP—2010–216447, September 2010.
- [9] *National Institute of Standards and Technology/SEMATECH e-Handbook of Statistical Methods*, <http://www.itl.nist.gov/div898/handbook/eda/section3/eda35e.htm>, (retrieved October 10, 2011).
- [10] Coles, S., *An Introduction to Statistical Modeling of Extreme Values*, Ch. 4, pp. 74–91, Springer-Verlag, London, 2001.
- [11] Luchinsky, D. G., Osipov, V. V., Smelyanskiy, V. N., Patterson-Hine, A., Hayashida, B., Watson, M., McMillin, J., Shook, D., Johnson, M., and Hyde, S., “Model-Based Diagnostics and Prognostics for Solid Rocket Motors,” Annual Conference of the Prognostics and Health Management Society 2009, San Diego, CA, Sept. 2009.
- [12] Luchinsky, D. G., Hafiychuk, V., Kulikov, I., Smelyanskiy, V., Patterson-Hine, A., Hanson, J., and Hill, A., “Stage Separation Failure: Model Based Diagnostics and Prognostics,” Annual Conference of the Prognostics and Health Management Society 2010, Portland, OR, Oct. 2010.




Fig. 1 Example of a GPD fit to pitch rate data, for the 5% of highest pitch rates.




Fig. 2 Example of comparison of trigger settings versus altitude for fitted and zero shape parameters, for a probability setting of 1×10^{-6} .




Fig. 3 A GPD fit to Monte Carlo data showing issues when there are outliers beyond the distribution.




Fig. 4 Example of an abort trigger design.




Fig. 5 RSS pitch and yaw trigger versus separate triggers for pitch and yaw.




Fig. 6 Phase plane trigger versus separate triggers for attitude rates and attitude errors.




Fig. 7 Behavior of data (attitude rates versus attitude errors) for possible phase plane trigger development is well behaved in some regions and not in others.

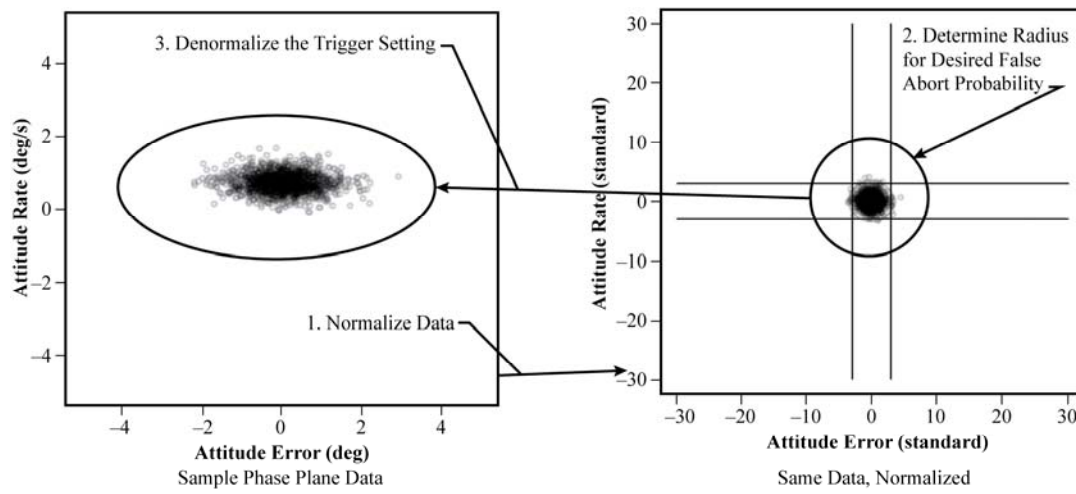


Fig. 8 Diagram of phase plane trigger design procedure.

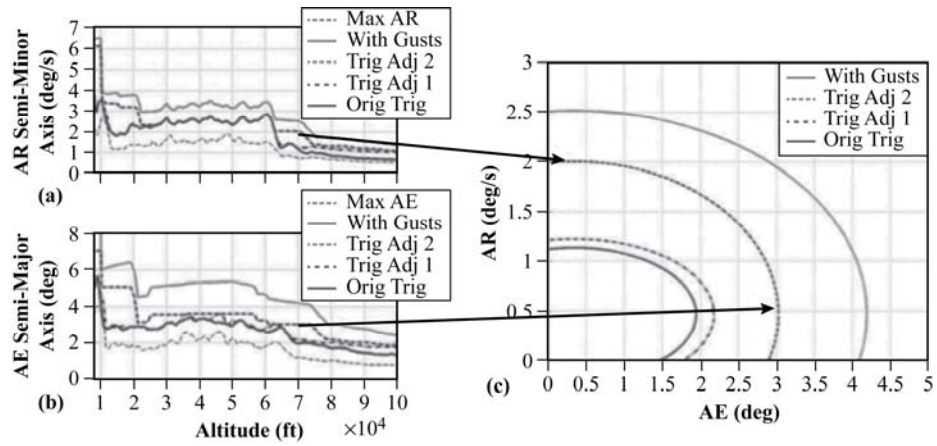


Fig. 9 Attitude rate (AR)-attitude error (AE) phase plane trigger adjustments using the same approach as in Fig. 4: (a) AR axis setting, (b) AE axis setting, and (c) AR-AE phase plane at 70,000 ft.




Fig. 10 Fraction of success resulting from Monte Carlo simulation (example plot).




Fig. 11 Time to escape for a particular failure type, as a function of when the failure occurs.




Fig. 12 Example of a vehicle load indicator (VLI) evaluated at locations on the vehicle.




Fig. 13 Attitude rates after actuator hardover failure, for failures occurring after launch abort system jettison.