Launch Vehicle Failure Dynamics and Abort Triggering Analysis

John M. Hanson[1]
NASA Marshall Space Flight Center

Ashley D. Hill[2]
Dynamic Concepts, Inc.

Bernard B. Beard[3]
ARES Corporation

Launch vehicle ascent is a time of high risk for an on-board crew. There are many types of failures that can kill the crew if the crew is still on-board when the failure becomes catastrophic. For some failure scenarios, there is plenty of time for the crew to be warned and to depart, whereas in some there is insufficient time for the crew to escape. 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 primarily on data already available from the GN&C system. This work is the result of failure analysis efforts performed during the Ares I launch vehicle development program. Derivation of attitude and attitude rate abort triggers to ensure that abort occurs as quickly as possible when needed, but that false positives are avoided, forms a major portion of the paper. Some of the potential failure modes requiring use of these triggers are described, along with analysis used to determine the success rate of getting the crew off prior to vehicle demise.

Nomenclature

Symbols

$F(x)$	=	CDF for a GPD
h	=	correlation height (or speed, etc.) for the data
H_M	=	total window of interest (altitude, speed, etc.)
P_M	=	desired false abort probability
p^*	=	probability setting for calculating the x^* trigger value
x	=	independent variable in a GPD
x^*	=	desired setting for a trigger as calculated by a GPD
μ	=	threshold setting for a GPD
σ	=	scale parameter for a GPD
ξ	=	shape parameter for a GPD

Acronyms

CDF	Cumulative distribution function
FPR	Flight Performance Reserve

---

[1]Lead for Flight Mechanics Analysis and Design, MSFC/EV40, Huntsville, Alabama 35812, AIAA Member.
[2]Senior Engineer, Huntsville, Alabama 35806.
[3]Senior Consultant, ARES Corporation Tennessee Valley Office, Huntsville, Alabama 35805.

GN&C Guidance, Navigation, and Control
GPD Generalized Pareto Distribution
RSS Root sum square, the square root of the sum of the squares
TVC Thrust Vector Control

I. Introduction

THERE are currently a number of potential humans-to-orbit developments being pursued. Some of these may result in new crew launch capabilities. Most of these potential 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 essential that the crew have abort capability to quickly escape a failing vehicle. Typically the abort capability early in flight is driven by two scenarios: pulling (or pushing) the crew away and to safety when the vehicle fails on the launch pad, and pulling (or pushing) the crew away and to safety at maximum drag and maximum dynamic pressure conditions. The first of these requires a substantial total impulse to take the crew far from the failing vehicle, and the second requires overcoming the drag and dynamic pressure effects to get the crew away from the (possibly still accelerating) vehicle with sufficient net relative acceleration. If the abort capability is successful in these regions, then other regions of flight are typically not as difficult. Later in flight, this escape system may be unavailable (it is typically jettisoned to avoid the payload penalty of taking it all the way to orbit), and engines would need to be shut down sufficiently rapidly so that the crew can safely abort the mission. In some cases, attitude rates after engine shut down late in flight can be large, potentially jeopardizing successful separation of the crew module. Issues can arise in all parts of flight, so all regions of flight should be investigated[1].

There are a large number of possible failures that might cause loss of control or loss of performance. 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 hard over (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 burn through might cause nozzle damage. A solid booster could suffer a case or joint burn through, where part of the plume escapes through a growing hole. Failure in the navigation system or an untested software issue could lead to making bad flight control commands. Some other failures (for example, 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.

In addition, a number of historic launch vehicle failures were ultimately a failure to control flight. Examples include the first Pegasus XL, the first Delta III, the Conestoga launch vehicle, and the first Ariane 5 launch[2, 3, 4]. Although a first flight of a new vehicle is typically the riskiest, things can and do go wrong on later flights. Examples in crewed flights include the Challenger and Columbia disasters, Apollo 8 being struck by lightning, Apollo 13 experiencing pogo problems and later catastrophic failure in the Service Module, and various Soyuz issues.

Other launch vehicles had engine failures that made it impossible to reach orbit. Cases where the vehicle is under control but will not reach orbit do not require immediate abort, so a calculation of the ability to reach orbit would be caution and warning information rather than an immediate abort trigger. The crew could wait and abort later or wait until the engine shuts down due to the low propellant condition.

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 all of them are detected if the objective is to understand where every failure originated. However, the 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.

This paper covers two primary topics related to abort recommendation using GN&C data, using analysis conducted during the Ares I development effort. First there is a discussion of loss-of-control abort trigger development. Then the paper covers the modeling of various failure modes in simulation and their use to measure the ability of the crew to escape the failing launch vehicle prior to reaching some "demise" limits.

II. Design of Loss of Control Abort Triggers

A. Trigger Lessons Learned

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. The general idea is to set the trigger values at points where they are beyond any values expected in normal (no-failure) dispersed flight, so that when the trigger is passed, something must be going wrong. On the other hand, setting the values too high is not a good idea since some failure situations allow for only a very short time before the crew would not survive. It is best to get the crew off the failing vehicle as soon as possible when failures are occurring that will shortly lead to vehicle demise.

An additional trigger possibility for detecting loss of control would be if a measurement of the actuator angle 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.

Following are some lessons learned from Ares I, which should be revisited for future applications in order to determine whether they apply to a new vehicle:

- Roll attitude error does not require abort; the vehicle can generally proceed to orbit successfully if the roll angle is not successfully controlled.
- 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. A reasonable value for a roll rate abort trigger might be somewhat below a value for which abort is no longer feasible; that is, the abort and successful recovery of the crew have their own roll rate limits. For stages with multiple engines, this may not be an issue since lack of roll control most likely means lack of 3-axis control in general. 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. 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 filtered data. The filtered data are more smooth, but the quickness of the response with the unfiltered data more than makes up for the more noisy data.
- 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 liftoff region (prior to tower clearance) should be designed separately from the rest of first stage ascent, since this region contains several unique qualities, including 1) flight control response to stabilize the vehicle once the nozzles clear the launch platform, 2) a nearby tower that is an additional demise condition, and 3) 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, use of a monotonic parameter such as altitude should provide better trigger behavior than time, since time does not vary with vehicle dispersions, whereas altitude does. 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.

- 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. This region of flight is usually benign from the standpoint of the rates and errors (disturbances in normal flight are small and no large maneuvers are needed), and thus somewhat insensitive to the choice of independent parameter.
- 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 liftoff to avoid the pad or a pitch maneuver just after staging). Figure 1 shows that a combination of pitch and yaw error or rate will be detected faster with a RSS trigger as opposed to separate pitch/yaw triggers.


Figure 1. RSS pitch and yaw trigger versus separate triggers for pitch and yaw.

- 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 2 shows an example that demonstrates how this trigger method can lead to faster detection of a problem. When there is a very rapid attitude rate buildup, such as would occur with an actuator hard over (at least for a single engine launch vehicle), the phase plane trigger does not offer much of an advantage. However, 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 two-engine vehicle with an actuator hard over, the phase plane trigger can be substantially faster. Even for a hard over failure, if the attitude error is already significant (one of the outliers in the black Monte Carlo scatter data), detection would be faster with the phase plane trigger.

Figure 2. Phase plane trigger versus separate triggers for attitude rates and attitude errors. Monte Carlo no-failure dispersed data are in black.

- The data used to generate the phase plane triggers must be well-behaved (somewhat like a bell-shaped curve of density) in order to generate a reasonable statistical ellipse (later, procedures for generating the appropriate ellipse will be discussed). Figure 3 shows how the data may be reasonable for this purpose in some flight phases and not in others. The data at 24 seconds and at 100 seconds are not good for designing a probability ellipse, whereas the data at 55 seconds is. An ellipse could be drawn around the 24-second data, but determining the statistical likelihood of exceeding the ellipse setting would be problematic.
- Abort success (defined here as getting the crew off prior to some demise condition) is typically the toughest when dynamic pressure is high. It's critical to design the triggers to provide for maximum warning in this region. Fortunately the data used to generate phase plane triggers is well-behaved in this region.
- 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 depending on the failure.
- Another challenging region of flight is just before orbit injection, when the vehicle weighs the least and has the most forward center of mass. A failure that leads to significant attitude rates, such as an actuator hard over, 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.
- If integral 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.

Figure 3. Behavior of data (attitude rates vs attitude errors) for possible phase plane trigger development is well-behaved in some regions and not in others.

B. 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 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.

Figure 4 shows that the behavior of the outlying values do not behave in a Gaussian manner. The black curves in the figure show attitude error Monte Carlo results, and during most of the flight the one or two extreme cases move around and in some places exceed all other values by a substantial factor. 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 2,000 samples). So clearly the trigger setting was not sufficiently high to ensure no false aborts. In the figure, the solid red curve shows the 95[th] percentile data points for high attitude error. It can be seen that the maximum dispersed values sometimes exceed the 95% data by a factor of more than two. Most of the Monte Carlo data fail the Anderson-Darling[5] 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. 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. In particular, the approach uses a peak-over-threshold method[6] and fits a Generalized Pareto Distribution (GPD) to the data. The peak-over-threshold method is displayed in Figure 5, which shows data for rainfall in a village in Scotland. The parameter that is calculated is the difference between the measured value and the threshold setting (the threshold is 5 in the figure). The approach is to find the "scale parameter" (σ) and "shape parameter" (ξ) that best fit the experimental cumulative distribution function (CDF) $F(x)$ given by

Figure 4. Monte Carlo no-failure dispersion results (attitude error versus altitude), showing non-Gaussian behavior. The red curve is at the 95% level.


Figure 5. Tail values from a distribution of rainfall in a Scottish village, illustrating the peak-over-threshold method, where the threshold is 5.

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

where μ is the threshold. Then, to find the value x^* associated with a certain low target probability p^*, use

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

The free statistics package "R"[7] has a function "fitgpd" for the purpose of determining σ and ξ. The GPD applies only asymptotically for a high threshold. 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 2,000 or top 1,000 out of 20,000 and perform the prescribed procedure). Much of the time, the fit works very well. An example is in Figure 6, 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. On the other hand, as Figure 4 showed, sometimes there are one or two data points that lie well beyond the rest. In this case, the fit does not work as well. Also, due to the random and nonlinear nature of the data as a function of altitude, using a fitted ξ leads to an abort trigger with significant spikes (see Figure 7). 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.

Notice how many altitude data points there are in the trigger settings calculated in Figure 7. Unlike the single sample of rainfall in Figure 5, flight control parameters are continuous functions of altitude (or speed, etc.). 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 is to calculate p* by using

Figure 6. Example of a GPD fit to pitch rate (PR) data.

$$p^* \sim P_M (h/H_M) \tag{3}$$

where P_M is the desired false abort probability (e.g. 1×10^{-6}), h is the correlation height (e.g. 2,000 ft), and H_M is the total time or altitude (or whatever) window. So, for example, if 2,000 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 Figure 8). So some additional margin must be added in these cases.

American Institute of Aeronautics and Astronautics

Figure 7. Example of comparison of trigger settings versus altitude for fitted and zero shape parameters, for a probability setting of 1x10^{-6}.


Figure 8. A GPD fit to Monte Carlo data showing issues when there are outliers beyond the distribution.

Figure 9 shows the development of final trigger settings. The maximum of the Monte Carlo data is shown in blue. The black data are what the GPD approach generates. For most of flight, the black curve is well above the blue data, but there are several issues. First, there are regions where the black candidate trigger setting is very close to the blue max value seen in simulation. Even worse, some cases resulted in the black candidate trigger level being less than

Figure 9. Example of an abort trigger design.

the worst-case blue value from the simulation. Cases like this can result where 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 (for example, in Figure 9 at higher altitudes the max attitude rate seen is about 0.5 deg/sec, 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 a hazard. Some margin on top of the max 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 blue curve just above 60,000 ft could be slightly later, and would thus exceed the black 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 aqua 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 (again based on engineering judgment; the data used for calculation is not the black trigger setting but what the aqua limit would be without the original black trigger). Finally, the green curve shows the final trigger setting, resulting from adding the wind gust effect as described earlier.

A similar approach may be used to design the phase plane triggers, by converting the two-dimensional phase plane parameters into a normalized radius (see Figure 10). The procedure is this:

- Normalize each parameter by its standard deviation
- Generate a radius (square root of the sum of squares)
- Generate the trigger setting for this one-dimensional radius
- De-normalize to get the individual axes back
- The center of the phase plane ellipse is at the mean of the two sets of data
- Apply the trigger adjustments (as in Figure 9) to the individual axes (see Figure 11)

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 (Figure 3), and the regions were more benign for abort timing, so the phase plane triggers were not used.
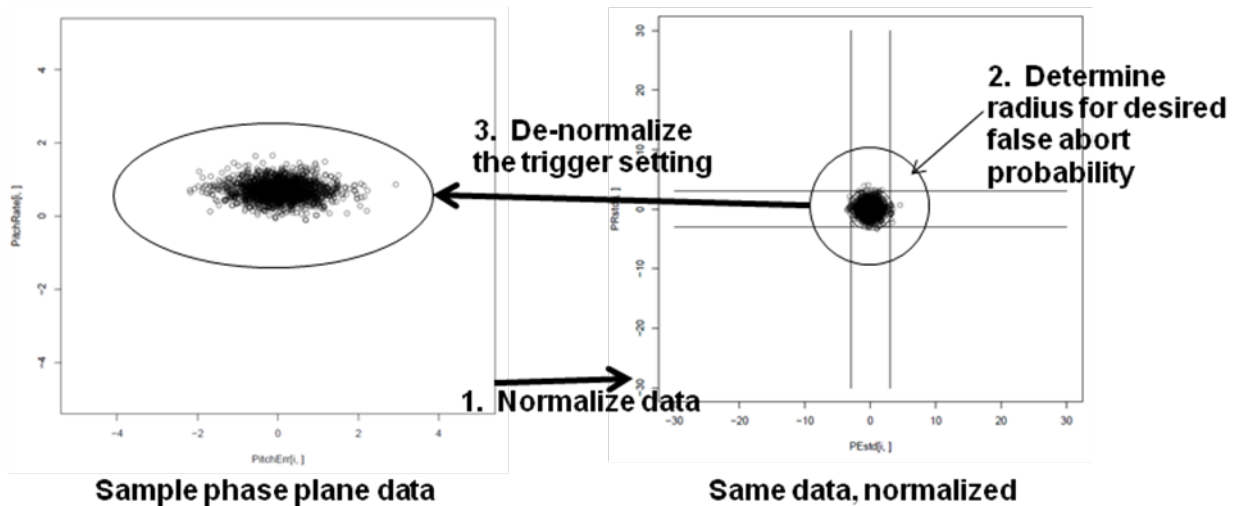
American Institute of Aeronautics and Astronautics

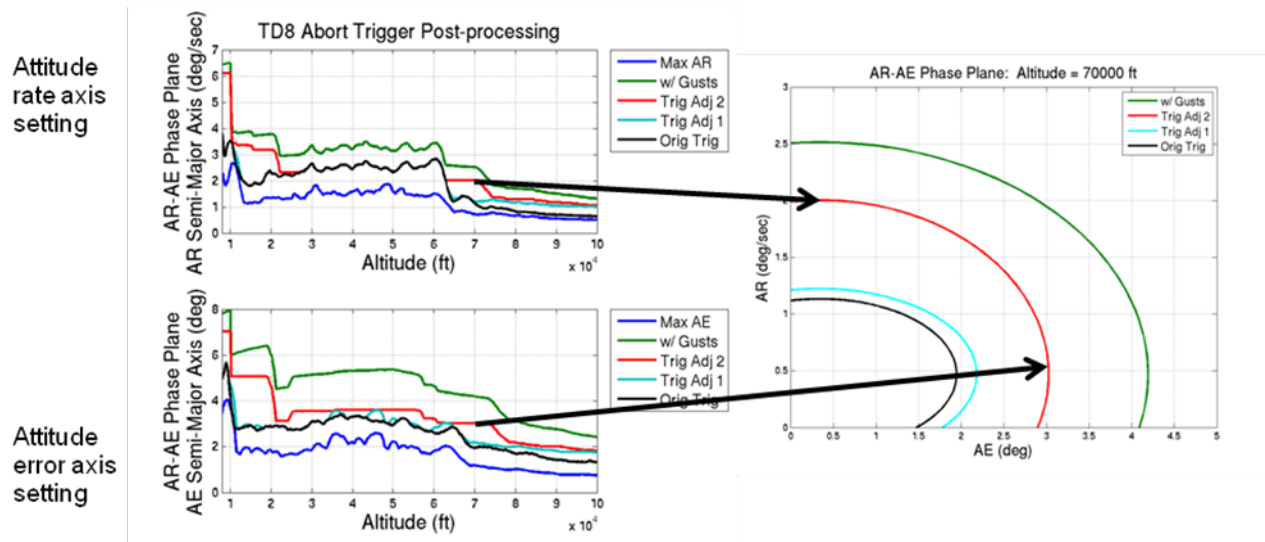**Figure 10. Diagram of phase plane trigger design procedure.**

Figure 11. Phase plane trigger adjustments using the same approach as in Fig. 9. Each individual parameter is adjusted as in Fig. 9 and the result is applied to the semi-major axes of the phase plane ellipse. The arrows show the adjustments to the phase plane trigger at 70,000 ft.

So the methods for designing the various abort triggers have 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, 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. Certainly test flights should be used to verify the model correctness (including such items as noise in the data) 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 (for example, 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 mis-modeling or for whatever reason) to clear up.

III. Using the Triggers to Determine Abort Success

Suppose that a high-fidelity simulation of the vehicle of interest is available, and the abort triggers have been determined using the procedures in section II. Then failures may be simulated and the ability of the system to abort prior to vehicle demise may be determined. The triggers developed using these methods are all for flight-control related measurements, so the failures of interest are ones that result in loss of control, rather than an explosion, for example. Some possible demise criteria include:

- 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, thereby allowing more time for abort. 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.
- A structural load indicator is exceeded that means the spacecraft will be unable to abort (due to structural reaction to the release of bending moment that occurs just after abort, for example). Again, a dynamic pressure times angle of attack indicator is only valid for trimmed flight, so a higher fidelity indicator is desirable.
- The launch tower is struck.
- An attitude rate limit is exceeded. This could be a limit for which spacecraft separation is no longer possible, or a limit for which abort will be unsuccessful.
- An attitude error limit is exceeded.
- Altitude rate becomes negative in parts of flight where it should be positive. 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.
- Crew physiological limits are about to be exceeded.

Other demise criteria are possible depending on the application. Some potential failure modes that could be modeled include:

- TVC fail to hard over, 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.
- TVC fail in place from a hardware failure, which means the starting condition after the failure would be trimmed flight.
- TVC fail to null from a hardware failure.
- Engine nozzle failure of various kinds; 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, 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).[8]
- 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, 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. [9]

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. Of course, if there are multiple engines, some of these failure types do not have as severe consequences as when there is only one engine providing the flight control.

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 after abort shutdown 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 liftoff 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 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 then 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 then for the abort motor to ramp up in thrust. Then the question to be answered is whether any of the demise criteria are exceeded before the crew leaves.

Success is defined as any case where the crew survives. This includes escaping the vehicle prior to demise, as well as cases where the vehicle did not experience demise. 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. Cases where abort did occur but the vehicle did not experience demise should be investigated, should any of these cases occur. 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. 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 of understanding results of Monte Carlo failure simulations are shown in Figures 12-15. In each of these simulations, a single failure type was modeled, with random time of failure, during a particular phase of flight (e.g. liftoff region, first stage ascent, upper stage prior to launch abort system jettison, and upper stage after launch abort system jettison). Figure 12 shows an example of resulting success fraction, with the different colors corresponding to different types of failure and the different groups corresponding to different combinations of abort triggers. Figure 13 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 full thrust. The time is how much time there is after zero prior to vehicle demise as defined by the various demise criteria. The colors correspond to the reason for demise in each case. Negative times should be investigated since they may correspond to loss of crew. Figure 14 shows the value of a particular load indicator from the various Monte Carlo


Figure 12. Fraction of success resulting from Monte Carlo simulation. The different colors are different failure types and the different groups are different combinations of abort triggers (e.g. phase plane, RSS attitude error, pitch/yaw attitude error separately, attitude rates, etc.)

Figure 13. Time to escape for a particular failure type, as a function of when the failure occurs. Zero time means the time when the crew departs, after the abort triggers are exceeded and after a suitable time delay. The time shown is when the demise criteria are exceeded, and the colors correspond to the reason for demise. Negative times should be investigated.

samples, evaluated at each vehicle location. This load indicator is normalized so that a value greater than 1.0 means the limits are exceeded. Finally, Figure 15 shows the value of attitude rate for failures that occur during upper stage flight after the launch abort system is jettisoned, for an actuator hard over failure. In these cases, the main engine would be shut down, and the issue is whether the spacecraft can be successfully separated at the rates that result. The different colors 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. Rates are not surprisingly highest late in flight, when the vehicle mass and inertia are least and the center of mass is most forward.

Figure 14. Example of a vehicle load indicator (VLI) evaluated at locations on the vehicle (zero is above the nose of the vehicle). A value higher than 1.0 is an indication that acceptable loads are exceeded.

Figure 15. Attitude rates after actuator hard over failure, for failures occurring after launch abort system jettison. The different colors correspond to different abort trigger combinations.

IV. 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 paper also describes how to use the trigger settings to generate abort success measures, which could lead to additional refinement of the trigger settings.

References

[1]David Shayler, Space Rescue: *Ensuring the Safety of Manned Spaceflight*, Springer, Nov. 2008.

[2]D. M. Harland and R. D. Lorenz, *Space Systems Failures: Disasters and Rescues of Satellites, Rockets and Space Probes*, Springer, May 2005.

[3] "The Wrong Stuff: A Catalog of Launch Vehicle Failures", http://www.astronautix.com/articles/thelures.htm.

[4]Ashley D. Hill, "An Investigation of the Flight Vibrations Experienced During the METEOR Mission of the Conestoga Launch Vehicle", Master's thesis, University of Alabama in Huntsville, 1996.

[5]See for example http://en.wikipedia.org/wiki/Anderson%E2%80%93Darling_test

[6]Stuart Coles, *An Introduction to Statistical Modeling of Extreme Values*, Springer-Verlag, London, 2001.

[7]R Statistics Package, located at http://www.r-project.org/.

[8]D. G. Luchinsky, V. V. Osipov, V. N. Smelyanskiy, A. Patterson-Hine, B. Hayashida, M. Watson, J. McMillin, D. Shook, M. Johnson, and S. Hyde " Model-Based Diagnostics and Prognostics for Solid Rocket Motors", Annual Conference of the Prognostics and Health Management Society 2009, San Diego, Sept. 2009.

[9]D. G. Luchinsky, V. Hafiychuk, I. Kulikov, V. Smelyanskiy, A. Patterson-Hine, J. Hanson, and A. Hill, "Stage Separation Failure: Model Based Diagnostics and Prognostics", submitted to the IEEE Aerospace Conference, 2011.