An Assessment of the Efficacy of Launch Collision Avoidance Analysis in Providing Risk Reduction for NASA

Brian A. Beaver
Kennedy Space Center

Mark E. Hametz
a.i. solutions, Inc.

Jarmaine C. Ollivierre
Kennedy Space Center

Lauri K. Newman
Goddard Space Flight Center

Matthew D. Hejduk
a.i. solutions, Inc.

March 2015

Introduction

The objective of this document is to assess the value of collision on launch assessment (COLA, often referred to as "launch COLA") practices and provide recommendations regarding its implementation for NASA robotic missions. The scope of this effort is limited to launch COLA screens against catalog objects that are either spacecraft or debris. No modifications to manned safety launch COLA practices are considered in this effort.

An assessment of the value of launch COLA can be broken down into two fundamental questions:

1) Does collision during launch represent a significant risk to either the payload being launched or the space environment?
2) Can launch collision mitigation be performed in a manner that provides meaningful risk reduction at an acceptable level of operational impact?

While it has been possible to piece together partial answers to these questions for some time, the first attempt to comprehensively address them is documented in reference (a), *Launch COLA Operations: an Examination of Data Products, Procedures, and Thresholds,* Revision A. This report is the product of an extensive study that addressed fundamental technical questions surrounding launch collision avoidance analysis and practice. The results provided in reference (a) will be cited throughout this document as these two questions are addressed.

The premise of this assessment is that in order to conclude that launch COLA is a value-added activity, the answer to both of these questions must be affirmative. A "no" answer to either of these questions points toward the conclusion that launch COLA provides little or

no risk mitigation benefit. The remainder of this assessment will focus on addressing these two questions.

Risk Posed by Collision During Launch

On the surface, the issue of a collision between a launch vehicle (and its associated payload) and existing on-orbit objects appears to be straightforward. It is widely recognized that space is crowded, with approximately 20,000 on-orbit objects currently cataloged by the US Air Force's Joint Space Operations Center (JSpOC). It is easy to visualize, then, that a launch vehicle heading to its prescribed orbit must fly through a buzzing cloud of operational satellites, spent rocket stages, and various other debris, and that collision with any of these objects would have serious, if not catastrophic, consequences. Furthermore, since the orbits of these objects are known fairly accurately, and the trajectory of the launch vehicle is also known in advance, it seems both possible and prudent to compare the launch vehicle's projected path for all possible launch times with the propagated states of the objects in the JSpOC's catalog to identify possible collisions. These potential collisions can then by mitigated by not launching at those times when an unacceptable collision risk presents itself.

In reality, however, there are numerous subtleties associated with this problem that render it non-trivial. The first, most basic of these, is the question of overall collision risk. If no mitigation steps are taken, what is the likelihood that a collision will take place? This question is fundamental in that it represents a characterization of the risk being mitigated, and its answer quantifies the importance of launch collision risk in relation to other risks that must be managed by NASA.

The study documented in reference (a) provides a comprehensive examination of this question. Systematic Monte Carlo analysis demonstrates that the odds of a launch vehicle striking a space object during launch are extremely small, even when no active mitigation measures are employed. Using a set of five representative launch trajectories analyzed over a total of 216,000 distinct screenings, this study showed that the median risk of collision (based on the higher precision SP object catalog) is approximately 4×10^{-8} for any given single launch, with a 5^{th}-95^{th} percentile range of 3×10^{-12} to 8×10^{-7}. If the trade space is expanded to include not just a given single launch, but a collection of launches over time, the risk that a collision will occur correspondingly increases. The cumulative median risk crosses the 10^{-6} threshold after about six launches, 10^{-5} after about 50 launches, and 10^{-4} after about 460 launches. These results are in general agreement with a much more rudimentary previous assessment documented in reference (b).

Of course, the cumulative probability of collision is dependent on the state of the on-orbit population, and will increase as that population grows over time; this effect is not explicitly captured in this analysis. It should be noted, however, that the probability of collision is roughly proportional to on-orbit object density, so unless the on-orbit population grows by roughly an order of magnitude, the overall conclusions presented here should not be substantially affected. Based on analyses published by the NASA Orbital Debris Program

Office as represented in references (c) and (d), an order of magnitude increase in on-orbit population is very unlikely, even projecting 100 years or more into the future.

It is readily apparent that the risk of collision during launch is much lower than the background risk of mission failure due to other causes (historically on the order of 10^{-2}, even for mature launch systems). Given the difference of four or more orders of magnitude between the risk of launch collision and overall risk of launch failure, and the plethora of other potential failure modes with a greater likelihood of occurrence than a collision during launch, it is difficult to cast launch collision risk as a significant mission success concern. Of course, there are specific instances where a risk assessment like this may not be appropriate. These situations can be identified explicitly and directly mitigated through prudent mission design practices. In order to mitigate risk to high-valued assets, in some cases insertion orbits can be selected that preclude the possibility of collision. For example, spacecraft whose ultimate destination is in NASA's Earth Observing System (EOS) constellation can be (and routinely are) intentionally inserted by the launch vehicle into an orbit below that of the constellation to eliminate the possibility of collision with existing EOS assets. These newly-launched spacecraft are then carefully maneuvered into final position using on-board propulsion capability.

It has also been argued that while collision during launch may not have a significant effect on launch mission success, it could present a long-term risk to the space environment. The reasoning behind this argument is that a collision between a launch vehicle and on-orbit object would produce a cloud of debris that would increase the on-orbit debris population and potentially place a large number of on-orbit assets at a perpetually elevated risk of collision. However, analysis of the evolution of the on-orbit debris environment performed by NASA's Orbital Debris Program Office and documented in reference (c) supports the contention that collision during launch in fact does not pose a significant risk to the space environment over the long term. This reference indicates that collisions between existing on-orbit objects are expected to occur at an average rate of approximately 1 every 5 years over the next 200 years, approximately half of which are projected to be "catastrophic" debris-producing events. By comparison, the risk of a single launch collision occurring at some point, even over the course of hundreds of launches, is at least three to four orders of magnitude below this background condition. Additionally, explicit analysis of the effect of active on-orbit collision avoidance documented in reference (c) demonstrates that even if all active payloads are removed from the pool of potential colliding bodies, the effect on the growth of the debris population is negligible. In a general sense, active on-orbit collision avoidance measures offer a very limited benefit from the standpoint of limiting long-term debris population growth in orbit (although it could provide shorter term risk reduction for groups of satellites flying in similar orbits), and since launch collision avoidance is by its very nature much more limited in scope that its on-orbit counterpart, its effect on the long-term debris population is even less significant. Potential near-term concerns over debris clouds that could be produced in orbits that coincide with high-valued payloads and/or constellations can be mitigated through careful launch vehicle target orbit selection, as mentioned previously.

Thus, from the standpoint of both the specific mission being launched and the on-orbit environment in general, extensive analysis demonstrates that launch collision risk is several orders of magnitude below known baseline on-orbit collision risk levels. *This conclusion supports the contention that the insignificant incremental risk presented by launch collision does not warrant active mitigation.*

Launch Collision Mitigation Effectiveness

The central question regarding the effectiveness of launch collision mitigation comes down to whether preflight analysis can result in a significant reduction in the risk of collision at an acceptable cost in terms of both resources and operational impacts. The following discussion will focus primarily on the tradeoff between risk reduction and operational impacts; financial and other resource requirements are beyond the scope of this document.

This discussion is based on the premise that technically sound analytical techniques, tools, and supporting data exist that can be brought to bear on the launch COLA problem. Reference (a) addresses this issue in great detail, concluding that:

1) The necessary input data to support probability-based screening is available; most significantly, existing methods of characterizing launch vehicle state uncertainty in flight are technically sound and consistent with flight data.
2) It is reasonable to use either the Special Perturbations (SP) or General Perturbations (GP) Department of Defense (DOD) catalog for probability-based screening, with the lower-fidelity GP catalog producing slightly more conservative results in general.
3) Miss-distance-based screening does not correlate with risk in any direct, practical way, and should not be used as a substitute for probability-based screening.

Given these three assertions, it can be concluded that launch collision probabilities indeed can be determined in a technically sound manner. These points will not be discussed in any detail here, and the reader is referred to reference (a) for extensive supporting information and discussion. However, it should be noted that assertion 1 above is of particular importance to subsequent discussion. This is because probability of collision is a function of three basic characteristics: the geometry of the encounter between the two objects, the sizes and shapes of the objects, and uncertainties in the objects' states. The first two contributing factors are reasonably well characterized and noncontroversial; likewise, the uncertainties associated with on-orbit objects are generally known and accepted. However, the analytically-produced state uncertainties associated with launch vehicle trajectories have historically been viewed with a certain amount of skepticism, and even mistrust, within the launch COLA community. As it happens, launch vehicle state uncertainties play the primary role in answering the question of launch COLA effectiveness due to the fact that they are typically much larger than those associated with on-orbit objects, and as a result they tend to be the dominant factor in the calculation of collision probability. Therefore, one of the most critical aspects of the study documented in reference (a) was to verify or refute the viability of analytically produced launch vehicle

uncertainties through statistical comparison with flight data. Reference (a) clearly demonstrates that for mature existing launch vehicle fleets, the trajectory uncertainties that feed collision probability calculations, while large, are indeed technically sound. Again, the reader is referred to reference (a) for a complete discussion of this topic.

Given the premise that collision probability can be calculated in a technically sound manner, it becomes possible to use existing analytical techniques to quantify the benefits of actively screening launches to identify and eliminate launch opportunities that involve conjunctions with probabilities of collision above a predetermined screening threshold. This is done by comparing the cumulative risk of collision incurred when no active measures are taken (discussed in the previous section) against the cumulative risk obtained by actively screening launches against the on-orbit catalog. Intuitively, the outcome of this sort of comparison will depend on the screening threshold selected, with the expectation that screening at tighter (i.e., lower) threshold values will result in lower cumulative risk. By sampling multiple trajectories at various launch COLA screening thresholds, statistics on the risk of collision as a function of the number of launch events and COLA screening threshold can be constructed; the results of this sort of approach are documented in reference (a) and presented in Figure 1 below. (These results are based on the higher fidelity SP catalog; results using the GP catalog are generally similar.) Note that the curves for no screening at all and a screening threshold of 10^{-5} are coincident; this illustrates the (initially) somewhat surprising conclusion that screening at a maximum probability of collision threshold of 10^{-5} or higher offers no benefit, from an overall risk standpoint, over not conducting screenings at all, regardless of the number of launch events considered. The dispersions (uncertainties) in the predicted launch trajectories, due both to uncertainties in rocket motor and guidance system performance, as well as winds and other unpredictable environmental factors on launch day, are large enough that they dominate the probability of collision calculation—given these uncertainties in the predicted position, it is simply not possible to generate a calculated likelihood greater than approximately 1 in 100,000 that a launching vehicle will intercept a resident space object. As the graph below shows, one must lower the threshold to 10^{-6} to begin to see some actual reduction in risk (a factor of two); and a threshold of 10^{-7} is required to begin to approach an order-of-magnitude benefit.

Screening at such low levels does not come without a cost in terms of operational impact. The data in reference (a) can be used to determine the relationship between risk improvement and launch window impact. This relationship is complex but as an example, screening at a threshold of 10^{-7} (which provides less than an order of magnitude improvement in cumulative risk) results in the loss of 65-70% of the nominal launch window (95th percentile) for launch windows between 10 and 60 minutes in duration using the SP catalog. This relationship is shown in Figure 2.
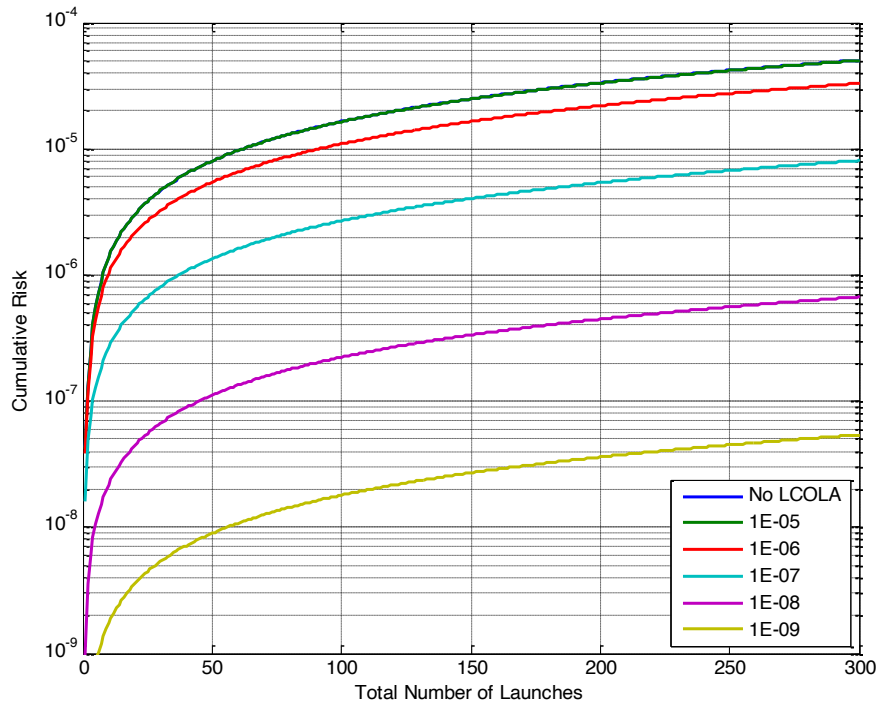
Figure 1: Cumulative risk of collision as a function of launch events and screening threshold
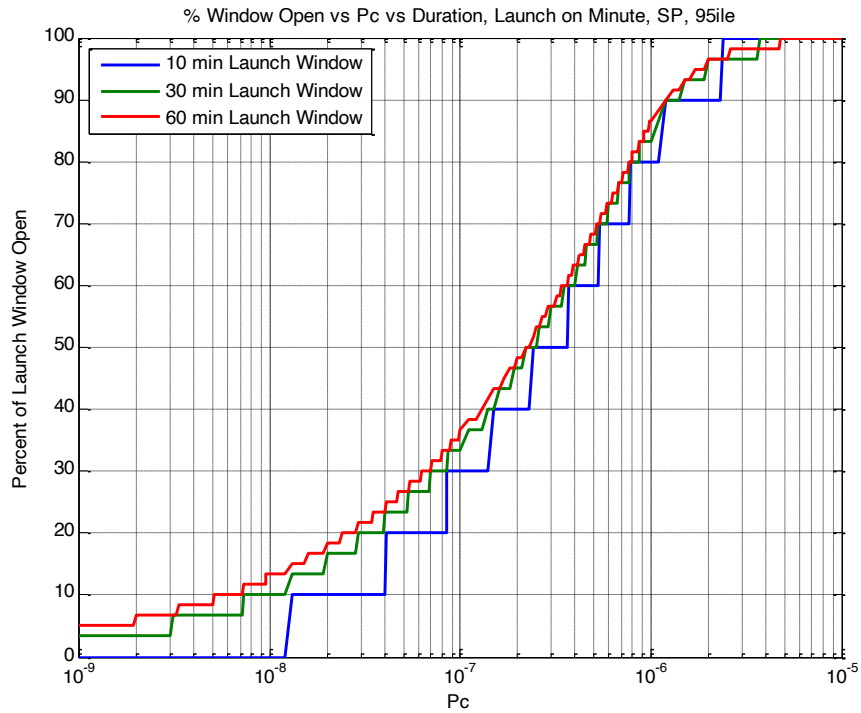

Figure 2: Launch window availability vs LCOLA screening level used

Thus, it appears that while technically sound methodologies do exist to compute collision probabilities, the risk reduction can be regarded as minimal at best, and then only at very low risk levels accompanied by potentially significant operational impacts. Therefore, some additional comparisons should be made to help put launch COLA into perspective with existing similar practices within the agency.

Risk Threshold Comparisons

Although no agency-wide launch COLA standard exists, there are other, related practices currently implemented within NASA that can be examined to gain insight into Agency risk tolerance for this type of issue.

One prominent source of insight here is the existing process employed by NASA for on-orbit conjunction assessment and collision avoidance for the International Space Station (ISS). The ISS is actively screened against the entire on-orbit catalog on an ongoing basis using catalog data, tools, and analytical methods that are similar or identical to those used for launch conjunction assessment. Reference (e) defines the maneuver criteria for ISS in response to on-orbit conjunctions. Of most relevance to this discussion is the fact that the document establishes a "green" probability of collision as 10^{-5} or less, meaning that under this condition debris avoidance maneuvers are not conducted. Furthermore, even for probabilities of collision between 10^{-5} and 10^{-4} ("yellow" condition) and greater than 10^{-4} ("red" condition), maneuvers are not mandatory, and can be waived under certain circumstances.

From a risk tolerance perspective, it is difficult to argue that collisions involving robotic launches are inherently less risk tolerant than those involving risk to human life aboard ISS. From a risk management perspective, then, it would seem consistent to screen launch conjunctions at a probability threshold that is no lower than that applied to ISS collision screening on orbit. If this approach – specifically, that no action is taken for probabilities of collision below 10^{-5} – is employed for launch screening, it can be concluded that active screening is not necessary because we now know that screening at that level or higher has no effect on overall cumulative risk of collision based on the results in reference (a). In order to see even a small reduction in cumulative risk, screenings must be conducted at a level below 10^{-6}, with all of the associated impacts to launch window.

A similar comparison with on-orbit screening practices for robotic payloads can also be performed, although the story is not quite as straightforward. NASA performs continuous conjunction assessment for approximately 70 specific unmanned assets. As part of this process, affected spacecraft operators are provided with notification information when the probability of collision exceeds approximately 10^{-7} and full analysis packages when the probability of collision exceeds approximately 10^{-4}. Unlike the ISS situation described above, there are no firmly established requirements for executing an avoidance maneuver for robotic spacecraft; instead, a number of event-specific aspects, such as conjunction geometry, quality of orbit determination solution(s), event parameter history, and space weather conditions, are all considered in addition to the calculated probability of collision in determining whether active remediation is appropriate. However, in the main, some

sort of risk mitigation action is typically pursued for probability of collision levels greater than 5×10^{-4}.

Potential basis for comparison can also be sought in NASA's existing orbital debris policy, as documented in reference (f). This document provides explicit requirements aimed at limiting the production of orbital debris by NASA space missions, as well as protecting the populace from risks associated with orbital debris. An examination of the various thresholds outlined in this standard provides some insight into historical risk levels accepted by NASA at the agency level; pertinent threshold values are listed in Table 1 below. The reader is referred to reference (f) for more details on each of these requirements.

Table 1: NASA Orbital Debris Requirement Risk Thresholds

Requirement	Description	Threshold
4.4-1	Accidental explosion during deployment and mission operations	10^{-3}
4.5-1	Risk of collision with large (i.e., tracked) objects over the lifetime of a space system	10^{-3}
4.7-1	Risk of human casualty associated with reentry	10^{-4}

These specific orbital debris requirements are useful in identifying precedent for the Agency's risk tolerance when it comes to matters of on-orbit collision and space debris generation. The first two table entries are closely related to the launch COLA problem in that their direct concern is limiting the generation of space debris. In particular, requirement 4.5-1, which applies to launch vehicle hardware as well as spacecraft, establishes a total limit on collision risk during the lifetime of a space system. If we consider the risk of collision during launch as a contributor to this total limit (even though reference (f) does not explicitly do so), we know from the earlier discussion based on reference (a) that its contribution is so small as to be insignificant even without active launch COLA measures. The limit on risk of human casualty (requirement 4.7-1) is included here to provide some general risk tolerance perspective. If NASA were to screen against launch collision at a level comparable to that already established for global human casualty risk there would be no reduction in the overall risk of collision; indeed, screening at a level three orders of magnitude more stringent than the risk level already accepted for human casualty would be required to have a meaningful effect on the overall risk of collision during launch. From an objective standpoint, it appears difficult to justify treating launch collision screening in a manner that is at least a thousand times more conservative than the established human casualty risk tolerance threshold.

Although it is difficult to establish an "apples-to-apples" equivalence between launch collision risk and existing risk thresholds already applied within NASA, three areas were identified that were similar enough in nature for consideration – collision avoidance for the International Space Station, on-orbit collision avoidance for NASA's robotic assets, and existing orbital debris mitigation policy. Based on a survey of these existing risk management policies and practices, it is recommended that a probability of collision threshold for NASA robotic mission launches should be no lower than 10^{-4} (1 in 10,000). This threshold is consistent with existing on-orbit screening COLA practices, and

recognizes a level of risk acceptance for unmanned assets that is no more stringent than that for human casualty, both in terms of humans on orbit (i.e., ISS inhabitants) and the general population on the ground.

Conclusions

In the introduction, two fundamental questions were posed whose answers would produce a framework for a technically sound launch COLA policy. The answers to these questions are summarized below.

Does collision during launch represent a significant risk to either the payload being launched or the space environment? Analysis shows that the risk posed by collision during launch is many orders of magnitude below existing background risks to both the specific mission being launched and the space environment, leading to the conclusion that launch collision risk is not a significant concern.

Can launch collision mitigation be performed in a manner that provides meaningful risk reduction at an acceptable cost? Sound analytical methods exist that allow the probability of collision to be reliably calculated prior to launch for conjunctions between a launch vehicle in flight and on-orbit objects, and such methods are in use today. However, the fundamental nature of the launch COLA problem (primarily the large uncertainties associated with launch vehicle flight) severely limits the ultimate benefit realized by performing launch COLA analysis. Screening launch conjunctions at risk levels that are consistent with similar existing NASA policies and practices provides no risk benefit, while screening at tighter collision probability thresholds provides minimal incremental benefit at the cost of significant impacts to launch windows.

Recommendations

A cumulative collision threshold of 10^{-4} for a given launch opportunity represents a risk level that is generally consistent with existing agency policies and practices regarding on-orbit collisions and orbital debris management. Furthermore, comprehensive analysis demonstrates that screening at this level offers no benefit over performing no active screening at all. This is because individual conjunctions with probabilities greater than 10^{-5} do not exist due to the large uncertainties associated with preflight launch vehicle trajectories. This position is supported by extensive statistical analysis, which has verified that large preflight predicted uncertainties are consistent with flight data, and are valid representations of how launch vehicles operate.

Based on the answers to the two questions posed in the introduction, it can be seen that across-the-board launch collision avoidance screening against the on-orbit catalog offers no real value from a risk reduction standpoint. Thus, it is recommended that this practice be discontinued for NASA robotic launches. This position should be periodically evaluated based on the evolution of launch vehicle guidance, navigation, and control systems as well as growth in the on-orbit population, although significant changes in these parameters (on

the order of an order of magnitude or more) would be required to significantly affect this recommendation.

References

(a) Hejduk, M., et. al., *Launch COLA Operations: an Examination of Data Products, Procedures, and Thresholds*, Revision A, FDSS-21-0182, 12 January 2014

(b) Beaver, B., *An Assessment of the Probability of Collision with On-Orbit Objects During Launch Operations*, ELVL-2001-0025354, 19 July 2001

(c) Liou, J.-C., *An Active Debris Removal Parametric Study for LEO Environment Remediation*, Advances in Space Research Vol. 47 p. 1865-1876, 2011

(d) Johnson, N., et. al., *The Characteristics and Consequences of the Break-up of the Fengyun-1C Spacecraft*, Acta Astronautica Vol. 63 p. 128-135, 2008

(e) *ISS Debris Avoidance Maneuver (DAM)/Pre-Determined DAM (PDAM) Criteria*, NASA International Space Station Flight Rule B4-101, 14 November 2011

(f) *Process for Limiting Orbital Debris*, NASA-STD-8719.14A, 8 December 2011