AAS 13-912

RECOMMENDED RISK ASSESSMENT TECHNIQUES AND THRESHOLDS FOR LAUNCH COLA OPERATIONS

M.D. Hejduk,[*] D. Plakalovic,[†] L.K. Newman,[‡] J.C. Ollivierre,[§] M.E. Hametz,[**] B.A. Beaver,[††] R.C. Thompson[‡‡]

This paper describes the second phase of a study to develop uniform launch collision avoidance and risk assessment (LCOLA) guidance and practices among the various NASA launch organizations and ranges. The first phase established the accuracy levels and covariance realism of predicted launch trajectories. This second phase used these trajectory data in a large screening experiment to examine the differences in results between general perturbations and special perturbations screenings. In addition, the trade space among probability of collision (Pc) screening threshold, required duration of the launch window, and average percent of the launch window remaining open was examined; the miss distance offset values that can be used as proxies for LCOLA screenings at different Pc values, and the difficulties that this approach presents, were explored; and the degree to which the risk posture is improved by having a LCOLA program in place, as compared to launching under a "big sky" assumption in which no conjunction prediction and mitigation are performed, was assessed.

INTRODUCTION

Launch collision avoidance and risk assessment (LCOLA)—the practice of screening launch trajectories against a catalogue of space objects to assess conjunction risk and remove launch times that exceed risk thresholds —is presently implemented in a non-uniform way among the various NASA launch organizations and ranges; both the screening and assessment techniques used and the particular thresholds applied vary substantially. NASA Headquarters thus commissioned a study to examine the current processes, recommend a standardized LCOLA approach, and provide a basis for choosing consistent risk thresholds for closing launch windows. This study has been completed and is presently in the approval process for NASA inter-agency

[*] Adjunct Scientist, Mission Services Division, a.i. solutions, Inc., 985 Space Center Drive, Suite 205, Colorado Springs, CO 80915, AIAA/AAS Senior Member.
[†] Aerospace Engineer, Mission Services Division, a.i. solutions, Inc., 985 Space Center Drive, Suite 205, Colorado Springs, CO 80915. AAS Member.
[‡] NASA Robotic Conjunction Assessment Manager, Robotic Systems Protection Program/Mail Code 595, NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771, AIAA Associate Fellow; AAS member.
[§] Mail Code AV-H1, Kennedy Space Center, FL 32899
[**] a.i. solutions, Inc., 8910 Astronaut Blvd., Suite 120, Cape Canaveral, FL 32920
[††] Mail Code AV-H1, Kennedy Space Center, FL 32899
[‡‡] Senior Engineering Specialist, The Aerospace Corporation, 15049 Conference Center Drive, CH4-500, Chantilly, VA 20151

release. For the purpose of conference presentation, the study has been divided into two phases, with this paper describing the results of the second part. The results of the first part, which examined the errors and covariance realism of the predicted launch trajectories provided by the launch vendors, were presented at a previous AAS/AIAA meeting[1]. This analysis established that these trajectories, while significantly more error-infused than a typical general perturbations (GP) satellite ephemeris, nonetheless have their errors properly represented by their accompanying covariances. Thus these trajectories meet the precision criteria for use in orbital safety calculations, as it is not so much the error values themselves but the ability properly to characterize those errors that is important. It was also shown that, because trajectory errors are about an order of magnitude greater than GP satellite catalogue error levels, orbital safety calculations can legitimately be performed assuming either a GP or a special perturbations (SP) paradigm.

With the predicted trajectories and their associated covariances determined to be a legitimate basis for LCOLA, the question of how precisely to use them in LCOLA analysis can be addressed. The differences between screenings conducted with a GP versus an SP catalogue is of interest, especially with regard to the use of the less-precise GP approach as a proxy for SP. Determining the specific ranges of Pc values that can serve as appropriate launch window closure thresholds was a primary goal of this analysis, although the purpose of this study is not simply to recommend hard values (or value ranges) for these thresholds. Rather, it is to give trade-space presentations, such as window closure percentages as a function of window size and Pc screening threshold, so that the practical effects of choosing different window closure (*i.e.,* screening) thresholds can be understood. A more detailed introduction to the concept of launch screening windows and how they are adjusted for LCOLA purposes is given in a subsequent section. Adjunct to this latter issue is an examination of the frequently-proposed use of miss distance as a substitute for Pc-based screening. While the miss distance that gives an equivalent result (at a certain percentile level) to a Pc-based screening can be calculated, the effect of this admittedly simplified procedure on launch window management must be quantified, as miss-distance-based screening is a heavy-handed approach that can exact a considerable toll in terms of launch window availability. Finally, an attempt is made to characterize the overall risk reduction secured by the LCOLA process; this was done by comparing the level of risk assumed by doing nothing at all to mitigate risk, to screening at a given Pc value and closing windows appropriately.

Additional issues are addressed in the full study report, such as the use of cumulative probability of collision of all items identified by a screening rather than that of merely the single conjunction with the largest Pc, and the possibility of screening against only defended assets rather than the entire known space catalogue. Interested readers are referred to the complete report[2] both for these items and for a more comprehensive treatment of the issues discussed in the remainder of this paper.

SCREENING EXPERIMENT DESIGN

To attempt to answer all of the questions outlined above, a large screening experiment was designed and performed to generate a dataset of screening results that could be characterized and mined. Five trajectories from the set analyzed in the first part of the study were chosen, representing different orbits and booster types; these trajectories are described in Table 1 below. The launch epoch for each of these Earth-centered, Earth-fixed trajectories was varied from 15 days prior to the nominal launch epoch to 15 days after, in one minute intervals. These trajectories were then screened against the catalogue of space objects current at each launch epoch, resulting in 43,200 screenings per launch trajectory, for a grand total of 216,000 screenings from all five trajectories.

Each of the trajectory sets was screened using two different methods. First, the trajectories were screened against the GP catalogue using the Aerospace Corporation's "Collision Vision" GP screening toolset;[3] this is the standard LCOLA approach and the toolset used by Kennedy Space Center (KSC). Because the GP space catalogue lacks covariance information, this utility uses component-based quadratic error growth curves for the appropriate satellite class to estimate component error as a function of time and thus synthesize a GP covariance at the time of closest approach (TCA) between the two objects.[4] The GP catalogue corresponding to the day of the trajectory launch was used; this scenario is slightly operationally unrealistic in that it can result in several hours' back-propagation of catalogue element sets in some cases, but this is not expected to alter the outcome notably. The screening is first conducted for a given keep-out volume around the trajectory (100-km box); for objects that penetrate this volume, a minimum miss distance and Pc are calculated at TCA. The typical value used for the hard-body radius, which indicates the approximate size of the spacecraft, is 10 meters (m); this value was used for the present screening set. The toolset imposes a reporting threshold of 1E-12, so conjunctions with Pc values smaller than this are not reported; occasionally, using this threshold for screening might produce no conjunctions, but this outcome was extremely rare for the GP screenings performed.

Table 1: Launch Trajectories Selected for Screening Experiment

Trajectory #	Satellite	Launch Vehicle	Target Orbit	Range
1	NOAA-N'	Delta II	LEO	Western
2	GPS II-R21	Delta II	LEO (parking)	Eastern
3	STSS Demo	Delta II	LEO	Eastern
4	DMSP-18	Atlas V	LEO	Western
5	SDO	Atlas V	HEO (transfer)	Eastern

The second method, screened the same trajectories and number and time-phasing of screenings but used the Air Force Space Command, Studies and Analysis Division's Computation of Miss Between Orbits (COMBO) astrodynamics standard software against the GP catalogue, with a large screening volume selected (100-km box), to identify all of the conjunctions. The large 100-km volume was selected in order to capture all the conjunctions that would emerge from an SP screening; typically the screening volume would be smaller. Once identified, this list of conjunctions was reprocessed using the SP catalogue; in each case the appropriate Vector Covariance Message (VCM) was retrieved and a covariance-enabled ephemeris was produced in the vicinity of the GP TCA. The VCM is the method of encapsulating all information about the SP state estimate; in this analysis the VCM with the epoch time closest but prior to TCA was used. The VCM-based ephemeris was compared to the trajectory using an interpolation technique to find the TCA and thus for that time calculate the SP miss distance and Pc. The standard two-dimensional Pc approximation was used, following the Foster calculation methodology.[5] This approach is intended to emulate the LCOLA approach run at the Joint Space Operations Center (JSpOC) at Vandenberg AFB, typically used in support of US Air Force launches. As with the Aerospace Corp.'s GP screenings, a 10-m hard-body radius was employed.

There is a mismatch that arises between the results of these two methods for a number of reasons. First, the JSpOC SP catalogue attempts to maintain a certain standard level of quality and as part of doing this requires a minimum amount of observational data in order to perform an SP update. Many objects do not receive this minimum observation and thus do not have an SP vector at all or are represented by an SP that is too old to be usable; as there is no such similar restriction on the GP catalogue, element sets do exist for these same objects. During the periods of time investigated, there were about 1000 such objects—those that are present in the GP but not the SP catalogue—at any given time, so one would expect conjunctions from a GP screening that are not seen in an SP screening. Second, the Aerospace Corp. screening software differs from the GP portion of the approach used for the COMBO screenings. Collision Vision is a different implementation from the COMBO astro standard; it does not use the same version of the SGP4 propagator; and, most significantly, it left-censors the results at 1E-12, whereas the SP results do not introduce any left-censoring (and thus produce a much larger return dataset). Given these factors, each results set (GP and SP) contains conjunctions not found in the other.[*]

This non-alignment of results is addressed in the following way. For investigations that compare SP and GP results directly to determine the degree of difference between the two methods, results for the subset of conjunctions that are found in both results sets are used. For all other investigations (which is most of the analyses of the remainder of this study), individual GP and SP results are each derived from the complete dataset that emerged from that particular approach (either the Collision Vision GP screening set or the COMBO GP-SP screening set). This allows the fullness of each screening approach's results to feed the analyses, restricting this only when it is necessary to have a direct comparison of the GP and SP results. For many of the results sets discussed in this paper, only the GP results are presented due to space limitations.

GP VERSUS SP SCREENING RESULTS

The first investigation seeks to compare the SP and GP Pc and miss distance results arising from the same conjunction events. SP is the more precise theory, and it natively produces a covariance (in contrast to the synthetic GP covariance produced by the Aerospace toolset); so it is typically and reasonably presumed to be the better and more believable result. The GP calculation, however, is a much faster computation and easier to obtain and work with (since the GP catalogue is more easily accessible). Furthermore, it is not necessarily a less reliable result. There is suspicion that the SP catalogue's covariances are often too small, and in fact there is at present an Aerospace Corporation analytical effort to investigate this;[6] if true, then the SP Pc values will at times be too large and at others too small, depending on the ratio of the miss distance to the covariance size. The GP synthetic covariance is not a true *a priori* covariance, that is, the inverse of the normal matrix; but because it is derived from error growth curves themselves constructed from accuracy information from a large number of element sets, it represents a durable error characterization representative of actual state estimate errors. So it is important to see how different the results are from these two approaches, both in a direct comparison and in general behavior over a range of different Pc values.

[*] It must be pointed out that the actual JSpOC process, when it lacks an SP vector for an object (for whatever reason) but does possess a GP solution, uses the propagated ephemeris from the GP solution to represent the position of the secondary during the collision avoidance screening. Because the JSpOC GP solution lacks a covariance, it cannot be used directly to compute a Pc; but it can be deployed to support the general screening activities and, if there is no other alternative, enable window closure decisions through a miss-distance examination.

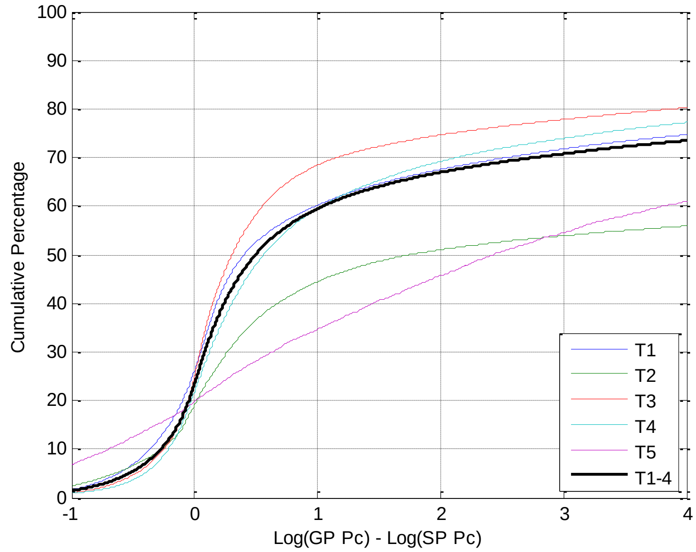

Figure 1: Comparison of GP and SP Pc Values

Figure 1 provides a cumulative distribution function (CDF) comparison of the Pc values for the portion of the experiment's event set that was common to both GP and SP runs; because the graph is dense and may require some orientation to read, an interpretive explanation is given here. One notes immediately the differences among the five trajectories. Three of the trajectories (T1, T3, and T4) clump within a fairly narrow range of behavior; T2's response gives the same general shape but in a manner well apart from the other three trajectories. T5, the one non-LEO trajectory, exhibits a very different behavior and also represents a much smaller response set—only between 3% and 5% of the number of conjunctions reported for each of the other four trajectories. The thick black line represents the CDF results for trajectories 1-4 taken together, as T5 does truly seem to represent a different distribution.

The particular construction of the x-axis of this CDF plot arises from the manner in which one determines whether a change in the Pc value is significant. The rule of thumb in the industry, arising from JSC experience with human space flight orbital safety work and subsequently broadly adopted, is that changes in Pc less than an order of magnitude typically do not result in a substantially altered risk posture and thus are not considered significant. This same standard is a reasonable threshold for defining a significant difference between GP and SP Pc values for the same conjunction event: differences greater than an order of magnitude constitute a significant departure. The actual data plotted are the difference between the base-10 logarithm of the GP Pc and the base-10 logarithm of the SP Pc. A value of zero represents no difference between the two Pc values, values falling within the region of -1 to +1 are within the single order-of-magnitude difference range and represent Pc values that would have a very similar operational interpretation, and values falling outside of this region do constitute significant differences between the two Pc values. A positive value means that the GP Pc is larger (*i.e.*, GP results indicate a conjunction is more likely) and a negative value the opposite (SP Pc represents a riskier scenario).

In interpreting these results, one notes first that 60% of the cases fall within the single-order-of-magnitude range, so a supermajority of cases essentially manifest equivalence between the GP and SP calculations. Of the remaining 40% that do not, nearly all of these cases fall to the right of the +1 boundary of the order-of-magnitude region, meaning that the GP Pc is larger. There is not ideal convergence of the CDF lines in this region—T3 shows 80% of cases to be smaller than four orders of magnitude, whereas T2's results show only 55% of cases to be better than four orders of magnitude; but on the left boundary (at -1) there is a surprising convergence: only about 2% of the cases, for all the main four trajectories and their average, produce an SP Pc that is more worrisome than a GP Pc. These results allow the significant conclusion that GP is a very good LCOLA proxy for SP; GP and SP Pc results do not differ significantly the majority of the time, and when they do, in nearly all of the cases the GP Pc is larger and thus the more conservative value. One drawback to using the GP Pc is that, because it is conservative, it is likely to result in unnecessary launch window closure activity.

Figure 2 provides an accompanying diagram of miss distance differences. The x-axis for this graph is the simple difference between the SP- and GP-computed miss distances for the paired conjunction events. A positive value indicates that the GP miss distance is smaller and thus (in general) a greater collision threat, making the GP calculation more conservative. Here one observes a fairly even distribution of the miss distance differences, unlike the Pc distribution shown in Figure 1. While in general the GP miss distance is smaller (60% of the cases) and thus more conservative, it is not at all the lopsided result observed for Pc differences, in which 80% of the cases GP is more conservative and in only 2% of the cases is SP more conservative in a significant way. This miss distance result is consistent with previous studies that have documented the relative lack of correlation of miss distance with Pc and thus with conjunction risk.[7]

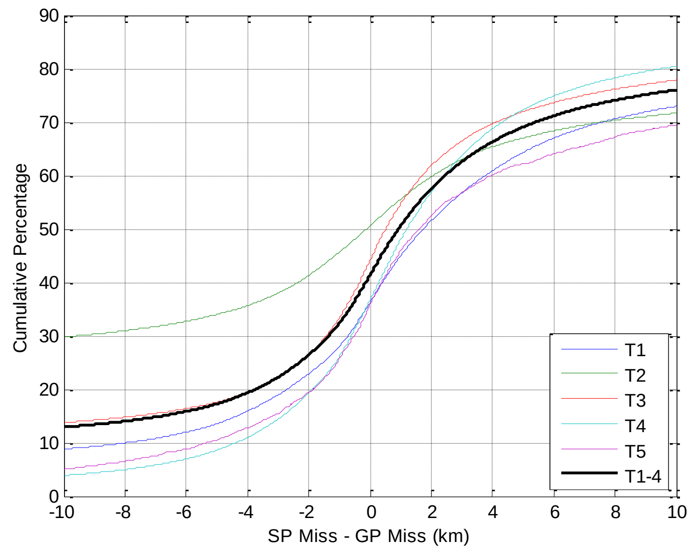

Figure 2: Comparison of GP and SP Miss Distances

Having established that GP Pc is a reasonable proxy for SP Pc, from this point forward all analyses consider GP and SP results separately, drawing from the two different executions of the screening experiment, so that behavior and thresholding for these different calculation types can be independently established. For the purposes of the present paper, only the GP results will be presented due to space limitations. Readers are referred to the full study report[2] for a presentation of the SP results and additional GP result permutations.

LAUNCH WINDOW CLOSURE VERSUS P_C

Satellite launch execution is managed through periods of time called *launch windows*. A launch window is a continuous period of time during which the payload, if launched, can obtain the desired orbit and is thus secondarily the time during which all launch support systems are fully prepared for a launch. For some final orbits, the launch window can be exceedingly long, as there are few, if any, time-based constraints that would prevent the vehicle from reaching the intended orbit. For other launches, especially space probe launches with a trajectory design that requires the gravity-assist of one or more additional luminaries, the launch opportunity may be a single instant, which, if missed, could force a launch delay of months or years. The LCOLA task is to determine which portions of a launch window are safe for launch activities and which are not.

The LCOLA approach differs in important respects from the on-orbit CA paradigm. When a potentially threatening conjunction is identified between two on-orbit objects, a multi-day process is initiated to acquire additional tracking data, refine the orbit(s), assess the risk through a multiplicity of different amplifying investigations, decide whether an avoidance maneuver is needed, plan the maneuver, and, updating the information as close as possible to the maneuver commit time, decide whether to execute the maneuver. When an LCOLA screening for a particular launch time indicates a potentially dangerous situation, there is no follow-up investigation to perform additional risk analysis; that launch time is simply removed as one of the viable launch times within the launch window. Typically, only certain time intervals within a launch window are proposed as candidate launch times. The majority of launches today are "launch on minute," meaning that within the launch window only the integer minutes (*i.e.,* once every sixty seconds) are actual candidate launch times; this simplifies the LCOLA screening process, as only those particular times need to be screened. In the present paper, only the "launch on minute" scenario, which is the most common, is presented. In all cases, the initial evaluation criterion for the desirability of a launch window is the percentage of the launch opportunities that are viable.

A useful contributor to the choosing of a Pc screening threshold is its expected effect on the portions of the launch window that will need to be closed or can be left open. This is not to say, of course, that the selection of a screening threshold should be merely for the convenience of launch operations, *i.e.,* choosing an easy-to-meet threshold so that the majority of candidate launch times within a launch window will be available so that launches can be conducted with as little encumbrance as possible. At the same time, it only makes sense as part of the threshold selection process to consider the effect of certain threshold choices on the ability to conduct launch operations. Risk mitigation is not a threshold entity but a continuum: as one makes the threshold more demanding, risk is reduced and operational encumbrance increases; but there is rarely a fixed reason to make any particular choice along the continuum. In the end, the selection of a threshold will be a subjective assessment of the desirable balance of risk mitigation and operational fluidity.

To this end, an investigation that attempts to establish the effect of choosing different screening thresholds on launch window closure is a helpful contribution to the overall threshold selection process; and the current screening experiment has produced data that are quite useful to such an investigation. Four of the five trajectories screened in the experiment are for LEO objects and are reasonably similar in behavior; thus, if one concatenates them into a single dataset, one will thus have 172,800 screening results at one-minute intervals, with only four seams in the dataset (at the points of transition from one trajectory to the next). Each such screening result will have a maximum Pc, a minimum miss distance, and a cumulative Pc.

One can thus construct launch windows of any desired duration and, through "moving window" shifting of the window through the entire dataset, examine statistically the portion of such a window that is open or closed. For example, suppose one were to choose a launch window of twenty minutes' length. The first twenty records (records 1-20) of this dataset could be used as a possible set of screening results for this window, and the percentage of these twenty records with a Pc smaller than a chosen threshold would represent the percent of the launch window that was "open" (for the "launch-on-minute" case). One could then shift this window down one record (examining records 2-21) and calculate this percentage again, and so on until the entire set of 173,000 records has been examined. One can then take percentile points from the result and thus report at, say, the 50th, 68th (1-sigma analogous equivalent),* 95th (2-sigma analogous equivalent), and 99th (3-sigma equivalent of sorts) percentiles the "percent open" of the twenty-minute launch window.

If this is performed for a large number of different window sizes and Pc thresholds, contour plots can be developed that give the percentage of launch window open as a function of launch window size and Pc threshold. This type of presentation can communicate in a more comprehensive way the impacts to launch mission operations of choosing different Pc thresholds; Figure 3 uses this approach to report results for the full GP screening set, arranged in quad chart format. For each of the quad sub-plots, the x-axis is the Pc threshold; the y-axis is the length of the launch window, in minutes (ranging from one minute to three hours); and the color indicates the percentage of the launch window that is open (blue indicates 0% open and dark red 100% open). Each quad subplot is a different percentile level, representing the four percentile levels above. In this presentation only the results for GP, max Pc, "launch on minute" screenings are given; but in the full report results are provided for maximum and cumulative Pc thresholds, both GP and SP screenings, and different densities of viable launch points within a window (launch on minute, launch on 30 seconds, and launch on second).

In examining the plots, one observes a discernible "chatter" along parts of certain contour lines, which produces localized instances of non-monotonic behavior. This is due to the taking of percentile points against short(er) launch windows which themselves are composed of discrete screening results described in a binary fashion of being either open or closed. While this chatter is not a major impediment to interpreting the graphs, it merits a more thorough explanation. The following example helps to illustrate the origin of the phenomenon.

One begins by postulating a launch window of one minute's duration (for this entire discussion "launch-on-minute" is assumed). Because only one screening event is incorporated in such a window, there are only two possibilities: it is either 0% open or 100% open. A launch window of two minutes' duration will offer three possibilities: 0% open, 50% open, or 100% open. A three-minute launch window will offer four: 0%, 33%, 67%, and 100%. As the size of the launch window increases, the set of possibilities will become more continuous, but for the smaller launch window sizes, there is a relatively small number of discrete levels for the "percent open" evaluation.

Table 2 gives a hypothetical set of 100 screening results (chosen specifically to illustrate these particular difficulties), 80% open (indicated by a 1) and 20% closed (indicated by a 0) when evaluated at a particular Pc threshold and arranged in the following way, essentially a 1-1-0 repeating sequence of open and closed opportunities until record 85, after which point the window is left entirely open.

---

* Because the range of values is bounded (0% to 100%) and the extremes of this range are encountered, the distribution is unlikely to have Gaussian properties; and even without this there is no real reason to suppose that the dataset would behave in a Gaussian manner. However, using percentile points that emulate those encountered for the 1-, 2-, and 3-sigma cases allow easier comparison with confidence levels typically encountered in statistical inference.

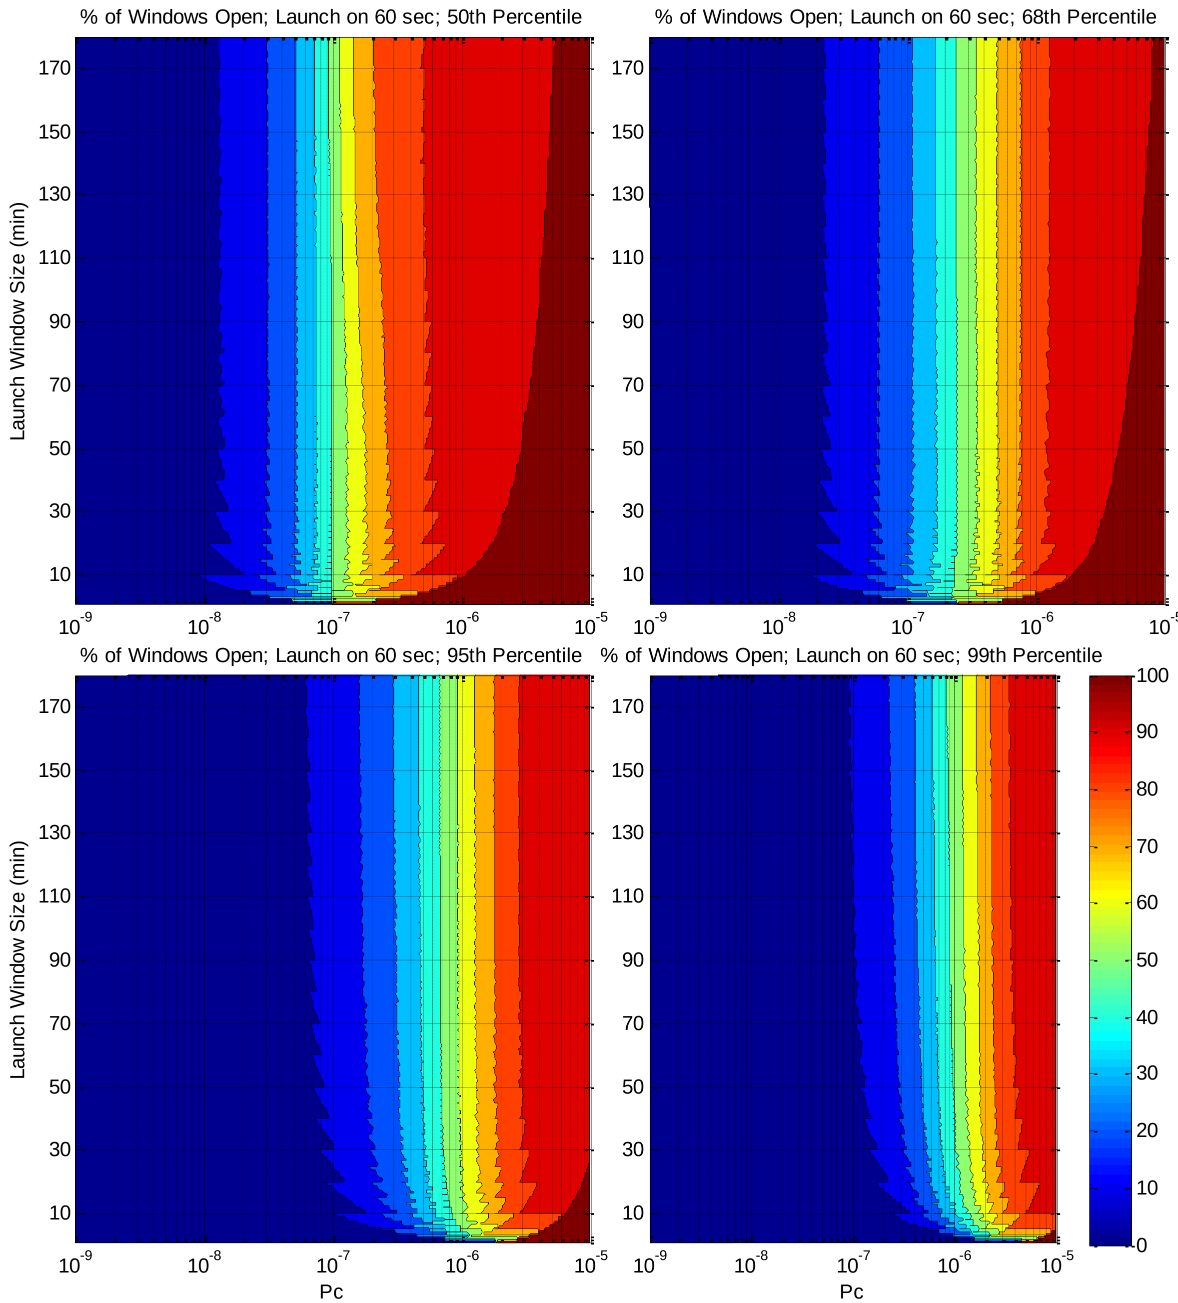

Figure 3: Window Open/Closed Percentage vs. Launch Window Size vs. Pc; GP Screenings, 60-sec Interval ("Launch on Minute"), Maximum Pc

For this hypothetical results set, moving windows of 1, 2, and 3 minutes' duration were run, producing groups of 100, 99, and 98 results, respectively. In summarizing these results with the percentile points suggested earlier (50^{th}, 68^{th}, 95^{th}, and 99th) for window sizes of 1, 2, and 3 minutes, the following summary values are observed:

Table 2: Hypothetical set of 100 screening results

1	1	11	1	21	0	31	1	41	1	51	0	61	1	71	1	81	0	91	1
2	1	12	0	22	1	32	1	42	0	52	1	62	1	72	0	82	1	92	1
3	0	13	1	23	1	33	0	43	1	53	1	63	0	73	1	83	1	93	1
4	1	14	1	24	0	34	1	44	1	54	0	64	1	74	1	84	0	94	1
5	1	15	0	25	1	35	1	45	0	55	1	65	1	75	0	85	1	95	1
6	0	16	1	26	1	36	0	46	1	56	1	66	0	76	1	86	1	96	1
7	1	17	1	27	0	37	1	47	1	57	0	67	1	77	1	87	1	97	1
8	1	18	0	28	1	38	1	48	0	58	1	68	1	78	0	88	1	98	1
9	0	19	1	29	1	39	0	49	1	59	1	69	0	79	1	89	1	99	1
10	1	20	1	30	0	40	1	50	1	60	0	70	1	80	1	90	1	100	1

Table 3: Results summary for hypothetical screening results

Window Size	50ile	68ile	95ile	99ile
1 minute	100%	100%	0%	0%
2 minutes	100%	50%	50%	50%
3 minutes	67%	67%	67%	67%

At different percentile points, one sees very different percentage results and different trends as the launch window is lengthened: at the 50^{th} percentile, the results are monotonically decreasing; at the 95^{th} and 99^{th} percentiles they are monotonically increasing; and at the 68^{th} percentile they follow a non-monotonic behavior. As the window size is increased, the impact of this behavior is lessened (and further smoothed by the graphics generation routines that produced Figure 3); but the basic problem of chattery results brought about by percentage evaluation of relatively small binary sets, coupled with percentile-point summary of larger groups of these percentile results, remains a feature of the overall dataset. As stated previously, this is not really an impediment to interpretation if one wishes general information by region, which is all that should appropriately be extracted from such a graph in any case.

Because the charts can be somewhat confusing to read, the following is given as a worked example, starting with the upper-left quad of Figure 3. If one chooses, say, 1E-06 as the Pc threshold value (x-axis) and a 70-minute launch window (y-axis), the corresponding color for that locus is red but not dark red, indicating that with those parameters about 90% of the launch window is open 50% of the time (because the upper-left quad is the 50^{th} percentile graph). At the 95^{th} percentile (bottom-left quad), for these same parameters perhaps about 55% of the window is open. Thus it can be said that, for a Pc threshold of 1E-06, half of the time 90% of a 70-minute window is open and nineteen out of twenty times 55% of it is open. Moving left to right along the 70-minute (horizontal) line, the percentage grows from 0% to 100% open, which makes sense in that as the Pc threshold is larger (*i.e.*, more lax), more of the launch window is open. Moving vertically on the 1E-06 line, one sees relatively little variation in the percentage open value—no more than ten percentage points or so. The results of the screenings seem to be more or less uniform this way, especially at the higher percentile points: once a launch window is of a reasonable size, the sampling appears to be large enough that the percentage closed does not vary as the launch window size is further increased.

The "U-shaping" of the contour plots is an expected result. In the upper-left quad of Figure 3, the dark blue and dark red actually meet along the launch window size = 1 line (right at the x-axis), without intermediate colors. Because with a one-minute size in a launch-on-minute situation the window must be either entirely open or entirely closed, it makes sense that only two colors appear for it; the transition takes place at a Pc threshold of about 1.5E-07. As the window size is increased, there is more room for different percentage values (rather than just 0% or 100%); and the U-shape moves quickly to the largely parallel contours that begin at launch window sizes of about 10 minutes.

It is very well to discuss verbally the response that is visible from the graphs themselves, but more helpful is an attempt at interpretation; and for this some guidance is needed on what would constitute acceptable operational performance with regard to launch window size and the portion of such windows that are open. The length of the launch window is governed by the orbit trajectory (and thus indirectly by destination) and by deployment issues specific to the spacecraft, such as trying to avoid earth-shadow or other unfavorable environmental conditions during component deployment. For some space probe launches, an "instantaneous" launch is required, meaning there is only one acceptable launch time, and for such situations there is no real issue of a launch window; but for more typical NASA launches windows of tens of minutes are possible, and launch operators are happier with an hour or so, becoming uncomfortable when the window shrinks to about twenty minutes. For USAF launches, the windows are longer, typically on the order of one to two hours; and it would be nice to preserve that length if possible. Fortunately, for most Pc thresholds, the response as a function of launch window size levels out quickly into straight contours; so choices that are appropriate for a twenty-minute window will probably be more or less appropriate for a 120-minute window and vice versa.

In terms of percentage of window open, it is typical to conduct operations with windows that are mostly open—80% or so; below this level operators begin to feel cramped, and this feeling becomes acute when the percentage decreases to the 50% level. Approximate Pc values, taken from Figure 3, for three different open-window percentages are given in Table 4 below. These values are for the 95th percentile and represent the situation over all but the smallest launch window sizes.

Table 4: Pc Values Corresponding to Open Window Percentages (approximate)

Percent of Window Open	GP Launch-on-Minute Pc, 95th Percentile
80%	2E-06
50%	7E-07
~33%	2E-07

If one wishes to remain within the more conservative levels for window closure, then screening thresholds that preserve an 80% "open-window" level seem appropriate. Such levels are already notably more demanding than the general US Air Force guidance of 1E-05. The next level of increase shown is to move all the way to the 50% level, a level at which one would expect to encounter some resistance from the launch community. However, in some respects more important than the overall open/closed status of a particular launch window is the ability to find, in any smallish sector of the window, some reasonable number of viable launch opportunities. For example, typically the initial launch opportunity is at the beginning of the launch window. If that opportunity is waved off due to a launch exigency, then perhaps an opportunity will be sought around the middle of the window; and if for some reason that opportunity also cannot be used, then one towards the end of the window will be selected. Since the contours in Figure 3 are reasonably straight, the "percent open" value for a given Pc contour (straight vertical line) is more or less constant until the extremely small windows are encountered; but even there the deviation is not substantial. This means that sub-parts of any given launch window are likely to have more or less the same "percent open" value as the overall window. For example, consulting Table 4, one can see that for the GP case the Pc threshold corresponding to an 80% availability figure is approximately 2E-06; this value can be taken from the bottom-left quad of Figure 3 at the 60-minute launch window point. However, one can additionally see that in descending to the 20-minute window size the openness level drops by perhaps only five percentage points: if the 60-minute window were to be divided into thirds, each of twenty minutes' duration, there would be about a 75% launch opportunity availability in each of these twenty-minute sectors.

If this is the case generally, then it may be possible to entertain more demanding Pc screening thresholds than this analysis shows *prima facie*. If the true concern is to ensure that a reasonable launch opportunity can be found in each of the sub-sections of a launch window, and it appears that availability within a sub-section is only slightly lower than that for the entire window, then a Pc that maps to an overall window-availability figure that is somewhat below what historically would be considered acceptable can perhaps be embraced. Given this dynamic, the Pc values that map to the 80% overall figure should certainly be acceptable without reservation; and it would seem that Pc values as low as those associated with 50% availability may be viable as well.

MISS DISTANCE SCREENINGS AS PROXY FOR P_C SCREENINGS

The only durable single calculation to assess the risk of any given conjunction is the probability of collision, or Pc. This calculation incorporates all of the relevant factors in that in uses the nominal miss distance, the size of the covariances of the primary and secondary satellites' state estimates, and the estimated combined sizes of the spacecraft; and it returns an actual probability as its result—a calculation appropriate to risk assessment. However, there are times when this calculation is not possible, such as when covariance information is not available and cannot reasonably be synthetically constructed. In such cases, it is a common practice simply to conduct conjunction screening and perform risk assessment using only the miss distance between the two objects at TCA. Indeed, with this approach there is rather little "risk assessment" that can be performed: either the conjunction has a miss distance less than the screening threshold, or it does not; but the overall directness of the approach—simply examining the minimum miss between the two objects—is quite appealing, and the fact that covariance information is not required promotes it all the more. Because of the interest in this alternative technique, not just for those cases where key information is not available for the Pc calculation but as a simplified procedure overall, it is useful to provide some analysis to indicate what can be done with a miss-distance-based paradigm and what cannot.

The first technical point that must be raised is that the miss distance cannot serve as a full substitutionary datum for the Pc: the correlation between the miss distance and the Pc is actually rather weak, or at best one-sided. While it is generally true that a smallish miss distance is required for a largish Pc, the logical entailment is only in one direction: a smallish miss distance may be a necessary condition (most of the time) for a largish Pc, but it is not a sufficient condition. The presence of a large covariance, and therefore uncertainty, about one or both of the satellites means that the miss distance (which is based merely on an estimate of the mean position for both objects) is not strongly determined; and if the computed miss distance is small, then the random selection of position points within an error ellipse about both objects would tend to generate a set of much larger miss distances, lowering the risk. If one uses a miss-distance-based screening, many—probably most—of the conjunctions will have a much smaller risk profile if evaluated with a Pc-based approach; so a large number of "false alarms" will be generated from such screenings. Furthermore, it is difficult to know what sort of risk level to ascribe to a conjunction with a given miss distance. By changing the associated covariance incrementally from a very small size to a very large one and calculating the Pc at each increment, one can obtain the maximum possible Pc for this miss distance[8]; but most of the Pc data associated with this miss distance will be much smaller than that, and some values will be essentially negligible. So when miss distance alone is used as a risk assessment criterion to determine launch window closures, it is not clear what risk level is actually being invoked.

However, while there may be no direct link between risk as evaluated by Pc and the miss distance, "probability profiling" of the miss distance against the Pc can be accomplished and is in fact rather straightforward; and the results of such profiling can be used to make more definitive statements about risk from the results of a miss-based evaluation. Such a profiling is conducted by collecting a large number of events with a Pc greater than or equal to a given threshold. The miss distance values from all of these events are then collected and their distribution characterized, either by a CDF or by percentile points. This approach can provide the context needed for miss-distance results to be mapped directly to a risk level, as one can now make a statement of the form "screening at a miss distance of x will identify 95% (or some other percentile point) of the events that would present a Pc risk of y." Because the higher percentile points (such as 99th, 99.9th, 99.99th percentile, &c.) are not well defined for the finite samples available for such profiling, if one wishes to push into the higher end of the distribution, the best approach is to fit a candidate distribution to the empirical distribution and use this candidate distribution, which has an analytical representation, to estimate the behavior of the empirical distribution at very high percentile points. This is also necessary for the present case because the experimental screenings were run with a 100 km x 100 km x 100 km miss-distance box and thus cannot report empirical data for miss distances greater than this.

The above rubric was followed with the experimental dataset of screening results. These results were arranged into groups (with replacement) of individual events in which the Pc was greater than 1E-05, 5E-06, 1E-06, &c., up to 5E-09. For each, a CDF was constructed; and additionally a two-parameter lognormal and two-parameter gamma distribution were fitted to each CDF, using maximum likelihood estimation for the distribution parameters. The two-parameter lognormal distribution has a probability density function (PDF) of the form

$$f(x;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}x}\exp\left(\frac{-[\ln(x)-\mu]^2}{2\sigma^2}\right) , \tag{1}$$

and parameter estimation is straightforward if one performs a logarithmic change of variable and estimates the mean (μ) and variance (σ) in logarithmic space, using standard methods. The two-parameter gamma distribution has a PDF of the form

$$f(x;\beta,m) = \frac{1}{\beta^m\Gamma(m)}x^{m-1}\exp\left(\frac{-x}{\beta}\right) , \tag{2}$$

in which β is the scale parameter and m is the shape parameter; the maximum likelihood estimation equations for these parameters are given in standard texts on distribution parameter estimation.[9] Figure 4 shows the three curves (CDF plot of actual data, lognormal PDF fit of data, and gamma PDF fit of data) for each of the Pc thresholds enumerated above. As can be seen, for nearly all cases the gamma fit is excellent—good enough that a formal goodness-of-fit test was judged not to be necessary, given the approximate nature of the enterprise. The quality of this fit is not entirely surprising, as Chan and others have remarked that the miss-distance distribution for a single event should follow a non-central chi-squared distribution,[10] which is part of the gamma distribution family; so the two-parameter gamma distribution suggested itself as a candidate from the beginning. Table 5, which appears after the figures, gives the actual miss distance values at the indicated percentile points for the Pc levels investigated.

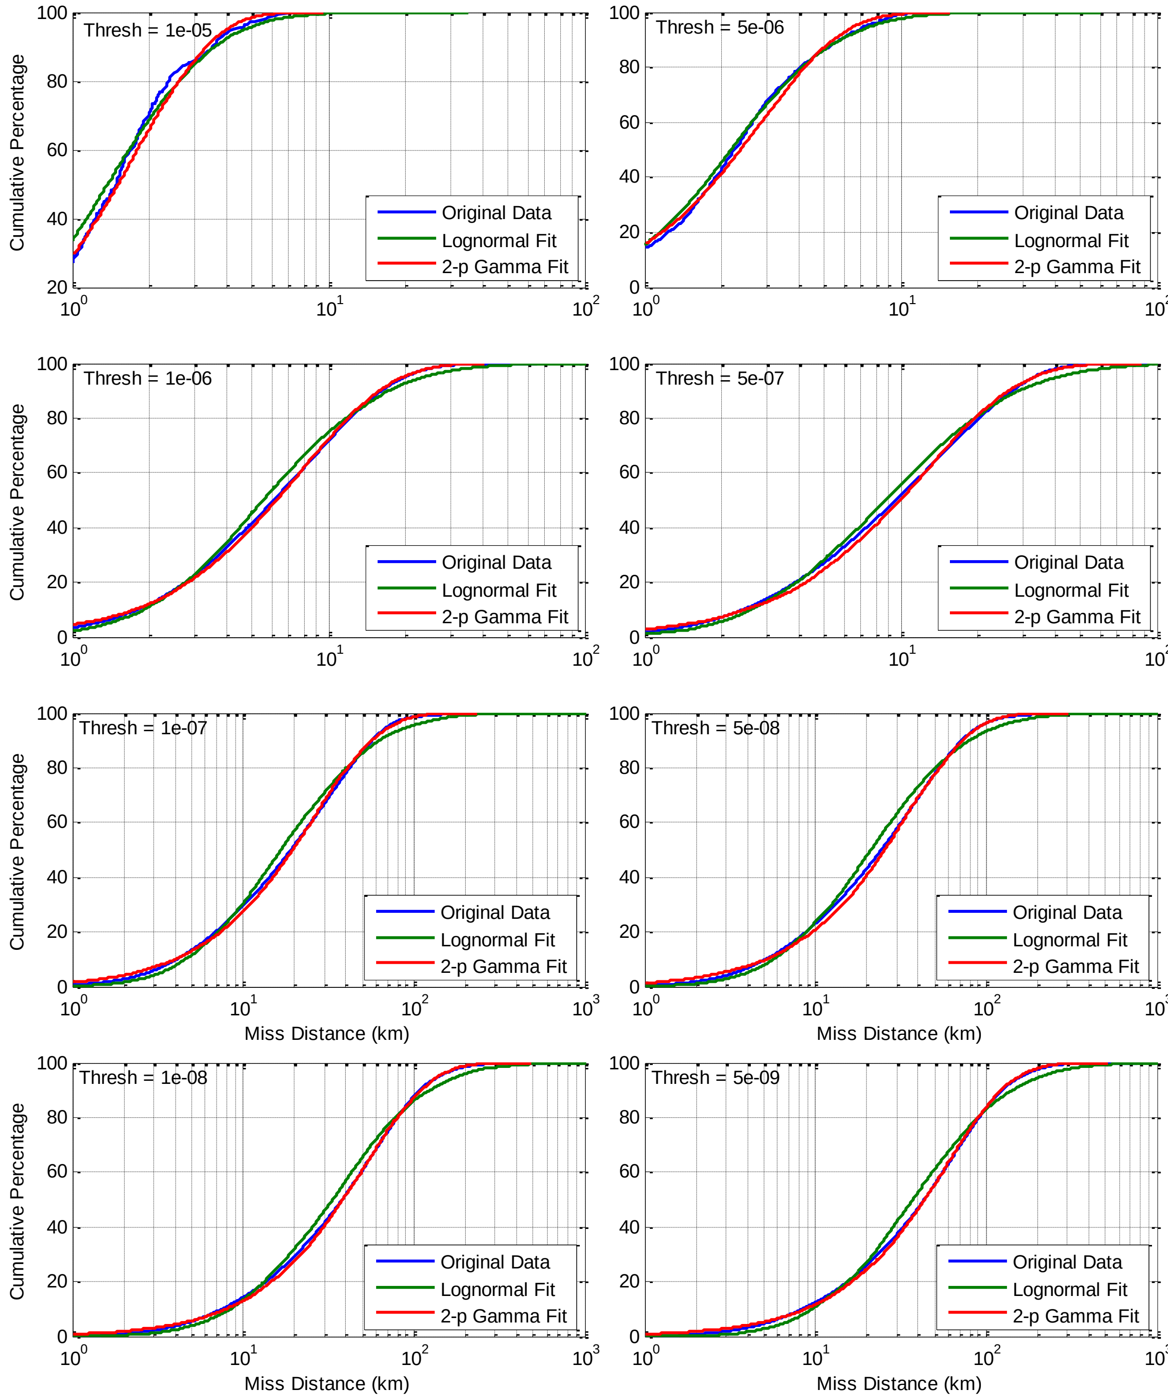

Figure 4: GP Miss Distance Distributions and Fits to Lognormal and Gamma Distributions for a Range of GP Pc Thresholds

Table 5: GP Miss Distance Screening Equivalents (km) for given Pc Values at Various Confidence Levels

	GP Pc							
Percentile	1.0E-05	5.0E-06	1.0E-06	5.0E-07	1.0E-07	5.0E-08	1.0E-08	5.0E-09
50	1.5	2.4	6.4	9.9	19.8	25.2	38.3	43.5
68	2.1	3.3	9.2	14.8	30.5	38.8	59.0	67.1
95	4.0	6.7	19.7	32.9	72.2	92.2	140.1	159.8
99	5.6	9.3	28.1	47.9	107.1	136.9	208.0	237.6
99.9	7.7	12.9	39.8	68.6	156.3	199.9	303.7	347.1
99.99	9.7	16.4	51.2	89.1	205.1	262.4	398.5	455.6
99.999	11.7	19.9	62.5	109.3	253.5	324.5	492.8	563.5
99.9999	13.6	23.3	73.7	129.5	301.7	386.3	586.7	671.0

Tables such as Table 5 can present a temptation to the CA risk analyst: one can perform a simple screening based on miss distance, free of the data-release hassle of covariance-enabled state estimate data and cognizant that it would be essentially as effective as the more cumbersome Pc calculation approach. Additionally, there are occasional situations in which the Pc approach, due to the lack of the proper data, is simply not possible; and in such cases, this sort of table can be helpful. However, as remarked earlier, the false-alarm rate associated with this approach is quite high; and such a quantity of false alarms will have a significant effect on the ability to keep launch windows open. To illustrate the difficulty, a quad-chart presentation similar to that used in the previous section, giving window open percentage as a function of window size and in this case miss distance, was generated and appears as Figure 5.

These charts make clear very quickly that screening at anything close to a largish miss distance will have very serious effects on launch window availability. Looking back at Table 5, one observes that for a Pc equivalent screening of 1E-06, a miss distance screening of 19.7 km is needed to meet a 95% confidence value and a miss distance screening of 28.1 km to meet a 99% confidence level. In taking these numbers to Figure 5, for the 19.7 km figure, the launch window would be 50% open at the 50th percentile and only about 25% open at the 95th percentile. If one wishes the 99% confidence level and moves to the 28.1km figure, one now obtains windows 25% open at the 50th percentile and only 10% open at the 95th percentile—values that are clearly operationally unacceptable. Moving to either a more demanding equivalent Pc threshold or a higher confidence value brings the situation to a point where the launch windows are almost entirely closed.

In summary, it is possible through this type of profiling to develop miss-distance screening threshold equivalents to Pc-based screening thresholds; but the operational use of these thresholds will come at a heavy price, namely a very large number of screening false alarms that quickly closes unacceptably large portions of launch windows. In occasional circumstances it may be necessary to perform screenings with this approach, but clearly such situations need to be rare.

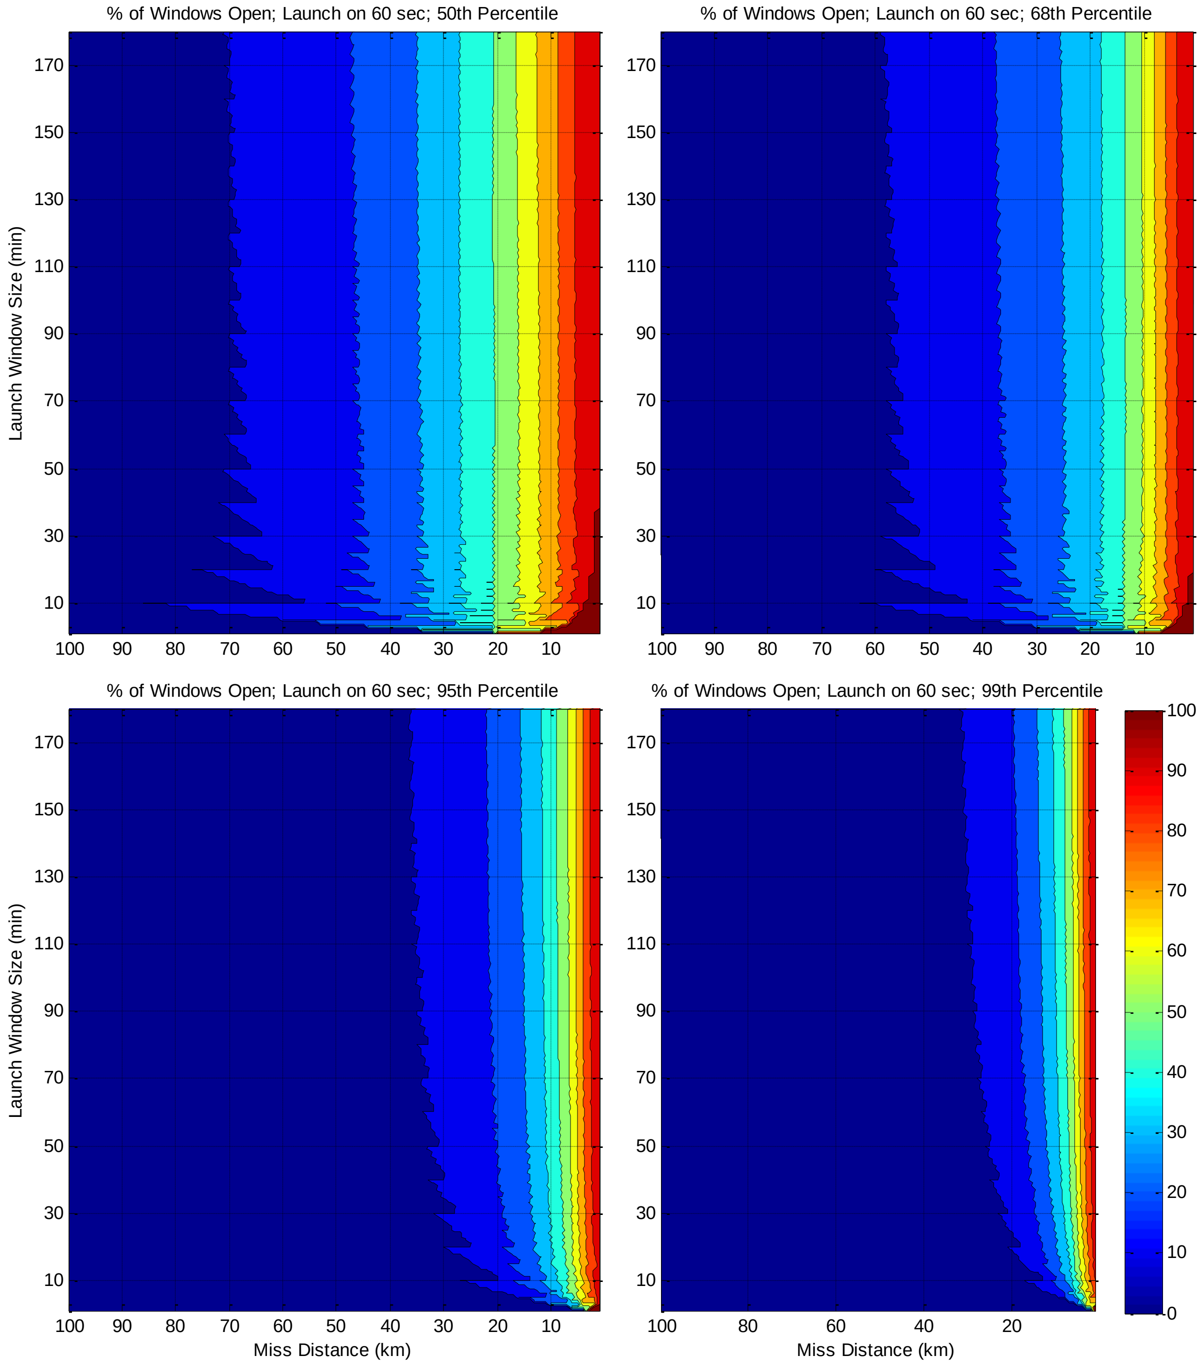

Figure 5: Window Open/Closed Percentage vs. Launch Window Size vs. Miss Distance; GP Screenings, 60-sec Interval ("Launch on Minute")

OVERALL RISK POSTURE IMPROVEMENT WITH LCOLA ACTIVITIES

A question that is perpetually present to the collision avoidance and risk assessment enterprise is whether all of the effort, attention, and cost of these activities produce any tangible benefit. There have been a few verified collisions, to be sure; and some have even resulted in loss of mission for one of the colliding objects. But the number of collisions over 50 years of space surveillance history is actually very small, the "sky" is relatively large, where one sets the Pc screening threshold is essentially an arbitrary decision, and remediation of even high-risk potential collisions can be waived by the proper authorities. All of this has led some observers to question the value of the entire process. What is needed is a way to evaluate the degree to which LCOLA practices improve the overall collision risk posture.

One approach to this that suggests itself is to try to determine the inherent risk of conducting launch operations with a "big sky" posture: essentially doing nothing at all about collision risk determination and abatement. Any given launch event will sustain a certain amount of collision risk, and two such events will sustain the cumulative risk of either launch sustaining a collision. If one could calculate this cumulative risk over a given number of launches, this risk could be compared to the risk level (*i.e.,* the Pc) at which screening and risk assessment activities are conducted. If risk assessment is calculated at a lower Pc (a smaller number) than the overall risk level experienced by doing nothing, then an improvement in the risk posture has been achieved—the screening activity is finding and remediating conjunctions with an even smaller likelihood than the latent risk of a certain number of launches. If, however, one is screening at a higher Pc (larger number) than the background risk, then one may identify a few high-probability events but has not appreciably improved the risk posture over simply doing nothing for conjunction assessment on the theory that the sky is "big."*

The following methodology can be used to determine the cumulative background risk of collision during a launch event as a function of the number of launch events. This approach requires the assumption that any single launch screening from the large screening experiment (that has provided all of the data for the last several sections of this paper) can be used to represent a typical screening result for any given launch. One draws 600 such events at random (with the number 600 representing the maximum launch density of 24 launches per year over a 25-year period) from the screening database (with replacement) and calculates the set of cumulative Pc values determined by sequentially stepping through the sample of 600: Pc of trajectory 1, Pc of trajectories 1 and 2 taken together, Pc of trajectories 1, 2, and 3 taken together, &c. This will yield a curve of the running cumulative probability of a launch-related collision as a function of the number of launch events. Such a cumulative probability can be computed in two ways: using the maximum Pc from each screening and convolving all of these into a cumulative probability ("max cumulative") or using the cumulative probability from each screening and convolving all of those into an overall cumulative probability ("cumulative cumulative"); Figure 6 gives results for both approaches. This experiment is then conducted as a series of 100,000 Monte Carlo trials of this calculation, and the Pc median and variability at each sequential point (1 through 600) is tabulated. The result is a series of curves giving the cumulative Pc as a function of number of launches. The number-of-launches datum can be turned into a time-frame if one will postulate the average number of launches expected to be sustained each year. Table 6

* To add precision to this statement, screening at a Pc below the cumulative lifetime latent Pc for a satellite certainly improves the risk posture, and the ratio of the two Pc values (lifetime Pc to Pc screening threshold) is a reasonable approximation of the improvement in risk posture. However, screening at a level greater than the lifetime Pc is not without any merit, as those high-probability events are tacitly included in the lifetime Pc calculation.

provides a matrix of these conversions, allowing the reader to specify a launch event frequency (per year) and a span of time over which LCOLA operations will be in place and producing the x-axis point (in terms of number of events) for this particular set of parameters.

Table 6: Matrix relating Sequence Number to Launches/Year and Time Span

Launch Events	Years of Launch Operations						
per Year	2	5	10	15	20	25	30
4	8	20	40	60	80	100	120
8	16	40	80	120	160	200	240
12	24	60	120	180	240	300	360
16	32	80	160	240	320	400	480
20	40	100	200	300	400	500	600
24	48	120	240	360	480	600	

To make the graphs somewhat more approachable, the following is a walk-through of how to interpret Figure 6. It is necessary to begin with an x-axis value of interest, so let this example suppose a twenty-year period of launch operation with twelve launch events per year; per Table 6 this equates to a total of 240 launch events, so this is the x-axis point to use. Because it is greater than 50, the subplots on the right side are the ones with the appropriate denominations. If one screens using the maximum Pc event within each screening (as opposed to the cumulative Pc for each screening) as the assessment of risk, then the upper-right subplot is the one of interest. One sees that for 250 total launch events, there is a cumulative risk of collision of about 1E-04 at the median value, with an overall 5^{th}-95^{th} percentile span running from 8E-05 to 1.5E-04—a fairly narrow range actually. If one thinks the 12 launches per year to be too aggressive, the same x-axis point can be reached if one supposes 8 launches per year and a 30-year launch operation period.

One sees that to screen at 1E-04 would leave one in essentially the same risk posture as doing nothing for LCOLA at all. To screen at a max Pc value of 1E-06, which seems to be a number that could not unreasonably be adopted given the preceding analyses, would yield a risk posture two orders of magnitude better than doing nothing for LCOLA. To this author’s thinking, this is a handsome payoff for the relatively minor investment in LCOLA screening activities and non-excessive closure of launch windows and thus disruption of launch operations.

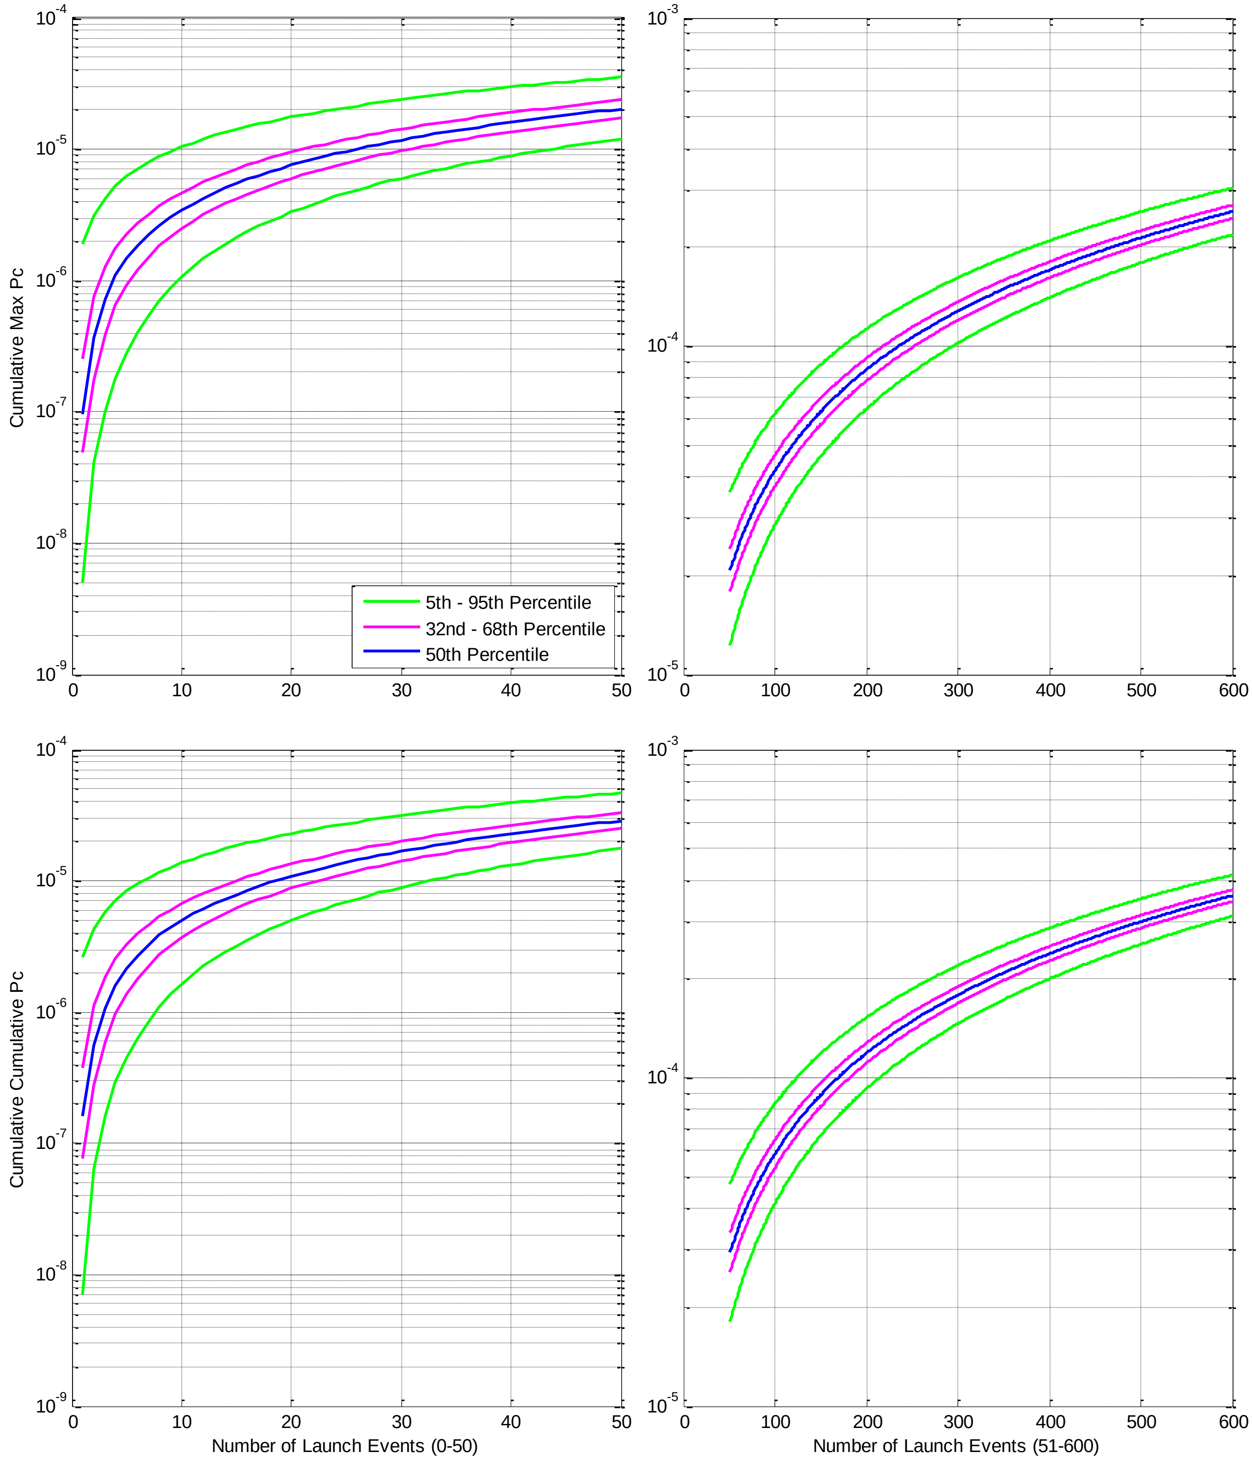

Figure 6: GP Cumulative Max and Cumulative Cumulative Pc as a Function of the Number of Launch Events

CONCLUSIONS

The present study has answered certain key questions, which should allow a standardized approach to LCOLA conduct to be assembled. GP screenings are a reasonable conservative proxy for SP screenings, as they produce equivalent results a supermajority of the time and conservative results for nearly all of the remainder. Plots of launch window closure as a function of Pc screening level and desired size of window can be constructed and are a good way of understanding the effects on LCOLA operations of choosing a particular Pc threshold. Miss distance proxies for Pc screening levels can be calculated from empirical data and extended to more extreme percentile points through curve-fitting, but conducting regular screening operations with this approach is heavy-handed and results in an unacceptable amount of window closure. Finally, the overall value of LCOLA operations can be assessed by comparing the risk posture produced by reasonable screening levels to that encountered were one to do nothing at all, and reasonable levels of screening produce a risk posture two orders of magnitude superior to a quiescent "big sky" approach—a more than adequate justification for a relatively modest LCOLA investment.

REFERENCES

[1] Hejduk, M.D., Plakalovic, D., Newman, L.K., Ollivierre, J.C., Hametz, M.E., Beaver, B.A., and Thompson, R.C. "Trajectory Error and Covariance Realism for Launch COLA Operations." 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI. February 2013.

[2] Hejduk, M.D., Plakalovic, D., Newman, L.K., Ollivierre, J.C., Hametz, M.E., Beaver, B.A., and Thompson, R.C. "Launch COLA Operations: Recommended Procedures and Thresholds." NASA/GSFC, /KSC, and /JSC Joint Technical Report, forthcoming.

[3] Gist, R.G. and Oltrogge, D.L. "Collision Vision: Covariance Modeling and Intersection Detection for Spacecraft Situational Awareness." 1999 AAS/AIAA Space Flight Mechanics Conference, Girdwood, Alaska, August 1999. See also Oltrogge, D.L. and Gist, R.G. "Collision Vision: Situational Awareness For Safe And Reliable Space Operations" Paper IAA-99-IAA.6.6.07, 50th International Astronautical Congress, Amsterdam, The Netherlands, 4-8 Oct 1999.

[4] Peterson, G.E., R. G. Gist, and D. L. Oltrogge. "Covariance Generation for Space Object Using Public Data." 2001 AAS/AIAA Space Flight Mechanics Meeting, Santa Barbara, CA, 11-15 February 2001.

[5] Foster, J.L. and Estes, H.S. "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles." NASA/JSC-25898 (August 1992).

[6] Cerven, W.T. "Covariance Error Assessment, Correction, and Impact on Probability of Collision." AAS/AIAA Space Flight Mechanics Meeting (San Diego, CA), February 2011.

[7] KSC Report ELVL-2008-0040593.

[8] Alfriend, K.T. et al. "Probability of Collision Error Analysis." *Space Debris*, Vol. 1, #1 (1999), pp. 21-35.

[9] Cohen, A.C. and Whitten, B.J. *Parameter Estimation in Reliability and Life Span Models*. New York: Marcel Dekker, Inc. 1988.

[10] Chan, F.K. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press, 2008.