Near-Elimination of Airspace Disruption from Commercial Space Traffic Using Compact Envelopes

Thomas J. Colvin[*] and Juan J. Alonso[†]

Stanford University, Stanford, CA 94305

Traditional methods for safely integrating space launch and reentry traffic into the National Airspace System (NAS) use hazard areas (i.e. no-fly zones) that restrict aircraft from larger areas, for longer times, than are necessary. We have previously proposed a new class of hazard area called *compact envelopes*,[1] which are dynamic in time, contoured in space as a function of altitude, and whose boundary represents a quantifiable level of safety. In this paper, we present a probabilistic analysis of the disruption to the NAS caused by using traditional hazard areas and compact envelopes during space vehicle operations. Quantities of interest include increased flight time, fuel burn, and distance flown for aircraft that must be rerouted around these hazard areas. We compare the results of using compact envelopes versus traditional methods of implementing no-fly zones to show a near complete elimination of airspace disruption on average and a dramatic reduction in the worst-case disruptions.

I. Introduction

During a rocket launch or space capsule reentry, it is imperative that the safety of aircraft in the National Airspace System (NAS) is assured. The FAA sets the criteria for acceptable levels of risk to non-participating aircraft and hazard areas are established to keep aircraft away from potentially dangerous areas of the airspace.[2,3] Because space vehicles travel vertically at extreme velocities through the NAS and have a non-trivial probability of exploding, the most realistic methods for protecting aircraft in accordance with these acceptable risk thresholds are to establish no-fly zones to keep aircraft away from potential danger. Current methods for constructing and implementing these no-fly zones produce new hazard areas called Temporary Flight Restrictions (TFRs) or activate historically-defined, generic hazard areas called Special Use Airspaces (SUAs).[4] Information about the shapes and timings of TFRs and SUAs are publicly communicated to pilots and air traffic controllers via Notices To Airmen (NOTAMs).

There are three main problems with these traditional methods. First, while the SUAs and TFRs are themselves in the public domain, their creation and analysis requires proprietary software. Second, their exact level of safety is unknown, necessitating a large degree of conservatism when creating the no-fly zones. TFRs are generated using conservative assumptions and safety factors, while SUAs are generic shapes that are not tailored to any specific space mission. It is likely that the resulting no-fly zones are substantially larger than is required to meet the FAA's stated safety requirements. Third, the no-fly zones are static shapes that lack the ability to respond to unfolding events, such as a debris-generating failure or even a successful launch. Thus, the modern methods for aircraft risk management not only restrict a larger volume of airspace than necessary, but their lack of dynamic capabilities make them restrict airspace for longer than necessary as well.

This is a problem; the bigger these no-fly zones are and the longer they are turned on, the more aircraft are disrupted by delays and reroutes to avoid them. These disruptions are costly; extra fuel is burned on the reroute, the crew must be paid for working longer hours, and passengers suffer from longer flights and missed connections. These expenses, in the ballpark of a few thousand dollars per flight affected, are simply unsustainable under a substantial increase in launch traffic. In the past, when launches were infrequent and only conducted by the government, airlines and passengers were willing to tolerate such delays and costs.

---

*PhD Candidate, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA.

†Professor, Department of Aeronautics and Astronautics, AIAA Associate Fellow.

American Institute of Aeronautics and Astronautics

However, the coming decade will bring a multitude of suborbital launch providers into the NAS who want to regularly use the airspace for commercial purposes. Some of these launch providers are targeting multiple launches per day from new spaceports that are close to highly-trafficked airspace. Given this potential spike in launch and reentry traffic as well as new spaceports in already-congested areas, new Air Traffic Management (ATM) procedures are necessary to allow space vehicles and aircraft to equitably and sustainably share the NAS.

We have previously proposed a new class of hazard areas called *compact envelopes*,[1] which are dynamic in time, contoured in space, and whose boundary represents a quantifiable level of safety. The creation of these compact envelopes incorporates projected improvements in air traffic management procedures from NextGen and assumes that agents within the NAS can dynamically react to off-nominal events if given a few minutes of advance warning. This method has the potential to offer hazard areas with a significantly decreased disruption to airline traffic as compared to traditional methods. Compact envelopes are one of the two primary methods currently under consideration by the FAA for their new Concept of Operations for Space Vehicle Operations.

Previous work on the compact envelope concept has demonstrated a dramatic reduction in disturbance to the NAS;[1] however, the examples were merely anecdotal. In this paper, we present a more comprehensive and probabilistic analysis of the airspace disruption caused by compact envelopes versus traditional methods of hazard area creation and implementation.

A. Compact Envelope Concept Of Operations

A compact envelope is a no-fly zone that blocks the smallest 4D volume of airspace as is *safely* possible. This is achieved by drawing the boundaries of the compact envelope to follow the time-evolving, probabilistic contours of acceptable risk tolerances. The shape of the compact envelope evolves with time, such that at every time step, the volume of restricted airspace corresponds only to areas that are hazarded at that moment. The moment that a space vehicle no longer endangers a portion of the airspace, that portion will be released for use by traversing aircraft. In this way, compact envelopes restrict less airspace for less time than traditional methods of separating launch and reentry traffic from airline traffic.

Implementation of the compact envelope concept leverages assumptions about the capabilities of the NAS under NextGen. The following assumptions are used in the remainder of this paper:

- Reaction Time: We make the assumption that pilots and air traffic controllers are able to *safely react* to dangerous events under relatively short notice. The amount of advance warning that they need to safely clear an unexpectedly-imposed hazard area will be called the "NAS reaction time", denoted by t_{react}. Because the NAS can react dynamically if given this amount of time, the only debris that needs to be protected against with hazard areas are those pieces of debris that will reach the aircraft before the pilot can enact a safe response.

- Vehicle Health Monitoring: As a launch or reentry operation progresses, the space vehicle operators have access to state information about the vehicle (positions, velocities, etc.). We assume that this information can be made available to air traffic controllers and pilots in nearly real-time. With this information, hazard areas can be turned on and off depending on the known location and state of the vehicle.

- Data Comm: We assume that NextGen will bring reliably-high-bandwidth data communication to all commercial aircraft in the NAS. Prior to taking off, pilots need not know the exact hazard areas they may encounter, because that information can be provided to them while they are in flight. Further, this allows the exact calculation of the aircraft risks to be delayed until moments before the launch, which can dramatically reduce the size of the hazard area by decreasing uncertainty in the atmospheric properties used to model potential debris dispersions. In the event of a rocket explosion, new hazard areas and rerouted flight plans can be quickly uploaded to the aircraft, thus decreasing the NAS reaction time. The level of data communication assumed in this paper, including accurate GPS position information, allows the pilot to understand and avoid more complex and dynamic shapes than previously possible.

American Institute of Aeronautics and Astronautics

B.   Compact Envlopes For A Commercial Suborbital Space Mission


Figure 1: Snapshots of a compact envelope evolving as a suborbital horizontal-takeoff-horizontal-landing space vehicle nominally carries out its mission. Time progresses from left to right, top to bottom. The vehicle travels to a height of 100km and spends approximately five minutes above the NAS before returning to Earth. The trails of the spacecraft are proportional to its velocity and the leading point of the trail specifies the vehicle's current location.

The compact envelope shown in Figure 1 demonstrates the dynamically evolving hazard areas that would be associated with an XCOR Lynx launching from Spaceport America. The compact envelopes were generated by constraining the cumulative probability of space-vehicle-debris colliding with an aircraft to be below a threshold of 1E-7, using a reaction time of 5 minutes, and assuming zero latency time for the vehicle health monitoring. The airspace has been vertically discretized into four equal sections of 15,000ft each, where the ceiling of the highest envelope corresponds to the upper edge of the NAS at 60,000ft. This example is explained in greater depth in the first paper on compact envelopes.[1]

We have made a standard – though very conservative – assumption in the calculation of these hazard areas regarding the presence of aircraft in the NAS. Specifically, we assume that at every point in space, there is an aircraft present and traveling at its cruising speed. This is obviously unrealistic, but is nevertheless useful for purposes of mission planning because it provides an upper bound for the risk to aircraft that might pass through that point. During the mission-planning phase of a space operation, or even when a NOTAM is generated the night before a launch or reentry, the 4D trajectories of aircraft that will be operating in the NAS

American Institute of Aeronautics and Astronautics

at the time of the space mission cannot be exactly known. This assumption removes the need to know aircraft 4D trajectories in advance and provides a generically safe volume of airspace to restrict from use. This assumption is used for all subsequent examples in this paper.

Figure 2 represents a different view of the same compact envelope from Figure 1. Instead of evolving with time, this figure shows all of the hazard areas activated simultaneously. Notice that the airspace directly above the launch location, from 15,000ft to the NAS ceiling, is *never* restricted from use; this is in stark contrast to traditional methods for mitigating risk to air traffic. Aircraft are free to overfly the launch site during launch operations. Another feature of the compact envelope concept

Figure 2: All hazard areas turned on simultaneously. The launch location is designated by the transparent rocket icon.

is that there exists a corridor between the launch and reentry shape that can be safely available for air traffic to pass through. These are both valuable features for launch operations at spaceports which are located near major airports or congested airways. Compact envelopes offer an air traffic management technique for handling space vehicle operations that can dramatically reduce disruption to the NAS while also meeting the FAA's stated level of acceptable risk.

II. Methodology

A. Future Commercial Space Traffic Estimate

We have collaborated with our colleagues at the FAA Office of Commercial Space Transportation (AST) and the FAA Office of NextGen: Advanced Operational Concepts Division to develop a suite of estimates for future space vehicle traffic.[5] The estimates are broadly grouped into near-term and long-term future traffic. Within each grouping, the estimates are parametrized by the probability of occurrence, with three distinct levels of traffic density: low, medium, and high. The "low" scenarios are thought to be very likely to occur, qualitatively corresponding to a probability of occurrence of 90%. "Medium" scenarios are believed to be reasonable traffic densities in the absence of major catastrophic setbacks – occurring with a probability of about 50%. Finally, "high" traffic scenarios are unlikely but possible, occurring with only a 10% chance. Our determination of these scenarios was guided by studies of the commercial space market published by FUTRON,[6] the Tauri Group,[7] and the FAA,[8] a study on sounding rocket launches,[9] as well as historical data, expert opinions, and other unpublished estimates of future space traffic.

While these estimates provide annual numbers of launch and reentry events, a simulation of the airspace during a space operation must necessarily occur on a single day. Thus, our colleagues at the Office of NextGen created a "representative day" from each possible estimate of future annual traffic. The representative days for the near-term traffic are shown in Figures 3-5. These are the missions, flown by seven distinct space vehicles, that this paper will simulate and analyze.

Near-Term Low-Traffic Scenario

Time	Time	Space Vehicle	Location	Azimuth	SUA Timing
16:00	9:00	SpaceShipTwo	Spaceport America, NM	0	15:50 - 16:20
16:20	11:20	Pegasus	Wallops, VA	150	16:05 - 16:47
20:15	14:15	Lynx	Midland, TX	120	20:05 - 20:35
(UTC)	(Local)			(Deg)	(UTC)

Figure 3: A representative day from the near-term low-traffic-density space operations estimate.

American Institute of Aeronautics and Astronautics

Near-Term Medium-Traffic Scenario

Time	Time	Space Vehicle	Location	Azimuth	SUA Timing
13:30	5:30	Atlas V	Vandenburg AFB, CA	191	5:13-9:28
15:00	10:00	Antares	Wallops, VA	110	9:53-11:23
18:45	11:45	SpaceShipTwo	Spaceport America, NM	0	11:35-12:05
23:23	15:23	Dragon Reentry	Pacific Ocean, CA	135	15:03-15:30
(UTC)	(Local)			(Deg)	(UTC)

Figure 4: A representative day from the near-term medium-traffic-density space operations estimate.

Near-Term High-Traffic Scenario

Time	Time	Space Vehicle	Location	Azimuth	SUA Timing
11:30	6:30	SpaceShipTwo	Titusville, FL	50	6:20-6:50
14:20	8:20	PM2	Van Horn , TX	0	8:10-8:40
16:10	11:10	Lynx	Cecil Field, FL	180	11:00-11:30
18:45	11:45	SpaceShipTwo	Spaceport America, NM	0	11:35-12:05
19:00	12:00	Sounding Rocket	White Sands, NM	0	11:00-14:00
19:50	13:50	Lynx	Midland, TX	0	13:40-14:10
23:45	16:45	Lynx	Front Range, CO	150	16:35-17:05
(UTC)	(Local)			(Deg)	(UTC)

Figure 5: A representative day from the near-term high-traffic-density space operations estimate.

B. Traditional No-Fly Zone Areas

This paper seeks to contrast the improved efficiency of compact envelopes versus traditional methods of implementing no-fly zones during space operations. Thus, we must first have reasonable approximations of a traditional hazard area for each space mission. For launches to orbit and capsule reentries, we use historical data to define the shape and timing of the "traditional" no-fly zones. These hazard areas have been provided to us by the FAA Office of Commercial Space Transportation and the Office of NextGen. For missions that have not yet been flown, we use estimates that were created in collaboration with the FAA offices listed above. Such missions include suborbital launches for horizontal-takeoff-horizontal-landing vehicles, as well as spaceships that are captive-carried to 40,000ft and released from a carrier aircraft before rocket ignition.

The traditional hazard areas associated with the missions described in Figures 3-5 can be seen in Figures A1-A3 of the appendix. For each of the seven vehicles simulated, a single hazard area is created. Space vehicles that launch from multiple locations are usually assumed to use the same traditional hazard area shape and timing at all locations. A brief discussion of the traditional hazard areas for each vehicle follows:

- Lynx, XCOR - This vehicle has yet to complete a successful full-burn mission; therefore, the estimate for traditional operations was derived from a combination of numerical analysis and the opinion of subject matter experts at the FAA offices of AST and NextGen. This hazard area was chosen to be a single rectangular shape – for simplicity – enclosing the areas that would be hazarded by debris from a small number of possible explosions. The hazard area becomes active ten minutes prior to rocket ignition and remains active for a total of thirty minutes.

- SpaceshipTwo, Virgin Galactic - Like the Lynx, this vehicle flies a suborbital mission and has yet to successfully launch. As a first-order approximation, the size and operation of the traditional hazard area is assumed to be identical to the Lynx hazard area described above.

- Pegasus, Orbital Sciences - Hazard area shapes and timing for Pegasus launches are based directly on a NOTAM from June 27, 2013 for a launch out of Vandenberg Air Force Base (VAFB). There were four relatively-small slender rectangular hazard areas that were activated to protect against the launch and the reentry of the first three stages. There were also many SUAs surrounding the launch area

American Institute of Aeronautics and Astronautics

that were activated. For this paper, the Pegasus launches instead from Wallops, so the traditional hazard area that is simulated activates a comparably-sized SUA that encloses the launch area, as well as directly translating / rotating the four rectangular hazard areas from the NOTAM.

- Atlas V, United Launch Alliance - We use the historical NOTAM for a launch from VAFB on February 11, 2013.

- Dragon, SpaceX - We use the historical NOTAM from March 26, 2013 for a Dragon reentry that splashes down in the Pacific Ocean near Los Angeles, CA.

- PM2, Blue Origin - This vehicle has never completed a launch; however, NOTAMs exist for its smaller-scale test flights. For this launch in the present study, the size of the operational hazard area is assumed to be similar to the test-flight hazard area. This is reasonable because the PM2 is expected to behave like a relatively low-altitude single-stage sounding rocket, where the rocket would reliably return to the launch pad. The simulated hazard area is thus circular, with a radius of 31.5km, and is active for thirty minutes, starting ten minutes prior to launch.

- Sounding Rocket - The hazard areas are based on published NOTAMs for this vehicle's current location at White Sands Missile Range; however, the activation times have been curtailed to make them more realistic for the specific mission being simulated. Nevertheless, this hazard area is still active for three hours. The long duration of this hazard area reflects the uncertainty in time of launch for a sounding rocket and also the remoteness of the launch location.

C. Compact Envelope Creation

We use the Stanford University Framework for Aircraft Risk Management tool (SU-FARM) to generate compact envelopes.[1] Since each compact envelope is tailored to a particular vehicle flying a particular mission, SU-FARM takes as input a collection of parameters that define the vehicle's nominal flight trajectory, the trajectory's probabilistic dispersion, the distributions governing the time of potential failure, etc. NASA's Global Reference Atmospheric Model[10] is used to provide probabilistic profiles for wind velocity and atmospheric density at the launch site. Debris catalogs are used to probabilistically model the debris that would be created by an explosion or aerodynamic breakup. A Monte Carlo simulation is run based on the RSAT model,[11] the results of which are used to infer 4D probability density functions for the location of debris in space and time. These density functions are combined with aircraft vulnerability models[12] to determine the risks posed to potential aircraft. Finally, hazard areas are generated to restrict aircraft from any location where the FAA's acceptable risk thresholds have been violated.

D. Historical Aircraft Traffic Data and Filters

This paper uses 90-days-worth of aircraft-track data obtained from the NASA Air Traffic Data Warehouse in the Aircraft Situation Display to Industry (ASDI) format. The data is for January, February, and March of 2013. This volume of data is too large to be directly used for simulating each of the 14 possible missions (Figures 3-5) 90 times apiece. There are simply too many aircraft in the sky for the simulation software to run in a timely manner.

To make the data more manageable for rapid simulation, each individual aircraft from the data is associated with the space missions – if any – that they might be affected by. An "affected area" is identified for each spaceport and mission-type, then the data is filtered to determine which aircraft are *ever* located within the affected area during the day. Those filtered aircraft are saved in separate daily traffic files that are specific to each spaceport-mission combination. When simulating the disruption of a space mission on the NAS, only those aircraft associated with that space mission are considered.


Figure 6: The affected area around Front Range Spaceport (teal circle) easily contains the approximate hazard area for a traditionally-created no fly zone for a suborbital mission (red box). Further, the traditional hazard area completely encloses the compact envelopes (solid red shapes).

For example, suborbital operations will have an "affected area" that is a circle with its center at the spaceport and a radius of 450km. Figure 6 shows the affected area that is specific to the Lynx launching from the spaceport at Front Range, CO. The figure also illustrates the size of the traditional hazard area and the dramatically smaller size of the compact envelope associated with the space mission. This specific compact envelope was seen earlier in Figures 1 and 2. It is obvious that the size of the filter area is sufficient to fully determine all aircraft that might be affected by this space mission.

Having associated each aircraft on each day with the space missions that might affect it, the final step in processing the data is to prepare it for simulation. Each aircraft's state vector is initialized according to its earliest appearance in the data. Its flight plan is set to the one that it was using when it entered the affected area.

E.   Simulating Disruption to the NAS

We use the Future ATM Concepts Evaluation Tool[13] (FACET), developed by NASA Ames, to quantitatively analyze the effect of hazard areas on the NAS. FACET is a physics-based simulation environment that models the take-off, cruise, and descent phases of aircraft flight. If an aircraft trajectory conflicts with an active hazard area, FACET can also simulate the aircraft's rerouted trajectory around the shape. Using FACET, the disruption to an aircraft from a hazard area can be measured in terms of added flight time, added distance flown, added fuel burn, etc. Modifications have been made in FACET to allow compact envelopes to be simulated for comparison against traditional hazard areas.

The specific NAS disruptions estimated in this paper are:

- Rerouted aircraft - An aircraft whose 4D trajectory conflicts with an active hazard area and which is also successfully rerouted around the danger.

- Added fuel burned [lbs] - For a rerouted aircraft, this is the amount of fuel used on the rerouted trajectory subtracted by the amount of fuel that would have been required to fly the nominal trajectory.

- Added flight time [min] - For a rerouted aircraft, this is the amount of time spent flying the rerouted trajectory subtracted by the amount of time that would have been required to fly the nominal trajectory.

- Added flight distance [n.mi] - For a rerouted aircraft, this is the amount extra distance that must be flown to avoid the hazard area.

- Delayed aircraft - An aircraft whose 4D trajectory conflicts with an active hazard area but which was *not able* to be successfully rerouted around the danger.

- Delayed time [min] - For a delayed aircraft, this is the amount of time spent flying inside an active hazard area.

III.   Results

It is assumed that 4D aircraft trajectories, in the first three months of a year, are probabilistically governed by an unknown distribution that accounts for variability in the daily number of flying aircraft, weather patterns, aircraft delays, Air Traffic Control reroutes, and all other possible sources of trajectory variability. Further, it is assumed that this distribution is stationary. Historical aircraft and their trajectories obviously represent samples of this distribution. For simulation, this distribution is *effectively* sampled from by assigning the initial conditions of each aircraft to be equal to those seen in the data and by assigning the flight plan to be an actual flight plan that was flown from those initial conditions.

Each run of a simulation is itself deterministic. For a given mission, the hazard areas will always turn on and off at exactly the same times, with exactly the same locations and orientations. For a given aircraft, its trajectory is deterministically propagated based on its initial state vector and flight plan. For an aircraft that will be rerouted, the rerouted trajectory is also deterministically calculated.

This setup is conceptually identical to a Monte Carlo analysis. A distribution is sampled from, each sample is fed through a simulation that calculates the disruption to the airspace, and the results of each simulation can be used to determine the expected values of disruption. Because the sampled distribution is assumed stationary and the simulation is deterministic, then the results of each simulation are thus governed

## Mean Values Of Airspace Disruptions Using Traditional Hazard Areas

Figure 7: Means and 95% confidence intervals of the disruption due to traditional hazard areas are calculated for each space mission with $N = 90$. Note that delayed aircraft, by definition, will have zero extra fuel burn and zero extra distance flown.

by a stationary distribution. The means and the 95% confidence intervals of the disruption-distributions are calculated from the simulation results with a sample size of $N = 90$ per space mission.

Figure 7 shows the mean and 95% confidence intervals for all quantities of interest from simulations of the traditional hazard area approach. The red bars indicate the aircraft that were rerouted around the hazard area during the operation. The blue bars indicate the aircraft that were unable to avoid the hazard area and entered the danger zone; this is not a failure of the rerouting algorithm. The phenomenon arises because these hazard areas are so large that they often include airports within their boundaries. FACET does not issue ground delays for aircraft about to take-off into an active hazard area, nor does it delay aircraft that are scheduled to land in an active hazard area. FACET simply lets them continue their nominal trajectory. In reality, these aircraft would have been delayed in order to deconflict them from the hazard area.

The amount of time that non-rerouted aircraft spend inside the active hazard areas is counted and provides a lower-bound on the amount of delay time that would have been required to deconflict. The most obvious reason this represents a lower-bound is that multiple aircraft at the same airport will be ground-delayed and when the nearby hazard area disappears, the affected aircraft will not all be able to take off at precisely the time the hazard area deactivates. Indeed, even the first aircraft to take off is unlikely to go at precisely this time. Alternatively, the situation may arise in simulation where an aircraft is descending to land as a large hazard area is activated around it. The simulated aircraft may only appear to spend a single minute inside the hazard area before it lands, but in reality it would have never been allowed to enter the area of the impending hazard and would thus have had to wait until the space operation was complete before landing.

The reduction in NAS disruption from using compact envelopes is so dramatic, that it is not even necessary to plot the mean values for the quantities of interest; the mean number of aircraft rerouted or delayed are nearly zero across the board. The greatest average number of rerouted aircraft occurs for Lynx launching from the Front Range Spaceport with a value of $\mu = 1.633 \pm 0.283$. Instead of mean values, Figure 8 shows the maximum value of each quantity of interest from simulations of the compact envelopes approach. Comparing Figures 7 and 8, it is easily seen that the maximum disruption from a compact envelope is generally one or two orders of magnitude less than the average disruption from a traditional hazard area! This is in addition to the fact that the average values of compact envelope disruption are nearly zero for every mission.

American Institute of Aeronautics and Astronautics

Figure 8: The maximum disturbance observed after 90 simulations using compact envelopes for each space mission is shown. The worst case disturbance for each mission across all 90 simulations is generally far less than the average disturbance of a traditional hazard area.

These results show that the disruption to NAS operations can be nearly eliminated, even in congested airspaces. For instance, The Front Range Spaceport in Colorado is only a few miles from the highly-trafficked Denver International Airport, yet it hopes to host regular launches of the XCOR Lynx vehicle. This paper has demonstrated that if traditional hazard areas are used at this spaceport, every launch would potentially impact a large number of aircraft with delays and reroutes. A launch at 4:45pm would cause an average of 37.53 ± 1.46 aircraft to be rerouted and 54.64 ± 4.97 aircraft to be delayed. However, the space operation becomes far more feasible if a compact envelope is used; the same launch would only cause an average of 1.63 aircraft rerouted and 1.27 aircraft delayed.

The improved efficiency of compact envelopes for the Front Range Spaceport suborbital mission is even more impressive when comparing the extreme values of disruption. The minimum and maximum number of aircraft rerouted using traditional hazard areas is observed to be 21 and 73, respectively. With compact envelopes, that interval shrinks to 0 to 7 aircraft rerouted, with 15 of the 90 simulations showing zero aircraft rerouted and 38 simulations showing only a single aircraft rerouted. Thus, if compact envelopes are used, it can be expected that more than half of the time, no more than one aircraft will need to be rerouted. With NAS disruptions this small, space operations from spaceports like Front Range become highly feasible.

The comparison between the tradition hazard area approach and the compact envelope approach can be seen more clearly by aggregating the space missions together into their "representative days" and then calculating the disruption statistics by day. The result of this aggregation for traditional hazard areas is shown in Figure 9, while Figure 10 shows the compact envelope results. Like before, the mean values for compact envelopes are so small that they are not worth plotting, so instead values of maximum disturbance are shown.

For the low-traffic, medium-traffic, and high-traffic representative days, the average number of aircraft affected (rerouted plus delayed) is 92.07, 98.85, and 381.01 aircraft respectively. If the days simulated are in-any-way representative of daily space traffic, then daily disruptions to the NAS of this magnitude are simply unsustainable. Conversely, the average daily impacts seen using compact envelopes are 0.75, 0.33, and 5.20 respectively. For all three days, the average disruptions due to compact envelopes are two-orders-of-magnitude smaller than than if traditional hazard areas are utilized. At levels of disruption this low, daily space vehicle operations can be efficiently merged into NAS operations.

American Institute of Aeronautics and Astronautics

## Mean Values Of Airspace Disruptions Using Traditional Hazard Areas

Figure 9: Means and 95% confidence intervals of the disruption due to traditional hazard areas are calculated for each representative day of space operations with $N = 90$.

## Maximum Values Of Airspace Disruptions Using Compact Envelopes

Figure 10: The maximum disturbance observed after 90 simulations using compact envelopes for each representative day of space operations is shown. The worst case disturbance for each day across all 90 simulations is generally far less than the average disturbance of a traditional hazard area.

IV.   Conclusions and Future Work

Comparing their efficacy against traditional methods for mitigating risk to air traffic from space vehicle operations, it is apparent that compact envelopes represent a tremendous potential savings in aggregate NAS disruption. The average number of aircraft that must be delayed or rerouted around the larger, static traditional hazard areas climbs into the hundreds for some of the space missions simulated. If compact envelopes are used instead, the number of aircraft that must be delayed or rerouted shrinks to nearly zero on average. In fact, the maximum disruptions from nominally operating compact envelopes tend to be an order-of-magnitude less than the average disruptions caused by traditional hazard areas.

The safety of non-participating aircraft operating the vicinity of a space launch or reentry is paramount and that is why compact envelopes are designed to conform to the FAA's acceptable risk thresholds. The creation of these compact envelopes incorporates projected improvements in air traffic management procedures from NextGen and makes reasonable assumptions about pilots' and air-traffic-controllers' ability to safely react to off-nominal events if given a few minutes advance warning. By making these few realistic assumptions, it is possible to nearly eliminate NAS disruption while maintaining aircraft safety.

The compact envelopes used in this study were generated by SU-FARM,[1] an open-source probabilistic risk analysis environment that we are developing. The tool is not yet ready for release, but it will be made freely available in the coming months at `www.github.com/physicsd00d/SU-FARM`.

We are also investigating the use of Partially Observable Markov Decision Processes (POMDPs) and Predictive State Representations (PSRs) within the compact envelope framework. At the moment, uncertainties due to launch and reentry timing delays and the possibility of suddenly appearing off-nominal hazard areas are handled in a rather simplistic way. A POMDP or PSR could be used to model an aircraft's decision to fly into or near a certain area that it believes may become hazardous at some unknown but near future time.

American Institute of Aeronautics and Astronautics

Acknowledgments

References

[1] Thomas J. Colvin and Juan J. Alonso, "Compact 4D Envelopes For Integrated Air-and-Space Traffic Management", *53rd AIAA Aerospace Sciences Meeting*. January 2015.

[2] Federal Aviation Administration, Department of Transportation, 14 CFR Parts 401, 406, 413, 415, 417 Licensing and Safety Requirements for Launch; Final Rule, Federal Register, Vol. 71, No. 165, August 25, 2006, p. 50508. Federal Aviation Administration, *Docket No. FAA-2000-7953, 71 FR 50537, Aug. 25, 2006,*

[3] Risk Committee, Range Safety Group, Range Commanders Council, Common Risk Criteria for National Test Ranges, RCC 321-07 Supplement, White Sands, NM, 2007.

[4] Federal Aviation Administration. Flight Safety Analysis Handbook Version 1.0. 2011

[5] Young, Jessica, Marie Kee, Christina Young. "Effects of Future Space Vehicle Operations On A Single Day In The National Airspace System: A Fast-Time Computer Simulation". *Technical Note* DOT/FAA/TC-TN15/14. April 2015.

[6] Futron Corporation. *Space Tourism Market Study: orbital space travel and destinations with suborbital space travel.* Futron, 2002.

[7] Tauri Group, The. "A 10-Year Forecast of Market Demand." *The Tauri Group under contract with the Federal Aviation Administration and Space Florida* 2012.

[8] FAA. "2013 Commercial Space Transportation Forecasts". *FAA Commercial Space Transportation (AST) and the Commercial Space Transportation Advisory Committee (COMSTAC)* May 2013

[9] NASA, "NASA Sounding Rockets Annual Report 2012", 2012

[10] Justus, Carl Gerald, and Dale Leroy Johnson. The NASA/MSFC Global Reference Atmospheric Model-1999 Version (GRAM-99). No. NASA-941. National Aeronautics and Space Administration, Huntsville AL, George C Marshall Space Flight Center, 1999

[11] Capristan, Francisco M., and Juan J. Alonso. "Range Safety Assessment Tool (RSAT): An analysis environment for safety assessment of launch and reentry vehicles." (2014).

[12] Wilde, Paul D., and Chris Draper. "Aircraft protection standards and implementation guidelines for range safety." Proc., 48th AIAA Aerosp. Sci. Meeting. 2010.

[13] Bilimoria, Karl D., et al. "FACET: Future ATM concepts evaluation tool." *Air Traffic Control Quarterly* 9.1 (2001).

American Institute of Aeronautics and Astronautics

Appendix
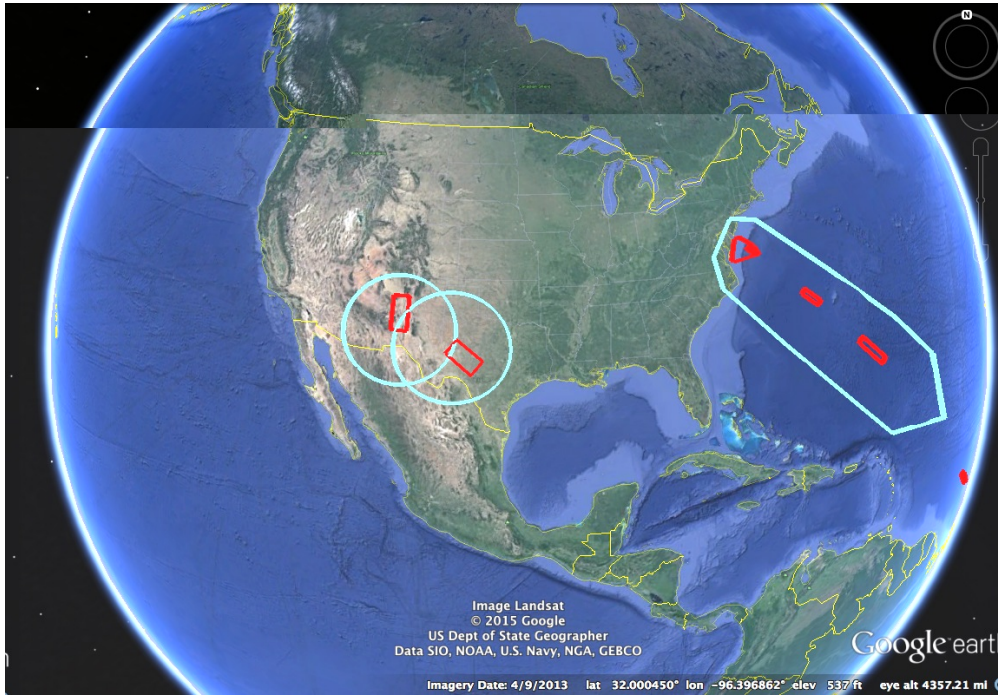
Filters and Traditional SUAs for Near-Term Low

Figure A1: The filters are shown in teal, the traditional hazard areas are in red. It may be noted that the Pegasus launch from Wallops shown here has a single down-range hazard area that is not enclosed in the filter for that spaceport. This is acceptable because we lack a meaningful amount of aircraft traffic data in that region.

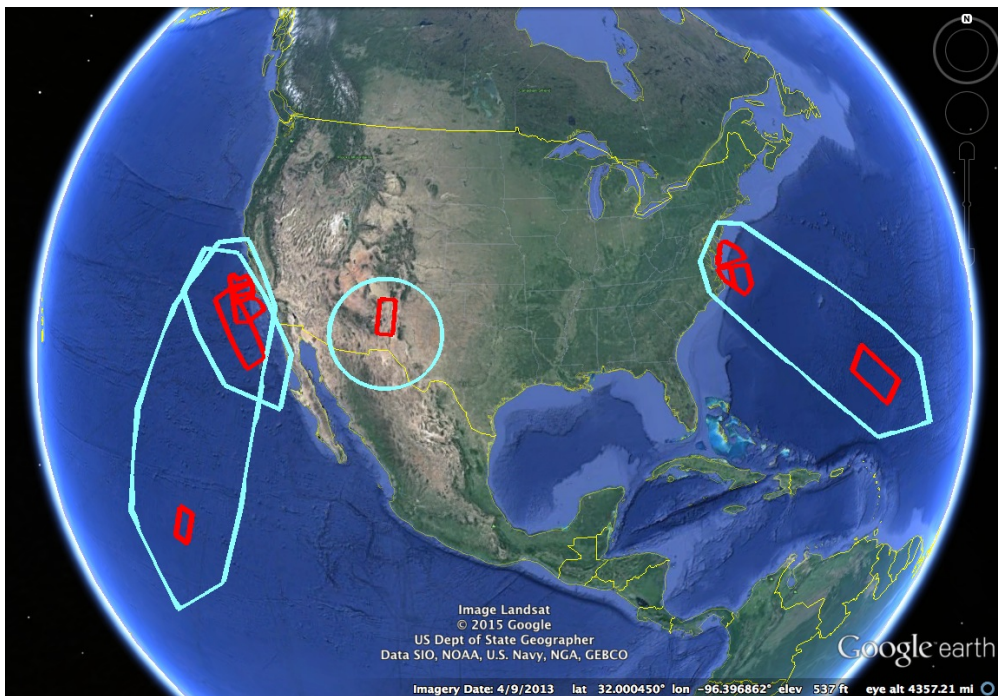## Filters and Traditional SUAs for Near-Term Medium

Figure A2: The filters are shown in teal, the traditional hazard areas are in red.

American Institute of Aeronautics and Astronautics
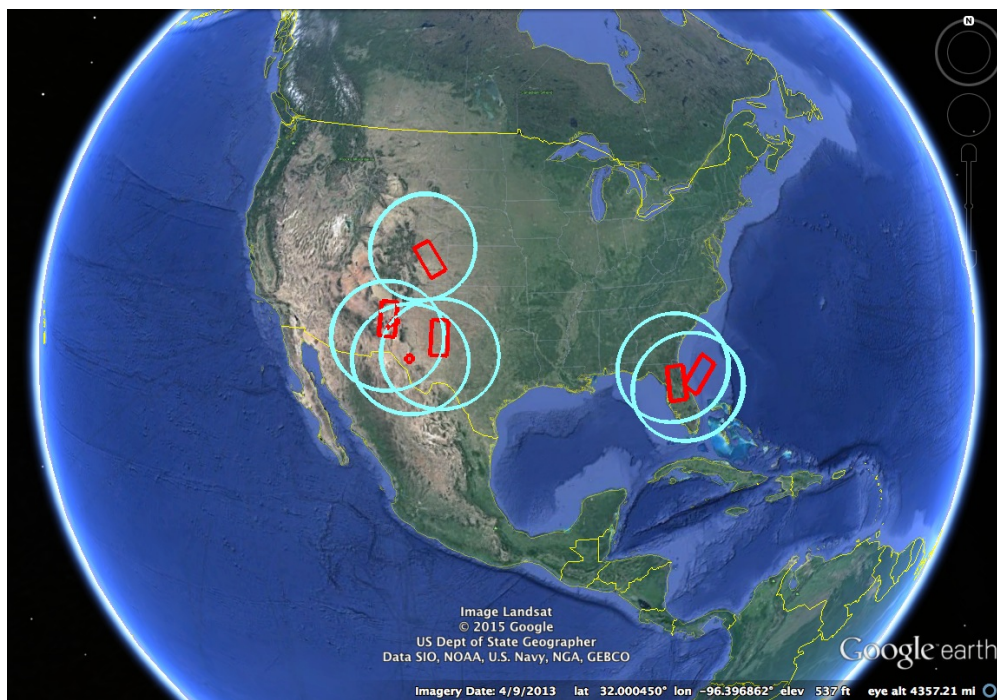
Filters and Traditional SUAs for Near-Term High

Figure A3: The filters are shown in teal, the traditional hazard areas are in red.

American Institute of Aeronautics and Astronautics