AIAA SPACE 2008 Conference & Exposition
9 - 11 September 2008, San Diego, California

AIAA 2008-7874

Partial Rocket Reuse Using Mid-Air Recovery

Mari Gravlee[*], Bernard Kutter[†], Frank Zegler[‡], Brooke Mosley[§]

United Launch Alliance
Denver, CO

Roy A. Haggard[**]

Vertigo
Lake Elsinore, CA

Finding rocket launch cost-reduction opportunities is becoming increasingly necessary in the current market environment due to projected US government budget cuts for space. United Launch Alliance (ULA) is pursuing numerous cost reduction initiatives, one of the more promising concepts being the reuse of the Atlas rocket booster's RD-180 engine. The RD-180 is derived from the reusable RD-170 engine and retains the ability to support multiple missions, offering the opportunity to reduce costs over a series of launches. Reusability of the Russian-built RD-180 also reduces the dependency on foreign hardware deliveries.

Many schemes for engine recovery have been considered in the past. Flyback boosters suffer from huge non-recurring cost and large performance impact. Parachute recovery of an engine module to the ocean suffers from high-impact G loading and exposure to harsh ocean environments which require a complex system to fully seal off the engine. ULA is investigating recovery of EELV engine modules, with an initial focus on the Atlas V's RD-180 engine. Using helicopter mid-air recovery as the engine module descends under a parafoil is a low-development-cost approach which brings back the booster engine with exposure to only benign environments.

Four elements are key in recovering the RD-180 after booster jettison – upgrading the booster to allow separation of the engine module from the tanks, designing both hypersonic and subsonic decelerators, and developing a helicopter mid-air recovery system with increased performance. To date the efforts have been focused on the subsonic parafoil and helicopter mid-air recovery system. A subscale system has been designed and tested to recover 750 lbs, with the intent to be scaled up to support the RD-180 recovery weight of 25,000 lbs. Components integral to the system are a helicopter, a steerable aerodynamic remote-controlled grapple, and a parafoil with stable drogue line attached to the engine module.

This paper outlines a system in development that can realize the cost and technological advantages of reusing high value elements of the launch vehicle, namely the booster aft thrust structure (ATS) containing the engine and other hardware.

[*] Sr. Mechanical Engineer, Advanced Programs, ULA, P. O. Box 277005 Denver Co 80127 MS U9115, AIAA Member
[†] Sr. Staff, Manager Advanced Programs, ULA, P. O. Box 277005 Denver Co 80127 MS U9115, AIAA Senior Member
[‡] Sr. Staff, Advanced Programs, ULA, P. O. Box 277005 Denver Co 80127 MS U9115, AIAA Member
[§] Mechanical Engineer, RD-180 Team, ULA, P. O. Box 277005 Denver Co 80127, AIAA Member
[**] CEO, Vertigo, Inc., 20590 Cereal St., Lake Elsinore, CA 92530, Associate Fellow AIAA

I. History of Launch Vehicle Reuse

Many people assume that fully or partially reusable launch vehicles will result in a significant reduction in the cost of space access. Over the years many methods to reuse rockets have been proposed and pursued. Fly-back boosters were considered during early trades leading to the Space Shuttle, Figure 1. In the 1990's Lockheed Martin and NASA pursued the fully reusable, single stage to orbit (SSTO) X-33, Figure 2. Both Boeing and General Dynamics were studying water recovery of the booster engine module during the Advanced Launch System (ALS) development. Numerous other methods of recovery and reuse have also been considered. Many of these studies showed reusability to only make sense at extremely high launch rates, for example 50 flights per year.

Many early shuttle concepts included fully-reusable boosters and orbiters that carried all of the propellant within the fuselage. Wings were either stowed during ascent and deployed for reentry or fixed outside the fuselage for both ascent and reentry. Wings severely detract from performance due to their weight, and wings fixed outside the fuselage impart huge lateral loads into the fuselage during ascent due to high dynamic pressure. The reusable fly back booster concept was eventually eliminated partially due to development costs, leaving the current Shuttle system with a reusable SRB, expendable propellant tank, and a reusable orbiter. The large refurbishment effort required prior to each shuttle flight has resulted in questionable cost effectiveness of reusing this hardware.

Fully reusable SSTO has been the holy grail for many space enthusiasts. The National Aerospace Plane (NASP) program funded by the US government in the late 80s and early 90s pursued the development of a craft known as the X-30, whose concept utilized bimodal air breathing rocket engines to overcome the harsh physics of reaching space. The VentureStar program funded by Lockheed Martin and NASA in the late 90s, with demonstration prototype dubbed the X-33, was intended to achieve SSTO using more near-term technologies than NASP. VentureStar would take off vertically like a rocket, land horizontally like an airplane, and have airplane-like refurbishment resulting in one-tenth of the recurring cost of the space shuttle. Initial concepts required extremely challenging mass fractions and relied on immature technologies such as lightweight composite hydrogen fuel tanks, which unfortunately could not be developed in the slated budget. The program was cancelled in 2001 after 5 years in development.

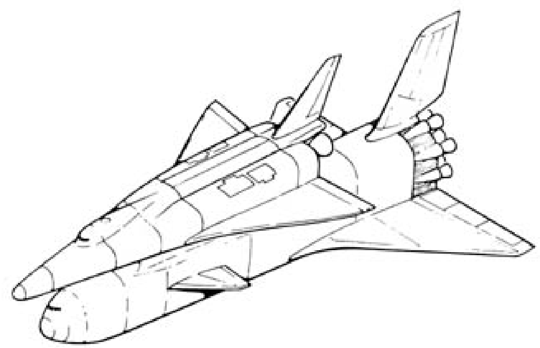

Figure 1 Reusable fly back booster concept considered during shuttle development, Credit NASA[1]

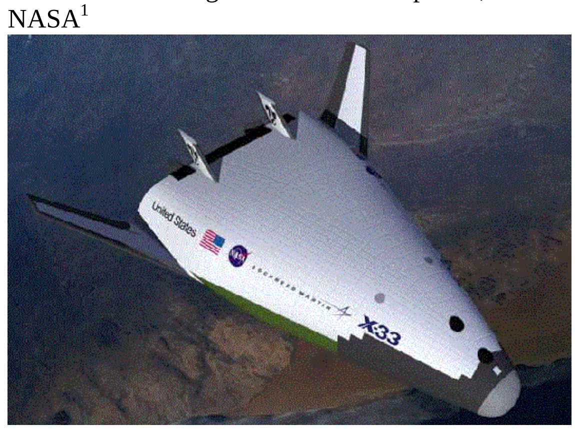

Figure 2 Fully reusable, single stage to orbit X-33 Concept pursued by Lockheed Martin and NASA, Credit NASA[2]

In an effort to realize the benefit of reusability at realistic launch rates, developers have considered partial rocket reuse. These studies primarily have focused on reuse of the booster engine since it typically represents the single most costly element. ALS planned on separating the engine module and using parachutes to "soft" land the engine in the ocean for recovery, using inflatable protective covers to minimize damage of sensitive hardware by the harsh sea salt environment. However, the high impact loads and likely environment-induced degradation made the magnitude of refurbishment extensive, likely more costly than a new engine.

ULA's focus on reusability parallels the ALS development, focusing on limited reuse of the high-value components using realistic flight rates of around 10 missions per year. The key enhancement to the ALS concept is mid-air recovery of the engine module to eliminate impact loads and avoid any chance of sea water contamination. The goal is to ensure that the environment experienced during flight are similar to what the engine already experiences during hot fire, simplifying the required maintenance between flights which is critical to realizing the potential cost savings.

II. Benefits of Engine Reuse

Rocket engines with especially robust designs are made up of components that can handle runtime durations well beyond that required for a single rocket launch. With a potential increase in launch demand for EELV-sized rockets, engine production rate can become a limiting factor. The ability to reuse rocket engines can support double or even triple launch capacity of current engine production capabilities, and the resulting decrease in launch cost could make reliable space flight more affordable.

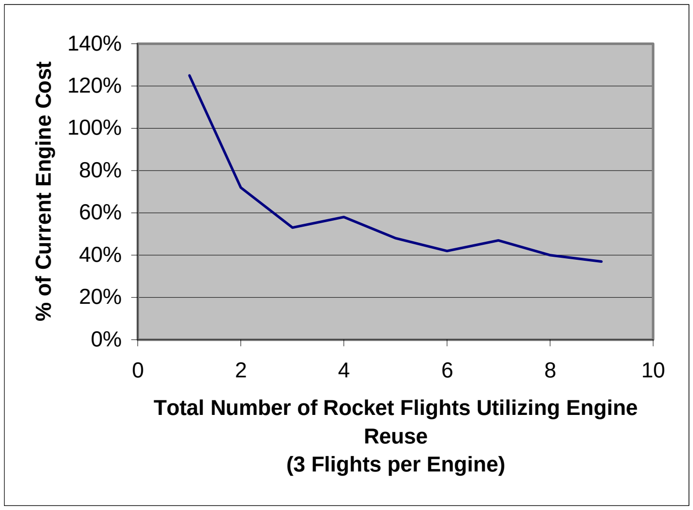

Figure 3 Reuse of the RD-180 can offer savings even if the engine is only used twice. Using the engine 3 times appears to provide maximum savings without pushing the engine into extreme run times.

The benefits of engine reuse have specifically been analyzed with respect to the Atlas V EELV's RD-180 engine. The cost savings for RD-180 engine reuse is projected to be realized in 2 flights, as shown in Figure 3. While previous partial reuse systems were potentially non-cost effective due to g loads where the engine impacted the ground or water, current development of a 3rd generation of mid-air recovery mitigates those g loads. This paper outlines a system in development that can realize the cost and technological advantages of reusing high value elements of the launch vehicle, namely the booster aft thrust section (ATS) containing the engine and other hardware. The minimal refurbishment required to reuse the ATS after being acquired in the air and gently retrieved to the surface allows booster engine reuse to be cost effective. It also facilitates inspection of the rocket engine after flight use, which has not yet been possible with the RD-180.

III. Engine Reuse Overview

ULA's approach to recovery and reuse of the RD-180 is summarized in Figure 4. This approach balances existing technology, realistic flight rates, and operational robustness to enable cost effective recovery and reuse of the RD-180.

The propellant tanks of the booster are large, fragile, and heavy, and only represent about 1/10 the total cost of booster. Therefore, ULA only intends to retrieve the aft end of the booster, Figure 5. The engine module with the adjacent Aft Transition Structure (ATS) weighs approximately 25,000 lb – well within demonstrated parafoil recovery capability. It is compact, strongly built, and can be readily severed from the remainder of the booster. The Aft Thrust Structure will be modified to ease separation and reintegration by adding a flange at the separation plane above the ATS structure to allow the ATS-to-tank interface to be preserved during separation. A flange at the separation plane will prevent the ordnance from cutting the tank which would allow propellant to escape and contaminate the engine.

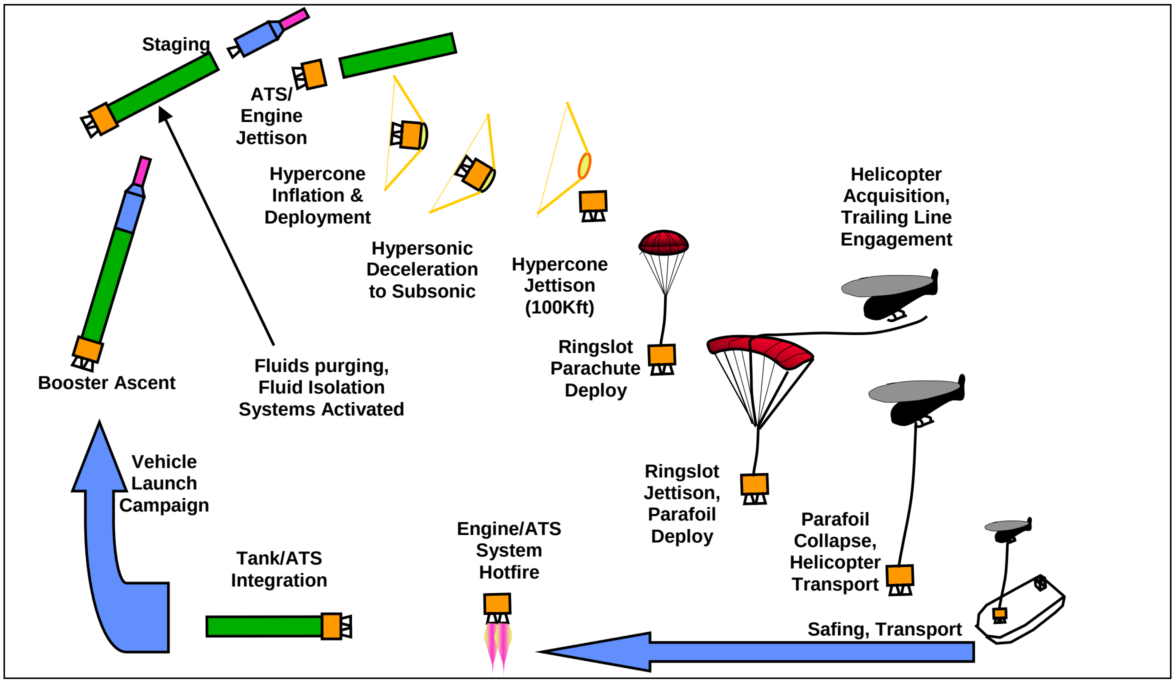

Figure 4 ULA's concept of reuse operations enable benign separation and recovery environments to make engine reuse cost effective to implement

After the ATS is separated from the booster tank structure, it will be decelerated with a hypercone in the upper atmosphere, Figure 6. The hypercone provides a moderate deceleration profile to minimize potential damage to the engine. Figure 7 shows the booster altitude nearing about 800,000 ft before commencing its descent. As the ATS descends through 400,000 ft, drag on the hypercone becomes noticeable.

Hypercones were originally conceptualized for decelerating large payloads at Mars but will function equally well in the Earth's upper atmosphere. The hypercone's shape, high-strength fiber structure, and inflatable torus backbone will permit the gradual deceleration of the ATS and protect it from intense aeroheating. Despite its low weight, the hypercone design creates very high drag under high Mach and micro-pressure conditions where other devices are ineffective. It will be jettisoned once in the lower atmosphere.

To continue deceleration, a ringslot parachute deployment may be utilized until low subsonic velocity is achieved. The chute is released and a parafoil with a trailing drogue line and internal GPS system, Figure 8, inflates above the ATS and steers toward the target area for helicopter

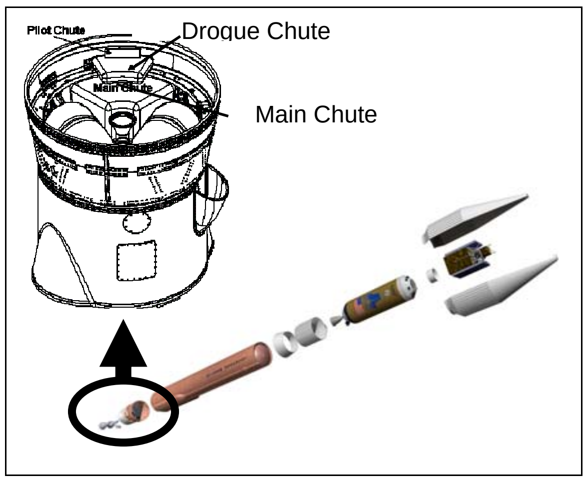

Figure 5 The ATS will be severed from the booster, enabling recovery of the ATS.

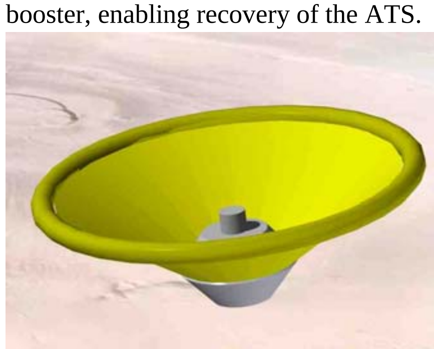

Figure 6 A hypercone is used to decelerate the ATS at hypersonic speeds

intercept. A self-guided parafoil that can currently handle the 25,000 lb load is Airborne Systems' MegaFly parafoil. The MegaFly[3] is designed to carry 30,000 lbs of cargo and can fly autonomously via GPS guidance for distances up to 40 kilometers to a designated point. When the descending parafoil-borne ATS location is determined, the helicopter and the ATS fly a cooperative intercept trajectory. Once the MAR helicopter is within visual range of the ATS, it approaches the parafoil, moving into a trailing formation where the grapple suspended below the helicopter is overlapping the ATS drogue capture line. The grappling hook is steered toward the drogue line for engagement.

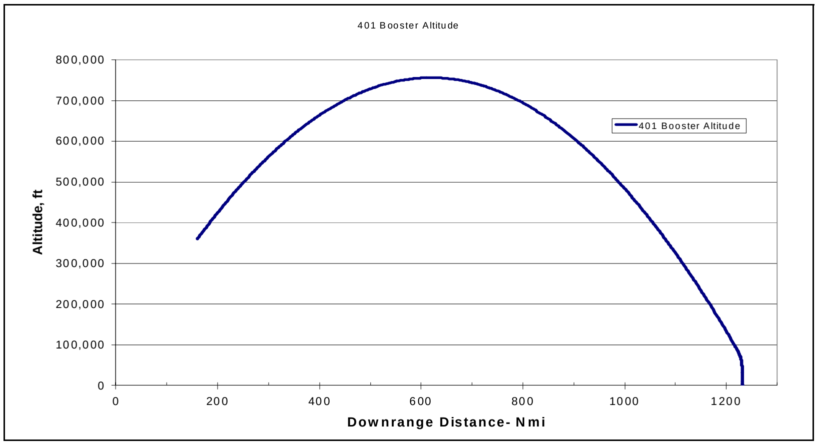

Figure 7 Altitude vs. down range distance, Atlas V 401, 750,000 ft max nominal altitude

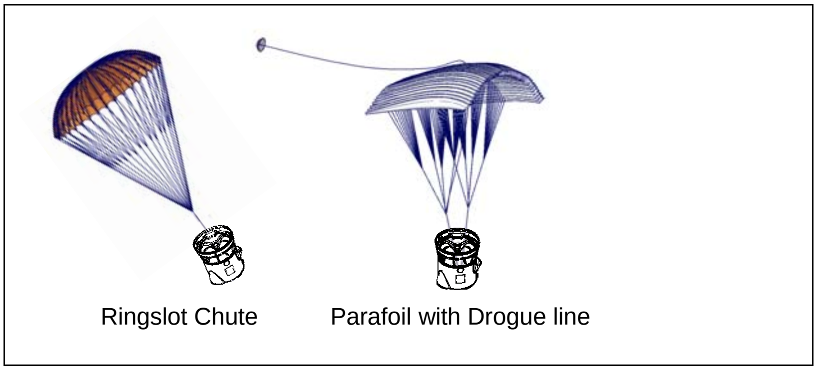

Figure 8 Ringslot chute may be used during transonic flight. Parafoil with drogue is used to fly in formation with helicopter until intercept is complete.

Once the grappling hook has engaged and captured the drogue line, the helicopter ascends and moves forward until the drogue-payload suspension line strips away from the parafoil. When the required tension load is achieved, the parafoil is released through the actuation of pyrotechnic cutters. This method provides smooth transfer of tension loads, minimal suspended payload drag during ferry under the helicopter, and eliminates the difficulty of managing the payload and the very large parafoil in the downwash of the helicopter during payload set down.

When the MAR is completed, the helicopter ferries the ATS as a suspended load to a benign recovery environment, such as to land or a barge. The ATS remains dry, and the setdown method is virtually zero impact, minimizing engine refurbishment scope and cost. 3rd generation MAR allows the helicopter to lift 80% of it's underslung capacity due to near-zero relative velocity between the helicopter and payload and the use of a parafoil instead of a parachute. The helicopter that is best suited to handle the load requirement of 25000 lbs with a transfer load of up to 1.2g's is the Sikorsky Super Stallion CH-53E (36,000-lb underslung load capability).

IV. Typical Trajectory Parameters

There are many factors to be considered in developing booster engine recovery parameters. In a flight using the Atlas V 401 configuration (which is the most common type of Atlas flown), the booster conveniently lands northeast of Puerto Rico, shown in Figure 9, allowing for a timely retrieval. However, more powerful trajectories of other configurations of Atlas may land the booster nearer to Africa than the US, reducing the cost-effectiveness and increasing the logistical complexity of recovery. Timing engine usage such that engines on their 3rd (and last) flight are flown on the more powerful versions of Atlas provides the greatest operational efficiency.

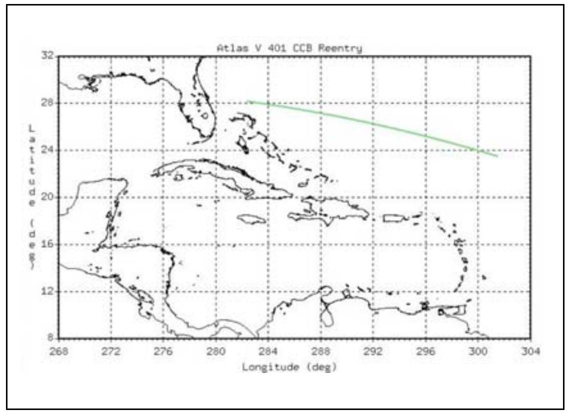

Figure 9 The Atlas V 401 ATS would be recovered northeast of Puerto Rico.

Figure 10 shows the calculated reentry ellipses for an Atlas 401 ATS. The red ellipses provide the 3-sigma impact area helicopters must stay clear of until the actual flight path of the booster is known at booster/ Centaur seperation (~4 min after liftoff). The yellow ellipse provides the ATS 3-sigma impact area. The green ellipse shows where an intact booster is expected to reach sea level 15 minutes after launch.

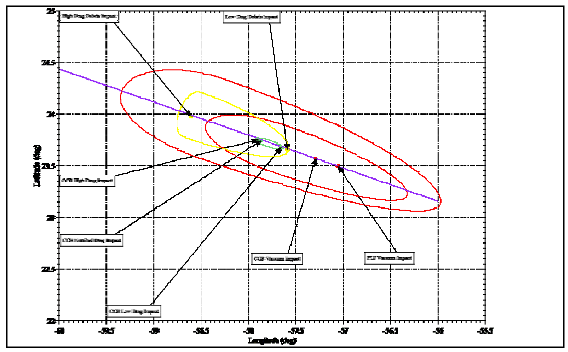
Figure 10 Anticipated impact ellipses for the Atlas 401 ATS

V. History of Mid-Air Recovery.

Mid-Air recovery (MAR) has been developed for many different purposes over the last 50 years, including the retrieval of unmanned air vehicles and air-launched cruise missiles. The first-generation of MAR utilized multiple engagement hooks attached to poles below aircraft. The hooks were connected to a special constant-tension winch that decreased loads and oscillation between the helicopter and payload by paying out line at a pre-set tension, generally about 1.25 times the weight of the payload. The payload floated beneath a round parachute, which, due to the round parachute’s static nature, had to be hit accurately just as the aircraft flew through the wake of the parachute.

Second-generation MAR upgraded to the use of a parafoil system that allowed the intercepting aircraft to fly in formation with the payload. Parafoils are designed to fly along a relatively straight path, facilitating a more controlled and smoother retrieval than round parachutes that hover in an unstable downward float.

3rd generation MAR improves on 2nd generation MAR with the addition of a trailing suspension line connected directly to the supported payload. This allows much heavier loads to be captured, limited by the helicopter performance. In 2005 the Atlas program partnered with Vertigo, Inc to demonstrate 3rd level MAR, capturing a ~200 lb person[5]. ULA is continuing this partnership to pursue component development for 3rd generation MAR, with the goal to eventually enable the capture of the 25,000 lb ATS. In 2007, ULA and Vertigo partnered to conduct two tests – the MAR of a 750-lb pod, demonstrating higher capture mass and transition load, and the MAR of a skydiver with an improved remote-controlled grapple hook.

MAR in the latest tested configuration includes a parafoil with trailing Tri-lobe drogue for maximum stability and minimal drag surfaces. A cable beneath the helicopter suspends the steerable grappling hook that utilizes remote-controlled surfaces to engage and secure the drogue line.

The basic grappling hook, facilitated by a liner actuator, is outlined in the paper, The Past, Present , and Future of Mid-Air Retrieval (AIAA 2005-1673). For most recent testing in November 2007 and February 2008, it was upgraded to handle the 750-lb load, and aerodynamic, steerable surfaces were formed around it to lessen capture time of the drogue line.

In The Past, Present, and Future of Mid-Air Retrieval (AIAA-2005-1673), the conclusion states that "it is conceivable that 3GMAR can be used to safely, reliably, and economically recover high-value payloads of up to 22,000 lb." With today's parafoil technology, recovery of up to 26,000 lbs is technically possible.

VI. Most Recent 3GMAR Testing

The design of the grappling hook has been through much iteration. The linear-actuating hook from early 3rd generation MAR has most recently been mounted to an aerodynamic "fish" design, whose back "fins" can be remote controlled from the helicopter to eliminate the pendulum effect of the hanging weight below the helicopter.

In November 2007, United Launch Alliance contracted with Vertigo to demonstrate the mid-air capture of a 750-lb pod, 550 lbs heavier than the last MAR demonstration load. Tests were performed near California City, California, Figure 14.

The pod with the parafoil stowed inside it was attached to the grapple hook assembly and carried to altitude by an Astar 350B2 helicopter. It was released, and subsequently captured using a grapple hook without aerodynamic surfaces.

Figure 11 C119 MAR near Hawaii: 1st generation MAR

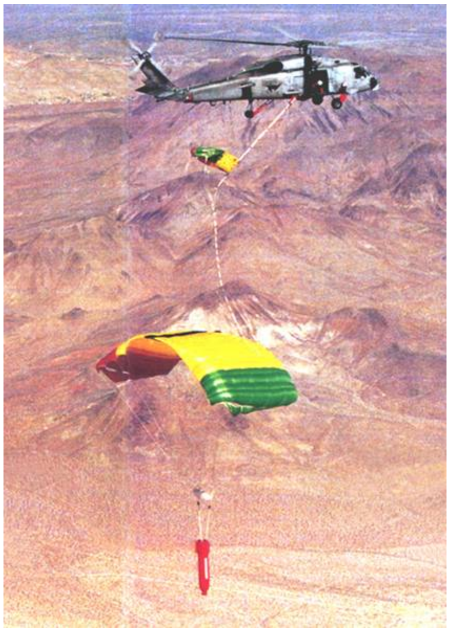

Figure 12 Tandem Parafoil MAR 1990: 2nd generation MAR

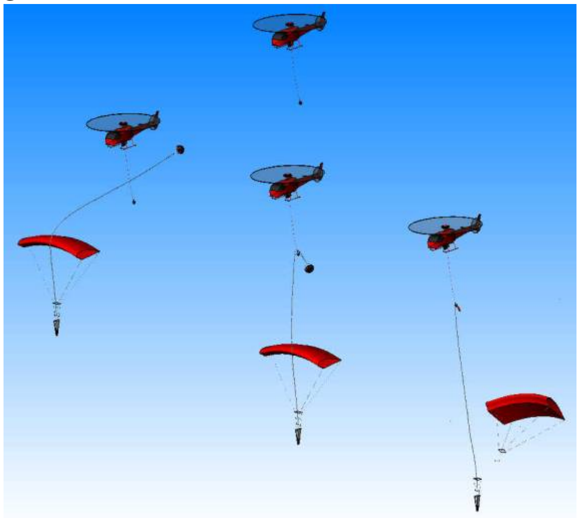

Figure 13 3rd generation MAR

Force data was acquired by examination of the energy attenuation incremental bridle system. The incremental bridle consisted of multiple legs of webbing sewn to each other with carefully sized bar-tack stitches. There were two sets of legs that provided two distinct levels of energy absorption. The first set provided a breaking force of 990 lbf for 10 feet. The second set provided 1,980 lbf for the next 10 feet of tear-out. The primary purpose of the incremental bridle was to act as a force limiter between the suspended load and the helicopter in the event of an overload condition, but because inspection showed that no bar tacks were broken, and none showed signs of significant stress, we can determine that the forces involved in the MAR were less than the measured breaking force of the bar tacks in the first leg of the incremental bridle (this method of data determination was used due to a malfunction of the electronic data acquisition system). The grapple and cargo capsule combined weight was 822 lbs. (70 and 752 lbs. respectively). Therefore, 990/822 equals 1.2 G's acceleration. Even without the force-time history record, the backup data established a very low dynamic load factor during the MAR pickup.

Figure 14 Test – November 2007

Figure 15 Test – February 2008

In February 2008, United Launch Alliance contracted with Vertigo to demonstrate the mid-air capture of a skydiver using a newly developed remote-control grappling hook, Figure 15. Tests were performed over Lake Elsinore in Lake Elsinore, California. The sky-diver jumped from a second helicopter (previous test revealed that skydiver jumping from helicopter that held the grapple was not ideal due to skydiver initial altitude loss). The skydiver's parafoil, a PD Sabre 150, held an 80-foot drogue line. The grapple hook with aerodynamic surfaces and remote-controlled steering, Figure 16, was suspended beneath the helicopter and engaged the drogue line extending from the top of the parafoil. With the remote-controlled grapple, engagement was significantly improved, allowing recovery to be completed on the first pass.

Figure 16 Grappling hook in open and closed positions

VII. Details of Refurbishment.

Figure 17 RD-180 Hot Fire

ULA's initial focus for engine recovery is on the Atlas V's RD-180, whose predecessor the RD-170 was qualified for 10 man-rated flights. If reuse is proven to be cost effective on the Atlas program ULA also plans to pursue recovery and reuse of Delta's RS68. As a 2-chamber derivative of the 4-chamber RD-170, the RD-180 uses 70% of the same flight proven components from the RD-170[5]. The remaining 30% are scaled versions of RD-170 components, and individual parts of the RD-180 have been tested to reusable levels, although full certification of the RD-180 for reuse has never been funded.

Customer requirements dictate that the engine be flown at a 4:1 ratio of certified runtime to flight time. The chamber of the RD-180 went through 21 test firings with a combined time of 4,000 seconds[6], the time equivalent of about 15 flights. The certification engine as a whole was tested 5 times, accumulating 1,024 seconds of runtime[5]. A standard production engine undergoes a 200-second acceptance test and up to 270 seconds of runtime per flight. Therefore, if each engine was certified to be reused twice (for a total of 3 flights), the certified runtime will need to be extended to 4,040 seconds to comply with customer requirements.

To protect the components of the engine during recovery, safeguards can be added to the airborne engine shutdown sequence for contamination control, equivalent to that employed during ground testing. The onboard residual helium supply would be augmented to supply purges to close the fuel inlet valve and the preburner inlet valve prior to separation. Propellant ducts and engine bells would be sealed with airbag-style closures. Many of these techniques were shown effective during prior efforts to prevent seawater entry during a recovery with a water landing. Smaller fluid connections would be sealed with traditional disconnects. Effectively the same processes used to safe the engine following hotfire testing would be imposed prior to atmospheric entry.

Following recovery, the engine would need to be secured. Securing would be focused on preventing possible damage due to changing internal or external engine environments, and preventing external contamination. All engine interfaces would be plugged and internal purges would prevent contamination during transportation.

To verify the engine was suitable for reuse, operations would be similar to the processes currently in place for engine post-abort processing. The engine would undergo a series of hydrocarbon mitigation procedures to rid the engine of hydrocarbons accrued during flight. The current procedure for an engine post-ignition abort involves a series of heated nitrogen purges applied to various ports on the engine via the Hot Gas Ducts followed by vacuuming. A post-flight version of this procedure may need to be more substantial.

The refurbishment would be facilitated by a borescope inspection of the turbine blades, preburner faceplate, and injector. Normal functional testing that follows an engine acceptance test would also be required. Standard testing includes: engine integrated pneumatic testing, engine installed electrical tests, Failure Response System verification, and engine hydraulic operations. These evaluations allow the engine to be qualified for reflight without an engine removal, breakdown, and rebuild. Once flight data review, inspections, and testing determined that the engine to be reused was in healthy working condition, the engine (still mounted in the original thrust structure) would be mated to a new booster tank.

A critical feature of this process is that the condition of the engine after flight can be documented and a database of the condition of critical elements could be built. The exact performance level of any particular engine would be known based on actual flight and not simply acceptance testing. As-installed component wear and performance data, so common to aircraft operations, would begin to be established for a high-thrust booster engine without resorting to costly high-duration hotfire tests. In essence a dynamic process combining flight and ground testing for evolving the engine towards an even higher reliability state would be enabled based on real-world flight environments. With a proven non-detrimental recovery system it may be perceived that flight on the second or third flights of a given engine represent the least risk.

VIII. Launch Architecture Effects

The present strategy for minimizing launch costs is to effectively saturate the launcher performance capability either by adding solid rocket boosters to match the payload weight or by combining payloads to match a fixed launch vehicle capability. The existence of low cost tankage and a cost effective reuse of high value engines and other components suggests an alternative strategy. Depending on the launch rate, manifest mix, payload mass, relative cost, and number of solid rocket motors, it may well be less costly to simply fly heavy payloads on a vehicle similar in configuration to the Delta HLV with large performance excess. This performance excess could enable rideshare opportunities or allow improved orbital insertion such as reduced inclination, that can extend the satellite life. Obviously with two LRB's the recovery process yields two engines per sortie with the energy state of the two liquid rocket booster engines being relatively low- resulting in a shorter downrange distance to the recovery zone easing recovery.

IX. Conclusion.

Extensive research has shown that current technologies and market based launch rates do not support the cost-effectiveness of the reuse of a rocket booster in its entirety. However, reuse of the booster's most costly components appears to be technically viable and cost effective. The booster recovery approach ULA is pursuing achieves the majority of the cost savings of fully reusable flyback booster concepts at a tiny fraction of the non-recurring investment.

ULA is pursuing partial rocket engine reuse to achieve numerous goals, which include: (1) producing cost savings at current launch rates, (2) mitigating dependence on foreign engines, (3) enhancing engine reliability through post-flight inspection, and (4) enabling higher rate launch rates without increased engine production rate and associated capital investment.

Practical rocket booster engine reuse is achievable by maintaining environments that are benign and avoiding contamination. A benign flight environment is enabled through the use of hypercones to decelerate the engine slowly in the upper atmosphere and 3rd generation Mid-Air Recovery. ULA and Vertigo have demonstrated the benign environments and reliable capture of 3rd generation MAR, which incorporates a combination of lessons learned from the extensive history of MAR systems. Current parafoil and helicopter technology already support the recovery of the 25-000 lb load required. Inflatable hypercone decelerators are already being pursued by NASA LaRC and industry. The next major steps to enable actual engine recovery include: (1) refinement of the hypercone to the specific needs of booster recovery, (2) increasing the demonstrated mass capture of 3rd generation MAR, (3) refinement and demonstration of the RD-180 recertification process, and (4) development of the ATS severance modifications.

X. Acronyms

3GMAR	3^{rd} Generation Mid-Air Recovery
ALS	Advanced Launch System
ATS	Aft Transition Structure
EELV	Evolved Expendable Launch Vehicle
LRB	Liquid Rocket Booster
MAR	Mid-Air Recovery
NASP	National Aerospace Plane
SSTO	Single Stage To Orbit
ULA	United Launch Alliance

XI. References:

---

[1] history.nasa.gov/SP-4221/ch8.htm #333
[2] http://apod.nasa.gov/apod/ap960702.html
[3] www.airborne-sys.com
[5] 3^{rd} generation MAR, as of 2005, is explained in “The Past, Present, and Future of Mid-Air Retrieval” AIAA-2005-1673
[5] Katorgin, Boris I., Chvanov, Vladamir K., Chelkis, Felix J. NPO Energomash Khimki, Russia; Popp, Michael, Tanner, Lawrence G., VanGeissen, Robert C., Connally, Scott J., Pratt & Whitney Space Propulsion, West Palm Beach, Florida; “RD-180 Engine Production and Flight Experience” AIAA-2004-3998 40th Joint Propulsion Conference Fort Lauderdale, Florida, July 11-14, 2004
[6] Fedorov V., Chvanov, Vladamir K., Polayansky A. Ivanov N. NPO Energomash Khimki, Russia; Lozino-Lozinskaya I. Keldysh Research Center, Russia; Buryak A. The Institute of Physical Chemistry and Electrochemistry of Russian Academy of Sciences, Russia; “The Chamber Cooling System of RD-170 Engine Family: Design, Parameters, and Hardware Investigation Data” AIAA-2006-4363 42nd Joint Propulsion Conference 9-12 July, 2006, Sacramento, CA