Autonomous Precision Landing of Space Rockets

Lars Blackmore, SpaceX

Opening statement:

Landing an autonomous spacecraft or rocket is very challenging, and landing one with precision close to a prescribed target even more so. Precision landing has the potential to improve exploration of our solar system, and to enable rockets that can be refueled and reused like an airplane. This paper will discuss the challenges of precision landing, recent advances that enabled precision landing on Earth for commercial reusable rockets, and what is required to extend this to landing on planets such as Mars.

A Brief History of Autonomous Space Landings:

Over the past 50 years, autonomous spacecraft have brought humans back from space, landed several rovers on the surface of Mars [1,2,3,4,5], got a probe onto Titan [6], landed on an asteroid [7], and more. Thanks to these missions, we have discovered that Mars was once warm with plenty of water and could likely have supported life. We now know that Titan has lakes of methane, an organic compound. Currently humans fly to and from the International Space Station on a regular basis. Steady progress has enabled heavier payloads to be landed in more exotic locations, and recent advances, such as advanced decelerator technologies [8], will further expand our reach in the Solar System.

Although these missions have aimed for a particular location on the surface of a target planet, the precision has varied. Precision is quantified using the *landing ellipse*, which is the region where it is 99% likely that the vehicle will land. Prior to flight, mission planners must choose a landing site such that everywhere within the landing ellipse is safe for touchdown. Figure 1 shows that the landing ellipse for Mars missions has steadily improved, but is still measured in kilometers, rather than meters.

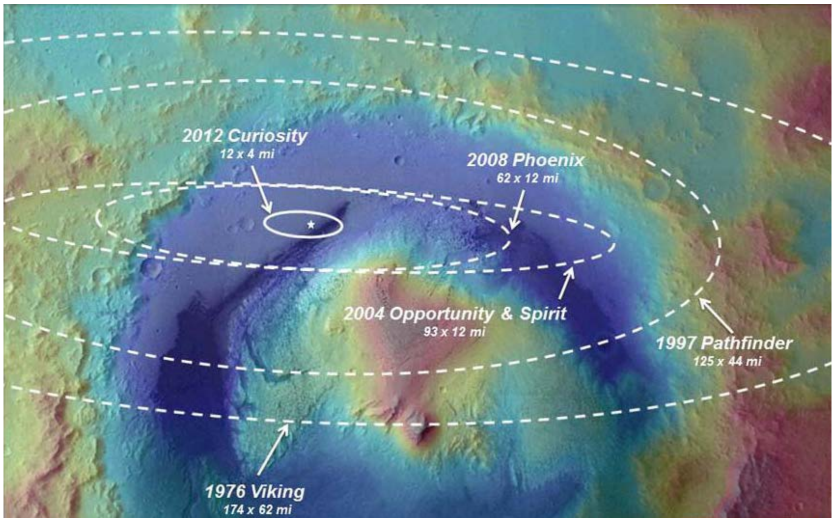

Figure 1: Landing ellipses (in miles) for successful Mars landings to date. Image credit: NASA

The Need for Precision:

When precision is measured in kilometers, missions must land in a desert (in the case of Mars) or in the ocean or on plains (in the case of Earth). If we could measure landing precision in meters instead of kilometers, a world of new opportunities would open up: we could explore Martian caves and valleys, return samples from other planets, and set up permanent outposts throughout the Solar System. We could also make rockets that, after putting a payload into orbit, can be refueled and reused like an airplane, instead of being thrown away after a single flight. This has the potential to decrease the cost of space travel dramatically.

The Challenges:

Precision landing on a planet has several challenges:

- **Extreme environment:** A vehicle entering an atmosphere from space goes through extreme conditions. The majority of the entry energy is dissipated through friction with the atmosphere, resulting in extreme heating that must be dissipated. For example, the leading edge of the Apollo heatshield reached over 2500 degrees Celsuis [9]. Drag causes enormous forces on the reentry vehicle. For example, SpaceX's Falcon 9 Reusable (F9R) weighs about 80 metric tonnes and has a peak deceleration of six times Earth gravity on reentry. Winds push around the reentry vehicle, with high-altitude winds at Earth regularly reaching over 100 miles per hour. Communication may be denied for all or part of reentry due to ionized air around the spacecraft interfering with radio communications. An example of this is the Apollo 13 return capsule, which endured a six minute blackout. And finally, a spacecraft operating outside of Earth orbit is subject to high radiation, which is potentially fatal for electronics. This is especially true of missions operating near Jupiter, where the radiation environment is particularly intense.

- **Small margin for error:** With most landings, the first attempt must be a success, or the vehicle will be destroyed on impact. Rarely is there additional propellant available for a second landing attempt. For large rocket engines, throttling down to a hover is technically challenging and inefficient – every second spent hovering is wasted propellant. For F9R, this means that the rocket has to hit zero velocity at exactly zero altitude. If it reaches zero velocity too low, it will crash; if it reaches zero too high, it will start going back up, at which point cutting the engines and falling is the only option. This requires precise knowledge and control of vertical position and velocity.
- **Touchdown is hard:** A dedicated system, such as landing legs, is usually used to attenuate the loads of landing, keep the rocket safe from rocks, and prevent it from tipping over after landing. Being able to design legs that can do this as mass- and space-efficiently as possible is a challenge, as is delivering the rocket to the upright and stationary position required to avoid overloading the legs' capabilities. For the Curiosity rover, the SkyCrane system enabled the dual use of the rover suspension as the landing attenuation system [10]. In addition, the landing environment may be unprepared and hazardous. For the Mars Exploration Rovers, the combination of rocks and high winds threatened to burst the landing airbags, so an autonomous vision and rocket system was added to detect and reduce lateral velocity [11].
- **Need to hit the target:** Achieving precision landing requires us to hit the target despite being pushed around by disturbances such as winds. For a space reentry vehicle, this is a unique problem, since it is neither a ballistic missile nor an airplane. A ballistic missile tries to hit its target at high speed, so (like a bullet) it uses a high ballistic coefficient and high velocity to avoid being affected by disturbances. An airplane does get pushed around by disturbances, but its wings give it the control authority to correct for those disturbances with ease. A rocket landing vertically has neither of these advantages, making precision landing highly challenging.

Recent Advances:

In the past two years, two commercial companies, SpaceX and Blue Origin, have sent rockets into space and landed them back on Earth within meters of their targets. Blue Origin's 'New Shepard' rocket has landed several times at their West Texas test site. SpaceX's Falcon 9 first stage has landed both on land at Cape Canaveral, and on a floating landing platform known as the Autonomous Spaceport Droneship (ASDS), shown in Figure 2. Some images from recent SpaceX landings are shown in Figure 3.

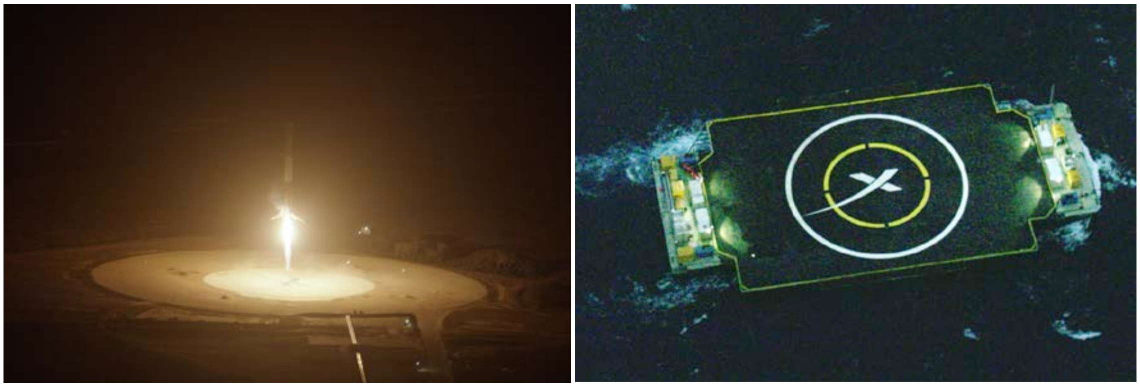

Figure 2: Left: SpaceX's Landing Zone 1 at Cape Canaveral. Right: The SpaceX autonomous spaceport droneship

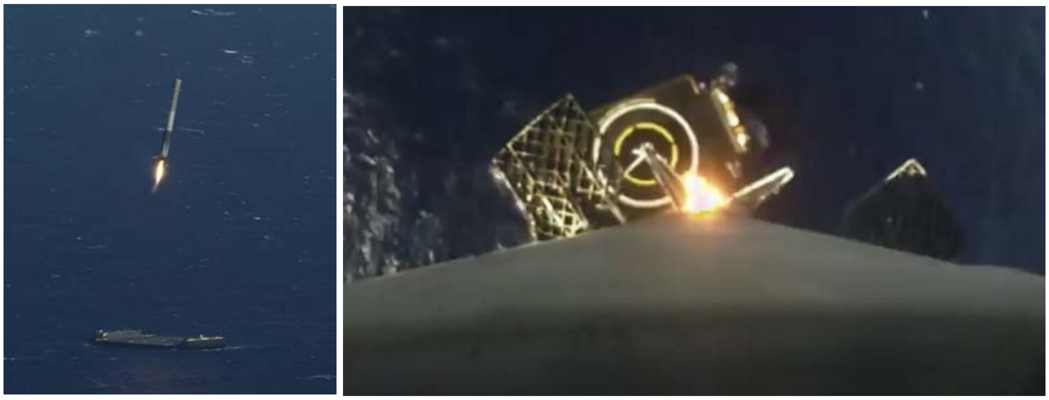

Figure 3: Space F9R approaching the droneship for landing

Central to achieving precision landing is the ability to control variations in the trajectory caused by environmental uncertainty, also known as 'dispersions'. To illustrate this, consider the example of Falcon 9's first stage returning from space. To achieve precision landing, we must control dispersions so that, at touchdown, all dispersions (or at least 99%) fit within the designated landing zone. For F9R, this means achieving dispersions in landing location of 15 meters or better for a droneship landing and 30 meters or better for a land landing at Cape Canaveral. Figure 4 shows the various phases of F9R's mission. On ascent, winds push the rocket around so that dispersions grow. The first opportunity to shrink dispersions is the boostback burn, which sends the rocket shooting back towards the launch pad. During atmospheric entry, winds and atmospheric uncertainties again act to increase dispersions. The landing burn is the last opportunity to reduce the dispersions, and requires the ability to *divert*, or move sideways.

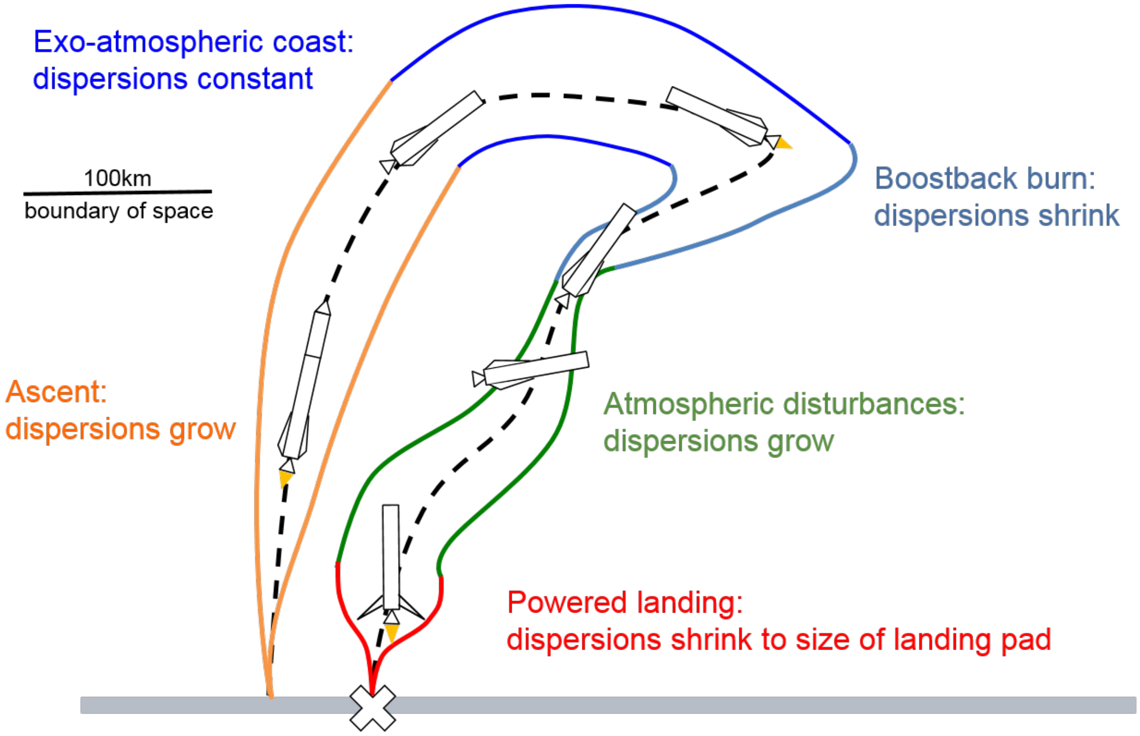

Figure 4: Phases of an F9R return-to-launch-site mission. The colored lines represent the largest possible variations in the trajectory, known as dispersions.

For F9R, controlling dispersions requires precision boostback burn targeting, endo-atmospheric control with fins (shown in Figure 5), and a landing burn with a divert. The latter is one of the most challenging aspects, and is also required for proposed precision landings on Mars [18]. The vehicle must compute a divert trajectory from the vehicle's current location to the target, that ends with the vehicle at rest and in a good orientation for landing, without exceeding the capabilities of the hardware. This computation must be done autonomously, in a fraction of a second. If we fail to find a feasible solution in time, we will crash the spacecraft into the ground. If we fail to find the optimal solution, we may use up our available propellant, with the same result. Finally, in the case of a hardware failure, we may have to replan the trajectory multiple times.

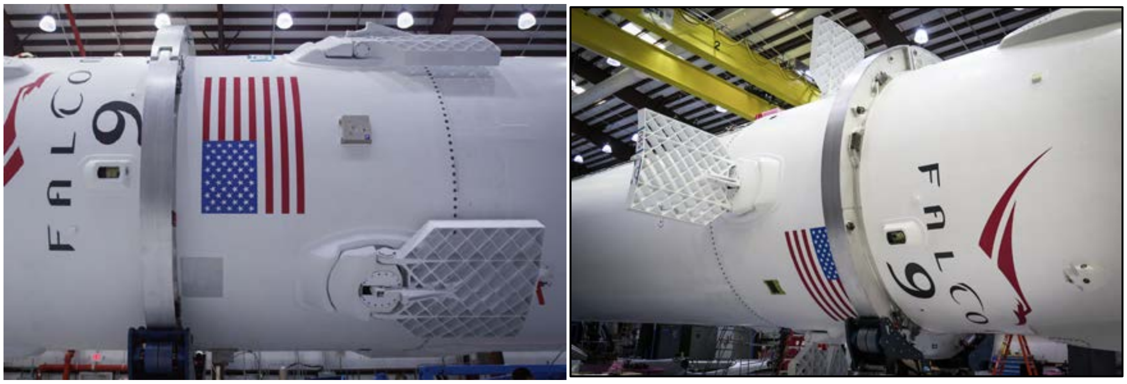

Figure 5: F9R's grid fins. Left: stowed configuration for launch. Right: deployed configuration for entry.

A general solution to such problems has existed in one dimension since the 1960's [12], but not in three dimensions. Over the past decade, research has shown how we can use modern mathematical optimization techniques to solve this problem for Mars landing, with guarantees that we will find the best solution in time [13,14]. Since Earth's atmosphere is one hundred times as dense as that of Mars, aerodynamic forces become the primary concern, rather than a disturbance to be neglected. As a result, Earth landing is a very different problem, but SpaceX and Blue Origin have shown that this too can be solved. SpaceX uses CVXGEN [15] to generate customized flight code, which enables very high-speed onboard convex optimization.

Next Steps:

Although high-precision landings from space have happened on Earth, additional challenges stand in the way of transferring this technology to landing on other bodies in our solar system. One example is navigation – precision landing requires that the rocket knows precisely where it is and how fast it's moving. While GPS is a great asset for Earth landing, everywhere else in the universe is a GPS-denied environment. Almost all planetary missions to date have relied exclusively on Earth-based navigation, where enormous radio antennas track the vehicle, compute its position and velocity, and uplink that information to the vehicle's flight computer. This is sufficient for landings that only need to be precise to many kilometers, but not for landings that need to be precise to many meters. Analogous to driving while looking in the rear view mirror, Earth-based tracking gets less and less accurate the further you get from your starting point. Instead, we need to look towards our destination planet, if we want to be able to land precisely on it.

Missions that used their target to navigate include Deep Impact [16] but this was an impactor mission, not a landing. Recent research to provide navigation accuracy on the order of tens of meters includes [17,18]. The idea behind this is to use *Terrain Relative Navigation,* where the lander images the surface of the planet during landing, and matches features with an onboard map to determine its location. This can be tested on Earth, at least in part, without the need to perform the entire reentry from space. Several companies have used experimental vehicles, some of which are shown in Figure 6, to prove out powered descent technology with low-altitude hops. Using these vehicles, Terrain Relative Navigation has been tested on Earth [17], and a demonstration on Mars is being considered for the Mars 2020 rover mission. If this is successful, combining Terrain Relative Navigation with already-demonstrated precision guidance and control could finally make precision landings on Mars, Europa and other bodies in our solar system a reality.

Figure 6: Various experimental Vertical Takeoff and Landing testbeds. Clockwise from top left: Morpheus, Mighty Eagle, Xombie, Grasshopper, DC-X and Mod

References:

1. Soffen, G. A., and C. W. Snyder, First Viking mission to Mars, *Science*, 193, 759–766, August 1976
2. Golombek, M. et al. 1997. Overview of the Mars Pathfinder Mission and Assessment of Landing Site Predictions. Science. Science: 278. pp. 1743–1748
3. Squyres, Steven W. (2005). Roving Mars: Spirit, Opportunity, and the Exploration of the Red Planet. New York City: Hyperion. pp. 128–138. ISBN 1401301495.
4. Eugene P. Bonfiglio; Douglas Adams; Lynn Craig; David Spencer; Raymond Arvidson; Tabatha Heet. Landing-Site Dispersion Analysis and Statistical Assessment for the Mars Phoenix Lander. Journal of Spacecraft and Rockets, Vol. 48, No. 5. September-October 2011.

5. Steltzner, A.; Kipp, D.; Chen, A.; et al. (2006). Mars Science Laboratory Entry, Descent, and Landing System. 2006 IEEE Aerospace Conference. Big Sky, Montana.
6. M G Tomasko; D Buchhauser; M Bushroe; L E Dafoe; L R Doose; A Eibl; C Fellows; E M Farlane; G M Prout; M J Pringle. The Descent Imager/Spectral Radiometer (DISR) Experiment on the Huygens Entry Probe of Titan. Space Science Reviews, 104, no. 1/2, (2002): 467-549
7. Bibring, J.-P.; Rosenbauer, H.; Boehnhardt, H.; Ulamec, S.; Biele, J.; et al. (2007). "The Rosetta Lander ("Philae") Investigations". Space Science Reviews. 128 (1–4): 205–220.
8. Brian Tibbits and Mark Ivanov. "Low Density Supersonic Decelerator Flight Dynamics Test-1 Flight Design and Targeting", 23rd AIAA Aerodynamic Decelerator Systems Technology Conference, Aerodynamic Decelerator Systems Technology Conferences, (AIAA 2015-2152)
9. Roger D. Launius and Dennis R. Jenkins. Coming Home: Reentry and Recovery from Space. NASA Aeronautics Book Series. (2012)
10. R. Prakash, P. D. Burkhart, A. Chen and D. W. Way. Mars Science Laboratory Entry, Descent, and Landing System Overview. Proceedings of the IEEE Aerospace Conference.
11. Johnson, A., Willson, R., Cheng, Y. et al. Design through Operation of an Image-Based Velocity Estimation System for Mars Landing. Int J Comput Vision (2007) 74: 319
12. Meditch, J. S. On the Problem of Optimal Thrust Programming for a Lunar Soft Landing. IEEE Transactions on Automatic Control, Vol. 9, No. 4, 1964, pp. 477–484.
13. Acikmese, B., and Ploen, S. R.. Convex Programming Approach to Powered Descent Guidance for Mars Landing. Journal of Guidance, Control, and Dynamics, Vol. 30, No. 5, 2007, pp. 1353–1366.
14. L. Blackmore, B. Acikmese and D. P. Scharf. Minimum Landing Error Powered Descent Guidance for Mars Landing using Convex Optimization. AIAA Journal of Guidance, Control and Dynamics, Vol. 33 No. 4, July-August 2010.
15. J. Mattingley and S. Boyd. CVXGEN: A Code Generator for Embedded Convex Optimization. Optimization and Engineering, 13(1):1–27, 2012.
16. Monte Henderson and William Blume. Deep Impact – A Review of the World's Pioneering Hypervelocity Impact Mission. Procedia Engineering 103 (2015) 165 – 172
17. Andrew E. Johnson, Yang Cheng, James F. Montgomery, Nikolas Trawny, Brent Tweddle, and Jason X. Zheng. "Real-Time Terrain Relative Navigation Test Results from a Relevant Environment for Mars Landing", AIAA Guidance, Navigation, and Control Conference, AIAA SciTech, (AIAA 2015-0851)
18. A. A. Wolf, B. Acikmese, Y. Cheng, J. Casoliva, J. M. Carson and M. C. Ivanov, "Toward improved landing precision on Mars," Aerospace Conference, 2011 IEEE, Big Sky, MT, 2011, pp. 1-8.