

Capabilities of Onboard, Convex Powered-Descent Guidance Algorithms for Pinpoint and Precision Landing

John M. Carson III, Behçet Açıkmeşe, Lars Blackmore and Aron A. Wolf
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena, CA 91109
{jmcarson, behcet, lars, awolf}@jpl.nasa.gov

Abstract—The PDG (Powered Descent Guidance) algorithm provides a numerical method for onboard generation of guidance profiles for use during the powered-descent phase of Mars pinpoint or precision landing. The algorithm incorporates both state and control constraints, including minimum and maximum thrust limits, glideslope constraints to avoid impacting the surface, and speed and attitude constraints. These constraints are particularly important for powered-descent scenarios requiring large-divert capabilities to achieve pinpoint or precision landing. Additionally, the constraints ensure that guidance profiles are physically achievable. For instance, the thrust limits are particularly relevant for spacecraft that implement rocket engines that cannot be throttled off after ignition. The formulation of PDG poses the problem as a SoCP (Second-order Cone Program) that can be solved with numerically-efficient interior-point solvers in a finite time to within a prescribed accuracy. This feature is ideal for onboard implementation during powered descent where total flight time is short, thus guidance methods must guarantee convergence to an achievable solution within a short time. If a spacecraft can physically perform maneuvers to achieve pinpoint or precision landing (i.e., the problem is feasible), then the SoCP formulation of PDG will find the solution. Further, this solution will satisfy the prescribed constraints on position, fuel, thrust, speed and attitude.

TABLE OF CONTENTS

1 INTRODUCTION	1
2 OVERVIEW OF ALGORITHM FORMULATION	2
3 SIMULATIONS OF CONVEX GUIDANCE CAPABILITIES	4
4 CONCLUSIONS	7
ACKNOWLEDGEMENTS	7
REFERENCES	7
BIOGRAPHY	7

1. INTRODUCTION

The EDL (Entry, Descent, and Landing) phase of Mars lander missions marks a mission event with significant risk and uncertainty. The development of robust and effective GN&C

(Guidance, Navigation and Control) algorithms, as well as reliable hardware and sensor systems, is essential for the success of EDL. The phases of a typical EDL system include a hypersonic entry phase with a heat shield to slow the entry vehicle, followed by a parachute phase (sub- and/or supersonic) and then a final descent phase either with air bags (as with Mars Exploration Rover and Pathfinder) or powered descent (Viking, Phoenix, and soon Mars Science Laboratory).

During EDL, significant uncertainty affects the descending spacecraft and landing precision. The parachute phase typically dominates the uncertainty, as the dynamics are influenced by the highly-variable atmospheric conditions on a given Martian day (the atmosphere is highly variable in both density and wind profiles). Active control of the parachute is not yet a component in a typical Mars EDL, so current options to null landing error include either long (and likely risky) rover traverses across the Martian surface or improved guidance during the powered descent phase. Powered-descent guidance algorithms reduce the landing error by vectoring the lander thrust to maneuver the lander toward the desired target.

The formulation of powered-descent guidance algorithms must ensure that distance to the desired target is minimized and that physical spacecraft constraints are considered in planning the thrust profile. The guidance problem must incorporate the basic physics of powered descent, along with constraints on the position and velocity, such as avoiding subsurface flight (or descending within a prescribed glide-slope) and limiting speed to avoid excessive drag or supersonic shock waves (if the system design does not handle such velocities). Lander systems typically utilize throttleable liquid thrusters that cannot be shut off (throttled to zero thrust) after ignition, hence, guidance algorithms must allow for both a maximum and minimum thrust magnitude. Additionally, attitude constraints are incorporated in guidance algorithms to allow designers to prevent the guidance solution from commanding undesirable thrust-pointing directions, such as downward thrust toward the surface of Mars or attitudes that violate the field of view of terrain-relative or other sensors necessary during powered descent. These various constraints result in the guidance algorithm trading required fuel and flight time to not only satisfy the constraints but to also achieve the overall objective of either minimum-fuel pinpoint landing [1] or minimum-landing-error precision landing [2].

¹ 978-1-4244-7351-9/11/\$26.00 ©2011 IEEE.

² IEEEAC Paper #1580, Version 1, Updated 12/22/2010.

The EDL phase is a short-duration event that requires guidance algorithms to be accurate and numerically efficient, guaranteeing that a feasible solution will be obtained in a limited amount of time and be accurate. These characteristics were considered in the development of the PDG (Powered Descent Guidance) algorithm, which is a convex formulation of the powered-descent guidance problem for pinpoint and precision landing [1], [2]. The algorithm is implemented as a SoCP (Second order Cone Program), which can be solved with numerically-efficient interior-point solvers that converge to within a desired accuracy in a specified number of steps [3], [4], [5]. The convexity of PDG ensures that the global optimal will be found for a feasible problem and that the solution will be found in finite time. This is in contrast to non-convex techniques that can become stuck in local minima, have no bounds on computation time, and rely upon human operators to provide good initial guesses.

Most of the powered-descent guidance constraints enforced within PDG are convex. However, the non-zero lower bound on thrust is a non-convex constraint. A primary contribution of the work in [1] implemented a relaxation on this constraint that provides a lossless convexification of the thrust bound. Additionally, for attitude limits greater than 90° the constraints are also non-convex, but an additional relaxation provides a further lossless convexification so that the PDG algorithm remains convex. These lossless convexifications ensure that feasible solutions of the relaxed guidance problem are also feasible solutions to the original guidance problem.

This paper summarizes the capabilities of the PDG algorithm and provides representative simulations to demonstrate various landing scenarios. Section 2 provides an overview of the original PDG algorithm from [1] with the relaxation on the thrust constraint, along with a discussion of the additional relaxation to maintain convexity with attitude constraints. Section 3 provides several simulations to demonstrate the effect of the available state and control constraints in the algorithm, and Section 4 provides some concluding remarks.

2. OVERVIEW OF ALGORITHM FORMULATION

This section will provide an overview of the PDG (Powered-Descent Guidance) algorithm for pinpoint landing. For further details of this formulation, including rigorous proofs and a related precision-landing version of the algorithm to minimize landing error, refer to the references [1], [2].

The PDG algorithm provides three-degree-of-freedom translational guidance for a spacecraft descending in a constant gravity field (\mathbf{g}) with negligible aerodynamic forces and a constant planetary rotation ($\boldsymbol{\omega}$). The attitude guidance for the spacecraft is not a part of PDG, and the algorithm assumes that separate attitude guidance and control algorithms are capable of providing the desired orientations for the translational thrust-vector directions provided by PDG. The thruster configuration is assumed to be in a symmetric arrangement

about the spacecraft, as illustrated in Figure 1, with identical canting angles (ϕ) and thrust magnitudes (T) per thruster.




Figure 1. Symmetrically positioned, identical thrusters provide a net thrust T_c along the spacecraft centerline.

The common thrust and cant angle for each thruster provides a net translation-only thrust along the spacecraft centerline, as shown in the side view of the spacecraft. The algorithm assumes that each thruster fires simultaneously to ensure translation-only thrust; a separate attitude guidance and control system could use differential thrust to command attitude (with yaw commanding achieved using separate attitude-control thrusters). The net translational thrust (T_c) is given by

$$T_c = n_{thr} T \cos \phi, \quad (1)$$

where ϕ is the thruster cant angle, T is the thrust magnitude per thruster, and n_{thr} is the number of thrusters. Note, n_{thr} is an integer multiple of 4 for the spacecraft illustrated. Multiple thrusters per corner could be needed to achieve desired thrust, as is the case for the 2011 Mars Science Laboratory.

The coordinate system used to describe the guidance problem herein is illustrated in Figure 2, which also includes descriptions for several variables within the PDG algorithm, including initial position \mathbf{r}_0 , initial velocity $\dot{\mathbf{r}}_0$, landing target surface location \mathbf{q} , and glideslope constraint angle γ_{gs} from the landing target. The unit vectors \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 associated with the coordinate frame are also defined within the figure; the frame is assumed to be a distance R_c from a Mars inertial frame (this will appear within the spacecraft dynamics).

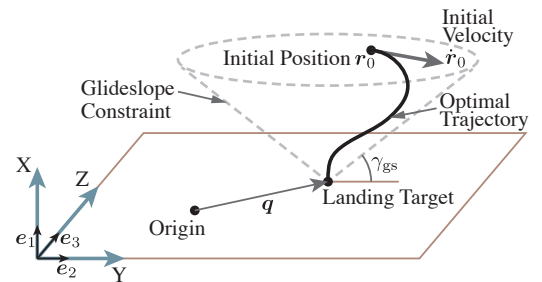


Figure 2. The glideslope constraint requires guidance profiles to remain above the surface and to remain within a cone centered at the landing target and defined by angle γ_{gs} .

The glideslope constraint shown in Figure 2 requires the guidance algorithm to maintain $\mathbf{r}(t) \in \mathbf{P}, \forall t \in [0, t_f]$; i.e., the position trajectory must remain within glideslope set \mathbf{P} during powered descent:

$$\mathbf{P} \triangleq \{ \mathbf{r} : \|S^T(\mathbf{r} - \mathbf{r}(t_f))\| - \mathbf{c}^T(\mathbf{r} - \mathbf{r}(t_f)) \leq 0 \}, \quad (2)$$

where

$$S = [e_2 \ e_3], \quad c = \frac{1}{\tan \gamma_{gs}} e_1. \quad (3)$$

The magnitude of the combined thrust vector \mathbf{T}_c has an upper bound, which is obvious for real systems. In addition, a lower bound is enforced on the magnitude of \mathbf{T}_c , which models thrusters that cannot be throttled to zero thrust after ignition; this is a common feature of liquid thrusters. This minimum and maximum thrust level is modeled as a constraint on the norm of the thrust vector:

$$0 < \rho_1 \leq \|\mathbf{T}_c(t)\| \leq \rho_2, \forall t \in [0, t_f]. \quad (4)$$

The upper bound on the norm of the thrust is a convex constraint, however, the convexity breaks down with the lower bound. As will be discussed, a relaxation of this constraint provides a lossless convexification for the PDG algorithm.

Thrust pointing constraints are incorporated in the PDG algorithm to ensure that the translational guidance solution does not generate thrust-vector directions that are outside of a desired attitude cone. This attitude cone allows the guidance to consider sensors and other systems (e.g., terrain-relative navigation systems) that require a specific field of view to obtain useful information. Further, mission requirements might also preclude certain thrust orientations, for instance thrusting toward the surface. These thrust pointing constraints are expressed as

$$\hat{\mathbf{n}}^T \mathbf{T}_c(t) \geq \|\mathbf{T}_c(t)\| \cos \theta, \quad (5)$$

where $\theta \in [0^\circ, 180^\circ]$ is the maximum allowable attitude. Constraint (5) is convex for $\theta \in [0^\circ, 90^\circ]$ but non-convex for $\theta \in (90^\circ, 180^\circ]$. As with the lower bound in (4), a relaxation will be discussed that provides a lossless convexification for the pointing constraint when $\theta \in (90^\circ, 180^\circ]$. Note, the pointing constraints are defined with respect to a reference unit vector $\hat{\mathbf{n}}$, which would typically be in the direction of the X axis (i.e., $\hat{\mathbf{n}} = e_1$).

Figure 3 illustrates the thrust constraints. The left image depicts the thrust magnitude bounds from (4), the center image depicts the pointing constraints from (5) for $\theta > 90^\circ$, and the right image depicts the allowable thrust region at the intersection of the two constraint sets.




Figure 3. Planar Representation of Thrust Constraints. The thrust magnitude constraint set is non-convex. The thrust pointing constraint is also non-convex for $\theta > 90^\circ$.

The pinpoint landing version of the PDG algorithm minimizes the fuel required to achieve pinpoint landing while simultaneously satisfying the powered-descent dynamics and the aforementioned state and control constraints (including maximum speed, available fuel, and initial and final conditions). Note, an initial and final thrust direction can be enforced as well [1]; the final pointing constraint is useful to ensure, for example, that the lander completes the PDG profile in an upright orientation. Mathematically, the pinpoint landing PDG algorithm is stated as the following optimization problem:

Problem 1: (Non-convex Minimum-Fuel Guidance)

$$\max_{t_f, \mathbf{T}_c(\cdot)} m(t_f) - m(0) = \min_{t_f, \mathbf{T}_c(\cdot)} \int_0^{t_f} \alpha \|\mathbf{T}_c(t)\| dt$$

subject to the following dynamics and trajectory constraints, $\forall t \in [0, t_f]$,

$$\begin{aligned} \ddot{\mathbf{r}}(t) &= -\boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{r}) - 2\boldsymbol{\omega} \times \dot{\mathbf{r}} + \mathbf{T}_c(t)/m(t) + \bar{\mathbf{g}}, \\ \dot{m}(t) &= -\alpha \|\mathbf{T}_c(t)\|, \\ 0 < \rho_1 &\leq \|\mathbf{T}_c(t)\| \leq \rho_2, \quad \hat{\mathbf{n}}^T \mathbf{T}_c(t) \geq \|\mathbf{T}_c(t)\| \cos \theta, \\ m(t_f) &\geq m_{\text{dry}}, \quad \mathbf{r}(t) \in \mathbf{P}, \quad \|\dot{\mathbf{r}}(t)\| \leq V_{\text{max}}, \end{aligned}$$

and the following boundary conditions

$$\begin{aligned} m(0) &= m_{\text{wet}}, \quad \mathbf{r}(0) = \mathbf{r}_0, \quad \dot{\mathbf{r}}(0) = \dot{\mathbf{r}}_0, \\ \mathbf{e}_1^T \mathbf{r}(t_f) &= 0, \quad S^T \mathbf{r}(t_f) = \mathbf{q}, \quad \dot{\mathbf{r}}(t_f) = 0, \\ \mathbf{T}_c(t_f) &= \|\mathbf{T}_c(t_f)\| \hat{\mathbf{n}}_f, \end{aligned}$$

where set \mathbf{P} and matrix S are defined in (2) and (3), respectively, $\bar{\mathbf{g}} = \mathbf{g} - \boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{R}_c)$, and parameters m_{wet} , m_{dry} , α , ρ_1 , ρ_2 , $\hat{\mathbf{n}}$, V_{max} , \mathbf{r}_0 , $\dot{\mathbf{r}}_0$, \mathbf{q} , $\hat{\mathbf{n}}_f$ are constant.

If pinpoint landing cannot be achieved with the available fuel, then a precision landing version of the above Problem 1 instead minimizes the landing error from the desired target \mathbf{q} . The precision landing, or fuel-limited targeting, version of this problem is formulated in [2].

The general guidance algorithm in Problem 1 is non-convex due to the constraints on the thrust vector. A significant contribution of the work in [1] and [2] implemented a relaxed, convex version of the thrust magnitude constraint (4) and rigorously proved that an optimal solution of the relaxed version of Problem 1 (without thrust pointing constraints) was also optimal for the unrelaxed problem. The relaxation provided a lossless convexification for the thrust magnitude constraint,

$$\|\mathbf{T}_c(t)\| \leq \Gamma(t), \quad (6)$$

$$\rho_1 \leq \Gamma(t) \leq \rho_2, \forall t \in [0, t_f]. \quad (7)$$

The slack variable Γ used in the thrust magnitude relaxation can also be used to relax the non-convex thrust pointing constraint. This new relaxation can also be rigorously proved to provide a lossless convexification such that an optimal solution of a relaxed version of Problem 1 with pointing con-

straints is also an optimal solution for the unrelaxed problem. The relaxed pointing constraints from (5) can be expressed as

$$\hat{\mathbf{n}}^T \mathbf{T}_c(t) \geq \Gamma(t) \cos \theta, \forall t \in [0, t_f]. \quad (8)$$

Making use of these relaxed constraints, the relaxed version of Problem 1 is expressed below. Note, the discrete version of this relaxed problem is an SoCP, which as discussed in the introduction can be solved with numerically-efficient interior-point methods that converge to within a desired accuracy in a specified number of steps [3], [4], [5]. The same holds true for a precision landing version of the relaxed problem.

Problem 2: (Relaxed Minimum-Fuel Guidance)

$$\min_{t_f, \mathbf{T}_c(\cdot), \Gamma(\cdot)} \int_0^{t_f} \Gamma(t) dt$$

subject to the following dynamics and trajectory constraints, $\forall t \in [0, t_f]$,

$$\begin{aligned} \dot{\mathbf{r}}(t) &= -\boldsymbol{\omega} \times \boldsymbol{\omega} \times \mathbf{r} - 2\boldsymbol{\omega} \times \dot{\mathbf{r}} + \bar{\mathbf{g}} + \mathbf{T}_c(t)/m(t), \\ \dot{m}(t) &= -\alpha \Gamma(t), \\ \|\mathbf{T}_c(t)\| &\leq \Gamma(t), \quad \hat{\mathbf{n}}^T \mathbf{T}_c(t) \geq \Gamma(t) \cos \theta, \\ 0 &< \rho_1 \leq \Gamma(t) \leq \rho_2, \\ m(t_f) &\geq m_{\text{dry}}, \quad \mathbf{r}(t) \in \mathbf{P}, \quad \|\dot{\mathbf{r}}(t)\| \leq V_{\text{max}}, \end{aligned}$$

and the following boundary conditions

$$\begin{aligned} m(0) &= m_{\text{wet}}, \quad \mathbf{r}(0) = \mathbf{r}_0, \quad \dot{\mathbf{r}}(0) = \dot{\mathbf{r}}_0, \\ \mathbf{e}_1^T \mathbf{r}(t_f) &= 0, \quad S^T \mathbf{r}(t_f) = \mathbf{q}, \quad \dot{\mathbf{r}}(t_f) = \mathbf{0}, \\ \mathbf{T}_c(t_f) &= \Gamma(t_f) \hat{\mathbf{n}}_f, \end{aligned}$$

where set \mathbf{P} and matrix S are defined in (2) and (3), respectively, $\bar{\mathbf{g}} = \mathbf{g} - \boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{R}_c)$, and parameters m_{wet} , m_{dry} , α , ρ_1 , ρ_2 , $\hat{\mathbf{n}}$, V_{max} , \mathbf{r}_0 , $\dot{\mathbf{r}}_0$, \mathbf{q} , $\hat{\mathbf{n}}_f$ are constant.

The solution to the relaxed guidance algorithm in Problem 2 comes from the Pontryagin Maximum Principle, which maximizes the Hamiltonian of Problem 2. The Hamiltonian for the relaxed problem is given by

$$H = R_1(t)\Gamma(t) + \mathbf{R}_2(t)^T \mathbf{T}_c(t) + R_0(t), \quad (9)$$

the details of which are discussed in [1], [2]. To remain valid for the original guidance problem (Problem 1) the solution of the relaxed problem (Problem 2) must not have feasible solutions that violate the original problem. Of concern is the relaxed thrust magnitude constraint in (6), which could allow for a thrust magnitude of zero, which clearly violates the original thrust magnitude constraint in (4) since $\rho_1 > 0$.

For the relaxed problem without the pointing constraint (8), the theory developed in [1], [2] rigorously proved that solutions to the relaxed problem were also valid solutions to the original, non-convex algorithm. For optimal solutions to Problem 2 with the relaxed pointing constraint, $\Gamma^*(t)$ will be strictly non-zero, as seen from constraint (7). Thus, the

relaxed attitude constraint (8) will ensure a non-zero $\mathbf{T}_c(t)$ for $\theta \in [0^\circ, 90^\circ) \cup (90^\circ, 180^\circ]$; for $\theta = 90^\circ$, $\mathbf{T}_c(t)$ will also be non-zero, but further theoretical justification, outside the scope of this paper, is required. Since \mathbf{T}_c will be non-zero, the maximization of H for relaxed Problem 2 will give $\|\mathbf{T}_c^*\| = \Gamma^*$ from constraint (6), and the bounds on Γ in (7) will ensure satisfaction of the original thrust magnitude bounds in (4). Thus, an optimal solution of relaxed guidance Problem 2 will also be an optimal solution of the original guidance Problem 1.

3. SIMULATIONS OF CONVEX GUIDANCE CAPABILITIES

Several example simulations for spacecraft performing pinpoint and precision landing were developed to highlight some of the capabilities of the PDG algorithm. This section will discuss these simulations.

The simulations used the following parameters:

$$\mathbf{g} = [-3.71 \ 0 \ 0]^T \text{ m/s}^2, \quad (10)$$

$$\boldsymbol{\omega} = [0 \ 0 \ 7.09 \times 10^{-5}]^T \text{ rad/s}, \quad (11)$$

$$\mathbf{R}_c = [3396.2 \times 10^3 \ 0 \ 0]^T \text{ m}, \quad (12)$$

$$\gamma_{\text{gs}} = 5^\circ, \quad (13)$$

$$m_{\text{wet}} = 1900 \text{ kg}, \quad m_{\text{dry}} = 1600 \text{ kg}, \quad (14)$$

$$\rho_1 = 0.2 T_{\text{max}}, \quad \rho_2 = 0.8 T_{\text{max}}, \quad (15)$$

$$T_{\text{max}} = 21500 \text{ N}, \quad \alpha = 5 \times 10^{-4} \text{ s/m}, \quad (16)$$

$$\mathbf{T}_c(t_f) = \|\mathbf{T}_c(t_f)\| \hat{\mathbf{n}}_f, \quad \hat{\mathbf{n}}_f = [1 \ 0 \ 0]^T, \quad (17)$$

where T_{max} is the maximum net-thrust magnitude of \mathbf{T}_c , and the thrust limits coincide with a throttle minimum and maximum of 20% and 80%, respectively. The values for \mathbf{r}_0 , $\dot{\mathbf{r}}_0$, and \mathbf{q} vary between the different simulations and can be inferred from the plots. Additionally, the speed constraint V_{max} and the attitude constraint θ vary between the different simulations. Vector $\hat{\mathbf{n}}_f$ is specified to ensure the final thrust vector points in the upright direction; the lander is upright at the end of the guidance profile. The thrust pointing constraint during the PDG profile is also constrained around this pointing vector (i.e., $\hat{\mathbf{n}} = \hat{\mathbf{n}}_f$).

The PDG algorithm design, as described in Section 2, explicitly enforces constraints on the thrust (minimum and maximum), available fuel, speed, attitude and glide slope (to avoid surface impact). An example pinpoint landing simulation that enforces all of these constraints is shown in Figure 4. The time of flight from powered-descent ignition to pinpoint landing is 35.4 seconds with a total fuel usage of 199 kg. As seen in the figure, the constraint limits are drawn alongside the PDG profiles.

The upper three plots in Figure 4 include mass, speed and thrust-pointing profiles versus the flight time. For each of these three plots, the PDG trajectory profile obeys the enforced constraint, including the new constraint on thrust pointing angle. Since the pinpoint landing algorithm mini-

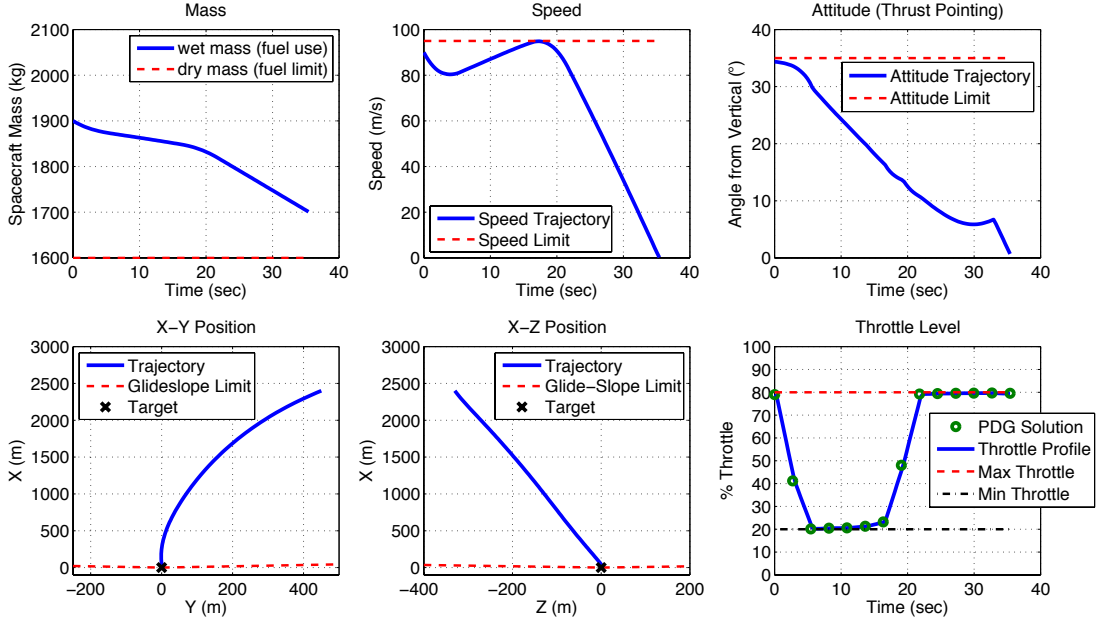


Figure 4. Typical constraints on a powered-descent guidance profile: fuel, speed, attitude, glide-slope and throttle.

mizes fuel usage, the spacecraft mass at the end of the flight time is well above the dry mass limit, indicating that excess fuel was available for powered descent. This excess fuel could be budgeted for final landing maneuvers required for certain types of landers or to account for sensor uncertainty that would require the final PDG target to be above the surface, such that a terminal vertical-descent phase would conduct landing.

The bottom three plots in Figure 4 include the X-Y and X-Z position trajectories and the thruster throttle level. The position trajectories include the glideslope constraint ($\gamma_{gs} = 5^\circ$), and the profiles are in the XYZ coordinates from Figure 2. The throttle plot includes both the minimum and maximum throttle limits alongside the throttle profile that comes from the PDG guidance solution from the discretized version of Problem 2. The throttle values are always within the bounds, as expected.

As constraints are tightened, there is generally a change in performance (Refer to Figure 5 and Table 1). For instance, in a shallow pinpoint-landing scenario with an initial downward velocity of 10 m/s and a lateral velocity of 40 m/s, the PDG algorithm optimization without any attitude constraints would choose a profile that initially thrusts towards the surface (at 125° to the local vertical); this profile minimizes flight time and fuel. However, thrusting toward the surface can be considered risky or can exceed the field of view of terrain-relative sensors. Since PDG is capable of applying general attitude constraints, the downward thrust can be avoided but at a trade off in guidance performance. As attitude constraints are tightened to 90° and to 45° , there is an increase in required fuel and time of flight. This is illustrated by the data in Table 1 which shows a small fuel and flight-time impact from the 90°

constraint but a much larger effect from the 45° constraint.

Table 1. Example case of varying attitude constraints and the effect on performance

Attitude	Required fuel (kg)	Flight time (sec)
Unconstrained	190	45.5
90° Constraint	192	47.7
45° Constraint	200	53.1

In addition to the required fuel and time of flight, the approach trajectory and thrust profile must change so that the attitude constraints are feasible. Figure 5 provides plots of the attitude profile, throttle profile and surface trajectory (Y-Z) for the three cases in Table 1. The increase in flight time seen in the table is also visible in the throttle profile plot. Notice that all throttle constraints are satisfied for the three different attitude constraint profiles. The surface trajectory is also significantly altered for the 45° attitude constraint and requires the lander to overshoot the target and then turn back toward it in order to satisfy the attitude constraint as well as the dynamics and other constraints during powered descent.

As discussed in the introduction and with the PDG algorithm overview, PDG can provide both pinpoint and precision landing guidance profiles. For precision landing, PDG minimizes the distance to the desired landing target given the fuel available. This is also known as fuel-limited targeting and was developed in [2]. Figure 6 provides a surface-trajectory plot for a precision landing case where initial ground velocity is in a direction 60° from the target direction. The PDG guidance profile minimizes the landing error with the available fuel and simultaneously satisfies the state and control constraints.




Figure 5. Constraints affect the performance. For instance, tighter attitude constraints can increase flight time, significantly alter the approach trajectory and increase required fuel.

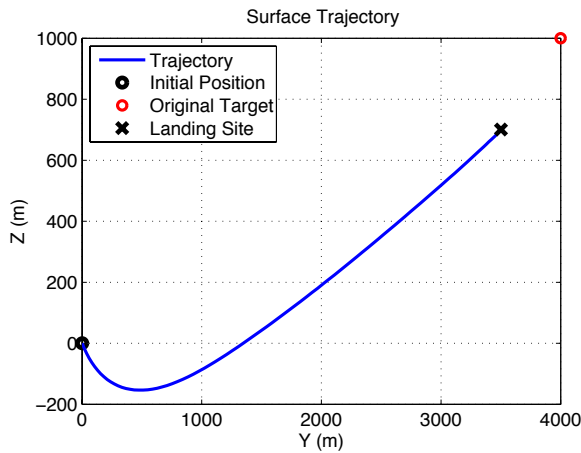


Figure 6. Precision landing with PDG minimizes the landing error from the original target.

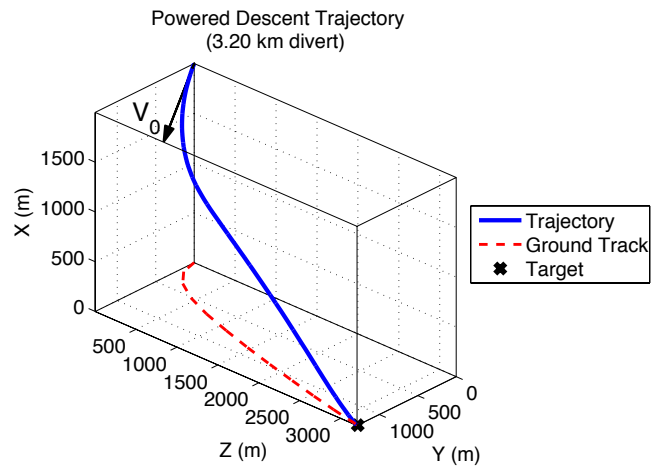


Figure 7. The PDG algorithm enables large divers in highly-constrained Mars pinpoint landing problems.

The divert required to reach the target in Figure 6 would be in excess of 4 km. Although this target is infeasible for a pinpoint landing with the particular example spacecraft, this figure illustrates that the PDG precision landing capability is able to minimize the landing error to the target. Additionally, this figure illustrates the capability of PDG to provide large-divert guidance profiles that avoid surface impact and satisfy the thrust and other state constraints. This capability is further highlighted by the three-dimensional landing trajectory and ground track in Figure 7 that illustrates a pinpoint landing case requiring a 3.2-km divert to reach the target, which is within the capabilities of the example spacecraft.

The large-divert pinpoint landing in Figure 7 illustrates a further capability of onboard PDG compared to heritage algorithms. Heritage guidance methods generally make use of analytic solutions to restricted powered-descent problems with a low-order polynomial-based guidance trajectory [6], [7]. These methods work suitably around well-known operating conditions and in environments, such as lunar, where dynamics and descent conditions are much slower than those associ-

ated with Mars landing. The polynomial guidance profiles are parameterized such that they can be tailored to obey state and control constraints around the operating region. This works well for small divers but can cause problems when larger divers are required, the reason being that the polynomial guidance profiles cannot guarantee state and control constraint satisfaction far from the design point.

Figure 8 provides one example of a comparison between PDG and a polynomial guidance scheme for a descent scenario with thrust constraints and a position constraint to avoid surface impact. In the figure, which is based on the work in [8], the downrange divert capability from different powered-descent ignition altitudes is compared between PDG and a second-order Apollo guidance algorithm. The simulation is for initial velocities of -50 m/s vertical and 20 m/s lateral, with the example spacecraft having the following parameters: $m_{wet} = 1308$ kg, $m_{dry} = 1100$ kg, $T_{max} = 16000$ N, and a minimum and maximum throttle of 15% and 95%, respectively. The comparison clearly indicates that the PDG algorithm has a significantly larger performance envelop for

pinpoint-landing diverts compared to the Apollo algorithm. The polynomial guidance, despite having adequate fuel to achieve the landing maneuvers, has no means of ensuring that guidance profiles satisfy the thrust constraints and the position constraints to avoid impacting the surface.

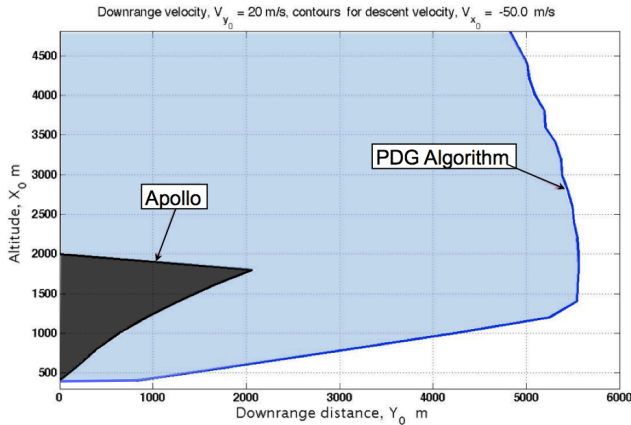


Figure 8. Envelope of target-relative initial positions from which Apollo and PDG can reach the desired landing target while satisfying all state and control constraints. PDG has a much wider envelope of divert capability.

4. CONCLUSIONS

The capabilities and strength of the PDG algorithm for onboard powered-descent guidance lies in the convexification of the optimization problem that explicitly enforces state and control constraints, including minimum and maximum thrust and general pointing constraints. The formulation of PDG ensures not only satisfaction of all the constraints but also that a solution will be found if the guidance problem and desired maneuvers are feasible for the spacecraft to perform. The numerical efficiency of convex algorithms makes PDG ideal for onboard implementation during Mars powered descent where the system is significantly constrained and time criticality and reliability dictate the need for rapid solutions and convergence guarantees. Further, the ability for PDG to provide large-divert guidance profiles while still ensuring state and control constraint satisfaction greatly enhances the overall landing capabilities for future Mars lander missions that will require pinpoint and precision landing.

ACKNOWLEDGMENTS

We gratefully acknowledge Mark Ivanov and Jordi Casoliva of the Jet Propulsion Laboratory for their valuable comments and suggestions. This research was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.


REFERENCES

- [1] B. Açıkmeşe and S. R. Ploen, “Convex programming approach to powered descent guidance for mars landing,” *AIAA Journal of Guidance, Control and Dynamics*, vol. 30, no. 5, pp. 1353–1366, 2007.
- [2] L. Blackmore, B. Açıkmeşe, 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, 2010.
- [3] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [4] Y. Nesterov and A. Nemirovsky, *Interior-point Polynomial Methods in Convex Programming*. SIAM, 1994.
- [5] K. Toh, M. Todd, and R. Tutuncu, “SDPT3 — a Matlab software package for semidefinite programming,” *Optimization Methods and Software*, vol. 11, no. 1, pp. 545–581, 1999.
- [6] A. R. Klumpp, “Apollo lunar descent guidance,” *Automatica*, vol. 10, pp. 133–146, 1974.
- [7] C. D’Souza, “An optimal guidance law for planetary landing,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference 1997*, 1997.
- [8] S. Ploen, B. Acikmese, and A. Wolf, “A comparison of powered descent guidance laws for mars pinpoint landing,” *AIAA Guidance, Navigation, and Control Conference, Keystone, CO*, 2006.

BIOGRAPHY



John M. Carson III received a B.S. in Aerospace Engineering (1997) and M.S. in Engineering Mechanics (1999) from the University of Texas at Austin, followed by a Ph.D. in Mechanical Engineering (2008) from the California Institute of Technology. From 1999 to 2000 he performed flight and structural dynamics testing and analysis on the Joint Strike Fighter prototype at the Lockheed Martin Skunk Works. From 2001-present he has worked in guidance, navigation, and control design and analysis at the NASA Jet Propulsion Laboratory for several missions, including Cassini/Huygens and Mars Exploration Rover, and in technology development programs for small-body and planetary landing systems.



Behçet Açıkmeşe received his Ph.D. in Aerospace Engineering in 2002 from Purdue University. He was a Visiting Assistant Professor of Aerospace Engineering at Purdue University before joining JPL in 2003. He is currently a technologist and a senior member of the Guidance and Control Analysis Group at Jet Propulsion Laboratory (JPL), where he is developing

guidance, control, and estimation algorithms for planetary landing, formation flying spacecraft, and asteroid and comet sample return. His research interests include robust and non-linear control, optimal control, model predictive control, convex optimization and linear matrix inequalities (LMIs) in guidance, control and estimation, and real-time optimization.



Lars Blackmore is currently a Technologist in the Guidance and Control Analysis group at JPL. His work focuses on guidance and control for planetary pinpoint and precision landing, guidance and control under stochastic uncertainty, and estimation for formation-flying spacecraft. Lars received his Ph.D. in Control and Estimation in 2007 from the Massachusetts Institute of Technology, and has an MEng in Electrical Engineering from the University of Cambridge (UK).



Aron A. Wolf is the Supervisor of the EDL / Aero Applications Group in the Guidance, Navigation, and Control Section at the Jet Propulsion Laboratory, where he has worked since 1982. In this capacity, he supervises mission and navigation design for atmospheric flight segments of JPL's missions including EDL, aerobraking, aerocapture, and aerogravity assist. He leads Mars Pinpoint Landing study efforts at JPL, and previously served as MSL Mission Design and Navigation Lead and managed the Advanced EDL work area in the Mars Technology Program. His previous work in the area of trajectory design has included the design of orbital tour trajectories for the Galileo mission to Jupiter and the Cassini mission to Saturn. He received his undergraduate degree from Brown University in 1977 and his Masters in Aerospace Engineering from the University of Michigan in 1980.