WALKING THE FILAMENT OF FEASIBILITY

GLOBAL OPTIMIZATION OF HIGHLY-CONSTRAINED, MULTI-MODAL INTERPLANETARY TRAJECTORIES USING A NOVEL STOCHASTIC SEARCH TECHNIQUE

Arnold C. Englander[1] and Jacob A. Englander[2]

Interplanetary trajectory optimization problems are highly complex and are characterized by a large number of decision variables and equality and inequality constraints as well as many locally optimal solutions. Stochastic global search techniques, coupled with a large-scale nonlinear programming (NLP) solver, have been shown to solve such problems but are inadequately robust when the problem constraints become very complex. In this work, we present a novel search algorithm that takes advantage of the fact that equality constraints effectively collapse the solution space to lower dimensionality. This new approach "walks the filament" of feasibility to efficiently find the global optimal solution.

INTRODUCTION

Interplanetary trajectory optimization problems are highly complex and often include a large number of decision variables and constraints. The more realistic the physics and spacecraft hardware model, the more constraints need to be accounted for. Typically, these constraints come in two forms: equality constraints that enforce trajectory continuity and inequality constraints that enforce feasibility of the mission from the perspective of spacecraft hardware, *e.g.* do not exceed the maximum available thrust or consume more propellant than is available in the tank, or from the perspective of operations, e.g. the mission flight time may not exceed some stated maximum value and one is not allowed to crash into any planets. Furthermore, these problems are often multi-modal, meaning that there are many locally optimal solutions. In fact, the number of locally optimal solutions is coupled to the number of constraint equations – the more constraints the problem has, the more local optima it has.

There are two general classes of mathematical models, or transcriptions, by which an interplanetary trajectory optimization problem may be described. The classical approach is the "indirect" method, in which one formulates the problem using the calculus of variations and then solves a multi-point boundary value problem [1]. A more recent approach is the "direct" family of transcriptions, in which the trajectory optimization problem is converted into a large-scale NLP characterized by a large number of variables that represent a time history of states and controls. This work focuses on a direct method, specifically the two-point direct shooting technique developed by Sims and Flanagan [2]. However, the optimization algorithm described in this work could be used with any transcription, or indeed any large, highly constrained optimization problem.

Classically, the direct optimization problem is solved using a large-scale gradient-based NLP solver such as Sparse Nonlinear Optimizer (SNOPT) [3]. This is an effective technique in finding a locally optimal solution if an initial guess is provided but suffers from significant limitations. First, the NLP solver can only converge to a solution in the vicinity of the initial guess and this may be only a local optimum. Second, the globally optimal solution is often non-intuitive and therefore selecting an initial guess near it is very challenging. Third, the NLP solver can fail if the constraint topology is "rippled," or non-monotonic, as is often the case in interplanetary design.

Recent research has explored the problem of finding the globally optimal solution to a complex trajectory optimization problem without an initial guess. One of the most promising techniques is

[1]Principal Consultant, Englander and Associates, LLC, Member AAS
[2]Aerospace Engineer, Navigation and Mission Design Branch, NASA Goddard Space Flight Center, Member AAS

monotonic basin hopping (MBH), which consists of a stochastic search step hybridized with a NLP solver [4]. This approach has been applied successfully to interplanetary trajectory optimization problems [5], and has recently been improved by modifying the stochastic search step to draw from a Pareto distribution [6]. The most recent version of MBH has been demonstrated to solve interplanetary problems subject to complex operational constraints [7] but is not adequately robust as the size of the problem, i.e. the number of variables and constraints, and therefore local optima, increases.

In this work, we develop the notion of a random walker that, as part of an MBH process, walks along one or more 1-dimensional filaments, representing the points that are feasible at least in the sense that they satisfy all of the equality constraints, in an n-dimensional solution space where MBH searches for global min (f(**x**)) subject to one or more equality constraints and one or more inequality constraints.

The benefits of the filament waking are:

1. It allows the gradient search sub-process to largely be unreliable; yet still provide robust global MBH performance.
2. It provides robust global MBH performance even for problems that involve a great many local minima that each satisfy the constraints.
3. It allows that, prior to its random walk along the filament(s), the MBH only needs to find a few points on the equality constraints. The filament walker will then find many more.
4. And perhaps most valuable, filament walking, if it is sufficiently "thorough," reduces over time, an N-dimensional search to a 1-dimsional search. This is important because:
   a) A stochastic search in 1 dimension is tremendously faster than stochastic searches in spaces of greater than 1 dimension. This can be shown analytically using First Passage Time theory, and by numerical experiments, and will be explored in a later paper.
   b) The recording of points **x** visited by random walks on intervals in 1 dimension, and the sampling of f(**x**) at those visited points, comprise lists. Such lists can be processed easily and quickly, *e.g.* their lengths can be evaluated, their minima and maxima can be evaluated, they can be "filtered," meaning that entries that violate inequality constraints can be deleted, *etc*.
   c) The "thoroughness" of stochastic search in 1 dimension is much easier to estimate and achieve than it is in spaces of greater than 1 dimension.

Properties 4a, 4b, and 4c, are important because they enable estimates to be made regarding how "thoroughly" filaments were sampled – meaning, how well the MBH explored the feasible solutions for global min (f(**x**)) after any period of time. In turn, that enables quantitative measures of confidence in the solutions found thus far.

Our ambition is to eventually develop quantitative confidence measures of whether all filaments were found, and whether all the filaments that were found were walked (sampled) with sufficient thoroughness, "almost surely", within some specifiable number of search steps. This is needed to address the question of when the optimization process can be stopped. We believe that, because filament walking has an explicit notion of filaments, filaments are 1-D, and the walking process tracks the positions along filament(s) where the **x⁺** are visited and the f(**x⁺**) are sampled, "thoroughness" can be stochastically controlled – and evaluated – in ways it cannot be (or much more easily than it can be) with prior versions of MBH.

Measurement and control of thoroughness is especially important in the context of "hybrid optimal control" (HOC), in which complex problems composed of both discrete and continuous variables may be solved [8, 9]. Such problems may be composed into an "outer-loop" that chooses the discrete variables and an "inner-loop" that solves a continuously valued sub-problem for each candidate choice of discrete variables. The "outer-loop" then chooses which sub-problem to create next based on the results

of the previous "inner-loop" run. For example, in interplanetary trajectory optimization, a trajectory is defined both by discrete variables such as the number of gravity assist maneuvers and the identities of the bodies at which those maneuvers are performed and also by continuous variables such as launch epoch, flight time, and the epochs, magnitudes, and directions of propulsive maneuvers. However, HOC only works if the "inner-loop" provides reliable information to the "outer-loop." For example, if the "inner-loop" is unable to find the globally optimal trajectory that travels to Jupiter via a flyby of the Earth, then the "outer-loop" will incorrectly discard that flyby sequence and this will affect the remainder of the "outer-loop" solution process. The notion of "thoroughness" is therefore very important to the "inner-loop" solution process because if "thoroughness" can be measured and controlled then the integrity of the HOC process can be guaranteed.

The algorithm described here is implemented in the Evolutionary Mission Trajectory Generator (EMTG) [6–10], which includes several suitable problem transcriptions, highly realistic spacecraft hardware modeling, and a great deal of flexibility that allows us to construct suitable benchmark problems for this work.

TRAJECTORY OPTIMIZATION BASED ON FILAMENT WALKING

We describe an approach to trajectory optimization that we call "filament walking." It is an extension of applications of MBH methods to trajectory optimization. We explain how it combines stochastic optimization with local gradient searches in a manner that addresses the challenges of finding the constrained global minimum among many local minima, complicated constraints, and problematic gradient surfaces. We describe how this approach speeds up MBH by dramatically shrinking the search space, and that it can be used to provide a quantitative criterion for good time to stop the MBH process because of a very high likelihood that the global minimum has already been found.

We discuss the main challenges of the trajectory optimization problem: The objective function $f(\mathbf{x})$ is nonlinear and has many local minima among the feasible $\mathbf{x}$; the equality constraints are generated by equations (surfaces) $g_i(\mathbf{x})$ that are non-linear, and, because feasible points must satisfy all of the equality constraints they can only exist at (oftentimes complicated) intersections of equality constraints; the inequality constraints are generated by equations (surfaces) $h_j(\mathbf{x})$ that are highly non-linear; and most of the search space is not worth exploring because most $\mathbf{x}$ are not feasible points.

We also explain that one of the most important issues and a key part of our on-going research is whether, in real trajectory optimization problems, filaments tend to be relatively long, continuous, and few, or short, discontinuous, dispersed and numerous. This is important because it determines the efficiency of the filament walking approach when it is applied to specific trajectory optimization problems.

MOTIVATION

Filament waking was developed as an improvement to MBH. Here, we clarify the fundamental differences between them, in order to explain the practical benefits of filament walking.

Filament walking searches for the location of $\mathbf{x^*}$, where $f(\mathbf{x^*})$ is the global minimum of $f(\mathbf{x^+})$ among the points that are feasible with respect to every equality constraint, namely the $\mathbf{x^+}$. During the process of performing "thorough" random walks on the $\mathbf{x^+}$, the filament walker "visits" $\mathbf{x^*}$ "almost surely". At every visit to any $\mathbf{x^+}$, the filament walker evaluates $f(\mathbf{x^+})$, and put the pair $(\mathbf{x^+}, f(\mathbf{x^+}))$ in a list of the points $\mathbf{x^+}$ it has visited. If the walk is sufficiently "thorough," $\mathbf{x^*}$ will be visited and the pair $(\mathbf{x^*}, f(\mathbf{x^*}))$ will be recorded in the list.

If one calls the set of points $\mathbf{x^+}$ that comprise a filament the domain, and $f(\mathbf{x^+})$ the range, it can be said that filament walking searches the domain in order to find the pair $(\mathbf{x}^*, f(\mathbf{x}^*))$. In contrast, MBH searches the range in order to find the pair $(\mathbf{x}^*, f(\mathbf{x}^*))$.

Here, it's important to recall that the domain of a filament is 1-dimensional whereas the domain of the space that needs to be explored in MBH is N-dimensional for large N. Attempting to thoroughly

search a bounded domain of R^N is nearly futile, whereas thoroughly searching a bounded domain of R^1 is quite reasonable. This can be shown using First Passage theory for random walks, but we leave that for a subsequent paper.

Here, we simply note that it is extremely hard to construct quantitative estimates that indicate how thoroughly MBH explored the N-dimensional space in which it searched for the global minimum. Therefore, for MBH, constructing measures of confidence in whether or not the optimizer is "done" by time t is very difficult. As a result, in a HOC type architecture such as EMTG which uses MBH in the inner loop, it is hard to have the MBH inner loop signal to the outer loop that it's "almost surely" done, and that the outer loop can "move-on." In optimizers which are entirely autonomous, such as EMTG, this kind of signaling is very important. Without such signaling, "thoroughness" can only be achieved by allocating a lot (usually excessive of amounts) of time to the MBH inner loop process. Or, one can allocate less time and risk the possibility that the MBH process was not sufficiently "thorough," and that the true f($\mathbf{x^*}$) was not really found.

In contrast, because filaments are 1D, making quantitative estimates of whether the optimizer has, by time t, has searched the domain with sufficient thoroughness, and thereby thoroughly searched the range, is not hard to do. First Passage Time theory for random walks can be used as a foundation for such estimates of "thoroughness".   In addition, First Passage Time theory also explains why "thoroughness" in 1D can be achieved in reasonable (practical) amounts of time "almost surely"; why a very long time is required in 2D; why an extremely long time is required in greater than 2D; and why for any number of dimensions much greater than 2, it gets futile very quickly as the number of dimensions increases.

This is not to say that the success of MBH in reasonable amounts of time is impossible or even unlikely. To the contrary, MBH usually finds the sought-after pair ($\mathbf{x^*}$, f($\mathbf{x^*}$)) so quickly and accurately that it is difficult to explain, in terms of mathematical analysis, exactly how it does it, although many heuristic explanations have been developed. MBH is not a random walk, but instead a sophisticated stochastic search process involving a "ratchet" – meaning that during in the process of optimization, the incumbent solution cannot "get worse" – and MBH searches the range, rather than the domain, to find the sought-after pair ($\mathbf{x^*}$, f($\mathbf{x^*}$)). We will comment further on the differences between filament walking and MBH later in this paper.

We did not set out to develop filament walking because MBH does not work well; it works very well. We set out to develop filament walking because it enables quantitative estimates for whether or not "thoroughness" has been achieved by step *s*.

In the paragraphs above, a number of words were put into quotes because they straddle, perhaps uncomfortably, the worlds of intuition and mathematical clarity: By "thoroughness" we mean that each point $\mathbf{x^+}$ on a filament is, with extremely high probability, visited at least once. This leads to a discussion to First Passage Time probabilities, and, although we will leave that for a subsequent paper, we are being careful to use words that will be consistent with that discussion. By "almost surely" we are using these words as they are used in probability theory. We mean that, for most practical purposes, such an outcome is extremely probable. And, by "ratchet" we refer to a ratchet like one finds in a watch or clock, meaning that progress, *i.e.* the time of day that is displayed, cannot slide backwards.

FORMAL STATEMENT OF THE PROBLEM

Many space-probe trajectory optimization problems (especially those that use "shooting" transcriptions), can be posed – and we propose, solved – as follows:

1. We define **x** as a point in a dense set of points **X** in solution space *S* in $\mathbf{R}^N$
    1.1. In typical problems, *N* is approximately 10 to 100
2. We seek global min(f(**x**)), subject to:

    2.1. $g_i(\mathbf{x})-c_i = 0$ for all I equality constraints

    2.2. $h_j(\mathbf{x})-k_j \leq 0$ for all J inequality constraints

3. In other words, we seek global min(f($\mathbf{x^{++}}$)), meaning global min(f($\mathbf{x}$)) only among the $\mathbf{x^{++}}$, where the $\mathbf{x^{++}}$ are strictly feasible points with respect to all of the equality and inequality constraints

4. We say that points $\mathbf{x^+}$ are feasible at least with respect to the set of equality constraints $g_i(\mathbf{x})-c_i$

5. The points $\mathbf{x^{++}}$, which also satisfy the inequality constraints, are a subset of $\mathbf{x^+}$.

6. Because we require that the $\mathbf{x^+}$ satisfy all of the equality constraints, they can only reside at the intersections of the $g_i(\mathbf{x})-c_i = 0$

7. Because the functions $g_i(\mathbf{x})$ that determine the equality constraints are known to be (at least) first-differentiable (at least) near $g_i(\mathbf{x})-c_i = 0$, and they are usually are reasonably "smooth," it is generally intuited that the gradients of the $g_i(\mathbf{x})$ near any $\mathbf{x^+}$ will usually "point" to that $\mathbf{x^+}$

8. Thus, it is commonly thought that the $\mathbf{x^+}$ reside in a "river" at the "bottom" of a "hyper-river-basin", meaning that a solver can often find an $\mathbf{x^+}$ by following the local gradient from a nearby point $\mathbf{x}$

9. We say the $\mathbf{x^+}$ comprise "filaments" in recognition of a protein transport model that inspired this work. We also refer to them as "filaments of feasibility" with respect to the equality constraints.

10. When the $\mathbf{x^+}$ form filaments, the filaments are 1-dimensional, even though they weave through multi-dimensional spaces.

11. When a problem has few enough and "long" enough filaments, and the gradients pointing to the filaments are "well-behaved", then the filaments can be found and "walked" along.

12. When a problem has a small number of filaments that can be found and walked along, then the search for global min(f($\mathbf{x^+}$)) collapses into one (or a few) 1-dimensional search problems, and the search among the $\mathbf{x^+}$ for global min(f($\mathbf{x^{++}}$)) becomes the search for min(f($\mathbf{x^+}$)) along the $\mathbf{x^+}$ that comprise the filament(s). The min(f($\mathbf{x^+}$)) found on each filament are then verified to not be in violation of any of the inequality constraints. Thereby, the search for global min(f($\mathbf{x^{++}}$)) can become very effective, efficient, and robust.

13. We call $\mathbf{x}$, such that f($\mathbf{x}$) = constrained global min(f($\mathbf{x}$)), $\mathbf{x^*}$.

14. Clearly, $\mathbf{x^*}$ is in the set of feasible points $\mathbf{x^{++}}$. Being a member of the set of points that are feasible with respect to the equality constraints, namely the $\mathbf{x^+}$, is a necessary condition for being in the set of feasible points $\mathbf{x^{++}}$. Finding min(f($\mathbf{x^+}$)) is the most challenging part of that process, and the work performed by filament walking.

15. In fact, filament walking provides a 1-dimensional array of f($\mathbf{x^+}$) per filament, sorted by size. A post-filament walking process merely eliminates those f($\mathbf{x^+}$) per filament for which that $\mathbf{x^+}$ violates any inequality constraint, resulting in the set of $\mathbf{x^{++}}$.

16. The smallest surviving f($\mathbf{x^{++}}$) is f($\mathbf{x^*}$), meaning constrained global min(f($\mathbf{x}$)), and that $\mathbf{x^+}$ is $\mathbf{x^*}$ <u>if</u> all the filaments were found and walked "thoroughly"

17. This "if" is a very important requirement, but before discussing it, we review some of the things that we were concerned might pose serious limitations in real applications (but apparently do not).

       We now illustrate these ideas with a simple example in which we purposely made the filament a straight line, without loss of generality, only because a straight-line filament is easy to simulate, visualize, and any curvy filament can be mapped into a straight line.

       Figure 1 shows a wireframe representation of the objective function f($\mathbf{x}$). The objective function is never seen in its totality by the optimizer, but we show it here *a priori* (having generated it for the sake of this example). Figure 2 is a heatmap representation of the same objective function f($\mathbf{x}$), but shown with

the filament as a light gray line. The heatmap representation of f(**x**) in Figure 2 is shown here because it will be a useful way to show the constrained global minimum f(**x***) later in Figure 6.

Figure 3 shows the objective function f along the filament, meaning f(**x**$^+$). The minimum found along the filament, min(f(**x**$^+$)), was found by the 2,000-step filament walk. Figure 4 shows a 2,000-step walk on the filament, including off-filament excursions from which the walker is returned to the filament by the gradient of g. The paths from the off-filament excursions, back onto the filament, are not drawn only because drawing them would make Figure 4 unreadable. The heatmap color coding in Figure 4 represents the value of the function g(**x**) that defines the equality constraint g(**x**$^+$) = 0 that, in turn, defines the filament.

Figure 5 shows the samples of f(**x**$^+$) taken during the 2,000-step walk on the filament. Here we see that, in this highly simplified problem, **x*** was found once among the x$^+$ (at step 1,740 out of 2,000). Finding **x*** at step 1,740 out of 2,000 does not allow much margin for error, and begs the question, "How many steps should one allow in order to assure that sampling f(**x**$^+$) by filament walking will assure the finding of **x***?" This gets to the heart of the "thoroughness" issue with respect to walking (rather than finding) the filament(s). We illustrate this issue further in Figures 7 and 8, but first we complete the cycle of Figures 1 through 5 by showing Figure 6 which shows that by finding the particular **x**$^+$ such that f(**x**$^+$) for that particular is min(f(**x**$^+$)), the pair (**x***, f(**x***)) was successfully found.

Figure 7 illustrates that min(f(**x**$^+$)) was found 13 times in 25,000 steps. We do not need min(f(**x**$^+$)) to be found 13 times; we need it to be found only once, but with extremely high probability of success. This is the "thoroughness" issue with respect to filament walking. We do not propose that the satisfactory answer is to "just take a great many steps." Rather, we were motivated to develop the filament walking methodology because we believe that it can support quantitative measures of "thoroughness" that will signal when to stop walking. That is work-in-progress and we intend for that to be the subject of a future paper. Figure 8 is simply the equivalent of Figure 4, but with 25,000 steps instead of 2,000.
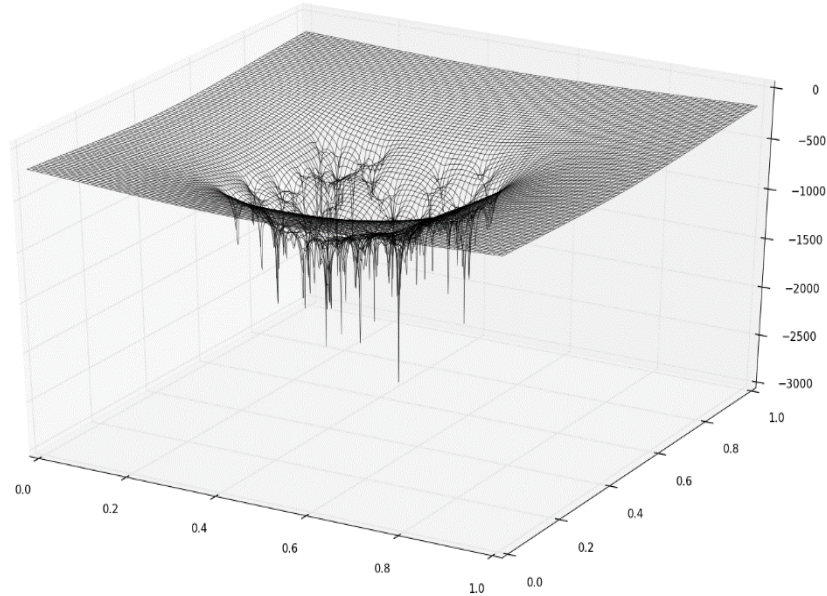

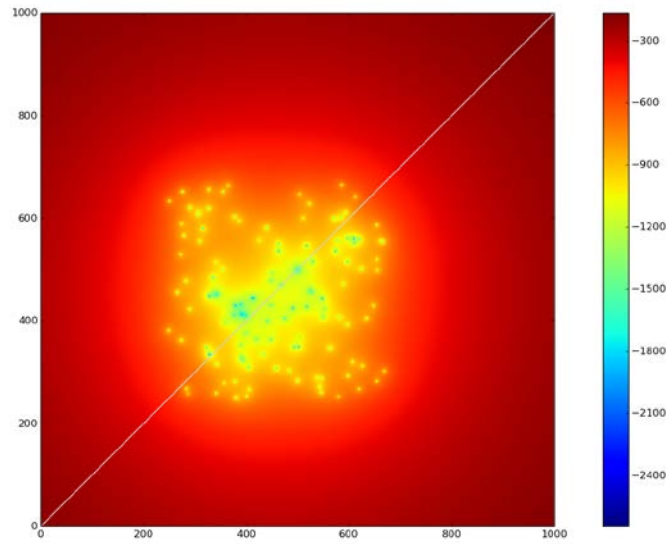Figure 1. A wireframe representation of an objective function f(**x**)

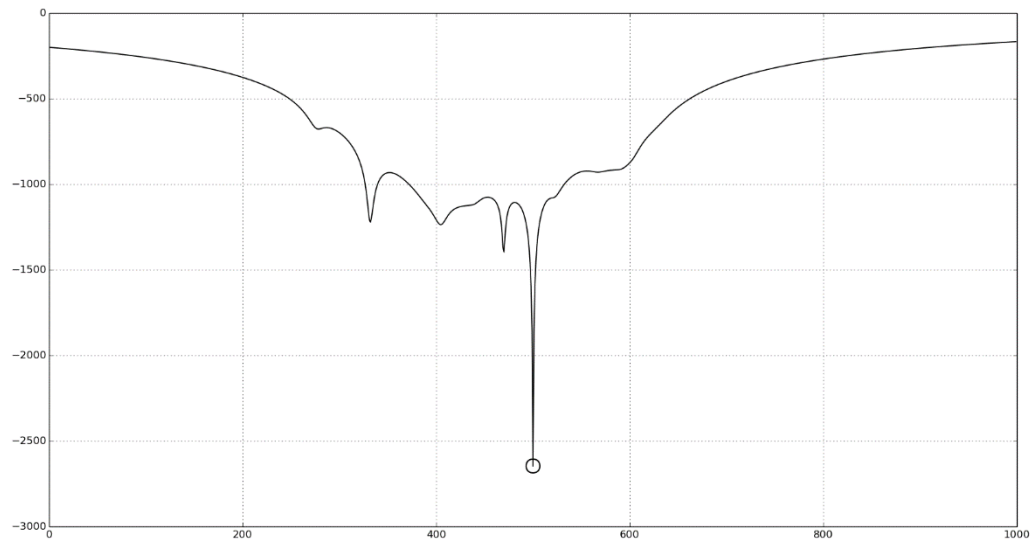Figure 2. A heatmap representation of an objective function f(**x**). The light-gray line represents the filament.

Figure 3. The objective function along the filament, meaning f(**x**⁺). The minimum found along the filament, min(f(**x**⁺)), was found by the 2,000-step filament walk.
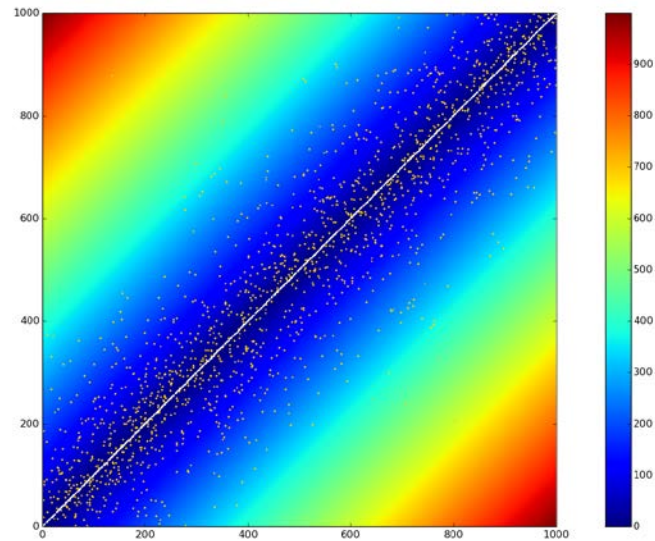
Figure 4. A 2,000-step walk on the filament (including off-filament excursions, from which the walker is returned to the filament by the gradient of g). The color coding represents the value of g.
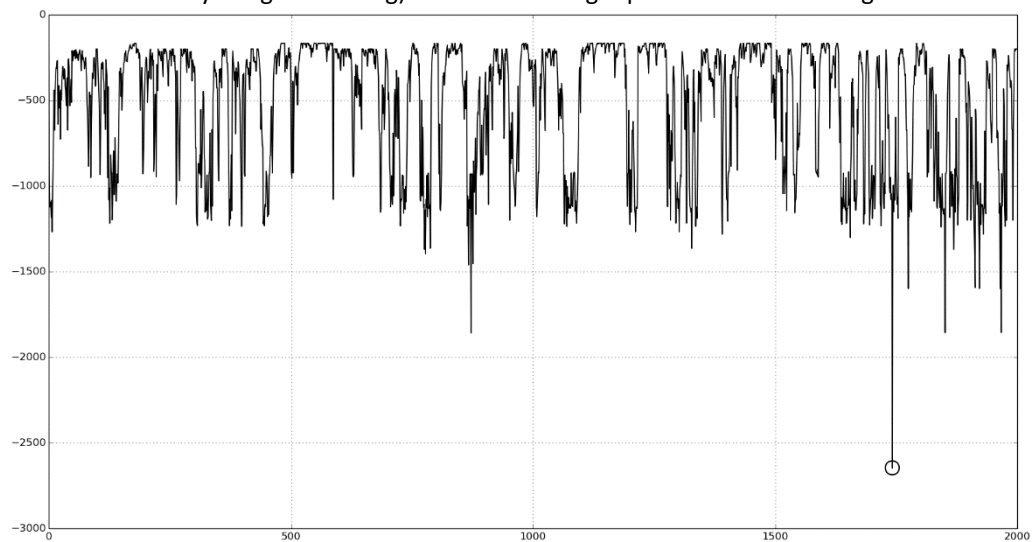

Figure 5. The samples of $\mathbf{x}^+$ taken during the 2,000-step walk on the filament. Here we see that, in this highly simplified problem, $\mathbf{x}^*$ was found once among the $\mathbf{x}^+$ (at step 1,740 out of 2,000)
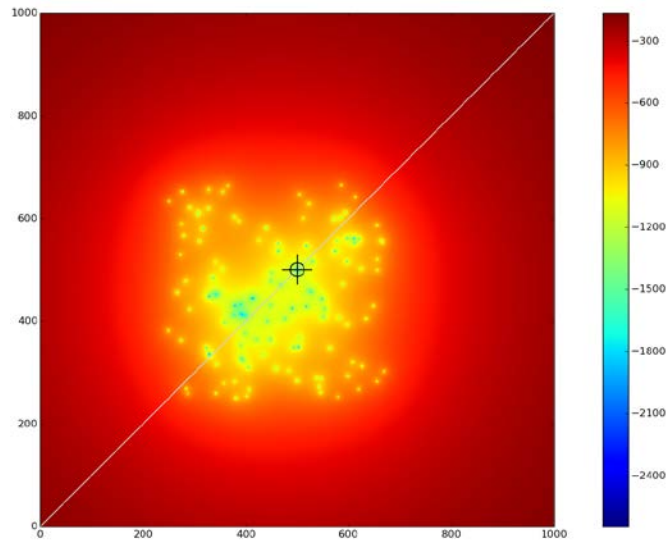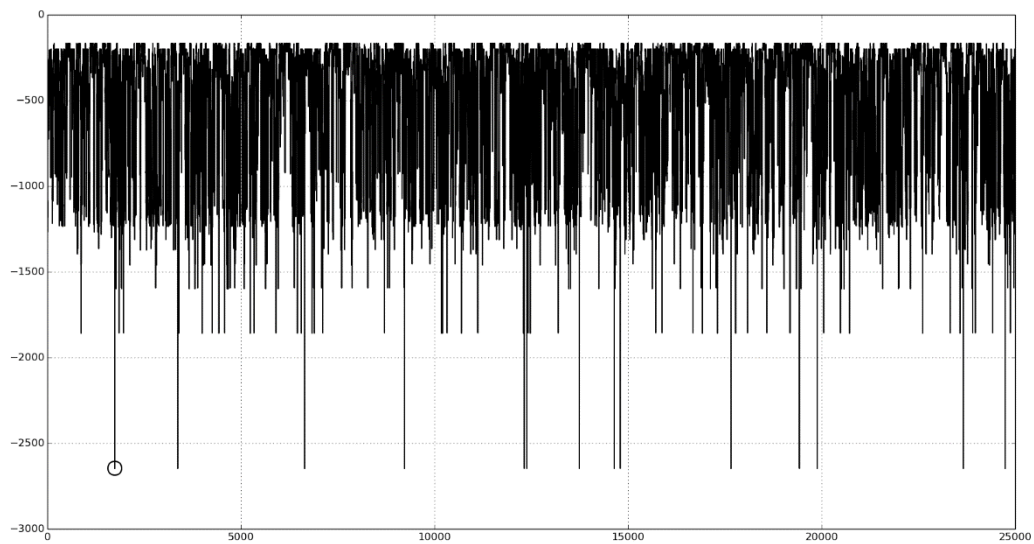
Figure 6. A heatmap representation of the same objective function f(**x**) as Figure 5 with min(f(**x**⁺)) found by the 2,000-step filament walk

Figure 7. The samples of f(**x**) taken during the 25,000-step walk on the filament. Here we see that, in this highly simplified problem, **x**⁺ was found 13 times in 25,000 steps.
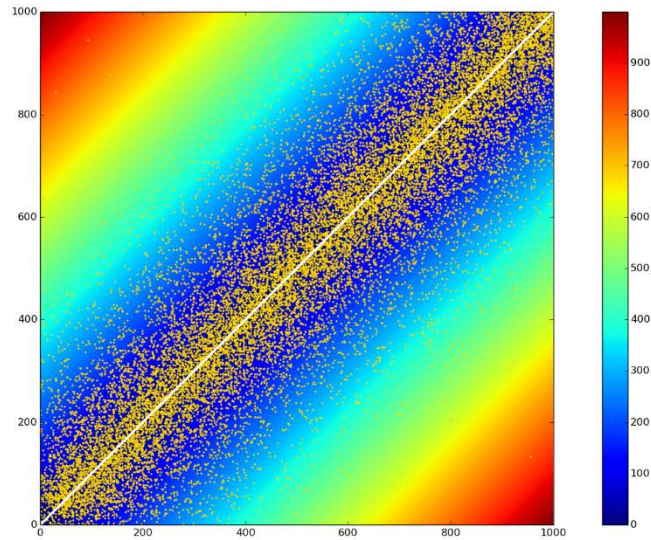
Figure 8. The 25,000-step walk on the filament (including off-filament excursions, from which the walker is returned to the filament by the gradient of g) that generated Figure 8.

INITIAL CONCERNS

At the start of this work, we had the following concerns that we believed would immediately reflect on the applicability of this work:

- Do filaments really exist in real and important space-probe trajectory optimization problems?
- Are the filaments few enough in number, and "long" enough, to be worth finding them and walking on them? It is easy to construct hypothetical cases in which filaments are too numerous and/or short for filament finding and walking to be worthwhile.
- Can filaments be found easily and quickly? This is a question about the extent and surface qualities of the gradients that surround the filaments. It is easy to construct hypothetical cases in which the gradients are unhelpful, or even destructive.
- If, while walking, walkers "fall off" of, or "stray from," filaments, can they be given the ability to quickly "find" their way back on to a filament, and resume walking in a productive direction? To some degree, this too is a question about the extent and surface qualities of gradients that surround and lead to the filaments. It's easy to construct hypothetical cases in which gradients are unhelpful, and even destructive, in this regard as well.
- How can we visualize filaments? Even though they are 1-dimensional, they are woven through high-dimensional spaces. We need ways to "unravel" them, in order to view and "measure" them, and in order to analyze the dynamics of walkers walking them – but we must unravel them in ways that preserve some of their properties.
- How can we assure that, "almost surely", all the filaments in a problem were found and walked?
- How can we assure that, "almost surely", each filament was walked so thoroughly that $\min(f(\mathbf{x}))$ along that filament was visited (sampled), rather than stepped-over ("overlooked")?
- How can we build quantitative confidence measures of whether all filaments were found, and all filaments that were found were walked sufficiently thoroughly, "almost surely," within some specifiable number of search steps?

Figure 9 is a 3-D illustration of the easiest case we imagined: There is at most one long filament, shown by the black line, and it is easy to find as shown by the brown random path terminating on the filament. Figure 10 is a 3-D illustration of a slightly harder case: There are at most a few filaments. At least the first and longest of them was easy to find. They all need to be found and walked "thoroughly". One is in the "corner" of the space, which may be somewhat challenging, and one is very short – a kind of "stutter" located at one end of the second longest filament in the center-right of the figure.

Figure 11 is an illustration of one of the very difficult cases we imagined. We were pleased and relieved to find that it is not characteristic of any of the real problems with which we've worked so far and may not be a property of multiple-shooting transcriptions. The surfaces generating the equality constraints $g_i(\mathbf{x})$ are such that the intersections of the equality constraints are discontinuous, and "chop" the filaments into many small, "randomly scattered" segments.
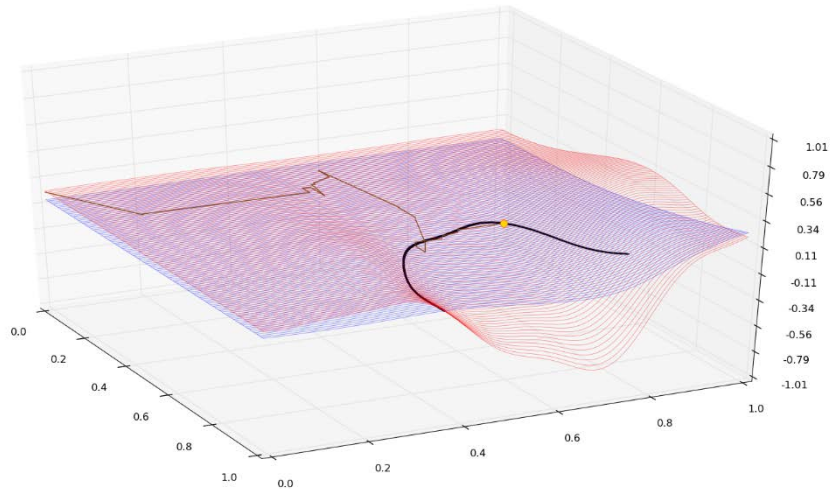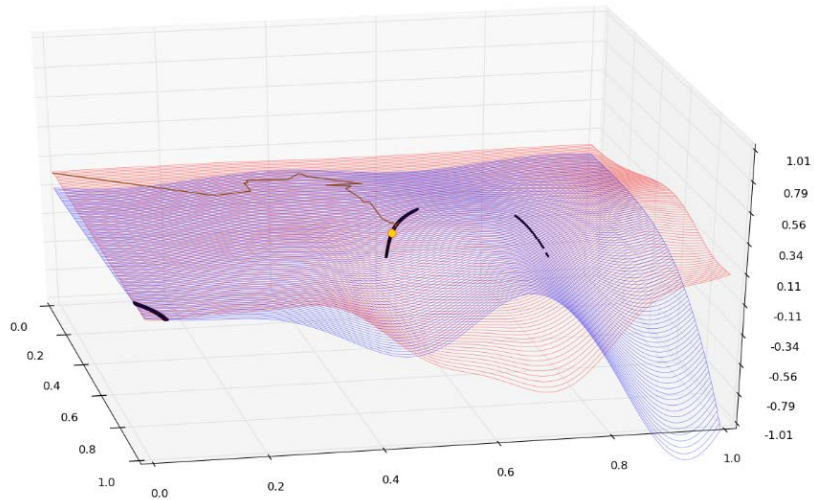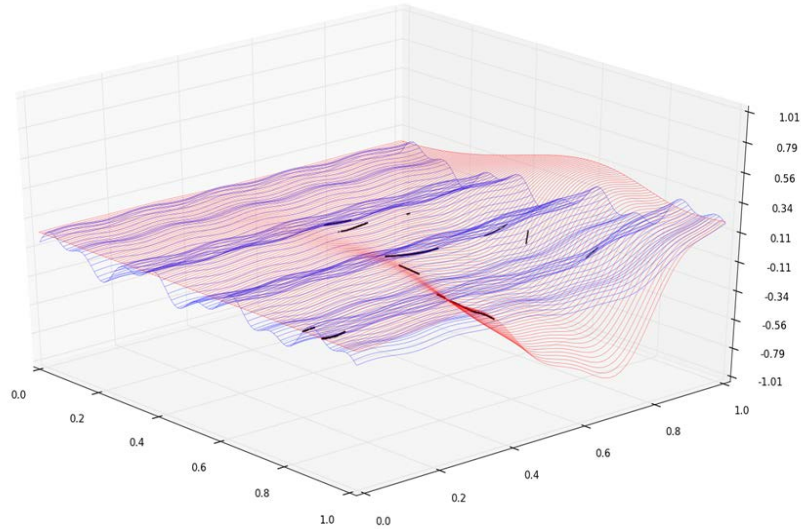

Figure 9

Figure 10

Figure 11

THE APPARENT NATURE OF REAL APPLICATIONS

What we found is that in real problems on which we've tested this so far:

A. Filaments being too numerous and/or too "short" does not seem to be a real problem. Filaments seem to be few enough and "long enough" to be worth finding and walking.

B. Gradients being too "rippled," "rugged," roughly textured, or locally flat, does not seem to be a real problem. In real problems, gradients seem to be "smooth" and "well-behaved." Gradients of the functions $g_i(\mathbf{x})$ near an $\mathbf{x}^+$ seem to "point" to that $\mathbf{x}^+$ from a nearby $\mathbf{x}$, "almost always." That, in turn, helps to make finding filaments fast and easy, and helps walkers – when they "fall off," or stray from, filaments – quickly "find" their way back on to the filament, and resume their walking in a productive direction.

C. Filaments can be "unraveled" using highly local (step-wise) distance-preserving transformations, thereby allowing them to be visualized, and the dynamics of walkers walking them to be studied

We next discuss a benchmark orbital mechanics problem – an interplanetary transfer from Earth to Mars. Here we make the assumption that the spacecraft launches from the Earth with a C3 of up to 48.5 km^2/s^2 and travels to Mars, where it performs a rendezvous maneuver to match velocity at the edge of the sphere of influence. This is of course a more expensive strategy than one would perform in real design because it does not take advantage of Mars gravity to aid the arrival maneuver but it is a good example for this paper. The spacecraft is allowed to perform up to one deep space maneuver (DSM) at any point during the trajectory and no gravity assists are allowed. The objective function is to minimize Δv, defined as the sum of the individual contributions from the DSM (if any) and the rendezvous maneuver.

This problem is constructed using the "multiple gravity assist with *n* deep space maneuvers using shooting," or MGAnDSMs transcription in EMTG [10]. MGAnDSMs, developed by Vavrina *et al.*, is a two-point direct shooting transcription in which the trajectory is propagated forward from the first body in each phase and backward from the second body. A set of "match point" constraints enforce trajectory continuity. In the context of this work, one can think of the set of decision vectors that satisfy the match

point constraints as the filament **x**⁺. Table 1 describes the assumptions and optimal objective function value for this example and Figure 13 is a plot of the optimal solution.

Figure 14 is the filament that we found in our first benchmark problem. We were very excited to find it, and see that it was much better "behaved" than we feared! Figure 14 is a projection onto 2-D a filament weaving through a space of 13 dimensions. It was easy to find, it is long, and – in this problem – we believe it is unique. It is easily walkable, and our interpretation of a filament "unraveling" technique we developed suggests that the gradient of g near the filament is "smooth" enough to quickly guide the walker back on to the filament when the walker makes a step that is an off-filament excursion, or "falls off." The filament is color coded by the value of the objective function f(**x**), where darker is better.

Figure 15 shows our first "unraveling" of a filament. The large jumps in which the distance travels resets to zero are simply multiple trials of the algorithm. The small jumps indicate where the walker "fell off" and was guided back on to the filament elsewhere from where it "fell off." Apart from the jumps, we see continuous walking along the filament. Note how the "slides," or NLP solves shown in red, onto the filament from off-filament excursions, contribute little to the filament walking. This is because the slides are typical normal to the filament.

This example is a relatively "easy" orbital mechanics problem because there is only a single phase – from Earth to Mars – and therefore the match point constraints only trace out one continuous filament. More difficult problems with multiple gravity assists (and therefore multiple phases) may have multiple filaments or a single discontinuous filament, but that is beyond the scope of this paper.

Table 1: Parameters of the Interplanetary Trajectory Benchmark Problem

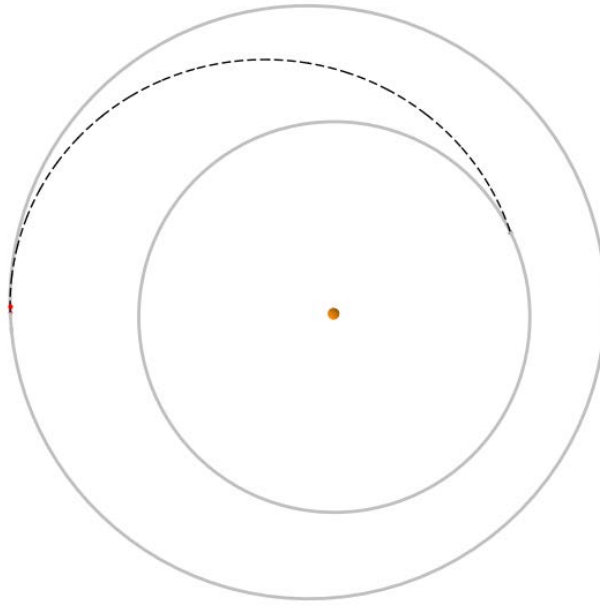|	
Transcription	MGAnDSMs [10]
Launch date bounds	2/16/2006 - 11/12/2008
Departure type	launch with $C_3 \leq 48.5$ km^2/s^2
Arrival type	chemical rendezvous
Objective function	minimize Δv
Flight time upper bound	500 days
Known global optimum (from MBH)	2.343 km/s

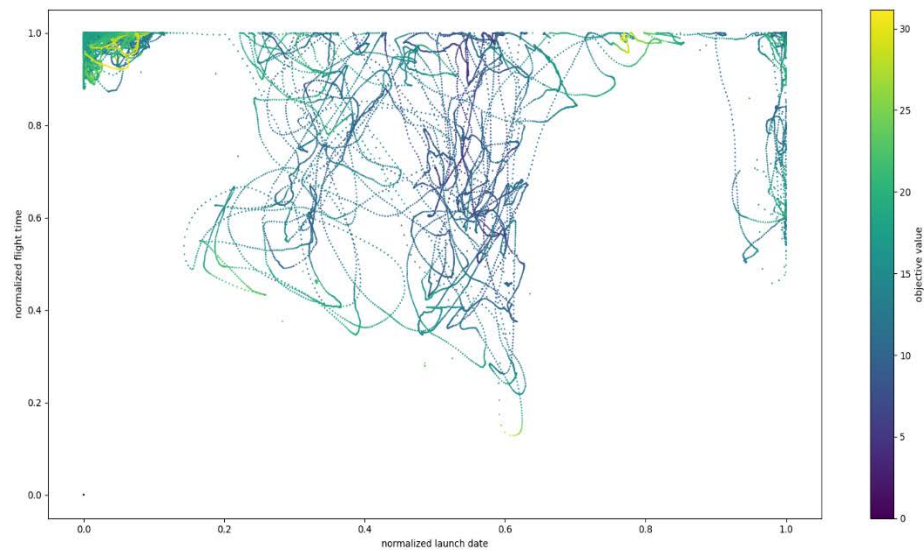Figure 13: Optimal Trajectory for the Earth to Mars example problem

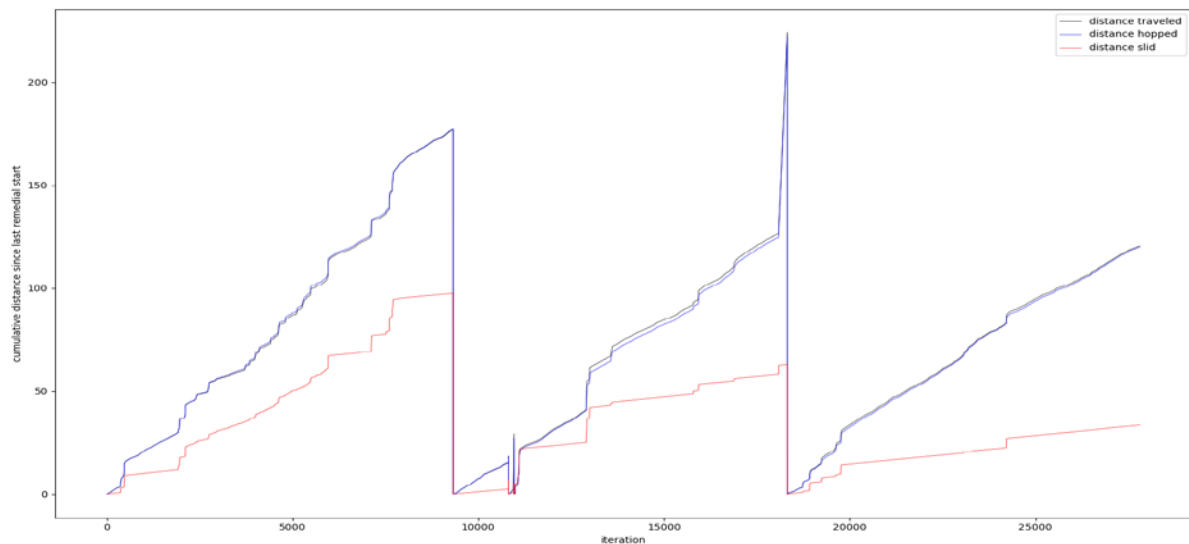Figure 14: Filament walked in the Earth to Mars example

Figure 15: Cumulative distance walked in each trial of the Earth to Mars example problem

REMAINING CONCERNS

While the apparent nature of real problems has encouraged us in terms of the applicability of this work, we are left with the following concerns (which we are pleased to now be facing):

A. How can we assure that, "almost surely," all of the filaments were found and walked?

B. How can we assure that, "almost surely," each filament was walked so thoroughly that $\min(f(\mathbf{x}^+))$ along that filament was visited, rather than stepped-over?

C. How can we build quantitative confidence measures for evaluating whether all filaments were found and walked with sufficient thoroughness, within some number of search steps?

D. We call these the three aspects of the "thoroughness" problem.

The remaining concerns relate to essentials in the theory of random walks, especially "probability of return to origin", and First Passage Time problems.

The key issues are:

- The dimensionality of the space being searched. In general, First Passage Time problems in 1-D are analytically straightforward, problems in 2-D are analytically difficult but solvable, but problems in dimension greater than 2 are analytically unsolvable. We are pleased to keep in mind that filaments, once found, are 1-D, even though they weave through spaces of very high dimensionality, *e.g.* 20-D.
    - Probability theorists joke that a drunken walker, trying to find their way home, would be fortunate to have gotten lost in 1-D. If they get lost in 2-D, they will find their way home eventually, but it could take a very long time. However, in 3-D, or any dimension higher than 2-D, it's just about hopeless. There are, of course, formal mathematical ways of saying this.
- The geometry of the neighborhood surrounding $x*$, such that $f(x*) = min(f(x^+))$, on a given filament. At one extreme, $f(x*)$ may be an infinitely deep singular point in an otherwise "flat" neighborhood, meaning that it is at the bottom of an infinitely deep funnel having an infinitesimally narrow mouth. At another extreme, $f(x*)$ may be at the bottom of a shallow, *i.e.* nearly flat, funnel that has a very wide mouth. Everywhere between these two extremes, $f(x*)$ may be at the bottom of a funnel of some depth and a mouth of some width.
- The nature of the stochastic dynamics, *e.g.* the particular kinds of randomness, or regime switching between randomness and gradient following, used to drive the stochastic search.
- In real problems, filaments are 1-D, and funnels are not infinitely deep and narrow, or extremely shallow and wide-mouthed. Therefore, we expect that the thoroughness problem for filament walking results will be more favorable than classic "probability of return to origin," and First Passage Time, problems are in general. We also expect, and already have empirical evidence indicating, that the nature of the stochastic dynamics used to drive the stochastic search will make a big difference. We suspect that for the kinds of funnels and clusters of funnels that are typical of real problems, the best kinds of stochastic dynamics for searching for filaments, and minima along filaments, will extend ideas the ideas first developed in a 2014 paper [6] by the current authors. However, this needs to be investigated, demonstrated, and if empirically shown to be true, explained analytically.
- The issues of "thoroughness" are critical and the focus of our continuing work.

PSEUDO-CODE

<u>**Find a filament**</u>
while filament_not_yet_found **== true***:*

      Use MBH to find an $\mathbf{x}^+$ such that $\sum_i^I g_i^2(\mathbf{x}^+) \leq \epsilon \,\forall\, i.$ By definition, such an $\mathbf{x}^+$ is on a filament.

<u>**Walk a filament**</u>
set $s = 0$
set reverse_flag = false
set slide_attempts = 0
while filament_not_yet_walked_thoroughly **== true:**

      ++s
      Sample f($\mathbf{x}_s^+$) and append the pair ($\mathbf{x}_s^+$, f($\mathbf{x}_s^+$)) to the list of points found on this filament
      if $s == 1$:

            Make a small zero-drift uniformly distributed random excursion
      else:

            Make a drift-biased uniformly distributed random excursion where:
            if reverse_flag == false:

                  drift equals k($\mathbf{x}_s^+$ **-** $\mathbf{x}_{s-1}^+$)
            else

                  drift equals **-**k($\mathbf{x}_s^+$ **-** $\mathbf{x}_{s-1}^+$)
                  set reverse_flag = false
            if $\sum_i^I g_i^2(\mathbf{x}_s^+) > \epsilon$ for any i:

                if s ≤ allowable_number_of_walk_attempts**:**

                      Slide along gradient (normal to filament) to get back onto the filament
                      ++slide_attempts
                      if slide finds a new $\mathbf{x}^+$, *i.e.* $\sum_i^I g_i^2(\mathbf{x}_s^+) \leq \epsilon \,\forall\, i$:

                            slide_attempts = 0
                      if slide_attempts > allowable_number_of_slide_attempts**:**

                            set slide_attempts = 0
                            set reverse_flag = true
                else:

                    Restart algorithm; find a new filament.

FUTURE WORK
Our future work will focus on the issues of:
- "Thoroughness"
- Constructing quantitative measures of confidence in the interim solution, so that the optimizer can autonomously determine when the current inner loop optimization "is done"
- Speeding-up the code so it can be benchmarked against the current production version of EMTG on identical benchmarking problems
- Massive (statistically significant) amounts of benchmarking against the current production version of EMTG

COMPARISION OF FILAMENT WALKING to MONOTONIC BASIN HOPPING

Monotonic Basin Hopping (MBH), as described and implemented by Englander and Englander [6] is the method used to find filaments, so the differences between filament walking and MBH concern the walking of filaments.

MBH does not have an explicit notion of filaments, *e.g.* finding them, counting them, estimating the likelihood that all filaments present in a problem have been found, evaluating whether they have been walked "thoroughly," *etc.*. However, MBH, at least as it is implemented in NASA Goddard's EMTG software, makes use of gradients to find points that are feasible with respect to the equality constraints. We interpret that as an implicit search for filaments, even though MBH does not attempt to walk every findable filament thoroughly, end-to-end, in search of min(f(**x**)) along that filament.

Therefore, we are optimistic that filament walking can be more efficient than MBH. We acknowledge that so far, this is just an intuition that needs to be demonstrated and explained analytically. Intuitively, we liken this to the difference between (filament walking) finding a very long strand of fine steel piano wire in a big haystack, and following that wire to find a small pearl on it, versus (MBH) looking for that small pearl directly. A better analogy is that MBH is looking for the small pearl directly with the help a metal detector, but the intuitive implications of these two analogies are not terribly different. Either way, imagine that big haystack is of large dimensionality.

We expect that, once we have assured sufficient "thoroughness" in finding and walking the filaments, filament walking will demonstrate itself to be more efficient, and at least as accurate, as MBH. We also expect that filament walking will be more robust than MBH in the sense that we believe filament walking will not get "stuck" in the ways that we have seen MBH get "stuck" when encountering neighborhoods in gradients that are locally too flat or rough.

Finally, our ambition is to eventually develop quantitative confidence measures of whether all filaments were found, and whether all the filaments that were found were walked with sufficient thoroughness, "almost surely," within some specifiable number of search steps. This is needed to address the question of when the optimization process can be stopped. We believe that, because filament walking has an explicit notion of filaments, filaments are 1-D, and the walking process tracks the positions along filament(s) where the x^+ are visited and the f(x^+) are sampled, "thoroughness" can be stochastically controlled – and evaluated – in ways that it cannot be with MBH. By contrast MBH has no known way to measure "thoroughness" at all. Indeed, we believe that this is the real value of filament walking.

CONCLUSION

In this work, we introduced a new method of global search known as "filament walking," which takes advantage of the 1-dimensional "filaments of feasibility" generated by the equality constraints in multiple shooting transcriptions. By collapsing the *N*-dimensional decision space into a 1-dimensional filament, we construct a framework in which it is possible to thoroughly explore the feasible subspace. We developed a "filament walker" that can first locate and then explore the filament. We demonstrated the walker on two benchmark problems, including one interplanetary trajectory design example.

The filament walker, unlike other global search methodologies such as MBH, provides a framework in which we can address the question of "thoroughness" – did we find all of the filaments, and, if so, did we walk each of them sufficiently "thoroughly" such that we can be "almost sure" that we have found the global optimum? This paper does not fully answer that question – rather, it is a progress report along the way and provides a framework for how we may measure and control "thoroughness" in future work.

REFERENCES

[1] A. Bryson and Y. Ho, Applied Optimal Control. Taylor and Francis, 1975.

[2] J. A. Sims and S. N. Flanagan, "Preliminary Design of Low-Thrust Interplanetary Missions," AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska, Paper AAS 99-338, August 1999.

[3] P. E. Gill, W. Murray, and M. A. Saunders, "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," SIAM J. Optim., Vol. 12, Jan 2002, pp. 979–1006, 10.1137/s1052623499350013.

[4] R. H. Leary, "Global Optimization on Funneling Landscapes," Journal of Global Optimization, Vol. 18, No. 4, 2000, pp. 367–383, 10.1023/A:1026500301312.

[5] C. H. Yam, D. D. Lorenzo, and D. Izzo, "Low-thrust trajectory design as a constrained global optimization problem," Vol. 225, SAGE Publications, Aug 2011, pp. 1243–1251, 10.1177/0954410011401686.

[6] J. A. Englander and A. C. Englander, "Tuning Monotonic Basin Hopping: Improving the Efficiency of Stochastic Search as Applied to Low-Thrust Trajectory Optimization," 24th International Symposium on Space Flight Dynamics, Laurel, MD, ISSFD, May 2014.

[7] J. Englander, M. Vavrina, D. Hinckley, and D. Ellison, "Global Optimization of Low-Thrust Interplanetary Trajectories Subject to Operational Constraints," AAS/AIAA Space Flight Mechanics Meeting, February 2016.

[8] J. A. Englander, B. A. Conway, and T. Williams, "Automated Mission Planning via Evolutionary Algorithms," Journal of Guidance, Control, and Dynamics, Vol. 35, Nov 2012, pp. 1878–1887, 10.2514/1.54101.

[9] J. A. Englander and B. A. Conway. "Automated Solution of the Low-Thrust Interplanetary Trajectory Problem", Journal of Guidance, Control, and Dynamics, Vol. 40, No. 1 (2017), pp. 15-27, https://doi.org/10.2514/1.G002124.

[10] M. Vavrina, J. Englander, and D. Ellison, "Global Optimization of N-Manuever, High-Thrust Trajectories Using Direct Multiple Shooting," AAS/AIAA Space Flight Mechanics Meeting, February 2016.