AIAA Aviation
13-17 June 2016, Washington, D.C.
16th AIAA Aviation Technology, Integration, and Operations Conference

Risk-Hedged Approach for Re-routing Air Traffic Under Weather Uncertainty

Alexander V. Sadovsky* and Karl D. Bilimoria†
NASA Ames Research Center, Moffett Field, CA 94035

This paper explores a new risk-hedged approach for re-routing air traffic around forecast convective weather. In this work, flying through a more likely weather instantiation is considered to pose a higher level of risk. Current operational practice strategically plans re-routes to avoid only the most likely (highest risk) weather instantiation, and then tactically makes any necessary adjustments as the weather evolves. The risk-hedged approach strategically plans re-routes by minimizing the risk-adjusted path length, incorporating multiple possible weather instantiations with associated likelihoods (risks). The resulting model is transparent and is readily analyzed for realism and treated with well-understood shortest-path algorithms. Risk-hedged re-routes are computed for some example weather instantiations. The main result is that in some scenarios, relative to an operational-practice proxy solution, the risk-hedged solution provides the benefits of lower risk as well as shorter path length. In other scenarios, the benefits of the risk-hedged solution are ambiguous, because the solution is characterized by a tradeoff between risk and path length. The risk-hedged solution can be executed in those scenarios where it provides a clear benefit over current operational practice.

I. Introduction

Nominal flight routes are typically determined by the aircraft operators (e.g., airline dispatchers) based on their preferences such as wind-favorable routing, while conforming with the structure of existing jet routes and pre-defined spatial constraints such as airspace reserved for military operations. This route, essentially a sequence of latitude/longitude waypoints, is transmitted to the air traffic service provider (ATSP) via the flight plan filed by the aircraft operator for each flight. If bad weather is forecast along the user-preferred route, the ATSP will re-route the affected aircraft around the weather-impacted airspace. Local re-routes may also be required to avoid congested airspace. This re-routing function of the ATSP is one aspect of Traffic Flow Management (TFM) [1, 2]. A more detailed description of TFM and current practices can be found in [3].

Research on various aspects of TFM has been conducted for the past several decades; see [4] for a high-level review in the broad context of Air Traffic Management (ATM). Works focused on addressing air traffic congestion include Refs. [5–8]. Works that include weather among the uncertainties affecting traffic management and congestion include [3] and [9–16]. The combinatorial nature of routing problems is manifest in most of these works through the use of discrete optimization models. An automated system called Dynamic Arrival Routes, designed for re-routing arrival flights around adverse weather regions that occur before the flight's freeze horizon [17], is described in [18].

This paper focuses on the problem of re-routing air traffic around strong convective weather affecting large volumes of airspace adversely. Standard TFM practice is to design such re-routes without explicit consideration of aircraft-to-aircraft separation requirements; the separation is implicitly enforced by ensuring that local airspace units (called *sectors*; see [4]) are not overloaded so that sector controllers are able to maintain aircraft-to-aircraft separation. A class of auto-regressive models for the automatic prediction of sector demands is proposed in [13].

Since a large number of aircraft must be re-routed a substantial distance away from their nominal (i.e., preferred) routes, traffic flow managers typically need to make re-routing decisions several hours in advance. However, due to the dynamic nature of convective weather events, weather forecast products with look-ahead times of several hours have significant uncertainties associated with their forecasts. An important objective of TFM is to maximize flight efficiency (e.g., minimize path stretching for re-routes) while maintaining flight safety. Hence current operational practice is to plan strategically for the most likely (highest-risk) weather event, and then tactically make any necessary adjustments as the weather evolves.

*Aerospace Engineer, Aviation Systems Division, Mail Stop 210-6, alexander.v.sadovsky@nasa.gov. Member, AIAA.
†Aerospace Engineer, Aviation Systems Division, Mail Stop 210-10, karl.bilimoria@nasa.gov. Associate Fellow, AIAA.

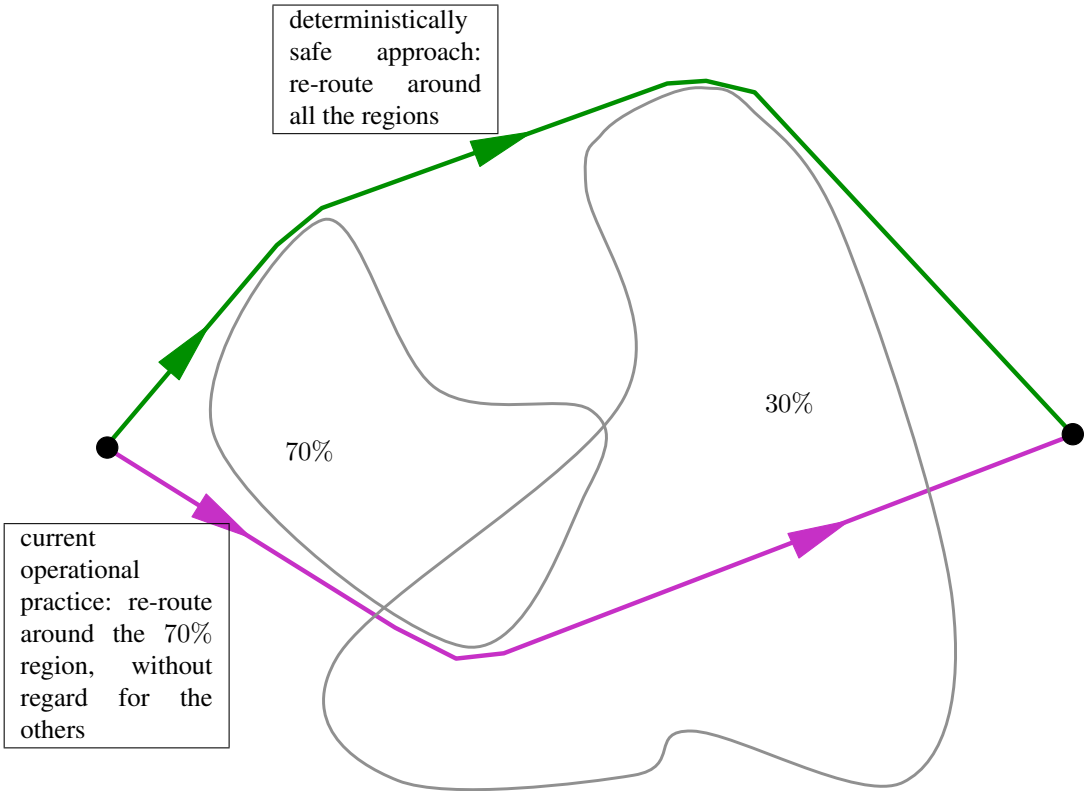

Figure 1: Two weather instantiations, with the respective occurrence risks 70% and 30%. Two possible re-routings are shown in two colors.

As a simple illustrative example (shown in Figure 1), consider two possible instantiations of convective weather forecasts, each represented by a region of airspace, with occurrence risks of 70% and 30%. As stated above, current operational practice (shown in magenta in the Figure) for this scenario would be to strategically re-route around the 70% region while penetrating the 30% region, and then tactically re-route around the 30% region if it materializes instead. At the other extreme, a "guaranteed", or deterministically safe, solution (shown in green) would be to re-route around both regions, but this would come at the expense of additional path length for this highly conservative re-routing. In this paper we explore an alternative re-routing strategy called risk hedging which seeks to minimize the risk-adjusted path length. In the context of the illustrative example above, this hedged re-routing may penetrate a smaller portion of the 30% region at the expense of a slightly longer path length, thereby softening the large path-length penalty of a tactical re-route should the 30% region materialize.

The model used herein, including the specification of the inputs (weather forecasts and the corresponding risk coefficients), the desired output, and the problem statement, are given in Section II. The analysis in Section III shows that the solution space of the problem turns out to have a well-understood geometric structure. This analysis includes a proof that there exist scenarios in which the risk-hedged solution outperforms current operational practice in both risk reduction and path length economy. The algorithm used herein for solving the problem is outlined and given as a flowchart in Section IV. Numerical solutions to a sample of instances of the problem are shown and analyzed in Section V.

II. The model and problem statement

A. Weather forecast characteristics

A forecast, several hours into the future, of the evolution of convective weather is needed for strategic planning of re-routes that avoid high-intensity weather regions. Currently available aviation weather forecast products typically provide a single (deterministic) instantiation of future weather at each of a series of look-ahead times. However, the Collaborative Convective Forecast Product currently provides up to two such instantiations: low-confidence (25 - 49%) and high-confidence (50 - 100%). Some aspects of air traffic management decision-making subject to weather uncertainties are covered in Refs. [19–22].

The far-term research reported here assumes the availability of a (yet to be developed) ensemble weather forecast product that provides multiple potential instantiations (their total number denoted here by N_w) of future weather, indexed by the symbol w. These instantiations can be thought to arise in the following manner. A currently known impassable weather region will evolve by moving and deforming. After some look-ahead time, T, the region will assume a new shape and a new spatial position, becoming a new region, which is regarded here as a possible future (at time T) weather instantiation. The evolution, however, cannot be predicted with certainty. Instead, a forecast will provide N_w mutually exclusive possible weather instantiations expected at T. Each instantiation w (the corresponding spatial region will be denoted by B_w) is accompanied by a nonnegative scalar (α_w) that is a measure of confidence that the associated instantiation will actually occur. Regions B_w can overlap.

In this work, flying through a more likely weather instantiation is considered to pose a higher level of risk. Hence, a forecast's confidence measure is considered to indicate the risk of flying through the associated weather region, and α_w is called the risk coefficient of w. For a given set of potential weather instantiations,

$$w = 1, 2, \ldots, N_W, \tag{1}$$

the associated risk coefficients may or may not add up to unity depending on the characteristics of the forecast methodology. If the "raw" risk coefficients do not add up to unity, they are normalized by their sum, and are henceforth assumed always to satisfy

$$\sum_{w=1}^{N_w} \alpha_w = 1, \text{ each } \alpha_w \text{ nonnegative.} \tag{2}$$

The role of the weather instantiation w in the model will be specified below. To summarize the above notation: for weather instantiation w, symbol B_w denotes the corresponding spatial region (impassable to flights), while symbol α_w denotes the corresponding risk coefficient.

B. Risk-adjusted cost of a re-route

The cost function will be defined by first defining a risk field which assigns a scalar $P(\mathbf{x})$ to each point $\mathbf{x}$ in the airspace (here $\mathbf{x}$ is a coordinate vector, of dimension 2 or 3, specifying a location in the airspace), as follows. For each weather instantiation w from among (1), let B_w denote the region of the airspace affected by that weather instantiation, and let $I_w(\mathbf{x})$ be the indicator function for B_w:

$$I_w(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ lies in } B_w \\ 0 & \text{otherwise.} \end{cases}$$

It follows that $\sum_{w=1}^{N_w} \alpha_w I_w(\mathbf{x})$ is the sum of all the risk coefficients of those regions B_w that contain $\mathbf{x}$ and can, therefore, be interpreted as the "total exposure to the risk of adverse weather at $\mathbf{x}$." This exposure is zero if $\mathbf{x}$ lies outside of all the B_w's.

Define the risk field $P(\mathbf{x})$ to have, at $\mathbf{x}$, the value

$$P(\mathbf{x}) = \frac{1}{1 - \sum_{w=1}^{N_w} \alpha_w I_w(\mathbf{x})}. \tag{3}$$

In particular, if $\mathbf{x}$ lies outside of all the B_w's, then $P(\mathbf{x}) = 1$; otherwise, $P(\mathbf{x}) > 1$. If all the B_w's have a common point $\mathbf{x}$, then $P(\mathbf{x}) = +\infty$. These properties of $P(\mathbf{x})$ are instrumental in the computation of optimal paths, described below, where $P(\mathbf{x})$ is used to assign cost to the candidate segments for a sought re-routing: the segments of infinite cost will be correctly interpreted as categorically unsuitable for flying. The quantity $\sum_{w=1}^{N_w} \alpha_w I_w(\mathbf{x})$ does not have this property. A notional example illustrating the behavior of P is shown in Figure 2; in particular, since there is no point common to all the weather regions, P never attains the infinite value. For the point $\mathbf{x}$ shown, $P(\mathbf{x}) = 1/(1 - (0.2 + 0.5)) = 3.3\overline{3}$.

If an aircraft is to fly a route, parameterized by arc length s in the form $\mathbf{x}(s)$, from an initial point $\mathbf{x}^{\text{in}}$ to a final point $\mathbf{x}^{\text{fi}}$ in the airspace, then the *risk-adjusted cost* of flying an elemental arc length $\mathrm{d}s$ at s is defined as

$$P(\mathbf{x}(s))\, \mathrm{d}s,$$

and the risk-adjusted cost of flying the re-routing path is, consequently, the integral

$$\int_{\text{path}} P(\mathbf{x}(s))\, \mathrm{d}s. \tag{4}$$

This definition of cost immediately implies that the aircraft's cost of flying the route:

- is at least as large as the geometric length of the route (when none of the route is affected by weather, one has $P = 1$ along the entire route), and
- is the higher, the more the risk of encountering adverse weather along the route.

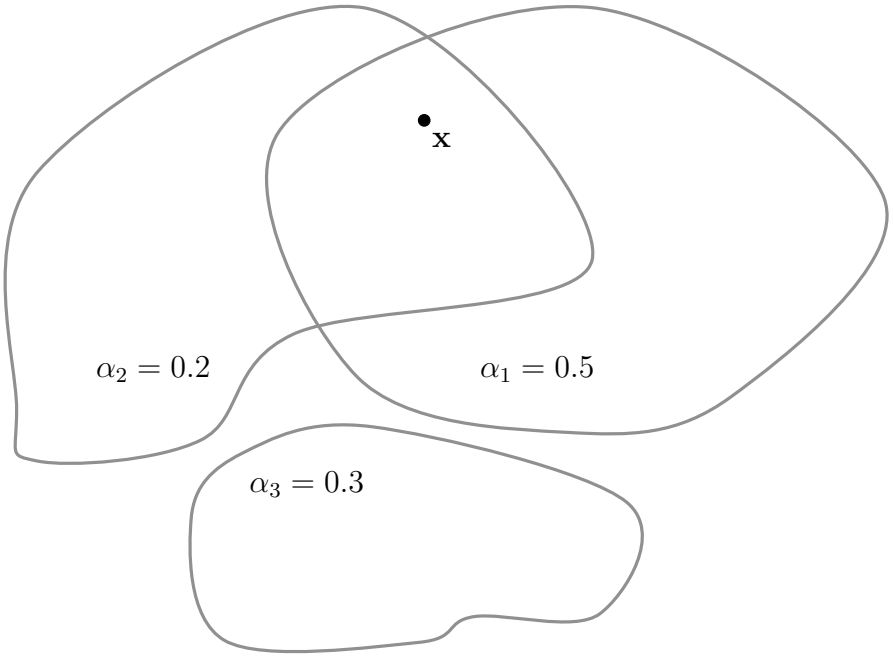

Figure 2: An illustration of the behavior of P for a case of $N_w = 3$ possible weather instantiations.

C. Problem statement

The central problem of this paper, to determine a re-routing that minimizes the adjusted risk, is as follows.

Risk-Hedged Re-routing Problem:
Given

- *prescribed initial and final positions in the airspace, $\mathbf{x}^{in}$ and $\mathbf{x}^{fi}$, and*
- *a collection (1) of the possible weather instantiations w, each accompanied by a risk coefficient α_w and the region B_w affected by weather w, and with the risk coefficients satisfying (2),*

find a path from $\mathbf{x}^{in}$ to $\mathbf{x}^{fi}$ that minimizes the cost, (4).

III. Insight into the structure of optimal solutions

The statements made in this section about the solutions to the problem of Section II.C hold if each weather region B_w has a sufficiently smooth (e.g., piecewise differentiable) boundary without self-intersections.

That P is piecewise constant, can be seen directly from its definition (3). Therefore, the problem of Section II.C is equivalent to the following optics problem: find the fastest route a ray of light will take from $\mathbf{x}^{\text{in}}$ to $\mathbf{x}^{\text{fi}}$ if the regions where P is constant represent media with different light permittivity indices; namely, the speed of light propagation in a constancy region that contains a point $\mathbf{x}$ is equal to $1/P(\mathbf{x})$. By Fermat's Principle of Least Time [23], such a route is piecewise linear. In particular, its portion contained in a region of the airspace in which P is constant is always a linear segment. Therefore, the solutions to the problem of Section II.C are piecewise linear trajectories, linear inside each constancy region. The piecewise linear geometry of the solutions will be instrumental in the numerical algorithm, given in Section IV, for solving the problem of Section II.C: the algorithm will at the outset seek solutions which are piecewise linear.

If a risk-hedged solution never outperformed the operational practice in both risk and path length, the benefits of considering risk-hedging would be highly questionable: one would at best be facing a tradeoff decision between risk and path length. The numerical results in Section V (and other numerical results, not included here) exhibit cases where the *approximate* (i.e., computed using the algorithm in Section IV) risk-hedged solution outperforms the operational practice in both risk and path length. This empirical evidence, however, does not prove conclusively that the *exact* risk-hedged solution can so outperform the operational practice. The rest of this section is such a proof.

A conceptual simplistic example, constructed solely for the purpose of the proof, is shown in Figure 3. It depicts a situation with $N_w = 2$ weather regions, with the respective risk coefficients α_1 and α_2 satisfying $\alpha_2 > \alpha_1$. The two regions (chosen rectangular for simplicity of exposition and computation) share a boundary segment and have dimensions l-by-$2m$ and a-by-$2m$. Furthermore, we choose the parameter m large enough that the risk-hedged solution (the one minimizing (4)) is the rectilinear segment from $\mathbf{x}^{\text{in}}$ to $\mathbf{x}^{\text{fi}}$ (shown in blue in Figure 3), as opposed to being a shortest path that goes around one or both weather regions.

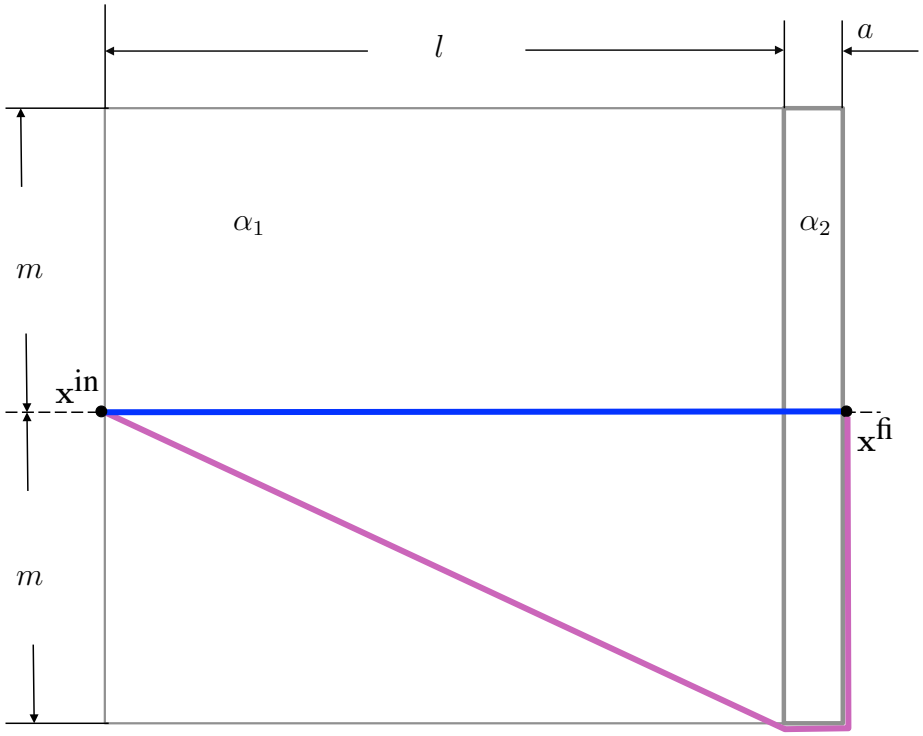

Figure 3: Proof that cases exist with risk-hedging having clear advantage.

The risk-hedged solution has path length

$$L_{\text{risk-hedged}} = l + a$$

and carries risk

$$R_{\text{risk-hedged}} = \alpha_1 l + \alpha_2 a.$$

Neither of the last two expressions involves m. A path corresponding to the current operational practice (shown in magenta in the Figure) has path length

$$L_{\text{operational practice}} = \sqrt{l^2 + m^2} + a + m$$

and carries risk (defined as the integral of $\sum_{w=1}^{N_w} \alpha_w I_w(\mathbf{x})$ along the path)

$$R_{\text{operational practice}} = \alpha_1 \sqrt{l^2 + m^2}.$$

A direct inspection shows that m can be chosen sufficiently large to have

$$L_{\text{risk-hedged}} < L_{\text{operational practice}} \quad \text{and} \quad R_{\text{risk-hedged}} < R_{\text{operational practice}}.$$

IV. An Algorithm to Solve the Risk-Hedged Re-routing Problem

A. The Computational Approach

The problem of Section II.C is a special case of the *Eikonal equation* [24],

$$\frac{1}{P(\mathbf{x})} \left| \text{grad} \left(\text{minimal cost of getting from the current state } \mathbf{x} \text{ to } \mathbf{x}^{\text{fi}} \right) \right| = 1,$$

which describes a number of physical propagation phenomena; for example, light propagation (see Section III). The problem belongs to a broader class of optimal control problems, which have been treated by many numerical algorithms, including those in Refs. [25–29]. The problem of Section II.C is a special case of a *minimal-time control problem* in that the adverse weather regions B_w are finitely many simple polygons. The algorithm given in Section IV.C, below, for solving it is, essentially, an approximation which is a special case of *Dynamic Programming* [30]. A more complicated structure of $P(\mathbf{x})$ can introduce additional computational difficulties into the use of Dynamic Programming; e.g., make the choice of a suitable mesh discretization more problematic.

The numerical approach here was chosen for its simplicity (aided further by the open-source availability of suitable software libraries), and not necessarily for best computational speed. Computational efficiency plays only a supporting role in this work, which is concerned mainly with a research investigation of the traffic re-routing problem under weather uncertainty. Nevertheless, as the approach will make direct use of Dijkstra's shortest-path algorithm [31, 32], shortest paths will be computed with the efficiency of Dijkstra's algorithm.

B. Assumptions and Their Implications for the Algorithm

A subset C of the airspace such that P is constant on C and non-constant on any bigger set that contains C, will be called a *constancy region* of P. Thus, P is constant on a constancy region, but not on any bigger set that contains a constancy region. The constancy region on which P is equal to 1 is the complement of the union of all the B_w's. The concept of a constancy region is illustrated notionally in Figure 4; P has five constancy regions, whose values are shown (to two decimal places).

If follows from the above that two constancy regions either have no common points or share one or more boundary points[a] (but cannot have overlaps). If a point $\mathbf{x}$ lies on the boundary of two or more constancy regions, then $P(\mathbf{x})$ is taken to be the smallest of all the constant values corresponding to those regions (because the aircraft can avoid the risk at $\mathbf{x}$ by going slightly into the safest adjacent constancy region).

Henceforth, *assume every B_w is a polygon*[b], generally non-convex, and with no two distinct sides having a common point except possibly a common endpoint. It follows that every constancy region of P is a polygonal region (a region with a piecewise linear boundary). A constancy region that is bounded will be referred to as a *weather-affected constancy polygon (of P)*. Weather-affected constancy polygons are illustrated in Figure 5, which shows $N_w = 3$ weather instantiations and are seven weather-affected constancy polygons, each shown in a shade of gray.

The following list sums up the aforementioned properties of the problem of Section II.C and adds some new ones:

- The risk field $P(\mathbf{x})$ is piecewise constant.
- Since the regions B_w are polygons and not, for instance, more complicated closed contours with many inflections, the risk field $P(\mathbf{x})$ has relatively (to N_w) few weather-affected constancy polygons.
- The weather-affected constancy polygons of $P(\mathbf{x})$ are, indeed, polygons (because the B_w's are) and are generally nonconvex.
- The polygonal shape of a weather-affected constancy polygon allows it to be triangulated.
- The problem is isotropic (i.e., the cost $P(\mathbf{x})$ depends on the current state $\mathbf{x}$ only, and not on the instantaneous direction of the path through $\mathbf{x}$), hence allows treatment by a Dijkstra-like algorithm.

[a] A boundary point of a set need not belong to the set.

[b] A bounded region with a "sufficiently nice" boundary can be, in a suitable sense, approximated by a polygon. Such approximations, however, require an analysis of numerical convergence and stability.

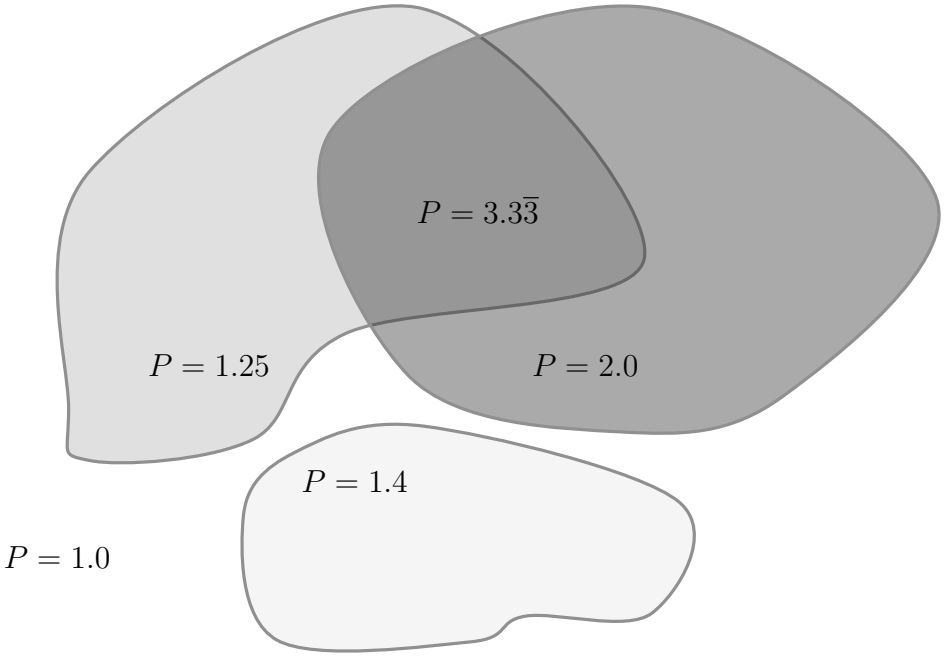

Figure 4: An illustration of the constancy regions of P for the example shown in Figure 2.

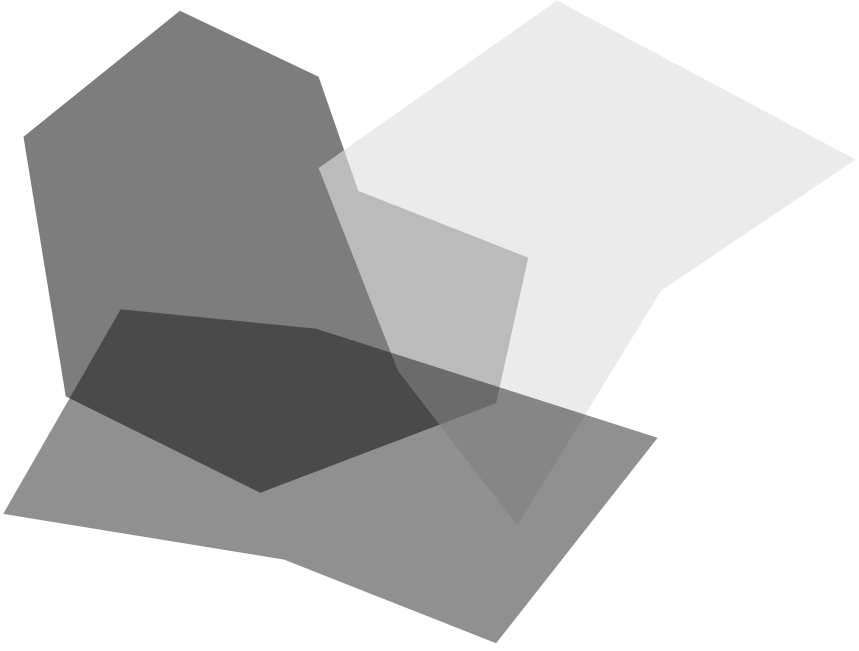

Figure 5: An example with $N_w = 3$ possible weather instantiations.

C. Algorithm Overview and Flowchart

The algorithm proceeds as follows. Given an instance of the problem of Section II.C, it first computes the weather-affected constancy polygons of the risk field P. It then computes the set of points where each point is either $\mathbf{x}^{\text{in}}$, or $\mathbf{x}^{\text{fi}}$, or a vertex of a weather-affected constancy polygon. This set is triangulated so as to retain all the boundary segments of all the weather-affected constancy polygons (a *conforming triangulation*). An example of such a triangulation is shown in Figure 6; the triangulation segments (shown dashed) corresponding to the polygon boundaries are kept slightly out from the polygon, to avoid computing routes that skirt an adverse weather region.

Depending on the geometry of the adverse weather regions B_w, the triangulation carried out in the algorithm (see the flowchart in Figure 7) may yield a mesh too crude to give a good approximation to the solution of the problem of Section II.C. This issue can be addressed in a number of ways, including:

- modifying the algorithm to include further refinements of the triangulation, which can be repeated until a suitable tolerance condition is met, and
- post-processing the approximately computed risk-hedged path to increase optimality and smoothness.

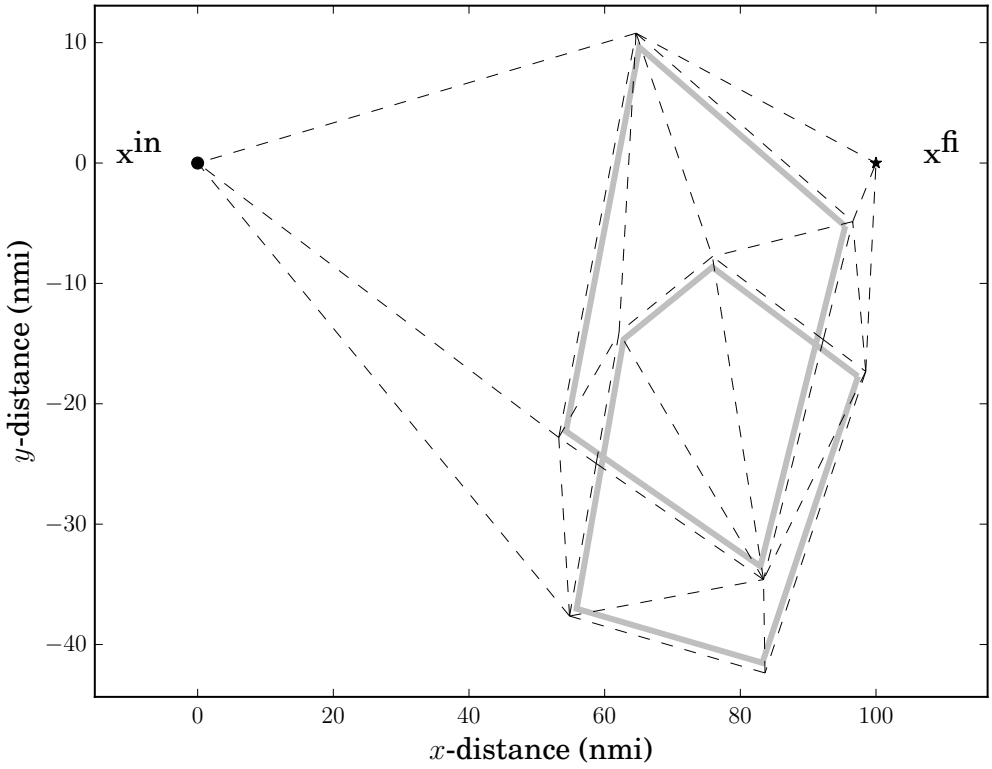

Figure 6: Two polygons B_w (solid gray) and the corresponding triangulation (dashed).

The calculated triangulation has two properties, essential to the algorithm: *(a)* the nodes include the points $\mathbf{x}^{\text{in}}$, $\mathbf{x}^{\text{fi}}$, and all vertices of all the weather-affected constancy polygons of P, and *(b)* the value of P remains constant along each link (side of a triangle) of the triangulation. By treating the triangulation as an undirected graph, in which node $\mathbf{x}^{\text{fi}}$ is reachable from node $\mathbf{x}^{\text{in}}$ (because of property *(a)*), and by assigning to each link the cost equal to the product

$$(\text{length of link})\,(\text{value of } P \text{ on link}), \tag{5}$$

which by property *(b)* is the exact value of integral (4) along the link, one obtains a *weighted graph* [32]. The problem of finding, in this graph, a path from $\mathbf{x}^{\text{in}}$ to $\mathbf{x}^{\text{fi}}$ of minimal total cost (weight) is a *shortest-path problem* [32] whose solutions approximate those of the problem of Section II.C. These solutions are approximate because they are based on a triangulated mesh which discretizes the infinitely-many waypoints available for re-routing and thereby restricts the class of possible paths. This shortest-path problem can be solved by a suitable algorithm.

The flowchart of the algorithm is shown in Figure 7. The two blocks shown in gray require a number of iterations, unknown at the outset, that depends on the layout of polygons B_w. Throughout the algorithm, maintain the initially

empty set of points called "`Nodes`", and the initially empty set called "`Links`" of unordered pairs of points from `Nodes`. In addition to these sets, initialize another empty one and call it "`Weather_Affected_Constancy_Vertices`."

D. Computational Costs of the Algorithm

Among the first segments inserted into `Links` are the sides of the weather-affected constancy polygons of P. Consequently, these segments (possibly broken up into subsegments) remain in `Links` throughout the algorithm and remain potential parts of an optimal routing.

The computational cost of the algorithm depends on the relative positions of the polygons B_w (the more intersections, the more iterations between the two blocks shown in gray in Figure 7) and on the amount of computation spent on the triangulation (which can vary depending on the type of triangulation sought). The shortest-path computation has well-known computational costs; e.g., $O(|\texttt{Nodes}|^2)$ if Dijkstra's algorithm is used.

V. Numerical Examples

A. Behavior of Risk-hedged Solutions with Risk Coefficients Varied

The two panels of Figure 8 show the same set of three weather regions B_w, with the risk coefficients varying from panel to panel, and the corresponding risk-hedged solutions (paths shown in blue) computed using the above algorithm. The border thickness shown for a region is proportional to the region's risk coefficient, whose value is shown in the interior of the region. The triangulation was refined locally to have each link of length at most 5 nmi; namely, each link longer than 5 nmi and with at least one endpoint on the last computed risk-hedged path was halved, and the entire network was re-triangulated, keeping the previously present edges. The discretization errors that result for this degree of coarseness (see the second paragraph of Section IV.C) are manifest in both panels: the computed paths have cusps that add length without reducing risk.

In Figure 8, the region with the highest risk is the polygon in the middle. A southerly deviation around this polygon has less path length than a northerly deviation does. In panel (A), the risk-hedged solution avoids the 65% polygon while penetrating a portion of the 5% polygon, because the southerly deviation is shorter and penetrates only a very-low-risk region. In panel (B), the risk-hedged solution avoids the 60% polygon, while penetrating a portion of the 10% polygon. Although a southerly deviation would be shorter, it would penetrate a substantial portion of the 30% polygon, resulting in a very high risk-adjusted path length.

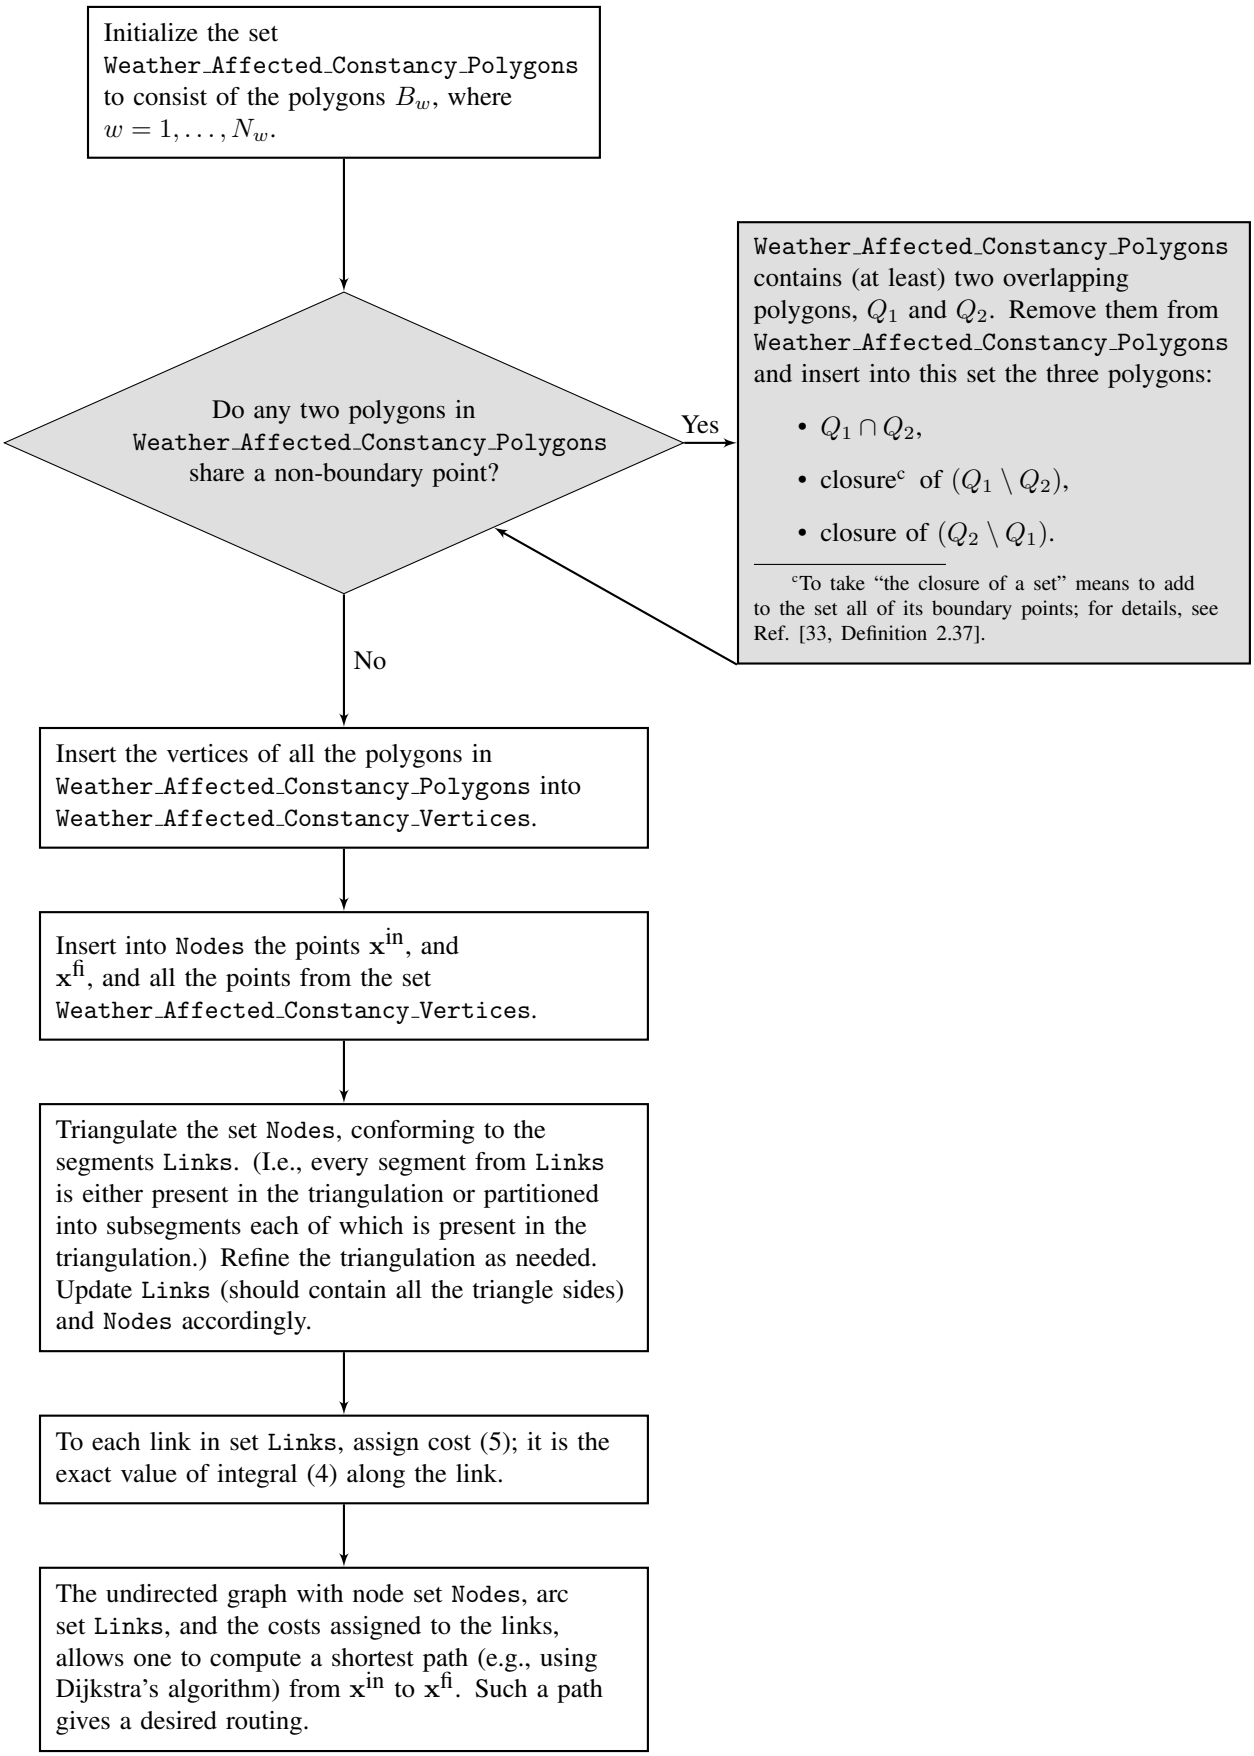

Figure 7: The flowchart for the algorithm. Gray blocks: iteration required.

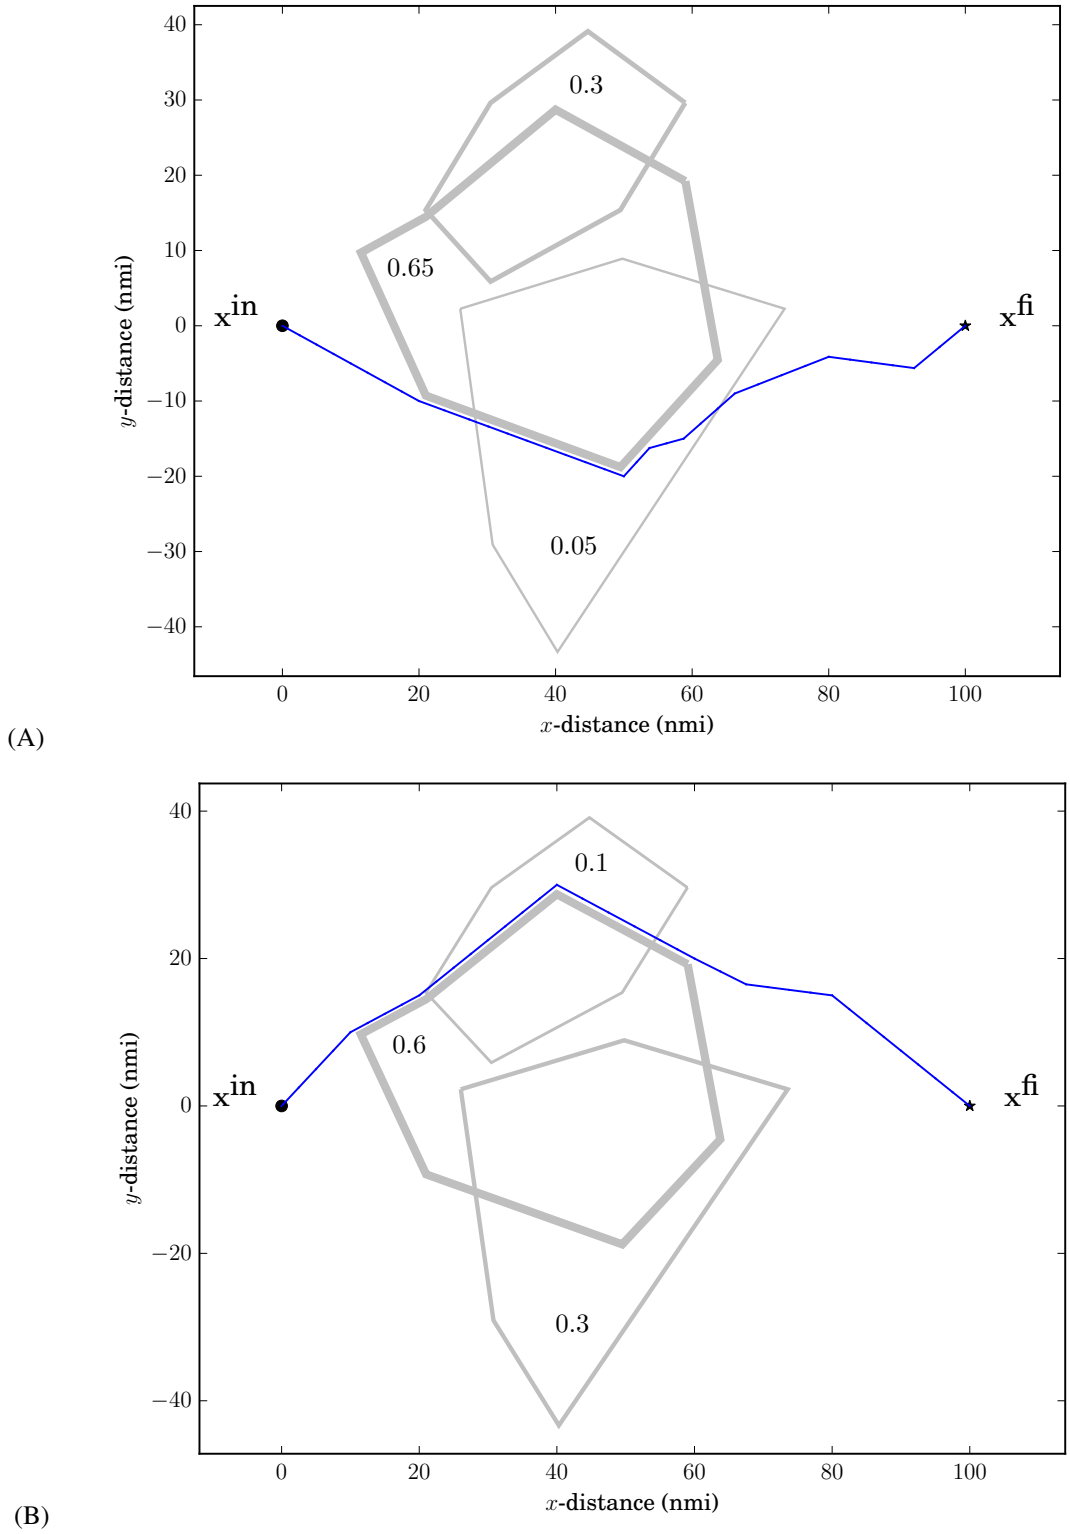

Figure 8: How a risk-hedged solution responds to changes in risk coefficients.

B. Cost-benefit analysis of risk-hedged solutions

The numerical results of this subsection were generated as follows. For an instance of the problem of Section II.C, three different optimizations were carried out to compute the following three paths: a shortest path clear of all the weather

Table 1: The three paths examined numerically in Section V.

path definition	path term	label in the Figure legend
a shortest path that avoids the interior of *all* the regions B_w	a *deterministically safe* path	`deterministically safe`
a path that minimizes the risk-adjusted re-routing cost (4)	the *risk-hedged* path	`risk-hedged`
a shortest path that avoids the interior of the region B_w with the highest risk coefficient α_w (assumed unique)	proxy for current operational practice	`operational practice`

regions B_w (the safest, but likely the most costly in path length), a path that minimizes risk-adjusted re-routing cost (4), and a shortest path that is clear of that weather region B_w with the highest risk coefficient α_w (the risk coefficients were chosen to be all distinct). These three types of paths are summarized in Table 1. The risk-hedged solutions are computed approximately using the algorithm given in Section IV. The triangulation was refined locally (near the solution path) to have each link of length at most 5 nmi, as described in Subsection A. The effects of the discretization error (see the second paragraph of Section IV.C) resulting from this degree of coarseness are discussed below. The deterministically safe and current-practice solutions are computed exactly.

Figures 9 through 11 show three instances of the problem of Section II.C. These instances differ in the set of the possible weather instantiations and in the values of the risk coefficients. Each of these Figures shows:

- in the top panel, the regions B_w (the border thickness shown for a region is proportional to the region's risk coefficient), their corresponding risk coefficients, and the three paths defined in Table 1;
- in the bottom panel, the three data points corresponding to the three paths defined in Table 1, each data point of the format:

$$\left.\begin{array}{ll} \textit{Ordinate:} & \text{the risk incurred by the path; i.e., the integral of } \sum_{w=1}^{N_w} \alpha_w I_w(\mathbf{x}) \text{ along the path,} \\ \textit{Abscissa:} & \text{savings in path length (defined below), relative to the deterministically safe path.} \end{array}\right\} \quad (6)$$

The length savings provided by a path relative to the deterministically safe path is defined as

$$\left(\frac{\text{the length of the deterministically safe path}}{\text{the length of the path in question}}\right) - 1.$$

The insights one gains from these Figures are as follows.

In Figure 9, the computed risk-hedged re-route is both shorter and less risky than the current operational practice (in the bottom panel, marker ▲ is below and to the right of marker ●), hence has an unambiguous advantage. The discretization error does not seem to have affected the optimality of the computed risk-hedged path.

In Figure 10, the computed risk-hedged route is less risky, but has a higher path length, than the current operational practice. A choice between the two paths involves a tradeoff. The discretization error that results from the coarseness of the mesh is manifest in the computed path (show in blue, in the top panel) in the form of clearly unnecessary cusps, which add path length without reducing risk. By replacing appropriate portions in the shown approximation of the risk-hedged path with rectilinear segments (e.g., the portion from the Eastern boundary of the 0.3 polygon to $\mathbf{x}^{\text{fi}}$), it is possible to obtain a new approximation with no greater risk and with shorter path length. Such an approximation would be an improvement on that shown. Consequently, for the exact risk-hedged solution, the bottom panel of Figure 10 would show marker ▲ farther to the right of the shown position, indicating an increase in the path-length savings.

In Figure 11, the computed risk-hedged route is more risky, but requires a shorter route, than the current operational practice. A choice between the two paths involves a tradeoff. The discretization error that results from the coarseness of the mesh is manifest in the computed path (show in blue, in the top panel) in the form of clearly unnecessary cusps, which add path length without reducing risk. By replacing appropriate portions in the shown approximation of the

risk-hedged path with rectilinear segments (e.g., the portion from the Eastern boundary of the 0.1 polygon to $\mathbf{x}^{\mathrm{fi}}$), it is possible to obtain a new approximation with no greater risk and with shorter path length. Such an approximation would be an improvement on that shown. Consequently, for the exact risk-hedged solution, the bottom panel of Figure 10 would show marker ▲ farther to the right of the shown position.

These results and insights can be summarized as follows:

- There exist instances of the problem of Section II.C in which (in the terminology of Table 1) the risk-hedged solution outperforms the current operational practice in both the path length (which approximates the flight cost) and in risk. (One such instance is shown in Figure 9.) In such instances, a risk-hedged path is clearly advantageous to the current operational practice.
- In all other instances (for a representative sample, Figures 10 and 11; compare the relative positions of the circle and the triangle in the bottom panel), one faces a tradeoff between risk and operational cost, for which the authors have no unambiguous interpretation.
- The preceding two bullets together suggest the following rule of thumb: choose the hedged path if it outperforms the current operational practice in both risk and operational cost, and choose the latter path otherwise.

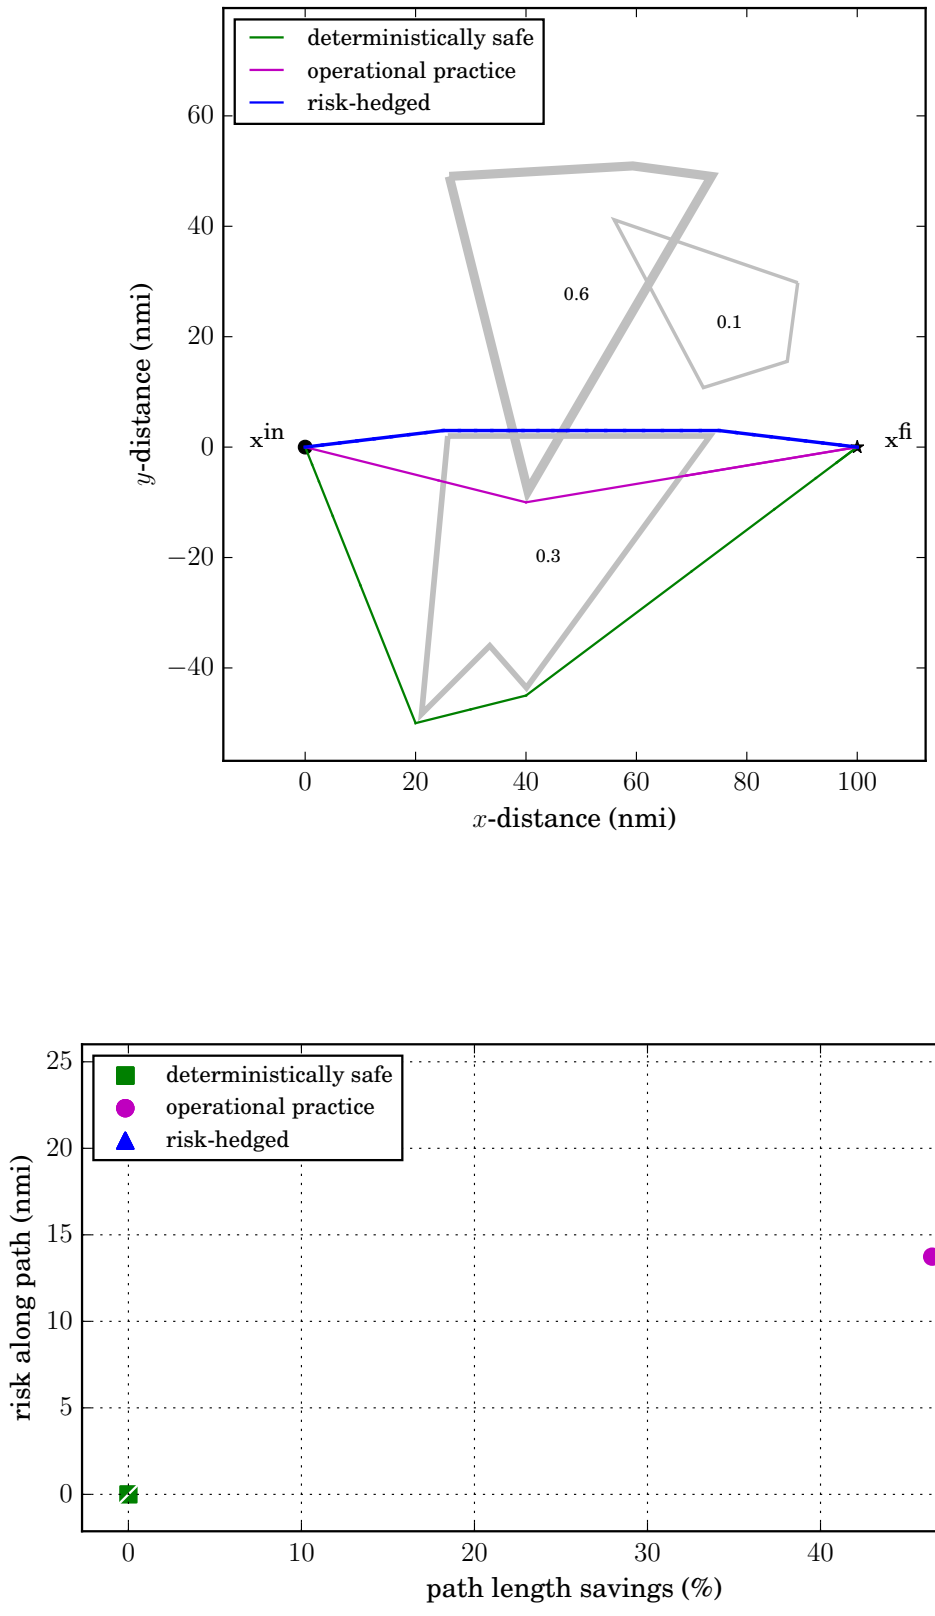

Figure 9: Top panel: a problem instance with the three paths described in Table 1. Bottom panel: a scatter plot of metrics for the three paths.

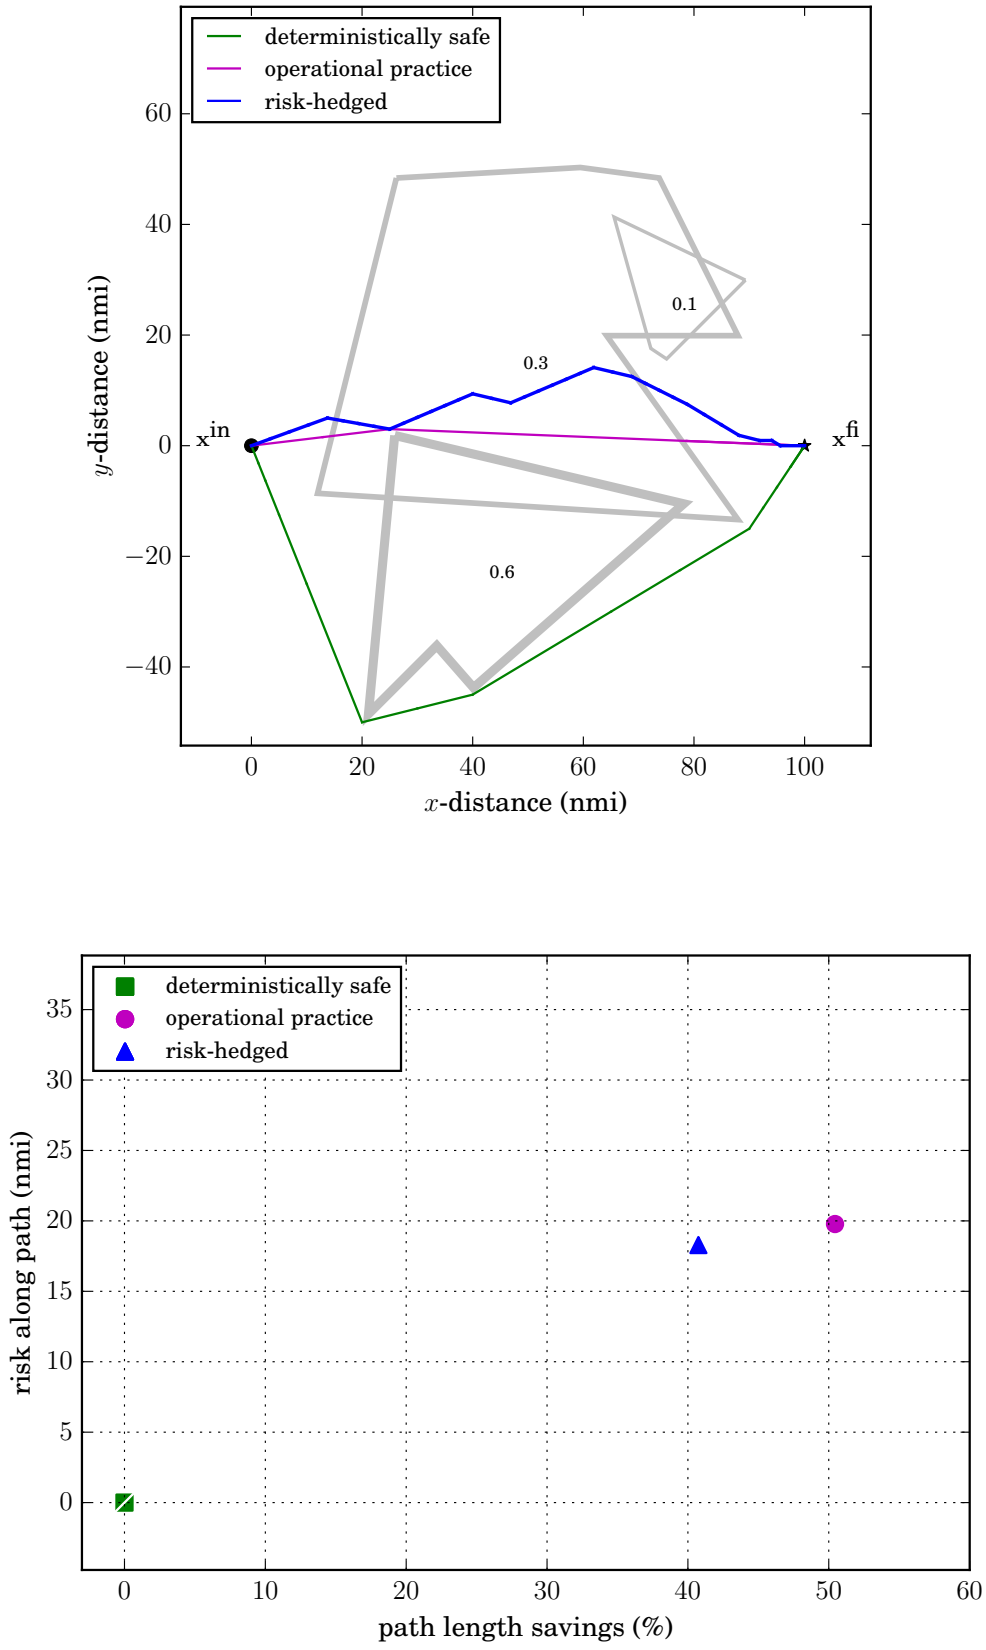

Figure 10: Top panel: a problem instance with the three paths described in Table 1. Bottom panel: a scatter plot of metrics for the three paths.

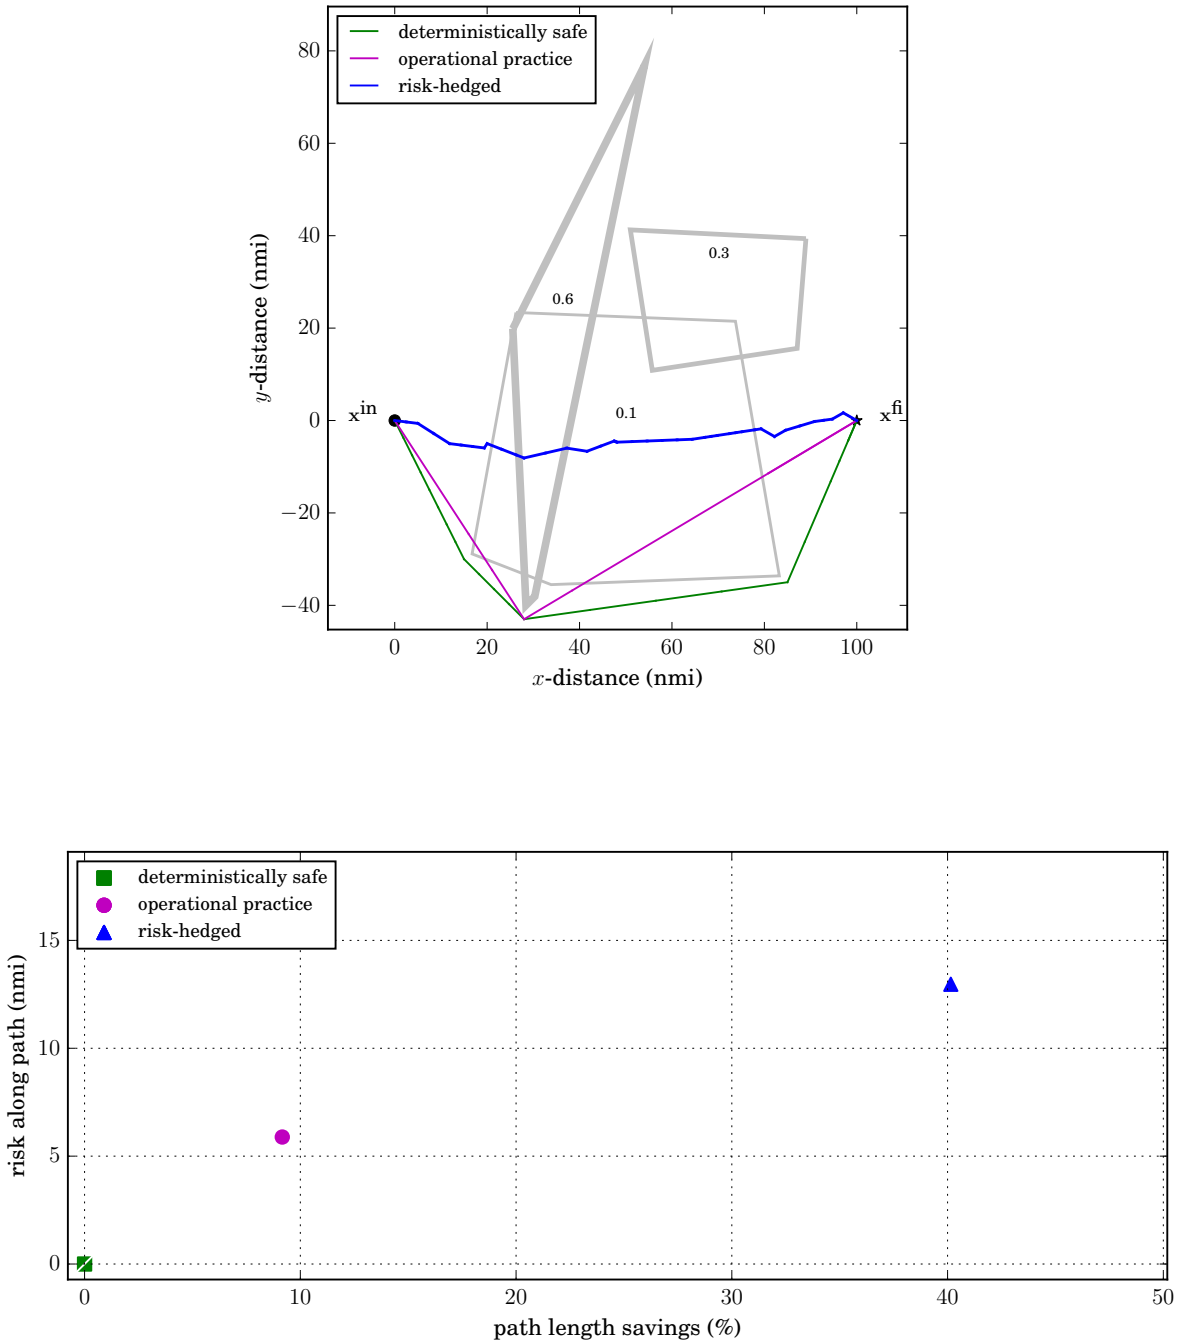

Figure 11: Top panel: a problem instance with the three paths described in Table 1. Bottom panel: a scatter plot of metrics for the three paths.

VI. Conclusion

A new risk-hedged approach for re-routing aircraft around convective weather has been presented. This approach seeks to minimize the risk-adjusted path length, given multiple possible weather instantiations with associated likelihoods/risks of occurrence. The resulting re-routing path generally avoids most of the highest-risk weather instantiation while possibly penetrating small regions of lower-risk weather instantiations. This is in contrast to current operational practice of designing re-routing paths to completely avoid the weather instantiation posing the substantially highest risk without regard for lower-risk weather instantiations. Both approaches provide a nominal strategic solution, with tactical adjustments made as necessary if the weather manifests itself differently than anticipated. However, the

hedged solution also anticipates and (partially) accounts for lower-risk weather instantiations.

Risk-hedged re-routes are analyzed for some example weather instantiations, using metrics of path length and path risk. The main result is that in some scenarios, relative to an operational-practice proxy solution, the risk-hedged solution provides the benefits of lower risk as well as shorter path length. In other scenarios, the benefits of the risk-hedged solution are ambiguous, because the solution is characterized by a tradeoff between risk and path length. The risk-hedged solution can be executed in those scenarios where it provides a clear benefit over current operational practice.

References

[1] Traffic Flow Management in the National Airspace System. *U.S.A. Federal Aviation Administration*. Air Traffic Organization, 2009.

[2] Air Traffic Control Order JO 7110.65U. *U.S.A. Federal Aviation Administration*. U.S. Dept. of Transportation, Washington, D.C., 2012.

[3] B. Sridhar, G. Chatterji, S. Grabbe, and K. Sheth. Integration of traffic flow management decisions. In *AIAA Guidance, Navigation, and Control Conference*. Monterey, California, 2002.

[4] D. R. Isaacson, A. V. Sadovsky, and D. Davis. Scheduling for Precision Air Traffic Operations: Problem Definition and Review of Prior Research. *Journal of Aerospace Information Systems*, 11:234–257, 2014.

[5] J. Prete and J.S.B. Mitchell. Safe routing of multiple aircraft flows in the presence of time-varying weather data. In *AIAA Guidance, Navigation, and Control Conf*, 2004.

[6] A. Mukherjee, S. Grabbe, and B. Sridhar. Arrival flight scheduling through departure delays and reroutes. *Air Traffic Control Quarterly*, 17(3), 2009.

[7] S. Grabbe, B. Sridhar, and A. Mukherjee. New York flow control with deterministic en route capacity constraints. *Air Traffic Control Quarterly*, 17(2), 2009.

[8] Z. Zhang, Z. Hao, and Z. Gao. A dynamic adjustment and distribution method of air traffic flow en-route. *Journal of Air Transport Management*, 42(C):15–20, 2015.

[9] A. Mukherjee, S. Grabbe, and B. Sridhar. Alleviating airspace restrictions through strategic control. In *AIAA Guidance, Navigation, and Control Conference, Honolulu, HI*, 2008.

[10] M. Bloem and B. Sridhar. Optimally and equitably distributing delays with the aggregate flow model. In *Digital Avionics Systems Conference*. IEEE, 2008.

[11] S. Grabbe, B. Sridhar, and A. Mukherjee. Sequential traffic flow optimization with tactical flight control heuristics. *Journal of Guidance, Control, and Dynamics*, 32(3):810–820, 2009.

[12] H. Arneson and M. Bloem. A method for scheduling air traffic with uncertain en route capacity constraints. In *AIAA Guidance, Navigation, and Control Conference, Chicago IL*, 2009.

[13] N. Y. Chen and B. Sridhar. Weather-weighted periodic auto regressive models for sector demand prediction. In *AIAA Guidance, Navigation, and Control Conference, Chicago, IL*, 2009.

[14] H. K. Ng, A. Morando, and S. Grabbe. Automated flight routing using stochastic dynamic programming. In *10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, page 9113, 2010.

[15] M. C. Drew. Dynamically evolving sectors for convective weather impact. In *10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, page 9161, 2010.

[16] H. Arneson. Initial analysis of and predictive model development for weather reroute advisory use. In *15th AIAA Aviation Technology, Integration, and Operations Conference*, page 3395, 2015.

[17] Air Traffic Management: Traffic Management Advisor. The NASA Ames Research Center website. `http://www.aviationsystemsdivision.arc.nasa.gov/research/foundations/tma.shtml`.

[18] C. Gong, D. McNally, and C. Lee. Dynamic arrival routes: A trajectory-based weather avoidance system for merging arrivals and metering. In *15th AIAA Aviation Technology, Integration, and Operations Conference*, page 3394, 2015.

[19] K. Sheth, B. Sridhar, and D. Mulfinger. Application of probabilistic convective weather forecasts for flight routing decisions. In *7th AIAA ATIO Conference, Belfast, UK*, 2007.

[20] C. Wanke and D. Greenbaum. Incremental, probabilistic decision making for en route traffic management. *Air Traffic Control Quarterly*, 15(4), 2007.

[21] M. Steiner, C. K. Mueller, G. Davidson, and J. A. Krozel. Integration of probabilistic weather information with air traffic management decision support tools: A conceptual vision for the future. In *American Meteorological Society Conference on Aviation, Range, and Aerospace Meteorology*, 2008.

[22] M. Steiner, R. Bateman, D. Megenhardt, Y. Lui, M. Xu, M. Pocernich, and J. Krozel. Translation of ensemble-based weather forecasts into probabilistic air traffic capacity impact. In *Digital Avionics Systems Conference, 2009. DASC'09. IEEE/AIAA 28th*. IEEE, 2009.

[23] R.P. Feynman, R.B. Leighton, and M.L. Sands. *The Feynman Lectures on Physics*. Number v. 1 in The Feynman Lectures on Physics. Addison-Wesley, 1963. URL: https://books.google.com/books?id=UtJEAAAAIAAJ.

[24] V. Arnold and R. Cooke. *Lectures on Partial Differential Equations*. Universitext. Springer Berlin Heidelberg, 2003. URL: https://books.google.com/books?id=qlNJAYwmfTcC.

[25] J. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40:1528–1538, 1995.

[26] L. C. Polymenakos, D. P. Bertsekas, and J. N. Tsitsiklis. Implementation of efficient algorithms for globally optimal trajectories. *Automatic Control, IEEE Transactions on*, 43(2):278–283, 1998.

[27] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

[28] R. Kimmel and J. A. Sethian. Fast marching methods for computing distance maps and shortest paths. *CPAM Report*, 669, 1996.

[29] J. Sethian and A. Vladimirsky. Ordered upwinds methods for hybrid control. *Proceedings Fifth International Conference on Hybrid Systems and Control, LCNS 2289*, 2002.

[30] R. Bellman. *Dynamic Programming*. Dover, 2003.

[31] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. doi:10.1007/BF01386390.

[32] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization; Algorithms and Complexity*. Dover Publications, 1998.

[33] J. A. Bærentzen, J. Gravesen, F. Anton, and H. Aanæs. *Guide to Computational Geometry Processing: Foundations, Algorithms, and Methods*. SpringerLink : Bücher. Springer London, 2012. URL: https://books.google.com/books?id=0lb4_pLIyP8C.