

Fast Graph-Search Algorithms for General-Aviation Flight Trajectory Generation

Eran Rippel,* Aharon Bar-Gill,† and Nahum Shimkin‡
Technion—Israel Institute of Technology, 32000 Haifa, Israel

We consider numerical algorithms for onboard flight trajectory generation and optimization in three-dimensional space. Our approach relies on graph-search algorithms, which perform a global search over the set of feasible trajectories. We start by formulating a simplified kinematic model that is appropriate for general-aviation aircraft. The cost function to be optimized includes accounts for position-dependent criteria such as the flight altitude and terrain clearance. Additional dynamic components that involve the angular velocities are introduced to account for riding qualities and pilot workload. Using an approximate grid-based discretization scheme, we transform the continuous optimization problem into a search problem over a finite graph and apply Dijkstra's shortest-path algorithm to this problem. To reduce the computation time to acceptable levels, we introduce a novel state reduction technique that leads to suboptimal search. Further speedup is achieved by heuristic search techniques and hierarchical methods. Performance of the proposed algorithms is evaluated for a trajectory optimization problem with terrain following over a 100×100 km area. Our experiments demonstrate the potential of these algorithms, when combined, to provide an onboard solution to realistic flight-trajectory-generation problems.

Nomenclature

$a_{\text{horizontal}}, \bar{a}_h$	= aircraft horizontal acceleration and its upper bound
$a_{\text{vertical}}, \bar{a}_v$	= aircraft vertical acceleration and its upper bound
$C(v, u)$	= cost of edge (v, u)
$c(\mathbf{X}, \mathbf{U})$	= immediate cost function
D	= set of destination vertices
$\text{DTM}(t)$	= ground level from digital terrain map
d_0	= trajectory accuracy
\mathcal{E}	= set of edges of a graph
G	= graph
g	= cost-so-far estimation function
H	= flight level
H_0	= clearance height
h	= cost-to-go heuristic function
h_x, h_y, h_z	= three-dimensional grid spacing
h_0	= safety clearance height
i, j, k	= integer grid indices
J	= cost functional
m	= maximal number of neighbors for each three-dimensional point
m, n, ℓ	= integer indices for neighboring grid points
N	= number of vertices in a graph
n	= number of three-dimensional points in the mesh
\mathbf{P}, \mathbf{P}_i	= point in three-dimensional Cartesian space
R_{\min}	= minimal aircraft turn radius
S	= set of vertices with known shortest-path values
S_0	= finite grid in Cartesian space
s	= source vertex
t_f	= first arrival time to a goal state
\mathbf{U}	= control signal
u, v	= vertices
\mathcal{V}	= set of vertices of a graph

$V(\mathbf{X})$	= aircraft velocity at state \mathbf{X}
w_κ	= cost criterion weight
\mathbf{X}	= aircraft state
x, y, z	= Cartesian coordinates
α	= tradeoff factor
γ	= vertical flight-path angle
$\Delta\gamma_0$	= largest separation between discrete angles
κ	= cost criterion index
\mathcal{X}_f	= set of goal states
ψ	= heading angle
Ω	= set of cost criteria

I. Introduction

GENERAL aviation (GA) constitutes a considerable part of today's saturated air traffic. In most instances, a GA aircraft is flown by a single pilot, who might find himself under adverse weather conditions, operating in single-pilot-instrument-flight-rule (SPIFR) mode; a faulty link to air traffic control (ATC) can complicate matters even more. The additional stress and workload cause many unfortunate accidents. Technological advancements, combined with increased affordability, open new venues for tackling this SPIFR safety problem. These include, for example, ground collision avoidance systems^{1,2} that employ geographic-information-system databases.^{3,4}

Three-dimensional flight trajectory generation is a complex task, involving various criteria that affect the resulting path. These include, among others, the flight duration, the degree to which different terrain features are utilized, the workload imposed on the pilot in order to follow the prescribed trajectory, and riding qualities (or passenger comfort), which are affected by the accelerations involved. Every generated trajectory has to accommodate for these often conflicting objectives. Evidently, the specified cost components and the tradeoff between them should reflect the specific requirements of the flight mission. For example, low altitude might be required to avoid flying within a cloud layer, which can add to pilot stress. "Sunday pilots," in particular, would prefer maximum dwell time below cloud ceiling, that is, flying as low as possible, subject—to course—to safety constraints. The weight given to the pilot workload (pumping the stick and the pedals) should depend on the pilot's preferences. And minimal exposure to high g levels can be critical if an injured passenger is carried onboard. All of these considerations bring out the need for a flexible flight-trajectory-generation capability.

Received 1 January 2004; revision received 24 May 2004; accepted for publication 24 May 2004. Copyright © 2004 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/05 \$10.00 in correspondence with the CCC.

*Graduate Student, Department of Electrical Engineering.

†Adjunct Faculty Member, Departments of Electrical Engineering and Aeronautical Engineering; also Senior Aerospace R&D Associate, H-E Systems Engineering, Inc., 69358 Tel-Aviv, Israel. Senior Member AIAA.

‡Associate Professor, Department of Electrical Engineering.

In this paper we address the problem of onboard generation and optimization of three-dimensional trajectories. GA applications call for a low-cost solution to this problem that would trace a flight path for rescuing a person, for example, seconds after destination has been decided on (hospital most suitable to treat the particular injury). Thus, by onboard algorithms we refer here to algorithms that can compute optimal flight paths within a few tens of seconds. Modern technology should allow for such a solution, in particular, under circumstances where communication with ATC fails and routine visual/instrumental flight rules cannot be maintained.

An implementation of such a trajectory generator includes an onboard computer, in which the digital terrain map (DTM) resides and an altimeter [e.g., a global positioning system receiver], on top of the standard navigational instrumentation normally available in GA airplanes. Laptop computers with DTMs are indeed increasingly being used by light aircraft and helicopters who fly rescue and other emergency missions without adhering to prescribed routes. Also, the provisioning of cues that enable the pilot to comfortably follow the planned trajectory is readily implementable even in today's low-cost aircraft. As stated, in the sequel we focus on the trajectory generation task.

A. Overview of the Proposed Approach

A control-oriented formulation of the trajectory optimization problem calls for finding a control signal $U(t)$ to minimize the cost functional:

$$J = \int_0^{t_f} c(\mathbf{X}(t), \mathbf{U}(t)) dt \quad (1)$$

subject to the nonlinear differential state equation

$$\dot{\mathbf{X}}(t) = f(\mathbf{X}(t), \mathbf{U}(t)), \quad 0 \leq t \leq t_f \quad (2)$$

In addition, various constraints can apply to the state and control variables. Theoretical solution approaches to this problem include dynamic programming, which iteratively computes the optimal value function from any state, and the calculus of variations (with the associated Euler–Lagrange equations), which traces the optimal path.^{5–7} Solution of practical problems requires numerical methods.^{5,8,9} State-of-the-art numerical methods related to the calculus of variations are essentially based on nonlinear optimization algorithms, such as gradient descent and Newton's method. These methods are basically of local nature, so that convergence to the global minimum cannot be guaranteed unless the problem is highly structured. This is not the case in the GA flight-generation problem, where rugged terrain and forbidden flight zones, for example, can create multiple locally optimal trajectories. Dynamic programming, on the other hand, is essentially a global method. However, because of its exhaustive nature, it is considered computationally feasible only for problems with low state dimension and coarse grid discretization (compare, Ref. 5, Chapter 7, and Ref. 8).

The present work aims to develop three-dimensional trajectory optimization methods, which are of global nature, yet computationally feasible, using the basic tool of graph search. Graph-search methods are a standard tool for shortest-path problems. However, existing work does not address problems with dynamic aspects, namely, when accelerations are taken into account as part of the problem constraints and cost. We note that graph-search methods are closely related to dynamic programming, but employ the special structure of the shortest-path problem to enhance computational efficiency.

To begin with, we set up a simplified model that captures the essential ingredients of the trajectory optimization problem for general-aviation aircraft. We apply to this model an approximate discretization scheme, which allows the application of efficient global search algorithms to find the optimal path. In particular, we discretize the model over a state space, which includes both spatial position and velocity, and formulate the trajectory optimization problem as a shortest-path search problem over a finite graph. This form is appropriate for the application of Dijkstra's shortest-path algorithms, which is an efficient single-pass algorithm. In essence, Dijkstra's

algorithm proceeds from the target (or initial) state and iteratively develops the front of equidistant states (or graph vertices). Thus, at each step one new node is tagged with its optimal value, until a target state is reached. As will be seen, the sequential structure of this scheme will allow us to introduce some important modifications to the algorithm.

The application of the basic Dijkstra algorithm to the discretized model is still short of providing onboard performance for problems of realistic size. We therefore explore algorithms that accelerate the process of trajectory generation. First, we reduce the dimensionality of the discrete state space by removing the velocity-related components. Nonetheless, the proposed reduced-state algorithm retains the ability to account for the acceleration cost, albeit in a suboptimal manner. Further enhancements that are considered include the addition of an A^* heuristic and a hierarchical (coarse-fine) search scheme. The performance of these algorithms is examined and compared in simulation.

B. Related Research

The nonlinear nature of both the aircraft state dynamics and the performance index J make the trajectory generation problem hard to solve. We briefly review here some of the solution approaches that have been proposed for this problem. These include definition of various cost function components, different aircraft model formulations, and algorithmic solution variants.

Funk¹⁰ performs vertical-plane optimization only, by means of cubic splines. The splines are generated by solving a nonlinear programming problem. He addresses only the altitude aspect of the cost. Another trajectory optimization solution, utilizing nonlinear programming, is given by Hargraves and Paris,¹¹ again only for the vertical plane.

Asseo¹² tackles the three-dimensional problem by looking separately at the horizontal and vertical components. His algorithm employs the gradient method, but restricts the gradient search to the horizontal plane only, avoiding the convergence problems associated with vertical constraints. He achieves vertical tracks by fitting parabola segments that clear critical points. This approach involves no vertical optimization. Also, only the kinematics of the aircraft is modeled, not accounting for aspects of flight dynamics.

Menon et al.¹³ use the optimal control approach to obtain trajectories that minimize a linear combination of flight time and terrain-following (TF) costs. They address the three-dimensional problem by employing the adjoint method^{6,8} and a one-dimensional search. The transformation allows the iterative algorithm to generate three-dimensional trajectories, while optimizing for a single parameter. The algorithm is complex, and its computational cost is high. Here too, only aircraft kinematics is considered, not accounting for flight dynamics.

Lu and Pierson¹⁴ use the optimal control approach and address off-line calculation of optimal vertical TF trajectories. They use a point-mass nonlinear dynamic model of the aircraft for their analysis and utilize the Hamiltonian function, with the inverse-dynamics approach.¹⁵

Spong et al.¹⁶ deal with the problem of bounded inputs, employing optimal decision strategy (ODS), which is a class of pointwise optimal control strategies. Accounting for the bounded nature of nonlinear aircraft dynamics (such as thrust and aerodynamics), Lee et al.¹⁷ apply the ODS method to the TF flight problem. Rehbock et al.¹⁸ convert the ODS method into a linear-programming problem, in contrast to the previous work, where it is posed as a quadratic-programming problem. Barnard¹⁹ combines the ODS method with the neural-networks technique.

Hess and Jung^{20,21} use a generalized predictive control algorithm that calculates an optimal control sequence several steps ahead. They then apply only the first control of the sequence and recalculate a new sequence (a receding horizon process). Lu²² Lu and Pierson^{22,23} also consider predictive algorithms.

Numeric solutions of continuous dynamic-programming methods are extensively treated in Ref. 9. However, these methods can require multiple passes over the entire state space. For the specific case of the shortest-path problem [which corresponds to a special

case of the problem in Eqs. (1) and (2) with $c(\mathbf{X}, \mathbf{U}) = c(\mathbf{X})$, $\dot{\mathbf{X}} = \mathbf{U}$ and $\|\mathbf{U}\| \leq 1$, more efficient alternatives are offered by appropriate discretization schemes coupled with efficient graph search procedures. An early example for the use of this approach can be found in Ref. 24. More recently, numerically consistent algorithms (which converge to the exact solution when the spatial resolution is increased) have been proposed in Refs. 25 and 26. All of these methods use a Dijkstra-like algorithm to perform a global search over the discretized space. Unfortunately, these algorithms do not accommodate at present the dynamic aspects of our problem, namely, the acceleration-related costs and constraints.

The remainder of this paper is organized as follows. The underlying optimization problem, the aircraft model, and the cost function are formulated in Sec. II. The transformation of the problem into a graph search and the application of Dijkstra's algorithm are described in Sec. III. Suboptimal enhancements are discussed in Sec. IV. Section V presents simulation results that demonstrate the performance of the proposed algorithms. Finally, in Sec. VI we conclude and indicate some topics of interest for further research.

II. Problem Definition

The problem we consider is of generating an optimal aerial trajectory between two points: the source and the destination. The trajectory has to be flyable, and hence attention is paid to the physics of the problem, including the flight model, acceleration constraints, and cost components that reflect passenger and pilot comfort. We proceed to formulate the continuous model, specify some important maneuverability constraints, and finally specify the proposed cost functional.

A. Model Formulation

We start with a simplified, kinematic model of an aircraft flying in a three-dimensional airspace. The aircraft is considered as a point in three-dimensional space, and its translational and angular states at time t are defined by the configuration vector $\mathbf{X}(t) = [x(t), y(t), z(t), \gamma(t), \psi(t)]^T \in G$. This is illustrated in Fig. 1. The set G is defined as $\mathfrak{R}^3 \times [\gamma_{\min}, \gamma_{\max}] \times [-\pi, +\pi]$. The coordinates x , y , and z are taken relative to an Earth frame of reference, such as WGS-84 (World Geodesic System). The extremal values γ_{\min} and γ_{\max} are derived from the flight model, detailed in the next subsection.

The flight speed $V(t)$ will be taken as given and will not be a controlled variable in our model. This is the common situation in GA flight. However, we do allow the speed to be a (static) function of the configuration vector $\mathbf{X}(t)$, namely, $V(t) = g(\mathbf{X}(t))$. In particular, V might depend on the flight altitude, as well on the climb angle $\gamma(t)$. We note that during climbs, even with throttle compensation, the aircraft can lose roughly 20% of its velocity, and during descends can gain up to 20% increased velocity. Wind-speed data can also be used to modify V according to the heading angle $\psi(t)$. We shall not proceed further here with a precise modeling of these effects.

We assume instantaneous control of the angular rates, so that the control signal $\mathbf{U}(t)$ coincides with $(\dot{\gamma}, \dot{\psi})$. The state equations




Fig. 1 Airspace position and orientation.

$\dot{\mathbf{X}}(t) = f(\mathbf{X}(t), \mathbf{U}(t))$ can now be written as

$$\begin{aligned} \dot{x}(t) &= V(t) \cos \gamma(t) \cos \psi(t), & \dot{y}(t) &= V(t) \cos \gamma(t) \sin \psi(t) \\ \dot{z}(t) &= V(t) \sin \gamma(t), & \dot{\gamma}(t) &= U_1(t), & \dot{\psi}(t) &= U_2(t) \end{aligned} \quad (3)$$

The control and state variables will be further constrained by the flight model, as detailed in the next subsection. Given an initial state X_0 and the goal set \mathcal{X}_f , feasible trajectories are continuous solution paths $\mathbf{X}(t)$ that start at $\mathbf{X}(0) = X_0$ and terminate at $\mathbf{X}(t_f) \in \mathcal{X}_f$. We shall be interested in optimal trajectories that minimize a cost functional to be defined in the following.

B. Flight Model Constraints

Aircraft performance is subject to physical constraints that affect its maneuverability. Additional restrictions can be applied to ensure passenger and pilot comfort. Assuming that upper bounds are specified for the lateral accelerations, we will find it convenient to represent them by corresponding bounds on the angular velocities and the aircraft turn radius.

Under normal flight condition of GA aircraft, decoupling of the motion plane components is justified, resulting in two approximate acceleration measures:

$$a_{\text{vertical}} = V \dot{\gamma}, \quad a_{\text{horizontal}} = V \dot{\psi} \quad (4)$$

These accelerations approximate the two components of the lateral acceleration relative to the aircraft frame. Accordingly, they should not exceed certain levels for comfort. Assuming hard bounds \bar{a}_v and \bar{a}_h on the vertical and horizontal accelerations, respectively, we obtain the following constraints on the angular rates:

$$|\dot{\gamma}| < \bar{a}_v / V, \quad |\dot{\psi}| < \bar{a}_h / V \quad (5)$$

The minimal turn radius is next determined from $V = \dot{\psi} R$:

$$R_{\min} \geq V / |\dot{\psi}|_{\max} = V^2 / \bar{a}_h \quad (6)$$

This minimal radius serves as a constraint imposed on aircraft trajectories. Note that these bounds can depend on the aircraft configuration \mathbf{X} through the velocity V . It might be convenient to use some nominal value V_0 for V in these equations, although this is not necessary.

The maximal climb angle, based on climb rate capability safety requirements, gives rise the additional constraint

$$|\gamma| < \gamma_{\max} \quad (7)$$

We note that different values can be assigned to the upper and lower values of γ and $\dot{\gamma}$ if so required. All of the preceding constraints need to be accounted for in the trajectory planning algorithm.

C. Cost Function

We next specify the class of cost functions that we consider. Referring back to Eq. (1), the cost functional will be of the form

$$J = \sum_{\kappa \in \Omega} w_{\kappa} J_{\kappa} \quad (8)$$

where $\Omega = \{\text{time, altitude, riding_qualities}\}$ is the set of criteria, possibly augmented by others as discussed next. The weights w_{κ} control the tradeoff between these sometimes conflicting criteria.

We first consider integral costs of the form

$$J_{\kappa} = \int_0^{t_f} c_{\kappa}(\mathbf{X}(t)) dt \quad (9)$$

where the running cost function $c_{\kappa}(\mathbf{X})$ is to be chosen appropriately. We next indicate some possible choices of special interest. Note that a general terminal cost of the form $\phi[\mathbf{X}(t_f), t_f]$ can be accommodated as well, but seems to be of limited use in the present context.

The flight duration t_f is an obvious performance measure. It can be expressed as a unit integral over the flight interval:

$$J_{\text{time}} = \int_0^{t_f} 1 \, dt \quad (10)$$

Next, the length of the trajectory is obtained by letting $c_\kappa(\mathbf{X}) = V$ (recall that V is a function of \mathbf{X} in our model). More generally, the fuel consumption rate can be approximated as a function $c_{\text{fuel}}(\mathbf{X})$ of the flight configuration \mathbf{X} , leading to the cost component J_{fuel} , which evaluates the total fuel consumption. In a typical application the fuel consumption is approximately proportional to the flight time, and in our experiments we only consider the latter.

Altitude cost components will be of the general form

$$J_{\text{altitude}} = \int_0^{t_f} F(z(t), \text{DTM}(t)) \, dt \quad (11)$$

where $z(t)$ is the aircraft altitude at time t , $\text{DTM}(t)$ is the ground height at the current (x, y) position, as can be read in the onboard digital terrain map, and F is an appropriate function thereof. For example, in terrain-following flight, we might require that the aircraft maintains a fixed clearance H_0 above the ground and penalize for deviations by

$$J_{\text{altitude}} = \int_0^{t_f} |z(t) - \text{DTM}(t) - H_0| \, dt \quad (12)$$

Another interesting case is when the flight altitude is required to remain below a certain level H , for example, to keep the aircraft below a cloud layer. An appropriate cost function can take the form

$$J_{\text{altitude}} = \int_0^{t_f} I\{z(t) > H\} \, dt \quad (13)$$

where $I\{\}$ is the indicator function [which equals one if $z(t) > H$, and zero otherwise]. In either case it is natural to impose a hard constraint of the form $z(t) \geq \text{DTM}(t) + h_0$, where h_0 is a safety clearance. This is easily incorporated in our graph model.

A second class of cost criteria that will be of interest to us concerns the passenger comfort, as affected by the lateral accelerations. We shall use the term riding qualities to refer to this criterion. Using Eq. (4), the proposed measure of the riding qualities is

$$J_{\text{riding_qualities}} = \int_0^{t_f} (|\dot{\gamma}(t)| + \alpha|\dot{\psi}(t)|)V(t) \, dt \quad (14)$$

This criterion quantifies the lateral acceleration components, its values being accumulated over the entire trajectory. These accelerations influence both the pilot and the passengers. The parameter α allows for relative weighting of the two component.

An additional concern for trajectory planning is the pilot workload. During the flight, the pilot is preoccupied with numerous tasks on top of the need to follow the outlined trajectory. Having a rapidly varying trajectory requires the pilot to constantly focus on the task of following it, paying less attention to other instrument and visual data. Furthermore, the accumulated workload can lead to excessive stress and fatigue and should therefore be an important consideration during the trajectory construction process, especially for low flight over irregular terrain. To some extent, the riding qualities cost as just defined accounts also for the pilot workload, as it quantifies the total angular change that the trajectory requires. Other measures of pilot workload can include the number of required maneuvers, the total time that the flight is nonstraight or nonlevel, and other related quantities. Obviously, the proper choice of cost here requires careful consideration of the human engineering factors involved. We shall therefore rely on the riding qualities cost to regulate the smoothness of the planned trajectory.

III. Basic Graph-Search Solution

In this section we propose an approximate solution of the trajectory optimization problem that was just presented. We begin by discretizing the continuous model. We then translate the problem into a weighted graph-search problem, which can be solved using Dijkstra's shortest-path algorithm.

A. Discretization

The proposed discretization scheme begins with a uniform discretization of the three-dimensional space of spatial positions. The discretized trajectory is constrained to go through the grid points, connected with straight line segments. The actual trajectory will of course be a smoothed version of the discrete one. The smoothed trajectory can be computed by a postprocessing method, or simply left to be determined by the pilot and the guidance filter.

Consider the discrete grid $\{(ih_x, jh_y, kh_z)\} \subset \mathbb{R}^3$, where the constants h_x, h_y, h_z determine the grid spacing in the respective coordinates, and i, j, k take integer values. A finite grid S_0 is obtained by intersecting this set with the set of feasible spatial points in continuous space. In particular, the definition of S_0 should incorporate all hard constraints concerning forbidden flight zones and altitudes. The horizontal grid spacings h_x and h_y will usually be identical, which we assume henceforth. The choice of these resolution parameters is discussed later.

We now create a mesh over the spatial grid by defining a neighborhood structure. The neighbors of a point $\mathbf{P} = (ih_x, jh_y, kh_z)$ are those grid points that are accessible from it in one step. The simplest definition that is useful for our purpose is given by

$$\text{Neighb}[(ih_x, jh_y, kh_z)] = \{[(i+m)h_x, (j+n)h_y, (k+\ell)h_z]$$

$$: m, n \in \{-1, 0, +1\}, |\ell| \leq \ell_{\text{max}}^{m,n}\} \quad (15)$$

with $n = m = 0$ excluded. Thus, the horizontal plane neighbors are the eight nearest points (Fig. 2a). Other possibilities will be discussed next then, ℓ_{max} is constrained by the allowed elevation angle γ_{max} [Eq. (7)] and should therefore obey the following bound:

$$h_z \ell_{\text{max}}^{m,n} \leq \tan(\gamma_{\text{max}}) \sqrt{(mh_x)^2 + (nh_y)^2} \quad (16)$$

For simplicity we assume next that $\ell_{\text{max}}^{m,n} = \ell_{\text{max}}$ is independent of m and n .

As the discrete path is restricted to neighboring mesh points, the specified neighborhood structure determines the discrete values that the flight-path angles γ and ψ can assume. Given two neighboring grid points $\mathbf{P}_1 = (x_1, y_1, z_1)$ and $\mathbf{P}_2 = (x_2, y_2, z_2)$, with $\mathbf{P}_2 - \mathbf{P}_1 = (mh_x, nh_y, \ell h_z)$, their relative path angles are given by

$$\psi_{1 \rightarrow 2} = \tan^{-1} \left(\frac{y_2 - y_1}{x_2 - x_1} \right) = \tan^{-1} \left(\frac{nh_y}{mh_x} \right) \quad (17)$$

$$\gamma_{1 \rightarrow 2} = \tan^{-1} \left(\frac{\Delta z}{\sqrt{\Delta x^2 + \Delta y^2}} \right) = \tan^{-1} \left(\frac{\ell h_z}{\sqrt{(mh_x)^2 + (nh_y)^2}} \right) \quad (18)$$

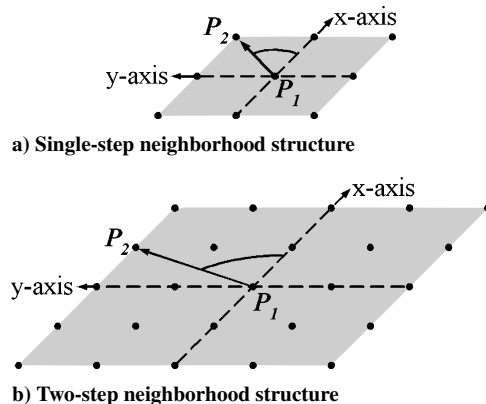


Fig. 2 Horizontal-plane neighbors of a grid point.

(where inverse tangent in the first equation should be interpreted as the four-quadrant inverse tangent, according to the signs of the nominator and denominator in its argument.)

Given the preceding neighborhood structure and $h_x = h_y$, the heading angle ψ is constrained to multiples of $\pi/4$. A finer angular resolution can be obtained by extending the definition of the neighboring points beyond the nearest grid points, as illustrated in Fig. 2b. Naturally, this extended definition will entail an increase in search time.

We next consider a convenient representation for the discretized state vector $\mathbf{X} = (x, y, z, \gamma, \psi)$. We shall use the representation $\mathbf{X} = (\mathbf{P}_1, \mathbf{P}_2)$, where the grid point $\mathbf{P}_1 = (x, y, z)$ provides the current position, and \mathbf{P}_2 gives the path angles according to Eqs. (17) and (18). In fact, we shall consider only such pairs $(\mathbf{P}_1, \mathbf{P}_2)$, where \mathbf{P}_2 is a neighbor of \mathbf{P}_1 , namely, $\mathbf{P}_2 \in \text{Neighb}(\mathbf{P}_1)$, so that γ and ψ are restricted to those values that point directly to neighboring points. This leads to the graph structure that we define later.

The choice of the grid resolution is a decisive factor in the complexity and accuracy of the resultant algorithm. We outline briefly some general guidelines for choosing the grid separation parameters. A specific example can be found in Sec. V. Starting with the horizontal parameters $h_x = h_y$, their choice should be directly dictated by the required trajectory resolution. In particular, if we require the trajectory to be specified to within an accuracy d_0 in the horizontal plane, we need to choose $\sqrt{2}h_x \leq d_0$. Moreover, to accommodate the allowed turning radius R_{\min} , and assuming that we do not allow discrete heading angle changes of more than 90 deg, h_x should not be larger than R_{\min} . When the accuracy requirement is the dominant one, as can be expected, we will have $h_x \ll R_{\min}$. In that case special care is required to enforce the turn radius constraint within the algorithm, as discussed in Sec. III.D. Finally, the resulting resolution parameter h_x will typically be larger than the available DTM grid separation, and for convenience can be chosen as a multiple thereof.

The vertical spacing h_z is chosen to satisfy several requirements. The obvious one is the required resolution in the flight elevation. Next is the required resolution in the vertical flight-path angle (or climb angle). The largest separation between two discrete angles can be seen to equal $\Delta\gamma_0 = \tan^{-1}(h_z/h_x)$ (this is just the first possible angle above zero when flying along horizontal grid lines), and so h_z should be chosen so that this value is small enough. Furthermore, h_z must be small enough so that a change in the (discrete) climb angles between adjacent grid points can be done within the allowed range of vertical acceleration. Specifically, suppose that the maximal vertical acceleration $a = \bar{a}_v$ is applied to the aircraft over a distance h_x , which is traversed in time $T = h_x/V$. Then the vertical increment (over a zero acceleration or straight line path) equals $aT^2/2$, so that we must have h_z not more than that:

$$h_z \leq \frac{1}{2}\bar{a}_v(h_x/V)^2 \quad (19)$$

Typically this bound should be fairly large, so that h_z can be chosen based on the required vertical and angular resolutions.

B. Graph Structure and Costs

We next define a graph structure for the discrete problem. This includes a finite set of vertices \mathcal{V} , a set of edges \mathcal{E} connecting pairs of vertices, and a cost for each edge. We shall also refer to \mathcal{V} as the (discrete) state space of the problem. Two vertices that are connected by an edge are said to be neighboring vertices. Note the distinction that is made here between neighboring vertices and neighboring points, as defined in the preceding subsection.

A vertex v corresponds to a discretized state \mathbf{X} as just defined, namely, a pair of neighboring points in the cartesian three-dimensional grid \mathcal{S}_0 :

$$v = (\mathbf{P}_1, \mathbf{P}_2) \in \mathcal{S}_0 \times \mathcal{S}_0 \quad \text{with} \quad \mathbf{P}_2 \in \text{Neighb}(\mathbf{P}_1) \quad (20)$$

Here $\mathbf{P}_1 = (x_1, y_1, z_1)$ is the current location of the aircraft, and $\mathbf{P}_2 - \mathbf{P}_1$ defines the direction of its velocity vector. Because the aircraft necessarily follows its specified velocity vector, \mathbf{P}_2 must be the aircraft's next location on the grid. A neighboring vertex u of a vertex $v = (\mathbf{P}_1, \mathbf{P}_2)$ is therefore of the form




Fig. 3 Vertex and its feasible neighbors.

$u = (\mathbf{P}_2, \mathbf{P}_3) = [(x_2, y_2, z_2), (x_3, y_3, z_3)]$, with \mathbf{P}_3 a neighbor of \mathbf{P}_2 as defined by Eq. (15). An edge in this graph is therefore of the form $[(\mathbf{P}_1, \mathbf{P}_2), (\mathbf{P}_2, \mathbf{P}_3)]$, with $\mathbf{P}_1 \in \mathcal{S}_0$, $\mathbf{P}_2 \in \text{Neighb}(\mathbf{P}_1)$, and $\mathbf{P}_3 \in \text{Neighb}(\mathbf{P}_2)$.

An example of a vertex $(\mathbf{P}_1, \mathbf{P}_2)$ and some of its neighbors $(\mathbf{P}_2, \mathbf{P}_3)$ is depicted in Fig. 3. Note that not all neighboring vertices will be feasible because of turn radius constraints, and indeed the figure does not depict neighboring vertices that require a turn of 180 deg or more. We will consider this constraint separately in Sec. III.D and use it to further restrict the set of edges of the searched graph.

Having defined the neighbors of a vertex, the cost $C(v, u)$ for each edge can now be specified. In accordance with Eq. (8), this cost will be in the form $C(v, u) = \sum_{\kappa} w_{\kappa} C_{\kappa}(v, u)$. We start with the cost components of the form [Eq. (9)]. Recall that each vertex v defines a configuration vector X_v of the form $X_v = (x_v, y_v, z_v, \gamma_v, \psi_v)$. The flight time between vertices v and u can be approximated as the distance divided by the average velocity, namely,

$$\Delta t_{vu} = \sqrt{(x_u - x_v)^2 + (y_u - y_v)^2 + (z_u - z_v)^2} / \bar{V}_{vu} \quad (21)$$

where $\bar{V}_{vu} = [V(X_v) + V(X_u)]/2$. Consequently, for a running cost function $c_{\kappa}(\mathbf{X})$ we let

$$C_{\kappa}(v, u) = \{[c(X_v) + c(X_u)]/2\} \Delta t_{vu} \quad (22)$$

In particular, for the time component from Eq. (10) we obtain

$$C_{\text{time}}(v, u) = \Delta t_{vu} \quad (23)$$

and for the altitude component from Eq. (11):

$$C_{\text{altitude}}(v, u) = \{[F(z_v, \text{DTM}_v) + F(z_u, \text{DTM}_u)]/2\} \Delta t_{vu} \quad (24)$$

where $\text{DTM}_v = \text{DTM}(x_v, y_v)$.

Next, using the notation of Eqs. (17) and (18) applied to the vertices $v = (\mathbf{P}_1, \mathbf{P}_2)$ and $u = (\mathbf{P}_2, \mathbf{P}_3)$, define

$$\Delta\gamma = \gamma_{2 \rightarrow 3} - \gamma_{1 \rightarrow 2} \quad (25)$$

$$\Delta\psi = \psi_{2 \rightarrow 3} - \psi_{1 \rightarrow 2} \quad (26)$$

The riding qualities cost component [Eq. (14)] is then given by

$$C_{\text{riding_qualities}}(v, u) = (|\Delta\gamma| + \alpha|\Delta\psi|) \bar{V}_{vu} \quad (27)$$

We have thus defined the graph $G = (\mathcal{V}, \mathcal{E})$ with nonnegative edge costs, $C(u, v) \geq 0$. The required solution is the shortest path on this graph, which can be found using standard graph-search algorithms. We next describe the basic Dijkstra algorithm as applies to this

problem and then consider certain modifications that are required in order to account for the turn radius constraints.

C. Dijkstra's Algorithm

Dijkstra's algorithm²⁷ provides an efficient solution to the minimal-path graph-search problem. Essentially, the algorithm recursively evolves the front of vertices that are closest to the source, until a target vertex is reached. Convergence is therefore guaranteed.

A pseudocode implementation of the Dijkstra algorithm is shown here:

```

Dijkstra ( $G = (\mathcal{V}, \mathcal{E}), C, s, d$ )
 $S = \{s\}$ 
PriorityQueue =  $\{s\}$ 
 $g(s) = 0$ 
for all  $v \in \mathcal{V} \setminus S$ ,
   $g(v) = \infty$ 
  predecessor( $v$ ) = none
  PriorityQueue.insert( $v$ )
while  $D \cap S = \emptyset$ ,
   $v =$  PriorityQueue.extract.minimum()
  retire  $v$  to  $S$ 
  for  $u \in$  neighbors( $v$ ),
    if ( $g(u) > g(v) + C(v, u)$ ) then
      PriorityQueue.decrease_key ( $u, g(v) + C(v, u)$ )
      PriorityQueue.predecessor ( $u$ ) =  $v$ 

```

Backtrack from D to s

$G = (\mathcal{V}, \mathcal{E})$ denotes the graph and C is the cost function over the edges \mathcal{E} . The algorithm stores for each vertex v an estimate $g(v)$ of the cost so far, namely, the cost of the shortest path from the source, and $g(v)$ is initialized to $+\infty$, except for the source where it is initialized to 0.

The algorithm also maintains a set $S \subset \mathcal{V}$ of vertices whose final shortest-path values have already been determined. This set is initialized with the source vertex $\{s\}$ and $g(s) = 0$. The algorithm then repeatedly selects (or retires) the vertex v in the complement of S , which is closest to the source. Whenever a vertex v is retired to S , we examine each of its neighbors that have not yet been retired. For each such neighbor u , we evaluate $C(u, v) + g(v)$, which is u 's distance to the destination if proceeding through v . If this is smaller than the current estimate $g(u)$ of u 's distance, we replace the latter by the former. After performing this update for all neighbors of v , the next vertex to be retired is the one with the smallest estimate $g(u)$ out of all unretired vertices. This selection is effectively handled by managing a priority queue, which maintains the best cost estimates so far for all neighbors of S . For any vertex u retired to S , we also save its predecessor v , so that the optimal path can be traced back.

It is readily verified that Dijkstra's algorithm does compute the shortest path, under the condition that the link costs are nonnegative.²⁸ The time complexity of the algorithm is $\mathcal{O}(N \log N)$, where N is the number of vertices in \mathcal{V} , provided that the degree of the graph (number of neighbors for each node) is bounded, and the queue is implemented with an efficient data structure. To realize that, note that the algorithm performs a single pass over the state space, with each vertex requiring at most one priority queue extraction and a number of priority queue updates bounded by the number of its neighbors. An efficient priority queue data structure, such as the binary heap, can handle extraction of a minimal element or value update in a list of N elements at $\mathcal{O}(\log N)$ each.

Returning to the flight trajectory optimization, Dijkstra's algorithm can be directly applied to solve this problem, provided that we ignore for the moment possible turn radius constraints. To estimate the time complexity of the algorithm, assume that the three-dimensional mesh has n points. Let m be the maximal number of neighbors for each point. For the neighborhood structure of Eq. (15), we have that $m \leq 8(2\ell_{\max} + 1)$. As a vertex in a pair of two neighboring points, the number of vertices is $N \leq mn$, and the basic Dijkstra algorithm runs at $\mathcal{O}(N \log N)$.

D. Constrained Dijkstra

An important factor that has been left out of the preceding algorithm is the turn radius constraint. In a typical situation, the minimal turn radius in the horizontal plane [as per Eq. (6)] can be significantly larger than the horizontal resolution h_x . In that case, in order to al-

low large single-step heading changes as required in our scheme the turn radius constraint should be enforced by considering a span of several consecutive vertices. For example, following a 90-deg turn the prescribed trajectory should not have another heading change (in the same direction) for a certain stretch of flight, so that the aircraft can complete the first maneuver.

An optimal solution to such multistep constraints would require to expand the state space into sequences of points of required length. This would easily lead to an increase of the state space by orders of magnitude, with a corresponding increase in the computation times. To avoid this problem, we propose a suboptimal modification of the basic Dijkstra algorithm, which retains the original state space.

The constrained Dijkstra algorithm utilizes the basic Dijkstra data structure, which allows tracing back the optimal path (to a required number of steps) for each retired vertex. Recall that whenever a state is retired (to the set S), it causes an update in the estimated cost so far of its (nonretired) neighbors. In the proposed modification, when a vertex is retired, it traces back the recorded path to the required number of steps, and then determines the set of its feasible neighbors in light of the added multistep constraints. Thus, the trajectory cannot proceed to vertices that violate these constraints. Other than that, the algorithm proceeds exactly as before. It is evident that the complexity of the algorithm remains the same, namely, $\mathcal{O}(N \log N)$.

To handle the turn radius constraint, it remains to specify the sequences of vertices that obey or violate this constraint. This can be done in different ways, and we specify next a simple heuristic criterion that will also be used in our experiments. With each (discrete) change in the heading angle of magnitude α , we associate a distance d before and after the turn point, where d designates the travel distance required to complete the turn under the turn radius constraint R_{\min} . Using the geometry in Fig. 4, we obtain $d = R_{\min} \tan(\theta/2)$. This rule can be easily converted into a "dictionary" of allowed turn sequences over the discrete grid.

As an illustration for the effect of the turn radius constraint, we depict in Fig. 5 two trajectories. One is created by the basic Dijkstra




Fig. 4 Trajectory segment.

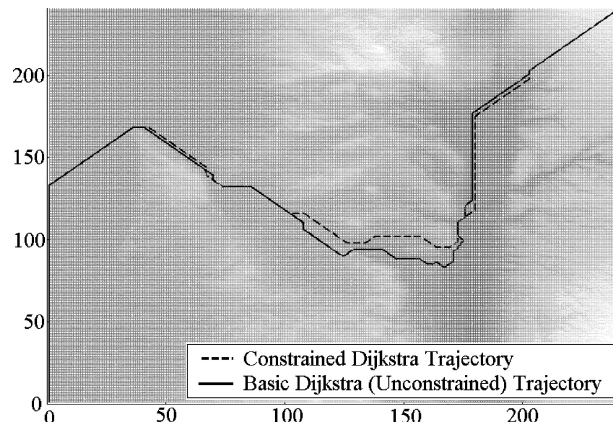


Fig. 5 Basic vs constrained Dijkstra trajectories: top view. Darker parts of the map indicate lower ground levels.

algorithm, while the other imposes turn radius constraints via the constrained Dijkstra algorithm. The general conditions of the simulation are as detailed in Sec. V. In this example, an altitude cost component penalized the aircraft for an altitude of more than 100 m above the ground, and the lateral acceleration cost components were removed in order to highlight the influence of the turn radius constraint. The smoothing affect of this constraint is apparent in the form of the resulting trajectory.

IV. Accelerating the Algorithm

The preceding section described a Dijkstra-based solution to our trajectory optimization problem. Despite the efficient search procedure, running times are still excessive because of the large cardinality of the discretized state space. For quantification metrics, see Sec. V.

In this section we present several approaches for accelerating the Dijkstra search. We begin by describing an A^* enhanced version, and explain why the acceleration we obtain is only minor.

We focus therefore on two additional Dijkstra-based algorithms that trade off accuracy for speed. The reduced-state Dijkstra algorithm cuts down the number of vertices in the search graph, by redefining a vertex as a spatial point without the angular (or next-point) components. Consideration of the angular components is retained through the backtracking property of the Dijkstra search. The hierarchical Dijkstra algorithm conducts an initial search at a coarser spatial resolution, which constrains the finer search to a smaller area of the map. Each algorithm by itself achieves a significant speedup compared to the basic Dijkstra algorithms of the preceding section. Combined, the two modified algorithms offer an onboard solution to the trajectory generation problem. As we demonstrate empirically in Sec. V, the performance degradation caused by these suboptimal modifications is minor.

A. Reduced-State Dijkstra

The major factor that slows down the Dijkstra algorithm is the large state space that needs to be explored. As was discussed before, our Dijkstra algorithm employs a state (or vertex) that comprises two neighboring spatial points, resulting in a state-space cardinality of mn , where n is the size of the spatial grid and m is the average number of neighbors to each point in that grid (see Sec. III). This two-point state decodes the flight velocity vector so that we can deduce the angular acceleration (and its associated costs) between two neighboring vertices. The underlying idea behind the reduced-state Dijkstra is to settle for a single-point vertex, while still accounting for all cost components, albeit in a suboptimal way. To take account of acceleration cost components, we shall use the trace-back feature that is inherent in the Dijkstra algorithm. The basic idea is very similar to the one used in Sec. III.D to enforce the turn radius constraints.

The reduced-state Dijkstra algorithm thus reduces the state space to n vertices. Its running time is therefore $\mathcal{O}(n \log n)$. A vertex is defined to be a single point in the three-dimensional mesh, that is, $v = \mathbf{X} = (x_1, y_1, z_1) \in \mathcal{S}_0$. A neighbor of vertex v is defined by Eq. (15).

Computation of the angular acceleration (and the associated riding qualities cost) requires three consecutive spatial points. Referring back to Dijkstra's algorithm, recall that for each retired vertex v (now a single spatial point) the algorithm stores its predecessor vertex in order to trace back the optimal path. Now, whenever a vertex is retired the algorithm is required to update the cost-to-go estimates for its neighbors. For such a vertex v we use the following notation: $\text{pre}(v) = (x_1, y_1, z_1)$ is its predecessor vertex; $v = (x_2, y_2, z_2)$ is the current vertex; and $u = (x_3, y_3, z_3)$ is the neighboring vertex. With these definitions, the edge cost $C(v, u)$ can be computed by using exactly the same equations (23–27) that were used before. Turn radius constraints can be incorporated into this algorithm exactly as described in Sec. III.D.

It should be apparent that the search algorithm proposed here is suboptimal when all components of the cost function are present. Specifically, it cannot fully account for the riding-quality cost that makes the optimal value function depend also on the directional

components of the full state. Instead of exploring all possible velocity vectors (or predecessor points) in a given spatial point, we freeze the velocity vector that leads to it the first time that it is reached. Other options (which could eventually turn up to be the optimal ones) are suppressed and excluded from further consideration. Nonetheless, the cost discrepancies between the results obtained by the full state Dijkstra solution and those obtained via the reduced-state Dijkstra solution turn out to be minor in all of our experiments. One way to explain this in the context of our problem is by noting that (small) deviations from the optimal velocity can usually be corrected later on with little penalty. The speedup achieved by the state reduction is, however, most substantial.

B. A^* Algorithm

The A^* algorithm, introduced by Nilsson,²⁹ accelerates graph-searching procedures by adding a heuristic cost component to direct the search towards an optimal route. Essentially, the heuristic cost function is an a priori estimate of the cost-to-go function from each vertex. In our problem, the search procedures used so far did not make use of any information concerning the location of the target state. It seems reasonable, for example, to give some priority to those vertices that point in the direction of the goal, rather than those that point in the opposite direction.

Recall that the Dijkstra algorithm used the cost-so-far function $g(v)$ to direct the search. This function is used to choose the vertex that is explored and retired at every step (by determining the order of vertices in the priority queue). The A^* algorithm uses for that purpose a modified cost $\tilde{g}(v)$, where $\tilde{g}(v) = g(v) + h(v)$ and $h(v)$ is a prespecified heuristic. The algorithm, as used here, is otherwise identical to the Dijkstra algorithm just described. Nilsson²⁹ proved that when $h(v)$ is an underestimate of the minimal cost to go between the current vertex v and the destination vertex, then A^* produces an optimal solution. In the ideal case that $h(v)$ equals the optimal cost to go, the algorithm directly finds the optimal path. Setting to $h(v) = 0$ recovers the basic Dijkstra algorithm.

The challenge in designing the A^* algorithm is in finding a good underestimating heuristic. We approach this task by examining separately each component of our cost function.

An underestimating time heuristic should be smaller than the (optimal trajectory) flight time from any given vertex v to the destination vertex d . A simple time heuristic is therefore the Euclidean distance from our current position to the destination, divided by the speed. The accuracy of this estimate depends of course on the deviation of the optimal trajectory from the straight line, which depends mainly on the other cost components.

On the other hand, finding a reasonable underestimate for the altitude cost or the riding-qualities cost is more challenging. If we look at each of these cost components separately while completely disregarding the others, we obtain an underestimate that tends to zero, which is obviously not helpful. For example, the altitude cost is an indication of how high above the terrain the aircraft is flying: the aircraft can simply fly at the zero-cost height all of the way to nullify that cost component. When only the riding qualities are considered, the best path is a direct course toward the destination, also resulting in near-zero cost.

So far, we have not been able to come up with an easily computable heuristic that adequately approximates these cost components in the “interesting” cases (namely, where the optimal flight path is not a straight line). Some initial attempts at analyzing the spatial regularity of the terrain (using Fourier transforms and wavelet methods) have not led to noticeable improvement in the algorithm. We therefore restrict further consideration to the just-mentioned straight line heuristics for the flight time, which will be examined in the experimental section.

C. Hierarchical Dijkstra

Although the reduced-state Dijkstra algorithm decreases the total number of vertices in the graph, the hierarchical Dijkstra restricts the detailed search to a specific segment in the overall graph. Hierarchical and multigrid schemes are of course widely used tools in numerical analysis and optimization, and we describe here a

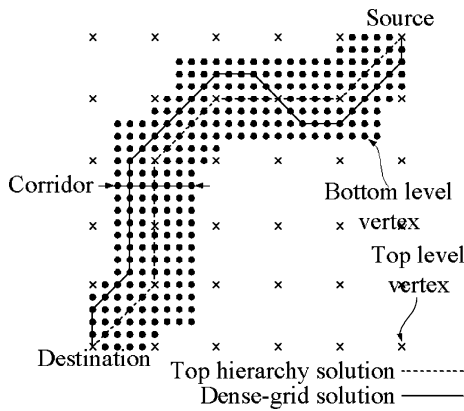


Fig. 6 Illustration of the hierarchical scheme.

simple coarse-fine scheme that is convenient for our application. The search problem is first solved at a coarser level, which is obtained by some downsampling of the original spatial grid.

The implementation considered here uses a downsampling by a factor of k along each axis, leading to a k^3 reduction in the spatial grid. The DTM was correspondingly downsampled. Because safety is a major concern, the ground elevation at each sampled point was taken as the maximal elevation over the map points it covers.

The algorithm calculates an optimal path at the coarse level, using the downsampled DTM. The resulting coarse-level trajectory defines a search corridor over the original (fine) spatial grid. The points in this corridor are simply those that are within a specified distance from the coarse trajectory. The fine-level search is then conducted inside this corridor (Fig. 6). The same algorithm can be used for both the coarse- and fine-level searches.

V. Experimental Results

This section describes some simulation experiments that demonstrate and compare the performance of the proposed algorithms. We consider the following issues: 1) performance of the reduced-state Dijkstra algorithm, as compared to the standard (full-state) algorithm; 2) effect of the cost parameters on the obtained trajectory; and 3) the A^* and Hierarchical enhancements to the reduced-state Dijkstra algorithm.

The DTM that was used is a 100×100 km digital map with a 50-m resolution grid. The map coordinates are as follows; bottom left (origin) is $N32^\circ 10'$, $E34^\circ 50'$ and top right (destination) is $N33^\circ 10'$, $E35^\circ 50'$.

The state at the origin point is specified by $\psi_0 = 45$ deg (heading toward the destination), $\gamma_0 = 0$. These variables are left free at the destination points. The initial and final heights are set to ground level. We assume a constant flight velocity of $V = 100$ m/s and a required bound of 5 m/s^2 ($g/2$) on the two lateral acceleration components. The cost function was taken according to Eqs. (8), (10), (12), and (14). The following weights were used, unless otherwise stated: $w_{\text{time}} = 0.2$, $w_{\text{altitude}} = 1$, and $w_{\text{riding-qualities}} = 1$. These weights were selected after some empirical experimentation that found them to give a good tradeoff between the various cost criteria.

The basic horizontal resolutions (h_x and h_y) were set to 800 m. This presents a reasonable trajectory resolution for general aviation. The DTM was accordingly downsampled. The minimal turn radius constraint from Eq. (6) computes as $R_{\min} = 2000$ m. As this is larger than h_x , the turn radius constraint will be enforced throughout by using the constrained Dijkstra method of Sec. III.D. The vertical resolution parameter h_z was set to 30 m (100 ft), with $\ell_{\max} = 2$ [see Eq. (15)]. This gives an angular resolution of $\tan^{-1}(h_z/h_x) = 2.15$ deg in the vertical flight-path angle γ , with $\gamma_{\max} = 4.3$ deg. The acceleration constraint from Eq. (19) requires $h_z \leq 160$ m, which is well satisfied.

The range of flight elevations considered was 1500 m, yielding 50 points in the vertical axis, and total number of states in the spatial grid is $n = 125^2 \times 50 = 625,000$. The number of neighbor points

in $m = 8 \times (2\ell_{\max} + 1) = 40$, which brings the number of vertices in the full-state formulation to $N = mn = 25 \times 10^6$. Simulations were conducted on an Intel® Pentium® 3,866 MHz with 0.75 GB RAM CPU. The algorithmic engine is a Microsoft® Visual C++ program, and the graphics engine is MathWorks MATLAB®.

The first experiment compares the full-state Dijkstra solution with the reduced-state Dijkstra algorithm. The results are shown in Fig. 7. The full-state Dijkstra algorithm takes some 2 h to complete, as opposed to 126 s for the reduced-state Dijkstra. The cost degradation of the latter is less than 3%. This gives close to 60-fold speedup, with a minor loss in performance. Recall that the reduced-state algorithms reduce the state-space cardinality by a nominal factor of $m = 40$, the number of neighbors of each point as just computed. This is directly reflected in the reduction in computation time. The two algorithms were further compared for various problem parameters, with similar conclusions. In view of its efficiency, the reduced-state Dijkstra is used as the underlying algorithm in the remainder of our simulation experiments.

We next examine briefly the effect of varying the weights of the cost function components. Fig. 8 depicts the results for different

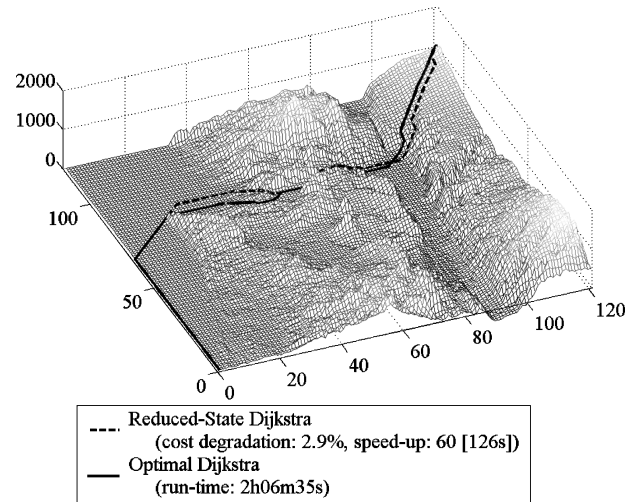


Fig. 7 Full vs reduced-state Dijkstra.

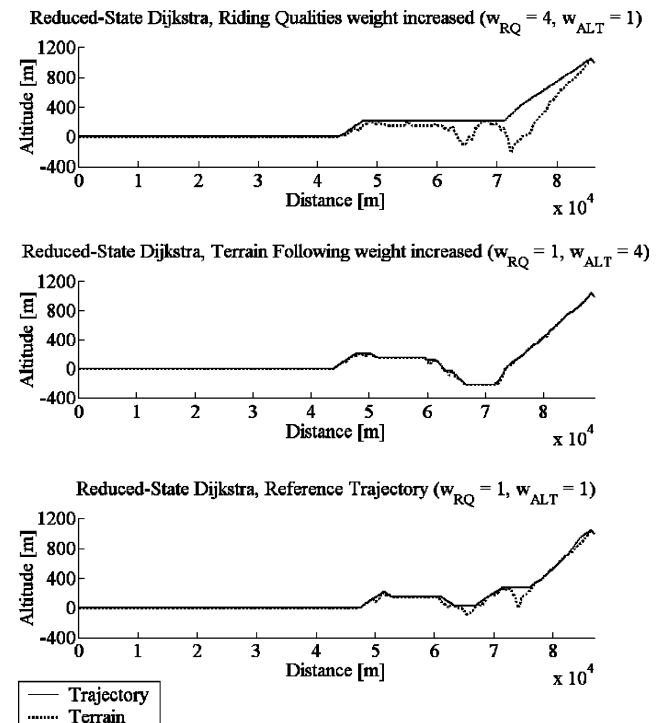


Fig. 8 Cost function criteria: trajectory smoothness tuning.

weights of the altitude cost as compared to the riding-qualities cost. The first run is conducted with unit weights to both cost criteria, resulting in the reference trajectory. The second run is conducted with the altitude weight w_{altitude} increased fourfold. Finally, the third run is conducted with the riding-qualities weight $w_{\text{riding-qualities}}$ increased fourfold.

Plotted in Fig. 9 are the unfolded sidewise projections (altitude vs flight distance) of the obtained trajectories. We can see that the riding-qualities-aware trajectory offers a smooth ride and does not follow closely the finer features of the terrain contour. The reference trajectory descends into such “valleys” to a certain degree, whereas the altitude-aware trajectory dives deep into them. We note that every trajectory passes through a slightly different area of the map; hence, the terrain under each trajectory is not the same. Overall, we achieve a smoothing effect, where the weights of the different cost criteria enable us to tune the degree of trajectory smoothness. Regarding the difference in run time, the altitude-aware trajectory barely probes into the vertical degree of freedom, thus concluding the search more rapidly.

Next, we consider the A^* variant of the algorithm with the time-to-go heuristic. Figure 9 shows the effect of different time-heuristic factors. The heuristic is taken to be $h(X) = wT_g(X)$, where T_g is the straight line underestimate of the remaining time to the goal position. The algorithm was tested with weights $w = 0.3, 1, 2, 4$, to which we refer as the slow, medium, fast, and faster heuristics. All runs were conducted with the reduced-state A^* algorithm. Figure 9 depicts the obtained trajectories, and Table 1 summarizes the run times and cost degradation (compared to the reduced-state Dijkstra reference) for these runs.

The two smaller weights (slow and medium) give rise to an underestimating heuristic and should not affect the optimal solution. However, they can affect the results of the (suboptimal) reduced-state Dijkstra. It is indeed seen that the slow heuristic did not modify the quality of the obtained solution, while the medium heuristic led to a slight degradation of the cost, and the computation time was about halved. The two more aggressive heuristics, on the other hand, led the trajectory to a different area of the map, with a considerable

Table 1 A^* heuristics comparison

Algorithm	Cost degradation	Run time, s
Reduced-state Dijkstra	Reference	126
Reduced-state A^* (slow)	0.0%	71
Reduced-state A^* (medium)	0.8%	56
Reduced-state A^* (fast)	22.1%	47
Reduced-state A^* (faster)	48.5%	14

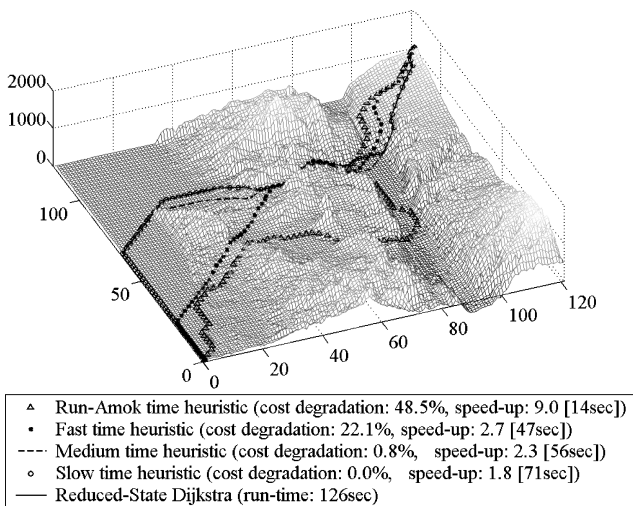


Fig. 9 A^* time heuristic within the reduced-state Dijkstra.

Table 2 Hierarchical Dijkstra: performance for different corridor parameters

Algorithm	Cost degradation	Run time, s
Reduced-state Dijkstra	Reference	126
Hierarchical reduced-state Dijkstra (20)	0.0%	83
Hierarchical reduced-state Dijkstra (10)	0.9%	63
Hierarchical reduced-state Dijkstra (7)	9.1%	57
Hierarchical reduced-state Dijkstra (3)	20.7%	40

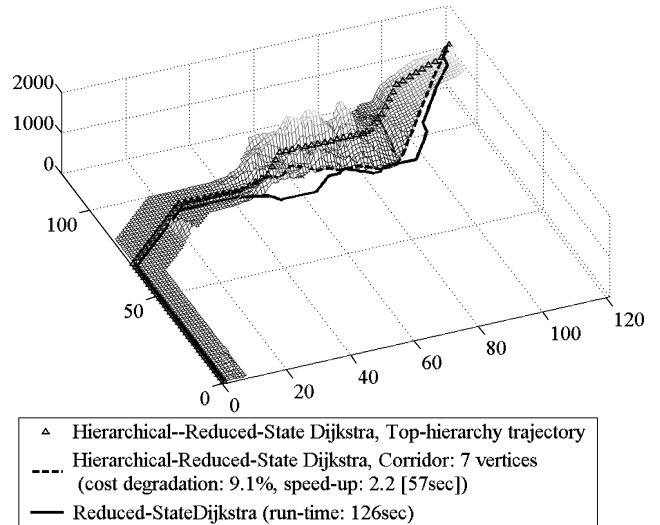


Fig. 10 Hierarchical algorithm: illustration of the corridor concept.

cost degradation. Thus, despite some added speedup, we consider them to be unacceptable.

The hierarchical variant of the Dijkstra algorithm is considered next. The hierarchical algorithm has two important parameters: the degree of downsampling of the original grid and the width of the corridor to explore around the top-level trajectory. The corridor concept is illustrated in Fig. 10. In Fig. 10 we use a downsampling of $k = 2$ in each axis. The algorithm first performs a search in the coarser ($60 \times 60 \times 20$) map. It then returns to the original map, opens a corridor seven vertices wide on each side of the coarse-level trajectory, and performs a detailed search in that corridor. The detailed search was not similarly constrained in the vertical axis. It is important to emphasize that the same basic algorithm is used for both hierarchical levels, namely, the reduced-state Dijkstra is used at both hierarchical levels.

The total run time of the hierarchical reduced-state Dijkstra algorithm (downsample, 2; corridor; 7) is 57 s, and the cost degradation is 9.1%. The top-level trajectory (diagonals) is found in just a few seconds, and the remaining time is used to find the detailed trajectory. It can be seen that the corridor around the coarse-level solution is not wide enough to encompass the optimal solution, so that the latter is not found by the finer search.

We next consider the effect of the two parameters of the hierarchical approach. First, we hold the downsampling of the map at a fixed value of two, and consider several corridor widths. This is shown in Fig. 11. As expected, increasing the width of the corridor improves the accuracy of the hierarchical algorithm and vice versa. In particular, a corridor of 10 vertices gives a close-to-optimal performance. Table 2 summarizes the run times and cost degradation (compared to the reduced-state Dijkstra reference) of different corridor parameters presented in Fig. 11.

Next, we consider different downsampling values while holding the corridor parameter constant at 10 vertices to each side. The results are shown in Fig. 12; the time improvement caused by increased downsampling is marginal, as most of the times is spent on the detailed find-level search. However, downsampling by four or more moves the search to a different area of the map altogether, and the solution deviates to an unacceptable degree from the optimal.

Table 3 Hierarchical Dijkstra: performance for different downsampling factors

Algorithm	Cost degradation	Run time, s
Hierarchical reduced-state Dijkstra (2)	Reference	63
Hierarchical reduced-state Dijkstra (3)	1.4%	55
Hierarchical reduced-state Dijkstra (4)	21.9%	67

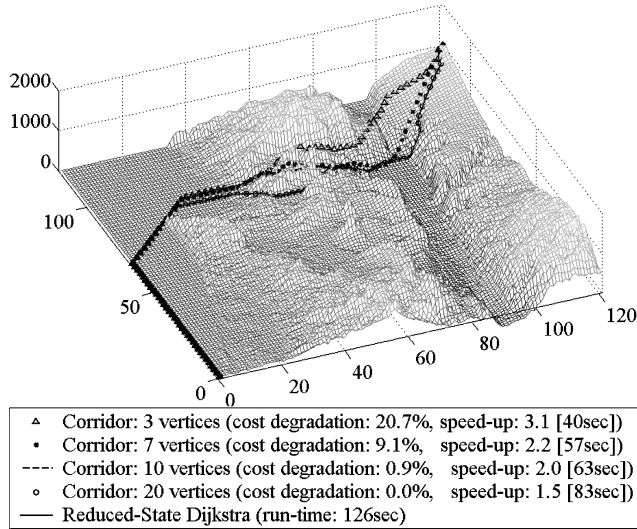
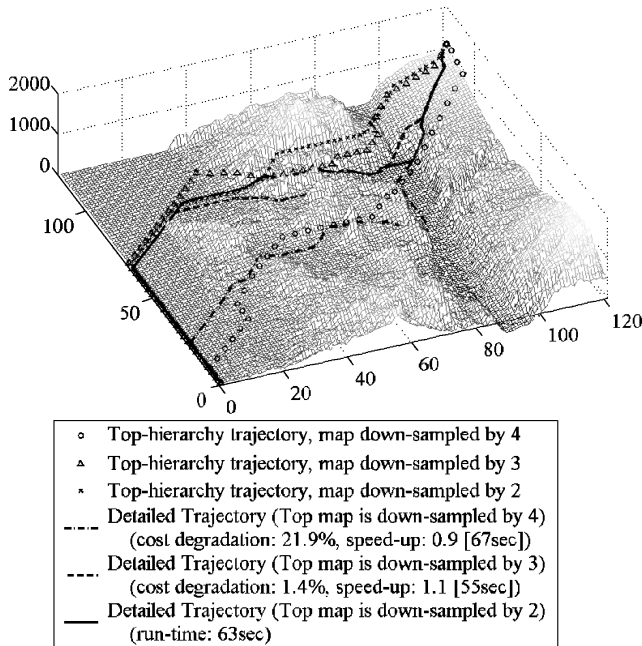
**Fig. 11 Hierarchical reduced-state Dijkstra: different corridor widths.****Fig. 12 Hierarchical reduced-state Dijkstra: different downsampling factors.**

Table 3 summarizes the run times and cost degradation (compared to the hierarchical reduced-state Dijkstra with downsampling of two) of the different downsampling parameters presented in Fig. 12.

The main results obtained in the just-described experiments are summarized in Table 4, to give a more complete perspective of the tradeoff of run-time speedup vs cost degradation. According to our previous findings, the A^* scheme utilizes the medium-time heuristic, and the hierarchical algorithm uses a corridor parameter of 10 vertices to each side and a downsampling factor of three. The last entry refers to a combined use of the A^* and hierarchical variants, yielding the fastest run time of 29 s, with mild performance degradation.

Table 4 Run time vs cost degradation comparison of main algorithms

Algorithm	Cost degradation	Run time, s
Full-state Dijkstra	Reference	7595
Reduced-state Dijkstra	2.9%	126
Reduced-state A^*	3.7%	56
Hierarchical reduced-state Dijkstra	4.3%	55
Hierarchical reduced-state A^*	4.4%	29

Obviously, the use of a more powerful computer and judicious programming will further improve the run time. These results clearly indicate the potential for real-time, onboard applicability.

VI. Conclusions

This paper examines the applicability of global graph-search techniques to realistic trajectory optimization problems for general-aviation applications. A central element of our model is the inclusion of acceleration-dependent cost, which accounts for passenger comfort (or riding qualities) and pilot workload. The complexity of the suggested algorithms was examined with respect to the goal of obtaining computation times that allow the use of the algorithm for onboard trajectory planning, namely, in the subminute area. Our results indicate that this goal is indeed feasible, by combining the computational acceleration methods that were considered.

Global-search techniques enjoy some distinct advantages. The basic Dijkstra algorithm indeed performs an exhaustive search over all possible paths and consequently obtains the globally optimal solution (up to the discretization error introduced in the discrete model setup). Even the reduced-state and other suboptimal acceleration schemes that were proposed perform an exhaustive search in Cartesian space, although with some restrictions on angular configurations. As a consequence, these algorithms do not need initialization with a candidate path or multiple runs with different initial condition, but rather obtain their best estimate for the global solution after a single run. Unstructured hard constraints that can induce multiple local optima, such as forbidden flight zones, are easily accommodated and merely reduce the search space. Moreover, the suggested algorithms are finite in nature, namely, they terminate with their solution after a finite number of steps. Thus, the computation time can be strictly bounded, and the question of convergence does not arise at all.

The reduction to a graph-search problem is obtained by discretizing the state space at some finite resolution. The selected resolution will obviously have a major impact on the accuracy of the obtained solution, as well as on computation time. We have outlined here some basic considerations that can lead to an appropriate choice of resolution for the intended application, so that the quantitative qualities of the continuous model are retained.

Many issues remain for further research. It would obviously be useful to obtain a theoretical handle on the approximation errors of the proposed algorithms, both those introduced as a result of the finite grid discretization and those that are caused by the suboptimality of the accelerated search algorithms. These should be accompanied by additional numerical experimentation. It would also be interesting to compare the results obtained here with standard solvers based on nonlinear programming, using direct or indirect methods. It should be appreciated, however, that these algorithms are essentially local in nature, so that they might not be suitable for problems with multiple local minima. The combination of global-search methods with local refinement using other (local) algorithms poses another potentially interesting topic that requires further inquiry.

The graph-search algorithms themselves can obviously be further enhanced. For example, the search can proceed simultaneously from both ends of the trajectory and terminate when these two parts meet. The effect of this and other possible enhancements should be evaluated empirically.

The discretization scheme that was used in this paper considers only trajectories that pass through the discrete grid points. As we mentioned in the Introduction, some recent Dijkstra-like algorithms

for the continuous shortest-path problems rely on more elaborate discretization schemes that allow paths that proceed in-between grid points^{25,26} and therefore have the potential for better accuracy for a given grid resolution. It would be therefore interesting to extend these methods to include the dynamic aspects of our problem, namely, the acceleration-related costs and constraints. From the end-user viewpoint, further work is required concerning the proper choice of cost function. This includes also the quantification of the pilot workload cost that was not elaborated in the present paper.

Acknowledgments

This research was supported by the Fund for the Promotion of Research at the Technion. We appreciate the valuable comments provided by J. L. Speyer and his research team and by M. Feron and his seminar group during informal presentations of this research at the University of California, Los Angeles, and Massachusetts Institute of Technology, respectively. We would like to thank the Associate Editor and the three referees for their helpful comments.

References

- ¹Hewitt, C., Henley, A. J., and Boyes, J. D., "A Ground and Obstacle Collision Avoidance Technique (GOCAT)," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 6, No. 8, 1991, pp. 13–20.
- ²Tymczyszyn, B., and Wilson, G., "Development and Flight Testing of a General/Corporate Aviation Terrain Awareness and Warning System," *2002 Report to the Aerospace Profession: Proceedings of the 46th Society of Experimental Test Pilots (SETP) Symposium*, Society of Experimental Test Pilots, Lancaster, CA, 2002, pp. 94–107.
- ³Clarke, K. C., *Getting Started with Geographic Information Systems*, Prentice-Hall, 1990, pp. 34–137.
- ⁴Fountain, J. R., "Digital Terrain Systems," *IEE Colloquium on Airborne Navigation Systems Workshop*, Digest No. 1997/169, Inst. of Electrical Engineers, London, 1997, pp. 4/1–4/6.
- ⁵Bryson, A. E., *Dynamic Optimization*, Addison Wesley Longman, Reading, MA, 1999.
- ⁶Bryson, A. E., and Ho, Y. C., *Applied Optimal Control*, Wiley, New York, 1975, pp. 1–127.
- ⁷Bertsekas, D. P., *Dynamic Programming and Optimal Control*, 2nd ed., Athena Scientific, Belmont, MA, 2000.
- ⁸Betts, J. T., "Survey of Numerical Methods for Trajectory Optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 21, No. 2, 1998, pp. 193–207.
- ⁹Kushner, H. J., and Dupuis, P. G., *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer-Verlag, New York, 1992.
- ¹⁰Funk, J. E., "Optimal-Path Precision Terrain-Following System," *Journal of Aircraft*, Vol. 14, No. 2, 1977, pp. 128–134.
- ¹¹Hargraves, C. R., and Paris, S. W., "Direct Trajectory Optimization Using Nonlinear Programming and Collocation," *Journal of Guidance, Control, and Dynamics*, Vol. 10, No. 4, 1998, pp. 338–342.
- ¹²Asseo, S. J., "Terrain Following/Terrain Avoidance Path Optimization Using the Method of Steepest Descent," *Proceedings of the IEEE National Aerospace and Electronics Conference*, Vol. 3, Inst. of Electrical and Electronics Engineers, New York, 1988, pp. 1128–1136.
- ¹³Menon, P. K. A., Kim, E., and Cheng, V. H. L., "Optimal Trajectory Synthesis for Terrain-Following Flight," *Journal of Guidance, Control, and Dynamics*, Vol. 14, No. 4, 1991, pp. 807–813.
- ¹⁴Lu, P., and Pierson, B. L., "Optimal Aircraft Terrain-Following Analysis and Trajectory Generation," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 3, 1995, pp. 555–560.
- ¹⁵Lu, P., "Inverse Dynamics Approach to Trajectory Optimization for an Aerospace Plane," *Journal of Guidance, Control, and Dynamics*, Vol. 16, No. 4, 1993, pp. 726–732.
- ¹⁶Spong, M. W., Thorp, J. S., and Kleinwaks, J. M. M., "The Control of Robot Manipulators with Bounded Input," *IEEE Transactions on Automatic Control*, Vol. AC-31, No. 6, 1986, pp. 483–490.
- ¹⁷Lee, S. M., Bien, Z., and Park, S. O., "On-Line Optimal Terrain-Tracking System," *Optimal Control Applications and Methods*, Vol. 11, No. 4, 1989, pp. 289–306.
- ¹⁸Rehbock, V., Teo, K. L., and Jennings, L. S., "A Linear Programming Approach to On-Line Constrained Optimal Terrain-Tracking Systems," *Optimal Control Applications and Methods*, Vol. 14, No. 4, 1993, pp. 229–241.
- ¹⁹Barnard, R., "Terrain Tracking Based on Optimal Aim Strategies," *Optimal Control Applications and Methods*, Vol. 15, No. 2, 1994, pp. 145–150.
- ²⁰Hess, R. A., and Jung, Y. C., "Generalized Predictive Control of Dynamic Systems," *Proceedings of the 1988 IEEE International Conference on Systems, Man and Cybernetics*, edited by J. Xinsong, Vol. 2, Elsevier, San Diego, CA, 1988, pp. 844–849.
- ²¹Jung, Y. C., and Hess, R. A., "Precise Flight-Path Control Using a Predictive Algorithm," *Journal of Guidance, Control, and Dynamics*, Vol. 14, No. 5, 1991, pp. 936–942.
- ²²Lu, P., "Nonlinear Predictive Controllers for Continuous Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 17, No. 3, 1994, pp. 553–560.
- ²³Lu, P., and Pierson, B. L., "Aircraft Terrain-Following Based on a Nonlinear Continuous Predictive Control Approach," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 4, 1995, pp. 817–823.
- ²⁴Mitchell, J. S. B., and Keirse, D. M., "Planning Strategic Paths Through Variable Terrain Data," *Applications of Artificial Intelligence, Proceedings of SPIE—International Society of Optical Engineering*, Vol. 456, Bellingham, WA, 1984, pp. 172–179.
- ²⁵Tsitsiklis, J. N., "Efficient Algorithms for Globally Optimal Trajectories," *IEEE Transactions on Automatic Control*, Vol. 40, No. 9, 1995, pp. 1528–1538.
- ²⁶Sethian, J. A., *Level Set Methods and Fast Marching Methods*, 2nd ed., Cambridge Univ. Press, Cambridge, England, U.K., 1999, Chap. 8.
- ²⁷Dijkstra, E. W., "A Note on Two Problems in Connection with Graphs," *Numerische Mathematik*, Vol. 1, 1959, pp. 269–271.
- ²⁸Cormen, T. H., Leiserson, C. E., and Rivest, R. L., *Introduction to Algorithms*, 2nd ed., MIT Press, Cambridge, MA, 2001, pp. 527–531.
- ²⁹Nilsson, N. J., *Principles of Artificial Intelligence*, Tioga Publishing Co., Palo Alto, CA, 1980, pp. 72–88.