aerospace

MDPI

Article

Incremental Simplicial Trajectory Replanning (ISTR*): Kinodynamic Search for Collision-Free Spacecraft Relative Motion

Tianji Ma [1], Bin Nan [1], Mingyao Sun [2] and Shunli Li [1,*]

[1] School of Astronautics, Harbin Institute of Technology, Harbin 150001, China; matianji1994@outlook.com (T.M.); nanbin12138@163.com (B.N.)
[2] Xi'an Modern Control Technology Research Institute, Xi'an 710065, China; mysunmtj@163.com
* Correspondence: lishunli@hit.edu.cn

Abstract

Collision-free relative motion planning for spacecraft in uncertain local orbital environments is critical for on-orbit servicing, inspection, and collision risk mitigation. Yet, existing graph- and tree-based planners typically restrict trajectory propagation to discrete connections, with guarantees tied to the maintained graph or sample set. This study proposes incremental simplicial trajectory replanning (ISTR*) with Directional Action-front Recursion (DAR) as its search core. DAR propagates arrival costs over directed state–time action domains and retains the realizing controls to construct dynamically feasible trajectories through continuous joining states. Based on the target- and time-indexed corrected-dynamics inferred-control (CDIC) posterior regions, the probabilistic occupancy of non-cooperative objects is modeled. Each posterior update modifies action admissibility and reopens the dependent search records. The theoretical analysis establishes finite-model feasibility and optimality, fixed-instance asymptotic optimality under feasible recovery and vanishing-error refinement, and probabilistic completeness for stochastic refinement. Numerical experiments demonstrate feasible trajectory recovery under updated collision constraints. The primary orbital experiments achieve mission-completion rates of at least 85%; active reoptimization reduces the median final executed cost by 55.9% on tasks completed by both policy configurations.

Keywords: spacecraft relative motion; kinodynamic search; simplicial complex; incremental replanning; posterior reachable region; asymptotic optimality

Check for updates

Academic Editor: Andrea Colagrossi

Received: 5 October 2026
Revised: 25 October 2026
Accepted: 26 October 2026
Published: 29 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Trajectory planning in a static environment seeks a feasible or minimum-cost motion from a prescribed initial state to a terminal set under fixed environmental and system constraints. Classical shortest-path and heuristic graph-search methods operate on explicit connectivity structures. Probabilistic roadmaps and rapidly exploring random trees construct sparse representations from sampled states and local connections instead of exhaustively discretizing the state space [1,2]. PRM* and RRT* establish asymptotic optimality under their respective sampling, connection, and steering assumptions [3]. Randomized kinodynamic planning extends sampling-based construction to systems with differential constraints [4]. Comparative studies further identify representation, sampling, local steering, and collision evaluation as distinct considerations in planner design [5]. These methods provide the principal algorithmic foundations for planning on a fixed problem instance.

Within this sampling-based framework, the sampling policy determines how accurately a sparse substrate represents the planning domain. Halton low-discrepancy sequences improve global dispersion, while adaptive refinement increases resolution in poorly represented or task-relevant regions [6,7]. Informed RRT* restricts sampling to the subset that can improve the current path-length solution [8]. Batch informed trees (BITs*) combine batch sampling with ordered search on progressively denser implicit random geometric graphs [9]. ABIT*, AIT*, and EIT* further introduce inflated and truncated search, adaptive heuristic construction, and effort-aware edge ordering [10–12]. These developments coordinate global coverage, local refinement, and search order on a reusable sparse substrate, improving the allocation of computational effort without constructing a complete discretization.

For systems with differential constraints, the local connection model is as important as the global search strategy. Dynamically feasible local connections allow the search to propagate the costs of admissible trajectories rather than geometric distances alone. Two-point boundary-value steering supports asymptotically optimal kinodynamic planning when the local solver enforces the governing dynamics and control bounds [13]. Robust-RRT incorporates forward reachable sets into sampling-based control synthesis and establishes probabilistic completeness for bounded uncertainty under its nonlinear-system assumptions [14]. Continuous-time Gaussian-process representations and planning-by-inference methods provide sparse trajectory descriptions and joint estimation–planning formulations [15–17]. These approaches motivate the distinction between a geometric connection and an executable trajectory.

Spacecraft relative-motion planning extends these formulations to inspection, formation flight, and maneuvering around exclusion zones. An adaptation of FMT* connects low-dispersion state samples by impulsive transfers and restricts the search to states admitting collision-avoidance maneuvers under specified thruster-failure conditions [18]. A complementary formulation represents periodic natural-motion trajectories as graph nodes and uses safe invariant tubes to determine transitions satisfying thrust and obstacle constraints [19]. Both approaches incorporate orbital dynamics and collision avoidance into the connections available to the planner. A related challenge is to propagate cost through the planning domain while allowing connections within simplex faces. Restricting propagation to graph edges can distort an optimal route. Theta* permits any-angle paths on grids, Field D* interpolates local costs, and simplicial Dijkstra and A* extend interpolation to simplicial decompositions [20–22]. Basic S* computes geometric shortest paths on a metric simplicial complex associated with an undirected graph [23]. Its continuous-predecessor construction improves path representation by allowing connections through simplex faces.

During mission execution, model mismatch, state-estimation uncertainty, and sensor error can invalidate a nominal trajectory. The motion of non-cooperative objects also changes the collision constraints. The planner must update the trajectory as observations become available, maintain kinodynamic feasibility, and reuse computations whose validity conditions remain satisfied.

When the environment or planning task changes, existing incremental planners generally retain search information that remains valid. Focussed D* and D* Lite address such changes by repairing affected shortest-path labels instead of recomputing a complete solution [24,25]. RRTX extends incremental maintenance to a sampling-based planning graph [26], while horizon-based lazy optimal RRT* (HL-RRT*) combines a finite planning horizon with lazy steering and adaptive sampling [27]. Lifelong-GLS (L-GLS) combines incremental repair with lazy edge evaluation to recover an optimal path on the current graph. Bounded L-GLS (B-LGLS) reduces search and evaluation effort under a user-selected suboptimality factor [28]. These methods demonstrate the value of reusable search struc-

tures and selective repair, but the search remains defined on the graph or tree maintained by each method. Optimization-based methods address spacecraft guidance and control through constrained finite-horizon optimization. Fast model predictive control (MPC) exploits quadratic-program structure for rendezvous guidance with moving-obstacle avoidance [29]. The TRACE pipeline combines target identification, offline-library-initialized nonlinear trajectory optimization, and tube-based MPC tracking under parameter uncertainty [30]. Distributed cascaded MPC coordinates multiple servicers through Lyapunov and barrier-function constraints [31]. Flexible-appendage regulation has also been studied using nonlinear model predictive control with nonlinear autoregressive exogenous-input networks for coupled attitude and antenna vibration dynamics [32].

Uncertain path constraints have also been incorporated into probabilistic trajectory optimization for close-proximity operations [33], while sampling-based spacecraft planning has been used to examine how orbital dynamics affect fuel use [34]. These studies address the uncertainty and dynamical cost associated with executing a planned path. In the non-cooperative-object setting considered here, a feasible trajectory prefix must also leave sufficient time and control authority to reach the terminal set. Each planning epoch therefore starts from the current spacecraft state, and a previously computed segment remains available when its dynamics, timing, and collision constraints hold under the updated posterior.

Planning around non-cooperative spacecraft also requires an updateable, set-valued description of their predicted occupancy. Zonotopes support efficient propagation under linear maps and Minkowski sums [35]. Conservative polynomialization encloses nonlinear reachable sets [36], and sampling-based reachability provides asymptotic and finite-sample approximation guarantees [37]. Hamilton–Jacobi reachability has been combined with receding-horizon planning through continuous-time forward reachable sets under bounded disturbances [38]. Posterior occupancy regions, forward disturbance-reachable sets, and backward terminal-reachable outer sets serve distinct roles in prediction, trajectory verification, and terminal screening.

Factor graphs modularly represent dynamics, measurements, priors, and local constraints, whereas incremental smoothing preserves sparse structure through factorization and Bayes-tree updates [39–42]. In our preceding work, a local Markov factor graph jointly estimated the shared chaser and augmented non-cooperative target states, retained observation-induced cross-target covariance, and propagated the target marginal posteriors into target- and time-indexed corrected-dynamics inferred-control (CDIC) probabilistic position regions [43]. The present study represents these regions as CDIC geometric objects that define equivalent infeasible regions and modify action admissibility without contributing to the accumulated trajectory cost.

The reviewed methods motivate a directed kinodynamic search formulation with continuous joining states. This paper presents ISTR* as an incremental replanning framework and DAR as its search core. Inspired by the continuous-predecessor idea of Basic S*, DAR uses a distinct search and reconstruction procedure: arrival-cost functions propagate over directed state–time action domains, and stored control realizations recover dynamically feasible state–control trajectories. ISTR* maintains the action archive and search dependencies across CDIC posterior updates, reconstructing affected domains and reusing valid records. CDIC provides the prediction model from our preceding work. The analysis establishes finite-model feasibility and optimality and, under feasible recovery and vanishing refinement and search errors, fixed-instance asymptotic optimality and probabilistic completeness of stochastic refinement.

The remainder of this paper is organized as follows. Section 2 formulates the planning problem. Section 3 presents ISTR* and DAR, and Section 4 establishes their theoretical

properties. Section 5 presents the simulation settings and comparative numerical analysis. Section 6 summarizes the contributions and directions for future work.

2. Problem Formulation

2.1. Nonlinear Spacecraft Dynamics

The scenario for motion planning consists of the mission spacecraft operating alongside non-cooperative space objects in a common local orbital frame. The spacecraft state and admissible control are defined as follows:

$$\mathbf{x} = (\mathbf{r}^T, \mathbf{v}^T)^T \in \mathcal{X} \subset \mathbb{R}^6, \qquad \mathbf{u} \in \mathcal{U} \subset \mathbb{R}^3, \tag{1}$$

and the governing model is

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t; \boldsymbol{\theta}_d), \tag{2}$$

where $\mathbf{r}$ and $\mathbf{v}$ denote the relative position and velocity, and $\boldsymbol{\theta}_d$ contains the declared orbital and perturbation parameters. The control set $\mathcal{U}$ is compact. The vector field $\mathbf{f}$ is continuous and locally Lipschitz continuous in $\mathbf{x}$, uniformly in $(\mathbf{u}, t)$ on an open neighborhood of the compact planning domain. The flow $\phi_{t,s}(y, u)$ is considered only on intervals where the corresponding solution exists and remains in that domain. These conditions ensure uniqueness and continuous dependence on the initial state. For a piecewise constant control block of duration h, the state propagates as

$$\mathbf{x}_{k+1} = \phi_{t_k, t_k+h}(\mathbf{x}_k, \mathbf{u}_k). \tag{3}$$

2.2. Continuous Trajectory Planning

Planning in a static environment is considered first. Let $\mathcal{X}_{\text{free}} \subset \mathcal{X}$ denote the admissible state domain and $\mathcal{O}^{\text{stat}} \subset \mathbb{R}^3$ the occupied region. The index i identifies a planning instance. Its occupancy input is $\mathcal{O}_i(t) = \mathcal{O}^{\text{stat}}$. Section 2.3 replaces this input with a time-dependent prediction while retaining the same objective.

With fixed planning constraints, planning begins from the specified initial state $\mathbf{x}_s$ at physical time t_s and terminates in $\mathcal{X}_g$ no later than $T_{\max}$. The terminal set $\mathcal{X}_g$ is closed in the compact state domain. The matrix $P_r = [I_3 \; 0_{3\times 3}]$ selects relative position from the six-dimensional state.

Problem 1 (Deadline-constrained spacecraft trajectory)**.** *Among all admissible terminal times t_f and measurable controls $\mathbf{u}(\cdot)$, the optimal value is defined by:*

$$\begin{aligned} J_i^\star(\mathbf{x}_s, t_s) = \inf_{\mathbf{u}(\cdot), t_f} \quad & \int_{t_s}^{t_f} \ell(\mathbf{x}(t), \mathbf{u}(t), t)\, \mathrm{d}t \\ \text{subject to} \quad & \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t; \boldsymbol{\theta}_d), \\ & \mathbf{x}(t_s) = \mathbf{x}_s, \qquad \mathbf{x}(t_f) \in \mathcal{X}_g, \\ & \mathbf{x}(t) \in \mathcal{X}_{\text{free}}, \qquad \mathbf{u}(t) \in \mathcal{U}, \\ & P_r \mathbf{x}(t) \notin \mathcal{O}_i(t), \qquad t \in [t_s, t_f], \\ & t_f \in [t_s, T_{\max}]. \end{aligned} \tag{4}$$

The integral in Problem 1 is denoted by $J[\Pi]$ for maneuver $\Pi = (\mathbf{x}, \mathbf{u}, t_f)$. The continuous nonnegative running cost ℓ is fixed across the compared planners. For control effort, take $\ell = \frac{1}{2}\mathbf{u}^T W_u \mathbf{u}$ with $W_u \succ 0$. The terminal set and deadline impose hard constraints. The value of Problem 1 is $J_i^\star = \inf_\Pi J[\Pi]$, with value $+\infty$ assigned to an empty feasible trajectory set.

DAR is the directed action-front search method within ISTR*. Its search records couple joining states and arrival costs to the controls that realize dynamically admissible trajectory segments.

2.3. Posterior-Driven Extension to Dynamic Environments

For moving non-cooperative objects, a time-indexed prediction replaces the fixed occupied region in Problem 1. At t_i, the updated posterior and propagated spacecraft state define the remaining-horizon problem on $[t_i, T_{\max}]$. Within that planning instance, the posterior is fixed, and its predicted occupied region varies with physical time.

The target regions are supplied by the CDIC method in our preceding work [43]. The planning input for target indices $\mathcal{J}$, query offsets $\mathcal{T}_i$, probability level α, and posterior identifier ν_i is as follows:

$$\mathcal{O}_i^{\text{post}} = \left\{ \left(j, t_i + \tau, \mathcal{E}_{j|i}^{\text{CDIC}}(t_i + \tau; \alpha), \nu_i \right) : j \in \mathcal{J},\ \tau \in \mathcal{T}_i \right\}, \tag{5}$$

where $\mathcal{E}_{j|i}^{\text{CDIC}}(t; \alpha)$ denotes the posterior position component defined by Equation (50) of our preceding work [43]. Under its local Gaussian predictive law, the region has probability content α for one target at one forecast time. The planning model herein employs this component directly, with residual-maneuver support and geometric dilation omitted. As such, the probability interpretation is attached to the selected position component. The full-interval trajectory feasibility below is defined by separation from its time-indexed numerical enclosure.

The discrete predictions must relate to the full-interval constraint in Problem 1. Accordingly, the planning model employs a continuous-time extension of the CDIC prediction over $[t_i, T_{\max}]$ that agrees with the query-time regions. Let $\mathcal{O}_i^{\text{CDIC}}(t; \alpha)$ be the union of the target regions at the same physical time. A numerical enclosure $\widehat{\mathcal{O}}_i(t; \alpha)$ for verifying trajectory separation is required to satisfy the following:

$$\begin{aligned} \mathcal{O}_i^{\text{CDIC}}(t; \alpha) := \bigcup_{j \in \mathcal{J}} \mathcal{E}_{j|i}^{\text{CDIC}}(t; \alpha) \subseteq \widehat{\mathcal{O}}_i(t; \alpha), \\ t \in [t_i, T_{\max}]. \end{aligned} \tag{6}$$

Equation (6) requires inclusion at every time in $[t_i, T_{\max}]$. Exact evaluation permits $\widehat{\mathcal{O}}_i = \mathcal{O}_i^{\text{CDIC}}$. For numerical propagation or interpolation, the enclosure accounts for the approximation error over each action interval while retaining the prescribed CDIC probability level and physical geometry.

The updated problem uses $t_s = t_i$, $\mathbf{x}_s = \widehat{\mathbf{x}}_A(t_i)$ and $\mathcal{O}_i(t) = \widehat{\mathcal{O}}_i(t; \alpha)$. Here A identifies the mission spacecraft, whose state is propagated within the incremental replanning procedure described in Section 3.2. This procedure advances the spacecraft state using the executed control sequence before replanning. The dynamics, running cost, terminal set and deadline remain fixed. The new root and occupancy prediction determine the feasible remaining maneuvers and their costs. The discretization of Section 3.1 applies to this instance with the posterior held fixed.

3. ISTR* Planning Method

ISTR* couples a reusable state–time action archive with DAR search. Section 3.1 develops the action representation, Bellman cost propagation, and trajectory recovery within DAR. Section 3.2 then presents archive construction, posterior-driven updates and the overall incremental replanning procedure.

3.1. DAR: Action Construction, Search and Recovery

3.1.1. Fixed-Endpoint Transfers

DAR draws inspiration from the continuous-predecessor idea of Basic S* [23], while defining its own search and trajectory-reconstruction procedures for P1. Basic S* propagates geometric distance through simplex faces and reconstructs a geometric path. DAR instead searches directed state–time action domains, combining arrival-cost functions with dynamically feasible local maneuver costs through Bellman propagation. Each accepted value retains its predecessor and the realizing control. Trajectory recovery follows these records and propagates the stored controls in physical time, producing a state–control trajectory with matching junction states. Simplex faces parameterize complete joining states; departure time, direction-dependent cost, and admissible controls determine the transfers. The construction proceeds from fixed-endpoint transfers to parameterized action families.

Let $V_{\delta_x} = \{v_1, \ldots, v_{N_{\delta_x}}\} \subset \mathcal{X}_{\text{free}}$ contain complete position–velocity states. A directed connection e joins neighboring states (v_a, v_b). It has departure time t_e and duration $\tau_e > 0$, with $[t_e, t_e + \tau_e] \subseteq [t_s, T_{\max}]$. Section 3.2.1 specifies the neighborhood rule. On a fixed control partition, $\mathbf{U}$ stacks the control blocks and $u_{\mathbf{U}}$ is their time history. The flow $\phi_{t_e,t}(v_a, u_{\mathbf{U}})$ follows P1. The compact domain $\mathcal{P}_e^{\text{safe}}$ contains coefficients satisfying the state, control and collision constraints throughout the interval. P2 adds the endpoint condition v_b.

Problem 2 (Local fixed-endpoint connection)**.** *The feasible control set, its minimum cost, and an attaining control for connection e are defined as:*

$$\begin{aligned} w_e &= \inf_{\mathbf{U}\in\mathcal{U}_e} J_e(\mathbf{U}), \qquad \inf \varnothing := +\infty, \\ J_e(\mathbf{U}) &= \int_{t_e}^{t_e+\tau_e} \ell(\phi_{t_e,t}(v_a, u_{\mathbf{U}}), u_{\mathbf{U}}(t), t)\, \mathrm{d}t, \\ \mathcal{U}_e &= \{\mathbf{U} \in \mathcal{P}_e^{\text{safe}} : \phi_{t_e,t_e+\tau_e}(v_a, u_{\mathbf{U}}) = v_b\}, \\ \mathbf{U}_e^\star &\in \arg\min_{\mathbf{U}\in\mathcal{U}_e} J_e(\mathbf{U}) \qquad (\mathcal{U}_e \neq \varnothing). \end{aligned} \tag{7}$$

P2 fixes the endpoints, duration, and control partition of P1. Compactness and continuity give an attaining control whenever $\mathcal{U}_e$ is nonempty. A verified empty domain gives $w_e = +\infty$. The retained connections form $\mathcal{G}_i^0 = (V_{\delta_x}, E_i^0, \{w_e\}_{e\in E_i^0})$. Each edge stores its realizing control and verified domain.

Since several transfers may connect the same vertices with different departure times or durations, each edge represents a timed maneuver with a specified time interval and control realization.

3.1.2. Simplicial Joining States and Action Families

Vertex transfers provide the initial maneuvers. DAR then allows their joining state to vary over a simplex face. The face parameterizes a complete state; the local control realizes the transfer. For $d+1$ affinely independent vertices in V_{δ_x}, define

$$\sigma = \operatorname{conv}\{v_{j_0}, \ldots, v_{j_d}\} = \left\{ \sum_{k=0}^{d} \lambda_k v_{j_k} \;\middle|\; \lambda_k \geq 0,\ \sum_{k=0}^{d} \lambda_k = 1 \right\}, \qquad 0 \leq d \leq 6. \tag{8}$$

The barycentric weights $\lambda_k = w(v_{j_k})$ follow the vertex-weight representation of Basic S* [23] (Section 3). In the present state space, $\mathbf{x} = \sum_k \lambda_k v_{j_k}$ is a complete position–velocity state.

By setting the selected barycentric coefficients to zero, the state is restricted to a face. A single unit coefficient recovers a vertex. These coordinates specify where two maneuvers meet. The associated control determines the motion between successive connections.

The state simplices are assembled into a conforming finite complex $\mathcal{K}_{\delta_x}$, with $|\mathcal{K}_{\delta_x}| = \bigcup_{\sigma\in\mathcal{K}_{\delta_x}} \sigma \subseteq \mathcal{X}_{\text{free}}$. The spatial resolution is denoted by δ_x, while h denotes a control block duration. To obtain consistent connection faces, the complex needs to satisfy the following:

$$\begin{aligned} &\sigma \in \mathcal{K}_{\delta_x},\ \rho \preceq \sigma \implies \rho \in \mathcal{K}_{\delta_x}, \\ &\sigma \cap \sigma' = \varnothing \quad \vee \quad \big(\sigma \cap \sigma' \preceq \sigma \ \wedge\ \sigma \cap \sigma' \preceq \sigma'\big). \end{aligned} \tag{9}$$

where $\rho \preceq \sigma$ denotes a face of σ. The participating simplices $\mathscr{C} = \{C_a : C_a \in \mathcal{K}_{\delta_x}\}$ define the geometric domains of local transfers and are termed carrier simplices. Conforming subdivision includes the specified initial state as a marked face and represents the selected terminal faces in $\mathcal{X}_g$. Thus, adjacent carriers meet on the connection face:

$$\mathcal{I}_{ab} := C_a \cap C_b. \tag{10}$$

Let $\mathscr{I}_i$ contain the nonempty connection faces, including the initial and terminal faces. Each face parameterizes joining states. Consecutive maneuvers share the same complete state and physical time. Their trajectories remain in the union of the participating carriers. Restriction to a common subface preserves state, arrival time, and cost. Representations with identical admissible continuations are identified.

An action family κ fixes the departure time t_κ, duration τ_κ and control partition. Choose $p+1$ affinely independent endpoint pairs $(y_j, q_j) \in \mathbb{R}^6 \times \mathbb{R}^6$, $j = 0, \ldots, p$. The weights $\lambda \in \mathbb{R}^{p+1}$ locate both endpoints. The variable $\zeta \in \mathbb{R}^{d_\zeta}$ contains the remaining control coordinates. On the compact verified domain Θ_κ, $\theta = (\lambda, \zeta)$ determines the endpoint maps, control coefficients and trajectory:

$$\begin{aligned} y_\kappa(\theta) &:= \sum_{j=0}^{p} \lambda_j y_j, & q_\kappa(\theta) &:= \sum_{j=0}^{p} \lambda_j q_j, \\ \Delta_p &= \{\lambda \geq 0 : \mathbf{1}^T\lambda = 1\}, & \theta &= (\lambda,\zeta) \in \Theta_\kappa \subseteq \Delta_p \times \mathbb{R}^{d_\zeta}, \\ X_\kappa(s;\theta) &= \phi_{t_\kappa,s}(y_\kappa(\theta), u_{\mathbf{U}_\kappa(\theta)}), & s &\in [t_\kappa, t_\kappa + \tau_\kappa], \\ X_\kappa(t_\kappa + \tau_\kappa;\theta) &= q_\kappa(\theta), \end{aligned} \tag{11}$$

Here, Δ_p is the standard simplex, $\mathbf{1}$ is the all-ones vector, and ζ is omitted when the endpoints determine the controls. The control map $\mathbf{U}_\kappa$ is continuous on each compact parameter domain; piecewise-continuous constructions are represented by separate continuous branches. Continuous dependence of the flow then makes J_κ continuous. Every retained parameter satisfies the endpoint and full-interval conditions in Equation (15).

The realizing state and control determine the action cost, namely:

$$J_\kappa(\theta) = \int_{t_\kappa}^{t_\kappa+\tau_\kappa} \ell\big(X_\kappa(s;\theta), u_{\mathbf{U}_\kappa(\theta)}(s), s\big)\, \mathrm{d}s. \tag{12}$$

For quadratic control effort, $J_\kappa(\theta) = \frac{1}{2}\mathbf{U}_\kappa(\theta)^T H_u \mathbf{U}_\kappa(\theta)$, where H_u is the block-diagonal matrix of control weights multiplied by the respective interval durations. For an affine construction $\mathbf{U}_\kappa(\lambda) = \sum_{j=0}^{p} \lambda_j \mathbf{U}_j$, each $\mathbf{U}_j$ is the control coefficient vector assigned to endpoint pair (y_j, q_j). The parameter cost is quadratic, $J_\kappa(\lambda) = \frac{1}{2}\lambda^T K_\kappa \lambda$, with $(K_\kappa)_{jk} = \mathbf{U}_j^T H_u \mathbf{U}_k$. These cross terms are not determined by the scalar vertex costs alone. General control maps are based on Equation (12).

The directed graph first specifies admissible state–time connections. Problem 2 then supplies controls realizing the endpoint pairs of each action family. For fixed departure time, duration, and control partition, an affine endpoint model has the form $q = \mathsf{F}_\kappa y + \mathsf{G}_\kappa \mathbf{U} + c_\kappa$, where the matrices and offset follow from the selected dynamics. Each retained vertex transfer satisfies $q_j = \mathsf{F}_\kappa y_j + \mathsf{G}_\kappa \mathbf{U}_j + c_\kappa$. Multiplying these identities by common barycentric weights and summing gives

$$q(\lambda) = \mathsf{F}_\kappa y(\lambda) + \mathsf{G}_\kappa \sum_{j=0}^{p} \lambda_j \mathbf{U}_j + c_\kappa, \qquad \sum_{j=0}^{p} \lambda_j = 1.$$

The control family $\mathbf{U}_\kappa(\lambda) = \sum_j \lambda_j \mathbf{U}_j$ therefore realizes the corresponding interpolated endpoints. Eliminating λ_0 yields the affine map $\mathbf{U}_\kappa(\lambda) = \mathbf{U}_0 + \sum_{j=1}^{p} \lambda_j (\mathbf{U}_j - \mathbf{U}_0)$. Each selected control remains piecewise constant in physical time. The retained domain verifies the full-interval constraints in Equation (15). Equation (13) minimizes cost over the retained controls for each endpoint pair. Nonlinear endpoint relations use the verified parameter domains in Equation (11).

For fixed endpoints (y, q), $\Theta_\kappa(y, q)$ contains the parameters realizing that pair. Let D_κ collect the pairs with nonempty parameter domains. Minimizing J_κ on each domain defines $A_\kappa(y, q)$. Compactness and continuity give a minimizing parameter θ_κ^* and control u_κ. The action record $\mathfrak{a}_\kappa$ stores these data. Its certificate χ_κ records the verified endpoint and full-interval constraints:

$$\begin{aligned} \Theta_\kappa(y, q) &= \{\theta \in \Theta_\kappa : y_\kappa(\theta) = y,\ q_\kappa(\theta) = q\}, \\ D_\kappa &= \{(y, q) : \Theta_\kappa(y, q) \neq \varnothing\}, \\ A_\kappa(y, q) &= \min_{\theta \in \Theta_\kappa(y,q)} J_\kappa(\theta), \\ \theta_\kappa^*(y, q) &\in \arg\min_{\theta \in \Theta_\kappa(y,q)} J_\kappa(\theta), \\ u_\kappa(\cdot; y, q) &= u_{\mathbf{U}_\kappa(\theta_\kappa^*(y,q))}(\cdot), \qquad \mathfrak{a}_\kappa = (\tau_\kappa, D_\kappa, u_\kappa, A_\kappa, \chi_\kappa). \end{aligned} \tag{13}$$

Each action is associated with entry and exit faces $\Gamma_\kappa^{\mathrm{in}}$ and $\Gamma_\kappa^{\mathrm{out}}$ and a closed carrier union $\mathcal{C}_\kappa$. Let $\mathbb{T}_i \subseteq [t_s, T_{\max}]$ be the finite set of represented departure and arrival times and A be the finite set of access labels. The labels $a_\kappa^{\mathrm{in}}, a_\kappa^{\mathrm{out}} \in \mathsf{A}$ distinguish permitted continuations not already determined by the state, time, and face. Otherwise, A is a singleton. The incidence conditions are as follows:

$$\begin{aligned} D_\kappa &\subseteq \Gamma_\kappa^{\mathrm{in}} \times \Gamma_\kappa^{\mathrm{out}}, \qquad \Gamma_\kappa^{\mathrm{in}}, \Gamma_\kappa^{\mathrm{out}} \in \mathscr{I}_i, \\ \Gamma_\kappa^{\mathrm{in}} \cup \Gamma_\kappa^{\mathrm{out}} &\subseteq \mathcal{C}_\kappa \subseteq |\mathcal{K}_{\delta_x}|, \quad t_\kappa, t_\kappa + \tau_\kappa \in \mathbb{T}_i. \end{aligned} \tag{14}$$

For each $(y, q) \in D_\kappa$, the selected parameter supplies a control satisfying the verified conditions throughout the maneuver. With $t = t_\kappa$ and u_κ in the flow map denoting $u_\kappa(\cdot; y, q)$, these conditions can be expressed as follows:

$$\begin{aligned} &\phi_{t,t+\tau_\kappa}(y, u_\kappa) = q, \\ &u_\kappa(s; y, q) \in \mathcal{U}, \quad \phi_{t,s}(y, u_\kappa) \in \mathcal{C}_\kappa, \\ &P_r \phi_{t,s}(y, u_\kappa) \notin \mathcal{O}_i(s), \quad s \in [t, t + \tau_\kappa]. \end{aligned} \tag{15}$$

By restricting the endpoints to a closed subface, the admissible controls and costs of each retained endpoint pair remain unchanged. Equation (12) provides the parameter-dependent cost, while Equation (13) defines its minimum for fixed endpoints.

Due to the compactness of the parameter relation and the continuity of the parameter cost, A_κ is lower-semicontinuous on D_κ, as established by the argument in Theorem 1. Let $\mathscr{A}_i$ index the finite collection of retained vertex transfers and parameterized local maneuver families. The directed cost graph has node set $\mathscr{V}_i$ and edge set $\mathscr{E}_i$. Each node records a connection face, time, and access label, with z_κ^- and z_κ^+ being the departure and arrival nodes of action κ:

$$\begin{aligned} \mathscr{G}_i^{\mathrm{cost}} &= (\mathscr{V}_i, \mathscr{E}_i, \{A_\kappa\}_{\kappa\in\mathscr{A}_i}), \\ \mathscr{V}_i &= \mathscr{I}_i \times \mathbb{T}_i \times \mathsf{A}, \\ z_\kappa^- &= (\Gamma_\kappa^{\mathrm{in}}, t_\kappa, a_\kappa^{\mathrm{in}}), \\ z_\kappa^+ &= (\Gamma_\kappa^{\mathrm{out}}, t_\kappa + \tau_\kappa, a_\kappa^{\mathrm{out}}), \\ \mathscr{E}_i &= \{(z_\kappa^-, \kappa, z_\kappa^+) : \kappa \in \mathscr{A}_i\}. \end{aligned} \tag{16}$$

The discrete model inherits $(\mathbf{x}_s, t_s)$ and $\mathcal{X}_g$ from P1, with $\mathbf{x}_s \in \mathcal{X}_{\mathrm{free}}$ outside the occupied region at t_s. The empty maneuver has zero duration and cost.

Retaining the same vertex transfers under identical carrier, control, collision, and timing constraints gives the comparison with vertex-restricted search in Section 4.

3.1.3. Action Chains and Finite-Model Objective

The declared finite model $\mathscr{M}_i$ comprises the carrier complex, action families, and represented times. Its objective is based on the complete admissible domains and remains fixed during a DAR call. The search progress is stored separately in $\mathscr{Z}_i^{(r)}$, as defined in Equation (22).

The principal symbols are collected in the nomenclature preceding the references.

In each declared family, $\Theta_\kappa^{\mathrm{adm}}$ denotes the complete compact admissible parameter domain. The superscript adm distinguishes the full domain from the verified subset Θ_κ. Applying Equation (13) to the full domain yields D_κ^{adm}, A_κ^{adm}, and u_κ^{adm}. Let $\Sigma_{\mathscr{M}_i}$ contain the complete feasible action sequences. A sequence $\varsigma = (\kappa_0, \ldots, \kappa_{m-1})$ has m actions whose exit and entry nodes match in face, time, and access label. It satisfies $t_0 = t_s$, $t_k = t_{\kappa_k}$, $t_{k+1} = t_k + \tau_{\kappa_k}$, and $t_m \le T_{\max}$. Here, $x_k \in \mathcal{X}$ is the full state at the k-th junction. The admissible state sequences are as follows:

$$\mathcal{D}_\varsigma = \left\{ (x_0, \ldots, x_m) \in \mathcal{X}^{m+1} \;\middle|\; \begin{array}{ll} x_0 = \mathbf{x}_s, & x_m \in \mathcal{X}_g, \\ (x_k, x_{k+1}) \in D_{\kappa_k}^{\mathrm{adm}}, & \forall k \in \{0, \ldots, m-1\} \end{array} \right\}. \tag{17}$$

For a fixed action sequence, the junction states determine the associated local controls. Their concatenation defines the maneuver $\Pi = (\mathbf{x}_\Pi, \mathbf{u}_\Pi)$ and its spatial path $\mathbf{r}_\Pi$:

$$\begin{aligned} \mathbf{x}_\Pi(t) &= \phi_{t_k,t}\big(x_k, u_{\kappa_k}^{\mathrm{adm}}(\cdot; x_k, x_{k+1})\big), \\ \mathbf{u}_\Pi(t) &= u_{\kappa_k}^{\mathrm{adm}}(t; x_k, x_{k+1}), \qquad t \in [t_k, t_{k+1}), \quad 0 \le k < m, \\ \mathbf{x}_\Pi(t_m) &= x_m, \\ \mathbf{r}_\Pi(t) &= P_r \mathbf{x}_\Pi(t). \end{aligned} \tag{18}$$

By matching full states and times, the trajectory is continuous across junctions. The full-interval action conditions enforce the dynamics and path constraints of P1. The additivity leads to $J[\Pi] = \sum_{k=0}^{m-1} A_{\kappa_k}^{\mathrm{adm}}(x_k, x_{k+1})$, so the action chain cost equals the actual maneuver cost.

For a fixed chain, $\mathcal{D}_\varsigma$ is a closed subset of a finite product of compact action domains. Minimizing over its junction states yields the following:

$$J_\varsigma = \begin{cases} \min\limits_{(x_0,\ldots,x_m)\in\mathcal{D}_\varsigma} \sum\limits_{k=0}^{m-1} A^{\mathrm{adm}}_{\kappa_k}(x_k, x_{k+1}), & \mathcal{D}_\varsigma \neq \varnothing, \\ +\infty, & \mathcal{D}_\varsigma = \varnothing. \end{cases} \tag{19}$$

The cost of the zero-action chain is zero when $\mathbf{x}_s \in \mathcal{X}_g$. Let $h_{\min} > 0$ be a lower bound on the durations of all declared actions in the finite model. Then $\tau_\kappa \geq h_{\min}$ limits the number of actions to $\lfloor (T_{\max} - t_s)/h_{\min} \rfloor$. The finite action collection consequently produces finitely many sequences, each retaining continuous junction variables. The model optimum is as follows:

$$J^\star_{\mathscr{M}_i} = \min_{\varsigma\in\Sigma_{\mathscr{M}_i}} J_\varsigma. \tag{20}$$

Equations (19) and (20) optimize junction states and action sequences. Feasible inclusion in P1 gives the following:

$$J^\star_i \leq J^\star_{\mathscr{M}_i}. \tag{21}$$

DAR evaluates the discretized maneuver family by the arrival-cost propagation in Section 3.1.6. Section 4 establishes the fixed-model optimality and convergence to the optimum of P1 along a feasible refinement sequence.

3.1.4. Search State and Initialization

DAR propagates arrival-cost (cost-to-come) functions over the constructed directed state–time actions. Each propagated value retains its predecessor and control realization for trajectory recovery.

The Bellman update minimizes accumulated cost over feasible predecessor states and control parameters, then forms the exit-value lower envelope. Domain records preserve unresolved alternatives and their dependencies. The four operations below separate cost propagation from search-record maintenance.

For a fixed planning instance, the inputs are a parent arrival-cost function and a compatible action with its control map. The outputs are an exit-value function and the predecessor and control maps required for trajectory recovery.

1. Restrict the parent and action domains to compatible states, times, and interfaces. Apply the terminal reachability screen to the exit domain.
2. Minimize the sum of parent arrival cost and action cost over the joining state and admissible control parameters. Use active-face reconstruction for polyhedral–quadratic families or certified domain refinement for nonlinear families.
3. Store each candidate value with its predecessor and control maps. Form the pointwise lower envelope of compatible candidates.
4. Retain unresolved domains and their lower bounds. Insert or reopen dependent successor and terminal records when a parent value improves.

Steps 1–3 perform Bellman cost propagation; Step 4 maintains unresolved search records. OPEN schedules record processing, and ISTR* updates action validity after each posterior update. Algorithm 1 specifies the complete search.

Algorithm 1 DAR on a fixed epoch model

Input: root fronts $\mathcal{B}^{\mathrm{root}}$, instance ι, carriers $\mathscr{C}$, actions $\mathscr{A}$, fronts $\mathscr{B}$, unresolved records $\mathscr{Q}$, dependencies $\mathscr{G}_D$, verified seed Π^{seed} or $\varnothing$, $(\mathcal{X}_g, T_{\max})$, budget B, tolerance ε, fair schedule $\prec_F$
Output: $(\Pi^\star, L, U, \mathscr{A}, \mathscr{B}, \mathscr{Q}, \mathscr{G}_D)$

// Compare compatible arrival costs (Equation (43)).
$\mathscr{B} \leftarrow \mathrm{Env}_{\iota,a,t}(\mathscr{B} \cup \mathcal{B}^{\mathrm{root}})$, $\mathrm{Pred}(\mathcal{B}^{\mathrm{root}}) \leftarrow \varnothing$

// First transfers and root terminal restrictions (Section 3.1.4).
$\mathscr{Q} \leftarrow \mathscr{Q} \cup \{R_{b\kappa} : \mathcal{B}_b \in \mathcal{B}^{\mathrm{root}},\ \kappa \in \mathscr{F},\ D_b \cap \Gamma^{\mathrm{in}}_\kappa \neq \varnothing,\ t_b = t_\kappa,\ a_b = a^{\mathrm{in}}_\kappa,\ t_b + \tau_\kappa \le T_{\max}\} \cup \{R^g_b : \mathcal{B}_b \in \mathcal{B}^{\mathrm{root}},\ D_b \cap \mathcal{X}_g \neq \varnothing\}$
$\mathrm{OPEN} \leftarrow \mathscr{Q}$, $\mathscr{T}_{\mathrm{rec}} \leftarrow \varnothing$, $(\Pi^\star, U) \leftarrow (\varnothing, +\infty)$
if $\exists \mathcal{B}_b \in \mathcal{B}^{\mathrm{root}} :\ D_b = \{q^g_0\} \subseteq \mathcal{X}_g$ **then**
  $(\Pi^\star, U) \leftarrow (\Pi_b(q^g_0), 0)$
end if
if $\Pi^{\mathrm{seed}} \neq \varnothing$ and $J[\Pi^{\mathrm{seed}}] < U$ **then**
  $(\Pi^\star, U) \leftarrow (\Pi^{\mathrm{seed}}, J[\Pi^{\mathrm{seed}}])$
end if
$L \leftarrow \min\{U, \min_{R \in \mathrm{OPEN}} J(R)\}$
while $B > 0$ and $\mathrm{OPEN} \neq \varnothing$ **do**
  $R \leftarrow \mathrm{select}_{\prec_F}(\mathrm{OPEN})$
  $w_R \leftarrow 0$
  if $R \in \mathscr{Q}^g$ **then**

    // Terminal restriction and remainder with bounds, cursors, and dependencies (Section 3.1.8).
    $(\mathscr{B}, \mathscr{T}_{\mathrm{rec}}, \mathcal{W}_R, \mathscr{G}_D, w_R) \leftarrow \textsc{ResolveTerminal}(R, \mathscr{B}, \mathscr{T}_{\mathrm{rec}}, \mathcal{X}_g, T_{\max}, B - w_R, \mathscr{G}_D)$
    $(\mathrm{OPEN}, \mathscr{G}_D) \leftarrow \mathrm{Rep}(\mathrm{OPEN}, \mathscr{G}_D; R, \mathcal{W}_R)$
  else

    // Compatible families and exit screening (Section 3.1.5).
    $\mathfrak{M}_R \leftarrow \{(b_R, \kappa, \tau_\kappa, \Gamma^{\mathrm{in}}_\kappa, \Gamma^{\mathrm{out}}_\kappa) : \kappa \in \mathscr{F},\ D_R \cap \Gamma^{\mathrm{in}}_\kappa \neq \varnothing,\ t_R = t_\kappa,\ a_R = a^{\mathrm{in}}_\kappa,\ C_R \sim C_\kappa\}$
    if $\mathfrak{M}_R = \varnothing \ \wedge\ c^{\mathrm{fam}}_R = 1$ **then**

      // Close only the verified empty component and retain its exclusion premises.
      $(\mathrm{OPEN}, \mathscr{G}_D) \leftarrow (\mathrm{OPEN} \setminus \{R\},\ \mathscr{G}_D \cup \mathrm{DepEdges}(\chi^{\mathrm{none}}_R))$
    end if
    for all $\omega = (b, \kappa, \tau, \Gamma^{\mathrm{in}}, \Gamma^{\mathrm{out}}) \in \mathfrak{M}_R$ **do**
      $D_\omega \leftarrow \begin{cases} \varnothing, & t_b + \tau > T_{\max}, \\ \mathcal{D}_{b\kappa}, & t_b + \tau \le T_{\max}. \end{cases}$
      if $t_b + \tau > T_{\max} \ \vee\ D_\omega \overset{\mathrm{v}}{=} \varnothing$ **then**
        $(\mathrm{OPEN}, \mathscr{G}_D) \leftarrow (\mathrm{OPEN} \setminus \{R|_\omega\},\ \mathscr{G}_D \cup \mathrm{DepEdges}(\chi^{\mathrm{empty}}_\omega))$
        continue
      end if
      $(\mathscr{P}_\omega, \mathcal{W}_\omega) \leftarrow (\varnothing, \varnothing)$
      if $\omega \in \mathfrak{M}_{\mathrm{PQ}}$ **then**
        form $K_{b\kappa}$ and $\Psi_{b\kappa}$ by Equations (33) and (34)
        for all $I \in \mathcal{I}_{\mathrm{act}}(R|_\omega, B - w_R)$ **do**
          $\mathcal{R}_I \leftarrow \{(\xi, \eta) \in K_{b\kappa} : A_I \xi = b_I + B_I \eta,\ N_I^T(H_{\xi\xi}\xi + H_{\xi\eta}\eta + h_\xi) = 0\}$
          $Q_I \leftarrow \mathrm{proj}_\eta\, \mathcal{R}_I$
          for all $\Delta = \mathrm{conv}\{\eta_j\}_{j=0}^r \in \mathrm{Tri}(Q_I)$ with $(I, \Delta) \notin \mathrm{Done}(R|_\omega)$ **do**
            $(\xi_j, \eta_j) \in \mathcal{R}_I, \quad j = 0, \dots, r$
            $\xi_{I,\Delta}(\eta) \leftarrow \sum_{j=0}^r \mu_j \xi_j$, with $\eta = \sum_{j=0}^r \mu_j \eta_j$
            $V_{I,\Delta} \leftarrow \Psi_{b\kappa}(\xi_{I,\Delta}, \cdot)$
            $\mathscr{P}_\omega \leftarrow \mathscr{P}_\omega \cup \{(\Delta, V_{I,\Delta}, \xi_{I,\Delta}, \kappa, \chi_\kappa)\}$
            $\mathrm{Done}(R|_\omega) \leftarrow \mathrm{Done}(R|_\omega) \cup \{(I, \Delta)\}$, $w_R \leftarrow w_R + 1$
          end for
        end for

Algorithm 1 *Cont.*

$\mathcal{E}_\omega^{\mathrm{res}} \leftarrow \{\eta \in \mathrm{proj}_\eta\, K_{b\kappa} : \mathcal{J}_\omega(\eta) \subseteq \mathrm{Done}(R|_\omega)\}$
$\mathcal{W}_\omega \leftarrow \mathrm{proj}_\eta\, K_{b\kappa} \setminus \mathcal{E}_\omega^{\mathrm{res}}$
else
// Validated nonlinear domain bounds (Equation (45)).
$(\mathscr{P}_\omega, \mathcal{W}_\omega, \Delta w_\omega) \leftarrow \textsc{CertifiedDomainSolve}(R|_\omega, V_b, \kappa, D_\omega, B - w_R)$
$w_R \leftarrow w_R + \Delta w_\omega$
end if
// Store each candidate's realizing control and immutable predecessor.
$\mathscr{A} \leftarrow \mathscr{A} \cup \{\kappa|_{\theta_P(\mathrm{dom}\,P)} : P \in \mathscr{P}_\omega\}$
for all $P \in \mathscr{P}_\omega$ **do**
$\mathrm{Pred}(P) \leftarrow (b, \kappa, y_P, \theta_P)$
$(\mathscr{B}, \mathscr{G}_D) \leftarrow (\mathrm{Env}_{\iota,a,t}(\mathscr{B} \cup \{P\}),\ \mathscr{G}_D \cup \mathrm{DepEdges}(P, \mathrm{Pred}(P)))$
end for
// Carry bounds, cursors, and dependencies and keep unprocessed siblings (Equation (42)).
$(\mathrm{OPEN}, \mathscr{G}_D) \leftarrow \mathrm{Rep}(\mathrm{OPEN}, \mathscr{G}_D; R|_\omega, \mathcal{W}_\omega)$
end for
end if
// Insert successors and reopen improved-parent records (Section 3.1.7).
$\mathrm{OPEN} \leftarrow \mathrm{OPEN} \cup \bigcup_{P \in \Delta\mathscr{B}} [\mathcal{S}(P) \cup \mathcal{T}(P)]$
// Terminal selection and stored-control recovery (Equation (49)).
$(b_g, q_g) \leftarrow \begin{cases} \min_{\prec_g}\left(\arg\min_{(b,q)\in\mathscr{T}_{\mathrm{rec}}} V_b(q)\right), & \mathscr{T}_{\mathrm{rec}} \neq \varnothing, \\ (\varnothing, \varnothing), & \mathscr{T}_{\mathrm{rec}} = \varnothing, \end{cases}$
if $b_g \neq \varnothing$ **then**
$\bar{\Pi} \leftarrow \bigoplus_{j=1}^{m_g} \mathrm{Seg}(q_{j-1}, \mathbf{U}_{\kappa_j}(\theta_j)), \quad \mathrm{Pred}_{b_j}(q_j) = (b_{j-1}, \kappa_j, q_{j-1}, \theta_j)$
$(\Pi^\star, U) \leftarrow (\bar{\Pi}, J[\bar{\Pi}])$ if $J[\bar{\Pi}] < U$
end if
$L \leftarrow \min\{U, \min_{R\in\mathrm{OPEN}} \underline{J}(R)\}$
$B \leftarrow B - w_R$
if $U < +\infty$ and $U - L \leq \varepsilon$ **then**
break
end if
end while
$\mathscr{Q} \leftarrow \mathrm{OPEN}$
// Keep dependencies of retained records and referenced histories (Section 3.1.8).
$\mathscr{G}_D \leftarrow \mathscr{G}_D[\mathrm{Anc}_{\mathscr{G}_D}(\mathrm{Ref}(\mathscr{A}, \mathscr{B}, \mathscr{Q}, \mathrm{Pred}))]$
return $(\Pi^\star, L, U, \mathscr{A}, \mathscr{B}, \mathscr{Q}, \mathscr{G}_D)$

Predecessor selection jointly determines a joining state and an admissible control realization. Directed action relations enforce state–time compatibility and dynamical feasibility between successive fronts.

A local action connects a source state to an exit state, both comprising the position and velocity. The source selected within a connection face is the continuous predecessor of the exit. An arrival-cost front records the costs and associated trajectories and control sequences of the states on its domain. These histories supply subsequent actions and the final trajectory recovery.

Each arrival-cost front is indexed by physical time, access class, and planning instance. Each retained value is accompanied by a realizing predecessor and control map. Unresolved search records retain competing transfers for subsequent evaluation.

The planning-instance identifier ι specifies the posterior, propagated root, terminal condition, carrier geometry, and action families of $\mathscr{M}_i$. The identifier ι is distinct from the local free-control coordinates ζ and the posterior version ν_i. Verification within a DAR

call resolves the subsets of this fixed model. Seed chains are based on the same model. Upon changes in these defining data, control parameterizations, or refinement level, a new instance starts.

For the fixed declared model $\mathscr{M}_i$, the search state after processing step r is as follows:

$$\mathscr{Z}_i^{(r)} = (\mathcal{K}_{\delta_x,i}, \mathbb{T}_i, \mathscr{A}_i^{(r)}, \mathscr{B}_i^{(r)}, \mathscr{Q}_i^{(r)}, \mathscr{G}_{D,i}^{(r)}), \tag{22}$$

where $\mathscr{A}_i^{(r)}$, $\mathscr{B}_i^{(r)}$, $\mathscr{Q}_i^{(r)}$, and $\mathscr{G}_{D,i}^{(r)}$ represent the current verified actions, fronts, unresolved records, and dependencies, and r is omitted in the algorithms. Verification enlarges $\Theta_\kappa \subseteq \Theta_\kappa^{\mathrm{adm}}$, yielding $D_\kappa \subseteq D_\kappa^{\mathrm{adm}}$ and $A_\kappa \geq A_\kappa^{\mathrm{adm}}$ on the verified domain. Resolving all admissible controls for each endpoint pair yields equality. These updates alter $\mathscr{Z}_i^{(r)}$ while preserving $J^\star_{\mathscr{M}_i}$.

A search node identifies a connection face, physical time, and access class. Its arrival-cost front is represented by the following:

$$\mathcal{B}_b = (D_b, t_b, a_b, V_b, \Pi_b), \qquad V_b(y) = J[\Pi_b(y)], \quad y \in D_b. \tag{23}$$

At a fixed posterior epoch, Algorithm 1 uses carriers $\mathscr{C}$, actions $\mathscr{A}$, cost fronts $\mathscr{B}$, root fronts $\mathcal{B}^{\mathrm{root}}$, unresolved records $\mathscr{Q}$ and dependency graph $\mathscr{G}_D$. Each scope records a transfer or terminal subproblem, domain, parent history, time, epoch, lower bound, and cursor. The search uses feasible incumbent cost U, lower bound L, work budget B, and stopping tolerance ε.

The family set $\mathscr{F}$ includes verified and unresolved declared action families; $\mathscr{Q}^g$ contains terminal-domain records, and $c_R^{\mathrm{fam}} = 1$ marks exhaustive family enumeration.

At the current admissible state $(\mathbf{x}_s, t_s)$, each allowed initial access label defines a singleton root front with zero arrival cost and a zero-duration history. The root collection is denoted by $\mathcal{B}_i^{\mathrm{root}}$.

Algorithm 1, line 2, creates first-transfer records by matching root domains to family entries, departure times, access classes, and the deadline. Intersections with $\mathcal{X}_g$ create terminal records. Each record stores its domain, epoch, bound, cursor and dependencies. ResolveTerminal evaluates a front's intersection with $\mathcal{X}_g$ at times up to $T_{\max}$, adds feasible pairs to $\mathscr{T}_{\mathrm{rec}}$, and returns updated fronts $\mathscr{B}$, unresolved records $\mathcal{W}_R$, dependencies and work count w_R.

The root-front insertion and first-transfer records initialize the search (Algorithm 1, lines 1–10). The operator $\mathrm{Env}_{\iota,a,t}$ evaluates Equation (43) only among records with the same instance, access class, and time, and each root receives an empty predecessor. A terminal root supplies a zero-cost incumbent. Otherwise, the incumbent becomes the seed verified by Rebase or remains empty.

3.1.5. Successor Selection and Terminal Screening

The schedule $\prec_F$ selects the OPEN records fairly if each persistent unresolved record is revisited after finitely many selections, including after successor insertion and reopening. Equation (29) supplies D_ω, following the application of the deadline and terminal reachability restrictions (Algorithm 1, lines 18–27). The unresolved subsets remain as candidates. Exhaustive empty family enumeration produces χ_R^{none}, and a verified empty relation produces $\chi_\omega^{\mathrm{empty}}$. Their premises are stored for reopening after prediction updates.

In a selected scope R, line 18 constructs $\mathfrak{M}_R$ based on the declared families $\mathscr{F}$. Its four predicates impose a nonempty entry-domain intersection, matching departure time and access class, and compatible carrier incidence. Thus, $\mathfrak{M}_R$ includes unresolved declared families as well as verified realizations in $\mathscr{A}$. Restrictions to common subfaces identify

compatible representations of the same source state. An empty result closes the scope only after examining all such families. Unvisited families are still recorded for a later visit.

The terminal set in P1 can describe either the arrival at a prescribed position with free velocity or the arrival at a prescribed complete state. Let $\mathcal{R}_g \subset \mathbb{R}^3$ be a closed target position set, and let $P_v = [0_{3\times 3}\ I_3]$ select the velocity. The prescribed terminal velocity is $\mathbf{v}_g \in \mathbb{R}^3$. The two terminal conditions are as follows:

$$\begin{aligned} \mathcal{X}_g^{\text{free}} &= \{x \in \mathcal{X} : P_r x \in \mathcal{R}_g\}, \\ \mathcal{X}_g^{\text{fix}} &= \{x \in \mathcal{X}_g^{\text{free}} : P_v x = \mathbf{v}_g\} \subseteq \mathcal{X}_g^{\text{free}}. \end{aligned} \tag{24}$$

The selected terminal condition is fixed within each planning instance and comparison. The free arrival velocity is implemented by intersecting the full-state front with $\mathcal{X}_g^{\text{free}}$ and retaining P1's state, control, and collision constraints. The internal joins match the complete position–velocity states. Arrivals with different velocities retain their respective histories and costs.

Terminal screening prevents locally feasible actions from consuming the remaining time without leading to a possible completion. Let $r = T_{\max} - t$. The obstacle and carrier constraints are removed while retaining the same nonlinear flow and a control family containing all admissible planning controls. Any replacement dynamics used to compute an outer approximation must be accompanied by a proved inclusion of the original reachable set. The relaxed predecessor set for a duration τ is as follows:

$$K_{t,\tau} = \{y : \exists u \in \mathcal{U}_{\text{rel}},\ \phi_{t,t+\tau}(y,u) \in \mathcal{X}_g\}, \qquad K_{t,\le r} = \bigcup_{0\le\tau\le r} K_{t,\tau}. \tag{25}$$

Herein, $\mathcal{U}_{\text{rel}}$ contains every admissible control of the planning problem. A computable outer set $\widehat{K}_{t,\le r} \supseteq K_{t,\le r}$ screens the candidate terminal completions. For $0 < \tau \le r$, let D_{out} denote the unscreened exit-state domain of the action. The exit must still admit a possible terminal completion, so its search domain is restricted to the following:

$$D_{\text{out}}^{\text{adm}} = D_{\text{out}} \cap \widehat{K}_{t+\tau,\le r-\tau}. \tag{26}$$

Actions with $\tau > r$ are excluded. Equation (26) retains the exits intersecting the terminal reachable outer set for subsequent action verification and search, while preserving the running cost objective.

For nonlinear dynamics, suppose a validated reachability calculation provides the following:

$$\sup_{u\in\mathcal{U}_{\text{rel}}} \mathbf{n}^T \phi_{t,t+\tau}(y,u) \le b_{t,\tau}(\mathbf{n},y), \qquad \mathbf{n} \in \mathbb{R}^6. \tag{27}$$

For a source domain D, the support gap is defined as follows:

$$m_{t,\tau}(D,\mathbf{n}) = \inf_{g\in\mathcal{X}_g} \mathbf{n}^T g - \sup_{y\in D} b_{t,\tau}(\mathbf{n},y) \tag{28}$$

A positive validated lower bound on $m_{t,\tau}$ excludes the completion at duration τ. Set propagation or linearization with a bounded remainder supplies the support bound [36,38]. Excluding the continuous arrival time union requires the coverage of every admissible duration, including bounds between query times. Domains awaiting that coverage remain UNCOVERED. In a discrete time model, the test covers all admissible arrival ticks.

Terminal screening and cost propagation are connected by retaining the endpoint pairs whose source lies in the parent front and entry face and whose exit lies in the exit face and passes the terminal test. The verified action supplies the dynamics relating these states. Together, these restrictions define the source–exit relation below.

Let $V_b : D_b \to \overline{\mathbb{R}}$ be a parent front compatible with action κ in the sense of Section 3.1. Let $\Gamma^{\mathrm{in}} = \Gamma^{\mathrm{in}}_\kappa$, $\Gamma^{\mathrm{out}} = \Gamma^{\mathrm{out}}_\kappa$, and $t_b = t_\kappa$. These interfaces and the action-indexed edge are those in Equations (14) and (16). The verified action relation D_κ contains endpoint pairs admitting at least one verified control realization of duration $\tau_\kappa > 0$. Restricting these pairs leads to the following:

$$\begin{aligned}\mathcal{D}_{b\kappa} = \{(y,q) \in D_\kappa : y \in D_b \cap \Gamma^{\mathrm{in}},\\ q \in \Gamma^{\mathrm{out}} \cap \widehat{K}_{t_b+\tau_\kappa, \le T_{\max}-t_b-\tau_\kappa}\}.\end{aligned} \tag{29}$$

Take the parent domain D_b and terminal outer set $\widehat{K}$ to be closed. Their intersection with the compact verified action relation and closed faces makes $\mathcal{D}_{b\kappa}$ compact. Bellman minimization on this relation selects a joining state and its realizing control.

3.1.6. Continuous-Predecessor Cost Propagation

The continuous joining state calculation uses the face interior predecessors and boundary face recursion [23] (Section 2.2). In Algorithm 1, lines 28–43 combine the parent front with the stored directed action cost.

For a compatible parent front b and action κ, fix an exit q and define its feasible predecessor set as $\mathcal{Y}^{\mathrm{loc}}_{b\kappa}(q) = \{y : (y,q) \in \mathcal{D}_{b\kappa}\}$. Combining the parent value $V_b(y)$ with the action cost $A_\kappa(y,q)$ gives the accumulated cost and its realizing predecessor:

$$\begin{aligned}V^{\mathrm{loc}}_{b\kappa}(q) &= \min_{y \in \mathcal{Y}^{\mathrm{loc}}_{b\kappa}(q)} [V_b(y) + A_\kappa(y,q)],\\ y^{\mathrm{loc}}_{b\kappa}(q) &\in \arg\min_{y \in \mathcal{Y}^{\mathrm{loc}}_{b\kappa}(q)} [V_b(y) + A_\kappa(y,q)].\end{aligned} \tag{30}$$

The compactness of the feasible relation and lower semicontinuity of the costs ensure attainment on each nonempty predecessor set. The empty set value is $+\infty$. The control parameters retained during reconstruction implement the same update in joint coordinates. Using the verified feasible parameter set for fixed endpoints $\Theta_\kappa(y,q)$ from Equation (13), the update satisfies

$$\begin{aligned}\min_{y \in \mathcal{Y}^{\mathrm{loc}}_{b\kappa}(q)} \left[V_b(y) + \min_{\theta \in \Theta_\kappa(y,q)} J_\kappa(\theta) \right]\\ = \min_{\substack{y \in \mathcal{Y}^{\mathrm{loc}}_{b\kappa}(q)\\ \theta \in \Theta_\kappa(y,q)}} [V_b(y) + J_\kappa(\theta)].\end{aligned} \tag{31}$$

To express the joining state on a source face $\rho = \mathrm{conv}\{v_0, \ldots, v_d\}$, let $B_\rho = [v_0 \ \cdots \ v_d]$ and $\Delta_d = \{\beta \ge 0 : \mathbf{1}^T\beta = 1\}$. The admissible source coordinates and recovered state are as follows:

$$\begin{aligned}\Lambda^\rho_{b\kappa}(q) &= \{\beta \in \Delta_d : (B_\rho\beta, q) \in \mathcal{D}_{b\kappa}\},\\ y^\star(q) &= B_\rho\beta^\star(q), \qquad \beta^\star(q) \in \Lambda^\rho_{b\kappa}(q).\end{aligned} \tag{32}$$

Herein, $\beta^\star$ minimizes the accumulated cost in Equation (30) after the substitution of $y = B_\rho\beta$. In a local model with an affine control parameterization, let η parameterize the exit, and let ξ contain independent source coordinates and remaining control parameters. After applying the affine coordinate maps, the polyhedral subcase has the following compact relation:

$$K_{b\kappa} = \{(\xi,\eta) : A_{\mathrm{eq}}\xi = b_{\mathrm{eq}} + B_{\mathrm{eq}}\eta,\ G_{\mathrm{in}}\xi \le w_{\mathrm{in}} + S_{\mathrm{in}}\eta,\ \eta \in Q\}. \tag{33}$$

The matrices and vectors $(A_{\mathrm{eq}}, b_{\mathrm{eq}}, B_{\mathrm{eq}})$ and $(G_{\mathrm{in}}, w_{\mathrm{in}}, S_{\mathrm{in}})$ describe the equality and inequality constraints, respectively, and Q is a compact exit polytope. Each pair in $K_{b\kappa}$

satisfies the verified action conditions. A local quadratic representation of the parent cost and quadratic control effort yields the following:

$$\Psi_{b\kappa}(\xi,\eta) = \tfrac{1}{2}\xi^T H_{\xi\xi}\xi + \xi^T(H_{\xi\eta}\eta + h_\xi) + \tfrac{1}{2}\eta^T H_{\eta\eta}\eta + h_\eta^T\eta + c_0. \tag{34}$$

The quadratic coefficients can be assembled directly based on the stored local parent cost representation and affine control map. On that domain, let s be independent parent coordinates, $y = y_{b,0} + T_b s$, and

$$\begin{aligned} V_b(y_{b,0} + T_b s) &= \tfrac{1}{2}s^T \mathsf{P}_b s + p_b^T s + r_b, \\ s &= E_\xi \xi + E_\eta \eta + e_b, \qquad \mathbf{U} = \mathsf{C}_\xi \xi + \mathsf{C}_\eta \eta + \mathbf{u}_0, \end{aligned}$$

where T_b maps the parent coordinates to the full state, E_ξ, E_η, e_b encode the selected joint coordinates, and $\mathsf{C}_\xi, \mathsf{C}_\eta, \mathbf{u}_0$ are the stored affine control coefficients. For the new segment with effort $\frac{1}{2}\mathbf{U}^T H_u \mathbf{U}$, the symmetric weight H_u includes the control weights and segment durations. Expanding the parent-plus-action cost leads to the following:

$$\begin{aligned} H_{\xi\xi} &= E_\xi^T \mathsf{P}_b E_\xi + \mathsf{C}_\xi^T H_u \mathsf{C}_\xi, \\ H_{\xi\eta} &= E_\xi^T \mathsf{P}_b E_\eta + \mathsf{C}_\xi^T H_u \mathsf{C}_\eta, \\ H_{\eta\eta} &= E_\eta^T \mathsf{P}_b E_\eta + \mathsf{C}_\eta^T H_u \mathsf{C}_\eta, \\ h_\xi &= E_\xi^T(\mathsf{P}_b e_b + p_b) + \mathsf{C}_\xi^T H_u \mathbf{u}_0, \\ h_\eta &= E_\eta^T(\mathsf{P}_b e_b + p_b) + \mathsf{C}_\eta^T H_u \mathbf{u}_0, \\ c_0 &= \tfrac{1}{2}e_b^T \mathsf{P}_b e_b + p_b^T e_b + r_b + \tfrac{1}{2}\mathbf{u}_0^T H_u \mathbf{u}_0. \end{aligned}$$

The coefficient expansion following Equation (34) is evaluated once for each fixed parent branch and affine control map. It retains the mixed terms between the joining-state and exit coordinates. A restriction or stationary lift substitutes an affine map into that same quadratic. In local coordinates, write its value as $\frac{1}{2}w^T H w + h^T w + c$ and the substitution as $w = Az + b$. The stored coefficients become

$$H_{\text{new}} = A^T H A, \qquad h_{\text{new}} = A^T(Hb + h), \qquad c_{\text{new}} = \tfrac{1}{2}b^T H b + h^T b + c.$$

Thus, a face restriction and its value representation use the same coordinate map. A change in the parent value or control map requires the affected coefficients to be updated; a change confined to action admissibility triggers the associated validity checks.

With each fixed exit, the stationary equations and face comparisons produce the candidate predecessor parameters. Varying η over the exit domain produces the associated cost and predecessor functions.

The equality constraints are eliminated before computing an interior stationary point. The process selects a particular solution $\xi_p(\eta) = g + L_\eta \eta$ satisfying $A_{\text{eq}}\xi_p(\eta) = b_{\text{eq}} + B_{\text{eq}}\eta$ on the compatible exit domain, and the columns of N form a basis of $\ker A_{\text{eq}}$. Thus, $\xi = \xi_p(\eta) + Nz$. Set $H_r = N^T H_{\xi\xi} N$ and $k_r(\eta) = N^T[H_{\xi\xi}\xi_p(\eta) + H_{\xi\eta}\eta + h_\xi]$. With $H_r \succ 0$, solving $H_r z^\star = -k_r(\eta)$ yields the following:

$$\begin{aligned} z^\star(\eta) &= -H_r^{-1} k_r(\eta), \\ \xi^\star(\eta) &= \xi_p(\eta) + N z^\star(\eta), \\ V_{\text{new}}(\eta) &= \Psi_{b\kappa}(\xi_p(\eta), \eta) - \tfrac{1}{2}k_r(\eta)^T H_r^{-1} k_r(\eta). \end{aligned} \tag{35}$$

This candidate is retained on exits satisfying $G_{\text{in}}\xi^\star(\eta) \le w_{\text{in}} + S_{\text{in}}\eta$. When the equalities determine ξ uniquely, ξ_p directly supplies the candidate. Boundary minima and semidefinite

source Hessians are covered by active face enumeration. An active face is a face of the joint feasible parameter set where selected inequalities hold as equalities. These equalities may represent a source-simplex boundary or a control bound. For an active set I, appending the selected inequality rows produces $A_I\xi = b_I + B_I\eta$. The columns of N_I span $\ker A_I$. With $H_{\xi\xi} \succeq 0$ on the current quadratic branch, the stationary relation and exit projection are as follows:

$$\begin{aligned} \mathcal{R}_I &= \{(\xi,\eta) \in K_{b\kappa} : A_I\xi = b_I + B_I\eta,\ N_I^T(H_{\xi\xi}\xi + H_{\xi\eta}\eta + h_\xi) = 0\}, \\ Q_I &= \operatorname{proj}_\eta \mathcal{R}_I. \end{aligned} \tag{36}$$

Each compatible local parent cost representation is processed separately. Each nonempty Q_I is triangulated in its affine hull, retaining the lower-dimensional cells. In an exit simplex $\Delta = \operatorname{conv}\{\eta_0, \ldots, \eta_r\}$, a feasible stationary lift $(\xi_k, \eta_k) \in \mathcal{R}_I$ is selected at each vertex. The weights μ below parameterize the exit simplex and its predecessor map, with β denoting parent coordinates and λ signifying the action endpoint weights:

$$\eta = \sum_{k=0}^{r} \mu_k \eta_k, \qquad \xi_{I,\Delta}(\eta) = \sum_{k=0}^{r} \mu_k \xi_k, \qquad \mu \in \Delta_r, \tag{37}$$

The convexity of $\mathcal{R}_I$ preserves the feasibility under this lift. For a fixed exit and active face, the source convexity makes the stationary candidates attain that face's minimum, including flat directions. Comparing the candidates across faces determines the branch minimum. The stored affine maps recover the joining state and control coefficients:

$$\begin{aligned} y^\star(\eta) &= y_{b,0} + T_b(E_\xi \xi^\star(\eta) + E_\eta \eta + e_b), \\ \mathbf{U}^\star(\eta) &= \mathsf{C}_\xi \xi^\star(\eta) + \mathsf{C}_\eta \eta + \mathbf{u}_0. \end{aligned}$$

For a face candidate, $\xi_{I,\Delta}$ is substituted into these maps. The declared control parameterization converts $\mathbf{U}^\star$ to $u^\star(t)$, whose forward flow starts at $(y^\star, t_b)$ and reaches q at $t_b + \tau_\kappa$. A complete comparison of the feasible candidates for the parent–action pair yields $y^{\mathrm{loc}}_{b\kappa}(q)$ and its associated trajectory segment π_κ. Their concatenation with the parent history leads to the following:

$$\begin{aligned} \Pi^{\mathrm{loc}}_{b\kappa}(q) &= \operatorname{Concat}\Big(\Pi_b(y^{\mathrm{loc}}_{b\kappa}(q)), \pi_\kappa(y^{\mathrm{loc}}_{b\kappa}(q), q)\Big), \\ J[\Pi^{\mathrm{loc}}_{b\kappa}(q)] &= V_b(y^{\mathrm{loc}}_{b\kappa}(q)) + A_\kappa(y^{\mathrm{loc}}_{b\kappa}(q), q) = V^{\mathrm{loc}}_{b\kappa}(q). \end{aligned} \tag{38}$$

The search stores each face candidate as a local arrival-cost representation containing its exit domain, realized cost, predecessor map, action index, and verification data:

$$P_{I,\Delta} = (\Delta, V_{I,\Delta}(\eta), \xi_{I,\Delta}(\eta), \kappa, \chi_\kappa), \qquad V_{I,\Delta}(\eta) = \Psi_{b\kappa}(\xi_{I,\Delta}(\eta), \eta). \tag{39}$$

Each candidate inherits the planning instance ι_P, arrival time t_P and access class a_P of its parent–action pair. Insertion requires a complete domain, value, and realizing map. Section 3.1.7 gives the nonlinear domain procedure (Algorithm 1, line 45).

3.1.7. Front Updates and Unresolved Domains

Front records retain domains, values, and realization histories; unresolved records additionally retain unprocessed domains, lower bounds, parent versions, and cursors. Finite implementation details are given in Section 5.1.

Lower-envelope updates and domain classification apply to both quadratic and nonlinear families.

The set $\mathscr{Q}_i$ retains the relations still to be resolved, while $\mathscr{G}_{D,i}$ records the premises of action and front validity. On a source–exit domain Ω, let $\mathcal{V}_\Omega$ contain pairs with a verified control, and let $\mathcal{X}_\Omega^{\mathrm{exc}}$ contain pairs for which all policies in the declared model are infeasible. The remaining pairs form $\mathcal{Q}_\Omega$, thus leading to the following:

$$\mathcal{V}_\Omega = \bigcup_{\kappa \in \mathscr{A}_\Omega} D_\kappa, \qquad \mathcal{Q}_\Omega = \Omega \setminus (\mathcal{V}_\Omega \cup \mathcal{X}_\Omega^{\mathrm{exc}}), \qquad \Omega = \mathcal{V}_\Omega \dot\cup \mathcal{X}_\Omega^{\mathrm{exc}} \dot\cup \mathcal{Q}_\Omega. \tag{40}$$

Competing control realizations remain in $\mathscr{Q}_i$ until value resolution or domain exclusion. Let $\mathfrak{P}(R)$ contain the complete trajectories represented by unresolved record R. The record retains its dependencies and a lower bound over this set:

$$\underline{J}(R) \le \inf_{\Pi \in \mathfrak{P}(R)} J[\Pi]. \tag{41}$$

A finite collection of polyhedral–quadratic models yields a finite exact front through exhaustive active-face enumeration and triangulation. Search processes all compatible families and retains unresolved domains; nonlinear families use the domain-resolution procedure below.

In a local transfer subproblem $\omega = (b, \kappa, \tau, \Gamma^{\mathrm{in}}, \Gamma^{\mathrm{out}})$, let $Q_\omega = \mathrm{proj}_q \mathcal{D}_{b\kappa}$ denote its complete feasible exit domain. In the quadratic coordinates, this domain is represented by $\mathrm{proj}_\eta K_{b\kappa}$. Let $\mathscr{P}_\omega$ contain feasible candidate cost representations. For the current record, $\mathcal{E}_\omega^{\mathrm{res}}$ contains exits whose minimum over all candidates of the local transfer subproblem has been established. Infeasible pairs are classified separately based on Equation (40).

$$\mathcal{W}_\omega = Q_\omega \setminus \mathcal{E}_\omega^{\mathrm{res}}, \qquad Q_\omega = \mathcal{E}_\omega^{\mathrm{res}} \dot\cup \mathcal{W}_\omega. \tag{42}$$

The union $\bigcup_{P \in \mathscr{P}_\omega} \mathrm{dom}(P)$ of feasible candidate domains may include exits whose local minima remain unresolved. After exhaustive exact face enumeration in the finite quadratic branch, the feasible candidate coverage resolves the local value. In the general branch, a candidate may improve an upper bound, while its domain remains in $\mathcal{W}_\omega$ until the matching lower bound closes the gap. No subset of $\mathcal{W}_\omega$ is removed merely because a feasible value has been found elsewhere. Across all compatible parent-cost representations and local action models, the propagated front is as follows:

$$V_\iota(a, q, t) = \min_{\substack{P:\, q \in \mathrm{dom}(P) \\ t_P = t,\ a_P = a,\ \iota_P = \iota}} V_P(q), \tag{43}$$

where an empty minimum is $+\infty$. The access class a, time t, planning-instance identifier ι, and immutable predecessor version identify compatible front records. During partial resolution, the lower envelope provides the least cost among the feasible prefix realizations obtained so far. The exhaustive resolution recovers the Bellman value over the declared admissible controls. Nonlinear or nonpolyhedral actions are based on the validated domain procedure below to return feasible arrival-cost representations, lower bounds, and unresolved domains.

On the convergent nonlinear branch, the lifted feasible relation is defined as follows:

$$\widetilde{\mathcal{D}}_{b\kappa} = \{(y, q, \theta) : (y, q) \in \mathcal{D}_{b\kappa},\ \theta \in \Theta_\kappa(y, q)\}.$$

Resolving this relation in a finite collection of coordinate domains leads to the following:

$$\psi_l : Q_l \times Z_l \longrightarrow \widetilde{\mathcal{D}}_{b\kappa}, \qquad \psi_l(q, z) = (y_l(q, z), q, \theta_l(q, z)), \tag{44}$$

where Q_l and Z_l are compact, q is the exit state, and z contains the remaining independent coordinates. The images cover the declared feasible relation, including its boundary and terminal restrictions. The endpoint elimination on a regular constraint branch supplies the local coordinates, and interval validation checks the feasibility throughout each retained coordinate domain. The convergent branch requires such a complete finite covering, and an unverified branch remains in the unresolved collection.

In a fixed parent-cost representation, let $g_l(q,z)$ be its arrival cost plus the realized action cost at $\psi_l(q,z)$. Assume that g_l is uniformly continuous in z, uniformly in q, with a computable modulus ϖ_l satisfying $\lim_{\delta\to 0^+}\varpi_l(\delta) = 0$. Z_l is partitioned into finitely many cells of diameter at most δ, and a representative z_{lj} is selected in each nonempty cell. Each representative defines a feasible source and control map on Q_l. Let $\widetilde{g}_{lj}(q)$ be a computed representative cost with $|\widetilde{g}_{lj}(q) - g_l(q,z_{lj})| \le \epsilon_{\mathrm{num}}$, uniformly on Q_l, where validated cost evaluation is performed to the prescribed tolerance ϵ_{num}. The refinement sequence decreases both the cell diameter and the numerical tolerance. Define

$$\begin{aligned} \overline{v}_l(q) &= \min_j \widetilde{g}_{lj}(q) + \epsilon_{\mathrm{num}}, \\ \underline{v}_l(q) &= \min_j \widetilde{g}_{lj}(q) - \varpi_l(\delta) - \epsilon_{\mathrm{num}}, \\ \underline{v}_\omega(q) &= \min_{l:q\in Q_l} \underline{v}_l(q), \qquad \overline{v}_\omega(q) = \min_{l:q\in Q_l} \overline{v}_l(q). \end{aligned} \tag{45}$$

The upper estimate is accompanied by its feasible realizing control and a validated bound on that control's cost. Every feasible coordinate lies within δ of a cell representative, thus leading to the following:

$$\begin{aligned} 0 &\le \overline{v}_\omega(q) - V^{\mathrm{loc}}_{b\kappa}(q) \le \overline{v}_\omega(q) - \underline{v}_\omega(q) \\ &\le \max_{l:q\in Q_l} \varpi_l(\delta) + 2\epsilon_{\mathrm{num}}, \qquad \lim_{\substack{\delta\to 0^+ \\ \epsilon_{\mathrm{num}}\to 0^+}} \left(\max_l \varpi_l(\delta) + 2\epsilon_{\mathrm{num}} \right) = 0. \end{aligned} \tag{46}$$

For nonlinear action families satisfying the covering and continuity conditions in Equations (44)–(46), the local Bellman value is approximated with a certified error bound. At fixed cell diameter and numerical tolerance, local optimality is guaranteed up to the upper–lower cost gap in Equation (46). Trajectory feasibility is enforced separately by full-interval verification of the realizing controls. Under the refinement conditions of Theorem 3, the local, terminal and search errors vanish along the refinement sequence, yielding the stated asymptotic optimality. The terminal minimization is based on the corresponding compact terminal coordinates. The polyhedral–quadratic actions retain the exact face construction of Section 3.1.6. CertifiedDomainSolve resumes the coordinate cells saved in $R|_\omega$, evaluates the representatives and bounds in Equation (45), and refines the unresolved cells within the remaining budget. It returns the feasible candidates in $\mathscr{P}_\omega$, unresolved exit records in $\mathcal{W}_\omega$ with their updated processing cursors, and the work increment Δw_ω.

Each feasible cost representation stores the actual parameter selector θ_P, source map y_P, and verification record χ_P. The latter verifies the realizing control on the entire retained domain and time interval. Its value is $V_b(y_P(q)) + J_\kappa(\theta_P(q))$, even when a smaller value remains unresolved. The restriction $\kappa|_{\theta_P(\mathrm{dom}\,P)}$ stores the same action family on the selected parameter image. Line 51 applies $\mathrm{Env}_{\iota,a,t}$ and adds the immutable predecessor dependencies, while line 53 applies Rep to preserve the unresolved alternatives (Algorithm 1, lines 48–53).

For each $P \in \Delta\mathscr{B}$, $\mathcal{S}(P)$ contains its compatible successor action records and $\mathcal{T}(P)$ its terminal restrictions. Their union in line 56 inserts records for a new front and reopens the affected records for an improved front with updated lower bounds. The scope identity

includes the parent version, action family, time, access class, and epoch. For each fixed parent version and action family in the exact quadratic branch, the enumerated faces, projected triangulation, selected lifts, and processing cursor are retained with this identity. Duplicate insertions therefore reuse the same candidates without losing alternatives. The terminal candidates are passed to the recovery step in Section 3.1.8. A replaced local lower envelope representation remains available to any immutable predecessor record that still refers to it.

3.1.8. Terminal Recovery and Stopping

RESOLVETERMINAL processes terminal restrictions and stores feasible pairs in $\mathscr{T}_{\mathrm{rec}}$, retaining unresolved subsets and their lower bounds in OPEN (Algorithm 1, lines 11–16). Line 57 selects the least-cost stored pair (b_g, q_g), using $\prec_g$ to break ties, or returns an empty record. For a chain of m_g segments, $\mathrm{Pred}_{b_j}(q_j) = (b_{j-1}, \kappa_j, q_{j-1}, \theta_j)$ supplies the state–control segments $\mathrm{Seg}(q_{j-1}, \mathbf{U}_{\kappa_j}(\theta_j))$ in Equation (49). Their concatenation $\bigoplus$ in increasing physical time gives the trajectory in line 59.

During search, only part of the feasible histories and terminal states have been verified. Let $\mathscr{B}_i^{\mathrm{cert}}$ contain the available fronts whose histories start at the current root and remain valid for the same instance. Let $\mathscr{T}_i$ collect the verified pairs (b, q) with $b \in \mathscr{B}_i^{\mathrm{cert}}$, $q \in D_b \cap \mathcal{X}_g$, and $t_b \le T_{\max}$, including the terminal record of a verified initial feasible trajectory. The incumbent is the least realized value among these candidates:

$$U_i = \min_{(b,q)\in\mathscr{T}_i} V_b(q), \qquad J_i^\star \le J_{\mathscr{M}_i}^\star \le U_i. \tag{47}$$

Under a position-only terminal requirement, DAR evaluates each terminal front over its admissible velocity range as well as its position range. Thus, Equation (47) specializes to the following:

$$U_i^{\mathrm{free}} = \min_{(b,q)\in\mathscr{T}_i^{\mathrm{free}},\ q=(\mathbf{r}^T,\mathbf{v}^T)^T} V_b(q), \tag{48}$$

where $\mathscr{T}_i^{\mathrm{free}}$ denotes the verified terminal pairs for $\mathcal{X}_g^{\mathrm{free}}$. The selected complete terminal state q_g includes its minimizing velocity. For a quadratic front on a polyhedral domain, the position constraint restricts the terminal parameter domain, and the same stationary face calculation used by DAR determines its candidates. For a general nonlinear front, the terminal restriction is included in the validated domain calculation. The unresolved terminal subsets retain their lower bounds in OPEN. The exact terminal minimization is required for exact finite-model optimality, whereas a validated unresolved gap supports the stated approximate stopping rule.

If no feasible terminal candidate is available, $U_i = +\infty$. Otherwise, the least-cost candidate and its stored predecessors provide the incumbent. The complete exact resolution yields $U_i = J_{\mathscr{M}_i}^\star$. During partial resolution, the feasible candidates supply the upper bounds, and the unresolved alternatives supply the lower bounds.

Equation (45) bounds continuations of a specified parent history. A feasible parent value supplies an upper bound; competing histories require the complete-trajectory scope bounds in Equation (41). The scopes retain every unresolved alternative, including parent histories whose prefix costs can still decrease. The minimum of their lower bounds and the incumbent gives the global bound. Zero is valid for the nonnegative objective; tighter bounds combine validated prefix and local-action bounds with an admissible terminal heuristic for the same instance.

The compatible fronts are compared at the same exit state, time, access class, and instance. Each accepted value stores its realizing predecessor and control during propagation.

Recovery reads these records, while Basic S* performs additional geometric predecessor queries [23] (Section 3).

For a selected terminal cost representation, each predecessor record retains the previous front, action, source state, and stored control parameter θ_j. Reading the records to the root and then applying the controls in forward time leads to the following:

$$\begin{aligned} \mathrm{Pred}_{b_j}(q_j) &= (b_{j-1}, \kappa_j, q_{j-1}, \theta_j), \\ u_\Pi(t) &= u_{\mathbf{U}_{\kappa_j}(\theta_j)}(t), \qquad t \in [t_{j-1}, t_j), \\ x_\Pi(t) &= \phi_{t_{j-1},t}\big(q_{j-1}, u_{\mathbf{U}_{\kappa_j}(\theta_j)}\big), \\ t_j &= t_{j-1} + \tau_{\kappa_j}. \end{aligned} \tag{49}$$

Equations (15) and (38) ensure segment feasibility and preservation of the accumulated cost.

Finally, the lower bounds of the unresolved domains and incumbent define the current gap (Algorithm 1, lines 62–70). The convention $\min \varnothing = +\infty$ applies. The work budget is debited after each selected scope, and the search stops on budget exhaustion, complete resolution, or the requested finite gap. The search budget B, denoted by B_i at replanning epoch i, allows an indivisible calculation started within the budget to finish, even if its completion exceeds the remaining allowance.

Persistent cursors retain unprocessed active faces, projected cells (including lower-dimensional cells), sibling families and competing values. An exit enters $\mathcal{E}_\omega^{\mathrm{res}}$ after all candidates in $\mathcal{J}_\omega(\eta)$ are processed. Counters w_R and Δw_ω measure scope and nonlinear-branch work.

Replacement retains unresolved domains with their bounds, cursors, parent versions, epochs, and dependencies. Exclusion premises remain available for reopening. At return, $\mathrm{Ref}(\mathscr{A}, \mathscr{B}, \mathscr{Q}, \mathrm{Pred})$ and line 69 preserve the records, exclusions, immutable histories and ancestor dependencies.

The unresolved domains carry the status UNCOVERED, UNFINISHED, or NUMERICALLYUNRESOLVED. A domain is removed by verified emptiness, exhaustive value resolution, or whole-domain dominance. The dominance rule compares records on the same joining interface, at the same physical time, access class, and planning instance. Record P dominates record Q when

$$D_Q \subseteq D_P, \qquad V_P(q) \le V_Q(q), \quad \forall q \in D_Q.$$

Common-face restrictions identify the same complete joining state q. From that state and time, the common interface and access class admit the same successor action relations. The realizing controls of any continuation ς of Q therefore generate the same future flow when appended to P. Additivity gives

$$J[P \oplus \varsigma] = V_P(q) + J[\varsigma] \le V_Q(q) + J[\varsigma] = J[Q \oplus \varsigma].$$

Thus, whole-domain dominance preserves every admissible terminal continuation at no greater cost.

Algorithm 1 assembles the preceding search and recovery steps. The fields below define the records used in Algorithm 1.

Scope fields $(b_R, D_R, t_R, a_R, C_R)$ specify the parent front, source domain, departure time, access class and carrier. The predicates $C_R \sim C_\kappa$ and $D_\omega \overset{\mathrm{v}}{=} \varnothing$ denote carrier compatibility and verified emptiness. The operator $\mathrm{Env}_{\iota,a,t}$ applies Equation (43). Restrictions

$R|_W$ and $R|_\omega$ select domain W and component ω, respectively. Replacement preserves unrestricted fields:

$$\mathrm{Rep}(\mathcal{O}, \mathcal{G}; R, \mathcal{W}) = \left((\mathcal{O} \setminus \{R\}) \cup \{R|_W : W \in \mathcal{W}\}, \mathcal{G} \cup \bigcup_{W \in \mathcal{W}} \mathrm{DepEdges}(R|_W) \right),$$

where $\mathrm{DepEdges}(\cdot)$ returns edges from supporting premises and predecessor records.

Selection follows the fair schedule $\mathrm{select}_{\prec_F}$; the fixed order $\prec_g$ breaks ties between terminal minimizers.

For the quadratic branch, $\mathfrak{M}_{\mathrm{PQ}}$ contains compact polyhedral transfer relations with quadratic costs and $H_{\xi\xi} \succeq 0$. The set $\mathcal{I}_{\mathrm{act}}(R|_\omega, B - w_R)$ stores the active sets still to be processed within the remaining budget. The unfinished projected simplices are $\{\Delta \in \mathrm{Tri}(Q_I) : (I, \Delta) \notin \mathrm{Done}(R|_\omega)\}$. For an exit η, let $\mathcal{J}_\omega(\eta) = \{(I, \Delta) : \Delta \in \mathrm{Tri}(Q_I),\ \eta \in \Delta\}$. Its local value is resolved after every incident candidate in this set has been processed.

For fronts $\Delta\mathscr{B}$ inserted or improved during the current visit, $\mathcal{S}(P)$ and $\mathcal{T}(P)$ contain successor-action and terminal records satisfying the time, access, carrier and terminal tests defined above. The operator Ref retains current records, exclusions and immutable predecessor histories; $\mathscr{G}_D[\mathrm{Anc}_{\mathscr{G}_D}(\cdot)]$ retains their ancestor dependency subgraph.

3.2. ISTR Incremental Replanning Framework*

3.2.1. Reusable Carrier and Action Archive

The model data $\mathfrak{m}$ comprise the dynamics, running cost, admissible domains, and coordinate convention. The construction settings $\mathfrak{b} = (N_H, N_D, N_G, r_{\mathrm{conn}}, \mathbb{T}, \mathbb{T}^{\mathrm{dep}}, \mathrm{g}_G)$ specify the sampling counts, growth limit, connection radius, action durations, departure times, and growth proposal policy for Algorithm 2. At epoch i, $\mathfrak{q}_i = (\mathbf{x}_i, t_i, \mathcal{X}_g, T_{\max})$ records the current planning boundary conditions. The prediction data $\mathfrak{p}_i = (\widehat{\mathcal{O}}_i, \mathbb{I}_i, \nu_i)$ comprise the occupied-region enclosure, its covered times, and the posterior identifier.

The carrier construction combines Halton coverage V_H, dispersion refinement V_D, and dynamics-guided growth V_G. Let $\underline{\mathbf{x}}, \overline{\mathbf{x}} \in \mathbb{R}^6$ be the lower and upper bounds of the complete state-coordinate box in $\mathfrak{m}$ and $\mathbf{s}_n^H \in [0,1]^6$ be its nth Halton point. The affine map in line 1 of Algorithm 2 places the Halton sequence in that box. DispersionRefinement adds N_D points within the same bounds based on the dispersion reduction criterion of Dispertio [7]. Action validation subsequently applies the spacecraft dynamics and constraints.

The declared growth policy g_G proposes (x, u, τ) from a current sample, an admissible control and $\tau \in \mathbb{T}$, or returns $\varnothing$. Algorithm 2 attempts N_G proposals. An endpoint is inserted when $(x, u) \in \mathrm{dom}\, \phi_{t_s, t_s+\tau}$ and the endpoint lies in $\mathcal{X}_{\mathrm{free}}$. Thus N_G bounds growth attempts. Time-indexed action verification subsequently enforces collision constraints.

The candidate vertices are as follows:

$$V = (V_H \cup V_D \cup V_G) \cap \mathcal{X}_{\mathrm{free}}. \tag{50}$$

The candidate set V is supplemented as needed to form the conforming carrier vertex set V_{δ_x} in Section 3.1. Sampling and local-connection principles follow established constructions [6,7,13]. Let $\mathscr{P}_C \subseteq \mathscr{C} \times \mathscr{C}$ be the declared ordered carrier pairs. For $(C_a, C_b) \in \mathscr{P}_C$, let $\mathcal{D}^{\mathfrak{m}}_{ab,t,\tau}$ be the source–exit relation permitted by dynamics and admissibility specification $\mathfrak{m}$ for departure time t and duration τ. This relation contains the flow, control, access-class, and interface conditions. Algorithm 2 associates each carrier interface and duration with verified control parameter domains and retains the exact complement that has not yet been resolved.

Posterior updates preserve every accepted sample in V, including $V_H \cap \mathcal{X}_{\mathrm{free}}$, as a carrier vertex. Collision screening restricts the time-indexed action domains. At a fixed

refinement level with unchanged $\mathcal{X}_{\text{free}}$, replacement triangulations include the retained sample coordinates. Consequently, any state coverage bound established for these samples is preserved. The refinement sequence conditions are specified in Section 4.

Algorithm 2 CONSTRUCTARCHIVE: reusable carriers and local actions

Input: construction settings $\mathfrak{b}$, model data $\mathfrak{m}$, interval prediction $\mathfrak{p}$, boundary data $\mathfrak{q} = (\mathbf{x}_s, t_s, \mathcal{X}_g, T_{\max})$
Output: $(\mathcal{K}_{\delta_x}, \mathscr{C}, \mathscr{A}, \mathscr{Q}, \mathscr{G}_D)$

// Halton coverage and Dispertio refinement (Section 3.2.1).
$V_H \leftarrow \{\underline{\mathbf{x}} + \operatorname{diag}(\bar{\mathbf{x}} - \underline{\mathbf{x}})\mathbf{s}_n^H\}_{n=1}^{N_H}$
$V_D \leftarrow$ DISPERSIONREFINEMENT$(V_H, N_D, \mathfrak{m})$
$V_G \leftarrow \varnothing$
for $k = 1, \ldots, N_G$ **do**
  $(x, u, \tau) \leftarrow \mathrm{g}_G(V_H \cup V_D \cup V_G, \mathfrak{m}, \mathbb{T}, t_s)$
  if $(x, u, \tau) \neq \varnothing \ \wedge \ (x, u) \in \operatorname{dom} \phi_{t_s, t_s+\tau}$ **then**
    $x^+ \leftarrow \phi_{t_s, t_s+\tau}(x, u)$
    $V_G \leftarrow V_G \cup \{x^+\}$ if $x^+ \in \mathcal{X}_{\text{free}}$
  end if
end for
$V \leftarrow [(V_H \cup V_D \cup V_G) \cap \mathcal{X}_{\text{free}}] \cup \{\mathbf{x}_s\}$

// Retain every accepted sample in V as a carrier vertex.
// Common-face assembly (Equations (9) and (10)).
$\mathcal{K}_{\delta_x} \leftarrow$ ASSEMBLECARRIERCOMPLEX$(V, r_{\text{conn}}, \mathfrak{q})$
$\mathscr{C} \leftarrow \{C : C \in \mathcal{K}_{\delta_x}\}$

// Insert the root exactly and represent terminal intersections and common subfaces.
write $\mathbb{T} = \{\tau_1, \ldots, \tau_{N_\tau}\}$
$\mathbb{T}_q \leftarrow \left\{ t + \sum_{j=1}^{N_\tau} m_j \tau_j : \ t \in \mathbb{T}^{\text{dep}} \cup \{t_s\}, \ \mathbf{m} \in \mathbb{N}_0^{N_\tau}, \ t + \sum_{j=1}^{N_\tau} m_j \tau_j \leq T_{\max} \right\}$
$(\mathscr{A}, \mathscr{Q}, \mathscr{X}) \leftarrow (\varnothing, \varnothing, \varnothing)$
for all $(C_a, C_b) \in \mathscr{P}_C$, $t \in \mathbb{T}_q$, $\tau \in \mathbb{T}$, $t + \tau \leq T_{\max}$ **do**
  $D_{ab,t,\tau} \leftarrow (C_a \times C_b) \cap \mathcal{D}^{\mathfrak{m}}_{ab,t,\tau}$

// Whole-interval action constraints (Equation (15)).
  $(\mathscr{A}_{ab,t,\tau}, \mathscr{Q}_{ab,t,\tau}, \mathscr{X}_{ab,t,\tau}) \leftarrow$ CERTIFYACTIONS$(D_{ab,t,\tau}, t, \tau, \mathfrak{m}, \mathfrak{p})$
  $(\mathscr{A}, \mathscr{Q}, \mathscr{X}) \leftarrow (\mathscr{A} \cup \mathscr{A}_{ab,t,\tau}, \mathscr{Q} \cup \mathscr{Q}_{ab,t,\tau}, \mathscr{X} \cup \mathscr{X}_{ab,t,\tau})$
end for

// Record the premises needed for the update in Section 3.2.2.
$\mathscr{G}_D \leftarrow \{(p, R) : R \in \mathscr{C} \cup \mathscr{A} \cup \mathscr{Q} \cup \mathscr{X}, \ p \in \operatorname{Prem}(R)\}$
return $(\mathcal{K}_{\delta_x}, \mathscr{C}, \mathscr{A}, \mathscr{Q}, \mathscr{G}_D)$

ASSEMBLECARRIERCOMPLEX combines the samples with the marked root and terminal intersections. The construction satisfies the subface and common-intersection conditions in Equations (9) and (10), preserving retained vertices and identical subdivisions on shared interfaces. The induced boundary subdivisions propagate to adjacent carriers. Then, actions are verified within these geometric domains.

Algorithm 2, line 15, closes departure times under sums of positive durations in $\mathbb{T}$ up to $T_{\max}$; $m_j \in \mathbb{N}_0$ counts duration τ_j. Finite initial sets and a positive minimum duration make the closure finite. Each carrier pair specifies entry and exit faces and an allowed carrier union, including required within-carrier transfers. The model restricts the source–exit relation before prediction-dependent certification (lines 17–21). P2 and Equations (11)–(13) supply the realizations, domains, and costs.

Refinement coordinates the state-cover radius, represented times and control partition through Proposition 1. The root is inserted exactly, and terminal states are selected continuously on intersections of exit domains with $\mathcal{X}_g$. Every transfer required by the recovery neighborhood is retained. Failed endpoint-correction tests leave the corresponding domains unresolved for further refinement.

CertifyActions applies Equation (40): acceptance verifies Equation (15) over every retained parameter and the full interval; exclusion rules out all declared control alternatives for a pair. Untested and inconclusive domains, together with domains whose action intervals are not fully covered by the available predictions, remain in $\mathscr{Q}$. Verified actions retain their controls and costs, reusing exact P2 solutions. Dependency edges (p, R) link supporting premises to carrier, action, unresolved-domain and exclusion records.

Unresolved records preserve the family, carriers, time, duration, parameter domain and cursor across updates and OPEN processing. Exclusion records preserve their domains and premises for reopening after premise changes.

3.2.2. Posterior-Driven Archive and Dependency Update

At posterior update i, let $\Delta_i = \{(r,t) : t \in \mathbb{I}_i \cap \mathbb{I}_{i-1},\ r \in \widehat{\mathcal{O}}_i(t) \triangle \widehat{\mathcal{O}}_{i-1}(t)\}$ describe the changed occupancy on the common temporal domain. For each premise $p \in \mathrm{Prem}(R)$, $\mathrm{stat}_i(p) \in \{\mathrm{true}, \mathrm{false}, \mathrm{uncovered}\}$ records validity and $\mathbb{I}(p)$ specifies temporal support. Forward closure from changed or uncovered premises gives the affected records $\mathscr{D}_i$. Dependency edges point from supporting premises or predecessor records to dependent records; $\mathrm{DepEdges}(R)$ and $\mathrm{AncRec}(R)$ denote supporting edges and ancestor records.

Each action certificate records the target–time queries and posterior version ν_i supporting Equation (5). The coordinate frame, probability level, and CDIC region definition remain fixed. These dependencies identify certificates requiring re-evaluation after a posterior update.

The verification for each action on $[t, t+\tau]$ is based on the continuous-time enclosure in Equation (6). A domain remains unresolved until its entire interval is covered by a valid prediction enclosure.

In Algorithm 3, Δ_i compares interval enclosures on their common temporal domain; stat_i and $\mathbb{I}(p)$ check premise validity and coverage. An unchanged region is reusable when its certificate remains valid. Algorithm 2 initializes the first epoch from the received prediction.

With an action κ, let R_κ be its verification record, and $\mathcal{S}_\kappa \subset \mathbb{R}^3 \times [t_\kappa, t_\kappa + \tau_\kappa]$ be its certified swept position support. Based on the record-valued dependency closure defined above, the retained family is as follows:

$$\mathscr{A}_i^- = \{\kappa \in \mathscr{A}_{i-1} : (\{R_\kappa\} \cup \mathrm{AncRec}(R_\kappa)) \cap \mathscr{D}_i = \varnothing, \quad \mathcal{S}_\kappa \cap \Delta_i = \varnothing\}. \tag{51}$$

The impacted carrier neighborhood contains the carriers referenced by affected records and the carriers with state–time support intersecting the changed occupied region enclosure. For a carrier C, $P_r(C) := \{P_r x : x \in C\}$ is its relative position image. The impacted set is

$$\mathscr{C}_i^{\mathrm{imp}} = \{C \in \mathscr{C}_{i-1} : C \in \mathrm{Car}(\mathscr{D}_i) \text{ or } (P_r(C) \times \mathbb{I}_C) \cap \Delta_i \neq \varnothing\}. \tag{52}$$

Here, $\mathbb{I}_C$ is the set of physical times represented by C, and $\mathrm{Car}(\mathscr{D}_i)$ is the set of carriers referenced by records in the affected dependency closure. A local neighborhood $\mathcal{N}_i$ is formed by $\mathscr{C}_i^{\mathrm{imp}}$ and its interface-adjacent carriers. Local refinement produces new carrier and action families satisfying the following:

$$\Omega_i^{\mathrm{loc}} = \mathcal{V}_i^{\mathrm{loc}} \dot\cup \mathcal{X}_i^{\mathrm{exc,loc}} \dot\cup \mathcal{Q}_i^{\mathrm{loc}}, \qquad \mathcal{V}_i^{\mathrm{loc}} = \bigcup_{\kappa \in \mathscr{A}_i^+} D_\kappa, \qquad \mathcal{Q}_i^{\mathrm{loc}} \subseteq \bigcup_{R \in \mathscr{Q}_i^+} \mathrm{dom}(R), \tag{53}$$

Here, Ω_i^{loc} is the source–exit scope induced by $\mathcal{N}_i$; $\mathcal{V}_i^{\mathrm{loc}}$, $\mathcal{X}_i^{\mathrm{exc,loc}}$ and $\mathcal{Q}_i^{\mathrm{loc}}$ are its verified, reliably excluded and unresolved parts, respectively, as in Equation (40). Search records

also retain unresolved values on verified domains. Exclusion records and dependencies support reopening after posterior contraction.

Algorithm 3 ISTR* replanning with simulated state propagation

Input: model data $\mathfrak{m}$, construction settings $\mathfrak{b}$, initial condition $(\mathbf{x}_s, t_s)$, terminal condition $(\mathcal{X}_g, T_{\max})$, simulated observation history $\mathcal{Z}$, replanning interval $T_{\text{exe}} > 0$, search budgets B_i, tolerances ε_i, fair schedule $\prec_F$
Output: applied-control history $\mathscr{H}^{\text{exe}}$ and task status

$(i, \mathbf{x}_0, t_0, \Pi^{\star}_{-1}, \mathscr{H}^{\text{exe}}) \leftarrow (0, \mathbf{x}_s, t_s, \varnothing, \varnothing)$
while true **do**
  // Prediction and root checks (Section 3.2.2).
  $(\mathfrak{q}_i, \mathfrak{p}_i) \leftarrow \textsc{UpdatePrediction}(\mathbf{x}_i, t_i, \mathcal{Z}_{\le t_i}, \mathfrak{m}, \mathcal{X}_g, T_{\max})$
  return $(\mathscr{H}^{\text{exe}}, \text{incomplete})$ **if** $\mathfrak{q}_i = \varnothing$
  return $(\mathscr{H}^{\text{exe}}, \text{reached})$ **if** $\mathbf{x}_i \in \mathcal{X}_g$
  return $(\mathscr{H}^{\text{exe}}, \text{incomplete})$ **if** $t_i = T_{\max}$
  // Construct once, then replace only the affected part.
  if $i = 0$ **then**
    // Carrier and action construction (Algorithm 2).
    $(\mathcal{K}_{\delta_x,i}, \mathscr{C}_i, \mathscr{A}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}) \leftarrow \textsc{ConstructArchive}(\mathfrak{b}, \mathfrak{m}, \mathfrak{p}_i, \mathfrak{q}_i)$
    $\mathscr{B}_i \leftarrow \varnothing$
  else
    $\Delta_i \leftarrow \{(r,t) : t \in \mathbb{I}_i \cap \mathbb{I}_{i-1},\ r \in \widehat{\mathcal{O}}_i(t) \triangle \widehat{\mathcal{O}}_{i-1}(t)\}$
    // Changed prediction premises (Section 3.2.2).
    $\mathscr{S}_i \leftarrow \{R : \exists p \in \text{Prem}(R),\ \text{stat}_i(p) \neq \text{stat}_{i-1}(p) \ \vee\ \mathbb{I}(p) \not\subseteq \mathbb{I}_i\}$
    $\mathscr{D}_i \leftarrow \mathscr{S}_i \cup \text{Reach}^{+}_{\mathscr{G}_{D,i-1}}(\mathscr{S}_i)$
    // Local replacement with common faces (Equation (57)).
    $(\mathcal{K}_{\delta_x,i}, \mathscr{C}_i, \mathscr{A}_i, \mathscr{B}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}) \leftarrow$
      $\textsc{LocalUpdate}(\mathcal{K}_{\delta_x,i-1}, \mathscr{C}_{i-1}, \mathscr{A}_{i-1}, \mathscr{B}_{i-1}, \mathscr{Q}_{i-1}, \mathscr{G}_{D,i-1},$
      $\mathscr{D}_i, \mathfrak{b}, \mathfrak{m}, \mathfrak{p}_i, \mathfrak{q}_i)$
  end if
  // Search from the current simulated state using valid retained information.
  $\iota_i \leftarrow i, \quad \mathcal{B}^{\text{root}}_i \leftarrow \{\mathcal{B}_{b_i(a)} : a \in \mathsf{A}^0_i\}$
  // Reset root-dependent values and verify the unexecuted suffix (see the text below).
  $(\mathscr{B}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}, \Pi^{\text{seed}}_i) \leftarrow \textsc{Rebase}$
    $(\mathcal{B}^{\text{root}}_i, \mathscr{A}_i, \mathscr{B}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}, \Pi^{\star}_{i-1}, \mathfrak{m}, \mathfrak{p}_i, \mathfrak{q}_i, \iota_i)$
  // Directed cost propagation and recovery (Algorithm 1).
  $(\Pi^{\star}_i, L_i, U_i, \mathscr{A}_i, \mathscr{B}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}) \leftarrow \text{DAR}$
    $(\mathcal{B}^{\text{root}}_i, \iota_i, \mathscr{C}_i, \mathscr{A}_i, \mathscr{B}_i, \mathscr{Q}_i, \mathscr{G}_{D,i}, \Pi^{\text{seed}}_i, (\mathcal{X}_g, T_{\max}), B_i, \varepsilon_i, \prec_F)$
  return $(\mathscr{H}^{\text{exe}}, \text{incomplete})$ **if** $\Pi^{\star}_i = \varnothing$
  // Integrate the prefix control to obtain the next planning state.
  $t^{\text{end}}_i \leftarrow \min\{t_i + T_{\text{exe}}, t_f(\Pi^{\star}_i)\}$
  $u^{\text{exe}}_i \leftarrow u_{\Pi^{\star}_i}|_{[t_i, t^{\text{end}}_i]}$
  $(\mathbf{x}_{i+1}, t_{i+1}) \leftarrow (\phi_{t_i, t^{\text{end}}_i}(\mathbf{x}_i, u^{\text{exe}}_i), t^{\text{end}}_i)$
  $\mathscr{H}^{\text{exe}} \leftarrow \mathscr{H}^{\text{exe}} \cup \{(t_i, t_{i+1}, u^{\text{exe}}_i)\}$
  $i \leftarrow i + 1$
end while

The neighborhood $\mathcal{N}_i$ is a collection of geometric carriers, and new states are sampled from its geometric realization,

$$\mathcal{X}^{\text{samp}}_i = \left(\bigcup_{C \in \mathcal{N}_i} C \right) \cap \mathcal{X}_{\text{free}}. \tag{54}$$

Local construction adds interface-adjacent carriers while preserving accepted samples and shared-face triangulations. Missing root or terminal connections trigger construction on

their incident carriers and interface neighbors, including when existing actions remain valid. Uncovered boundary portions use Algorithm 2. The generated relations are classified by Equation (40). For an unchanged occupied-region enclosure, write $\mathbf{x}_i = \widehat{\mathbf{x}}_A(t_i)$ and let $\mathscr{A}_i^{\mathrm{keep}}$ denote the retained valid actions and define the root-compatible outgoing set

$$\begin{aligned} \mathscr{A}_i^{\rightarrow}(\mathbf{x}_i, t_i) &= \{\kappa \in \mathscr{A}_i^{\mathrm{keep}} : t_\kappa = t_i,\ \exists\, \mathbf{y},\ (\mathbf{x}_i, \mathbf{y}) \in D_\kappa\}, \\ \mathscr{C}_i(\mathbf{x}_i) &= \{C \in \mathscr{C}_{i-1} : \mathbf{x}_i \in C\}. \end{aligned} \tag{55}$$

The outgoing set includes only actions with a permitted access label. Retained-action validity and root connectivity are separate conditions: $\mathscr{A}_i^{\mathrm{keep}} \neq \varnothing$ can hold while $\mathscr{A}_i^{\rightarrow}(\mathbf{x}_i, t_i) = \varnothing$. In that case, construct the missing root connections on

$$\mathcal{N}_i \leftarrow \mathcal{N}_i \cup \mathscr{C}_i(\mathbf{x}_i) \cup \mathrm{Nbr}_{\mathrm{face}}(\mathscr{C}_i(\mathbf{x}_i)), \tag{56}$$

where $\mathrm{Nbr}_{\mathrm{face}}$ returns interface-adjacent carriers. Valid actions outside this neighborhood remain available.

Let $\mathcal{K}_i^-$ be the retained subcomplex, $\mathcal{L}_i$ the subcomplex to replace, and $\mathcal{B}_i$ their shared boundary subcomplex. A replacement complex $\mathcal{R}_i$ satisfies the following:

$$\begin{aligned} \mathcal{K}_{\delta_x, i-1} = \mathcal{K}_i^- \cup \mathcal{L}_i, \quad \mathcal{K}_i^- \cap \mathcal{L}_i = \mathcal{B}_i, \\ |\mathcal{K}_i^-| \cap |\mathcal{R}_i| = |\mathcal{B}_i|, \qquad |\mathcal{R}_i| \subseteq \mathcal{X}_{\mathrm{free}}, \end{aligned} \qquad \mathcal{K}_{\delta_x, i} = \mathcal{K}_i^- \cup \mathcal{R}_i. \tag{57}$$

The notation $|\cdot|$ denotes geometric realization. The retained and replacement complexes induce the same triangulation on $|\mathcal{B}_i|$, and their simplices intersect only along a common boundary face:

$$\sigma \in \mathcal{K}_i^-,\ \rho \in \mathcal{R}_i \quad \Longrightarrow \quad \sigma \cap \rho \in \{\varnothing\} \cup \{\beta : \beta \in \mathcal{B}_i,\ \beta \preceq \sigma,\ \beta \preceq \rho\}. \tag{58}$$

The subface and common-intersection conditions preserve the simplicial-complex structure.

Let $\mathscr{C}_i^-$ contain carriers outside the replaced neighborhood. The retained front and unresolved domain records in $\mathscr{B}_i^-$ and $\mathscr{Q}_i^-$ satisfy $(\{R\} \cup \mathrm{AncRec}(R)) \cap \mathscr{D}_i = \varnothing$. The epoch transition is

$$\begin{aligned} \mathfrak{U}_i : \quad \mathscr{C}_i &= \mathscr{C}_i^- \cup \mathscr{C}_i^+, \quad \mathscr{A}_i = \mathscr{A}_i^- \cup \mathscr{A}_i^+, \\ \mathscr{B}_i &= \mathscr{B}_i^-, \qquad\quad \mathscr{Q}_i = \mathscr{Q}_i^- \cup \mathscr{Q}_i^+ \cup \mathrm{Reopen}(\mathscr{D}_i). \end{aligned} \tag{59}$$

For a fixed posterior, initial state, terminal condition, and cost, augmenting the represented feasible trajectory set from $\mathscr{P}$ to $\mathscr{P}^+$ preserves the following order:

$$\mathscr{P} \subseteq \mathscr{P}^+ \quad \Longrightarrow \quad \inf_{\Pi \in \mathscr{P}^+} J[\Pi] \leq \inf_{\Pi \in \mathscr{P}} J[\Pi].$$

The update defines $\mathscr{M}_i$, retains valid records, and rebuilds affected action domains and missing root or terminal connections. Algorithm 2 preserves conforming interfaces and departure-time closure; REBASE recomputes root-dependent costs and bounds.

A contraction of the occupied region can admit a previously excluded transfer. Collision-excluded domains intersecting Δ_i are therefore reopened together with invalidated domains. Exclusions depending only on unchanged dynamics, terminal conditions and deadline remain valid.

3.2.3. Overall Replanning Procedure

Algorithm 3 coordinates posterior updates, propagation of the current spacecraft state, and replanning. Algorithm 2 constructs the carrier–action archive. DAR then searches the current action model and returns a trajectory with its control sequence through Algorithm 1. The archive and dependency records retain reusable information between planning epochs.

Each planning query solves an open-loop trajectory-search problem and returns a nominal state trajectory with its realizing control sequence. Prediction updates initiate replanning from the propagated state. Subsequent queries revalidate affected actions and descendant fronts, reuse records with unchanged supporting conditions, and update arrival costs at the new root.

The initial state and time are assigned in line 1.

UPDATEPREDICTION updates the occupied regions that define the collision constraints by using the simulated observations available up to t_i, denoted by $\mathcal{Z}_{\leq t_i}$ (Algorithm 3, lines 3–6). The current state $\mathbf{x}_i$ is derived from the preceding simulated propagation. Before planning, the current data must satisfy $t_i \leq T_{\max}$, $\mathbf{x}_i \in \mathcal{X}_{\text{free}}$, $t_i \in \mathbb{I}_i$, and $P_r\mathbf{x}_i \notin \widehat{\mathcal{O}}_i(t_i)$. Unavailable data or an unverified root gives $\mathfrak{q}_i = \varnothing$. A valid state in $\mathcal{X}_g$ completes the task, including arrival at $T_{\max}$. Otherwise, the cycle continues only with remaining maneuver time.

The first prediction initializes CONSTRUCTARCHIVE (Algorithm 2). Subsequent calls apply LOCALUPDATE from Section 3.2.2 to the affected representation and current boundary data.

REBASE initializes zero-cost root fronts $\mathcal{B}_i^{\text{root}}$ at the propagated state. Each $b_i(a)$ corresponds to a permitted root access label $a \in \mathsf{A}_i^0 \subseteq \mathsf{A}$. The collection starts at $\mathcal{B}_0^{\text{root}}$ and is passed to DAR as $\mathcal{B}^{\text{root}}$. Here, i counts posterior updates and ι_i identifies the fixed instance of the current call. Old arrival values and bounds are retained if valid based on this root. The previous unexecuted suffix supplies Π_i^{seed} when it starts at $(\mathbf{x}_i, t_i)$, reaches $\mathcal{X}_g$ by $T_{\max}$, and passes the current full-interval constraints. Otherwise, the seed is empty. DAR returns the remaining plan, its realizing controls, and bounds L_i, U_i, while preserving reusable records for the next epoch.

A certified terminal plan permits execution up to $t_i^{\text{end}} = \min\{t_i + T_{\text{exe}}, t_f(\Pi_i^\star)\}$, where T_{exe} is the replanning interval. Its restricted control u_i^{exe} propagates the next root through $\phi_{t_i, t_i^{\text{end}}}$, and the interval and control are appended to $\mathscr{H}^{\text{exe}}$. The next prediction uses observations available at $t_{i+1} = t_i^{\text{end}}$. Budget exhaustion without a certified terminal plan returns INCOMPLETE and suspends the execution. The asymptotic guarantees apply to the separate fixed-instance refinement sequence in Section 4.

4. Theoretical Analysis

The posterior, current root, terminal set, and deadline are fixed, and the resulting Problem 1 instance is denoted by P_i. Its declared action model $\mathscr{M}$ represents the feasible root-to-goal chains $\Sigma_{\mathscr{M}}$ with value $J^\star_{\mathscr{M}} = \inf_{\varsigma \in \Sigma_{\mathscr{M}}} J_\varsigma$. The construction verifies its full-interval feasibility on compact predecessor relations with nonnegative lower-semicontinuous costs and positive durations $\tau_\kappa \geq h_{\min} > 0$. Algorithm 1 preserves the complete candidate coverage through unresolved domain retention and successor reopening. The refinement additionally requires feasible trajectory approximation and vanishing representation and search errors.

Bellman propagation uses compact relations and lower-semicontinuous costs. Exact active-face evaluation additionally uses affine control maps, polyhedral domains, and quadratic effort; asymptotic convergence requires feasible recovery, coverage, and vanishing search errors.

Compact feasible parameter domains and closed face restrictions ensure compactness; positive durations and finite represented times bound chain depth. The finite exact result applies to exhaustively resolved polyhedral–quadratic domains. For nonlinear domains, Equations (44)–(46) specify the covering and continuity conditions for convergent approximate resolution. In the flow-based arguments, let n_x denote the complete state dimension; the spacecraft formulation has $n_x = 6$. The dynamics enter through $\phi_{t,s}$ and the stated endpoint-recovery estimates. All state and control errors below are based on fixed dimensionless coordinates defined using positive reference scales.

For the source relation $\mathcal{D}_{b\kappa}$, let $\mathcal{Y}_{b\kappa}(q) = \{y : (y, q) \in \mathcal{D}_{b\kappa}\}$ and define the following:

$$(\mathcal{T}_{b\kappa} V)(q) = \inf_{y \in \mathcal{Y}_{b\kappa}(q)} \{V(y) + A_\kappa(y, q)\}. \tag{60}$$

Theorem 1 (Bellman-front construction)**.** *On a compact relation with finite nonnegative lower-semicontinuous parent and action costs, the minimum is attained at every feasible exit, and the exit value is lower-semicontinuous. For two parent values on the same relation, the operator satisfies*

$$\begin{aligned} V \le W \implies \mathcal{T}V \le \mathcal{T}W, \\ \|\mathcal{T}V - \mathcal{T}W\|_\infty \le \|V - W\|_\infty < \infty. \end{aligned} \tag{61}$$

The norms are restricted to finite values on the common feasible domains. The exhaustive exact active-face construction in the convex polyhedral–quadratic branch produces the local value as a finite lower envelope.

Proof. Each nonempty compact predecessor set admits a minimum of its lower-semicontinuous objective. For $q_n \to q$, select a subsequence attaining the lower limit of the optimal values. Compactness gives convergent minimizing pairs, hence

$$\mathcal{T}V(q) \le V(y) + A_\kappa(y, q) \le \liminf_n \mathcal{T}V(q_n). \tag{62}$$

Monotonicity follows by taking the infima. If $c = \|V - W\|_\infty$, then $W - c \le V \le W + c$, hence $\mathcal{T}W - c \le \mathcal{T}V \le \mathcal{T}W + c$. In the quadratic branch, every minimizer is stationary on the relative interior of an active face. Enumerating all faces covers that minimizer. Convex feasible lifting preserves each candidate, thus

$$\operatorname{proj}_q \mathcal{D}_{b\kappa} = \bigcup_{I,\Delta} \operatorname{dom} P_{I,\Delta}, \qquad \mathcal{T}V = \min_{I,\Delta} V_{I,\Delta}. \tag{63}$$

The endpoint-compatible action construction following Equation (12) supplies affine control maps. On each projected simplex, Equation (37) gives an affine predecessor map. Composition preserves affine dependence. Starting from the zero-cost root, induction over the actions gives $\mathbf{U}_j(s) = C_j s + d_j$ on each candidate cell, where s denotes the independent parent coordinates. With the quadratic control-effort weights $H_{u,j} \succeq 0$, the Hessians satisfy

$$\mathsf{P}_b = \sum_j C_j^T H_{u,j} C_j \succeq 0, \qquad H_{\xi\xi} = E_\xi^T \mathsf{P}_b E_\xi + \mathsf{C}_\xi^T H_u \mathsf{C}_\xi \succeq 0.$$

The Hessian calculation applies to each candidate separately; the arrival-cost front is their pointwise lower envelope. General nonlinear control maps are handled by the validated domain procedure in Section 3.1.7.

For stationary lifts ξ_1, ξ_2 on face I at a fixed exit η, set $d = \xi_2 - \xi_1 = N_I c$, since $A_I d = 0$. Stationarity at ξ_1 and subtraction of the two stationarity equations give

$$\nabla_{\xi}\Psi(\xi_1,\eta)^T d = c^T N_I^T \nabla_{\xi}\Psi(\xi_1,\eta) = 0,$$
$$N_I^T H_{\xi\xi} d = 0, \qquad d^T H_{\xi\xi} d = c^T N_I^T H_{\xi\xi} d = 0.$$

The exact quadratic expansion therefore yields

$$\Psi(\xi_2,\eta) - \Psi(\xi_1,\eta) = \nabla_{\xi}\Psi(\xi_1,\eta)^T d + \tfrac{1}{2} d^T H_{\xi\xi} d = 0. \tag{64}$$

Quadraticity and stationarity give equal values; positive semidefiniteness gives minimality on the affine face. The compact polyhedral stationary relation preserves feasibility and stationarity under the barycentric lifts in Equation (37). Exhaustive face enumeration therefore yields the exact lower envelope, including singular $H_{\xi\xi}$. □

With static reversible metric actions and the Basic S* parent front $V(y) = d(o_b, y)$ on a face F, the same operator reduces to the following:

$$\mathcal{T}V(q) = \min_{y \in F} \{ d(o_b, y) + d(y, q) \}. \tag{65}$$

Equation (65) recovers the Basic S* distance-through-simplex update with proper-face restriction [23]. General spacecraft actions retain their direction-dependent costs. The inclusion $K_{t,\le r} \subseteq \widehat{K}_{t,\le r}$ makes terminal screening a sound exclusion test, with surviving domains passed to action verification.

Theorem 2 (Finite-model feasibility and optimality)**.** *In the exact polyhedral–quadratic branch, assume the declared source–exit scopes are finite, all the feasible continuations are represented or soundly excluded, and each local and terminal domain is resolved exhaustively in finitely many operations. Assume further that every dominance removal preserves all the admissible continuations at no greater cost. The DAR run to exhaustive resolution with fair processing and zero stopping tolerance terminates at the finite-model optimum when a terminal chain exists. Before termination, its bounds, once a finite feasible incumbent is available, satisfy the following:*

$$L^{(r)} \le J^{\star}_{\mathscr{M}} \le U^{(r)}, \qquad 0 \le U^{(r)} - J^{\star}_{\mathscr{M}} \le U^{(r)} - L^{(r)}. \tag{66}$$

Every finite incumbent is feasible for P_i*. The exhaustive exclusion establishes the infeasibility of the declared model* $\mathscr{M}$*.*

Proof. Let θ_k be the stored parameter of segment k, and let $u_k = u_{\mathbf{U}_{\kappa_k}(\theta_k)}$ denote its realizing control. Its realized cost is $J_{\kappa_k}(\theta_k)$. The controls produce matching state–time endpoints and additive costs:

$$\begin{aligned} x_{k+1} &= \phi_{t_k, t_{k+1}}(x_k, u_k), \quad t_{k+1} = t_k + \tau_{\kappa_k}, \\ J[\Pi] &= \sum_{k=0}^{m-1} J_{\kappa_k}(\theta_k), \quad m \le M_{\max} := \left\lfloor \frac{T_{\max} - t_s}{h_{\min}} \right\rfloor. \end{aligned} \tag{67}$$

At exact local resolution, the selected control satisfies $J_{\kappa_k}(\theta_k) = A_{\kappa_k}(x_k, x_{k+1})$. During partial resolution, the stored control and its actual cost supply a feasible candidate. The full-interval action verification and the matching root, intermediate, and terminal states lead to a feasible concatenation. Initially, the root scopes cover all first actions. The processing replaces a scope only with resolved candidates, sound exclusions, and its unresolved remainder. The scope activation inserts all compatible successors and terminal restrictions,

while an improved parent reopens the dependent scopes. Induction over these operations preserves the coverage of every not-yet-excluded complete chain. The whole-domain dominance argument above preserves every admissible continuation at no greater total cost. Thus, OPEN covers every unresolved continuation with a valid lower bound. As such,

$$L^{(r)} = \min\{U^{(r)}, \inf_{R\in \mathrm{OPEN}^{(r)}} \underline{J}(R)\} \le \inf_{\varsigma\in\Sigma_{\mathscr{M}}} J_\varsigma \le U^{(r)}. \tag{68}$$

Theorem 1 supplies exact local values. With $N_{\mathrm{fam}} = |\mathscr{F}|$ declared action families and chain length at most $M_{\max}$, the number of action-index sequences is bounded by

$$\sum_{m=0}^{M_{\max}} {N_{\mathrm{fam}}}^m < \infty$$

Let $\mathcal{C}_m$ be the candidate family at depth m. The root family is finite, and each parent–action pair produces finitely many candidates by active-face enumeration and projected triangulation. Induction to depth $M_{\max}$ makes $\mathcal{C} = \bigcup_m \mathcal{C}_m$ finite. Each decrease in the cost front introduces a previously unavailable member of $\mathcal{C}$; duplicate identities reuse stored results and dominance removal preserves the envelope. Thus only finitely many improving versions occur. Forward-time dependencies then admit finitely many reopenings following cost improvements. Finite local resolution and fair processing establish termination, while the coverage invariant gives $L = U = J^\star_{\mathscr{M}}$ or excludes every represented terminal chain. □

For a completed lower- and upper-value sweep on common feasible domains, choose δ_r uniformly over every propagation stage and parent record on the compact relations $\mathcal{D}_{b\kappa}$, or their feasible coordinate domains $Q_l \times Z_l$. Choose ε_r uniformly over the compact terminal restrictions. Refining Equation (46) to these common tolerances gives $\delta_r, \varepsilon_r \to 0$, with every upper value retaining its realizing control and validated cost.

With exact root initialization, let $E_{k,r}$ be the maximum lower- or upper-value error after k propagation stages. Monotonicity and nonexpansiveness give

$$E_{0,r} = 0, \qquad E_{k+1,r} \le E_{k,r} + \delta_r, \qquad E_{k,r} \le k\delta_r.$$

After terminal resolution, each one-sided error is at most $\varepsilon^{(r)}_{\mathrm{loc}} = M_{\max}\delta_r + \varepsilon_r$. For the fixed model,

$$0 \le U^{(r)} - L^{(r)} \le 2\varepsilon^{(r)}_{\mathrm{loc}}, \qquad \varepsilon^{(r)}_{\mathrm{loc}} \longrightarrow 0. \tag{69}$$

For a changing model sequence, assign the local and terminal tolerances from the positive target $\epsilon_{S,n} \to 0$ as

$$\delta_n \le \frac{\epsilon_{S,n}}{4(M_{\max,n}+1)}, \qquad \varepsilon_n \le \frac{\epsilon_{S,n}}{4}.$$

The same completed-sweep estimate then gives $2(M_{\max,n}\delta_n + \varepsilon_n) \le \epsilon_{S,n} \to 0$, including the growth in the number of propagation stages.

Lemma 1 (Convergence of the represented optimum)**.** *Assume $J^\star_i < \infty$, and let $\mathscr{M}_n$ be a sequence of the feasible trajectory representations. For every $\epsilon > 0$, suppose a feasible reference Π_ϵ satisfies $J[\Pi_\epsilon] \le J^\star_i + \epsilon$ and has represented feasible realizations $\mathfrak{R}_{\mathscr{M}_n}\Pi_\epsilon$ whose costs converge to $J[\Pi_\epsilon]$. Here, $\mathfrak{R}_{\mathscr{M}_n}$ denotes the feasible representation map at level n. Then,*

$$J^\star_{\mathscr{M}_n} \ge J^\star_i, \qquad \lim_{n\to\infty} J^\star_{\mathscr{M}_n} = J^\star_i. \tag{70}$$

Proof. For each fixed $\epsilon > 0$, a feasible inclusion produces $J_i^\star \le J_{\mathscr{M}_n}^\star \le J[\mathfrak{R}_{\mathscr{M}_n}\Pi_\epsilon]$. Taking the upper limit yields the following:

$$J_i^\star \le \liminf_{n\to\infty} J_{\mathscr{M}_n}^\star \le \limsup_{n\to\infty} J_{\mathscr{M}_n}^\star \le J[\Pi_\epsilon] \le J_i^\star + \epsilon. \tag{71}$$

Taking the limit as $\epsilon \to 0^+$ proves the claim. A cost-preserving representation of an attained optimum also leads to equality at a finite level. □

In particular, let $\mathscr{P}_{\mathrm{v}}$ contain trajectories restricted to the vertex joins, $\mathscr{P}_{\mathscr{M}}$ those represented by continuous joins, and $\mathscr{P}_i$ the feasible trajectories of P_i. Retaining the same vertex actions under the same constraints leads to the following:

$$\begin{gathered} \mathscr{P}_{\mathrm{v}} \subseteq \mathscr{P}_{\mathscr{M}} \subseteq \mathscr{P}_i, \\ J_i^\star \le J_{\mathscr{M}}^\star \le J_{\mathrm{v}}^\star, \qquad J_{\mathrm{v}}^\star := \inf_{\Pi\in\mathscr{P}_{\mathrm{v}}} J[\Pi]. \end{gathered} \tag{72}$$

Taking the infima over the nested trajectory families yields the stated model-optimum ordering.

For matched feasible trajectories on a common compact domain with identical initial state and terminal time, the Lipschitz continuity of $\mathbf{f}$ and ℓ produces continuous dependence on the control. If $\eta_{\mathscr{M}} = \|\mathbf{u}_{\mathscr{M}} - \mathbf{u}_\epsilon\|_{L^1} \to 0$, Gronwall's inequality and integration of the running cost imply that

$$\|\mathbf{x}_{\mathscr{M}} - \mathbf{x}_\epsilon\|_\infty \to 0, \qquad J[\mathfrak{R}_{\mathscr{M}}\Pi_\epsilon] \to J[\Pi_\epsilon]. \tag{73}$$

For a Lipschitz reference control, the piecewise-constant approximation with maximum interval h has L^1 error $O(h)$. Thus, a feasible endpoint correction whose L^1 norm also tends to zero produces $\eta_{\mathscr{M}} \to 0$.

Then, Equation (72) extends the approximation result to continuous-predecessor search.

Proposition 1 (Feasible trajectory approximation)**.** *Fix a feasible reference trajectory* $(\mathbf{x}_\epsilon, \mathbf{u}_\epsilon)$ *on* $[t_s, t_f]$ *with Lipschitz control and positive margins to the control boundary and forbidden state–time set. On a compact neighborhood of this trajectory, write the Lipschitz conditions for the declared dynamics and running cost as*

$$\begin{aligned} \|\mathbf{f}(x,u,t) - \mathbf{f}(\widetilde{x},\widetilde{u},t)\| &\le L_x\|x-\widetilde{x}\| + L_u\|u-\widetilde{u}\|, \\ |\ell(x,u,t) - \ell(\widetilde{x},\widetilde{u},t)| &\le L_{\ell x}\|x-\widetilde{x}\| + L_{\ell u}\|u-\widetilde{u}\|, \end{aligned}$$

uniformly in t. Let the endpoint maps be continuously differentiable with Lipschitz derivatives. At level n, represented times are selected with durations $\tau_{n,k} \in [c_h h_n, h_n]$, $c_h > 0$, *and represented states* $y_{n,k}$ *satisfying the following:*

$$\max_k \|y_{n,k} - \mathbf{x}_\epsilon(t_{n,k})\| \le d_n, \quad y_{n,0} = \mathbf{x}_s, \quad y_{n,m_n} = \mathbf{x}_\epsilon(t_f) \in \mathcal{X}_g. \tag{74}$$

For each reference segment, let $\rho_n > 0$ *be a uniform lower bound on its distance to the complement of the declared carrier union. The local action model includes the corresponding endpoint pairs and control partitions. The terminal point is represented on an explicit terminal domain, where the last connection is resolved continuously. Let* $\mathbf{U}_{n,k}^0$ *generate a piecewise-constant approximation with uniform control error at most* a_n. *For the endpoint map* $E_{n,k}(\mathbf{U}) = \phi_{t_{n,k},t_{n,k+1}}(y_{n,k}, u_{\mathbf{U}})$, *suppose:*

$$\mathrm{D}E_{n,k}(\mathbf{U}_{n,k}^0)\mathsf{R}_{n,k} = I_{n_x}, \qquad \|\mathsf{R}_{n,k}\| \le \gamma_n, \tag{75}$$

where $\mathsf{R}_{n,k}$ *is a right inverse of the endpoint map derivative. The coefficient norm is the maximum norm of the control blocks, and* γ_n *is uniform over the finitely many segments at level* n*. Let* M_n *bound the Lipschitz constants of these derivatives on coefficient balls of radius* b_n *about* $\mathbf{U}^0_{n,k}$*. The declared control family includes every feasible coefficient vector in these balls. With* C_0 *independent of* n, k*, define*

$$\begin{gathered} r_n = C_0(d_n + h_n a_n), \qquad 2\gamma_n r_n \le b_n, \qquad 2M_n\gamma_n^2 r_n \le \tfrac{1}{2}, \\ e_n = d_n + a_n + 2\gamma_n r_n \to 0, \qquad e_n/\rho_n \to 0. \end{gathered} \tag{76}$$

The right-inverse and derivative bounds, together with Equation (76), define the endpoint-correction test for the selected flow and control partition. For all sufficiently large n*, there is a represented feasible chain* $\Pi_{n,\epsilon}$ *with the exact prescribed endpoints, and*

$$\begin{gathered} \|\mathbf{x}_{n,\epsilon} - \mathbf{x}_\epsilon\|_\infty = O(e_n), \qquad \|\mathbf{u}_{n,\epsilon} - \mathbf{u}_\epsilon\|_\infty \le e_n, \\ J[\Pi_{n,\epsilon}] \to J[\Pi_\epsilon]. \end{gathered} \tag{77}$$

The implicit constant may depend on the fixed reference but not on n*.*

Proof. The flow estimate for the declared dynamics gives

$$\begin{aligned} \|E_{n,k}(\mathbf{U}^0_{n,k}) - y_{n,k+1}\| &\le (1 + e^{L_x h_n})d_n + L_u h_n e^{L_x h_n} a_n \\ &\le C_0(d_n + h_n a_n) = r_n, \end{aligned}$$

with C_0 independent of n, k. Suppress the indices and write $\mathbf{e}_0 = E(\mathbf{U}^0) - y^+$. For $\mathbf{e}_0 = 0$, take $\mathbf{U}^* = \mathbf{U}^0$. Otherwise define, on $\|z\| \le 2\|\mathbf{e}_0\|$,

$$\mathcal{T}_{\mathrm{corr}}(z) = z - [E(\mathbf{U}^0 + \mathsf{R}z) - y^+]. \tag{78}$$

The coefficient-ball condition keeps $\mathbf{U}^0 + \mathsf{R}z$ in the domain of the derivative bound. Since $\mathrm{D}E(\mathbf{U}^0)\mathsf{R} = I_{n_x}$,

$$\begin{aligned} \mathrm{D}\mathcal{T}_{\mathrm{corr}}(z) &= [\mathrm{D}E(\mathbf{U}^0) - \mathrm{D}E(\mathbf{U}^0 + \mathsf{R}z)]\mathsf{R}, \\ \|\mathrm{D}\mathcal{T}_{\mathrm{corr}}(z)\| &\le M_n\gamma_n^2\|z\| \le \tfrac{1}{2}, \\ \|\mathcal{T}_{\mathrm{corr}}(z)\| &\le \|\mathbf{e}_0\| + \tfrac{1}{2}\|z\| \le 2\|\mathbf{e}_0\|. \end{aligned}$$

Thus, $\mathcal{T}_{\mathrm{corr}}$ is a contraction of the closed ball into itself. Its fixed point gives $\mathbf{U}^* = \mathbf{U}^0 + \mathsf{R}z^*$ and

$$E(\mathbf{U}^*) = y^+, \qquad \|\mathbf{U}^* - \mathbf{U}^0\| \le 2\gamma_n r_n. \tag{79}$$

Set

$$\Delta u_n = a_n + 2\gamma_n r_n, \qquad \Delta x_n = e^{L_x h_n}[d_n + L_u h_n \Delta u_n].$$

The block-maximum coefficient norm bounds the uniform control error by Δu_n, and Gronwall's inequality bounds each segment's state error by Δx_n. Exact endpoint matching uses the same joining state in both neighboring segments, so these estimates hold on the concatenated trajectory. Let $\mu_u > 0$ and $\mu_x > 0$ denote the reference control and state–time margins, with μ_x also chosen within the compact-neighborhood reserve. Since $e_n \to 0$ and $e_n/\rho_n \to 0$, for sufficiently large n,

$$\Delta u_n < \mu_u/2, \qquad \Delta x_n < \tfrac{1}{2}\min\{\mu_x, \rho_n\}.$$

Applying the estimate up to a possible first exit from the neighborhood excludes that exit. The same inequalities preserve the control constraints, avoidance of $\mathcal{O}_i(t)$, and carrier

containment. The endpoint correction has L^1 norm at most $2(t_f - t_s)\gamma_n r_n \to 0$, and the total control error satisfies $\|\mathbf{u}_{n,\epsilon} - \mathbf{u}_\epsilon\|_{L^1} \le (t_f - t_s)\Delta u_n \to 0$. Finally,

$$|J[\Pi_{n,\epsilon}] - J[\Pi_\epsilon]| \le (t_f - t_s)(L_{\ell x}\Delta x_n + L_{\ell u}\Delta u_n) \longrightarrow 0.$$

The retained neighboring actions and explicit terminal connection therefore realize the feasible chain and the bounds in Equation (77). □

The refinement parameters are linked through the endpoint sensitivity of the declared flow. Fix a reference trajectory and its compact neighborhood, and choose the uniform constants on that neighborhood. Let d_n be the state-cover radius, $h_n \to 0$ the maximum action duration, and s_n the independent control-partition width. The reference control has $a_n \le C_a s_n$, and the endpoint bounds satisfy $\gamma_n \le C_\gamma h_n^{-p}$ and $M_n \le \overline{M}_n$, where $p > 0$ belongs to the chosen endpoint-map class. Choose the carrier-clearance lower bound with $0 < \rho_n \le \bar{\rho} < \infty$. The sufficient resolution rule is

$$\begin{aligned} d_n = o(\rho_n h_n^p), \qquad s_n = o\Big(\rho_n \min\{1, h_n^{p-1}\}\Big), \\ \overline{M}_n h_n^{-2p}(d_n + h_n s_n) \to 0, \qquad 2\gamma_n r_n \le b_n. \end{aligned} \tag{80}$$

Substituting $a_n \le C_a s_n$ and $\gamma_n \le C_\gamma h_n^{-p}$ into the definition of e_n gives

$$\frac{e_n}{\rho_n} \le \frac{d_n}{\rho_n} + C_a\frac{s_n}{\rho_n} + C\left(\frac{d_n}{\rho_n h_n^p} + \frac{h_n^{1-p}s_n}{\rho_n}\right) \longrightarrow 0.$$

The first two conditions in Equation (80) make all four terms vanish, because both $\min\{1, h_n^{p-1}\}$ and $h_n^{1-p}\min\{1, h_n^{p-1}\}$ are at most one for $h_n \le 1$. The constants are uniform in n for the fixed reference, and $\rho_n \le \bar{\rho}$ also yields $e_n \to 0$. Moreover,

$$2M_n\gamma_n^2 r_n \le C''\overline{M}_n h_n^{-2p}(d_n + h_n s_n) \longrightarrow 0.$$

Together with the explicit coefficient-ball condition, this establishes the smallness conditions of Proposition 1 for sufficiently fine levels.

All neighboring transfers within the reference segment reach are retained with connection radius at least $c_f h_n + 2d_n$, where c_f bounds $\|\mathbf{f}\|$ in the scaled coordinates. Their carrier unions retain the prescribed clearance ρ_n, and terminal intersections are inserted explicitly. State coverage is refined relative to this clearance, with the endpoint sensitivity coordinating the time and control resolutions.

Application example: endpoint regularity for fully actuated relative motion. The planar and orbital problems considered here have the fully actuated form $\dot{\mathbf{r}} = \mathbf{v}$, $\dot{\mathbf{v}} = g(\mathbf{r}, \mathbf{v}, t) + \mathbf{u}$. Let E_h map the control coefficients to the complete endpoint after duration h. For the planar double integrator in Equation (88), $g = 0$. Independent two-axis perturbations a, b on the two half-intervals give the exact derivative

$$\frac{\partial E_h}{\partial(a,b)} = \begin{bmatrix} \frac{3}{8}h^2 I_2 & \frac{1}{8}h^2 I_2 \\ \frac{1}{2}hI_2 & \frac{1}{2}hI_2 \end{bmatrix}.$$

Its rank is four for every $h > 0$, and the affine endpoint map has $M_n = 0$. For orbital motion, the CW and time-varying J_2-gradient drifts are smooth and affine in the relative state. Nonlinear central-gravity–J_2 relative dynamics are twice continuously differentiable on compact orbital neighborhoods separated from the central-body singularity. Their

derivatives are bounded on the finite planning horizon. With independent three-axis perturbations a, b over the two half-intervals, variation of constants gives

$$\begin{aligned} \frac{\partial E_h}{\partial(a,b)} &= \begin{bmatrix} \frac{3}{8}h^2 I_3 + O(h^3) & \frac{1}{8}h^2 I_3 + O(h^3) \\ \frac{1}{2}h I_3 + O(h^2) & \frac{1}{2}h I_3 + O(h^2) \end{bmatrix}, \\ \operatorname{diag}(h^{-2}I_3, h^{-1}I_3)\frac{\partial E_h}{\partial(a,b)} &= \begin{bmatrix} \frac{3}{8}I_3 & \frac{1}{8}I_3 \\ \frac{1}{2}I_3 & \frac{1}{2}I_3 \end{bmatrix} + O(h). \end{aligned} \tag{81}$$

Integrating the leading position and velocity kernels, $(h-s)I_3$ and I_3, over the two half-intervals gives these coefficients. Bounded derivatives of g control the remainders. Each limiting scalar block has determinant $1/8$. Hence, for sufficiently small h, the endpoint derivative has row rank six and a right inverse with $\gamma_h = O(h^{-2})$ in the fixed dimensionless coordinates. The two independent control blocks supply the six perturbation coordinates required for full-state correction.

For the fully actuated class, suppose M_n is uniformly bounded, $0 < \rho_n \le \bar{\rho}$, and the coefficient balls retain a radius $b_n \ge b_0 > 0$ within the control reserve. For a carrier-cover sequence satisfying Proposition 1, one explicit choice in the scaled coordinates is

$$d_n = \rho_n h_n^5, \qquad s_n = \rho_n h_n^4, \qquad h_n \longrightarrow 0.$$

Since $a_n \le C_a s_n$ and $\gamma_n = O(h_n^{-2})$, the endpoint residual bound satisfies $r_n = O(\rho_n h_n^5)$. Consequently,

$$\begin{aligned} \frac{e_n}{\rho_n} &= O(h_n^3) \longrightarrow 0, \\ 2M_n\gamma_n^2 r_n &= O(\rho_n h_n) \longrightarrow 0, \\ \frac{2\gamma_n r_n}{b_n} &= O(\rho_n h_n^3/b_0) \longrightarrow 0. \end{aligned}$$

Bounded carrier clearance also gives $e_n \to 0$. The half-interval perturbations remain available under control-block subdivision. Both planar and orbital dynamics therefore satisfy the local rank condition with $\gamma_h = O(h^{-2})$. For robust reference trajectories, the prescribed carrier cover, terminal connections and refinement rule give feasible approximations. A robust near-optimal reference family and vanishing search gaps then yield Theorem 3. Section 5.7 reports finite-resolution orbital endpoint checks.

The following theorem uses one refinement sequence $\mathscr{M}_n$ for all feasible comparison trajectories Π_ϵ.

Theorem 3 (Asymptotic optimality under feasible refinement). *P_i is fixed with finite optimum. For every $\epsilon > 0$, let a reference trajectory satisfy $J[\Pi_\epsilon] \le J_i^\star + \epsilon$ and the regularity, margin, and endpoint-rank conditions of Proposition 1. Assume that the declared refinement satisfies the coverage and correction conditions. For every fixed $\epsilon > 0$, the refinement errors and finite feasible DAR bounds satisfy the following relations as $n \to \infty$:*

$$\begin{aligned} h_n \to 0, \qquad e_{n,\epsilon} \to 0, \\ 0 \le U_n - L_n \le \epsilon_{S,n}, \qquad \lim_{n\to\infty} \epsilon_{S,n} = 0. \end{aligned} \tag{82}$$

Then $U_n \to J_i^\star$.

Proof. For fixed $\epsilon > 0$, Proposition 1 supplies a represented feasible chain $\Pi_{n,\epsilon}$ with recovery error

$$e_{n,\epsilon} = d_n + a_n + 2C_0\gamma_n(d_n + h_n a_n) \longrightarrow 0.$$

The recovery conditions in Equation (76) give this convergence. Equation (80) provides a sufficient refinement rule under the stated endpoint-sensitivity bounds. The Lipschitz cost estimate gives

$$|J[\Pi_{n,\epsilon}] - J[\Pi_\epsilon]| \le C_\epsilon e_{n,\epsilon},$$

where C_ϵ is independent of n. Feasible inclusion yields

$$J_i^\star \le J_{\mathscr{M}_n}^\star \le J[\Pi_{n,\epsilon}] \le J_i^\star + \epsilon + C_\epsilon e_{n,\epsilon}.$$

The approximation condition of Lemma 1 and the search interval therefore give

$$\begin{aligned} 0 \le U_n - J_{\mathscr{M}_n}^\star \le U_n - L_n \le \epsilon_{S,n}, \\ 0 \le U_n - J_i^\star \le \epsilon_{S,n} + \epsilon + C_\epsilon e_{n,\epsilon}. \end{aligned} \tag{83}$$

Taking $\limsup_{n\to\infty}$ gives an error at most ϵ for every $\epsilon > 0$ along the same sequence. Since $U_n \ge J_i^\star$, letting $\epsilon \to 0$ proves the claim. □

Let $\mathcal{E}_{n,m}^{\mathrm{cov}}$ denote coverage and successful resolution for the reference with $\epsilon = 1/m$. If $\sum_n \Pr((\mathcal{E}_{n,m}^{\mathrm{cov}})^c) < \infty$ for every positive integer m, Borel–Cantelli and the countable intersection over m give eventual recovery almost surely. Equation (83) then gives almost-sure cost convergence. The following proposition supplies a sufficient coverage test.

Proposition 2 (Probabilistic completeness of stochastic refinement)**.** *Let $\mathcal{N}_{n,k}^{\mathrm{cov}}$, $k = 1, \dots, K_n$, be ordered positive-measure neighborhoods of a robust feasible reference in the regular state and free-control coordinates at represented times. Use state radii at most d_n and free-control neighborhoods whose represented approximations have error at most a_n. The endpoint-correction and carrier-margin bounds of Proposition 1 hold uniformly for neighboring selections. Retain the corresponding neighboring transfers and resolve their whole-interval verification. The randomized construction samples these coordinates with density bounded below by a positive constant. The exact root and terminal intersections are inserted deterministically, and the covered models are resolved to the prescribed tolerance. If $\Pr\{\mathcal{N}_{n,k}^{\mathrm{cov}}$ absent$\} \le \alpha_{n,k}$ and $\lim_{n\to\infty} \sum_k \alpha_{n,k} = 0$, the exhaustive finite search succeeds with probability tending to one.*

Proof. On the coverage event, select one state $y_{n,k}$ within d_n of the fixed reference at each represented time and share that state between adjacent segments. Use the exact root and the reference endpoint on the represented terminal domain, as specified in Equation (74). With reference state speed bounded by c_f,

$$\|y_{n,k+1} - y_{n,k}\| \le c_f h_n + 2d_n,$$

so the construction retains every required neighboring transfer. The shared nodes, represented control approximations, and uniform neighborhood bounds satisfy Proposition 1, which supplies controls $u_{n,k}$ such that

$$\phi_{t_{n,k},t_{n,k+1}}(y_{n,k}, u_{n,k}) = y_{n,k+1}.$$

The margin estimates preserve full-interval admissibility, and the stipulated verification records the corrected segments as actions. Matching state–time endpoints then form a certified terminal chain recovered by the covered-model search. Consequently, for sufficiently fine levels,

$$\Pr\{\text{search fails}\} \le \Pr\{\text{some } \mathcal{N}_{n,k}^{\mathrm{cov}} \text{ is absent}\} \le \sum_k \alpha_{n,k} \longrightarrow 0. \tag{84}$$

With N_n independent draws at refinement level n in a declared joint parameter domain, a sufficient neighborhood of probability mass $p_{n,k}$ is missed with probability at most $(1-p_{n,k})^{N_n}$. Selecting $\alpha_{n,k} = (1-p_{n,k})^{N_n}$ and $p_{n,\min} = \min_k p_{n,k}$ leads to the following:

$$\sum_{k=1}^{K_n} \alpha_{n,k} \le K_n e^{-N_n p_{n,\min}}, \qquad \lim_{n\to\infty} \left(N_n p_{n,\min} - \log K_n\right) = +\infty. \tag{85}$$

The stated limit makes the bound $\exp[-(N_n p_{n,\min} - \log K_n)]$ vanish as the neighborhoods shrink and their number grows. For a fixed recovery reference, a sufficient summable schedule is

$$N_n p_{n,\min} \ge \log K_n + (2+\gamma)\log(n+1), \qquad \gamma > 0. \tag{86}$$

The failure bound $(n+1)^{-2-\gamma}$ is summable, thus providing the events required for almost-sure convergence after complete resolution at each level. The deterministic Halton and dispersion refinement adopt the state coverage radius d_n and the carrier and control conditions in Proposition 1, and the random sampling bound applies to the randomized variant. □

The terminal velocity comparison follows directly from the feasible set inclusion. For identical dynamics, actions, and initial data, $\mathcal{X}_g^{\text{fix}} \subseteq \mathcal{X}_g^{\text{free}}$ produces the following:

$$J_i^{\star,\text{free}} \le J_i^{\star,\text{fix}}, \qquad J_{\mathcal{M}_i}^{\star,\text{free}} \le J_{\mathcal{M}_i}^{\star,\text{fix}}. \tag{87}$$

The planning feasibility is based on the CDIC occupied-region model from our previous work and its statistical interpretation [43].

As in Basic S*, the queue operations and label reopening contribute to DAR's computational cost [23] (Section 4.3). The local action evaluation and front reconstruction add problem-dependent work, which is reduced by reusing valid records.

5. Results

5.1. Simulation Models and Evaluation Protocol

5.1.1. Dynamics and Scenario Models

ISTR* is evaluated in a four-dimensional state space for planar motion and in the six-dimensional Counterflow (CF) and Moderate orbital scenarios. The planar experiment examines trajectory reconstruction following a prescribed uncertainty update. The orbital experiments evaluate mission-completion rate, final executed cost, and replanning time. Orbital target motion and CDIC posterior histories are adopted from our preceding study [43]. Both scenarios use P1's objective and constraint structure with piecewise-constant controls; the planar instance has four state and two control coordinates, and the orbital instance has six state and three control coordinates.

The planar scenario (P2D-D) uses a double-integrator model with two position coordinates and a four-dimensional position–velocity state. Let $\mathbf{r} = (x,y)^T$, $\mathbf{v} = (v_x, v_y)^T$, $\mathbf{x}_p = (\mathbf{r}^T, \mathbf{v}^T)^T$, and $\mathbf{u} = (u_x, u_y)^T$. The control input is the translational acceleration:

$$\dot{\mathbf{r}} = \mathbf{v}, \qquad \dot{\mathbf{v}} = \mathbf{u}, \qquad \dot{\mathbf{x}}_p = \begin{bmatrix} \mathbf{0}_2 & \mathbf{I}_2 \\ \mathbf{0}_2 & \mathbf{0}_2 \end{bmatrix} \mathbf{x}_p + \begin{bmatrix} \mathbf{0}_2 \\ \mathbf{I}_2 \end{bmatrix} \mathbf{u}. \tag{88}$$

Here, $\mathbf{0}_2$ and $\mathbf{I}_2$ are the 2×2 zero and identity matrices. With control held constant over $h = 0.25$ s, the state is propagated exactly as

$$\mathbf{r}_{k+1} = \mathbf{r}_k + h\mathbf{v}_k + \tfrac{1}{2}h^2\mathbf{u}_k, \qquad \mathbf{v}_{k+1} = \mathbf{v}_k + h\mathbf{u}_k. \tag{89}$$

Spacecraft relative motion is expressed in the local-vertical/local-horizontal (LVLH) orbital frame, with x, y, and z along the radial, along-track, and orbit-normal directions, respectively. The relative-motion equations combine the Clohessy–Wiltshire (CW) reference dynamics [44] with the differential acceleration $\delta\mathbf{a}_{J_2} = (\delta a_x, \delta a_y, \delta a_z)^T$ associated with Earth's second zonal gravitational harmonic coefficient, J_2:

$$\begin{aligned} \ddot{x} - 2n\dot{y} - 3n^2 x &= u_x + \delta a_x, \\ \ddot{y} + 2n\dot{x} &= u_y + \delta a_y, \qquad n = \sqrt{\mu / r_c^3}, \\ \ddot{z} + n^2 z &= u_z + \delta a_z, \end{aligned} \tag{90}$$

Here, μ is Earth's gravitational parameter and r_c is the reference-orbit radius. Linearized relative-motion models incorporating J_2 provide a useful perturbation description [45]. The diagnostic model here uses the instantaneous J_2 gravity gradient on a circular reference orbit, obtained by differentiating the inertial J_2 acceleration and applying it to $\mathbf{r} = (x, y, z)^T$:

$$\begin{aligned} \delta\mathbf{a}_{J_2} &= \frac{6\mu J_2 R_e^2}{r_c^5}\, \mathbf{M}_{J_2}(i_{\mathrm{orb}}, \vartheta)\, \mathbf{r}, \\ \mathbf{M}_{J_2} &= \begin{bmatrix} 1 - 3s_{i_{\mathrm{orb}}}^2 s_\vartheta^2 & s_{i_{\mathrm{orb}}}^2 s_{2\vartheta} & s_{2i_{\mathrm{orb}}} s_\vartheta \\ s_{i_{\mathrm{orb}}}^2 s_{2\vartheta} & -\frac{1}{4} - s_{i_{\mathrm{orb}}}^2 \left(\frac{1}{2} - \frac{7}{4} s_\vartheta^2\right) & -\frac{1}{4} s_{2i_{\mathrm{orb}}} c_\vartheta \\ s_{2i_{\mathrm{orb}}} s_\vartheta & -\frac{1}{4} s_{2i_{\mathrm{orb}}} c_\vartheta & -\frac{3}{4} + s_{i_{\mathrm{orb}}}^2 \left(\frac{1}{2} + \frac{5}{4} s_\vartheta^2\right) \end{bmatrix}, \end{aligned} \tag{91}$$

Here, R_e is Earth's equatorial radius, i_{orb} is the reference-orbit inclination, $\vartheta(t) = \vartheta_0 + nt$ is the reference argument of latitude, and $s_{(\cdot)} = \sin(\cdot)$ and $c_{(\cdot)} = \cos(\cdot)$. The zero-perturbation limit of Equation (90) recovers the time-invariant CW model. The J_2 contribution introduces time dependence through $\vartheta(t)$.

The archived orbital mission batches use the CW model with zero-order-hold (ZOH) controls for local steering and state propagation. The completion, cost, and timing statistics correspond to that configuration. Section 5.7 evaluates the full J_2 gradient model and nonlinear central-gravity–J_2 propagation through separate fidelity and local endpoint-correction tests. For each model, action generation, endpoint recovery, path verification, and terminal screening must use consistent dynamics and control partitions.

5.1.2. Baseline Policies and Evaluation Protocol

The primary comparison evaluates ISTR*, RRTX, HL-RRT*, L-GLS, and B-LGLS on the same 40 boundary-condition tasks in each scenario, giving 600 method–task evaluations across the three scenarios. A task identifier fixes the initial and terminal states, target motion, and prediction history across methods. A separate CF comparison evaluates two ISTR* reoptimization policies on the same 40 tasks, adding 80 evaluations. A mission is complete when the prescribed terminal state is reached, and the executed trajectory passes independent replay verification. Completion rates include all attempted tasks; final-cost statistics use completed missions. The empirical completion rate is $100N_{\mathrm{succ}}/40$ for each method and scenario. Two-sided 95% Wilson score intervals are computed without continuity correction. Bar heights show the empirical rates; error bars and bracketed limits quantify uncertainty in the completion probability.

At each posterior update, every method plans from the propagated spacecraft state and verifies the complete remaining trajectory against the update's full future forbidden-region union. Simulated flight pauses during planning.

All five methods use the local connection problem in Problem 2, with the same scenario dynamics, quadratic running cost, and admissibility limits. Orbital connections use the CW/ZOH endpoint model in Equation (93); planar connections use Equation (89). The

common local steering formulation supplies fixed-state transfers to the baselines and vertex realizations to DAR action families. Each policy schedules state, control, and prediction-separation checks; DAR also optimizes continuous joining states. For orbital candidates exceeding the control bound, RRTX and the lazy roadmaps invoke a bounded endpoint solve. HL-RRT* uses the bounded solve for terminal connections and accepts already bounded controls for ordinary rewiring and root reconnection.

Initial queries include online structure preparation; later queries reuse retained records. The primary orbital runs use four-way parallel local action generation for ISTR* and serial execution for the baselines. The reported completion and timing comparisons concern these implementation configurations. Offline archive construction and shared posterior estimation are accounted for separately. The methods retain their own sampling and neighborhood policies within the common scenario bounds; Table 1 and Section 5.2 specify the policy parameters.

Table 1. Baseline policy settings.

Method	Parameter	Value
HL-RRT*	$H_{\text{near}}^{\text{orbital}}$ (s)	100
HL-RRT*	$H_{\text{near}}^{\text{planar}}$ (s)	10
HL-RRT*	$(\Delta t_{\text{near}}, \Delta t_{\text{far}})$ (s)	$(20, 60)$
HL-RRT*	(p_u, p_a, p_g)	$(0.22, 0.70, 0.08)$
B-LGLS	$(w_{\text{inf}}, w_{\text{trunc}})$	$(1.2, 1.2)$

RRTX repairs affected edges and goal-rooted costs [26], with orbital radius $r_N = \min\{3, 6(\log N/N)^{1/7}\}$ for $N \geq 3$ in normalized state–time coordinates. HL-RRT* combines near-horizon eager verification, deferred distant checks, and cost-guided adaptive sampling [27]; every accepted suffix reaches the original terminal state within the mission deadline. L-GLS repairs lazily evaluated time-layered roadmaps, while B-LGLS adds estimate inflation and bounded-search truncation [28].

Table 1 lists the eager-check horizon H_{near}, near/far extension durations $(\Delta t_{\text{near}}, \Delta t_{\text{far}})$, uniform/adaptive/goal proposal probabilities (p_u, p_a, p_g), and B-LGLS inflation/truncation factors $(w_{\text{inf}}, w_{\text{trunc}})$. The proposal probabilities apply after the adaptive distribution is available. Planar root reconnection considers up to 12 retained neighbors for RRTX and HL-RRT*, while the lazy roadmaps retain up to eight neighbors in the next time layer.

Under the common running cost defined after Equation (4), the executed piecewise-constant control sequence has cost

$$\begin{aligned} J_R = J_{\text{exec}} &= \tfrac{1}{2} \sum_k \Delta t_k \mathbf{u}_k^T W_u \mathbf{u}_k \\ &= \tfrac{1}{2} \sum_k \Delta t_k \|\mathbf{u}_k\|_2^2. \end{aligned} \tag{92}$$

The planar model uses $\mathbf{u}_k \in \mathbb{R}^2$ and $W_u = I_2$; the orbital model uses $\mathbf{u}_k \in \mathbb{R}^3$ and $W_u = I_3$. Each Δt_k is the duration executed under $\mathbf{u}_k$.

An exact two-sided McNemar test compares paired completion outcomes for each orbital scenario, with Holm adjustment across the four baseline comparisons. All adjusted p-values for the archived 40-task batches are at most 1.54×10^{-8}. The tests quantify differences under the reported configurations.

For method m, final-cost medians use its completed-task set $\mathcal{S}_m$ and the final J_R in Equation (92). Paired comparisons use the common completed-task set $\mathcal{S}_a \cap \mathcal{S}_b$, matched by task identifier.

Planning-query time includes prediction preparation, search, and candidate verification, covering initial and failed calls. Report generation and post-run replay are timed separately. The software specification is C++20 with Python 3.10.11 for analysis; the reference platform is an Intel Core i9-13900HX (2.20 GHz) with 32 GB RAM.

5.1.3. Finite Implementation and Numerical Settings

The experiments use finite action families and control-based recovery from Section 3.1. Each orbital patch stores its source state and time, duration, endpoint vertices, weights, controls, and effort. Both ISTR* and the orbital baselines propagate held controls with the same CW/ZOH transition:

$$\mathbf{x}_{k+1} = \Phi_h \mathbf{x}_k + \Gamma_h \mathbf{u}_k, \qquad \mathbf{x}_{N_u} = \Phi_h^{N_u} \mathbf{x}_0 + \sum_{k=0}^{N_u-1} \Phi_h^{N_u-1-k} \Gamma_h \mathbf{u}_k. \tag{93}$$

Here, Φ_h and Γ_h are the state-transition and held-input matrices, and N_u is the number of control blocks. For the constant CW system matrix A_{CW} and acceleration-input matrix B, $\Phi_h = \exp(A_{\mathrm{CW}} h)$ and $\Gamma_h = \int_0^h \exp(A_{\mathrm{CW}}(h-s))B\,\mathrm{d}s$. Equation (93) is the exact ZOH transition of that linear model. Minimum-effort endpoint controls are checked against control, endpoint, and path constraints. For a fixed duration and partition, the affine dependence of state and control histories on endpoint data supports vertex-based family verification. Orbital search uses finite source-face updates and a deferred frontier; planar search uses detour and timing hypotheses with projected weight updates.

For a hold interval h with acceleration bound $\bar{a}$, the position chord allowance is $\epsilon_{\mathrm{chord}} = \bar{a}h^2/8 + \epsilon_{\mathrm{geom}}$, where ϵ_{geom} is the numerical position allowance in the same units as $\mathbf{r}$. A common separating plane must cover all vertex traces at both endpoints, the chord allowance, and prescribed clearance. Position and speed checks include their interval allowances. Changes in dynamics, control limits, or action domains require recomputation. General nonlinear families use the domain-based reconstruction and full-interval verification in Section 3.1.

For vertex control sequences U_j and hold duration h, the stored Gram matrix is $H_{jk} = h\sum_\ell U_j(\ell)^T U_k(\ell)$, giving effort $\frac{1}{2}\lambda^T H \lambda$. The terminal quadratic estimate enters through the endpoint map. Search nodes store state, time, accumulated cost, parent, and selected patch; a time-keyed map indexes deferred nodes. The face kernel uses at most seven weights, a 10^{-10} Hessian-symmetry tolerance and a 10^{-7} KKT-residual tolerance. Weights below -10^{-12} trigger face restriction; remaining negative weights are also restricted. Endpoint replay uses tolerance 10^{-8}. Tangent-space positive definiteness enables the fast factorization; vertex fallback retains a verified candidate if factorization fails.

Orbital search orders nodes by accumulated control effort plus a terminal control-effort estimate, with default weight one. The estimate uses a constraint relaxation with the same dynamics and endpoints. Discovery and diversity priorities further organize the search frontier. Certificates are reused for unchanged action families and recomputed for changed families. Mission runs use the finite quadratic kernel; Section 5.7 evaluates local recovery under nonlinear dynamics.

Nonlinear domain refinement follows Algorithm 1.

For a quadratic family with $N_\lambda \le 7$ weights, N_u control blocks and $N_{\mathrm{face}} \le 2^{N_\lambda} - 1$ visited faces, Gram assembly costs $O(N_u N_\lambda^2)$ and dense face solves cost $O(N_{\mathrm{face}}(N_\lambda + 1)^3)$. Verification work scales with tested intervals and supporting planes; dependency updates traverse affected records and edges.

5.2. Planar Double-Integrator Reconstruction

The planar control and speed limits are $\|\mathbf{u}\|_2 \leq 0.96$ m/s^2 and $\|\mathbf{v}\|_2 \leq 12$ m/s. The nominal initial and terminal states are $\mathbf{x}_p(0) = (-120, -60, 1.4, 0.7)^T$ and $\mathbf{x}_p(60) = (120, 60, 1.4, 0.7)^T$, with positions in meters and velocities in meters per second. Each endpoint-position component is perturbed independently within $[-1.4, 1.4]$ m. Endpoint velocities remain fixed, giving a prescribed terminal speed of approximately 1.565 m/s. Each selected plan is executed for 5 s before replanning from the propagated state. Two targets move across the workspace according to

$$\mathbf{r}_0^B(t) = \begin{bmatrix} 2\sin(0.04t + \phi) \\ -48 + 1.6t \end{bmatrix}, \qquad \mathbf{r}_1^B(t) = \begin{bmatrix} 24 + 2\cos(0.04t + \phi) \\ 48 - 1.6t \end{bmatrix}. \tag{94}$$

The target trajectories evolve smoothly. At each update time $t_j = 5j$ s, $j = 0, \ldots, 11$, the predictor extrapolates the currently observed target velocity:

$$\widehat{\mathbf{r}}_{b,j}(\tau) = \mathbf{r}_b^B(t_j) + \tau \dot{\mathbf{r}}_b^B(t_j), \qquad 0 \leq \tau \leq 60 - t_j, \quad b \in \{0, 1\}. \tag{95}$$

The prediction radius includes a 7.2 m target radius, a 4 m separation margin, Gaussian position and velocity uncertainty, and a bounded-acceleration remainder:

$$R_j(\tau) = 7.2 + 4 + \sqrt{5.991\left(0.32^2 + \sigma_{v,j}^2 \tau^2\right)} + \tfrac{1}{2} a_{B,j} \tau^2, \tag{96}$$

Here, distances are in meters, $\sigma_{v,j}$ is in meters per second, and $a_{B,j}$ is in meters per second squared. The coefficient 5.991 is the 95% quantile of the chi-square distribution with two degrees of freedom, defining the two-dimensional Gaussian probability contour. At $t_j = 15$ s, the prescribed uncertainty update changes $(\sigma_{v,j}, a_{B,j})$ from $(0.016, 0.020)$ to $(0.08, 0.036)$ while the targets continue along Equation (94). This change provides a reproducible constraint-update event.

Each 5 s prediction interval is enclosed by a circumscribed polygon constructed from the two predicted center endpoints, the maximum interval radius, and 24 uniformly spaced support directions. The union of these polygons for both targets defines the current forbidden set. The complete remaining trajectory must avoid this set. Under constant control, the position trajectory is quadratic, allowing continuous support-plane separation tests over each control interval.

The 40 predefined tasks vary endpoint positions and target phase ϕ under the stated scenario settings. Task conditions are paired across methods, while each method generates its own candidate connections. All five methods receive identical task states and predictions and use a 2 s query allowance. DAR evaluates 90 detour and timing hypotheses through certified local action families. The planar RRTX and HL-RRT* implementations adapt the cited graph-repair and horizon-based lazy-tree mechanisms to control search in a four-dimensional state space, with at most 48 independently generated full-horizon control candidates per query. L-GLS and B-LGLS each construct 48 waypoint-transfer families. Selected controls satisfy the common full-trajectory feasibility criterion and undergo independent execution replay.

At $t = 15$ s, increased uncertainty invalidates the previous remaining-plan certificate. ISTR* reconstructs a dynamically feasible trajectory from the propagated state to the original terminal state under the updated occupancy constraints (Figure 1a). Across the 40 tasks, ISTR* completes 100.0%, compared with 10.0–47.5% for the baselines, and its 95th-percentile query time is 0.535 s within the 2 s allowance (Figure 1b and Table 2). RRTX and HL-RRT* obtain initial complete plans but lose certified continuations after subsequent updates. The results demonstrate sustained planar trajectory repair under the stated query allowance.

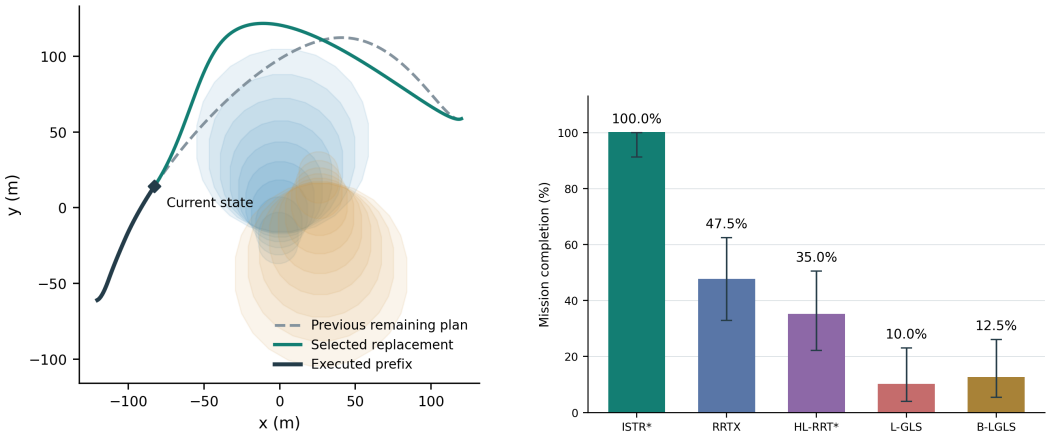

(**a**) (**b**)

Figure 1. Planar trajectory reconstruction at 15 s (**a**) and mission completion (**b**). Error bars indicate Wilson 95% confidence intervals.

Table 2. Planar mission-completion rates and median final executed costs. Wilson 95% confidence intervals are given in brackets; costs are evaluated on completed missions.

Method	Completion [95% CI] (%)	Median Final J_R	P95 Query (s)
ISTR*	100.0 [91.2, 100.0]	7.337282	0.535
RRTX	47.5 [32.9, 62.5]	5.490200	0.024
HL-RRT*	35.0 [22.1, 50.5]	5.759226	0.009
L-GLS	10.0 [4.0, 23.1]	5.444285	0.062
B-LGLS	12.5 [5.5, 26.1]	5.140902	0.059

The baseline final-cost medians are 5.140902–5.759226, compared with 7.337282 for ISTR*. Each median uses its method's completed-task set. The planar results therefore show higher mission completion for ISTR* alongside a higher completion-conditioned median cost.

5.3. Six-Dimensional Orbital Scenarios

The planar experiment demonstrates the reconstruction mechanism, while the orbital scenarios provide the principal evaluation of mission completion and executed cost. Both orbital scenarios use the CW/ZOH model in Equation (93) and the parameters in Table 3. Flight-event time $t = 0$ corresponds to posterior time $H = 100$ s, with a 500 s prediction horizon. After each 20 s flight segment, the updated posterior defines the remaining-horizon forbidden union with an additional 36 m verification margin, replacing the preceding constraint set. Each selected plan starts from the propagated state and reaches the prescribed terminal state. ISTR* uses active reoptimization in CF and continuation search in Moderate. The search allowance is 8 s, and the query deadline is 10 s.

The archived runs used acceptance thresholds of 10^{-3} m and 10^{-4} m/s. Their completion labels and final costs are retained under the reporting tolerances in Table 3.

ISTR* achieves completion rates of 100.0% in CF and 85.0% in Moderate, exceeding the four comparison methods in both task ensembles (Table 4 and Figure 2). Comparator completion rates range from 0.0% to 27.5% in CF and from 0.0% to 10.0% in Moderate. In CF, initialization failures account for 95.0%, 97.5%, 55.0%, 55.0% of the tasks, and subsequent replanning failures account for 2.5%, 2.5%, 17.5%, 17.5%, in the order RRTX, HL-RRT*, L-GLS, and B-LGLS. CF therefore primarily assesses the availability of initially certified connections, while Moderate examines recovery as the remaining feasible region changes.

Table 3. Parameters for the six-dimensional orbital experiments. Perturbation bounds apply to each component.

Parameter	Value
Nominal start position (m)	$(-369.648, -762.626, -151.262)$
Nominal goal position (m)	$(-235.153, 508.969, 521.767)$
Nominal endpoint velocity (m/s)	$(0.075725, 0.384334, 0.310729)$
Mission duration (s)	500
Posterior update period (s)	20
Reference mean motion (rad/s)	0.0011
Control hold (s)	0.5
Uniform endpoint-position perturbation (m)	$[-5, 5]$
Uniform initial-velocity perturbation (m/s)	$[-0.015, 0.015]$
Control norm limit (m/s^2)	0.036
Control component limit (m/s^2)	0.036
Speed limit (m/s)	6
Fixed terminal speed (m/s)	0.5
Terminal position reporting tolerance (m)	0.1
Terminal velocity reporting tolerance (m/s)	0.01

Table 4. Orbital mission-completion rates and median final executed costs. Wilson 95% confidence intervals are given in brackets; costs are evaluated on completed missions.

Scenario	Method	Completion [95% CI] (%)	Median Final J_R
CF	ISTR*	100.0 [91.2, 100.0]	0.123175
	RRTX	2.5 [0.4, 12.9]	0.121940
	HL-RRT*	0.0 [0.0, 8.8]	–
	L-GLS	27.5 [16.1, 42.8]	0.126483
	B-LGLS	27.5 [16.1, 42.8]	0.131064
Moderate	ISTR*	85.0 [70.9, 92.9]	0.187638
	RRTX	5.0 [1.4, 16.5]	0.194503
	HL-RRT*	0.0 [0.0, 8.8]	–
	L-GLS	10.0 [4.0, 23.1]	0.184859
	B-LGLS	5.0 [1.4, 16.5]	0.184714

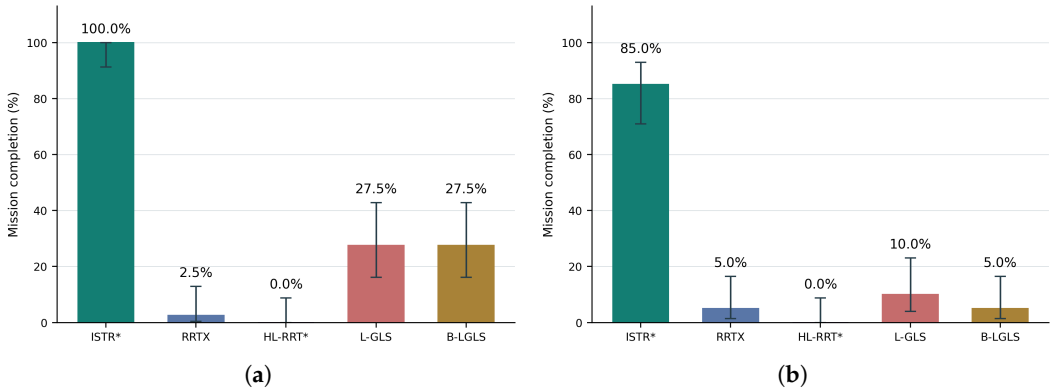

Figure 2. Mission completion in CF (**a**) and Moderate (**b**). Error bars indicate Wilson 95% confidence intervals.

Figure 3 shows the fraction of the original 40 tasks retaining a certified complete continuation after every update. A downward step records the loss of that continuation, and the value at 500 s equals the mission-completion rate. Losses occur mainly during initialization in CF and during subsequent updates in Moderate. Geometric intersection is assessed separately during replay.

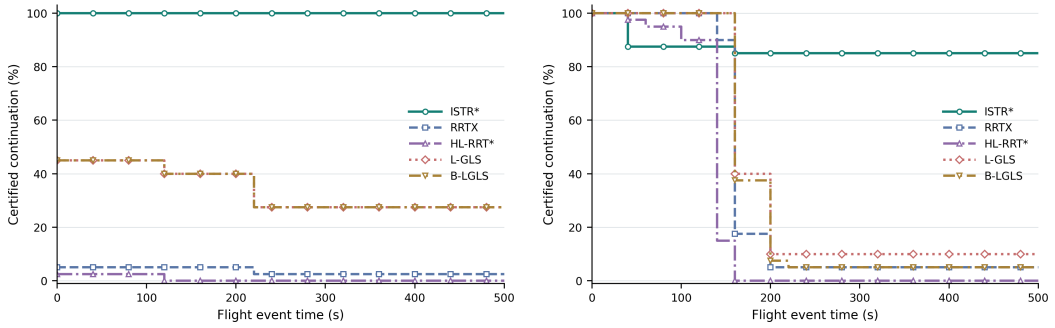

Figure 3. Certified mission continuation in CF (**a**) and Moderate (**b**). The terminal proportions equal the mission-completion rates.

The median final executed costs of ISTR* are 0.123175 in CF and 0.187638 in Moderate (Figure 4). Because the sets of completed tasks differ, paired cost comparisons use tasks completed by both methods. In CF, the median paired cost differences, defined as ISTR* minus comparator cost, are -0.006737 for L-GLS and −0.007516 for B-LGLS. ISTR* therefore incurs lower executed costs on these common tasks. RRTX provides one final CF cost. In Moderate, the comparator medians of 0.184714–0.194503 are close to the ISTR* median, whereas mission-completion rates differ substantially.

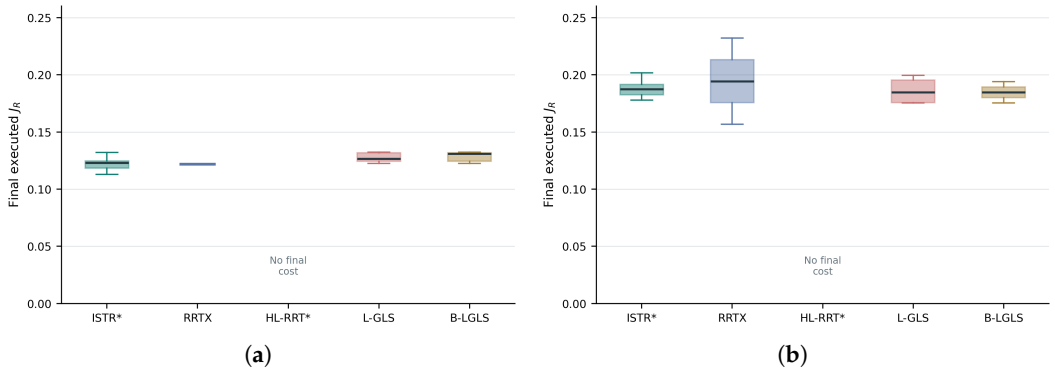

Figure 4. Final executed cost of completed missions in CF (**a**) and Moderate (**b**). Boxes show medians and interquartile ranges; whiskers follow the 1.5-IQR rule. A single observation is shown as a horizontal line.

5.4. Trajectory Reconstruction and Constraint Certification

Representative missions illustrate the effects of prediction updates on trajectory feasibility and cost under the settings in Section 5.3. The nominal CF example in Figure 5 is separate from the statistical batches and uses a 20 s search allowance. The cost-history comparison uses a batch task completed by ISTR*, L-GLS, and B-LGLS. Subsequent examples illustrate recovery of the separation condition and a complete Moderate mission.

Each remaining plan in Figure 5 starts from a propagated state on the executed trajectory. The initial planned cost is 0.255470, and the final executed cost is 0.115715. Each plan is verified against the future union associated with its own posterior update.

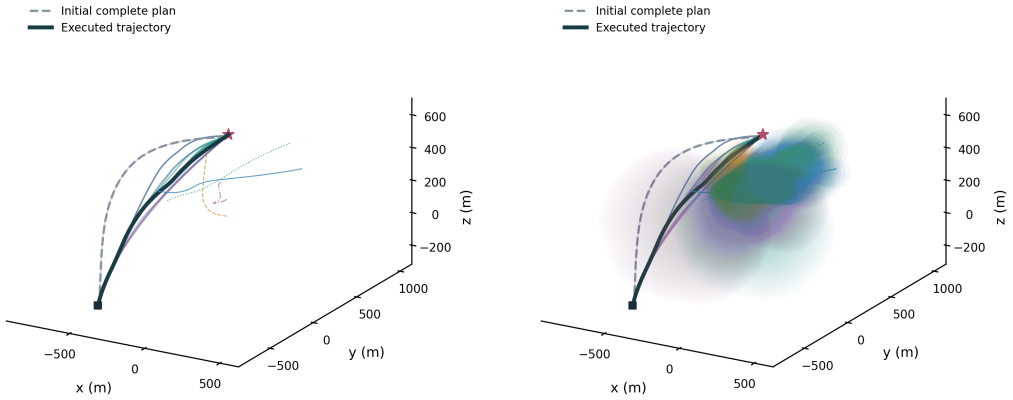

(**a**) (**b**)

Figure 5. CF trajectory history (**a**) and its overlay with six reconstructed posterior unions (**b**). Gray dashed and thick dark-teal curves denote the initial plan and executed trajectory, respectively. Remaining-plan colors progress from purple through blue and green to yellow with update time. Blue solid, orange dashed, green dotted, and purple dash-dotted curves identify the four targets; posterior shells use matching colors. The square and star mark the initial and terminal positions. Earlier plans and shells are more transparent.

Figure 6a shows the rolling total mission-cost estimate.

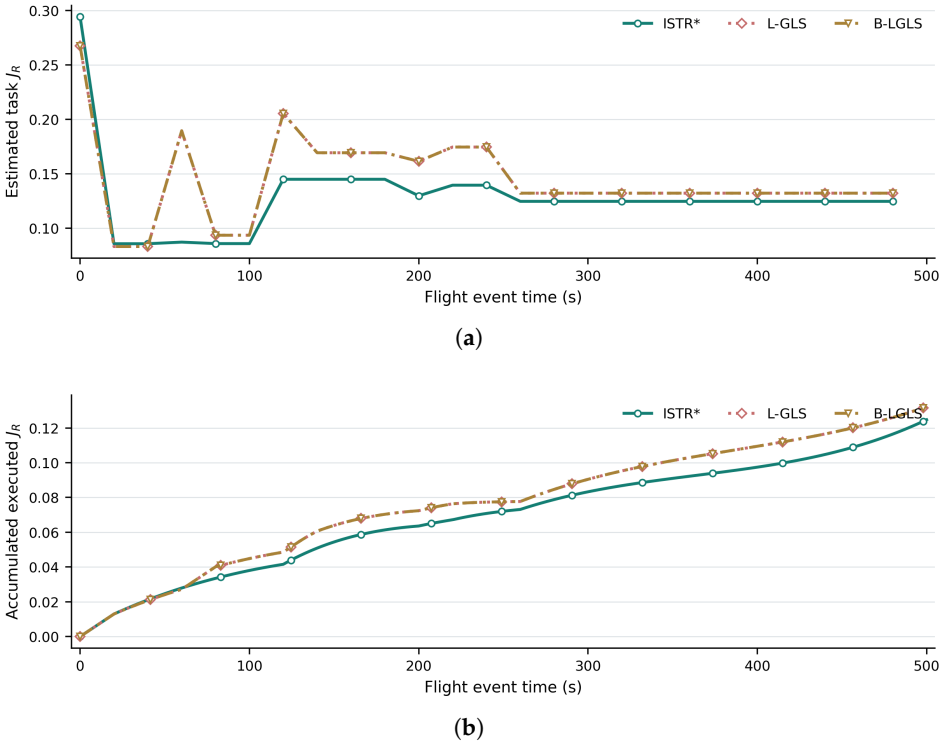

(**a**)

(**b**)

Figure 6. CF cost histories: accumulated executed cost plus remaining-plan cost (**a**), and accumulated executed cost (**b**).

At update time t_i, $J_{\text{roll}}(t_i) = J_{\text{exec}}(0, t_i) + J[\Pi_i^{\text{rem}}]$: the first term is the accumulated executed cost, and the second is the cost of the selected remaining plan. Prediction updates affect the remaining plan while preserving the executed prefix. All three methods reduce it at the first update. At 60 s, the two graph-planner estimates rise to approximately 0.19, while the ISTR* estimate remains near 0.09. At 120 s, the graph-planner estimates

are approximately 0.205, and the ISTR* estimate is approximately 0.145. The estimate increases accompany changes in the feasible set. In Figure 6b, executed cost accumulates monotonically, reaching approximately 0.125 for ISTR* and 0.132 for both graph planners.

Figure 7 compares the previous remaining plan with its selected replacement under the same updated posterior over the entire remaining mission. A positive separation lower bound certifies the separation condition. The previous plan loses this certificate at several updates, including 60 s and 120 s, while the replacement achieves a positive bound before execution resumes. At 120 s, the previous-plan bound is approximately −105 m. A negative bound indicates loss of the sufficient separation certificate; geometric intersection requires a separate test.

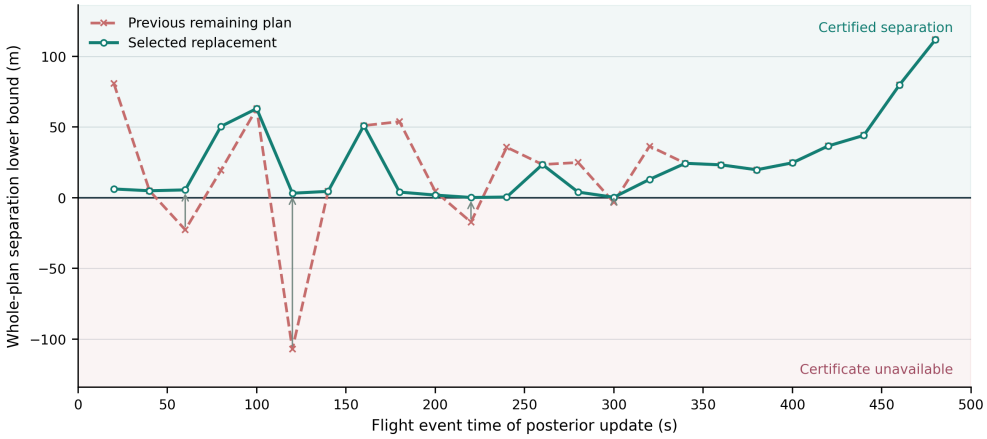

Figure 7. Full-trajectory separation lower bounds before and after replanning under the same updated posterior. Arrows indicate recovery of certification.

The Moderate mission in Figure 8 maintains a connection from each propagated state to the fixed terminal state through successive reconstructions. The shells depict the initial posterior geometry, and each subsequent plan satisfies the constraints associated with its own update.

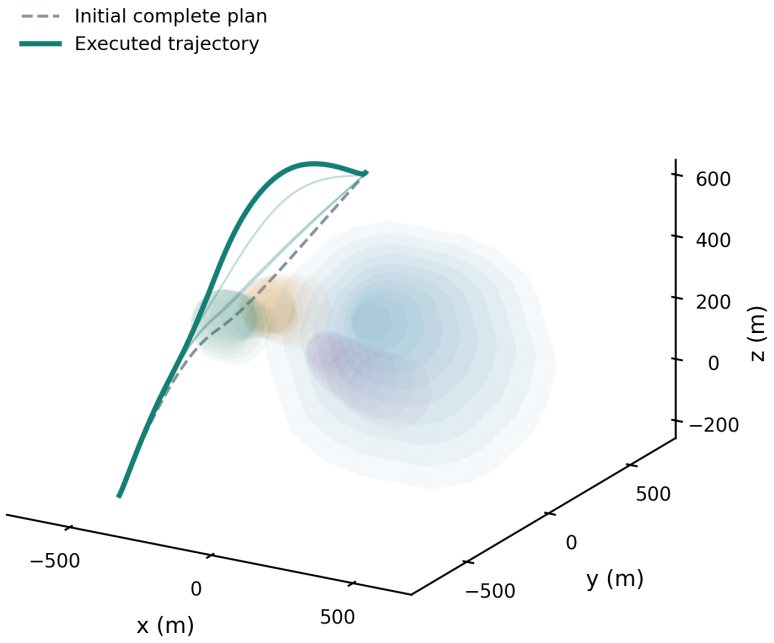

Figure 8. Initial plan, successive reconstructions, and executed trajectory in Moderate, overlaid with the initial posterior shells.

5.5. Computational Performance and Search-Policy Assessment

Initialization and replanning queries. Table 5 separates the 40 initialization queries per method and scenario from all subsequent queries, including unsuccessful attempts. Replanning-query counts depend on task stopping times. The failure columns distinguish initial-plan failure from loss of a complete plan after an update; both percentages use the original 40-task denominator. All orbital queries retain the original terminal state and 500 s mission deadline.

Table 5. Phase-specific query times (s) and task failures (%).

Scene	Method	Init. Med.	Init. P95	Replan Med.	Replan P95	Init. Fail	Later Fail
Planar	ISTR*	0.537	0.585	0.176	0.477	0.0	0.0
	RRTX	0.014	0.015	0.019	0.024	0.0	52.5
	HL-RRT*	0.009	0.010	0.001	0.006	0.0	65.0
	L-GLS	0.056	0.077	0.015	0.032	15.0	75.0
	B-LGLS	0.054	0.073	0.014	0.030	15.0	72.5
CF	ISTR*	5.769	6.209	1.947	6.388	0.0	0.0
	RRTX	8.002	8.006	8.009	8.012	95.0	2.5
	HL-RRT*	8.057	8.062	8.129	8.166	97.5	2.5
	L-GLS	1.051	1.405	0.325	0.667	55.0	17.5
	B-LGLS	1.135	1.558	0.310	0.689	55.0	17.5
Moderate	ISTR*	5.849	6.392	1.176	3.403	0.0	15.0
	RRTX	8.010	8.023	8.010	8.014	0.0	95.0
	HL-RRT*	8.054	8.071	8.120	8.207	0.0	100.0
	L-GLS	2.946	3.394	0.717	1.224	0.0	90.0
	B-LGLS	3.103	3.477	0.741	1.446	0.0	95.0

In CF, initial planning accounts for most baseline task failures. RRTX and HL-RRT* queries typically consume the search allowance, whereas the lazy roadmaps can exhaust their retained connections earlier. In Moderate, all methods obtain an initial plan; later updates account for task failures. The median replanning time of ISTR* decreases from its initialization value in both orbital scenes. The lazy methods return queries faster but complete fewer tasks. The timing and failure columns therefore describe complementary aspects of online replanning. A failure to obtain a certified complete suffix identifies the search outcome under the recorded representation and budget.

Figure 9 compares planning-query times for the five methods. Figure 10 evaluates the final executed cost and computational effort of two ISTR* reoptimization policies in CF. The policies share the same 40 tasks, dynamics, prediction histories, action representation, 8 s search allowance, and 10 s query deadline. With a certified remaining plan, the restricted policy permits 1.5 s of guided improvement; the active policy retains the plan as its feasible incumbent and permits independent search refresh within 8 s. Both policies use the full recovery allowance after certification is lost. The comparison evaluates the combined choice of search refresh and improvement allowance.

The cumulative distributions in Figure 9 show the proportion of queries returned within a given wall-clock time, including initial, successful, and failed queries. The median, 95th-percentile, and maximum ISTR* query times are 2.110, 6.388, and 8.523 s in CF, and 1.266, 5.424, and 8.605 s in Moderate. L-GLS and B-LGLS return queries faster but complete fewer missions, resulting in fewer subsequent queries. Most RRTX and HL-RRT* queries approach the 8 s search allowance. These measurements characterize the recorded implementations.

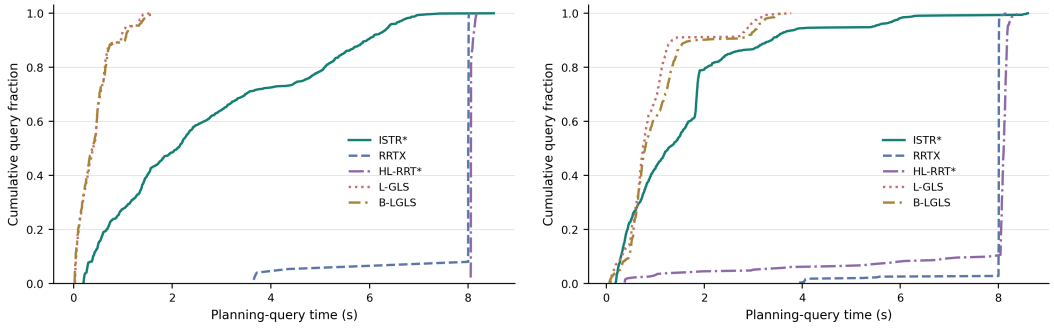

Figure 9. Cumulative query-time distributions in CF (**a**) and Moderate (**b**), including initial and failed queries.

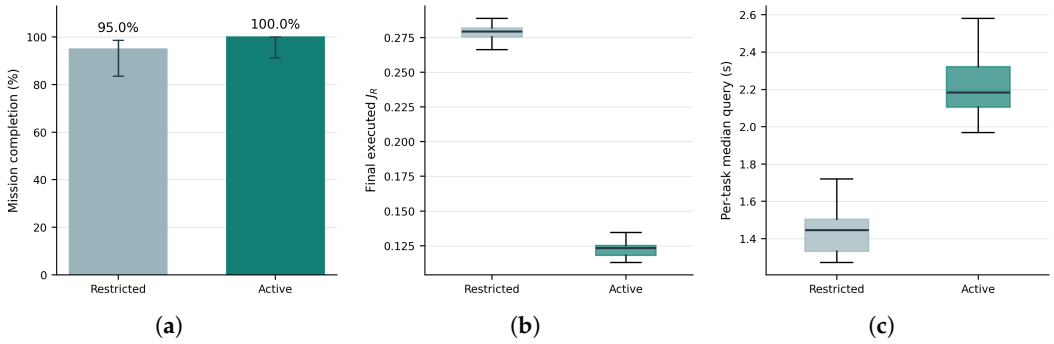

Figure 10. CF policy comparison: mission completion (**a**), final executed cost (**b**), and task-level median query time (**c**). Panel (**a**) includes Wilson 95% confidence intervals; (**b**,**c**) use tasks completed by both policies.

The restricted and active policies achieve CF completion rates of 95.0% $[83.5, 98.6]$ and 100.0% $[91.2, 100.0]$, respectively (Figure 10a). Final executed costs are compared on the same subset of tasks completed by both policies. On this common subset, active search reduces the median final executed cost from 0.279498 to 0.123354, a reduction of 55.9% between the sample medians (panel b). The reduction is $100(1 - \widetilde{J}_{\text{active}}/\widetilde{J}_{\text{restricted}})$, with both medians evaluated on the same 38 tasks completed by both policies. The overall median of task-level median query times increases from 1.446 to 2.184 s (panel c), and the 95th-percentile time across all queries increases from 2.907 to 6.888 s. Active reoptimization thus obtains a substantial executed-cost reduction through additional computation while retaining the certified incumbent during valid-plan updates.

Figure 11 compares baseline candidate allowances of 48, 96 and 192 on the same 40 planar tasks per method; the ISTR* configuration remains fixed. These sensitivity runs are analyzed separately from the primary comparison.

The archive contains 2733 carriers. Table 6 reports generated action patches, cumulative active-face visits, and query times from the same diagnostic runs with recorded operation counts, separately for initialization and replanning. The diagnostic runs use four predefined tasks per scenario under the primary action-construction settings. Face visits accumulate over all local quadratic solves within a query, including repeated solves. Replanning has lower medians for patch generation, face visits, and query time in both scenarios. The CF replanning P95 of 5.933 s also captures occasional computationally demanding repairs. Table 5 reports timing for the separate primary batch.

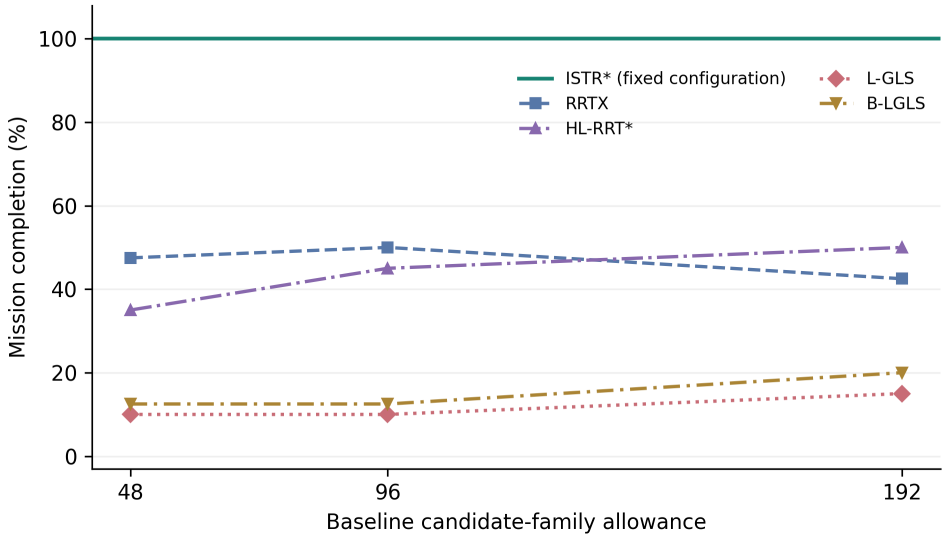

Figure 11. Planar completion under increased baseline candidate allowances. The ISTR* configuration remains fixed.

Table 6. Query workload and timing from diagnostic runs.

Scene	Phase	$N_{gen,50}$ (10^3)	$F_{tot,50}$ (10^6)	$F_{tot,95}$ (10^6)	T_{50} (s)	T_{95} (s)
CF	Initial	242.8	4.65	4.79	5.557	5.629
CF	Replan	97.2	1.64	4.17	1.780	5.933
Moderate	Initial	226.2	3.85	3.87	5.519	5.639
Moderate	Replan	40.5	0.68	1.27	1.036	2.972

Here, N_{gen} counts generated action patches, F_{tot} counts cumulative active-face visits, and T is complete-query time. Subscripts 50 and 95 denote the median and P95 over queries in the indicated phase. Independent enumeration checks cover 6720 local quadratic problems. For seven weights, the median, P95 and maximum visited-face counts are 18, 32 and 58, respectively, below the combinatorial bound of 127.

For the orbital queries, the nominal 8 s search allowance lies within the 10 s complete-query deadline. Preparation, trajectory recovery and feasibility verification add to search time, giving a maximum recorded ISTR* query time of 8.605 s in the primary batches. The workload measurements identify action construction, Gram-matrix assembly, and active-face solves as repeated computations. Reuse of unchanged certificates, matrix factorizations for unchanged endpoint maps, and processing cursors provides a basis for reducing this work.

5.6. Directed-Transfer Example

For the scalar double integrator $\dot{r} = v$, $\dot{v} = u$ over a normalized unit duration, minimizing $\frac{1}{2}\int_0^1 u(t)^2\,dt$ gives

$$J^*(x_0, x_1) = \tfrac{1}{2} d^T W^{-1} d, \qquad d = \begin{bmatrix} r_1 - r_0 - v_0 \\ v_1 - v_0 \end{bmatrix}, \quad W = \begin{bmatrix} 1/3 & 1/2 \\ 1/2 & 1 \end{bmatrix}. \tag{97}$$

The forward and reverse transfers between $(0,0)$ and $(1,1)$ cost 2 and 14, respectively, illustrating the directional action cost retained by DAR.

5.7. Dynamics Fidelity and Local Recovery Checks

Nonlinear central-gravity-plus-J_2 propagation provides the reference for comparing identical initial states and inputs in a common chief-orbit frame. The CW-plus-J_2 model in Equations (90) and (91) uses the instantaneous gravity gradient on the nominal circular orbit; the variational model follows the nonlinear chief trajectory and its rotating frame. The constants are $\mu = 3.986004418 \times 10^{14}$ m^3/s^2, $R_e = 6378137$ m, $J_2 = 1.08262668 \times 10^{-3}$ and $n = 0.0011$ rad/s. Three inclinations (0°, 51.6°, 97.4°) and four phases spaced by 90° define twelve orbital orientations. Figure 12 summarizes 576 free-flight cases and 24 mission-spacecraft control replays over 500 s. Over 500 s, CW-plus-J_2 discrepancies reach 0.957 m in CF and 1.045 m in Moderate; variational propagation reduces both to approximately 0.01 m. The free-flight cases use initial distances of 100, 500, and 1000 m, eight position directions, and initial speeds of 0 or 0.374 m/s, with nonzero initial velocity directed along the normalized $(0.3, -0.2, 0.1)$.

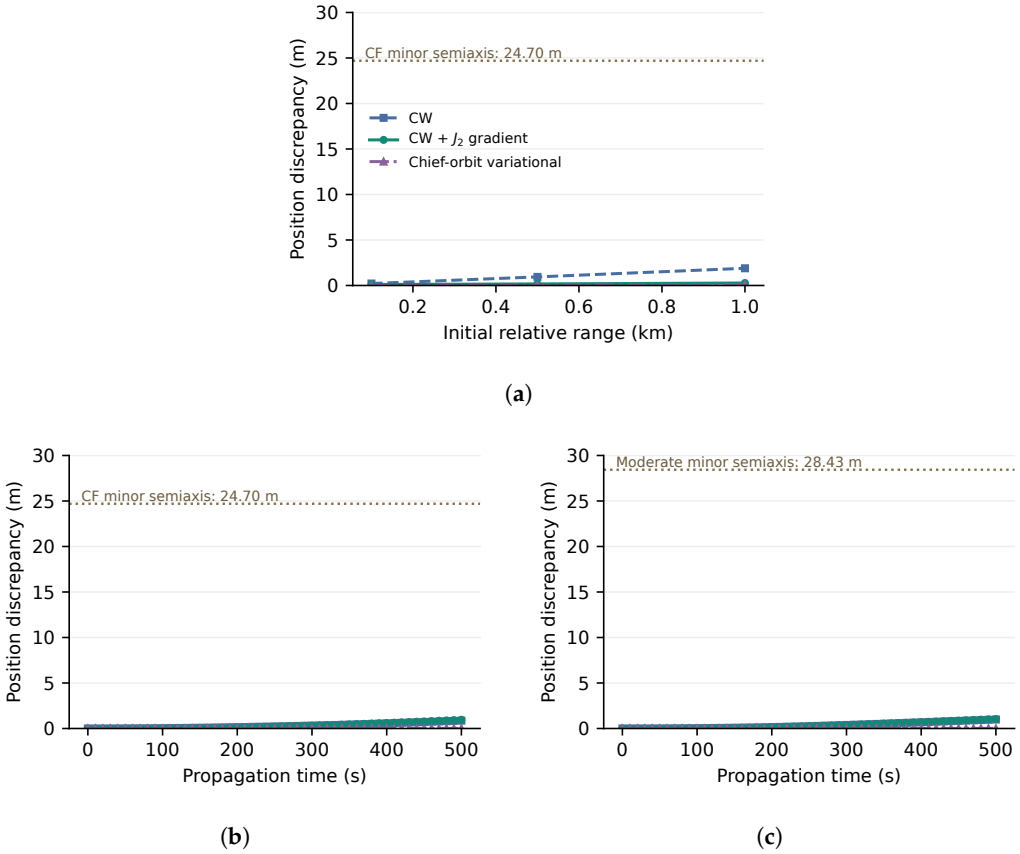

Figure 12. Model discrepancies over 500 s: (**a**) free flight; (**b**,**c**) CF and Moderate mission-spacecraft replays. Curves show sampled maxima. Reference lines mark the 24.70 and 28.43 m minor semiaxes forecast at 550 s for 600 s.

Table 7 characterizes propagation accuracy over an expanded range of distances, durations, and inputs using 22,464 trajectories. Initial speeds are selected from $\{0, 0.374, 1, 5, 10\}$ m/s and constant RTN input magnitudes from $\{0, 0.014, 0.036\}$ m/s^2, up to the sampled maxima listed in Table 7. Eight position directions are combined with three nonzero-velocity directions: normalized $(0.3, -0.2, 0.1)$, parallel to the initial position, and antiparallel. Nonzero inputs align with the initial position. Initial distances and maximum nonlinear-reference distances are measured from the chief; discrepancies compare model positions at identical times. Recorded discrepancies are 1.110 m at 20 km over 100 s and 1.076 m at 40 km over 50 s; extending the latter window to 500 s gives 145.539 m. The

selected distance–duration pairs characterize propagation accuracy for the stated sampled inputs and orbital orientations.

Table 7. Selected range–duration tests for CW plus J_2. Distance and discrepancy maxima are taken separately over sampled cases at 1 s intervals.

$\|r_0\|$ (km)	T (s)	$v_{0,\max}$ (m/s)	$a_{\text{test,max}}$ (m/s^2)	max $\|r\|$ (km)	$e_{\max}$ (m)
10	100	10	0.036	11.361	0.305
20	100	10	0.036	21.543	1.110
40	50	10	0.036	40.728	1.076
40	500	10	0.036	67.896	145.539
2	500	0.374	0.014	4.888	0.968

Eight archived non-cooperative-target maneuver histories provide 96 continuous 500 s replays over the same orientations. Both models propagate identical initial states and controls without state resets; position discrepancies are sampled at 0.25 s. The targets travel 275.2–831.1 m in CF and 263.0–434.0 m in Moderate, with maximum CW-plus-J_2 discrepancies of 0.377 and 0.134 m. These cumulative path lengths and propagation discrepancies characterize the target motion histories, complementing the mission-spacecraft replays in Figure 12.

The terminal outer approximation in Equation (26) is evaluated for the archived CW/ZOH configuration and prescribed terminal time by intersecting an endpoint box with an energy necessary condition. For $\tau = N_u h$, Equation (93) gives $G_j = \Phi_h^{N_u-1-j}\Gamma_h$, $j = 0, \ldots, N_u - 1$, and $d = \mathbf{x}_g - \Phi_h^{N_u}\mathbf{x}$. With componentwise control bound u_c and terminal radii ϵ_p, ϵ_v, Equation (98) bounds the endpoint coordinates; e_i is the ith coordinate unit vector.

$$|d_i| \le b_i + \epsilon_i, \qquad b_i = u_c \sum_{j=0}^{N_u-1} \|G_j^T e_i\|_1, \qquad \epsilon_i = \begin{cases} \epsilon_p, & i \le 3, \\ \epsilon_v, & i > 3. \end{cases} \tag{98}$$

Full row rank of $[G_0 \ \cdots \ G_{N_u-1}]$ makes W_{N_u} positive definite. Let $g = (g_r^T, g_v^T)^T$ partition the gradient into position and velocity components. The minimum-energy quadratic supplies the affine supporting bound over the terminal position and velocity balls in Equation (99); the energy ceiling uses the mission-spacecraft control bound $u_{A,\max} = 0.036\ \text{m/s}^2$.

$$\begin{aligned} W_{N_u} &= \sum_{j=0}^{N_u-1} G_j G_j^T / h, \qquad g = W_{N_u}^{-1} d, \\ \underline{E}(\mathbf{x}) &= \max\{0, \tfrac{1}{2} d^T W_{N_u}^{-1} d - \epsilon_p \|g_r\|_2 - \epsilon_v \|g_v\|_2\} \le \tfrac{1}{2} u_{A,\max}^2 \tau. \end{aligned} \tag{99}$$

The componentwise box contains norm-admissible controls, so intersecting the two necessary conditions preserves feasible terminal transfers. Recorded endpoints are screened with energy/box tolerances $10^{-10}/10^{-8}$; exclusion of an entire action domain requires verification over that domain.

With $h = 0.5$ s and $u_c = 0.036\ \text{m/s}^2$, 4000 feasible sequences over 50, 100, 250 and 500 s are each tested with exact terminals and with terminal radii $(\epsilon_p, \epsilon_v) = (10^{-3}\ \text{m}, 10^{-4}\ \text{m/s})$, yielding 8000 endpoint checks. Table 8 records no false exclusions and reproduces every archived rejection. Early CF records use exact endpoints; Moderate checks include the terminal neighborhood. A further 240 local-recovery tests cover both CW-plus-J_2 and nonlinear-J_2 propagation at a 500 m range, with durations $\tau \in \{0.5, 1, 2, 5, 10\}$ s, three inclinations, two phases and four position directions. Two independent half-interval

controls correct perturbations $(10^{-4}\tau^2\ \mathrm{m}, 10^{-4}\tau\ \mathrm{m/s})$. All corrections converge, with maximum position and velocity residuals of 1.97×10^{-11} m and 5.57×10^{-14} m/s. The minimum scaled Jacobian singular value, 0.15729, supports the endpoint-rank condition in Equation (81); the maximum control norm, 5.02×10^{-4} m/s^2, retains control reserve. Obstacle separation is verified over each action interval.

Table 8. Terminal-screening checks and replay agreement. N counts checks or records; support/box differences use state-coordinate units, and energy differences use the cost units of Equation (92).

Quantity	Value
N_{endpoint}	8000
$N_{\mathrm{false\ exclusion}}$	0
N_{support}	48
$\Delta^{\max}_{\mathrm{support}}$	1.82×10^{-12}
N_{CF}	21,238
N_{Moderate}	16,654
$\Delta^{\max}_{E}$	3.93×10^{-14}
$\Delta^{\max}_{\mathrm{box}}$	4.64×10^{-12}

6. Conclusions

This paper presents the ISTR* framework for spacecraft motion under updated probabilistic occupancy constraints. The framework integrates hybrid state sampling, reusable offline action representations, and online simplicial reconstruction within a unified planning procedure. DAR provides a directed kinodynamic search method that combines simplicial joining-state domains with Bellman propagation over certified state–time actions. Each action retains its realizing control, cost, and validity conditions, linking continuous predecessor selection to dynamical reachability and trajectory feasibility. Following posterior region updates using the CDIC method, the recorded dependencies identify the affected carrier simplices and actions, allowing local reconstruction to reuse the valid search records. Theoretical analysis establishes the feasibility and optimality on the specified finite action model and proves the fixed-instance asymptotic optimality and probabilistic completeness under the stated feasible-recovery, refinement, and coverage conditions.

Numerical experiments validate the reconstruction mechanism for planar motion in a four-dimensional state space and full-mission replanning in six-dimensional orbital scenarios. The reconstructed remaining trajectories satisfy the updated occupancy constraints and preserve the prescribed terminal state. In the primary orbital experiments, ISTR* achieves a mission-completion rate of at least 85% in each tested orbital task ensemble, higher than the rates obtained by the comparison methods. Recorded query times remain within the prescribed orbital computation deadline, with median query times on the scale of seconds. On the common subset of tasks completed by both configurations in the paired CF policy comparison, active reoptimization reduces the median final executed cost by 55.9%. For active reoptimization, the median of per-task median query times is 2.184 s. The experiments demonstrate incremental trajectory recovery and continued cost improvement during mission execution.

Future work will focus on adaptive state and control refinement for strongly maneuvering targets and on incorporating computation latency into the planning-execution cycle. Evaluation across independent prediction histories, together with consistent higher-fidelity orbital propagation and onboard implementation, can further characterize the method's robustness and computational requirements for collision-free spacecraft relative motion.

Author Contributions: Conceptualization, T.M. and S.L.; methodology, T.M.; software, T.M.; validation, T.M., B.N. and M.S.; formal analysis, T.M. and B.N.; writing—original draft preparation, T.M.; writing—review and editing, T.M., B.N., M.S. and S.L.; supervision, S.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The datasets presented in this article are not readily available because the datasets are part of an ongoing study. Requests to access the datasets should be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

Symbol	Meaning
$\mathbf{x}$; $\mathbf{r}, \mathbf{v}$	Complete state; relative position and velocity.
$\mathbf{u}, \mathcal{U}$	Control acceleration and admissible control set.
$\mathbf{f}, \boldsymbol{\theta}_d$	Dynamics vector field and declared dynamics parameters.
$\phi_{t,s}(y,u)$	State at time s propagated from state y at time t under control u.
$\mathcal{X}, \mathcal{X}_{\text{free}}$	Planning state domain and admissible state domain.
P_r, P_v	Position and velocity selection matrices.
$\mathbf{x}_s, t_s$; $\mathcal{X}_g$	Initial state and time; terminal state set.
$t_f, T_{\max}$	Arrival time and mission deadline.
i, t_i, T_{exe}	Posterior-update index, update time, and replanning interval.
ι, ι_P	Planning-instance identifier and the identifier stored with candidate P.
$h, h_{\min}$	Control-block duration and positive lower bound on declared action durations.
$\ell, J[\Pi]$	Running cost and integrated cost of maneuver Π.
$A, \mathcal{J}$	Mission spacecraft and set of non-cooperative target indices.
$\widehat{\mathbf{x}}_A(t_i)$	Mission-spacecraft state propagated to update time t_i.
$\mathcal{T}_i, \tau$	Forecast query offsets and an individual offset.
α, ν_i	Local-Gaussian probability content and posterior-version identifier.
$\mathcal{E}^{\text{CDIC}}_{j\|i}(t;\alpha)$	Target posterior position component conditioned on the prediction at update i.
$\mathcal{O}^{\text{post}}_i$	Collection of target-time regions and their posterior identifiers.
$\mathcal{O}^{\text{CDIC}}_i(t;\alpha)$	Union of target posterior position regions at physical time t.
$\widehat{\mathcal{O}}_i(t;\alpha)$	Continuous-time numerical enclosure used for separation verification.
$\mathcal{O}_i(t)$	Occupied-region input of the current planning problem.
$\mathbb{I}_i, \mathfrak{p}_i$	Prediction coverage interval and collected prediction data.
V_{δ_x}, δ_x	Sampled full-state set and state-space resolution.
$\mathcal{K}_{\delta_x}, \mathscr{C}$	Conforming simplicial complex and carrier simplices.
$\mathscr{I}_i, \mathbb{T}_i, \mathsf{A}$	Connection faces, represented physical times, and access labels.
e, t_e, τ_e	Fixed-endpoint connection, departure time, and duration.
σ, ς	Geometric simplex and ordered action sequence, respectively.
A^0_i	Permitted root access labels, a subset of A.
$\text{stat}_i(p)$	Validity status of premise p at update i.
$\mathbf{U}, u_{\mathbf{U}}$	Stacked control coefficients and their piecewise-constant control history.
$\kappa, t_\kappa, \tau_\kappa$	Action-family index, departure time, and duration.
λ, ζ, θ	Endpoint barycentric weights, free-control coordinates, and joint parameter (λ, ζ).
$\Theta^{\text{adm}}_\kappa, \Theta_\kappa$	Complete admissible parameter domain and its currently verified subset.
$y_\kappa, q_\kappa, X_\kappa$	Source map, exit map, and realizing state trajectory.
$\Theta_\kappa(y,q), D_\kappa$	Verified feasible parameter set for fixed endpoints and verified endpoint relation.
$J_\kappa(\theta), A_\kappa(y,q)$	Realized parameter cost and minimum cost for fixed endpoints.

$\theta_\kappa^*, u_\kappa, \chi_\kappa$	Minimizing parameter, realizing control, and verification data.
$\mathfrak{a}_\kappa$	Stored directed action data.
$\Gamma_\kappa^{\mathrm{in}}, \Gamma_\kappa^{\mathrm{out}}$	Entry and exit connection faces.
$\mathscr{M}_i, \Sigma_{\mathscr{M}_i}$	Declared finite action model and its feasible complete action sequences.
$J_i^\star, J_{\mathscr{M}_i}^\star$	Continuous-problem optimum and declared finite-model optimum.
$\mathscr{A}_i, \mathscr{B}_i, \mathscr{Q}_i$	Verified actions, arrival-cost fronts, and unresolved search records.
$\mathscr{G}_{D,i}, \mathscr{D}_i$	Dependency graph and affected-record closure.
$\mathscr{N}_i$	Local carrier neighborhood selected for reconstruction.
$\mathcal{B}_i^{\mathrm{root}}, \mathcal{B}^{\mathrm{root}}$	Zero-cost root-front collection at update i and in the current DAR call.
$\mathscr{Z}_i^{(r)}$	Search state after processing step r for the fixed declared model.
$D_b, V_b, \mathcal{B}_b$	Front domain, arrival-cost function, and stored front record.
$\mathcal{D}_{b\kappa}$	Source-exit relation restricted by the parent front and action constraints.
$\mathcal{T}_{b\kappa} V_b$	Bellman propagation of the parent value through action κ.
$\widehat{K}_{t,\le r}$	Outer set for relaxed terminal predecessors within remaining duration r.
$L_i, U_i, B_i, \varepsilon_i$	Search lower bound, feasible incumbent cost, work budget, and stopping tolerance.
$\Pi_i^{\mathrm{seed}}, \Pi_i^\star$	Verified seed trajectory and selected remaining trajectory.
J_R, J_{roll}	Final executed control-effort cost and rolling total mission-cost estimate.
$\mathcal{S}_m$	Task identifiers of missions completed by method m.
$N_{\mathrm{succ}}, N_{\mathrm{task}}$	Integer counts of successfully completed and attempted tasks.
$N_\lambda, N_u, N_{\mathrm{face}}$	Counts of weights, control blocks and visited faces in a local quadratic solve.
N_{fam}	Number of declared action families in the finite-termination proof.
i_{orb}	Reference-orbit inclination.
r_{conn}, r_c	Carrier-connection radius and reference-orbit radius, respectively.
$\epsilon_{\mathrm{num}}, \epsilon_{\mathrm{geom}}$	Numerical cost-error bound and numerical position allowance, respectively.
$\mathbf{e}_0, r_n$	Local endpoint-residual vector and its scalar norm bound.

References

1. Kavraki, L.E.; Svestka, P.; Latombe, J.C.; Overmars, M.H. Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE Trans. Robot. Autom.* **1996**, *12*, 566–580. [CrossRef]
2. LaValle, S.M. *Rapidly-Exploring Random Trees: A New Tool for Path Planning*; Technical Report TR 98-11; Department of Computer Science, Iowa State University: Ames, IA, USA, 1998. Available online: https://msl.cs.illinois.edu/~lavalle/papers/Lav98c.pdf (accessed on 6 October 2026).
3. Karaman, S.; Frazzoli, E. Sampling-Based Algorithms for Optimal Motion Planning. *Int. J. Robot. Res.* **2011**, *30*, 846–894. [CrossRef]
4. LaValle, S.M.; Kuffner, J.J. Randomized Kinodynamic Planning. *Int. J. Robot. Res.* **2001**, *20*, 378–400. [CrossRef]
5. Orthey, A.; Chamzas, C.; Kavraki, L.E. Sampling-Based Motion Planning: A Comparative Review. *Annu. Rev. Control Robot. Auton. Syst.* **2024**, *7*, 285–310. [CrossRef]
6. Halton, J.H. On the Efficiency of Certain Quasi-Random Sequences of Points in Evaluating Multi-Dimensional Integrals. *Numer. Math.* **1960**, *2*, 84–90. [CrossRef]
7. Palmieri, L.; Bruns, L.; Meurer, M.; Arras, K.O. Dispertio: Optimal Sampling for Safe Deterministic Motion Planning. *IEEE Robot. Autom. Lett.* **2020**, *5*, 362–368. [CrossRef]
8. Gammell, J.D.; Srinivasa, S.S.; Barfoot, T.D. Informed RRT*: Optimal Sampling-Based Path Planning Focused via Direct Sampling of an Admissible Ellipsoidal Heuristic. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Chicago, IL, USA, 14–18 September 2014; pp. 2997–3004. [CrossRef]
9. Gammell, J.D.; Barfoot, T.D.; Srinivasa, S.S. Batch Informed Trees (BIT*): Informed Asymptotically Optimal Anytime Search. *Int. J. Robot. Res.* **2020**, *39*, 543–567. [CrossRef]
10. Strub, M.P.; Gammell, J.D. Advanced BIT* (ABIT*): Sampling-Based Planning with Advanced Graph-Search Techniques. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 130–136. [CrossRef]
11. Strub, M.P.; Gammell, J.D. Adaptively Informed Trees (AIT*): Fast Asymptotically Optimal Path Planning Through Adaptive Heuristics. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 3191–3198. [CrossRef]
12. Strub, M.P.; Gammell, J.D. Adaptively Informed Trees (AIT*) and Effort Informed Trees (EIT*): Asymmetric Bidirectional Sampling-Based Path Planning. *Int. J. Robot. Res.* **2022**, *41*, 390–417. [CrossRef]

13. Xie, C.; van den Berg, J.; Patil, S.; Abbeel, P. Toward Asymptotically Optimal Motion Planning for Kinodynamic Systems Using a Two-Point Boundary Value Problem Solver. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Seattle, WA, USA, 26–30 May 2015; pp. 4187–4194. [CrossRef]
14. Wu, A.; Lew, T.; Solovey, K.; Schmerling, E.; Pavone, M. Robust-RRT: Probabilistically-Complete Motion Planning for Uncertain Nonlinear Systems. In *Robotics Research*; Springer Proceedings in Advanced Robotics; Springer Nature: Cham, Switzerland, 2023; Volume 27, pp. 538–554. [CrossRef]
15. Barfoot, T.D.; Tong, C.H.; Särkkä, S. Batch Continuous-Time Trajectory Estimation as Exactly Sparse Gaussian Process Regression. In Proceedings of the Robotics: Science and Systems X. Robotics: Science and Systems Foundation, Berkeley, CA, USA, 12–16 July 2014. [CrossRef]
16. Mukadam, M.; Dong, J.; Yan, X.; Dellaert, F.; Boots, B. Continuous-Time Gaussian Process Motion Planning via Probabilistic Inference. *Int. J. Robot. Res.* **2018**, *37*, 1319–1340. [CrossRef]
17. Mukadam, M.; Dong, J.; Dellaert, F.; Boots, B. STEAP: Simultaneous Trajectory Estimation and Planning. *Auton. Robot.* **2019**, *43*, 415–434. [CrossRef]
18. Starek, J.A.; Schmerling, E.; Maher, G.D.; Barbee, B.W.; Pavone, M. Fast, Safe, Propellant-Efficient Spacecraft Motion Planning Under Clohessy–Wiltshire–Hill Dynamics. *J. Guid. Control Dyn.* **2017**, *40*, 418–438. [CrossRef]
19. Frey, G.R.; Petersen, C.D.; Leve, F.A.; Kolmanovsky, I.V.; Girard, A.R. Constrained Spacecraft Relative Motion Planning Exploiting Periodic Natural Motion Trajectories and Invariance. *J. Guid. Control Dyn.* **2017**, *40*, 3100–3115. [CrossRef]
20. Daniel, K.; Nash, A.; Koenig, S.; Felner, A. Theta*: Any-Angle Path Planning on Grids. *J. Artif. Intell. Res.* **2010**, *39*, 533–579. [CrossRef]
21. Ferguson, D.; Stentz, A. Field D*: An Interpolation-Based Path Planner and Replanner. In *Robotics Research*; Thrun, S., Brooks, R., Durrant-Whyte, H., Eds.; Springer Tracts in Advanced Robotics; Springer: Berlin/Heidelberg, Germany, 2007; Volume 28, pp. 239–253. [CrossRef]
22. Yershov, D.S.; LaValle, S.M. Simplicial Dijkstra and A* Algorithms for Optimal Feedback Planning. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Francisco, CA, USA, 25–30 September 2011; pp. 3862–3867. [CrossRef]
23. Bhattacharya, S. Towards Optimal Path Computation in a Simplicial Complex. *Int. J. Robot. Res.* **2019**, *38*, 981–1009. [CrossRef]
24. Stentz, A. The Focussed D* Algorithm for Real-Time Replanning. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Montreal, QC, Canada, 20–25 August 1995; pp. 1652–1659. Available online: https://www.ijcai.org/Proceedings/95-2/Papers/082.pdf (accessed on 6 October 2026).
25. Koenig, S.; Likhachev, M. Fast Replanning for Navigation in Unknown Terrain. *IEEE Trans. Robot.* **2005**, *21*, 354–363. [CrossRef]
26. Otte, M.; Frazzoli, E. RRTX: Asymptotically Optimal Single-Query Sampling-Based Motion Planning with Quick Replanning. *Int. J. Robot. Res.* **2016**, *35*, 797–822. [CrossRef]
27. Chen, Y.; He, Z.; Li, S. Horizon-Based Lazy Optimal RRT for Fast, Efficient Replanning in Dynamic Environment. *Auton. Robot.* **2019**, *43*, 2271–2292. [CrossRef]
28. Lim, J.; Ghanei, M.; Lawson, R.C.; Srinivasa, S.; Tsiotras, P. Lazy Incremental Search for Efficient Replanning with Bounded Suboptimality Guarantees. *Int. J. Robot. Res.* **2024**, *43*, 1175–1207. [CrossRef]
29. Bashnick, C.; Ulrich, S. Fast Model Predictive Control for Spacecraft Rendezvous and Docking with Obstacle Avoidance. *J. Guid. Control Dyn.* **2023**, *46*, 998–1007. [CrossRef]
30. Specht, C.; Bishnoi, A.; Lampariello, R. Autonomous Spacecraft Rendezvous Using Tube-Based Model Predictive Control: Design and Application. *J. Guid. Control Dyn.* **2023**, *46*, 1243–1261. [CrossRef]
31. Xu, C.; Kang, G.; Li, Z.; Wu, J.; Zhou, J.; Hao, C.; Wu, J. Distributed cascaded model predictive control for multi-spacecraft capture of a tumbling target. *Adv. Space Res.* **2026**, *78*, 1195–1219. [CrossRef]
32. Kalaycioglu, S.; de Ruiter, A. Vibration Control of Satellite Antennas via NMPC and NARX Neural Networks. *IEEE Trans. Aerosp. Electron. Syst.* **2025**, *61*, 9406–9433. [CrossRef]
33. Jewison, C.; Miller, D.W. Probabilistic Trajectory Optimization Under Uncertain Path Constraints for Close Proximity Operations. *J. Guid. Control Dyn.* **2018**, *41*, 1843–1858. [CrossRef]
34. Apodaca, B.K.; Stirling, L.; Atkins, E.M. Orbital Dynamic Effects on Fuel Use in Sampling-Based Plans for Proximity Operations. *J. Spacecr. Rocket.* **2025**, *62*, 1018–1029. [CrossRef]
35. Girard, A. Reachability of Uncertain Linear Systems Using Zonotopes. In Proceedings of the Hybrid Systems: Computation and Control (HSCC), Zurich, Switzerland, 9–11 March 2005; pp. 291–305. [CrossRef]
36. Althoff, M. Reachability Analysis of Nonlinear Systems Using Conservative Polynomialization and Non-Convex Sets. In Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control (HSCC), Philadelphia, PA, USA, 8–11 April 2013; pp. 173–182. [CrossRef]

37. Lew, T.; Janson, L.; Bonalli, R.; Pavone, M. A Simple and Efficient Sampling-Based Algorithm for General Reachability Analysis. In *Proceedings of the 4th Annual Learning for Dynamics and Control Conference, Stanford University, Stanford, CA, USA, 23–24 June 2022*; Firoozi, R., Mehr, N., Yel, E., Antonova, R., Bohg, J., Schwager, M., Kochenderfer, M., Eds.; Proceedings of Machine Learning Research; PMLR: New York, NY, USA, 2022; Volume 168, pp. 1086–1099. Available online: https://proceedings.mlr.press/v168/lew22a.html (accessed on 6 October 2026).
38. Seo, H.; Lee, D.; Son, C.Y.; Jang, I.; Tomlin, C.J.; Kim, H.J. Real-Time Robust Receding Horizon Planning Using Hamilton–Jacobi Reachability Analysis. *IEEE Trans. Robot.* **2023**, *39*, 90–109. [CrossRef]
39. Kschischang, F.R.; Frey, B.J.; Loeliger, H.A. Factor Graphs and the Sum-Product Algorithm. *IEEE Trans. Inf. Theory* **2001**, *47*, 498–519. [CrossRef]
40. Dellaert, F.; Kaess, M. Square Root SAM: Simultaneous Localization and Mapping via Square Root Information Smoothing. *Int. J. Robot. Res.* **2006**, *25*, 1181–1203. [CrossRef]
41. Kaess, M.; Ranganathan, A.; Dellaert, F. iSAM: Incremental Smoothing and Mapping. *IEEE Trans. Robot.* **2008**, *24*, 1365–1378. [CrossRef]
42. Kaess, M.; Johannsson, H.; Roberts, R.; Ila, V.; Leonard, J.J.; Dellaert, F. iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree. *Int. J. Robot. Res.* **2012**, *31*, 216–235. [CrossRef]
43. Ma, T.; Nan, B.; Sun, M.; Li, S. Joint Posterior Reachable-Region Prediction via Local Markov Factor Graphs. *Mathematics* **2026**, *14*, 3383. [CrossRef]
44. Clohessy, W.H.; Wiltshire, R.S. Terminal Guidance System for Satellite Rendezvous. *J. Aerosp. Sci.* **1960**, *27*, 653–658. [CrossRef]
45. Schweighart, S.A.; Sedwick, R.J. High-Fidelity Linearized J_2 Model for Satellite Formation Flight. *J. Guid. Control Dyn.* **2002**, *25*, 1073–1080. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.