

DUAL BÉZIER SPACE–TIME OPTIMIZATION UNDER HARD OBSERVATION-WINDOW CONSTRAINTS: A FEASIBILITY-ORIENTED JERK-CAP COMPARISON

Heewon Kim[1] & Suwon Lee[2]

[1]Department of Automotive and Mobility Engineering, Kookmin University, 77 Jeongneung-ro, Seongbuk-gu, Seoul 02707, Republic of Korea
[2]Department of Future Mobility, Kookmin University, 77 Jeongneung-ro, Seongbuk-gu, Seoul 02707, Republic of Korea

Abstract

Satellite retargeting maneuvers for Earth observation are constrained by spatial feasibility, time allocation, and viewing geometry within a finite observation window. Even when a geometric path is feasible, a uniform time parameterization can induce sharp acceleration and jerk peaks and may violate a viewing-angle constraint enforced as a hard inequality over the window. This paper presents a dual Bézier space–time optimization framework that co-optimizes a smooth three-dimensional trajectory and a monotone time schedule using Bernstein polynomials. The time-scaling profile is optimized directly with analytic derivatives and a compact set of decision variables, mitigating jerk peaks caused by fixed uniform timing. Validation is conducted in two stages. Stage 1 compares a baseline (space-only optimization with uniform timing) with the proposed dual formulation (space–time co-optimization). Stage 2 enforces a hard viewing-angle requirement within an automatically placed observation window; SQP robustness is improved using fixed-length, time-uniform sampling inside the window together with continuation and backtracking on the viewing-angle constraint. Under a fixed observation window and viewing-angle constraint, we perform an integral-cost Pareto sweep over five weight combinations of acceleration and jerk energies, complemented by a cross-degree monotonicity verification for Bézier degrees $n = 5$ through 8 using exact degree-elevation warm starts. Under stringent viewing constraints, the uniform-time baseline becomes infeasible or incurs substantially larger jerk-energy budgets, whereas the dual formulation remains feasible with lower jerk allowance across the Pareto front and reproduces the monotonic cost decrease theoretically guaranteed by Bézier-space nesting. Replicating the integral-cost sweep on the baseline further reveals that under uniform timing the weight-driven trade-off collapses and increased Bézier degree benefits only the acceleration-only case, isolating time co-optimization as the enabler of both the trade-off space and the cross-degree benefit.

Keywords: Bézier curve, Bernstein basis, time reparameterization, observation window, Pareto sweep, degree monotonicity, jerk-cap feasibility

1. Introduction

Earth-observation satellites frequently perform retargeting maneuvers to redirect their line-of-sight from one ground target to another. During such transfers, the spacecraft must not only reach the desired final state but also maintain a favorable viewing geometry within a finite observation window, a short time interval during which the satellite passes over the target region and imaging or surveillance data must be acquired. Satisfying boundary conditions alone is therefore insufficient: the transfer must also remain trackable and observation-stable, which motivates limiting peak-like motion quantities such as acceleration and jerk.

This becomes particularly challenging when the viewing requirement is enforced only during a short observation window, because the effective feasibility is governed by how time is allocated around that

window rather than by geometry alone. Consider a transfer trajectory that is geometrically smooth: if the time parameterization is uniform, the satellite may traverse the observation-critical portion of its path too quickly, producing high acceleration or jerk peaks precisely where the viewing constraint is active. In other words, even a geometrically plausible path can become infeasible if its timing concentrates motion variation inside the constrained window.

The jerk magnitude, the time derivative of acceleration, is particularly relevant in this context. High jerk peaks indicate rapid changes in the thrust vector, which stress attitude-control actuators (e.g., reaction wheels or control-moment gyroscopes) [1], induce structural vibrations in flexible appendages such as solar panels [2], and can degrade pointing stability during the observation window [3]. In agile Earth-observation missions where rapid retargeting is combined with high-resolution imaging, managing jerk is therefore a critical design consideration [4]. A formulation that can satisfy the viewing constraint while minimizing the required jerk allowance directly translates to relaxed hardware requirements, longer actuator lifetimes, and improved image quality.

Furthermore, the observation window is typically short relative to the total transfer duration, often less than 5% of the orbit period [5, 6]. This asymmetry creates a strong incentive to redistribute time: by allocating more physical time to the window (slowing down there) and less time elsewhere (speeding up), the trajectory can simultaneously reduce jerk peaks within the window and satisfy the viewing-angle constraint. This is precisely the capability that the proposed dual Bézier formulation provides.

Bézier curves and Bernstein polynomials have become popular parameterizations for trajectory optimization due to their favorable numerical properties. Lee and Kim [7] demonstrated Bézier-based output trajectory shaping for guidance problems, exploiting the convex-hull property of Bernstein polynomials. Kielas-Jensen et al. [8] and Cichella et al. [9] developed Bernstein-polynomial methods for multi-vehicle trajectory generation with collision avoidance, while Walton et al. [10] addressed endpoint-constrained trajectory optimization. Sun et al. [11] proposed a bilevel optimization framework with analytical gradients for fast UAV trajectory planning, and Hobbs and Kaneshige [12] specifically addressed jerk-limited trajectory generation for unmanned systems. On the spacecraft side, optimal control approaches using polynomial parameterizations have been applied to formation flying [13], proximity operations [14], and general nonlinear programming formulations [15].

However, most of these formulations primarily optimize the spatial shape of the trajectory under a fixed (typically uniform) time parameterization. When hard operational constraints, such as a viewing-angle constraint within an observation window, are introduced, the fixed-time assumption can be overly restrictive. The optimizer has no ability to redistribute physical time along the path, so it cannot slow down during the constrained window (to reduce acceleration/jerk there) and speed up elsewhere. This limitation becomes critical as the viewing constraint tightens: the space-only formulation may become infeasible even though a solution exists if time is treated as a free variable.

To address this gap, we make three contributions. First, we propose a dual Bézier space–time formulation that jointly optimizes a smooth 3D path and a monotone time-scaling profile within a Bernstein-polynomial parameterization, explicitly treating time allocation as a decision variable. The spatial path and the time schedule are each represented as Bézier curves, and both are co-optimized by sequential quadratic programming (SQP). Second, we develop an SQP-stable enforcement scheme for a windowed hard viewing-angle constraint using a window-aligned discretization and a continuation-based tightening procedure, which avoids the convergence difficulties that arise from variable-dimension constraint vectors. Third, we introduce an integral-form cost function combining acceleration and jerk energies, and validate the space–time trade-off and cross-degree cost monotonicity through a Pareto sweep over five representative weight combinations and four Bézier degrees. For consistent cross-degree comparison, we warm-start each higher-degree optimization with the previous degree's solution lifted via exact degree elevation, numerically reproducing the monotonic decrease theoretically guaranteed by the fact that every degree-$(n-1)$ Bézier curve admits an exact degree-n representation through degree elevation.

The remainder of this paper is organized as follows. Section 2. introduces the dual Bézier space–time parameterization and derives the chain-rule expressions for physical-time derivatives. Section 3. formulates the optimization problem, including decision variables, objective function, and constraints.

Section 4. presents the hard observation constraint, the two-stage solution strategy, and the continuation procedure. Section 5. describes the experimental design for the weight Pareto sweep and cross-degree monotonicity verification. Section 6. reports numerical experiments. Section 7. discusses the mechanism behind the observed advantages, the role of Bézier degree, and the limitations and practical implications of the approach, and Section 8. concludes.

2. Dual Bézier Space–Time Parameterization

Throughout this section, a dot ($\dot{}$) denotes differentiation with respect to physical time t, while a prime ($'$) denotes differentiation with respect to the Bézier parameter $\tau \in [0,1]$.
For $\tau \in [0,1]$, the Bernstein basis of degree n is [16]

$$B_i^n(\tau) = \binom{n}{i}(1-\tau)^{n-i}\tau^i, \quad i = 0,\dots,n. \tag{1}$$

2.1 Spatial path

The three-dimensional transfer trajectory is represented as a Bézier curve of degree n:

$$\mathbf{r}(\tau) = \sum_{i=0}^{n} \mathbf{P}_i B_i^n(\tau), \tag{2}$$

where $\mathbf{P}_i \in \mathbb{R}^3$ are the spatial control points. Endpoint positions and velocities are enforced by fixing the first two and last two control points, leaving only the interior control points $\mathbf{P}_2,\dots,\mathbf{P}_{n-2}$ as free variables.

2.2 Time schedule

Physical time is also parameterized as a Bézier curve of degree m:

$$t(\tau) = \sum_{j=0}^{m} T_j B_j^m(\tau), \quad t'(\tau) = \frac{\mathrm{d}t}{\mathrm{d}\tau} > 0 \ \ (\forall\,\tau), \tag{3}$$

where $T_0 = 0$, $T_m = t_f$, and monotonicity is enforced so that $t'(\tau) > 0$ everywhere. The endpoint time derivatives are fixed by constraining T_1 and T_{m-1} to preserve the velocity boundary conditions under the chain rule.
Physically, $t'(\tau) = \mathrm{d}t/\mathrm{d}\tau$ represents the local time rate: regions where t' is large correspond to intervals where physical time passes slowly relative to the curve parameter (the satellite lingers), while small t' means rapid traversal. For a quintic time Bézier ($m = 5$), the six control points $T_0,\dots,T_5$ are arranged as follows: $T_0 = 0$ and $T_5 = t_f$ fix the total duration; $T_1 = t_f/5$ and $T_4 = 4t_f/5$ are fixed to match the endpoint velocity boundary conditions; and only T_2 and T_3 remain as free optimization variables. Despite having only two free time parameters, this representation provides sufficient flexibility to redistribute time across the transfer, as demonstrated in the numerical experiments.

2.3 Chain rule and physical-time derivatives

The chain rule relates physical-time derivatives to τ-derivatives:

$$\frac{\mathrm{d}}{\mathrm{d}t}(\cdot) = \frac{1}{t'(\tau)}\frac{\mathrm{d}}{\mathrm{d}\tau}(\cdot). \tag{4}$$

Applying the chain rule once yields the physical velocity:

$$\dot{\mathbf{r}} = \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} = \frac{\mathbf{r}'(\tau)}{t'(\tau)}. \tag{5}$$

For the acceleration, we differentiate again and apply the quotient rule. Writing $\dot{\mathbf{r}} = \mathbf{r}'/t'$ and using $\mathrm{d}/\mathrm{d}t = (1/t')\,\mathrm{d}/\mathrm{d}\tau$:

$$\ddot{\mathbf{r}} = \frac{1}{t'}\frac{\mathrm{d}}{\mathrm{d}\tau}\left(\frac{\mathbf{r}'}{t'}\right) = \frac{1}{t'}\cdot\frac{\mathbf{r}''t' - \mathbf{r}'t''}{(t')^2} = \frac{\mathbf{r}''t' - \mathbf{r}'t''}{(t')^3}. \tag{6}$$

The jerk is obtained by one further differentiation. Denoting $\mathbf{a}(\tau) \equiv \mathbf{r}''t' - \mathbf{r}'t''$ and $b(\tau) \equiv (t')^3$, we compute $\dddot{\mathbf{r}} = (1/t')(\mathbf{a}'b - \mathbf{a}b')/b^2$, which after simplification gives:

$$\dddot{\mathbf{r}} = \frac{\mathbf{r}'''(t')^2 - 3\,\mathbf{r}''t't'' - \mathbf{r}'\left[t'''t' - 3(t'')^2\right]}{(t')^5}. \tag{7}$$

The τ-derivatives of the Bézier curves are themselves Bézier curves of reduced degree: for a degree-n spatial curve, $\mathbf{r}'(\tau)$ is a degree-$(n-1)$ curve with control points $\Delta\mathbf{P}_i = n(\mathbf{P}_{i+1} - \mathbf{P}_i)$, and similarly for higher derivatives and the time curve. This structure enables efficient analytic evaluation of all quantities in Eqs. (5)–(7) without numerical differentiation.

For thrust-equivalent mode, the control acceleration is defined as $\mathbf{u}(t) = \ddot{\mathbf{r}}(t) - \mathbf{a}_{\text{grav}}(t)$, where $\mathbf{a}_{\text{grav}} = -\mu\,\mathbf{r}/\|\mathbf{r}\|^3$ is the two-body gravitational acceleration. The thrust-equivalent jerk is then $\dot{\mathbf{u}}(t) = \dddot{\mathbf{r}}(t) - \dot{\mathbf{a}}_{\text{grav}}(t)$. This metric isolates the propulsive contribution from the natural gravitational field, providing a physically meaningful measure of actuator effort.

Thus, when time is a decision variable, the physical acceleration and jerk depend on both $\mathbf{r}(\tau)$ and $t(\tau)$, which is the key mechanism by which the dual formulation can redistribute motion effort across the transfer.

To build intuition, consider a segment of the path where the spatial curvature $\mathbf{r}''$ is large. Under uniform timing, this segment necessarily produces a large physical acceleration $\ddot{\mathbf{r}}$. However, if the optimizer increases $t'(\tau)$ in that region, effectively allocating more physical time to traverse the high-curvature segment, the denominator $(t')^3$ in Eq. (6) grows, reducing the peak acceleration. The same principle applies to the jerk expression (7) with an even stronger $(t')^5$ dependence. This is precisely the advantage exploited by the dual formulation.

3. Optimization Problem Formulation

3.1 Decision variables

For the baseline formulation, only the free interior spatial control points $\{\mathbf{P}_i\}_{i=2}^{n-2}$ are optimized, while time is uniform ($t(\tau) = t_f\tau$). For the dual formulation, the decision vector additionally includes the free interior time control points $\{T_j\}_{j=2}^{m-2}$. In our implementation with quintic ($n = m = 5$) Bézier curves, this corresponds to 6 spatial variables and 2 time variables for a total of 8 decision variables for the dual case.

3.2 Objective function design

We minimize an integral-form cost function that weights two aspects of trajectory smoothness, namely acceleration energy and jerk energy:

$$J(w_a, w_j) = w_a I_{a,\text{ref}}^{-1} \int_0^{t_f} \|\mathbf{a}(t)\|^2\,\mathrm{d}t + w_j I_{j,\text{ref}}^{-1} \int_0^{t_f} \|\mathbf{j}(t)\|^2\,\mathrm{d}t, \tag{8}$$

where (w_a, w_j) is a pair of normalized weights, and $I_{a,\text{ref}}^{-1}$ and $I_{j,\text{ref}}^{-1}$ are normalizing coefficients whose values are fixed by setting $I_{a,\text{ref}}$ and $I_{j,\text{ref}}$ equal to the integrals $\int \|\mathbf{a}\|^2\,\mathrm{d}t$ and $\int \|\mathbf{j}\|^2\,\mathrm{d}t$, respectively, obtained at the reference configuration (Case A and Case E with $n = 5$). This normalization keeps J of $\mathcal{O}(1)$ magnitude across weight cases, enabling direct comparison.

To explore the acceleration–jerk trade-off, we examine five representative weight combinations: Case A, $(w_a, w_j) = (1, 0)$, pure acceleration-energy minimization; Case B, $(0.75, 0.25)$, acceleration-dominant; Case C, $(0.5, 0.5)$, balanced; Case D, $(0.25, 0.75)$, jerk-dominant; and Case E, $(0, 1)$, pure jerk-energy minimization.

Unlike the alternative of minimizing $\max_t \|\mathbf{j}(t)\|$, the integral objective directly minimizes smooth quantities. This yields two benefits: (i) the SQP-based nonlinear solver fmincon (MATLAB Optimization Toolbox) reaches deeper local optima without the numerical difficulties associated with non-smooth peak quantities, and (ii) cost monotonicity with Bézier degree, theoretically guaranteed because every degree-$(n-1)$ Bézier curve admits an exact degree-n representation through degree elevation, is confirmed numerically in Section 6.8, where all 35 of 35 combinations exhibit monotonic decrease across $n = 5 \to 8$.

3.3 Constraints

The optimization is subject to four hard constraints. The endpoint conditions fix the first and last two spatial and time control points to enforce position and velocity boundary conditions. Time monotonicity requires $t'(\tau) > 0$ for all $\tau \in [0,1]$, which is equivalently guaranteed by the control-point ordering $T_1 < T_2 < \cdots < T_{m-1}$. The actuator acceleration limit $\|\mathbf{a}(t)\| \le a_{\mathrm{cap}}$ reflects realistic thruster capability. Finally, the observation-window viewing-angle constraint $\psi(t) \le \psi_{\max}$ holds for $t \in [t^{\mathrm{obs}}_{\mathrm{start}}, t^{\mathrm{obs}}_{\mathrm{end}}]$ (see Section 4.).

4. Hard Observation Constraint and Two-Stage Strategy

4.1 Observation window and viewing angle

Let $\psi(t)$ denote the central angle between the satellite position vector $\mathbf{r}_{\mathrm{sat}}(t)$ and the ground-target direction $\mathbf{r}_{\mathrm{tgt}}(t)$, both referenced to the Earth center (Fig. 1); the satellite-to-target line-of-sight (LOS) is shown for reference, and a small ψ corresponds to favorable viewing geometry. The hard observation constraint enforces:

$$\psi(t) \le \psi_{\max}, \quad t \in [t^{\mathrm{obs}}_{\mathrm{start}}, t^{\mathrm{obs}}_{\mathrm{end}}]. \tag{9}$$

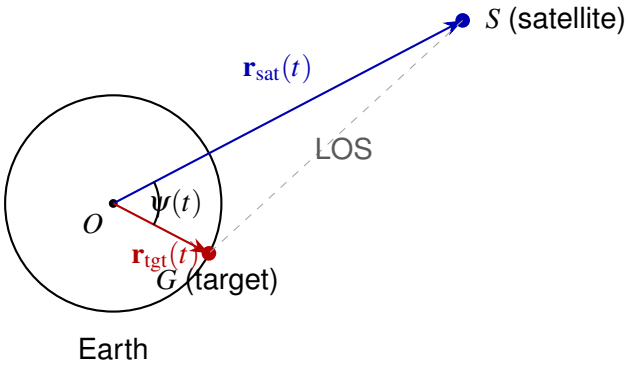


Figure 1 – Viewing central angle $\psi(t)$.

Physically, the observation window is the time interval during which the satellite has line-of-sight to the ground target with a viewing angle smaller than the required threshold; imaging or surveillance data must be acquired within this interval. Because the satellite is continuously moving along the transfer trajectory, the viewing angle $\psi(t)$ varies over time, reaching its minimum near the closest approach to the target (see Fig. 7). The window of fixed half-width w should therefore be placed around this minimum so that the viewing constraint is as easy to satisfy as possible.

To automate this placement, we adopt a minimax criterion: the center t_c is chosen to minimize the peak viewing angle over the window,

$$t_c = \arg\min_{t} \max_{t' \in [t-w, t+w]} \psi(t'), \tag{10}$$

which places the window at the portion of the trajectory where the satellite is most favorably aligned with the target.

4.2 Two-stage solution strategy

Stage 1 optimizes trajectory smoothness without any observation constraint. This stage is itself divided into two sub-cases for comparison. In the baseline sub-case (1A), only the spatial control points are optimized while the time schedule is held uniform. In the dual sub-case (1B), the spatial and time control points are optimized jointly. The Stage 1 solutions serve as initial guesses for the subsequent constrained optimization.

Stage 2 introduces the hard observation constraint (9) using the corresponding Stage 1 solution as initial guess. Because hard observation-window inequalities may degrade SQP convergence when the effective constraint dimension varies or when peak violations fall between coarse samples, two implementation choices stabilize convergence. First, the viewing angle is evaluated at a fixed number

of points uniformly spaced in physical time t within the window (not uniformly in τ), yielding a constraint vector of constant length across iterations; sampling in physical time is essential because the observation constraint is defined on a physical-time interval $[t_{\text{start}}^{\text{obs}}, t_{\text{end}}^{\text{obs}}]$, and under the dual formulation the mapping $t(\tau)$ is itself a decision variable, so a τ-uniform grid would correspond to different time spacings at each iteration and would not reliably capture peak violations. Second, the viewing-angle constraint is enforced via continuation: starting from a relaxed bound $\psi_{\text{start}} > \psi_{\max}$, the limit is progressively tightened toward the target $\psi_{\max}$; if a tightening step renders the problem infeasible, the step size is halved (backtracking) until progress resumes.

5. Integral-Cost Pareto Sweep and Cross-Degree Monotonicity

This section presents the experimental design for evaluating the proposed cost function (8). We perform a systematic sweep along two axes: a weight sweep that explores the acceleration–jerk trade-off, and a degree sweep that verifies the effect of increasing Bézier flexibility.

5.1 Weight Pareto sweep

For each of the five cases A–E defined in Section 3., we record the optimal trajectory and the two integrals $\int_0^{t_f} \|\mathbf{a}\|^2 \,\mathrm{d}t$ and $\int_0^{t_f} \|\mathbf{j}\|^2 \,\mathrm{d}t$. The extreme cases A and E yield the individual minima of the two axes, while the intermediate cases B/C/D sample intermediate points on the Pareto curve.

5.2 Degree sweep with degree-elevation warm start

For degrees $n \in \{5,6,7,8\}$, we solve an otherwise identical problem (same w_a, w_j, same $\psi_{\max}$, same observation window). Because every degree-$(n-1)$ Bézier curve admits an exact degree-n representation through degree elevation, the optimal cost at degree n cannot exceed that at degree $n-1$. To numerically guarantee this monotonicity, the optimization at degree n uses as its initial point the solution from degree $n-1$ lifted to the n-th degree space via exact degree elevation. Because the lifted point represents the same Bézier curve as the original (hence feasible with the same cost), any further optimization by fmincon can only preserve or reduce the cost.

5.3 Cost normalization and per-case window locking

Because the two integrals differ by several orders of magnitude ($\int \|\mathbf{a}\|^2 \sim \mathcal{O}(10^5)$, $\int \|\mathbf{j}\|^2 \sim \mathcal{O}(10^{-1})$), a direct weighted sum would be dominated by whichever unit is larger. For each $\psi_{\max}$ value, we therefore measure reference integrals $I_{a,\text{ref}}$ and $I_{j,\text{ref}}$ from Cases A and E at the first degree ($n = 5$), and reuse these as normalization coefficients across all subsequent degrees and cases. This fixed normalization enables direct comparison of J across degrees.

The observation window center for each case is determined from a pre-pass trajectory (minimizing the case's own cost, no observation constraint) as the time of closest approach to the ground target; once determined, this window is reused at all higher degrees of the same case. Per-case window locking is a prerequisite for within-case cross-degree monotonicity: only if the same window is used does the degree-elevation argument apply to the actual problem.

5.4 Success criteria

We evaluate the effectiveness of the proposed formulation by two conditions: (i) degree monotonicity, where $J(n=5) \geq J(n=6) \geq J(n=7) \geq J(n=8)$ for every case; and (ii) Pareto consistency, where at each degree, moving from Case A to Case E causes $\int \|\mathbf{a}\|^2$ to increase and $\int \|\mathbf{j}\|^2$ to decrease. Satisfying both conditions simultaneously demonstrates that the integral cost resolves the degree-monotonicity violation observed with peak-form formulations while providing a meaningful trade-off.

6. Numerical Experiments

6.1 Scenario and parameters

Table1 summarizes the nominal LEO retargeting scenario. The scenario represents a realistic sun-synchronous surveillance satellite transferring between two circular orbits at slightly different altitudes while securing observation capability of a ground target in the Korean peninsula. The transfer spans three orbital revolutions (≈4.8 hours), during which the satellite completes the altitude change while maintaining smooth, low-jerk motion. The observation window is placed near the closest approach

to the target, as determined automatically by the minimax criterion described in Section 4.. All optimizations are solved using MATLAB's fmincon with the SQP algorithm and central finite differences for gradient estimation.

Table 1 – Nominal scenario parameters.

Item	Value	Notes
μ / R_e	3.986×10^{14} m^3/s^2 / $6,378,137$ m	Earth
Initial / target altitude	550 / 520 km	Circular LEO
Inclination	97.6°	Sun-synchronous
Transfer duration	$t_f = 3$ orbits	$\approx 17,217$ s
Bézier degree	$n = m = 5$	6 control points
Observation target	$(35^\circ\text{N}, 127^\circ\text{E})$	Ground point
Window half-width	$0.01\,t_f$	≈ 344 s full width

6.2 Analysis of smoothness without observation constraint

Both formulations satisfy the same boundary conditions and produce broadly similar 3D paths under the same initial guess. The free interior spatial control points (the optimizable subset $\mathbf{P}_2, \ldots, \mathbf{P}_{n-2}$) settle at slightly different positions in the dual case, because its cost depends on time-derivative quantities computed through the chain-rule relations (5)–(7); these relations couple the spatial control points to the time schedule $t(\tau)$, so the cost gradient with respect to each spatial control point differs between the baseline (where $t(\tau) = t_f\,\tau$ is fixed) and the dual (where $t(\tau)$ is co-optimized). The most prominent difference between the two formulations lies in the temporal profile rather than in the spatial profile: the dual formulation actively redistributes physical time, as shown in Fig. 2. The baseline maintains a constant $\mathrm{d}t/\mathrm{d}\tau = t_f$ (uniform time), whereas the dual solution exhibits a wave-like profile: $\mathrm{d}t/\mathrm{d}\tau$ increases near $\tau \approx 0.15$ and $\tau \approx 0.85$ (where the satellite lingers) and decreases near $\tau \approx 0.5$ (where the satellite traverses more quickly). This redistribution shifts motion effort away from the mid-course region, precisely where the observation window will be placed in Stage 2.
The resulting acceleration and jerk profiles (Figs. 3–4) exhibit similar global trends, with endpoint peaks dominated by gravitational effects. The dual formulation, however, produces a smoother profile across all weight cases. Table 2 quantifies this advantage using the integral-cost sweep of Section 6. at $\psi_{\max} = 15^\circ$ (a loose viewing-angle setting for which the constraint is inactive: the largest $\psi(t)$ attained over the trajectory is $\max_t \psi(t) = 2.14^\circ$ for the baseline and at most 5.70° for the dual, both well below the 15° bound) and quintic Bézier degree ($n = 5$). Across the five weight cases, the dual formulation reduces $\int \|\mathbf{a}\|^2\,\mathrm{d}t$ by 53–55%, $\int \|\mathbf{j}\|^2\,\mathrm{d}t$ by 15–21%, and peak $\|\mathbf{a}\|$ by about 30%. The baseline yields identical solutions across all five cases because the spatial-only degrees of freedom cannot exploit weight variation; the same collapse phenomenon under tighter $\psi_{\max}$ is documented in detail in Section 6.8.

Table 2 – Stage 1 smoothness comparison: integral-cost metrics at $\psi_{\max} = 15^\circ$ (observation constraint inactive) and quintic Bézier degree.

	$\int \|\mathbf{a}\|^2\,\mathrm{d}t$		$\int \|\mathbf{j}\|^2\,\mathrm{d}t$		Peak $\|\mathbf{a}\|$	
Weight	Baseline	Dual	Baseline	Dual	Baseline	Dual
A $(1.00, 0.00)$	1.19×10^5	5.35×10^4	0.122	0.104	9.25	6.50
B $(0.75, 0.25)$	1.19×10^5	5.39×10^4	0.122	0.099	9.25	6.54
C $(0.50, 0.50)$	1.19×10^5	5.46×10^4	0.122	0.097	9.25	6.58
D $(0.25, 0.75)$	1.19×10^5	5.56×10^4	0.122	0.096	9.25	6.63
E $(0.00, 1.00)$	1.19×10^5	5.57×10^4	0.122	0.096	9.25	6.63

Importantly, the Stage 1 comparison serves primarily to expose the dual mechanism in isolation; the main advantage of time co-optimization becomes even more decisive under the hard constraints of

Stage 2.

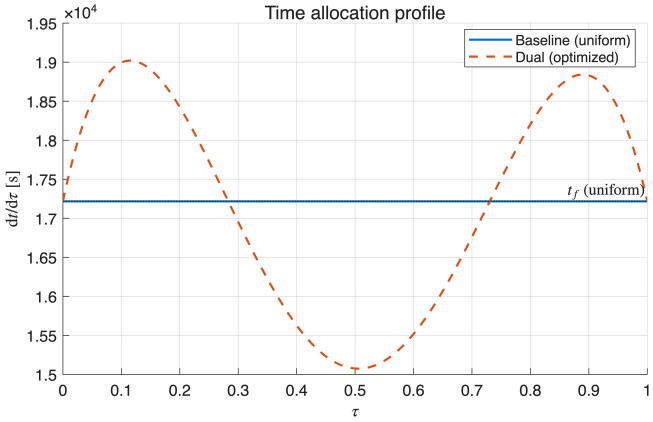


Figure 2 – Time allocation profile $\mathrm{d}t/\mathrm{d}\tau$.

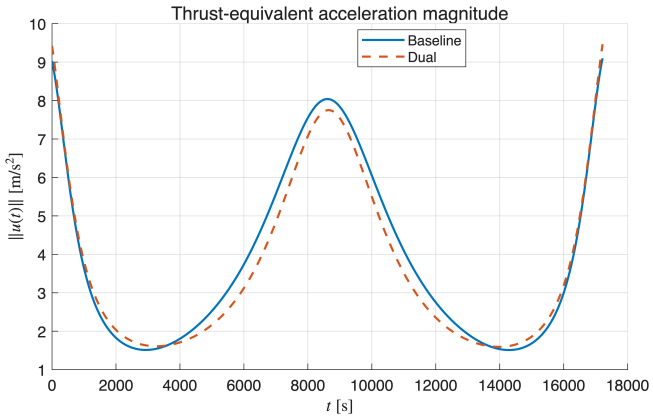


Figure 3 – Thrust-equivalent acceleration magnitude.

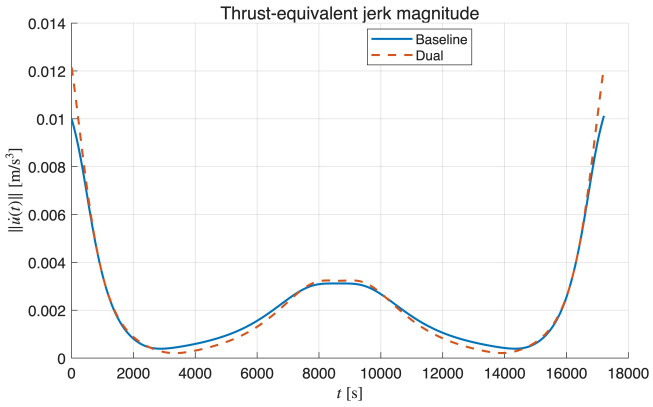


Figure 4 – Thrust-equivalent jerk magnitude.

6.3 Analysis of minimum feasible jerk cap under varying viewing-angle constraint

This subsection and the following three (Sections 6.4–6.6) employ a peak-jerk-cap hard constraint $\|\mathbf{j}(t)\| \leq j_{\max}$, the original formulation explored prior to introducing the integral-cost objective of Section 5.. The integral-cost analysis is revisited in Section 6.7; both perspectives yield consistent conclusions on dual-vs-baseline feasibility.

Figure 5 and Table 3 present the main result: the minimum feasible jerk cap $j^*_{\max}$ as the viewing-angle constraint $\psi_{\max}$ is tightened from 15° to 1.5° across nine evaluation points.

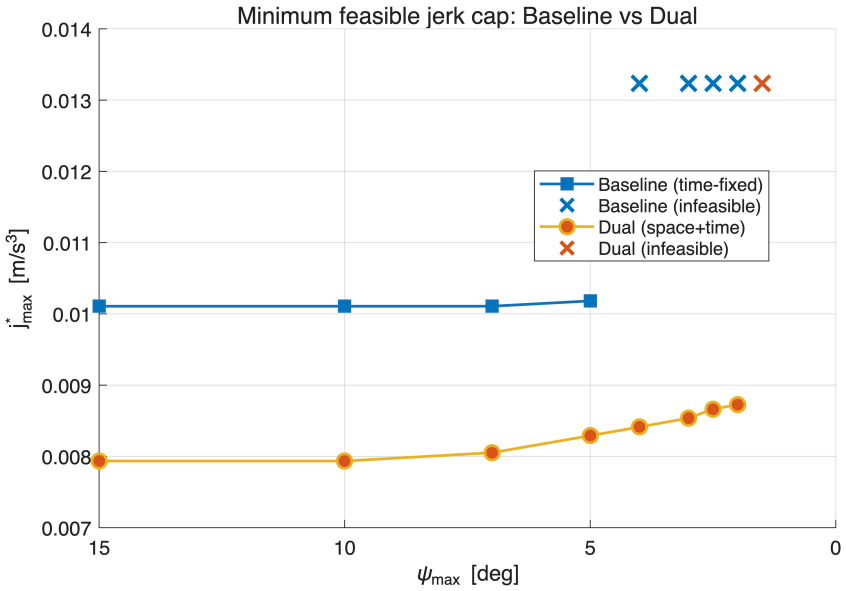


Figure 5 – Minimum feasible jerk cap $j^*_{\max}$ vs. viewing-angle constraint $\psi_{\max}$. Triangles indicate infeasible cases.

Table 3 – Minimum feasible jerk cap $j^*_{\max}$ under varying $\psi_{\max}$ (window half-width $= 0.01\,t_f$).

$\psi_{\max}$ [deg]	Baseline $j^*_{\max}$	Dual $j^*_{\max}$	Improvement
15	1.011×10^{-2}	7.936×10^{-3}	21.5%
10	1.011×10^{-2}	7.936×10^{-3}	21.5%
7	1.011×10^{-2}	8.052×10^{-3}	20.3%
5	1.018×10^{-2}	8.291×10^{-3}	18.6%
4	infeasible	8.414×10^{-3}	dual only
3	infeasible	8.537×10^{-3}	dual only
2.5	infeasible	8.663×10^{-3}	dual only
2	infeasible	8.727×10^{-3}	dual only
1.5	infeasible	infeasible	both infeasible

Three distinct regimes are visible. For $\psi_{\max} \geq 5^\circ$, both formulations are feasible, with the dual achieving 18–22% lower $j^*_{\max}$. The baseline's $j^*_{\max}$ remains nearly constant across this range ($\approx 1.01 \times 10^{-2}$), because the viewing constraint is loose enough that it does not bind; the jerk cap is determined by the unconstrained smoothness problem. In contrast, the dual's $j^*_{\max}$ gradually rises as $\psi_{\max}$ tightens, reflecting the increasing cost of satisfying the observation constraint.

For $4^\circ \leq \psi_{\max} \leq 2^\circ$, the baseline becomes infeasible: the uniform-time formulation cannot simultaneously satisfy the viewing-angle constraint and the jerk cap, regardless of spatial reshaping. The dual formulation remains feasible throughout this range, with $j^*_{\max}$ increasing from 8.41×10^{-3} to 8.73×10^{-3} as the bound tightens. This demonstrates the critical role of time-allocation freedom: by redistributing physical time so that the satellite traverses the observation window more slowly, the dual formulation can satisfy stringent viewing constraints without excessive jerk.

At $\psi_{\max} = 1.5^\circ$, both formulations become infeasible, indicating the absolute feasibility limit for this scenario and quintic Bézier degree. This limit could potentially be extended by increasing the spatial or temporal degrees of freedom (e.g., higher-degree Bézier curves or multi-segment parameterizations), which is left for future work.

6.4 Analysis of the effect of observation-window size

Figure 6 and Table 4 show the minimum feasible jerk cap as the observation window half-width increases from $0.3\%\,t_f$ to $3.0\%\,t_f$, with $\psi_{\max} = 3^\circ$ fixed.

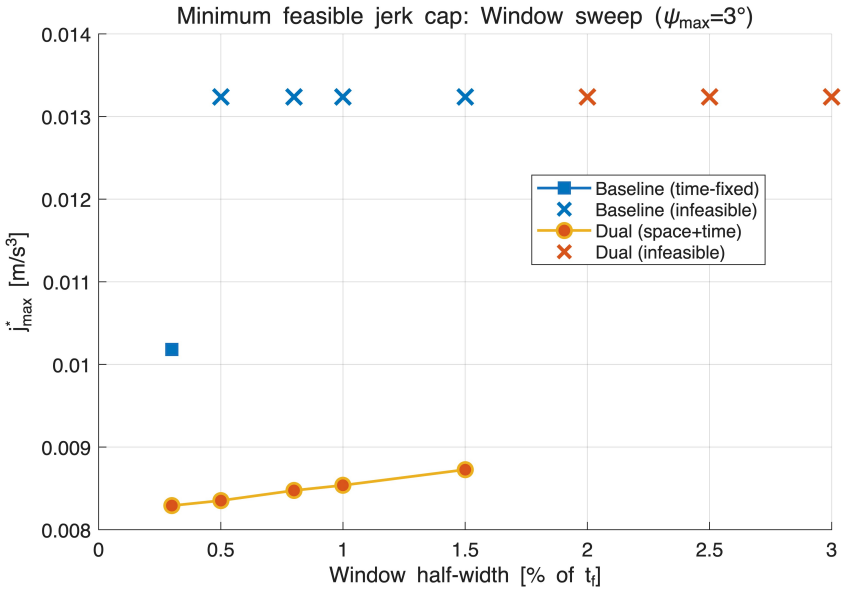


Figure 6 – Effect of observation window size on $j^*_{\max}$ ($\psi_{\max} = 3°$).

Table 4 – Window size sweep ($\psi_{\max} = 3°$): minimum feasible jerk cap $j^*_{\max}$.

Half-width [t_f]	Width [s]	Baseline $j^*_{\max}$	Dual $j^*_{\max}$	Impr. [%]
0.003	103	1.018×10^{-2}	8.291×10^{-3}	18.6
0.005	172	infeasible	8.352×10^{-3}	dual only
0.008	276	infeasible	8.475×10^{-3}	dual only
0.010	344	infeasible	8.537×10^{-3}	dual only
0.015	517	infeasible	8.727×10^{-3}	dual only
0.020	689	infeasible	infeasible	both infeas.

The baseline is feasible only at the narrowest window ($0.3\%t_f$ half-width), while the dual formulation remains feasible up to $1.5\%t_f$, roughly five times wider. Beyond $2.0\%t_f$, both formulations become infeasible, indicating the feasibility limit of the current quintic Bézier parameterization. This confirms that the time-allocation freedom of the dual formulation provides substantial margin for accommodating wider observation windows.

6.5 Robustness across orbital scenarios

Table 5 evaluates both formulations under ten orbital scenarios with $\psi_{\max} = 3°$ and window half-width $= 0.01t_f$. The first six scenarios use circular Keplerian orbits, while the last four include J_2 secular perturbations (RAAN and argument-of-perigee drift) and nonzero eccentricity to assess robustness under more realistic orbital conditions.

Table 5 – Scenario robustness ($\psi_{\max} = 3^\circ$, window $= 0.01\,t_f$).

Scenario	Baseline $j^*_{\max}$	Dual $j^*_{\max}$	Impr. [%]
Circular Keplerian			
2 orb, 550–520 km	1.129×10^{-2}	7.804×10^{-3}	30.9
3 orb, 550–520 km (nominal)	infeasible	8.537×10^{-3}	n/a
4 orb, 550–520 km	infeasible	8.703×10^{-3}	n/a
3 orb, 600–550 km	infeasible	8.273×10^{-3}	n/a
3 orb, $i = 85^\circ$	infeasible	8.475×10^{-3}	n/a
3 orb, $i = 51.6^\circ$ (ISS-like)	1.011×10^{-2}	8.792×10^{-3}	13.0
With J_2 perturbation and/or eccentricity			
3 orb, circular $+J_2$	infeasible	8.542×10^{-3}	n/a
3 orb, $e = 0.001$ $+J_2$	infeasible	8.430×10^{-3}	n/a
3 orb, $e = 0.01$ $+J_2$	infeasible	8.879×10^{-3}	n/a
3 orb, ISS, $e = 0.005$ $+J_2$	infeasible	9.140×10^{-3}	n/a

The dual formulation is feasible in all ten scenarios, whereas the baseline fails in eight out of ten. When both are feasible (short transfer or low inclination), the dual achieves 13–31% improvement. For the J_2 and elliptical scenarios, the secular RAAN drift shifts the observation geometry over the multi-orbit transfer, and eccentricity introduces altitude variation along the orbit. Despite these additional complexities, the dual formulation consistently maintains feasibility. The most challenging case, ISS-like inclination with $e = 0.005$ and J_2, requires $j^*_{\max} = 9.14 \times 10^{-3}$, which is higher than the nominal circular case but still well within the feasible range. This confirms that the proposed approach is robust not only across different Keplerian scenarios but also under more realistic orbital perturbation conditions.

6.6 Analysis of viewing-angle behavior under the hard constraint

To provide context for the observation window placement, Fig. 7 shows the viewing angle $\psi(t)$ over the full transfer duration, equal to three orbital revolutions ($t_f = 3\,T_{\text{orb}} \approx 17{,}217$ s), during which the satellite completes three revolutions around the Earth. The angle varies between approximately 0° (closest approach, best viewing) and 90° (worst viewing), with the minimum occurring near $t \approx 7{,}000$ s. The minimax criterion places the observation window at this minimum, ensuring the most favorable viewing geometry is captured.

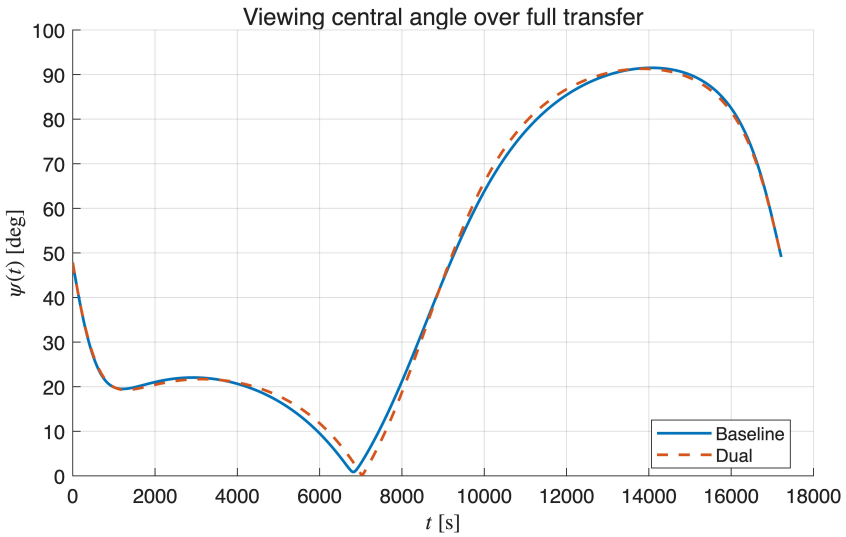


Figure 7 – Viewing central angle $\psi(t)$ over the full transfer. The observation window is placed near the global minimum ($t \approx 7{,}000$ s).

Figure 8 zooms into the observation window, comparing the Stage 1 (unconstrained) and Stage 2

(hard $\psi \leq 3^\circ$) dual solutions. Without the hard constraint, the Stage 1 solution slightly exceeds 3° near the end of the window. Stage 2 successfully suppresses $\psi(t)$ below the bound throughout the window, while the two profiles remain nearly identical outside the constrained region. The difference is subtle but critical: this small violation is precisely what makes the baseline infeasible under the hard constraint, whereas the dual formulation resolves it through minor time redistribution.

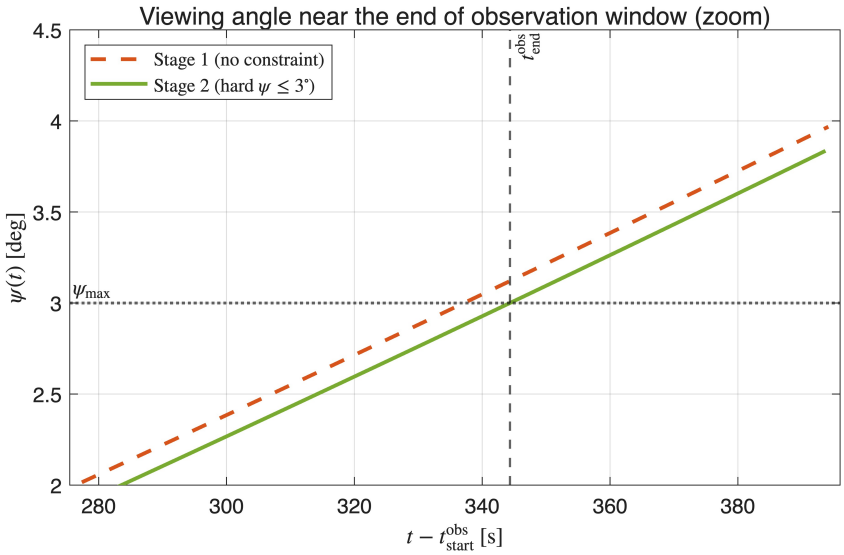


Figure 8 – Viewing angle $\psi(t)$ near the observation window.

6.7 Weight sweep and cross-degree monotonicity

Table 6 and Fig. 9 report the outcome of the integral-cost weight sweep across four Bézier degrees for two observation-window constraints ($\psi_{\max} = 1.5^\circ$ and 3°). For each $(n, w_a, w_j, \psi_{\max})$ combination, the optimization uses degree-elevation warm-start from the next-lower degree's solution. All reported values use the normalizing coefficients $I_{a,\text{ref}}^{-1}$ and $I_{j,\text{ref}}^{-1}$ from Eq. (8).

Figure 9 summarizes the sweep visually. Panels (a) and (b) plot all five cases A–E in the $(\int \|\mathbf{a}\|^2, \int \|\mathbf{j}\|^2)$ plane at $n = 8$ for the two $\psi_{\max}$ settings, under the dual formulation: Case A occupies the lower-left corner of each plot (lowest in both metrics), while Cases B–E cluster in the upper-right. Panels (c) and (d) report the cross-degree convergence of J_{norm}, contrasting the dual (solid lines) and baseline (dashed lines) formulations. Under the dual every case decreases monotonically with degree; under the baseline only Case A improves, and Cases B–E remain effectively flat near $J_{\text{norm}} = 1$, with the four B–E curves coinciding to within a 0.01 spread (highlighted in panel (c) by the ellipse and annotation).

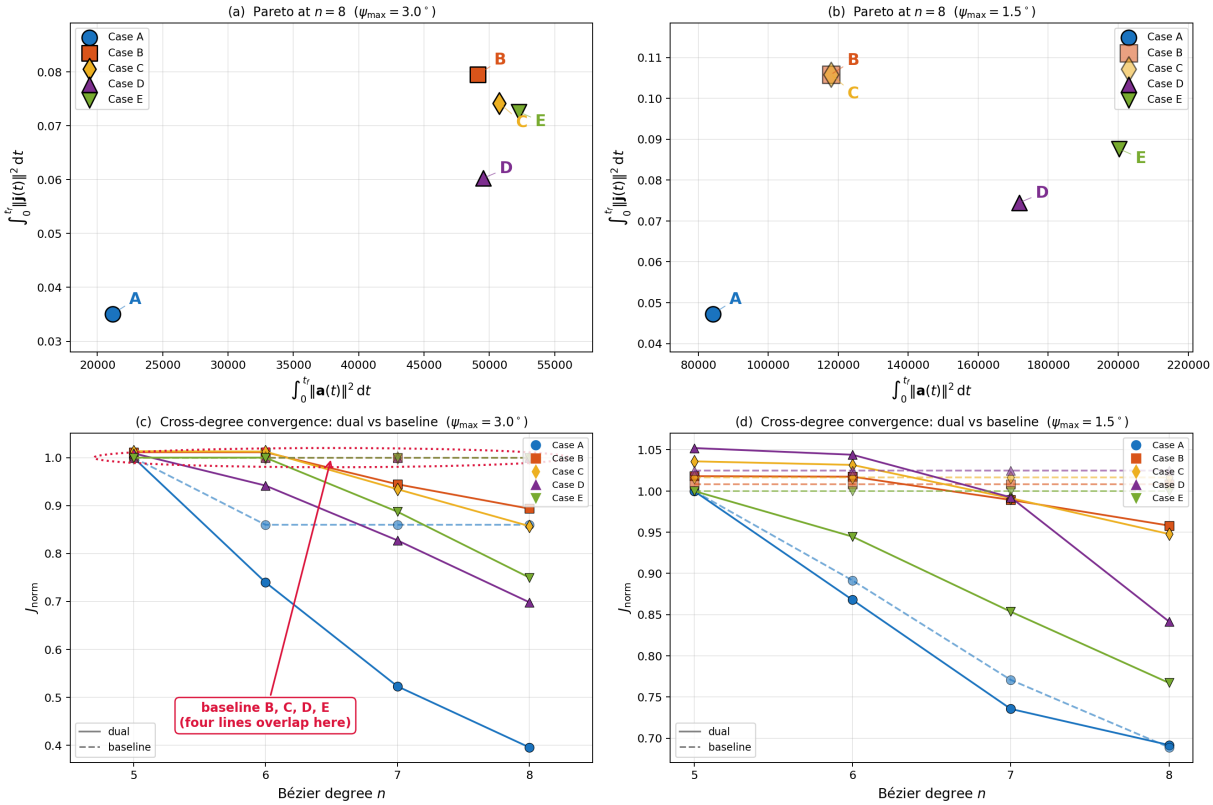


Figure 9 – Pareto distribution at $n = 8$ (a, b) and cross-degree convergence of J_{norm} (c, d), at $\psi_{\max} = 3^\circ$ (left column) and $\psi_{\max} = 1.5^\circ$ (right column).

Table 6 – Normalized cost J_{norm} across Bézier degree, weight case, and $\psi_{\max}$.

			J_{norm} by weight case				
$\psi_{\max}$	n	Free vars	A $(1, 0)$	B $(0.75, 0.25)$	C $(0.5, 0.5)$	D $(0.25, 0.75)$	E $(0, 1)$
3.0°	5	9	1.000	5.011	3.788	2.510	1.000
	6	13	0.740	4.794	3.721	2.510	0.884
	7	17	0.523	4.761	3.688	2.407	0.811
	8	21	0.395	2.107	1.779	1.219	0.753
1.5°	5	9	1.000	1.556	1.391	1.184	1.000
	6	13	0.868	1.544	1.358	1.145	0.970
	7	17	0.736	1.487	1.280	1.062	0.882
	8	21	0.692	1.392	1.047	0.869	0.820

Two observations support the effectiveness of the proposed integral formulation.

Degree monotonicity holds across all $(\psi_{\max}, \text{case})$ combinations. For every combination of weight case and $\psi_{\max}$, $J_{\text{norm}}(n) \leq J_{\text{norm}}(n-1)$ is satisfied for $n = 6, 7, 8$, with strict decrease in all but two (tied) cells. This is the behavior theoretically guaranteed by the exact degree-elevation property of Bézier curves. In contrast, the peak-form (min-max) formulation previously exhibited non-monotone behavior at $\psi_{\max} = 1.5^\circ$ between $n = 7$ and $n = 8$; the integral cost resolves this violation.

Feasibility extends to the lowest degree under the tightest window. At $\psi_{\max} = 1.5^\circ$, all five cases are feasible starting from $n = 5$. This contrasts with the peak-form formulation, which was infeasible at $n = 5$ under $\psi_{\max} = 1.5^\circ$ and required $n \geq 6$ for any feasible solution. The integral cost thus broadens the applicable design space, enabling narrow-window surveillance missions with the simplest quintic parameterization.

As the weight shifts from Case A to Case E, $\int \|\mathbf{a}\|^2 \, dt$ increases while $\int \|\mathbf{j}\|^2 \, dt$ decreases, the expected trade-off. Intermediate cases B/C/D occupy the interior of the Pareto frontier. At each degree, the

frontier shifts inward toward the origin, consistent with the larger feasible space of higher-degree curves.

6.8 Aggregate analysis

The per-case examination above reports outcomes at individual $(\psi_{\max}, n)$ operating points under the dual formulation. We complement it with aggregate statistics across the full $7 \times 4 \times 5$ sweep grid, replicated for both formulations: 140 dual runs and 140 baseline runs (uniform timing, otherwise identical optimization setup including $\psi_{\max}$ values, degrees, weight cases, and per-case window locking). Table 7 compares the two formulations across several diagnostic measures.

Table 7 – Aggregate sweep statistics: dual vs. baseline (140 + 140 = 280 runs, all converged).

Property	Dual	Baseline
Cross-degree monotonicity (J_{norm} decreases $n{=}5 \to 8$)	$35/35$ (100%)	$9/35$ (26%)
Mean J_{norm} improvement, Case A	50.8%	28.1%
Mean J_{norm} improvement, Cases B–E	18.2%	$\approx 0\%$
Pareto cleanliness ($\int \|\mathbf{a}\|^2 \uparrow$, $\int \|\mathbf{j}\|^2 \downarrow$ along A$\to$E)	$8/28$ (29%)	$0/28$ (0%)
Mean relative spread of $\int \|\mathbf{a}\|^2$ across Cases B–E	29.2%	9.3%
Mean relative spread of $\int \|\mathbf{j}\|^2$ across Cases B–E	15.5%	0.7%

The contrast is sharp. Under the dual formulation, every $(\psi_{\max}, \text{case})$ combination achieves cross-degree monotonicity from $n{=}5$ to $n{=}8$, confirming the theoretical nesting guarantee of Section 5. in numerical practice; the corresponding mean J_{norm} improvement is 50.8% for Case A and 18.2% averaged over Cases B–E. Under the baseline, only 9 of 35 combinations are monotonic, and Cases B–E exhibit essentially zero J_{norm} improvement on average as the degree increases: the higher-degree spatial control points are not exploited when the time schedule is fixed. Only Case A benefits from increased Bézier degree under baseline, because its objective is most aligned with the spatial path-constraint structure.

The same collapse pattern appears in the A$\to$E Pareto distribution. Cleanliness is achieved in 8 of 28 $(\psi_{\max}, n)$ slices under the dual formulation but in 0 of 28 slices under the baseline. To quantify how distinct Cases B–E are from each other, we measure the relative spread $(\max - \min)/\text{mean}$ of $\int \|\mathbf{a}\|^2$ and $\int \|\mathbf{j}\|^2$ across B–E within each slice. The dual formulation produces a mean spread of 29% in $\int \|\mathbf{a}\|^2$ and 15% in $\int \|\mathbf{j}\|^2$, whereas the baseline produces 9% and 0.7%, respectively: under baseline, Cases B/C/D/E land at essentially the same $(\int \|\mathbf{a}\|^2, \int \|\mathbf{j}\|^2)$ point, with no meaningful trade-off representation possible. This is consistent with the few decision degrees of freedom available to a baseline quintic optimization under tight path constraints; the feasible region narrows to a single trajectory neighborhood, and weight variation alone cannot pull the solution to distinct smoothness trade-offs.

These observations elevate the dual formulation's advantage from a local quantitative gain to a qualitative one: time co-optimization is what activates the weight-driven trade-off space and what allows higher Bézier degree to benefit every case rather than only the acceleration-only objective. Within the dual results, a degree-dependent residual asymmetry remains: the A-dominance ratio (mean of B–E over A) grows from 1.05 at $n{=}5$ to 2.73 at $n{=}8$ for $\int \|\mathbf{a}\|^2$, indicating that even with time co-optimization, the degree-elevation warm-start procedure leaves B–E in the basin associated with A's solution while A pulls away.

7. Discussion

The central insight unifying the preceding analyses is that time allocation, not spatial shape alone, governs both feasibility under a hard viewing-angle constraint and the expressivity of the weighted-cost trade-off. Under uniform timing the binding path constraint narrows the spatial feasible set to a neighborhood in which weight variation produces near-identical trajectories and degree elevation provides no benefit beyond Case A; the dual formulation enlarges the feasible region along the orthogonal axis of physical-time redistribution, simultaneously rescuing feasibility at tight $\psi_{\max}$, opening

a meaningful A-to-E Pareto trade-off, and giving every weight case room to improve as the Bézier degree increases when paired with degree-elevation warm-starting.

The choice of Bernstein/Bézier polynomials over alternative parameterizations such as B-splines, pseudospectral methods, or Fourier series reflects three properties favorable to small-scale, smoothness-oriented optimization: the convex-hull property bounds the trajectory geometrically without curve evaluation, the endpoint interpolation and derivative-at-endpoint structure make boundary-condition enforcement trivial, and only $n-3$ interior control points (per dimension) remain free after fixing endpoint positions and velocities (8 variables in the quintic dual case) so that SQP solves the problem efficiently with finite-difference gradients. B-splines offer similar properties with the additional advantage of local support, while pseudospectral methods directly enforce dynamics as equality constraints; both become preferable in regimes with many control points, multi-segment paths, or high-fidelity dynamics not addressed here.

Several limitations remain. The trajectory parameterization is purely kinematic, so the Bézier curves are not constrained to satisfy orbital dynamics and the framework should be understood as a trajectory-shaping tool rather than a flight-ready trajectory generator; the J_2 perturbation is modeled only through secular effects on RAAN and argument of perigee, which is reasonable for the multi-orbit LEO transfers considered here but may require refinement for longer-duration or eccentric scenarios; and while every dual run achieves cross-degree monotonicity, the A-to-E Pareto front retains a residual local-minimum-clustering pattern in which Cases B–E settle in the basin reached via Case-A warm-starting, with the case-D anomaly $\int \|\mathbf{j}\|_D^2 < \int \|\mathbf{j}\|_E^2$ appearing in roughly half of the $(\psi_{\max}, n)$ slices. Natural follow-ups include higher-fidelity dynamics with equations of motion imposed as equality constraints, multi-window and multi-target retargeting, multi-segment Bézier parameterizations for wider observation windows, and a multi-start procedure that initializes Cases B–E from solutions other than Case A to clean up the weight-Pareto representation.

Despite the kinematic nature of the parameterization, the framework supports several mission-design tasks. In early-phase design, the dual Bézier formulation maps feasibility boundaries (such as the maximum achievable viewing-angle constraint for a given jerk cap) and acceleration–jerk Pareto curves at modest computational cost, informing actuator sizing and observation-quality budgeting without running a full dynamic trajectory optimization. The smooth, constraint-satisfying reference it produces also serves as a useful warm start for high-fidelity optimizers such as pseudospectral or direct-collocation solvers, which are sensitive to initialization. The systematic sweeps reported in Sections 6.3–6.5 demonstrate the resulting capability for trade-space exploration over observation-window width, viewing-angle constraint, transfer duration, and orbital geometry.

8. Conclusions

We presented a dual Bézier framework that jointly optimizes spatial shape and time allocation under smoothness criteria and a hard observation-window inequality. A robust two-stage strategy, namely fixed-length time-uniform window sampling and continuation on $\psi_{\max}$ with backtracking, stabilizes SQP convergence.

Two complementary analyses corroborate the dual formulation. Under a peak-jerk-cap hard constraint (Sections 6.3–6.6), the uniform-time baseline becomes infeasible under stringent viewing constraints while the dual formulation remains feasible with a smaller required jerk allowance; the dual formulation also accommodates substantially wider observation windows, and the advantage is qualitatively consistent across orbital scenarios including cases with J_2 secular perturbations and nonzero eccentricity. Under an integral-cost objective combining acceleration and jerk energies (Section 5. , Section 6.7), a weight Pareto sweep over five representative cases traces the acceleration–jerk trade-off, and a cross-degree comparison ($n = 5$ through 8) using exact degree-elevation warm starts reproduces the monotonic cost decrease theoretically guaranteed by Bézier-space nesting. Replicating the same sweep on the baseline (Section 6.8) yields no clean Pareto slices and monotonicity in only 9 of 35 $(\psi_{\max}, \text{case})$ combinations, with Cases B–E producing essentially identical solutions: under uniform timing, the weight-driven trade-off space collapses and increased Bézier degree benefits only the acceleration-only objective.

Together, these results show that the decisive role of time redistribution holds across multiple smoothness criteria: not only in feasibility under windowed viewing constraints (peak-jerk-cap analysis), but

also in the activation of the weight-driven Pareto trade-off and in the per-case benefit of Bézier degree elevation (integral-cost analysis). The integral-cost formulation provides a stable, theoretically grounded trade-off representation compatible with the dual parameterization.
Future work will address multi-window and multi-target retargeting, higher-fidelity dynamics (e.g., drag, solar radiation pressure), and multi-segment Bézier parameterizations for further spatial flexibility under wider observation windows.

9. Acknowledgments

This work was supported by the Korea Research Institute for Defense Technology Planning and Advancement (KRIT) grant funded by the Korea government (Defense Acquisition Program Administration, DAPA) (No. KRIT-CT-22-030, ReUSV-41, 2026).

10. Contact Author Email Address

Corresponding author: Suwon Lee (Tel: +82-10-8501-5086, suwon.lee@kookmin.ac.kr)

11. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] F. L. Markley and J. L. Crassidis. *Fundamentals of Spacecraft Attitude Determination and Control*. Springer, 2014.

[2] P. W. Likins. *Dynamics and Control of Flexible Space Vehicles*. JPL Technical Report 32-1329, Jet Propulsion Laboratory, 1971.

[3] P. C. Hughes. *Spacecraft Attitude Dynamics*. Dover Publications, 2004.

[4] V. Lappas and V. Kostopoulos. A survey on small satellite technologies and space missions for geodetic applications. *Progress in Aerospace Sciences*, 98:1–14, 2018.

[5] Y. Ulybyshev. Satellite constellation design for complex coverage. *Journal of Spacecraft and Rockets*, 45(4):843–849, 2008.

[6] J. R. Wertz, D. F. Everett, and J. J. Puschell. *Space Mission Engineering: The New SMAD*. Microcosm Press, 2011.

[7] S. Lee and Y. Kim. Optimal output trajectory shaping using Bézier curves. *Journal of Guidance, Control, and Dynamics*, 44(5):1027–1035, 2021.

[8] C. Kielas-Jensen, V. Cichella, T. Berry, I. Kaminer, C. Walton, and A. Pascoal. Bernstein polynomial-based method for solving optimal trajectory generation problems. *Sensors*, 22(5):1869, 2022.

[9] V. Cichella, I. Kaminer, C. Walton, N. Hovakimyan, and A. Pascoal. Optimal multivehicle motion planning using Bernstein approximants. *IEEE Transactions on Automatic Control*, 63(11):4000–4015, 2018.

[10] C. Walton, I. Kaminer, and V. Cichella. Trajectory optimization using Bernstein basis polynomials with endpoint constraints. *IEEE Control Systems Letters*, 6:3014–3019, 2022.

[11] W. Sun, G. Tang, and K. Hauser. Fast UAV trajectory optimization using bilevel optimization with analytical gradients. *IEEE Transactions on Robotics*, 37(6):2010–2024, 2021.

[12] K. L. Hobbs and J. T. Kaneshige. Jerk-limited trajectory generation for unmanned aerial systems. In *AIAA Scitech 2019 Forum*, pages AIAA 2019–0919, 2019.

[13] C. Louembet, D. Arzelier, and G. Deaconu. Minimum fuel deployment for spacecraft formations via optimal control. *Journal of Guidance, Control, and Dynamics*, 38(12):2371–2380, 2015.

[14] J. Ventura, M. Ciarcià, M. Romano, and U. Walter. Fast and near-optimal guidance for docking to uncontrolled spacecraft. *Journal of Guidance, Control, and Dynamics*, 40(12):3138–3154, 2017.

[15] J. T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. SIAM, 2nd edition, 2010.

[16] R. T. Farouki. The Bernstein polynomial basis: a centennial retrospective. *Computer Aided Geometric Design*, 29(6):379–419, 2012.