Constrained Satellite Rendezvous with HCW Dynamics: LQR, Standard MPC, and Economic MPC

Raaghav Thirumaligai

March 18, 2026

Abstract

This report studies relative satellite rendezvous in the Hill frame using the Hill-Clohessy-Wiltshire (HCW) linearized model with actuation bounds. Three controllers are compared on the same discrete-time model: (i) saturated LQR, (ii) a state-focused constrained MPC (MPCv2), and (iii) an economic MPC (MPCv1) with fuel-prioritized stage cost. The main findings are that constrained MPC expands the practically stable set of initial conditions compared with saturated LQR, while economic MPC reduces control-effort proxy at the cost of slower convergence. A key observation during implementation with MPCTools [5] is that hard terminal sets can improve system dynamics and control effort but may induce intermittent infeasibility when thrust authority is tight.

1 Introduction

Autonomous relative motion control is central to satellite flight operations such as rendezvous, docking, and inspection. In this setting, fuel is expensive, and actuator authority is limited, making constrained control design more important than unconstrained optimality. MPC is a natural choice because constraints on thrust can be enforced directly during optimization.

The goal of this project is to compare a traditional controller and MPC variants on a single satellite relative-motion benchmark. The comparison is intentionally practical: all controllers use the same discretized HCW dynamics and are evaluated under similar bounded-input assumptions. The three controllers are:

1. Saturated LQR baseline,
2. State-focused constrained MPC (fast convergence target),
3. Economic MPC (fuel-prioritized objective).

2 Problem Statement

HCW Dynamics

The relative state is

$$\mathbf{x} = [x,\ y,\ z,\ \dot{x},\ \dot{y},\ \dot{z}]^\top,$$

with control acceleration

$$\mathbf{u} = [u_x,\ u_y,\ u_z]^\top.$$

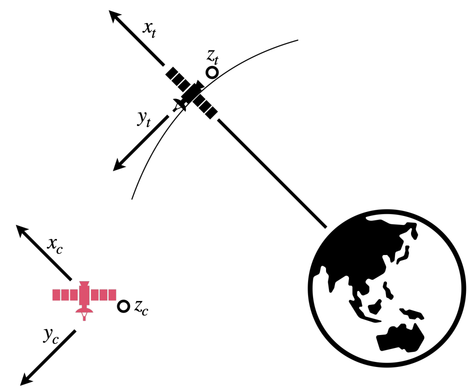

Figure 1: Relative-motion setting in the Hill frame. The directions of the coordinates are set by the target satellite. Then devieated variables are constructed to drive the difference between the chaser satellite(pink) and target to zero.

The continuous-time model is

$$\dot{\mathbf{x}} = A_c\mathbf{x} + B_c\mathbf{u}, \tag{1}$$

where

$$A_c = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 3n^2 & 0 & 0 & 0 & 2n & 0 \\ 0 & 0 & 0 & -2n & 0 & 0 \\ 0 & 0 & -n^2 & 0 & 0 & 0 \end{bmatrix}, \quad B_c = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{2}$$

The orbital-rate parameter, n, is dependent on a, the target's radius and μ the standard gravitational parameter: (parameter values from [4])

$$\mu = 3.398 \times 10^{14}\ \mathrm{m}^3/\mathrm{s}^2, \quad a = 6{,}793{,}137\ \mathrm{m}, \quad n = \sqrt{\frac{\mu}{a^3}}\ \mathrm{rad/s}.$$

In simulation, a fixed sampling time is used and zero-order hold (ZOH) discretization yields

$$\mathbf{x}_{k+1} = A_d\mathbf{x}_k + B_d\mathbf{u}_k.$$

Then converting from the continuous to the discrete system (with $T_s = 10s$) is as follows:

$$\begin{aligned} A_d &= e^{A_c T_s} \\ B_d &= \int_0^{T_s} e^{A_c\tau} B_c \, d\tau \end{aligned} \quad \Rightarrow \quad \exp\left(\begin{bmatrix} A_c & B_c \\ 0 & 0 \end{bmatrix} T_s\right) = \begin{bmatrix} A_d & B_d \\ 0 & I \end{bmatrix}.$$

Constraints and Objective

All controller designs enforce bounded acceleration

$$|u_i| \le a_{\max}, \quad i \in \{x, y, z\},$$

with $a_{\max} = 0.5\,\mathrm{mm/s}^2$ in this study, chosen from the constraint scale in [3]. The state is driven toward rendezvous near the origin. The comparison emphasizes two competing goals:

- rapid convergence to a small terminal neighborhood,
- low cumulative control effort (fuel proxy).

3 Controller Formulations

Saturated LQR (Baseline)

A discrete-time LQR gain K is computed from (A_d, B_d, Q, R), then clipped to satisfy actuator limits:

$$\mathbf{u}_k = \mathrm{sat}(-K\mathbf{x}_k, a_{\max}).$$

This baseline is simple and fast but does not optimize future behavior under constraints.

State-Focused MPC (MPCv2)

MPCv2 solves, at each step, a finite-horizon constrained OCP:

$$\min_{\{\mathbf{u}_k\}} \sum_{j=0}^{N_t-1} \left(\mathbf{x}_{k+j}^\top Q \mathbf{x}_{k+j} + \mathbf{u}_{k+j}^\top R \mathbf{u}_{k+j} \right) + \mathbf{x}_{k+N_t}^\top Q_f \mathbf{x}_{k+N_t} \tag{3}$$

subject to dynamics and input bounds, with terminal-set constraints used in selected runs. This controller is tuned for aggressive state reduction.

Economic MPC (MPCv1)

MPCv1 uses a fuel-prioritized stage cost, approximating fuel usage with an L1-type control term:

$$\ell(\mathbf{x}, \mathbf{u}) = w_x \, \mathbf{x}^\top \mathbf{x} + w_{u,1} \sum_i \sqrt{u_i^2 + \varepsilon^2}. \tag{4}$$

Compared with MPCv2, this formulation accepts slower transient response to reduce control effort.

4 Simulation Setup

The same ZOH-discretized HCW model is used for all controllers. Monte Carlo experiments sample random initial conditions over a bounded position-velocity box, and each controller is run on the same seeds/initial-condition set.

Reported metrics are:

1. convergence rate within the simulation horizon,
2. time-to-convergence (steps),
3. cumulative control effort proxy, $\sum_k \|\mathbf{u}_k\|_1$,
4. solver success ratio for MPC runs.

5 Results

Qualitative Trajectory Behavior

Single-case trajectories show the expected behavior: LQR can require more time and may fail under saturation; MPCv2 drives faster toward rendezvous; MPCv1 remains stable but is intentionally less aggressive when fuel is prioritized.

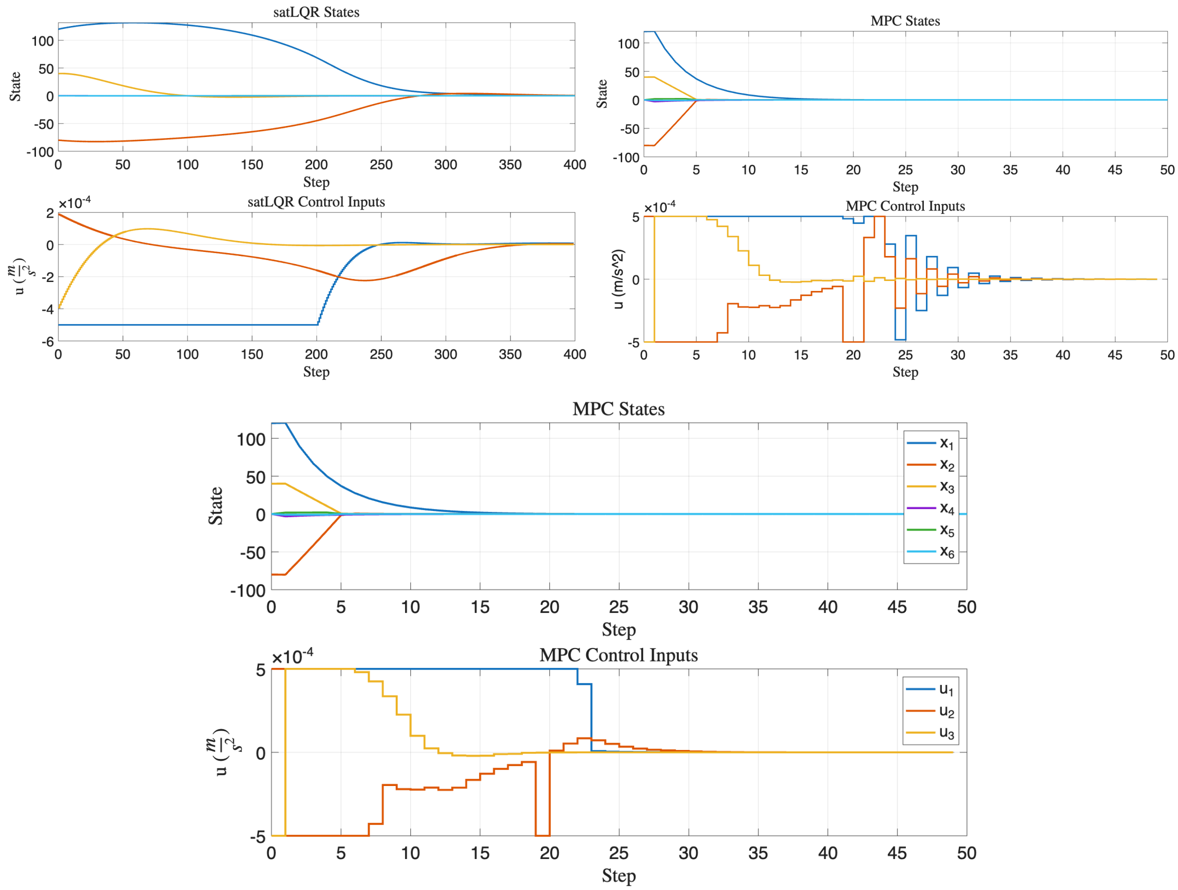

Figure 2: Representative single-case responses: saturated LQR (left-top), state-focused MPCv2 (right-top), and economic MPCv1 (bot). The difference in time-scale is worth pointing out as LQR takes much longer to stabilize. The oscillating input signature in MPCv2 is not in MPCv1 which may be where the controller is able to eek out a drop of extra fuel.

Multi-Case Comparison

The aggregate trajectory overlay (Fig. 3) and histogram comparisons support the controller tradeoff story: MPC formulations improve constrained behavior relative to saturated LQR, and economic MPC reduces effort relative to the state-focused MPC in many cases.

Observed numerical trend. The project experiments indicate that economic MPC reduced input-effort proxy compared with the state-focused MPC and LQR baselines, while convergence was generally slower. From the current run logs, the decrease in 'fuel' was 11.06% (MPCv2 to MPCv1) and 79.93% (LQR to MPCv1).

6 Discussion

A key practical lesson is the sensitivity of hard terminal sets under tight thrust limits. Hard terminal constraints can improve targeting accuracy but can also generate intermittent infeasibility or hit solver iteration-limits. Softening terminal requirements (or using strong terminal penalties with realistic

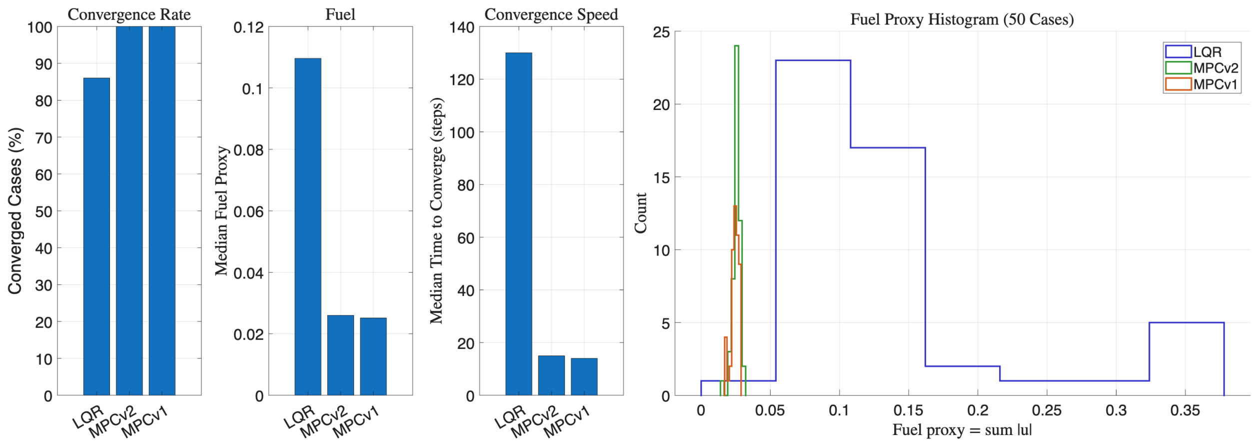

Figure 3: Key issues to notice here are slower and even lack of convergence to the equilibrium for LQR and much larger cost to stabilize.

terminal regions) improves practical robustness. (In this example, I needed to allow a much larger control input and much larger terminal set to always get feasibility.)

Another lesson is objective-shaping: if the cost heavily penalizes input, shorter horizons may appear better because long horizons accumulate large effort penalties, creating overly conservative control. This does not necessarily imply better closed-loop optimality; feasibility and convergence metrics must be reported jointly.

7 Conclusions and Future Work

This work implemented and compared three controllers for constrained HCW rendezvous using the same discretized model. The results support MPC for constrained satellite tracking, and show that economic MPC can further reduce control effort when slower convergence is acceptable.

Immediate next steps are to add estimation and robustness layers: moving-horizon estimation (MHE) with angle-only measurement models, and direct comparison to EKF/KF state estimation under measurement noise.

References

[1] W. H. Clohessy and R. S. Wiltshire, "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, 1960.

[2] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed., Nob Hill Publishing, 2017.

[3] X. Lin, G. Zhang, and R. Vazquez, "Spacecraft Constant-Thrust Minimum-Time Rendezvous via Reachable Set Theory," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 61, no. 5, pp. 13224–13232, 2025, doi: 10.1109/TAES.2025.3576316.

[4] "Clohessy–Wiltshire equations," Wikipedia, https://en.wikipedia.org/wiki/Clohessy%E2%80%93Wiltshire_equations, accessed Mar. 9, 2026.

[5] J. B. Rawlings Group, Risbeck, M.J., Rawlings, J.B., 2016. MPCTools: Nonlinear model predictive control tools for CasADi (Octave interface).

A Trajectory Sweep

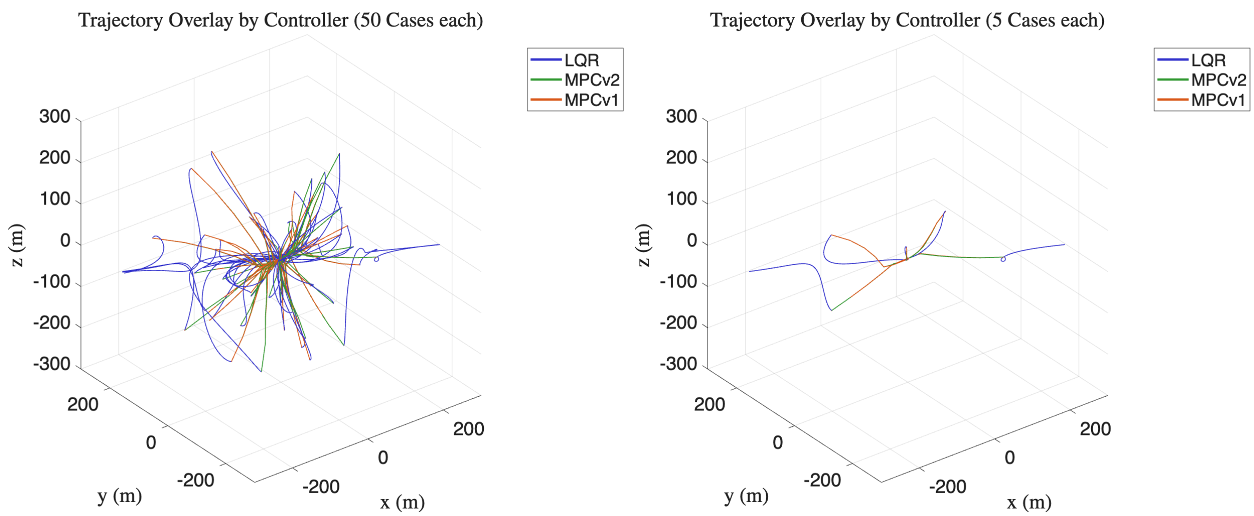

Figure 4: Trajectories from the simulation. All controllers run from the same initial condition; the right panel highlights five runs, with two LQR trajectories diverging.

B Additional LQR–MPC Comparison

This appendix summarizes the extra comparison study of separate two effects: (i) controller structure (LQR vs MPC) and (ii) hard input limits. The first pair of plots shows the unconstrained-style baseline behavior, where both controllers are free to use aggressive actuation. The second pair shows the same setup after imposing acceleration limits, highlighting how saturation changes response shape and convergence quality. The initial tuning is chosen so that after saturating the LQR, the trajectory is stable from this initial condition. Clearly in Figure 5 we see controllers produce the same solutions.

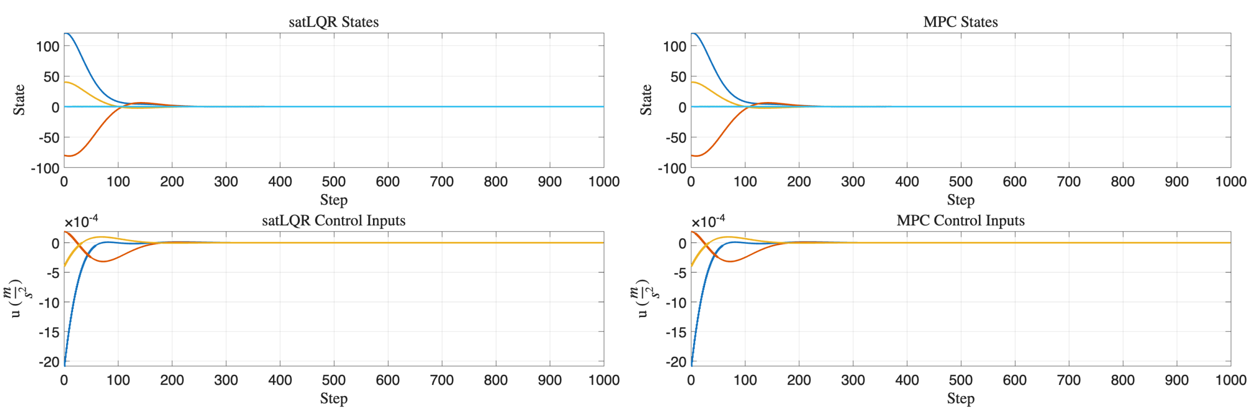

Figure 5: Unconstrained-style baseline comparison from the bonus analysis: LQR (left) and MPC (right).

Under bounded control, the distinction becomes more practically important: MPC keeps constraint handling inside the optimization problem at each step, while saturated LQR relies on clipping after

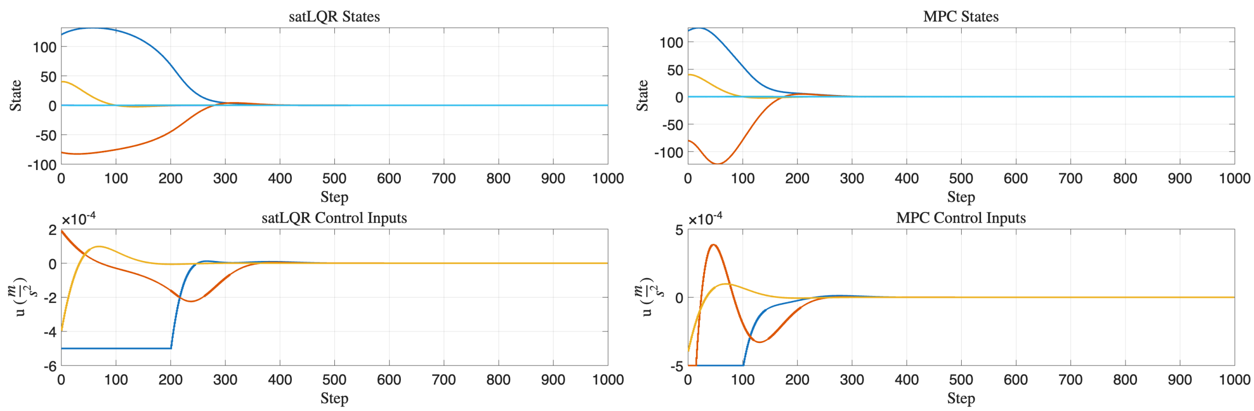

Figure 6: Constrained comparison with bounded input: constrained LQR (left) and constrained MPC (right).

gain application. In this study, the constrained MPC formulation remained the more consistent option for driving states toward the rendezvous region under the same actuation bound.

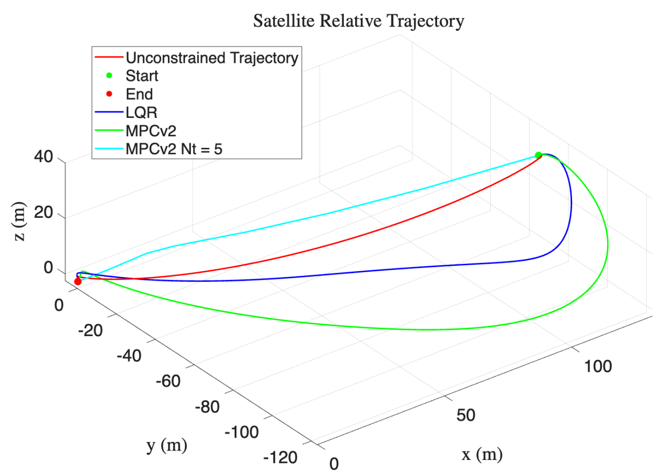

Figure 7: Final trajectory comparison from the bonus study across formulation settings and horizon choice.

The final comparison plot (Figure 7) reinforces the central report conclusion: tuning choices (especially constraints and horizon) materially affect both convergence behavior and total control effort. This is consistent with the main text, where economic MPC reduced fuel proxy at the cost of slower convergence.