

CONTROLLED TRANSPORT AND STATION KEEPING BETWEEN PERIODIC CYCLES IN THE STANDARD MAP

Dirk M. Luchtenburg* and Shane D. Ross†

We study bounded controlled transport between periodic cycles of the standard map as a testbed for constructing directed transport graphs between invariant structures. Periodic-cycle phase points are treated as nodes, and feasible bounded-control transfers are edges labeled by transfer time, actuation effort, and post-arrival station-keeping quality. For prescribed source and target phase points, we compute nominal open-loop finite-horizon minimum-effort transfers into a target neighborhood under an additive control bound, then use bounded linear time-varying model predictive control for closed-loop tracking and station keeping. Numerical examples for hyperbolic and elliptic target cycles compare natural and controlled transport, quantify effort–time and handoff tradeoffs, and demonstrate feedback correction under a small initial-condition perturbation. The standard-map results provide a controlled setting for methods intended for later extension to Keplerian maps and astrodynamics transfer design.

INTRODUCTION

Invariant structures such as periodic orbits, invariant manifolds, resonant regions, and mixed phase-space structures organize transport in many nonlinear dynamical systems.¹ In astrodynamics, analogous structures are used to design low-energy transfers and station-keeping strategies near desired orbits. A central trajectory-design question is how limited actuation can modify natural transport between such structures while still exploiting the underlying dynamics.

The standard map provides a compact discrete-time setting in which this question can be studied with high numerical resolution.² Although the standard map is not itself a spacecraft model, it contains many of the phase-space features relevant to transport design, including invariant curves, resonant islands, hyperbolic periodic cycles, elliptic periodic cycles, and transport barriers.³ For this reason, it is useful as a testbed for control methods intended to transfer between dynamically meaningful structures. From an astrodynamics viewpoint, the standard-map phase-space structures play a role analogous to resonant regions, libration-point orbits, and manifold tubes in multi-body dynamics.

This paper studies bounded controlled transport between selected phase points on periodic cycles of the standard map. Given a source phase point and a target neighborhood around a selected phase point on another cycle, the objective is to reach the target neighborhood in few map iterates while respecting a pointwise control bound and limiting accumulated control effort. The resulting

*Associate Professor, Mechanical Engineering, Cooper Union, New York, NY, USA, dirk.luchtenburg@cooper.edu

†Professor, Aerospace and Ocean Engineering, Virginia Tech, Blacksburg, VA, USA, sdross@vt.edu

performance quantities include feasibility, first-entry time, required control bound, accumulated effort, and robustness under perturbations.

A central distinction in the paper is between *transport* and *station keeping*. The transport problem ends when the trajectory first enters the target neighborhood. The station-keeping problem begins after arrival and asks whether feedback can maintain the trajectory near the full target periodic cycle. These objectives are related but not equivalent: a trajectory may reach the target neighborhood without being well positioned for subsequent maintenance, while a local station-keeping controller does not by itself solve the global transport problem.

The computational pipeline reflects this separation. First, uncontrolled manifold-seeded trajectories are computed to establish natural transport benchmarks. The main transfer-generation step solves direct finite-horizon bounded optimal-control problems⁴ to compute nominal open-loop transfers, consisting of reference trajectories and associated feedforward control sequences. These nominal transfers are then executed using bounded linear time-varying model predictive control (LTV MPC) as a receding-horizon feedback controller.⁵ During transport, the controller tracks the nominal transfer; after first entry into the target neighborhood, it switches to cyclic station keeping on the target cycle.

We use MPC for station keeping because the same feedback formulation applies to both hyperbolic and elliptic target cycles. By contrast, OGY-style local stabilization relies on local hyperbolic structure⁶ and is therefore not directly applicable to elliptic target cycles. This is important for the present study because we compare hyperbolic-to-hyperbolic and hyperbolic-to-elliptic transfers using the same online feedback framework.

The broader motivation is to organize controlled transfers into a directed transport graph. In such a graph, periodic-cycle phase points serve as nodes, and feasible bounded-control transfers serve as directed edges labeled by first-entry time, control effort, required control authority, terminal accuracy, and post-arrival station-keeping quality. The standard map provides a controlled setting in which this graph-building logic can be developed before extension to maps arising in astrodynamics, such as Keplerian maps or Poincaré maps associated with higher-fidelity multi-body models.⁷

PROBLEM FORMULATION

Uncontrolled and controlled standard maps

We consider the standard map with parameter K .² The uncontrolled dynamics are

$$p_{k+1} = p_k + K \sin(\theta_k), \quad (1)$$

$$\theta_{k+1} = \theta_k + p_{k+1}. \quad (2)$$

The controlled map uses an additive bounded input in the p -update:

$$p_{k+1} = p_k + K \sin(\theta_k) + u_k, \quad (3)$$

$$\theta_{k+1} = \theta_k + p_{k+1}. \quad (4)$$

Here $u_k \in \mathbb{R}$ is the scalar control input and satisfies

$$|u_k| \leq u_{\max}. \quad (5)$$

We define the state as $z_k := (p_k, \theta_k)^\top$, and write the controlled map compactly as

$$z_{k+1} = F_K(z_k, u_k). \quad (6)$$

Periodic cycles and phase-point transfers

Let

$$\mathcal{P}^{(m)} = \{z_0^{(m)}, z_1^{(m)}, \dots, z_{m-1}^{(m)}\} \quad (7)$$

denote a period- m cycle of the uncontrolled standard map, where

$$F_K(z_i^{(m)}, 0) = z_{i+1}^{(m)}, \quad i = 0, \dots, m-2,$$

and

$$F_K(z_{m-1}^{(m)}, 0) = z_0^{(m)}.$$

The index $i \in \{0, \dots, m-1\}$ labels the phase point on the cycle. A phase-point-to-phase-point transfer is specified by

$$z_{\text{src}} = z_i^{(m)}, \quad z_{\text{goal}} = z_j^{(n)}. \quad (8)$$

Thus, a period-2 source cycle and a period-3 target cycle define $2 \times 3 = 6$ possible phase-point transfer cases.

In this paper we study two source–target transfer cases:

- (i) transport from a hyperbolic period-2 source cycle to a hyperbolic period-3 target cycle, denoted $\text{P2H} \rightarrow \text{P3H}$;
- (ii) transport from a hyperbolic period-2 source cycle to an elliptic period-3 target cycle, denoted $\text{P2H} \rightarrow \text{P3E}$.

These cases use the same source-cycle geometry but different target stability types. This comparison is important for the station-keeping analysis because the same cyclic LTV MPC formulation is used for both hyperbolic and elliptic target cycles.

Figure 1 shows the selected source and target cycles on the uncontrolled phase portrait for $K = 1.25$.




Figure 1: Selected source and target cycles for the two transfer cases at $K = 1.25$. Both cases use the same hyperbolic period-2 source cycle. The target is a hyperbolic period-3 cycle in (a) and an elliptic period-3 cycle in (b). The cycles are shown on the uncontrolled standard-map phase portrait.

Distance and target neighborhood

For $z = (p, \theta)^\top$ and $z' = (p', \theta')^\top$, we use the cylinder metric

$$d(z, z') = \left[(p - p')^2 + \text{wrap}_{[-\pi, \pi)}(\theta - \theta')^2 \right]^{1/2}, \quad (9)$$

where $\text{wrap}_{[-\pi, \pi)}$ denotes the wrapped angular difference. The target neighborhood of radius $r > 0$ about the target phase point is

$$\mathcal{G}_r(z_{\text{goal}}) = \{z : d(z, z_{\text{goal}}) \leq r\}. \quad (10)$$

Transport objective

The transport problem is to find a bounded control sequence $\{u_0, \dots, u_{N-1}\}$ such that

$$z_0 = z_{\text{src}}, \quad z_{k+1} = F_K(z_k, u_k), \quad |u_k| \leq u_{\text{max}}, \quad (11)$$

and

$$z_N \in \mathcal{G}_r(z_{\text{goal}}). \quad (12)$$

The accumulated control effort is measured by

$$J = \sum_{k=0}^{N-1} u_k^2. \quad (13)$$

In spacecraft terms, N is analogous to a discrete time-of-flight or number of map passages, while J is a surrogate for maneuvering effort. The pair (N, J) therefore represents a discrete analogue of the familiar time-of-flight versus fuel tradeoff in trajectory design. A shorter transfer may require larger control authority or accumulated effort, whereas a longer transfer may exploit the natural dynamics more efficiently.

Station-keeping objective

The station-keeping problem begins after the controlled trajectory reaches the target neighborhood. Its objective is no longer to minimize transfer time, but to keep the trajectory close to the target periodic cycle. We denote the target cycle for a given transfer case by

$$\mathcal{P}_{\text{goal}} = \{z_0^*, z_1^*, \dots, z_{n-1}^*\}. \quad (14)$$

A station-keeping performance measure is the distance to the nearest phase point on the target cycle,

$$e_k = \min_{\ell \in \{0, \dots, n-1\}} d(z_k, z_\ell^*). \quad (15)$$

Equivalently, we write $e_k = d(z_k, \mathcal{P}_{\text{goal}})$.

The transport objective is defined relative to a neighborhood of the selected target phase point z_{goal} , whereas the station-keeping objective is defined relative to the full target cycle $\mathcal{P}_{\text{goal}}$. This distinction is essential. A trajectory that reaches $\mathcal{G}_r(z_{\text{goal}})$ solves the finite-time transport problem, but it does not necessarily solve the station-keeping problem.

Stage	Role	Purpose
1	Natural transport benchmark	Seed trajectories near the source unstable manifold and record uncontrolled first hits of the target neighborhood.
2	Offline open-loop transfer generation	Solve direct constrained finite-horizon minimum-effort targeting problems to compute nominal transfers; record feasible terminal horizons, control bounds, and accumulated effort.
3	Online feedback tracking and station keeping	Track selected nominal transfers using bounded receding-horizon LTV MPC; after first entry into the target neighborhood, switch to cyclic LTV MPC station keeping on the exact target cycle.

Table 1: Computational pipeline used to generate nominal transfers and evaluate closed-loop tracking and station keeping.

COMPUTATIONAL PIPELINE

The computations follow the three-stage workflow summarized in Table 1. Stage 1 establishes natural transport benchmarks, Stage 2 generates nominal open-loop transfers, and Stage 3 evaluates closed-loop tracking and post-arrival station keeping.

This separation mirrors standard mission-design practice. A nominal transfer may be designed offline using a computationally intensive method, while online guidance and control execute the transfer with feedback correction. In this interpretation, the finite-horizon optimal-control stage generates candidate transfer arcs for a precomputed library, and the LTV MPC stage provides the online feedback layer used to correct state errors, mitigate model mismatch, and perform post-arrival station keeping.

NATURAL TRANSPORT BENCHMARK

Manifold-based seeding

Natural transport is first studied under the uncontrolled map,

$$z_{k+1} = F_K(z_k, 0).$$

For each source–target phase pair, initial conditions are placed on or near the unstable manifold associated with the selected phase point of the hyperbolic period-2 source cycle and are then iterated forward.

The source-phase label identifies the periodic-orbit phase point used to generate the corresponding unstable-manifold branch. The seeded initial condition does not coincide exactly with that phase point; rather, it is displaced along or slightly transverse to the associated unstable manifold. An initial condition placed exactly on the source cycle would remain on that cycle under the uncontrolled map and would not produce transport away from it.

A trajectory is counted as a successful natural hit if it enters the prescribed target neighborhood within $N_{\max} = 6000$ map iterates. The primary results use $r = 0.1$, while $r = 0.05$ is used as a

sensitivity check.

Phase-sweep hit statistics

For each seeded trajectory, the natural first-hit time is defined as

$$N_{\text{hit}}^{\text{nat}} = \inf \left\{ k \geq 0 : F_K^k(z_0, 0) \in \mathcal{G}_r(z_{\text{goal}}) \right\}. \quad (16)$$

If the trajectory does not enter $\mathcal{G}_r(z_{\text{goal}})$ within N_{max} iterates, it is counted as a non-hit.

Tables 2 and 3 summarize the phase-sweep results for the hyperbolic and elliptic target cycles, respectively. Representative first-hit-time distributions are shown in Fig. 2.

Table 2: Natural phase-sweep benchmark for the P2H $\rightarrow$ P3H transfer case at $K = 1.25$ and $r = 0.1$. Each phase pair uses 32,000 uncontrolled manifold-seeded trials.

Source phase	Target phase	Hit rate	Fastest hit time
0	0	64.3%	24
0	1	65.5%	22
0	2	64.8%	23
1	0	63.2%	20
1	1	64.6%	18
1	2	63.9%	19


Table 3: Natural phase-sweep benchmark for the P2H $\rightarrow$ P3E transfer case at $K = 1.25$ and $r = 0.1$. Each phase pair uses 32,000 uncontrolled manifold-seeded trials.

Source phase	Target phase	Hit rate	Fastest hit time
0	0	0.006%	–
0	1	16.3%	61
0	2	0.131%	–
1	0	0.034%	–
1	1	15.9%	66
1	2	0.159%	–


The fastest representative trajectories for the two transfer cases are shown in Fig. 3. In both panels, the marked start point is a seeded state associated with the selected source-phase manifold branch, not the periodic-orbit phase point itself. This displacement is necessary for the trajectory to depart from the source cycle and follow the unstable manifold into the surrounding phase space.

Interpretation

The natural benchmark identifies dynamically accessible transport corridors and provides baseline first-hit times against which bounded controlled transfers can be compared. For the P2H $\rightarrow$ P3H case, Table 2 shows that all six source–target phase combinations have comparable hit rates, near 65%, with fastest hits between 18 and 24 iterates. The corresponding broad hit-time distribution is illustrated in Fig. 2a.




(a) P2H $\rightarrow$ P3H.




(b) P2H $\rightarrow$ P3E.

Figure 2: Natural hit-time distributions for uncontrolled manifold-seeded trajectories at $K = 1.25$ and $r = 0.1$: 20,732 hits out of 32,000 trials in (a) and 5,224 hits out of 32,000 trials in (b).



(a) P2H $\rightarrow$ P3H.



(b) P2H $\rightarrow$ P3E.

Figure 3: Fastest natural hits for uncontrolled manifold-seeded trajectories at $K = 1.25$ and $r = 0.1$: $N_{\text{hit}}^{\text{nat}} = 23$ in (a) and $N_{\text{hit}}^{\text{nat}} = 61$ in (b).

The $P2H \rightarrow P3E$ case is substantially more phase selective. As shown in Table 3 and Fig. 2b, successful trajectories are concentrated primarily on target phase 1, whereas the other target phases are reached only rarely.

The representative trajectories in Fig. 3 demonstrate that the unstable manifold of the source cycle can provide natural transport toward both target types. These trajectories nevertheless remain uncontrolled benchmarks: they do not provide prescribed transfer times, robustness to perturbations, or post-arrival station keeping near the target cycle.

OFFLINE OPEN-LOOP TRANSFER GENERATION

Direct constrained minimum-effort transfer formulation

The natural-hit benchmarks identify dynamically plausible transport corridors, but they do not minimize actuation effort under a prescribed control bound. The offline stage computes a nominal open-loop transfer, consisting of a reference trajectory and an associated feedforward control sequence. This transfer is generated by solving a direct constrained finite-horizon minimum-effort problem under the prescribed input bound. The resulting trajectory and control sequence are not a feedback controller; they define the nominal transfer that will later be tracked by the online receding-horizon controller. The objective is to reach the prescribed target neighborhood $\mathcal{G}_r(z_{\text{goal}})$ with minimum actuation effort, not necessarily to arrive exactly at the selected target phase point z_{goal} .

For a fixed terminal horizon N , source phase point z_{src} , target phase point z_{goal} , target radius r , and control bound u_{max} , the offline optimizer selects only the feedforward control sequence

$$U_N = [u_0 \quad u_1 \quad \cdots \quad u_{N-1}]^\top.$$

The associated reference trajectory is generated by exact forward simulation:

$$z_0(U_N) = z_{\text{src}}, \tag{17a}$$

$$z_{k+1}(U_N) = F_K(z_k(U_N), u_k), \quad k = 0, \dots, N-1. \tag{17b}$$

Thus, the controlled dynamics are enforced by construction rather than by introducing the states as independent optimization variables with equality constraints.

For each fixed pair (N, u_{max}) , we solve

$$\min_{U_N} J_N(U_N) = \sum_{k=0}^{N-1} u_k^2 \tag{18a}$$

$$\text{subject to } d(z_N(U_N), z_{\text{goal}})^2 \leq r^2, \tag{18b}$$

$$|u_k| \leq u_{\text{max}}, \quad k = 0, \dots, N-1. \tag{18c}$$

The objective J_N is the fixed-horizon version of the accumulated effort J defined in Eq. (13). The terminal constraint enforces membership in the target neighborhood. Because the objective minimizes actuation effort rather than terminal distance to the selected phase point, a constrained minimum-effort solution is only required to terminate inside $\mathcal{G}_r(z_{\text{goal}})$. Consequently, the terminal state may lie near the boundary of the target neighborhood rather than near z_{goal} .

The constrained nonlinear program is solved using sequential least-squares quadratic programming (SLSQP)⁸ with box constraints on the controls and a nonlinear terminal inequality constraint.

Because the problem is nonlinear and nonconvex, the reported results should be interpreted as best observed feasible solutions from the numerical sweep, not as certified global optima.

Direct minimum-effort capture sweep for the hyperbolic-to-hyperbolic P2H $\rightarrow$ P3H case

Table 4 and Fig. 4 summarize the direct constrained minimum-effort sweep for the P2H $\rightarrow$ P3H transfer from source phase 0 to target phase 1, at $K = 1.25$ and $r = 0.1$. For each terminal horizon N , we solved the constrained minimum-effort problem in Eq. (18) for a prescribed set of candidate control bounds $u_{\max}$. The table reports, for each N , the feasible run with the smallest optimized fixed-horizon effort J_N , where J_N is the objective in Eq. (18a). Thus, the listed $u_{\max}$ is the control bound associated with the best observed feasible run at that horizon.

Table 4: Best observed direct constrained minimum-effort transfers for the P2H $\rightarrow$ P3H transfer from phase 0 to phase 1, at $K = 1.25$ and $r = 0.1$. The reported effort is the optimized fixed-horizon objective J_N in Eq. (18a).

N	$u_{\max}$	J_N	Status
8	–	–	no feasible run found
9	0.060	3.029×10^{-3}	feasible
10	–	–	no feasible run found
11	0.035	2.303×10^{-3}	feasible
12	0.050	3.584×10^{-3}	feasible
13	0.060	1.028×10^{-2}	feasible
14	0.022	2.534×10^{-4}	feasible
15	0.050	3.676×10^{-3}	feasible
16	0.060	1.624×10^{-3}	feasible
17	0.018	4.714×10^{-4}	feasible
18	0.060	3.481×10^{-3}	feasible
19	0.060	1.154×10^{-4}	feasible
20	0.018	5.111×10^{-4}	feasible
21	0.050	1.293×10^{-3}	feasible
22	0.003	0	raw natural hit




Figure 4: Direct constrained minimum-effort sweep diagnostic for the P2H $\rightarrow$ P3H transfer from phase 0 to phase 1, at $K = 1.25$ and $r = 0.1$. The figure reports the optimized fixed-horizon effort J_N , see Eq. (18a), for the best observed feasible solution at each terminal horizon.

Compared with the natural benchmark in Table 2, the controlled sweep in Table 4 and Fig. 4 shows that bounded control can produce terminal capture substantially earlier than the 22-iterate natural hit for this phase pair. The shortest feasible terminal horizon found is $N = 9$, while several longer horizons require substantially smaller effort. In particular, the $N = 14$ and $N = 19$ cases illustrate lower-effort transfers that arise at favorable phase alignments.

The nonmonotone dependence of J_N on N in Fig. 4 reflects both the underlying phase-sensitive dynamics and the local nature of the nonconvex numerical search. Different terminal horizons produce different alignments with the target phase, so a longer allowable transfer does not necessarily yield a lower-effort feasible solution for the selected phase-point target. Because the solutions are obtained from a finite multistart search, the curves should not be interpreted as certified global value functions.

Direct minimum-effort capture sweep for the hyperbolic-to-elliptic P2H $\rightarrow$ P3E case

Table 5 and Fig. 5 summarize the corresponding direct constrained minimum-effort sweep for the P2H $\rightarrow$ P3E transfer from source phase 0 to target phase 1, at $K = 1.25$ and $r = 0.1$. This case is more phase selective than the P2H $\rightarrow$ P3H case. In the natural benchmark, successful uncontrolled hits to the elliptic target are concentrated primarily on target phase 1, and the fastest natural hit for this source–target phase pair occurs at $N_{\text{hit}}^{\text{nat}} = 61$; see Table 3 and Fig. 3. The same reporting convention is used: for each terminal horizon N , the table lists the candidate control bound u_{max} associated with the lowest-effort feasible run found at that horizon.

Table 5: Best observed direct constrained minimum-effort transfers for the P2H $\rightarrow$ P3E transfer from phase 0 to phase 1, at $K = 1.25$ and $r = 0.1$. The reported effort is the optimized fixed-horizon objective J_N in Eq. (18a).

N	u_{max}	J_N	Status
8	–	–	no feasible run found
9	–	–	no feasible run found
10	0.060	7.196×10^{-3}	feasible
11	–	–	no feasible run found
12	0.007	1.535×10^{-5}	feasible
13	0.060	9.975×10^{-3}	feasible
14	0.060	2.620×10^{-3}	feasible
15	0.022	5.361×10^{-4}	feasible
16	0.060	6.347×10^{-3}	feasible
17	0.025	1.238×10^{-5}	feasible
18	0.022	1.098×10^{-3}	feasible
19	0.050	8.116×10^{-4}	feasible
20	0.025	1.454×10^{-6}	feasible
21	0.035	1.294×10^{-3}	feasible
22	0.060	2.449×10^{-5}	feasible

The controlled sweep produces feasible terminal capture much earlier than the fastest natural hit. The shortest feasible terminal horizon found is $N = 10$. However, Table 5 and Fig. 5 show that several longer-horizon transfers require much smaller effort. The $N = 12$ and $N = 20$ cases are representative examples of lower-effort transfers that arise at more favorable phase alignments.




Figure 5: Direct constrained minimum-effort sweep diagnostic for the P2H $\rightarrow$ P3E transfer from phase 0 to phase 1, at $K = 1.25$ and $r = 0.1$. The figure reports the optimized fixed-horizon effort J_N , Eq. (18a), for the best observed feasible solution at each terminal horizon.

Thus, the P2H $\rightarrow$ P3E case illustrates the same time–effort tradeoff as the hyperbolic target case, but against a much slower natural-transport baseline. Short controlled captures are possible, but longer horizons can exploit the natural transport geometry more efficiently and require substantially less actuation.

Warm starts and computational cost

The direct minimum-effort sweep uses backward continuation in horizon. For each fixed $u_{\max}$ branch, the solver begins at the longest horizon and then uses the solution at horizon $N + 1$, when available, as a warm start for horizon N . Additional warm starts are drawn from earlier terminal-penalty targeting runs, interpolated reference trajectories, the zero-control sequence, and random bounded control sequences. This strategy improves the probability of finding feasible solutions, but it does not remove the nonconvexity of the problem. All reported transfer solutions should therefore be interpreted as computed feasible transfers from the stated numerical search procedure.

ONLINE FEEDBACK TRACKING AND STATION KEEPING

The offline stage produces nominal open-loop transfers: a reference trajectory and an associated feedforward control sequence. In an astrodynamics setting, this is analogous to a precomputed transfer arc or maneuver plan. The online stage plays the role of a feedback guidance layer: it uses the current state to correct deviations from the nominal transfer and, after arrival, to maintain the trajectory near the target cycle.

We use a two-mode bounded LTV MPC controller. The first mode tracks the finite nominal transfer. The second mode, activated after first entry into the selected target-phase neighborhood, performs cyclic station keeping on the exact target cycle. Because the controller is recomputed at each map iterate using the current state, the resulting receding-horizon law is a feedback controller.

LTV MPC tracking of nominal transfers

Let the offline transfer-generation stage provide the nominal transfer

$$z_0^{\text{ref}}, z_1^{\text{ref}}, \dots, z_{N_{\text{ref}}}^{\text{ref}}, \quad u_0^{\text{ref}}, u_1^{\text{ref}}, \dots, u_{N_{\text{ref}}-1}^{\text{ref}},$$

where z_k^{ref} is the reference trajectory and u_k^{ref} is the feedforward control sequence. At closed-loop time k , the controller constructs a local prediction model by linearizing the controlled standard map in Eq. (6) about a nominal rollout, and then solves a finite-horizon tracking problem over a horizon H . Let

$$\hat{z}_{j|k} = (\hat{p}_{j|k}, \hat{\theta}_{j|k})^\top, \quad \hat{u}_{j|k},$$

denote the nominal state and input about which the map in Eq. (6) is linearized at prediction step j .

The resulting local affine prediction model is

$$z_{j+1|k} = A_{j|k} z_{j|k} + B u_{j|k} + c_{j|k}, \quad j = 0, \dots, H-1, \quad (19)$$

with

$$A_{j|k} = \begin{bmatrix} 1 & K \cos \hat{\theta}_{j|k} \\ 1 & 1 + K \cos \hat{\theta}_{j|k} \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \quad (20)$$

The affine term $c_{j|k}$ is chosen so that the local model reproduces the nonlinear map at $(\hat{z}_{j|k}, \hat{u}_{j|k})$. After stacking the predicted states and eliminating the state recursion, the prediction can be written in condensed form as

$$Z_k = T_k U_k + b_k, \quad (21)$$

where Z_k stacks the predicted states over the horizon, U_k stacks the predicted control inputs, and T_k and b_k are determined by the local affine model and the current measured state.

During transfer tracking, the reference stacks Z_k^{ref} and U_k^{ref} are obtained from the nominal open-loop transfer. The bounded tracking problem is

$$\min_{U_k} \left\| W_k \left(T_k U_k + b_k - Z_k^{\text{ref}} \right) \right\|_2^2 + r_u \|U_k\|_2^2 + w_{\Delta u} \|DU_k - d_{\Delta u, k}\|_2^2 + w_{\text{ff}} \|U_k - U_k^{\text{ref}}\|_2^2 \quad (22a)$$

$$\text{subject to } |u_{j|k}| \leq u_{\text{max}}, \quad j = 0, \dots, H-1. \quad (22b)$$

The first term penalizes deviation of the predicted trajectory from the nominal transfer reference, with W_k defining the state-tracking weights. The second term, weighted by r_u , penalizes control effort. The third term penalizes input variation, where D is the finite-difference matrix and $d_{\Delta u, k}$ accounts for the previously applied input at the beginning of the horizon. The final term, weighted by w_{ff} , keeps the optimized input sequence close to the nominal feedforward sequence computed offline. Together, these terms trade off reference tracking, maneuver size, input smoothness, and adherence to the offline maneuver plan.

Let

$$U_k^* = \begin{bmatrix} u_{0|k}^* & u_{1|k}^* & \cdots & u_{H-1|k}^* \end{bmatrix}^\top$$

denote the optimal solution of Eq. (22). Only the first optimized input $u_{0|k}^*$ is applied to the nonlinear controlled map; the problem is then rebuilt and solved again at the next map iterate. This receding-horizon implementation provides the feedback correction that distinguishes closed-loop execution from open-loop application of the nominal feedforward sequence.

Cyclic LTV MPC station keeping

After first entry into the selected target-phase neighborhood, the controller switches from finite-transfer tracking to station keeping. The station-keeping reference is the exact target cycle

$$\mathcal{P}_{\text{goal}} = \{z_0^*, \dots, z_{n-1}^*\}.$$

At each step, the controller advances a cyclic reference around this period- n cycle. This is analogous to maintaining a spacecraft near a selected periodic or resonant structure after arrival, rather than continuing to track a finite transfer arc.

The same local affine prediction model is used, but the finite-transfer feedforward term is omitted. Let $Z_{\text{cyc},k}^*$ denote the cyclic target-cycle reference over the prediction horizon. The station-keeping problem is

$$\min_{U_k} \|W_k (T_k U_k + b_k - Z_{\text{cyc},k}^*)\|_2^2 + r_u \|U_k\|_2^2 + w_{\Delta u} \|DU_k - d_{\Delta u,k}\|_2^2 \quad (23a)$$

$$\text{subject to } |u_{j|k}| \leq u_{\text{max}}, \quad j = 0, \dots, H-1. \quad (23b)$$

As in the transfer-tracking mode, only the first optimized input $u_{0|k}^*$ is applied, and the problem is resolved at the next closed-loop step.

Mode switching and station-keeping metrics

The switch from transfer tracking to station keeping is triggered by first entry into the selected target-phase neighborhood:

$$d(z_k, z_{\text{goal}}) \leq r. \quad (24)$$

Before this event, the controller tracks the finite nominal transfer. After this event, it tracks the exact target cycle using cyclic LTV MPC station keeping.

Station-keeping performance is measured by the distance to the exact target cycle,


$$d(z_k, \mathcal{P}_{\text{goal}}) = \min_{z_i^* \in \mathcal{P}_{\text{goal}}} d(z_k, z_i^*). \quad (25)$$

The reported diagnostics are the first-entry index, the post-entry distance to the exact cycle, late-window distance, final cycle distance, and applied-input history.


Online tracking and station keeping for the P2H $\rightarrow$ P3H case

We first evaluate closed-loop tracking and station keeping for the P2H $\rightarrow$ P3H transfer from source phase 0 to target phase 1. The nominal reference is the $N = 11$ direct constrained minimum-effort transfer reported in Table 4. This reference is tracked using the bounded LTV MPC transport controller with maximum prediction horizon $H = 8$ and input bound $|u_k| \leq 0.06$. Because the transfer reference is finite, the effective transport horizon is shortened near the terminal point according to $H_k = \min(H, N - k)$. After first entry into the selected target-phase neighborhood, the controller discards the finite transfer reference and switches to cyclic exact-cycle station keeping.

Figure 6 shows the applied-input and distance-to-cycle histories for this run. The closed-loop trajectory enters the target neighborhood at $k = 11$, matching the terminal horizon of the nominal transfer. At the switching time, the arrival state is already close to the selected exact-cycle phase, with distance 4.48×10^{-2} . After switching, the distance-to-cycle history in Fig. 6b shows a short transient followed by convergence to the exact target cycle to numerical precision.



(a) Applied control input u_k .



(b) Distance to the target cycle $d(z_k, \mathcal{P}_{\text{goal}})$.

Figure 6: Closed-loop P2H $\rightarrow$ P3H diagnostics for the $N = 11$, $r = 0.1$ reference. The switch from transfer tracking to cyclic station keeping occurs at $k = 11$.

Online tracking and station keeping for the P2H $\rightarrow$ P3E case

We next evaluate closed-loop tracking and station keeping for the P2H $\rightarrow$ P3E transfer. The target is the elliptic period-3 cycle used in the direct constrained sweep of Table 5. Because the target cycle is elliptic, post-entry maintenance is performed using cyclic LTV MPC station keeping.

The comparison uses two $N = 12$ nominal open-loop transfer references computed with different terminal capture radii. The first reference uses $r = 0.1$ and corresponds to the low-effort $N = 12$ solution in Table 5. The second reference tightens the terminal capture radius to $r = 0.05$. This tighter terminal condition increases the nominal open-loop transfer effort, but it places the arrival state closer to the exact target cycle and therefore improves the subsequent handoff to cyclic station keeping.

The results show a clear tradeoff between nominal open-loop transfer effort and closed-loop handoff quality. For the $r = 0.1$ nominal transfer, the required feedforward effort is very small, but the trajectory enters station keeping farther from the exact target cycle. For the tightened $r = 0.05$ nominal transfer, the feedforward effort increases, but the maximum post-entry distance to the exact target cycle is reduced and the trajectory converges to the exact target cycle to numerical precision.

Figures 7 and 8 show the corresponding applied-input and distance-to-cycle histories. In both cases, the controller first tracks the finite nominal transfer reference and then switches to cyclic station keeping after first entry into the selected target-phase neighborhood. The tightened terminal-radius case produces a smoother handoff in the distance-to-cycle diagnostic.

Feedback correction under an initial-condition perturbation

The preceding closed-loop results show that the LTV MPC controller can track nominal transfer references and then maintain the trajectory near the exact target cycle. To distinguish feedback correction from open-loop application of the nominal feedforward control sequence, we also test sensitivity to a small initial-condition perturbation.


The test uses the P2H $\rightarrow$ P3H transfer with the same $N = 11$, $r = 0.1$ direct constrained minimum-effort reference used in Fig. 6. The initial condition is perturbed in the momentum coordinate by

$$\delta p = 10^{-3}, \quad \delta \theta = 0.$$


Two executions are compared: open-loop application of the nominal feedforward control sequence and closed-loop LTV MPC tracking from the same perturbed initial condition.

Under the perturbation, the open-loop feedforward sequence no longer reaches the $r = 0.1$ target neighborhood and subsequently remains separated from the exact target cycle. In contrast, the closed-loop LTV MPC controller corrects the perturbed state, enters the target neighborhood at the nominal time $k = 11$, and captures the exact target cycle to numerical precision.

Figure 9 compares the corresponding exact-cycle distance histories. The result demonstrates the role of feedback in the pipeline: the offline computation supplies a feasible nominal open-loop transfer, while the online feedback controller corrects state error and completes the transport-to-station-keeping transition.




(a) Applied control input u_k .




(b) Distance to the target cycle $d(z_k, \mathcal{P}_{\text{goal}})$.

Figure 7: Closed-loop P2H $\rightarrow$ P3E diagnostics for the $N = 12$, $r = 0.1$ reference. The switch from transfer tracking to cyclic station keeping occurs at $k = 12$.




(a) Applied control input u_k .




(b) Distance to the target cycle $d(z_k, \mathcal{P}_{\text{goal}})$.

Figure 8: Closed-loop P2H $\rightarrow$ P3E diagnostics for the $N = 12$, $r = 0.05$ reference. The switch from transfer tracking to cyclic station keeping occurs at $k = 12$.



(a) Open-loop application of the nominal feedforward sequence.



(b) Closed-loop LTV MPC tracking and cyclic station keeping.

Figure 9: Effect of the initial-condition perturbation $\delta p = 10^{-3}$, $\delta \theta = 0$ on the P2H $\rightarrow$ P3H transfer with $N = 11$ and $r = 0.1$. The open-loop feedforward sequence misses the target neighborhood, whereas closed-loop LTV MPC corrects the perturbation, reaches the target neighborhood at $k = 11$, and converges to the target cycle.

EXTENSION TOWARD KEPLERIAN MAPS

The standard map is used here as a low-dimensional testbed for constructing controlled transport graphs between dynamically meaningful structures. In this viewpoint, selected phase points on periodic cycles define graph nodes, and feasible bounded-control transfers define directed edges labeled by first-entry time, accumulated effort, required control bound, terminal accuracy, and post-arrival station-keeping quality.

This viewpoint has a natural astrodynamics analogue. For Keplerian maps, CR3BP Poincaré maps, or higher-fidelity trajectory models, the nodes may represent resonant states, periodic orbits, or selected points on invariant structures, while the edges represent maneuver-assisted transfers between them. The same offline–online structure then applies: compute nominal open-loop transfer edges, track selected edges with online feedback, and perform local station keeping near the target structure.

A natural next step is the Keplerian map, which provides a reduced-order setting for resonance-to-resonance transport in perturbed two-body motion.⁷ Unlike the standard map, the variables in a Keplerian map have direct orbital interpretation, and bounded control can be related more directly to small velocity or energy changes applied at discrete orbital passages. Future work will apply the graph-building pipeline developed here to Keplerian maps and then to higher-fidelity astrodynamics models.

CONCLUSIONS

This paper presented a computational pipeline for bounded controlled transport between periodic-cycle phase points in the standard map. The formulation is motivated by the construction of directed transport graphs, in which periodic-cycle phase points are nodes and feasible bounded-control transfers are edges labeled by transfer time, required control authority, accumulated effort, and station-keeping quality.

The approach separates finite-time transport from post-arrival station keeping. Natural manifold-seeded trajectories provide baseline transport statistics and reveal phase-dependent access to the

target cycles. The main transfer-generation method is direct constrained finite-horizon minimum-effort targeting, which produces nominal open-loop transfers for subsequent feedback tracking.

For the P2H $\rightarrow$ P3H case, the direct constrained sweep produces feasible terminal capture at horizons as short as $N = 9$, compared with a 22-iterate natural benchmark for the same source–target phase pair. For the P2H $\rightarrow$ P3E case, feasible terminal capture is obtained at horizons as short as $N = 10$, compared with a 61-iterate natural benchmark. The online results show that bounded LTV MPC can track selected nominal transfer references and switch after first entry to cyclic station keeping on the exact target cycle. A perturbation test further shows that feedback is essential: open-loop application of the nominal feedforward sequence can miss the target neighborhood, whereas the closed-loop feedback controller corrects the state error and completes the transport-to-station-keeping transition.

From the astrodynamics viewpoint, the principal result is not the specific standard-map transfer, but the demonstrated workflow: identify natural transport opportunities, quantify the uncontrolled benchmark, compute bounded-control accelerations of that transport, and then close the loop with feedback tracking and station keeping. This provides a prototype for constructing maneuver libraries or transport graphs in more realistic orbital maps, where each edge is labeled by transfer time, control effort, terminal accuracy, and post-arrival maintenance cost.

Future work will extend this graph-building pipeline to Keplerian maps and higher-fidelity astrodynamics models, where analogous questions arise in low-energy transfer design and local maintenance near resonant or periodic structures.

ACKNOWLEDGMENTS

D. M. Luchtenburg acknowledges travel support from the Dean’s Office at The Cooper Union, and S. D. Ross acknowledges support from the Air Force Office of Scientific Research under Grant No. FA9550-24-1-0194.

REFERENCES

- [1] J. Meiss, “Thirty years of turnstiles and transport,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, Vol. 25, No. 9, 2015.
- [2] B. V. Chirikov, “A universal instability of many-dimensional oscillator systems,” *Physics reports*, Vol. 52, No. 5, 1979, pp. 263–379.
- [3] N. Hiraiwa, M. Bando, I. Nisoli, and Y. Sato, “Designing robust trajectories by lobe dynamics in low-dimensional Hamiltonian systems,” *Physical Review Research*, Vol. 6, No. 2, 2024, p. L022046.
- [4] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*. SIAM, 2010.
- [5] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.
- [6] E. Ott, C. Grebogi, and J. A. Yorke, “Controlling chaos,” *Physical Review Letters*, Vol. 64, No. 11, 1990, pp. 1196–1199, 10.1103/PhysRevLett.64.1196.

- [7] S. D. Ross and D. J. Scheeres, “Multiple Gravity Assists, Capture, and Escape in the Restricted Three-Body Problem,” *SIAM Journal on Applied Dynamical Systems*, Vol. 6, No. 3, 2007, pp. 576–596, 10.1137/060663374.
- [8] D. Kraft, “A software package for sequential quadratic programming,” Tech. Rep. DFVLR-FB 88-28, DLR German Aerospace Center, Oberpfaffenhofen, Germany, 1988.