

Dynamics-Aware Control of Electromagnetic Levitation for Hyperloop Pods Under Motion, Track Irregularity, and Transient Loading

Anubhav Nigam

Cornell University, Ithaca, NY
adn38@cornell.edu

June 15, 2026

Abstract

Electromagnetic levitation concepts for high-speed Hyperloop pods are often introduced through a static force balance between pod weight and magnetic lift. That framing is useful for preliminary actuator sizing, but it is not sufficient for assessing whether a moving pod can maintain clearance over a non-ideal guideway. Once the vehicle moves, accelerates, brakes, and passes over track irregularities, the levitation gap becomes a dynamic state coupled to pod heave, pitch, actuator current, guideway geometry, eddy-current effects, and longitudinal load transfer. This paper develops a reduced-order, dynamics-aware levitation model for a representative Hyperloop pod with front and rear electromagnetic levitation modules. The model combines nonlinear electromagnetic force, finite current dynamics, vertical heave-pitch motion, speed-dependent guideway excitation, eddy-current damping and drag surrogates, and acceleration-induced redistribution of lift demand. Linearization around operating points is used to identify when classical linear gap control is defensible, while a constraint-aware nonlinear controller is formulated for cases in which gap, current, voltage, and load-transfer constraints become active. The central result is conceptual rather than architecture-specific: static levitation can establish a nominal operating point, but closed-loop dynamic simulation determines the safe speed-disturbance envelope. Hyperloop levitation should therefore be treated as a coupled vehicle-control-guideway problem rather than as a fixed magnetic force-balance calculation.

Keywords: Hyperloop, magnetic levitation, electromagnetic suspension, nonlinear control, track irregularity, vehicle dynamics, eddy-current damping

1 Introduction

Hyperloop vehicles are commonly motivated by the possibility of high-speed, low-contact ground transportation in a controlled tube environment. In that setting, magnetic levitation is attractive because it can remove rolling contact, reduce mechanical wear, and decouple the vehicle from many of the frictional limitations of wheel-rail systems [2]. However, the removal of mechanical contact also removes a passive source of constraint. The levitation system must continuously maintain the clearance between the pod and guideway while the pod moves through a spatially imperfect structure and experiences changing longitudinal operating conditions.

A common introductory levitation model treats this problem as a static balance,

$$mg = F(I, g), \quad (1)$$

where m is the supported mass, g is the nominal air gap, I is actuator current, and F is magnetic lift. This equation is useful because it gives a first estimate of the current and magnetic geometry required to support the vehicle. It also makes clear that levitation is not free: lift authority depends on current, gap, magnetic geometry, and available power. But as a design argument for a moving Hyperloop pod, the static balance is incomplete. It assumes that the gap is fixed, that lift demand is evenly distributed, that the guideway is geometrically ideal, and that magnetic force can be changed instantaneously.

Those assumptions fail during operation. The pod does not hover above a fixed reference surface, but rather travels over a guideway whose height varies with position. A spatial guideway defect becomes a time-domain disturbance whose frequency content depends on vehicle speed. As the pod body heaves and pitches, a disturbance at one levitation station changes local clearance as well as front-rear load distribution. Longitudinal acceleration and braking produce also load transfer through the height of the center of mass, changing the nominal force required at the front and rear levitation modules. Conductive guideway or reaction elements can introduce eddy-current damping, drag, and heating [3, 5]. Additionally, the current in an electromagnetic actuator has finite electrical dynamics, so the commanded force cannot be realized without bandwidth, voltage, and thermal limits.

The relevant vertical dynamics are therefore closer to

$$m\ddot{z} = F(I, g, v, \dot{g}) - mg + d(t), \quad (2)$$

where z is the vertical pod coordinate, v is longitudinal speed, $\dot{g}$ is gap rate, and $d(t)$ represents disturbances and unmodeled coupling. In a multi-module pod, even this equation is only a local simplification because each levitation module sees a different guideway height and a different pitch-induced displacement.

This creates a design challenge for magnetic levitation. Levitation systems must be designed with sufficient current and voltage margin to reject track irregularities, as well as damping and front-rear coordination to remain stable against oscillations and braking forces. Static levitation design is not adequate to deal with these issues.

We address this by developing a reduced-order, dynamics-aware levitation framework for a representative Hyperloop pod. The goal is to isolate the smallest set of coupled dynamics needed to show why static force balance is not an adequate operational model. The contributions of this paper are:

1. a nonlinear vertical levitation model with front and rear electromagnetic modules, finite current dynamics, heave, and pitch;
2. a speed-dependent representation of track irregularity as base excitation entering the levitation gaps;
3. a reduced-order eddy-current damping and drag surrogate that captures motion-dependent electromagnetic effects;
4. an acceleration and braking load-transfer model showing how longitudinal motion shifts front-rear lift demand;

5. an operating-point linearization used to define when linear gap control is valid and when nonlinear or constrained control becomes necessary;
6. a controller comparison showing how local PID, coordinated linear body feedback, and constrained nonlinear control occupy different regions of the same levitation operating envelope.

The paper is organized as follows. Section II motivates the dynamics-aware formulation and defines the operational questions that a levitation model must answer. Section III reviews the relevant background in electromagnetic suspension, guideway excitation, and motion-dependent electromagnetic effects. Section IV develops the reduced-order pod and levitation-module model. Section V derives the operating-point linearization and discusses validity limits. Section VI presents candidate control formulations. Section VII defines the operating envelope and design constraints that should be evaluated before simulation. Section VIII defines simulation scenarios for transient loading and guideway irregularity. Section IX discusses design implications for Hyperloop levitation systems.

2 Motivation

The goal of this study is to determine how motion, guideway irregularity, and transient loading change the levitation problem relative to a nominal static design. This is a deliberately narrower question than full Hyperloop vehicle design. We are not attempting to specify a final actuator geometry, guideway material stack, power electronics architecture, or certification envelope. Rather, we ask which dynamic couplings must appear in a levitation model before the model can support useful claims about stability, clearance, current margin, and disturbance rejection.

This distinction matters because static magnetic sizing can look deceptively complete. If a levitation module can generate a force greater than its share of vehicle weight at a chosen nominal gap, it is tempting to conclude that the module is sufficient. But the vehicle never operates only at that nominal point. The pod mass may vary with payload. The local guideway height changes with position. The pod pitches during disturbances. The force available from an electromagnetic actuator changes nonlinearly with gap. The current command is filtered by coil inductance and voltage limits. Longitudinal acceleration changes the distribution of required lift between modules. These effects are not small implementation details; they define the actual operating problem.

A dynamics-aware levitation assessment should therefore answer four questions that the static model cannot answer.

First, it should determine whether the nominal levitation equilibrium is stable under closed-loop control. Electromagnetic attraction is strongly gap-dependent, and the sign and magnitude of the gap stiffness matter for stability. A controller that is adequate at one gap and payload may not retain the same margin at another operating point.

Second, it should determine whether guideway irregularities can be rejected at operating speed. Track geometry is usually specified spatially, but the controller experiences it temporally through $r(x(t))$. A sinusoidal guideway error with wavelength λ appears at temporal frequency v/λ . The same physical defect is therefore a low-frequency disturbance during slow travel and a high-frequency disturbance during fast travel.

Third, it should determine whether transient longitudinal maneuvers change the vertical problem. Launch and braking are often treated as propulsion events, but for levitation they are also load-transfer events. If the center of mass is above the levitation plane, longitudinal acceleration creates a pitch moment and shifts the required support force between front and rear modules. The resulting operating-point shift changes the local plant seen by each controller.

Fourth, it should determine whether the levitation system retains actuator margin under combined disturbances. The most important case is not necessarily a single isolated input. A pod may brake while crossing a guideway dip, or accelerate while carrying a heavier payload, or pass over periodic support-span variation at a speed that excites a lightly damped heave-pitch mode. The design metric is not simply nominal lift, but the minimum clearance and maximum actuator effort over these combined events.

For these reasons, this paper treats levitation as a state-aware control problem. The state of interest includes not only vertical position and current, but also guideway-relative gap, pitch, speed, acceleration, and disturbance history. This framing is useful even when a final system uses a more detailed controller, because it makes the simplifying assumptions visible. Static design, linear control, gain-scheduled control, and nonlinear constrained control are not competing narratives; they are approximations valid over different parts of the operating envelope.

3 Background

3.1 Electromagnetic Levitation as a Control Problem

Electromagnetic suspension systems use controlled magnetic force to maintain a small air gap between a vehicle-mounted actuator and a ferromagnetic or otherwise magnetically coupled guideway [1]. The central advantage is that lift can be produced without mechanical contact. The central control challenge is that magnetic force is a nonlinear function of current and gap. In a simplified attractive actuator, lift can be written as

$$F(I, g) = k \frac{I^2}{g^2}, \quad (3)$$

where k absorbs magnetic geometry and permeability effects. This form is not intended to replace finite-element design, but it captures the important local feature: force rises rapidly as the gap decreases.

This gap dependence makes the static equilibrium fragile in open loop. If the gap decreases slightly while current is held fixed, the attractive force increases, which tends to decrease the gap further. If the gap increases, force falls, and the supported body tends to drop away from the equilibrium. Active feedback is therefore not an accessory to electromagnetic suspension; it is part of the suspension itself.

Most practical gap controllers use some combination of feedforward current, gap feedback, current feedback, and damping injection through gap-rate or velocity estimates. Linear controllers can be effective when the vehicle remains close to a chosen operating point. Recent maglev-control work has also emphasized model-predictive and nonlinear approaches when ride response, actuator dynamics, and guideway excitation are included in the control problem [6, 7]. However, the derivatives of the force law change with both current and gap:

$$\frac{\partial F}{\partial I} = 2k \frac{I_0}{g_0^2}, \quad (4)$$

$$\frac{\partial F}{\partial g} = -2k \frac{I_0^2}{g_0^3}. \quad (5)$$

The plant seen by the controller therefore changes as the nominal gap, payload, or lift demand changes. A controller designed for one operating point implicitly assumes a limited range of gap excursions and actuator effort.

3.2 Guideway Irregularity

The guideway is not only a structure; it is also an input to the levitation controller. A vertical height profile $r(x)$ enters the local gap measurement directly. If the pod travels with longitudinal position $x(t)$, then the same spatial profile becomes the time-domain disturbance $r(x(t))$. For a periodic guideway component

$$r(x) = A_r \sin\left(\frac{2\pi x}{\lambda}\right), \quad (6)$$

the disturbance experienced by a pod traveling at constant speed v is

$$r(t) = A_r \sin\left(\frac{2\pi vt}{\lambda}\right). \quad (7)$$

Thus the excitation frequency is

$$f = \frac{v}{\lambda}. \quad (8)$$

This speed scaling is one of the simplest reasons static clearance analysis is insufficient. A long-wavelength settlement error may appear as a slow bias at low speed, while a support-span periodicity or construction waviness may enter the bandwidth of the levitation controller at high speed. The relevant guideway tolerance is therefore not only an allowable vertical deviation in millimeters, but a spatial spectrum interpreted through the operating speed envelope.

3.3 Transient Loading and Load Transfer

Longitudinal acceleration also enters the vertical levitation problem. If the vehicle center of mass is located a height h above the effective levitation plane, a longitudinal acceleration a_x produces a pitch moment. In a two-module representation with front and rear stations separated by wheelbase-like distance L , the quasi-static change in front-rear support demand can be approximated by

$$\Delta F = \frac{ma_x h}{L}. \quad (9)$$

The sign depends on the convention for positive acceleration and the location of the coordinate origin, but the physical consequence is invariant: acceleration and braking shift support demand between the front and rear levitation modules.

This matters because lift redistribution changes both equilibrium current and local stiffness. During a braking event, the front module may require more lift at the same time that pitch motion changes its air gap. During launch, the opposite module may become more heavily loaded depending on the geometry and sign convention. A levitation controller that treats each module as an independent single-input single-output gap regulator can miss this body-level coupling.

3.4 Eddy-Current Effects

Motion through magnetic fields can induce currents in conductive guideway or reaction elements. These currents create additional electromagnetic forces that may appear as vertical damping, lift modification, longitudinal drag, and heat generation [3, 4]. The details depend on guideway conductivity, magnetic geometry, speed, frequency content, and material thickness. A full treatment requires electromagnetic field simulation. For control-oriented vehicle dynamics, however, it is useful to retain a reduced-order representation of the effect:

$$F_{e,z} = -c_e(v, g)\dot{g} + F_{e,lift}(v, g), \quad (10)$$

Table 1: Principal variables used in the reduced-order levitation model.

Symbol	Meaning
$z, \dot{z}$	pod heave and heave velocity
$\theta, \dot{\theta}$	pod pitch and pitch rate
g_f, g_r	front and rear levitation gaps
I_f, I_r	front and rear actuator currents
u_f, u_r	front and rear actuator commands
$r(x)$	guideway vertical profile
v, a_x	longitudinal speed and acceleration
m, J_y	pod mass and pitch inertia
l_f, l_r	module offsets from center of mass
h	center-of-mass height above levitation plane

$$D_{e,x} = D_e(v, g, I). \quad (11)$$

This form separates the vertical damping-like contribution from the longitudinal drag contribution. It also makes the main tradeoff visible. Eddy-current coupling can help suppress gap-rate motion, but it can also increase energy loss and thermal load. For Hyperloop pods operating in low-pressure environments, where aerodynamic damping may be reduced, this electromagnetic damping and drag tradeoff becomes especially relevant.

4 Reduced-Order Levitation Model

4.1 Pod Geometry and States

We model the pod as a rigid body with vertical heave z and pitch θ . The body is supported by two equivalent levitation stations, one at the front and one at the rear. Each station represents the aggregate action of several physical levitation actuators. The front and rear stations are located at longitudinal distances l_f and l_r from the center of mass. The longitudinal position of the center of mass is $x(t)$, with speed $v(t) = \dot{x}(t)$ and acceleration $a_x(t) = \dot{v}(t)$.

Let $r(x)$ denote the vertical guideway profile in the inertial frame. Under a small-pitch approximation, the front and rear air gaps are

$$g_f = r(x + l_f) - z - l_f\theta, \quad (12)$$

$$g_r = r(x - l_r) - z + l_r\theta. \quad (13)$$

These equations are the point at which guideway geometry, body motion, and module location enter the same variable. A local controller may measure only g_f or g_r , but the physical gap contains both rigid-body motion and track input.

The state vector used for the reduced-order model is

$$\mathbf{x} = [z \quad \dot{z} \quad \theta \quad \dot{\theta} \quad I_f \quad I_r]^T. \quad (14)$$

The commanded actuator voltages or normalized drive commands are collected in

$$\mathbf{u} = [u_f \quad u_r]^T. \quad (15)$$

The disturbance input includes guideway height and slope terms at the front and rear modules, along with longitudinal acceleration:

$$\mathbf{w} = [r_f \quad r_r \quad \dot{r}_f \quad \dot{r}_r \quad a_x]^T. \quad (16)$$

4.2 Electromagnetic Force Model

Each levitation module produces a vertical force

$$F_i = F_{m,i}(I_i, g_i) + F_{e,i}(v, g_i, \dot{g}_i), \quad (17)$$

where $i \in \{f, r\}$. The magnetostatic part is represented by

$$F_{m,i}(I_i, g_i) = k_i \frac{I_i^2}{g_i^2}. \quad (18)$$

This expression is a control-oriented surrogate. In a design study with a specific actuator and guideway, $F_{m,i}$ should be replaced by a fitted force map from finite-element electromagnetic analysis or experimental calibration. The reduced-order form is retained here because it exposes the nonlinear dependence that drives the control problem.

The motion-dependent eddy-current contribution is represented as

$$F_{e,i}(v, g_i, \dot{g}_i) = -c_{e,i}(v, g_i)\dot{g}_i + F_{e,lift,i}(v, g_i). \quad (19)$$

For simulation studies, a bounded damping surrogate can be used,

$$c_{e,i}(v, g_i) = c_{0,i} \frac{v}{g_i^{n_i} (1 + (v/v_{c,i})^2)}, \quad (20)$$

where $c_{0,i}$ sets the damping scale, n_i captures the gap dependence, and $v_{c,i}$ represents a characteristic speed beyond which the simple low-speed proportionality no longer holds. This is not a universal eddy-current law. It is a compact surrogate used to preserve the fact that electromagnetic damping is speed and gap dependent.

The associated longitudinal drag is tracked separately as

$$D_{e,i} = \alpha_i \frac{v I_i^2}{g_i^2 + \epsilon_i}, \quad (21)$$

where ϵ_i prevents singular behavior in the surrogate. This term is not fed directly into the vertical equations, but it is important for energy accounting and for understanding the propulsion cost of levitation-induced electromagnetic coupling.

4.3 Current Dynamics

The current in each actuator is not assumed to follow the command instantaneously. A physically motivated coil equation is

$$L_i(g_i)\dot{I}_i + R_i I_i + I_i \frac{\partial L_i}{\partial g_i} \dot{g}_i = u_i, \quad (22)$$

where $L_i(g_i)$ is the gap-dependent inductance, R_i is coil resistance, and u_i is applied voltage. The third term represents electromechanical coupling through changing inductance. It becomes important when gap rate is large or when the actuator operates in a region with strong inductance variation.

For controller comparison, this can be reduced to a first-order current response,

$$\tau_i \dot{I}_i = -I_i + K_{u,i} u_i, \quad (23)$$

with current and voltage saturation applied separately. The simplified form is less physically complete, but it preserves the main limitation needed for a control study: force cannot be changed with infinite bandwidth.

4.4 Heave and Pitch Dynamics

The vertical rigid-body equations are

$$m\ddot{z} = F_f + F_r - mg + d_z(t), \quad (24)$$

$$J_y \ddot{\theta} = l_f F_f - l_r F_r + M_d(t), \quad (25)$$

where J_y is the pitch moment of inertia, $d_z(t)$ is an aggregate vertical disturbance, and $M_d(t)$ is an aggregate pitch disturbance. These equations show why the two levitation modules cannot be treated as completely independent actuators. A force change at either station modifies both heave and pitch, and pitch in turn changes both gaps.

The nominal lift required at the front and rear stations depends on static weight distribution and longitudinal acceleration. With $L = l_f + l_r$, the quasi-static front and rear force demands can be written as

$$F_{f,0} = \frac{mgl_r}{L} + \frac{ma_x h}{L}, \quad (26)$$

$$F_{r,0} = \frac{mgl_f}{L} - \frac{ma_x h}{L}, \quad (27)$$

for the sign convention in which positive a_x shifts support demand toward the front module. If the opposite convention is used, the signs of the acceleration terms reverse. The important feature is the dependence on $a_x h/L$: larger acceleration, higher center of mass, and shorter module spacing all increase the load-transfer effect.

5 Operating-Point Linearization

The nonlinear model can be written compactly as

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}, \mathbf{w}; v). \quad (28)$$

For a chosen operating point $(\mathbf{x}_0, \mathbf{u}_0, \mathbf{w}_0, v_0)$, small perturbations satisfy

$$\delta \dot{\mathbf{x}} = A \delta \mathbf{x} + B \delta \mathbf{u} + E \delta \mathbf{w}, \quad (29)$$

where

$$A = \left. \frac{\partial f}{\partial \mathbf{x}} \right|_0, \quad B = \left. \frac{\partial f}{\partial \mathbf{u}} \right|_0, \quad E = \left. \frac{\partial f}{\partial \mathbf{w}} \right|_0. \quad (30)$$

This linearization is useful because it gives a local plant for classical control design and stability analysis. It also makes visible why a single fixed model is not sufficient across the full operating envelope. The matrix A changes with nominal gap, nominal current, speed, acceleration, and eddy-current damping. In particular, the gap stiffness term inherited from $\partial F / \partial g$ scales with I_0^2 / g_0^3 . A smaller operating gap may reduce current demand, but it also increases sensitivity to

gap perturbations. Similarly, higher lift demand during braking raises the nominal current at one module and changes the local force-current gain.

The linear model is therefore best interpreted as an operating-point approximation rather than a global representation. It is expected to be accurate when gap perturbations are small relative to g_0 , current commands remain unsaturated, the guideway disturbance remains inside controller bandwidth, and acceleration remains close to the value used in the linearization. It is expected to degrade when the pod encounters large guideway defects, when braking shifts the front-rear force demand substantially, when voltage saturation limits current response, or when eddy-current forces change significantly over the speed range of interest.

This operating-point view provides a practical way to organize controller design. A single linear controller may be adequate for low-speed testing near a fixed payload and guideway condition. A gain-scheduled controller may be needed when speed, payload, and acceleration vary over a wider range. A nonlinear constrained controller becomes more relevant when the design objective includes explicit clearance, current, voltage, and ride-comfort constraints during combined disturbance events.

6 Control Formulations

6.1 Static Feedforward with Gap Feedback

The simplest baseline controller combines a static feedforward current with local gap feedback. The feedforward term is computed from the nominal force balance,

$$I_{0,i} = g_{0,i} \sqrt{\frac{F_{0,i}}{k_i}}, \quad (31)$$

where $F_{0,i}$ is the nominal support force at module i . A feedback law then corrects deviations from the desired gap:

$$\begin{aligned} u_i = & u_{0,i} + K_{p,i}(g_i - g_{0,i}) \\ & + K_{d,i}(\dot{g}_i - \dot{g}_{0,i}) + K_{I,i} \int (g_i - g_{0,i}) dt. \end{aligned} \quad (32)$$

This controller is useful as a reference case because it closely follows the static-design intuition. Its limitation is that it treats each module locally and does not explicitly account for pitch coupling, current constraints, or acceleration-dependent load transfer except through updated feedforward terms.

6.2 State Feedback and Gain Scheduling

A more structured linear design uses the operating-point model to synthesize state feedback,

$$\delta \mathbf{u} = -K \delta \mathbf{x}. \quad (33)$$

The gain K may be selected by pole placement, LQR, or robust loop-shaping. This approach directly controls heave and pitch modes rather than only local gaps, which allows the front and rear modules to coordinate their force commands.

Because the linearized plant varies with operating condition, the gain can be scheduled over speed, acceleration, and payload:

$$K = K(v, a_x, m). \quad (34)$$

This is a practical middle ground. It retains the interpretability and computational simplicity of linear control while acknowledging that the levitation plant is not fixed. The main engineering burden is validation: the scheduled controller must be checked not only at grid points, but also between grid points and during transitions.

6.3 Constraint-Aware Nonlinear Control

When disturbances are large enough to activate constraints, nonlinear control becomes more attractive. A useful comparison controller should not simply invert the force law. It should use the instantaneous gap geometry and also respect module force limits. The control law used here first computes desired front and rear gap accelerations,

$$\ddot{g}_f^* = -k_g(g_f - g_{f,ref}) - c_g\dot{g}_f + b_f(g_f, \dot{g}_f), \quad (35)$$

$$\ddot{g}_r^* = -k_g(g_r - g_{r,ref}) - c_g\dot{g}_r + b_r(g_r, \dot{g}_r), \quad (36)$$

where b_i is a soft safety-barrier term that activates near the hard minimum gap. For example,

$$b_i = \begin{cases} 0, & g_i \geq g_{soft}, \\ k_b(g_{soft} - g_i) + c_b \max(-\dot{g}_i, 0), & g_i < g_{soft}. \end{cases} \quad (37)$$

This term does not replace the nominal controller. It adds emergency opening acceleration when a module approaches the clearance limit, especially when the gap is still closing.

The desired gap accelerations are mapped into desired body accelerations using the front/rear gap geometry. Neglecting guideway second derivatives in the controller model,

$$\ddot{g}_f \approx -\ddot{z} - l_f\ddot{\theta}, \quad (38)$$

$$\ddot{g}_r \approx -\ddot{z} + l_r\ddot{\theta}. \quad (39)$$

Solving these two equations gives desired heave and pitch accelerations, with additional pitch trim used to prevent persistent body attitude. The required total force and pitch moment are then

$$F_{tot}^* = m(g + \ddot{z}^*), \quad (40)$$

$$M^* = J_y\ddot{\theta}^* - m a_x h, \quad (41)$$

where the second term cancels the inertial load-transfer moment. The front and rear forces are allocated by

$$\begin{bmatrix} 1 & 1 \\ l_f & -l_r \end{bmatrix} \begin{bmatrix} F_f^* \\ F_r^* \end{bmatrix} = \begin{bmatrix} F_{tot}^* \\ M^* \end{bmatrix}. \quad (42)$$

The allocated forces are then clamped to current-feasible limits,

$$0 \leq F_i^* \leq \rho k_i \frac{I_{max}^2}{g_i^2}, \quad (43)$$

where $\rho < 1$ reserves margin for current dynamics and command saturation. Finally, the nonlinear magnetic force law is inverted,

$$I_i^* = g_i \sqrt{\frac{F_i^*}{k_i}}, \quad (44)$$

subject to current and command limits. This controller is not a full nonlinear model predictive controller, but it captures the key features needed for the comparison: instantaneous gap-coordinate

regulation, heave-pitch force allocation, load-transfer cancellation, and explicit current-feasible force limits.

Nonlinear model predictive control provides the most direct way to represent the operational levitation problem. At each time step, the controller solves

$$\min_{\mathbf{u}_{0:N-1}} \sum_{j=0}^N \left(\|\mathbf{g}_j - \mathbf{g}_{ref}\|_{Q_g}^2 + \|\ddot{z}_j\|_{Q_z}^2 + \|\theta_j\|_{Q_\theta}^2 + \|\mathbf{u}_j\|_R^2 \right), \quad (45)$$

subject to the nonlinear dynamics and constraints

$$g_{min} \leq g_i \leq g_{max}, \quad (46)$$

$$0 \leq I_i \leq I_{max}, \quad (47)$$

$$|u_i| \leq u_{max}. \quad (48)$$

This framing is computationally heavier, but it matches the physical problem: the controller must regulate gap, limit acceleration for ride quality, avoid pitch excursions, and keep current and voltage within allowable bounds. The constrained nonlinear controller above can be interpreted as a lightweight approximation to this idea. Its value should be judged by whether it expands the safe operating envelope, not simply by whether it reduces small-signal gap error.

7 Operating Envelope and Design Constraints

The reduced-order model is most useful when it is used to define an operating envelope rather than a single nominal operating point. A static levitation design produces a point: a selected gap, current, mass, and lift force. A dynamics-aware design produces a region: the set of speeds, accelerations, payloads, guideway inputs, and actuator states over which the closed-loop system maintains clearance and remains within current, voltage, thermal, and ride-quality limits. This section defines the constraints that should be checked before interpreting any simulation as evidence of levitation feasibility.

7.1 Clearance Margin

The most direct constraint is the minimum allowable air gap. Let g_{nom} be the nominal operating gap and g_{hard} be the minimum physically allowable clearance before contact, scrape, or unacceptable field concentration occurs. A practical controller should maintain a dynamic margin

$$M_g(t) = g_i(t) - g_{hard} \quad (49)$$

for every levitation module. The design condition is not simply $g_{nom} > g_{hard}$, but

$$\min_{t,i} M_g(t) > M_{req}, \quad (50)$$

where M_{req} is a reserved margin for sensing error, guideway uncertainty, unmodeled structural motion, and emergency shutdown transients.

This framing matters because a controller can track the nominal gap well in RMS terms while still producing a small number of dangerous low-clearance events. For levitation, rare minima are often more important than average tracking error. A useful operating envelope should therefore report both continuous ride metrics and worst-case clearance metrics.

7.2 Current, Voltage, and Slew Limits

The electromagnetic actuator is constrained by current, voltage, and current slew. The force law may imply that a recovery force is mathematically available, but the coil and power electronics may not be able to realize it quickly enough. We write the actuator constraints as

$$0 \leq I_i(t) \leq I_{max}, \quad (51)$$

$$|u_i(t)| \leq u_{max}, \quad (52)$$

$$|\dot{I}_i(t)| \leq \dot{I}_{max}. \quad (53)$$

These constraints distinguish force authority from force bandwidth. Static sizing mostly checks the former. Disturbance rejection depends strongly on the latter.

Voltage margin is especially important when the guideway disturbance has high spatial frequency or when speed maps a moderate spatial disturbance into a high temporal frequency. From the current dynamics,

$$u_i = L_i(g_i)\dot{I}_i + R_i I_i + I_i \frac{\partial L_i}{\partial g_i} \dot{g}_i, \quad (54)$$

large current slew and large gap rate both consume voltage. Thus an actuator can be current-capable at a given gap but voltage-limited during a fast transient. This is one of the ways a static force calculation can overstate usable levitation authority.

7.3 Bandwidth and Guideway Spectrum

Guideway quality should be evaluated as a speed-dependent input spectrum. If the guideway profile contains spatial components with wavelength λ , then at speed v those components enter the levitation controller near frequency v/λ . Let ω_b denote a representative closed-loop bandwidth. A simple screening condition is

$$\frac{2\pi v}{\lambda} < \eta \omega_b, \quad (55)$$

where $\eta < 1$ is a design factor that reserves bandwidth for phase margin, unmodeled dynamics, sensing delay, and actuator lag.

This condition should not be interpreted as a hard physical limit, because a controller may attenuate some disturbances above bandwidth or rely on mechanical filtering. It is still useful as a first-pass compatibility check between guideway construction tolerance and levitation control. A guideway profile that is acceptable at low speed may become unacceptable at high speed because its spatial content is shifted into a frequency range where the actuator has little phase margin.

7.4 Ride and Structural Constraints

The controller should not maintain gap by injecting unacceptable body acceleration or pitch motion. A levitation system that preserves clearance while producing large vertical acceleration may be technically stable but operationally poor. We therefore distinguish clearance survival from ride acceptability. Representative ride and body-motion constraints are

$$a_{z,rms} \leq a_{rms,max}, \quad (56)$$

$$|\theta(t)| \leq \theta_{max}, \quad (57)$$

$$|\dot{\theta}(t)| \leq \dot{\theta}_{max}. \quad (58)$$

Table 2: Constraint classes used to define the levitation operating envelope.

Constraint	Metric	Design meaning
Clearance	$\min g_i$	contact margin
Tracking	$\max g_i - g_{ref} $	gap regulation
Current	$\max I_i$	magnetic/thermal margin
Voltage	$\max u_i $	transient authority
Slew	$\max \dot{I}_i $	force bandwidth
Ride	$a_{z,rms}$	passenger/body response
Pitch	$\max \theta $	load-transfer coupling
Drag	$\int D_e v dt$	energy penalty

These constraints are particularly relevant when comparing controllers. A stiff local gap controller may reduce gap error but transmit more guideway disturbance into the pod body. A softer controller may improve ride but consume more clearance margin. The correct tradeoff depends on the available nominal gap, guideway quality, and actuator reserve.

7.5 Envelope Representation

For a given controller, the safe operating region can be represented as a set

$$\mathcal{S} = \{v, a_x, m, A_r, \lambda : \mathcal{C}(\mathbf{x}, \mathbf{u}, \mathbf{w}) = 1\}, \quad (59)$$

where $\mathcal{C}$ is an indicator that all clearance, current, voltage, slew, and ride constraints are satisfied over the simulated event. This representation is more informative than a single success/failure case because it shows which variable closes the envelope first.

For example, one controller may be limited by current saturation during braking, another by voltage saturation over short-wavelength guideway waviness, and another by ride acceleration before clearance is lost. These are different design failures and should lead to different engineering responses. Current saturation suggests larger magnets, more modules, or lower acceleration. Voltage saturation suggests different coil design, power electronics, or guideway spectral limits. Ride limitation suggests different control weighting or mechanical isolation.

This envelope view is the bridge between the model and the simulation study. The simulations should not merely show time histories. They should identify which constraints become active, under which combinations of speed and disturbance, and how those active constraints differ across control strategies.

8 Simulation Study Design

The simulation study is organized around operating cases that static levitation cannot evaluate. Each case uses the same reduced-order pod and guideway model and compares three controller classes. The first is static feedforward with local PID gap feedback. This is the baseline corresponding most closely to the static-design intuition. The second is coordinated body feedback, synthesized from a linearized heave-pitch model. This represents the improvement obtained by treating the front and rear modules as coupled force actuators rather than independent local gap devices. The third is the constraint-aware nonlinear controller described above. This controller operates directly in instantaneous gap coordinates, allocates front/rear forces through the nonlinear geometry, and clamps force commands to current-feasible limits.

Table 3: Simulation cases used to separate static, linear, and nonlinear levitation effects.

Case	Primary input	Main observed limit
Hover	none	small-signal stability
Sinusoid	$r(x)$	bandwidth and speed scaling
Dip	isolated $r(x)$	clearance recovery
Launch	$a_x > 0$	rear/front load shift
Braking	$a_x < 0$	opposite load shift
Brake + dip	$a_x, r(x)$	combined margin
Speed sweep	v, λ	resonance proximity
Uncertainty	m, k_i	robustness

This comparison is designed to test a specific claim. The claim is not that nonlinear control is universally superior. Near nominal hover and under small disturbances, the linearized model should be adequate. The expected separation appears when guideway disturbance, speed, braking, and actuator limits push the system away from the small-signal region. In that regime, nonlinear constrained control should expand the safe operating envelope.

8.1 Scenario Matrix

The baseline cases are:

1. nominal hover at the design gap;
2. constant-speed travel over a sinusoidal guideway profile;
3. constant-speed travel over an isolated guideway dip;
4. acceleration launch over a smooth guideway;
5. emergency braking over a smooth guideway;
6. braking while crossing an isolated guideway dip;
7. speed sweep over periodic guideway waviness;
8. payload and force-map uncertainty under the combined worst case.

The combined braking-and-dip case is expected to be the most informative because it exercises several couplings at once. The guideway dip increases gap demand at one or both modules, braking shifts the nominal front-rear force distribution, pitch modifies the local gap geometry, and current dynamics limit how quickly the levitation modules can respond.

8.2 Envelope Sweep

In addition to the individual time-domain cases, the simulation performs an operating-envelope sweep over speed and isolated guideway-dip amplitude. For each controller, the pod crosses the dip near the middle of the simulated interval so that low-speed cases are not artificially easy. A case is marked unsafe if it violates minimum clearance, current, command, pitch, or ride-acceleration constraints. The resulting map is not a tuning artifact; it is the central dynamic design object. It

shows, for the same static magnetic design, which combinations of speed and guideway disturbance remain admissible under each controller.

The most compact representation is the safe-envelope boundary in the speed-amplitude plane. A controller whose boundary lies at larger disturbance amplitude for the same speed has preserved more dynamic clearance margin. This boundary is the direct simulation counterpart to the thesis of the paper: static force balance gives a nominal point, whereas closed-loop dynamics define a region of feasible operation.

8.3 Output Metrics

The primary clearance metric is minimum gap,

$$g_{min,case} = \min_{t,i} g_i(t). \quad (60)$$

This is the most direct measure of contact risk. The corresponding tracking metric is maximum gap error,

$$e_{g,max} = \max_{t,i} |g_i(t) - g_{ref,i}|. \quad (61)$$

Ride response is evaluated using RMS vertical acceleration and peak pitch angle:

$$a_{z,rms} = \sqrt{\frac{1}{T} \int_0^T \ddot{z}(t)^2 dt}, \quad (62)$$

$$\theta_{max} = \max_t |\theta(t)|. \quad (63)$$

Actuator burden is evaluated through peak current, RMS current, and current slew:

$$I_{max,case} = \max_{t,i} I_i(t), \quad (64)$$

$$\dot{I}_{max,case} = \max_{t,i} |\dot{I}_i(t)|. \quad (65)$$

Energy penalty from motion-dependent electromagnetic effects is tracked through

$$E_{drag} = \int_0^T \sum_i D_{e,i}(t) v(t) dt. \quad (66)$$

These metrics intentionally separate global clearance, ride response, actuator effort, and energy cost. A controller that minimizes gap error may increase acceleration or current. A controller that allows more gap motion may improve comfort but reduce clearance margin. The severity of a levitation event is therefore metric dependent, and the controller comparison should preserve that distinction.

9 Simulation Outputs and Discussion

The first result is intentionally uneventful: nominal hover does not distinguish the controllers. All three controllers satisfy the static lift condition and remain near the design gap when the guideway is flat and no transient loading is applied. This is important because it shows that hover is a weak validation case. A levitation system can satisfy $mg = F(I, g)$ and still have unknown dynamic margin.

The second result is that guideway excitation creates the first meaningful separation between controllers. A sinusoidal profile enters the gap through $r(x(t))$, so increasing speed shifts the same spatial disturbance to higher temporal frequency. In the speed-sweep case, the local PID controller is expected to lose clearance margin or demand increasing current as the disturbance frequency approaches the effective actuator bandwidth. Coordinated body feedback improves this response because it regulates heave and pitch together rather than reacting independently at each module. This confirms that front and rear levitation stations should be treated as coupled body actuators.

The third result is that braking by itself is not necessarily the worst levitation case when load-transfer feedforward or body coordination is included. Smooth braking shifts the front/rear lift demand, but the resulting moment can be largely canceled by coordinated force allocation. The critical case is braking while crossing a guideway defect. That event combines load transfer, pitch coupling, finite current response, and a local gap disturbance. It is therefore the clearest test of whether the controller is merely stabilizing a nominal operating point or actively preserving dynamic clearance margin.

The fourth and most important result is the operating-envelope map. For the same static actuator sizing and nominal gap, the admissible combinations of speed and guideway-dip amplitude differ across controllers. Static feedforward with local PID has the narrowest interpretation: it can regulate small local errors, but it does not explicitly coordinate body motion or account for force feasibility. Coordinated linear body feedback extends the envelope by controlling heave and pitch modes, but it remains tied to a local operating approximation. The constrained nonlinear controller expands the admissible region when large gap excursions and actuator limits become active because it regulates instantaneous front/rear gaps, allocates feasible module forces, and activates a soft clearance barrier near the hard gap.

This result should be interpreted carefully. It does not imply that nonlinear control is always required. Near the nominal operating point, linear control can be adequate and simpler to validate. The nonlinear controller becomes valuable when the pod leaves the small-signal region or approaches hard constraints. In that regime, the relevant comparison is the size of the safe envelope, not whether one controller produces the smallest RMS gap error in mild cases.

The results also clarify how guideway specifications should be written. A maximum spatial deviation alone is not enough, because disturbance severity depends on speed and wavelength. A guideway defect with spatial wavelength λ appears near temporal frequency v/λ . The safe-envelope map therefore links civil tolerance, vehicle speed, and controller bandwidth in a single design object.

Finally, the eddy-current surrogate exposes an energy-stability tradeoff. Increasing electromagnetic coupling can add vertical damping and reduce gap-rate excursions, but it also increases longitudinal drag and heating. The simulation framework therefore records both clearance and drag-energy metrics. This prevents a controller or actuator configuration from appearing favorable only because it hides vertical disturbance rejection inside untracked electromagnetic losses.

10 Design Implications

Several implications follow from the dynamics-aware formulation.

First, levitation actuator sizing should include dynamic current and voltage margins, not only static lift margin. A module that can produce $mg/2$ at the nominal gap may still saturate during a guideway defect or braking event.

Second, front and rear levitation modules should be coordinated through body-level heave and pitch dynamics. Independent local gap controllers can regulate local clearance, but they can also inject pitch moments if they respond to disturbances without accounting for body motion.




Figure 1: Safe-envelope boundary in the speed–guideway-dip-amplitude plane. The key comparison is not small-signal gap tracking, but the size of the admissible operating region before clearance, current, command, pitch, or ride constraints become active.

Third, acceleration and braking profiles are levitation design inputs. The longitudinal motion planner changes the vertical support problem through load transfer. A levitation-aware pod controller should therefore coordinate propulsion, braking, and suspension commands rather than treating them as independent subsystems.

Fourth, guideway tolerance should be evaluated in the frequency domain seen by the moving pod. A millimeter-level spatial deviation does not have a single dynamic meaning; its severity depends on wavelength, vehicle speed, and the closed-loop bandwidth of the levitation system.

Fifth, eddy-current effects should be included in both vertical dynamics and energy accounting. Treating them only as damping ignores their drag and thermal consequences. Treating them only as drag ignores their possible contribution to vertical stability.

Taken together, these implications support a shift in perspective. Hyperloop levitation should be designed as a coupled electromechanical control system with guideway and motion inputs, rather than as a static magnetic lift device that is later stabilized by feedback.

11 Limitations and Next Steps

Several limitations define the scope of the present model. The electromagnetic force law is reduced order and should be replaced by finite-element force maps for a specific actuator geometry. The pod body is modeled as rigid, so structural flexibility and guideway flexibility are not represented. The eddy-current model is a surrogate intended to preserve qualitative speed and gap dependence, not a




Figure 2: Scenario-level metric comparison. The same controller can look favorable under one metric and unfavorable under another, so clearance margin, ride acceleration, pitch, and actuator burden are reported separately.

substitute for electromagnetic diffusion analysis. Thermal dynamics are represented only through current and energy metrics. Lateral, yaw, and guidance dynamics are omitted. The propulsion and braking system is treated as an exogenous source of speed and acceleration rather than as a coupled controller.

These limitations point to natural extensions. A higher-fidelity study should couple this reduced-order control model to electromagnetic force maps, flexible guideway modes, power-electronics limits, and thermal constraints. Multiple levitation modules per side should be represented explicitly so that local guideway defects and sensor noise can be studied. The model should also be extended laterally to capture guidance and roll-yaw coupling. Finally, the most useful system-level extension would be co-design: selecting guideway tolerance, levitation actuator margin, controller bandwidth, and acceleration profile together.

12 Conclusion

This paper presents a reduced-order dynamics-aware formulation for electromagnetic levitation of Hyperloop pods under motion, track irregularity, and transient loading. The central result is that static force balance is a necessary sizing condition but not an operational guarantee. Once the pod moves, the levitation gap becomes a dynamic state shaped by guideway geometry, heave, pitch, current dynamics, speed, eddy-current effects, and longitudinal load transfer.

The model shows why levitation design must be evaluated through closed-loop dynamic metrics: minimum clearance, gap error, ride acceleration, pitch response, peak current, current slew,

voltage demand, and drag energy. It also shows why linearized control is useful but operating-point dependent. Linear models are appropriate when perturbations remain small and constraints are inactive. Nonlinear and predictive formulations become important when guideway disturbances, braking, payload variation, or actuator limits move the pod outside the small-signal envelope.

The broader conclusion is that Hyperloop levitation should be treated as a coupled vehicle-control-guideway problem. Static lift calculations can start the design, but they cannot finish it.

References

- [1] R. M. Goodall, “The theory of electromagnetic levitation,” *Physics in Technology*, 1985.
- [2] H.-W. Lee, K.-C. Kim, and J. Lee, “Review of maglev train technologies,” *IEEE Transactions on Magnetics*, 2006.
- [3] M. T. Thompson, “Flux-canceling electrodynamic maglev suspension: Part II test results and scaling laws,” *IEEE Transactions on Magnetics*, 1999.
- [4] Y. Cai and D. M. Rote, “A review of dynamic stability of repulsive-force maglev suspension system,” Argonne National Laboratory, Tech. Rep.
- [5] J. L. He, D. M. Rote, and H. T. Coffey, “Study of Japanese electrodynamic-suspension maglev systems,” Argonne National Laboratory, Tech. Rep., 1994.
- [6] M. Hermle, A. Kargl, and P. Eberhard, “Vertical vibration reduction of maglev vehicles using nonlinear MPC,” 2025.
- [7] J. Paul, K. N. van Dalen, A. B. Faragau, R. van Leijden, B. Carboni, and A. V. Metrikine, “Principal and combination parametric resonances of an electromagnetically suspended vehicle subject to base excitation,” 2025.