Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Stabilizing unstable periodic orbits of underactuated systems with low-harmonic periodic feedback

Mustafa Mert Ankarali

mertan@metu.edu.tr

Middle East Technical University

Research Article

Keywords:

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11668411/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Stabilizing unstable periodic orbits of underactuated systems with low-harmonic periodic feedback

Abstract

Many rhythmic motions are unstable and must be stabilised by feedback. Along a periodic orbit the linearised dynamics are periodic, so the optimal feedback is periodic too. We ask how much of this time variation is actually needed. We study unstable orbits of three underactuated mechanical systems, a cart-pole, a pendubot and an acrobot, whose linearisations are strongly time varying. Constant gains are the usual choice in practice, and two shortcuts make them look justified. Averaging predicts the open-loop Floquet exponents within a few percent, yet constant gains designed on the averaged model destabilise four of ten orbits. Constant gains designed against the true Floquet multipliers are linearly stable but have almost no basin of attraction on some orbits. Truncated harmonic (lifted) models, the second shortcut, predict stable loops that are unstable on the true system. Taking periodic LQR as the benchmark, we show that the gap between constant and fully periodic feedback is closed by a single harmonic. A gain at twice the orbit frequency, designed against the true Floquet multipliers, comes within 0.1–28 % of the periodic-LQR cost on all symmetric orbits and has, on average, a basin of the same size. A symmetry argument explains why only even harmonics are needed. The optimal periodic gain itself is not low-harmonic: keeping only its leading harmonics destabilises three of six orbits. Finally, indexing the gain by an orbit phase with a phase observer reduces the deviation from the orbit under timing errors.

Keywords: Unstable periodic orbits, Linear time-periodic systems, Floquet theory, Averaging, Harmonic balance, Underactuated systems

1 Introduction

Many tasks of underactuated mechanical systems are rhythmic and unstable. A gymnast swinging on a bar, a robot that pumps a pendulum, a juggling or hopping machine and a walking robot near an upright posture all move along periodic orbits that are unstable without feedback [1–6]. Humans and animals stabilise such rhythmic movements with sensory feedback as well, and the feedback is often phase dependent [7–9]. Designing this feedback is a problem of dynamics and control at the same time [10, 11]. The orbit comes from the nonlinear dynamics, the feedback acts on a linearisation that changes along the orbit, and the closed loop must be checked on the nonlinear system. The question has a long history. Ott, Grebogi and Yorke stabilised unstable orbits embedded in a chaotic attractor with small, well-timed corrections [12], and Pyragas did the same with continuous delayed feedback [13]. In vibrational control, fast periodic forcing changes the stability of an equilibrium, as in the Kapitza pendulum [14–17].

Linearising around a periodic orbit gives a linear time-periodic (LTP) system, and its stability is set by the Floquet multipliers. Such models can also be identified from input–output data of rhythmic systems, through harmonic transfer functions [18–20], and phase-dependent responses have been identified in human walking and in muscle during rhythmic motion [21, 22]. The natural feedback for an LTP system is itself periodic. Two lines of work design it. Harmonic-domain methods lift the LTP system to an infinite-dimensional time-invariant model built from Fourier coefficients [23–30], including controllers designed directly from harmonic transfer functions and tested on an LTP laboratory bench [31, 32]. They are rigorous, but have been applied mainly to mildly

nonlinear systems with stable periodic regimes, such as power converters [33]. Orbital stabilisation of underactuated mechanical systems uses transverse coordinates and periodic Riccati equations or sliding surfaces [34–39]. These methods handle strongly nonlinear, unstable orbits, but produce fully time-varying gains. Periodic and constant output-feedback gains for periodic systems have also been designed by optimisation [40, 41]. A preliminary study [42] explored harmonic-domain gains for a cart-pole around a periodic orbit in a limited setting: it approximated the orbit by a single sinusoid, designed constant gains on a truncated harmonic model and assessed stability on that model. Between a constant gain and a fully periodic one lies a simple question that, to our knowledge, has not been asked for unstable orbits of strongly nonlinear systems: *how much time variation does the feedback actually need?*

Two shortcuts suggest that the answer is "very little". The first is averaging: replace the LTP system by its time average and design a constant gain. The second is truncation: keep a few harmonics of the lifted model and design on that finite model. Both are standard, and both look justified on our systems. The averaged model predicts the open-loop Floquet exponents within a few percent, even when most of the power of the linearisation lies in its harmonics.

This paper shows that both shortcuts mislead in closed loop, and that the answer to the question is nevertheless "very little", provided the little is of the right kind. We use periodic LQR as the benchmark: it is the optimal periodic gain and, as we show, it is also robust. The question is how close a gain with almost no time variation can come to it. On ten unstable orbits of a cart-pole, a pendubot and an acrobot, constant gains designed on the averaged model destabilise four orbits. The frozen-time closed loop is unstable for most of the period, although the averaged closed loop is strongly stable. Truncated lifted models predict finite costs for loops that are unstable on the true system, and an optimiser exploits this. When the feedback is designed against a truncation-free cost and the true Floquet multipliers, with exact gradients, a gain with a single second harmonic comes within 0.1–28 % of the optimal periodic gain on all symmetric orbits. A symmetry argument explains why only even harmonics appear. The low-harmonic gain is not the low-harmonic part of the optimal gain: truncating the optimal gain to the same harmonics destabilises three of the six orbits. In nonlinear simulation the one-harmonic gain has the basin of attraction of periodic LQR. On the two-link orbits, constant gains that are linearly stable have basins that are three times smaller or almost empty; on the milder cart-pole orbit a constant gain is enough. How much time variation is needed thus depends on the orbit, and where a constant gain fails, the step that matters is from a constant gain to a single harmonic.

We then index the gain by an orbit phase, with a phase observer. This makes the loop autonomous and reduces the deviation from the orbit under timing errors.

The answer also matters for practice and teaching. Constant gains are what most implementations and laboratory courses use, and on these orbits they are not enough. A gain with one harmonic is the smallest structure that works: it is evaluated with one sinusoid and a dozen multiply–adds per control step, and it can be read and adjusted by hand. The cart-pole, the pendubot and the acrobot are standard teaching setups [1, 43]. Floquet multipliers, parametric effects, the limits of averaging and the notion of an orbit phase are usually taught on paper, and a controller this simple brings them within reach of a laboratory course.

Section 2 introduces the systems, the orbits and the linearised dynamics. Section 3 shows what averaging and truncation get right and wrong. Section 4 presents low-harmonic feedback, the symmetry result and the nonlinear tests. Section 5 indexes the feedback by phase. Section 6 discusses the results and their limits, and concludes.

2 Systems, orbits and mathematical background

2.1 Systems and periodic orbits

We consider mechanical systems $\dot{x} = f(x, u)$ with configuration q, state $x = (q, \dot{q}) \in \mathbb{R}^n$ and input $u \in \mathbb{R}^m$. The three systems in Fig. 1 have one input and two degrees of freedom ($n = 4$, $m = 1$). The cart-pole is driven by a motor voltage on the cart. The pendubot is actuated at the shoulder and the acrobot at the elbow. All orbits lie

around the upright equilibrium, which is unstable. Parameters are listed in Appendix B.

2.2 Periodic orbits by harmonic balance

The equations of motion are written in manipulator form,

$$M(q)\,\ddot{q} + h(q,\dot{q}) = S\,u, \tag{1}$$

where M is the mass matrix, h collects Coriolis, gravity and damping terms, and S selects the actuated coordinates. A T-periodic orbit $(q_0(t), u_0(t))$ with $\omega = 2\pi/T$ is represented by truncated Fourier series,

$$q_0(t) = \sum_{k=-H}^{H} Q_k\, e^{jk\omega t}, \quad u_0(t) = \sum_{k=-H}^{H} U_k\, e^{jk\omega t}, \\ Q_{-k} = \bar{Q}_k, \qquad U_{-k} = \bar{U}_k. \tag{2}$$

Equivalently, each signal is sampled at the $N = 2H+1$ collocation points $t_i = iT/N$, and derivatives are taken with the spectral differentiation matrix D, which is exact for trigonometric polynomials of degree H. Harmonic balance (trigonometric collocation) [44–46] then requires (1) to hold at every collocation point:

$$r_i = M(q_{0,i})\,\big(D^2 q_0\big)_i + h\big(q_{0,i}, (Dq_0)_i\big) \\ - S\,u_{0,i} = 0, \qquad i = 0,\dots,N-1. \tag{3}$$

With one input and two degrees of freedom, (3) has $2N$ equations in $3N$ samples, so one coordinate is prescribed. We prescribe $q_1(t) = a\sin\omega t$ for the pendubot, the elbow angle $q_2(t) = a\sin\omega t$ for the acrobot, and the pendulum angle $\theta(t) = \pi + a\sin\omega t$ for the cart-pole. The remaining coordinate and u_0 are unknown. A coordinate that enters the dynamics only through its derivatives, such as the cart position, has a free mean, which is fixed by a gauge condition. The nonlinear equations are solved by a Levenberg–Marquardt method with an analytic Jacobian and amplitude continuation. With $H = 30$ the residual $|\dot{x}_0 - f(x_0,u_0)|$ falls below 10^{-9} (Fig. 2). A single-sine approximation of the orbit leaves residuals of 6–420, which is too inaccurate for the analysis that follows. We study ten symmetric orbits with amplitudes $a = 0.3$–$1.2\,\mathrm{rad}$ and frequencies 1–3 Hz, and one asymmetric pendubot orbit, $q_1(t) = 0.15 + a\sin\omega t + 0.3a\sin(2\omega t + 0.7)$.

2.3 Linear time-periodic dynamics

With $x = (q, \dot{q})$, (1) reads $\dot{x} = f(x,u)$. Let $e = x - x_0(t)$ and $v = u - u_0(t)$. Linearising along the orbit gives

$$\dot{e} = A(t)\,e + B(t)\,v, \\ A(t) = \frac{\partial f}{\partial x}\Big|_{x_0(t),u_0(t)}, \qquad B(t) = \frac{\partial f}{\partial u}\Big|_{x_0(t),u_0(t)}, \tag{4}$$

with $A(t+T) = A(t)$ and $B(t+T) = B(t)$. This is a linear time-periodic (LTP) system. It has two properties that an LTI system lacks. Its response depends on when, within the period, an input arrives. And it moves energy between frequencies: an input at one frequency produces outputs at that frequency shifted by every multiple of ω. LTP models are standard for rotating machinery, power converters and helicopter rotors [23, 47]. They also describe rhythmic biological and robotic behaviour linearised about a limit cycle, from human rhythmic motor control and walking to legged robots and muscle mechanics [7–9, 18, 19, 22]. An LTP system can be studied in three equivalent ways: in the time domain through Floquet theory, in the frequency domain through harmonic transfer functions, and as an infinite-dimensional LTI system in the lifted (harmonic) coordinates. Figures 3 and 4 show these views for the pendubot orbit ($a = 1.2\,\mathrm{rad}$, $f = 2\,\mathrm{Hz}$). We use all three in this paper.

2.4 Time domain: Floquet theory

For a periodic feedback $v = -K(t)e$, the closed-loop matrix $A_{\mathrm{cl}}(t) = A(t) - B(t)K(t)$ is also T-periodic. Let $\Phi(t,\tau)$ be its state-transition matrix, $\partial_t\Phi(t,\tau) = A_{\mathrm{cl}}(t)\Phi(t,\tau)$, $\Phi(\tau,\tau) = I$. The monodromy matrix $\Phi_T = \Phi(T,0)$ maps the error over one period. By Floquet's theorem [47–50],

$$\Phi(t,0) = P(t)\,e^{Ft}, \quad P(t+T) = P(t), \quad \Phi_T = e^{FT}, \tag{5}$$

so the periodic system is equivalent, through the periodic change of coordinates $P(t)$, to a time-invariant system with matrix F (Lyapunov–Floquet transformation). Sampling once per

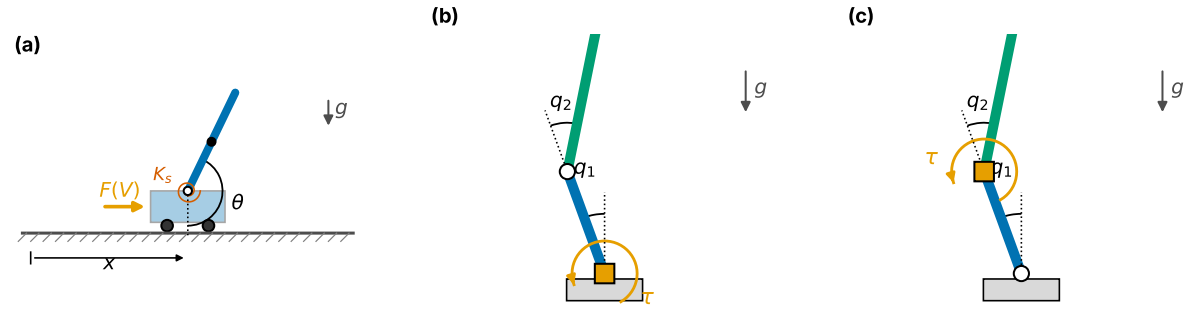

orange square: actuated joint; white circle: passive joint; angles measured from the upright (two-link) / downward (cart-pole)

Fig. 1 The three underactuated systems: (a) cart-pole driven by a DC motor (input: motor voltage; upright at $\theta = \pi$), (b) pendubot (torque at the shoulder, elbow passive), (c) acrobot (torque at the elbow, shoulder passive). Orange squares: actuated joints; white circles: passive joints.

period gives the stroboscopic map

$$\begin{aligned} e\big((k+1)T+\tau\big) &= \Phi(T+\tau,\tau)\, e(kT+\tau), \\ \Phi(T+\tau,\tau) &= P(\tau)\,\Phi_T\,P(\tau)^{-1}, \end{aligned} \tag{6}$$

a discrete-time LTI system. If $e(0)$ is an eigenvector of Φ_T with multiplier $\mu = e^{\lambda T}$, then $e(t) = e^{\lambda t}p(t)$ with p T-periodic: an exponential times a periodic waveform. Its matrix depends on the sampling instant τ, but only through a similarity transformation, so its eigenvalues do not. These eigenvalues μ_i of Φ_T are the Floquet multipliers, and $\lambda_i = \frac{1}{T}\ln\mu_i$ are the Floquet exponents. The logarithm is defined only modulo $j\omega$: if λ is an exponent, so is $\lambda + jk\omega$ for every integer k. A negative real multiplier gives an exponent with imaginary part $\omega/2$; it corresponds to a mode that changes sign every period. The loop is asymptotically stable if and only if $|\mu_i| < 1$ for all i, that is, if the spectral radius $\rho(\Phi_T) < 1$. We call $\max_i \operatorname{Re}\lambda_i$ the Floquet abscissa. Figure 3a shows this decomposition on the pendubot loop. The error does not decay monotonically; within each period it rises and falls with $p(t)$. Sampled once per period, however, it decays exactly as $e^{\lambda t}$. From a generic initial error the faster modes die out, and samples taken at any phase τ decay at the rate of the slowest exponent (Fig. 3b). The open-loop orbits have Floquet abscissae of 4.4–12.9 s^{-1} (Table 1), so every orbit is strongly unstable.

The instantaneous (frozen-time) eigenvalues of $A_{\mathrm{cl}}(t)$ do not decide stability. A classical example of Markus and Yamabe [51] has stable eigenvalues at every instant and growing solutions. The converse also occurs; Section 3 shows stabilised loops whose frozen-time matrices are unstable for most of the period.

2.5 Lifted representation and harmonic transfer functions

Write $A(t) = \sum_k A_k e^{jk\omega t}$, and similarly B and K. Following Wereley [23], consider exponentially modulated periodic (EMP) signals $e(t) = \sum_k e_k\, e^{(s+jk\omega)t}$, the LTP counterpart of e^{st}. Substituting into (4) and matching the coefficients of $e^{(s+jk\omega)t}$ gives, for the stacked vector $\mathbf{e} = (\dots, e_{-1}, e_0, e_1, \dots)$,

$$\begin{aligned} s\,\mathbf{e} &= (\mathcal{A} - \mathcal{N})\,\mathbf{e} + \mathcal{B}\,\mathbf{v}, \\ \mathcal{A} = \big[A_{k-l}\big]_{k,l}, &\qquad \mathcal{B} = \big[B_{k-l}\big]_{k,l}, \\ \mathcal{N} &= \operatorname{blkdiag}(jk\omega I), \end{aligned} \tag{7}$$

where $\mathcal{A}$ and $\mathcal{B}$ are block Toeplitz: block (k,l) depends only on $k-l$. Block A_{k-l} tells how much of the error at harmonic l drives the error at harmonic k. This harmonic state-space model is time invariant but infinite dimensional. The eigenvalues of $\mathcal{A} - \mathcal{N}$ are the Floquet exponents repeated at every shift $jk\omega$, $k \in \mathbb{Z}$ [23, 26], which is the non-uniqueness of the logarithm seen from the frequency side (Fig. 4a).

With an output $y = C(t)e$, the input–output map in EMP coordinates is the harmonic transfer

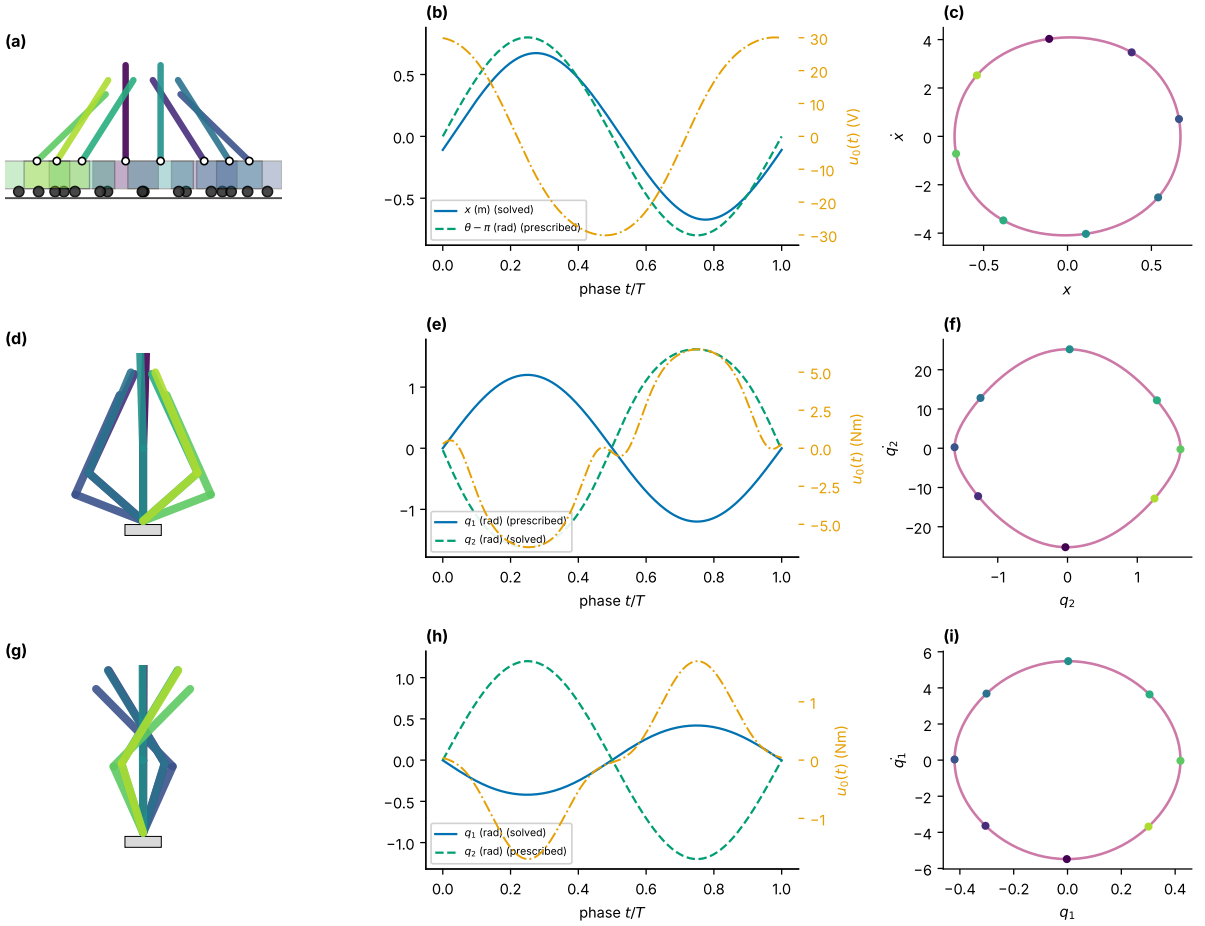

Fig. 2 Periodic orbits computed by harmonic balance: (a–c) cart-pole, (d–f) pendubot, (g–i) acrobot (parameters in Table 1). Left column: configurations at eight phases of one period (dark to light: phase 0 to $7T/8$). Middle column: coordinates and the open-loop input $u_0(t)$. Right column: phase portrait of the solved coordinate.

function (HTF) [23, 27]

$$\mathbf{y} = \mathcal{G}(s)\,\mathbf{v}, \qquad \mathcal{G}(s) = \mathcal{C}\left(sI - (\mathcal{A} - \mathcal{N})\right)^{-1}\mathcal{B}. \quad (8)$$

Its entry $\mathcal{G}_{k,l}$ is the gain from the input at $s + jl\omega$ to the output at $s + jk\omega$. For an LTI system $\mathcal{G}$ is block diagonal, so each frequency maps only to itself. For an LTP system the off-diagonal entries couple frequencies (Fig. 4c). HTFs are the basis of frequency-domain analysis [27] and identification of LTP systems, including rhythmic hybrid systems, hopping robots and human rhythmic motion [18–20].

Feedback acts as $\mathcal{A} - \mathcal{N} - \mathcal{B}\mathcal{K}$, with $\mathcal{K}$ the Toeplitz matrix of the K_k. The structure of the gain is visible in $\mathcal{K}$ (Fig. 4d). A constant gain gives a block-diagonal $\mathcal{K}$: it does not couple harmonics by itself, but the loop is still LTP because $\mathcal{A}$ and $\mathcal{B}$ do. A gain with harmonics in a set S fills only the block diagonals $k - l \in \{0\} \cup S \cup (-S)$. The question of this paper, how much time variation the feedback needs, is therefore how many diagonals of $\mathcal{K}$ must be non-zero. On symmetric orbits the plant fills only the even diagonals; Section 4 shows why and what this implies for the gain.

In practice (7) is truncated to $|k| \le h$. The truncated matrix is no longer exactly Toeplitz at its edges. Its spectrum approximates the shifted Floquet exponents near the centre of the band, but it also contains spurious eigenvalues near the edges, and the highest retained harmonics are poorly resolved. Figure 4b shows this for the periodic-LQR loop with $h = 12$: the eigenvalues for $|k| \le 4$ are within $0.2\,\mathrm{s}^{-1}$ of the true

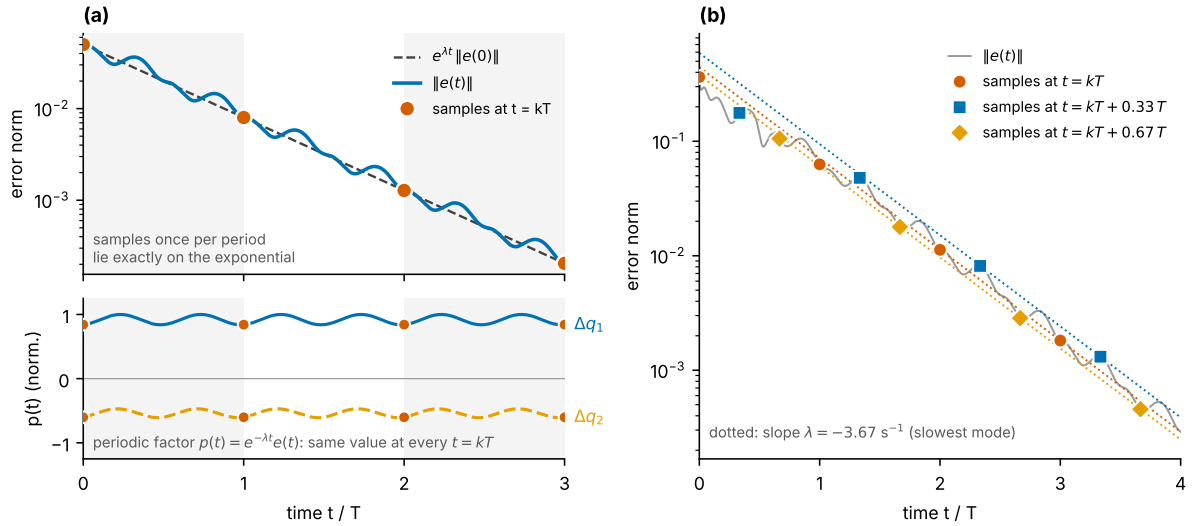

Fig. 3 Floquet theory in the time domain (pendubot, $a = 1.2\,\mathrm{rad}$, $f = 2\,\mathrm{Hz}$, periodic-LQR loop). (a) Start on the slowest Floquet mode ($\lambda = -3.67\,\mathrm{s}^{-1}$). Top: the error norm oscillates within each period, but the samples at $t = kT$ (dots) lie exactly on $e^{\lambda t}\|e(0)\|$ (dashed). Bottom: removing the exponential leaves the periodic factor $p(t) = e^{-\lambda t}e(t)$ (angle components, normalised), which takes the same value at every $t = kT$. (b) Generic initial error. Samples taken once per period at three different phases τ (markers) all decay at the same rate, set by the slowest exponent (dotted), because the multipliers do not depend on τ; see (6).

ones, those beyond drift, and a spurious pair sits on the imaginary axis at the edge of the band. With $h = 4$ the same stable loop even has a pair at $+0.22\,\mathrm{s}^{-1}$. Section 3 shows the opposite error: truncated models that report stability that is not there.

2.6 Averaging, and when time variation matters

The averaged model keeps only A_0 and B_0, that is, only the central block of the lifted model. Its justification is classical averaging theory [52], or, for linear systems, the Magnus expansion [53] of the monodromy matrix, $\Phi_T = \exp(T\,\Omega)$ with

$$\Omega = \Omega_1 + \Omega_2 + \dots, \qquad \Omega_1 = \frac{1}{T}\int_0^T A_{\mathrm{cl}}(t)\,dt,$$
$$\Omega_2 = \frac{1}{2T}\int_0^T\int_0^{t_1} \big[A_{\mathrm{cl}}(t_1), A_{\mathrm{cl}}(t_2)\big]\,dt_2\,dt_1, \tag{9}$$

where $[X, Y] = XY - YX$; the series converges when $\int_0^T \|A_{\mathrm{cl}}\|\,dt < \pi$ [54]. Averaging keeps Ω_1. In open loop $\Omega_1 = A_0$. With a constant gain K, $\Omega_1 = A_0 - B_0K$, which is what an averaged design places. The second-order term Ω_2 consists of commutators that involve the oscillating part $\tilde{A}_{\mathrm{cl}} = A_{\mathrm{cl}} - \Omega_1$, that is, products of the harmonics A_k and B_kK with each other and with the mean. A gain chosen on Ω_1 alone changes Ω_2 and the higher terms without controlling them. Classical examples show that the oscillating part can both destabilise a stable mean (parametric resonance in the Mathieu equation [55]) and stabilise an unstable one (the Kapitza pendulum) [14, 17].

Figure 5 asks when this happens for our systems. For every orbit on a grid of amplitudes a and frequencies f we computed the open-loop Floquet abscissa, the largest real eigenvalue of A_0, and the Floquet abscissa of the true loop under the averaged LQR gain (designed on A_0, B_0 with the weights of Appendix B). In open loop the averaged model is accurate: on the pendubot it is within $1.3\,\mathrm{s}^{-1}$ of the Floquet abscissa on the whole grid, and within $2.8\,\mathrm{s}^{-1}$ on the acrobot (Fig. 5a). In closed loop it is not. The averaged prediction $\max \mathrm{Re}\,\lambda(A_0 - B_0K)$ stays at -3.2 to $-3.4\,\mathrm{s}^{-1}$ on every orbit, while the true abscissa ranges from -6.6 to above $+25\,\mathrm{s}^{-1}$. The averaged gain fails on 36 % of the pendubot grid and 14 % of the acrobot grid. On the pendubot it is safe below $a \approx 0.8$ rad at every frequency and fails above $a \approx 1.1\,\mathrm{rad}$ at most frequencies (Fig. 5b); on the acrobot it fails at large amplitude and high

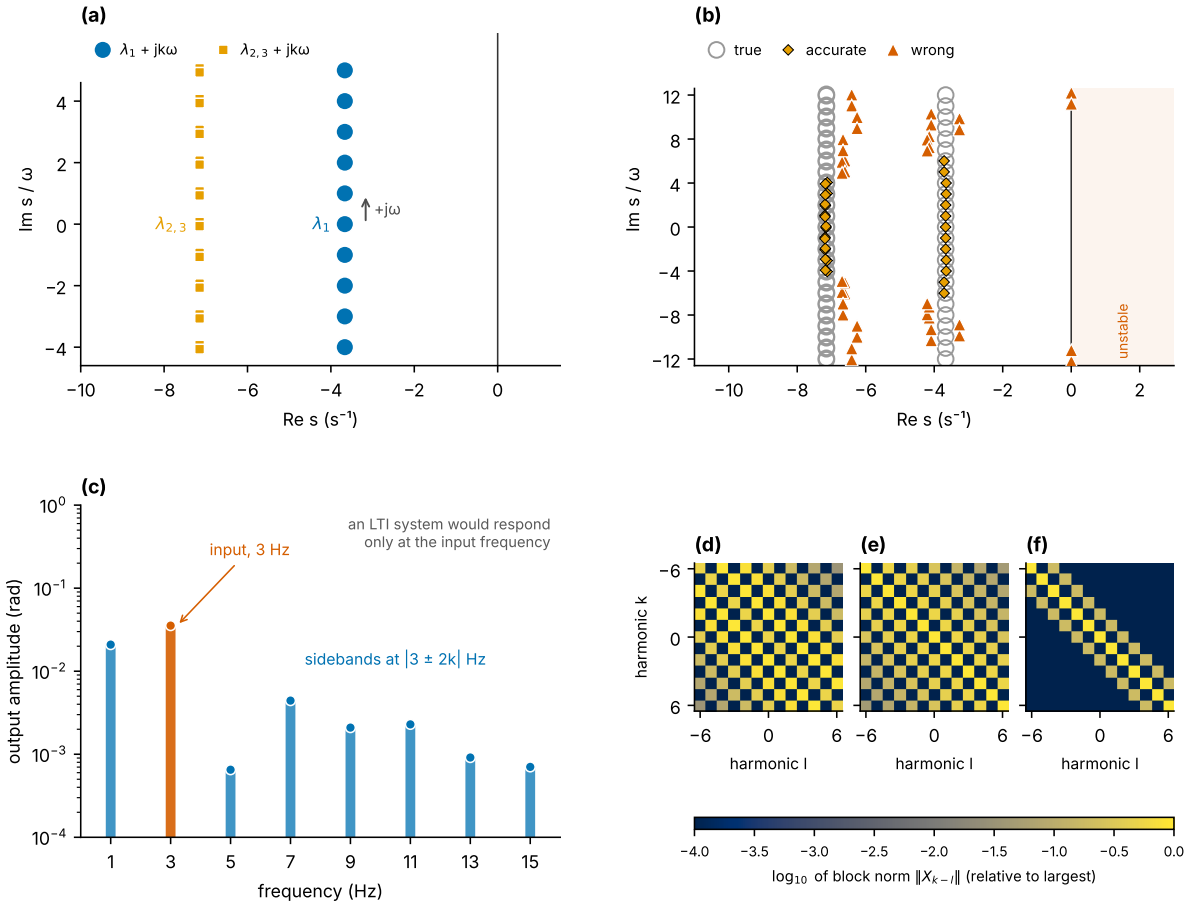

Fig. 4 The lifted (harmonic) view of the same loop. (a) Eigenvalues of $\mathcal{A}-\mathcal{N}-\mathcal{B}\mathcal{K}$: each Floquet exponent appears once in every strip of height ω (circles: the real exponent λ_1; squares: the complex pair $\lambda_{2,3}$; a fourth exponent at $-67\,\mathrm{s}^{-1}$ is off-scale). (b) Eigenvalues of the model truncated at $|k| \le 12$ against the true ones (grey). Near the centre they are accurate (diamonds, error $< 0.3\,\mathrm{s}^{-1}$). Towards the edges they drift (triangles), and a spurious pair appears on the imaginary axis although the loop is stable. (c) Frequency coupling: a sinusoidal input at 3 Hz through $B(t)$ produces outputs at $|3+2k|$ Hz (passive angle). (d–f) Block norms of the lifted matrices on a log scale: (d) the plant matrix $\mathcal{A}$, (e) the periodic-LQR gain $\mathcal{K}$, (f) the one-harmonic gain $\mathcal{K}$. The first two fill every even diagonal; the one-harmonic gain fills only the diagonals 0 and ± 2.

frequency (Fig. 5c). The Magnus convergence condition gives no warning: $\int_0^T \|A_{\mathrm{cl}}\|\,dt$ is at least 49 in open loop, where averaging works, and above 10^3 in closed loop. Truncating the series after Ω_2 does not help either; for the averaged-LQR loop on the pendubot orbit ($a = 1.2\,\mathrm{rad}$, $f = 2\,\mathrm{Hz}$) the eigenvalues of $\Omega_1+\Omega_2$ reach $+243\,\mathrm{s}^{-1}$, against a true abscissa of $+8.7\,\mathrm{s}^{-1}$. Only the monodromy matrix decides. Section 3 analyses the mechanism.

2.7 Quadratic cost and periodic LQR

As a performance measure we use the period-averaged quadratic cost of the loop driven by unit-intensity white noise w, $\dot{e} = A_{\mathrm{cl}}(t)e + w$:

$$J(K) = \frac{1}{T}\int_0^T \mathrm{tr}\big[(Q + K^\top R K)\,S(t)\big]\,dt, \qquad (10)$$

where the state covariance $S(t)$ is the T-periodic solution of the periodic Lyapunov equation

$$\dot{S} = A_{\mathrm{cl}}(t)\,S + S\,A_{\mathrm{cl}}(t)^\top + I, \qquad S(t+T) = S(t). \qquad (11)$$

The periodic solution exists and is unique if and only if the loop is asymptotically stable, so J is finite only for stabilising gains. Its initial value solves the discrete Lyapunov equation $S(0) =$

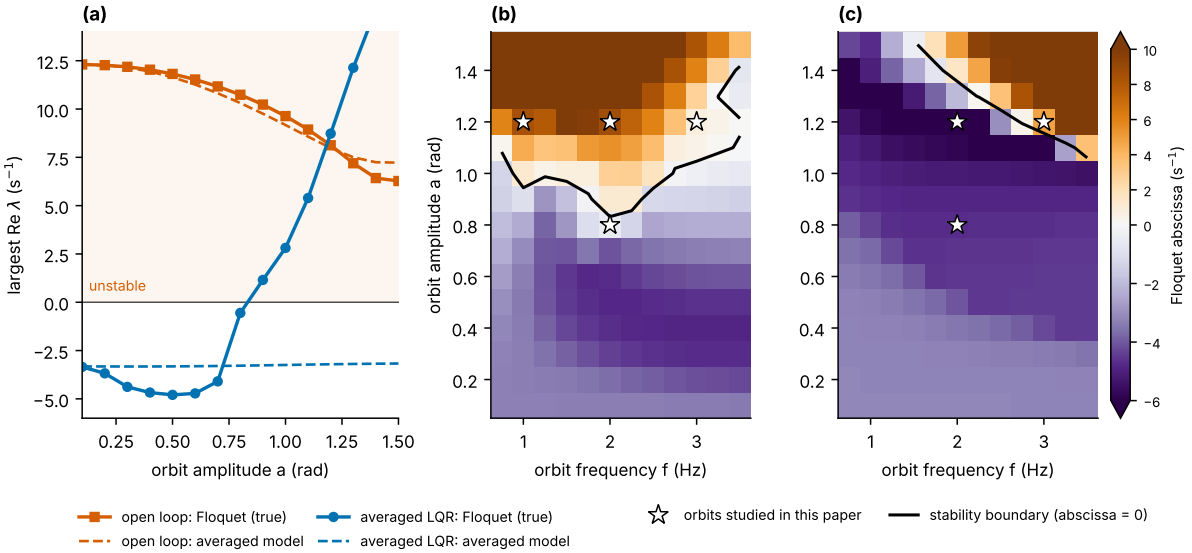

Fig. 5 When does averaging fail? (a) Pendubot, $f = 2\,\mathrm{Hz}$: largest real part of the Floquet exponents (solid) and of the averaged model (dashed), in open loop and with the averaged LQR gain, versus orbit amplitude. The averaged model tracks the open loop but predicts a fixed decay rate in closed loop, while the true loop loses stability between $a = 0.8$ and $0.9\,\mathrm{rad}$. (b,c) Floquet abscissa of the true loop under the averaged LQR gain over orbit amplitude and frequency, for (b) the pendubot and (c) the acrobot (purple: stable; orange: unstable). The averaged prediction is between -3.4 and $-3.2\,\mathrm{s}^{-1}$ everywhere. Stars: orbits of Table 1.

$\Phi_T S(0)\Phi_T^\top + \int_0^T \Phi(T,\tau)\Phi(T,\tau)^\top d\tau$, and J is the squared H_2 norm of the periodic system from w to the performance output $(Q^{1/2}e,\ R^{1/2}Ke)$ [47]. The minimiser of J over all periodic gains is periodic LQR,

$$-\dot{P} = A^\top P + PA - PBR^{-1}B^\top P + Q,$$
$$P(t+T) = P(t), \quad K_{\mathrm{pLQR}}(t) = R^{-1}B(t)^\top P(t), \tag{12}$$

where P is the unique symmetric, positive semi-definite, stabilising T-periodic solution of the periodic Riccati equation, which exists when (A, B) is stabilisable and $(Q^{1/2}, A)$ detectable [47, 56]; numerical methods are reviewed in [57, 58]. We use periodic LQR as the benchmark. All time-indexed designs below minimise the same cost (10), with the same weights Q and R and hence the same input penalty, but over a restricted class of gains. The ratio $J/J_{\mathrm{pLQR}} \geq 1$ therefore measures the price of restricting the structure of the gain, not a difference in objectives. Properties outside the cost, such as the basin of attraction or the tolerance to delay, are targeted by neither design; we report them for both. We compute (11) in the time domain without truncation (Appendix A).

3 What averaging and truncation get right and wrong

3.1 Averaging is accurate in open loop

For every orbit, the open-loop Floquet exponents agree with the eigenvalues of A_0 to within a few percent. This holds even when 97 % of the power of $[A\ B]$ lies in the harmonics. In open loop the oscillating parts of A average out. Averaging is therefore a tempting basis for design.

3.2 Averaging fails in closed loop

Table 1 compares controllers on all orbits. LQR designed on the averaged model places the eigenvalues of $A_0 - B_0K$ near $-3.2\,\mathrm{s}^{-1}$ on every orbit. On the true system it destabilises four orbits, with Floquet abscissae up to $+8.7\,\mathrm{s}^{-1}$. Where it stabilises, its cost is up to 20 times that of periodic LQR. In nonlinear simulation of the pendubot orbit ($a = 1.2$, $f = 2\,\mathrm{Hz}$) the averaged-LQR loop leaves the orbit within 0.14 s.

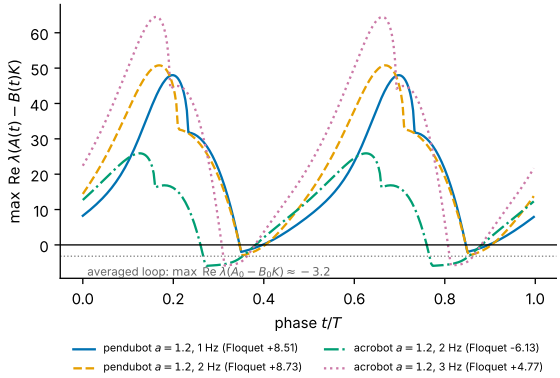

Fig. 6 Frozen-time stability of the averaged-LQR loop: largest real part of the eigenvalues of $A(t)-B(t)K$ over one period. The averaged closed loop has all eigenvalues near $-3.2\,\mathrm{s}^{-1}$ (dotted line), yet the frozen-time loop is unstable for most of the period. Legend: true Floquet abscissa.

The frozen-time picture shows why (Fig. 6). The matrix $A(t) - B(t)K$ has eigenvalues with real parts up to $65\,\mathrm{s}^{-1}$ and is unstable for 65–90 % of the period on the two-link orbits. The averaged matrix $A_0 - B_0K$ hides this. In open loop the oscillations of A cancel over a period. With feedback, the product $B(t)K$ adds oscillating terms that interact with the oscillations of A. These products survive averaging at second order [17], and on large orbits they dominate. Averaging commutes with the open-loop dynamics but not with feedback.

The same mechanism can also create slow dynamics. On a cart-pole with a torsional spring, whose upright is weakly stable, the fast horizontal motion of the pivot nearly cancels the net stiffness, as in the Kapitza pendulum. The orbit has a Floquet exponent of $-0.007\,\mathrm{s}^{-1}$, and this mode is almost uncontrollable through B_0: its modal input gain is 7.6×10^{-4}, against 0.14–8 for the other modes. A controller designed on the averaged model sees a slow, nearly uncontrollable drift whose origin is the orbit, not the plant.

3.3 Truncation creates stability that is not there

The lifted model must be truncated. With low truncation orders the truncated model reports stable loops with finite cost that are unstable on the true system (Fig. 7). On the pendubot orbit, a constant gain designed with $h = 1$ has a predicted cost of 381 but a true Floquet abscissa of $+12.4\,\mathrm{s}^{-1}$. The truncated model also has spurious eigenvalues at the edge of the retained band, and the edge harmonics are poorly resolved. When the cost or a spectral measure of the truncated model is optimised directly, the optimiser uses these artefacts. In one case it reported a strongly negative spectral abscissa for a loop that was unstable. Reliable designs need $h \geq 4$ here, and even then the design must be checked against the true monodromy. This is the reason for the truncation-free design in the next section.

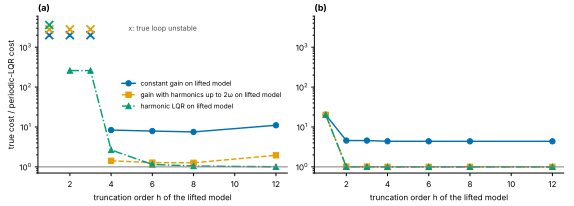

Fig. 7 True cost of gains designed on truncated lifted models, relative to periodic LQR, versus the truncation order h: (a) pendubot and (b) acrobot, both $a = 1.2\,\mathrm{rad}$ and $f = 2\,\mathrm{Hz}$. Crosses: the true loop is unstable although the truncated model predicted a finite cost.

4 Low-harmonic feedback

4.1 Gain class and truncation-free design

We restrict the gain to a set S of harmonics,

$$K(t) = K_0 + \sum_{k \in S} \left(K_k e^{jk\omega t} + \bar{K}_k e^{-jk\omega t} \right), \quad (13)$$

so that $S = \emptyset$ is a constant gain. For a single input and $n = 4$ states, $S = \{2\}$ has 12 real parameters. The design minimises

$$\log J(K) \; - \; \mu \, \log\!\left(1 - \frac{\rho(\Phi_T)}{e^{-\sigma T}}\right), \quad (14)$$

where $\rho(\Phi_T)$ is the spectral radius of the true closed-loop monodromy. The barrier keeps every iterate stable with decay rate at least σ. Both J and Φ_T are computed in the time domain on a grid of 512 points per period, with an exponential midpoint rule, and their exact gradients are obtained by automatic differentiation [59] (Appendix A). Gradient-based design of structured periodic gains goes back to periodic output feedback [40]. No truncation enters. The harmonic set is chosen

Table 1 Closed-loop results on the ten symmetric orbits. Each entry: largest real Floquet exponent (s^{-1}; **bold**: unstable) and, below it, the cost J relative to periodic LQR (— : not stabilising). "OL": open-loop Floquet abscissa. Lifted-model designs use truncation order $h = 8$.

orbit	OL	averaged LQR	lifted, constant	lifted, $\leq 2\omega$	lifted, $\leq 4\omega$	periodic LQR
cart-pole a=0.3, 1 Hz	4.36	−1.31 / 1	−1.33 / 1	−1.33 / 1	−1.33 / 1	−1.33 / 1
cart-pole a=0.8, 1 Hz	5.77	−1.12 / 1.16	−1.29 / 1.09	−1.32 / 1.01	−1.31 / 1	−1.31 / 1
cart-pole a=0.8, 2 Hz	8.48	−1.19 / 1.29	−1.37 / 1.21	−1.30 / 1.01	−1.31 / 1	−1.31 / 1
pendubot a=0.8, 2 Hz	10.7	−0.55 / 17.6	−1.58 / 1.84	−3.17 / 1.02	−3.43 / 1	−3.40 / 1
pendubot a=1.2, 1 Hz	8.19	**+8.51** / —	−0.61 / 752	−2.90 / 1.06	−3.20 / 1.02	−3.28 / 1
pendubot a=1.2, 2 Hz	8.13	**+8.73** / —	−3.70 / 7.52	−3.11 / 1.27	−4.13 / 1.07	−3.67 / 1
pendubot a=1.2, 3 Hz	8.10	**+0.82** / —	−2.77 / 4.75	−3.39 / 1.28	−3.80 / 1.07	−3.97 / 1
acrobot a=0.8, 2 Hz	11.9	−4.69 / 1.99	−1.23 / 1.61	−3.21 / 1	−3.21 / 1	−3.21 / 1
acrobot a=1.2, 2 Hz	11.2	−6.13 / 20.2	−0.70 / 4.37	−3.29 / 1	−3.27 / 1	−3.27 / 1
acrobot a=1.2, 3 Hz	12.9	**+4.77** / —	−1.09 / 8.59	−2.65 / 1.03	−3.45 / 1	−3.53 / 1

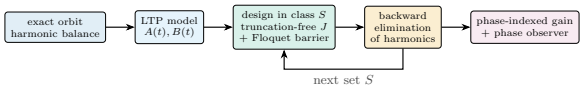

Fig. 8 Pipeline. Sections 2–4 cover the first four blocks and Section 5 the phase-indexed gain.

by greedy backward elimination: starting from $S = \{1, \ldots, 6\}$, the harmonic whose removal increases the cost least is removed, and the design is repeated. Figure 8 summarises the procedure.

The design differs from periodic LQR in kind. Periodic LQR follows from a periodic Riccati equation, which returns the optimal gain but cannot impose any structure on it. Here the gain is a parametrised class, and the cost and the stability constraint are evaluated on the true periodic system, with exact gradients. Any parametrisation can be used in the same way: a set of harmonics, a sparsity pattern, or the phase-indexed gain and phase observer of Section 5, for which no Riccati solution exists. The barrier also enforces a minimum decay rate σ, which periodic LQR does not.

Large language model tools (Claude, Anthropic) were used to help write simulation code, produce figures and draft text. The authors reviewed, verified and takes full responsibility for all content.

4.2 Symmetry selects the even harmonics

All orbits in Table 1 have a half-wave symmetry. The following result shows what this implies for the feedback.

Proposition 1. *Let the coordinates be centred at the upright equilibrium, and suppose that* $f(-x, -u) = -f(x, u)$ *and that the orbit satisfies* $x_0(t + T/2) = -x_0(t)$ *and* $u_0(t + T/2) = -u_0(t)$. *Then*

(i) *A and B have period $T/2$; their Fourier series contain only even harmonics of ω;*

(ii) *the periodic LQR gain has period $T/2$ and contains only even harmonics;*

(iii) *the cost (10) is invariant under the shift $K(t) \mapsto K(t + T/2)$, which reverses the sign of the odd harmonics of K. Hence a gradient-based design started from a gain with only even harmonics keeps only even harmonics.*

Proof Differentiating $f(-x, -u) = -f(x, u)$ gives $f_x(-x, -u) = f_x(x, u)$ and $f_u(-x, -u) = f_u(x, u)$. Along the orbit, $A(t + T/2) = f_x(-x_0(t), -u_0(t)) = A(t)$, and the same holds for B. This proves (i). The stabilising T-periodic solution $P(t)$ of the periodic Riccati equation is unique [47]. Since A, B, Q, R are invariant under the shift by $T/2$, $P(t + T/2)$ is also a stabilising T-periodic solution, so $P(t + T/2) = P(t)$ and $K = R^{-1}B^{\top}P$ has period $T/2$. This proves (ii). For (iii), the loop with gain $K(\cdot + T/2)$ is the loop with gain K shifted in time by $T/2$, because A and B are unchanged by the shift. The period-averaged cost is invariant under time shifts. The shift multiplies the harmonic K_k by $(-1)^k$, so the gradient of J at a gain with only even harmonics has no odd components. □

The assumption holds for the pendubot and acrobot with angles measured from the upright, and for the cart-pole after shifting the pendulum angle by π and centring the cart. We verified (i) and (ii) numerically on the symmetric orbits: A, B and the periodic LQR gain repeat after $T/2$ to within 10^{-8} relative error. On the asymmetric orbit they do not. One consequence is that the monodromy matrix is the square of the half-period transition matrix, so the Floquet multipliers come in squares.

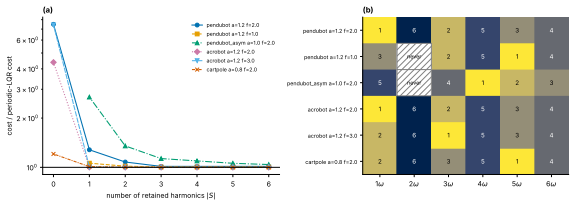

Fig. 9 Backward elimination of harmonics. (a) Cost relative to periodic LQR versus the number of retained harmonics. (b) The order in which harmonics are removed (1 = first; "never": no feasible constant gain, so this harmonic is always kept). On symmetric orbits the odd harmonics go first and 2ω is kept last; on the asymmetric orbit ω and 2ω are kept.

4.3 Results

The backward elimination confirms the proposition (Fig. 9 and Table 2). On every symmetric orbit the odd harmonics are removed first, at no cost, and $\{2\}$ is the last harmonic to be kept. A single 2ω harmonic brings the cost within 0.1–28 % of periodic LQR. On the asymmetric orbit the odd harmonics matter: $\{1, 2\}$ gives 35 % and $\{1, 2, 3\}$ gives 13 %. A constant gain is 1.2–7.5 times worse than periodic LQR, and on two orbits no constant gain reaches the required decay rate.

The low-harmonic gain is not the low-harmonic part of the optimal gain. Truncating periodic LQR to $\{2\}$ destabilises both large pendubot orbits (Floquet abscissae +8.12 and $+3.34\,\mathrm{s}^{-1}$), and truncating it to $\{1, 2\}$ destabilises the asymmetric orbit ($+3.53\,\mathrm{s}^{-1}$). On the acrobot and cart-pole orbits the truncated gain works, so truncation is unreliable rather than always wrong. The optimal periodic gain has large higher harmonics (Fig. 10), and removing them changes the closed loop strongly. The designed gains are smooth and look quite different. Feedback with little time variation exists, but it must be designed as such.

4.4 Basins of attraction and implementation

Linear stability does not guarantee a useful basin of attraction [37, 60], so we test the gains on the nonlinear system. Figure 11 shows slices of the basins in the plane of the two configuration deviations. Within the plotted window, the one-harmonic gain has a basin as large as that of periodic LQR on the acrobot. On the pendubot it contains almost all of the periodic-LQR basin (164 of 170 grid cells; the other six are isolated edge cells) and extends further along the long direction of the basin (121 % of the area). The constant gain keeps 29 % of the area on the pendubot. On the acrobot its basin is too small for the grid: bisection along six directions in this plane gives a radius of only 0.003–0.012 rad, although the gain is linearly stable. All basins are long and thin. They are narrow in the absolute angle of the link that hangs on the passive joint ($q_1 + q_2$ for the pendubot, q_1 for the acrobot) and long in the direction that the actuated joint can correct.

The cart-pole reverses the ranking (bottom row of Fig. 11). Its orbit is the mildest: a constant gain is within 21 % of the periodic-LQR cost. In the slice of cart position and pendulum angle, the constant gain has the largest basin (253 % of the periodic-LQR area), the one-harmonic gain 221 %, and periodic LQR the smallest, in agreement with its smaller radius in Table 3. Here no time variation is needed, and the design method shows this, since the constant gain is already close to the optimum.

Table 3 extends the test to all orbits and to all directions in state space. The region of attraction is measured by bisection along eight random directions in state space. The low-harmonic gain keeps the region of attraction of periodic LQR. Its mean radius is within 5 % of periodic LQR, or larger, on every orbit. Its minimum radius is equal or larger on four orbits and smaller by 13 % and 33 % on the other two, both pendubot orbits. On the pendubot, periodic LQR therefore keeps an advantage in the worst direction of the full state space, which the angle slice of Fig. 11 does not show; on average the two basins are equal. A constant gain shrinks it threefold on the pendubot and almost to zero on the acrobot, although it is linearly stable. On the cart-pole all designed gains reach the largest radius tested.

The optimal periodic gain is not low-harmonic. On the symmetric pendubot orbits, periodic LQR must keep harmonics up to 10ω or 12ω before its truncation stays within 5 % of its cost (44–52 numbers), and on the asymmetric orbit it needs harmonics up to 13ω (108 numbers). Truncating it earlier degrades or destabilises the loop. A gain designed with one harmonic needs 12 numbers. On the acrobot and cart-pole orbits truncated periodic LQR is already near optimal at 2ω, and the two gains coincide in structure. The design itself is slower than solving the periodic Riccati equation

Table 2 Low-harmonic gains designed against the true Floquet multipliers. Cost relative to periodic LQR. "Infeasible": no constant gain reaches the decay rate $\sigma = 0.5\,\mathrm{s}^{-1}$. Truncated pLQR: the periodic LQR gain truncated to $\{2\}$ (Floquet abscissa in s^{-1} in brackets when unstable).

orbit	constant	$\{2\}$	$\{2,4\}$	selected ≤ 3	truncated pLQR $\{2\}$
pendubot a=1.2, 2 Hz	7.49	1.28	1.08	1.01 $\{2,4,6\}$	unstable (+8.12)
pendubot a=1.2, 1 Hz	infeasible	1.06	1.02	1.00 $\{2,4,6\}$	unstable (+3.34)
pendubot asymmetric	infeasible	2.69	—	1.13 $\{1,2,3\}$	158; $\{1,2\}$: unstable (+3.53)
acrobot a=1.2, 2 Hz	4.39	1.001	1.000	1.000 $\{2,4,6\}$	1.00
acrobot a=1.2, 3 Hz	7.43	1.004	1.000	1.000 $\{2,4,6\}$	1.00
cart-pole a=0.8, 2 Hz	1.21	1.010	1.002	1.001 $\{2,4,6\}$	1.02

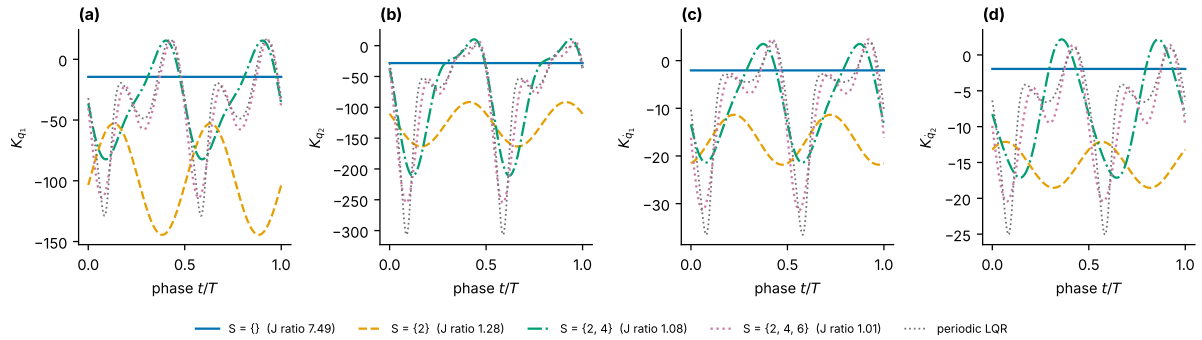

Fig. 10 Gains along the pendubot orbit ($a = 1.2$, $f = 2\,\mathrm{Hz}$) for the four state components (a–d): constant, $\{2\}$, $\{2,4\}$, $\{2,4,6\}$ and periodic LQR.

(6–10 s against 0.7–1.9 s per gain on one CPU core), but this is a one-time, offline cost.

4.5 Input clipping, delay and sampling

Periodic LQR is a robust benchmark, and the one-harmonic gain does not improve on it. With the input clipped at $\pm 7\,\mathrm{Nm}$ on the pendubot orbit ($a = 1.2$, 2 Hz), where the open-loop input already peaks at 6.51 Nm, both gains keep their basins in the slice of Fig. 11: periodic LQR keeps 99 % and the one-harmonic gain 100 % of the area of the unclipped periodic-LQR basin. Clipping is active but brief. Half of the recoverable starts command more than 7 Nm, up to 17 Nm, yet the clipped input sits at the limit for a median of 0.1 % and at most 2.6 % of the run, all at the start, when the error is largest. Even a limit of 6.55 Nm leaves the periodic-LQR basin unchanged. Recovery needs a short burst of input, and the margin left by the open-loop input is large over most of the period.

Table 4 reports linear tolerances to timing errors, sampling and input delay. We expected the gain with fewer harmonics to be less sensitive to timing, since a time shift δ rotates the k-th harmonic of the gain by $k\omega\delta$. The opposite holds on the pendubot: periodic LQR tolerates twice the input delay and half the sampling rate of the one-harmonic gain. On the acrobot and the cart-pole the two are equal. The one-harmonic gain thus matches periodic LQR in cost and basin, and is less tolerant of delay on the hardest orbits. Its value is in showing how little structure is needed, and in replacing the constant gain, not in replacing periodic LQR when the full model and its Riccati solution are available.

5 Phase-indexed feedback

A gain indexed by time assumes that the system is on schedule. Rhythmic systems in biology are usually described instead by a phase that follows the state, as in phase-oscillator models of central pattern generators [61]. Indexing it by an orbit phase φ instead makes the loop autonomous. We parametrise the orbit by its phase, $x_0(\varphi)$, $u_0(\varphi)$, with period 2π in φ, so that $x_0(\omega t)$ is the orbit of

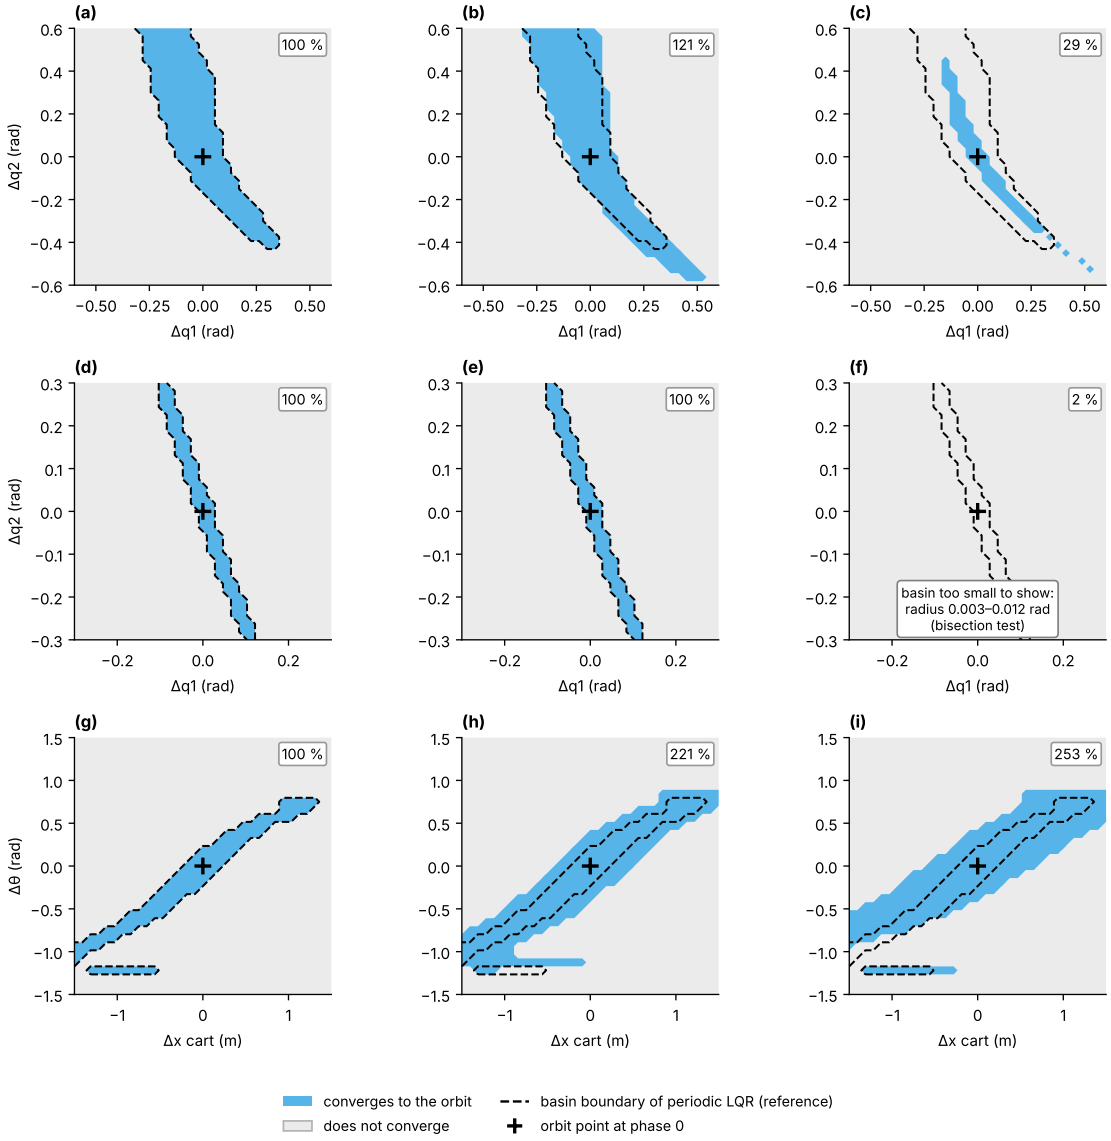

Fig. 11 Basins of attraction of the orbit in nonlinear simulation, for initial deviations of the two coordinates at phase 0 with zero velocity deviation: the two angles for the pendubot and the acrobot, cart position and pendulum angle for the cart-pole. Columns: periodic LQR (a,d,g), the one-harmonic gain {2} (b,e,h) and the constant gain (c,f,i). Blue: initial conditions from which that gain brings the system to the orbit. Grey: initial conditions that do not converge. The dashed line repeats the boundary of the periodic-LQR basin in every panel as a reference. The number in each panel is the basin area relative to periodic LQR, measured within the plotted window. Rows: pendubot ($a = 1.2$, 2 Hz), acrobot ($a = 1.2$, 3 Hz) and cart-pole ($a = 0.8$, 2 Hz).

Section 2. The controller is

$$u = u_0(\varphi) - K(\varphi)\, e, \qquad e = x - x_0(\varphi), \\ \dot{\varphi} = \omega + L(\varphi)^{\top} e, \tag{15}$$

where the phase observer L moves the phase towards the state. Estimating the phase of a rhythmic system from its state is a standard tool in the analysis of oscillators and of animal and robot locomotion [62–64]. Differentiating $x = x_0(\varphi) + e$ gives $\dot{x} = x_0'(\varphi)\dot{\varphi} + \dot{e}$ with $x_0' = dx_0/d\varphi$. On the orbit, $\omega\, x_0'(\varphi) = f(x_0(\varphi), u_0(\varphi))$.

Table 3 Region of attraction (radius over eight random unit directions in state space; minimum / mean) in nonlinear simulation, and implementation size. "pLQR needs": highest harmonic that truncated periodic LQR must keep to be within 5 % of its full cost. Parameters: real numbers that define the gain ($m = 1$, $n = 4$); for periodic LQR, only its nonzero harmonics are counted.

orbit	periodic LQR	low-harmonic	constant	pLQR needs	parameters
pendubot a=1.2, 2 Hz	0.094 / 0.23	0.082 / 0.23	0.035 / 0.08	12ω	12 vs 52
pendubot a=1.2, 1 Hz	0.035 / 0.08	0.035 / 0.08	infeasible	10ω	12 vs 44
pendubot asymmetric	0.070 / 0.34	0.047 / 0.33	infeasible	13ω	20 vs 108
acrobot a=1.2, 2 Hz	0.012 / 0.06	0.012 / 0.06	< 0.01	2ω	12 vs 12
acrobot a=1.2, 3 Hz	0.012 / 0.06	0.012 / 0.06	< 0.01	2ω	12 vs 12
cart-pole a=0.8, 2 Hz	0.89 / 1.41	≥ 1.5	≥ 1.5	2ω	12 vs 12

Table 4 Linear tolerances (true Floquet multipliers). Clock offset: largest constant error δ in the gain's clock, $u = -K(t + \delta)e$, as a fraction of T. Sampling: fewest samples per period with sample-and-hold of e and K. Delay: largest input delay. Constant gains are time invariant and therefore insensitive to clock offsets.

orbit	gain	clock offset	samples/period	delay (ms)
pendubot a=1.2, 2 Hz	periodic LQR	0.015	32	14.6
	one harmonic	0.015	64	7.8
	constant	—	32	12.2
pendubot a=1.2, 1 Hz	periodic LQR	0.015	64	13.7
	one harmonic	0.022	128	8.8
pendubot asymmetric	periodic LQR	0.022	32	13.7
	$\{1, 2\}$	0.015	> 256	0.5
acrobot a=1.2, 2 Hz	periodic LQR	0.152	64	5.4
	one harmonic	0.158	64	5.4
	constant	—	256	1.5
acrobot a=1.2, 3 Hz	periodic LQR	0.165	64	5.5
	one harmonic	0.168	64	5.5
	constant	—	256	0.7
cart-pole a=0.8, 2 Hz	periodic LQR	0.5	8	39.1
	one harmonic	0.5	16	38.1

Expanding f to first order in e,

$$\begin{aligned} \dot{e} &= f\big(x_0 + e,\, u_0 - Ke\big) - x_0'\big(\omega + L^\top e\big) \\ &= \big(A - BK - x_0' L^\top\big)(\varphi)\, e + O(|e|^2), \end{aligned} \tag{16}$$

because the zeroth-order terms cancel. With $\delta = \varphi - \omega t$, $\dot{\delta} = L^\top e$. To first order the coefficient matrix in (16) is evaluated along $\varphi = \omega t + \delta$, a time shift of a periodic matrix, which does not change its Floquet multipliers. Orbital stability therefore reduces to stability of the n-dimensional periodic system $\dot{e} = (A - BK - x_0' L^\top)\, e$. When e decays exponentially, δ converges to a constant: the system returns to the orbit with a phase offset. With $L = 0$ the time-indexed loop is recovered and δ stays at its initial value. We design K with harmonics $\{2\}$ and L with harmonics $\{1, 2\}$ by the same method, adding a term $r_L L^\top L$ to the cost. Periodic LQR offers no counterpart for L, so this design is only possible with a structured method.

The cost is now evaluated on $e = x - x_0(\varphi)$ rather than on $x - x_0(t)$, so ratios to the periodic-LQR cost below one are possible.

Table 5 compares the two forms. The phase-indexed loop has a lower cost on e (it can absorb part of the error by retiming) and a 30–45 % smaller distance to the orbit under timing errors. Its effect on robustness to large errors is mixed. On the pendubot it roughly halves the region of attraction and lowers the tolerated initial phase offset. On the acrobot both grow slightly. In both cases the tolerated phase offset stays below a tenth of a period.

6 Discussion and conclusions

The answer to our question is that feedback around these unstable orbits needs little time variation, but the right kind. One harmonic at twice the orbit frequency recovers almost all of the

Table 5 Time-indexed and phase-indexed low-harmonic feedback (nonlinear simulation). Jitter: smooth timing error of amplitude $0.03T$ in the controller's clock; RMS distance to the orbit. RoA: radius of attraction over eight random directions (minimum/mean), the same test as Table 3. Orbits: pendubot $a = 1.2$, 2 Hz; acrobot $a = 1.2$, 3 Hz.

orbit	feedback	J/J_{pLQR}	phase offset tolerated	jitter: distance	RoA
pendubot	time, $K(t)$	1.28	$0.027\,T$	0.275	0.082/0.227
pendubot	phase, (K, L)	0.87	$0.021\,T$	0.153	0.047/0.120
acrobot	time, $K(t)$	1.00	$0.066\,T$	0.163	0.012/0.062
acrobot	phase, (K, L)	0.66	$0.084\,T$	0.112	0.012/0.070

performance of the optimal periodic gain on the symmetric orbits. Which harmonic is needed is not a tuning choice: it follows from the symmetry of the orbit, which makes the linearisation repeat every half period. This is in the spirit of vibrational control, where a fast oscillation changes stability through its interaction with the slow dynamics [15, 17]. Here the oscillation belongs to the plant, and the feedback needs to match only its leading parametric component.

The failure of averaging has a simple form. Averaging is accurate for the open loop of every orbit we studied, so a designer who checks the averaged model against the true Floquet exponents will trust it. It fails once feedback is added, because the products of the oscillating parts of $B(t)$ and $A(t)$ with the gain do not average out. The frozen-time loop is unstable for most of the period, and the averaged design cannot see this. A practical check is therefore to evaluate any averaged design on the true monodromy, not on the averaged model. A criterion that predicts the failure remains open. In our tests the Magnus convergence bound never applied, and truncating the expansion at second order overestimated the instability by more than an order of magnitude (Section 2).

Truncation fails for a related reason. A truncated lifted model is a finite approximation of an operator whose spectrum has a periodic structure, and the truncation adds eigenvalues at the edge of the band. An optimiser that sees only the truncated model finds gains that exploit these artefacts. The remedy we used, a time-domain cost and a barrier on the true monodromy, is simple and generic. It is also a design method in its own right: it optimises any structured periodic gain against the true Floquet multipliers, and it can be used for harmonic-domain designs in other settings as well.

The low-harmonic gain sits between two established approaches. Unlike harmonic-domain control [28, 30], it targets strongly nonlinear systems with unstable orbits and does not depend on a truncation order. Compared with orbital stabilisation by transverse linearisation [36, 38], it gives up optimality in exchange for a gain that is described by a few numbers and a single sinusoid. A direct comparison is beyond the scope of this paper. Transverse linearisation solves a related but different problem: it stabilises the orbit as a set, through $n - 1$ transverse coordinates and a projection onto the orbit, and its gain comes from a periodic Riccati equation, so it is fully time varying by construction. Our question concerns how much time variation a gain needs for a given timing of the orbit, and periodic LQR is the matching benchmark. Asking the same question inside the transverse framework, with low-harmonic transverse gains, is a natural extension; the phase-indexed controller of Section 5 is a first step in that direction.

Periodic LQR remains the reference. It is optimal by construction, and our tests show that it is also robust: it keeps its basin under input clipping and tolerates more delay than the one-harmonic gain on the pendubot. We do not propose to replace it when the full periodic model and its Riccati solution are available. The comparison that matters in practice is with the constant gain, which is what most implementations and laboratory courses use. Constant gains designed on the averaged model can be unstable, and constant gains designed on the true model can have almost no basin. A single harmonic removes both problems. It can be written down, plotted and adjusted by hand, which makes it a natural next step for embedded controllers and for teaching. A course can use these setups to show, on hardware, that an unstable orbit is stabilised by feedback with a single harmonic, that a controller designed on the

averaged model can fail, and that the orbit's symmetry decides which harmonic is needed. We have not yet tested this on hardware.

Phase indexing changes the nature of the loop. With a phase, the controller follows the system instead of a clock, and timing errors are absorbed rather than fought. Timing noise is known to change the optimal feedback itself [65]; a phase-indexed gain is one way to make the loop insensitive to it. On our orbits this reduced the deviation under timing jitter but did not consistently enlarge the basin.

The design needs only the periodic matrices $A(t)$ and $B(t)$. These can also be estimated from data through harmonic transfer functions [18–20], as has been done for human walking and muscle [21, 22]. Because a low-harmonic gain has few parameters, it could also be tuned from such models or directly from closed-loop experiments, where a full periodic gain would be hard to identify. We have not tested this.

The study has limits. All results are numerical, on three planar systems with one input. We did not study model mismatch or measurement noise beyond timing errors, delay and sampling. Input limits that bind for long periods, as in commanded changes of the orbit's timing, call for constraint-aware schemes that are beyond the scope of this paper.

In summary, unstable periodic orbits of strongly nonlinear underactuated systems can be stabilised with feedback that contains one or two harmonics, provided that the feedback is designed against the true Floquet multipliers. Averaging and truncation, the two usual shortcuts, mislead in closed loop although averaging is accurate in open loop. The symmetry of the orbit determines which harmonics are needed. Where a constant gain is not enough, as on the two-link orbits, a single harmonic closes the gap to periodic LQR in cost and basin of attraction; on the mild cart-pole orbit a constant gain already suffices, while periodic LQR remains the more tolerant of delay on the hardest orbits. Indexing the feedback by phase makes the loop autonomous and reduces the effect of timing errors.

Statements and Declarations

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

Data availability. The code and data that support the findings of this study will be made available later.

References

[1] M. W. Spong, The swing up control problem for the Acrobot, IEEE Control Syst. Mag. 15 (1995) 49–55.

[2] K. J. Åström, K. Furuta, Swinging up a pendulum by energy control, Automatica 36 (2000) 287–295.

[3] M. M. Ankaralı, U. Saranlı, Stride-to-stride energy regulation for robust self-stability of a torque-actuated dissipative spring-mass hopper, Chaos 20 (2010) 033121.

[4] T. McGeer, Passive dynamic walking, Int. J. Robot. Res. 9 (1990) 62–82.

[5] J. W. Grizzle, G. Abba, F. Plestan, Asymptotically stable walking for biped robots: analysis via systems with impulse effects, IEEE Trans. Autom. Control 46 (2001) 51–64.

[6] E. R. Westervelt, J. W. Grizzle, D. E. Koditschek, Hybrid zero dynamics of planar biped walkers, IEEE Trans. Autom. Control 48 (2003) 42–56.

[7] M. M. Ankaralı, H. T. Sen, A. De, A. M. Okamura, N. J. Cowan, Haptic feedback enhances rhythmic motor control by reducing variability, not improving convergence rate, J. Neurophysiol. 111 (2014) 1286–1299.

[8] M. M. Ankaralı, S. Sefati, M. S. Madhav, A. Long, A. J. Bastian, N. J. Cowan, Walking dynamics are symmetric (enough), J. R. Soc. Interface 12 (2015) 20150209.

[9] D. Logan, Y. P. Ivanenko, T. Kiemel, G. Cappellini, F. Sylos-Labini, F. Lacquaniti, J. J.

Jeka, Function dictates the phase dependence of vision during human locomotion, J. Neurophysiol. 112 (2014) 165–180.

[10] M. W. Spong, Underactuated mechanical systems, in: B. Siciliano, K. P. Valavanis (Eds.), Control Problems in Robotics and Automation, Lecture Notes in Control and Information Sciences 230, Springer, London, 1998, pp. 135–150.

[11] I. Fantoni, R. Lozano, Non-linear Control for Underactuated Mechanical Systems, Springer, London, 2002.

[12] E. Ott, C. Grebogi, J. A. Yorke, Controlling chaos, Phys. Rev. Lett. 64 (1990) 1196–1199.

[13] K. Pyragas, Continuous control of chaos by self-controlling feedback, Phys. Lett. A 170 (1992) 421–428.

[14] P. L. Kapitza, Dynamic stability of a pendulum when its point of suspension vibrates, Sov. Phys. JETP 21 (1951) 588–592.

[15] S. M. Meerkov, Principle of vibrational control: theory and applications, IEEE Trans. Autom. Control 25 (1980) 755–762.

[16] R. Bellman, J. Bentsman, S. M. Meerkov, Vibrational control of nonlinear systems: vibrational stabilizability, IEEE Trans. Autom. Control 31 (1986) 710–716.

[17] F. Bullo, Averaging and vibrational control of mechanical systems, SIAM J. Control Optim. 41 (2002) 542–562.

[18] M. M. Ankaralı, N. J. Cowan, System identification of rhythmic hybrid dynamical systems via discrete time harmonic transfer functions, in: Proc. 53rd IEEE Conf. Decision and Control, 2014, pp. 1017–1022.

[19] İ. Uyanık, M. M. Ankaralı, N. J. Cowan, U. Saranlı, Ö. Morgül, Identification of a vertical hopping robot model via harmonic transfer functions, Trans. Inst. Meas. Control 38 (2016) 501–511.

[20] İ. Uyanık, U. Saranlı, M. M. Ankaralı, N. J. Cowan, Ö. Morgül, Frequency-domain subspace identification of linear time-periodic (LTP) systems, IEEE Trans. Autom. Control 64 (2019) 2529–2536.

[21] D. Logan, T. Kiemel, J. J. Jeka, Using a system identification approach to investigate subtask control during human locomotion, Front. Comput. Neurosci. 10 (2017) 146.

[22] E. D. Tytell, J. A. Carr, N. Danos, C. Wagenbach, C. M. Sullivan, T. Kiemel, N. J. Cowan, M. M. Ankaralı, Body stiffness and damping depend sensitively on the timing of muscle activation in lampreys, Integr. Comp. Biol. 58 (2018) 860–873.

[23] N. M. Wereley, Analysis and control of linear periodically time varying systems, PhD thesis, MIT, 1990.

[24] N. M. Wereley, S. R. Hall, Frequency response of linear time periodic systems, in: Proc. 29th IEEE Conf. Decision and Control, 1990, pp. 3650–3655.

[25] E. Möllerstedt, B. Bernhardsson, Out of control because of harmonics: an analysis of the harmonic response of an inverter locomotive, IEEE Control Syst. Mag. 20 (2000) 70–81.

[26] J. Zhou, T. Hagiwara, H_2 and H_∞ norm computations of linear continuous-time periodic systems via the skew analysis of frequency response operators, Automatica 38 (2002) 1381–1387.

[27] H. Sandberg, E. Möllerstedt, B. Bernhardsson, Frequency-domain analysis of linear time-periodic systems, IEEE Trans. Autom. Control 50 (2005) 1971–1983.

[28] N. Blin, P. Riedinger, J. Daafouz, L. Grimaud, P. Feyel, Necessary and sufficient conditions for harmonic control in continuous time, IEEE Trans. Autom. Control 67 (2022) 4013–4028.

[29] P. Riedinger, J. Daafouz, Solving infinite-dimensional harmonic Lyapunov and Riccati equations, arXiv:2203.09774 (2022).

[30] F. Vernerey, P. Riedinger, J. Daafouz, A TBLMI framework for harmonic robust control, IEEE Trans. Autom. Control 70 (2025) 2537–2551.

[31] E. Hidir, İ. Uyanık, Ö. Morgül, Harmonic transfer functions based controllers for linear time-periodic systems, Trans. Inst. Meas. Control 41 (2019) 2171–2184.

[32] B. Sert, B. Çatalbaş, İ. Uyanık, Identification and control of a linear time-periodic test bench using harmonic transfer functions, in: Proc. 9th Int. Conf. Control, Decision and Information Technologies (CoDIT), 2023, pp. 2366–2371.

[33] M. Grosso, P. Riedinger, J. Daafouz, S. Pierfederici, H. Janati-Idrissi, B. Lapôtre, Harmonic control of three-phase AC/DC converter with time-domain guarantees, IEEE Trans. Control Syst. Technol. 33 (2025) 887–902.

[34] A. Shiriaev, J. W. Perram, C. Canudas-de-Wit, Constructive tool for orbital stabilization of underactuated nonlinear systems: virtual constraints approach, IEEE Trans. Autom. Control 50 (2005) 1164–1176.

[35] L. Freidovich, A. Robertsson, A. Shiriaev, R. Johansson, Periodic motions of the Pendubot via virtual holonomic constraints: theory and experiments, Automatica 44 (2008) 785–791.

[36] A. S. Shiriaev, L. B. Freidovich, S. V. Gusev, Transverse linearization for controlled mechanical systems with several passive degrees of freedom, IEEE Trans. Autom. Control 55 (2010) 893–906.

[37] I. R. Manchester, U. Mettin, F. Iida, R. Tedrake, Stable dynamic walking over rough terrain, in: Robotics Research, Springer Tracts Adv. Robot. 70 (2011).

[38] C. F. Sætre, A. S. Shiriaev, L. B. Freidovich, S. V. Gusev, L. M. Fridman, Robust orbital stabilization: a Floquet theory-based approach, Int. J. Robust Nonlinear Control 31 (2021) 8075–8108.

[39] M. Surov, Orbital stabilization and time synchronization of unstable periodic motions in underactuated robots, arXiv:2509.20082 (2025).

[40] A. Varga, S. Pieters, Gradient-based approach to solve optimal periodic output feedback control problems, Automatica 34 (1998) 477–481.

[41] L. Viganò, M. Bergamasco, M. Lovera, A. Varga, Optimal periodic output feedback control: a continuous-time approach and a case study, Int. J. Control 83 (2010) 897–914.

[42] A. D. Duyul, Control of an underactuated system around a periodic orbit, MSc thesis, Middle East Technical University, 2018.

[43] M. W. Spong, D. J. Block, The Pendubot: a mechatronic system for control research and education, in: Proc. 34th IEEE Conf. Decision and Control, 1995, pp. 555–556.

[44] M. S. Nakhla, J. Vlach, A piecewise harmonic balance technique for determination of periodic response of nonlinear systems, IEEE Trans. Circuits Syst. 23 (1976) 85–91.

[45] M. Krack, J. Gross, Harmonic Balance for Nonlinear Vibration Problems, Springer, Cham, 2019.

[46] T. Detroux, L. Renson, L. Masset, G. Kerschen, The harmonic balance method for bifurcation analysis of large-scale nonlinear mechanical systems, Comput. Methods Appl. Mech. Eng. 296 (2015) 18–38.

[47] S. Bittanti, P. Colaneri, Periodic Systems: Filtering and Control, Springer, London, 2009.

[48] G. Floquet, Sur les équations différentielles linéaires à coefficients périodiques, Ann. Sci. Éc. Norm. Supér. 12 (1883) 47–88.

[49] V. A. Yakubovich, V. M. Starzhinskii, Linear Differential Equations with Periodic Coefficients, Wiley, New York, 1975.

[50] J. A. Richards, Analysis of Periodically Time-Varying Systems, Springer, Berlin, 1983.

[51] L. Markus, H. Yamabe, Global stability criteria for differential systems, Osaka Math. J. 12 (1960) 305–317.

[52] J. A. Sanders, F. Verhulst, J. Murdock, Averaging Methods in Nonlinear Dynamical Systems, 2nd ed., Springer, New York, 2007.

[53] W. Magnus, On the exponential solution of differential equations for a linear operator, Commun. Pure Appl. Math. 7 (1954) 649–673.

[54] S. Blanes, F. Casas, J. A. Oteo, J. Ros, The Magnus expansion and some of its applications, Phys. Rep. 470 (2009) 151–238.

[55] I. Kovacic, R. Rand, S. M. Sah, Mathieu's equation and its generalizations: overview of stability charts and their features, Appl. Mech. Rev. 70 (2018) 020802.

[56] S. Bittanti, P. Colaneri, G. De Nicolao, The periodic Riccati equation, in: S. Bittanti, A. J. Laub, J. C. Willems (Eds.), The Riccati Equation, Springer, Berlin, 1991, pp. 127–162.

[57] J. J. Hench, A. J. Laub, Numerical solution of the discrete-time periodic Riccati equation, IEEE Trans. Autom. Control 39 (1994) 1197–1210.

[58] A. Varga, On solving periodic Riccati equations, Numer. Linear Algebra Appl. 15 (2008) 809–835.

[59] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of Python+NumPy programs, software, 2018, https://github.com/jax-ml/jax.

[60] G. Chesi, Domain of Attraction: Analysis and Control via SOS Programming, Springer, London, 2011.

[61] P. L. Várkonyi, T. Kiemel, K. Hoffman, A. H. Cohen, P. Holmes, On the derivation and tuning of phase oscillator models for lamprey central pattern generators, J. Comput. Neurosci. 25 (2008) 245–261.

[62] S. Revzen, J. M. Guckenheimer, Estimating the phase of synchronized oscillators, Phys. Rev. E 78 (2008) 051907.

[63] S. Revzen, J. M. Guckenheimer, Finding the dimension of slow dynamics in a rhythmic system, J. R. Soc. Interface 9 (2012) 957–971.

[64] H. Nakao, Phase reduction approach to synchronisation of nonlinear oscillators, Contemp. Phys. 57 (2016) 188–214.

[65] A. Lamperski, N. J. Cowan, Optimal control with noisy time, IEEE Trans. Autom. Control 61 (2016) 319–333.

[66] J. Nocedal, S. J. Wright, Numerical Optimization, 2nd ed., Springer, New York, 2006.

Appendix A Truncation-free cost and Floquet barrier

The closed-loop matrix $A_{\mathrm{cl}}(t) = A(t) - B(t)K(t)$ is sampled at the midpoints of $N = 512$ intervals of length $\Delta = T/N$. Each interval uses the transition $E_k = \exp(A_{\mathrm{cl}}(t_k)\Delta)$, computed as the square of $\exp(A_{\mathrm{cl}}(t_k)\Delta/2)$. The monodromy is $\Phi_T = E_{N-1}\cdots E_0$. The covariance obeys $\dot{S} = A_{\mathrm{cl}}S + SA_{\mathrm{cl}}^{\top} + I$. Over one interval $S_{k+1} = E_k S_k E_k^{\top} + N_k$ with $N_k = \Delta\, E_k^{1/2} E_k^{1/2\top}$. The periodic solution follows from $S_0 = \Phi_T S_0 \Phi_T^{\top} + \sum_k (\cdots)$, solved as a linear system of size n^2, and the cost (10) is evaluated at the interval midpoints. All operations are differentiable, and the gradient of J with respect to the gain parameters is obtained by reverse-mode automatic differentiation.

For the barrier, let μ be the multiplier of largest modulus with right and left eigenvectors v and w. Then

$$\partial\rho = \mathrm{Re}\big(\bar{\mu}\, w^H\, \partial\Phi_T\, v\big) \big/ \big(|\mu|\, w^H v\big), \tag{A1}$$

and $w^H \partial\Phi_T v$ is computed by a vector–Jacobian product through the same chain of transitions. We use $\sigma = 0.5\,\mathrm{s}^{-1}$ and $\mu = 10^{-2}$ in (14), minimise with BFGS [66] in scaled coordinates, and start from the averaged-LQR gain or, if that is unstable, from a first phase that minimises $\log\rho$ alone.

Appendix B Model parameters

Cart-pole (Quanser IP02-type, input: motor voltage): cart mass 0.57 kg, pendulum mass 0.21 kg, pendulum half-length 0.3 m (uniform rod), cart viscous friction 4.3 N s/m, pendulum viscous friction 0.05 N m s/rad, gearbox ratio 3.71, motor constants $K_t = K_m = 7.67 \times 10^{-3}$, armature resistance 2.6 Ω, pinion radius 6.35 mm. The optional torsional spring is $K_s = 0$ except in the weakly stable example of Section 3 ($K_s = 0.74$).

Pendubot and acrobot (two-link arm, angles from the upright): $m_1 = 0.8\,\mathrm{kg}$, $m_2 = 0.4\,\mathrm{kg}$, $l_1 = 0.2\,\mathrm{m}$, $l_{c1} = 0.1\,\mathrm{m}$, $l_{c2} = 0.15\,\mathrm{m}$, $I_1 = 0.0027\,\mathrm{kg\,m^2}$, $I_2 = 0.003\,\mathrm{kg\,m^2}$, joint damping $b_1 = 0.005$ and $b_2 = 0.002\,\mathrm{N\,m\,s/rad}$, $g = 9.81\,\mathrm{m/s^2}$.

Weights: $Q = \mathrm{diag}(10, 10, 1, 1)$; $R = 0.1$ for the cart-pole and $R = 1$ for the two-link systems.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.mp4