HHS Public Access

Author manuscript

Automatica (Oxf). Author manuscript; available in PMC 2020 February 01.

Published in final edited form as:

Automatica (Oxf). 2019 February ; 100: 336–348. doi:10.1016/j.automatica.2018.11.012.

Pharmaceutical-based entrainment of circadian phase via nonlinear model predictive control

John H. Abel[a,b,e], **Ankush Chakrabarty**[c], **Elizabeth B. Klerman**[b,d], **Francis J. Doyle III**[b,c,*]

[a]Department of Systems Biology, Harvard Medical School, Boston, MA 02115, USA.

[b]Division of Sleep Medicine, Harvard Medical School, Boston, MA 02115, USA.

[c]Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA.

[d]Division of Sleep and Circadian Disorders, Department of Medicine, Brigham and Women's Hospital, Boston, MA 02115, USA.

[e]Present address: Department of Anesthesiology, Critical Care and Pain Medicine, Massachusetts General Hospital, Boston, MA 02114; Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, MA 02139.

Abstract

The widespread adoption of closed-loop control in systems biology has resulted from improvements in sensors, computing, actuation, and the discovery of alternative sites of targeted drug delivery. Most control algorithms for circadian phase resetting exploit light inputs. However, recently identified small-molecule pharmaceuticals offer advantages in terms of invasiveness and potency of actuation. Herein, we develop a systematic method to control the phase of biological oscillations motivated by the recently identified small molecule circadian pharmaceutical KL001. The model-based control architecture exploits an infinitesimal parametric phase response curve (ipPRC) that is used to predict the effect of control inputs on future phase trajectories of the oscillator. The continuous time optimal control policy is first derived for phase resetting, based on the ipPRC and Pontryagin's maximum principle. Owing to practical challenges in implementing a continuous time optimal control policy, we investigate the effect of implementing the continuous time policy in a sampled time format. Specifically, we provide bounds on the errors incurred by the physiologically tractable sampled time control law. We use these results to select directions of resetting (i.e. phase advance or delay), sampling intervals, and prediction horizons for a nonlinear model predictive control (MPC) algorithm for phase resetting. The potential of this ipPRC-informed pharmaceutical nonlinear MPC is then demonstrated *in silico* using real-world scenarios of jet lag or rotating shift work.

* Corresponding author. johnhabel@g.harvard.edu (John H. Abel), achakrabarty@seas.harvard.edu (Ankush Chakrabarty), ebklerman@research.bwh.harvard.edu (Elizabeth B. Klerman), frank_doyle@seas.harvard.edu (Francis J. Doyle III).

Keywords

Phase response curve; nonlinear dynamics; circadian oscillators; systems biology; optimal control; model predictive control; phase; entrainment

1. Introduction

A coalescing of control engineering and systems biology has resulted in widespread use of open-loop and feedback control approaches in biological and biomedical applications [1, 2]. These include identifying and modulating cellular behavior [3–5], treating diseases [6, 7], constructing synthetic biological circuits [8, 9], optimizing biomanufacturing productivity [10, 11], and formulating targeted drug delivery systems [12–15]. Unlike open-loop systems that require manual intervention at critical times to prevent deleterious outcomes, closed-loop drug delivery systems enable effective regulation of targeted biological pathways by leveraging control-relevant models, systematic prediction, or decision-making based on clinical targets. Therefore, these approaches have found use in manufacturing and medical devices. Additional improvements in the quality of closed-loop drug delivery and adherence within society has been fueled by the invention of wearable sensors [16], minimally invasive actuators, and embedded decision-making platforms [17], along with novel drug delivery mechanisms [18] or input-output pairs [19]. This paper demonstrates how selecting control input (from light-based methods to small-molecule pharmaceuticals) or controller design parameters plays an essential role in resetting the phase of mammalian circadian rhythms.

Circadian rhythms are endogenous daily oscillations in gene expression or metabolism driving temporal adaptations in most organisms. In mammals, these oscillations are generated by genetic feedback loops within each cell of the organism [20]. Since environmental signals (such as ambient light) set the time of this biological clock, mistimed environmental cues may result in adverse changes to these rhythms, and consequently, a loss of temporal regulation of the genetic architecture, and psychological and physiological pathologies [21]. In recent years, small-molecule pharmaceuticals have gained significant interest as a path toward modulating the circadian clock to reduce the effects of circadian disturbances [22, 23]. Small-molecule pharmaceuticals present critical benefits over the use of light for clock resetting including avoiding the day-time “dead-zone” where light evokes a minimal phase response, and reducing the impractical and burdensome challenge of attempting to tightly control one’s light environment. For example, light-based regulation would necessitate sustained periods of time wearing low-transmission glasses or light visors [24]. Most importantly, light dosing strategies typically take several days to reset circadian rhythms: hence the prevalence of sustained jet lag following long trips across time zones. Even state-of-the-art optimal control policies involving light necessitate as much as eight days to complete a phase resetting when allowing for regular sleep timing [25].

Pharmaceuticals are expected to enable direct manipulation of sensitive control targets within the clock, allowing more rapid clock resetting [26]. The complex nature of the circadian oscillator necessitates a control approach toward dosing strategies for these drugs, as identical stimuli applied at different times of the day may have drastically disparate

effects on the clock. Modeling approaches such as those described in [27] have been employed to identify the underlying mechanistic action of these drugs. Such models may, in turn, be used to inform advanced model-based control strategies such as model predictive control (MPC) for circadian regulation.

Prior studies employing models of circadian dynamics have independently investigated optimal control-based approaches for manipulating the circadian clock. For example, nonlinear MPC and multi-target MPC for light-resetting a *Drosophila* circadian clock model was demonstrated in [26, 28], and complementary phase response curves (PRCs) were leveraged to identify potential control targets. Other studies using nonlinear *Drosophila* clock models [29, 30] proposed more efficient procedures to obtain light-based optimal control trajectories by relaxing the nonlinear control problem into a mixed-integer formulation. Recent work has extended the concept of using light-based feedback control to investigate re-entrainment of clocks in humans. Attempts have been made to pose this complex problem as a scheduling problem to obtain the best bang-off-bang strategy [24] for light inputs. Recently, the authors in [31] used lower harmonics to formulate an approximate model of the circadian oscillations in humans, and exploited the structure of the co-state equation to propose an efficient line-search algorithm to obtain optimal control sequences. Although these methods are elegant and have been demonstrated to be effective via numerous simulation studies, a significant drawback of optimal control approaches is that they are rarely immune to measurement inaccuracies, exogenous noise, or plant-model mismatch (i.e. discrepancies between the system under control and the model describing it). Further-more, without model approximation, deriving optimal control sequences remain computationally prohibitive, if not intractable, for complex high-dimensional nonlinear models [30].

The use of feedback in circadian entrainment remains a relatively unexplored problem. A notable exception is [32], where the authors developed a framework for control of circadian rhythms using short-duration pulses to shift the phase, via proportional-integral-derivative (PID) control by exploiting a low-dimensional approximate predictive model of the phase dynamics, known as a *phase response curve* (PRC). This low-dimensional representation of pertinent phase information and the predictive power of a ipPRC enables the construction of MPC algorithms that are well-known to be inherently robust to noise and plant-model mismatch. A prominent recent study derived a feedback control law for circadian entrainment using lookup tables from the optimal control [33]. This study focused on a light input which may be switched on and off at will, and although highly effective for light, pharmaceutical delivery under this policy would incur errors. More importantly, the main result proven herein bounds the nearness of their feedback control approach to their optimal control law. A preliminary investigation of the potential of MPC in circadian control with pharmaceuticals was undertaken in [34, 35], which demonstrated MPC of mammalian circadian rhythms via KL001, a small-molecule pharmaceutical that was recently used to manipulate (*in vivo*) the circadian clock in mammals [22, 23]. Although previous investigations have revealed the importance of carefully selecting prediction horizon and controller sampling time of the MPC to reset the circadian oscillator, no systematic method of selecting these design variables has been derived.

This study seeks to answer three fundamental questions:

1. What does the optimal control policy for pharmaceutical-driven circadian phase resetting look like?
2. What design considerations should be made in implementing feedback control of circadian phase? More specifically, how does the selection of sampling or switching times and the prediction horizon affect the optimality of the applied control?
3. How can these design considerations shape our selection of circadian therapeutics?

In answering these questions, this study presents a systematic framework for constructing ipPRC-based nonlinear MPC formalisms capable of manipulating mammalian clocks using small-molecule pharmaceuticals.

2. Phase-Reduced Model and Optimal Phase Shifting

First, the circadian oscillator model and its accompanying notation are presented.

2.1. Circadian oscillator model

Circadian dynamics are most commonly represented by limit-cycle oscillators. These oscillators are generally modeled using a class of smooth nonlinear dynamical systems of the form

$$\frac{dx}{dt} = f(x,\ p,\ u), \tag{1}$$

where $x(t) \in \mathbb{R}^{\mathrm{n}}$ denote states such as mRNA or protein concentrations, $\mathrm{p} \in \mathbb{R}^{m}$ denote kinetic parameters, and $\mathrm{u}(t) \in \mathbb{R}^{\ell}$ denote control inputs. The zero-input system $\dot{x} = f(x,p,0)$ comprises an exponentially-attractive limit cycle $\Gamma \subset \mathbb{R}^{\mathrm{n}}$ that satisfies

$$\lim_{t \to \infty} [x(t) - x(t - T)] = 0, \tag{2}$$

with period T and angular frequency $\omega = 2\pi/T$. A two-dimensional limit cycle with states $x = [x_1, x_2]$ is visualized in Fig. 1A.

This study is concerned with dynamics of the phase of oscillation, and so we first map every unique point on the limit cycle $x_0 \in \Gamma$ to a unique scalar phase $\phi_0 \in \mathbb{S}^1 = [0,\ 2\pi)$. Phase is defined such that for an oscillator on the limit cycle, phase advances linearly with time. Let $\Phi_\Gamma : \Gamma \to \mathbb{S}^1$ denote the corresponding map relating x_0 to oscillator phase $\phi_0 = \Phi_\Gamma(x_0)$.

Let $\gamma(t; x_0)$ denote the solution of system (1) with initial condition $x(0) = x_0 \in \Gamma$. Thus a time-dependent phase variable

$$\phi_\Gamma(t) = \Phi_\Gamma\left[\gamma(t, x_0)\right] \tag{3}$$

may be established. Considering that the unperturbed oscillator traverses the limit cycle at the constant rate ω, the system phase evolves in time according to:

$$\phi_\Gamma(t) = \omega t + \phi_0 \bmod 2\pi\,. \tag{4}$$

The application of an exogenous input (such as a control action) will divert the state of an oscillatory system away from the limit cycle Γ. In order to ascertain the phase shift incurred by such an input, we must assign a phase to points on $\mathbb{R}^n \backslash \Gamma$ that are not on the limit cycle Γ, that are exponentially attracted to Γ. This can be done by assigning them the phase of the trajectory to which they ultimately converge. The trajectory denoted $\theta(t; x_a, u)$ is the time evolution of point $x_a \notin \Gamma$, with dynamics given by (1) and control input u. For the zero-input case ($u = 0$), the state x_a may be assigned the phase of the initial condition of the trajectory $\gamma(t; x_0)$ to which the trajectory $\theta(t; x_a, 0)$ ultimately converges. The asymptotic phase of x_a, denoted ϕ_a, is equal to the phase ϕ_0 of point $x_0 \in \Gamma$, subject to

$$0 = \lim_{t \to \infty} \| \theta(t, x_a, 0) - \gamma(t, x_0) \| \,. \tag{5}$$

This relationship yields the asymptotic phase map $\phi_a = \Phi(x_a)$. The asymptotic phase map may be used to formulate a time-dependent asymptotic phase variable

$$\phi(t, 0) = \Phi\left[\theta(t, x_a, 0)\right] \tag{6}$$

for the zero-input case. The trajectories γ and θ with identical phase, and their asymptotic convergence, are shown schematically in Fig. 1B.

This formulation yields phase dynamics identical to the case for ϕ, with particular solutions of the form $\phi(t; 0) = \omega t + \phi_0$ with initial phase ϕ_0. The following subsection describes how the concept of an asymptotic phase variable can be extended for the more challenging and useful case of non-zero control input u.

2.2. Infinitesimal parametric phase response curves and the phase-reduced model

For an oscillator in the neighborhood of Γ, the phase dynamics resulting from a control input may be derived by the chain rule:

$$\frac{d\phi(t, u)}{dt} = \omega + \frac{\partial}{\partial t} \frac{\partial \Phi\left[\gamma(t, x_0)\right]}{\partial u} u(t) \tag{7}$$

which yields a first-order approximation of phase dynamics for a nonzero control input.

Pharmaceutical manipulations of the clock are generally considered to be mediated by temporary changes in one or more parameter values (e.g. increased degradation of a protein, or decreased transcriptional rate) [23, 36]. Input $u(t)$ is incorporated into the ODE model as a time-dependent modification of parameters p

$$p(t) = p_0 + u(t). \tag{8}$$

Substituting this into (7) yields

$$\frac{d\phi}{dt} = \omega + \frac{\partial}{\partial t}\frac{\partial \Phi[\gamma(t, x_0)]}{\partial p} u(t), \tag{9}$$

where $\frac{\partial}{\partial t}\frac{\partial \Phi[\gamma]}{\partial p}$ is the infinitesimal parametric PRC (ipPRC) for the oscillator on the limit cycle [37]. The formulation (9) is valid up to the first-order approximation in both duration and magnitude of the control, and is referred to as the *phase-reduced model*.

The one-dimensional phase equation is therefore given by:

$$\frac{d\phi}{dt} = \omega + B(t, \phi) \cdot u(t), \tag{10}$$

where $B(t; \phi)$ is the ipPRC and $u(t)$ is the parametric perturbation.

Assumption 1.—The ipPRC B(·) is fully known and dependent only on the asymptotic phase ϕ(t).

The ipPRC can be calculated numerically from mechanistic models of the clock for a known control input as above, or alternately could be characterized experimentally [37, 38]. Since the ipPRC is generated from a clock model, it may exhibit minor differences with the ipPRC of the mammalian circadian system. This contributed to our decision to later apply control via MPC, as it is inherently robust to plant-model mismatch. The ipPRC may be quantified directly via impulse response tests, though these experiments are challenging to perform.

A prior method of formulating asymptotic phase dynamics assumed short pulses of control inputs that relax completely to the limit cycle after each pulse [32], and thus the phase sensitivity on the limit cycle may be used. Assumption 1, rather, implies that the phase response dynamics are constant along isochrons (state-space regions of constant phase). This assumption is violated as the oscillator becomes more distant from its limit cycle, or closer to the fixed point at the center of the limit cycle. Identifying violations of this assumption may be studied in the future, and in these cases, the model should be simulated in full. This assumption has been used in prior studies as well [25, 33]. Our assumption yields a few practical advantages when implementing pharmaceutical inputs. Short pulses are less able to evoke significant phase shifts, as the realized phase shift is determined by integration of the ipPRC. Thus, the ipPRC must be of extreme magnitude to yield appreciable shifts for short pulses. Pharmaceutical inputs to the clock are not likely to have the extremely rapid pharmacokinetics needed to be approximated as short pulses.

Assumption 2.—The ipPRC map is continuous on [0, 2π] and has exactly two zero crosses in that interval; the derivatives of the ipPRC at the zero crosses have opposite sign.

This PRC form is readily found in the literature, and for stimuli of a finite duration is commonly referred to as a type 1 phase response [38, 39]. Methods for calculation of the ipPRC using limit-cycle oscillator models are described elsewhere [37]. The other phase response observed (solely) for finite-duration stimuli, a type 0 phase response, contains a discontinuity. This form is associated with a strong and extended stimulus or multiple stimuli over several days, and therefore does not correspond to the ipPRC [38–40]. Furthermore, a type 0 phase response may be generated from an ipPRC that conforms to Assumption 2 by applying a long-duration stimulus.

2.3. The phase resetting problem and optimal phase resetting via Pontryagin's maximum principle

The phase-resetting problem involves tracking a reference oscillator with phase given by $\phi_r(t) \in [0, 2\pi)$, with the oscillator $\phi(t; u)$ which obeys (10), and with control input bounded by $u \in [u_{\min}, u_{\max}]$. The phase of the reference oscillator is given by:

$$\phi_r(t) = \omega_r t + \phi_r(0) \mod 2\pi.$$

Assumption 3.—The reference angular velocity is identical to the angular velocity of the oscillator, that is, $\omega \equiv \omega_r$.

Assumption 4.—The oscillator is not perturbed by exogenous inputs (e.g. light, metabolic changes) other than control actions.

Through these assumptions we seek to achieve a phase shift rather than entrain to a dynamic phase angle or an environment that also perturbs the system under control. Future studies may approach the more complicated case in which small differences between the day length and human circadian periods (see [41]), and dynamic light input, are accounted for. This simple case of achieving a static phase shift more simply yields optimal control policies that are sufficient to provide understanding for feedback controller design.

The phase difference between the reference oscillator and the oscillator under control is:

$$\chi(t) = \phi(t) - \phi_r(t) \mod 2\pi.$$

The total desired phase shift is such that $\chi(t_f) = 0$ where t_f is the time at which the control action is finished, and therefore

$$\Delta\phi_f = -\chi(0) \mod 2\pi.$$

In an ideal case where pharmaceutical inputs can be continuously dosed, one can derive an optimal control trajectory for entrainment via Pontryagin's Maximum/Minimum Principle [42]. The optimal control problem may be formulated with state dynamics

$$\dot{\phi} = \omega + B(\phi) \cdot u(t) \tag{11}$$

$$\phi(0) = \phi_0,$$

and cost functional

$$J[u(t)] = \int_0^{t_f} 1 \, dt, \tag{12}$$

for minimum-time entrainment. This cost functional does not include a penalty for introducing a drug which may have side effects or be otherwise undesirable. If the drug selected elicits dose-dependent off-target effects, the cost functional could be modified to penalize dose as well, and the performance between cases could be compared.

This yields a free-time fixed-endpoint problem, which can be solved using standard optimal control approaches [33, 42]. The Hamiltonian

$$\mathrm{H}(\phi, \lambda, u) = (\omega + B(\phi) \cdot u) \cdot \lambda + 1 \tag{13}$$

is formulated with λ is the dynamic costate with terminal condition $\lambda(t_f) = 0$. Using Pontryagin's Maximum Principle with constrained control inputs in the range $[u_{\min}, u_{\max}]$ we can write

$$\begin{aligned} \max_u \mathrm{H}(\phi, \lambda, u) &= \max_{u_{\min} \le u \le u_{\min}} \{(\omega + B(\phi) \cdot u) \cdot \lambda + 1\} \\ &= \omega \cdot \lambda + 1 + \max_{u_{\min} \le u \le u_{\min}} \{B(\phi) \cdot u \cdot \lambda\} \end{aligned}$$

which yields the optimal "bang-bang" control policy

$$\mathrm{u}^{\star}(\mathrm{t}) = \begin{cases} u_{\max} & \text{if } B(\phi) \cdot \lambda > 0 \\ u_{\min} & \text{if } B(\phi) \cdot \lambda \le 0 \end{cases}. \tag{14}$$

Here, the terminal condition of the costate, $\lambda(t_f) = 0$, forces the costate to be either positive or negative for any given initial condition for all time, and so for a specific initial condition the control input depends exclusively on the sign of $B(\phi)$, c.f. [33].

This implies that the oscillator should be shifted in only either the negative or positive direction toward its final phase shift, but not both. For an oscillator with desired phase shift $\Delta\phi f$ and initial phase ϕ_0, the optimal control input $u^{\star}$ is one of two admissible trajectories: either that for achieving the positive phase shift

$$\mathrm{u}_{+}^{*}(\mathrm{t}) = \begin{cases} u_{\max} & \text{if } B(\phi) > 0 \\ u_{\min} & \text{if } B(\phi) \le 0 \end{cases}. \tag{15}$$

or the negative phase shift

$$u^{*}_{-}(t) = \begin{cases} u_{\min} & \text{if } B(\phi) \geq 0 \\ u_{\min} & \text{if } B(\phi) < 0 \end{cases}. \tag{16}$$

Which of these two trajectories is optimal depends upon which reaches $\Delta\phi_f$ more rapidly, where $\Delta\phi$ accumulates over time according to

$$\Delta\phi = \int_0^t B(\phi(t')) \cdot u(t')dt' \mod 2\pi . \tag{17}$$

A decision boundary between advances and delays is denoted $\Delta\phi_b$ (i.e., $\Delta\phi_f > \Delta\phi_b$ indicates that delays are preferable, and $\Delta\phi_f < \Delta\phi_b$ indicates that advances are preferable). Correspondingly, that boundary for the continuous time optimal control is denoted $\Delta\phi_B^{\star}$. Thus, the two admissible trajectories may be calculated and compared to determine optimality.

This concept is illustrated for the small-molecule pharmaceutical KL001 input in Fig. 3. For KL001 drug action, $u_{\max} = 0.06$ to reflect a 60% decrease in degradation rate, and $u_{\min} = 0$ to reflect that decreasing the degradation rate cannot result in the synthesis of new biomolecules. The ipPRC is shown in Fig. 2 Here, we visualize the two admissible trajectories as two oscillators $\phi^{+}(t)$ and $\phi^{-}(t)$ receiving control inputs $u_{+}^{\star}(t)$ and $u_{-}^{\star}(t)$, respectively, causing the accumulation of phase shifts $\Delta\phi^{+}(t)$ and $\Delta\phi^{-}(t)$ moving in opposite directions around the circle of all phase shifts $\Delta\phi \in [0, 2\pi)$ starting at $\Delta\phi = 0$. The time $t_B^{\star}$ at which these oscillators meet therefore achieves any phase shift on $[0, 2\pi)$. This time may be found for any initial phase ϕ_0 by solving the coupled equations (10) and (17) for inputs $u_{+}^{\star}(t)$ and $u_{-}^{\star}(t)$, with the terminal condition of $\Delta\phi^{+}(t^{\star}) = \Delta\phi^{-}(t^{\star})$.

Thus, $\Delta\phi_b^{\star}(\phi_0)$ is defined as the phase shift at which $\phi^{-}(t)$ and $\phi^{+}(t)$ meet. The maximum time to achieve *any* phase shift starting from ϕ_0 is given by $t_b^{*}(\phi_0)$, and $\Delta\phi_b^{*}(\phi_0)$ (depicted as a solid black line in Fig. 3D) yields the boundary between where a positive shift and negative shift are optimal to reach a desired $\Delta\phi_f$ from phase ϕ_0. The optimal control input for a given ϕ_0 and $\Delta\phi_f$ is therefore defined as:

$$u^{\star}(t) = \begin{cases} u_{+}^{\star}(t) \text{ if } \Delta\phi_f < \Delta\phi_b^{\star}(\phi_0) \\ u_{-}^{\star}(t) \text{ if } \Delta\phi_f \geq \Delta\phi_b^{\star}(\phi_0) \end{cases}. \tag{18}$$

Therefore the optimal shift time $t_f^{\text{opt}}(\phi_0, \Delta\phi_f)$ may be calculated by applying this continuous time optimal control until $\Delta\phi = \Delta\phi_f$ One can formally write $t_b^{\star}$ as

$$t_b^{\star}(\phi_0) = \sup_{\Delta\phi_f} t_f^{\text{opt}}(\phi_0, \Delta\phi_f). \tag{19}$$

Remark. A bang-bang optimal control policy for light entrainment was shown previously in [24] and [31]. Applying these solutions, however, necessitated either a line search or a search for switching times due to the nonlinear temporal relationship between light input and driving force on the oscillator. Because pharmaceutical control of the clock can be parameterized directly within the limit cycle model, this formulation provides a simple way to compute the theoretical optimal phase resetting policy. Therefore, this provides a means for determining the relative efficacy of any applied control policy.

3. Closing the loop: designing an ipPRC-based nonlinear MPC

3.1. Motivation

Model predictive control is generally implemented by solving a finite-horizon optimal control problem at discrete sampling times, that is, for $t = 0, \tau, 2\tau, \ldots$, where τ is the sampling time. This section will demonstrate that time discretization causes a loss of optimal control in regions where the step includes a zero cross of the ipPRC, as the input cannot be adjusted during the step and the continuous-time optimal control policy will be violated (visualized in Fig. 5). Additionally, care must be taken to ensure that a controller selects the optimal direction of phase shifting (e.g., 8h advance vs. 16h delay) when only observing a finite portion of the future ipPRC. For the continuous time optimal control, that boundary is known to be $\Delta\phi_b^\star$, however, this is unlikely to be the correct boundary for the sampled-time case. Further, there is no guarantee that this boundary will appear from simply solving the finite-horizon optimal control.

This section explores how the formulation of the finite-horizon optimal control problem for MPC affects our ability to control the clock. For example, consider an ipPRC with a small positive region. A large sampling time in this case may result in a controller unable to access the positive ipPRC region due to the surrounding negative region being included in each step, and the inadvertent delays canceling the desired advance. This would force such a controller to use phase delays to achieve even small positive phase shifts.

It is therefore useful to compare the performance of a controller with piecewise constant control and uniformly spaced steps in sampling and switching time (the *sampling-time control*) to the continuous-time optimal control case to identify where inefficiency accumulates. This technique can be used to find a balance between high performance and the computational cost and physical impracticality of frequently updating the algorithm via measurement of physiological phase markers.

Additionally, the controller design will inform the suitability of a specific pharmaceutical for use. For example, a drug with a small positive ipPRC region will not be an appropriate choice for achieving large phase advances if an alternative drug with a larger positive ipPRC region exists. Alternately, the drug with the larger positive ipPRC region may be unsuitable if its pharmacokinetics are slow, such that sampling times must be very long to allow drug clearance between adjusting the dose.

3.2. Exploiting the ipPRC to choose MPC sampling time

Consider the more practical scenario of implementing control inputs parameterized as piecewise constant signals with equispaced time intervals between the knots (i.e., the switching times) of the control parameterization trajectory. This is commonly done when solving and applying MPC in continuous-time. The control input is defined as constant across time steps of uniform duration in a finite horizon, as commonly done for MPC:

$$u(t) = u_{\mathrm{i}} \text{ for } t \in [t_{\mathrm{i}-1}, t_{\mathrm{i}}), \tag{20}$$

where t_i are the discrete sampling times, with $i = 1, \ldots, N_p$ where N_p is the number of steps in the prediction horizon (and control horizon, in this paper). Then, $u_i \in [u_{\min}, u_{\max}]$ is a scalar constant, and $\tau = t_i - t_{i-1}$ is the sampling time for the controller.

Next, we are interested in quantifying the error incurred due to piecewise constant control actions at discrete sampling times. To better understand this error, we derive bounds on two metrics: the residual phase error, and the number of additional cycles to complete reset. Definitions for and proofs of bounds on these metrics follow.

Definition 1. *The* ***residual phase error*** $\left(\mathrm{E}_{\tau}^{\phi}\right)$ *is the phase error remaining at sampling time* t_k*, the first sampling time for which the continuous-time optimal control would have first completed the reset.*

Recall that $t_{\mathrm{f}}^{\mathrm{opt}}$ is the time at which $\chi(t) = 0$ under the optimal control and so

$$\mathrm{k} = \arg\min_{i \in \mathbb{N}} i \text{ subject to: } t_i \geq t_{\mathrm{f}}^{\mathrm{opt}}. \tag{21}$$

Based on Definition 1, the residual phase error is given by

$$\mathrm{E}_{\tau}^{\phi} = \chi\left(t_{\mathrm{k}}\right). \tag{22}$$

First, we provide bounds on E_{τ}^{ϕ} and demonstrate that these bounds are solely dependent on the sampling time τ and the ipPRC specific to the pharmaceutical being used. To do so, let n_{cyc} be the number of 2π cycles required to achieve the continuous time optimal control resetting using the optimal control policy (18). Furthermore, we denote $\mathrm{E}_{\tau}^{0,-}$,$\mathrm{E}_{\tau}^{0,+}$, and E^s to be the errors incurred at the step including the negative zero cross, the step including the positive zero cross, and a correction for phase advance cases, respectively.

Theorem 1 (Bounding of E_{τ}^{ϕ})**.** *Suppose Assumptions 1–4 hold. The residual phase error incurred by the sampled- time control* (20) *satisfies*

$$\mathrm{E}_{\tau}^{\phi} \leq \left(\mathrm{n}_{\mathrm{cyc}} + 1\right)\left(\mathrm{E}_{\tau}^{0,-} + \mathrm{E}_{\tau}^{0,+} + \mathrm{E}^{s}\right).$$

Proof. Note that the sampling-time control is able to identically track the continuous-time optimal control except where consecutive sampling times occur on either side of a zero-cross of the ipPRC.

By Assumption 2, the ipPRC has two zero-crosses in each $[0, 2\pi)$ cycle and so this occurs at most twice per cycle. A zero cross occurs at $t = t^0$ (where $t_{i-1} < t^0 < t_i$). The phase shift accumulated for a single step of the controller from t_{i-1} to t_i is given by

$$\Delta\phi_i(u_i) = \int_{t_{i-1}}^{t_i} B(\phi(t')) \cdot u_i dt', \tag{23}$$

which must be solved numerically with (10) to provide $\phi(t')$. Based on the arguments made in the previous section using Pontryagin's principle, we know that the optimal control policy selects either $u_{\max}$ or $u_{\min}$ until the zero-cross, then the opposite following it. Thus, the maximal loss of phase advance or delay is incurred if the switching times are positioned such that the shift incurred over that time step $\Delta\phi_i(u_i) = 0$ for the maximal input case, as this indicates that phase advances and delays cancel. The worst-case loss is therefore incurred where the controller takes no action on either side of the zero-cross, as the maximal action would also incur no shift (as shown in Fig. 5).

The maximal loss in phase shift is of equal magnitude in the positive and negative directions because competing advances and delays of identical magnitudes cancel. The maximal error incurred at a given zero-cross is

$$\mathrm{E}_\tau^0 = \max_{t_{i-1}, u} \left| \int_{t_{i-1}}^{t^0} B(\phi(t')) \cdot u dt' \right| \tag{24}$$

and the error for the positive or negative zero-cross is denoted $\mathrm{E}_\tau^{0,+}$ or $\mathrm{E}_\tau^{0,-}$, respectively.

The maximal phase error incurred for each cycle is the sum of the maximal error from each zero-cross, all other steps are identical to the continuous-time optimal control. A demonstration of the phase-shift cancellation by competing advances and delays is shown in Fig. 5.

For phase advances, additional error is incurred by phase advances accumulating less due to the zero crosses, thus, the oscillator phase will reach the next positive ipPRC region later than it would if it could achieve the full phase advance. The additional error accumulated is equal to the maximal shift that could be achieved in the extra time that it takes to reach the positive region. The time lost per cycle for phase advances is:

$$t_E = \frac{T}{2\pi} \cdot \left(E_\tau^{0,+} + E_\tau^{0,-} \right) \tag{25}$$

where T is the oscillatory period. This results in an additional residual phase error per cycle of

$$E^s = \begin{cases} \max\limits_{0 \le t \le T} \int_t^{t+\ t_E} B(\phi(t'))dt' & \text{if } \Delta\phi_f < \Delta\phi_b \\ 0 & \text{if } \Delta\phi_f \ge \Delta\phi_b \end{cases} \tag{26}$$

where $\Delta\phi_b$ is the decision boundary between phase advances and delays. In the case of phase delays, this quantity is 0, since the oscillator phase will reach the next negative region more rapidly due to achieving less phase delay.

These errors $\left(E_\tau^{0,-} + E_\tau^{0,+}, \text{ and } E^s\right)$ are caused by deviations from the optimal control policy at ipPRC zero crosses, and the sampling-time optimal control otherwise exactly follows the continuous time optimal control. Each of these errors occurs at a rate of once per cycle. Thus, we can bound the residual time phase error

$$E_\tau^\phi \le \left(n_{\text{cyc}}\left(\Delta\phi_f, \phi_0\right) + 1\right)\left(E_\tau^{0,-} + E_\tau^{0,+} + E^s\right) \tag{27}$$

for a given τ, which concludes the proof.

The bound E_τ^ϕ may be thought of as the maximal value of $\chi(t_k)$ for optimal control with sampling time τ.

Remark. At first glance, it may seem that this bound is loose for the case of arbitrarily large n_{cyc}, the number of cycles to achieve any reset is bounded by the finite scalar

$$n_{\text{cyc}}^{\max} = \max_{\Delta\phi_f,\ \phi_0} n_{\text{cyc}}\left(\Delta\phi_f, \phi_0\right)$$

the maximum number of number of cycles to achieve any continuous time optimal reset. Applying this value in (27), we may bound the residual time phase error for any sampling-time optimal reset.

Fig. 6 shows how the residual phase error bounds evolve over time for phase advances or delays, $\phi_0 = 0$, and $\tau = 2$ h. This corresponds to the actual residual phase error incurred, by numerically calculating the optimal control. Although the error bound increased identically for each cycle, the actual residual phase error varies each cycle, depending upon how each zero-cross aligns with the sampling times.

Fig. 7A-B demonstrates this bound for KL001 under $\tau = 1, 2, 4$ h, for $\phi_0 = 0$. The number of cycles to achieve each reset for phase advances or delays $\Delta\phi_f \in [-2\pi, 2\pi]$ (to explicitly show advances and delays) and $\phi_0 = 0$ under the continuous-time optimal control for KL001 is shown in Fig. 7A. The region of this plot that corresponds to the phase resetting directions selected by the continuous-time optimal control (advances for $\Delta\phi_f < \Delta\phi_B^\star$, delays for $\Delta\phi_f \ge \Delta\phi_B^\star$) are denoted between the dashed lines. Notably, the number of cycles to achieve

phase advances accumulates more rapidly due to the smaller positive region of the ipPRC, and as such, error will incur more quickly in phase advance regions. Fig. 7B shows the bounds on the residual phase error for $\tau = 1, 2, 4$ h, and the actual residual phase error incurred for numerically calculating the optimal control for each shift and sampling time. As observed, the calculated residual phase errors fall within the computed bounds.

In addition to the residual phase error, it is useful to determine how much additional time it would take to reach a complete reset. Effectively, we aim to nd t^{add} for $\chi\left(t_f^{\text{opt}} + t^{\text{add}}\right) = 0.$ Finding or bounding t^{add} is complicated, as it necessitates knowledge of the entire discrete-time optimal control trajectory for each ϕ_0 and $\Delta\phi_f$ How-ever, the most additional time to reset is incurred when a

phase advance requires waiting through the negative ipPRC region (or the inverse case, a phase delay forcing a wait through the positive ipPRC region) to complete the resetting. Conveniently, we may bound the number of additional cycles to complete a phase reset without explicit knowledge of the full discrete-time optimal control.

Definition 2. *The number of additional cycles to complete reset* $\left(\mathrm{n}_{\tau}^{\text{add}} \in \mathbb{N}\right)$ *is a positive scalar that enforces the constraint*

$$\chi\left(t_f + \mathrm{n}_{\text{cyc}}^{\text{add}} T^{+,\,-}\right) = 0$$

where $T^{+,-}$ *is the accelerated/decelerated period of the oscillator induced by a phase advance/delay, respectively.*

Remark. We note that $T^{+} < T < T^{-}$, and by virtue of being on a limit cycle yields

$$\lim_{t \to \infty} \left[x\left(t,\ p, u_{+}\right) - x\left(t - t^{+}, p, u_{+}\right)\right] = 0$$

and

$$\lim_{t \to \infty} \left[x(t,\ p, u\_) - x(t - t^{-}, p, u\_)\right] = 0$$

which hold true for any T^{+}-periodic control input u_{+} that phase advances the oscillator and shortens the period, or T $^{-}$-periodic control input u_{-} that phase delays the oscillator and lengthens the period. In sum, phase advances result in transient shortening of the period, and phase delays result in a transient lengthening of the period.

The following upper bound is proposed on the number of additional cycles to complete reset.

Theorem 2 (Bounding of $\mathrm{n}_{\text{cyc}}^{\text{add}}$).

Let $\Delta\phi^{-}_{g,\,\tau}$ *and* $\Delta\phi^{+}_{g,\,\tau}$ *be the maximum phase delay and advance that is guaranteed to be achievable in a cycle of* 2 *respectively. Recall that* $\Delta\phi_b$ *is the decision boundary between phase advances or delays. The number of additional cycles to correct for the residual phase error incurred by the sampled-time control* (20) *is bounded above by*

$$n^{\text{add}}_{\text{cyc}} \leq \begin{cases} \left\lfloor E^{\phi}_{\tau}/\Delta\phi^{+}_{g,\,\tau} \right\rfloor + 1 \;\; if\; \Delta\phi_f \leq \Delta\phi_b \\ \left\lfloor E^{\phi}_{\tau}/\Delta\phi^{-}_{g,\,\tau} \right\rfloor + 1 \;\; if\; \Delta\phi_f \leq \Delta\phi_b \end{cases}, \tag{28}$$

where $\lfloor \cdot \rfloor$ denotes the floor function.

Proof. Let $\Delta\phi^{+}_{\text{g},\,\tau}$ denote the phase advance that are guaranteed to be achievable in a single cycle. This is more simply thought of as the maximum total phase advance or delay achievable in a cycle minus the worst-case zero-cross error. For advances, the maximum phase advance achieved in a cycle is achieved by the continuous time optimal control policy for advances:

$$\Delta\,\phi^{+}_{\text{cyc}} \;=\; \int_0^{T^{+}} B(\phi(t'))\;\cdot\; u^{*}_{\_}(\text{t'})dt'. \tag{29}$$

The phase shift that is guaranteed to be achievable for each cycle for a step size of τ is therefore

$$\Delta\,\phi^{+}_{\text{g},\,\tau} \;=\; \Delta\,\phi^{+}_{\text{cyc}} \;-\; E^{0,+}_{\tau} \;-\; E^{0,\,-}_{\tau}, \tag{30}$$

and an equivalent method can be used to calculate $\Delta\phi^{+}_{\text{g},\,\tau}$.

The bound may be calculated by dividing the remaining error E^{ϕ}_{τ} by the advance or delay that is guaranteed to be achieved each cycle for a sampling time of τ This implies (28).

This bound is plotted for the use of KL001 and τ = 1, 2, 4 h in Fig. 7C. For sampling times of 1 or 2 h, the residual phase error may corrected in 1–2 cycles, with the vast majority of cases needing only an additional cycle. Furthermore, because this method assumes the worst possible alignment of the control switching times for every cycle, it is unlikely that the full two additional cycles will be needed. Alternately, the 4 h sampling time potentially results in many additional cycles needed to achieve a complete reset. This is because, as shown in Fig. 5, a 4 h step could be aligned such that nearly the entire phase advance region of the ipPRC is lost. Correspondingly, $\Delta\phi^{+}_{\text{g},\,\tau}$ is small in value. Thus, we may conclude that a sampling time of 2 h is approximately as effective as 1 h, and provides a 50% reduction in computational and sampling costs over 1 h sampling times, given our assumptions. Our controller designed in Section 4 therefore uses τ = 2 h. However, if a controller was designed to only produce phase delays, a 4 h sampling time is sufficiently near to the optimal control. This metric is therefore useful in designing controllers for a given problem and potential pharmaceutical given its ipPRC.

For these plots, dashed lines bound the phase shifts that are performed by the continuous-time optimal control to achieve complete coverage of phase shifts from 0 to 2π This implicitly assumes that the direction of optimal phase shifting is known for our choice of τ and is identical to the continuous time optimal control bound $\Delta\phi^\star$. The following subsection addresses the choice of resetting direction (given $\Delta\phi^\star$ for the continuous optimal case).

While numerical calculation of the in finite-horizon optimal control may seem advantageous, the optimizations involved are costly, as the selection of each step affects the region of the ipPRC available in the ensuing steps. Rather, we propose to formulate the MPC problem using the bounds derived from the optimal control as guidelines, and perform phase shifts by repeatedly solving the finite-horizon optimal control. It must be emphasized that because the bound we have derived is for the maximum residual phase error incurred at each zero-cross, a controller need not solve the in finite-time optimal sampling time control to be governed by the bound. Solving a finite-horizon optimal control problem is sufficient for these bounds on accumulated error to apply.

3.3. On MPC formulation and the selection of a prediction and control horizon

MPC involves repeatedly solving a finite horizon optimal control problem. For the continuous-time optimal control (in Section 2.3) the direction of phase shifting was determined by calculation of the equal-time phase shift boundary $\Delta\phi^\star$. That boundary may differ for the discrete-time optimal control for a given step size. However, in MPC, whether the oscillator should advance or delay is implicitly solved in the finite horizon optimization. This subsection demonstrates via simulations that the finite-horizon optimal control in MPC is dependent on the number of steps. Because the finite-horizon optimal phase shifting direction can differ from the true optimal direction, we propose a multi-staged optimization for solving the finite-horizon optimal control.

The MPC problem is formulated in a similar fashion as [34], where the prediction and control horizons are set to be identical. The prediction horizon consists of N_p steps with sampling time τ For an oscillator at $t = t_k$ this yields:

$$U \triangleq \left[u(t_k+\tau) \quad \cdots \quad u(t_k+N_p\tau)\right]^\top \tag{31}$$

as the knots of the control trajectory defined at each of the N_p steps. Eqn. (10) is used as the MPC predictor, which estimates the oscillator phase at the end of each of $\ell \in \left[1, \cdots, N_p\right]$ steps. The predictor formulation of (10) is

$$\hat{\phi}(t_i+\ell\tau) = \phi(t_i) + \omega\tau + \sum_{k=1}^{\ell}\int_{t_i+(k-1)\tau}^{t_i+k\tau} B(\phi)\cdot u_k\, dt\,. \tag{32}$$

Here, i is the index of the applied control steps, and k is the index of the predictor. The predicted phase error $\hat{\chi}(t)$ is de ned as the magnitude of the phase difference between the predicted phase $\hat{\phi}$ and the environmental phase:

$$\hat{\chi}(t) \triangleq \hat{\phi}(t) - \phi_r(t) \mod 2\pi \tag{33}$$

We want the phase error to reflect the absolute phase distance between reference and controlled oscillators, and so we define

$$h_\phi(t) = \min|(\hat{\chi}(t), 2\pi - \hat{\chi}(t))|$$

To avoid overcorrecting to for small amounts of noise in the phase estimate, we ignore phase error below a constant δ_ϕ:

$$g_\phi(t) \triangleq \begin{cases} 0 & \text{if } h_\phi(t) < \delta_\phi \\ h_\phi(t) & \text{otherwise} \end{cases} \tag{34}$$

so that imprecisions in calculating phase do not result in controller action (similar to the zone-MPC approach of [7]). Thus, the finite-horizon optimal control problem at each time t_i may be solved for the optimal trajectory $u^\star_{\text{MPC}}$:

$$u^\star_{\text{MPC}} = \arg\min_U \Sigma_{l=1}^{N_p} \omega^\phi \, g_\phi^2(t_i + \ell\tau) + \omega^u u_\ell^2 \tag{35}$$

Subject to: Eqn. (32) Eqn. (33)

$$u_{\min} \le u_\ell \le u_{\max},$$

for all $\ell = 1, \ldots, N_p$, where w^ϕ and w^u are positive weighting scalars evaluated at the end of the time step and start of the time step, respectively, as phase error is calculated after the control is applied for that step. Weights w_u are nominal so that the control input is 0 when input would incur no phase shift. After identifying the optimal piecewise control trajectory $u^\star$, the initial step $u_1^\star$ is applied to the system (in this case the full 8-state ODE model) for $t \in (t_i, t_i + \tau]$ as is standard in model predictive control.

For the ensuing simulations, a sampling time of $\tau = 2$ h was used, as described in the prior section. Additionally, weights $w^\phi = 10$ and $w^u = 1$, so that no control is delivered for marginal gains in phase shifting, and $\delta_\phi = 0.1$ as in [35].

Simulations demonstrate that the optimization step in the MPC is susceptible to errors induced by choosing the slower direction to achieve a phase reset. The example in Fig. 8 shows how the finite horizon optimal control $u^\star_{\text{MPC}}$ calculated for $\Delta\phi_f = \pi$, $\phi_0 = 0$ varies with N_p. For short prediction horizons, the controller applies a phase advance, because more of the positive ipPRC region is immediately visible to the predictor, suggesting that this will achieve most rapid reset. By extending the prediction horizon to $N_p = 10$ (20 h), more of the negative region of the ipPRC is available, and so the controller targets phase delays, consistent with the in finite-time optimal control for $\tau = 2$ h. In some cases, performing the

optimization with both advances and delays available results in an “indecisive” controller that changes directions during the phase shifting, causing excessive lags to reset.

One solution to the problem of inconsistent finite-horizon optimal control policies is to extend the prediction horizon such that it always observes the entirety of the reset, to guarantee that the finite-horizon optimal control calculated during MPC is equal to the in finite-horizon optimal control, that is: $u^{\star}_{\mathrm{MPC}} = u^{\star}_{\tau}$ However, this leads to excessive computational costs, especially for small sampling times. Furthermore, the optimization formulation itself poses challenges. Minimizing the time to complete the reset is challenging, as there is no gradient along which to move unless the reset is completed during the prediction horizon, leading to cases with an in finite cost. Minimizing phase error at each step may be suboptimal as well: the fastest reset achievable may involve temporarily increasing $\chi(t)$ for cases where $\Delta\phi_f < \pi$ and $\Delta\phi_f > \Delta\phi^{\star}_{\tau}$.

An alternate approach, which is employed in the following section, is to use a two-stage optimization, where the direction of the reset (advance or delay) is chosen first using a pre-selected bound $\Delta\phi_b$. Furthermore, we propose using $\Delta\phi_b = \Delta\phi^{\star}$, as this will ensure that the bounds on residual phase error and additional cycles to reset hold, as derived in the prior section, and thus provide guarantees on controller performance in the absence of plant-model mismatch or noise. One further important consideration is that a controller with $N_p = 1$ will be unable to avoid ac-cumulating error at the zero-crosses, whereas a controller with a larger N_p provides exibility to adjust the control inputs preceding a zero-cross to minimize the accumulated error. For this reason, the design parameters were set to $N_p = 5$ and $\Delta\phi_b = \Delta\phi^{\star}$ for our controller. We note, however, that the derived bounds hold even for $N_p = 1$.

4. In silico application of designed MPC controller for phase resetting

To demonstrate the efficacy of this control approach, this section presents two control scenarios in our *in silico* circadian simulator (shown schematically in Fig. 4). We note here that the predictor in our MPC formulation was the phase only model, but the system that we applied the control input to is the full 8-state ODE model. Thus, we are testing Assumption 1, since control inputs will perturb the system from the limit cycle, and control will only be viable if the phase-only model reduction accurately describes phase dynamics in state space. We parameterized the controller with $\tau = 2$h, $N_p = 5$, and $\Delta\phi_b = \Delta\phi^{\star}$, as prescribed by the design considerations in the prior sections.

The first scenario (Fig. 9) is an example of extreme jet lag, in which the environment undergoes a 5 h phase advance, followed three days later by a phase delay of 11 h This is the equivalent of a flight from Boston to London, then London to Honolulu. Here, the controller first waits for a positive region of the ipPRC rather than immediately attempting a delay, as $\Delta\phi_f < \Delta\phi_b$. Thus, it completes this phase advance in approximately 30 h. The 11 h delay occurs during a region where $B(\phi) < 0$, and thus control input is provided immediately. This shift is completed in a total of two continuous doses, in approximately 36 h. Without a control approach, light input alone achieves a shift of 2–3h per day at most, thus, this represents a significant speed up in phase resetting. If light is delivered under optimal

control, these shifts are predicted to take several days [24], thus demonstrating the potential for improvement by a pharmaceutical input.

The second scenario (Fig. 10) involves the delivery of KL001 to align an individual working on a rotating shift pattern, from [43]. Here, the time on the x-axis is local time, and the phase of the oscillator is initialized for an individual entrained to a normal light-dark cycle. In this protocol, the individual works two days of morning shift, followed by two days of evening shift, followed by two days of night shift (work shifts denoted by gray boxes. This perpetually phase-delaying schedule is easily tracked by KL001 input, as KL001 more easily achieves phase delays. Notably, by taking one large dose after the conclusion of each pair of shifts, the individual rapidly shifts to the new work phase. Each shift worked in this protocol occurs while the individual is exactly phase-aligned with their work environment. This scenario exclusively involves phase delays, and so under the analysis of Section 3 a larger sampling time of $\tau = 4$ h (or possibly even larger) would be similarly effective.

Importantly, the possibly competing effects of light was neglected in these scenarios. This could be achieved simply by restricting light input during the phase-shifting steps, or treating light as a complementary control input, using an approach such as [25]. This is a possible topic for future study.

5. Conclusion

This paper presents tools for designing feedback control of circadian rhythms based upon an optimal control policy. Importantly, this paper provides a bound on the errors incurred by discrete sampling and update times during pharmaceutical control of circadian oscillation. These results may then be applied to compare against the efficacy of a light-based approach; see, for example, [33]. Thus, we implicitly considered the non-instantaneous pharmacokinetics of circadian pharmaceuticals. The optimal control problem in the case of known pharmacokinetics will likely necessitate computational, rather than analytic, solutions. In our simpler discrete-time formulation, one may think of the time step of each input as an approximation of the pharmacokinetics by a square wave. While this approximation is clearly non-physiological, it yields analytic solutions that can guide the search for viable pharmaceutical actuators. The principle that imprecision in following the optimal control policy is incurred when drug action occurs during ipPRC zero-crosses points indicates, for example, that a drug with slow pharmacokinetics and a short ipPRC positive region will be inefficient at achieving phase advances. Future studies will directly incorporate pharmacokinetics, which may be easily added to an MPC formulation, and then MPC solutions may be compared to performance under the optimal control policy. Furthermore, we have shown the advantages provided by pharmacological actuation, as any phase reset is predicted to be achievable in under 60 h. Uniquely, we propose a formulation of the nonlinear MPC based on the bounds derived from the optimal control problem.

Since drug discovery for the manipulation of circadian rhythms is ongoing, this paper presents a perspective on which pharmaceuticals will enable control. Specifically, an ideal circadian pharmaceutical will have (i) large-magnitude positive and negative ipPRC regions, to enable bidirectional phase resetting, and (ii) rapid pharmacokinetics, to enable precise

timing of its effect and minimal losses during ipPRC zero-crosses. This is in contrast to traditional pharmaceutical approaches that select for a long drug half-life to ensure its effect is sustained. Future studies should explicitly incorporate simple pharmacokinetic profiles, rather than piecewise-constant control. However, achieving closed-form solutions or error bounds for these problems is challenging.

While this study has approached control of circadian phase, circadian amplitude has been also been recognized as an important physiological parameter. Previous studies have predicted that light resetting of the clock is associated with diminished amplitudes [24, 44]. Changes in amplitude of circadian parameters may be due to deviation from the limit cycle within individual cells, or due to desynchrony among cellular oscillators during phase shifting [45]. Deviation from the limit cycle may be predicted through amplitude response curves (ARCs), and changes in synchrony may be predicted by judicious use of the ipPRC. Elsewhere, we have begun the development of MPC that maintains circadian synchrony through phase shifting [46] Future studies of circadian control may explicitly include control of circadian amplitudes in addition to phase.

Finally, the remaining link in achieving closed-loop circadian control is the development of a noninvasive circadian phase sensor. While phase may be assessed relatively simply from cellular bioluminescent reporters, noninvasive *in vivo* phase inference in mammals is an open challenge. Recent work in assessing phase from ambulatory monitoring has shown promise and reasonably accurate predictions [47], however, the propagation of measurement errors through such a control algorithm also remains an open question. An ideal phase sensor would rely on simple data such as solely actigraphy, rather than additional biophysical sensors (e.g. skin temperature, salivary melatonin), and remains in development. The abilities of phase inference methods and controller robustness will be of consequence for the implementation of any form of circadian feedback control mentioned herein.

Acknowledgments

1 This work was funded by: National Institutes of Health (NIH) T32-HLO7901, R01-GM096873–01, and K24-HL105664.

Biographies

John H. Abel is a Postdoctoral Research Fellow in the Department of Anesthesiology, Critical Care and Pain Medicine at Massachusetts General Hospital, and an affiliated researcher at the Picower Institute for Learning and Memory at the Massachusetts Institute of Technology (MIT). He received the B.S. degree in Chemical Engineering from Tufts University (2013), the M.S. degree in Chemical Engineering from the University of California, Santa Barbara (2015), and the Ph.D. degree in Systems Biology from Harvard University (2018). His research interests are in the application of dynamical systems and

control theory to neuroscience, systems biology of circadian rhythms, and physiological closed-loop control.

Ankush Chakrabarty received a B.E. with first-class honors from Jadavpur University, India (2011), and a Ph.D. from Purdue University, West Lafayette, USA (2016), both in Electrical Engineering. Since 2018, he is a Visiting Research Scientist in the Control and Dynamical Systems Group at Mitsubishi Electric Research Laboratories (MERL), prior to which he was a Postdoctoral Fellow at Harvard University. His work leverages artificial intelligence and approximation/relaxation methods to provide computationally efficient and provably safe solutions to challenging problems arising in control and estimation, with special emphasis on model predictive control and unknown input estimation.

Elizabeth B. Klerman is an Associate Professor of Medicine at Harvard Medical School, and a Physician in the Department of Medicine at the Brigham and Women's Hospital. She received the B.S. in Interdisciplinary Science from the Massachusetts Institute of Technology (MIT), and the PhD in Physiology and M.D. degrees from Harvard University. She is a member of the Sleep Research Society, the Society for Research in Biological Rhythms, the Association for Clinical and Translational Science, and the International Academy of Astronautics. She was received Phi Beta Kappa for her undergraduate work and a NASA JSC Director’s Innovation Award as member of International Space Station Flexible Lighting Team. Dr. Klerman's current research focuses on the interaction of endocrine, circadian and sleep rhythms in normal and pathological states; and mathematical modeling of the sleep and circadian system and markers of its function.

Francis J. Doyle III is the John A. Paulson Dean of the Harvard John A. Paulson School of Engineering and Applied Sciences at Harvard University, where he also is the John A. & Elizabeth S. Armstrong Professor. He received a B.S.E. degree from Princeton, C.P.G.S. from Cambridge, and Ph.D. from Caltech, all in Chemical Engineering. He has held faculty appointments at Purdue University, the University of Delaware, and UCSB. In addition he

has held visiting positions at DuPont, Weyerhaeuser, and Stuttgart University. He has been recognized as a Fellow of multiple professional organizations including: IEEE, IFAC, AIMBE, and the AAAS. He was the President for the IEEE Control Systems Society in 2015, and was the Vice President of the International Federation of Automatic Control from 2014 to 2017. In 2005, he was awarded the Computing in Chemical Engineering Award from the AIChE for his innovative work in systems biology, and in 2015 received the Control Engineering Practice Award from the American Automatic Control Council for his development of the artificial pancreas. In 2016, he was inducted as a Fellow into the National Academy of Medicine for his work on biomedical control. His research interests are in systems biology, network science, modeling and analysis of circadian rhythms, and drug delivery for diabetes.

References

[1]. Lauffenburger DA, Cell signaling pathways as control modules: Complexity for simplicity?, Proc. Natl. Acad. Sci. U. S. A 97 (10) (2000) 5031–5033. [PubMed: 10805765]

[2]. Bailey J, Haddad W, Drug dosing control in clinical pharmacology, IEEE Control Syst. Mag 25 (2) (2005) 35–51.

[3]. Julius AA, Halasz A, Sakar MS, Rubin H, Kumar V, Pappas GJ, Stochastic modeling and control of biological systems: The lactose regulation system of escherichia coli, IEEE Trans. Automat. Contr 53 (2008) 51–65.

[4]. Perley JP, Mikolajczak J, Harrison ML, Buzzard GT, Rundell AE, Multiple model-informed open-loop control of uncertain intracellular signaling dynamics, PLoS Comput. Biol 10 (4).

[5]. Yan G, Vertes PE, Towlson EK, Chew YL, Walker DS, Schafer WR, Barabasi A-L, Network control principles pre-dict neuron function in the caenorhabditis elegans connectome, Nature 550 (7677) (2017) 519–523. [PubMed: 29045391]

[6]. El-Samad H, Kurata H, Doyle JC, Gross CA, Khammash M, Surviving heat shock: Control strategies for robustness and performance, Proc. Natl. Acad. Sci. U. S. A 102 (8) (2005) 2736–2741. [PubMed: 15668395]

[7]. Gondhalekar R, Dassau E, Doyle III FJ, Periodic zone-mpc with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes, Automatica 71 (2016) 237– 246. [PubMed: 27695131]

[8]. Chen S, Harrigan P, Heineike B, Stewart-Ornstein J, El-Samad H, Building robust functionality in synthetic circuits using engineered feedback regulation, Curr. Opin. Biotechnol 24 (4) (2013) 790–796. [PubMed: 23566378]

[9]. Del Vecchio D, Abdallah H, Qian Y, Collins JJ, A blueprint for a synthetic genetic feedback controller to reprogram cell fate, Cell Syst 4 (1) (2017) 109–120. [PubMed: 28065574]

[10]. Craven S, Whelan J, Glennon B, Glucose concentration control of a fed-batch mammalian cell bioprocess using a nonlinear model predictive controller, J. Process Control 24 (4) (2014) 344–357.

[11]. Giordano N, Mairet F, Gouze JL, Geiselmann J, de Jong H, Dynamical allocation of cellular resources as an optimal control problem: Novel insights into microbial growth strategies, PLoS Comput. Biol 12 (3) (2016) 1–28.

[12]. Haddad WM, Hayakawa T, Bailey JM, Adaptive control for non-negative and compartmental dynamical systems with applications to general anesthesia, Int. J. Adapt. Control Signal Process 17 (3) (2003) 209–235.

[13]. Luni C, Doyle III FJ, Robust multi-drug therapy design and application to insulin resistance in type 2 diabetes, Int. J. Robust Nonlinear Control 21 (15) (2011) 1730–1741.

[14]. Yu J, Zhang Y, Yan J, Kahkoska AR, Gu Z, Advances in bioresponsive closed-loop drug delivery systems, Int. J. Pharm (2017) (online).

[15]. Jin X, Kim C-S, Dumont GA, Ansermino JM, Hahn J-O, A semi-adaptive control approach to closed-loop medication infusion, Int. J. Adapt. Control Signal Process 31 (2) (2017) 240–254.

[16]. Amendola S, Lodato R, Manzari S, Occhiuzzi C, Marrocco G, R d technology for iot-based personal healthcare in smart spaces, IEEE Internet Things J 1 (2) (2014) 144–152.

[17]. Zavitsanou S, Chakrabarty A, Dassau E, Doyle F III, Embedded control in wearable medical devices: Application to the artificial pancreas, Processes 4 (4) (2016) 35.

[18]. Qiu WY, Tong S, Zhang L, Sakurai Y, Myers DR, Hong L, Lam WA, Bao G, Magnetic forces enable controlled drug delivery by disrupting endothelial cell-cell junctions, Nat. Commun 8 (2017) 15594. [PubMed: 28593939]

[19]. Bakh NA, Cortinas AB, Weiss MA, Langer RS, Anderson DG, Gu Z, Dutta S, Strano MS, Glucose-responsive insulin by molecular and physical design, Nat. Chem 9 (10) (2017) 937–943. [PubMed: 28937662]

[20]. Mohawk JA, Green CB, Takahashi JS, Central and peripheral circadian clocks in mammals, Annu. Rev. Neurosci 35 (1) (2012) 445–462. [PubMed: 22483041]

[21]. Marquie J-C, Tucker P, Folkard S, Gentil C, Ansiau D, Chronic effects of shift work on cognition: findings from the visat longitudinal study, Occup. Environ. Med 72 (4) (2015) 258–264. [PubMed: 25367246]

[22]. Hirota T, Lee JW, Lewis WG, Zhang EE, Breton G, Liu X, Garcia M, Peters EC, Etchegaray J-P, Traver D, Schultz PG, Kay SA, High-throughput chemical screen identifies a novel potent modulator of cellular circadian rhythms and reveals ckiα as a clock regulatory kinase, PLoS Biol 8 (12) (2010) e1000559. [PubMed: 21179498]

[23]. Hirota T, Lee JW, John PC St., Sawa M, Iwaisako K, Noguchi T, Pongsawakul PY, Sonntag T, Welsh DK, Brenner DA, Doyle FJ III, Schultz PG, Kay SA, Identfication of small molecule activators of cryptochrome, Science 337 (6098) (2012) 1094–1097. [PubMed: 22798407]

[24]. Serkh K, Forger DB, Optimal schedules of light exposure for rapidly correcting circadian misalignment, PLoS Comput. Biol 10 (4) (2014) e1003523. [PubMed: 24722195]

[25]. Julius A, Yin J, Wen JT, Time-optimal control for circadian entrainment for a model with circadian and sleep dynamics, in: IEEE Conf. Decis. Control 2017, pp. 4709–4714.

[26]. Bagheri N, Stelling J, Doyle III FJ, Circadian phase resetting via single and multiple control targets, PLoS Comput. Biol 4 (7) (2008) e1000104. [PubMed: 18795146]

[27]. John PC St., Hirota T, Kay SA, Doyle FJ III, Spatiotemporal separation of per and cry posttranslational regulation in the mammalian circadian clock, Proc. Natl. Acad. Sci. U. S. A 111 (5) (2014) 2040–2045. [PubMed: 24449901]

[28]. Bagheri N, Stelling J, Doyle III FJ, Circadian phase entrainment via nonlinear model predictive control, Int. J. Robust Nonlinear Control 17 (2007) 1555–1571.

[29]. Slaby O, Sager S, Shaik OS, Kummer U, Lebiedz D, Optimal control of self-organized dynamics in cellular signal transduction, Math. Comput. Model. Dyn. Syst 13 (5) (2007) 487–502.

[30]. Shaik O, Sager S, Slaby O, Lebiedz D, Phase tracking and restoration of circadian rhythms by model-based optimal control, IET Syst. Biol 2 (1) (2008) 16–23. [PubMed: 18248082]

[31]. Zhang J, Qiao W, Wen JT, Julius A, Light-based circadian rhythm control: Entrainment and optimization, Automatica 68 (2016) 44–55.

[32]. E mov D, Sacre P, Sepulchre R, Controlling the phase of an oscillator: A phase response curve approach, in: IEEE Conf. Decis. Control 2009, pp. 7692–7697.

[33]. Qiao W, Wen JT, Julius A, Entrainment control of phase dynamics, IEEE Trans. Automat. Contr 62 (1) (2017) 445–450.

[34]. Abel JH, Doyle III FJ, A systems theoretic approach to analysis and control of mammalian circadian dynamics, Chem. Eng. Res. Des 116 (2016) 48–60. [PubMed: 28496287]

[35]. Abel JH, Chakrabarty A, Doyle FJ III, Nonlinear model predictive control for circadian entrainment using small-molecule pharmaceuticals, IFAC-PapersOnLine 50 (1) (2017) 9864–9870.

[36]. Kim JK, Forger DB, Marconi M, Wood D, Doran A, Wager T, Chang C, Walton KM, Modeling and validating chronic pharmacological manipulation of circadian rhythms, CPT Pharmacometrics Syst. Pharmacol 2 (7) (2013) e57. [PubMed: 23863866]

[37]. Taylor SR, Gunawan R, Petzold LR, Doyle FJ III, Sensitivity measures for oscillating systems: Application to mammalian circadian gene network, IEEE Trans. Automat. Contr 53 (2008) 177–188. [PubMed: 19593456]

[38]. Johnson CH, Forty years of prcs-what have we learned?, Chronobiol. Int 16 (6) (1999) 711–743. [PubMed: 10584173]

[39]. Winfree AT, The geometry of biological time, Springer-Verlag, 2001.

[40]. Kronauer RE, Forger DB, Jewett ME, Quantifying human circadian pacemaker response to brief, extended, and repeated light stimuli over the phototopic range, J. Biol. Rhythms 14 (6) (1999) 500–515. [PubMed: 10643747]

[41]. Duffy JF, Cain SW, Chang A-M, Phillips AJK, Munch MY, Gron er C, Wyatt JK, Dijk D-J, Wright KP, Czeisler CA, Sex difference in the near-24-hour intrinsic period of the human circadian timing system, Proc. Natl. Acad. Sci. U. S. A 108 (Supplement 3) (2011) 15602–15608. [PubMed: 21536890]

[42]. Kirk DE, Optimal Control Theory: An Introduction, Dover Publications, 2012.

[43]. Vetter C, Fischer D, Matera JL, Roenneberg T, Aligning work and circadian time in shift workers improves sleep and reduces circadian disruption, Curr. Biol 25 (7) (2015) 907–911. [PubMed: 25772446]

[44]. Diekman CO, Bose A, Reentrainment of the circadian pacemaker during jet lag: East-west asymmetry and the effects of north-south travel, J. Theor. Biol 437 (2018) 261–285. [PubMed: 28987464]

[45]. John PC St., Taylor SR, Abel JH, Doyle FJ III, Amplitude metrics for cellular circadian bioluminescence reporters, Biophys. J 107 (11) (2014) 2712–2722. [PubMed: 25468350]

[46]. Abel JH, Chakrabarty A, Doyle FJ III, Controlling biological time: Nonlinear model predictive control for populations of circadian oscillators, in: Emerg. Appl. Control Syst. Theory. Lect. Notes Control Inf. Sci. Proceedings, Springer, 2018, pp. 123–138.

[47]. Kolodyazhniy V, Spati J, Frey S, Gotz T, Wirz-Justice A, Krauchi K, Cajochen C, Wilhelm FH, An improved method for estimating human circadian phase derived from multichannel ambulatory monitoring and artificial neural networks, Chronobiol. Int 29 (8) (2012) 1078–1097. [PubMed: 22891656]

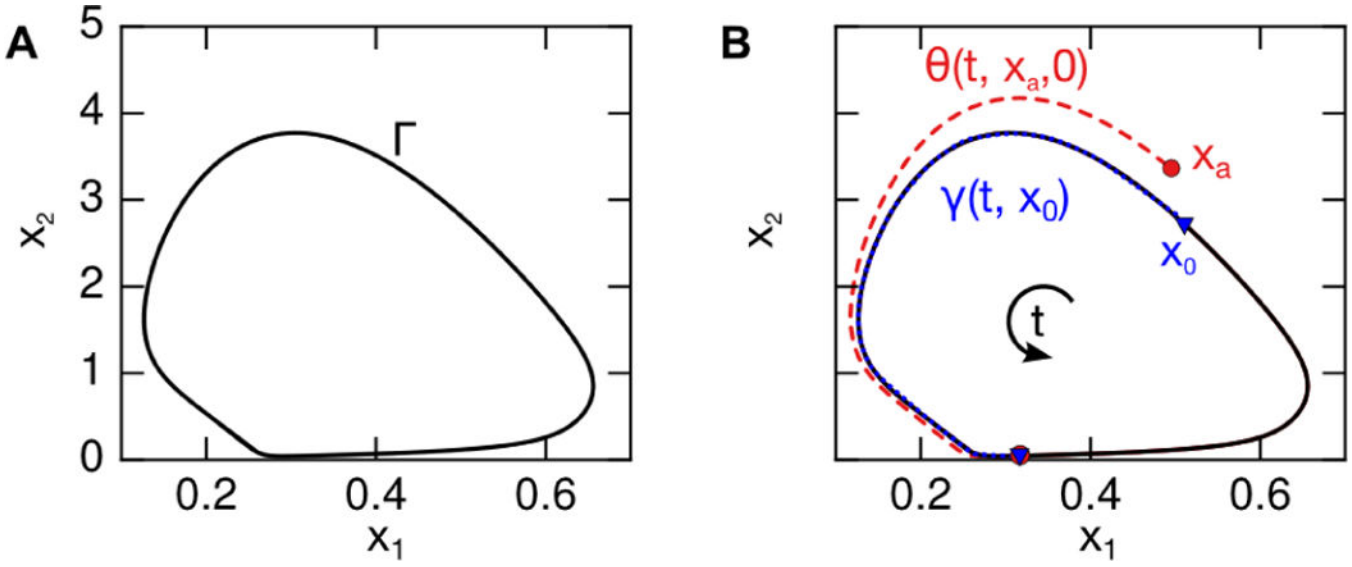

Figure 1:

Schematic of limit cycle dynamics for arbitrary states x_1 and x_2. (A) The limit cycle Γ in state space. (B) Schematic of trajectories γ and θ These points may be assigned identical phases ($\phi_a = \phi_0$), as the trajectories originating at these points converge asymptotically.

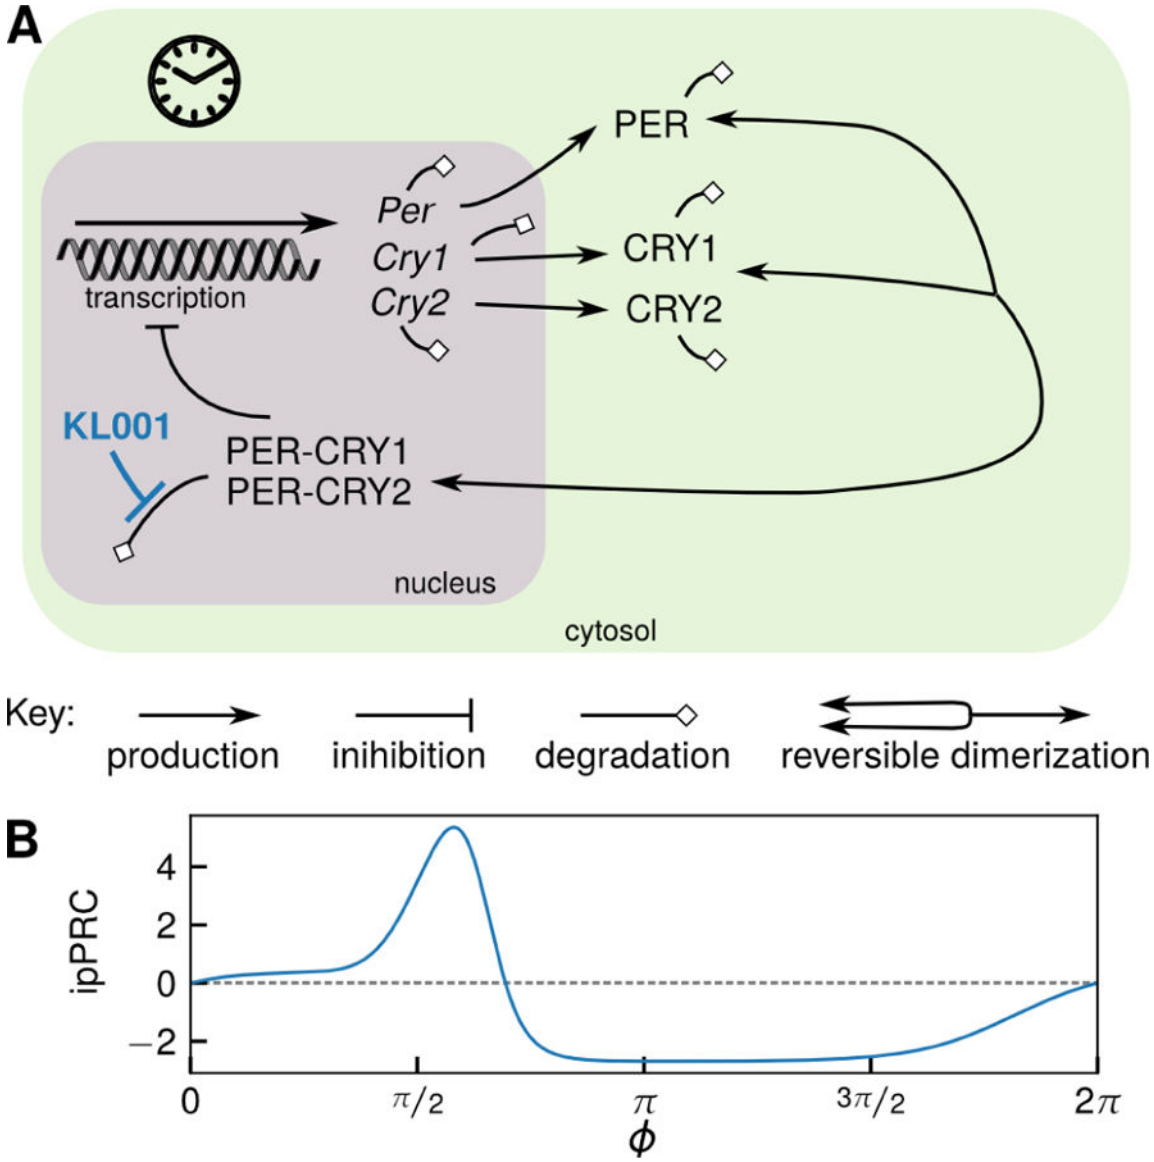

Figure 2:

Schematic of the core circadian gene regulatory network and effect of small molecule KL001. (A) The core circadian negative feedback loop. KL001 stabilizes nuclear CRY by reducing its degradation rate, as shown in blue. (B) Infinitesimal parametric phase response curve (ipPRC) for KL001-mediated stabilization of nuclear CRYs.

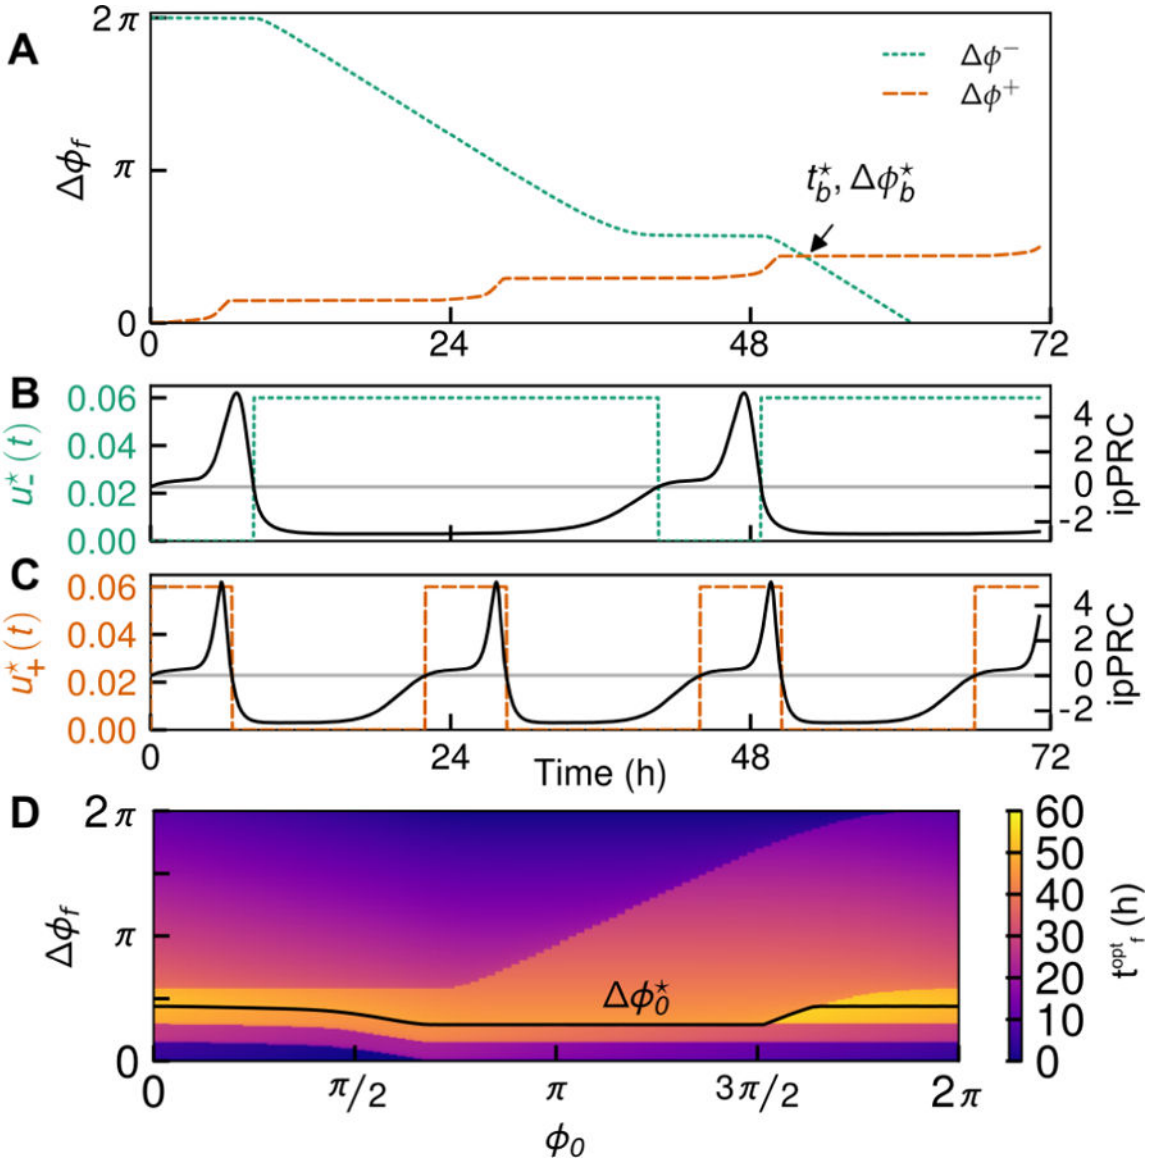

Figure 3:

Optimal control of the circadian clock by use of the ipPRC shown in Fig. 2B. (A) Maximal phase advances $\Delta\phi_+^\star(t)$ and delays $\Delta\phi_-^\star(t)$ for the ipPRC, with $\phi_0 = 0$ and $u_{max} = 0.06$, $u_{min} = 0.00$. The respective maximal phase shifts meet at time $t_B^\star$ and phase shift $\Delta\phi_B^\star$ as marked on the graph. All phase shifts may be achieved in time less than or equal to $t_B^\star$. Note that $t_f^{opt}(0, \Delta\phi_f)$ can be found on this plot by seeking the minimum time at which the oscillator crosses the desired $\Delta\phi_f$ (for example, $t_f^{opt}(0, 3\pi/2) \approx 20$ h, and is achieved by a phase delay). (B,C) Plot of optimal inputs $u_-^\star(t)$ and $u_+^\star(t)$, respectively, for the phase delay and advance shown in A. (D) Heatmap of optimal time to reset the clock t_f^{opt} as a function of initial phase

as desired shift. Note that $\Delta\phi_B^\star$, the boundary between advances and delays, is dependent on ϕ_0.

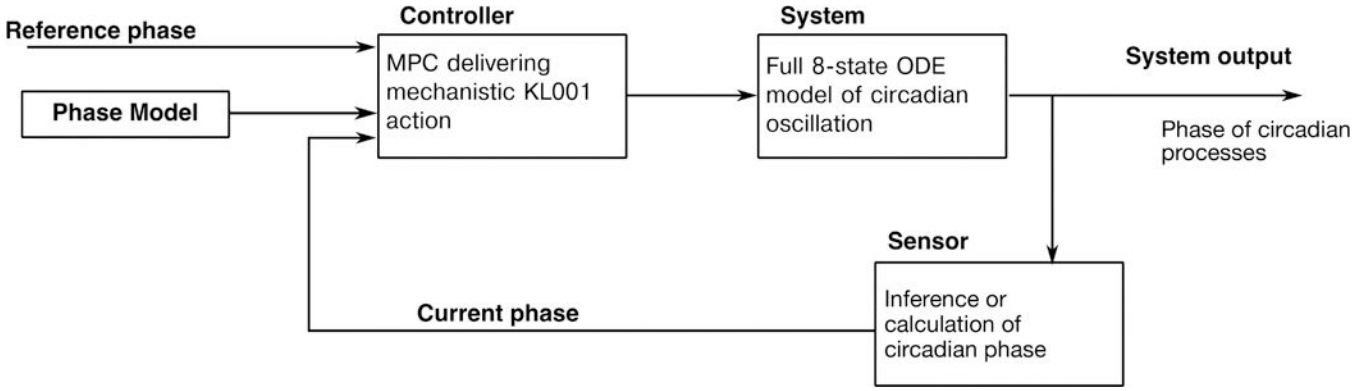

Figure 4:

In silico scheme for MPC of circadian rhythms. In experimental or clinical application, a human, organism, or cell culture would replace the 8-state ODE model as the system under control, and phase would be estimated rather than calculated.

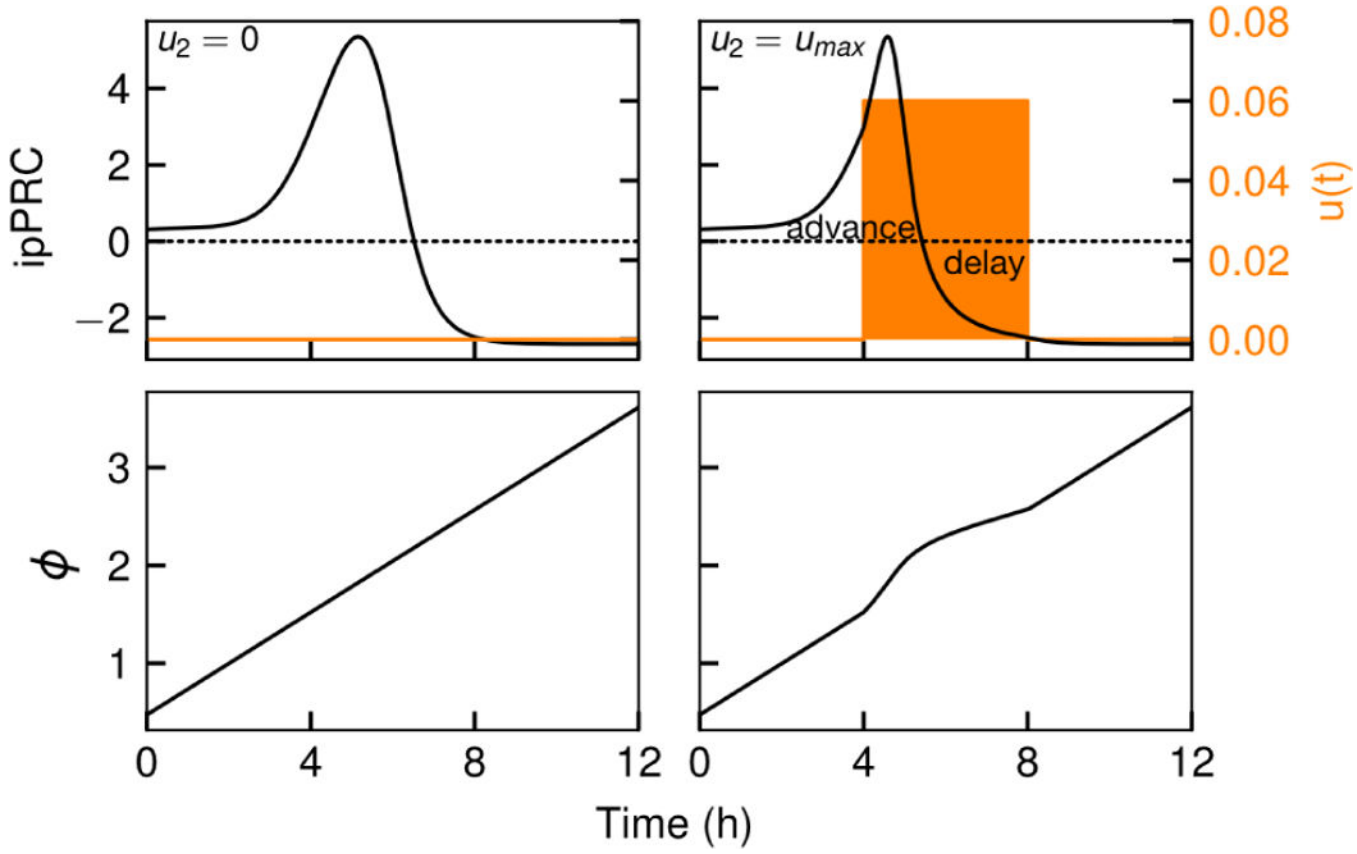

Figure 5:

Error in applying the optimal control policy occurs at the zero-crosses of the ipPRC. For a sampling time of 4 h, the zero-cross occurs in the second step $t_1 < t^0 < t_2$, and so u_2 will necessarily violate the continuous-time optimal control law. For the cases shown here, $\Delta\phi_2(u_{\max}) = \Delta\phi_2(u_{\min}) = 0$, and the resulting phase at the end of step 2 is identical due to cancellation of the advance by the delay. Finally, it is noteworthy that this single pathological step consumes most of the positive region of the ipPRC, suggesting that $\tau = 4$ h is a poor choice for sampling time if phase advances are desirable.

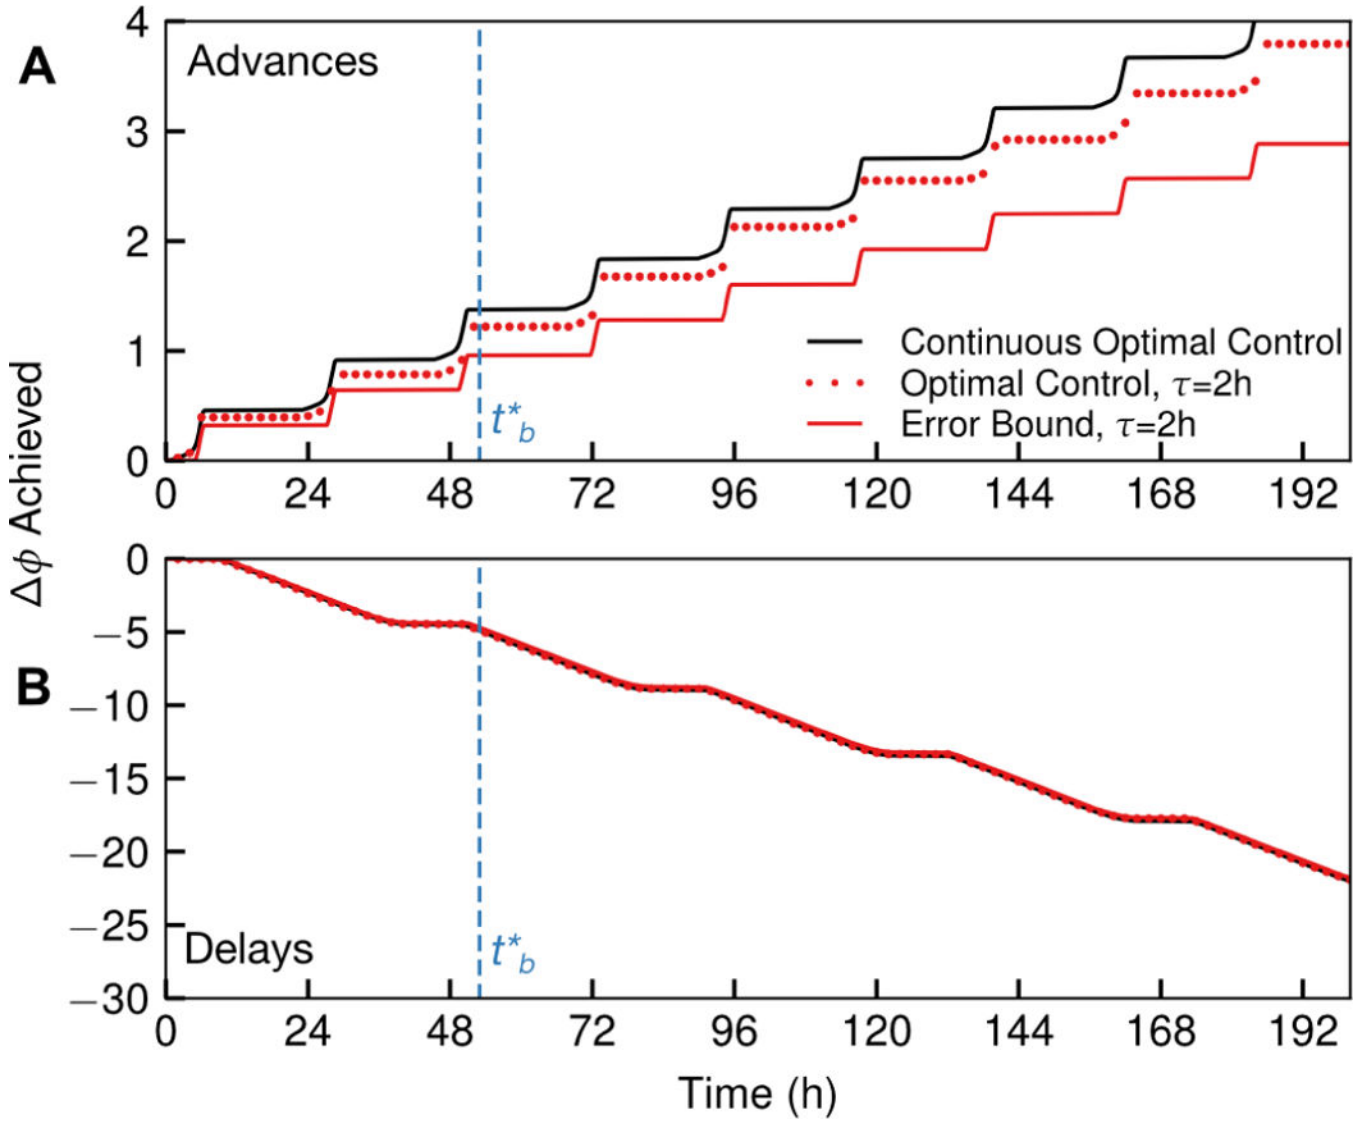

Figure 6:

The phase shift accumulated in the continuous-time optimal control (black line) plotted as a function of time, in comparison to the lower bound on phase shift accumulated (red line) and the actual discrete-time phase shift accumulated (red dots), for $\tau = 2$ h and $\phi_0 = 0$. The residual phase error is the difference between the continuous time optimal control and the actual discrete-time control. Phase advances (A) and delays (B) are shown explicitly. The bound deviates further from the optimal control at each completed cycle, however, only the shifts that occur before $t_B^\star$ (dashed line) are attempted under the optimal control policy.

Additionally, the error is more severe for phase advances, due to the relatively smaller positive region of the ipPRC, the fact that phase advances incur zero crosses more rapidly due to shortening the period, and the proximity of that positive region to a large negative region. For panel (B), these lines are nearly on top of one another. The numerically-calculated optimal control for $\tau = 2$ h indeed obeys the bounds derived in Theorem 1 from the continuous time optimal control.

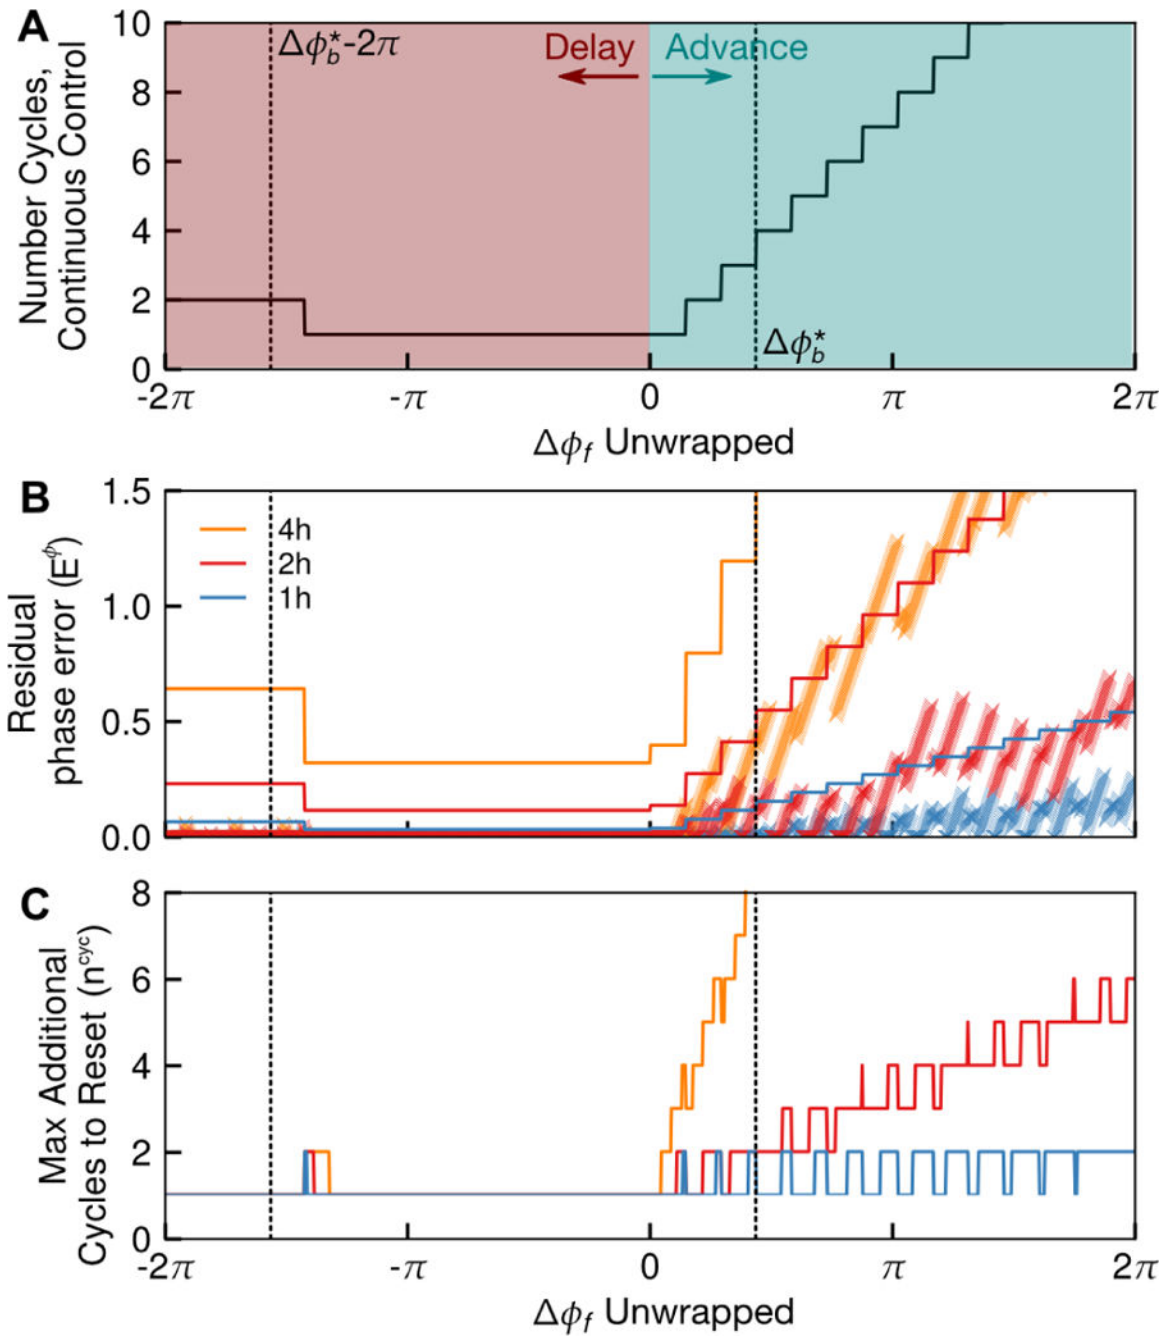

Figure 7:

Sampling time τ affects circadian phase resetting for optimal control with evenly-spaced switching times, a common feature of feedback control and MPC. (A) Plot of number of cycles required to achieve each phase advance or delay under continuous-time optimal control. Errors due to switching timing occurs where the ipPRC crosses 0, and thus is residual phase error is a function of the number of cycles required for continuous-time optimal control. (B) A bound on E_τ^ϕ, the distance to the nal phase that will remain at $t = t_f^{\text{opt}}$ under optimal control with a constant sampling time is derived from the continuous optimal

control trajectory. Here, the derived bound is shown as a line, and the E_τ^ϕ for numerically calculated optimal control solutions are shown as x markers for each $\boldsymbol{\tau}$ (plotted as discrete points for ease of visualization). The residual phase error in each case indeed obeys the theoretical bounds. (C) A bound on n_{cyc}^{add} is also derived. For this example, 1 h and 2 h shifts yield similar residual phase error and allow a complete reset in at most two additional cycles, under $\Delta\boldsymbol{\phi}_b = \Delta\boldsymbol{\phi}^\star$. Thus, either is a reasonable choice for sampling time, though a 2 h sampling time will reduce the cost of sampling by half. Alternately, if this drug were used only to achieve phase delays, a 4 h sampling rate would be similarly suitable.

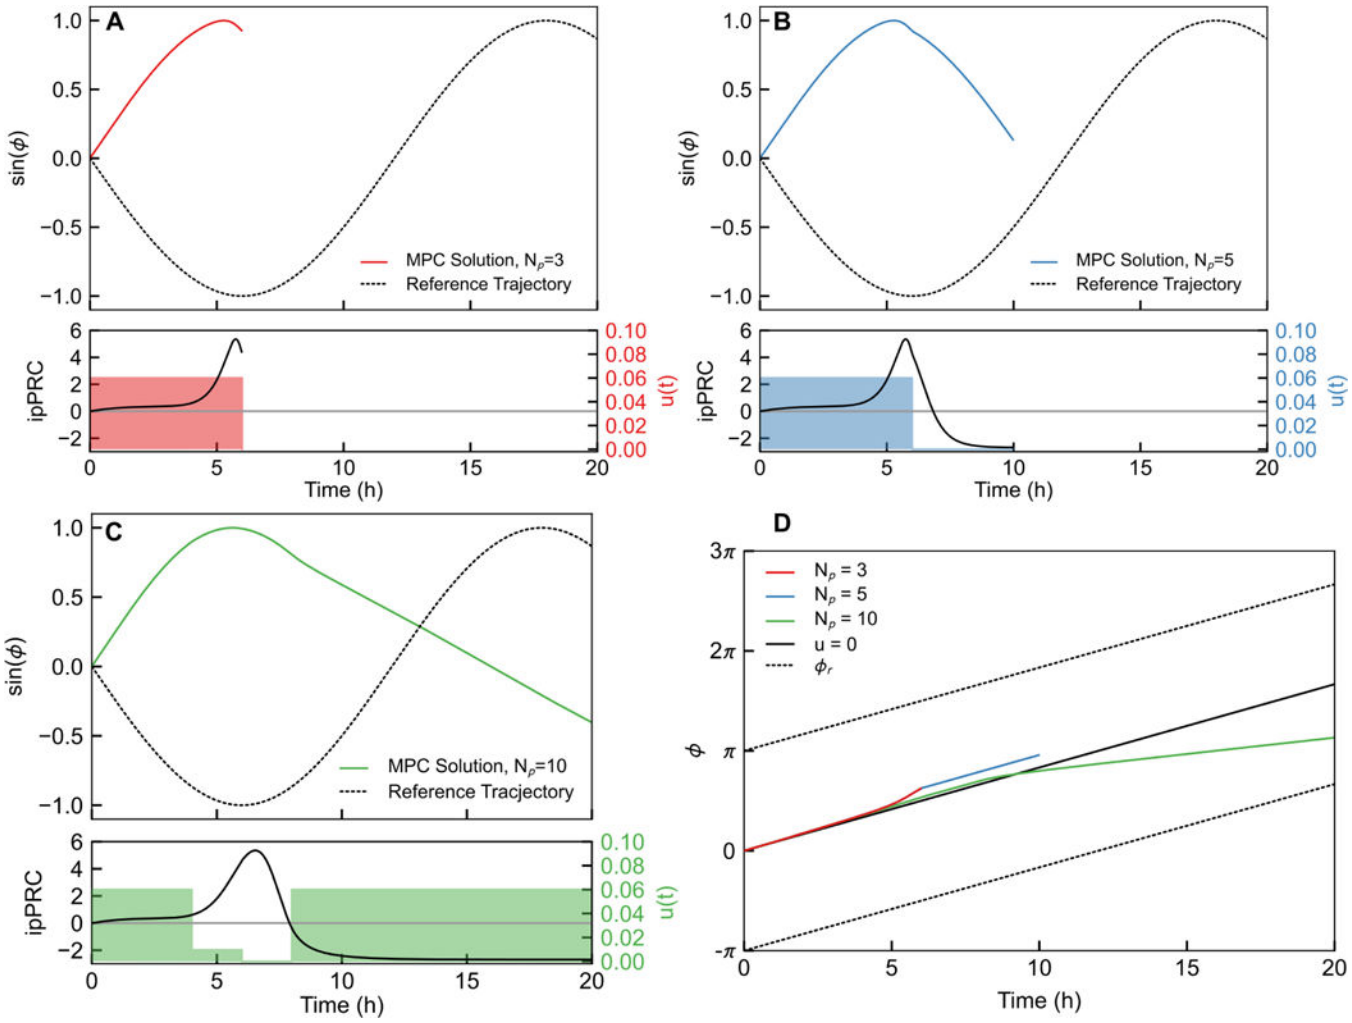

Figure 8:

Changes in the prediction horizon affect the finite-horizon optimal control trajectory based on the observable ipPRC. This is demonstrated by computing the first finite-horizon optimal control trajectory and varying that horizon. (A-C) Finite-horizon optimal control trajectories for $N_p = 3, 5,$ and 10, respectively. In all cases, $\tau = 2$h, $\Delta\phi_f =$ and $\phi_0 = 0$. Note that not only do the finite horizon optimal controls computed by the MPC differ, they seek to achieve the same shift by either advances (A,B) or delays (C). The in finite-horizon optimal control is achieved via delays ($\Delta\phi_f > \Delta\phi^*$), and so A or B would lead to excessively-long resetting by either selecting advances to achieve the shift, or first selecting advances then later choosing delays as in C. (D) This result is visualized by plotting the phase progression for each MPC case and for the 0-input case. Phase advances evidently yield slower progress toward the reference phase, and thus short prediction horizons choose ineffectively. This complication in MPC problem formulation may be circumvented by providing the controller with the optimal resetting direction *a priori*.

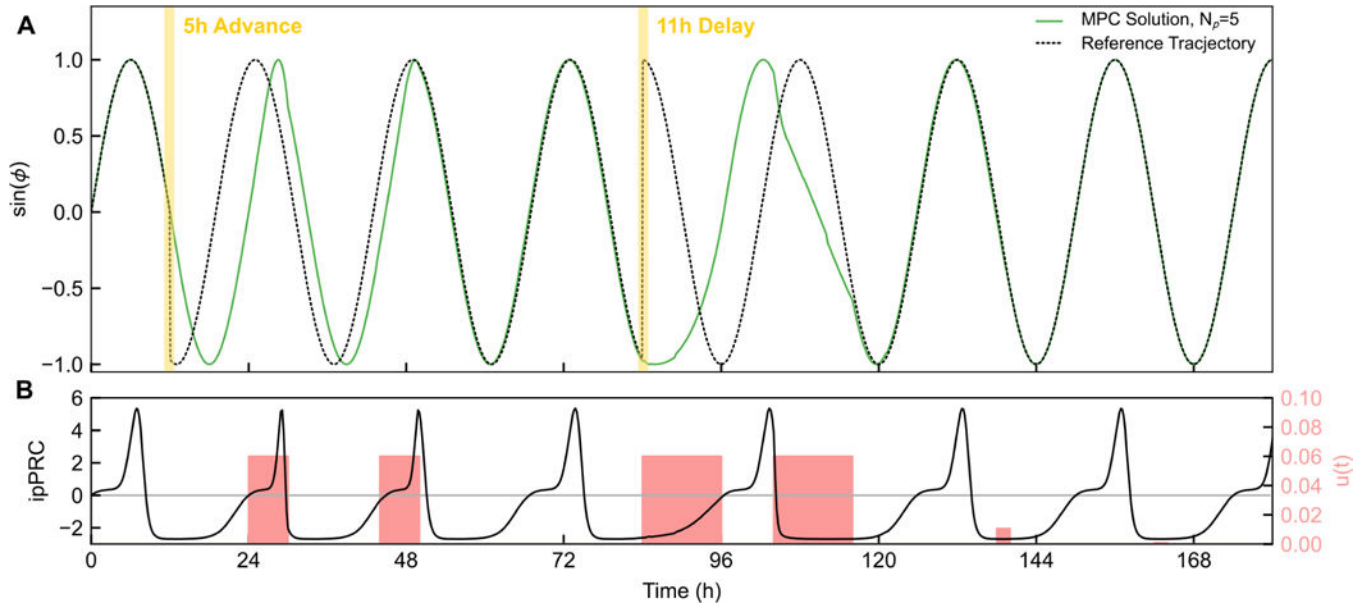

Figure 9:

Demonstration of MPC for phase resetting in response to jet lag. This scenario involves a 5h (0.417π rad) phase advance followed by a 11h (0.917π rad) phase delay, the equivalent of flying from Boston to London, then London to Honolulu three days later. (A) Phase of oscillator under MPC compared to reference phase (local time) for the jet lag problem. For this problem, $\tau = 2$h and $N_p = 5$. (B) Timing of control inputs and ipPRC throughout the simulation. In both cases, the reset is completely achieved in less than 48h, a drastic speedup in comparison to the untreated case or the light-input case [33]. For simplicity, no light input to the clock was assumed in this case, though it may be incorporated as an additional disturbance input to the phase model, or applied as part of multi-input control.

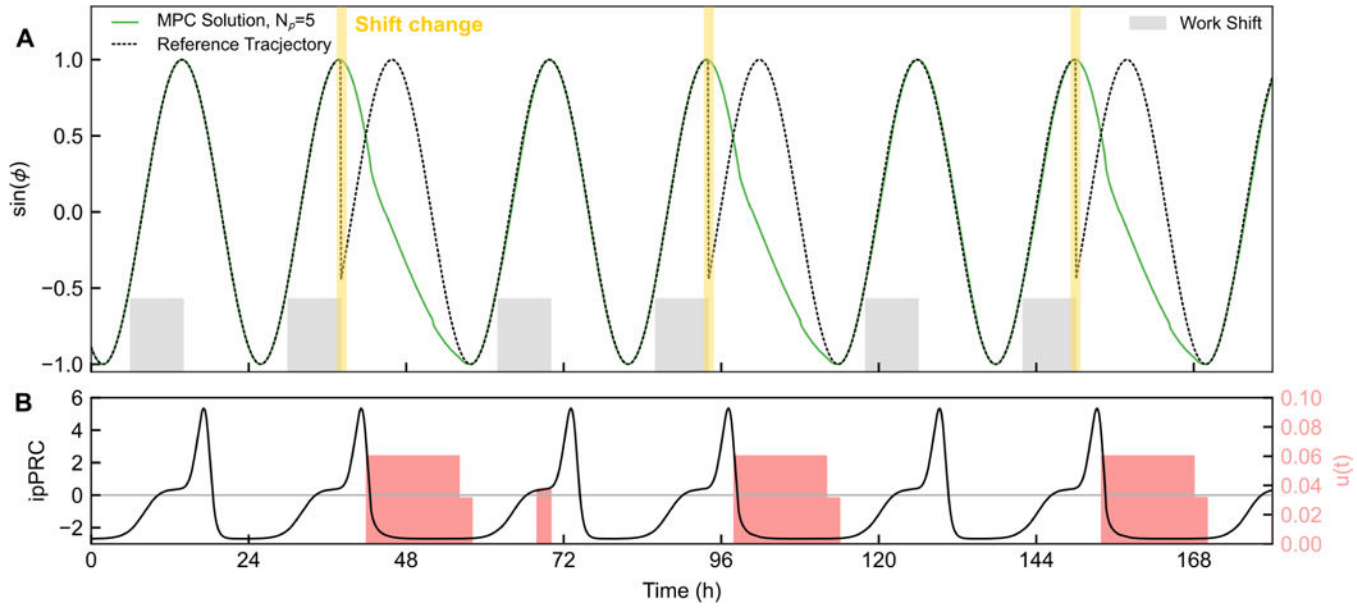

Figure 10:

Demonstration of MPC for phase resetting in response to a weeklong rotating shift work schedule, described in [43]. This scenario involves 8h phase delays every two days when rotating from morning (06:00–14:00), to evening (14:00–22:00) to night (22:00–06:00) shifts, followed by days off. (A) Phase of oscillator under MPC compared to reference phase for the shift work problem. In this example, time corresponds to clock time, such that *Per* expression in an individual entrained for morning work will occur near dusk. The reference phase was adjusted to the next rotating shift following the last work cycle of the prior rotating shift. For this problem, $\tau = 2$h and $N_p = 5$. (B) Timing of control inputs and ipPRC throughout the simulation. Because all shifts are achieved by delays, a slower sampling rate (longer duration between dosage changes) would have a similar efficacy. For simplicity, no light input to the clock was assumed in this case, though it may be incorporated as an additional disturbance input to the phase model, or applied as part of multi-input control.