

Chapter 9

Controlling Biological Time: Nonlinear Model Predictive Control for Populations of Circadian Oscillators

John H. Abel, Ankush Chakrabarty and Francis J. Doyle III

Abstract In mammals, circadian regulation of gene expression is accomplished within each cell through a transcriptional oscillator commonly modeled by a limit cycle. There has been recent interest in regulating this oscillator by delivering doses of pharmaceuticals or light in a systematic manner. Generally, controller design for circadian manipulation has been formulated by considering the dynamics of a single oscillator representing the average dynamics of the population. We illustrate in this paper that such an approximation can result in desynchronization of circadian oscillators even if the mean dynamics attain desired behavior, due to the range of dynamic responses elicited among oscillators in a population with nonidentical phases. To address this issue, we present herein nonlinear MPC for control of phase and synchrony within a population of uncoupled circadian oscillators, by explicitly predicting the evolution of the phase probability density function. We then demonstrate in silico phase shifting of an example oscillator population while maintaining a high degree of synchrony. The MPC strategy formulated herein is a step toward a detailed, systems approach integrating population effects, pharmacokinetics and pharmacodynamics, and interactions between different oscillator populations.

J. H. Abel
Department of Systems Biology, Harvard Medical School,
Boston, MA 02115, USA
e-mail: johnhabel@g.harvard.edu

A. Chakrabarty
Harvard John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, USA
e-mail: achakrabarty@g.harvard.edu

F. J. Doyle III (✉)
Harvard John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, USA
e-mail: frank_doyle@seas.harvard.edu

F. J. Doyle III
Division of Sleep Medicine, Harvard Medical School, Boston, MA 02115, USA

9.1 Introduction

Disruption to the circadian clock, such as misaligned light or feeding, or extended night shift work, has been associated with a wide range of disorders such as decreased cognitive function [6, 13, 15]. Circadian rhythms are endogenous oscillatory processes involved in biological timekeeping and temporal regulation of biological function in nearly all forms of life [7]. Although the mechanism driving these oscillations varies between organisms, circadian rhythms share the characteristics of being endogenously generated, entrainable to environmental rhythms, and maintain a near-24h periodicity. In mammals, a transcriptional–translational oscillator present in nearly every cell type regulates transcriptional architecture, and along with further posttranscriptional regulation, intricately coordinates cellular and tissue function [23]. These cell-autonomous oscillators regulate gene expression within local tissue and, in turn, are coordinated by signals from the hypothalamic suprachiasmatic nucleus (SCN) which responds to light cues to entrain to the environment [26]. The sleep-wake cycle is the most immediately recognizable feature of circadian regulation, however, it is estimated that 5–20% of all transcripts within any mammalian cell type oscillate with a circadian frequency [28], and circadian regulation is broadly integrated with metabolic function [6, 16].

Low-amplitude circadian oscillation within tissue may result from perturbations to the oscillator driving lower amplitudes within individual cells, or driving desynchronization of the oscillators comprising the tissue population [22, 25]. Meanwhile, high-amplitude circadian oscillations are associated with good metabolic health [17]. Thus, there is a need to develop therapies to mitigate the effects of disruption to the circadian clock, to rapidly realign circadian phase with the environment following jet lag or shift work to avoid cognitive difficulty, and to promote high-amplitude circadian oscillation. Light, and more recently, small-molecule pharmaceuticals [10] provide possible two paths toward control of the clock. Since circadian rhythms are highly complex phenomena, phase and amplitude control of the clock may be achieved by a control theoretic approach to timing the delivery of drugs or light [5, 18, 27].

Control strategies devised for the manipulation of mammalian circadian rhythms may analogously be applied to other oscillatory biological systems, provided they are well described by limit cycle oscillators. Circadian oscillators in other species such as the KaiABC system in the cyanobacterium *Synechococcus* [12], or the *Period-Timeless* oscillator in *Drosophila* [8] have been modeled as limit cycle oscillators for more than a decade. More generally, genetic or phosphorylation-driven oscillators are a ubiquitous biological motif involved in the metabolic processes of numerous organisms from prokaryotes to mammals [9, 11], and the development of strategies for manipulating these systems is broadly desirable.

Several previous studies have examined the phase control of circadian rhythms. Several recent studies have focused on the application of light to drive phase shifts in the *Drosophila* [19, 20] or mammalian [18, 27] clock using an optimal control approach. However, solving the resulting optimal control problem is computationally

prohibitive due to the inherent nonlinearity of the model, and the computed optimal control trajectory is susceptible to modeling errors and disturbances. Iteratively solving finite-horizon optimal control problems (as in model predictive control (MPC)) has provided demonstrably superior control performance. For example, MPC was used for shifting of the *Drosophila* clock [4, 5] or the mammalian clock [2]. This work uses nonlinear MPC as in [2] as a starting point for developing control of an oscillator population. These studies of the control of biological oscillators have focused on control of a single oscillator despite notable mismatch between single-cell and population dynamics [14]. Additionally, these studies have primarily focused on the application of light to the clock. Recently, pharmaceuticals have presented advantages for control of the clock, as they are much less invasive than strict control of an individual's light environment, and may be delivered at any time of day. A more detailed approach to control of circadian dynamics would integrate tissue or population-scale effects, pharmacokinetics and pharmacodynamics, and interactions between different oscillator populations.

Herein, we begin to approach this detailed formulation by presenting a MPC framework for manipulating phase and synchrony within populations of uncoupled biological oscillators using a pharmaceutical agent, and demonstrate the efficacy of such an approach by *in silico* simulations of phase shifting in the mammalian clock. We do so by describing the population of oscillators as a phase probability density function (PDF), as in [22], and using the parametric infinitesimal phase response curve (PRC) to calculate how the phase PDF evolves in response to control input. By using a predictor that accounts explicitly for the variability in phase within a synchronized population, we are able to maintain synchrony of oscillators while performing a phase shift, unlike the single-oscillator case.

Chapter Overview

In this chapter, we present a framework for the control of a population of biological oscillators, motivated by the example of the mammalian circadian clock. In Sect. 9.2.1, we present an established model of the mammalian circadian oscillator and its response to the small molecule KL001, an early [10, 21]. In Sect. 9.2.2, we motivate and formulate the phase shifting control problem. We use a previously derived simplification of oscillator dynamics using the parametric phase response to a control input to reduce the dynamical model to a phase-only representation [2]: this is advantageous in reducing the dimensionality of the MPC problem, thereby curbing computational effort. In Sect. 9.2.3, we describe a controller for the case for a single oscillator, and demonstrate its function *in silico* in Sect. 9.2.4. In Sect. 9.3.1, we demonstrate that although this formulation may successfully control a single oscillator, applying this controller to manipulate mean phase of a population of oscillators may effect a desynchronization detrimental to biological function. We then modify the MPC problem in Sect. 9.3.2 using a probability density function of population phase in conjunction with the simplified dynamics to exert simultaneous control over phase and synchrony within an oscillator population. We demonstrate in Sect. 9.3.3

the *in silico* efficacy of this approach in maintaining synchrony throughout a phase realignment. Finally in Sect. 9.4, we discuss limitations of this approach, and challenges to its implementation *in vivo* and *in vitro*.

9.2 Control of a Single Circadian Oscillator

A standard approximation in the control of the circadian oscillator is describing the targeted circadian system (i.e., the population of cells comprising clocks in brain or peripheral tissues such as the liver) as a single limit cycle oscillator [2, 4, 5, 18–20, 27]. We begin our treatment of applying control to phase shift the circadian oscillator by exploiting this approximation, as in our previous work [2], to illustrate where it fails to capture population-scale phenomena.

9.2.1 Modeling the Circadian Oscillator

The mammalian circadian oscillator within an individual cell is comprised of interlocked transcriptional–translational feedback loops. The core negative feedback loop involves isoforms of the genes *Period* (*Per*) and *Cryptochrome* (*Cry*), which form PER-CRY heterodimers and enter the nucleus to bind to BMAL1-CLOCK E-box activators to repress their own transcription. As these repressors are degraded, BMAL1-CLOCK dimers activate transcription of *Pers* and *Crys*, resulting in a self-sustained oscillation. Downstream of the circadian feedback loops, clock components regulate transcriptional architecture through D-box, E-box, and ROR-binding elements. For an excellent review of the genetic components of the mammalian oscillator, see [23].

Numerous dynamical models of the circadian oscillator have been proposed. In this work, we selected the model from [10, 21], as it was created to identify the effect of the small-molecule KL001 on the mammalian oscillator. This model consists of 8 nonlinear coupled ODEs and 21 kinetic parameters, fully described in the supplement to [21] and not reproduced here due to space considerations. Because KL001 was found to stabilize nuclear PER-CRY transcription factors, control is implemented in the model by modifying the ODEs describing the degradation of PER-CRY dimers as follows:

$$\frac{dC1P}{dt} = v_{a,CP}P \cdot C1 - v_{d,CP}C1P - \frac{(vdCn - u(t))C1P}{k_{deg,Cn} + C1P + C2P} \quad (9.1a)$$

$$\frac{dC2P}{dt} = v_{a,CP}P \cdot C2 - v_{d,CP}C2P - \frac{(vdCn - u(t))m_{C2N}C2P}{k_{deg,Cn} + C2P + C1P}, \quad (9.1b)$$

where $u(t) \in [0, \bar{u}]$ is control input at time t , which reduces degradation rate $vdCn$. Generally \bar{u} should not be greater than $vdCn$, as a degradation reaction cannot be reversed so far as to synthesize new PER-CRY. Throughout, we set this value to 0.08. The states present in this model, and the effect of KL001, are shown in the schematic in Fig. 9.1a.

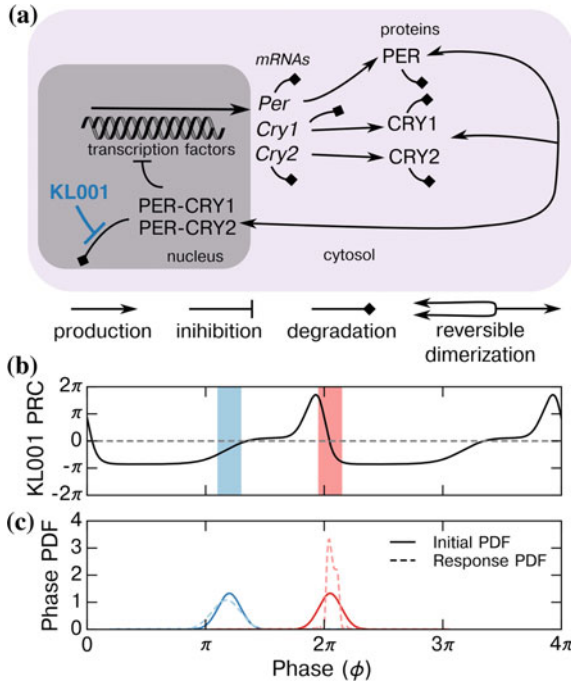


Fig. 9.1 Schematic of mammalian circadian oscillator and effect of KL001. **a** Diagram of the core negative feedback loop driving mammalian circadian rhythms. All states and reactions shown are included explicitly in the model [21]. **b** Parametric infinitesimal PRC describing response to KL001 (reduction in parameter $vdCn$). Note that this is double plotted to allow visualization. **c** Synchrony of a population is affected by timing of KL001 application. Two example probability density functions (PDFs) of phase are plotted before (solid) and after (dashed) application of KL001 in silico at the phases shown. In regions of positive slope (blue), the phase PDF is dispersed, whereas in regions of negative slope (red), the PDF is condensed, even though the mean phase is unchanged

9.2.2 Model Reduction and the Phase Control Problem

Control of the circadian clock is primarily focused on shifting the phase of the circadian oscillator. A unique phase $\phi \in [0, 2\pi)$ may be assigned to each unique point on the limit cycle. An unperturbed oscillator on the limit cycle will advance in phase at a constant rate. Thus, points in state space that are not on the limit cycle may be assigned the phase of the point on the limit cycle to which they converge asymptotically in time [2].

The 8-ODE oscillator may be reduced to a single-ODE describing phase [24]:

$$\frac{d\phi}{dt} = \frac{2\pi}{\tau} - u(t) \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{x^v}, \tag{9.2}$$

where τ is the period of oscillation (set to 24 h here), $u(t)$ is the control input, and $-\frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{\mathbf{x}^\nu}$ is the parametric infinitesimal phase response curve (PRC), the first-order Taylor approximation of the phase response with respect to changing parameter $vdCn$, evaluated on the limit cycle trajectory \mathbf{x}^ν . The negative sign here is included in the PRC to reflect that $u(t)$ is defined to have a positive value. Importantly, the PRC is itself a function of phase, resulting in a nonlinear ODE. This function may be calculated numerically in advance from the 8-ODE model by previously defined methods [2, 24]. The PRC is plotted in Fig. 9.1b. One may consider the control input to be either “speeding up” or “slowing down” the oscillation depending on the sign of the PRC. We opted to set the reference $\phi = 0$ where the concentration of *Per* mRNA is at a maximum.

As in [2], we are aiming to align the phase of the circadian oscillator (ϕ) with an external tracking phase (denoted ϕ_e), for example, the phase of the environment before and after a plane flight through multiple time zones or before and after beginning a shift work cycle. This externally imposed phase may be captured by a single ODE as well:

$$\frac{d\phi_e}{dt} = \frac{2\pi}{\tau_e} + \Delta\phi \delta(t_{shift}), \quad (9.3)$$

where τ_e is the period of the environment (set to 24 h), and $\Delta\phi \in [-\pi, \pi)$ is the phase shift of the environment that occurs at time t_{shift} , for example, when disembarking from the flight or starting a period of shift work. We are searching for a control trajectory that will drive the oscillator ϕ to external phase ϕ_e and maintain it at ϕ_e for all subsequent time. That is:

$$\lim_{t \rightarrow \infty} \|\phi(t, u) - \phi_e(t)\| = 0.$$

9.2.3 Control of Single-Oscillator Phase

While light may be applied or removed instantaneously, pharmaceutical perturbation of the clock is constrained by pharmacokinetics. To alleviate numerical complications arising due to unidentified (and likely nonlinear) pharmacokinetics, we selected a piecewise constant parameterization of the control:

$$u(t) = u(t_k) \forall t_k \leq t < t_{k+1}, \quad (9.4)$$

and denote the sampling time $t_{k+1} - t_k$ as τ_u . For a predictive horizon of N_p steps of duration τ_u , we define

$$U \triangleq [u(t_k) \ u(t_k + \tau_u) \ \cdots \ u(t_k + (N_p - 1)\tau_u)]^\top \quad (9.5)$$

as the knots of the control trajectory defined at each of the N_p steps. To estimate the phase at the end of each step, the phase dynamics in (9.2) may be integrated over each of the $\ell \in [1 \cdots N_p]$ steps to yield:

$$\hat{\phi}(t_i + \ell\tau_u) = \phi(t_i) + \frac{2\pi\ell\tau_u}{\tau} - \sum_{k=0}^{\ell-1} \int_{t_{k+i}}^{t_{k+1+i}} u(t_k) \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{\mathbf{x}^v} dt, \quad (9.6)$$

where $\hat{\phi}$ is the predicted phase, t_i is the current time, $\phi(t_i)$ is measured, and the PRC is a function of $\hat{\phi}$. We define the phase error to be the magnitude of the phase difference between the predicted phase and the environmental phase:

$$e_\phi(\cdot) \triangleq \left| \angle \left(\exp(i\hat{\phi}(\cdot)) - i\phi_e(\cdot) \right) \right|, \quad (9.7a)$$

where $i = \sqrt{-1}$. Computationally, driving phase error to 0 is numerically unstable, so we relax the terminal constraint by ignoring phase error below a constant δ_ϕ :

$$g_\phi(\cdot) \triangleq \begin{cases} 0 & \text{if } e_\phi(\cdot) < \delta_\phi \\ e_\phi(\cdot) & \text{otherwise} \end{cases} \quad (9.7b)$$

so that numerical imprecisions do not result in controller action. Thus, the finite-horizon optimal control problem at each time t_j may be solved for the optimal trajectory u^* :

$$u^* = \arg \min_U \sum_{\ell=1}^{N_p} w_\ell^\phi g_\phi^2(t_i + \ell\tau_u) + w_{\ell-1}^u u^2(t_i + (\ell-1)\tau_u)$$

subject to:

$$\hat{\phi}(t_i + \ell\tau_u) = \phi(t_i) + \frac{2\pi\ell\tau_u}{\tau} - \sum_{k=0}^{\ell-1} \int_{t_{k+i}}^{t_{k+1+i}} u(t_k) \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{\mathbf{x}^v} dt, \quad (9.8)$$

$$0 \leq u_{\ell-1} \leq \bar{u},$$

for all $\ell = 1, \dots, N_p$, where w_ℓ^ϕ and $w_{\ell-1}^u$ 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. After identifying the optimal piecewise control trajectory u^* , we applied $u^*(t_i)$ to the full 8-state ODE model for $t \in (t_i, t_i + \tau_u]$ as is standard in model predictive control.

For this MPC formulation, w^ϕ , w^u , N_p , and τ_u are design parameters which may vary. We have selected:

$$\begin{aligned} w_\ell^\phi &= \ell, \\ w_{\ell-1}^u &= 1, \end{aligned}$$


for $\ell = 1, \dots, N_p$. We have elsewhere investigated optimal selection of design variables N_p and τ_u for this exact formulation of MPC [1]. Based on these findings, we set $N_p = 3$ (and $N_u = N_p$ where N_u is the number of length of the control horizon) and $\tau_u = 2$ h, and instead studied the how phase control of a single oscillator differs from an oscillator population.

9.2.4 Case Study #1: Nonlinear MPC for a Single Oscillator

To demonstrate the behavior of this controller, we applied it for the phase shifting control problem where a phase delay of 6 h occurs at 12 h (a phase delay of $\pi/4$ at a phase of π). The environmental phase for this example was given by (9.3), where $\Delta\phi = -\pi/2$, and $t_{shift} = 12$ h. We used the Python language to solve the MPC problem, specifically, we used CasADi [3] and SciPy for formulation, and PySwarm to solve the nonlinear programming optimization problem. Here, the numerical error in calculating phase was $\mathcal{O}(10^{-2})$ and so design variable $\delta_\phi = 0.1$.

The results of this simulation are shown in Fig. 9.2. Briefly, the controller began its action at $t = 12$ h to achieve a phase delay. The full 6 h phase delay could not be achieved within the initial negative region, and so the remainder of the shift was accomplished when the PRC returned to a negative value near 30 h. A decrease in amplitude occurred near $t = 30$ h due to transient deviation from the limit cycle, and the full amplitude returned for the following peak.

Fig. 9.2 Application of the controller detailed in Sect. 9.2.3 for a phase delay of 6 h ($\Delta\phi = -\pi/2$) with $t_{shift} = 12$ h. The controller acts immediately beginning at $t = 12$ h, and completes the phase delay to align with the desired reference trajectory after the PRC returns to a negative value. Note that the initial negative region of the PRC is elongated due to control slowing the advance of phase and maintaining the oscillator in a phase with a negative PRC for longer



9.3 Control of Population Phase and Synchrony

The circadian oscillator is cell-autonomous, and each tissue is comprised of many thousands of individual oscillators. While the SCN master pacemaker maintains its synchrony through intercellular communication, other tissues such as the liver that lack paracrine signaling are kept synchronized through a variety of identified and as-yet unidentified means, as guided by the SCN [26]. The application of a pharmaceutical to any of these populations will affect this synchrony, and thus the amplitude of oscillation [22, 25].

While transient deviations from the limit cycle eventually return to the limit cycle amplitude, reduction in amplitude due to desynchrony will persist in the absence of signaling, making populations with weak coupling susceptible to long-term desynchrony from mistimed control. The change in synchrony in response to perturbation may be calculated from the PRC [22]. In Fig. 9.1c, we show how this may be intuitively understood based on the slope of the PRC: oscillator populations that lie on regions of positive slope result in the cells which are ahead in phase advancing further (or being delayed less) than oscillators which lag to begin with, broadening the probability density function (PDF) of phase. Inversely, populations lying on a region of negative slope are condensed in phase by a similar argument.

Here, we first apply the previously described controller to a population of uncoupled oscillators to demonstrate the deleterious effect of a population-agnostic nonlinear MPC on synchrony. We then modify the MPC controller to explicitly penalize population desynchronization and demonstrate the ability to simultaneously control phase and synchrony of a population.

9.3.1 Case Study #2: Limitation of Single-Oscillator Assumption


We first applied the controller from Sect. 9.2.3 to a population of oscillators, with slight modification for tracking the mean phase of the population. We modified the predictor in (9.6) to use the mean phase of the population $\bar{\phi}$ rather than the phase of an individual oscillator:

$$\hat{\phi}(t_i + \ell\tau_u) = \bar{\phi}(t_i) + \frac{2\pi\ell\tau_u}{\tau} - \sum_{k=0}^{\ell-1} \int_{t_{k+i}}^{t_{k+1+i}} u(t_k) \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{x^y} dt. \quad (9.9)$$

Here, we calculated $\bar{\phi}(t_i)$ from the phases of each oscillator ϕ_n in the population by the parameter $z \in \mathbb{C}$ describing the population:

$$z = \rho \exp(i\bar{\phi}) = \frac{1}{N} \sum_{n=1}^N \exp(i\phi_n), \quad (9.10)$$

Fig. 9.3 Application of the controller from Sect. 9.2.3 for a phase delay of 6 h ($\Delta\phi = -\pi/2$) with $t_{shift} = 12$, for shifting mean phase of a population of 200 circadian oscillators (individual trajectories plotted in gray). While control is applied in nearly the same regions as Fig. 9.1 and a -6 h shift was attained, this resulted in a dispersion of phase, a decrease in synchrony index, and a reduction in mean oscillatory amplitude



where ρ is the Kuramoto order parameter, or colloquially, the synchrony index. We emphasize that the controller did not observe any information about the population aside from its mean phase, and as such, the predictor (9.6) used in the finite horizon optimal control problem was imprecise due to slight differences between single-cell and population mean phase response [22]. The controller was otherwise parameterized identically to that in Sect. 9.2.3.

Figure 9.3 shows the result of applying this controller to a population of 200 identical uncoupled oscillators with initial phases sampled from the PDF:

$$p(\phi) = f_{WN}(\phi; \phi_0, \sigma), \quad (9.11)$$

where f_{WN} indicates a wrapped normal distribution with mean ϕ_0 (set to 0, to match the mean with the single oscillator case), and standard deviation σ (set to $\pi/12$, approximately capturing the distribution of phases observed) [22]. As in the single-cell case, control was applied immediately starting at $t = 12$ h to begin correcting for the 6 h phase delay. In the population case, this corresponds to a region of positive slope of the PRC, and intuitively resulted in a desynchronization of phase as evidenced by the decline in synchrony index for the duration of the KL001 pulse. Despite the amplitudes of the individual oscillators (gray) returning after a transient to their pre-pulse levels, the population mean [PER] amplitude was reduced by approximately one-third due to desynchronization of the population. Because there was no intercellular or external communication driving synchrony, the amplitude of the mean remained diminished for the duration of the simulation.

9.3.2 Population MPC Algorithm for Phase Coherence

A more sophisticated approach to control of an oscillator population involves predicting the evolution of the PDF itself rather than only mean phase. Methods have been developed previously to compute the change in PDF directly in response to stimulation [22]. A change in variables allows the numerical computation of the new phase PDF:

$$\hat{p}(\phi, t)dh(\phi) = p(\phi, t)d\phi, \quad (9.12)$$

where $\hat{p}(\phi, t)$ is the predicted PDF at time t for phases ϕ , and $h(\phi) = \phi + \Delta\phi$, called the phase transition curve (PTC), is the total response to perturbation. We may therefore revise the prediction model to explicitly involve calculation of the phase PDF at each step of the predictive horizon. The PTC for each step within the predictive horizon may be calculated in a similar fashion to the predictor for the single-oscillator case:

$$h_k(\phi) = \phi + \frac{2\pi\tau_u}{\tau} - u(t_k) \int_{t_k}^{t_{k+1}} \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{x^*} dt, \quad (9.13)$$

where u is piecewise constant and the integrand is a function of ϕ which may be calculated numerically in advance as previously defined. This function may be calculated numerically for each step, and used to calculate the evolution of the PDF:

$$\hat{p}(\phi, t_{k+1})dh_k(\phi) = \hat{p}(\phi, t_k)d\phi. \quad (9.14)$$

To calculate the first step, $\hat{p}(\phi, t) = p(\phi, t)$ is measured from the population under control. The population predicted mean phase $\hat{\phi}$ and synchrony index $\hat{\rho}$ may then be calculated from the predicted PDF:

$$\hat{\rho}(t_k) \exp(i\hat{\phi}(t_k)) = \int_0^{2\pi} \hat{p}(\theta, t_k) \exp(i\theta) d\theta, \quad (9.15)$$

where θ is a dummy variable of phase. The phase error term g_ϕ may remain as defined in (9.7), and we define a similar error term penalizing desynchrony

$$g_\rho(\cdot) \triangleq \begin{cases} 0 & \text{if } \rho > \delta_\rho \\ (1 - \rho(\cdot)) & \text{otherwise} \end{cases} \quad (9.16)$$

which reduces to 0 for a satisfactorily synchronized population.

Thus, the population control problem is:

$$\begin{aligned}
 u^* &= \arg \min_U \mathcal{J} \\
 \text{subject to:} \\
 \mathcal{J} &= \sum_{\ell=1}^{N_p} w_\ell^\phi g_\phi^2(t_i + \ell\tau_u) + w_\ell^\rho g_\rho^2(t_i + \ell\tau_u) + w_{\ell-1}^u u^2(t_i + (\ell-1)\tau_u), \\
 h_k(\phi) &= \frac{2\pi\tau_u}{\tau} - u(t_k) \int_{t_k}^{t_{k+1}} \frac{d}{dt} \frac{d\phi}{d(vdCn)} \Big|_{x^y} dt, \\
 \hat{p}(\phi, t_{k+1}) dh_k(\phi) &= \hat{p}(\phi, t_k) d\phi \\
 \hat{p}(t_k) \exp(i\hat{\phi}(t_k)) &= \int_0^{2\pi} \hat{p}(\theta, t_k) \exp(i\theta) d\theta \\
 0 &\leq u_{\ell-1} \leq \bar{u},
 \end{aligned} \tag{9.17}$$

where phase error and synchrony are calculated at the end of each step within the predictive horizon.

9.3.3 Case Study #3: Implementation of Population Nonlinear MPC Controller

As before, w^ϕ , w^u , w^ρ , N_p and τ_u are design parameters. These parameters were set as in Sect. 9.2.3, with the exception of the new weighting of synchrony:

$$w^\rho = 10(\ell + 1),$$

which was set such that the synchrony term is of the same order of magnitude as the phase term, and increases to allow flexibility of synchrony early in the horizon. Tuning this parameter will adjust the sensitivity of the controller to temporary desynchrony. As $w^\rho \rightarrow \infty$, the controller will take no action unless it results in no loss of synchrony, i.e., the population lies completely on a region of negative slope of the PRC. For lower but nonzero w^ρ , as in the case here, some flexibility is permitted, in that the controller may apply action that desynchronizes the population if it results in a large reduction in phase error. Correspondingly, the controller will later resynchronize the population to account for this early desynchrony. For $w^\rho = 0$, the controller will behave as the population-agnostic case. After calculating the finite-horizon optimal control u^* , we apply the first step $u^*(t_j)$ to all oscillators within the population for $t \in [t_j, t_j + \tau_u)$, and repeat this process.

We applied the controller described in Sect. 9.3.2 to the same phase shifting problem as the previous controllers: $t_{shift} = 12$ h, $\Delta\phi = -6$ h. The initial phase of the 200 uncoupled oscillators comprising the population was sampled from the PDF:




Fig. 9.4 Application of the controller from Sect. 9.3.2 for a phase delay of 6h ($\Delta\phi = -\pi/2$) at $t_{shift} = 12$ for a population of 200 oscillators (individual trajectories plotted in gray). This controller explicitly accounted for synchrony of the population. After a brief input to begin the shift, the controller delayed the majority of its input to find a region where population synchrony would be maintained. Indeed, synchrony is slightly improved by the control action, and a phase delay of 6h was achieved

$$p(\phi) = f_{WN}(\phi; \phi_0, \sigma), \tag{9.18}$$

where f_{WN} indicates a wrapped normal distribution with mean $\phi_0 = 0$ and standard deviation $\sigma = \pi/12$.

Results from this simulation are shown in Fig. 9.4. The controller first applied input briefly in a slight desynchronizing region of the PRC, then paused for the remainder of the first negative PRC region due to the likelihood of further desynchronizing the population. Once the population PDF returned to a region of negative PRC and a negative first derivative of the PRC, the controller resumed input driving the population to its 6h phase delay and restoring full synchrony. Strikingly, changing the time of input resulted in a slight increase in the synchrony of the population in comparison to its original state, evidenced by an increase in the synchrony index. Visually, this is evident from the clear alignment of individual oscillators within the population (plotted in gray) in comparison to the dispersion in phase evident in Fig. 9.3 where synchrony was ignored.

9.4 Conclusion

We have presented a modification of nonlinear MPC for phase manipulation of circadian oscillator populations, in which a PDF of phase is used in solving the finite-horizon optimal control problem, allowing mean phase and population synchrony to be regulated simultaneously. For many PRCs that have been calculated, there exists a region where the PRC or first derivative of the PRC is zero [7], it is therefore possible to manipulate phase and synchrony independently for a population of circadian oscillators through a single control input. In reality, the ability to target these regions is limited by precision of the measurements of the PRC and population phase PDF.

One significant challenge in implementing this control algorithm in vitro or in vivo is the construction of an observer with sufficient accuracy to accurately reconstruct the phase PDF. Current methods of assessing the phase of circadian oscillators rely on either noisy single-cell bioluminescent markers in vitro or system level metrics in vivo such as melatonin, activity, or body temperature. However, even a simplistic assumption such as a wrapped normal PDF with an arbitrary estimate of standard deviation of phase could help avoid delivering control inputs where the slope of the PRC is expected to be positive, and thus help avoid desynchrony. In this case, a long predictive horizon would quickly become inaccurate, however, necessitating careful selection of design variables τ_u and N_p .

Another challenge toward implementing such an algorithm is incorporating the as-yet uncharacterized pharmacokinetics and pharmacodynamics (PK/PD) for a small molecule such as KL001. This could potentially be achieved by including these terms directly in the prediction step of the MPC. Uncertainty and individual variability in PK/PD measurements may reduce the accuracy of such an approach, however, further study is necessary to determine the extent of this variability, and how these inaccuracies affect controller performance.

It is our hope that these and similar control theoretic methods will inform the discovery of circadian therapies, and enable novel experimental design toward better understanding the dynamics of cellular populations and communication.

References

1. Abel, J.H., Chakrabarty, A., Doyle III, F.J.: Nonlinear model predictive control for circadian entrainment using small-molecule pharmaceuticals. Proceedings of 20th IFAC World Congress. 9864–9870 (2017)
2. Abel, J.H., Doyle III, F.J.: A systems theoretic approach to analysis and control of mammalian circadian dynamics. Chem. Eng. Res. Des. **116**, 48–60 (2016)
3. Andersson, J., Akesson, J., Diehl, M.: Recent Advances in Algorithmic Differentiation, vol. 87. Springer, Berlin Heidelberg (2012)
4. Bagheri, N., Stelling, J., Doyle III, F.J.: Circadian phase entrainment via nonlinear model predictive control. Int. J. Robust Nonlinear Control **17**(May), 1555–1571 (2007)
5. Bagheri, N., Taylor, S.R., Meeker, K., Petzold, L.R., Doyle III, F.J.: Synchrony and entrainment properties of robust circadian oscillators. J. R. Soc. Interface **5**, S17–S28 (2008)

6. Bass, J., Takahashi, J.S.: Circadian integration of metabolism and energetics. *Science* **330**(6009), 1349–1354 (2010)
7. Dunlap, J.C., Loros, J.J., DeCoursey, P.J.: *Chronobiology*. Sinauer Associates, Inc. (2004)
8. Glossop, N.R., Lyons, L.C., Hardin, P.E.: Interlocked feedback loops within the drosophila circadian oscillator. *Science* **286**(5440), 766–768 (1999)
9. Goodwin, B.C.: Oscillatory behavior in enzymatic control processes. *Adv. Enzyme Regul.* **3**, 425–437 (1965)
10. Hirota, T., Lee, J.W., St. John, P.C., Sawa, M., Iwaisako, K., Noguchi, T., Pongsawakul, P.Y., Sonntag, T., Welsh, D.K., Brenner, D.A., Doyle III, F.J., Schultz, P.G., Kay, S.A.: Identification of small molecule activators of cryptochrome. *Science* **337**(6098), 1094–1097 (2012)
11. Huang, K.C., Meir, Y., Wingreen, N.S.: Dynamic structures in escherichia coli: Spontaneous formation of mine rings and mind polar zones. *Proc. Natl. Acad. Sci. USA* **100**(22), 12724–12728 (2003)
12. Ishiura, M., Kutsuna, S., Aoki, S., Iwasaki, H., Andersson, C.R., Tanabe, A., Golden, S.S., Johnson, C.H., Kondo, T.: Expression of a gene cluster kaiABC as a circadian feedback process in cyanobacteria. *Science* **281**(5382), 1519–1523 (1998)
13. Marquié, 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), 258–264 (2015)
14. Mirsky, H.P., Liu, A.C., Welsh, D.K., Kay, S.A., Doyle III, F.J.: A model of the cell-autonomous mammalian circadian clock. *Proc. Natl. Acad. Sci. USA* **106**(27), 11107–11112 (2009)
15. Mukherji, A., Kobiita, A., Damara, M., Misra, N., Meziane, H., Champy, M.F., Chambon, P.: Shifting eating to the circadian rest phase misaligns the peripheral clocks with the master scn clock and leads to a metabolic syndrome. *Proc. Natl. Acad. Sci. USA* **112**(48), E6691–E6698 (2015)
16. Panda, S., Antoch, M.P., Miller, B.H., Su, A.I., Schook, A.B., Straume, M., Schultz, P.G., Kay, S.A., Takahashi, J.S., Hogenesch, J.B.: Coordinated transcription of key pathways in the mouse by the circadian clock. *Cell* **109**(3), 307–320 (2002)
17. Ramkisoensing, A., Meijer, J.H.: Synchronization of biological clock neurons by light and peripheral feedback systems promotes circadian rhythms and health. *Front. Neurol.* **6**(MAY) (2015)
18. Serkh, K., Forger, D.B.: Optimal schedules of light exposure for rapidly correcting circadian misalignment. *PLoS Comput. Biol.* **10**(4), e1003523 (2014)
19. Shaik, O., Sager, S., Slaby, O., Lebedez, D.: Phase tracking and restoration of circadian rhythms by model-based optimal control. *IET Syst. Biol.* **2**(1), 16–23 (2008)
20. Slaby, O., Sager, S., Shaik, O.S., Kummer, U., Lebedez, D.: Optimal control of self-organized dynamics in cellular signal transduction. *Math. Comput. Model. Dyn. Syst.* **13**(5), 487–502 (2007)
21. St. John, P.C., Hirota, T., Kay, S.A., Doyle III, F.J.: Spatiotemporal separation of per and cry posttranslational regulation in the mammalian circadian clock. *Proc. Natl. Acad. Sci. USA* **111**(5), 2040–2045 (2014)
22. St. John, P.C., Taylor, S.R., Abel, J.H., Doyle III, F.J.: Amplitude metrics for cellular circadian bioluminescence reporters. *Biophys. J.* **107**(11), 2712–2722 (2014)
23. Takahashi, J.S.: Transcriptional architecture of the mammalian circadian clock. *Nat. Rev. Genet.* (2016)
24. Taylor, S.R., Doyle III, F.J., Petzold, L.R.: Oscillator model reduction preserving the phase response: application to the circadian clock. *Biophys. J.* **95**(4), 1658–1673 (2008)
25. Ukai, H., Kobayashi, T.J., Nagano, M., Masumoto, K.H., Sujino, M., Kondo, T., Yagita, K., Shigeyoshi, Y., Ueda, H.R.: Melanopsin-dependent photo-perturbation reveals desynchronization underlying the singularity of mammalian circadian clocks singularity behaviour in circadian clocks. *Nat. Cell Biol.* **9**(11), 1327–1334 (2007)
26. Welsh, D.K., Takahashi, J.S., Kay, S.A.: Suprachiasmatic nucleus: cell autonomy and network properties. *Annu. Rev. Physiol.* **72**, 551–577 (2010)

27. Zhang, J., Qiao, W., Wen, J.T., Julius, A.: Light-based circadian rhythm control: entrainment and optimization. *Automatica* **68**, 44–55 (2016)
28. Zhang, R., Lahens, N.F., Ballance, H.I., Hughes, M.E., Hogenesch, J.B.: A circadian gene expression atlas in mammals: Implications for biology and medicine. *Proc. Natl. Acad. Sci. U. S. A.* **111**(45), 16219–16224 (2014)