Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Average Measures of Phase and Synchrony in Inhomogeneous Populations of Circadian Oscillators ⋆

Lindsey S. Brown * **Harman Mehta**[1] * **Francis J. Doyle III** *,**

* *Harvard John A. Paulson School of Engineering and Applied Sciences, Cambridge, MA 02138 USA (e-mail: lindsey_brown@g.harvard.edu, harmanmehta98@gmail.com, frank_doyle@seas.harvard.edu)*
** *Division of Sleep Medicine, Harvard Medical School, Boston, MA 02115 USA*

Abstract: Circadian misalignment between the phase of molecular circadian oscillators and the external environment has been linked to adverse health effects, including cardiovascular disease, obesity, cancer, and psychiatric disorders. Desynchrony of the circadian clock within populations of oscillators has also been linked to these conditions. For this reason, we wish to develop a control strategy to shift molecular circadian phase while also controlling for population synchrony. Previous work has demonstrated that model predictive control can effectively solve this problem when synchrony is determined directly by measuring the phase of each cell in a homogeneous population. Such cell-specific phase measurements are rarely possible *in vivo*, and such homogeneity is not biologically plausible. For these reasons, we wish to design an observer that can accurately determine the mean phase of the population and derive a proxy for population synchrony. We show that a model parameterized with the average parameter set of the population can be used to sense phase from mean population expression levels, despite inaccuracies when sensing the phase of a single cell. Similarly, we are able to use the average amplitude of the population in comparison to the amplitude of the average population oscillator as a measure of synchrony within the population. Taken together, these two metrics, based on the average behavior of the cell, allow us to control the phase and synchrony of the population of cellular oscillators without measuring the phase of individual cells.

Keywords: Synchronization, oscillators, biomedical control, predictive control

1. INTRODUCTION

The body's internal circadian phase may become misaligned with the external environment as a result of jetlag, shiftwork, or disease. Such misalignment has been implicated as a risk factor in a range of other conditions including cardiovascular disease, diabetes, and depression (Baron and Reid (2014)). For this reason, we wish to develop methods to shift internal circadian phase to reduce circadian misalignment with the environment.

A number of simulations have shown that model predictive control can determine the optimal timing and dosing of either light (Shaik et al. (2008); Zhang et al. (2016)) or small molecules (Abel and Doyle III (2016); Abel et al. (2019)) to shift the phase of the circadian oscillator. Such models assume that a model of a single oscillator can be used to represent organismal or population level rhythms despite the fact that each cell in the population contains its own molecular circadian oscillator driven by two interlocked transcription translation feedback loops (Ko and Takahashi (2006)). Because this approach considers only a single oscillator, it neglects the question of circadian desynchrony, in which individual oscillators become out of phase with each other internally. As in the case of circadian misalignment with the external environment, this internal desynchrony within the population has also been linked to negative health outcomes, including obesity and metabolic disorders (Kolbe et al. (2018); Zhang and Sehgal (2019)). In fact, previous work has shown that applying model predictive control to shift the mean phase of an oscillator population without controlling for population synchrony may result in desynchrony. Instead, by incorporating feedback of both mean circadian phase and measures of population synchrony, the resulting model predictive control strategy is able to shift the mean phase of the population to a reference phase while maintaining or improving synchrony within the population (Abel et al. (2018)).

Such an approach is limited in its biological plausibility. First, the feedback assumes that the phase of each cell can be sensed individually to determine the mean phase and synchrony index of the population. Second, the

⋆ This work was supported by NIH T32-HL007901.
1. Current address: Department of Biochemical Engineering and Biotechnology, Indian Institute of Technology Delhi

work considered homogeneous populations of cells which differed only in initial phase but were governed by the same parameters. In this study, we consider the control of inhomogeneous populations where our measurements are limited to population level averages. We show that we are able to estimate mean circadian phase (Section 3) and synchrony (Section 4) from these population level measurements, thus allowing for model predictive control of both circadian phase and synchrony for inhomogeneous cellular populations without requiring cellular level measurements (Section 5).

2. MODEL PREDICTIVE CONTROL FOR PHASE AND SYNCHRONY

Consider a population of N cells. The circadian oscillator in cell i can be modeled as a system of ordinary differential equations,

$$\frac{dx_i}{dt} = f(x_i(t), p_i, u(t)) \tag{1}$$

with $x_i(t) \in \mathbf{R}^n$ representing the model states of the ith cell, $p_i \in \mathbf{R}^m$ the model parameters governing the ith cell, and $u(t) \in \mathbf{R}^k$ an input control stimulus, where the model dynamics are such that for $u(t) = 0$, the oscillator has an attractive limit cycle, Γ_i, with period T_i,

$$\lim_{t\to\infty} \|x_i(t) - x_i(t + T_i)\| = 0. \tag{2}$$

We define the phase, $\phi_i(t) \in [0, 2\pi)$, of a point on the limit cycle, $x_i^0(t) \in \Gamma_i$ by

$$\phi_i(t) = \omega_i t + \phi_i(x_i^0(0)) \tag{3}$$

for $\omega_i = \frac{2\pi}{T_i}$, so that points on the limit cycle are equidistant in both time and phase space. For simplicity, we will assume that every cell has the same period, T, and correspondingly the same natural frequency, ω. Because we assume that the limit cycle is attractive, we can define an asymptotic phase mapping for any point in state space by taking $\phi_i^a(x_i(t)) = \phi_i(x_i^0(t))$ where

$$\lim_{t\to\infty} \left\| f(x_i(t), p_i, 0) - f(x_i^0(t), p_i, 0) \right\| = 0. \tag{4}$$

Such a model is useful because it allows us to approximate the effect of a control stimulus on phase to first order, with the assumption that $u(t)$ can be reflected in a change in the jth model parameter, by

$$\frac{d\widehat{\phi_i}(t)}{dt} = \omega + \frac{\partial}{\partial t}\frac{\partial \phi_i(x)}{\partial u}u(t) = \omega + \frac{\partial}{\partial t}\frac{\partial \phi_i(x)}{\partial p_{i,j}}u(t), \tag{5}$$

thus providing us a phase-reduced model (Taylor et al. (2008)) for use in our model predictive control algorithm.

2.1 Generating an Inhomogeneous Oscillator Population

In our simulations, we use a 14 state ordinary differential equations model, where each state represents the expression level of a core clock gene or protein concentration, to describe the dynamics of the molecular circadian clock (Brown and Doyle III (2018)). We model KL001 as our control input, which acts by decreasing the degradation rate of nuclear cryptochrome (Hirota et al. (2012)). To generate an inhomogeneous population of 100 cells, we define a parameter set p_i for the ith cell by drawing each of the 46 model parameters uniformly in a range about the baseline that produced oscillations when the parameter was manipulated individually. For each parameter set, we require that the model is a limit cycle and has the correct period sensitivities to the knockout of the cryptochrome mRNAs, *Cry1* and *Cry2*, as the corresponding proteins are targeted by our control input. For simplicity, we rescale the parameters such that all models have the same period. These different parameter sets correspond to intercellular variability, where different cells will have slightly different rates of transcription, translation, and degradation of the core clock genes and proteins as would be expected *in vivo*.

For a cellular population consisting of N cells, we can define mean phase and synchrony index using the kuramoto order parameter. Let

$$\bar{z}(t) = \frac{1}{N}\sum_{i=1}^{N} e^{i\phi_i(x_i(t))}. \tag{6}$$

The mean phase of the population is $\bar{\phi} = \angle(\bar{z})$, and the population synchrony index is $\rho = |\bar{z}|$, where $\rho \in [0, 1]$ with $\rho = 1$ corresponding to perfect synchrony. The synchrony index also defines the circular standard deviation,

$$\sigma(\rho) = \sqrt{-2\ln(\rho)}. \tag{7}$$

From this cellular population, we define the *average model*, parameterized by the average parameter set,

$$\frac{d\bar{x}}{dt} = f(\bar{x}(t), \bar{p}, u(t)) \tag{8}$$

where $\bar{p_j} = \frac{1}{N}\sum_{i=1}^{N} p_{i,j}$. This model is also a limit cycle, so we can similarly define an asymptotic phase mapping $\phi_{\bar{p}}$ and corresponding phase-reduced model,

$$\frac{d\widehat{\phi_{\bar{p}}}(t)}{dt} = \omega + \frac{\partial}{\partial t}\frac{\partial \phi_{\bar{p}}(x)}{\partial \bar{p_j}}u(t). \tag{9}$$

2.2 A Model Predictive Control Algorithm

We assume that different control inputs cannot be targeted to individual cells, so each cell receives the same control input, $u(t)$. This means that due to the inhomogeneity of the oscillator population, the phase shift in response to the same control input will be different for different cells. Despite this challenge, we wish to find a series of control inputs which will minimize t^* such that for all $t > t^*$, the mean phase of the population has been aligned to a reference, $\bar{\phi}(t) = \phi_r(t)$, while controlling for synchrony.

Our ideal nonlinear model predictive control (MPC) algorithm would determine this series of inputs through the following steps:

(1) Find a series of N_p control inputs u for a predictive horizon of N_p timesteps which minimizes

$$L(u) = \sum_{j=1}^{N_p} w_j \left| \widehat{\bar{\phi}}(t_j) - \phi_r(t_j) \right|^2 + q_j |u_j|^2 + r_j \left| 1 - \widehat{\rho}(t_j) \right|^2$$

(2) Apply u_1 to each cell and calculate $\phi_i(t)$ for each cell using the full model dynamics.
(3) Repeat steps 1-2 for the next timestep.

However, this cost function, $L(u)$, requires knowing the phase of each individual cell in order to compute the mean phase, $\bar{\phi}$ and synchrony ρ of the population, which is not biologically plausible. We also require a model of the evolution of these estimates for each timestep. Thus, we wish to determine estimates for $\hat{\bar{\phi}}(t_j)$(section 3) and $\hat{\rho}(t_j)$ (section 4) from population level observations, which will allow us to define a biologically plausible cost function for our MPC algorithm (section 5).

3. PHASE SENSING IN INHOMOGENEOUS POPULATIONS

Because we assume that we do not know the parameters of individual cells exactly, it is impossible to sense phase using the individual $\phi_i(x_i)$ models. Instead, we could use the average model to sense the phase of individual cells. For each timepoint, we can calculate $\phi_{\bar{p}}(x_i)$ to get an estimate for the phase of the ith cell. As shown in Figure 1a, sensing at the individual level results in estimates which are on average accurate but with some small variance for the majority of the period. However, for the true phase in $(\frac{-\pi}{2}, \frac{\pi}{2})$, there are inaccuracies in the mean with increased variance of the estimate. This increased variance corresponds to a region of state space where the expression levels of many of the genes and proteins attain a maximum, so there is relatively slower rates of change of expression across the phases in this region, potentially explaining this greater variance (Brown and Doyle III (2018)). Corresponding to the greater variance in estimates, we find that if we compute the synchrony index for the population at these phases, our estimate drops below 0.2 despite the cells being in perfect phase (Figure 1b). Therefore, even if it were biologically practical to sense gene and protein expression at the level of individual cells, these resulting inaccuracies in estimates of mean phase and synchrony motivate using a different measurement.

We propose using the average population expression levels and the average model to determine the mean phase of the population. In particular, we predict the mean phase by

$$\bar{\phi}(t) = \phi_{\bar{p}}(\bar{x}(t)) \tag{10}$$

where $\bar{x}(t) = \frac{1}{N}\sum_{i=1}^{N} x_i(t)$, the mean expression levels, could be inferred from population recordings (for example, PERIOD2::LUCIFERASE recordings). In our simulations, we consider a population of 100 cells and show that this method results in accurate predictions of the phase of the population (Figure 1a). Since we no longer have estimates of the phases of the individual cells we can no longer compute the synchrony index of the population. Without incorporating the synchrony term, the MPC algorithm may successfully shift the mean phase of the population but result in desynchrony among the individual cells (Figure 4a and Abel et al. (2018)).

4. MEASURING SYNCHRONY IMPLICITLY

From just a single estimate of phase, we cannot estimate synchrony in the traditional manner using the synchrony index. Based on the observation that desynchrony results in lower amplitude average rhythms, we propose comparing the mean expression levels of the population to the

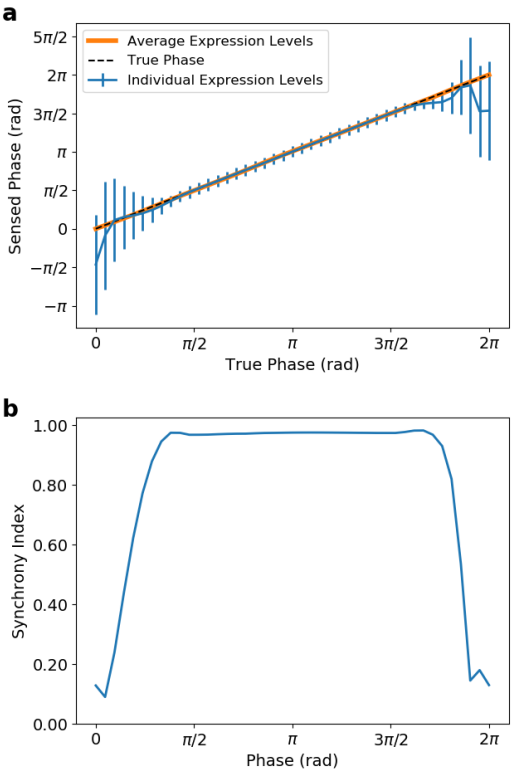

Fig. 1. a. Phase sensed using the average model from individual cells (blue, bars show circular standard deviation) and average population expression levels (orange) compared to the true model phase. b. Synchrony index for phases of individual cells sensed by the average model.

expression levels of the average model to implicitly infer synchrony of the population. We note that even in the case of perfect synchrony, the mean expression levels of the population are not always equal to the expression levels of the average model (Figure 2), demonstrating that the average parameter set does not reproduce the average dynamics of the model exactly.

The distribution of mean expression levels varies depending on the species. For example, in Figure 2, the expression of *Cry1* is tightly distributed about the mean population while the expression of BMAL1 varies more widely. For this reason, our metric, detailed below, compares estimates of synchrony across all the model species.

Figure 2 also shows different behavior in different model species when control is applied; some species show a phase advance in expression (*Cry1* and BMAL), while others show delays (*Rev-erbα*). Similarly, we see an increase in amplitude of the expression of some species (*Cry1* and BMAL) and a decrease of others (PER and *Rev-erbα*). The change in PER is particularly striking in that continuously applied maximal control drives the amplitude below the typical levels. Given the changes that occur while control is applied, it is natural to ask whether our prediction of mean phase of the population and our estimate of synchrony will be valid during the transients of a phase shift, but our simulations show that these measurements still provide the necessary feedback to achieve a phase shift.

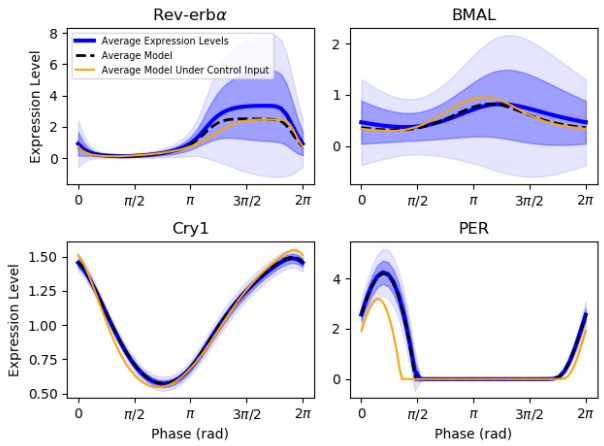

Fig. 2. The mean expression levels (blue, shading indicates ± 1 and 2 standard deviations) of the population of cells compared with the expression levels of the average model at baseline (black dashed) and when maximal control has been applied continuously for the entire period (orange).

Let $y(t) \in \mathbf{R}$ represent the expression level of a single gene or protein (i.e., $y_i(t)$ is a single component of $x_i(t)$). For a single species, we estimate synchrony by

$$\rho_y(\bar{\phi}(t)) = \frac{\bar{y}(t)}{c(\bar{\phi})y_{\bar{p}}(\bar{\phi}(t))} \tag{11}$$

where $c(\phi)$ is chosen such that $\rho_y(\phi) = 1$ when the oscillator population is in perfect synchrony to account for differences in the average dynamics and the average model. In regions where the expression level curves are strictly convex and concave down, $\bar{y}(t)$ will be strictly less than $y_{\bar{p}}(\bar{\phi}(t))$. In regions that are not strictly convex, synchrony may decrease but preserve the same mean expression level for a single model species. The regions where the expression level curves are strictly convex differ for each model species as different model species peak at different phases. Thus assuming that the range of phases in the cellular population is not too large, at least one of the model species will show strictly convex behavior allowing us to define synchrony of the population by

$$\bar{\rho}(\bar{\phi}(t)) = \min_y \rho_y(\bar{\phi}(t)). \tag{12}$$

Figure 3 shows the implicit synchrony compared to the true synchrony index of the population over a period. For a population that is perfectly synchronized, the implicit synchrony is designed to be 1 in agreement with the synchrony index. For populations where phase is distributed about the mean, the our synchrony estimates are much lower than the synchrony index. Still, we see that our implicitly defined synchrony measure has the property that for a population with a higher synchrony index, the implicit synchrony is also generally higher. This property is all that is needed for our measure to be useful in control because an increase in implicit synchrony corresponds to an increase in population synchrony, our desired performance.

Unlike the true synchrony index which is constant throughout the period, the implicit synchrony index fluctuates slightly depending on the mean phase of the population.

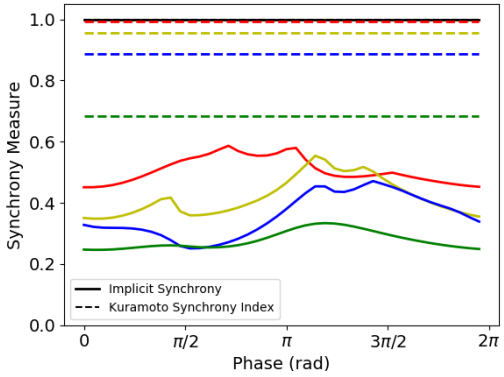

Fig. 3. The synchrony index computed using the Kuramoto order parameter (dashed lines) compared with the implicit synchrony computed from the mean expression levels (solid lines) with each color corresponding to a population with different phase distributions.

This means that the same population may appear to be more or less synchronized depending on the phase at which it is sampled. While this is not a desired feature of the measurement, it does not negatively impact our control algorithm for two reasons. First, MPC considers a limited future time window, and the changes in measured synchrony are gradual. Second, and more importantly, our measure of implicit synchrony is not explicitly used to predict future synchrony but only as a way of initializing our predictions as described in the next paragraph.

To be used in our control algorithm, we need to be able to predict the evolution of synchrony following our control input at each step in the algorithm from our initial synchrony measure. We can describe the changes in phase as a result of control input through a phase transition curve (St. John et al. (2015)),

$$\Delta(\phi, t_k) = \omega + u(t_k) \int_{t_k}^{t_{k+1}} \frac{\partial}{\partial t} \frac{\partial \phi_{\bar{p}}}{\partial p} dt. \tag{13}$$

Let $f(\phi, t)$ be the probability distribution of phases, ϕ, in the population at time t. Using the phase transition curve, we find the distribution of phases at a future timestep by

$$f(\phi, t_{k+1}) d\Delta(\phi, t_k) = f(\phi, t_k) d\phi. \tag{14}$$

For any phase distribution, we can find the population synchrony, $\hat{\rho}$, by

$$\hat{\rho}(t) e^{i\bar{\phi}(t)} = \int_0^{2\pi} e^{i\theta} f(\theta, t) d\theta. \tag{15}$$

Thus, we can use the phase transition curve to predict the evolution of population synchrony from an initial distribution. In our MPC algorithm, we only use the implicit measured synchrony to define an initial distribution and the phase transition curve to predict synchrony at future timesteps, so that the changes in measured synchrony do not affect our choice of control which is based only on the cost function of our predicted synchrony. In particular, we define an initial starting phase distribution by

$$f(\phi, t) = \mathcal{N}(\phi_{\bar{p}}(\bar{x}), \sigma(\bar{\rho}(t))). \tag{16}$$

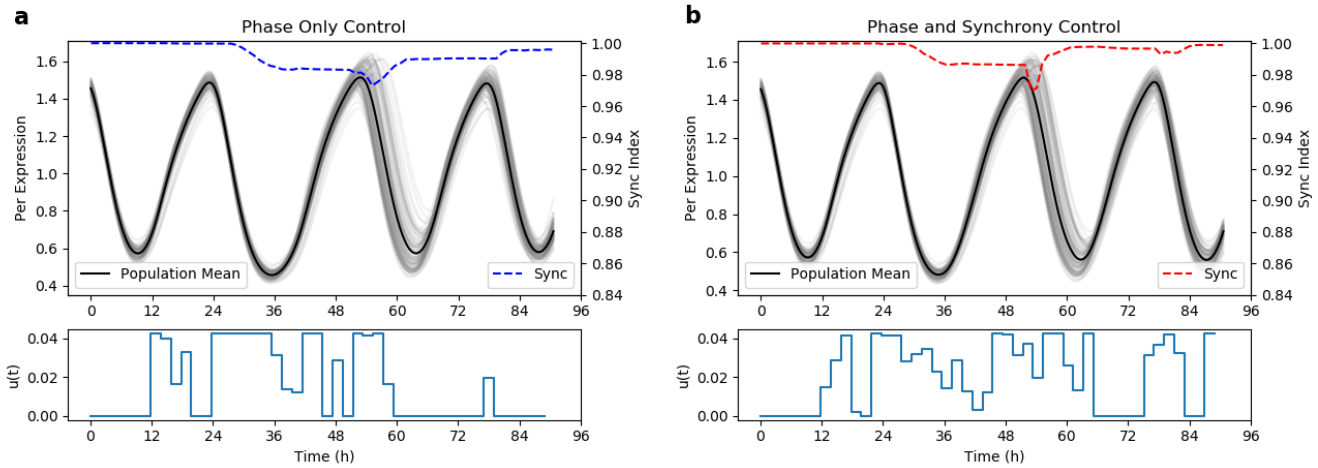

Fig. 4. a. Using the MPC algorithm with average population sensing without accounting for population synchrony drives the population out of phase. b. The MPC algorithm using average population sensing for mean phase and synchrony shifts the phase of the population while maintaining synchrony.

5. COMBINING AVERAGE PHASE SENSING AND IMPLICIT SYNCHRONY MEASUREMENTS

In sections 3 and 4, we proposed measures which used only the mean expression levels of the population to estimate the mean phase and synchrony of the population. This allows us to modify our MPC algorithm to a population average level, an algorithm which is potentially practical biologically, dependent on experimental ability to sense mean expression levels of our model states. The steps of the resulting MPC algorithm are the following:

(1) Sense $\bar{x}$, and use this to compute $\bar{\phi}$ (equation 10) and $\bar{\rho}$ (equation 12) for the population.
(2) Find a series of N_p control steps which minimize

$$L(u) = \sum_{j=1}^{N_p} w_j \left|\widehat{\phi}_{\bar{P}}(t_j) - \phi_r(t_j)\right|^2 + q_j |u_j|^2 + r_j \left|1 - \widehat{\bar{\rho}}(t_j)\right|^2$$

(3) Apply u_1 to each cell and calculate $x_i(t)$ for each cell using the full model dynamics.
(4) Repeat steps 1-3 for the next timestep.

We simulate this MPC algorithm with $N_p = 3$, shown to be optimal for a single cell in Abel et al. (2019), for a population of cells that begins perfectly in phase with a phase shift of -6 h occurring at 12 h into our simulation. In Figure 4a, we see that if we do not control for population synchrony, the population becomes desynchronized by the phase shift (most clearly seen between hours 48 and 72 of the simulation), as in the case of homogeneous populations (Abel et al. (2018)). In order to shift the phase of the oscillator, some loss of synchrony is unavoidable because the inhomogeneity of the population makes it so that the phase response to an input stimulus will not be identical for each cell. The results of our simulation show that when we control for both mean phase and synchrony, we can still shift the mean phase of the population while minimizing the loss of synchrony (Figure 4b).

We see a tradeoff between controlling for synchrony and the time to achieve the desired phase shift (Figure 5a). When we also control for synchrony, it takes slightly longer for the mean phase of the oscillator population to achieve the phase shift because control that would shift the mean phase of the oscillator will not be applied if it decreases the synchrony too much. Control also continues to be delivered much later in the simulation because more sustained manipulation is required to correct the synchrony. Our measured synchrony is typically much less than the true synchrony index, so that the evolution of phases is from a broader distribution, covering a broader range of phase responses, which will result in more conservative control being applied.

Figure 5a also shows that the phase sensed from the mean expression levels of the population closely tracks the true mean phase of the population, allowing us to make good predictions for input which will shift the population. This demonstrates that our phase sensing algorithm is still valid while control is being applied despite resulting changes in average expression level. Unsurprisingly, when the population is least synchronized, we see the highest discrepancies between our sensed phase and the true phase because these mean values will differ the most from the unperturbed oscillator.

In addition to shifting the phase of the oscillator, these simulations show that we can also choose control inputs to maintain high levels of synchrony. In particular, when we explicitly incorporate synchrony into the cost function of the model, the synchrony of the population throughout the phase shift and after the phase shift have occurred are higher (Figure 5b). Future work might explore what levels of synchrony are biologically meaningful to maintain.

6. DISCUSSION

These simulations demonstrate that there exist a choice of MPC parameters which can achieve population level control, leaving the optimization of these parameter choices as a topic for future work. Such work should consider the balance of our choices of w, q, and r which reflect the relative importance of phase errors, control applied, and population synchrony, all of which may have consequences

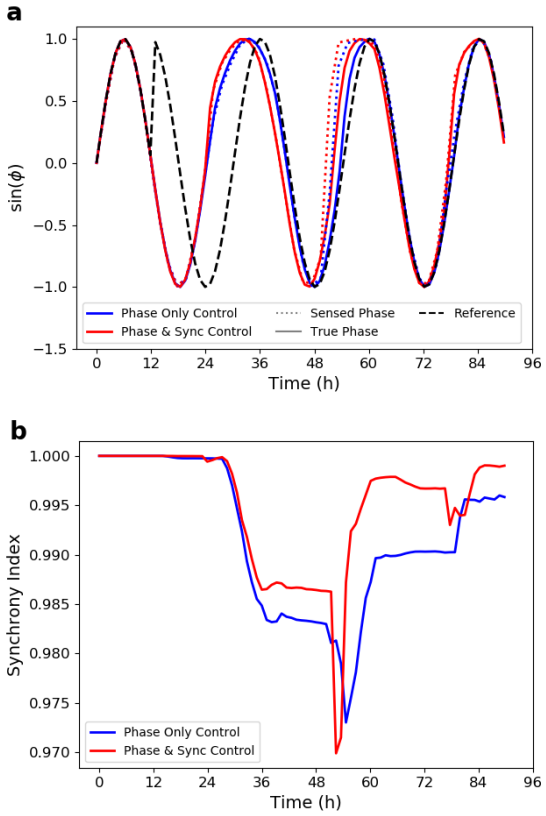

Fig. 5. a. A comparison between the true mean phase of the system (solid lines) and the sensed phase of the system (dotted lines) as the system is aligned to the reference trajectory (black dashed line) and b. a comparison of the true population synchrony index when we control for phase only (blue) and when we control for both phase and synchrony (red).

for health in terms of circadian misalignment, stimulus toxicity, and circadian desynchrony.

In our proposed measures of mean phase and synchrony, we use all the states of the model, which may not be efficient or possible biologically. Future work should explore whether such measurements can be estimated from a reduced state space, thus requiring the measurement of fewer genes and proteins. Relatedly, future work may consider subsampling from a population of cells to infer the mean phase and synchrony of the entire population. Initial results suggest that using only subsets of the population for computing mean population phase may still be effective, but a more formal analysis on expected accuracy as a function of the proportion of the population subsampled is warranted. Furthermore, we assumed that all of the cells in the population have the same period, and future work should consider how differences in underlying oscillator period may impact these results. We also assumed perfect sensing of the mean population expression levels, but *in vivo*, we would expect noise in these measurements, so it should be considered how errors in these measurements impact the sensed phase and hence lead to errors in our MPC algorithm (Brown et al. (2019)).

While our simulations explicitly model populations of cells, this work demonstrates that using the average model to predict both the phase of the population and the response to a control input may adequately capture the dynamics of cellular populations when the control algorithm also incorporates an implicit measure of synchrony.

REFERENCES

Abel, J.H., Chakrabarty, A., and Doyle III, F.J. (2018). Controlling Biological Time: Nonlinear Model Predictive Control for Populations of Circadian Oscillators. In *Emerging Applications of Control and Systems Theory*, 123–138.

Abel, J.H., Chakrabarty, A., Klerman, E.B., and Doyle III, F.J. (2019). Pharmaceutical-based entrainment of circadian phase via nonlinear model predictive control. *Automatica*, 100, 336–348.

Abel, J.H. and Doyle III, F.J. (2016). A systems theoretic approach to analysis and control of mammalian circadian dynamics. *Chemical Engineering and Research Design*, 116, 48–60.

Baron, K.G. and Reid, K.J. (2014). Circadian Misalignment and Health. *Int Rev Psychiatry*, 26(2), 139–154.

Brown, L.S. and Doyle III, F.J. (2018). Toward Multi-Input Control : A Dual-Feedback Loop Model of the Mammalian Circadian Clock. *IFAC-PapersOnLine*, 51(19), 24–27.

Brown, L.S., Klerman, E.B., and Doyle III, F.J. (2019). Compensating for Sensor Error in the Model Predictive Control of Circadian Clock Phase. *IEEE Control Systems Letters*, 3(4), 853–858.

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., and Kay, S.A. (2012). Identification of Small Molecule Activators of Cryptochrome. *Science*, 337(6098), 1094–1097.

Ko, C.H. and Takahashi, J.S. (2006). Molecular components of the mammalian circadian clock. *Human Molecular Genetics*, 15(SUPPL. 2), 271–277. doi: 10.1093/hmg/ddl207.

Kolbe, I., Brehm, N., and Oster, H. (2018). Interplay of central and peripheral circadian clocks in energy metabolism regulation. *Journal of Neuroendocrinology*, (September), 1–8.

Shaik, O.S., Sager, S., Slaby, O., and Lebiedz, D. (2008). Phase tracking and restoration of circadian rhythms by model-based optimal control. *IET Systems Biology*, 2(1), 16–23.

St. John, P.C., Taylor, S.R., Abel, J.H., and Doyle, F.J. (2015). Amplitude metrics for cellular circadian bioluminescence reporters. *Biophysical Journal*, 107(11), 2712–2722.

Taylor, S.R., Doyle III, F.J., and Petzold, L.R. (2008). Oscillator model reduction preserving the phase response: application to the circadian clock. *Biophysical journal*, 95(4), 1658–1673. doi:10.1529/biophysj.107.128678.

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

Zhang, S.L. and Sehgal, A. (2019). Circadian Rhythms and Disease. In R.E. Pyeritz, B.R. Korf, and W.W. Grody (eds.), *Emery and Rimoin's Principles and Practice of Medical Genetics and Genomics: Clinical Principles and Applications*, 299–314. Academic Press, 7 edition.