

State-Estimation and Cooperative Control with Uncertain Time

Sean G. Carver¹, Eric S. Fortune², and Noah J. Cowan¹

Abstract—State estimation and control typically depend on the implicit assumption that actuation and measurement occur at known points in time. This assumption is predicated on sufficiently precise timekeeping afforded by engineered clocks. Biological control systems do not have access to quartz-crystal technology and yet animals perform behaviors that engineered systems cannot. Here, we examine the problem of state estimation using imprecisely timed measurements, with known temporal statistics. We consider the case that there are two controllers, each with its own imperfect clock, performing a cooperative task—formulated in a leader–follower paradigm. In our problem, the follower estimates the state of the leader’s noisy clock along with other state variables. An example of a complex cooperative leader–follower task is ballroom dancing.

I. INTRODUCTION

No known clock tells time perfectly [1].¹ We, therefore, expect no objection to the statement that biological timekeeping falls short of perfection. However the consequences this unobjectionable statement have hardly been explored for motor control and sensory processing.

How does imperfect chronometry affect estimation and control? First, we point out that there are two problems: the initial synchronization and the ongoing problem of temporal drift. We focus our efforts here on a particular problem in which two agents perform a simplified cooperative control task that is loosely analogous to ballroom dancing and more closely analogous to “cooperative sawing.” Cooperative sawing might be performed by two lumberjacks, pushing and pulling on opposite ends of a saw. This task is illustrated schematically in Figure 1. The goal is to drive $y(t)$ to follow the reference signal $r(\cdot)$, but since neither agent knows real time, t , they must synchronize their chronometers in order to achieve the cooperative behavior. That is, the agents must agree on what time it is relative to the task, $r(\cdot)$. Here we assume the clocks start synchronized, and consider the ongoing problem of staying synchronized, despite the fact that one or both partners have imperfect timekeeping. Clearly synchrony means, in this case, synchronization of clocks between agents without regard to some hypothetical external clock. Even for a linear plant, if $r(\cdot)$ is nonlinear (e.g. periodic), synchrony requires nonlinear estimation, and in this paper we resort to an Extended Kalman Filter (EKF).

^{*}This material is based upon work supported by the National Science Foundation (NSF) under grant #0845749

¹S. G. Carver and N. J. Cowan are with the Department of Mechanical Engineering, Johns Hopkins University, USA scarver2@jhu.edu

²E. S. Fortune is with the Department of Biological Sciences, New Jersey Institute of Technology, USA

¹http://www.nist.gov/pml/div688/logicclock_020410.cfm

A. Biological Time Keeping

Neuroscientists have long recognized that the brain must keep or process time—from circannual rhythms to precise moment-to-moment motor control—to perform behaviors [3]. Temporal processing has been implicated [4] in sensory perception such as sound localization, motion processing, music and speech perception, locomotion, breathing, vocalizations, circadian rhythms, and many other functions. Further, many species of animals coordinate their behavior (among two or more individuals) over wide temporal scales, from milliseconds to years.

Biological clocks have been found in all tissues, and implemented using a wide variety of strategies [5]. Time keeping for the control of behavior, however, is generally the province of excitable tissues, particularly the nervous system. Within the nervous system of an organism, we commonly find multiple clocks using multiple independent mechanisms for timekeeping [6]. For example, central pattern generators (CPGs) are collections of neurons that can oscillate at a preferred frequency and are used in motor control [7]. Another strategy for temporal processing is the use of labeled lines [8]. In this case, time-varying sensory information is passed through filters that have differences in their temporal dynamics – these differences are used to determine the temporal structure of salient sensory stimuli.

Labeled lines have been implicated in barn owl sound localization [8] where they are known to operate on the scale of microseconds. More exotic mechanisms have been

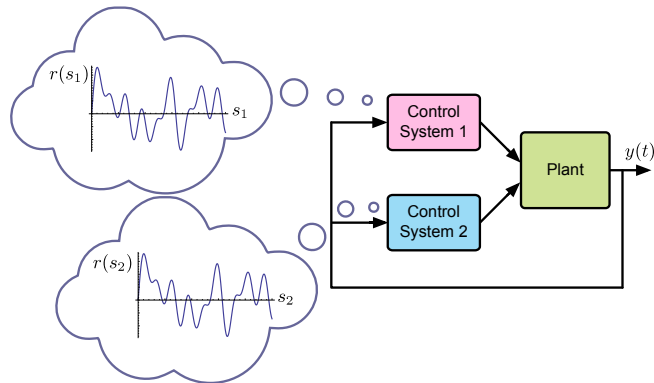


Fig. 1. A simplified block-diagram of two controllers cooperatively controlling a single plant. Inspired by cooperating (duetting) songbirds [2], we assume both controllers have a representation of the combined cooperative performance; this is illustrated here as a copy of desired trajectory of the plant output in the memory of each controller. However, each controller is assumed to have its own estimate of time, s_1 and s_2 , that are not intrinsically synchronized, and thus the synchronization of control inputs requires synchronization of the clocks.

hypothesized for the range of microseconds and seconds such as the internal clock model [9], the population clock model [10] or the multiple oscillators model [11]. In these timescales, neural mechanisms remain least understood [12]. Researchers have a better handle on some clocks with larger timescales, such as the one for circadian rhythms [13].

To be sure, the model described here does not replicate neural mechanisms of timing, but rather our interest is in understanding the consequences of *imperfect time keeping* on the control of behavior, and on the estimation of quantities relevant to the control of behavior.

B. Our imperfect clock

There are undoubtedly many choices for a model of the control agent’s clock, particularly in light of different hypothesized neural mechanisms. To begin to explore the consequences of imperfect timekeeping, we choose to start with only one clock—the simplest possible imperfect clock model that we can find to suit our purpose. Here we choose a “ticking clock” with random, but approximately regular ticks. Our choice is based entirely on simplicity and parsimony and not on an assessment of biological relevance.

We have a notion of “real time,” i.e. the time a perfect clock would read if one existed. The agent’s imperfect clock “ticks” at real times t_k , for $k = 1, 2, \dots$. These t_k are random variables. We will also assign an initial tick with certain value $t_0 = 0$.

As a further simplification of our model we will assume that the inter-tick times $\Delta_k = t_k - t_{k-1}$ for $k = 1, 2, \dots$, are independent and identically distributed (IID) random variables with common probability density function (PDF) $f(\Delta)$. Here we choose the PDF of an inverse Gaussian for f , although other choices could be made. Let δ denote the mean of the distribution given by f and let σ denote its standard deviation. Note that while δ and σ parameterize the inverse Gaussian distribution, they are not the most common parameterization. We assume that these statistics are accurately known to the agent.

For our purposes, we want the clock to have a continuous state s that evolves in time, but still ticks at events, with independent intertick intervals. This desire is our motivation for the choice of the inverse Gaussian distribution for f . Indeed, let the clock’s state follow a drift-diffusion equation:

$$\begin{aligned} ds(t) &= dt + b_1 dW(t), \\ s(0) &= 0, \text{ with probability } 1. \end{aligned} \quad (1)$$

Here $W(t)$ is a (standard) Wiener random process, and b_1 is a positive number called the clock’s diffusion coefficient. For simplicity, by choice of its initial condition, the diffusive clock starts synchronized with real time. We call this a diffusive clock with unit drift. The n th tick occurs when the diffusion first crosses level $n\delta$. It follows that the intertick times are distributed Inverse Gaussian and with mean δ and standard deviation $\sigma = b_1\sqrt{\delta}$.

C. Literature on control with imperfect timekeeping

We found very little in the literature concerning the consequences of imperfect timing for estimation and control.

Indeed our literature search pulled up only one conference proceeding by LaValle & Egerstedt [14]. The authors open their abstract with the following strong and provocative statements: “this paper addresses the peculiar treatment that time receives when studying control systems. For example, why is the ability to perfectly observe time assumed in virtually all control formulations?” They cite no papers which violate this assumption, and moreover we have found no additional papers that do so.

LaValle and Egerstedt’s most prescient results concern the implications for imperfect timing on open-loop stability. Three examples show a range of possible stability behaviors for an imperfectly timed controlled system. Their first example shows that it is possible for a system to be stabilized by a controller with perfect timing, but the same system with the same control law will suffer catastrophic instability if there is *any* imperfection in the timing. Their second example concerns a system with three stable states. Which stable point attracts the system’s trajectory depends on when (and if) the controller delivers a stereotyped impulse. All three stable states are possible limit points (and the only possible limit points). In the last example, the controller throws a switch which changes the unstable system into a stable system. LaValle Egerstedt call this “strongly open loop stable,” because stability—in this case, convergence to the origin—does not depend on when the controller throws the switch—only, that, in fact, the controller does throw the switch. LaValle and Egerstedt clearly see their paper as the first of many on this topic. They close with a section on “open questions and issues” and “argue that many fascinating directions for future controls research emerge.” We agree.

II. COOPERATION PROBLEM

We start with a plant of the form

$$\dot{y}(t) = u_1(t) + u_2(t), \quad (2)$$

where y is the state variable of the plant, and u_1 and u_2 are the inputs from two controllers. The shared goal of the controllers is to make the plant produce a desired output. The desired output is a function—called the task representation—denoted $r(t)$, e.g. $r(t) = \sin(t)$. In other words, together, the controllers aim to make $y(t) \approx r(t)$. To accomplish this goal, we assume controllers know the functions $r(\cdot)$, $\dot{r}(\cdot)$, and $\ddot{r}(\cdot)$, have clocks (perfect or imperfect), and can measure the present plant state $y(t)$ (with or without measurement noise). Initially, we assume the clocks and plant measurements are perfect, then consider the case that really interests us: imperfect observations of both time and plant state.

In what follows, we assume that Controller 1 and Controller 2 are posed together, with knowledge of each other’s control law. We do not solve the adaptive control problem in which each Controller is naive to the other’s control law.

A. Single Actor, Perfect Clock

We assume one actor; specifically, $u_2(t) = 0$, for all t . We pose Controller 1 (the actor) as follows:

$$u_1(t) = -Ky(t) + 2f(t), \quad (3)$$

where

$$f(t) = \frac{K}{2}r(t) + \frac{1}{2}\dot{r}(t), \quad (4)$$

K is a positive gain, and $r(t)$ is the actor's representation of the task. Controller 1's ability to implement this control law is predicated on the assumption that Controller 1, at every moment, knows the present t (has a perfect clock) and knows the present y , (perfectly observes the plant).

Why does this control law lead the plant to follow $r(t)$? The control law implies, (plugging $u_1(t)$ and $u_2(t)$ into (2)), that

$$\dot{y}(t) = -Ky(t) + Kr + \dot{r}, \quad (5)$$

and so

$$\frac{d}{dt}(y(t) - r(t)) = -K(y(t) - r(t)). \quad (6)$$

Using $K > 0$, we conclude that $y(t) \rightarrow r(t)$ exponentially.

B. Symmetric Actors, Perfect Clock

Here we assume that each controller can, at every time t , perfectly observe t and the $y(t)$. We set

$$\begin{aligned} u_1(t) &= -\frac{K}{2}y(t) + f(t), \\ u_2(t) &= -\frac{K}{2}y(t) + f(t), \end{aligned} \quad (7)$$

where f is given by (4). The same conclusion follows.

C. Single Actor, Imperfect Clock

We return to the case in which Controller 1 (the actor) sees alone ($u_2 = 0$), but we now relax the assumption that the actor knows the present t . For the purposes of simplicity, we still assume that the actor, at all times t , can perfectly observe $y(t)$. When we later move to the leader-follower paradigm, only the follower will rely on a Kalman filter, while both will have imperfect clocks. Assuming the leader measures the plant continuously and accurately will simplify the follower's model for the leader.

We assume the actor's clock is a diffusive clock, with state $s(t)$ that satisfies stochastic differential equation (1). The clock's state $s(t)$ is the actor's estimate of time. Because the drift (coefficient on dt in (1)) is unity, $s(t)$ is an unbiased estimate of t .

The actor's control law, based on s , is then

$$u_1(t) = -Ky(t) + 2f(s(t)), \quad (8)$$

where $f(\cdot)$ is given as in (4). Again, in (8) we assume the actor has access to $y(t)$ at time t , even though the actor does not know t .

D. Single Actor, Separate Observer, Imperfect Clocks

Here, the actor, Controller 1, remains the same (see subsection II-C, above), but we let other agent *observe* the process, rather than just lie dormant. As before, the actor sees solo, i.e. we still have $u_2(t) = 0$, for all t . This situation is analogous to the one considered by Wolpert and coauthors in [15].

The observer measures $y(t)$, at discrete times, and with measurement noise, and has its own diffusive clock, $q(t)$. The evolution of q is given by (1), with q replacing s , and an independent Wiener process, (i.e. independent from W), replacing W . Here $q(t)$ serves as the follower's crude estimate of $s(t)$. The observer, derives a better (lower variance error) estimate of $s(t)$, called $\hat{s}(t)$, by combining the plant measurements with q , and knowledge of the actor's task representation r . Our derivation is based on an EKF, as follows.

The evolution of the plant state $y(t)$ is given by (2), (8), and $u_2 = 0$, as follows:

$$dy(t) = (-Ky(t) + 2f(s(t))) dt. \quad (9)$$

For simplicity, we choose $r(0)$ as the initial condition for y —the task starts as represented by r . We combine (1) and (9) into the observer's model of the leader and plant—the following a single two-dimensional coupled SDE for the state $x(t) = (s(t), y(t))$:

$$\begin{aligned} dx(t) &= a(x(t)) dt + b dW(t), \\ x(0) &= (0, r(0)), \end{aligned} \quad (10)$$

where

$$a(x(t)) = a(s(t), y(t)) = \begin{bmatrix} 1 \\ -Ky(t) + 2f(s(t)) \end{bmatrix}, \quad (11)$$

$$b = \begin{bmatrix} b_1 \\ 0 \end{bmatrix}. \quad (12)$$

These diffusive clocks have positive (unit) drift, so they reach all positive levels, with probability one. Let δ be a positive number. We say the observer's diffusive clock *ticks* when the diffusion level first reaches $n\delta$, for all integers n . It follows that δ is the expected inter-tick interval. Now let N be a positive integer. We say that the observer's clock *chimes* when the diffusion level first reaches $kN\delta$ for all integers k . Chimes happen every N ticks, and the expected inter-chime interval is $N\delta$.

Measurements are taken at chimes. Ticks times determine the discretization of the observer's SDE model for the system using the Euler-Maruyama method. The observer's Kalman predict step generates state estimates corresponding to these discretization times. But the estimates are only based on measurements up to the previous chime, and are therefore available to the observer immediately after the previous chime. Thus while the observer does something at chimes (takes measurements), it does nothing of consequence at ticks. This situation will change when the observer becomes a follower: the ticks will trigger a discrete change in the actuation of the follower's control input $u_2(t)$.

Let η^k be the time of the k th chime. Then the measurement taken at that time is $y(\eta^k) + \zeta^k$, where ζ^k is a normal random variable with mean zero and variance R .

1) *Extended Kalman filter operation*: Let X_n denote the model's state at tick n —this quantity is a random variable because it depends on the actor's imperfect clock. For $n \in \{0, 1, \dots, N\}$, let $X_{n|k}$ denote X_{n+Nk} conditioned on all measurements up through chime k . The extended Kalman

filter computes $E[X_{n|k}]$ and $\text{Cov}[X_{n|k}]$, for all chimes k and for all inter-chime ticks $n \in \{0, 1, \dots, N\}$.

The initial condition $X_{0|0}$ has value $(0, r(0))$, with probability one. The Kalman filter alternates between predict steps and update steps, one each for each chime. The $k + 1$ st predict step computes, by recursion, the means and covariances of the random variables $X_{0|k}, X_{1|k}, \dots, X_{N|k}$. The $k + 1$ st update step derives the statistics of $X_{0|k+1}$ from the statistics of $X_{N|k}$. Both of these random variables (respectively, posterior and prior) correspond to the same tick—namely, chime $k + 1$, when a measurement occurs.

2) *Predict Step*: Given values for $E[X_{n-1|k}]$ and $\text{Cov}[X_{n-1|k}]$, we want to compute values for $E[X_{n|k}]$ and $\text{Cov}[X_{n|k}]$. Note that the chime number, k , remains constant throughout a predict step. In an abuse of notation, we will drop our reference to k and refer to $X_{n|k}$ as X_n .

We name the components of the expected state vector as follows:

$$E[X_n] = \hat{X}_n = \begin{bmatrix} \hat{s}_n \\ \hat{y}_n \end{bmatrix}. \quad (13)$$

We recall the nonlinear recursion given by the Euler-Maruyama stochastic integration method for time step e_{n-1} . Here e_{n-1} is the inter-tick interval, which is known in the classical application of the method. In our case, the time step is random and unknown, but its statistics are known. The recursion follows:

$$X_n = X_{n-1} + a(X_{n-1})e_{n-1} + b\sqrt{e_{n-1}}\nu_{n-1}. \quad (14)$$

Here ν_{n-1} is a standard normal random variable. We know the full distribution for e_{n-1} (inverse Gaussian), but we will only make reference to its first two moments—its mean, δ , and its standard deviation, σ . If we were to assume any other distribution for e_{n-1} , our derivation would still hold, upon substitution of the correct statistics. Regardless of the distribution chosen, we assume the inter-tick intervals are independent and identically distributed, and e_{n-1} , X_{n-1} and ν_{n-1} are mutually independent.

Now rewrite (14) as

$$X_n = g(X_{n-1}, e_{n-1}, \nu_{n-1}), \quad (15)$$

where

$$g(x, e, \nu) = x + a(x)e + b\sqrt{e}\nu. \quad (16)$$

Here $x \in \mathbb{R}^2$, $a(x) \in \mathbb{R}^2$, $b \in \mathbb{R}^2$, $e \in \mathbb{R}$, $\nu \in \mathbb{R}$. Now we linearize g around an arbitrary point (x_0, e_0, ν_0) :

$$\begin{aligned} g(x, e, \nu) &\approx g(x_0, e_0, \nu_0) \\ &+ \left. \frac{\partial g}{\partial x} \right|_{(x_0, e_0, \nu_0)} (x - x_0) \\ &+ \left. \frac{\partial g}{\partial e} \right|_{(x_0, e_0, \nu_0)} (e - e_0) \\ &+ \left. \frac{\partial g}{\partial \nu} \right|_{(x_0, e_0, \nu_0)} (\nu - \nu_0). \end{aligned} \quad (17)$$

Now, instead of an arbitrary point, we choose to linearize g around the expected value of $(X_{n-1}, e_{n-1}, \nu_{n-1})$ which is

$(\hat{X}_{n-1}, \delta, 0)$. If we evaluate the terms of (17), put everything together and simplify we obtain:

$$\begin{aligned} X_n &= (I + A_{n-1}^1 \delta)X_{n-1} - A_{n-1}^1 \delta \hat{X}_{n-1} \\ &\quad + A_{n-1}^0 e_{n-1} + b\sqrt{\delta}\nu_{n-1}, \end{aligned} \quad (18)$$

where

$$A_{n-1}^0 = a(\hat{X}_{n-1}) = \begin{bmatrix} 1 \\ -K\hat{y}_{n-1} + 2f(\hat{s}_{n-1}) \end{bmatrix}, \quad (19)$$

$$A_{n-1}^1 = \frac{\partial}{\partial x} a(\hat{X}_{n-1}) = \begin{bmatrix} 0 & 0 \\ 2f'(\hat{s}_{n-1}) & -K \end{bmatrix}. \quad (20)$$

Now replace the nonlinear recursion (14) with the linear recursion (18). We treat matrices A_{n-1}^0 and A_{n-1}^1 as constant (nonrandom) in computing the following expectations. This is the linearization approximation. Now find recursion for mean:

$$E[X_n] = E[X_{n-1}] + A_{n-1}^0 \delta. \quad (21)$$

Now find recursion for covariance, using the fact that the terms of (18) are mutually independent:

$$\begin{aligned} \text{Cov}[X_n] &= (I + A_{n-1}^1 \delta) \text{Cov}[X_{n-1}] \left(I + A_{n-1}^1 \delta \right)^T \\ &\quad + A_{n-1}^0 \text{Cov}[e_{n-1}] A_{n-1}^{0T} \\ &\quad + b\sqrt{\delta} \text{Cov}[\nu_{n-1}] \sqrt{\delta} b^T. \end{aligned} \quad (22)$$

Simplifying yields

$$\begin{aligned} \text{Cov}[X_n] &= \text{Cov}[X_{n-1}] \\ &\quad + \delta \left(A_{n-1}^1 \text{Cov}[X_{n-1}] + \text{Cov}[X_{n-1}] A_{n-1}^{1T} \right) \\ &\quad + \delta^2 A_{n-1}^1 \text{Cov}[X_{n-1}] A_{n-1}^{1T} \\ &\quad + \sigma^2 A_{n-1}^0 A_{n-1}^{0T} + \delta b b^T. \end{aligned} \quad (23)$$

3) *Update Step*: At the k th chime, we have arrived at real time η^k with state $x(\eta^k)$. Before the measurement, we estimate the state's statistics as $E[X_{N|k-1}]$ and $\text{Cov}[X_{N|k-1}]$. We now derive the Bayes update yielding $E[X_{0|k}]$ and $\text{Cov}[X_{0|k}]$. Call the prior statistics \bar{m} , and \bar{P} , respectively, and the posterior statistics m , P , respectively.

At the k th chime, the observer takes measurement y^k . The model of this measurement process is given by:

$$y^k = h(x(\eta^k), \zeta^k) = y(\eta^k) + \zeta^k, \quad (24)$$

where h is the function that evaluates the y -component of the state and adds ζ^k , assumed to be a normal random variable with mean 0 and variance R . Recall that $x = (s, y)$.

For the purposes of the extended Kalman filter, we model the measurement as follows:

$$\bar{y}^k = h(\bar{m}^k, \zeta^k = 0) = \bar{m}_y^k. \quad (25)$$

Now define the linearization matrices

$$H = \frac{\partial h(x, \zeta)}{\partial x} \quad (26)$$

$$= \begin{bmatrix} \frac{\partial h}{\partial s} & \frac{\partial h}{\partial y} \end{bmatrix} \quad (27)$$

$$= \begin{bmatrix} 0 & 1 \end{bmatrix}, \quad (28)$$

$$V = \frac{\partial h(x, \zeta)}{\partial \zeta} \quad (29)$$

$$= 1. \quad (30)$$

Because the measurement equation is linear, the matrices H and V do not depend on k .

We calculate the innovation mean:

$$v^k = y^k - \bar{y}^k. \quad (31)$$

We calculate the innovation covariance:

$$S^k = H\bar{P}^k H^T + V R V^T. \quad (32)$$

We calculate the Kalman gain:

$$K^k = \bar{P}^k H^T \{S^k\}^{-1}. \quad (33)$$

We calculate the updated mean:

$$m^k = \bar{m}^k + K^k v^k. \quad (34)$$

We calculate the updated covariance:

$$P^k = \bar{P}^k - K^k S^k \{K^k\}^T. \quad (35)$$

E. Both Saw, One Leads, The Other Follows

When the observer becomes a follower, its EKF-based estimate of the leader's time s together with the task representation through time $r(\cdot)$ will determine the follower's sawing. In this case we have three clocks: the leader's motor clock $s(t)$, the follower's sensory clock $q(t)$, and the follower's motor clock. Here the follower's motor clock is its estimate of the leader's motor clock $\hat{s}(t)$. All clocks start synchronized with real time at real time zero.

The leader's motor clock, $s(t)$, obeys SDE (1), and is never updated by feedback from the follower. The leader produces motor command $u_1(t)$, sawing away according to the shared task representation $r(\cdot)$ and directed by its own autonomous clock $s(t)$ – with no feedback from the plant or follower.

$$u_1(t) = -\frac{K}{2}y(t) + f(s(t)), \quad (36)$$

where $f(\cdot)$ is given by (4).

The follower produces motor command $u_2(t)$, according to the shared task representation $r(\cdot)$ but instead of using an autonomous clock to direct its sawing, it cooperates with the leader by using its estimate of the leader's clock $\hat{s}(t)$ to direct its sawing:

$$u_2(t) = -\frac{K}{2}\hat{y}(t) + f(\hat{s}(t)). \quad (37)$$

Now the plant equation is given by

$$\begin{aligned} \dot{y}(t) &= u_1(t) + u_2(t) \\ &= -\frac{K}{2}(y(t) + \hat{y}(t)) + f(s(t)) + f(\hat{s}(t)). \end{aligned} \quad (38)$$

The sensory clock q emits a tick at real times $\{t_0, t_1, \dots\}$, where $t_0 = 0$, where the inter-tick intervals are inverse Gaussian distributed with $E[t_n - t_{n-1}] = \delta$ and $\text{Var}[t_n - t_{n-1}] = \sigma^2$. The ticks time feedforward updates to $\hat{y}(t)$ and $\hat{s}(t)$, where the zero-order-hold rule is enforced. In other words, the follower's estimate, $\hat{y}(t)$ and $\hat{s}(t)$, are constant between ticks of the sensory clock $q(t)$. The follower's system model is still given by (10), (11) and (12). Chimes of the sensory clock trigger measurements of the plant and feedback updates to $\hat{y}(t)$ and $\hat{s}(t)$. A numerical example for following a single cycle of a sinusoidal reference is shown in Figure 2(A).

Figure 2(B) shows simulation for which the leader's clock and the follower's model of the leader's clock have different drift rates. This difference clock drift rates is equivalent to a different time-scaling of $r(t)$. This error in the model leads to a constant DC offset (mod 2π) that, eventually, although remarkably the cooperative performance is still achieved for long bouts. Future work will test a clock model that specifies $\dot{s}(t) = \text{white noise}$, allowing for constant, but unspecified drift, so that the eventually the follower's phase rate would adapt to that of the leader.

III. DISCUSSION

Do the imperfections of biological clocks matter? Are biological clocks, like quartz-crystal clocks, so precise that control theory can ignore² their imprecisions. The question is both experimental and theoretical. The question is experimental because we need to know just how precise biological clocks actually are. And the question is theoretical because we need to know just how precise biological clocks need to be to ignore their imprecisions. Of course there is no one answer to either of these questions. Biological clocks exist on many timescales [12] employing different mechanisms and undoubtedly possessing varying precision. Particularly on the timescales of motor control, these mechanisms are poorly understood, and a wide range of hypothetical mechanisms remain conceivable. For example, our follower synchronizes its clock by observing the continuous plant output rather than observing explicit synchronization events. It would be interesting to see how Bayes-optimal updates (of the follower's estimates of state) with observations of synchronization events, would improve performance. There is also an issue of uncertainty in producing and measuring such events—i.e. imprecisely timed events and imprecisely delayed measurements of such events may have the same effect.

Just as there are many clocks, there are many behaviors that rely on clocks. Different behaviors undoubtedly need different clock precisions and indeed use clocks in different ways. For example, pure state feedback around an equilibrium may work well without accurate timekeeping, whereas

²We thank an anonymous reviewer for pointing out that an accounting for uncertain time-keeping may also be needed in engineering applications in which fast time scales and/or long durations render significant the imperfections of engineered clocks.

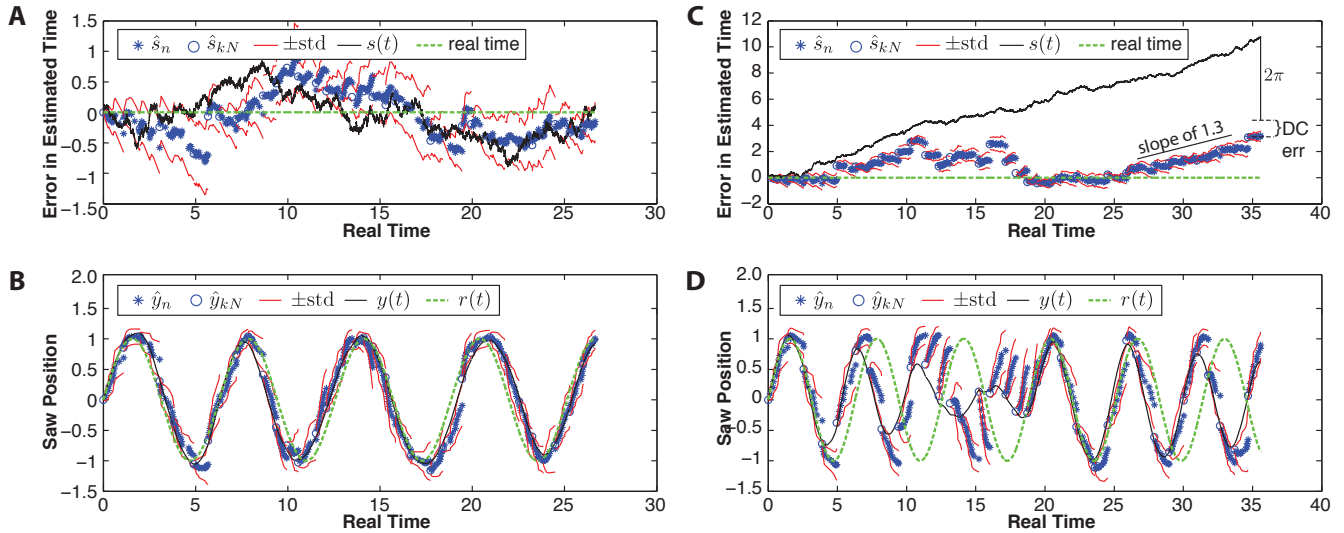


Fig. 2. (A,B): Simulation of a cooperative sawing task in which both agents have a shared representation, $r(\cdot)$ of the desired task, which is to follow a simple sine wave. The control law is such that both agents apply half the required control signal necessary to achieve perfect tracking (see Equation (7)), but when the two work together the output can be tracked. However, since both agents have noisy clocks, the “follower” must estimate the leader’s internal clock, to keep the two agents synchronized. The top panel shows the evolution of the second controller’s estimate of the first controller’s time, and the bottom panel shows that the pair of cooperating agents succeed in achieving this simple task. (C,D): Simulation of a cooperative sawing task in which the leader has a clock that follows $ds = 1.3 dt + b dW$, with a drift of 1.3, but the follower’s model still obeys (1) (i.e. the follower assumes that the drift rate is still 1.0, not 1.3). Between 0 and 5 seconds (real-time), the leader gets “ahead” of the follower and then the follower begins to keep pace. From 10 to 20 seconds the follower is out of sync with the leader, but there is phase reset at around real-time 20 s, at which time the follower locks in about 2π behind the leader, and it is back in sync. Notice that between 5 and 10 s, and again between 25 and 35, the follower’s clock advances with roughly the same slope (1.3) as the leader’s clock, although a persistent “DC” error is maintained, as expected since the follower does not adapt to discover the leader’s phase rate.

following a memorized trajectory may require accurate time-keeping in order to produce appropriately timed forces. Still it remains to be seen to what extent the imperfections in biological timekeeping make testable predictions about animal motor control, or constrain the possible controller schemes that can be used to explain animal behavior.

Concerning the different types of proposed mechanisms, a natural question arises: do different clocks lead to different predictions about estimation and control? This possibility should be seen as positive, for it would open the door to experiments that could test which type of neural clock (or clocks) underlie animal motor control. On the contrary, the discovery that different imperfect mechanisms make the same predictions would mean that we can continue developing our motor control theory with the simplest possible model.

IV. ACKNOWLEDGMENTS

Thanks to Dan Koditschek for thoughtful discussions on the topic of imperfect chronometry.

REFERENCES

- [1] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband, “Frequency comparison of two high-accuracy Al^+ optical clocks,” *Phys. Rev. Lett.*, vol. 104, p. 070802, Feb 2010. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevLett.104.070802>
- [2] E. S. Fortune, C. Rodríguez, D. Li, G. F. Ball, and M. J. Coleman, “Neural mechanisms for the coordination of duet singing in wrens.” *Science*, vol. 334, no. 6056, pp. 666–670, Nov 2011. [Online]. Available: <http://dx.doi.org/10.1126/science.1209867>
- [3] K. S. Lashley, “The problem of serial order in behavior,” in *Cerebral Mechanisms in Behavior*, L. A. Jeffress, Ed. New York: John Wiley & Sons, 1951.
- [4] D. V. Buonomano and U. R. Karmarkar, “How do we tell time?” *The Neuroscientist*, vol. 8, no. 1, pp. 42–51, 2002.
- [5] J. L. Barclay, A. H. Tsang, and H. Oster, “Interaction of central and peripheral clocks in physiological regulation.” *Prog Brain Res*, vol. 199, pp. 163–181, 2012. [Online]. Available: <http://dx.doi.org/10.1016/B978-0-444-59427-3.00030-7>
- [6] S. M. Reppert and D. R. Weaver, “Coordination of circadian timing in mammals.” *Nature*, vol. 418, no. 6901, pp. 935–941, Aug 2002. [Online]. Available: <http://dx.doi.org/10.1038/nature00965>
- [7] M. Zheng, W. O. Friesen, and T. Iwasaki, “Systems-level modeling of neuronal circuits for leech swimming.” *J Comput Neurosci*, vol. 22, no. 1, pp. 21–38, 2007. [Online]. Available: <http://dx.doi.org/10.1007/s10827-006-9648-7>
- [8] C. E. Carr, “Processing of temporal information in the brain.” *Annu Rev Neurosci*, vol. 16, pp. 223–43, 1993.
- [9] J. Gibbon, C. Malapani, C. L. Dale, and C. Gallistel, “Toward a neurobiology of temporal cognition: advances and challenges.” *Curr Opin Neurobiol*, vol. 7, no. 2, pp. 170–84, 1997.
- [10] D. V. Buonomano and R. Laje, “Population clocks: motor timing with neural dynamics.” *Trends Cogn Sci*, vol. 14, no. 12, pp. 520–7, 2010.
- [11] M. S. Matell and W. H. Meck, “Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes.” *Brain Res Cogn Brain Res*, vol. 21, no. 2, pp. 139–70, 2004.
- [12] D. V. Buonomano, “The biology of time across different scales.” *Nat Chem Biol*, vol. 3, no. 10, pp. 594–7, 2007.
- [13] D. P. King and J. S. Takahashi, “Molecular genetics of circadian rhythms in mammals.” *Annu Rev Neurosci*, vol. 23, pp. 713–42, 2000.
- [14] S. M. LaValle and M. B. Egerstedt, “On time: Clocks, chronometers, and open-loop control,” in *Proceedings of the 46th IEEE Conference on Decision and Control*, 2007, pp. 1916–1922.
- [15] D. M. Wolpert, K. Doya, and M. Kawato, “A unifying computational framework for motor control and social interaction.” *Philos Trans R Soc Lond B Biol Sci*, vol. 358, no. 1431, pp. 593–602, 3 2003.