

Estimating the phase of synchronized oscillators

Shai Revzen*

*PolyPEDAL Lab, Integrative Biology Department, University of California Berkeley,
3060 Valley Life Sciences, Building 3140, Berkeley, California 94720-3140, USA*

John M. Guckenheimer

Mathematics Department, Malott Hall, Cornell University, Ithaca, New York 14853-2401, USA

(Received 23 June 2008; revised manuscript received 21 August 2008; published 10 November 2008)

The state of a collection of phase-locked oscillators is determined by a single *phase* variable or cyclic coordinate. This paper presents a computational method, *Phaser*, for estimating the phase of phase-locked oscillators from limited amounts of multivariate data in the presence of noise and measurement errors. Measurements are assumed to be a collection of multidimensional time series. Each series consists of several cycles of the same or similar systems. The oscillators within each system are not assumed to be identical. Using measurements of the noise covariance for the multivariate input, data from the individual oscillators in the system are combined to reduce the variance of phase estimates for the whole system. The efficacy of the algorithm is demonstrated on experimental and model data from biomechanics of cockroach running and on simulated oscillators with varying levels of noise.

DOI: [10.1103/PhysRevE.78.051907](https://doi.org/10.1103/PhysRevE.78.051907)

PACS number(s): 87.10.-e, 87.19.lm, 87.85.gj

I. INTRODUCTION

Periodicity appears throughout the sciences and engineering. The simplest example of periodic dynamics is uniform rotation of the unit circle: $(x(t), y(t)) = (\cos(t), \sin(t))$. Any periodic orbit of a deterministic dynamical system can be mapped to the circle so that the angular velocity of the image is constant. We then call the angular variable on the circle a *phase* coordinate φ for the periodic orbit. We also call a real coordinate x a phase if $\varphi = x \bmod 2\pi$. Phase coordinates of a periodic orbit are determined up to a constant that one can regard as the point on the orbit with zero phase. Coupled systems, each of which has a periodic orbit, frequently *phase lock* to a dynamical state that itself has a periodic orbit. The periodic orbit of the coupled system projects onto the phase spaces of the component oscillators. When the systems are *weakly coupled*, these projections approximate the periodic orbits of the component oscillators. This paper presents a method to estimate phase variables for a broad class of phase-locked coupled oscillators from short spans of noisy multivariate data.

After providing some background from dynamics and biology, we specify the class of phase estimation problems we wish to solve and describe the particular difficulties other methods encounter in such a regime. We then describe our algorithm, and show the results of applying this algorithm to both real data and simulated data at various noise levels.

A. Background: Dynamics

The theory of nonlinear oscillators has been a useful tool of modern science and engineering for studying synchronized, periodic dynamics of physical and biological systems, with a large body of established work [1–4]. Within the con-

text of linear systems, frequency domain analysis based upon the fast Fourier transform (FFT) algorithm has become the primary tool for investigation of oscillatory physical dynamical systems. Complementary methods of “nonlinear time series analysis” rooted in dynamical systems theory have been developed that exploit qualitative features of periodic, quasi-periodic, and chaotic systems [5,6].

Noise reduction based upon fitting data to low-dimensional deterministic models [7,8] has played an important role in nonlinear time series analysis. The focus of nonlinear time series analysis has been on long time series, often of a single observed quantity. We examine data sets that are multivariate and shorter than those customarily considered in this literature. In addition to treating measurement noise superimposed on deterministic dynamical models, some recent work addresses identification of stochastic differential equation (SDE) models from data [9,10]. Our algorithm was developed with dynamical noise in mind, but our treatment of the topic does not utilize techniques of SDE.

Empirical methods have been developed for decomposing time series into generalized Fourier series [11] by using the Hilbert transform. These provide insight into using the Hilbert transform for phase recovery (see Sec. III B for a description of the transform). Some recent methods characterize phase and oscillator coupling in weakly coupled oscillators [12,13] by fitting a trigonometric series to the flow on the torus representing the state of the oscillators. This fitting problem is well-posed only when phase locking is sufficiently weak for the observations to cover the torus. The method requires data that contains all possible relative phases of the oscillators, in contrast to systems of phase-locked oscillators like those considered here.

Trajectories of phase-locked oscillations only sample states close to a closed curve that constitutes the periodic orbit of the system. Nonetheless, an essential feature of our method is that the measurement data are multivariate time series incorporating observations of all the oscillators within the system. With phase-locked oscillations, the different os-

*shrevz@berkeley.edu; <http://polypedal.berkeley.edu>

cillators operate coherently and have a single, global phase variable. The system can be regarded as a single “master oscillator” expressing itself through multiple nonlinear projections onto the phase spaces of the individual oscillators. The goal of the methods presented here is to estimate the phase of this master oscillator by combining the observations of the individual oscillators to reduce their measurement noise.

B. Background: Biology

Biological organisms rely upon complicated oscillatory dynamics for many vital processes. The heartbeat, respiration, the cell cycle, electrical activity of the nervous system, reproduction, and locomotion are all examples. The phase relations among coupled oscillators are critical to the function of these processes, as in the coordination of locomotion and the rhythmic pumping of the heart. At every scale there are important biological dynamics that exhibit periodic behavior and are modulated by noisy inputs. As in all physical sciences, our ability to measure the state of these oscillating systems is hampered by limited measurement precision.

Winfree’s classical book on biological time [4] utilized the mathematical theory of periodic systems to investigate biological problems. He argued for the value of studying phase response and phase resetting as a tool for experimental biology [14]. Cohen, Holmes, and Rand [15] presented a mathematical model for phase reduction of a network of neural oscillators generating fictive locomotion in a lamprey [16]. More recently, Golubitsky *et al.* [17] predicted constraints on the structure of central pattern generators in the spinal cords of vertebrates due to symmetries in systems of coupled oscillators. Together these point to the potentially large payoff of having an effective means of characterizing phase.

The *Phaser* algorithm presented here grew from our attempts to model the locomotion of a running cockroach (*Blaberus discoidalis*) as a system of coupled nonlinear oscillators. Cockroaches running on treadmills were filmed from below, and motion capture techniques were used to measure the moving positions of all *tarsi* (feet) relative to the cockroach body. When running, cockroaches use their six legs in two sets of three legs moving alternately. This alternating tripod gait can be viewed as six nonidentical oscillators cycling in a phase locked mode. The oscillators are non-identical as leg geometries and motions are different in front, middle, and hind legs and some individual animals even demonstrate differences between their right and left side on the same body segment.

A rich literature on mathematically modeling the dynamics of legged locomotion in general, and cockroach locomotion in particular, is reviewed in [18]. Revzen, Koditschek, and Full [19] proposed an assay of perturbation tests for deciding among possible neuromechanical control architectures employed in a given rhythmic behavior. This assay is based on the ability to estimate the phase of the animal in its locomotive cycle, and our *Phaser* algorithm developed out of attempts to improve this estimation procedure.

Most previous work in biology, dating as far back as the invention of stop motion photography [20], describes gaits

used by running animals in terms of discrete events such as touch-down and lift-off. Phase response curves of stick insects, one of the best understood arthropod locomotor systems, were developed using *anterior extreme position* (AEP) events as a measure of phase [21,22]. However, animal motions are mechanical and governed by dynamics that obey piecewise continuous (or smoother) equations of motion, suggesting a finer approach to phase estimation is possible.

Averaging data to produce a *typical cycle* may introduce systematic errors unless this averaging is done with respect to phase [23]. A good typical cycle with confidence intervals allows perturbations away from the typical cycle to be studied, and used for analysis of the control maintaining the cyclic behavior. The common practice of averaging trajectories based on some start-of-cycle boundary event and linearly scaling cycles to match durations implicitly assumes that state changes at a uniform rate, otherwise the size of the bins will be unequal in a systematic way in each cycle, leading to potential statistical errors. Furthermore, if the rate of oscillation is variable then bins far from a cycle boundary will have larger uncertainty in the independent parameter than a bin close to the event, making valid statistical inference more difficult.

Full and Koditschek [24] hypothesize that control of animal locomotion is based on the stabilization of an isolated cycle as the attractor of animal dynamics. Their approach is unusual in that very few investigators have treated the nervous system and body mechanics as a single neuromechanical oscillator, despite the mathematical elegance of this approach. Our work also models the behavior of the freely running organism as a simple dynamical system—a network of synchronized oscillators with an isolated limit cycle.

II. PHASE ESTIMATION PROBLEM

Our phase estimation problem is shaped by limits on sample sizes, durations, and signal-to-noise ratios of empirical data. Other methods of phase estimation may be more effective in different ranges of these parameters than those considered here. We discuss the importance of the various data constraints we face as we introduce them.

A. Dynamical system

The starting point for the analysis presented here is a dynamical system defined by a piecewise smooth vector field F on a manifold \mathcal{X} that depends on parameters $\mu \in \mathbf{B}^m \subset \mathbb{R}^m$. The equations of motion are

$$\dot{\mathbf{x}} = F(\mathbf{x}, \mu); \quad (\mathbf{x}, \mu) \in \mathcal{X} \times \mathbf{B}^m. \quad (1)$$

The parameters μ are included to account for the differences between individuals and changes between experiments with the same individual.

The flow constructed from the solutions of this equation will be denoted $\Phi: \mathcal{X} \times \mathbb{R} \times \mathbf{B}^m \rightarrow \mathcal{X} \times \mathbf{B}^m$, so that $\Phi(\mathbf{x}_0, t; \mu)$ is the solution to Eq. (1) for the initial condition $(\mathbf{x}_0; \mu)$. We are interested in the case where throughout the parameter domain $\mu \in \mathbf{B}^m$, the system equation (1) has an isolated asymptotically stable periodic orbit (cycle) $\mathbf{p}(t, \mu)$ of period

$T(\mu)$ for every choice of μ . We assume further that these \mathbf{p} are stable in the sense that the return maps of the cycles have eigenvalues bounded in magnitude by $1-\varepsilon$ (for some $\varepsilon>0$ valid for all $\mu \in \mathbf{B}^m$). For clarity, we elide the μ in the sequel wherever convenient.

In specifying further constraints, we choose coordinates that position and scale the cycles so that

$$\int_{t=0}^T \mathbf{p}(t;0) = 0, \quad (2)$$

$$\frac{1}{T} \int_{t=0}^T \|\mathbf{p}(t;0)\| = 1. \quad (3)$$

Denote by \mathcal{C} the locus of all these cycles in $\mathcal{X} \times \mathbf{B}^m$. Since \mathbf{B}^m is simply connected, \mathcal{C} is homeomorphic to $S^1 \times \mathbf{B}^m$, i.e., is shaped like a cylinder.

B. Definition of phase

There is a natural notion of *asymptotic phase* $\phi_\infty: \mathcal{X} \rightarrow \mathcal{C}$ satisfying $\lim_{t \rightarrow \infty} \|\Phi(\mathbf{x}_0, t; \mu) - \Phi[\phi_\infty(\mathbf{x}_0; \mu), t; \mu]\| \rightarrow 0$ defined for initial values \mathbf{x}_0 and parameters μ close to \mathcal{C} . Φ maps every initial value to a point in \mathcal{C} whose trajectory it approaches with increasing time. The inverse images of points on \mathcal{C} under this mapping are called *isochrons* [25,4].

We want to define phase as a complex number of modulus 1 rather than an element of \mathcal{C} . Our phase variable should satisfy the property that equal phase angles are traversed in equal proportions of the periods of the periodic orbits. Since the phase variable of a single periodic orbit is only determined up to a constant, the definition of this phase requires the choice of a cross section of \mathcal{C} transverse to the flow. We assume that the section consists of the points of \mathcal{C} that are chosen to have phase zero. Any function $\mathcal{C} \rightarrow S^1$ which satisfies the traversal rate requirement and maps the zero phase section to 1 may be extended to a phase coordinate $\phi: X \subseteq \mathcal{X} \times \mathbf{R}^m \rightarrow S^1$ for the entire stability basin X of \mathcal{C} by requiring that it is constant on isochrons.

C. Measurement model

We assume that empirical data from which we estimate a phase variable consists of N multivariate time series $\{\mathbf{y}^{(k)}(t)\}$, $k=1, \dots, N$, that are stochastic perturbations of trajectories $\{\mathbf{x}^{(k)}(t)\}$ of the vector field F . The trajectories $\{\mathbf{x}^{(k)}(t)\}$ are each defined on an interval $[t_L^{(k)}, t_U^{(k)}]$ and lie close to \mathcal{C} . For clarity we suppress the $^{(k)}$ superscripts in the sequel. Each time series is sampled at regular, small intervals $\Delta t < 0.1T$. The cockroach locomotion trajectories that we analyze extend over fewer than 100 cycles, and the signal-to-noise ratio of the measurements is roughly 10:

$$\mathbf{y}(t_j) = \mathbf{x}(t_L + j\Delta t) + \eta, \quad (4)$$

$$\eta \sim \mathcal{N}(0, Q),$$

$$\Delta t < 0.1T, \quad (5)$$

$$t_U - t_L < 100T. \quad (6)$$

Note that this statistical model of the data does not include stochastic terms in the vector field F , though we are interested in extending our methods to that case and others described below.

Noise reduction is often tied to the selection of a metric for the space of data points. A metric can be used to define an orthogonal projection onto a set of “noiseless” states, or it can be used to define a centroid for a cluster of related noisy measurements; it is used in both these capacities in [26]. Kantz and Schreiber [5] note that the choice of a metric for multivariate datasets is by no means obvious.

We propose that the covariance matrix Q of the measurement noise η can be used to compute a useful metric. If measurement noise is Gaussian, the induced Mahalanobis distance [27] is a metric related by a simple monotone change of variables to the probability density. Likelihood maximization with respect to the measurement noise distribution is distance minimization with respect to the Mahalanobis metric. A linear change of variables converts the Mahalanobis metric to the Euclidean (l_2) norm. The transformation is given by multiplication by a square root of Q^{-1} . The new variables are *scores*—they represent likelihood levels similarly to the well-known z *scores* used with scalar Gaussians. The covariance of measurement errors for the scores is the identity matrix, so they are linearly uncorrelated. To the extent that this decorrelation carries through to phase estimate errors, if we now obtain a phase estimate from each score we may have reasonable faith that averaging these estimates will decrease the overall phase estimation error.

It should be noted that the Q matrix is typically diagonal as measurement errors of separate instruments are not usually correlated, but treatment of the more general case shown here is no more complicated than that of the diagonal covariance matrix case.

The conditions of Eqs. (4)–(6) are far from the typical conditions in which Fourier transform methods are traditionally recommended for estimating periodic structure. First, the length limit on the samples implies the appearance of significant windowing artifacts in Fourier methods. Second, there is no guarantee that the parameters $\mu^{(k)}$ are “close” for different values of k —only that the ensemble is “close” to \mathcal{C} —providing no obvious way of averaging or otherwise combining samples into a single dataset for phase coordinate estimation. Third, the dependence of period T on the parameters μ means that phase signals will lose coherence with a Fourier basis of fixed period, diminishing the largest Fourier series components.

D. Estimation problem specification

The problem we wish to address is to estimate a phase function ϕ from the data. We have a dataset consisting of an ensemble of sequences of multivariate measurements $\mathbf{y}(t_j)$, which we believe to have been derived from some (unobserved) state space trajectories $\mathbf{x}(t)$, and therefore have an associated sequence of phase values $\phi[\mathbf{x}(t_j)]$. We would like to compute values $\hat{\phi}(t_j)$ that are “good” estimates of $\phi[\mathbf{x}(t_j)]$

for an arbitrarily selected phase coordinate function ϕ . The arbitrary part of ϕ is a choice of gauge, and has no physical significance for the dynamics. For any two phase coordinate functions ϕ and $\tilde{\phi}$, the difference is a constant on any trajectory, and thus the gauge is fully specified by values ϕ takes on a transverse section of \mathcal{C} .

The gold standard for statistical estimators is to find a maximum-likelihood estimate (MLE). However finding an exact MLE of the phase is a very difficult problem. Instead, we seek what is merely a “good” estimate $\hat{\phi}$, possessing the following properties.

- (i) $\hat{\phi}$ has little or no bias as an estimate for ϕ .
- (ii) The residual $\hat{\phi} - \phi$ has small variance compared to the noise.
- (iii) The distribution of the residual $\hat{\phi} - \phi$ is Gaussian.

The first two properties are required for all consistent estimators if the underlying probability distributions are not pathological [28]. The third property is important for both practical and theoretical reasons. The practical motivation is that we estimate phases for the purpose of making a decision regarding some hypothesis, and the tools for making decisions with Gaussian uncertainty are well-understood, making estimators with Gaussian errors more attractive. The theoretical motivation is that an MLE of a random variable depending on a smooth probability distribution always has a Gaussian error distribution at the low noise limit, and so any MLE we find is sure to have this property.

III. PHASER ALGORITHM

Our phase estimation algorithm Phaser builds on the work of Kraleman *et al.* [12] but applies to the estimation of phase from multidimensional data produced by a phase locked system of coupled oscillators, a different regime from that investigated by Kraleman *et al.* The work of [12] focuses on the coupling of angle variables of an invariant torus, given a univariate measurement for each of the participating oscillators. Their approach is in line with much of the non-linear analysis described in [5], which is focused on dynamics whose dimension may be equal or higher than the dimension of measurements, and therefore some thought is required in the application of delay coordinate selection. However, we are concerned with a problem in which the invariant set is known to be S^1 and nonlinear embeddings of this object in the space of measurements are easy to come by. Rather than finding an embedding and a statistic of the system attractor (as in [5]), our problem is one of locally inverting the embedding in a way that reduces the variance of the phase residuals due to noise.

Our algorithm can be used both to fit a phase estimator to training data and to obtain a phase estimate of novel data after the estimation parameters for a system have been fit. The algorithm has the following steps.

- (1) Correlated measurements are transformed to independent scores with equal variance.
- (2) Protophases are computed from the component univariant time series of the scores.
- (3) A Fourier series based correction is applied to the protophases producing phase estimates for the individual oscillators.

(4) Principal component analysis of the phase estimates of the individual oscillators yields a single protophase for the coupled system.

(5) The Fourier series correction of the protophase is used to compensate for remaining systematic errors.

The following sections describe these steps in detail when used with training data. When applying the estimator to compute phases for novel data, the process is similar, except that steps (3)–(5) use the transformations determined from the training data.

A. Metrization

We assume that the covariance matrix Q of the measurement noise has been computed. This symmetric matrix is used to define an affine transformation of the time series data into its z scores via

$$\mathbf{z} \triangleq Q^{-1/2}(\mathbf{y} - \langle \mathbf{y} \rangle) = U^T \Lambda^{-1/2} U(\mathbf{y} - \langle \mathbf{y} \rangle), \quad (7)$$

where $U^T \Lambda U = \text{svd}(Q)$ is a singular value decomposition of the symmetric positive definite matrix Q . The norm $\|\mathbf{z}\| = \sqrt{\mathbf{z}^T \cdot \mathbf{z}} = \sqrt{\mathbf{y}^T Q^{-1} \mathbf{y}}$ is the Mahalanobis distance defined by the Gaussian with covariance Q . The transformation to z scores yields coordinates that have uncorrelated measurement noise, of similar variance.

B. Protophases

Given a d -dimensional real time series of z scores $\mathbf{z}(t_j)$, we define their protophases $\theta_k(t_j)$ as the complex argument of their Hilbert transform:

$$\zeta_k(t_j) \triangleq \mathcal{H}[\mathbf{z}_k(t_j)], \quad (8)$$

$$\rho_k \triangleq \langle |\zeta_k| \rangle, \quad (9)$$

$$\theta_k(t_j) \triangleq \arg[\zeta_k(t_j)]. \quad (10)$$

Recall that the Hilbert transform [11] is a linear operator which creates a complex analytic time series from a real time series so that the original series is the real part of the complex series. It converts sines and cosines into complex exponentials of the appropriate phase and frequency. The Hilbert transform of a periodically oscillating signal of mean 0 is a path in the complex plane that winds around 0. If the signal has the same number of zero crossings and extrema (up to ± 1) and alternate extrema are nearly equidistant from zero, then the Hilbert transform will always wind in the counter-clockwise direction and our protophase will be an increasing function of time. Huang *et al.* [11] define the notion of an *intrinsic mode function* (IMF) with these guarantees, but their definition of an IMF requires a specification of an envelope of local maxima and minima which is immaterial to our purpose. We use the term IMF loosely to represent functions whose Hilbert transform winds around the origin at a rate bounded from below above zero.

The mean modulus ρ_k of the Hilbert transform is of particular importance in cases where \mathbf{z} series are not IMF. In those cases, ζ_k loops back: θ_k decreases at some times and there is an associated decrease in $|\zeta_k|$. The decrease in ρ_k

induced by this effect penalizes non-IMF components of \mathbf{z} when the phase estimates are combined (see below), making the algorithm more robust.

C. Series correction operator $\mathcal{P}[\cdot]$

As this section describes a computation which applies to a single component θ_k , we will use the abridged notation $\theta_j \triangleq \theta_k(t_j)$. The next step of the Phaser algorithm is a nonlinear transformation of the protophase coordinates to variables that increase at a constant rate. There is no reason for a protophase variable θ to have increments proportional to the time steps. If it did, phase estimation would only require the trivial computation of a constant $\omega = \langle \frac{d\theta}{dt} \rangle$. For a true phase coordinate φ we have $\dot{\varphi} = \omega$, some constant.

For each protophase variable θ , we seek a function $\phi: S^1 \rightarrow S^1$ such that $\frac{d}{dt}\phi[\theta(t)] = \omega$, implying

$$\frac{d\phi}{d\theta}[\theta(t)] = \omega \frac{dt}{d\theta}(\theta), \quad (11)$$

$$\phi(\theta) = \omega \int_{\alpha=0}^{\theta} \frac{dt}{d\theta}(\alpha) d\alpha, \quad (12)$$

following a procedure similar to that leading to Kraleman *et al.* [[12], Equation (16)]. Eq. (12) may be used to find ϕ , but we do not compute this integral directly as in [12] because the integration of a discrete approximation to $\frac{dt}{d\theta}$,

$$\frac{dt}{d\theta} \left(\frac{\theta_j + \theta_{j+1}}{2} \right) \approx \frac{\Delta t}{\Delta \theta_j} = \frac{\Delta t}{\theta_{j+1} - \theta_j}, \quad (13)$$

has poor statistical properties, as follows.

The measurement noise of θ may be of magnitude comparable to $\Delta\theta$, implying that the values of Eq. (13) may be distributed in a *ratio distribution* with heavy tails. Such distributions have slowly convergent or even divergent first moments, rendering averaging useless for improving the estimate. To avoid this problem, we estimate and smooth $\frac{d\theta}{dt}$ instead of $\frac{dt}{d\theta}$.

We want to express $\frac{d\theta}{dt}$ as a function of θ rather than t . Denoting this function by f and its truncated Fourier series approximation of order N by $\mathcal{F}_N[f]$, we estimate the Fourier coefficients of $\mathcal{F}_N[f]$ by first sorting all $\langle s, f(s) \rangle$ pairs in the input by s and then integrating over consecutive pairs using the trapezoidal rule

$$\int_{s \in S^1} e^{-iks} f(s) ds \approx \sum_j e^{-ik(s_j + s_{j+1})/2} \frac{f(s_j) + f(s_{j+1})}{2} (s_{j+1} - s_j). \quad (14)$$

We then invert the Fourier series, obtaining coefficients $\{C_k \in \mathbb{C}\}_{k=-N, \dots, N}$ such that

$$f(\theta) \triangleq \mathcal{F}_N \left[\frac{d\theta}{dt}(\theta) \right], \quad (15)$$

$$\frac{d\theta}{dt}(\theta) = \sum_{k=-N}^N C_k e^{ik\theta} = \mathcal{F}_N \left[\frac{1}{f(\theta)} \right]. \quad (16)$$

The measurement noise was reduced in the averaging implicit in computing the Fourier coefficients for $f(\theta)$.

Substitution of Eq. (16) into Eq. (12) gives

$$\hat{\phi}(\theta) = \omega C_0 \theta - \sum_{k=1}^N \frac{\omega}{ik} (C_k e^{-ik\alpha} - C_{-k} e^{ik\alpha}) \Big|_{\alpha=0}^{\theta}. \quad (17)$$

We have assumed that the protophase θ does not lose or gain cycles relative to the actual phase ϕ , so we conclude $C_0 = \omega^{-1}$. $\mathcal{F}_N[\cdot]$ was applied to a real time series, thus $C_{-k} = C_k^*$ giving $C_{-k} e^{ik\alpha} = (C_k e^{-ik\alpha})^*$ yielding the following expression for the phase of oscillator:

$$\hat{\phi}(\theta) = \theta - \frac{1}{C_0} \sum_{k=1}^N \frac{2}{k} \text{Im} (C_k e^{-ik\alpha}) \Big|_{\alpha=0}^{\theta}, \quad (18)$$

with the sum correcting for systematic errors in $\hat{\phi}(\theta)$ as a function of θ using coefficients that we can readily compute. The equation is similar to Eq. (16) of [12], except that the coefficients C_k were obtained via a process better conditioned for dealing with measurement noise. We denote the entire process taking θ to $\hat{\phi}$ by the operator $\mathcal{P}[\cdot]$, writing $\hat{\phi} = \mathcal{P}[\theta]$.

D. Combining multiple estimates of phase

The final step of our algorithm is to combine the phase estimates of individual oscillators into a single, improved estimate for a phase variable ϕ of the phase-locked coupled oscillator system. We assume that the phase estimate $\hat{\phi}_k$ obtained from the \mathbf{z}_k component oscillator is noisy and seek to average these phases so as to improve the signal-to-noise ratio (SNR). We do this by first reconstructing a trajectory in \mathbb{R}^{2N} from the $\hat{\phi}_k$. The trajectory will be an ellipse in the ideal case that the $\hat{\phi}_k$ are functions of time with constant slope ω . We then introduce a procedure that is designed to be optimal in the case that the residual of $\hat{\phi}_k - \phi_k$ is Gaussian:

$$\hat{\phi}_k(t_j) = \hat{\phi}_k(t_L) - \phi(t_L) + \phi(t_j) + \nu = \hat{\phi}_k(t_L) + \omega(t_j - t_L) + \nu, \quad \nu \sim N(0, \sigma). \quad (19)$$

Since the noise ν is related to the instrument measurement noise η [of Eq. (4)] by a nonlinear transformation, its distribution is unlikely to be Gaussian. Nonetheless, the distribution will be approximately Gaussian if the variance of ν is small.

We wish to reduce the noise ν in $\hat{\phi}$. We cannot simply average the complex phases of the component oscillators because the values of $\hat{\phi}_k(t_L)$ are not known, and these relative phase shifts may cause phase to interfere destructively when the averages are computed, decreasing the signal-to-noise ratio. As an extreme example, with only two measurements having $\hat{\phi}_1(t_L) = 0$ and $\hat{\phi}_2(t_L) = \pi$, in their average the signal will be completely lost and only noise will remain.

A practical approach, suggested by the computation use in an example in [12], is to average the *unwrapped* phases, i.e., the phases taken as continuous angles in \mathbb{R} (not $[-\pi, \pi]$). Such averaging implicitly assumes that the estimation error variance of these phase estimates is equal, and may thus be improved upon by judicious choice of weights (see below regarding weights for maximum-likelihood estimates from simultaneous measurements). In addition, averaging is susceptible to any correlations in the estimation errors. To filter out the noise even further and remove some of the correlated errors, we suggest the projection method below. Comparison of these three approaches is found at the end of Sect. IV.

We represent each phase estimate time series $\hat{\phi}_k(t)$ by two real time series $\hat{q}_{2k}(t), \hat{q}_{2k+1}(t)$ at phase angles of $\frac{\pi}{2}$, with a magnitude $\rho_k = \sqrt{\hat{q}_{2k}^2(t) + \hat{q}_{2k+1}^2(t)}$ previously obtained from Eq. (9).

$$\begin{aligned}\hat{q}_{2k} &\triangleq \rho_k \sin(\hat{\phi}_k), \\ \hat{q}_{2k+1} &\triangleq \rho_k \cos(\hat{\phi}_k).\end{aligned}\quad (20)$$

The relative magnitudes of ρ_k represent our confidence that the associated $\hat{\phi}_k$ is a “good” protophase variable, i.e., winds around the origin with a strictly positive rate.

The choice of weights ρ_k is based on the formula for the maximum-likelihood estimate (MLE) for a quantity obtained from independent measurements with Gaussian error distributions of different variances. Let $Z = \{\mathbf{z}_k\}_{k=1}^n$, $\mathbf{z}_k \in \mathbb{R}^d$ be n independent noisy measurements of some unknown constant $\mathbf{x} \in \mathbb{R}^d$. Assume these measurements each have isotropic Gaussian measurement errors of variance $\boldsymbol{\sigma} = \{\sigma_k\}_{k=1}^n$. The joint distribution $P(Z|\mathbf{x}, \boldsymbol{\sigma}) = \prod_{k=1}^n \mathcal{N}(\mathbf{z}_k|\mathbf{x}, \sigma_k)$ may be used to obtain the MLE for the value of \mathbf{x} . This \mathbf{x}_{MLE} is a critical point of the probability density, and thus also of the log-likelihood:

$$\ln P(Z|\mathbf{x}, \boldsymbol{\sigma}) = - \sum_{k=1}^n \frac{\|\mathbf{z}_k - \mathbf{x}\|^2}{\sigma_k^2} \quad (21)$$

$$- \frac{\partial}{\partial x_j} \ln P(Z|\mathbf{x}, \boldsymbol{\sigma}) = \left(x_j \sum_{k=1}^n \frac{2}{\sigma_k^2} \right) - \left(\sum_{k=1}^n \frac{2[\mathbf{z}_k]_j}{\sigma_k^2} \right), \quad (22)$$

$$\mathbf{x}_{\text{MLE}} \sum_{k=1}^n \frac{1}{\sigma_k^2} = \sum_{k=1}^n \frac{\mathbf{z}_k}{\sigma_k^2}. \quad (23)$$

Our ζ_k of Eq. (8) are distributed in an annulus around the origin, scaled so that the width of the annulus is equal across k . Let us imagine the path ζ_k traces to be that of an ideal phase variable (a circular path) with additive, radially symmetric measurement noise. Under these assumptions the average radius ρ_k of Eq. (9) is inversely proportional to the magnitude of noise in the polar angle θ_k which constitutes the protophase. We conclude that under these idealized conditions the choice of weights for the q_k would give the maximum-likelihood estimate of the quantity $e^{i\varphi}$.

When θ_k is nonmonotone, as is typically the case for non-IMF \mathbf{z}_k , the segments in which θ_k decreases correspond to loops with a reduced $|\zeta_k|$. Thus whenever the signals \mathbf{z}_k are not IMF, that signal’s ρ_k is penalized because the average $\langle |\zeta_k| \rangle$ is smaller than one would obtain for a sinusoid of similar peak-to-peak amplitude.

Returning to Eq. (19), one may conclude that the noiseless version of Eq. (20) comprises sinusoids of the same frequency ω at various phase offsets and amplitudes. The space of all such sinusoids is spanned by $\sin(\omega t)$ and $\cos(\omega t)$, and therefore the points occupy a two-dimensional linear subspace independent of the number of oscillators N .

For our actual, noisy $\hat{\mathbf{q}}$, we identify a candidate two-dimensional subspace by projecting onto the subspace spanned by the two largest principal components. We regard the projection of $\hat{\mathbf{q}}$ as a combined protophase for the full system. This protophase is then series-corrected with the $\mathcal{P}[\cdot]$ operator to obtain our final resulting phase estimate

$$\hat{\phi}(t_j) = \mathcal{P}\{\arg[(\mathbf{pc}_1 + i\mathbf{pc}_2)^T \hat{\mathbf{q}}(t_j)]\}, \quad (24)$$

with \mathbf{pc}_1 and \mathbf{pc}_2 the first two principal components of $\hat{\mathbf{q}}$.

IV. EXAMPLES

We tested the performance of our Phaser algorithm on three data sets from both experimental and synthetic sources. These three data sets explore different aspects of the efficacy of Phaser.

The first example is derived from motion capture of a running cockroach. We compare the results of Phaser with other methods used in biological and biomechanical studies. We do not know whether the cockroach is an ensemble of synchronized oscillators nor do we have a reference phase to use for computing estimation errors of this data set. The results show some of the qualitative features of analysis with each method, motivating our development of Phaser.

A second example is obtained by adding noise to a deterministic periodic trajectory of a Fourier series model of the cockroach data from the first example. The results show that for systems with little or no phase drift a slight improvement to the commonly used AEP method can provide comparable results to Phaser, but naive AEP estimation is subject to systematic errors.

The third example is a stochastic differential equation that introduces noise to a well-studied dynamical system—the Hopf oscillator. The resulting data is used to compare Phaser to the use of a naive phase estimator and the phase estimator of [12] which takes a univariate input, allowing us to demonstrate how Phaser takes advantage of the extra information available to it from a multivariate input. This synthetic class of systems is also used to demonstrate the significance of the choice of weights ρ_k and the use of the projection in Eq. (24).

We present the results of our comparative analysis with two types of plots: plots of the *residual phase* and error density histograms. Residual phase is obtained by first computing a linear regression fit to the measured or known phase as in Revzen *et al.* [19]. This fitted constant frequency model is then subtracted from the measured phase leaving a plot of the fitting residual. The residual plot shows how the recov-

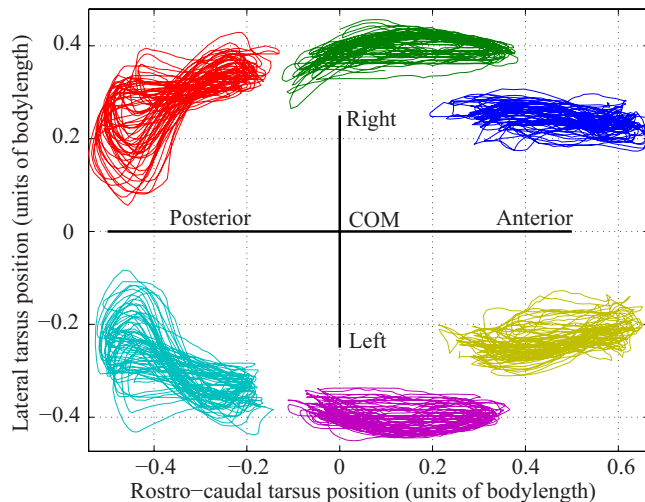


FIG. 1. (Color online) Tracks of the horizontal projection of cockroach (*Blaberus discoidalis*) foot motions in the body frame of reference. The origin is placed at the animal's center of mass. This running trial was recorded at 500 frames per second. It is 2045 frames long and includes over 37 strides of the animal running in characteristic tripod gait. Measurement precision of foot positions is approximately 0.04 body lengths.

ered phase differs from an idealized constant frequency oscillator. For the cockroach data, the fitting residuals combine phase estimation errors and drifts of the phase due to slow changes in the frequency of the cockroach leg motions. For the synthetic models where a reference phase governing the deterministic (nonstochastic) part of the system dynamics is known, the linear regression is taken against this reference phase rather than against the phase estimates.

Error density histograms plot the log of frequency of each bin of phase estimation errors relative to the reference phase. This is only possible in the synthetic examples where the reference phase is known. Our requirements of an estimator to have a low bias, small variance, and Gaussian residual are expressed by error histogram curves with a maximum near 0, a narrow peak, and a parabolic shape, respectively.

A. Cockroach empirical data

Revzen and colleagues developed an assay that helps identify the neuromechanical control architecture of a locomotor behavior [19] from perturbation experiments, and applied these procedures to cockroach locomotion data [29,30]. These studies provided the initial impetus to develop the Phaser algorithm. We present the results of applying Phaser and several other methods of phase estimation to empirical data. The raw empirical data consisted of two-dimensional tarsus (foot) positions for all six legs of a running cockroach (*Blaberus discoidalis*) relative to the animal's body frame of reference (see Fig. 1).

In the case of real animals we have no model limit cycle and reference phase to compare to our data. Instead we compare Phaser output to other methods from the literature. It should be noted that although we treat each leg of the running animal as a single oscillator, evidence shows that the




FIG. 2. (Color online) Residual phase of the running cockroach from Fig. 1 showing the deviation of the various phase estimates from a linear regression fit. The duration of the plot is slightly over 37 strides (cycles). The residual phases estimated with different methods are offset from each other for clarity. No importance should be attributed to the choice of phase at time 0. All phase estimates show very similar large scale structures, but differ greatly in their short time-scale details. Phases estimated from anterior extreme position (AEP) events show their characteristic polygonal line structure, with a six-legged mean AEP showing far superior temporal resolution to the estimate derived from middle right leg AEP. The estimate based on PCA and our Phaser show high frequency structure, but the PCA high frequency effects seem far less cross-correlated at short lags—suggesting noise rather than dynamic structure as the cause.

motor signals for each joint are neurally independent, at least in some arthropods [31], and coupled only via mechanosensory feedback. The motions within a leg are highly correlated while cockroaches are running, so our simplification does not sacrifice much fidelity of representation.

We compare Phaser with three other methods of obtaining phase, described in detail below.

(1) Phase from the anterior extreme position events (see below) of a single leg, a method used in many neuroethological studies.

(2) Phase from projection on principal components, an approach in line with methods currently used by clinical biomechanists [32].

(3) Average of the anterior extreme position phase estimates from all legs.

Much of the work in biology has used the anterior extreme position (AEP) events of a single leg (e.g., Fig. 2-AEP Rgt Mid uses middle right leg AEP) as a source of phase estimates. An AEP event is defined as the time when a leg reaches the anteriormost position in a given cycle. Phase between AEP events is interpolated linearly, providing a piecewise linear estimate. However, a simple argument shows that despite its visual convenience, the AEP event is a poor choice for phase estimation. AEP is defined by a local maximum in the anterior-posterior position $r(t)$ of a foot relative to the body. As such, the time derivative of the position $\frac{dr}{dt}(t)=0$, thus $\frac{dr}{d\phi} \cdot \frac{d\phi}{dt}=0$. Since at all times $\frac{d\phi}{dt}=\omega > 0$

we have $\frac{dr}{d\varphi} = 0$, so that $\frac{d\varphi}{dr}$ is locally unbounded. Because the first order sensitivity of phase to measurement error is infinite at the AEP, phase estimation errors are governed by the second and higher derivatives of $r(\varphi)$ and grow as $\sqrt{\Delta r}$ or worse (see [33]).

A slight improvement can be gained by averaging the AEP derived polygonal lines expressing the phase estimates from all six legs (averaging phases from all component oscillators; see Fig. 2-mean AEP). With only six events per cycle the number of phase estimate values per unit time still gives a substantially coarser temporal estimate than that obtained from our algorithm with sample rates larger than 50 samples per cycle.

Phase can also be estimated by a dimension reduction to two dimensions (2D) using principal component analysis (PCA) followed by a polar decomposition. The use of PCA for analysis of kinematic data has recently gained popularity in the clinical biomechanics community [32].

It has been known for a while that without the velocities cockroach kinematic data has only one significant principal component [34], thus we used $(\mathbf{y}, \dot{\mathbf{y}})$ as the input vectors to the PCA computation. The necessity of including velocities is not surprising. Both positions and the velocities (or momenta) are necessary to have a complete phase space state for any mechanical system.

On the plane spanned by the largest two principal components, the cockroach data exhibits an annular path and the angle of the polar representation of points on this path can be used as a phase estimate. As the next example shows, the phase thus obtained is very noisy. While it seemed a good candidate as a protophase source instead of the Hilbert transform, it performed much more poorly in practice, leading to our adoption of the Hilbert transform as the protophase source for Phaser.

Figure 2 displays residual phase estimates for our cockroach locomotion data, using each of the four methods described above. The residual is relative to a linear regression fit to the data. It is apparent that the Phaser, PCA, and AEP methods produce different estimates at time scales of a cycle or less. The PCA estimates have much larger short-term fluctuations than the others while the AEP estimates are linear between events. The Phaser estimates are intermediate. We believe that the Phaser algorithm provides better estimates of phase within cycles than is the case for the other methods, but there is no apparent way to test this with the empirical data. Instead, we use synthetic data constructed with a known reference phase.

B. Cockroach synthetic data

We next tested phase estimation methods for synthetic data from a periodic orbit with small measurement noise. The orbit was generated by a Fourier series model fit to cockroach kinematic data similar to that of Fig. 1, with measurement noise added at a SNR of approximately 20. The foot traces of the signals generated by the Fourier series model are displayed in Fig. 3.

The regularity of the Fourier model illustrates how the sensitivity of the AEP based measures depends upon the ex-




FIG. 3. (Color online) Tracks of simulated cockroach foot motions in the horizontal plane, in a body frame of reference. The origin is placed at the animal's center of mass. Tracks were generated from a seventh order Fourier series fitted to animal data and consisted of 50 cycles at 50 samples per cycle. The measurement noise added was generated from Gaussian random numbers at a SNR of 20, filtered with a second order Butterworth low-pass filter at a cutoff of 0.1 samples.

act details of how an “extreme position” is measured. With a naive method—identifying the AEP as the sample at which a local extremum is reached—performance is quite poor (see AEP sample in Figs. 4 and 5). The use of an interpolator, based on fitting a fifth order polynomial to an 11 sample




FIG. 4. (Color online) Residual phase plots for four estimation methods applied to data from Fig. 3 over a duration of several cycles. The PCA method cannot fully remove nonlinearities and therefore has an oscillatory residual. AEP sample—identifying AEP by the sample at which the anterior extreme is reached—shows systematic errors due to the period not being commensurate with the sample rate. AEP fifth order—identifying AEP by interpolation with a fifth order polynomial—gives results similar in error to those obtained from our Phaser algorithm. All the methods provide excellent estimates with root mean squared errors of less than 0.07 rad, perhaps because the system is so remarkably regular.




FIG. 5. (Color online) Marginal probability density of the phase estimation residual form in the complete sequences from which Fig. 4 was taken. PCA ($\sigma=0.039$) and AEP sample ($\sigma=0.039$) have large phase estimation errors with nearly uniform distribution over the range of observed $\Delta\phi$. AEP fifth order ($\sigma=0.017$) and Phaser ($\sigma=0.015$) provide a Gaussian-like residual with much smaller variance.

window with five samples on each side, and locating the local extremum by finding the roots of its derivative with a root solver (see AEP fifth order in Figs. 4 and 5) gives results similar in quality to Phaser. Such good results are in part an artifact of the system having no phase drift, and thus estimating phase accurately at a point in time gives small residuals for all other times.

C. Hopf oscillator model

We further tested the phase estimation algorithm with simulation data from a dynamical system with an easy to determine phase variable—the Hopf oscillator. Written as an ordinary differential equation (ODE) for a complex variable $z \in \mathbb{C}$, a Hopf oscillator with dynamical noise η is given by

$$\frac{d}{dt}z = [\gamma(r^2 - |z|^2) + i\omega]z + \eta \quad (25)$$

and exhibits a limit cycle at radius r , with frequency ω and convergence rate governed by γ . Rotational symmetry of the equations ensures that phase always coincides with the complex argument (angle of polar representation) of a point. We used $\omega=1$, $r=1$, and $\gamma=0.1$. The dynamical noise term η was modeled with a cubic spline having one random knot point every cycle. The knot points were generated with a Gaussian distribution $\mathcal{N}(0,0.02)$, corresponding to a signal-to-noise ratio (SNR) of 50. By introducing smooth dynamical noise we could integrate the equations with a standard ODE integrator instead of requiring a specialized SDE integrator.

We then mapped the Hopf oscillator trajectory into \mathbb{C}^6 by warping it using six randomly generated smooth mappings of \mathbb{C} into itself. The choice of six derived oscillators yields data

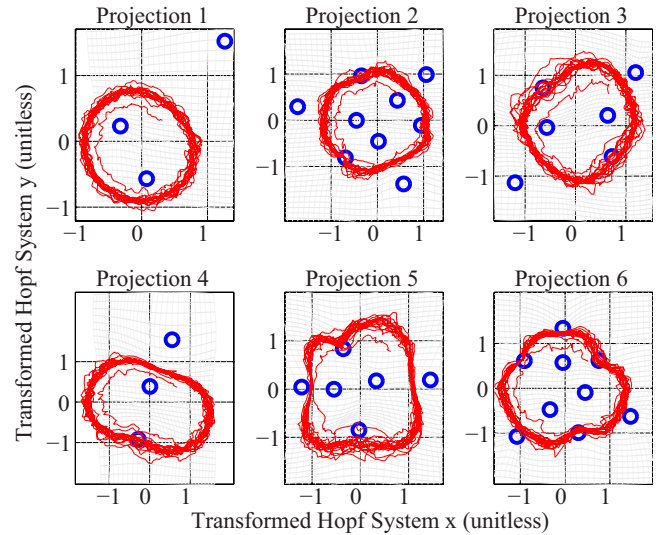


FIG. 6. (Color online) A 100 cycle, 12-dimensional trajectory generated by the Hopf oscillator model with a SNR of 15 and plotted in each of the six copies of \mathbb{C} that comprise our embedding. The background grid shows warping of the Cartesian coordinates in each case. The centers of the radial basis warping functions are marked on each plot. Decay of the initial condition $(0.5,0.5)$ to the orbit shows how different the six observations are from each other.

of comparable dimension to that of the other two examples. Each mapping was generated using radial basis functions with centers for the basis functions placed around the circles at radii 0.5, 1, and 1.15 and mapping parameters such as radius of effect and attraction or repulsion chosen at random (see Fig. 6).

Prior to warping we added synthetic measurement noise to each of the six trajectories separately. We chose to have the measurement noise warped to mimic the state dependent nature of the noise distribution in experimental data. For a given SNR value, the measurement noise was the sum of “pink” noise generated by low-pass filtering white noise with a second order Butterworth filter at a 0.2 sample cutoff, and Gaussian white noise ten times weaker. This particular choice of noise model—pink noise over a baseline of white noise—mimics the power spectral density found in the cockroach experimental data described above.

We applied three phase estimation methods to the data.

- (1) Angle: a naive method, which uses the complex argument of the first projection as the phase.
- (2) 1-dim: phase estimation of [12] (using our $\mathcal{P}[\cdot]$ instead of their series, improving resilience to noise) applied to the first coordinate of the first projection.
- (3) Phaser: our proposed algorithm, which utilizes multivariate data for the estimation procedure.

The AEP based methods are omitted from this lineup of methods, since their performance does not depend on the noise model. Their quality only depends on the noise characteristics of the AEP event itself. In a constructed model, AEP statistical properties can be manipulated to make the AEP method overperform or underperform other methods almost arbitrarily [35]. Instead, our comparisons demonstrate the improvements gained by using increasingly advanced methods for extracting phase information from multiple channels and combining it.




FIG. 7. (Color online) The residual phase relative to the reference phase regression line. The reference phase is the complex argument of the Hopf system which is dynamically perturbed away from constant frequency. Results are shown at measurement noise SNR values of 5, 10, and 15. Plots are zoomed in on illustrative 20 cycle segments of each 100 cycle simulation. The phase excursion event in the bottom (SNR=15) plot demonstrates the value of having improved phase estimation, as the event is all but impossible to resolve with the 1-Dim and Angle methods.

Because measurement noise is of equal magnitude in all coordinates, the simple approach of averaging phases used as an example in [12] leads to results comparable to Phaser. See Fig. 9 and text for a more detailed examination of how averaging compares with Phaser under more general conditions.

To make the measured phases comparable we chose the common transverse section $x=0$ and $y>0$ in the first projection as zero phase. We identified the sample indices of points on this section and subtracted the mean phase of these samples from all phase measures used.

The precision of phase estimation varies with the level of measurement noise (see Figs. 7 and 8, Table I). Estimation of phase based on the entire multivariate dataset was always superior to using partial information, suggesting that our method for fusing the phase information from multiple sources is effective.

As Table I reveals, even at extremely high levels of noise—a signal-to-measurement noise ratio of 5—our algorithm recovers a fairly accurate phase estimate. The estimated accuracy improves at lower levels of measurement noise, and in this particular simulation is about twice as accurate as application of the phase recovery algorithm to only one projection of the six. This 1-Dim method is an improved version of [12] with regards to noise rejection, so one may conclude at least as great an improvement relative to the method of Kraleman *et al.* Examination of Fig. 9 further reveals that both our Phaser method and the 1-Dim method seem to have a gaussian error distribution (the latter only at larger SNR values)—as the curves peak in an inverted parabola. The naive Angle estimate not only suffers from large errors, it is also non-gaussian and biased, suggesting that *ad hoc* techniques may provide very poor statistical quality indeed.




FIG. 8. (Color online) Phase error densities at measurement noise SNR values of 5, 10, and 15, for similar oscillators. Errors are the instantaneous difference of estimated phase and the reference phase. Angle: naive phase estimation based on the complex argument (angle) of the first embedding; 1-Dim: phase estimation based on recovery from a scalar time series, similar to the method of [12]; Phaser: phase estimation based on our proposed multivariate method.

The benefits of our weighted combination of phasors [weights ρ of Eq. (9)] and the projection method [PCA and Eq. (24)] for their combination are illustrated in Fig. 9. The figure shows the rms phase estimation errors obtained when combining phasors in three ways: avg—circular averaging [36] (similar to an example worked out in [12]); wgt—weighted circular averaging using the weights obtained from ρ ; and phr—using the projection method suggested for Phaser with weights ρ .

All three methods were applied to an ensemble of random Hopf oscillator models similar to the model in Fig. 6. The ensemble provides test cases with a variety of nonlinearities while maintaining known noise characteristics. Measurement covariance matrices Q were estimated dynamically under the assumption that measurements have independent errors (Q diagonal) by taking the variance of a high-pass filtered version of each channel as its error variance. The filtering kernel was $[1, -2, 1]$. These dynamically estimated covariances produced much better results than the actual covariances used by the simulation, possibly due to the warping that the noise undergoes. The distribution of the rms phase estimation errors clearly demonstrates the superiority of weighted versus uniform averaging, and of projection versus weighted averaging (note the scale is logarithmic).

The large disparity between these methods shown in Fig. 9 only appears when measurement coordinates have different

TABLE I. Variance of the errors in phase estimation, as a function of measurement SNR and estimation method.

SNR	5	10	15
Angle	0.15	0.13	0.12
1-Dim	0.16	0.06	0.06
Phaser	0.04	0.03	0.03

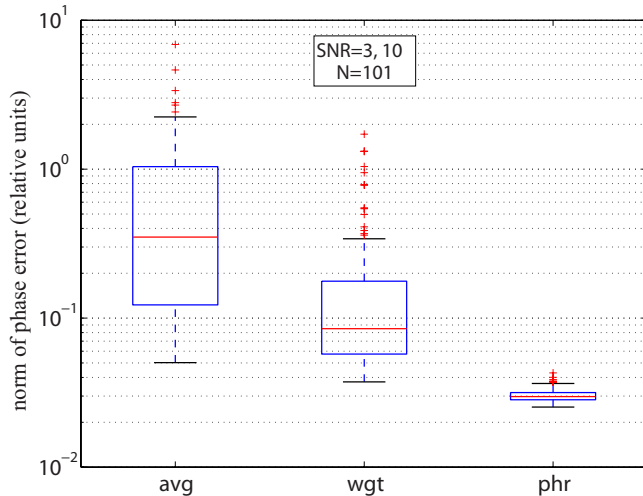


FIG. 9. (Color online) Distribution of rms phase estimation errors in 101 randomly generated Hopf systems with half of the measurement coordinates at SNR 3 and half at SNR 10. The measurement covariance was dynamically estimated for each replicate. Replicates differ in the nonlinear coordinate changes but not in the underlying deterministic dynamics. Phase was estimated by combining $\mathcal{P}[\cdot]$ corrected protophases in three ways: avg—by circular averaging; wgt—by circular averaging with contribution weighted by ρ ; and phr—using the projection method of Phaser.

measurement errors. In this case, half of the projections had measurement error at a SNR of 10 and half at a SNR of 3. In our tests with isotropic measurement errors the performance of all three approaches was indistinguishable.

V. DISCUSSION

We have developed a method for estimating phase from multivariate time series observations of coupled phase-locked oscillators. Our method was motivated by tests of neuromechanical models for cockroach locomotion. The limitations of relatively short data sets with substantial noise, variation among individuals, and changes in speed require phase estimates that reduce both dynamical and measurement noise in the presence of nonlinearities of the motion. Our goal has been to simultaneously reduce the noise and transform the signal to a phase—an angular coordinate that advances with constant speed. We assert that the transformation of data to such phases prior to averaging improves the statistical analysis of system state.

The results of our empirical tests of Phaser are promising. However, we have not been able to provide a complete theoretical analysis of Phaser’s efficacy and correctness. In lieu of such a complete analysis we have provided a description of the mathematical intuitions that lead to various design choices and produced supporting evidence from numerical simulations. We would like to encourage further investigation into these theoretical questions.

The tests we performed compared our method with alternatives based upon principal components analysis, methods that utilize anterior extreme positions and other methods that utilize Hilbert transforms. In stochastic perturbations of a

deterministic system with a stable limit cycle, variance of the residual between estimated phase and the phase of the underlying deterministic cycle provides a quantitative measure of the various methods. In all of the cases that we tested our method performs as well or better than the others. With synthetic data that only added a small Gaussian “measurement” error to a deterministic system, careful use of anterior extreme positions gave comparable results to our method.

From our results the experimental biologist may conclude that the existing practice of using AEP events is highly questionable if the hypotheses being tested involve small differences in phase or quantitative estimation of phase response curves. At the very least, a local polynomial fit similar to our AEP fifth order should be used. If the investigator wishes to be able to detect changes in timing that persist for less than a complete cycle, the use of a phase estimation technique such as Phaser is advisable.

Further empirical tests that delineate the effectiveness of different phase estimation methods in coping with different dynamical phenomena would be helpful. We think that it is important to investigate how methods cope with slowly varying systems. Slow changes in a limit cycle representing the motion of the coupled system are likely to be present in many examples, resulting in phase drift relative to a cycle of fixed frequency. Singular perturbation theory provides a formal approach to the study of slowly changing systems. In the setting of slow-fast systems of the form

$$\dot{x} = f(x, y), \quad (26)$$

$$\dot{y} = \varepsilon g(x, y), \quad (27)$$

ε is the ratio of time scales. In the singular limit $\varepsilon=0$, the system becomes a family of vector fields in x parametrized by y . Stable limit cycles of the singular limit perturb to attracting invariant manifolds for small $\varepsilon>0$. Since the periods of the limit cycles in the singular limit typically depend upon y , defining the phase of the slowly varying system is problematic.

A second direction in which it is important to further test phase estimation methods is with relaxation oscillators, in which abrupt changes occur at particular places on a limit cycle. Examples include recordings from spiking neurons, electrical activity of muscles, and mechanical systems with impacts. All of these time series share the property that phase space velocity is highly nonuniform, typically with one or a few short intervals of high speed as the system traverses the cycle. In our tests Fig. 2—mean AEP linear interpolation provided a reasonable means of extracting phase information from these sorts of variables. However, the topic bears further investigation.

We have examined discrete systems of coupled oscillators. Continuously coupled systems like the bodies of fish or snakes also appear as examples of animal locomotion. We are currently investigating how to include these higher dimensional data types in the phase estimation computation.

Beyond phase, the behavior of a dynamical system near a stable limit cycle can be characterized by Lyapunov exponents and Floquet coordinates that characterize how nearby trajectories approach the limit cycle. Excitation of a trajec-

tory away from the limit cycle will undergo transients as the trajectory returns to the limit cycle. We are comparing our cockroach locomotion data with a model derived from Floquet analysis. We assume that the empirical data comes from a stochastic perturbation of a stable limit cycle and that inherent fluctuations allow the system to sample a neighborhood of the limit cycle in its phase space. Since the quantities of interest come from differences of trajectory segments, accurate phase estimates that align the trajectories relative to the limit cycle are extremely important in this analysis. The AEP methods do not appear to have the short time resolution needed to obtain these estimates.

Source code for Phaser is available for download from the corresponding author [37,38].

ACKNOWLEDGMENTS

The authors would like to thank the colleagues who provided many useful comments and corrections and the anonymous referees for their valuable critique. This research was funded by the National Science Foundation program Frontiers in Integrative Biology Research (FIBR; Grant No. 0425878-Neuromechanical Systems Biology).

-
- [1] R. Abraham and J. E. Marsden, *Foundations of Mechanics: Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields* (Addison-Wesley, Reading, MA, 1978).
- [2] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields* (Springer-Verlag, Berlin, 1983).
- [3] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering* (Perseus Books Group, Cambridge, MA, 2000).
- [4] A. T. Winfree, *The Geometry of Biological Time* (Springer-Verlag, New York, 1980).
- [5] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, U.K., 2004).
- [6] T. Schreiber, *Phys. Rep.* **308**, 1 (1999).
- [7] J. Brocker, U. Parlitz, and M. Ogorzalek, *Proc. IEEE* **90**, 898 (2002).
- [8] E. J. Kostelich and J. A. Yorke, *Physica D* **41**, 183 (1990).
- [9] L. Battiston and M. Rossi, *Nuovo Cimento Soc. Ital. Fis., C* **25**, 305 (2002).
- [10] F. Böttcher, J. Peinke, D. Kleinhans, R. Friedrich, P. G. Lind, and M. Haase, *Phys. Rev. Lett.* **97**, 090603 (2006).
- [11] N. E. Huang *et al.*, *Proc. R. Soc. London, Ser. A* **454**, 903 (1998).
- [12] B. Kralemann, L. Cimponeriu, M. Rosenblum, A. Pikovsky, and R. Mrowka, *Phys. Rev. E* **77**, 066205 (2008).
- [13] M. Rosenblum, L. Cimponeriu, and A. Pikovsky, *Handbook of Time Series Analysis* (Wiley, New York, 2006), pp. 159–180.
- [14] L. Glass and A. T. Winfree, *Am. J. Physiol. Regulatory Integrative Comp. Physiol.* **246**, R251 (1984).
- [15] A. Cohen, P. J. Holmes, and R. H. Rand, *J. Math. Biol.* **13**, 345 (1982).
- [16] *Fictive locomotion* is the generation of rhythmic undulatory motor patterns from a spinal cord that was removed from the animal and is stimulated pharmacologically *in vitro*.
- [17] M. Golubitsky, I. Stewart, P. L. Buono, and J. J. Collins, *Nature (London)* **401**, 693 (1999).
- [18] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, *SIAM Rev.* **48**, 207 (2006).
- [19] S. Revzen, D. E. Koditschek, and R. J. Full, *Progress in Motor Control—A Multidisciplinary Perspective* (Springer, New York, 2008).
- [20] H. Allen, *Proc. Acad. Natural Sci. Philadelphia* **39**, 60 (1887).
- [21] H. Cruse and A. Knauth, *J. Exp. Biol.* **144**, 199 (1989).
- [22] H. Cruse and W. Schwarze, *J. Exp. Biol.* **138**, 455 (1988).
- [23] To be independent of the choice of coordinates the averaging for computing the typical cycle has to be done relative to the unique *invariant measure* of the dynamical system on the cycle. The derivative of the phase is the reciprocal of the distribution of the invariant measure.
- [24] R. J. Full and D. E. Koditschek, *J. Exp. Biol.* **202**, 3325 (1999).
- [25] J. Guckenheimer, *J. Math. Biol.* **1**, 259 (1975).
- [26] M. E. Mera and M. Morna, *Chaos* **16**, 013116 (2006).
- [27] P. C. Mahalanobis, *Proc. Natl. Inst. Sci. India* **2**, 49 (1936).
- [28] P. H. Garthwaite, I. T. Jolliffe, and B. Jones, *Statistical Inference* (Oxford University Press, New York, 2002).
- [29] S. Revzen, J. Bishop-Moser, A. J. Spence, and R. J. Full, *Integr. Comp. Biol.* **46**, e1-162 (2006).
- [30] S. Revzen, D. E. Koditschek, and R. J. Full, *Integr. Comp. Biol.* **47**, e1-152 (2007).
- [31] A. Büschges, *J. Neurophysiol.* **93**, 1127 (2005).
- [32] A. Daffertshofer, C. J. C. Lamoth, O. G. Meijer, and P. J. Beek, *Clin. Biomech. (Bristol, Avon)* **19**, 415 (2004).
- [33] Alternatively, one may take the view that animal locomotion is governed by a *hybrid* system that undergoes a discontinuity, or *hybrid transition* at AEP due to foot touchdown and therefore the sensitivity analysis cannot rely on the existence of derivatives. However, stability considerations [39] suggest that the discontinuity which animals undergo due to foot touchdown must occur after the AEP, when the legs begin retraction. If a hybrid transition does occur, the transition surface is a natural choice as a Poincaré section.
- [34] T. M. Kubow and R. J. Full (personal communication).
- [35] S. Revzen, Tutorial 2 presentation, Workshop 4, Mathematical Biosciences Institute, March 2008 (unpublished).
- [36] N. I. Fisher, *Statistical Analysis of Circular Data* (Cambridge University Press, Cambridge, England, 1993).
- [37] Polypedal laboratory website at University of California at Berkeley <http://polypedal.berkeley.edu>.
- [38] Website for National Science Foundation program for Frontiers of Integrative Biological Research activities at polypedal laboratory, Berkeley <http://fibr.berkeley.edu>.
- [39] A. Seyfarth, H. Geyer, and H. Herr, *J. Exp. Biol.* **206**, 2547 (2003).