Featured in Physics

Locating the Source of Forced Oscillations in Transmission Power Grids

Robin Delabays,[1,*] Andrey Y. Lokhov,[2] Melvyn Tyloo,[2,3] and Marc Vuffray[2]

[1] *University of Applied Sciences of Western Switzerland, Sion, Switzerland*

[2] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico, USA*

[3] *Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico, USA*

(Received 20 February 2023; revised 7 May 2023; accepted 12 May 2023; published 7 June 2023)

A forced oscillation event in power grids refers to a state where malfunctioning or abnormally operating equipment causes persisting periodic disturbances in the system. While power grids are designed to damp most perturbations during standard operations, some of them can excite normal modes of the system and cause significant energy transfers across the system, creating large oscillations thousands of miles away from the source. Localization of the source of such disturbances remains an outstanding challenge due to a limited knowledge of the system parameters outside of the zone of responsibility of system operators. Here, we propose a new method for locating the source of forced oscillations that addresses this challenge by performing a simultaneous dynamic model identification using a principled maximum likelihood approach. We illustrate the validity of the algorithm on a variety of examples where forcing leads to resonance conditions in the system dynamics. Our results establish that an accurate knowledge of system parameters is not required for a successful inference of the source and frequency of a forced oscillation. We anticipate that our method will find a broader application in general dynamical systems that can be well described by their linearized dynamics over short periods of time.

DOI: 10.1103/PRXEnergy.2.023009

I. INTRODUCTION

The power grid is indubitably one if not the greatest engineering achievement of the past century, as recognized by the National Academy of Engineering [1]. It is an intricate and complex network that is the size of a continent with thousands of constituents that require constant supervision. An important aspect of this surveillance is to ensure that voltage frequencies remain within a narrow band (50 ± 0.05 Hz in Europe [2] and 60 ± 0.05 Hz in the United States [3]), as the violation of this requirement can cause significant damage to vital assets and results in blackouts [4]. Voltage frequency fluctuations are primarily caused by real-time imbalances between production and consumption, such as variations in the charging of electric vehicles, or a sudden gust at a wind farm. The power grid is designed to damp these random electromechanical oscillations appearing during standard operating conditions before it creates a resonance with one of the normal modes of the system. However, malfunctioning equipment or abnormal operating conditions can cause periodic disturbances that would persist over time, creating an undesirable transfer of energy across the system, an effect referred to as forced oscillations.

Whereas most forced oscillations are localized to a particular area, some may be close in frequency to one of the dominant normal modes, resulting in a system-wide response and significant energy transfers [5]. Potential impacts of these wide-area oscillations include equipment failure, inadvertent tripping or control actions, and problems with the automatic generation control. This is why fast and reliable location of the source of forced oscillations is crucial in ensuring the safety and reliability of power grids. However, it remains an outstanding challenge, even when forced oscillation events are detected on the network. For instance, on January 11, 2019, a forced oscillation event happened across the entire Eastern Interconnection in the U.S. power grid that was promptly noticed by the reliability coordinators. Nonetheless, existing tools were ineffective at identifying the source location, and a wide-area operator action did not contribute to mitigating the event [6]. The root cause was later fortuitously identified as a faulty input from a steam turbine at a combined-cycle power plant in Florida that forced the system to oscillate for around 18 min before local plant personnel removed the unit from service. The forced oscillations created by the faulty turbine had a peak-to-peak

*robin.delabays@hevs.ch

amplitude of 200 MW at the generating unit, with power swings of about 50 MW observed as far as the New England area [6], which shows how these disturbances have the power to affect an entire continent. In the case of the November 29, 2005 western American oscillation event, a forced oscillation with amplitude 20 MW originating from Alberta in Canada created a resonance effect across the entire Western Interconnection. This led to oscillations of amplitude 200 MW registered on the California-Oregon interface, thousands of miles away from the source [7]. This ability of forced oscillations to cause disturbances at long distances and to be amplified by the grid dynamics seriously complicates the search of their source. Forced oscillations pose a permanent threat to the power grid with more than 20 large-scale events in the past 30 years documented in the United States, with some of them still lacking a well-identified root cause [8].

The increasing deployment of time-synchronous and distributed frequency sensors in the grid, such as phasor measurement units (PMUs) [9], presents an opportunity for developing advanced data-driven and automated detection and localization techniques of forced oscillations. In the majority of cases, *detection* of forced oscillations poses a limited challenge as they can be directly observed from the Fourier peaks in the signal spectrum [10]. However, forced oscillations need to be differentiated from weakly damped normal modes, or free oscillations [11–13]. Weakly damped modes are typically analyzed through Prony analysis [14], and are mitigated via preventative measures by power system operators. On the other hand, the *localization* of forced oscillations constitutes a much tougher challenge. A complete and perfect knowledge of the system and its dynamics would allow one to locate the source of a forcing [15–17]. However, an accurate instantaneous knowledge of the power grid dynamics appears as too strong of an assumption given that the system parameters can fluctuate on the scale of tens of minutes due to local feedback control or temperature variations [18], while the details on the system topology may be unavailable outside of the zone of responsibility for reliability coordinators. Different methods have been proposed to circumvent this lack of information about the system. Some techniques are based on local physical properties such as the monitoring of front arrival times [19], the evaluation of energy flow [20,21], signal decomposition [22], or verification of the linear relation between voltage and current by multiple PMUs [23,24]. Unfortunately, these methods can be very sensitive to modeling errors such as an inaccurate assessment of the fluctuation propagation speed, or can fail to localize perturbations that are amplified by normal modes. Black-box machine learning methods [25–27] have been developed with the aim of being fully model agnostic, but suffer from a prohibitive requirement in training examples of forced oscillation events. These various shortcomings motivated recent calls from system operators and regulators to develop robust tools for performing forced oscillation analysis and localization [6,28], which led to a further exploration [29,30]. Nevertheless, a correct localization of the source in the case where dominant system modes are excited due to the resonance phenomenon remains an outstanding challenge.

In this paper, we propose a new principled method of detection and localization of forced oscillations that is agnostic to knowledge of the system topology and parameters, fully capable of identifying the source of distant normal mode excitation, and which does not rely on any offline training. Our method operates within a much broader framework that extends rather universally to any dynamical system that can be well described over a short period of time by its linearized dynamics. This makes it a method of choice for the detection and localization of energy transfers created by small disturbances on a large class of complex networks. The key insight in our method consists in leveraging random frequency fluctuations naturally present in the system to build in real time an effective dynamic model of the network. The dynamic model identification and source localization are performed at the same time using a principled maximum likelihood approach. We illustrate the performance of our approach on a number of examples where forcing excites the natural system modes and creates a resonance phenomenon, as well as on a real-world PMU data set.

II. CHALLENGES AND SOLUTIONS FOR LOCATING THE OSCILLATION SOURCE

Learning of the dynamics of linear stochastic systems lies at the foundation of our approach. In the ambient regime of fluctuations, a complex system such as a power grid with a general nonlinear dynamics is typically found in a state that is close to a stable equilibrium. In this ambient regime of small perturbations, the dynamics of the system is well described by a linear stochastic equation corresponding to a celebrated model of coupled harmonic oscillators. Without surprise, this family of models includes the popular swing equations describing the ambient dynamics of generators in a transmission power grid, which are derived under the assumption of small deviations from the steady state [31,32] (see Sec. S1 within the Supplemental Material [33] for more details).

We consider the linear stochastic network dynamics with an additional forcing at a single node l, indexed by $\mathbf{e}_l$, the canonical basis vector with nonzero lth component. Mathematically, this dynamics is described by the linear stochastic differential equation

$$\mathbf{M}d\mathbf{p}_t = \mathbf{D}\mathbf{p}_t dt + \mathbf{L}\mathbf{x}_t dt + \gamma\,\mathbf{e}_l \cos(2\pi(ft+\phi))dt + d\mathbf{W}_t, \tag{1}$$

where $\mathbf{x}_t$ represents the network state variables (for a power grid, they correspond to deviations of phases from the

steady-state values), $\mathbf{p}_t dt = d\mathbf{x}_t$ is the generalized momentum (deviation of frequencies from the steady-state values in a power grid), $\mathbf{M}$, $\mathbf{D}$, $\mathbf{L}$ are generalized mass, damping, and network coupling parameters, γ, f, and ϕ are the amplitude, frequency and phase of forcing, respectively, and $\mathbf{W}_t$ is a Wiener noise process describing random fluctuations.

Our goal is to reconstruct the oscillation frequency and the location of the forcing source from the measured time series of $\mathbf{x}_t$ and $\mathbf{p}_t$ of length T at a sequence of N discrete time steps $t \in \{t_1, \ldots, t_N\}$. We do not assume any knowledge of the system parameters or topology, which represents a realistic scenario in power grids: instantaneous awareness of system parameters is almost never available to system operators [18]. Hence, we do not suppose any knowledge on the inertia $\mathbf{M}$, damping $\mathbf{D}$, or grid Laplacian matrix $\mathbf{L}$. Obviously, we also do not assume any knowledge of the parameters related to the forcing, i.e., γ, l, f, or ϕ. All these parameters of interest need to be recovered from the noisy data $\{\mathbf{x}_t\}$ and $\{\mathbf{p}_t\}$.

In Fig. 1, we illustrate with a toy example of a three-node network the key mechanism that renders the localization of forced oscillations elusive to standard signal processing analysis. Namely, a forcing with a frequency close to a natural frequency of the system may excite the natural modes of the system peaked at a different node. As a result, neither the correct forcing frequency nor the source of the oscillations are evident from the Fourier spectrum in Fig. 1(d). This effect is reminiscent of forced oscillation events in power grids such as those of November 29, 2005 and January 11, 2019 discussed above, where large perturbations can be seen far from the source and at a very different frequency.

To address this challenge, we develop a principled method for determining the oscillation frequency f and locating the source of forced oscillations based on a maximum likelihood approach (see Sec. V below and Sec. S3 within the Supplemental Material [33] for a detailed derivation). A naïve real-space estimator leads to a complicated nonconvex optimization problem in frequency, as we explain in Sec. S4 within the Supplemental Material [33]. A key insight that leads to an efficient solution consists in realizing that the finite length of the time series imposes a finite resolution on the frequencies. This leads to a tractable formulation featuring the Fourier-transformed quantities, which can be efficiently solved with state-of-the-art interior-point method solvers (see Sec. V below).

Knowledge of the number of sources leads to a discrete formulation of the problem, where optimization is run for every node (or every pair of nodes if two sources of forcing are present, etc.). We refer to this method as to the *System Agnostic Localization of Oscillations* (SALO) algorithm. The SALO approach is fully parallelizable over all nodes in the network, making it the method of choice for a high-performance computing system. An example of an application of the SALO framework is given in Fig. 1(e) for our toy resonance example: both the source and the frequency of forced oscillations are unambiguously and correctly identified. For streaming applications where a faster identification is desired, we consider a computationally advantageous relaxation of the problem. In this SALO-relaxed version of the algorithm, all nodes are formally allowed to be a source with a respective amplitude γ_i for each node i (see Sec. V below for more details). In what follows, we benchmark the method on a number of simulated and real use cases.

III. TESTS ON SYNTHETIC SYSTEMS AND REAL DATA

An effect where interaction between the forcing frequency and one of the natural modes of the system leads to a peak response away from the forcing source may arise in networks with a much more complex structure compared to the three-node system illustrated in Fig. 1. Such a resonance behavior represents an outstanding challenge for detection algorithms whereby the peaks in the Fourier spectrum may not only involve nodes far away from the source, but also point to frequencies related to the natural modes of the system rather than to the frequency of the forcing. We showcase this phenomenon on a synthetically generated data set according to model (1) on a network inspired by the UK high-voltage grid; see Fig. 2. This test case demonstrates some of the features observed in the 2005 Western Interconnection and in the 2019 Eastern Interconnection oscillation events. In particular, the largest amplitudes in the Fourier components can be located very far away from the source, at distances comparable to the diameter of the network. On the other hand, we see that the maximum-likelihood-based SALO method identifies the correct forcing frequency and the correct source, for which the Fourier signals are otherwise completely hidden among the responses of other nodes in the network; see Figs. 2(b)–2(d).

In Sec. S5 within the Supplemental Material [33], we further illustrate the challenges of Fourier-based source localization under the resonance phenomenon on a standard IEEE test case topology with 57 nodes. Yet in this case again, as shown in the Fig. S3 within the Supplemental Material, the SALO method precisely and unequivocally identifies the correct frequency and location of the forcing source, without exploiting any prior knowledge on the system topology and parameters. We also use this test case to demonstrate the performance of the SALO-relaxed version of the algorithm, which accurately points to the correct location of the source of forced oscillations, while benefiting from the computational complexity of a single source verification under the full maximum likelihood approach. This shows that the SALO-relaxed version

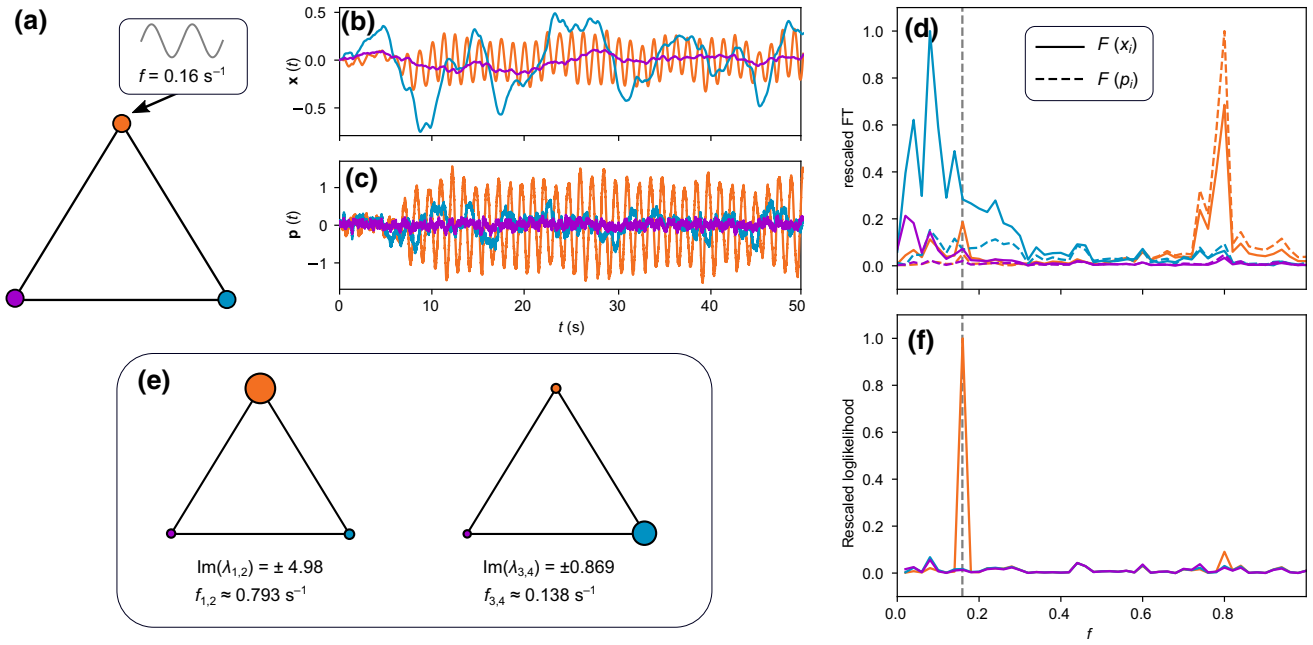

FIG. 1. A toy example with three nodes illustrating the challenges and solutions for locating the source of forced oscillations in a network of coupled oscillators. (a) A forced oscillation is induced on the orange node with frequency $f = 0.16\ s^{-1}$. Time series of the network states corresponding to generalized positions $\mathbf{x}_t$ and momentum $\mathbf{p}_t$ (corresponding to phase deviations $\boldsymbol{\theta}_t$ and frequency deviations $\boldsymbol{\omega}_t$ in power grids; see Sec. S1 within the Supplemental Material [33]) presented in panels (b) and (c), respectively, are generated using Eq. (1) and the parameters detailed in Sec. S2 within the Supplemental Material [33]. (d) A naïve Fourier analysis does not allow one to identify either the forcing source or location: the Fourier transform (FT) of the position time series displays its largest peak at the blue node with frequency $f = 0.08\ s^{-1}$, and the Fourier transform of the momentum time series displays peaks at the orange node, with frequency $f = 0.8\ s^{-1}$. (e) An eigendecomposition of the dynamic state matrix shows the natural frequencies (eigenvalues) and modes (eigenvectors, with the sizes of the nodes indicating the magnitudes of the respective components) of the system. This analysis reveals the reason for the observed behavior of the Fourier transforms: the forcing frequency is close to a natural frequency of the system, thus creating a resonance effect and exciting the natural modes of the system. (f) Our source localization algorithm confidently points to the correct source (orange node) and frequency (indicated with a dashed gray line) of the forced oscillation. The accuracy of the determination of the forcing frequency f is fundamentally defined by the duration of the available time series.

can be used for a quicker assessment of the forced oscillations once they have been detected in the system, prior to running parallelized computations under the SALO framework. We further show this computational advantage on a series of synthetic instances of increasing size in Fig. S4 within the Supplemental Material, whereby the ratio of run times of the SALO and SALO-relaxed versions scales linearly with the system size.

In the tests described above, we assumed that the time series have been produced using model (1), albeit the system parameters are not known to the reconstruction algorithm. In particular, it is assumed that there is a *single* and *observed* source node l, and that the forcing of type $\cos(2\pi(ft + \phi))$ is associated with a single frequency f. In Fig. 3 we look at the results produced by the SALO algorithm in situations where these assumptions are violated. In Fig. 3(a), the source node is outside of the observable system. We see that in this case, the SALO algorithm is still able to correctly identify the forcing frequency, and points to the immediate neighbors of the source node inside the observable system. Fig. 3(b) demonstrates the case where two sources of oscillations at different nodes and frequencies are simultaneously present in the system. Notably, two peaks corresponding to both forcings appear in the rescaled likelihood score. Finally, in Fig. 3(c), we show the rescaled likelihood scores for the input forcing signals of different types. Even in this case, the SALO algorithm correctly identifies the source location, showing up as several peaks in the likelihood at different harmonics of the input forcing signal. This makes the algorithm remarkably robust to the assumptions behind the model. In Sec. S6 within the Supplemental Material [33], we show an example of an application of the SALO algorithm to real data that display a combination of features observed in Fig. 3.

Another possible misspecification is the assumption of linearity of the dynamics. For instance, real power systems do not exactly follow the linear model of type (1)

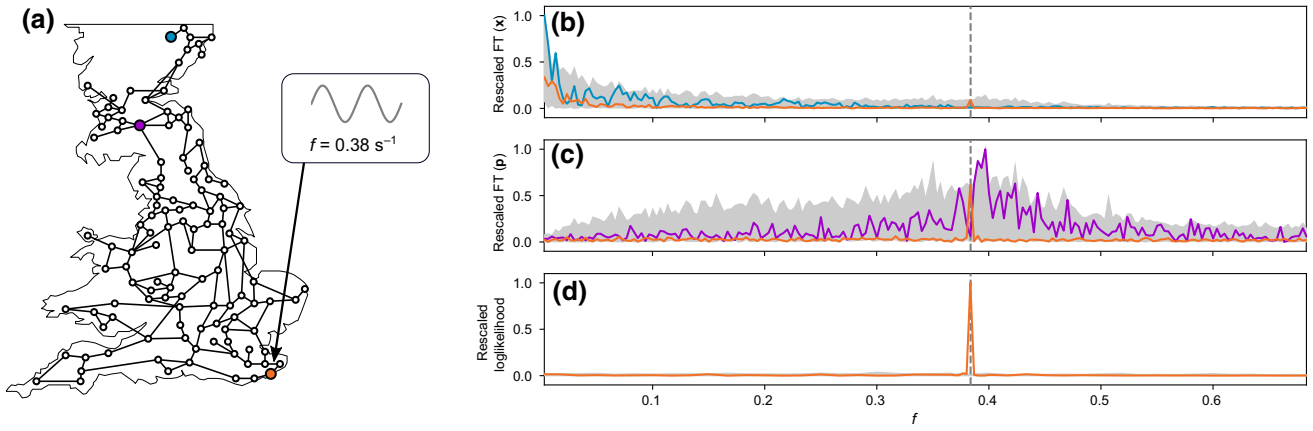

FIG. 2. Detection and localization of forced oscillations under the resonance conditions. (a) Synthetic test case with a topology inspired by the UK high-voltage grid, designed to reproduce features observed in real oscillatory events where oscillations interacted with the system modes. Details on the network parameters are given in Sec. S2 within the Supplemental Material [33]. The forcing at the orange node results in a highest response in Fourier spectra at the opposite side of the network, as shown for the Fourier components of the (b) generalized state and of the (c) generalized momentum. (d) SALO algorithm confidently identifies the correct forcing frequency and source without any knowledge of the system topology or parameters. The envelope of scores for nonhighlighted nodes is shown in gray.

with constant system parameters. Instead, a linear swing model that falls within class (1), as discussed in Sec. S1 within the Supplemental Material [33], represents an approximation to a general nonlinear dynamics of generators under small deviations from the steady state, valid over finite periods of time [18,34]. However, with the understanding that Eq. (1) only serves as an *effective* model providing an adequate description of the complex system

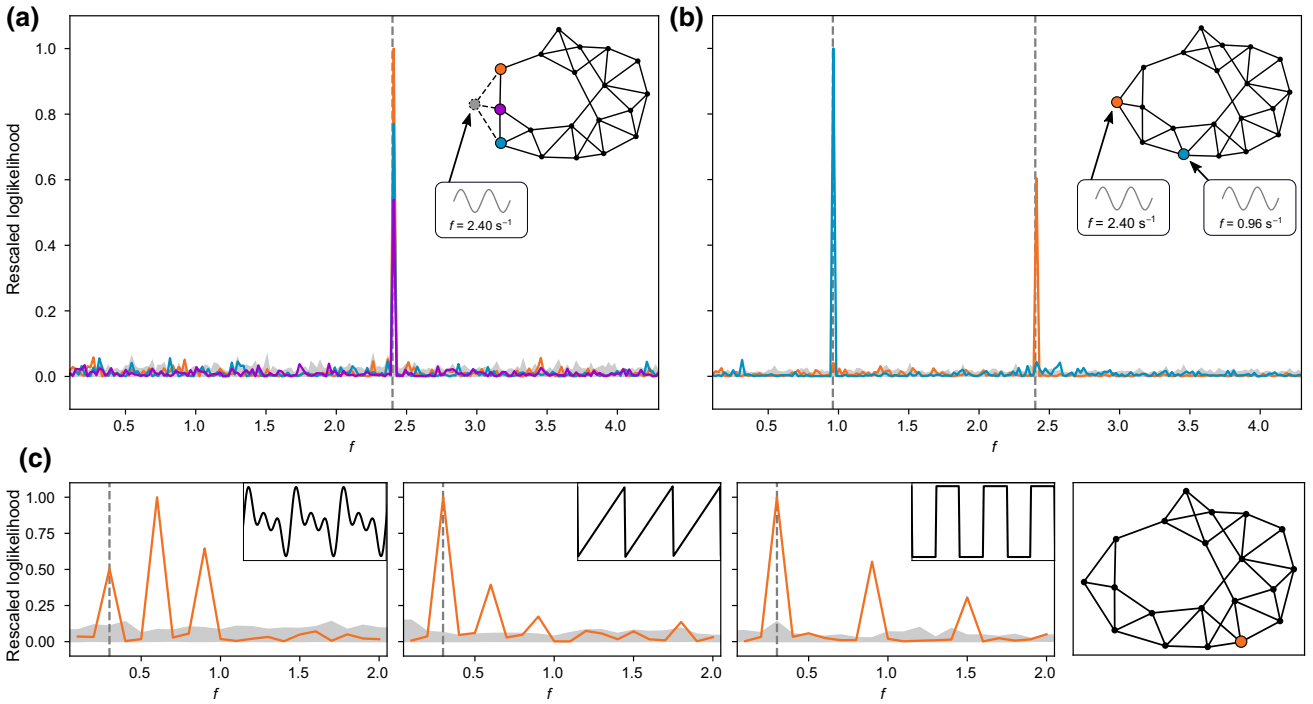

FIG. 3. Robustness of the SALO algorithm to model misspecification. (a) In the case where the source of forced oscillations is outside of the observed system (gray node), the algorithm still correctly identifies the forcing frequency and the neighbors of the hidden source inside the visible system (three overlapping peaks in likelihood at the forcing frequency correspond to orange, purple, and blue nodes). (b) In the case of several sources, both locations appear as peaks in the rescaled likelihood score at their respective forcing frequencies. (c) In the case of nonsinusoidal forcing injected into the system, the forcing location is still correctly identified, while the complex nature of the forcing shows up as likelihood peaks at different harmonics of the forcing signal. The envelope of scores for nonhighlighted nodes is shown in gray in all panels.

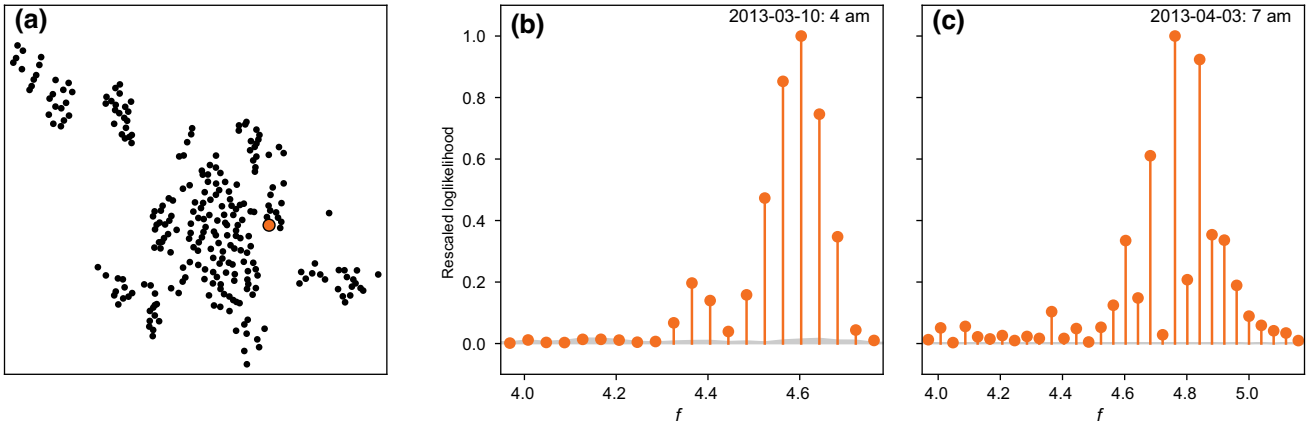

FIG. 4. Identifying the location of forced oscillations for PMU data from a U.S. power system operator. (a) Geographical layout of PMU sensors in the system providing time series data (anonymized and modified coordinates). The node identified as the most likely source of forced oscillations is highlighted in orange. (b) According to the SALO algorithm, the identified node has a much higher likelihood for being a source for a range of frequencies in the vicinity of $f = 4.6\ \mathrm{s}^{-1}$ compared to the rest of the nodes in the network (corresponding likelihoods depicted in gray). A finite range of candidate frequencies is a realistic feature that may emerge in real systems due to the finite length of the collected data. (c) The algorithm consistently points to the same candidate source in a similar frequency range even for time series collected over two time periods separated by almost a month (March 2013 versus April 2013). This consistency strongly indicates that sustained oscillations may originate from a single faulty component of the system. The envelope of scores for other nonhighlighted nodes in the system is shown in gray.

dynamics over short time scales, we show that the SALO method is still applicable to data from real transmission power grids with the purpose of identifying the conditions pertaining to forced oscillations. Here, we benchmark the SALO algorithm on an instance of real PMU data from a U.S. transmission power grid with 200 nodes, where a presence of sustained oscillations has been suggested in previous studies. A Fourier-type analysis of the PMU data from a U.S. independent system operator showed a presence of sustained oscillations with a frequency in the 4–6 $[\mathrm{s}^{-1}]$ range, responsible for the emergence of correlations between several node clusters [35]. The analysis of these time series with the SALO method, see Fig. 4, confidently points to a single source of oscillations, with a frequency close to the range previously identified in Ref. [35]. Incidentally, although the ground truth for this system is unknown, the identified node is consistently pointed to as the most likely source of sustained oscillations even for PMU data separated by a time interval of about a month; see Figs. 4(b) and 4(c). An extended range of candidate frequencies is likely to be connected to the fundamental limits on the data resolution, as exemplified in Sec. S4 within the Supplemental Material [33], where shorter time series lead to a wider log-likelihood objective function in the frequency domain (see Fig. S1 within the Supplemental Material).

IV. CONCLUSIONS

We proposed a rigorous maximum-likelihood-based framework enabling simultaneous system identification and localization of the source of forced oscillations, without any prior knowledge of the system topology or parameters. In particular, our method is able to perfectly locate the oscillation source, even in the case where the forcing excites one of the natural modes of the system and creates an amplitude peak at a far-away node and at a different frequency. This scenario is reminiscent of some of the real-world historical events, such as the 2005 Western Interconnection event, which turned out to be the most challenging from the oscillation source localization perspective. The ease of parallelization and a relaxed version of the algorithm makes the method scalable to large network instances and multiple sources, while robustness to modeling assumptions makes the algorithm applicable to data produced by real-world dynamical systems.

The SALO algorithm can be naturally adapted to the situation where additional prior information is available. For instance, matrices $\mathbf{M}$, $\mathbf{D}$, and $\mathbf{L}$ do not need to be reconstructed from data if a prior knowledge on the network structure and parameters is available. In Sec. S7 within the Supplemental Material [33], we show that, under this scenario, the SALO algorithm is able to identify the correct source and frequency of the forcing using significantly less data compared to the most challenging setting considered in this work, where no prior information on the system is available. Similarly, knowledge of the disturbance type or number of sources can also be directly incorporated into the algorithm. In the context of power grid applications, it could be practical to extend the SALO method to the case of partial sensor coverage, as well as to a source localization under the model that specifically takes into account different types of generators, buses, and PMUs in the grid. For instance, the primary setting considered in our work

assumes that all buses have nonzero inertia, which is true for generators in the power grid, but not necessarily for other buses, e.g., those representing loads.

Because of a wider applicability of our approach to general stochastic linear dynamics and coupled harmonic oscillators, we anticipate that our methods will find a broader range of applications beyond power grids, e.g., vehicular platoon subjected to malfunctioning elements or malicious attacks, and more broadly multiagent systems where automated units must reach an overall consensus [36], as well as forced oscillations in wave propagation dynamics.

V. METHODS

A. Model

The dynamic equation (1) can be reformulated as a first-order linear dynamic system

$$d\mathbf{p}_t = \mathbf{A}\mathbf{X}_t dt + \gamma \mathbf{e}_l \,\mathrm{Re}(e^{2\pi i(ft+\phi)})dt + d\mathbf{W}_t, \tag{2}$$

where we have used $\mathbf{p}_t dt = d\mathbf{x}_t$ and defined $\mathbf{X}_t = (\mathbf{x}_t, \mathbf{p}_t)$. Discretizing Eq. (2) with an Euler-Maruyama approximation scheme over N points yields the finite difference equation for $j = 0, \ldots, N-1$,

$$\boldsymbol{\Delta}_{t_j} = \mathbf{A}\mathbf{X}_{t_j} + \gamma \mathbf{e}_l \,\mathrm{Re}(e^{2\pi i(kj/N+\phi)}) + \boldsymbol{\xi}_j, \tag{3}$$

where $\boldsymbol{\Delta}_{t_j} = (\mathbf{X}_{t_{j+1}} - \mathbf{X}_{t_j})/\tau$ for $\tau = T/N$, we assumed that $t_j = j\tau$, and the frequency relates to the integer $0 < k < N/2$ as $k = fT$. As discussed in the problem formulation above, the measurements $\{\mathbf{X}_{j\tau}\}_{j=0,\ldots,N-1}$ are assumed to be available, but the parameters of system $\mathbf{A}$, γ, l, f, and ϕ, are unknown and need to be estimated from the data.

B. Reconstruction algorithm

Given that, on finite time intervals, the contribution of the Wiener process is an independent and identically distributed Gaussian variable with zero mean, i.e., $\boldsymbol{\xi}_{t_j} \sim \mathcal{N}(0, \tau)$, the negative loglikelihood reads

$$L(\mathbf{A}, \gamma, l, k, \phi | \{\mathbf{X}_{t_j}\}_{j=1}^N) = \frac{1}{N}\sum_{j=0}^{N-1} \|\boldsymbol{\Delta}_{t_j} - \mathbf{A}\mathbf{X}_{t_j} - \gamma \mathbf{e}_l \mathrm{Re}(e^{2\pi i(kj/N+\phi)})\|^2. \tag{4}$$

Note that a discretized set of frequencies appearing in the forcing term as a result of finiteness in T and N is crucial, because an optimization over a continuous variable f leads to a hard nonlinear optimization problem, as exemplified in Sec. S4 within the Supplemental Material [33]. For a fixed frequency k and node l, the joint minimization over $\mathbf{A}$, γ, and ϕ remains challenging to solve as the negative loglikelihood possesses numerous local minima. Nevertheless, the global minimization of Eq. (4) over phase ϕ, as a function of the remaining parameters, can be performed exactly, leading to the following expression for the partial negative loglikelihood:

$$\begin{aligned} &L_{\mathrm{SALO}}(\mathbf{A}, \gamma, l, k | \{\mathbf{X}_{t_j}\}_{j=0}^{N-1}) \\ &\quad = \mathrm{Tr}(\mathbf{A}^\top \mathbf{A}\Sigma_0) - 2\mathrm{Tr}(\mathbf{A}\Sigma_1) + \frac{1}{2}\gamma^2 \\ &\quad\quad - \frac{2\gamma}{\sqrt{N}}\sqrt{\mathrm{Tr}(\mathbf{A}_{l,\cdot}^\top \mathbf{A}_{l,\cdot} F(k)) - 2f_l(k)\mathbf{A}_{l,\cdot} + g_l(k)}; \end{aligned} \tag{5}$$

see Sec. S3 within the Supplemental Material [33] for details of the derivation and the definitions of all the quantities that can be directly computed from the available time series $\{\mathbf{X}_{j\tau}\}_{j=1,\ldots,N}$. For a fixed l and k, expression (5) remains a nonconvex function in the arguments $\mathbf{A}$ and γ. However, we have observed that this cost function, unlike its naive counterpart (4), seems to always possess a single minimum that can be efficiently found using state-of-the-art interior point methods. In this work, we use the optimization software Ipopt [37] within JULIA's JuMP modeling framework for mathematical optimization [38]. Note that the optimization is fully parallelizable over k and l. We refer to the minimization of the partial negative loglikelihood as the SALO method.

We also consider an accelerated method referred to as the SALO-relaxed method, which exploits the spatial relaxation of the problem in variable γ. Under this variant of the algorithm, all nodes are formally allowed to be a source, and their relative likelihood to be a source is encoded by a respective amplitude γ_i for each node i (so that the regression now runs over the vector $\boldsymbol{\gamma}$):

$$L_{\mathrm{SALOr}}(\mathbf{A}, \boldsymbol{\gamma}, k | \{\mathbf{X}_{t_j}\}_{j=1}^N) = \mathrm{Tr}(\mathbf{A}^\top \mathbf{A}\Sigma_0) - 2\mathrm{Tr}(\mathbf{A}\Sigma_1) + \sum_{l=1}^n \left[\frac{\gamma_i^2}{2} - \frac{2\gamma_i}{\sqrt{N}}\sqrt{\mathrm{Tr}(\mathbf{A}_{l,\cdot}^\top \mathbf{A}_{l,\cdot} F(k)) - f_l(k)\mathbf{A}_{l,\cdot} + g_l(k)}\right]. \tag{6}$$

We benchmark this method in Sec. S5 within the Supplemental Material [33], where we also discuss the computational speed-up of the SALO-relaxed estimator compared to the SALO estimator.

All data that support the plots within this paper and other findings of this study are available from the authors on reasonable request. The code implementing the SALO and SALO-relaxed algorithms in JULIA is available from GitHub [39].

ACKNOWLEDGMENTS

This work is supported by the U.S. DOE/OE as part of the DOE Advanced Sensor and Data Analytics Program and by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Projects No. 20220797PRD2, No. 20210078DR, and No. 20220774ER. R.D. is supported by the Swiss National Science Foundation, under Grant No. P400P2_194359. The authors are grateful to Professor Daniel Bienstock from Columbia University, and Dr. Yilu Liu and Dr. Wenpeng "Wayne" Yu from the Power IT Lab at Oak Ridge National Lab and UT Knoxville for sharing real PMU data used in this work. The authors also thank Dr. Slava Maslennikov from ISO New England and Professor Michael Chertkov from the University of Arizona for fruitful discussions.

All authors designed and performed the research, wrote the manuscript, reviewed and edited the paper. Correspondence and requests for materials should be addressed to M.V. The authors declare no competing interests.

[1] G. Constable and B. Somerville, *A Century of Innovation: Twenty Engineering Achievements That Transformed Our Lives* (Joseph Henry Press, Washington, DC, USA, 2003).

[2] Standard EN 50160—Voltage Characteristics of Public Distribution Systems, https://www.se.com/ww/library/SCHNEIDER˙ELECTRIC/SE˙LOCAL/APS/204836˙1312/DraftStandard0026rev2-DraftEN501602005-05.pdf (accessed 2022-10-25).

[3] B. J. Kirby, J. Dyer, C. Martinez, R. A. Shoureshi, R. Guttromson, and J. Dagle, *Frequency Control Concerns in the North American Electric Power System* (Department of Energy, United States, 2003).

[4] O. Alizadeh Mousavi, R. Cherkaoui, and M. Bozorg, Blackouts risk evaluation by Monte Carlo simulation regarding cascading outages and system frequency deviation, Electr. Power Syst. Res. **89**, 157 (2012).

[5] S. A. N. Sarmadi and V. Venkatasubramanian, Inter-area resonance in power systems from forced oscillations, IEEE Trans. Power Syst. **31**, 378 (2015).

[6] NERC report: Eastern interconnection oscillation disturbance January 11, 2019 forced oscillation event, https://www.nerc.com/pa/rrm/ea/Documents/January˙11˙Oscillation˙Event˙Report.pdf (accessed 2022-04-26).

[7] S. A. N. Sarmadi, V. Venkatasubramanian, and A. Salazar, Analysis of November 29, 2005 western American oscillation event, IEEE Trans. Power Syst. **31**, 5210 (2016).

[8] M. Ghorbaniparvar, Survey of forced oscillations in power systems, J. Mod. Syst. Clean Energy **5**, 671 (2017).

[9] P. W. Sauer, M. A. Pai, and J. H. Chow, *Power System Dynamics and Stability: With Synchrophasor Measurement and Power System Toolbox* (John Wiley & Sons, Hoboken, NJ, USA, 2017).

[10] N. Zhou and J. Dagle, Initial results in using a self-coherence method for detecting sustained oscillations, IEEE Trans. Power Syst. **30**, 522 (2015).

[11] X. Wang and K. Turitsyn, Data-driven diagnostics of mechanism and source of sustained oscillations, IEEE Trans. Power Syst. **31**, 4036 (2016).

[12] R. Xie and D. J. Trudnowski, in *Proc. of the IEEE PESGM* (IEEE, Chicago, IL, USA, 2017).

[13] H. Ye, Y. Lie, P. Zhang, and Z. Du, Analysis and detection of forced oscillation in power system, IEEE Trans. Power Syst. **32**, 1149 (2017).

[14] J. F. Hauer, C. J. Demeure, and L. L. Scharf, Initial results in prony analysis of power system response signals, IEEE Trans. Power Syst. **5**, 80 (1990).

[15] T. R. Nudell and A. Chakrabortty, in *2013 Amercian Control Conference (ACC)* (IEEE, Washington, DC, USA, 2013), p. 3467.

[16] I. R. Cabrera, B. Wang, and K. Sun, in *Proc. of the IEEE PESGM* (IEEE, Chicago, IL, USA, 2017).

[17] R. Delabays, L. Pagnier, and M. Tyloo, Locating line and node disturbances in networks of diffusively coupled dynamical agents, New J. Phys. **23**, 043037 (2021).

[18] A. Y. Lokhov, M. Vuffray, D. Shemetov, D. Deka, and M. Chertkov, in *2018 Power Systems Computation Conference (PSCC)* (IEEE, Dublin, Ireland, 2018), p. 1.

[19] A. Semerow, S. Horn, B. Schwarz, and M. Luther, in *IEEE International Conference on Power System Technology (POWERCON)* (IEEE, Wollongong, NSW, Australia, 2016).

[20] L. Chen, Y. Min, and W. Hu, An energy-based method for location of power system oscillation source, IEEE Trans. Power Syst. **28**, 828 (2013).

[21] S. Maslennikov, B. Wang, and E. Litvinov, Dissipating energy flow method for locating the source of sustained oscillations, Int. J. Electr. Power Energy Syst. **88**, 55 (2017).

[22] T. Huang, N. M. Freris, P. R. Kumar, and L. Xie, in *Proc. of the IEEE CISS* (IEEE, Shanghai, China, 2018).

[23] S. Chevalier, P. Vorobev, and K. Turitsyn, Using effective generator impedance for forced oscillation source location, IEEE Trans. Power Syst. **33**, 6264 (2018).

[24] S. Chevalier, P. Vorobev, and K. Turitsyn, A Bayesian approach to forced oscillation source location given uncertain generator parameters, IEEE Trans. Power Syst. **34**, 1641 (2019).

[25] Z. Chen and J.-C. Maun, Artifical neural network approach to single-ended fault locator for transmission lines, IEEE Trans. Power Syst. **15**, 370 (2000).

[26] G. Cardoso, J. G. Rolim, and H. H. Zürn, Application of neural-network modules to electric power system fault section estimation, IEEE Trans. Power Delivery **19**, 1034 (2004).

[27] H.-W. Lee, J. Zhang, and E. Modiano, in *IEEE International Conference on Communications, Control, and Computing Techologies for Smart Grids (SmartGridComm)* (IEEE, Aalborg, Denmark, 2018).

[28] 2021 IEEE-NASPI oscillation source location contest, http://web.eecs.utk.edu/kaisun/Oscillation/2021Contest/ (accessed 2022-04-26).

[29] P. G. Estevez, P. Marchi, C. Galarza, and M. Elizondo, Non-stationary power system forced oscillation analysis using synchrosqueezing transform, IEEE Trans. Power Syst. **36**, 1583 (2021).

[30] P. G. Estevez, P. Marchi, F. Messina, and C. Galarza, Forced oscillation identification and filtering from multi-channel time-frequency representation, Preprint ArXiv:2108.08736 (2021).

[31] A. R. Bergen and V. Vittal, *Power Systems Analysis* (Prentice Hall, Upper Saddle River, NJ, USA, 2000).

[32] P. Kundur, N. J. Balu, and M. G. Lauby, *Power System Stability and Control* (McGraw-Hill, New York, 1994), Vol. 7.

[33] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PRXEnergy.2.023009 for the dynamical model used to approximate the system; the parameters of the test cases; the mathematical derivations of the SALO and SALO-relaxed algorithms; details about the shape of the likelihood function; performance analysis of the SALO algorithm; further results on real PMU data; improvement of the algorithm, using prior information on the system.

[34] C. Hannon, D. Deka, D. Jin, M. Vuffray, and A. Y. Lokhov, in *2021 IEEE Madrid PowerTech* (IEEE, Madrid, Spain, 2021), p. 1.

[35] M. Escobar, D. Bienstock, and M. Chertkov, in *Proc. of the IEEE PowerTech* (IEEE, Milano, Italy, 2019).

[36] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, in *42nd IEEE International Conference on Decision and Control*, Vol. 2 (IEEE, Maui, HI, USA, 2003), p. 2010.

[37] A. Wächter and L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Math. Progr. **106**, 25 (2006).

[38] I. Dunning, J. Huchette, and M. Lubin, Jump: A modeling language for mathematical optimization, SIAM Rev. **59**, 295 (2017).

[39] https://github.com/lanl-ansi/SALO.