Persistent-Homology-based Detection of Power System Low-frequency Oscillations using PMUs

Yang Chen, *Member, IEEE*, Harish Chintakunta, *Member, IEEE*, Le Xie, *Senior Member, IEEE*, Yuliy M. Baryshnikov, and P. R. Kumar, *Fellow, IEEE*

***Abstract*—This paper presents a new methodology to detect low-frequency oscillations in power grids by use of time-synchronized data from phasor measurement units (PMUs). Principal component analysis (PCA) is first applied to the massive PMU data to extract the low-dimensional features, i.e., the principal components (PCs). Then, based on persistent homology, a *cyclicity response function* is proposed to detect low-frequency oscillations through the use of PCs. Whenever the cyclicity response exceeds a numerically robust threshold, a low-frequency oscillation can be detected instantly. Such swift detection can then be followed by modal analysis tools for more detailed information about the oscillation. Numerical examples using real data illustrate the effectiveness of the proposed methodology for quick detection of oscillations during operations.**

***Index Terms*—Persistent homology, phasor measurement unit, principal component analysis, low-frequency oscillation, detection.**

I. Introduction

ONE of the prime examples of improving grid monitoring via synchrophasors is the monitoring and control of low-frequency oscillations. Since the early 1960s, low-frequency oscillations have led to many system-wide failures, such as the 1996 Western Electricity Coordinating Council (WECC) Blackout induced by a 0.25 Hz oscillation [1]. Traditional approaches to detecting low-frequency oscillations require a detailed dynamical representation of the system to conduct modal analysis [2]–[5]. However, the increasing penetration of distributed and variable resources poses difficulties in obtaining an accurate dynamic model of the system.

Alternatively, one could obtain detailed information about such oscillations through monitoring and analysis of online measurement data such as from synchrophasors. Synchrophasors have shown great potential for improving wide-area monitoring, protection and control (WAMPAC) [6]–[8] since the 1980s: The time synchronization by the global positioning system (GPS) enables the WAMPAC functionality, and the 30 Hz (and higher) sampling rate further provides the potential for analyzing low-frequency oscillations [9]–[11]. Fourier spectral analysis has been applied to synchrophasor data to estimate the eigenvalues for monitoring inter-area oscillation [12]. In [13], an adaptive stochastic subspace identification algorithm is proposed to estimate the mode and damping of low-frequency oscillations via fast computation.

In addition to theoretical work, there are many PMU applications for online oscillation detection: The PSGuard wide-area monitoring and control system [14] is capable of monitoring phase angle stability, voltage stability, and oscillation, etc. The Mode Meter has been tested with field measurement data [15] in WECC, and has been utilized as a supplementary online oscillation monitoring tool. California Independent System Operator (CAISO) has already started since 2008 to utilize synchrophasor data in the control room and to use a realtime dynamics monitoring system (RTDMS) for WECC wide-area visualization and monitoring [16].

Although much progress has been made in data-driven monitoring and control of low frequency oscillations, one of the big remaining challenges in today's practice is the early detection of the *oscillation starting point*. Most existing approaches iteratively perform PMU-data-driven modal analysis with a moving window. However, such a scheme may increase not only the computational burden by use of the ambient data, but also the false alarm rate.

This paper directly addresses the issue of detecting the starting point of an oscillation. The proposed method is purely data-driven. First, principal component analysis (PCA) is applied to extract the low-dimensional features, PCs, from the high-volume raw measurements. Then, a persistent-homology-based cyclicity response function is proposed by using the PCs for oscillation detection. A pre-defined threshold can be determined through statistical offline training with historical eventful PMU data. Cyclicity response can detect the transient oscillatory behavior, and a value exceeding the threshold indicates the occurrence of an oscillation. The detected occurrence time can serve as a triggering signal for some existing modal analysis tools.

This paper is organized as follows. Section II introduces the PCA-based dimensionality reduction of synchrophasor data. The persistent-homology-based cyclicity response is presented in Section III for oscillation detection. The implementation of the proposed methodology is provided in Section IV. In Section V, numerical examples utilizing real PMU data are presented to demonstrate the effectiveness of the proposed methodology. Conclusions and possible future research direc-

This work is supported in part by NSF Contracts ECCS-1546682, ECCS-1150944, CNS-1646449, and NSF Science & Technology Center Grant CCF-0939370.

Y. Chen is with PJM Interconnection, Audubon, PA, 19403 USA (email: yang.chen@pjm.com).

H. Chintakunta is with Department of Electrical Engineering, Florida Polytechnic University, Lakeland, FL, 33805 USA (email: hchintakunta@flpoly.org).

L. Xie, and P. R. Kumar are with Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, 77843 USA (email: le.xie@tamu.edu, prk@tamu.edu).

Y. M. Baryshnikov is with Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61821 USA (email: ymb@illinois.edu).

tions are summarized in Section VI.

II. Dimensionality Reduction of PMU Data

With the increasing deployment of PMUs with high sampling rate in the power grid, the resulting large amount of PMU data raises new challenges to online applications in terms of the processing and computational burden.

In the proposed detection methodology, PCA is applied to the raw PMU measurements to reduce the dimensionality. PCA, as a linear dimensionality reduction technique, is attractive [17], [18] for its fast computational feature. Mathematically, PCA aims at finding a low-dimensional embedding from a high-dimensional space by preserving the most variance.

A measurement matrix is defined as $\mathbf{Y_e} := [\mathbf{y^{(1)}}, \ldots, \mathbf{y^{(N)}}] \in \mathbb{R}^{n \times N}$, which includes a total number of N measurements. Each measurement has n samples constituting a time history, i.e., $\mathbf{y^{(i)}} := [y_1^{(i)}, \ldots, y_n^{(i)}]^T$, $i = 1, \ldots, N$.

One can apply PCA on $\mathbf{Y_e}$ to extract the PCs as $\mathbf{PC_i} = (\mathbf{Y_e}(t) - \mu)\,\mathbf{u_i}$, where $\mathbf{PC_i}$ is the ith PC, and the $\mathbf{u_i}$'s, $i = 1, \ldots, N$ are the orthonormal eigenvectors for the covariance matrix of $\mathbf{Y_e}$ corresponding to the nonnegative eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_N$. Let $\mu = E[\mathbf{Y_e}]$.

III. Low-frequency Oscillation Detection based on Cyclicity Response

In this section, we describe *cyclicity response*, which is used to detect the oscillatory behavior of the power grid. The oscillation detection procedure takes the PCs as inputs, and calculates the corresponding cyclicity responses as outputs. Cyclicity response is based on:

1. Takens' delay embedding theorem [19], which implies that a mapping $g : X \to \mathbb{R}^2$ from the phase space [20] of an oscillatory system to $\mathbb{R}^2$ contains a cyclic structure, and
2. Persistent homology [21], which quantifies the cyclic structure of a point cloud.

In the context of the application considered here, X represents the phase space of the power grid, and the mapping g represents the projection of PMU measurements onto the two significant principal components, which yields a scatter plot. The premise here is that when the power grid shifts into an oscillatory condition, the projection of the PMU measurements will start to behave cyclically (either circularly or elliptically), which can be detected using persistent homology computation as described in Section III-B.

A. Delay Embedding Theorem

Consider a discrete-time dynamical system $f : X \to X$, which has a strange attractor [22] . Then for sufficiently large k, and for any generic function $g : X \to \mathbb{R}^k$, the delay embedding theorem states that the map

$$G(x) = \left(g(x), g\left(\phi(x)\right), g\left(\phi^2(x)\right), \ldots, g\left(\phi^{k-1}(x)\right)\right) \quad (1)$$

is an embedding. The dimension k is determined by the structure of the attractor. For an attractor with dimension d, $k > 2d$ is usually sufficient. Readers may refer to [19], [23] for a formal treatment of the delay embedding theorem.

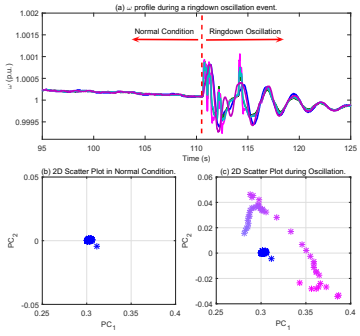

Fig. 1. Bus frequency profile and scatter plot during a ringdown oscillation.

We model the power grid as a discrete-time dynamical system f, where the time points are determined by the sampling rate of the PMUs. When the power grid is in an oscillatory state, the phase of the system exhibits cyclic behavior, and by the delay embedding theorem, so does the mapping g. Fig. 1 illustrates this phenomenon in the power grid. Fig. 1(b) shows the scatter plot when the power grid is not undergoing oscillations, while Fig. 1(c) shows the case of an oscillation. $\mathbf{PC_1}$ and $\mathbf{PC_2}$ are the first two PCs from PCA, and the reason for using only the first two PCs is demonstrated in [17]. The task then is to distinguish between these two cases automatically, which is what we aim to accomplish using persistent homology computation.

B. Persistent Homology and Cyclicity Response

We now provide a brief introduction to persistent homology. The purpose here is to provide the reader an intuition for what the tool computes. A formal introduction to the theory is beyond the scope of this article, and we refer the reader to [21] for a thorough exposition. Persistent homology may be viewed as a generalization of traditional hierarchical clustering to other topological features. We first briefly describe hierarchical clustering, followed by how this is paralleled by persistent homology to cyclic structures.

1) Hierarchical Clustering: Given a point cloud, hierarchical clustering is a traditional tool for obtaining a multi-scale summary of how the points are clustered. Fig. 2 illustrates single linkage clustering where the output, either a "dendrogram" or a "barcode" as described below, provides a multi-scale summary of the clustering. The algorithm is as follows: for each value ϵ, draw an edge between any two points (v_i, v_j), where the distance between the two points $d(v_i, v_j) \leq \epsilon$. The vertices corresponding to connected components in the resulting graph are the clusters. One can then follow how these connected components merge as we increase the value of ϵ. This pattern in which the connected components merge can be described by either a dendrogram (as in Fig. 2(b)) or by a barcode (as in Fig. 2(c)). Each vertical line segment along the x-axis in Fig. 2(b) corresponds to a connected component, and these segments are joined when the connected components merge as ϵ increases along the y-axis. Equivalently, we may view the points as connected components "born" at $\epsilon = 0$, and whenever two components merge together, we may say one

of the connected components has "died" at that ϵ value. This birth-death pattern is represented in a "barcode" in Fig. 2(c), where each bar represents the birth and death of a connected component, and the birth and death points can be read from the x-axis. For example, for the point cloud in Fig. 2(a), it is apparent from the output that there are two clusters.

2) Persistent Homology: We now describe how to generalize the above hierarchical clustering to cyclic structures. For this purpose, we use the example of the point cloud shown in Fig. 3(a). In the case of clustering described above, we have followed the "birth" and "death" of the connected components in the graph. Here, we instead follow the birth and death of cycles.

As in the case of clustering, we increase the value of ϵ, and for each ϵ value we add all the edges in the Delaunay triangulation [24] whose lengths are less than or equal to ϵ.

At ϵ_1, we see the birth of two cycles shown in red lines in Fig. 3(b): c_1 and c_2. We will consider a cycle "dead" if it gets "filled in" by 3-sided triangles, i.e., when it bounds a set of 3-sided triangles. We see that c_1 dies at ϵ_2, and c_2 dies at ϵ_3, resulting in the barcode shown in Fig. 3(e). It is straightforward to infer from the barcode that the point cloud exhibits two cyclic structures, one of which is larger than the other.

Given a point cloud, the *cyclicity response* is defined as the length of the largest bar in the barcode. As evidenced in the above example, the longest bar represents the biggest cyclic structure in the point cloud, and we see a cyclic structure when the system is undergoing oscillations. For the sake of concreteness, we now provide some algebraic background which gives rise to these bars.

For a triangulation K_ϵ, as in Figure 3, one can define an abstract vector space $H_1(K_\epsilon)$, called the first homology space, where the basis elements are equivalence classes of cycles in K_ϵ. All cycles which bound triangles are considered to be equivalent to 0. For example, the equivalence class of $[c_1] \neq 0$ in K_{ϵ_1}, whereas $[c_1] = 0$ in K_{ϵ_2}. Also, since ϵ_2 is the smallest value for which $[c_1] = 0$, we say that $[c_1]$ persists in the interval $[\epsilon_1, \epsilon_2)$, and this is represented as a bar from ϵ_1 to ϵ_2 in the output. More generally, the homology which persists in the interval $[a, b)$ is given by $Img\left(H_1(K_a) \xrightarrow{i_*} H_1(K_b)\right)$, where i_* is the map induced by the inclusion $i : K_a \to K_b$. In other words, the number of bars in any interval $[a, b)$ in the barcode is equal to the dimension of $Img\left(H_1(K_a) \xrightarrow{i_*} H_1(K_b)\right)$.

From a data analysis perspective, each bar of interval $[a, b)$ in the barcode corresponds to a cyclic structure which is present in the triangulations corresponding to thresholds in that interval. The length of the longest bar corresponds to the largest cyclic structure present in the point cloud, which is the basis of cyclicity response.

For further technical details behind the computation of the barcode, we refer readers to [21].

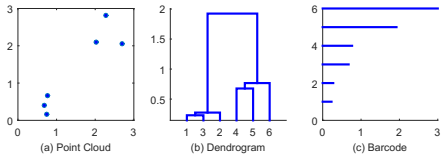

Fig. 2. Demonstration of hierarchical clustering.

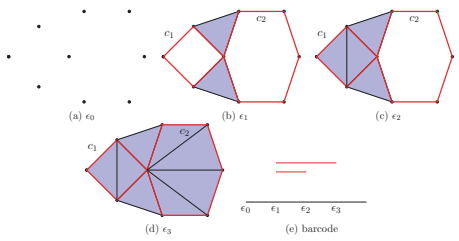

Fig. 3. Demonstration of the computation of persistent homology.

IV. Implementation of Algorithm

In this section, we propose an algorithm to detect low-frequency oscillations in power grids. The flowchart is presented in Fig. 4. We will not address the issue of data quality. The algorithm can be described as follows:

Offline Learning

(1) Use historically eventful PMU data, including both oscillations and other non-oscillatory events, to form $\mathbf{Y_e^{off}}$.

(2) For each collected event, perform PCA on $\mathbf{Y_e^{off}}$ to extract the PCs.

(3) Calculate the cyclicity response function using the current PCs.

(4) Based on statistical analysis on all historical events, calculate the detection threshold.

Online Monitoring

1) Form $\mathbf{Y_e}$ with real-time PMU measurements, and perform PCA on $\mathbf{Y_e}$ to extract the PCs.

2) Calculate the cyclicity response using the current PCs.

3) If the resulting cyclicity response has a value exceeding the threshold at time t, then an oscillation is detected. (The time t can serve as the triggering time point of existing modal analysis tools, such as [9], [15]). Otherwise, return to **Step 1**, update $\mathbf{Y_e}$ and repeat.

V. Numerical Examples

This section illustrates the use of the proposed methodology in detecting oscillations from other types of events. The detection accuracy is discussed first, with the detection capability shown next using real PMU data [25].

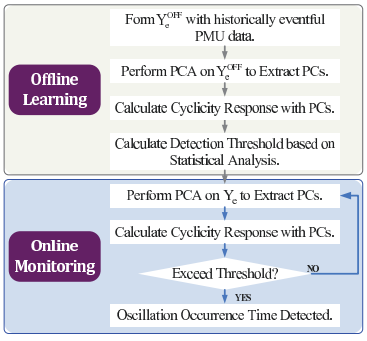

Fig. 4. Implementation of the algorithm.

A. Oscillation Detection Accuracy

Real eventful PMU data from the Western and Texas Interconnection are processed, and 30 events are selected, including oscillation, unit tripping, switching events, etc. A threshold of 2 is determined based on the statistical analysis of these events.

With such a threshold, the apparent error of oscillation detection is 13.33%, including a mis-detection rate of 10%, and a false-alarm rate of 3%.

B. Cyclicity Response for Oscillation Detection

Fig. 5 demonstrates the voltage magnitude profile and corresponding cyclicity response for a sustained oscillation. From Fig. 5(b), the cyclicity response is seen to be able to detect the sustained oscillation around the same time as it occurs. So we see that with the cyclicity response one does not need lengthy buffering of data for oscillation detection.

Fig. 6 illustrates the voltage magnitude profile (Fig. 6(a.1)) and corresponding cyclicity responses (Fig. 6(b.1)) for several ringdown oscillations. The second oscillation is highlighted in Figs. 6(a.2) and 6(b.2). As can be observed, with the threshold of 2, the cyclicity response is able to detect the starting time of the ringdown oscillation around 0.17s earlier than if only using the voltage magnitude profile in Fig. 6(a.2).

Figs. 7 and 8 demonstrate two non-oscillatory events: unit tripping and switching events. From the corresponding cyclicity responses in Figs. 7(b) and 8(b), it can be observed that during the two non-oscillatory events, the cyclicity responses are below the threshold. Therefore, no oscillation is detected.

Through Figs. 5-8, the capability to detect low-frequency oscillations by use of the proposed algorithm is demonstrated.

VI. Concluding Remarks

We have proposed a new data-driven methodology to detection of low-frequency oscillations. PCA is first applied on raw PMU measurements to extract PCs. Then, oscillation detection is achieved by use of the persistent-homology-based cyclicity response on the PCs. Tests based on real PMU data suggest the effectiveness of the proposed algorithm for quick detection of such oscillations.

Our future work includes testing the robustness of the proposed algorithm by use of more real-world examples, and using the results for training. Considering the availability of more eventful data in the future, we expect that the detection threshold can be subsequently refined to achieve higher accuracy of oscillation detection.

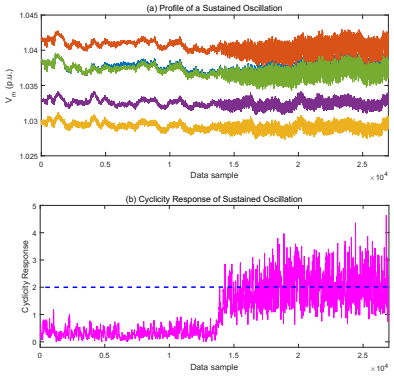

Fig. 5. Detection of a sustained oscillation.

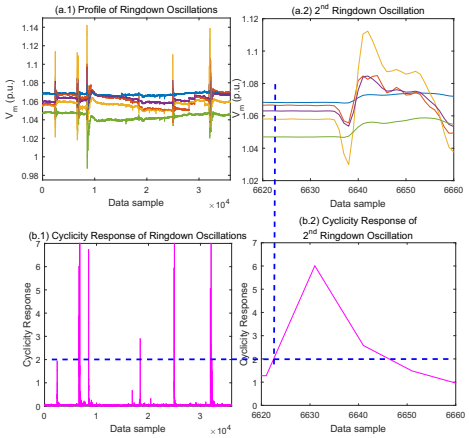

Fig. 6. Detection of ringdown oscillations.

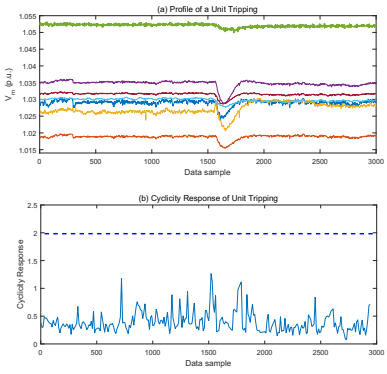

Fig. 7. Profile and cyclicity response of a unit tripping event.

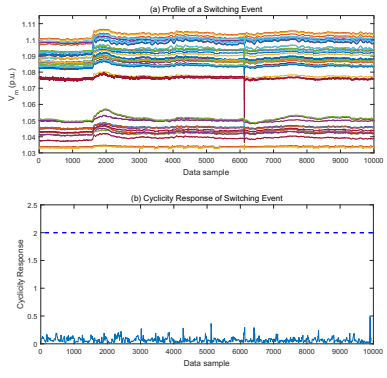

Fig. 8. Profile and cyclicity response of a switching event.

REFERENCES

[1] North American Electric Reliability Council, "1996 system disturbances," August 2002. [Online]. Available: http://www.nerc.com/pa/rrm/ea/System%20Disturbance%20Reports%20DL/1996SystemDisturbance.pdf

[2] M. Klein, G. Rogers, P. Kundur *et al.*, "A fundamental study of inter-area oscillations in power systems," *IEEE Transactions on Power Systems*, vol. 6, no. 3, pp. 914–921, 1991.

[3] P. Kundur, N. J. Balu, and M. G. Lauby, *Power system stability and control*. McGraw-hill New York, 1994, vol. 4, no. 2.

[4] A. R. Messina, J. Ramirez, C. Canedo *et al.*, "An investigation on the use of power system stabilizers for damping inter-area oscillations in longitudinal power systems," *Power Systems, IEEE Transactions on*, vol. 13, no. 2, pp. 552–559, 1998.

[5] H. Vu and J. Agee, "Comparison of power system stabilizers for damping local mode oscillations," *Energy Conversion, IEEE Transactions on*, vol. 8, no. 3, pp. 533–538, 1993.

[6] A. Phadke and R. M. de Moraes, "The wide world of wide-area measurement," *Power and Energy Magazine, IEEE*, vol. 6, no. 5, pp. 52–65, 2008.

[7] V. Terzija, G. Valverde, D. Cai, P. Regulski, V. Madani, J. Fitch, S. Skok, M. M. Begovic, and A. Phadke, "Wide-area monitoring, protection, and control of future electric power networks," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 80–93, 2011.

[8] P. Gao, M. Wang, S. G. Ghiocel, J. H. Chow, B. Fardanesh, and G. Stefopoulos, "Missing data recovery by exploiting low-dimensionality in power system synchrophasor measurements," *Power Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1–8, 2015.

[9] D. J. Trudnowski, J. W. Pierre, N. Zhou, J. F. Hauer, and M. Parashar, "Performance of three mode-meter block-processing algorithms for automated dynamic stability assessment," *Power Systems, IEEE Transactions on*, vol. 23, no. 2, pp. 680–690, 2008.

[10] J. Ma, P. Zhang, H.-j. Fu, B. Bo, and Z.-y. Dong, "Application of phasor measurement unit on locating disturbance source for low-frequency oscillation," *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 340–346, 2010.

[11] J. Ma, T. Wang, Z. Wang, and J. S. Thorp, "Adaptive damping control of inter-area oscillations based on federated kalman filter using wide area signals," *Power Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 1627–1635, 2013.

[12] N. Kakimoto, M. Sugumi, T. Makino, and K. Tomiyama, "Monitoring of interarea oscillation mode by synchronized phasor measurement," *Power Systems, IEEE Transactions on*, vol. 21, no. 1, pp. 260–268, 2006.

[13] S. Nezam Sarmadi and V. Venkatasubramanian, "Electromechanical mode estimation using recursive adaptive stochastic subspace identification," *Power Systems, IEEE Transactions on*, vol. 29, no. 1, pp. 349–358, 2014.

[14] ABB, "Wide area monitoring systems: Portfolio, applications and experiences," April 2012. [Online]. Available: http://www05.abb.com/global/scot/scot221.nsf/veritydisplay/3d85757b8c7f3bb6c125784d0056a586/$file/1KHL501042%20PSGuard%20WAMS%20Overview%202012-04.pdf

[15] Z. Huang, N. Zhou, F. K. Tuffner, Y. Chen, D. J. Trudnowski, R. Diao, J. C. Fuller, W. A. Mittelstadt, J. F. Hauer, and J. E. Dagle, *MANGO: Modal Analysis for Grid Operation: a Method for Damping Improvement Through Operating Point Adjustment*. Pacific Northwest National Laboratory, 2010.

[16] North American SynchroPhasor Initiative, "Synchrophasor update success story," June 2008. [Online]. Available: https://www.naspi.org/meetingarchives

[17] L. Xie, Y. Chen, and P. R. Kumar, "Dimensionality reduction of synchrophasor data for early anomaly detection: Linearized analysis," *Power Systems, IEEE Transactions on*, vol. 29, no. 4, pp. 2784–2794, 2014.

[18] Y. Chen, L. Xie, and P. R. Kumar, "Power system event classification via dimensionality reduction of synchrophasor data," in *Sensor Array and Multichannel Signal Processing Workshop, 2014. SAM 2014. 8th IEEE*. IEEE, 2014, pp. 57–60.

[19] L. Noakes, "The takens embedding theorem," *International Journal of Bifurcation and Chaos*, vol. 1, no. 04, pp. 867–872, 1991.

[20] "Phase space," available: https://en.wikipedia.org/wiki/Phase_space.

[21] A. Zomorodian and G. Carlsson, "Computing persistent homology," *Discrete & Computational Geometry*, vol. 33, no. 2, pp. 249–274, 2005.

[22] "Attractor," available: https://en.wikipedia.org/wiki/Attractor.

[23] F. Takens, *Detecting strange attractors in turbulence*. Springer, 1981.

[24] P.-L. George and H. Borouchaki, "Delaunay triangulation and meshing," 1998.

[25] [Online]. Available: https://www.naspi.org/site/Module/Meeting/Forms/General.aspx?m__ID=MEETING&meetingid=347