Power Oscillation Monitoring using Statistical Learning Methods

Hallvar Haugdal and Kjetil Uhlen
Department of Electric Power Engineering
Norwegian University of Science and Technology
Trondheim, Norway

Abstract—This paper describes development and testing of a method for estimation of electromechanic modes and corresponding mode shapes from frequency or voltage angle measurements in Wide Area Monitioring Systems. The method uses Complex Principal Component Analysis to perform a decomposition of the dynamics captured in the measurements, reducing any detected oscillations into sets of parameters. The numerous sets of parameters generated after a time period can be interpreted as points or observations in high-dimensional space, on which a clustering algorithm can be applied to pinpoint areas where high densities of points are accumulated. Results show that highly accurate estimates of oscillatory modes in the power system and their mode shapes can be derived by averaging the point observations belonging to each cluster.

Important development described in this paper includes the use of voltage angle as input and testing on a medium size power system simulation model with synchronized measurements from 44 nodes. Furthermore, introduction of the DBSCAN clustering algorithm shows very promising results when applying the method to recorded phasor measurements from the Nordic power system.

Index Terms—Clustering, Mode Shapes, Phasor Measurement Units, Power Oscillations, Complex Principal Component Analysis

I. INTRODUCTION

New methods for estimation of modes and mode shapes are frequently reported in the literature, motivated by blackouts in the past that could possibly have been avoided given sufficient information about oscillatory modes. Most methods are either aimed at providing estimates during ambient conditions, or from more pronounced oscillations during transient responses or ringdowns. In [1], a monitoring scheme based on two algorithms is proposed: Ambient conditions are analyzed using a wavelet-based method, while Prony's method, triggered by a pre-defined rule, is used to analyze ringdowns. This method, however, is not able to capture information about mode shapes. The method described in [2] also captures mode shapes from ringdowns, but is not applicable during ambient conditions. In [3], a Multichannel AutoRegressive Moving Average eXogenous (ARMAX) model is used to extract information on modes and mode shapes from ambient conditions. A slight deterioration of most ambient methods is reported when an oscillatory event is introduced in the analyzed time series [4]. In [5], a Bayesian approach to mode estimation from ambient conditions is proposed, where specific frequency bands are targeted to reduce the complexity of the problem.

The method described in this paper was developed for analyzing short time series of pronounced ringdown events. However, results indicate that mode estimates can also be produced from longer time spans of ambient conditions. In this respect, this method can be considered more general than most of the above mentioned methods.

The method consists of two main parts: In Part I, a sliding time window containing the samples from the last 5 to 10 seconds is analyzed using Complex Principal Component Analysis (CPCA). The oscillatory dynamics in the time window are decomposed into monofrequency components and a set of coefficients that describe the amplitude and phase of each frequency in the measurements. This information is stored as a set of point observations. In Part II of the method, the DBSCAN clustering algorithm is applied to all the observations from a given period (on the scale of minutes). The observations in densely populated areas of the input space are associated with dynamics that occur frequently, and are grouped together. Finally, averaged modes and mode shapes can be computed as the average of all the observations belonging to each cluster.

CPCA is described in [6], intended for analysis of geophysical phenomena. Similar work on analysis of electromechanical modes is reported in [7]–[10], where the term Empirical Orthogonal Functions (EOF) is used instead of Principal Component Analysis (PCA). An important contribution from the proposed method described in this paper is the second part, where the observations are clustered. The purpose of this part is to present the information from the CPCA dynamics decomposition in a more meaningful and understandable way to the operators.

The basic structure of the method (CPCA dynamics decomposition combined with clustering for averaging) has been described previously in [11], where it is applied to analysis of large disturbances/ringdowns. Continuing this work, we propose further development of the method by introducing the DBSCAN clustering algorithm, in addition to further testing on a medium size grid and on ambient conditions.

Section II gives a description of the method, based on the more thorough description in [11], however with some modifications. Section III describes results from applying the method to simulated and recorded PMU data, followed by discussions and conclusions in Sections IV and V.

II. EMPIRICAL METHOD FOR ESTIMATION OF MODES AND MODE SHAPES

The proposed method consists of two main parts: In Part I, a sliding time window is analyzed using CPCA. The dynamics in the input series are decomposed into monofrequency components, allowing the dynamics occurring to be quantified in terms of frequency and correlation. The dynamics are reduced to a set of parameters, which are stored and used in the further analysis. One set of parameters, describing an oscillation with a specific frequency and amplitude- and phase correlation between all input measurements, is referred to as an *observation*.

As the time window slides through time, numerous observations are generated. In Part II of the method, the point observations are clustered, such that similar observations are grouped together. Finally, the averaged modes and mode shapes are computed as the centroids of the clusters.

Part I: Dynamics Decomposition using Complex Principal Component Analysis

The sliding time window can be considered as an $M \times N$-matrix, where M is the number of measurement series and N is the number of samples per series:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_M \end{bmatrix} = \begin{bmatrix} x_1(t_1) & x_1(t_2) & \cdots & x_1(t_N) \\ x_2(t_1) & x_2(t_2) & \cdots & x_2(t_N) \\ \vdots & & \ddots & \\ x_M(t_1) & x_M(t_2) & & x_M(t_N) \end{bmatrix} \quad (1)$$

Applying PCA, as described in [12], we effectively transform the correlated series above into a set of uncorrelated series, or Principal Components (PCs),

$$\mathbf{S} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \\ \vdots \\ \mathbf{s}_{M_{PC}} \end{bmatrix} = \mathbf{U}^T \mathbf{X} \quad (2)$$

The matrix $\mathbf{U}$ is constituted by the eigenvectors of the covariance matrix of $\mathbf{X}$. The PCs are sorted such that the first PC captures the most variance in the input series, the second component captures second most variance, and so on. The number of components M_{PC} is chosen such that the significant variance in the input series is captured (i.e. $\approx 99\%$ etc.). Given that the input series are highly correlated, we can analyze the dynamics following a disturbance by considering only a few PCs, rather than the full set of input series.

If phase shifted oscillations occur in two different measurements, we can not capture the oscillations using only one component. This makes the case for employing a complex variant of PCA. Complex Principal Component Analysis (CPCA), as described in [6], is very similar to PCA, but prior to applying the linear transform in (2), the Hilbert transform is applied to the input series to generate a complex time series. Applying PCA (as described above) to the complex series yields a complex covariance matrix and a complex transformation matrix.

Ensuring a well posed Hilbert Transform requires that the input series contains only one frequency, and that any non-oscillatory trends are subtracted prior to applying the transform. This can be achieved by adopting a two-layer structure of PCA: In the first layer, PCA is applied to the input measurement series. Given some correlation, this significantly reduces the number of components that needs to be analyzed by CPCA.

In the second layer, CPCA is applied: First, the non-oscillatory trend of each signal is computed as the residual from Empirical Mode Decomposition, as described in [13]. Further, the Hilbert transform is computed before the second layer of PCA equations is applied.

Proceeding with the PCs computed as shown above, we subtract the residuals $\mathbf{R}$ computed using EMD:

$$\mathbf{S}' = \mathbf{S} - \mathbf{R} \quad (3)$$

Applying the Hilbert Transform yields the complex series:

$$\mathbf{Y} = \mathbf{S}' + jH(\mathbf{S}') \quad (4)$$

Here, $H(\mathbf{S}')$ denotes series-wise application of the Hilbert transform. Also, in the practical implementation roughly 10% should be tapered in each end of the series at this stage, to limit end-effects [6].

The complex series $\mathbf{Y}$ is the input to the second layer, where CPCA is applied:

$$\mathbf{Z} = \mathbf{V}^H \mathbf{Y} = \mathbf{V}^H \left(\mathbf{S}' + jH(\mathbf{S}') \right) \quad (5)$$

The superscript $(\cdot)^H$ denotes the conjugate transpose. $\mathbf{V}$ is the CPCA transformation matrix, analogously to the PCA transformation matrix $\mathbf{U}$ in (2).

It can be shown that $\mathbf{U}$ is orthonormal [12] and $\mathbf{V}$ is unitary [6], such that their inverses are simply the transpose and the conjugate transpose, respectively. Inverting the above transformation matrices, we can write the input measurement series in terms of PCs and CPCs:

$$\begin{aligned} \mathbf{X} &= \mathbf{US} \\ &= \mathbf{U}(\mathbf{S}' + \mathbf{R}) \\ &= \mathbf{U}(\mathrm{Re}(\mathbf{Y}) + \mathbf{R}) \\ &= \mathbf{U}(\mathrm{Re}(\mathbf{VZ}) + \mathbf{R}) \end{aligned} \quad (6)$$

Neglecting the residual and defining the matrix $\mathbf{W} = \mathbf{UV}$, we get

$$\mathbf{X} = \mathrm{Re}(\mathbf{UVZ}) = \mathrm{Re}(\mathbf{WZ}) \quad (7)$$

Finally, we can write

$$\mathbf{x}_i = \sum_{j=1}^{M} \mathrm{Re}(w_{ij}\mathbf{z}_j) \approx \sum_{j=1}^{M_{CPC}} \mathrm{Re}(w_{ij}\mathbf{z}_j) \quad (8)$$

This gives the contribution of the Complex Principal Component (CPC) $\mathbf{z}_j$, to the measurement $\mathbf{x}_i$ as the coefficient w_{ij}, i.e. element (i, j) of the matrix $\mathbf{W}$.

Insignificant components are discarded during the decomposition, such that only M_{PC} components are kept from the first

layer and M_{CPC} from the second layer. Thus, the matrices $\mathbf{U}$, $\mathbf{V}$ and $\mathbf{W}$ will be of dimension $M \times M_{PC}$, $M_{PC} \times M_{CPC}$ and $M \times M_{CPC}$, respectively.

At the end of the dynamics decomposition of each time window, a point on the following form is stored for each of the CPCs:

$$\mathbf{p} = [f_j, \text{Re}(w_{1j}), \text{Im}(w_{1j}), \text{Re}(w_{2j}), \text{Im}(w_{2j})... \\ ...\text{Re}(w_{Mj}), \text{Im}(w_{Mj})]) \tag{9}$$

All the point observations from a given period are analyzed further in Part II of the method, discussed next.

Part II: Mode averaging using DBSCAN Clustering

In the second part of the method, we cluster the observations generated by the dynamics decomposition in Part I. This allows us to filter out bad observations/noise, and to compute averages of the frequency and mode shape of different modes.

The clustering is carried out using the DBSCAN clustering algorithm, described in [14]. The algorithm clusters observations in densely populated areas together, and classifies observations in areas with low densities as noise. This algorithm requires the user to input two parameters, which determines the density of points requred to form clusters, described thoroughly in [14]. In the results reproduced below, the parameter values Eps = 0.2 and MinPts = 75 are found by trial and error to give good results.

Clustering the observations allows us to compute the average of all the observations in each cluster, resulting in the final, averaged modes and corresponding mode shapes.

III. RESULTS

The method is tested on a 44-generator model of the Nordic Power System [15], and on recorded PMU data from an oscillatory event in the real Nordic Power System.

A. The Nordic 44 test network

A short circuit with a clearing time of $10\,\text{ms}$ is simulated in `DIgSILENT PowerFactory`. A $10\,\text{s}$ excerpt from the resulting rotor angle time series (of total length $20\,\text{s}$) is shown in Fig. 1. Applying the method to these time series results in two averaged modes and corresponding mode shapes; the first with a frequency of $0.70\,\text{Hz}$, the second $0.89\,\text{Hz}$. Examining Fig. 1, the $0.70\,\text{Hz}$-mode is the one most clearly visible. One of the time series with a high observability of this mode is highlighted in blue. The $0.89\,\text{Hz}$-mode is slightly less excited, but is also clearly visible in the time series highlighted in red. The observations contributing to each of the modes are indicated in Fig. 2.

Further, the averaged mode estimates are compared with results from modal analysis on the linearized model, conducted in `DIgSILENT PowerFactory`: It is found that both estimates computed using the proposed method resemble modes from linear analysis quite accurately. The two most poorly damped modes computed using modal analysis have frequencies of $0.66\,\text{Hz}$ and $0.93\,\text{Hz}$ respectively, which matches the above results well. Here, one has to keep in mind that since
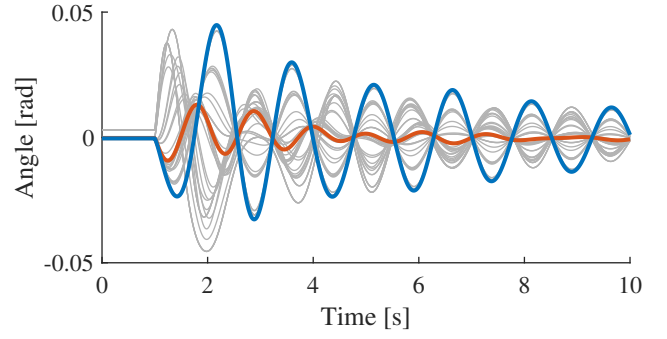

Fig. 1: Rotor angle time series generated by simulating a short circuit near one of the generators in the Nordic 44 test network. The mean of each time series is subtracted to clearly indicate the characteristics of the oscillations. The time series highlighted in blue and red has high content of the $0.67\,\text{Hz}$ and $0.93\,\text{Hz}$ modes, respectively.
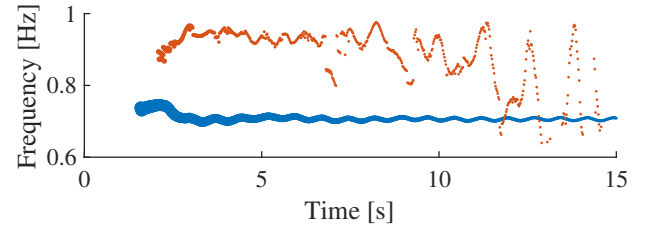

Fig. 2: The observations shown in blue are grouped together in one cluster, similarly for the observations shown in red. Averaging the observations in the clusters yields mode estimates of $0.70\,\text{Hz}$ and $0.89\,\text{Hz}$, which are close to the modes produced by modal analysis of the linearized model. The variance in the observations shown in red is larger than for the blue observations, caused by higher excitation and lower damping of the mode indicated in blue.

the system is not linear, one should not expect the oscillations following a large disturbance to match perfectly the results from small signal (modal) analysis.

A comparison of the mode shapes is shown in Figs. 3 and 4, for the $0.67\,\text{Hz}$-mode and the $0.93\,\text{Hz}$-mode respectively. For the $0.67\,\text{Hz}$-mode, the mode shape computed using modal analysis is very similar to the mode shape produced by the proposed method when using the rotor angle as input. The $0.93\,\text{Hz}$-mode is slightly less accurately reproduced by the proposed method, probably due to this mode being less excited and having higher damping.

The rotor angles of generators are not directly available from the data provided by PMUs. Instead, the voltage angle at the bus nearest to each generator could be used as an approximation to the rotor angle. The mode shapes computed using voltage angles as input are also shown in the two figures, revealing that this is slightly less accurate, but still well capable of capturing the main characteristics of the rotor angle dynamics.
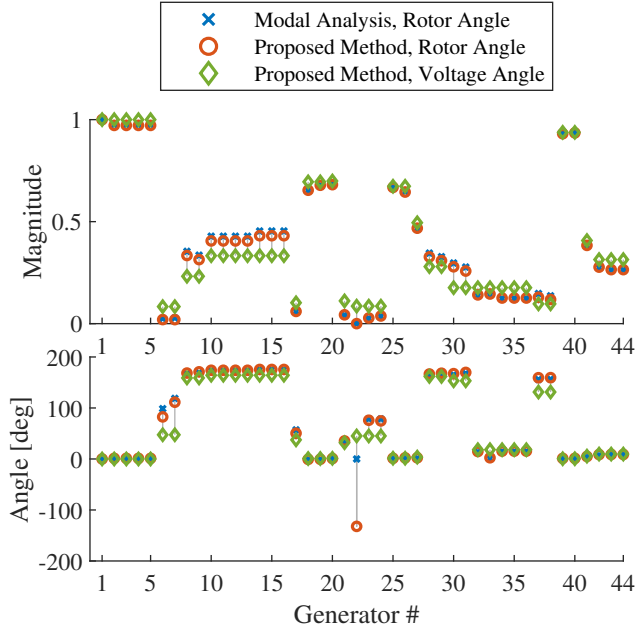
Fig. 3: Estimation accuracy of the 0.67 Hz-mode for each of the 44 measurements. The results produced by the method applied to rotor angle measurements is very accurate, while using the voltage angle as input gives slightly less accurate results.
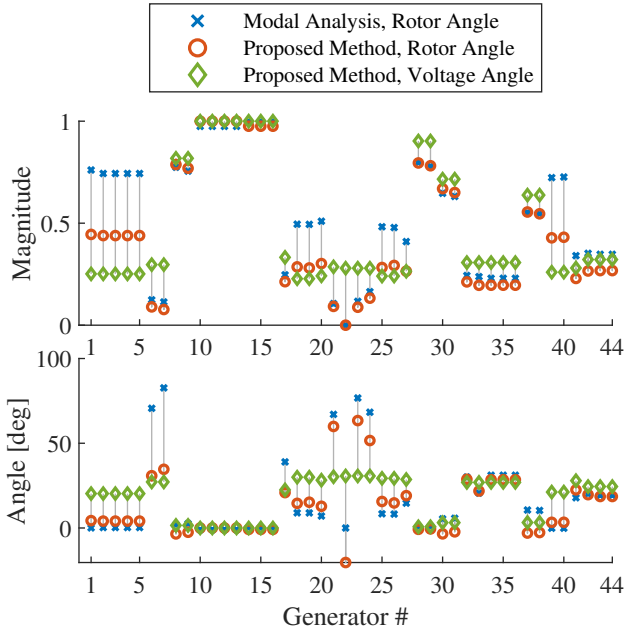

Fig. 4: Estimation accuracy of the 0.93 Hz-mode for each of the 44 measurements. The accuracy of this estimate is poorer than the accuracy for the more excited, less damped 0.67 Hz mode.
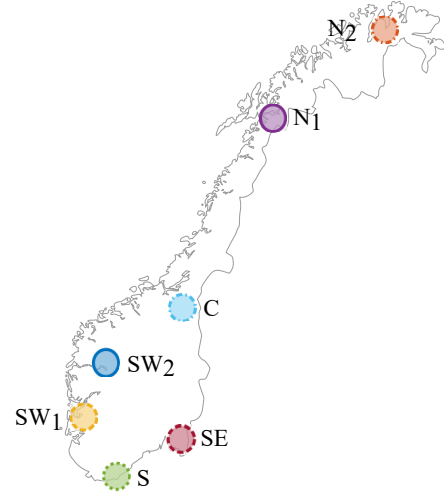

Fig. 5: The approximate locations of the seven PMUs recording the oscillatory event in the Nordic Power System.
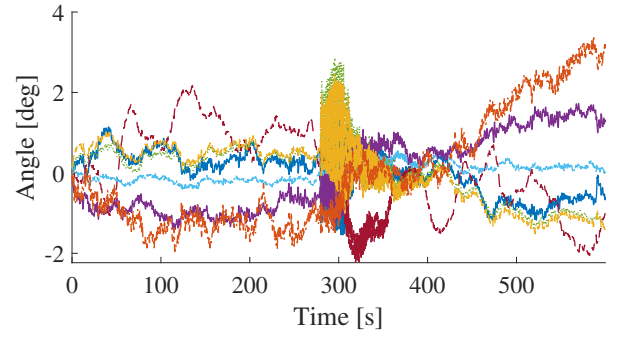

Fig. 6: Voltage angle time series measured by PMUs at the seven locations indicated in Fig. 5, using the same colors and styles for the locations.
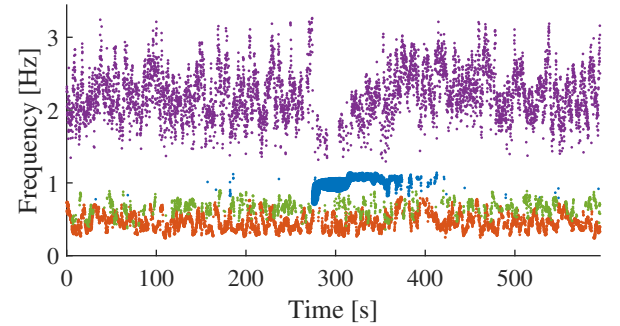

Fig. 7: The time instants and frequencies of the observations contributing to each of the four clusters are indicated by the coloring of the markers. The same colors are used in Figs. 8 and 9 to indicate the clustering division. Further, the size of each marker indicates the severity of the oscillations.

Mode estimate 1
f=1.02 Hz

Mode estimate 2
f=0.461 Hz
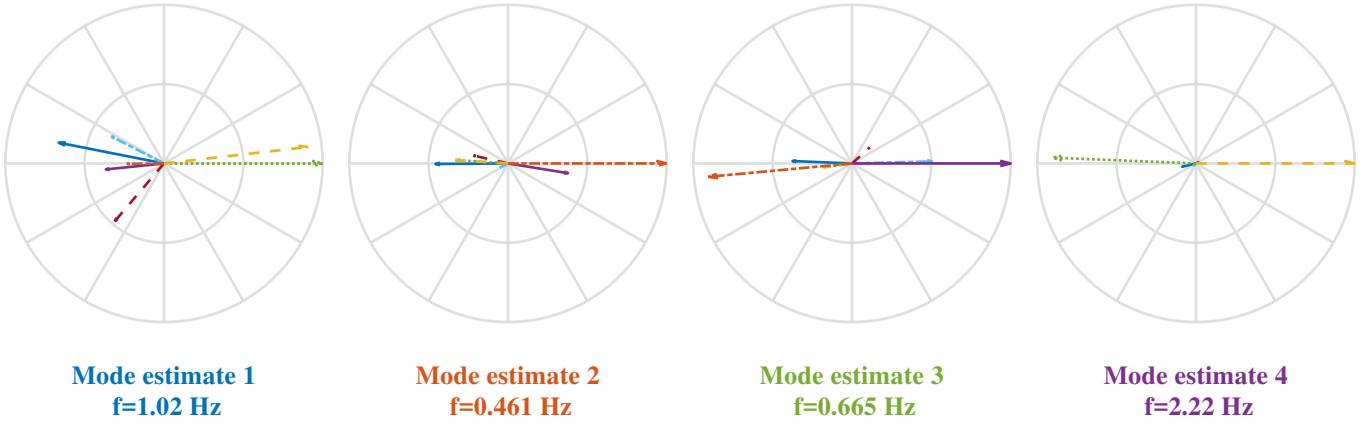
Mode estimate 3
f=0.665 Hz

Mode estimate 4
f=2.22 Hz

Fig. 8: The four resulting averaged mode shapes resulting from applying the method to the time series measured by PMUs during an oscillatory event in the Nordic Power System.
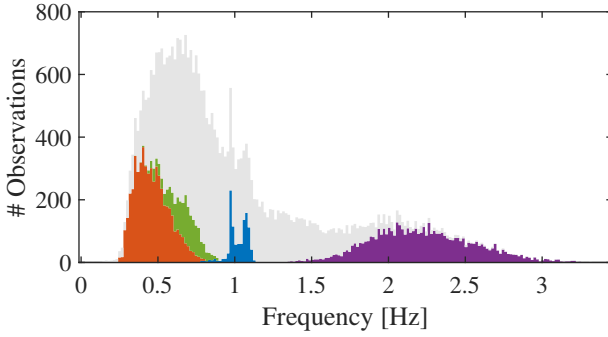

Fig. 9: The histogram indicates the number of observations with a given frequency. The coloring indicates the number of observations assigned to the different clusters.
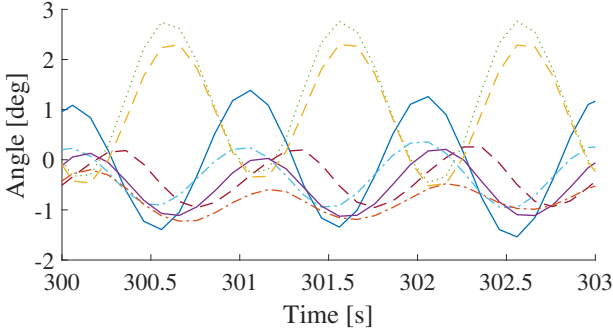

Fig. 10: An excerpt of the time series in Fig. 6 is shown, clearly indicating the amplitude and phase of the oscillations in each of the measurements. The same colors and styles are used to indicate the locations as in Fig. 5.

B. PMU-data from the Nordic Power System

Fig. 5 indicates approximate locations of seven PMUs in the Nordic Power System, where time series of voltage angles are recorded during a disturbance event that created system oscillations; as indicated in Fig. 6. Applying the proposed method to these time series results in four averaged modes and

mode shapes. The observations contributing to each cluster are indicated in Fig. 7, and the final averaged mode shapes are shown in Fig. 8. The histogram in Fig. 9 indicates the number of observations with a given frequency assigned to the different clusters.

The first mode estimate of about $1\,\text{Hz}$, indicated in blue in Figs. 7, 8 and 9, describes the severe oscillations starting around $280\,\text{s}$ and lasting until about $400\,\text{s}$. Observing the excerpt of the recorded time series in Fig. 10, it is clear that this mode shape accurately describes the oscillations. The second mode estimate of about $0.45\,\text{Hz}$, shown in red, corresponds to a well-known inter-area mode in the Nordic power system; here observed as oscillations between the north and south of Norway. The third estimate of about $0.67\,\text{Hz}$, appears to describe oscillations between the two northern locations N_1 and N_2, but also involving locations further south. The fourth estimate of about $2.2\,\text{Hz}$ (in average) could possibly describe a local mode involving generators in the south of Norway oscillation against each other. However, the sample rate in this recording is $10\,\text{Hz}$, which is too low to capture oscillations at $2\,\text{Hz}$ with a high degree of certainty. Moreover, the magnitude of these oscillations is very low. It is therefore difficult to say whether the fourth mode estimate is just noise, or if it really describes an electromechanical mode.

IV. DISCUSSION

The results from testing the method on the simulated data indicates that the proposed method is able to produce relatively accurate estimates of modes and mode shapes. The two most poorly damped modes of the Nordic 44 test network are detected quite accurately from a $20\,\text{s}$-measurement series following a short circuit, where both modes are excited at the same time.

Testing the method on recorded data from PMUs reveals that the proposed method is able to produce time averaged estimates of modes and mode shapes from ambient data, but also accurate descriptions of pronounced oscillations and ringdowns.

In a practical implementation, the CPCA-dynamics decomposition in Part I of the method would run continuously, producing mode observations for each new set of synchronized samples received from the PMUs. The clustering algorithm could be run a bit less frequent, but including all the observations from a given period back in time in the clustering.

Fig. 7 indicates the results that would constitute the detailed information to be included in an operator information system. The figure gives information about the certainty and accuracy of each of the four estimates, and gives a clear indication of the severity of the oscillations. Before the 1 Hz oscillations started, the operators would be presented with only the three estimates produced from ambient data (estimates 2, 3 and 4). As the oscillations started the operators would be presented with the 1 Hz mode estimate potentially seconds or tens of seconds after the oscillations started. In this case, the mode shape indicated the approximate location of the generator causing the instability (which was in the proximity of locations S and SW_1 in Fig. 5), and would have given the operators valuable and useful information for determining the proper remedial action. This example illustrates how critical information could be obtained very quickly, and as such contribute to increase operators' situational awareness.

To be able to provide an early warning of initiated oscillations, it would be of interest to perform the clustering as fast and often as possible, while at the same time having a clustering window of at least some minutes to be able to capture information of well-damped modes with low excitation from ambient data. The 35 000 observations from the 600 s-period in the PMU-example above takes approximately 50 s to cluster on a laptop with a Intel Core i7 2.80 GHz processor and 16.0 GB of RAM. However, this was carried out using a very naive implementation of the DBSCAN algorithm. The computation time could potentially be reduced significantly using a more sophisticated implementation; in [16] and [17], it is shown that the computation time for the clustering can be reduced by factors in the range 10 to 100 by taking advantage of parallel computing and graphics computing units (GPUs). If this was achievable in our case, the operators could be given an alert about the 1 Hz oscillations discussed above only seconds after they started, with an accurate mode shape supplementing the alert.

Regarding the computational complexity of the dynamics decomposition in Part I of the method, this is presumably less of a challenge. In the case with 44 simulated measurements from the Nordic 44 test network, the average and maximum computation time for the decomposition of one time window was about 20 ms and 70 ms, respectively, on the same computer as described above. Since the decomposition of one time window is independent on the previous and the next, this could be paralleled to allow the method to run online.

V. CONCLUSION

The results from applying the proposed method to both simulated data and recorded PMU data from the Nordic Power System are promising. The results from the case with simulated data indicate that relatively accurate estimates of simultaneously excited modes and corresponding mode shapes can be produced from ringdowns. Further, the case with recorded PMU data indicates that the method is able to estimate modes from ambient conditions at the same time as severe ongoing oscillations are captured; this is a comparative advantage to many other methods. Further work needs to be done to quantify the accuracy of the method, as well as improving the computational efficiency to a level required by online monitoring applications.

REFERENCES

[1] D. T. Duong and K. Uhlen, "An Empirical Method for Online Detection of Power Oscillations in Power Systems," in *Proc. IEEE Innovative Smart Grid Technologies (ISGT)*, 2018, pp. 758–763.

[2] N. R. Chaudhuri and B. Chaudhuri, "Damping and relative mode-shape estimation in near real-time through phasor approach," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 364–373, 2011.

[3] L. Dosiek and J. W. Pierre, "Estimating electromechanical modes and mode shapes using the multichannel ARMAX model," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1950–1959, 2013.

[4] L. Dosiek, N. Zhou, J. W. Pierre, Z. Huang, and D. J. Trudnowski, "Mode shape estimation algorithms under ambient conditions: A comparative review," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 779–787, 2013.

[5] J. Seppanen, S. K. Au, J. Turunen, and L. Haarla, "Bayesian Approach in the Modal Analysis of Electromechanical Oscillations," *IEEE Transactions on Power Systems*, vol. 32, no. 1, pp. 316–325, 2017.

[6] J. D. Horel, "Complex Principal Component Analysis: Theory and Examples," *Journal of Climate and Applied Meteorology*, vol. 23, no. 12, pp. 1660–1673, 1984.

[7] A. R. Messina and V. Vittal, "Extraction of Dynamic Patterns From Wide-Area Measurements Using Empirical Orthogonal Functions," *IEEE Transactions on Power Systems*, vol. 22, no. 2, pp. 682–692, 2007.

[8] P. Esquivel and A. R. Messina, "Complex Empirical Orthogonal Function analysis of wide-area system dynamics," in *IEEE Power and Energy Society 2008 General Meeting: Conversion and Delivery of Electrical Energy in the 21st Century*, 2008, pp. 1–7.

[9] P. Esquivel, "Wide-area wave motion analysis using complex empirical orthogonal functions," in *CCE International Conference on Electrical Engineering, Computing Science and Automatic Control*, vol. 52, no. 33, 2009, pp. 1–6.

[10] A. R. Messina, P. Esquivel, and F. Lezama, "Wide-area PMU data monitoring using spatio-temporal statistical models," in *2011 IEEE/PES Power Systems Conference and Exposition*. IEEE, 3 2011, pp. 1–7.

[11] H. Haugdal and K. Uhlen, "Mode Shape Estimation using Complex Principal Component Analysis and k-Means Clustering," *arXiv:1812.02966*, 2018.

[12] J. E. Jackson, *A User's Guide to Principal Components*, ser. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc., 1991.

[13] P. Flandrin, P. Goncalves, and G. Rilling, "Detrending and denoising with empirical mode decompositions," in *Proc. European Signal Processing Conf.*, no. 3, 2015, pp. 1581–1584.

[14] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proc. Second International Conf. on Knowledge Discovery and Data Mining (KDD)*, E. Simoudis, J. Han, and U. M. Fayyad, Eds. AAAI Press, 1996, pp. 226–231.

[15] S. H. Jakobsen and E. H. Solvang, "The Nordic 44 test network," 2018. [Online]. Available: https://figshare.com/articles/The_Nordic_44_test_Network/7464386

[16] C. Böhm, R. Noll, C. Plant, and B. Wackersreuther, "Density-based clustering using graphics processors," in *Proc. 18th ACM conference on Information and knowledge management (CIKM)*, 2009, p. 661.

[17] G. Andrade, G. Ramos, D. Madeira, R. Sachetto, R. Ferreira, and L. Rocha, "G-DBSCAN: A GPU accelerated algorithm for density-based clustering," *Procedia Computer Science*, vol. 18, pp. 369–378, 2013.