

A dynamic mode decomposition approach for large and arbitrarily sampled systems

Florimond Guéniat, Lionel Mathelin, and Luc R. Pastur

Citation: *Physics of Fluids* (1994-present) **27**, 025113 (2015); doi: 10.1063/1.4908073

View online: <http://dx.doi.org/10.1063/1.4908073>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/pof2/27/2?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Molecular dynamics investigation of nanoscale cavitation dynamics](#)

J. Chem. Phys. **141**, 234508 (2014); 10.1063/1.4903783

[A Lagrangian subgrid-scale model with dynamic estimation of Lagrangian time scale for large eddy simulation of complex flows](#)

Phys. Fluids **24**, 085101 (2012); 10.1063/1.4737656

[The effects of non-normality and nonlinearity of the Navier–Stokes operator on the dynamics of a large laminar separation bubble](#)

Phys. Fluids **22**, 014102 (2010); 10.1063/1.3276903

[A dynamic global-coefficient subgrid-scale model for large-eddy simulation of turbulent scalar transport in complex geometries](#)

Phys. Fluids **21**, 045109 (2009); 10.1063/1.3115068

[On the bimodal dynamics of the turbulent horseshoe vortex system in a wing-body junction](#)

Phys. Fluids **19**, 045107 (2007); 10.1063/1.2716813



A dynamic mode decomposition approach for large and arbitrarily sampled systems

Florimond Guéniat,^{1,2} Lionel Mathelin,¹ and Luc R. Pastur^{1,2}

¹LIMSI-CNRS, 91403 Orsay Cedex, France

²Paris Sud University, 91400 Orsay Cedex, France

(Received 6 June 2014; accepted 2 February 2015; published online 19 February 2015)

Detection of coherent structures is of crucial importance for understanding the dynamics of a fluid flow. In this regard, the recently introduced Dynamic Mode Decomposition (DMD) has raised an increasing interest in the community. It allows to efficiently determine the dominant spatial modes, and their associated growth rate and frequency in time, responsible for describing the time-evolution of an observation of the physical system at hand. However, the underlying algorithm requires uniformly sampled and time-resolved data, which may limit its usability in practical situations. Further, the computational cost associated with the DMD analysis of a large dataset is high, both in terms of central processing unit and memory. In this contribution, we present an alternative algorithm to achieve this decomposition, overcoming the above-mentioned limitations. A synthetic case, a two-dimensional restriction of an experimental flow over an open cavity, and a large-scale three-dimensional simulation, provide examples to illustrate the method. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4908073>]

I. INTRODUCTION

Identification of coherent structures is nowadays a pivotal instrument for understanding the phenomenology and dynamics of a fluid flow. A powerful analysis tool, the dynamic mode decomposition (DMD), has recently been introduced.^{19,20} Dynamic modes reveal spatial coherent structures associated with temporal spectral components, including temporal growth rates. It has been shown in Ref. 3 that Dynamic Mode Decomposition is connected to discrete Fourier transform, while¹⁴ discussed the connection between DMD and Koopman analysis, inherited from the dynamical system theory. Dynamic modes are therefore informative of the dynamical skeleton of a flow. DMD has been successfully applied to a wide range of systems, in particular, fluid flows such as cavity flows,^{9,20,22} or jet flows,^{21,23} among others. Some recent developments include error analysis of the identified growth rates,⁷ and improvements on the approximation method.^{3,11}

Dynamic Mode Decomposition requires time-resolved data, uniformly sampled in time. The time interval must be chosen small enough so as to resolve all time-scales of interest of the underlying dynamics. This sampling strategy brings severe constraints on the measurement workflow. As an example, consider the typical situation where the observable (sample) is a two-dimensional (2-D), two-component (2-C), velocity field acquired with a Particle Imagery Velocimetry (PIV) technique. Standard in PIV are fields of 1000×1000 pixels. Suppose the highest frequency of interest in the flow field is 200 Hz, a mild assumption, the Shannon-Nyquist criterion imposes a sampling frequency above 400 Hz to resolve the fast time scales of the flow. With 12-bit images, the resulting data rate is then already above 1 GB/s. Further, if the Fourier spectrum is wide-banded, the timespan of the acquisition procedure has to be large to capture the low frequency components. The combination of a high sampling frequency and a long acquisition sequence then quickly results in intractable constraints, both on the measurement chain hardware and the computational resources. This, of course, also applies when considering time-resolved datasets from numerical computations. Another limitation is that observation data could be corrupted from external or intrinsic sources, say from the experimental setup. As a typical example, one can think of acquisition failure or

corrupted images in PIV. Physically irrelevant information is then introduced when correcting such perturbations and adversely affects the standard DMD accuracy.

In this paper, we address these above-mentioned limitations. The method we present can essentially rely on a *subset* of the observation data, both in the spatial and time domains, hence naturally handling corrupted or missing data and arbitrary sampling. The resulting small and scattered dataset is suitably handled by techniques focusing the scarce available information on the dominant approximation modes by exploiting the compressibility of the decomposition in the retained format. The selection of the subset could rely on Fourier spectra estimated with a compressed sensing technique, see Ref. 5, among many others for an introduction to the theory. To avoid the computational cost associated with the compressed sensing problem, a suboptimal criterion based on statistics of the time-series is employed. Thanks to working on a subset of the original data, the computational burden is, sometimes drastically, alleviated compared to the standard DMD method, consequently allowing very large flow fields, and/or systems with a wide spectrum, to be analyzed. The key of the proposed method is to decouple the identification of the dominant temporal scales from their associated spatial modes.

Preliminary results of the present contribution were introduced in Ref. 10. A recent effort,¹¹ discusses the use of compressed sensing to select dominant DMD modes. Our method is different in the sense that it is not an *a posteriori* approach relying on a standard DMD, and hence subjected to its limitations, but instead directly determines the dominant dynamic modes from an (almost) arbitrary dataset.

After the standard DMD algorithm is briefly recalled in Sec. II, the different aspects of the proposed method are introduced in Sec. III. They are illustrated and compared with a standard DMD analysis on a synthetic 2-D dataset and a 2-D space- and time-resolved experimental dataset from the flow over an open cavity in Sec. IV. A three-dimensional numerical simulation, involving more than 2×10^6 spatial degrees-of-freedom, is also considered to illustrate the applicability of the present approach to large-scale situations. Closing remarks conclude the paper in Sec. V.

II. DYNAMICAL MODE DECOMPOSITION

In this section, we briefly recall the dynamic mode decomposition algorithm. More details can be found in Refs. 19 and 20.

We are interested in characterizing the linear, time-invariant, operator A mapping a real-valued observation vector $\mathbf{u}_n \in \mathbb{R}^{n_p}$ of the physical system at hand at time t_n to the observation \mathbf{u}_{n+1} when a time Δt has elapsed, $t_{n+1} = t_n + \Delta t$. Rearranging a sequence $K := (\mathbf{u}_1 \dots \mathbf{u}_N) \in \mathbb{R}^{n_p \times N}$ of N snapshots in matrices $K_1 := (\mathbf{u}_1 \dots \mathbf{u}_{N-1})$ and $K_2 := (\mathbf{u}_2 \dots \mathbf{u}_N) \in \mathbb{R}^{n_p \times (N-1)}$ results in characterizing the properties of A such that

$$AK_1 = K_2. \quad (1)$$

Introducing the economy-size singular value decomposition (SVD) of $K_1 \equiv U\Sigma W^*$, with W^* the Hermitian transpose of W , Eq. (1) leads to

$$\widehat{A}U := UU^*AU = U \underbrace{U^*K_2W\Sigma^{-1}}_{=:S}. \quad (2)$$

The matrix $S := U^*K_2W\Sigma^{-1}$ is similar to \widehat{A} , the approximation of A in the column space of K_1 , so that its eigenvalues $\{\lambda_k\}_k$ are eigenvalues of \widehat{A} , and eigenvectors $\{\boldsymbol{\varphi}_k\}_k$ of S are related to eigenvectors $\{\boldsymbol{\phi}_k\}_k$ (i.e., dynamic modes) of \widehat{A} via $\boldsymbol{\phi}_k \equiv U\boldsymbol{\varphi}_k$, $\forall k$. Since λ_s are complex, they express as

$$\lambda_k = \rho_k \exp\left(\sqrt{-1} \omega_k \Delta t\right), \quad \forall k. \quad (3)$$

Any eigenvalue is hence related to a frequency $f_k = \omega_k/2\pi$ and each spatial mode $\boldsymbol{\phi}_k$ is, therefore, oscillating in time at a single frequency. Equation (3) also implies that transient states can be captured by the DMD decomposition: when an eigenvalue has a modulus ρ larger than unity, the associated mode amplitude grows over time. Conversely, a modulus less than unity denotes a decaying behavior in time.

III. IDENTIFICATION OF DYNAMIC MODES FROM ARBITRARY DATA

A. Efficient use of the available data

As discussed in the Introduction, the standard dynamic mode decomposition approach presented above is restricted to data available in a strict format since the sampling must be uniform in time (so that the snapshots matrix K can be associated with a Krylov matrix) and the entire spatial field is considered. In many situations, the resulting decomposition exhibits some dynamic modes which are dominant with respect to some others and one is typically interested in characterizing these dominant modes *only*—in a nutshell, the real part $\Re[\phi_k]$ of a spatial mode ϕ_k considered as dominant lies in a cone of small mean angle $\angle_k := N^{-1} \sum_n^N \angle_{k,n}$ from the snapshots, with $\angle_{k,n} \in [-\pi, \pi]$, $\cos(\angle_{k,n}) = \Re[\phi_k]^* \mathbf{u}_n \|\mathbf{u}_n\|_2^{-1} \|\Re[\phi_k]\|_2^{-1}$. When the spectrum of the decomposition is sparse, in the sense that only a limited number of modes is responsible for describing most of the dynamics, the number of dominant modes is small and they hence require only a small amount of information to be determined. It results that the standard Krylov matrix K contains significantly more information than actually needed for this objective. This also suggests a route for a tractable DMD when the dataset is very large and cannot be processed by a standard DMD.

The key of the method we now expose is to extract information relevant for the dominant modes from scarce data (small dataset). To this end, the determination of the spectral features of the dominant modes is decoupled from the approximation of their spatial description, in contrast with the standard DMD where the two aspects are intricately coupled in the eigenproblem issued from Eq. (2).

To proceed, we rely on an alternative description of the dataset in a similar spirit as in Ref. 3. The DMD analysis leads to a decomposition of the observable on a spatial basis with complex temporal coefficients,

$$\mathbf{u}(t_n) =: \mathbf{u}_n \approx \sum_{k=1}^{N_{\text{md}}} \rho_k^n e^{\sqrt{-1} \omega_k n \Delta t} \phi_k \equiv \sum_{k=1}^{N_{\text{md}}} \lambda_k^n \phi_k, \tag{4}$$

with N_{md} is the number of modes retained for the approximation and t_n given as $t_n = n \Delta t, n \in \mathbb{N}$. The dataset K can then be approximated in terms of spatial modes $\{\phi_k\}_{k=1}^{N_{\text{md}}} \in \mathbb{C}^{n_p}$ and temporal coefficients $\{\lambda_k\}_{k=1}^{N_{\text{md}}} \in \mathbb{C}$. From Eq. (4), the set of snapshots K may be expressed as

$$K = M \Lambda + R_{\text{es}}, \tag{5}$$

where $R_{\text{es}} \in \mathbb{C}^{n_p \times N}$ is a residual, $M \in \mathbb{C}^{n_p \times N_{\text{md}}}$ contains the N_{md} dominant spatial modes,

$$M := (\phi_1 \dots \phi_{N_{\text{md}}}), \tag{6}$$

and $\Lambda \in \mathbb{C}^{N_{\text{md}} \times N}$ is a pseudo-Vandermonde matrix containing the associated N_{md} temporal coefficients,

$$\Lambda := \begin{pmatrix} \lambda_1^1 & \lambda_1^2 & \dots & \lambda_1^N \\ \lambda_2^1 & \lambda_2^2 & \dots & \lambda_2^N \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{N_{\text{md}}}^1 & \lambda_{N_{\text{md}}}^2 & \dots & \lambda_{N_{\text{md}}}^N \end{pmatrix}. \tag{7}$$

B. Non-uniform sampling

One can relax the constraint on the uniform sampling in time and account for snapshots sampled at arbitrary times. When dealing with a non-uniform sampling, the standard DMD algorithm fails since Eq. (1) does then not hold. However, for an arbitrary time $t_n \in \mathbb{R}$, the approximation (4) writes

$$\mathbf{u}_{t_n} \approx \lambda_1^{t_n} \phi_1 + \lambda_2^{t_n} \phi_2 + \dots + \lambda_{N_{\text{md}}}^{t_n} \phi_{N_{\text{md}}}, \quad \forall t_n \in \mathbb{R}, \tag{8}$$

and the, now alternant, matrix Λ now reads

$$\Lambda = \begin{pmatrix} \lambda_1^{t_1} & \lambda_1^{t_2} & \dots & \lambda_1^{t_N} \\ \lambda_2^{t_1} & \lambda_2^{t_2} & \dots & \lambda_2^{t_N} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{N_{\text{md}}}^{t_1} & \lambda_{N_{\text{md}}}^{t_2} & \dots & \lambda_{N_{\text{md}}}^{t_N} \end{pmatrix}, \tag{9}$$

where $\{t_1, t_2, \dots, t_N\} \in \mathbb{R}^N$ is arbitrary. Without loss of generality, we set the time reference to t_1 so that $t_{\text{ref}} \leftarrow t_1, t_n \leftarrow t_n - t_{\text{ref}}, 1 \leq n \leq N$.

The approximation of the dynamics of the system can then be derived from matrices Λ and M . Provided the norm, in the Frobenius sense, of the residual R_{es} is small, M can be well approximated from Eq. (5) by

$$M \approx K \Lambda^+. \tag{10}$$

Substituting M from (10) in (5) leads to $K \approx K \Lambda^+ \Lambda + R_{es}$, with Λ^+ the Moore-Penrose pseudo-inverse of Λ . Λ is then determined by minimizing the Frobenius norm of the residual R_{es} ,

$$R_{es} \approx K (I_N - \Lambda^+ \Lambda), \tag{11}$$

and spatial modes $\{\phi_k\}_{k=1}^{N_{\text{md}}} \in \mathbb{C}^{n_p}$ follow immediately from Eq. (10).

The temporal information of the N_{md} dominant modes is entirely given by $\lambda := (\lambda_1 \dots \lambda_{N_{\text{md}}}) \in \mathbb{C}^{N_{\text{md}}}$ and then requires no more than $2 N_{\text{md}}$ pieces of information to be estimated. A very low number N of snapshots is then sufficient in principle.

Since a low number of modes is supposed dominant, the temporal spectrum of the system at hand is sparse in a Fourier basis (few dominant frequencies) and can then be well approximated from a limited number of snapshots via a residual norm minimization under a cardinality constraint as justified by the compressed sensing theory for the linear situation.^{1,2,5} In particular, the Shannon-Nyquist sampling limitations do not apply in this context and an observable with a wide-banded, but sparse or even just compressible, spectrum can be effectively retrieved via this technique as the provided examples will demonstrate in Sec. IV. Instead of a penalty on the effective cardinality of the set of dominant DMD modes, we here directly choose the number of approximation modes N_{md} . It results in an optimization problem such as

$$\lambda \in \arg \min_{\hat{\lambda} \in \mathbb{C}^{N_{\text{md}}}} \left\| K \left(I_N - \Lambda(\hat{\lambda})^+ \Lambda(\hat{\lambda}) \right) \right\|_F, \tag{12}$$

with entries of Λ computed from λ , cf. Eq. (9). The modes are then given by $M = K \Lambda^+$. Notice that, since $\Lambda \in \mathbb{C}^{N_{\text{md}} \times N}$, the matrix $\Lambda^+ \Lambda$ is a square, rank- N_{md} , matrix of size $N \times N$. Since $N_{\text{md}} \ll N$, $\Lambda^+ \Lambda$ is highly rank-deficient and cannot get close to I_N in the sense of $\|\cdot\|_F$ for any λ . Upon evaluation of the economy-size QR decomposition of the snapshot matrix, $QR \stackrel{\text{QR}}{=} K$, $Q \in \mathbb{R}^{n_p \times N}$ and $R \in \mathbb{R}^{N \times N}$, and since left-multiplication of a matrix by a unit-normed column matrix does not change its Frobenius norm, the optimization problem (12) finally reduces to

$$\lambda \in \arg \min_{\hat{\lambda} \in \mathbb{C}^{N_{\text{md}}}} \left\| R \left(I_N - \Lambda(\hat{\lambda})^+ \Lambda(\hat{\lambda}) \right) \right\|_F, \tag{13}$$

which only involves matrices of the size of the number of snapshots N , rather than the size n_p of the snapshots themselves.

This approach allows to rely on scarce data, requiring only a low number of snapshots N and hence resulting in tractable matrices R , I_N , and Λ . Since N_{md} is low, the optimization problem is low-dimensional and can be efficiently solved, either in $\mathbb{R}^{2 N_{\text{md}}}$ with algorithms possibly as simple as the derivative-free Nelder-Mead method,¹² used in the present work or directly in $\mathbb{C}^{N_{\text{md}}}$, see Ref. 25. The approach described in this section is hereafter referred to as the Non-Uniform DMD (NU-DMD) algorithm.

C. Large dataset

1. General comments

Similarly to datasets with a large number of snapshots, performing a standard DMD analysis on a dataset with large snapshot vectors such as those from, say, 3-D flow fields, is computationally challenging. In the typical case where $n_p \gg N$, the overall computational cost of the SVD, as required by the standard DMD algorithm, is $O(n_p N^2)$, which may be prohibitive in case of 3-D space and time-resolved datasets from direct numerical simulations. Similarly, the amount of memory needed for the SVD may limit the use of standard DMD.

Further, while the QR decomposition of the Krylov matrix K makes optimization problem (13) independent of the size of the snapshots, and hence potentially suitable for large fields, an $n_p \times N$ -matrix still needs to be stored in memory to achieve the QR decomposition, even using incremental algorithms.

The Takens' theorem²⁸ suggests that the time series of a well observable quantity may contain the full dynamics of the underlying dynamical system. Since the spatial support associated with a given time evolution pattern is often large, as shown by Robinson,¹⁷ a space-decimated observable vector may then be a relevant approximation to capture the temporal features of the flow. As already discussed in Secs. III A and III B, the dominant features of the flow are often described by a few degrees of freedom and scarce data may suffice to accurately approximate them. This idea is further supported by the fact that the Koopman operator—which has been shown to be related to the DMD in the sense that the DMD algorithm results in the estimation of the eigenvalues of the Koopman operator associated with the dynamics^{14,19}—is defined only by the dynamical flow and should therefore be independent from the choice of the observable. Changing the observable does then not change the eigenvalues of the evolution operator A .

We now further improve upon the NU-DMD algorithm introduced in Sec. III B. Our strategy is to decimate the original dataset K by selecting a low number \tilde{n}_p of components, $\tilde{n}_p < n_p$. Consequently, space-decimated observables $\tilde{u}_n \in \mathbb{R}^{\tilde{n}_p}$ are constructed, and a new, smaller, dataset $\tilde{K}_{\tilde{n}_p} \in \mathbb{R}^{\tilde{n}_p \times N}$ follows. A DMD, or NU-DMD, analysis is then carried-out on the resulting small dataset. The NU-DMD optimization problem (Eq. (13)) now formulates as³¹

$$\lambda \in \arg \min_{\tilde{\lambda} \in \mathbb{C}^{N_{\text{md}}}} \left\| \tilde{K}_{\tilde{n}_p} (I_N - \Lambda^+ \Lambda) \right\|_F, \quad (14)$$

$$M = K \Lambda^+. \quad (15)$$

Since the optimization problem (14) involves small matrices (both N and \tilde{n}_p are small), this algorithm allows the computation of a few dominant modes with low computational requirements.

2. Subset selection

The selection of the spatial components of K retained in the space-decimated $\tilde{K}_{\tilde{n}_p}$ must be such that their cardinality \tilde{n}_p is low, allowing for a computationally efficient optimization problem (14), while retaining relevant temporal information so that the frequency and growth rates of the dominant modes can be accurately estimated. The subset selection problem can then formulate as finding the set of indices $\mathcal{J} = \{j_1, j_2, \dots, j_{\tilde{n}_p}\}$, $1 \leq j_i \in \mathbb{N} \leq n_p$, so that λ derived from $\tilde{K}_{\tilde{n}_p}$ via (14) is a good approximation of λ derived from K via (12). The idea is to cluster spatial components whose dynamics are similar. One must then determine the dominant Fourier spectrum of the discrete time series of components of the observable,

$$\mathbf{u}^{(i)}(t) = \left(u^{(i)}(t_1) u^{(i)}(t_2) \dots u^{(i)}(t_N) \right)^T \in \mathbb{R}^N, \quad (16)$$

with $1 \leq i \leq n_p$. In case the sampling is non-uniform in time, as considered in Sec. III B, a discrete Fourier transform cannot be applied. One then resorts to a minimization problem to determine the N_f dominant Fourier modes that best approximate the time-series in the form $u^{(i)}(t) \approx \sum_k^{N_a} a_k^{(i)} \exp(\sqrt{-1} \omega_k t)$, $\forall 1 \leq i \leq n_p$,

$$\begin{aligned} \mathbf{a}^{(i)} \in \arg \min_{\widehat{\mathbf{a}}^{(i)} \in \mathbb{C}^{N_a}} \left\| \mathbf{u}^{(i)}(t) - \mathfrak{R} \left[\Psi \widehat{\mathbf{a}}^{(i)} \right] \right\|_2, \\ \text{s.t. } \text{card} \left(\widehat{\mathbf{a}}^{(i)} \right) = N_f, \end{aligned} \quad (17)$$

with $N_f < \min \{N_a, N\}$ typically chosen as $N_f = N_{\text{md}}$, N_a the cardinality of the retained Fourier basis, $\mathbf{a}^{(i)} := (a_1^{(i)} \dots a_{N_a}^{(i)})^T$, $\Psi \in \mathbb{C}^{N \times N_a}$, and $\Psi_{n,k} := \exp(\sqrt{-1} \omega_k t_n)$ the sensing matrix. This optimization problem may be solved with an orthogonal matching pursuit algorithm, see for instance,²⁹ or its more sophisticated variants. Once the \tilde{n}_p dominant spectra $\{\mathbf{a}^{(j_i)}\}_{i=1}^{\tilde{n}_p}$ are estimated (possibly in parallel since the n_p problems (17) are non-coupled), a clustering algorithm can be employed to classify and cluster the collection $\{\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \dots, \mathbf{a}^{(n_p)}\}$ in $\{\mathbf{a}^{(j_1)}, \mathbf{a}^{(j_2)}, \dots, \mathbf{a}^{(j_{\tilde{n}_p})}\}$. The retained components are hence given by the coefficient vectors closest to their cluster centroid. In this work, the standard K-means algorithm is used to this aim.⁸

3. A suboptimal relaxation

The above approach allows to rigorously select a subset of spatial components of the snapshots such that the dominant temporal dynamics of the original dataset K is best preserved. It handles situations where the sampling is arbitrary in time and where a discrete Fourier transform can then not be applied. However, while the resulting compressed-sensing-based problems (17) are non-coupled, such an approach may be computationally intensive in cases the time series are long (large N). A fast, while still accurate, alternative is hence proposed.

Frequencies denote the typical number of oscillating events that statistically occur during a given time. The identification of frequencies of interest from complex (e.g., noisy) time series, by mean of statistical techniques, has been intensely discussed in the signal processing literature.^{15,26,30} We propose to cluster spatial components with similar statistical features in time. The time series of a given spatial component is treated as realizations of a random variable and its dynamics are characterized by the corresponding probability density function (pdf).^{6,24} By clustering spatial components with similar pdfs, components with similar dynamics, in a statistical sense, are identified. This idea is further supported by the fact that the DMD is linked to the Koopman operator, originally introduced to quantify statistical properties of dynamical systems.^{4,13}

A large amount of realizations (i.e., snapshots) is needed to estimate the pdf while computing the first \tilde{N}_m central moments, $\tilde{N}_m \ll N$, only requires limited information and can be estimated with embarrassingly parallelizable computations. The strategy we then adopt is to cluster spatial components with respect to their estimated statistical moments. Two criteria assess the relevance of the identification. For \tilde{N}_m large enough, there is a unique pdf associated with the moments of the time series if Krein's condition holds.²⁷ Alternatively, one can estimate Pisarenko's dominant frequencies,¹⁵ and compare with the clustering resulting from the classification step.

This clustering scheme allows to seamlessly handle the case of non-uniform sampling in time. Note that it is suboptimal in the sense that two time-series exhibiting the same Fourier spectrum may have different first statistical moments, potentially preventing them to be assigned to the same cluster. However, the above criteria bring some guarantees that two series of different dynamics will not be assigned to the same cluster.

IV. RESULTS

We now illustrate the above methods with two situations. In Sec. IV A, the NU-DMD algorithm is applied to a synthetic velocity field used in Ref. 7, representing a 1-D, linear instability. This field allows a precise control on frequencies and growth rates and is hence suitable for a detailed assessment of our method. In Sec. IV B, the NU-DMD algorithm is applied to experimental and numerical cases. In Sec. IV B 1 (respectively IV B 2), the NU-DMD algorithm is tested on two-dimensional, experimental (respectively three-dimensional, numerical) velocity fields of an open cavity flow. To allow comparison with a standard DMD, snapshots are sampled uniformly in time but are randomly reordered to mimic a non-uniform sampling for NU-DMD.

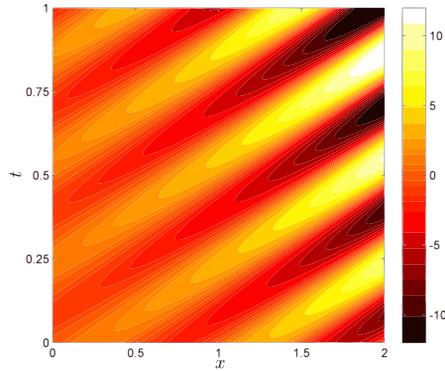


FIG. 1. Spatial-temporal diagram of the synthetic velocity field. Colors indicate the value of the velocity.

A. Synthetic system

A synthetic field is used as a benchmark for assessing the performance of the NU-DMD algorithm. The chosen synthetic field, introduced in Duke *et al.*,⁷ is a model for a 1-D, linear instability and allows full control on the frequency $f \equiv \omega/2\pi$ and on the spatial γ and temporal σ growth rates of the resulting dynamics,

$$u(x, t) = u_0(1 + \xi) \sin(2\pi\kappa x - \omega t) \exp(\sigma t + \gamma x). \quad (18)$$

The observable u is made of n_p points uniformly sampled in space in $[0, 2]$, with a step δx . The dataset consists of N realizations of the observable, sampled every $\delta t = 1/(N - 1)$ within $[0, 1]$, i.e., $t_n := (n - 1)\delta t$. The (x, t) domain is discretized on a $n_p \times N$ grid. The pulsation ω is 20, i.e., the frequency $f = \omega/2\pi \approx 3.18$. The wavenumber κ , the initial amplitude u_0 , and spatial growth rate γ are all unitary. The temporal growth rate is set to $\sigma = 0.75$. The white, uniformly distributed, multiplicative noise $\xi \sim U([-1, 1])$ is introduced to mimic an actual observable from an experiment. The number of components of the velocity field is fixed to $n_p = 2001$, and the number of snapshots is $N = 2000$. The synthetic flow is represented in Fig. 1 in a (x, t) -diagram.

We set the Noise to Signal Ratio (NSR = $\max|\xi|$) to 5% and randomly select N_r (ranging from 10 to N) snapshots from the original dataset (with no specific order). The subset selection step is then applied to this dataset, in order to keep \tilde{n}_p (ranging from 10 to n_p) spatial points. We compare the results of our NU-DMD algorithm with a standard DMD applied to the N_r first snapshots of the original dataset, and to the same $n_p^{\text{DMD}} = \tilde{n}_p$ spatial points as taken for the NU-DMD algorithm. As can be appreciated from Figs. 2 and 3, the frequency $f_{\text{NU-DMD}}$ and the temporal growth $\sigma_{\text{NU-DMD}}$ identified by the NU-DMD are usually significantly more accurate. For small datasets ($\tilde{n}_p/n_p \ll 1$ and/or $N_r/N \ll 1$), the relative error in the identified f and σ by NU-DMD may be several orders of magnitude lower than that computed with the standard DMD algorithm. For some degenerated cases, i.e., when $\tilde{n}_p < N_r$ or when N_r is too small, the DMD algorithm fails to identify

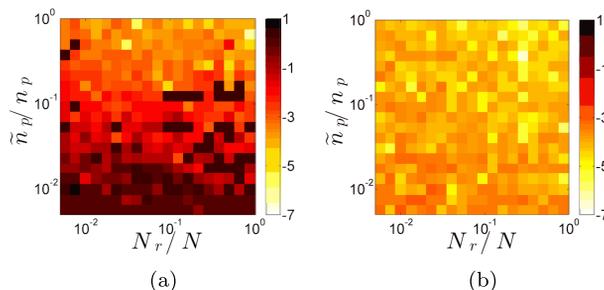


FIG. 2. Relative error $\epsilon_f = |f - \hat{f}|/f$ on the identification of the frequency f , in a \log_{10} scale, of the dominant mode with DMD (left) and NU-DMD (right), both relying on a $\tilde{n}_p \times N_r$ dataset.

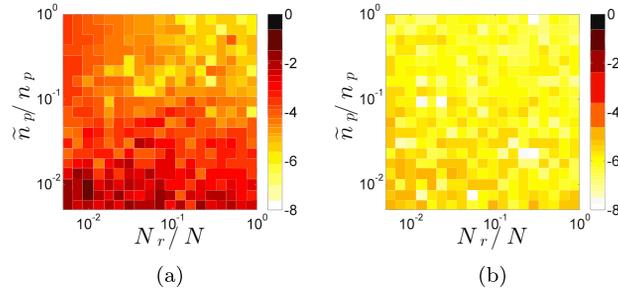


FIG. 3. Relative error $\epsilon_\sigma = |\sigma - \hat{\sigma}|/\sigma$ on the identification of the growth rate σ , in a \log_{10} scale, of the dominant mode with DMD (left) and NU-DMD (right), both relying on a $\tilde{n}_p \times N_r$ dataset.

frequencies and growth rates while the NU-DMD captures the right frequency and growth rate with as low as $N_r = 10$ snapshots and relying on $\tilde{n}_p = 10$ spatial points only.

Once the dominant frequency is identified, the corresponding spatial mode is reconstructed via Eq. (10). The mode associated with the main frequency is recovered with a good accuracy even when relying on a drastically reduced subset, see Fig. 4 where it is plotted as estimated from datasets of different sizes.

The stability of both the DMD and the NU-DMD algorithm under noise corrupting the data is also investigated. The accuracy of the identified dominant frequency when the noise to signal ratio varies is given in Table I. The NU-DMD appears to be very robust with respect to noise and more accurate than DMD in the estimation of the main frequency.

B. Application to a cavity flow

The flow over an open cavity exhibits dynamical features of interest for illustrating our method. The most energetic phenomenon is due to self-sustained oscillations of the impinging mixing layer.¹⁸ Such oscillations result from the feedback loop formed, on one hand, by the Kelvin-Helmholtz vortices convected downstream to the cavity trailing corner, and, on the other hand, by the instantaneous feedback from the pressure field, from trailing to leading edge, induced by the impingement of vortices. As a result, power spectra are organized around a few narrow peaks (most often a single peak at frequency f), with Strouhal numbers fL/U_∞ close to multiples of $1/2$, where L is the cavity length and U_∞ is the incoming velocity.

1. Two-dimensional experimental flow

We first consider an experimental incompressible open cavity flow with aspect ratio $L/H = 1.5$ (depth $H = 50$ mm), Reynolds number $\text{Re} = UL/\nu \simeq 8800$, and span-wise aspect ratio of 6. A typical power spectrum, in the configuration of the flow, is shown in Fig. 5(b). PIV measurements

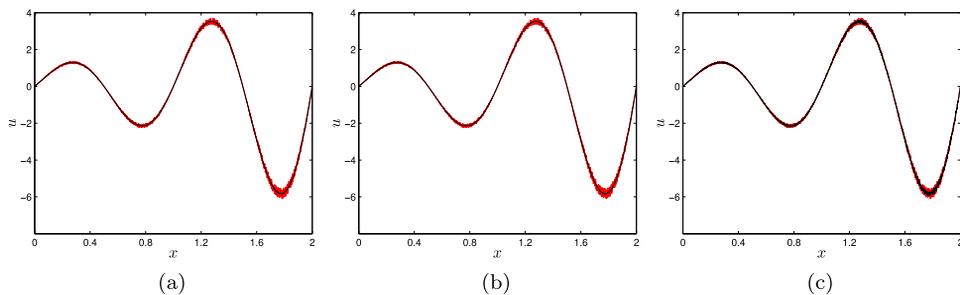


FIG. 4. In black, real part of the dominant spatial mode identified by (a) a standard DMD of the whole dataset. (b) NU-DMD algorithm, with $N_r/N = 5.0 \times 10^{-2}$ and $\tilde{n}_p/n_p = 5.0 \times 10^{-2}$. (c) NU-DMD, $N_r/N = 5.0 \times 10^{-4}$, $\tilde{n}_p/n_p = 5.0 \times 10^{-4}$. Spatial modes are in phase with the velocity field at $t = 0$, in red for comparison.

TABLE I. Error $\epsilon_f = |f - \hat{f}|/f$ on the identification of the frequency with a standard DMD applied to the N_r first snapshots of the original dataset, and with a NU-DMD applied to N_r randomly selected snapshots, with respect to the noise to signal ratio.

NSR	5.0×10^{-3}		5.0×10^{-2}		5.0×10^{-1}	
	DMD	NU-DMD	DMD	NU-DMD	DMD	NU-DMD
$N_r/N = 0.01$	5.2×10^{-4}	5.9×10^{-5}	1.3×10^{-2}	7.3×10^{-4}	2.7×10^{-1}	4.7×10^{-3}
$N_r/N = 0.1$	3.6×10^{-4}	2.3×10^{-5}	2.1×10^{-3}	8.8×10^{-5}	3.8×10^{-1}	2.7×10^{-3}
$N_r/N = 1.0$	5.8×10^{-5}	3.1×10^{-6}	9.1×10^{-4}	7.7×10^{-5}	4.1×10^{-2}	5.8×10^{-4}

in a vertical (x, y) -plane provide time-resolved 2-D, 2-C velocity fields (see Fig. 5(a) for a visualization of the flow in the observation plane). The dataset is formed of $N = 5000$ fields with a spatial grid of $n_p = 7018$ points.

To illustrate the NU-DMD algorithm, N_r snapshots taken at random from the full PIV dataset, hence at random times, are considered. The algorithm is applied to the fluctuating part (zero time-averaged) of this N_r -snapshot dataset. The influence of the number of spatial components \tilde{n}_p used in the NU-DMD algorithm is also investigated. The spatial mode associated with the frequency of the shear layer from both a standard DMD (N samples uniform in time) and NU-DMD (with varying number of snapshots and retained spatial components) is shown in Fig. 6. When the snapshots are not sampled uniformly in time, the DMD method cannot be applied while the NU-DMD still correctly identifies the flow field dominant modes. Further, very few snapshots are necessary for the dominant features to be accurately identified with this technique since retaining only one snapshot among 500 ($N_r/N = 2.0 \times 10^{-3}$ with $N_r = 10$) and one spatial point over 700 ($\tilde{n}_p/n_p = 1.4 \times 10^{-3}$ with $\tilde{n}_p = 10$) still results in a decent approximation of the dominant shear-layer mode, see Fig. 6(e). In order to quantify the discrepancy, the absolute value of the relative difference between the amplitude of the NU-DMD and DMD modes is plotted Fig. 7 for the dominant mode. The modes have been set in phase and the error ϵ_{mode} has been normalized by the norm of the corresponding component of the DMD mode. The error for the i th component is hence defined, in the general case, as

$$\epsilon_{\text{mode}_i}(x, y, z) := \frac{|\phi_{i,\text{NU-DMD}}(x, y, z) - \phi_{i,\text{DMD}}(x, y, z)|}{\|\phi_{i,\text{DMD}}\|_2}, \quad i \in \{x, y, z\}. \quad (19)$$

The NU-DMD spatial modes are given by the pseudo-inverse solution Eq. (14) and spatial fields in the null-space of Λ can then not be resolved. Further, a small frequency difference δf between, say, the exact dominant frequency and that estimated via NU-DMD will affect the spatial modes M by a contribution associated with that low frequency δf via Eq. (14). It results that the discrepancy is most visible where the energy of the flow at these low frequencies is large. However, it is seen that it remains below a few percents throughout the spatial domain.

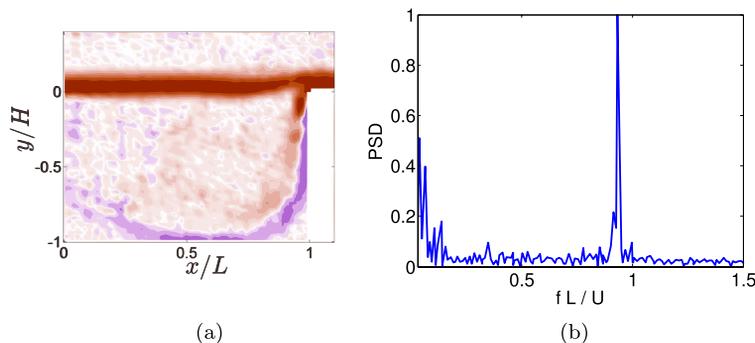


FIG. 5. (a) Snapshot of the flow in the observation plane (vorticity field). (b) Power spectral density of the cavity flow produced by averaging the individual spectra at each spatial points of the velocity field.

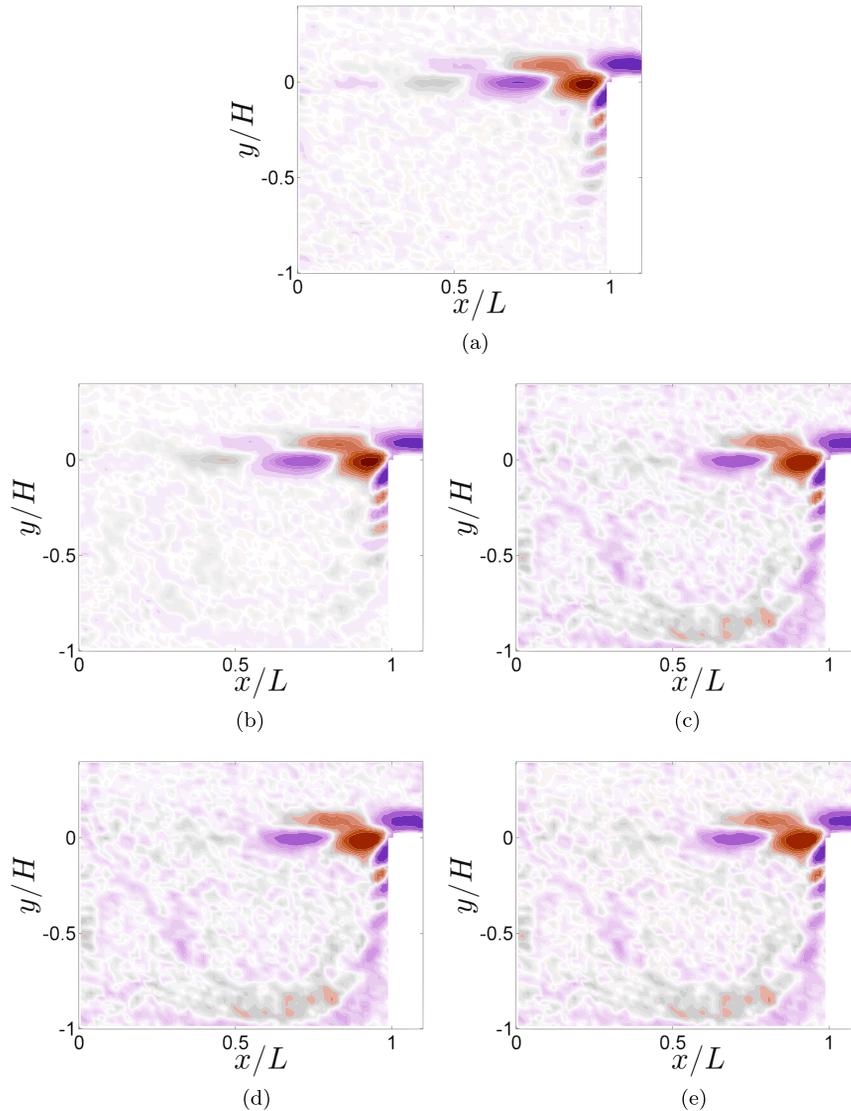


FIG. 6. Dominant mode identified by (a) standard DMD (N snapshots of n_p components). (b) NU-DMD algorithm, with $N_r/N = 1.0 \times 10^{-1}$ and $\tilde{n}_p/n_p = 1.4 \times 10^{-1}$. (c) NU-DMD, $N_r/N = 1.0 \times 10^{-1}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-3}$. (d) NU-DMD, $N_r/N = 2.0 \times 10^{-3}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-1}$. (e) NU-DMD, $N_r/N = 2.0 \times 10^{-3}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-3}$. Colors indicate the value of the vorticity.

The retained spatial points, with their associated cluster supports, are plotted in Fig. 8 for $\tilde{n}_p = 10$, and in Fig. 9 for $\tilde{n}_p = 20$. As expected, physically relevant areas are discriminated, i.e., the shear-layer, the incoming flow, the flow along the downstream wall, and the inner-flow recirculation. Fourier spectra associated with centroids are, by construction, different, see Fig. 10 where they are plotted for some clusters. The supports of the identified clusters correspond to regions of space where the relative influence of the shear layer (high frequency) and the inner-flow (low frequencies) varies.

The clustering step is robust in the sense that the supports of the identified clusters for $\tilde{n}_p = 10$, Fig. 8, foliate into smaller and nested supports when $\tilde{n}_p = 20$, Fig. 9. In particular, the topology of the clusters is comparable and no significant new region arises in the clusters support when classifying with a finer description.

For validation purpose, the supports of clustered points identified using spectra directly estimated from a Fourier analysis of snapshots sampled uniformly in time, hence serving as the reference, are plotted in Fig. 11. They are seen to be rather similar to the clusters found using the

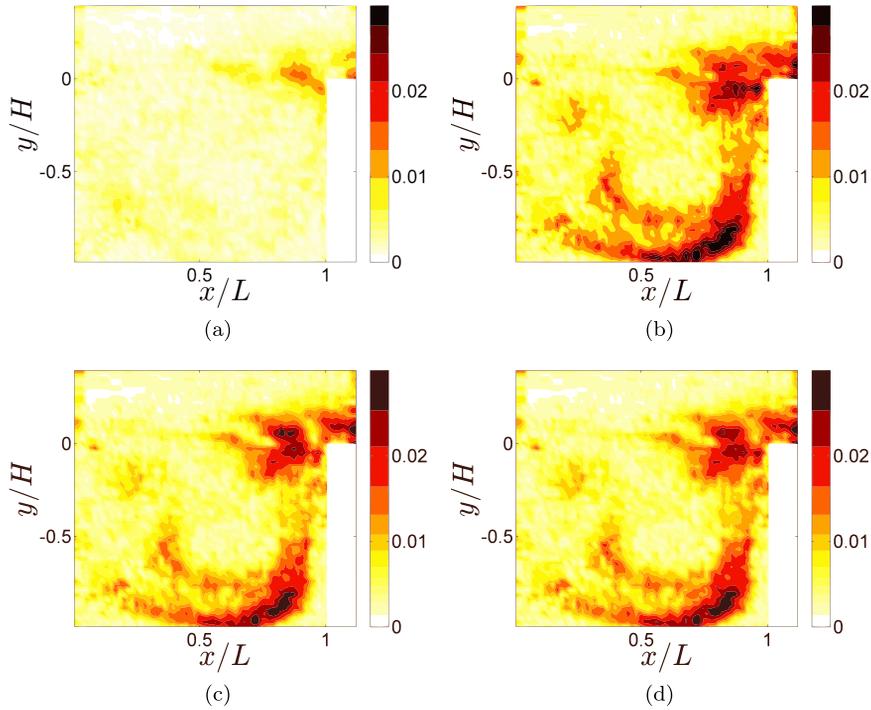


FIG. 7. Relative difference ϵ_{mode} , see Eq. (19), between the dominant mode identified by a standard DMD and the NU-DMD algorithm. (a) NU-DMD algorithm, with $N_r/N = 1.0 \times 10^{-1}$ and $\tilde{n}_p/n_p = 1.4 \times 10^{-1}$. (b) NU-DMD, $N_r/N = 1.0 \times 10^{-1}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-3}$. (c) NU-DMD, $N_r/N = 2.0 \times 10^{-3}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-1}$. (d) NU-DMD, $N_r/N = 2.0 \times 10^{-3}$, $\tilde{n}_p/n_p = 1.4 \times 10^{-3}$.

suboptimal method presented in Sec. III C 3, see Fig. 8, hence justifying the proposed suboptimal clustering step scheme as a suitable and relevant proxy for an accurate clustering, robust with respect to the sampling. Our suboptimal method took about 60 s of central processing unit (CPU) time to classify the whole flow field while the minimization of the n_p cost functions in Eq. (17) of a compressed sensing approach would require a very significant computational effort.

The accuracy of the estimation of the dominant frequency f by the NU-DMD algorithm when varying the amount of available information is given in Table II in terms of relative error $\epsilon_f = |f_{\text{NU-DMD}} - f_{\text{DMD}}| / f_{\text{DMD}}$. The quality of the identification of spatial modes ϕ can be estimated from the quantification of their main feature, i.e., the amplified travelling waves in the shear layer.

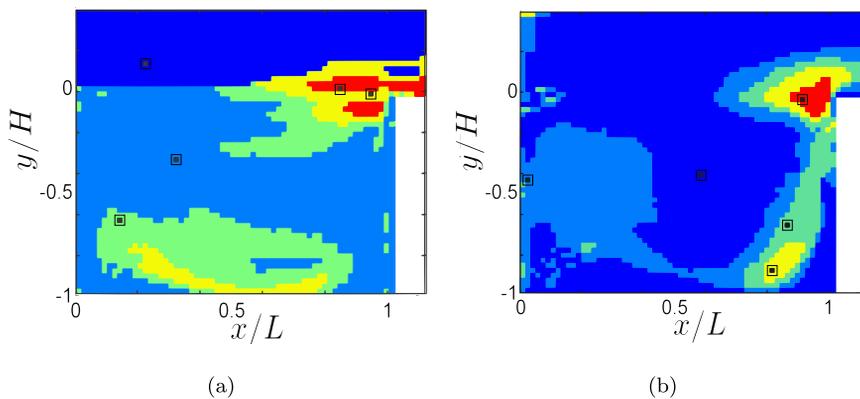


FIG. 8. Illustration of the clustering step proposed in Sec. III C 3. $\tilde{n}_p = 10$. Black dots indicate spatial components associated to centroids selected during the clustering process. Colors indicate the different clusters. (a) x -velocity component. (b) y -velocity component.

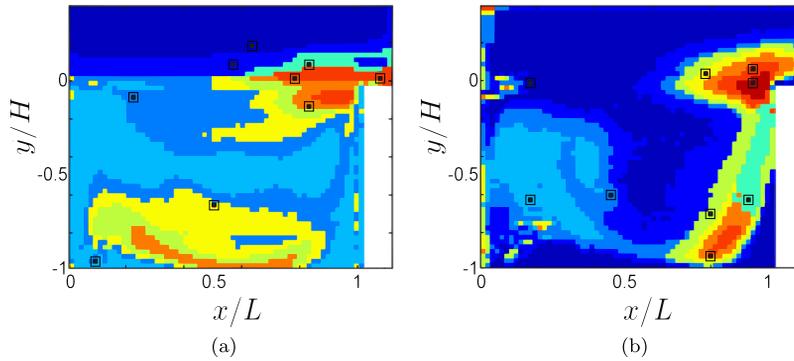


FIG. 9. Illustration of the clustering step proposed in Sec. III C 3. $\tilde{n}_p = 20$. Black dots indicate spatial components associated to centroids selected during the clustering process. Colors indicate the different clusters. (a) x -velocity component. (b) y -velocity component.

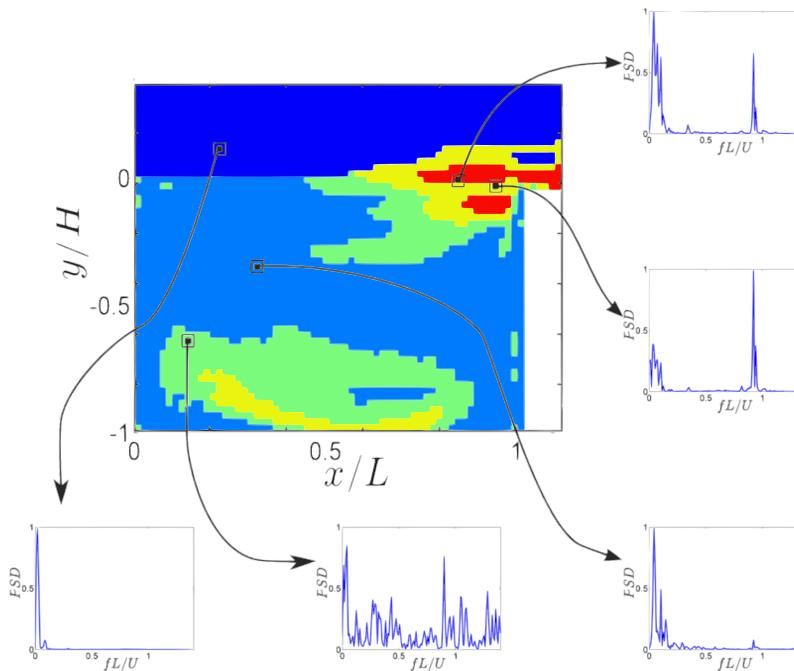


FIG. 10. x -velocity component power spectral densities associated with the cluster centroids. $\tilde{n}_p = 10$.

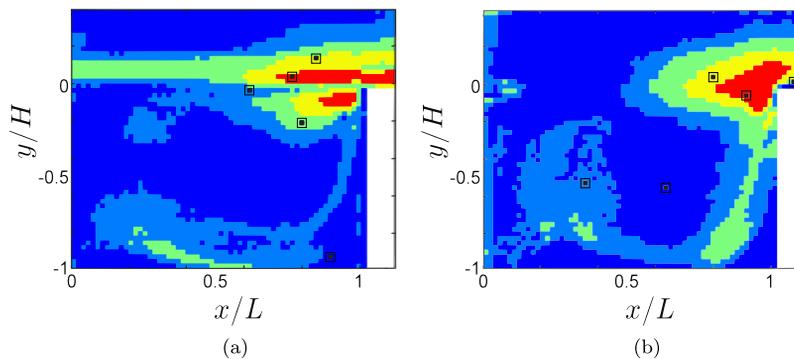


FIG. 11. Fourier-based subset selection. Illustration of the clustering step. $\tilde{n}_p = 10$. Black dots indicate spatial components associated to centroids selected during the clustering process. Colors indicate the different clusters. (a) x -velocity component. (b) y -velocity component.

TABLE II. Accuracy of the dominant frequency and the wavenumber identification of the shear layer mode by the NU-DMD algorithm. $\epsilon_f := |f_{\text{NU-DMD}} - f_{\text{DMD}}| / f_{\text{DMD}}$, $\epsilon_\kappa := |\kappa_{\text{NU-DMD}}^{-1} - \kappa_{\text{DMD}}^{-1}| / \kappa_{\text{DMD}}^{-1}$. The DMD is here computed with the maximum information (N snapshots of n_p components) and hence serves as reference.

	$\tilde{n}_p / n_p = 0.14$		$\tilde{n}_p / n_p = 0.0014$	
	ϵ_f	ϵ_κ	ϵ_f	ϵ_κ
$N_r / N = 0.1$	8.0×10^{-4}	8.0×10^{-3}	3.2×10^{-3}	4.0×10^{-2}
$N_r / N = 0.002$	6.9×10^{-4}	4.0×10^{-2}	6.9×10^{-4}	4.1×10^{-2}

The main characteristic of these impacting vortices is their wavenumber κ , which can be determined by fitting the velocity profile $\mathbf{u}(x, y)$, at some given $y = y_0$, with the theoretical profile of a spatially developing instability,

$$\mathbf{u}_{th}(x, y_0) = \mathbf{v}_0 + \mathbf{v}_1 \exp(\sigma x) \cos(\kappa x + \varphi), \quad (20)$$

where σ is the spatial growth, φ is a reference phase, and $\{\mathbf{v}_0, \mathbf{v}_1\}$ are normalization factors. It results in identification of the wavenumber by the NU-DMD algorithm with an error $\epsilon_\kappa = |\kappa_{\text{NU-DMD}}^{-1} - \kappa_{\text{DMD}}^{-1}| / \kappa_{\text{DMD}}^{-1}$ of about 4%, see Table II.

2. Three-dimensional numerical simulation

We now consider the numerical simulation of a three-dimensional incompressible open cavity flow with aspect ratio $L/H = 2$ (depth $H = 50$ mm), Reynolds number $\text{Re} = UL/\nu \simeq 4000$ and span-wise aspect ratio of 2, with periodic spanwise boundary conditions. The numerical simulation provides time-resolved three-dimensional (3-D) three-component (3-C) velocity fields (see Fig. 12 for a snapshot of the flow in the symmetry (x, y) observation plane. For the sake of readability, colors have been saturated). The dataset is formed of $N = 500$ snapshots of the velocity field, with a spatial grid of $n_p = 2033$ 130 points. Details about the numerical simulation can be found in Ref. 16.

The spatially averaged velocity field power spectrum is shown in Fig. 13. It exhibits a series of dominant features around a reduced frequency of about 1. In order to assess the applicability and accuracy of the NU-DMD algorithm, $N_r = 50$ ($N_r/N = 0.1$) snapshots of the velocity field are taken at random from the full DNS dataset, hence at random times. The algorithm is applied to the fluctuating part (zero time-averaged) of this N_r -snapshot dataset. The number of spatial points kept for the clustering step is $\tilde{n}_p = 100$ ($\tilde{n}_p/n_p = 4.9 \times 10^{-5}$).

The main computational bottleneck of the NU-DMD is the identification of the clusters since the classification step in the K-means algorithm (see Sec. III C 2) takes approximately 40 000 s of total CPU time on a desktop computer, and about 4 GB of memory. While the minimization of the n_p cost functions in Eq. (17) in a compressed sensing approach would constitute an intractable computational effort, the evaluation of the central moments in the suboptimal approach proposed in

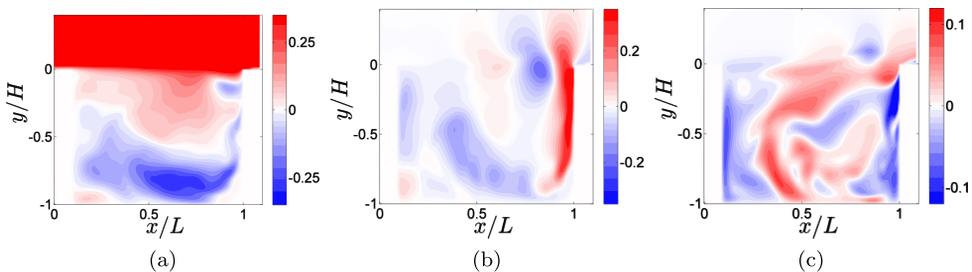


FIG. 12. Snapshot of the flow in the symmetry plane. Colors encode the velocity field magnitude. (a) x -velocity component. (b) y -velocity component. (c) z -velocity component. Colors have been saturated in figure (a) in order to make the cavity internal flow more visible.

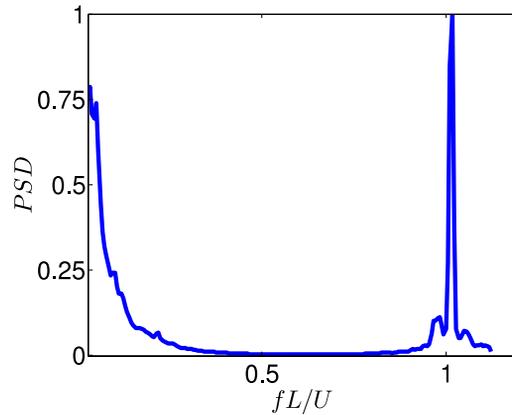


FIG. 13. Spatially averaged power spectrum density of the velocity field, from the DNS of the 3-D cavity flow.

Sec. III C 3 only requires a few seconds. For illustration and comparison purpose, a DMD using the standard method^{19,20} is also computed. Its evaluation requires over 40 GB of memory and more than 320 000 s of total CPU time. It is hence out-of-reach for a desktop computer and had to be computed on an HPC machine.

The supports of the clusters are plotted in Fig. 14 in the (x, y) symmetry plane of the cavity. Clusters of the x - and y -components of the velocity field are reminiscent of those of the two-dimensional case discussed in Sec. IV B 1 (compare with Fig. 8). Both the shear layer and the incoming flow are dynamically well separated (cluster-wise) from the rest of the flow.

Similarly to the two-dimensional case, Sec. IV B 1, dominant DMD modes are associated to oscillations of the shear-layer. The dominant mode, as derived both from the standard DMD algorithm and the present NU-DMD approach, is plotted in Fig. 15 for illustration in terms of isosurfaces of the Q -criterion for $Q/(U/L)^2 \in \{1, 2\}$ (left) and $Q/(U/L)^2 \in \{-1, -2\}$ (right). A very good agreement between the reference, computationally very intensive, DMD method and the NU-DMD approach is observed. A few snapshots, as well as a dramatically reduced number of points in space, are hence sufficient to compute a decent approximation of the dominant DMD mode. The mode associated with the dynamics of the shear layer is plotted in Fig. 16 in a 2-D plane both from the NU-DMD and the DMD approach. The relative error $\epsilon_f := |f_{\text{NU-DMD}} - f_{\text{DMD}}| / f_{\text{DMD}}$ in the identification of the frequency is less than 1.9×10^{-2} while the relative error $\epsilon_\kappa = |\kappa_{\text{DMD}}^{-1} - \kappa_{\text{NU-DMD}}^{-1}| / \kappa_{\text{DMD}}^{-1}$ in the wavenumber of the shear-layer, evaluated via Eq. (20) as in Sec. IV B 1, is less than 4.4×10^{-3} , again illustrating the accuracy of the present method.

In order to quantify the discrepancies between DMD and NU-DMD, the absolute value of the difference between the amplitude of the dominant DMD and NU-DMD mode as defined in Eq. (19) is plotted Fig. 17. The modes have been set in phase and the error has been normalized by the norm of the corresponding component of the DMD mode. Differences remain inferior to a few percents and are mainly observed within the cavity for the z -component. As in the 2-D case, patterns typical of inner-flow structures at low frequency are discernible within the cavity on the spatial mode, as

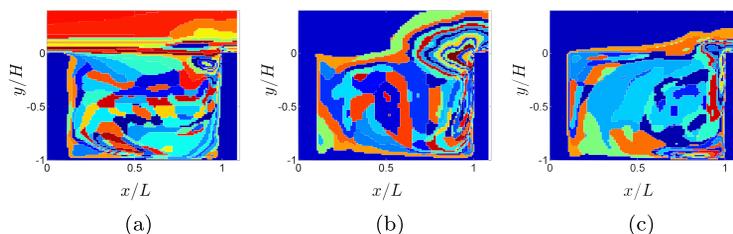


FIG. 14. Illustration of the clustering step introduced in Sec. III C 3. $\bar{n}_p = 100$. Colors indicate the different clusters. (a) x -velocity component. (b) y -velocity component. (c) z -velocity component.

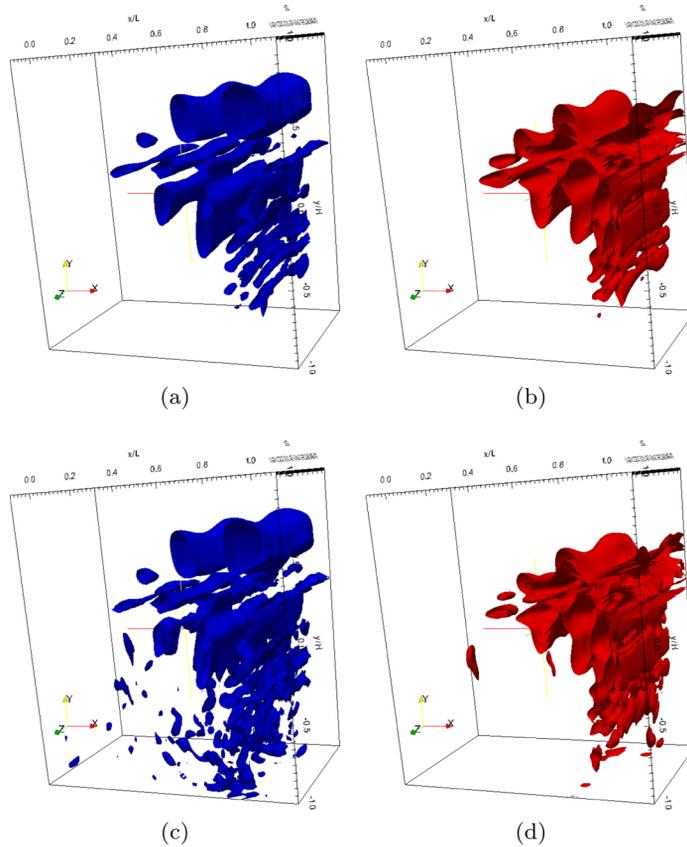


FIG. 15. Isosurfaces of positive (left column) and negative (right column) Q criterion. Top row (respectively, bottom row) corresponds to the DMD (respectively, NU-DMD) dominant shear-layer mode.

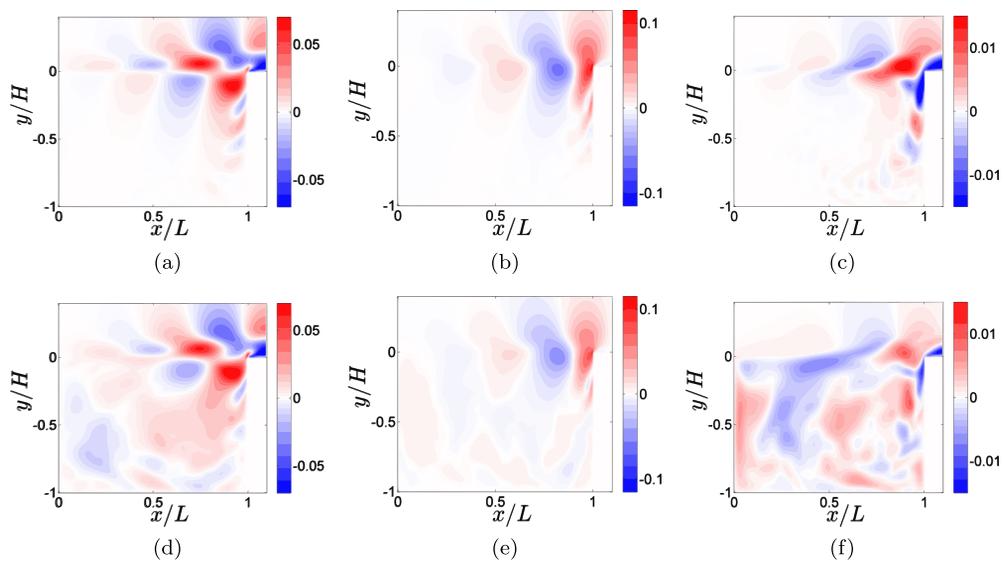


FIG. 16. Slice of the dominant mode, identified by a standard DMD (top line, real part of the mode, N snapshots of n_p components) and the NU-DMD algorithm (bottom line, real part of the mode, with $N_r/N = 0.1$ and $\bar{n}_p/n_p = 4.9 \times 10^{-5}$). First, second and last column correspond to x -, y - and z -velocity components respectively. Colors encode the magnitude of the velocity field.

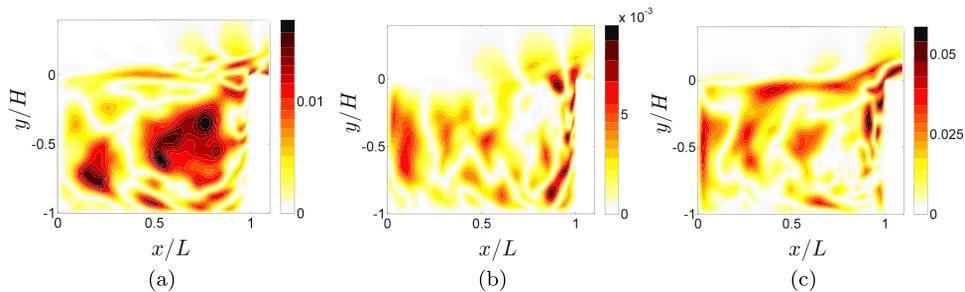


FIG. 17. Slice of relative difference ϵ_{mode} , see Eq. (19), between the dominant mode identified by a standard DMD and the NU-DMD algorithm. Left, middle and right plots correspond to the x -, y -, and z -velocity component respectively.

discussed in Sec. IV B 1. Further, since the z -component of the velocity fields stored in K is small and carries little energy, the affected part of the mode has a larger relative impact compared to that of the x - and y -component. While not done here, a suitable relative normalization of the different components of the snapshots would balance their contribution to the Frobenius norm in Eq. (12) and result in a better relative accuracy of the z -component.

V. CONCLUSION

We have presented a method for efficiently computing dynamic modes in case of datasets of very large size, where direct computations on the resulting Krylov matrix are out-of-reach, and/or from a system arbitrarily sampled in time, where a standard DMD analysis does not apply. The method essentially formulates the problem in an optimization setting and decouples the estimation of the temporal description from the spatial description. This allows to handle very general datasets, with essentially no restriction on the sampling strategy or the size of the snapshots. With an appropriate sampling strategy, both very low and very high frequency contents can be retrieved without the need of a large number of snapshots. Further, spatial correlations in space of the observable are exploited to estimate frequencies and growth rates of the dynamic modes from a smaller dataset obtained via a clustering algorithm based on the estimated Fourier spectrum of each component of the spatial field. While the estimation of the dominant Fourier spectra, expected to be compressible, naturally lends itself to a compressed sensing framework, a more computationally friendly implementation relying on statistics of the time series is here used. The effective dataset is small and the resulting NU-DMD method is computationally efficient.

The method has been illustrated on a synthetic flow, exhibiting a 1-D linear instability, and on the experimental dataset of 2-D PIV fields of a 3-D open cavity flow. A large-scale three-dimensional configuration, involving more than 2×10^6 degrees-of-freedom for each snapshot, was also considered. In all of these cases, spatial modes, frequencies, and growth rates were accurately recovered, even when relying on as low as about 0.1% of the original spatial field and 0.2% of the original snapshots as needed by the standard DMD.

The present NU-DMD method hence constitutes a valuable and widely applicable tool for analyzing physical systems from a dataset of observables with very mild constraints (sampling strategy, missing snapshots, small dataset, etc.).

ACKNOWLEDGMENTS

This work was partially funded by the Digiteo FLUCTUS project.

- ¹ E. Candès and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Found. Comput. Math.* **6**(2), 227–254 (2006).
- ² E. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inf. Theory* **52**(12), 5406–5425 (2006).
- ³ K. Chen, J. Tu, and C. Rowley, "Variants of dynamic mode decomposition: Boundary condition, Koopman, and Fourier analyses," *J. Nonlinear Sci.* **22**, 887–915 (2012).

- ⁴ J. Ding and A. Zhou, *Statistical Properties of Deterministic Systems* (Springer, 2010).
- ⁵ D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006).
- ⁶ T. Dreeben and S. Pope, "Probability density function/monte carlo simulation of near-wall turbulent flows," *J. Fluid Mech.* **357**, 141–166 (1998).
- ⁷ D. Duke, J. Soria, and D. Honnery, "An error analysis of the dynamic mode decomposition," *Exp. Fluids* **52**, 529–542 (2012).
- ⁸ A. Gersho and R. Gray, *Vector Quantization and Signal Compression* (Springer, 1992).
- ⁹ F. Guéniat, L. Pastur, and F. Lusseyran, "Investigating mode competition and three-dimensional features from two-dimensional velocity fields in an open cavity flow by modal decompositions," *Phys. Fluids* **26**, 085101 (2014).
- ¹⁰ F. Guéniat, L. Pastur, and L. Mathelin, "Snapshot-based flow analysis with arbitrary sampling," *Procedia IUTAM* **10**, 59 (2014).
- ¹¹ M. Jovanović, P. Schmid, and J. Nichols, "Sparsity-promoting dynamic mode decomposition," *Phys. Fluids* **26**(2), 024103 (2014).
- ¹² J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder-Mead simplex method in low dimensions," *SIAM J. Optim.* **9**, 112–147 (1998).
- ¹³ I. Mezić, "Spectral properties of dynamical systems, model reduction and decompositions," *Nonlinear Dyn.* **41**, 309–325 (2005).
- ¹⁴ I. Mezić, "Analysis of fluid flows via spectral properties of the Koopman operator," *Annu. Rev. Fluid Mech.* **45**, 357–378 (2013).
- ¹⁵ V. Pisarenko, "The retrieval of harmonics from a covariance function," *Geophys. J. Int.* **33**(3), 347–366 (1973).
- ¹⁶ B. Podvin, Y. Fraigneau, F. Lusseyran, and P. Gougat, "A reconstruction method for the flow past an open cavity," *J. Fluids Eng.* **3**(128), 531–540 (2006).
- ¹⁷ J. Robinson, "A rigorous treatment of 'experimental' observations for the two-dimensional Navier-Stokes equations," *Proc. R. Soc. A* **457**(2008), 1007–1020 (2001).
- ¹⁸ D. Rockwell and E. Naudascher, "Self-sustained oscillations of impinging free shear layers," *Annu. Rev. Fluid Mech.* **11**, 67 (1979).
- ¹⁹ C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *J. Fluid Mech.* **641**, 115–127 (2009).
- ²⁰ P. Schmid, "Dynamic mode decomposition of numerical and experimental data," *J. Fluid Mech.* **656**, 6–28 (2010).
- ²¹ P. Schmid, "Application of the dynamic mode decomposition to experimental data," *Exp. Fluids* **50**(4), 1123–1130 (2011).
- ²² A. Seena and J. Sung, "Dynamic mode decomposition of turbulent cavity flows for self-sustained oscillations," *Int. J. Heat Fluid Flow* **32**(6), 1098–1110 (2011).
- ²³ O. Semeraro, G. Bellani, and F. Lundell, "Analysis of time-resolved PIV measurements of a confined turbulent jet using pod and Koopman modes," *Exp. Fluids* **53**(5), 1203–1220 (2012).
- ²⁴ P. V. Slooten and S. Pope, "Advances in pdf modeling for inhomogeneous turbulent flows," *Phys. Fluids* **10**(1), 246–265 (1998).
- ²⁵ L. Sorber, M. Van Barel, and L. De Lathauwer, "Unconstrained optimization of real functions in complex variables," *SIAM J. Optim.* **22**(3), 879–898 (2012).
- ²⁶ P. Stoica and T. Soderstrom, "Statistical analysis of music and subspace rotation estimates of sinusoidal frequencies," *IEEE Trans. Signal Process.* **39**(8), 1836–1847 (1991).
- ²⁷ J. Stoyanov, "Krein condition in probabilistic moment problems," *Bernoulli* **6**, 939–949 (2000).
- ²⁸ F. Takens, "Detecting strange attractors in turbulence dynamical systems and turbulence," *Lect. Notes Math.* **898**, 366–381 (1981).
- ²⁹ J. Tropp and A. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory* **53**(12), 4655–4666 (2007).
- ³⁰ D. Tufts and R. Kumaresan, "Estimation of frequencies of multiple sinusoids: Making linear prediction perform like maximum likelihood," *Proc. IEEE* **70**(9), 975–989 (1982).
- ³¹ Whenever $N < \tilde{n}_p$, the optimization problem (14) may be advantageously substituted with $\lambda \in \arg \min_{\lambda \in \mathbb{C}^N} \|\tilde{R}_{\tilde{n}_p} (I_N - \Lambda^+ \Lambda)\|_F$, with $\tilde{Q}_{\tilde{n}_p} \tilde{R}_{\tilde{n}_p}^{\text{QR}} \equiv \tilde{K}_{\tilde{n}_p}$.