Eurika Kaiser | Bernd R. Noack | Laurent Cordier
Andreas Spohn | Marc Segond | Markus Abel
Guillaume Daviller | Jan Osth | Sinisa Krajnovic | Robert K. Niven

Cluster-based reduced-order modelling of a mixing layer

Suggested citation referring to the original publication:
Journal of Fluid Mechanics 754 (2014), pp. 365–414
DOI https://doi.org/10.1017/jfm.2014.355
ISSN (print) 0022-1120
ISSN (online) 1469-7645; 1750-6859

J. Fluid Mech. (2014), *vol.* 754, *pp.* 365–414.
doi:10.1017/jfm.2014.355

Cluster-based reduced-order modelling of a mixing layer

Eurika Kaiser[1,†], Bernd R. Noack[1], Laurent Cordier[1], Andreas Spohn[1], Marc Segond[2], Markus Abel[2,3,4], Guillaume Daviller[5], Jan Östh[6], Siniša Krajnović[6] and Robert K. Niven[7]

[1]Institut PPRIME, CNRS – Université de Poitiers – ENSMA, UPR 3346, Département Fluides, Thermique, Combustion, CEAT, 43, rue de l'Aérodrome, F-86036 Poitiers CEDEX, France

[2]Ambrosys GmbH, Albert Einstein Strasse 1–5, D-14469 Potsdam, Germany

[3]LEMTA, 2 Avenue de la Forêt de Haye, F-54518 Vandoeuvre-lès-Nancy, France

[4]Potsdam University, Institute for Physics and Astrophysics, Karl-Liebknecht Strasse 24–25, D-14476 Potsdam, Germany

[5]CERFACS, 42 Avenue Gaspard Coriolis, F-31057 Toulouse CEDEX 01, France

[6]Division of Fluid Dynamics, Department of Applied Mechanics, Chalmers University of Technology, SE-412 96 Göteborg, Sweden

[7]School of Engineering and Information Technology, The University of New South Wales at ADFA, Canberra, Australian Capital Territory, 2600, Australia

(Received 27 August 2013; revised 7 April 2014; accepted 19 June 2014; first published online 6 August 2014)

We propose a novel cluster-based reduced-order modelling (CROM) strategy for unsteady flows. CROM combines the cluster analysis pioneered in Gunzburger's group (Burkardt, Gunzburger & Lee, *Comput. Meth. Appl. Mech. Engng*, vol. 196, 2006*a*, pp. 337–355) and transition matrix models introduced in fluid dynamics in Eckhardt's group (Schneider, Eckhardt & Vollmer, *Phys. Rev.* E, vol. 75, 2007, art. 066313). CROM constitutes a potential alternative to POD models and generalises the Ulam–Galerkin method classically used in dynamical systems to determine a finite-rank approximation of the Perron–Frobenius operator. The proposed strategy processes a time-resolved sequence of flow snapshots in two steps. First, the snapshot data are clustered into a small number of representative states, called centroids, in the state space. These centroids partition the state space in complementary non-overlapping regions (centroidal Voronoi cells). Departing from the standard algorithm, the probabilities of the clusters are determined, and the states are sorted by analysis of the transition matrix. Second, the transitions between the states are dynamically modelled using a Markov process. Physical mechanisms are then distilled by a refined analysis of the Markov process, e.g. using finite-time Lyapunov exponent (FTLE) and entropic methods. This CROM framework is applied to the Lorenz attractor (as illustrative example), to velocity fields of the spatially evolving incompressible mixing layer and the three-dimensional turbulent wake of a bluff body. For these examples, CROM is shown to identify non-trivial quasi-attractors and transition processes in an unsupervised manner. CROM has numerous potential applications for the systematic identification of physical mechanisms of complex dynamics, for comparison of flow

† Email address for correspondence: eurika.kaiser@univ-poitiers.fr

evolution models, for the identification of precursors to desirable and undesirable events, and for flow control applications exploiting nonlinear actuation dynamics.

Key words: low-dimensional models, nonlinear dynamical systems, shear layers

1. Introduction

Multiscale phenomena, such as turbulence, exhibit a large number of degrees of freedom. The dimensionality of corresponding detailed flow models poses a challenge for understanding the physical mechanisms, for optimising realistic configurations and for controlling turbulent flows in real time, even though computational power and resources are continually increasing. The general strategy of coarse-graining, system reduction or model reduction – the term depends on the community – is to reduce the number of degrees of freedom through simplification of the used models, by keeping only those degrees of freedom that are relevant to the particular phenomena of interest, and also to the purpose of the model. Coarse-graining was introduced in the physics community more than a century ago, and is deeply connected to statistical physics and thermodynamics. Model reduction was more recently introduced for fluid flow analysis and control. A wide range of reduced-order models (ROMs) have been proposed and applied. Here, focus is placed on a data-driven ROM. Model reduction typically involves two separate but related steps: (i) selection of the coordinates comprising the reduced system and (ii) dynamical modelling of their interactions.

A wide class of model reduction techniques are linked to the governing evolution equations via projection techniques. These methods rest on the assumption that in many cases trajectories in the high-fidelity phase space are contained in low-dimensional subspaces. In the Petrov–Galerkin framework (Antoulas 2005), a low-dimensional subspace is first used to introduce a reduced state vector, and the reduced-order equations are then obtained by projection of the residual of the full-order model equations onto another low-dimensional subspace. Proper orthogonal decomposition (POD) models of fluid flows provide one elegant example of this class (see, e.g., Noack, Morzyński & Tadmor 2011; Holmes *et al.* 2012). Kinematically, POD yields an 'optimal' Galerkin expansion which minimises the averaged residual in the energy norm with respect to all other expansions of the same order. Dynamically, the POD coefficient space can be described by an evolution equation, here called a Galerkin system, determined by the projection of the Navier–Stokes equation onto the POD modes, obtained via a model identification technique (Cacuci, Navon & Ionescu-Bujor 2013), or a combination of the two approaches (Cordier *et al.* 2013). For control purposes, an actuation term can be added to the Galerkin system and exploited for control design (Bergmann & Cordier 2008; Luchtenburg *et al.* 2009). Twenty-five years after the pioneering work on wall turbulence by Aubry *et al.* (1988), POD models have become a widely used class of ROM, and numerous variations and generalisations have been suggested (weighted POD, POD with derivatives, adaptive POD, etc.). For the control of linear systems, one of the most important is balanced POD (Rowley 2005) where approximate balanced truncation (Moore 1981) is computed through a POD approximation of the product of the controllability and observability Gramians. By construction, those ROMs are optimal in terms of input–output description of the system, and error bounds exist that are close to the lowest error possible from any ROM. Another way of deriving a ROM is to project a linearised version of the governing equations around some base flow onto the

global and adjoint global modes. This approach was popularised for the control of flow instabilities observed in some aerodynamic configurations by Sipp *et al.* (2010). Those linearised models are particularly appropriate for model-based control since many open-loop and closed-loop control strategies were developed years ago in the control theory community (see Bagheri *et al.* 2008, for a review).

When projection methods are not used to derive ROMs, system-identification techniques relying on Markov parameters and Hankel matrix (Huang & Kim 2008), or auto-regressive models, e.g. ARMAX (Hervé *et al.* 2012), are mostly employed. One exciting recent direction is based on Koopman or dynamic mode decomposition (DMD) (Rowley *et al.* 2009; Schmid 2010) in which pure frequency modes are distilled.

In this study, we propose a novel cluster-based reduced-order modelling (CROM) strategy for unsteady flows. CROM combines a cluster analysis of an ensemble of observations and a Markov model for transitions between the different flow states educed by the clustering. Recently, clustering principles have also been applied to fluid dynamics problems in the group of Gunzburger. In Burkardt, Gunzburger & Lee (2006*a*), a centroidal Voronoi tessellation (CVT) approach is introduced for producing a ROM of the Navier–Stokes equation. Starting with a snapshot set as is done in POD, CVT relies on the construction of a special Voronoi clustering of the snapshots for which the means of the clusters are also the generators of the corresponding Voronoi cells (Du, Faber & Gunzburger 1999). In the CVT setting, a dynamical model is then derived by projection of the governing equations onto the clusters' centroids. This approach can then be interpreted as a generalisation of POD-based reduced-order modelling. In the CROM framework, the dynamical model does not rest on a Galerkin projection. On the contrary, we follow inspiring developments in the group of Eckhardt (Schneider, Eckhardt & Vollmer 2007) and model the transition process between the identified clusters with a Markov process (Norris 1998). This dynamical transition model proves to be simple and powerful.

Cluster analysis is an integral part of machine learning (see, e.g. Bishop 2007, for an introduction), a branch of artificial intelligence dedicated to the development of methods that can learn automatically from data. The starting point of clustering is – like POD – a sequence of snapshots, called observations, defined in some state space and a predefined distance measure. The distance measure employs a feature vector, quantifying important characteristics of the snapshot space. In our case, we define the original state space as feature and use the standard Euclidean distance for defining similarity between observations. The objective of clustering is to organise data so that the inner-cluster similarity is maximised while the inter-cluster similarity is minimised. Clustering partitions a large amount of snapshots into a given small amount of geometrically close subsets, called clusters, and determines a representative centre of mass, named centroid. The clustering approach identifies relatively homogeneous groups of observations according to the chosen metric. For solving the cluster analysis problem, an efficient heuristic algorithm known as k-means is used here (Lloyd 1956; Steinhaus 1956; Ball & Hall 1965; MacQueen 1967). This algorithm is probably the most popular technique of clustering due to its simplicity of implementation and fast execution. It is used in numerous areas ranging from market segmentation, computer vision, geostatistics and astronomy to agriculture (Murphy 2012). Clustering is sometimes used as a preprocessing step for other algorithms, for example to find a starting configuration. The close relation between CVTs and k-means algorithm was already mentioned in Burkardt *et al.* (2006*a*).

The CROM strategy can be interpreted as a generalisation of the Ulam–Galerkin method for the finite-rank approximation of the Perron–Frobenius operator. Complex

systems can be studied by the evolution of a single trajectory of a dynamical system or by adopting a probabilistic viewpoint and considering a swarm of trajectories evolving in the state space. The Perron–Frobenius operator is associated with such evolution equations, i.e. a Liouville equation, and describes the evolution of a probability density for a trajectory ensemble in the state space. In contrast to the Ulam–Galerkin method, CROM rests on a cluster analysis which allows to code additional, e.g. entropic, information in the distance metric.

CROM is first applied to the Lorenz attractor as a well-known introductory example (Lorenz 1963; see Sparrow 1982 for details). The second application is a two-dimensional incompressible mixing layer flow undergoing several successive vortex pairings. The instability mechanism of spatially developing mixing layers and the spatial organisation of their dominant large-scale structures have been the subject of many studies (see, e.g., Brown & Roshko 1974; Browand & Weidman 1976; Dimotakis & Brown 1976; Monkewitz 1988). The last example is the three-dimensional turbulent wake of a bluff body, the so-called Ahmed body. The Ahmed body is a generic car model to study the relationship between characteristic flow structures typically arising at passenger cars and their impact on the drag and lift forces (Ahmed, Ramm & Faltin 1984; Lienhart & Becker 2003; Pastoor *et al.* 2008). Recently, Grandemange, Gohlke & Cadot (2013) showed that the wake of a square-back body as in our example is characterised by a bimodal behaviour, i.e. the flow switches between two asymmetric flow states over large time scales $T \sim 10^2$–$10^3 H/U_\infty$ for Reynolds numbers ranging from $Re \sim 10^2$ to $Re \sim 10^6$ (Grandemange, Cadot & Gohlke 2012).

These examples show complex dynamical behaviour associated with so-called ‘quasi-attractors’. We shall employ the term quasi-attractor as an isolated region in the state space with a long average residence time of the entering trajectory. For instance, the two ears of the Lorenz attractor can be considered as quasi-attractors since the switching between them is slow as compared with the oscillatory time scale. Other considered examples of quasi-attractors are the harmonic and subharmonic vortex shedding regime of a mixing layer, and the bimodal behaviour of the Ahmed body wake. More formally, a quasi-attractor (see, e.g., Afraimovich & Shil’nikov 1983; Gonchenko, Shil’nikov & Turaev 1997) is an attractive limit set of an attractor which contains non-trivial hyperbolic subsets and which may contain attractive periodic orbits of extremely long periods. The modelling of such quasi-attractors is challenging due to the complex interaction between long and short time scales.

This paper is organised as follows. In § 2 the methodology of the CROM is outlined. Section 3 presents a CROM for the Lorenz attractor as an introductive example, while in § 4 CROM distils physical mechanisms of the mixing layer. An interpretation of CROM as a generalised Ulam–Galerkin method is presented in § 5. The relative merits and challenges of CROM benchmarked against POD models are discussed in § 6. Finally, the main findings are summarised and future directions are given (§ 7). In appendix A, CROM is applied to a three-dimensional wake with broadband frequency dynamics. A CROM-based control strategy is outlined in appendix B. CROM’s connection to the Perron–Frobenius and Koopman operators is elaborated in appendix C. A simple visualisation method for the cluster topology as explained in appendix D provides additional insights on the CROM analysis.

2. Cluster-based reduced-order modelling

In this section, the CROM methodology is presented. CROM combines a low-order kinematic description of the state space as obtained by a cluster analysis of an

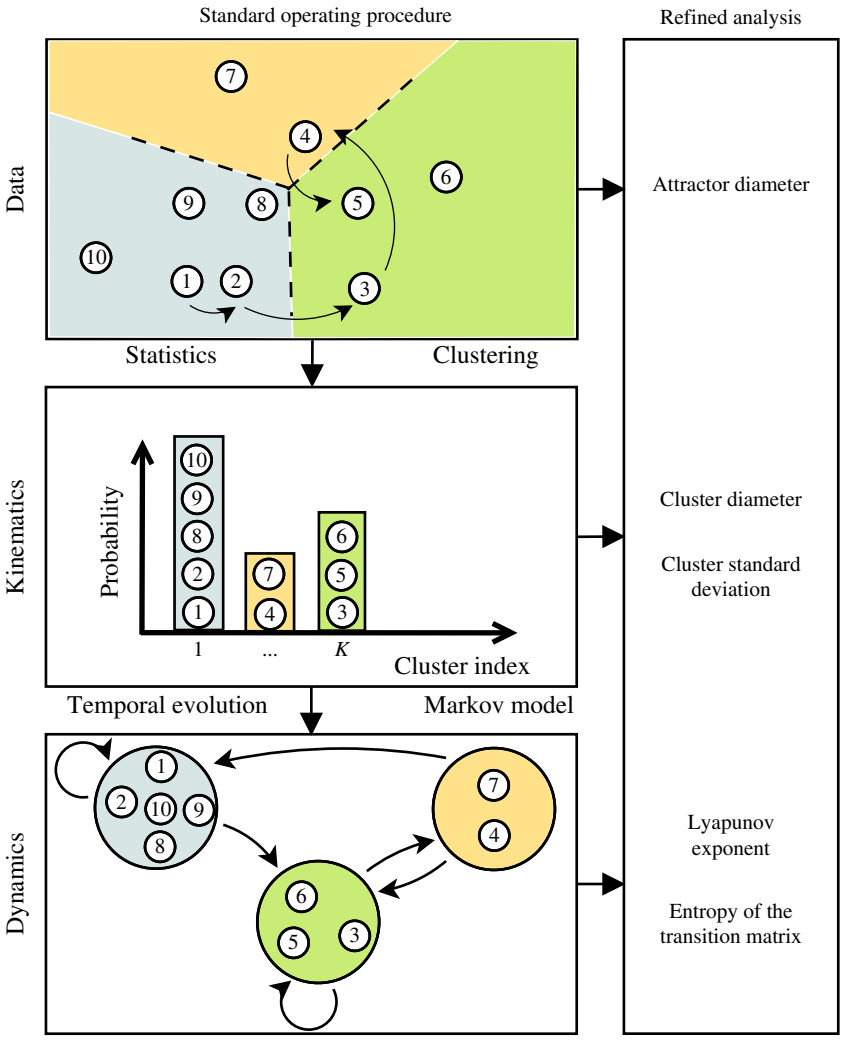

FIGURE 1. (Colour online) Scheme of CROM methodology applied to a sequence of observations. The standard operating procedure employs: (i) a partitioning of the observations into clusters of kinematically similar states (k-means algorithm); (ii) an estimation of the probability distribution of being in a given state is provided by a statistical analysis (this estimation is based on the assumption that a representative state of each cluster is supplied by its centroid); (iii) an exploitation of the sequential information of the observations by assuming a Markov process for the cluster transitions (this step yields to a Markov model for the evolution of the probability distribution). For all steps (i)–(iii), a refined analysis provides additional information on the attractor. For details see § 2.

ensemble of snapshots (see § 2.1), and a dynamical model for the cluster states (see § 2.2). In § 2.3, different criteria are defined which help to characterise the attractor educed by the CROM. These criteria will be later employed in § 3 and § 4 for the Lorenz attractor and mixing layer, respectively. A schematic road map for the CROM is displayed in figure 1.

2.1. *Kinematic analysis*

CROM rests on the definition of a metric (see § 2.1.1) which measures the similarity between snapshots. The state space is then partitioned with a clustering algorithm called k-means (§ 2.1.2). The means of the resulting clusters define a low-dimensional description of the state space and produce an optimal reduced basis in the sense that the total cluster variance is minimised.

2.1.1. *Metric*

We consider velocity fields in a steady domain Ω. The location is described in a Cartesian coordinate system x, y, z with the location vector $\boldsymbol{x} = (x, y, z)$. The velocity field $\boldsymbol{u}(\boldsymbol{x}, t)$ depends on the position $\boldsymbol{x}$ and the time t. Basis for CROM is an ensemble of M snapshots $\{\boldsymbol{u}^m(\boldsymbol{x})\}_{m=1}^M$, where $\boldsymbol{u}^m(\boldsymbol{x}) := \boldsymbol{u}(\boldsymbol{x}, t_m)$ is the mth flow realisation at discrete time t_m. Thus, the employed data are similar as for reduced-order representations via POD (see, e.g., Holmes *et al.* 2012) or DMD (Rowley *et al.* 2009; Schmid 2010). The flow may be obtained experimentally or from numerical simulations.

In principle, the snapshots may arise from transient processes or from a post-transient stationary state (attractor). In the later case, the snapshots can be expected to approximate first and second statistical moments of the data, like for POD. In the following, focus is placed on such attractor data, including the applications to the Lorenz attractor and mixing-layer simulation.

For later reference, we define a Hilbert space, an inner product and an associated metric. Let $\boldsymbol{f}$ and $\boldsymbol{g}$ be two velocity fields in the Hilbert space $\mathscr{L}^2(\Omega)$ of square-integrable vector fields in the domain Ω. Their inner product is defined by

$$(\boldsymbol{f}, \boldsymbol{g})_\Omega := \int_\Omega \mathrm{d}\boldsymbol{x}\, \boldsymbol{f} \cdot \boldsymbol{g} \tag{2.1}$$

with the corresponding norm

$$\|\boldsymbol{f}\|_\Omega^2 := (\boldsymbol{f}, \boldsymbol{f})_\Omega. \tag{2.2}$$

The metric for the cluster analysis is chosen to be the distance between two velocity fields according to the above norm. In particular, the distance between snapshots $\boldsymbol{u}^m$ and $\boldsymbol{u}^n$ is given by

$$D_{mn}^\bullet := \|\boldsymbol{u}^m - \boldsymbol{u}^n\|_\Omega. \tag{2.3}$$

Here and in the following, the superscript '$\bullet$' refers to the raw snapshot data.

Burkardt *et al.* (2006*a*) showed that the k-means clustering technique leads naturally to the definition of basis functions that can be used to derive ROMs as classically done with the POD modes. For this reason, we decide to approximate the velocity fields by POD to illustrate in the results (see § 4) the relation between the modes obtained by POD and by clustering. Moreover, neglecting the preprocessing costs of POD for a moment, the computational and input load of the cluster analysis is significantly reduced in the POD subspace since the evaluation of $D_{mn}^\bullet$ by (2.3) implies an expensive volume integration over all grid points.

POD rests on a time or ensemble average denoted by $\langle\cdot\rangle$. For the considered snapshot ensemble, the conventional average is employed. POD generalises the Reynolds decomposition of a flow into a mean flow $\boldsymbol{u}_0(\boldsymbol{x}) := \langle \boldsymbol{u}(\boldsymbol{x}, t)\rangle$ and a fluctuation $\boldsymbol{v}(\boldsymbol{x}, t)$. The POD expansion (Noack *et al.* 2003; Holmes *et al.* 2012) reads

$$\boldsymbol{v}(\boldsymbol{x}, t) := \boldsymbol{u}(\boldsymbol{x}, t) - \boldsymbol{u}_0(\boldsymbol{x}) \approx \sum_{i=1}^{N} a_i(t)\, \boldsymbol{u}_i(\boldsymbol{x}), \tag{2.4}$$

where the $\boldsymbol{u}_i$ are the POD modes, and $a_i := (\boldsymbol{v}, \boldsymbol{u}_i)_\Omega$, $i = 1, \ldots, N$, are the mode coefficients. Note that M snapshots can span a space of maximum dimension $M - 1$. Hence, the number of modes in the expansion is restricted to $N \leqslant M - 1$.

The POD expansion enables the cluster analysis of the flow snapshots in terms of the mode coefficients. Each snapshot is represented by a POD coefficient vector $\boldsymbol{a}^m = [a_1(t_m),\ a_2(t_m),\ \ldots,\ a_N(t_m)]^{\mathrm{T}}$, where the superscript T denotes the transpose. If the Galerkin expansion (2.4) were exact, the distance between two snapshots would be expressed by the Euclidean distance between the two corresponding POD coefficient vectors, i.e.

$$D^{\bullet}_{mn} := \|\boldsymbol{v}^m - \boldsymbol{v}^n\|_\Omega = \|\boldsymbol{a}^m - \boldsymbol{a}^n\|. \tag{2.5}$$

The relation (2.5) is also true for any orthonormal basis. The evaluation of the right-hand side requires about M multiplications, subtractions and additions. Hence, the computational savings by employing the mode coefficient metric for the following CROM operations are enormous. However, the preprocessing cost of the POD requires the computation of $M(M + 1)/2$ correlation integrals, which have a similar computational count as the distance integrals. For comparison, one iteration of the clustering algorithm will effect KM distance integrals between M snapshots and K clusters. The total cost of the cluster algorithm scales with IKM, I being the total number of iterations until convergence. Evidently, POD becomes computationally advantageous if $IK > (M + 1)/2$.

We emphasise that the POD coefficient vector is only used for logistic convenience and is not necessary for the clustering algorithm. Thus, we keep the notation of snapshots for the methodological part and only refer to mode coefficients when necessary.

2.1.2. *Clustering algorithm*

CROM assumes a sequence of fluctuations $\{\boldsymbol{v}^m\}_{m=1}^M$ at discrete times t_m. In the machine learning context (see Hastie, Tibshirani & Friedman 2009, for instance), these snapshots are also called observations or objects. We wish to partition these snapshots into K clusters $\mathscr{C}_k$, $k = 1, \ldots, K$, characterised in terms of similarity of their elements. Here, similarity of objects is defined based on the measure $D^{\bullet}_{mn}$ introduced in (2.5).

In clustering, each snapshot is connected to the nearest cluster by using a characteristic function defined as

$$T_k^m := \begin{cases} 1, & \text{if } \boldsymbol{v}^m \in \mathscr{C}_k, \\ 0, & \text{otherwise.} \end{cases} \tag{2.6}$$

The number of observations n_k in cluster k is given by

$$n_k := \sum_{m=1}^{M} T_k^m. \tag{2.7}$$

The total number of snapshots can be derived from

$$M = \sum_{k=1}^{K} n_k = \sum_{k=1}^{K} \sum_{m=1}^{M} T_k^m. \tag{2.8}$$

The cluster mean or cluster centroid $\boldsymbol{c}_k$ of $\mathscr{C}_k$ is here defined as the average of the snapshots belonging to the cluster. This centroid can be conveniently expressed by the characteristic function:

$$\boldsymbol{c}_k = \frac{1}{n_k} \sum_{\boldsymbol{v}^m \in \mathscr{C}_k} \boldsymbol{v}^m = \frac{1}{n_k} \sum_{m=1}^{M} T_k^m \boldsymbol{v}^m. \tag{2.9}$$

Thus, all snapshots have the same weight like in Du *et al.* (1999) and Burkardt, Gunzburger & Lee (2006*b*) for the CVT approach.

The k-means algorithm partitions the data space into K centroidal Voronoi cells which are defined as particular Voronoi cells for which the generating points of the Voronoi tessellation are equal to the mass centroids of the Voronoi regions (Du *et al.* 1999). The quality of the algorithm is monitored by the total cluster variance, which sums up the variances of all the clusters $\mathscr{C}_k$:

$$J(\boldsymbol{c}_1, \ldots, \boldsymbol{c}_K) = \sum_{k=1}^{K} \sum_{\boldsymbol{v}^m \in \mathscr{C}_k} \|\boldsymbol{v}^m - \boldsymbol{c}_k\|_\Omega^2 = \sum_{k=1}^{K} \sum_{m=1}^{M} T_k^m \, \|\boldsymbol{v}^m - \boldsymbol{c}_k\|_\Omega^2 \,. \tag{2.10}$$

The optimal locations of the centroids are determined as solutions of the minimisation problem based on J, i.e.

$$\boldsymbol{c}_1^{opt}, \ldots, \boldsymbol{c}_K^{opt} = \underset{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_K}{\operatorname{argmin}} J(\boldsymbol{c}_1, \ldots, \boldsymbol{c}_K), \tag{2.11}$$

where $\boldsymbol{c}_k^{opt}$ are given by (2.9). In this sense, the ensemble of M snapshots is optimally represented by the K clusters and their centroids. In the following, the centroids are assumed to be optimal according to (2.11) and the superscript 'opt' will be dropped.

The k-means algorithm contains three steps. First, K centroids are initialised. Then, each observation is assigned to its closest centroid using the chosen metric. This is called the assignment step. In the update step, the cluster centres are recalculated as the mean of all observations belonging to this cluster. The assignment step and the update step are repeated until convergence is reached. Convergence means that either the centroids do not move or do not change below a threshold. Since k-means has a sensitivity to initial conditions, a cluster centre initialisation using k-means++ (Arthur & Vassilvitskii 2007) is applied here.

From a machine learning perspective, the clustering algorithm k-means is an unsupervised classification algorithm. This means that the observations are classified based on their similarity without any prior assumption or categorisation of the data. For a strictly periodic flow, this results in a generalisation of phase averaging. For a sufficient number of clusters K, each cluster centroid represents the phase average of the snapshots in each cluster.

Each of steps 2 and 3 of the k-means algorithm reduces the value of the criterion J, so that convergence is assured when the number of clusters K is increased. In terms of numerical accuracy, the choice $K = M$ is then optimal since in this case the centroids become the snapshots and the total cluster variance vanishes. On the other hand, in terms of compression of the data, we would like to employ as few clusters as possible. The choice of the number of clusters is then a matter of compromise between numerical accuracy and compression of the data. In general, K is fixed at a value $K_c \ll M$ corresponding to the number of clusters for which the variation of the variance J with K becomes small. This method is known in machine learning as

the 'elbow criterion' (Hartigan 1975; Tibshirani, Walther & Hastie 2001). Recently, Chiang & Mirkin (2010) compared in an extensive study different measures to determine a suitable number of clusters that represent hidden groups in the data. In our study, we consider K_c small to reveal the main physical transition mechanisms with a sufficient numerical accuracy. In §§ 3 and 4, we adopt K_c equal to 10.

From a system point of view, the cluster analysis can be seen as a method of dimension reduction. A high-dimensional state space is discretised into K centroidal Voronoi cells. Thus, the probability distribution over the state space can be discretised as a probability distribution over the clusters. Each cluster $\mathscr{C}_k$ is represented by its centroid $\boldsymbol{c}_k$, which approximates the state of the set of snapshots in this cluster. For this reason, in the following sections, we decide by convention to refer to the cluster centroids $\boldsymbol{c}_k$ as states.

2.2. *Dynamical model*

In this section, the cluster analysis is augmented with a dynamical model for the cluster states. In § 2.2.1, the cluster probability distribution is obtained directly from the snapshot data. The major step is the determination of the cluster transition matrix (CTM) which serves as the propagator in terms of probability. This matrix is the constitutive element of a dynamical model elaborated in § 2.2.2. The derived dynamical model has the Markov property which assumes that the system is memoryless. A discussion of the applicability of the Markov property is provided in § 2.2.3.

2.2.1. *Identification of the CTM*

In the kinematic part of CROM (§ 2.1), the snapshots are assumed to fill the state space sufficiently densely for some representative statistics. For a statistical analysis such as clustering, the advantage of having uncorrelated snapshots is to minimise redundancy. For the dynamical component of CROM, the CTM, we have to request a small time step. In other words, the time step $\Delta t = t^{m+1} - t^m$, $m = 1, \ldots, M-1$ is constant and small as compared with the characteristic time scale of the flow.

The clustering algorithm yields state space coarse-grained in centroidal Voronoi cells. Different points of a cluster have different trajectories. Thus, on the level of cluster resolution, we can only make statements about 'the most probable' trajectory. Hence, we search for a propagator of a probability distribution discretised by K clusters. More precisely, we will describe the dynamics of clusters as a Markov chain (Norris 1998). In this paper, the convention is made to use a left stochastic matrix (each column sum is equal to one) for the dynamical propagator and column vectors of probabilities.

First, a probabilistic representation of the system, equivalent to the weighted average realisation of the ensemble (see, e.g., Fowler 1929; Niven 2009), can be defined by assigning a probability q_k to each cluster $\mathscr{C}_k$. Each component q_k is approximated by the number of snapshots in cluster k normalised with the total number of snapshots:

$$q_k := \frac{n_k}{M} = \frac{1}{M} \sum_{m=1}^{M} T_k^m. \tag{2.12}$$

Since every cluster contains at least one element ($n_k > 0$), the probabilities q_k are strictly positive ($q_k > 0$). In addition, (2.12) fulfils the normalisation condition $\sum_{k=1}^{K} q_k = 1$. These probabilities are summarised in the cluster probability vector

$\boldsymbol{q} = [q_1, \ldots, q_K]^{\mathrm{T}}$. This vector shall approximate the fixed point of the dynamical model developed below.

In the following, the temporal evolution of a general cluster probability vector $\boldsymbol{p}$ is pursued. For computational convenience, we describe the evolution with the same time step Δt as the original data, i.e. pursue a time-discrete formulation. The transition to a continuous-time formalism is described below.

The CTM $\boldsymbol{P}$ is identified as the one-step transition probability matrix. Let P_{jk} be the probability of moving from cluster $\mathscr{C}_k$ to cluster $\mathscr{C}_j$ in one forward time-step, the elements of the resulting CTM can be inferred by

$$P_{jk} = \frac{n_{jk}}{n_k}, \tag{2.13}$$

where n_{jk}, the number of snapshots that move from $\mathscr{C}_k$ to $\mathscr{C}_j$, are given by

$$n_{jk} := \sum_{m=1}^{M-1} T_k^m \, T_j^{m+1}. \tag{2.14}$$

Note that M consecutive snapshots define only $M - 1$ transitions.

All elements of this matrix are non-negative,

$$P_{jk} \geqslant 0, \quad j, k = 1, \ldots, K. \tag{2.15}$$

The elements of each column sum up to unity

$$\sum_{j=1}^{K} P_{jk} = 1 \tag{2.16}$$

as the total probability for all transitions is one. This property is necessary to preserve the normalisation condition of the probability vector.

The diagonal element P_{jj} gives probability to stay in the cluster $\mathscr{C}_j$ during one time step. In order to extract the most probable path, the clusters are ordered so that the first cluster has the highest state probability q_1. The second cluster is chosen that the highest transition probability is to move from state $\boldsymbol{c}_1$ to state $\boldsymbol{c}_2$, etc. If one cluster has already been utilised, the cluster with the second highest transition probability is used, etc. The eigenvalues of the CTM are sorted by descending moduli, $|\lambda_1| \geqslant |\lambda_2| \geqslant \cdots \geqslant |\lambda_K|$.

2.2.2. *CTM model*

The above cluster transition analysis leads naturally to Markov chains. A Markov chain is characterised by the Markov property that the probability to move from the current state $\boldsymbol{c}_k$ at time step l to another state $\boldsymbol{c}_j$ in the next time step $l + 1$ depends only on the current state and not on past values. The autonomous nature of the Navier–Stokes equation suggests a similar time-independent property for the propagator $\boldsymbol{P}$ of the cluster probability vector. For sufficiently smooth initial conditions and steady boundary conditions, the uniqueness of the solution can be proven for two-dimensional flows (Ladyzhenskaya 1963), and, in general, can be assumed for most three-dimensional numerical simulations. However, the solution of the Navier–Stokes equation may depend on the history if, e.g., the boundary

conditions are discontinuous as shown by Rempfer (2006). We refer to § 2.2.3 for a more detailed discussion of the Markov assumption.

The cluster probability vector with probabilities p_k^l to be in state $\boldsymbol{c}_k$ at time l is denoted by

$$\boldsymbol{p}^l = \left[p_1^l, \ldots, p_K^l\right]^{\mathrm{T}}. \tag{2.17}$$

This vector respects that only non-negative probabilities

$$p_k^l \geqslant 0 \tag{2.18}$$

have a meaning and the normalisation condition

$$\sum_{k=1}^{K} p_k^l = 1 \tag{2.19}$$

holds for each time step l. The conditional probability to be in state $\boldsymbol{c}_j$ if the previous state has been $\boldsymbol{c}_k$ is given by the CTM defined in (2.13).

The evolution of the cluster probability vector can be described as follows. Let $\boldsymbol{p}^0$ be the initial probability distribution. Then, consecutive distributions are described by the iteration formula

$$\boldsymbol{p}^{l+1} = \boldsymbol{P}\boldsymbol{p}^l, \quad l = 0, 1, 2, \ldots. \tag{2.20}$$

The CTM properties (2.15) and (2.16) guarantee properties (2.18) and (2.19) of the cluster probability vector. The cluster probability vector at time l is compactly given by

$$\boldsymbol{p}^l = \boldsymbol{P}^l \boldsymbol{p}^0. \tag{2.21}$$

The propagator $\boldsymbol{P}$ defines a time-homogeneous Markov chain with well-known properties (Meyer 2000):

(i) The propagator $\boldsymbol{P}$ is a stochastic matrix with the properties (2.15) and (2.16).

(ii) The sequence of probability vectors $\boldsymbol{p}^l$, $l = 0, 1, 2, \ldots$ has no long-term memory. The state at iteration $l+1$ only depends on the lth state and not on any previous iterations.

(iii) The absolute values of all eigenvalues of this matrix do not exceed unity. This excludes a diverging vector sequence.

(iv) It exists an eigenvalue $\lambda_1(\boldsymbol{P}) = 1$ with algebraic multiplicity 1 and all other eigenvalues satisfy $|\lambda_k(\boldsymbol{P})| < 1$ for $k = 2, \ldots, K$. This is a consequence of the Perron–Frobenius theory for non-negative matrices (Meyer 2000). The eigenvector $\boldsymbol{p}_1^{ev}$ associated with the dominant eigenvalue λ_1 fulfils $\boldsymbol{P}\boldsymbol{p}_1^{ev} = \boldsymbol{p}_1^{ev}$. Since $|\lambda_k(\boldsymbol{P})| < 1$ for $k = 2, \ldots, K$, the vector $\boldsymbol{p}_1^{ev}$ is the only one that survives infinite iterations.

The long-term behaviour can be studied by powers of the CTM as defined in (2.21). The asymptotic probability distribution is obtained by

$$\boldsymbol{p}^\infty := \lim_{l\to\infty} \boldsymbol{P}^l \boldsymbol{p}^0 \tag{2.22}$$

and the infinite-time CTM $\boldsymbol{P}^\infty$ is defined as

$$\boldsymbol{P}^\infty := \lim_{l\to\infty} \boldsymbol{P}^l. \tag{2.23}$$

If $\boldsymbol{P}^l$ converges to a unique, stationary matrix $\boldsymbol{P}^\infty$, the system can be said to be *ergodic*, in the sense that it will be probabilistically reproducible: regardless of the initial region of state space in which it is sampled, the ensemble mean will converge in the infinite-time limit to the time mean. Mathematically, the stationary probability vector must be identical with the eigenvector $\boldsymbol{p}_1^{ev}$ associated with the dominant eigenvalue $\lambda_1 = 1$, and thus is a fixed point to (2.21) for any l and with respect to iteration with (2.23). If, however, $\boldsymbol{P}^\infty$ is oscillatory or non-stationary, the system will not be probabilistically reproducible, displaying a more complicated connection between the initial sampling region and its convergence properties. The asymptotic rate of convergence at which $\boldsymbol{P}^l$ approaches $\boldsymbol{P}^\infty$ is given by its second largest eigenvalue modulus $|\lambda_2|$. This is similar to the power method where the rate of convergence is determined by the ratio $|\lambda_2/\lambda_1|^l$ and how fast it approaches zero (Meyer 2000).

Iterating $\boldsymbol{P}$ amounts to determining the 'most significant' eigenvector. This procedure is used in many fields, the most prominent example being the PageRank algorithm used by the Google search engine (Langville & Meyer 2012). This reveals the underlying system dynamics, in the sense that a cluster is not simply important by itself, but rather by the clusters which transition to it. This provides a powerful approach to determine the *dynamically* important clusters (or structures).

The accuracy of the CTM can be assessed when it is benchmarked against the continuous-time form of the evolution equation (2.20):

$$\frac{\mathrm{d}\boldsymbol{p}(t)}{\mathrm{d}t} = \boldsymbol{P}^{cont}\boldsymbol{p}(t). \tag{2.24}$$

The discrete-time Markov model (2.20) can be derived from (2.24) by an analytical integration over the time Δt. The resulting discrete-time transition matrix reads

$$\boldsymbol{P} = \exp\left(\boldsymbol{P}^{cont}\Delta t\right) \Rightarrow \boldsymbol{P}^{cont} = \frac{1}{\Delta t}\log \boldsymbol{P} \approx \frac{\boldsymbol{P}-\boldsymbol{I}}{\Delta t} + O(\Delta t), \tag{2.25}$$

where $\boldsymbol{I}$ is the identity matrix. The approximate term on the right-hand side coincides with the Euler integration scheme for (2.24). Note that the transition matrix $\boldsymbol{P}$ is an approximation of the Perron–Frobenius operator for discrete states and discrete times (Lasota & Mackey 1994; Cvitanović *et al.* 2012). We refer to § 5 for more details.

The reverse operation, i.e. the derivation of the continuous-time Markov model (2.24) from the discrete-time version (2.20), requires a carefully chosen limiting process $\Delta t \to 0$. Let us consider a continuous-time trajectory over the time interval $[0, T]$. Let us further assume that the state space is partitioned into K fixed clusters. From this trajectory, M snapshots are taken at constant sampling rate, i.e. with time step $\Delta t = T/M$. Each of these snapshots can be attributed to one cluster $k = 1, \ldots, K$ and the discrete-time transition matrix $\boldsymbol{P}_M$ can be computed via (2.13) as function of M. As $M \to \infty$ or, equivalently, $\Delta t \to 0$, $(1/\Delta t)\log \boldsymbol{P}_M \to \boldsymbol{P}^{cont}$. From $P_{jk} = n_{jk}/n_k$, the possible values of a transition matrix element are discrete as n_k and n_{jk} are integers. The maximum resolution is controlled by the denominator $1/n_k \geqslant 1/M$. Thus, increasing M improves the maximum accuracy of $\boldsymbol{P}_M$. The time horizon T for the trajectory also plays an important role. The number of non-trivial cluster transitions scales with T and are well defined in the limit $\Delta t \to 0$. Thus, the accuracy of the continuous-time transition matrix also increases with the time horizon.

2.2.3. *Validity of the Markov model*

The assumed Markov property of the model requires a plausibility argument: the Markov property assumes that a system is memoryless, i.e. future states of the system are solely determined by the current state and do not depend on the previous history of the system. Counterexamples are differential equations with history terms (time integral over past states), like for some non-Newtonian fluids. Also, the solution of the incompressible Navier–Stokes equation may indeed depend on the history, particularly for time-discontinuous boundary conditions (Rempfer 2006).

In the following, we consider a steady domain with steady boundary conditions and additional mathematical smoothness properties of the domain and the initial condition. Under such premises, the uniqueness of the Navier–Stokes solution can be proven for two-dimensional flows (Ladyzhenskaya 1963). For three-dimensional flows, the uniqueness is generally assumed in an adequate numerical Navier–Stokes discretisation, i.e. the computed numerical solution is expected to be essentially unique up to the prediction horizon of the initial boundary value problem. The resulting grid-based or spectral methods yield an autonomous dynamical system in a finite-dimensional state space. Traditional Galerkin methods (Fletcher 1984), with POD models as a subset, yield the same type of dynamical systems. In all of these spatial Navier–Stokes discretisations, the corresponding evolution of the probability density in the state space is described by a Liouville equation. As will be shown in § 5, the state space discretisation of the Liouville equation (5.2) leads naturally to a continuous-time autonomous Markov model (5.8) with the Markov property.

Under similar premises, smoothly varying boundary conditions can be expected to lead to a time-dependent Markov model. We refer to Hopf (1952) and Vishik & Fursikov (1988) for a discussion of the Liouville equation for the (undiscretised) Navier–Stokes equation.

2.3. *Attractor characterisation*

In the following sections, the CROM is used to educe geometric and dynamic properties of the flow dynamics for the Lorenz attractor (§ 3), a mixing layer (§ 4) and an Ahmed body (appendix A), respectively. The properties are quantified by parameters introduced in this section.

First, the size of the attractor is represented by the snapshot data and the size of the clusters is quantified. The diameter of the attractor represented by the snapshot ensemble is given by

$$D_A^\bullet := \sup_{m,n} \left\{ D_{mn}^\bullet \,:\, \boldsymbol{v}^m, \boldsymbol{v}^n \in \text{attractor} \right\}, \tag{2.26}$$

where $D_{mn}^\bullet$ is defined by (2.5). Here, $D_A^\bullet$ constitutes an upper bound and reference level to all considered distances.

Analogously, we introduce the diameter of the cluster $\mathscr{C}_k$ as

$$D_k := \sup_{m,n} \left\{ D_{mn}^\bullet \,:\, \boldsymbol{v}^m, \boldsymbol{v}^n \in \mathscr{C}_k \right\}. \tag{2.27}$$

Similar diameters for all clusters indicate a relatively homogeneous partition of the state space.

The standard deviation R_k of the cluster $\mathscr{C}_k$ is given by

$$R_k = \sqrt{\frac{1}{n_k} \sum_{\boldsymbol{v}^m \in \mathscr{C}_k} \|\boldsymbol{v}^m - \boldsymbol{c}_k\|_\Omega^2} = \sqrt{\frac{1}{n_k} \sum_{m=1}^{M} T_k^m \, \|\boldsymbol{v}^m - \boldsymbol{c}_k\|_\Omega^2}. \tag{2.28}$$

This quantity allows to assess the homogeneity of the clusters.

In analogy to the snapshot distance matrix, we introduce the cluster distance matrix $\boldsymbol{D} = (D_{jk})$, which incorporates the distances between the cluster centroids:

$$D_{jk} := \left\| \boldsymbol{c}_j - \boldsymbol{c}_k \right\|_{\Omega}, \quad j, k = 1, \ldots, K. \tag{2.29}$$

The cluster distance matrix allows, for instance, to estimate the trajectory length for cluster transitions, e.g., D_{jk} for the direct transition from cluster k to j.

In addition, a characteristic flow norm variance represented by the cluster probability distribution $\boldsymbol{p} = [p_1, \ldots, p_K]^{\mathrm{T}}$ is naturally defined by

$$R^2(\boldsymbol{p}) := \sum_{j=1}^{K} \sum_{k=1}^{K} p_j \, p_k \, D_{jk}^2. \tag{2.30}$$

The mixing property of the propagator $\boldsymbol{P}$ is characterised by the finite-time Lyapunov exponent (FTLE)

$$\lambda_{FTLE}^{l} = \frac{1}{l} \ln \left\| \frac{\boldsymbol{P}^l \, \delta \boldsymbol{p}^0}{\delta \boldsymbol{p}^0} \right\|. \tag{2.31}$$

This exponent measures the divergence of the distance $\delta \boldsymbol{p}^l = \boldsymbol{p}_a^l - \boldsymbol{p}_b^l$ of initially close probability distributions $\boldsymbol{p}_a^0$ and $\boldsymbol{p}_b^0$ via (2.21). The dependence of the FTLE on the particular initial condition decreases rapidly for increasing l. Eventually, all trajectories converge to the fixed point (2.22) and the Lyapunov exponent vanishes

$$\lambda_{FTLE} = \lim_{l \to \infty} \lambda_{FTLE}^{l} = 0. \tag{2.32}$$

Finally, the transition matrix of the propagator is assessed. The infinite-time propagator (2.23) with stable asymptotic probability distribution (2.22) can be compared with a transition matrix $\boldsymbol{Q} = (Q_{jk})$ constructed from the cluster analysis, with elements

$$Q_{jk} = q_k, \quad j, k = 1, \ldots, K. \tag{2.33}$$

Note that $\boldsymbol{P}^{\infty}$ is obtained from a dynamically asymptotic realisation (mode) and $\boldsymbol{Q}$ from a weighted average realisation (mean). Hence, one cannot expect them to be identical. However, in an ergodic system in the limit of a very large number of samples, the two matrices should approach a common value. The difference between the two can be quantified by the Kullback–Leibler entropy (Kullback & Leibler 1951; Kullback 1959; Noack & Niven 2012)

$$H(\boldsymbol{P}, \boldsymbol{Q}) = -\mathfrak{D}(\boldsymbol{P}, \boldsymbol{Q}) := -\sum_{j=1}^{K} \sum_{k=1}^{K} P_{jk} \ln \frac{P_{jk}}{Q_{jk}}. \tag{2.34}$$

Alternatively, the divergence $\mathfrak{D}$ is given by the negative Kullback–Leibler entropy. For $\boldsymbol{P} = \boldsymbol{Q}$, the divergence vanishes. For $\boldsymbol{P} \neq \boldsymbol{Q}$, the divergence assumes positive values. This entropy measures the uncertainty of P_{jk} relative to Q_{jk} as prior (least-informative) expectation, while the divergence can be considered as the information gained from the quantification of the finite-time transitions. More importantly, the entropy provides a measure of the system ergodicity, reflected by convergence towards $H = 0$ in the infinite sampling limit $\boldsymbol{P} \to \boldsymbol{Q}$.

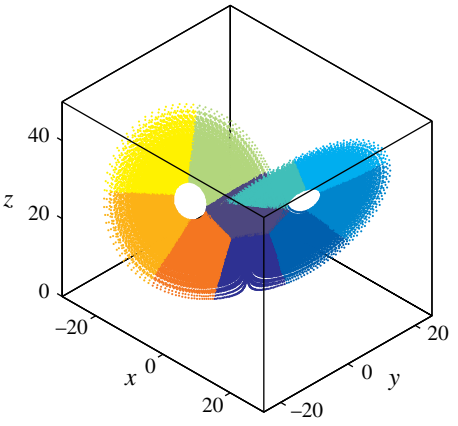

FIGURE 2. (Colour online) Phase portrait of the clustered Lorenz attractor for $K_c = 10$ clusters. Points with the same colour belong to the same cluster. Note that the 'ears' of the Lorenz attractor are partitioned into clean sectors.

3. Lorenz attractor as an illustrating example

We illustrate CROM for a nonlinear system of ODEs, the Lorenz equations,

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \sigma\ (y - x)\,, \tag{3.1a}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = x\ (r - z) - y, \tag{3.1b}$$

$$\frac{\mathrm{d}z}{\mathrm{d}t} = x\,y - b\,z, \tag{3.1c}$$

with the system parameters $\sigma = 10$, $b = 8/3$ and $r = 28$. These equations represent a model for Rayleigh–Bénard convection as proposed by Lorenz (1963). For these values of the parameters, the attractor possesses three unstable fixed points at $(0,\ 0,\ 0)$ and $(\pm\sqrt{72},\ \pm\sqrt{72},\ 27)$. The trajectory is most of the time in the neighbourhood of the weakly unstable oscillatory fixed points $(\pm\sqrt{72},\ \pm\sqrt{72},\ 27)$. These neighbourhoods are loosely referred to as 'ears' of the Lorenz attractor. Starting from an initial state close to one of the non-zero fixed points, the trajectory oscillates for several periods around this point before it moves to the other ear. Here, it oscillates in a similar manner before it transitions back again. The Lorenz system is solved using an explicit fourth-order Runge–Kutta formula with an initial condition chosen to be on the attractor. The sampling time step is $\Delta t = 0.005$. The cluster analysis is performed on snapshot data $\{\boldsymbol{x}(t_m)\}_{m=1}^{M}$ with $\boldsymbol{x} = [x, y, z]^{\mathrm{T}}$ for $M = 100\,000$ time units.

Figure 2 displays the phase portrait of the Lorenz attractor and the segmentation according to the clustering. We chose $K_c = 10$ clusters. This value is in good agreement with the 'elbow criterion' (see § 2.1.2). This number is large enough to resolve the main transition mechanism and small enough to obtain a simple structure for the model. A qualitative discussion on the choice of the number of clusters, its lower and upper bounds as well as its influence on the state and transition resolution is provided at the end of § 4. Each colour in the figure represents one cluster. Note that the sawtooth at the outer edge of each cluster is a visualisation artefact. The

cluster analysis divides the attractor into clean sectors around the two repelling fixed points. The centroids of the clusters associated with the fixed points can be considered as phase average with very coarse phase bins. The distance measure for the clustering algorithm is directly based on the Euclidean norm $\|\cdot\|$ as defined in (2.5) for the POD coefficients.

The evolution of the variables is displayed in figure 3(*b–d*). The colour marking in the time series plots corresponds to different groups of clusters (T in orange, E^+ in black and E^- in grey) which are explained below. Note that these do not represent each cluster as in figure 2. The first two variables, x and y, clearly show the location of the trajectory on the attractor. Extrema for positive x and y mark one ear, and for negative values the other ear. In figure 3(*a*), the cluster assignment of each snapshot is visualised. It reveals three groups of clusters, namely clusters $k = 1, 2$ (group 'T'), $k = 3, 4, 5, 6$ (group 'E^+') and $k = 7, 8, 9, 10$ (group 'E^-'). Within a group, the transitions are from k to $k+1$ for T and E^+ and from k to $k-1$ for E^-. The two ears around the fixed points are partitioned in two different groups of clusters, E^+ and E^-. After each circulation (E^+ or E^-), the trajectory returns to group T. Group T serves as a branching region between the two ears. The colour marking confirms this conclusion. These visualisations of the clustering provide revealing insights on the nature of the attractor without assuming any prior knowledge of it.

The distance matrix (see figure 4*a*) also suggests three groups, $k = 1, 2$, $k = 3, 4, 5, 6$ and $k = 7, 8, 9, 10$ in which the centroids within each region are comparably close. These regions correspond to the groups, T, E^+ and E^-, identified above and visualised with square boxes.

The dynamical behaviour between those groups and each cluster can be studied via the CTM displayed in figure 4(*b*). The high probabilities in the principal diagonal indicate a lingering in the cluster. This is due to the small time step Δt and the large size of the clusters. For instance, let the initial probability distribution be $\boldsymbol{p}^0 = [1, 0, \ldots, 0]^{\mathrm{T}}$, i.e. only occupies state $k = 1$. Then, the probability $P_{11} \approx 1$ is high that a snapshot will remain in this cluster. With lower probability it can move to the states $k = 2$, 3 or 10. Assuming now a unit probability of state 3 for $\boldsymbol{p}^0$, in the next time step, it can stay there or it moves forward in the direction $3 \to 4$, then $4 \to 5$ and $5 \to 6$. In $k = 6$ the probability is large for returning to the initial state 1. We have already identified above the states $k = 3, 4, 5, 6$ as the oscillation around one ear of the Lorenz attractor. The other ear resembled by $k = 10, 9, 8, 7$ can be reached by moving from $k = 1$ to $k = 10$. Note that we use the reverse numbering to accentuate the direction of the change of states. From the last state $k = 7$ it returns to the first state which is always passed after each orbit. An intuitive picture of the CTM and its state transitions is provided in figure 5. States are displayed as colour-marked bullets, small colour-marked dots are those instants where the trajectory is in the corresponding cluster.

In general, the CTM gives information on the nature of the states and their relation. Circulation regions are represented by non-zero main P_{jj} and subdiagonal entries $P_{j+1,j}$ (or $P_{j,j+1}$ depending on the ear). An additional non-zero entry that connects the last state of a group with state $k = 1$ is necessary for a full orbit. Distinct states that allow crossing from different groups such as E^+ and E^- are also identifiable. The state $k = 1$ serves as the connecting point between the two ears. From this state, the trajectories can branch out to the orbits, therefore this state is called a branching point.

The dynamical behaviour of the cluster probability vector can be studied from figure 6 via powers of the CTM (2.21). The cluster dynamics yield information on how the trajectory evolves in time on the attractor. The entries of P^l_{jk} give the

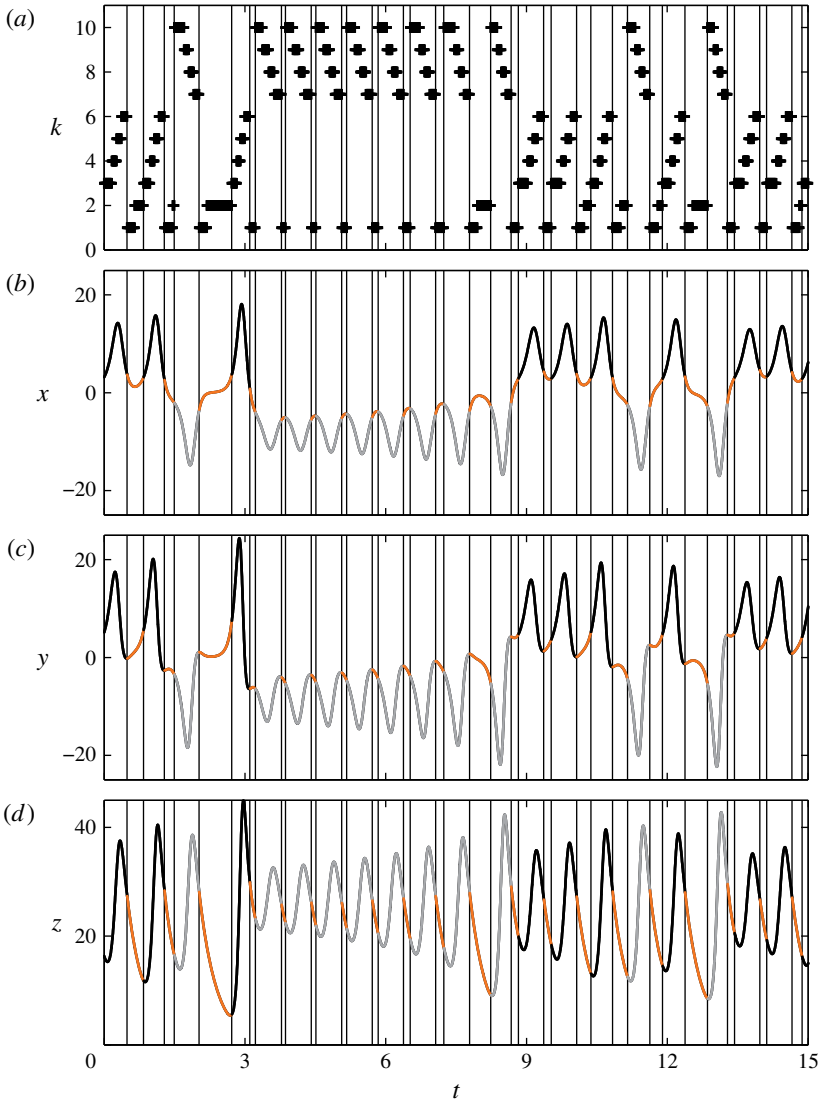

FIGURE 3. (Colour online) Temporal cluster analysis of the Lorenz attractor: (*a*) cluster assignment of each snapshot to a cluster k; (*b–d*) the time series for all three variables x, y, and z. The vertical lines separate between the branching clusters T ($k = 1, 2$), and the two ears (group E^+ with clusters $k = 3, 4, 5, 6$ and group E^- with $k = 7, 8, 9, 10$). In the time series plots, T is colour marked in orange, E^+ in black and E^- in grey.

probability to move from state k to state j in l time steps. After $l = 10$ iterations of the CTM, the different regions are still visible. As an example, we assume unit probability in state $k = 3$, i.e. $p_k^0 = \delta_{k3}$, which is the first cluster of orbit E^+. After $l = 10$ time steps the final transitions $3 \to 4$, $3 \to 5$ and $3 \to 6$ are possible. Since it is not possible to transition to the intermediate state $k = 1$, $l = 10$ is not sufficient to complete a full orbit. On the other hand, if we assume $p_k^0 = \delta_{k5}$, transitions are possible to state $k = 1$, to state $k = 3$ (the begin of another oscillation around the same orbit) and to state $k = 10$ which marks the starting point of the other orbit E^-.

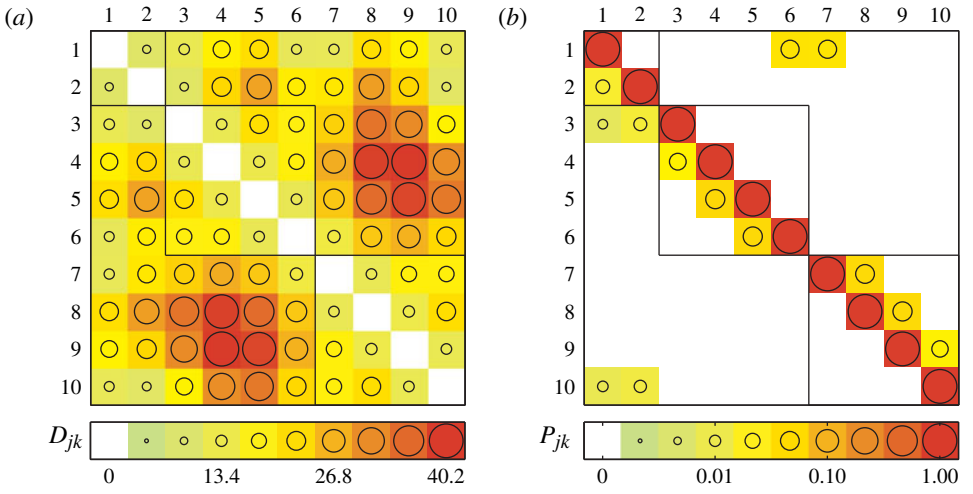

FIGURE 4. (Colour online) Kinematics and dynamics of the clustered Lorenz attractor. The geometry of the clusters is depicted with (*a*) a distance matrix $\boldsymbol{D}$ and the dynamical evolution is quantified by (*b*) the CTM $\boldsymbol{P}$. All matrix elements are non-negative. The value is depicted by background colour and increases with the radius of the corresponding circle. Note that the scale for the distance matrix is linear, while it is logarithmic for the transition matrix. The squares mark distances and transitions in the cluster groups T ($k = 1, 2$), E^+ ($k = 3, 4, 5, 6$) and E^- ($k = 7, 8, 9, 10$).

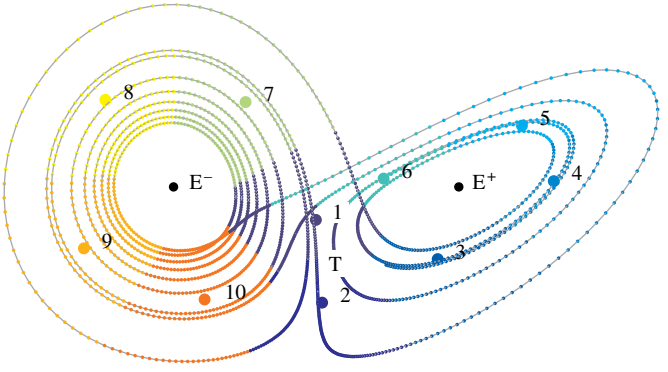

FIGURE 5. (Colour online) Cluster transitions of the Lorenz attractor. Cluster centroids are displayed as coloured bullets, the black bullets are the two unstable fixed points. The trajectory is displayed as a grey line with coloured dots as discrete states. For a better visibility of the clusters a non-symmetric perspective has been chosen. Three qualitatively different groups of clusters can be distinguished: the cycle groups of the ears E^+ and E^- as well as the transition clusters T.

Thus, we can conclude, a full circulation around one fixed point needs more than 10 time steps. In general, the probabilities on the principal diagonal have decreased as a benefit to the subdiagonals along the possible paths and scaled with the distance. With an increasing number of iterations the transition probabilities smooth out, and the transition matrix eventually converges to $\boldsymbol{P}^\infty$ (see figure 6*b*).

In figure 7(*a*) the asymptotic cluster probabilities $\boldsymbol{p}^\infty$ are compared with the distribution from the data $\boldsymbol{q}$. The asymptotic probability distribution is inhomogeneous.

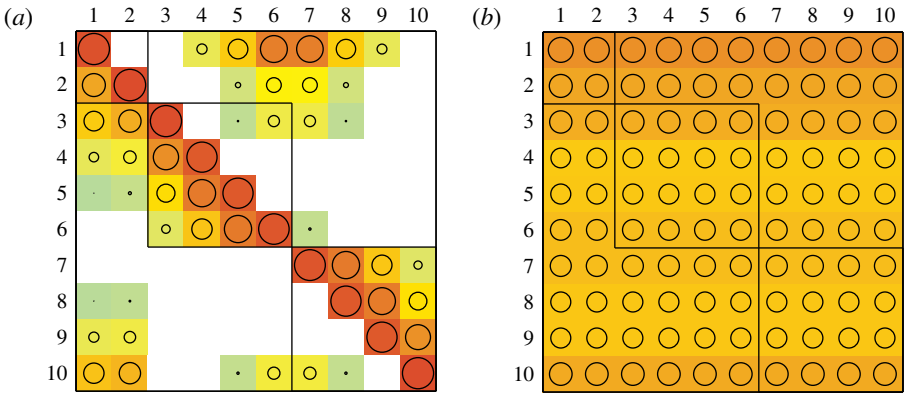

FIGURE 6. (Colour online) Dynamics of the clustered Lorenz attractor illustrating (*a*) $l=10$ and (*b*) $l=1000$ iterations of the transition matrix $\boldsymbol{P}$. The matrices are visualised like in figure 4(*b*). Note that the increasing iterations make the matrices $\boldsymbol{P}^l$ more uniform.

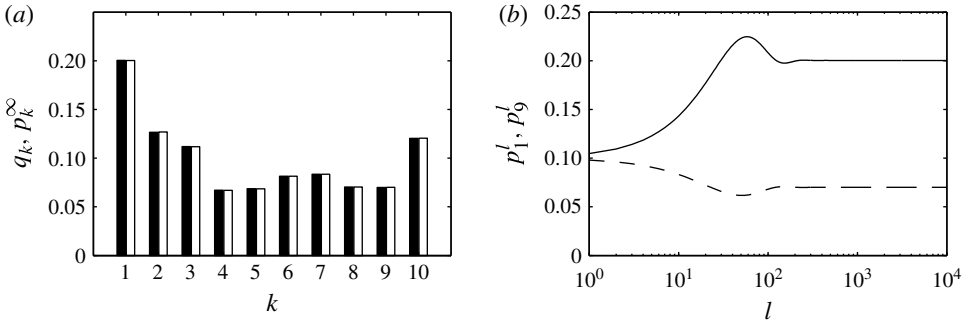

FIGURE 7. Cluster probability distributions of the Lorenz attractor. (*a*) The probability distributions from the data $\boldsymbol{q}$ (2.12) (solid rectangles) and from converged iteration of the transition dynamics $\boldsymbol{p}^\infty$ (open rectangles) (2.22). (*b*) The convergence of cluster probabilities p_1^l (solid line) and p_9^l (dashed line) with number of iterations l. The initial condition $\boldsymbol{p}^0$ was from equipartition of probabilities.

The largest probability q_1 results from the distinct position this cluster has. As a connection between the two ears this state should be passed about twice as often than the others. The lowest ones, q_4, q_5 and q_8, q_9, correspond to the outer clusters of the two orbits. *Pars pro toto*, the evolution of two cluster probabilities, the worst defined as $\max|p_k^{10}-q_k|$ and the best defined as $\min|p_k^{10}-q_k|$, p_1^l and p_9^l, respectively, are displayed in figure 7(*b*). Although the initial difference $|p_k^{10}-q_k|$ is larger for p_1, both converge at $l=10^2$.

Another way to access the asymptotic probability distribution is the stability analysis of the CTM (see figure 8). As expected by the theory (see § 2.2.2), the complex spectrum yields a dominant eigenvalue $+1$. The corresponding eigenvector $\boldsymbol{p}_1^{ev}$ is the only one that survives after sufficiently long iterations and, since it is the fixed point of the propagator, $\boldsymbol{p}^\infty$ has to be equal to $\boldsymbol{p}_1^{ev}$. Moreover, the rate of convergence at which $\boldsymbol{p}^\infty$ is approached is given by the second largest eigenvalue modulus as mentioned in § 2.2.2 and shown below.

The geometric nature of the discretised attractor allows us to assess the quality of the discretisation and to identify cluster groups and transition regions. The similar

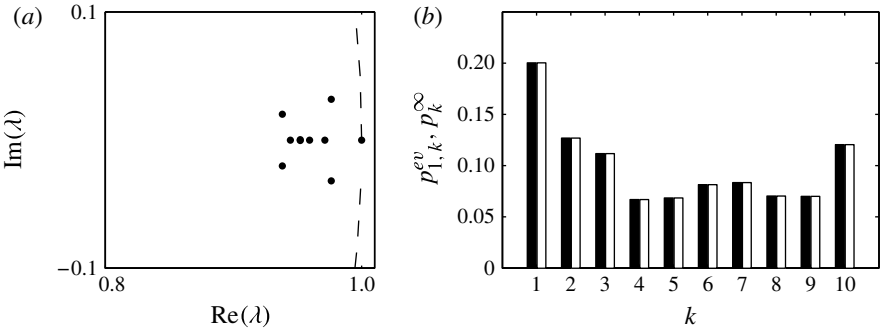

FIGURE 8. Stability analysis of the transition matrix associated with the clustered Lorenz attractor: (*a*) complex spectrum and (*b*) eigenvector $\boldsymbol{p}_1^{ev}$ associated with the dominant eigenvalue $+1$. Note that all eigenvalues are real or complex conjugate pairs on or within the unit circle (dashed line). The components of the eigenvector $\boldsymbol{p}_1^{ev}$ are all non-negative, normalised to unit probability and visualised in a histogram (*b*) (solid rectangles). Note the (expected) agreement between this probability distribution and the converged distribution $\boldsymbol{p}^{\infty}$ (open rectangles).

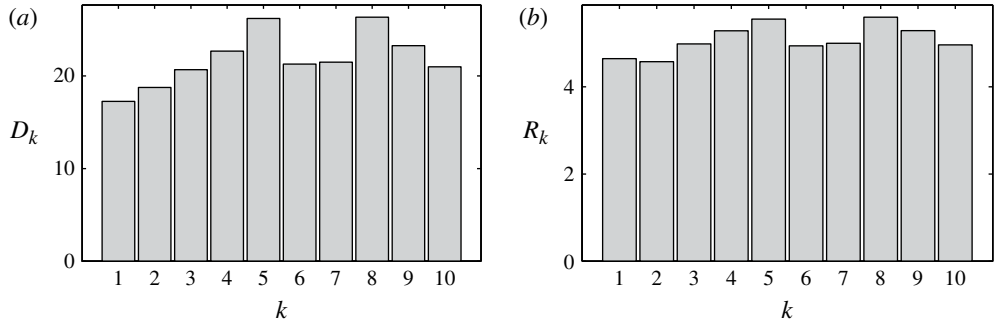

FIGURE 9. (*a*) Diameters D_k defined in (2.27) and (*b*) standard deviation R_k as defined in (2.28) of clustered Lorenz attractor. Evidently, the clusters have similar sizes.

diameters and standard deviations of the clusters, as shown in figure 9, confirm a nearly homogeneous partition of the state space.

The dynamics on the attractor can be further analysed by inspecting the evolution of attractor properties (see figure 10). As mentioned above, the rate of convergence is given by the second largest eigenvalue modulus. The power iterations of this quantity are shown in figure 10(*a*), it converges to zero at approximately $l = 10^2$.

In figure 10(*b*) the variance is displayed for different initial probability conditions. Initial distributions $p_k^0 = \delta_{k1}$, $p_k^0 = \delta_{k2}$, etc., are displayed in solid lines, and dashed lines indicates equipartition. For all of these initial conditions except for the initial equipartition, the spreading increases until it reaches a finite value. Initial distributions with $p_k^0 = \delta_{k1}$ or δ_{k2} (dark colour) exhibit a large spreading and overshoot which can be linked to their special role as a branching region. All curves converge after $l = 10^2$ iterations.

The evolution of the Kullback–Leibler entropy (see figure 10*c*) offers important insights about the Markov model. As expected, the Kullback–Leibler entropy (2.34) of the iterated CTM $\boldsymbol{P}^l$ increases with the iteration l, i.e. information is lost by

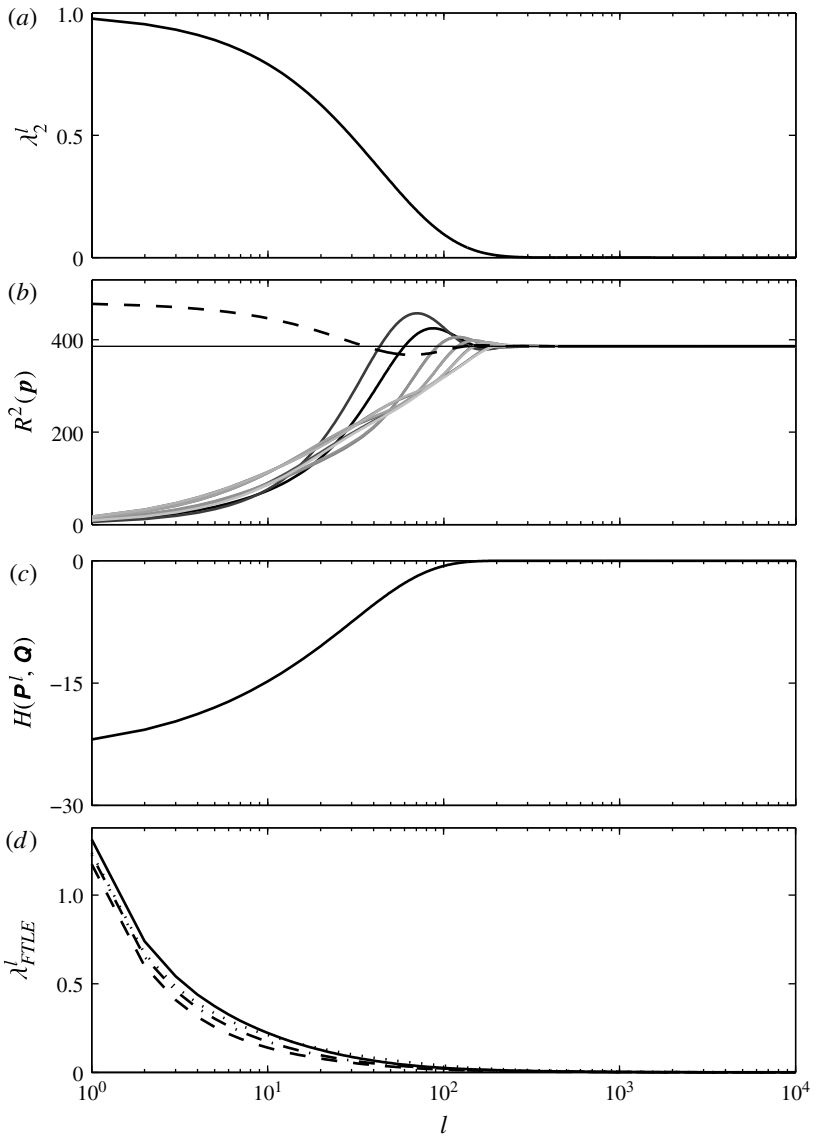

FIGURE 10. CROM convergence study for the Lorenz attractor. Convergence with respect to the number of iterations l is illustrated (*a*) for the second eigenvalue modulus $|\lambda|_2^l = |0.9762 + 0.0318i|^l$ of $\boldsymbol{P}$, (*b*) for the variance $R^2(\boldsymbol{p})$ (2.30) for different initial distributions of (*i*) the equipartition (dashed line) or (*ii*) a pure cluster $p_k^0 = \delta_{jk}$ with $j = 1$ or $j = 2$ is a cluster of the transition region (black) and $j = 3, \ldots, 10$ is a cluster of E^- or E^+ (light grey), (*c*) for the Kullback–Leibler entropy H (2.34), and (*d*) for the FTLE λ_{FTLE}^l (2.31) for different initial distances of the Lorenz attractor. Evidently, all quantities converge at approximately $l = 10^2$, i.e. each of the measures can be used as a convergence indicator.

irreversible diffusion. The entropy converges around $l \approx 10^2$, when the iterated CTM describes the l-independent asymptotic probability $\boldsymbol{p}^\infty$. A small negative value of H characterises the small difference between the inferred cluster probability distribution $\boldsymbol{q}$ of the data and the fixed point $\boldsymbol{p}^\infty$. Intriguingly, there exists an iteration interval

$l < 10$ with roughly constant H. In this regime, the iterated CTM is well approximated by a linear Taylor expansion from the continuous-time Markov model (2.24), i.e.

$$\boldsymbol{P}^l = \exp(\boldsymbol{P}^{cont} l\Delta t) \approx \boldsymbol{I} + \boldsymbol{P}^{cont} l\Delta t + O(\Delta t^2). \tag{3.2}$$

This implies that the information gained from $\boldsymbol{P}^l$ are comparable for different l. In other words, the Kullback–Leibler entropy does not only indicate the converged state of the iterations and a difference between Markov model and data. The entropy also indicates the validity regime of the Euler scheme for the continuous-time Markov model (2.24).

The evolution of the FTLE is depicted in figure 10(*d*). Initially distant trajectories eventually converge to the asymptotic distribution $\boldsymbol{p}^\infty$. All four properties converge at $l \approx 10^2$ and any of them can serve as an adequate measure of convergence.

Summarising, the CROM strategy segments the Lorenz attractor into nearly homogeneous clusters which are arranged like pieces of a pie around the unstable fixed points. Without employing any prior knowledge of the nature of the attractor, the CTM yields the two characteristic orbits, and identifies the states that behave as a branching region. The convergence behaviour is in good agreement with the result of the statistical analysis of the clustered data. Increasing the number K of clusters does not change the main conclusions. The number of clusters in the two ears (state resolution) does increase and fine-structure details of the branching process can be observed. Note that in the special case $K = 2$ CROM yields two clusters, each resolving one ear. Variations of the snapshot time step Δt have no strong effect on the observed oscillation and transition mechanisms. The transition probabilities between different clusters depend, in the first approximation, linearly on the time step.

4. Mixing layer model

CROM is applied to a two-dimensional incompressible mixing layer with Kelvin–Helmholtz vortices undergoing vortex pairing. This flow is computed with a direct numerical simulation (DNS) of the Navier–Stokes equation. The velocity ratio is $r = U_1/U_2 = 3$ where U_1 and U_2 denote the upper (fast) and lower (slow) stream velocity. The Reynolds number $Re = \Delta U \delta_\omega/\nu = 500$ is based on the velocity difference $\Delta U = U_1 - U_2$, the initial vorticity thickness δ_ω and the kinematic viscosity ν. A tanh profile with stochastic perturbation is chosen as the inlet profile. A snapshot ensemble of $M = 2000$ with a sampling time of $\Delta t = 1$, which is non-dimensionalised with respect to U_1 and δ_ω, is provided by a finite difference Navier–Stokes solver. Details of this solver can be found in Daviller (2010) and Cavalieri *et al.* (2011). The computational domain is $140\,\delta_\omega$ long and $56\,\delta_\omega$ high. The resolution of the spatial discretisation increases in the mixing region.

The mixing layer exhibits the typical roll-up of initial Kelvin–Helmholtz vortices and vortex pairing events further downstream. A vorticity snapshot is displayed in figure 11. In the flow simulations, vortex merging of two and three vortices occur which results into a large spreading of the shear layer. Moreover, four vortices may merge into two pairs which will form a single structure further downstream.

The flow data is compressed with a POD from the snapshot ensemble. The POD modes are visualised in figure 12 via isocontours of the transversal velocity component and colour plots of the vorticity. The first two mode pairs ($i = 1, 2, 3, 4$) represent the convection of nearly periodic vortex structures in the streamwise direction (Rajaee, Karlsson & Sirovich 1994). The first mode pair ($i = 1, 2$) represents flow structures

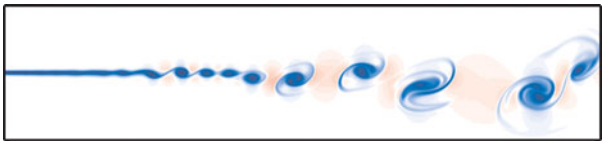

FIGURE 11. (Colour online) Two-dimensional DNS simulation of the mixing layer: colour plot of an instantaneous vorticity distribution.

associated with the subharmonic instability, while the second mode pair ($i = 3, 4$) corresponds to Kelvin–Helmholtz vortices associated with a higher frequency and increasing wavenumber in streamwise direction (Noack *et al.* 2004; Laizet, Lardeau & Lamballais 2010). Higher-order modes contain a mixture of different frequencies.

The first four temporal POD coefficients are displayed in figure 13(*b–e*). A simultaneous decrease in the amplitudes of the first mode pair and increase in the amplitudes of the second mode pair is an indicator for a change of the flow from vortex pairing to predominant Kelvin–Helmholtz vortices.

For the reasons discussed in § 2, the cluster analysis is performed on the POD coefficient vectors $\boldsymbol{a}^m$, $m = 1, \ldots, M$, with $\boldsymbol{a} = [a_1, \ldots, a_{M-1}]^{\mathrm{T}}$. We chose the number of clusters $K_c = 10$, like in the Lorenz attractor. As for the Lorenz attractor, the choice of $K_c = 10$ is in good agreement with the 'elbow criterion' as explained in § 2.1.2. The assignment of each snapshot to the clusters is shown in figure 13(*a*). Vortex-pairing-dominated flows and Kelvin–Helmholtz vortices are clearly separated into different groups of clusters, called VP with $k = 5, 6, 7, 8, 9, 10$ for the former and KH with $k = 2, 3, 4$ for the latter. The two dynamical regimes are mainly connected by cluster $k = 1$. In this case, the flow has to pass this state for the transition from one dynamical regime to the other.

The centroids associated with each cluster are displayed in figure 14. In the subharmonic group VP, the centroids $k + 1$ lead the centroids k by a sixth of the wavelength of the merged vortices. In the group KH, a similar behaviour can be observed but with shifts corresponding to about a third of the wavelength of the Kelvin–Helmholtz structures. The centroid $k = 1$ represents an intermediate state between those two groups. The wavenumber increases but is somewhat between the wavenumber corresponding to the subharmonic and the fundamental frequency. The order of the centroids is clearly aligned with the dynamical evolution of the flow structures. In contrast, the first four POD modes represent sine- and cosine-like wave pairs associated with the two shedding regimes. In the case of CROM, flow structures are resolved not by energy content, frequency and wavelength, but by phase. Moreover, transient states can be captured explicitly.

The distance matrix and the CTM (for both see figure 15) allow us to analyse the geometric relation of the centroids and their temporal transitions. The cluster groups VP, KH and the transition cluster $k = 1$ are visible in the distance matrix and in the CTM as will become clearer below. Furthermore, the inner cluster distances between the centroids of group VP are more distant to each other than the corresponding distances of group KH, since vortex pairing is associated with larger fluctuation levels than the more monochromatic Kelvin–Helmholtz shedding. The CTM shows an inner-group circulation for both, for example $5 \to 6 \to \cdots \to 10 \to 5 \to \cdots$, etc. Cluster $k = 1$ has a special significance as the 'flipper cluster' which acts as a switch between these groups. As elaborated above, this cluster resembles an intermediate state and links both regimes. All Kelvin–Helmholtz states can be accessed from the

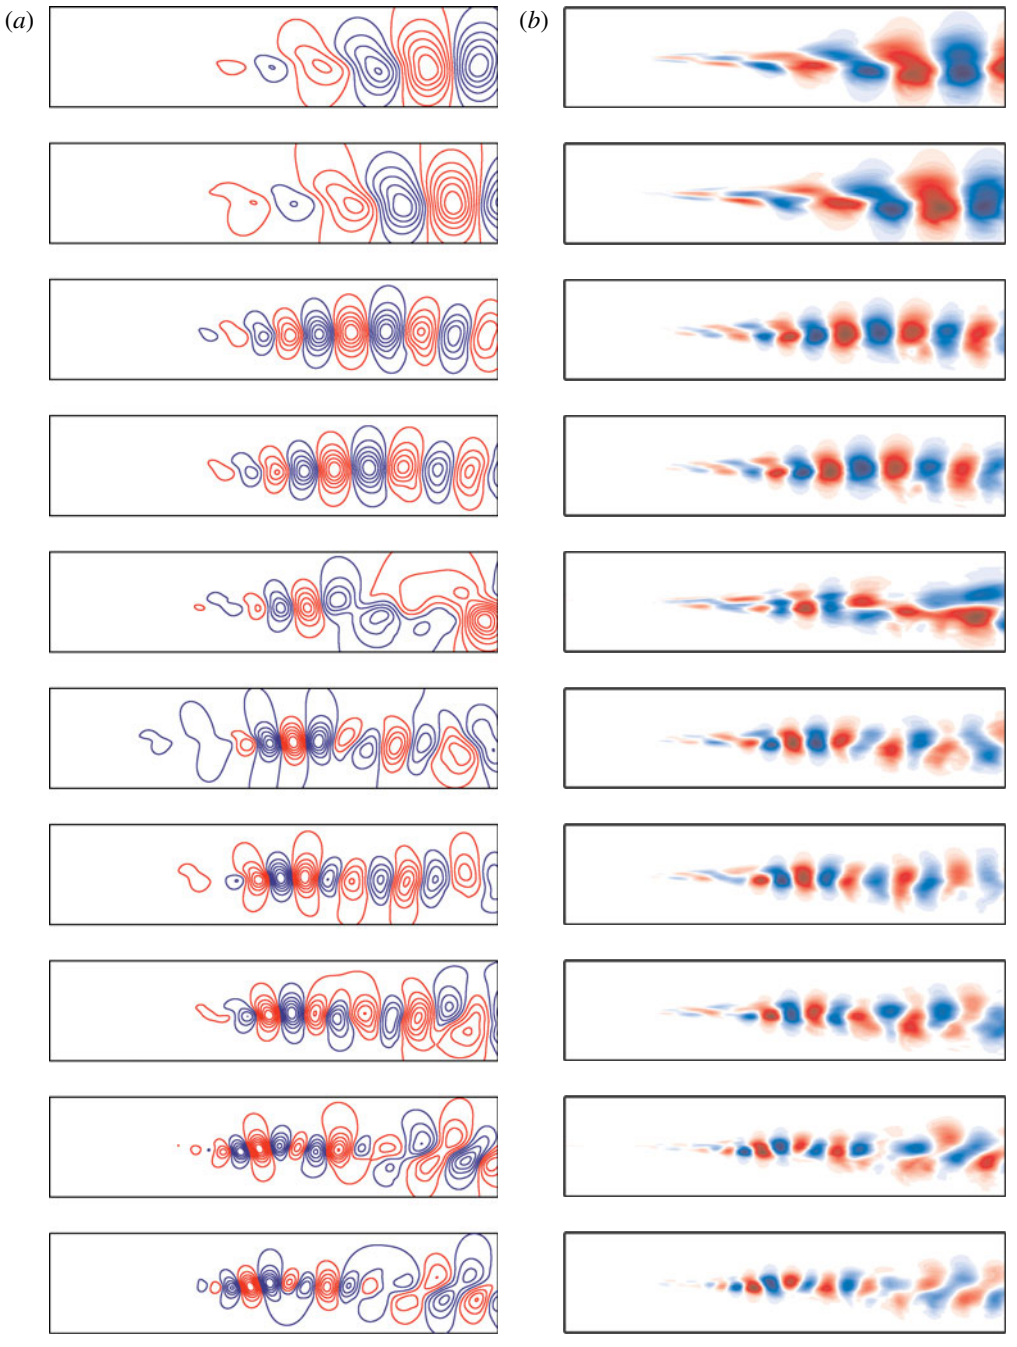

FIGURE 12. (Colour online) POD modes $\boldsymbol{u}_i$, $i = 1, \ldots, 10$ (top to bottom) of the mixing layer. The flow is visualised by isocontours of transversal velocity component (*a*) and by colour plots of the vorticity (*b*).

flipper cluster. Moreover, $k = 1$ can only be reached by states $k = 3$ from KH and $k = 8$ of group VP. A direct connection between the groups is solely possible by $2 \rightarrow 9$ in one direction. The sketch in figure 16 represents a simplified picture of

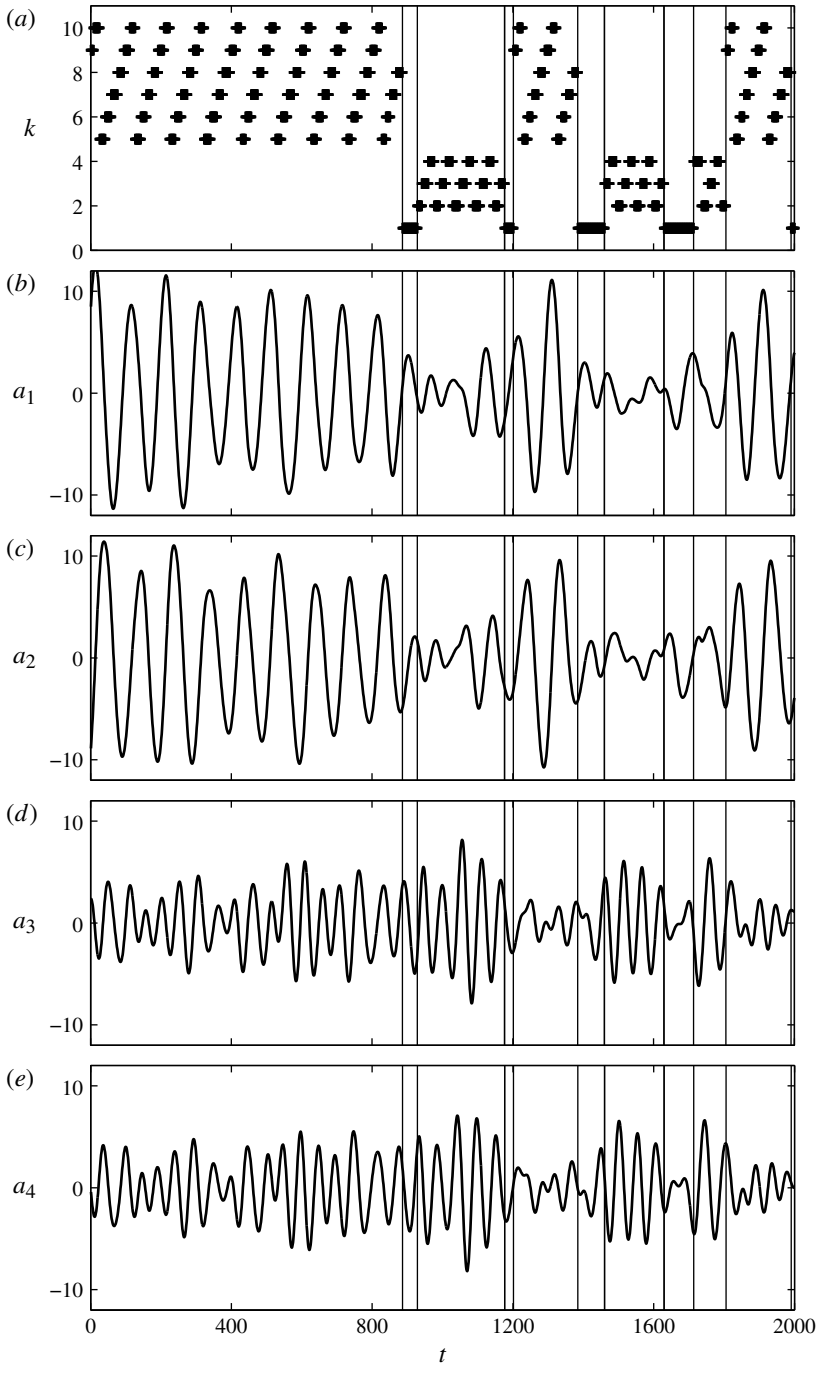

FIGURE 13. Temporal cluster analysis of the mixing layer. (*a*) The cluster assignment of each snapshot to a cluster; (*b–d*). POD coefficients of the first and second POD mode pairs a_i, $i = 1, 2, 3, 4$ as a function of time t. The vertical lines separate between Kelvin–Helmholtz vortices (cluster $k = 2, 3, 4$) and vortex pairing dominated flows (remaining clusters).

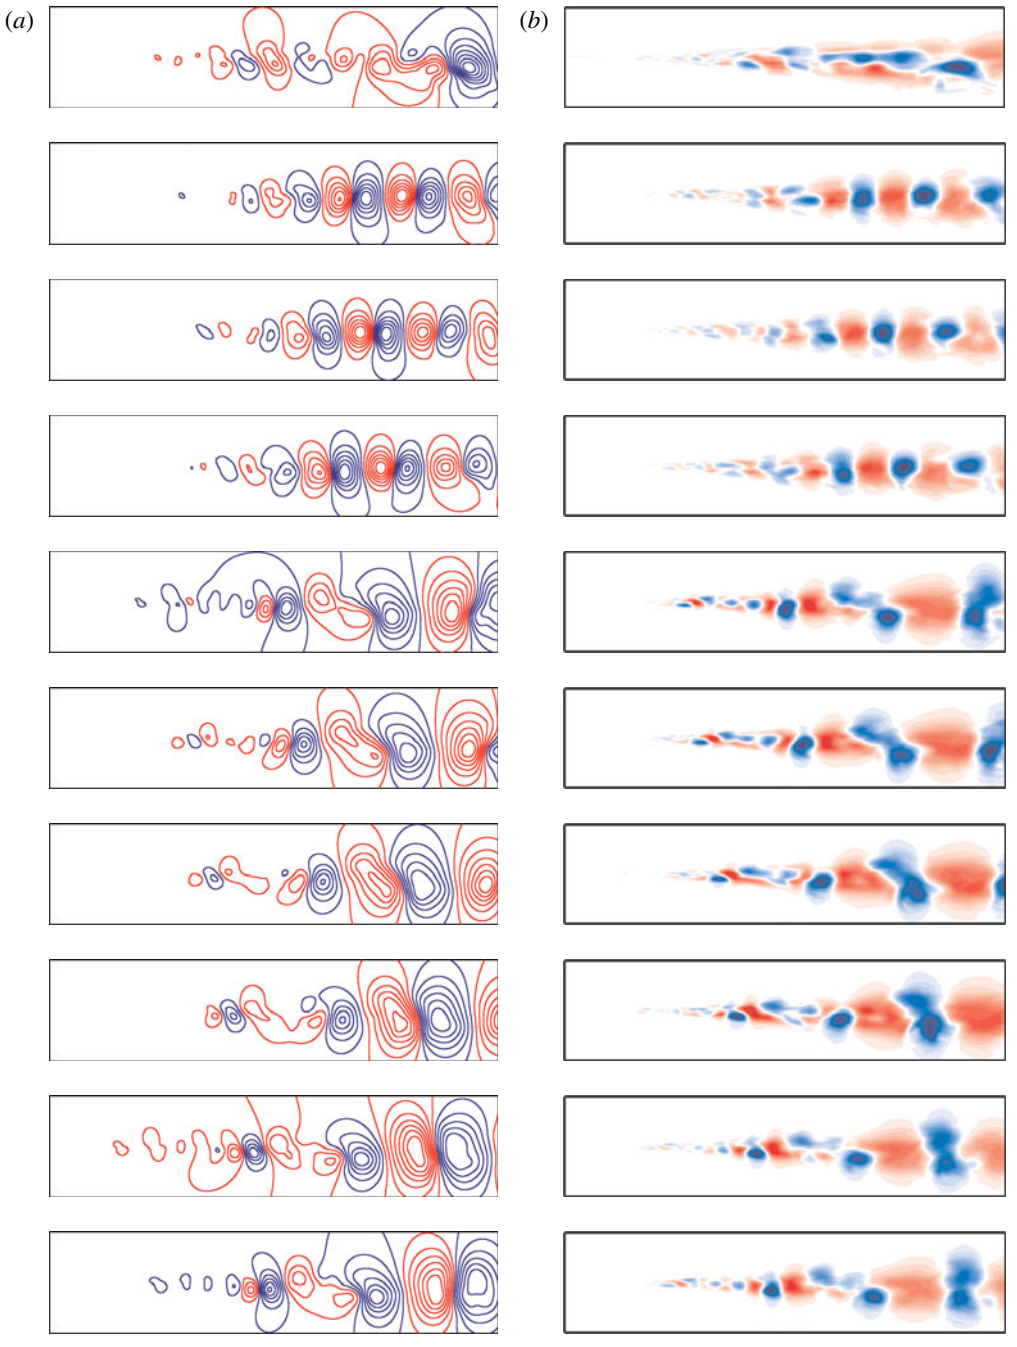

FIGURE 14. (Colour online) Cluster centroids $\boldsymbol{c}_k$, $k = 1, \ldots, 10$ (top to bottom) of the mixing layer. The flow is visualised like in figure 12.

the CTM. The most dominant transitions are displayed. The threshold for selection of P_{jk} is set to a level which is just large enough to fulfil the connectivity of all states. Summarising, the change from Kelvin–Helmholtz dominated flows to vortex pairing is possible by the crossing of $k = 1$ or by the direct transition to $2 \to 9$.

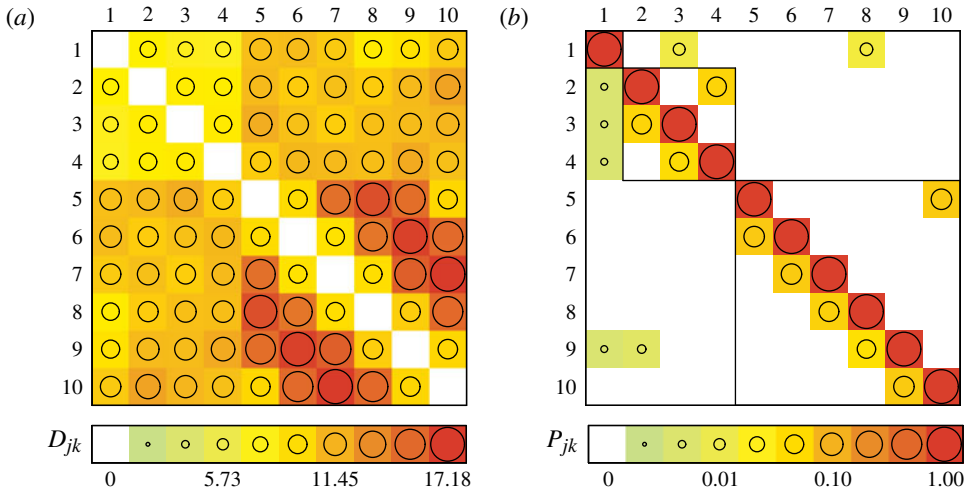

FIGURE 15. (Colour online) Kinematics and dynamics: (*a*) distance matrix $\boldsymbol{D}$ and (*b*) transition matrix $\boldsymbol{P}$ of the clustered mixing layer data (see figure 13). Values are displayed as in figure 4.

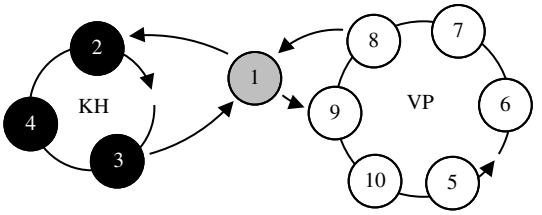

FIGURE 16. Simplified cluster transitions of the mixing layer. The arrows indicate possible transitions above a certain threshold. This threshold is chosen small enough to guarantee full connectivity. The graph highlights two cyclic cluster groups associated with Kelvin–Helmholtz shedding (black) and with vortex pairing (white). The grey cluster $k = 1$ serves as bifurcation point.

These observations may be important for control. If one wants to enforce or prevent pure Kelvin–Helmholtz vortices, the manipulation of the flipper cluster as effective bifurcation point appears as a particularly effective tool for control. Its manipulation or use as a precursor requires further analysis which is beyond the scope of this article. Nevertheless, a general CROM-based control strategy provided in appendix B seems promising for the exploitation of nonlinearities.

From the evolution in time of the transitions (see figure 17), we can deduce periodical behaviour in a certain direction similar to the Lorenz attractor. With increasing power l, the entries of $\boldsymbol{P}$ in each column approach the long-term probability. Each iteration of the CTM corresponds to the time step Δt.

For the sake of completeness, we give a short description of convergence behaviour and attractor properties. The cluster probabilities (see figure 18*a*) are not strongly biased to any state. The converged cluster probabilities belonging to the vortex pairing are slightly underestimated, whereas p_k^∞, $k = 1, 2, 3, 4$, are larger than the cluster probabilities based on the data. This difference can be attributed to discretisation error of the CTM. Only 1999 snapshot transitions are used to estimate 10×10 matrix elements.

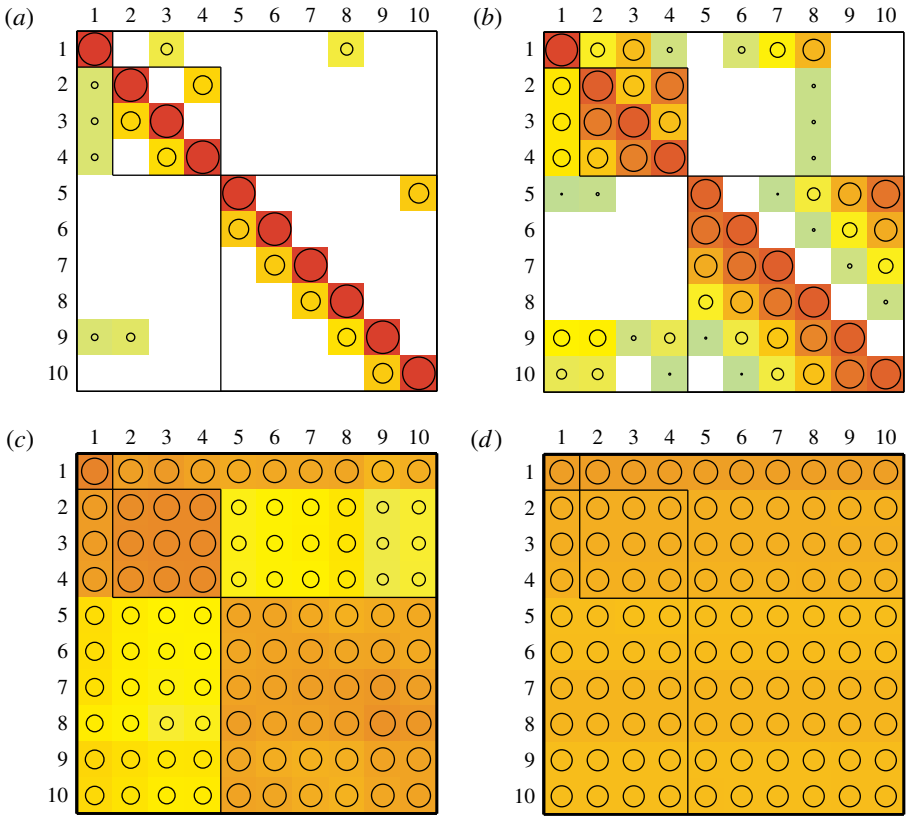

FIGURE 17. (Colour online) Cluster dynamics of the mixing layer. The evolution is quantified by $\boldsymbol{P}^l$ for (*a*) $l = 10^0$, (*b*) $l = 10^1$, (*c*) $l = 10^2$ and (*d*) $l = 10^3$.

In analogy to the CROM of the Lorenz attractor, the best cluster $k = 4$ is determined as the cluster with the minimum value of $|p_k^{10} - q_k|$, and the worst cluster $k = 6$ is defined as the cluster with the maximum value. For the mixing layer, these clusters, $k = 4$ and $k = 6$ in figure 18(*b*), converge both at $l \approx 10^3$.

As for the Lorenz attractor, the eigenvector associated with the dominant eigenvalue yields the converged cluster probabilities (see figure 19*b*).

The clustering of the snapshots gives a relatively homogeneous partition of the state space as confirmed by the cluster diameters and the cluster standard deviations (see figure 20*a*,*b*).

Figure 21 displays the convergence of the second eigenvalue modulus, the variance, the Kullback–Leibler entropy and the FTLE. All converge at approximately $l = 10^3$. The grey lines in figure 21(*b*) correspond to initial distributions with unit probability, $p_k^0 = \delta_{kj}$ for $j = 5, 6, 7, 8, 9, 10$, in one of the vortex pairing states. These distributions spread out faster and have a larger variance at an intermediate stage than those belonging to Kelvin–Helmholtz states which are displayed in black lines. The variance of an initial equipartition (dashed lines) is already close to the converged value. Nevertheless, all distributions converge after the same number of iterations. The above results confirm a good state space partition and convergence behaviour of the transition matrix.

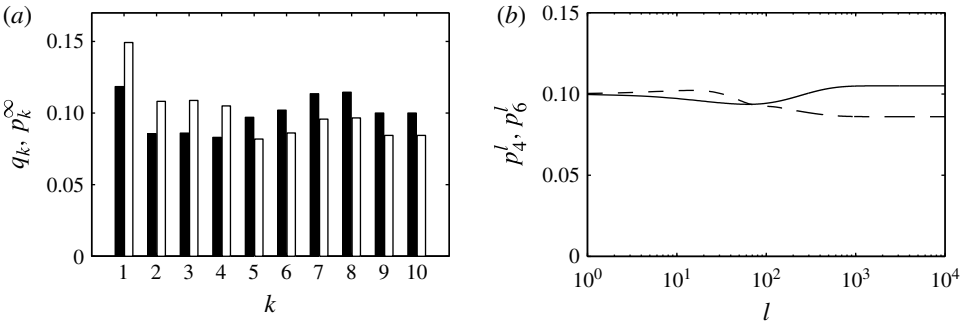

FIGURE 18. Cluster probability distributions of the mixing layer (like figure 7 for the Lorenz attractor). (*a*) The probability distributions from the data $\boldsymbol{q}$ (2.12) (solid rectangles) and from converged iteration of the transition dynamics $\boldsymbol{p}^\infty$ (open rectangles) (2.22). (*b*) The convergence of cluster probabilities p_4^l (solid line) and p_6^l (dashed line).

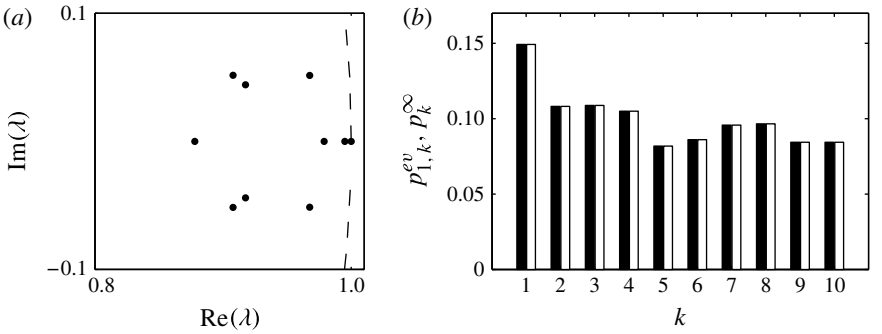

FIGURE 19. Stability analysis of the transition matrix associated with the mixing layer: (*a*) complex spectrum and (*b*) eigenvector $\boldsymbol{p}_1^{ev}$ associated with the dominant eigenvalue $+1$. Note that all eigenvalues are real or complex conjugate pairs on or within the unit circle (dashed line). The components of the eigenvector $\boldsymbol{p}_1^{ev}$ are all non-negative, normalised to unit probability and visualised in a histogram (solid rectangles). Note the (expected) similarity between this probability distribution and the converged distribution $\boldsymbol{p}^\infty$ (open rectangles).

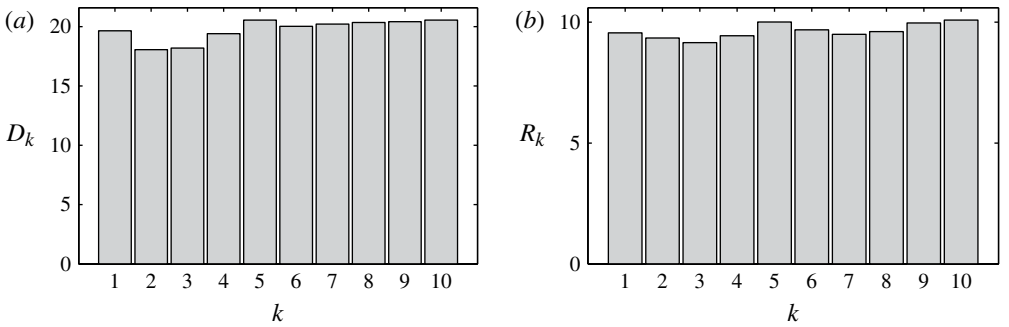

FIGURE 20. (*a*) Diameters D_k defined in (2.27) and (*b*) standard deviation R_k as defined in (2.28) like figure 9, but for the mixing layer. Note that the state space partition is relatively homogeneous.

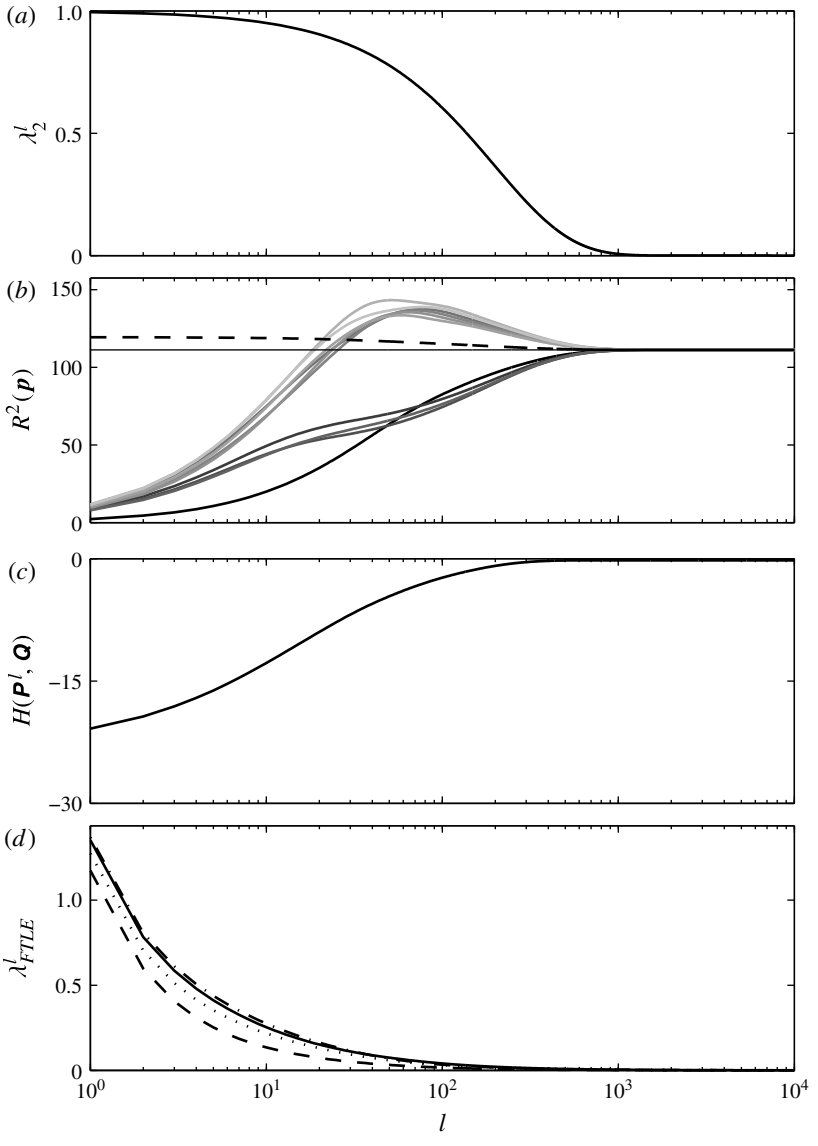

FIGURE 21. Analogously to figure 10, CROM convergence study for the mixing layer. Convergence with respect to the number of iterations l is illustrated (a) for the second eigenvalue modulus $|\lambda_2|^l = |0.995|^l$ of $\boldsymbol{P}$, (b) for the variance R^2 (2.30) for different initial distributions of a pure cluster $p_k^0 = \delta_{jk}$ with $j = 1$ (black), j is a cluster of 'VP' (light grey) and j is a cluster of 'KH' (dark grey), (c) for the Kullback–Leibler entropy H (2.34) and (d) for the FTLE λ_{FTLE}^l (2.31) for different initial distances of the mixing layer. Note that the converged value for R^2 is indicated with the horizontal thin line. Evidently all quantities converge at approximately $l = 10^3$, i.e. each of the measures can be used as a convergence indicator.

In general, the performance of a model can be evaluated by how well and how long it is able to predict future states. Here, we are interested in how long the model is capable of predicting future states from the initial conditions. We can conclude from

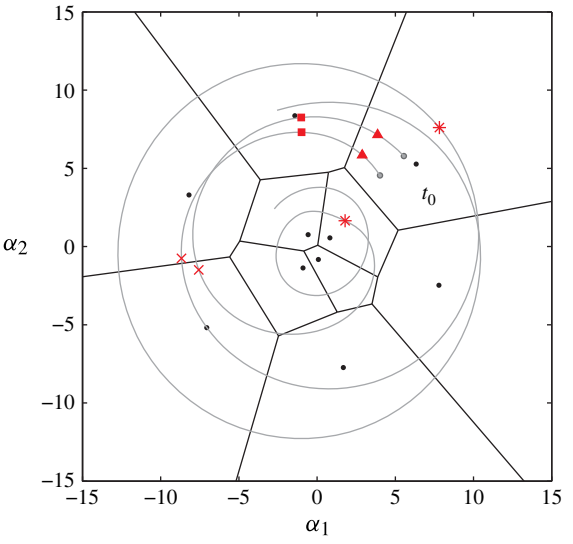

FIGURE 22. (Colour online) Prediction horizon of the mixing layer: two different initially close trajectories (grey lines) obtained from the integration of the Galerkin system with a nonlinear eddy-viscosity model as in Cordier *et al.* (2013) for $N = 20$ are displayed. The initial conditions are visualised as ●. The prediction horizon t_H defined in (4.1) is displayed as $*$. Other values correspond to $(1/4)\,t_H$(▲), $(1/2)\,t_H$(■) and $(3/4)\,t_H$(×). The clusters are displayed as a Voronoi diagram (black lines) where ● corresponds to a centroid. The two-dimensional visualisation is obtained as explained in appendix D. The diffusion of the probability distribution is related to the divergence of trajectories of the dynamical system.

the convergence study that the prediction horizon must be related to the number of iterations until the asymptotic state is reached. In the following, we define the model's prediction horizon in terms of the Kullback–Leibler entropy as we have shown it is a suitable measure for the convergence of the model for the Lorenz attractor as well as for the mixing layer (compare figures 10 and 21).

In figure 22, as an example, two trajectories (grey lines) are displayed for two different but closely located initial conditions (grey bullets) obtained from the integration of the Galerkin system for the mixing layer (Cordier *et al.* 2013). The thin black lines display the borders of the Voronoi cells as obtained from the clustering of the data and the black bullets represent the corresponding centroids. For the two-dimensional visualisation, the centroids and the data points are projected onto the first two POD modes obtained from the centroids' covariance matrix as described in appendix D.

Let the prediction horizon $t_H := l_{0.9}$ be defined by

$$H(\boldsymbol{P}^{t_H}, \boldsymbol{P}^{\infty}) - H(\boldsymbol{P}^{1}, \boldsymbol{P}^{\infty}) \approx 0.9\,(H(\boldsymbol{P}^{\infty}, \boldsymbol{P}^{\infty}) - H(\boldsymbol{P}^{1}, \boldsymbol{P}^{\infty})) \tag{4.1}$$

in terms of the Kullback–Leibler entropy $H(\boldsymbol{P}^{l}, \boldsymbol{P}^{\infty})$ with the transition matrix $\boldsymbol{P}$ as determined for the mixing layer (see figure 15*b*). Then, the red (grey) markers correspond to $\gamma\, t_H$ with $\gamma = 1/4, 1/2, 3/4, 1$.

The initial conditions of the two trajectories are in the same cluster which is associated with one shedding regime. The trajectories evolve similarly in the state

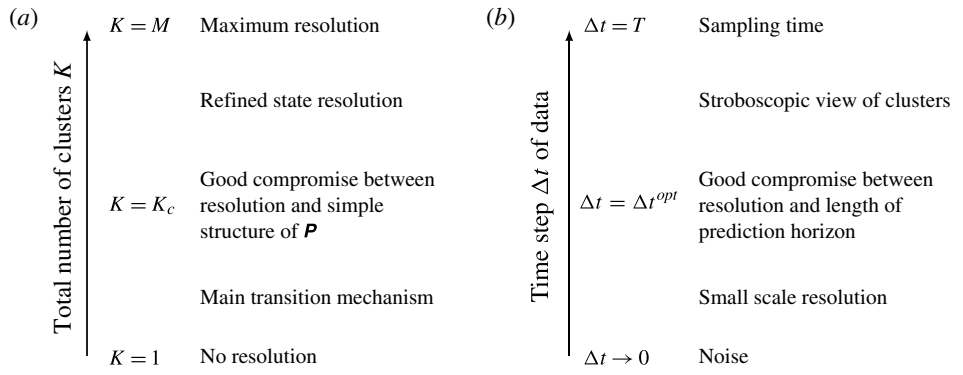

FIGURE 23. Impact of parameter variation: the influence of (*a*) the number of clusters K and (*b*) the time step Δt of the data on the resulting CROM. The time length T (physical time length of measurement) of the data is assumed to be constant which allows a variation of the number of observations M as a result of a variation of Δt.

space until about $(3/4)\, t_H$ when they start to diverge and one trajectory switches to the other shedding regime. This divergence of trajectories is associated with a diffusion of the probability distribution. After exceeding the prediction horizon, the probability distribution has converged to the model's asymptotic state and initial conditions, i.e. the initial cluster cannot be recovered anymore from backwards iteration. In the asymptotic state, the invariant distribution corresponds to the fraction of time the trajectory spends in each cluster.

The diffusion of probability, which is directly linked to the prediction horizon, is induced by the cluster size and the time step. Figure 23 serves as a guide for the following discussion.

For a fixed number of clusters K the prediction horizon can be optimised by varying the time step Δt of the data. The trajectory spends a long time in each cluster and only few transitions to other clusters occur. The large ratio of inter-cluster transitions to intra-cluster transitions introduces uncertainty in the model. This ratio can be decreased by a suitable choice for the time step such that, ideally, a stroboscopic view of the clusters is obtained. We consider as example a limit cycle which is discretised into K phase bins of equal size with centroids $\boldsymbol{c}_k$ corresponding to the phase $2\pi k/K$, $k = 1, \ldots, K$. Then, the ideal time step would be $2\pi/K$ for which CROM yields a transition matrix with minimised diffusion ($P_{jj} \to 0$ while $P_{jj+1} \to 1$ for $j = 1, \ldots, K-1$) and maximised prediction horizon. A necessary upper bound for Δt is determined by the required time resolution of the data for the estimation of the CTM. The situation is more complex if several frequencies occur in the data for which the optimal time step Δt is a good compromise between a stroboscopic view for one frequency and a sufficient resolution of the higher frequency.

The prediction horizon can also be optimised with respect to the chosen number of clusters K. A larger number of clusters reduces the ratio of inter-cluster transitions to intra-cluster transitions resulting in a longer prediction horizon. The lower bound is given by $K = 1$ for which all snapshots $\boldsymbol{u}^m$ are in one cluster and its centroid represents the mean flow $\boldsymbol{u}_0$. It should be noted that K is limited by the number of snapshots M, as $K \times K$ CTM elements are estimated from $M - 1$ snapshot transitions. The resulting CTM for $K = M$ is completely deterministic having entries $P_{jj+1} = 1$ for $j = 1, \ldots, K-1$. Its last state $\boldsymbol{c}_K = \boldsymbol{c}_M$ is absorbing, i.e. it is not possible to leave this

state, since there is no transition to the first snapshot of the data. A good compromise between the resolution of the main transition mechanism and a refined state resolution can be found by the elbow criterion.

In summary, CROM resolves the two prominent dynamical regimes of the mixing layer, Kelvin–Helmholtz vortices and a flow governed by vortex pairing. The centroids give a revealing picture of the evolution of the flow structures and the intermediate role of a 'flipper cluster' is discovered. The CROM naturally filters out the dominant wavenumbers associated with the characteristic features of the flow. The data partitioning is relatively homogeneous and the convergence behaviour is in good agreement with the result of the statistical analysis of the clustered data. Increasing the cluster number K does not change the principal picture of two dynamical regimes with a transition region. However, one may obtain a refined resolution of the metamorphosis from one state to the other. CROM has a built-in uncertainty due to the cluster size which results in a diffusion of probability. As a consequence, the model converges to an asymptotic state which defines its prediction horizon. The snapshot time step Δt has no strong effect on the observed oscillation and transition mechanisms as long as it is small compared with the characteristic period. However, adjusting the time step may improve the prediction horizon. In contrast, POD models can be integrated infinitely long after exceeding their prediction horizon leading to erroneous model predictions. CROM's asymptotic state resembles the discretised ergodic measure on the attractor. As such it has a physical meaning and provides additional information on the flow data such as the identification of transients.

5. Interpretation of CROM as a generalised Ulam–Galerkin method

For a complex dynamical system, a probabilistic description corresponding to the evolution of a swarm of points in the state space is often more insightful than the track of an individual trajectory (Cvitanović *et al.* 2012). One of the reasons is that the time evolution of probability densities is governed by a linear evolution operator called Perron–Frobenius (Lasota & Mackey 1994). It can be shown (Gaspard *et al.* 1995) that this operator is directly related to a Liouville equation which describes the evolution of a probability density function (p.d.f.) of the state variables. In this section, we demonstrate that, from a theoretical point of view, CROM can be considered as a generalisation of the Ulam–Galerkin method (Li 1976). This method is classically utilised for the approximation of the Perron–Frobenius operator associated with the Liouville equation (Bollt & Santitissadeekorn 2013). For more information on the relationship between the Liouville equation and the Perron–Frobenius operator as well as its adjoint, the Koopman operator, the reader is referred to appendix C.

The evolution equation for the mode coefficient vector $\boldsymbol{a}$ can be derived from a Galerkin projection onto the Navier–Stokes equation (Noack *et al.* 2011). This projection leads to the system

$$\frac{\mathrm{d}\boldsymbol{a}}{\mathrm{d}t} = \boldsymbol{f}(\boldsymbol{a}) \tag{5.1}$$

with a constant-linear-quadratic propagator $\boldsymbol{f} := [f_1, f_2, \ldots, f_N]^{\mathrm{T}}$ with $f_i = c_i + \sum_{j=1}^{N} l_{ij} a_j + \sum_{j,k=1}^{N} q_{ijk} a_j a_k$. As is well known (Gaspard *et al.* 1995), (5.1) induces a Liouville equation for the evolution of the probability density of the POD coefficients

$$\frac{\partial}{\partial t} p(\boldsymbol{a}, t) = -\nabla_{\boldsymbol{a}} \cdot \left[\boldsymbol{f}(\boldsymbol{a}) p(\boldsymbol{a}, t)\right] = \hat{\boldsymbol{L}}\, p(\boldsymbol{a}, t) \tag{5.2}$$

where $\hat{\boldsymbol{L}}$ is the Liouville operator. Here and in appendix C, we adopt the symbol $\hat{}$ to denote operators.

The aim of the Ulam–Galerkin method (Bollt & Santitissadeekorn 2013) is to determine a finite-rank approximation of the Perron–Frobenius operator via a Galerkin projection of the Liouville equation (5.2) onto a set of basis functions $\{\Psi_k\}_{k=1}^K$. Let Ω_k, $k=1,\ldots,K$, represent a Voronoi tessellation of the state space such that $\Omega = \cup_{k=1}^K \Omega_k$ with $\Omega_k \cap \Omega_j = \emptyset$ for $k \neq j$. The basis functions for the Ulam–Galerkin method are a family of characteristic functions

$$\Psi_k(\boldsymbol{a}) = \begin{cases} 1/|\Omega_k|^{1/2} & \text{if } \boldsymbol{a} \in \Omega_k, \\ 0 & \text{otherwise,} \end{cases} \tag{5.3}$$

where $|\Omega_k| := \int_{\Omega_k} \mathrm{d}\boldsymbol{a}$ denotes the volume of the kth cell.

The p.d.f. $p(\boldsymbol{a}, t)$ shall be approximated by a modal expansion

$$p(\boldsymbol{a}, t) = \sum_{k=1}^{K} \rho_k(t)\, \Psi_k(\boldsymbol{a}), \tag{5.4}$$

where $\rho_k(t) = \int_{\Omega_k} \mathrm{d}\boldsymbol{a}\, p(\boldsymbol{a}, t) \Psi_k(\boldsymbol{a})$ so that $\sum_{k=1}^K |\Omega_k|^{1/2} \rho_k = 1$.

Substituting (5.4) in (5.2), multiplicating with the basis function $\Psi_j(\boldsymbol{a})$ and integrating over Ω yields

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_j = \sum_{k=1}^{K} \mathscr{P}_{jk}^{cont}\, \rho_k, \quad j = 1, \ldots, K, \tag{5.5}$$

where

$$\mathscr{P}_{jk}^{cont} = \int_{\Omega} \mathrm{d}\boldsymbol{a}\, \Psi_j(\boldsymbol{a}) \hat{\boldsymbol{L}}\, \Psi_k(\boldsymbol{a}). \tag{5.6}$$

Here $\mathscr{P}^{cont} = (\mathscr{P}_{jk}^{cont})$ corresponds to the finite-rank approximation of the Perron–Frobenius operator which is normalised such that $\sum_{j=1}^K \mathscr{P}_{jk}^{cont} = 1$. If only observations $\{\boldsymbol{a}^m\}_{m=1}^M$ of the discrete map

$$\boldsymbol{a}^{m+1} = \boldsymbol{F}(\boldsymbol{a}^m) \tag{5.7}$$

are available, it is commonly accepted in the Ulam–Galerkin framework that $\mathscr{P}_{jk}^{cont}$ can be approximated by

$$\mathscr{P}_{jk} = \frac{\operatorname{card}\left\{\boldsymbol{a}^m \,|\, \boldsymbol{a}^m \in \Omega_k \text{ and } \boldsymbol{F}(\boldsymbol{a}^m) \in \Omega_j\right\}}{\operatorname{card}\left\{\boldsymbol{a}^m \in \Omega_k\right\}} \tag{5.8}$$

where $\operatorname{card}\{\cdot\}$ denotes the cardinality of a set. Then, the jkth entry of the matrix $\mathscr{P}^{cont}$ can be interpreted as the ratio of the fraction of the cell volume Ω_k that will be mapped inside the cell Ω_j after one action of the map $\boldsymbol{F}$ to the cell volume of Ω_k (Santitissadeekorn & Bollt 2007; Froyland, Junge & Koltai 2013). In appendix A of Bollt & Santitissadeekorn (2013), an implementation of the Ulam–Galerkin method is described where (5.8) is used in conjunction with a tessellation of the state space to approximate the Perron–Frobenius operator.

The CTM $\boldsymbol{P}$ as determined by (2.13) is another finite-rank approximation of the Perron–Frobenius operator. From the above description, CROM can be considered as an extension of the classical Ulam–Galerkin method for which the tessellation of the state space is not obtained from a Delaunay triangulation relying only on geometrical

	CROM	POD Galerkin model
Dynamics	Statistical closure (ergodic measure) with transients and built-in uncertainty (cluster size)	Nonlinear phase dynamics (fixed points, limit cycle, etc.; stability analysis)
Transients	Suitable for transient phenomenology	Challenges for transient dynamics
Control	Control design for ensemble of trajectories (ensemble control)	Standard control design (controller and observer)
Resolution	Complete probabilistic resolution of the whole fluctuation energy	Complete deterministic resolution in low-energy subspace
Sensitivity	Sensitive to cluster size and time step Computationally robust for long-term	Sensitive to subscale closures integration

TABLE 1. Capabilities and challenges of CROM and POD models.

principles but from a cluster analysis based on the k-means algorithm. Thus, CROM is more flexible since the cluster analysis can be easily modified by changing the definition of the distance metric used in the analysis.

6. Discussion

The derivation of low-dimensional models in fluid dynamics is driven by the need to understand nonlinear mechanisms in turbulent flows and their promising potential for control purposes. Different perspectives, e.g. a deterministic and a probabilistic viewpoint, can complement each other to give a broader picture on the governing dynamical system. CROM yields a probabilistic dynamical model for the evolution of the state variables. By benchmarking it against a deterministic POD model, we can identify their respective capabilities and challenges as a consequence from their methodological difference. The POD model serves – pars pro toto – as an example for many other data-driven models (Hyvärinen 2012; Kutz 2013). Figure 24 and table 1 serve to guide the following discussion.

POD models and CROM share many methodological similarities suggesting a critical assessment of their relative strengths and weaknesses. The similarities include the following: (i) both ROMs are based on velocity snapshots. (ii) The ROMs yield a small number of modes from which the full state is approximated. (iii) These modes (POD or centroids) are linear combinations of the original snapshots. (iv) The temporal evolution of the coarse-grained state is described by an autonomous system. For a better comparability, the continuous-time Markov process (2.24) is considered for CROM.

Yet, there are also important differences between POD models and CROM. POD models minimise the averaged residual of the corresponding Galerkin expansion. This expansion defines a linear subspace parametrised by mode coefficients $\boldsymbol{a} = [a_1, \ldots, a_N]^{\mathrm{T}}$. The individual mode $\boldsymbol{u}_i$ is just a basis vector. As such it has *per se* no physical meaning and may not be close to any flow realisation. The uncertainty, i.e. the residual of a POD expansion, is in the orthogonal complement of the linear POD space. The notion of a modal expansion and orthogonality implies that POD models are rooted in Hilbert spaces with an inner product.

In contrast, CROM centroids $\boldsymbol{c}_k$ minimise the total cluster variance, i.e. each centroid shall approximate all snapshots of the corresponding cluster. The minimal uncertainty of the centroidal flow representation is quantified by the 'radius' of

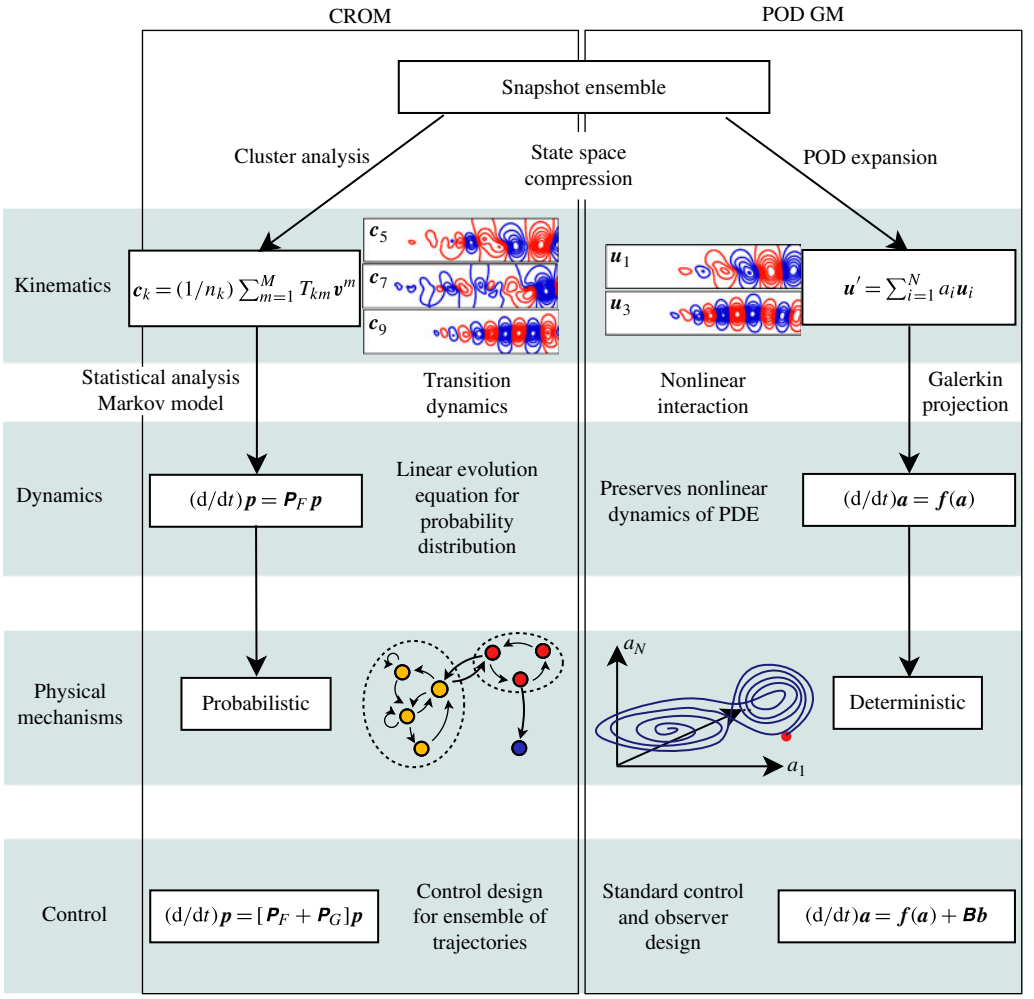

FIGURE 24. (Colour online) CROM versus POD models. For details see text.

the cluster, i.e. the maximum distance between the centroid and all snapshots of the corresponding cluster. Evidently, CROM requires a metric space with distance definition and no Hilbert space with an inner product.

The difference between a POD model and CROM becomes very apparent for periodic flow. For such data, POD generally yields cosine- and sine-like wave pairs for the first and higher frequency harmonics (Deane *et al.* 1991; Rempfer & Fasel 1994; Noack, Papas & Monkewitz 2005). In contrast, CROM discretises the limit cycle into sectors, i.e. phase bins.

POD and CROM have a different weighting of rare short-term events. POD performs a statistical weighting of events, i.e. rare events tend to be ignored if their contribution to the fluctuation energy is small. In contrast, CROM is better conceptualised as a geometric partitioning covering all snapshots and less dependent on the frequency of events.

Another important difference lies in the state spaces spanned by POD coefficients or by cluster probabilities. A POD model yields a deterministic description in the

corresponding subspace. In principle, the mode coefficient vector $\boldsymbol{a}$ can assume arbitrary values by construction. In contrast, CROM state descriptions range from the specification of a pure cluster, e.g. $\boldsymbol{p} = [\delta_{k1}, \ldots, \delta_{kK}]^{\mathrm{T}}$ for cluster k, to the ergodic measure for all clusters, i.e. a statistically stationary state. By construction, the cluster probabilities are bounded: by $0 \leqslant p_k$ and $\sum_{k=1}^{K} p_k = 1$.

The Galerkin system (5.1) can be derived from a projection onto the Navier–Stokes equation. This projection leads to a constant-linear-quadratic propagator $\boldsymbol{f}$ in the dynamical system (Holmes *et al.* 2012). The propagator may also be identified directly from data (Cordier, Abou El Majd & Favier 2010), like in CROM. In contrast, the Markov model of CROM is a linear evolution equation.

A linear evolution equation describing an unsteady nonlinear Navier–Stokes equation may appear initially surprising. The key to the explanation is that a Markov model describes the evolution of a probability distribution (ensemble of trajectories), not the evolution of a single trajectory (see appendix C). Such evolution equations can be derived from the Navier–Stokes equation, starting with the linear Liouville equation of a suitable probability space. The Hopf (1952) formalism for the Navier–Stokes equation is a prominent example. A simpler version constitutes the Liouville equation for a Galerkin system as in (5.2). The reader is referred to Noack & Niven (2012) for a detailed discussion. Thus, CROM is closely aligned with closure schemes, in which a stable fixed point represents the ergodic measure for the unsteady attractor in velocity space.

As a final note, we wish to highlight that CROM can also be utilised for flow control purposes. A CROM-based control strategy with the goal to bring the probability distribution as close as possible to a predefined distribution is outlined in appendix B.

7. Conclusions

In the present study, a novel CROM strategy is proposed for identifying physical mechanisms in an unsupervised manner. CROM is a generalisation of the Ulam–Galerkin method for the approximation of the Perron–Frobenius operator and yields a dynamical model for the evolution of an ensemble of trajectories in terms of a probability distribution. In contrast to the Ulam–Galerkin method, which is based on a Delaunay triangulation of the state space, CROM rests on a cluster analysis of the data, e.g. via the k-means algorithm. The definition of the distance metric enters the approach as a new parameter. One example is the observable-inferred decomposition (OID) generalising POD for aerodynamic or aeroacoustic observables (Schlegel *et al.* 2012). CROM can be formulated with the distance metric of OID. Then, the distance definition is based on the observable and not the flow itself. In principle, one can also think of an entropic distance metric taking into account information flux (Simovici & Djeraba 2008) between clusters.

CROM is applied to the Lorenz attractor (in § 3), the two-dimensional incompressible mixing layer undergoing vortex pairing (in § 4), and the three-dimensional incompressible turbulent wake of an Ahmed body (in appendix A). For these examples, CROM has been shown: (i) to identify quasi-attractors such as the two shedding regimes of the mixing layer or the bimodal states of the Ahmed body, and their intrinsic oscillatory or periodical behaviour; (ii) for periodic flows, the cluster analysis of CROM is a generalisation of phase averaging for which the average phase is represented by a centroid and the phase bins are of similar size; (iii) main transition processes between those quasi-attractors are characterised by branching

regions (Lorenz attractor), periodic transition regions (Ahmed body) or by a flipper cluster like for the mixing layer (note that for the mixing layer the transition between vortex pairing and Kelvin–Helmholtz shedding occurs only during one particular shedding phase); (iv) an analysis of the spectrum provides additional information on characteristic frequencies of the data and the convergence of the model to its asymptotic state, which determines the model's prediction horizon (see appendix A); (v) the prediction horizon can be defined in terms of the Kullback–Leibler entropy measuring the uncertainty of the model compared with its least-informative state. The convergence of the Kullback–Leibler entropy is validated with the rate of convergence obtained from the spectral decomposition.

CROM has a built-in uncertainty due to the cluster size which results in a diffusion of probability. Thus, the dynamical model converges to an asymptotic state which resembles the discretised ergodic measure on the attractor and yields additional statistical information on the data. Care has to be taken in the choice of data. Evidently, the model can only reveal the dynamics that are covered by the data ensemble. If the data ensemble contains transients, they can be easily detected as their asymptotic probability is zero.

A brief discussion shall give an overview of CROM's connection to other efforts with regard to nonlinear flow analysis and state space partitioning. CROM as an approximation of the Perron–Frobenius operator is strongly connected to its adjoint, the Koopman operator. Recently, Bagheri (2013) studied the conditions under which the DMD (Schmid 2010) approximates the Koopman modes. This analytical approach is clearly unthinkable for a complex dynamical system, and new approaches must be considered where CROM can play an important role.

Also, there exists an intriguing analogy between the cluster transition model and the ϵ-entropy for symbolic dynamics (Abel *et al.* 2000*a*,*b*). The symbol S_k in these references corresponds to the cluster $\boldsymbol{c}_k$, ϵ to the average distance between the centroids, and Δt to the sampling time τ_r in (14) of Abel *et al.* (2000*b*). The exit times relate to the transitions between the clusters.

CROM is also connected to ergodic partitioning. As far back as the 1960s, Kolmogorov & Tikhomirov (1959) developed entropic partitioning methods for the state space which also include ergodic partitioning. Ergodic partitioning as used by Mezic's group (see e.g. Mezic & Wiggins 1999) is a concept related to measure-preserving flows which can be divided into different sets inside which the flow is ergodic. If the attractors are separated in the state space, the cluster analysis of CROM is identical to ergodic partitioning. If not, ergodic partitioning will yield different results.

Summarising, CROM sufficiently exposes the characteristic features of the attractor and identifies quasi-attractors from the data. The dynamical connection between different flow states can be uncovered and singular events, meta-stable states, bimodality, hysteresis of several states, multi-attractor behaviour etc. may be detected. As such, CROM is not restricted to fluid flows. CROM is just based on a sequence of observations with distance measure. The distance metric enters as an additional parameter allowing, e.g., entropic or OID-based formulations of CROM. There are no restrictions on the type of data, e.g. time series, snapshots, etc., or the system from which they originate. CROM provides tools to identify transition regions, which can be crucial for mixing enhancement, and precursors that might lead to undesirable events. Thus, critical states can be determined which have to be manipulated in order to avoid those events. A promising CROM-based flow control strategy employs linear control laws taking into account nonlinear actuation dynamics. In future work, the authors will test CROM for flow control applications.

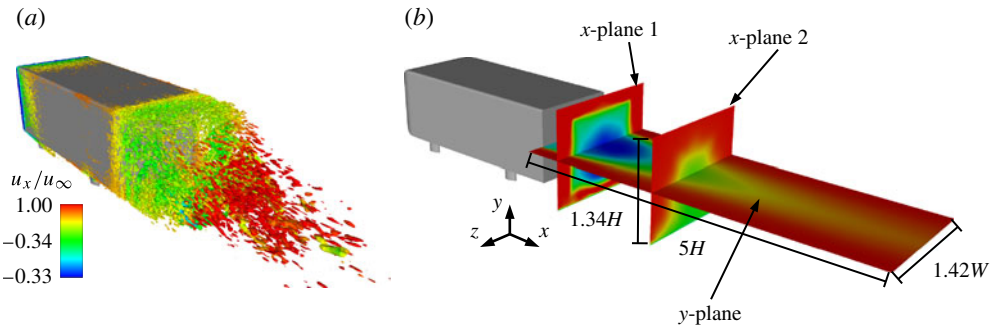

FIGURE 25. (Colour online) Bluff body wake. (*a*) Instantaneous isosurface of $Q = 150\,000$ which is coloured by the pressure coefficient $c_p = \Delta p/(1/2)\rho U_\infty^2$ with $\Delta p = p - p_\infty$ where p_∞ is the pressure in the freestream. The wake is characterised by a broadband spectrum exhibiting no obvious large-scale structures. (*b*) Principal sketch of the bluff body model. Results are visualised in the plane $y = 0$ which is coloured here by the streamwise velocity component.

Acknowledgements

The authors acknowledge the funding and excellent working conditions of the Chair of Excellence 'Closed-loop control of turbulent shear flows using reduced-order models' (TUCOROM) supported by the French Agence Nationale de la Recherche (ANR) and hosted by Institute PPRIME. This work is also supported by the NSF PIRE grant OISE-0968313. E.K. also thanks for the support by the region Poitou-Charentes, France. M.S. would like to acknowledge the support of the LINC project (no. 289447) funded by EC's Marie-Curie ITN program (FP7-PEOPLE-2011-ITN). We thank the Ambrosys GmbH (Society for Complex Systems Management) and the Bernd Noack Cybernetics Foundation for additional support.

We appreciate valuable stimulating discussions with the TUCOROM team and others: D. Barros, J. Borée, J.-P. Bonnet, S. Brunton, J. Delville, T. Duriez, N. Kamenzky, N. Kutz, J. Lewalle, J.-C. Laurentie, M. Morzyński, M. Schlegel, V. Parezanović and G. Tissot. Last but not least, we thank the referees for many important suggestions.

Special thanks are due to N. Maamar for a wonderful job in hosting the TUCOROM visitors.

Appendix A. Example with broadband turbulence and simple CROM structure

In this section, CROM is applied to an incompressible three-dimensional turbulent wake flow of a vehicle bluff body with the Reynolds number $Re_H = U_\infty H/\nu = 3 \times 10^5$ based on the freestream velocity U_∞, the height of the body H and the kinematic viscosity ν. The vehicle bluff body is a generic car model, also known as Ahmed body, and widely used to elucidate the relationship between observable flow structures and drag and lift forces arising typically on passenger cars. The employed model has a square back and is placed on four cylindrical supports (see figure 25*b*).

In this study, the flow is highly turbulent and exhibits a broadband spectrum in the wake. The dominant frequency is $St = fH/U \approx 0.2$ corresponding to the global shedding frequency of the wake. An instantaneous realisation of an isosurface of the second invariant of the velocity gradient tensor $\boldsymbol{Q} = -(1/2)(\partial\overline{u_i}\partial\overline{u_j}/\partial x_i \partial x_j)$ colour-marked by the pressure gradient is shown in figure 25(*a*).

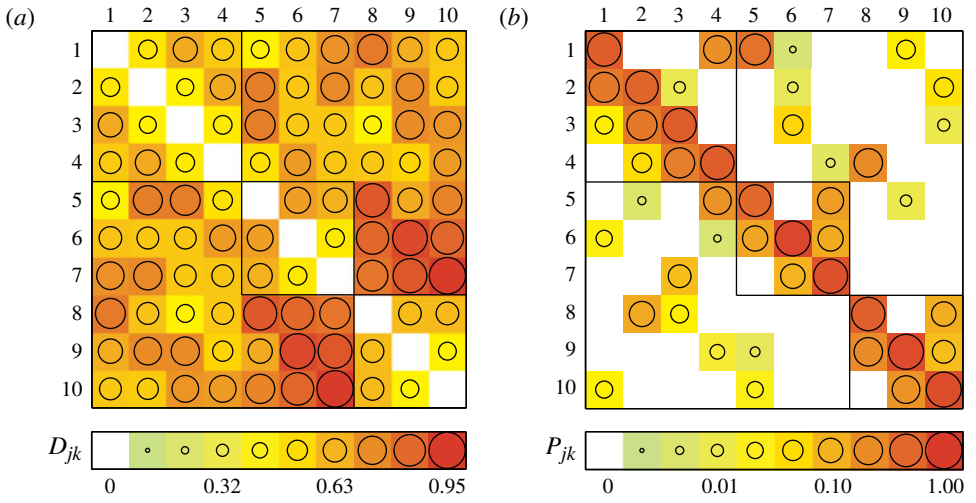

FIGURE 26. (Colour online) Kinematics and dynamics: (*a*) distance matrix $\boldsymbol{D}$ and (*b*) transition matrix $\boldsymbol{P}$ of the clustered Ahmed body data. Values are displayed as in figure 4. Three groups of clusters are distilled, namely groups 'T' ($k = 1, 2, 3, 4$), 'B1' ($k = 5, 6, 7$) and 'B2' ($k = 8, 9, 10$).

Recently, Grandemange *et al.* (2013) showed that the wake of a square-back body is characterised by a bimodal behaviour. Over large time scales of the order of $T \sim 10^2$–$10^3 H/U_\infty$, the flow switches between two asymmetric states which possess an asymmetric pressure distribution at the base with respect to the z direction. The presence of this bimodal behaviour has been found from Reynolds numbers ranging from $Re \sim 10^2$ to $Re \sim 10^6$ (Grandemange *et al.* 2012, 2013). The employed database exhibits this low-frequency base flow modulation as elaborated by Östh *et al.* (2013). In their study, they derived successfully a POD Galerkin model which is capable of predicting these events.

The database of $M = 1000$ velocity snapshots with sampling time $\Delta t = 0.5$, which is the normalised time step with respect to U_∞ and H, is computed with a large eddy simulation (LES). The computational domain comprises $[L_x \times L_y \times L_y] = [(8H + L + 20H) \times 5.33H \times 8.05H$ where L is the length of the car model. The snapshot data is preprocessed by a statistical symmetrisation with regard to the z-plane and then compressed with POD. A detailed description of the solver and the employed data as well as its validation with experimental data can be found in Östh *et al.* (2013). All results below are visualised with regard to the horizontal plane $y = 0$ as displayed in figure 25(*b*).

The cluster analysis is performed in the full POD space for $K_c = 10$ clusters. The distance matrix is visualised in figure 26(*a*) displaying three groups of clusters. Two groups, namely clusters $k = 5, 6, 7$ ('B1') and $k = 8, 9, 10$ ('B2'), are comparably far away from each other, while they are equally close to the group of clusters $k = 1, 2, 3, 4$ ('T'). Moreover, these groups are characterised by an intrinsic periodical behaviour as concluded from the CTM (see figure 26*b*). Similar to the mixing layer, high probabilities in the first subdiagonal and a closing transition from the last cluster to the first cluster inside the group indicate a periodical oscillation. While the groups 'B1' and 'B2' show nearly no connectivity, the group 'T' serves as a transition region between them. The most significant transitions are $3 \to 7$ (T $\to$ B1) and $2 \to 8$ (T $\to$ B2) as well as $5 \to 1$ (B1 $\to$ T) and $8 \to 4$ (B2 $\to$ T).

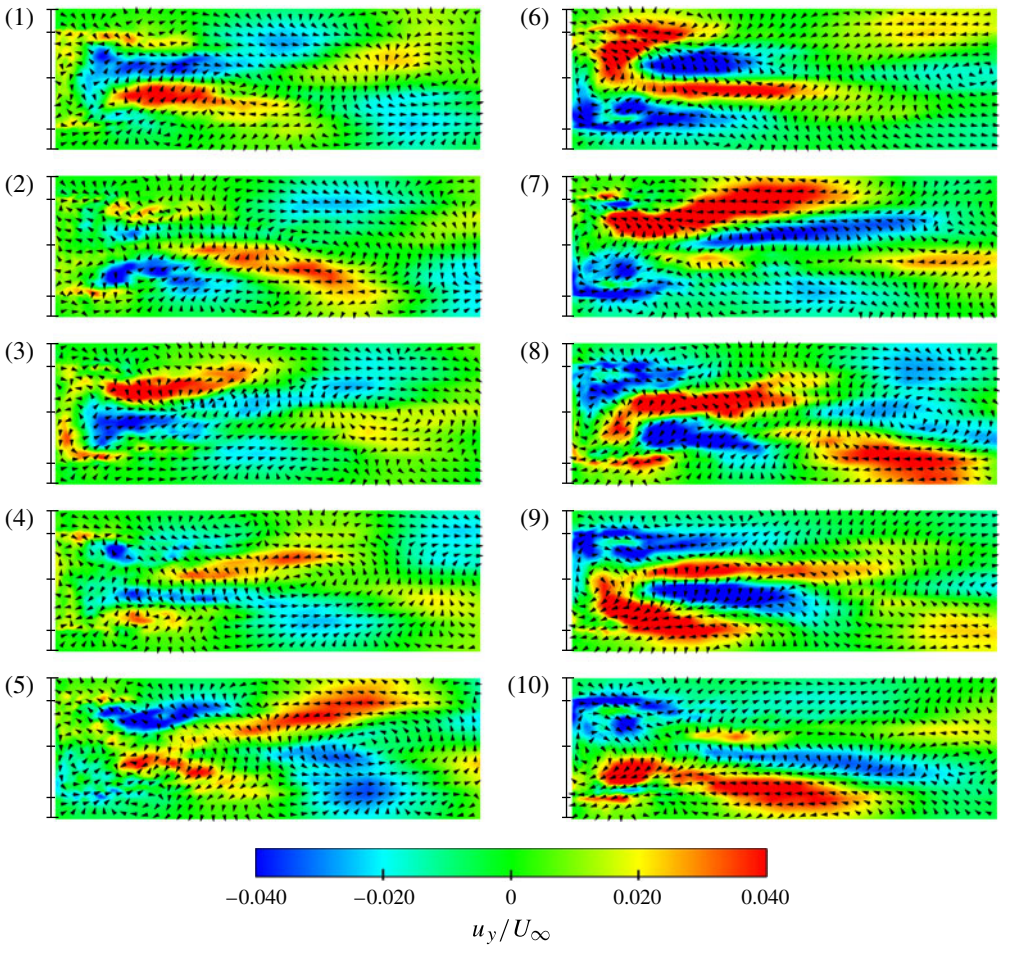

FIGURE 27. (Colour online) Cluster centroids of the Ahmed body. Visualisations of the y-planes coloured by u_y/U_∞. Cluster group 'T' ($k = 1, 2, 3, 4$) resembles the symmetric transition region, cluster groups 'B1' ($k = 5, 6, 7$) and 'B2' ($k = 8, 9, 10$) represent the two semimodal asymmetric states.

The cluster centroids computed as the mean of the velocity fluctuations in each cluster are displayed in figure 27. The clusters in 'B1' and 'B2' are characterised by strong vortical structures in opposite direction resembling the two semimodal flow states. Cluster $k = 5$ is an exception which is addressed below. The strength of the clusters in 'T' is weaker and the vector fields seem to be more symmetric with regard to the centreline $z = 0$. The principal branching clusters, similar as the flipper cluster of the mixing layer, are $k = 5$ and $k = 8$ connecting group 'T' with 'B1' and 'B2', respectively. These are intermediate states which explains why $k = 5$ seems more similar to the states in 'T' and $k = 8$ shares more dominant features with 'B2'.

An intuitive picture of the cluster transitions and the corresponding states is provided in figure 28. The simplified cluster transitions determined as for the mixing layer in § 4 and neglecting inner-cluster transitions are displayed in figure 28(*a*). The large cluster group 'T' connects the two semimodal states characterised by 'B1' and 'B2'. Clusters $k = 5$ and $k = 8$ have a special role and serve as branching

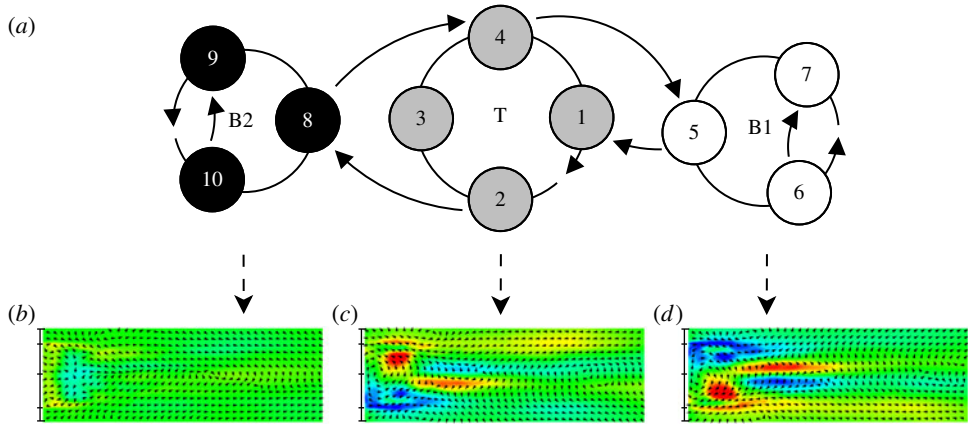

FIGURE 28. (Colour online) Simplified cluster transitions (analogous to figure 16). The arrows indicate possible transitions above a certain threshold. This threshold is chosen small enough to guarantee full connectivity. The graph highlights three cyclic groups associated with the two asymmetric bimodal states and a transition region between them. Two branching clusters, $k=5$ and $k=8$, connect these groups. Mean velocity fields of the three cluster groups of Ahmed body are shown below: (*b*) bimodal group 'B2' of clusters $k=8, 9, 10$, (*c*) symmetric group 'T' of clusters $k=1, 2, 3, 4$, (*d*) bimodal group 'B1' of clusters $k=5, 6, 7$. The visualisation of the mean flows corresponds to those in figure 27.

clusters. Apparently, the plots of the mean of each cluster group (figure 28*b–d*) confirms that 'B1' and 'B2' represent the two asymmetric states while 'T' represents a low-amplitude oscillation around the symmetric base flow (figure 28*c*).

A further analysis of the forces reveals that the two semimodal states are associated with an increase in the magnitude of the side force which was also found by Grandemange *et al.* (2013) and Östh *et al.* (2013) for the same data. In figure 29 the Voronoi diagram of the clusters, the cluster centroids (coloured bullets) and the data points (coloured dots) are displayed as described in appendix D. The mean forces associated with each cluster are shown as bar plots of which the first bar corresponds to the drag coefficient $c_D := F_x/((1/2)\rho U_\infty HW)$, the second to the lift coefficient $c_L := F_y/((1/2)\rho U_\infty HW)$, and the third to the side force coefficient $c_S := F_z/((1/2)\rho U_\infty HW)$. While the drag and the lift do not show any significant drift, the side force coefficient changes signs when the flow switches from one side to another. The lowest values in magnitude correspond to the clusters in 'T' which is consistent since they are tendentially symmetric.

Finally, the spectrum of the CTM displayed in figure 30 is analysed. In contrast to §§ 3 and 4, a different visualisation is employed displaying the real and imaginary part of $\mu = 1/\Delta t \log(\lambda)$ where λ is an eigenvalue of the CTM as defined in § 2. This transformation yields the growth rates and frequencies of the corresponding probability eigenvectors which can be directly linked to the growth rates and frequencies of observables, e.g. velocity field or pressure measurements (for details see appendix C). In addition to the invariant distribution corresponding to $\lambda = 1$, there exist three oscillatory modes. The frequency of the oscillatory mode with the smallest damping is $St = \omega_4/2\pi \approx 0.17$ (displayed in red bullets) close to $St = 0.2$ of the global wake shedding. Note that the oscillations are only resolved by three or four clusters. Better estimations can be obtained by a refined resolution, i.e. increasing the number of clusters K.

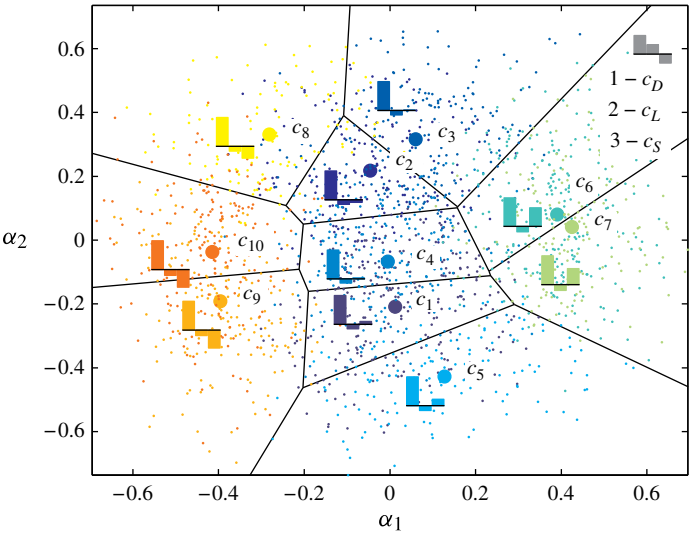

FIGURE 29. (Colour online) Cluster forces of the Ahmed body: visualisation of the Voronoi diagram of the clusters, the cluster centroids (coloured bullets), the data points (coloured dots) and bar plots of the mean force-related coefficients of each cluster in the space of the first two POD modes α_1 and α_2 associated with the centroids (as explained in appendix D). The coefficients are normalised with respect to the drag coefficient. The side coefficient is $\times 10$ enlarged for visualisation purposes. The colours indicate the different clusters. The clusters associated with the largest magnitude of the side force correspond to the outer clusters $k = 5, 6, 7$ and $k = 8, 9, 10$, respectively, which represent the two asymmetric flow states.

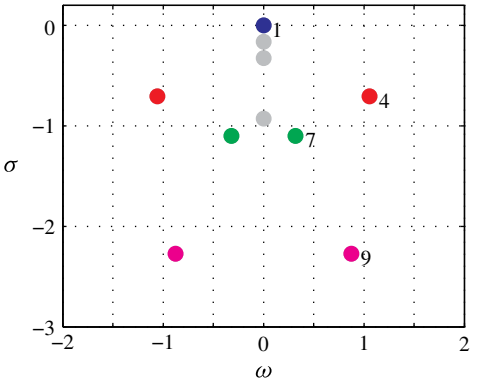

FIGURE 30. (Colour online) Stability analysis of the transition matrix associated with the Ahmed body: in contrast to figure 19, the growth rate $\sigma = \mathrm{Re}(\mu)$ and the frequency $\omega = \mathrm{Im}(\mu)$ are shown according to the transformation $\mu = 1/\Delta t \log(\lambda)$. The marginal stable eigenvalue μ_1, or $\lambda = 1$, has vanishing growth rate and frequency (blue bullet) and is associated with the asymptotic probability distribution. Other non-oscillatory modes are visualised with grey bullets. The frequency corresponding to the oscillatory mode with the smallest damping (red bullets) is $St \approx 0.17$.

In summary, CROM distils the two bimodal states of the base flow and identifies the main transition mechanism. These asymmetric quasi-attractors are connected to the comparably symmetric transition region solely via the two branching clusters $k=5$ and $k=8$. All groups are intrinsically periodic and the dominant frequency is determined as $St=0.17$. A larger number of clusters increases the resolution of the bimodal states and the branching clusters but not change the main mechanism.

Appendix B. Towards CROM-based flow control

In this section, a control design is proposed for the Markov model. The goal is to bring the probability distribution $\boldsymbol{p}$ as close as possibly to the desired distribution $\boldsymbol{p}^{target}$ with the corresponding Markov matrix $\mathscr{P}^{target}$. In the following, we drop the superscript 'cont' to simplify the notation.

We start with the Galerkin system

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{a}=\boldsymbol{f}(\boldsymbol{a})+\boldsymbol{g}(\boldsymbol{a}), \tag{B 1}$$

where $\boldsymbol{f}$ is the propagator of the unforced dynamics and $\boldsymbol{g}$ the actuation term incorporating the control law. The corresponding Liouville equation for the probability distribution $p(\boldsymbol{a}, t)$ reads

$$\frac{\partial}{\partial t}p(\boldsymbol{a}, t)+\nabla_{\boldsymbol{a}}\cdot\left[p(\boldsymbol{a}, t)\,(\boldsymbol{f}(\boldsymbol{a})+\boldsymbol{g}(\boldsymbol{a}))\right]=0. \tag{B 2}$$

Let $\boldsymbol{\rho}=[\rho_1, \ldots, \rho_K]^{\mathrm{T}}$ be the Galerkin coefficients of $p(\boldsymbol{a}, t)$ on the basis $\{\Psi_k(\boldsymbol{a})\}_{k=1}^{K}$, the Markov model can be derived in complete analogy to § 5

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{\rho}=\left[\mathscr{P}_F+\mathscr{P}_G\right]\boldsymbol{\rho}, \tag{B 3}$$

where $\mathscr{P}_F=(\mathscr{P}_{jk}^{F})$ is the transition matrix for the unactuated propagator $\boldsymbol{f}$ while $\mathscr{P}_G=(\mathscr{P}_{jk}^{G})$ represents the closed-loop actuation term $\boldsymbol{g}$. Note that $\sum_{k=1}^{K}|\Omega_k|^{1/2}\rho_k=1$ is guaranteed by

$$\sum_{j=1}^{K}\mathscr{P}_{jk}^{F}=0 \quad \text{for all } k=1, \ldots, K, \tag{B 4a}$$

$$\sum_{j=1}^{K}\mathscr{P}_{jk}^{G}=0 \quad \text{for all } k=1, \ldots, K. \tag{B 4b}$$

The actuation term comprises a fixed actuation hardware and a free control law to be optimised. For simplicity, we assume a single volume force actuation. Then, the Galerkin system (B 1) has the actuation term $\boldsymbol{g}=\boldsymbol{B}b$, where b is the actuation command and $\boldsymbol{B}$ a column vector representing the shape of the volume force. Again, for simplicity, the control law is searched in the form of a linear full-state feedback ansatz, i.e.

$$b=\boldsymbol{K}\boldsymbol{a}, \tag{B 5}$$

where $\boldsymbol{K}$ is a row vector. Summarising,

$$\boldsymbol{g}=\boldsymbol{B}\boldsymbol{K}\boldsymbol{a}. \tag{B 6}$$

For each cluster $k = 1, \ldots, K$, the actuation term can be estimated by the corresponding centroid $\boldsymbol{a} = \boldsymbol{c}_k$, i.e. $\boldsymbol{g}_k = \boldsymbol{B}\boldsymbol{K}\boldsymbol{c}_k$. Following, the derivation of Markov model (5.5) the corresponding actuation transition matrix reads

$$\mathscr{P}_G = \mathscr{P}_B \boldsymbol{K}, \tag{B 7}$$

where $\mathscr{P}_B$ corresponds to the volume force.

The solution of (B 3) should be as close as possible to the post-transient goal distribution, i.e.

$$\boldsymbol{\rho}(t) = \mathrm{e}^{\mathscr{P}_F t}\boldsymbol{\rho}_0 + \int_0^t \mathrm{d}\tau\, \mathrm{e}^{\mathscr{P}_F(t-\tau)}\mathscr{P}_B \boldsymbol{K}\boldsymbol{\rho}(\tau) \approx \boldsymbol{\rho}^{target}. \tag{B 8}$$

The corresponding optimality condition reads

$$\|\boldsymbol{\rho}(t) - \boldsymbol{\rho}^{target}\| \stackrel{!}{=} \min, \tag{B 9}$$

where $\|\cdot\|$ may denote any suitable vector norm.

The control design task is to determine $\boldsymbol{K}$ from this optimality condition. A small (large) time horizon t corresponds to an aggressive (conservative) control law.

It should be noted that this linear control law takes into account the nonlinear actuation dynamics resolved by the clustered state space. In principle, the CROM-based control strategy can be expected to exploit strong nonlinearities, such as frequency cross-talk. We actively pursue this direction for our experimental flow control plants.

Appendix C. Perron–Frobenius and Koopman operators

In this section, relations of the Liouville equation with the Perron–Frobenius and Koopman operators are elaborated. Intriguingly, these are linear operators although the system dynamics is nonlinear which explains their important role in dynamical systems theory. If the system is noisy, the Perron–Frobenius operator is replaced by the Fokker–Planck operator. We refer the reader to Lasota & Mackey (1994) for a mathematically rigorous treatment, and give here a more intuitive picture.

The Perron–Frobenius operator describes the temporal evolution of a probability density. Let $\hat{\boldsymbol{L}}$ be the Liouville operator as in (5.2), the solution $p(\boldsymbol{a}, t)$ of the Liouville equation

$$\frac{\partial}{\partial t}p(\boldsymbol{a}, t) = \hat{\boldsymbol{L}}p(\boldsymbol{a}, t) \tag{C 1}$$

can be written formally as $p(\boldsymbol{a}, t) = \hat{\boldsymbol{P}}_t p_0(\boldsymbol{a})$ where $p_0(\boldsymbol{a})$ is the initial probability distribution. The Perron–Frobenius operator $\hat{\boldsymbol{P}}_t$ maps the p.d.f. $p(\boldsymbol{a}, t)$ forward in time. Since $\hat{\boldsymbol{P}}_t(\alpha p_1 + \beta p_2) = \alpha\hat{\boldsymbol{P}}_t p_1 + \beta\hat{\boldsymbol{P}}_t p_2$ for any functions p_1, p_2 and scalars α, β, this operator is linear. From (C 1), we can infer that the Perron–Frobenius operator is also given by $\hat{\boldsymbol{P}}_t = \exp(t\hat{\boldsymbol{L}})$ where $\hat{\boldsymbol{L}}$ is the Liouville operator defined by $\hat{\boldsymbol{L}}p := -\nabla_{\boldsymbol{a}} \cdot (\boldsymbol{f}\, p)$ with the system dynamics $\boldsymbol{f}$ as in (5.1). The p.d.f. at time t can also be written as an integral (Gaspard *et al.* 1995)

$$p(\boldsymbol{a}, t) = \hat{\boldsymbol{P}}_t p_0(\boldsymbol{a}) = \int \mathrm{d}\boldsymbol{a}_0\, \delta\left(\boldsymbol{a} - \boldsymbol{\Phi}_t\left(\boldsymbol{a}_0\right)\right)\, p_0(\boldsymbol{a}_0) \tag{C 2}$$

where $\boldsymbol{\Phi}_t$ is the flow map associated with the Galerkin system (5.1). Analogous to the Perron–Frobenius operator for a p.d.f., the flow map $\boldsymbol{\Phi}_t$ shifts the state variable $\boldsymbol{a}$ forward in time with $\boldsymbol{a}(t) = \boldsymbol{\Phi}_t \boldsymbol{a}_0$ where $\boldsymbol{a}_0 := \boldsymbol{a}(0)$ is the initial condition. The kernel $\delta(\boldsymbol{a} - \boldsymbol{\Phi}_t(\boldsymbol{a}_0))$ of the Perron–Frobenius operator can be interpreted as a conditional probability density for the trajectory to be at point $\boldsymbol{a}$ if it had been initially at point $\boldsymbol{a}_0$. A finite-rank approximation of the Perron–Frobenious operator can be obtained from the Ulam–Galerkin method which is described in § 5.

In general, one is interested in observables of a system such as measurements of the pressure or the velocity field. The average or the expected value of a scalar observable g at time t is given by

$$\langle g \rangle(t) = \int \mathrm{d}\boldsymbol{a}\, g(\boldsymbol{a}) p(\boldsymbol{a}, t) = \iint \mathrm{d}\boldsymbol{a}\, \mathrm{d}\boldsymbol{a}_0\, g(\boldsymbol{a}) \delta\left(\boldsymbol{a} - \boldsymbol{\Phi}_t\left(\boldsymbol{a}_0\right)\right) p_0(\boldsymbol{a}_0). \tag{C 3}$$

This expression and the definition (C 2) of $\hat{\boldsymbol{P}}_t$ can be exploited (see Gaspard *et al.* 1995, for a proof) to define the Koopman operator, denoted by $\hat{\boldsymbol{K}}_t$, which is adjoint to the Perron–Frobenius operator according to

$$\langle g \rangle(t) = \langle g,\ \hat{\boldsymbol{P}}_t p_0 \rangle = \langle \hat{\boldsymbol{K}}_t g,\ p_0 \rangle \tag{C 4}$$

where the inner product is defined by $\langle g, f \rangle := \int \mathrm{d}\boldsymbol{a}\, g(\boldsymbol{a}) f(\boldsymbol{a})$ for two functions f and g. The observable g at time t is then given by

$$g\left(\boldsymbol{a}(t)\right) = \hat{\boldsymbol{K}}_t g(\boldsymbol{a}_0) = \int \mathrm{d}\boldsymbol{a}\, \delta\left(\boldsymbol{a} - \boldsymbol{\Phi}_t\left(\boldsymbol{a}_0\right)\right) g(\boldsymbol{a}). \tag{C 5}$$

Thus, the time evolution of the observable is ruled by the Koopman operator. The corresponding adjoint Liouville equation reads

$$\frac{\partial}{\partial t} g(\boldsymbol{a}) = (\boldsymbol{f} \cdot \boldsymbol{\nabla}_{\boldsymbol{a}})\ g(\boldsymbol{a}) = \hat{\boldsymbol{L}}^{\dagger} g(\boldsymbol{a}), \tag{C 6}$$

where $\hat{\boldsymbol{L}}^{\dagger}$ is the adjoint Liouville operator. In general, the Liouville operator and its adjoint are related by $\hat{\boldsymbol{L}} + \hat{\boldsymbol{L}}^{\dagger} = -\boldsymbol{\nabla} \cdot \boldsymbol{f}$ where $\boldsymbol{\nabla} \cdot \boldsymbol{f}$ is not vanishing for dissipative systems.

In recent years, the spectral analysis of complex nonlinear flows has received increasing attention. Rowley *et al.* (2009) determined the eigenfunctions and eigenvalues of the Koopman operator $\hat{\boldsymbol{K}}_t$ from only an ensemble of observables via the DMD which was introduced in Schmid (2010).

More recently, Bagheri (2013) studied the conditions under which the DMD algorithm approximates the Koopman modes. For this purpose, he considered the first Hopf bifurcation of the flow past a circular cylinder and constructed analytically the Koopman modes. This analytical approach is clearly unthinkable for a complex dynamical system, and new approaches must be considered where CROM can play an important role. Indeed, the Koopman analysis and CROM are strongly connected. As discussed in § 5, CROM can be interpreted as a finite-rank approximation of the Perron–Frobenius operator which is adjoint to the Koopman operator.

Appendix D. Visualisation of the cluster topology

The visualisation of trajectories in the state space can give an intuitive picture of the underlying dynamics and the state space structure. It can become a challenging and tedious task if one considers high-dimensional data such as a velocity snapshot ensemble or POD coefficients and wants to keep certain properties of the data.

In the case of CROM, the two-dimensional visualisation of the cluster arrangement in the state space can contribute to a better understanding of the dynamics on the attractor. A simple method that optimally preserves the centroids' pointwise distances in a least-mean-square-error sense employs the POD as mentioned in § 2. Instead of the velocity snapshots we consider here the centroids.

The matrix $\boldsymbol{C}$ contains the centred centroids $\boldsymbol{c}'_k = \boldsymbol{c}_k - 1/K \sum_{k=1}^{K} \boldsymbol{c}_k$ as columns, i.e. $\boldsymbol{C} = [\boldsymbol{c}'_1 \cdots \boldsymbol{c}'_K]$. The eigenvalues ν_k, $k = 1, \ldots, K$, are obtained from a spectral decomposition of the covariance matrix $\boldsymbol{C}^{\mathrm{T}}\boldsymbol{C}$, where the superscript 'T' denotes the transpose. They are ordered, $\nu_1 \geqslant \nu_2 \geqslant \cdots \geqslant \nu_k \geqslant 0$. The matrix $\boldsymbol{W}_r$ contains the first r eigenvectors of the covariance matrix in the columns. Then, a projection of the original centroids $\boldsymbol{c}_k$ onto $\boldsymbol{W}_r$ yields points

$$\left(\boldsymbol{c}_k^r\right)^{\mathrm{T}} = (\boldsymbol{c}_k)^{\mathrm{T}}\, \boldsymbol{W}_r, \quad k = 1, \ldots, K, \tag{D 1}$$

in r dimensions where the distances between the points $\boldsymbol{c}_k^r$ approximate the distances between the centroids $\boldsymbol{c}_k$. Analogously, the projection of the snapshots is given by

$$(\boldsymbol{u}^{m,r})^{\mathrm{T}} = (\boldsymbol{u}^{m})^{\mathrm{T}}\, \boldsymbol{W}_r. \tag{D 2}$$

For visualisation purposes in appendix A, we choose $r = 2$. The POD mode amplitudes associated with the cluster centroids are denoted by α_k.

In a more general context, for a given distance matrix according to a (non-Euclidean) distance metric, multidimensional scaling (MDS) (Mardia, Kent & Bibby 1979; Cox & Cox 2000) aims to find corresponding points in a low-dimensional subspace so that the distances between the points are preserved. This is referred to as classical scaling. The solution can vary in terms of a translation, a rotation and reflections. In the case where the distance is measured via the Euclidean metric, this method coincides with the POD, and the mean is at the origin and the axes are the POD eigenvectors (Cox & Cox 2000).

REFERENCES

Abel, M., Biferale, L., Cencini, M., Falcioni, M., Vergni, D. & Vulpiani, A. 2000*a* Exit-time approach to ϵ-entropy. *Phys. Rev. Lett.* **84**, 6002–6005.

Abel, M., Biferale, L., Cencini, M., Falcioni, M., Vergni, D. & Vulpiani, A. 2000*b* Exit-times and ϵ-entropy for dynamical systems, stochastic processes, and turbulence. *Physica* **147**, 12–35.

Afraimovich, V. S. & Shil'nikov, L. P. 1983 *Nonlinear Dynamics and Turbulence*. Pitmen.

Ahmed, S. R., Ramm, G. & Faltin, G. 1984 *Some Salient Features of the Time Averaged Ground Vehicle Wake*. Society of Automotive Engineers; 840300.

Antoulas, A. C. 2005 *Approximation of Large-Scale Dynamical Systems*. Society for Industrial and Applied Mathematics.

Arthur, D. & Vassilvitskii, S. 2007 *k*-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM–SIAM Symposium on Discrete Algorithms, Philadelphia, PA, USA*, pp. 1027–1035. Society for Industrial and Applied Mathematics.

AUBRY, N., HOLMES, P., LUMLEY, J. L. & STONE, E. 1988 The dynamics of coherent structures in the wall region of a turbulent boundary layer. *J. Fluid Mech.* **192**, 115–173.

BAGHERI, S. 2013 Koopman-mode decomposition of the cylinder wake. *J. Fluid Mech.* **726**, 596–623.

BAGHERI, S., HOEPFFNER, J., SCHMID, P. J. & HENNINGSON, D. S. 2008 Input–output analysis and control design applied to a linear model of spatially developing flows. *Appl. Mech. Rev.* **62**, 1–27.

BALL, G. & HALL, D. 1965 ISODATA, a novel method of data anlysis and pattern classification. *Tech. Rep.* NTIS AD 699616. Stanford Research Institute, Stanford, CA.

BERGMANN, M. & CORDIER, L. 2008 Optimal control of the cylinder wake in the laminar regime by trust-region methods and POD reduced order models. *J. Comput. Phys.* **227**, 7813–7840.

BISHOP, C. M. 2007 *Pattern Recognition and Machine Learning*. Springer.

BOLLT, E. M. & SANTITISSADEEKORN, N. 2013 *Applied and Computational Measurable Dynamics*. Society for Industrial and Applied Mathematics.

BROWAND, F. K. & WEIDMAN, P. D. 1976 Large scales in the developing mixing layer. *J. Fluid Mech.* **76** (1), 127–144.

BROWN, G. L. & ROSHKO, A. 1974 On density effects and large structure in turbulent mixing layers. *J. Fluid Mech.* **64**, 775–816.

BURKARDT, J., GUNZBURGER, M. & LEE, H. C. 2006*a* POD and CVT-based reduced-order modeling of Navier–Stokes flows. *Comput. Meth. Appl. Mech. Engng* **196**, 337–355.

BURKARDT, J., GUNZBURGER, M. & LEE, H.-C. 2006*b* Centroidal Voronoi tessellation-based reduced-order modeling of complex systems. *SIAM J. Sci. Comput.* **28** (2), 459–484.

CACUCI, D. G., NAVON, I. M. & IONESCU-BUJOR, M. 2013 *Computational Methods for Data Evaluation and Assimilation*. Chapman and Hall/CRC.

CAVALIERI, A., DAVILLER, G., COMTE, P., JORDAN, P., TADMOR, G. & GERVAIS, Y. 2011 Using large eddy simulation to explore sound-source mechanisms in jets. *J. Sound Vib.* **330** (17), 4098–4113.

CHIANG, M. M.-T. & MIRKIN, B. 2010 Intelligent choice of the number of clusters in K-means clustering: an experimental study with different cluster spreads. *J. Classification* **27**, 3–40.

CORDIER, L., ABOU EL MAJD, B. & FAVIER, J. 2010 Calibration of POD reduced-order models using Tikhonov regularization. *Intl J. Numer. Meth. Fluids* **63** (2), 269–296.

CORDIER, L., NOACK, B. R., TISSOT, G., LEHNASCH, G., DELVILE, J., BALAJEWICZ, M., DAVILLER, G. & NIVEN, R. K. 2013 Control-oriented model identification strategy. *Exp. Fluids* **54** (8), 1580; Invited paper for the Special Issue 'Flow Control' (ed. Bonnet, J.-P. & Cattafesta, L.).

COX, T. F. & COX, M. A. A. 2000 *Multidimensional Scaling*, 2nd edn, Monographs on Statistics and Applied Probability, vol. 88. Chapman and Hall.

CVITANOVIĆ, P., ARTUSO, R., MAINIERI, R., TANNER, G. & VATTAY, G. 2012 *Chaos: Classical and Quantum*. Niels Bohr Institute.

DAVILLER, G. 2010 Étude numérique des effets de température dans les jets simples et coaxiaux. PhD thesis, École Nationale Supérieure de Mécanique et d'Aérotechnique.

DEANE, A. E., KEVREKIDIS, I. G., KARNIADAKIS, G. E. & ORSZAG, S. A. 1991 Low-dimensional models for complex geometry flows: application to grooved channels and circular cylinders. *Phys. Fluids* A **3**, 2337–2354.

DIMOTAKIS, P. E. & BROWN, G. L. 1976 The mixing layer at high Reynolds number: large-structure dynamics and entrainment. *J. Fluid Mech.* **78** (3), 535–560.

DU, Q., FABER, V. & GUNZBURGER, M. 1999 Centroidal Voronoi tessellations: applications and algorithms. *SIAM Rev.* **41**, 637–676.

FLETCHER, C. A. J. 1984 *Computational Galerkin Methods*, 1st edn. Springer.

FOWLER, R. A. 1929 *Statistical Mechanics*. Cambridge University Press.

FROYLAND, G., JUNGE, O. & KOLTAI, P. 2013 Estimating long-term behavior of flows without trajectory integration: the infinitesimal generator approach. *SIAM J. Numer. Anal.* **51** (1), 223–247.

GASPARD, P., NICOLIS, G., PROVATA, A. & TASAKI, S. 1995 Spectral signature of the pitchfork bifurcation: Liouville equation approach. *Phys. Rev.* E **51**, 74–94.

GONCHENKO, S. V., SHIL'NIKOV, L. P. & TURAEV, D. V. 1997 Quasiattractors and homoclinic tangencies. *Comput. Maths Applics.* **34**, 195–227.

GRANDEMANGE, M., CADOT, O. & GOHLKE, M. 2012 Reflectional symmetry breaking of the separated flow over three-dimensional bluff bodies. *Phys. Rev.* E **86**, 035302.

GRANDEMANGE, M., GOHLKE, M. & CADOT, O. 2013 Turbulent wake past a three-dimensional blunt body. Part 1. Global modes and bi-stability. *J. Fluid Mech.* **722**, 51–84.

HARTIGAN, J. A. 1975 *Clustering Algorithms*. John Wiley & Sons.

HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. 2009 *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*, 2nd edn. Springer.

HERVÉ, A., SIPP, D., SCHMID, P. J. & SAMUELIDES, M. 2012 A physics-based approach to flow control using system identification. *J. Fluid Mech.* **702**, 26–58.

HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.

HOPF, E. 1952 Statistical hydromechanics and functional analysis. *J. Rat. Mech. Anal.* **1**, 87–123.

HUANG, S.-C. & KIM, J. 2008 Control and system identification of a separated flow. *Phys. Fluids* **20** (10), 101509.

HYVÄRINEN, A. 2012 Independent component analysis: recent advances. *Phil. Trans. R. Soc.* A **371** (1984).

KOLMOGOROV, A. N. & TIKHOMIROV, V. M. 1959 ε-entropy and ε-capacity of sets in function spaces. *Usp. Mat. Nauk* **86**, 3–86.

KULLBACK, S. 1959 *Information Theory and Statistics*, 1st edn. John Wiley & Sons.

KULLBACK, S. & LEIBLER, R. A. 1951 On information and sufficiency. *Ann. Math. Statist.* **22**, 79–86.

KUTZ, J. N. 2013 *Data-Driven Modeling and Scientific Computation: Methods for Complex Systems and Big Data*. Oxford University Press.

LADYZHENSKAYA, O. A. 1963 *The Mathematical Theory of Viscous Incompressible Flow*, 1st edn. Gordon and Breach.

LAIZET, S., LARDEAU, L. & LAMBALLAIS, E. 2010 Direct numerical simulation of a mixing layer downstream a thick splitter plate. *Phys. Fluids* **22** (1).

LANGVILLE, A. N. & MEYER, C. D. 2012 *Google's PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press.

LASOTA, A. & MACKEY, M. C. 1994 *Chaos, Fractals, and Noise*, 2nd edn. Springer.

LI, T. Y. 1976 Finite approximation for the Frobenius–Perron operator: a solution to Ulam's conjecture. *J. Approx. Theory* **17** (2), 177–186.

LIENHART, H. & BECKER, S. 2003 Flow and turbulent structure in the wake of a simplified car model. *SAE Paper* (2003-01-0656).

LLOYD, S. 1956 Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**, 129–137; originally as an unpublished Bell laboratories Technical Note (1957).

LORENZ, E. N. 1963 Deterministic non-periodic flow. *J. Atmos. Sci.* **20**, 130–141.

LUCHTENBURG, D. M., GÜNTER, B., NOACK, B. R., KING, R. & TADMOR, G. 2009 A generalized mean-field model of the natural and actuated flows around a high-lift configuration. *J. Fluid Mech.* **623**, 283–316.

MACQUEEN, J. 1967 Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Math. Stat. and Prob.*, Vol. 1, pp. 281–297.

MARDIA, K. V., KENT, J. T. & BIBBY, J. M. 1979 *Multivariate Analysis*. Academic Press.

MEYER, C. D. 2000 *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics.

MEZIC, I. & WIGGINS, S. 1999 A method for visualization of invariant sets of dynamical systems based on the ergodic partition. *Chaos* **9** (1), 213–218.

MONKEWITZ, P. A. 1988 Subharmonic resonance, pairing and shredding in the mixing layer. *J. Fluid Mech.* **188**, 223–252.

MOORE, B. 1981 Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control* **26**, 17–32.

MURPHY, K. P. 2012 *Machine Learning: A Probabilistic Perspective*. MIT Press.

NIVEN, R. K. 2009 Combinatorial entropies and statistics. *Eur. Phys. J.* B **70** (1), 49–63.

NOACK, B. R., AFANASIEV, K., MORZYŃSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.

NOACK, B. R., MORZYŃSKI, M. & TADMOR, G. 2011 *Reduced-Order Modelling for Flow Control*, CISM Courses and Lectures, vol. 528. Springer-Verlag.

NOACK, B. R. & NIVEN, R. K. 2012 Maximum-entropy closure for a Galerkin system of incompressible shear flow. *J. Fluid Mech.* **700**, 187–213.

NOACK, B. R., PAPAS, P. & MONKEWITZ, P. A. 2005 The need for a pressure-term representation in empirical Galerkin models of incompressible shear flows. *J. Fluid Mech.* **523**, 339–365.

NOACK, B. R., PELIVAN, I., TADMOR, G., MORZYŃSKI, M. & COMTE, P. 2004 Robust low-dimensional Galerkin models of natural and actuated flows. In *Fourth Aeroacoustics Workshop*, pp. 0001–0012. RWTH Aachen, 26–27 February 2004.

NORRIS, J. R. 1998 *Markov Chains*. Cambridge University Press.

ÖSTH, J., NOACK, B. R., KRAJNOVIĆ, S., BARROS, D. & BOREÉ, J. 2014 On the need for a nonlinear subscale turbulence term in pod models as exemplified for a high Reynolds number flow over an ahmed body. *J. Fluid Mech.* **747**, 518–544.

PASTOOR, M., HENNING, L., NOACK, B. R., KING, R. & TADMOR, G. 2008 Feedback shear layer control for bluff body drag reduction. *J. Fluid Mech.* **608**, 161–196.

RAJAEE, M., KARLSSON, S. K. F. & SIROVICH, L. 1994 Low-dimensional description of free-shear-flow coherent structures and their dynamical behaviour. *J. Fluid Mech.* **258**, 1–29.

REMPFER, D. 2006 On boundary conditions for incompressible Navier–Stokes problems. *Appl. Mech. Rev.* **59**, 107–125.

REMPFER, D. & FASEL, F. H. 1994 Dynamics of three-dimensional coherent structures in a flat-plate boundary-layer. *J. Fluid Mech.* **275**, 257–283.

ROWLEY, C. W. 2005 Model reduction for fluids using balanced proper orthogonal decomposition. *Intl J. Bifurcation Chaos* **15** (3), 997–1013.

ROWLEY, C. W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **645**, 115–127.

SANTITISSADEEKORN, N. & BOLLT, E. M. 2007 Identifying stochastic basin hopping by partitioning with graph modularity. *Physica* D **231**, 95–107.

SCHLEGEL, M., NOACK, B. R., JORDAN, P., DILLMANN, A., GRÖSCHEL, E., SCHRÖDER, W., WEI, M., FREUND, J. B., LEHMANN, O. & TADMOR, G. 2012 On least-order flow representations for aerodynamics and aeroacoustics. *J. Fluid Mech.* **697**, 367–398.

SCHMID, P. J. 2010 Dynamic mode decomposition for numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.

SCHNEIDER, T. M., ECKHARDT, B. & VOLLMER, J. 2007 Statistical analysis of coherent structures in transitional pipe flow. *Phys. Rev.* E **75**, 066313.

SIMOVICI, D. A. & DJERABA, C. 2008 *Mathematical Tools for Data Mining – Set Theory, Partial Orders, Combinatorics*. Springer-Verlag.

SIPP, D., MARQUET, O., MELIGA, P. & BARBAGALLO, A. 2010 Dynamics and control of global instabilities in open-flows: a linearized approach. *Appl. Rev. Mech.* **63**, 251–276.

SPARROW, C. 1982 *The Lorenz Equations: Bifurcations, Chaos, and Strange Attractors*, 1st edn. Applied Mathematical Sciences, vol. 41. Springer-Verlag.

STEINHAUS, H. 1956 Sur la division des corps matériels en parties. *Bull. Acad. Polon. Sci.* **4** (12), 801–804.

TIBSHIRANI, R., WALTHER, G. & HASTIE, T. 2001 Estimating the number of clusters in a data set via the gap statistics. *J. R. Stat. Soc.* B **63**, 411–423.

VISHIK, M. I. & FURSIKOV, A. V. 1988 *Mathematical Problems of Statistical Hydrodynamics*. Kluwer.