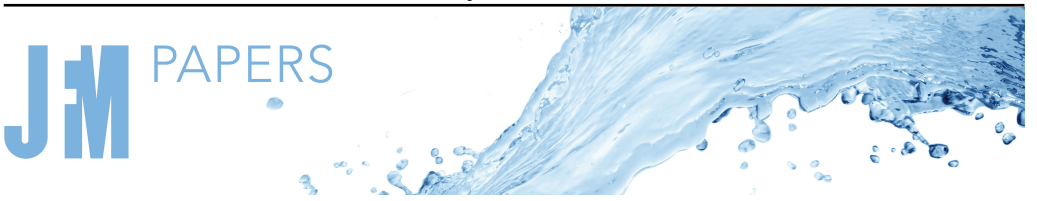
J. Fluid Mech. (2026), *vol.* 1032, A62, doi:10.1017/jfm.2026.11407



Enhancing modal synchronisation in wall-bounded turbulence via sequential variational data assimilation

Chuangxin He[1,2], **Peng Wang**[1,2], **Hyung Jin Sung**[3], **Wenwu Zhou**[1,2], **Di Peng**[1,2] **and Yingzheng Liu**[1,2]

[1]Key Laboratory of Education Ministry for Power Machinery and Engineering, School of Mechanical Engineering, Shanghai Jiao Tong University, 800 Dongchuan Road, Shanghai 200240, PR China

[2]Gas Turbine Research Institute, Shanghai Jiao Tong University, 800 Dongchuan Road, Shanghai 200240, PR China

[3]Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, 291 Daehak-Ro, Yuseong-Gu, Daejeon 34141, Korea

Corresponding author: Yingzheng Liu, yzliu@sjtu.edu.cn

(Received 15 May 2025; revised 3 March 2026; accepted 4 March 2026)

This study explores partial synchronisation in turbulent channel flows using sequential variational data assimilation with sparse observations, emphasising the roles of model and observation uncertainties. Unlike previous work that focused on synchronisation using direct numerical simulation, this study considers synchronisation under imperfect models and noisy data. In the first part, a synchronisation map is constructed, revealing invariance with respect to variations in the predictive model, Reynolds number and mesh resolution. Full synchronisation emerges above a critical level of equivalent observation density. At lower observation densities, modal synchronisation is observed, where the energies of dominant modes evolve independently of initial conditions. As data become sparser, the system transitions to a non-synchronisation regime, with assimilated flows exhibiting minimal correlation with the observations. The second part of this study uses the master flow interpolated from down-sampled sparse observations. The delay-coordinate strategy is introduced to enhance the modal synchronisation. Results indicate that the optimal σ lies near the threshold between modal synchronisation and non-synchronisation. This demonstrates that the modal synchronisation serves as a critical prerequisite for leveraging historical information in data assimilation to improve the accuracy of turbulence reconstruction. These findings extend the scope of synchronisation theory and provide valuable guidance for advancing data assimilation methodologies.

Key words: variational methods, channel flow, Navier–Stokes equations

© The Author(s), 2026. Published by Cambridge University Press.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

1. Introduction

Recovering global flow fields from sparse observations is a central challenge in turbulence research, particularly in experimental settings where even advanced techniques such as tomographic particle image velocimetry (tomo-PIV) and three-dimensional Lagrangian particle tracking (3D-LPT) resolve only a limited range of turbulence scales. Synchronisation theory offers a data assimilation (DA) framework capable of reconstructing the full spectrum of turbulence when large-scale modes are observed up to a critical wavenumber. This criterion is often too restrictive for experimental data, where the resolved scales fall short of the synchronisation threshold. This limitation poses a fundamental question: to what extent can partial synchronisation (only partial turbulence information has been synchronised and the slave systems develop their own chaotic dynamics) be achieved from sparse observations? Addressing this question could pave the way for a new theoretical foundation for reconstructing fine-scale flow structures in turbulent systems under realistic experimental constraints.

Flow-field reconstruction is now a standard procedure in volumetric measurements, commonly used for interpolating velocity fields from discrete particle trajectories and for data augmentation with error reduction. Being different from the reconstruction of mean flow quantities (Deng *et al.* 2018; Li *et al.* 2023, 2024), this process aims to reproduce smaller-scale structures, often for the purpose of achieving super-resolution in turbulent flow analysis (Abe *et al.* 2025). The rapid development of the data-driven DA techniques enables the optimal interpolation of the flow fields from noisy data under specified physical constraints. Radial basis functions (Casa & Krueger 2013; Sperotto *et al.* 2022; Li & Pan 2024), including variants such as FlowFit (Gesemann *et al.* 2016; Huhn *et al.* 2018) and vortex-in-cell-based methods (Schneiders & Scarano 2016; Jeon, Müller & Michaelis 2022), have attracted significant attention in the context of 3D-LPT. These methods approximate fields as linear combinations of a single univariate function and are typically employed for interpolating or smoothing multivariate functions known only at sparse points, while also enforcing compliance with governing equations. The rapid growth of artificial intelligence enables the direct learning of the flow fields in the age of data. Convolutional neural networks can be established to map low-resolution fields to high-resolution ones once trained using both low- and high-resolution databases (Fukami, Fukagata & Taira 2019; Liu *et al.* 2020; Kim *et al.* 2021). High-sampling-rate training data can also be employed to achieve super-temporal resolution (Fukami, Fukagata & Taira 2021). Physics-informed neural networks (Karniadakis *et al.* 2021; Zhu *et al.* 2024) are a class of deep learning algorithms that seamlessly integrate data and the governing equations. The core of flow-field reconstruction involves solving a least-squares problem that balances observational data with physical constraints. These methods fall within the category of physics-constrained data-driven approaches, where the observations define the solution space for optimisation. This implies that the training dataset must encompass all relevant events, limiting the potential to capture new, unobserved dynamics during reconstruction. Consequently, achieving super-resolution solely by refining the grid is challenging, as fine-scale structures cannot be reproduced without additional input or information beyond what the initial observations provide.

A novel concept considered here is termed data-constraint physics-driven, which relies on the physical equations for the reproduction of field data that have an optimal match with the limited observations. The most obvious advantage is the ability for the generation of new event data spans outside the observation space. A typical realisation in turbulent flows utilises the spatiotemporal correlation properties to increase the resolution in either spatial or temporal domain from the existing information in temporal or spatial domain

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

(Scarano & Moore 2012; Scarano *et al.* 2022). Representative work has been conducted by Schneiders *et al.* (2014) for the time-super-sampling of tomo-PIV results based on the weighting averaging of forward and backward time integration using the vorticity equations. Alternatively, this philosophy can be better realised using variational methods. Li *et al.* (2023) reconstructed fine-scale structures to achieve super-spatial-resolution in the flow over an aerofoil from spatially sparse observations using sliding-window four-dimension variational (4DVar) DA. Small-scale vortices are faithfully recovered from the convection of the sparse velocity signals with high sampling frequency. Subsequent work was done by He *et al.* (2024) for the reconstruction of super-temporal-resolution fields from two given instantaneous fields with large time interval, achieving a temporal resolution enhancement exceeding 20-fold the Nyquist sampling constraint. Herein, a high-sampling-rate (or high-spatial-resolution) database is not required in the training for super-temporal-resolution (or super-spatial-resolution) purposes, differing significantly from that in machine learning. Physical information is converted between the spatial and temporal domains, making use of the convective properties of the Navier–Stokes equations.

When fine-scale information is unavailable in both spatial and temporal domains, an alternative physics-driven approach leverages the modulation between different turbulence scales. Prior research has demonstrated that large-scale structures modulate small-scale phenomena in both amplitude and frequency (Mathis, Hutchins & Marusic 2009; Ganapathisubramani *et al.* 2012), with additional evidence of preferential alignment between vortices across scales (Fiscaletti *et al.* 2016). A key strategy for improving flow reconstruction involves introducing external perturbations that are modulated to align with realistic small-scale turbulence dynamics by existing large-scale vortical structures. This approach has been implemented in studies such as those of Li *et al.* (2020) and He *et al.* (2024), where small-scale features were incorporated through initial conditions or observations via the 4DVar method. This has catalysed research on turbulence synchronisation (Yoshida, Yamaguchi & Kaneda 2005), which examines two independently evolving turbulence systems with identical configurations. The synchronisation process involves substituting large-scale vortices below a specific wavenumber in a 'slave' system with corresponding structures from a 'master' system, thereby achieving synchronisation of small-scale vortices. This demonstrates that the complete turbulence spectrum can be reconstructed once large-scale structures are accurately captured. The critical wavenumber for synchronisation determined via direct substitution is approximately $k_c \approx 0.2\eta_K^{-1}$ in both isotropic and rotating turbulence, where η_K denotes the Kolmogorov length scale (Lalescu, Meneveau & Eyink 2013; Li *et al.* 2024). This critical wavenumber corresponds to the threshold at which perturbations with wavenumber $k > k_c$ exhibit negative characteristic Lyapunov exponents in time-dependent turbulent flows (Nikolaidis & Ioannou 2022). For wall-bounded turbulence in physical space, Wang & Zaki (2022) proposed the Taylor microscale as a more suitable metric for defining the synchronisation threshold. Notably, Vela-Martín (2021) observed that partial synchronisation persists even when the wavenumber falls significantly below k_c, particularly in regions of intense vorticity. This finding holds significant practical relevance, as experimental techniques often fail to resolve scales finer than five times the Kolmogorov scale, yet partial synchronisation still remains achievable under such conditions.

This study addresses flow reconstruction from observations which are much sparser than those typically required for synchronisation. To distinguish between levels of synchronisation, the complete alignment of all turbulence scales and kinetic energy

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

observed in prior studies is referred to as full synchronisation, while alignment limited to certain scales is termed partial synchronisation. We additionally observed that numerical inconsistencies between the master and slave simulations, such as those caused by differing mesh configurations or computational parameters, prevent the slave flow from converging to a state that is perfectly identical to the master flow. Nevertheless, multiple slave systems can still converge to mutually identical states, even when their initial conditions differ completely. We also describe this behaviour as full synchronisation. The extended spectral proper orthogonal decomposition (ESPOD) method is used to identify synchronised modes embedded within chaotic turbulence, which remain invariant to initial conditions. This framework provides a clear criterion for assessing the degree of partial synchronisation achievable at different observation densities. Unlike previous methods that directly substitute large-scale structures from the master system into the slave system in spectral space, this work performs partial synchronisation of turbulent channel flows in physical space using a sequential variational DA (SeVar) scheme (He, Liu & Gan 2020, 2022*a*). It is worth noting that synchronisation in physical space is more involved than that in spectral space as there is no clear separation for the large- and small-scale structures, but it is of more practical significance in dealing with complex wall-bounded turbulence. Compared with the 4DVar algorithm, SeVar avoids the costly coupling of forward and backward integrations by decomposing the time domain for the adjoint system. This allows the forward equations to be solved across all time steps, while adjoint iterations are confined to subdomains where observations are available. The approach shares similarities with the nudging technique (Zauner *et al.* 2022), where a feedback term is added to the momentum equations to align predictions with observational trends. Unlike nudging, the present DA method circumvents the challenge of selecting a nudging coefficient by automatically determining the optimal body force through the adjoint loop.

2. Fundamentals and set-up

2.1. *Sequential variational data assimilation*

The SeVar method was originally proposed by He *et al.* (2020) and has since been successfully applied to tomo-PIV measurements in a circular jet (He *et al.* 2022*a*) and planar PIV measurements of turbulent flow around an aerofoil (He *et al.* 2022*b*). The method is briefly recapped here for clarity. In the SeVar framework, a body force $\boldsymbol{F}$ is added to the momentum equations and is optimised to minimise the discrepancy between model predictions and observational data:

$$\frac{\partial \boldsymbol{U}}{\partial t} + (\boldsymbol{U}\cdot\nabla)\,\boldsymbol{U} = -\frac{1}{\rho}\nabla p + \nabla\cdot\left[\nu_{eff}\left(\nabla\boldsymbol{U} + (\nabla\boldsymbol{U})^{T}\right)\right] + \boldsymbol{F}, \tag{2.1}$$

$$\nabla\cdot\boldsymbol{U} = 0. \tag{2.2}$$

The above equations are also named as primary equations in which ν_{eff} denotes the effective viscosity used in either direct numerical simulation (DNS) or large-eddy simulation (LES). The flow velocity $\boldsymbol{U}$ and pressure p are known as primary variables distinguishing from those in the adjoint system. The objective of DA is to minimise the cost function $\mathcal{J}$ subject to the governing equations by optimising the body force distribution:

$$\mathcal{J} = \xi \int_{t_{obs}}^{t_{obs}+\delta t} \int_{\Omega} M \left(\frac{\boldsymbol{U}-\boldsymbol{U}_{obs}}{U_0}\right)^2 \mathrm{d}\Omega\,\mathrm{d}t. \tag{2.3}$$

Here, t_{obs} and δt are the instant when the observations exist and the time step of the computation respectively. Term $\boldsymbol{U}_{obs}$ is the velocity observation. It is noted that

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

optimisation is performed in a temporal window with only one time step when the observations exist, rather than in a long window used by He *et al.* (2024). This yields a sequential process that avoids the backward integration of the adjoint system. In fact, the time integral in the cost function above can be omitted, but it is retained here in integral form to maintain consistency with the formulations used in 4DVar and to facilitate the understanding of the derivation of the SeVar equations. Function M is the mask function used to specify the spatial sites where the observations exist and the weight to count for their relative confidence level. Term ξ is a dimension converter with dimension $[\mathrm{L}^2\ \mathrm{T}^{-3}]$ (power per unit mass) and value unity to address the dimensional inconsistency. Parameter U_0 which can be usually set as the free-stream velocity is employed to adjust the Courant number of the adjoint equations so that the same time step can be used in the computation of both primary and adjoint systems. Term Ω represents the computational domain while $\mathcal{R} = (\mathcal{R}_{NS}, \mathcal{R}_C)$ are the incompressible unsteady Navier–Stokes equations and the continuity equation. The Lagrange function $\mathcal{L}$ can thus be expressed as

$$\mathcal{L} = \mathcal{J} + \int_{t_{obs}}^{t_{obs}+\delta t} \int_{\Omega} (\boldsymbol{V}, q)\, \mathcal{R}\, \mathrm{d}\Omega \mathrm{d}t, \tag{2.4}$$

where the adjoint velocity $\boldsymbol{V}$ and the adjoint pressure q are introduced as the Lagrange multipliers. The resultant adjoint equations are presented below directly while readers are referred to Appendix A and He *et al.* (2020) for a detailed derivation:

$$\frac{\boldsymbol{V}}{\delta t} + \nabla \boldsymbol{U} \cdot \boldsymbol{V} - (\boldsymbol{U} \cdot \nabla)\, \boldsymbol{V} - \nabla \cdot \left[\nu_{eff}\left(\nabla \boldsymbol{V} + (\nabla \boldsymbol{V})^T\right)\right] + \frac{1}{\rho}\nabla q + 2\xi M \frac{\boldsymbol{U} - \boldsymbol{U}_{obs}}{U_0^2} = 0, \tag{2.5}$$

$$\nabla \cdot \boldsymbol{V} = 0. \tag{2.6}$$

It is noted that $\boldsymbol{V}/\delta t$ appears in the first term of (2.5) instead of the time derivative as the first-order Euler scheme and the zero condition of the adjoint velocity at the terminal time are used. After the adjoint state variables are obtained, the body force can be updated gradually using the deepest descending algorithm according to

$$\boldsymbol{F} = \boldsymbol{F} - \lambda_F \frac{\partial \mathcal{L}}{\partial \boldsymbol{F}} = \boldsymbol{F} + \lambda_F \boldsymbol{V}. \tag{2.7}$$

The step size λ_F can be readily adjusted to achieve the maximum convergence speed in the first round of iteration. It is important to note that although the parameter λ_F is also used for tuning, its value solely influences the iteration's convergence speed. The principle of the SeVar algorithm is illustrated in figure 1. In the temporal subdomain, the linear adjoint equations are computed at the physical instants where observations are available. When there are no observations, the conventional simulation with $\boldsymbol{F} = 0$ is obtained. Each iteration round involves computing the primary and adjoint equations just once, followed by the update of the body force. In this study, only 2–3 rounds of iteration are needed to decrease the relative value of the objective function to less than 1 %, significantly reducing the computational cost compared with the traditional 4DVar algorithm. When observations are incorporated every 10 time steps, the overall computational cost increases by approximately 30 % compared with conventional simulations. Although this approach is more time-consuming than the nudging technique, it is currently the preferred basis as it can be readily extended for further advanced 4DVar research.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

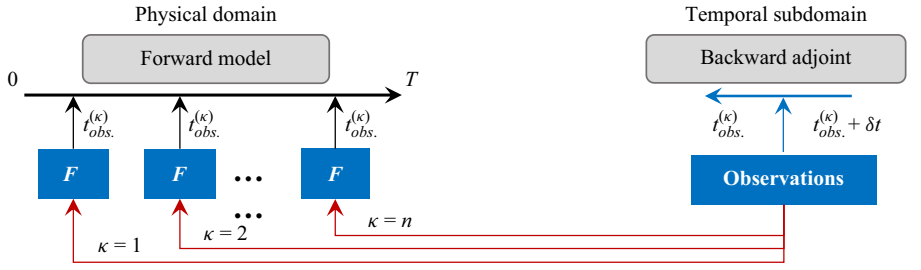


Figure 1. Principle of the SeVar algorithm.

2.2. *Extended spectral proper orthogonal decomposition*

The synchronisation experiments conducted by Yoshida *et al.* (2005), Nikolaidis & Ioannou (2022) and Li *et al.* (2024) involve two turbulence systems, i.e. a master and a slave, which evolve independently prior to the DA implementation. A wavenumber truncation k_t is defined to separate the large- and small-scale vortical structures in both systems. When k_t exceeds a critical value k_c, replacing the large scales in the slave system with those from the master system ensures identical spatiotemporal evolution of all vortices in the slave. In the present study, physical space is employed, allowing the truncation to be naturally determined by the density of observations extracted from the master system. Here, the SeVar method is applied to replace sparse observational data in the slave domain with information from the master system. Consequently, the global flow field in the slave system begins to resemble the master flow, even when the observation density falls below the threshold required for full synchronisation. To quantify this behaviour, a criterion for assessing the degree of partial synchronisation is introduced, based on ESPOD analysis.

The SPOD method, introduced by Towne, Schmidt & Colonius (2018), is a variation of the classical POD method (Lumley 1967). It is designed to extract energetic, orthogonal modes at individual frequencies by applying singular value decomposition to the cross-spectral density (CSD) matrix. The CSD is usually obtained from time-series data using Welch's method (Welch 1967), by dividing the entire dataset into numerous segments in time with a specified overlap between two successive segments. The discrete Fourier transform (DFT) of the sparse observations in time direction yields

$$\hat{\boldsymbol{U}}_{obs}\left(\boldsymbol{x}, \tau, f\right) = \left[\hat{\boldsymbol{u}}_{obs}\left(\boldsymbol{x}, 1, f\right), \hat{\boldsymbol{u}}_{obs}\left(\boldsymbol{x}, 2, f\right), \ldots, \hat{\boldsymbol{u}}_{obs}\left(\boldsymbol{x}, N_s, f\right)\right], \tag{2.8}$$

where $\tau = 1, 2, \ldots, N_s$ is the index of the data segment and $\hat{\boldsymbol{u}}_{obs}(\boldsymbol{x}, n, f)$ is the DFT of the observations in the nth segment. Term f is the frequency which is omitted for simplification in all the following formulations. For each frequency f the CSD is defined as

$$\hat{\boldsymbol{S}}\left(\tau, \tau'\right) = \sum_{\boldsymbol{x}} \sum_{\boldsymbol{x}'} \hat{\boldsymbol{U}}_{obs}^{*}\left(\boldsymbol{x}, \tau\right) w\left(\boldsymbol{x}, \boldsymbol{x}'\right) \hat{\boldsymbol{U}}_{obs}\left(\boldsymbol{x}', \tau'\right). \tag{2.9}$$

Here, w denotes spatial weight of the observations and is unity in the present study. The asterisk denotes the conjugate transpose of a matrix. It is noted here that the CSD is defined in a form similar to that in the snapshot POD method (Sirovich 1987). The eigenvalue problem reads

$$\sum_{\tau'=1}^{N_s} \hat{\boldsymbol{S}}\left(\tau, \tau'\right) \hat{\boldsymbol{\Phi}}\left(\tau', i\right) = \hat{\lambda}_i \hat{\boldsymbol{\Phi}}\left(\tau, i\right), \tag{2.10}$$

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

with the ith ($i = 1, 2, \ldots, N_s$) SPOD mode at the specified frequency derived as

$$\hat{\boldsymbol{\varphi}}\,(\boldsymbol{x}, i) = \sum\nolimits_{\tau=1}^{N_s} \hat{\boldsymbol{U}}_{obs}\,(\boldsymbol{x}, \tau)\,\hat{\boldsymbol{\Phi}}\,(\tau, i). \tag{2.11}$$

The corresponding ith mode coefficient is determined as

$$\hat{\alpha}\,(i, \tau) = \sum\nolimits_{\boldsymbol{x}} \hat{\boldsymbol{\varphi}}^{*}\,(\boldsymbol{x}, i)\,\hat{\boldsymbol{U}}_{obs}\,(\boldsymbol{x}, \tau). \tag{2.12}$$

This enables the restoration of the DFT modes via

$$\hat{\boldsymbol{u}}_{obs}\,(\boldsymbol{x}, \tau) = \sum\nolimits_{i=1}^{N_s} \hat{\boldsymbol{\varphi}}\,(\boldsymbol{x}, i)\,\hat{\alpha}\,(i, \tau). \tag{2.13}$$

The physical representation of the SPOD modes can be obtained by the inverse Fourier transform applied to $\hat{\boldsymbol{\varphi}}(\boldsymbol{x}, i)$. Before doing this, we would rather like to know how the global flow field is driven by these sparse observations. To this end, the SPOD modes are extended to each wall-parallel slice of the global flow in the master system using the method proposed by Boree (2003). Similar to (2.8), the DFTs of the flow slice in the master system for each frequency are arranged as

$$\hat{\boldsymbol{U}}_m\,(\boldsymbol{x}, \tau) = \left[\hat{\boldsymbol{u}}_m\,(\boldsymbol{x}, 1)\,, \hat{\boldsymbol{u}}_m\,(\boldsymbol{x}, 2)\,, \ldots, \hat{\boldsymbol{u}}_m\,(\boldsymbol{x}, N_s)\right], \tag{2.14}$$

where the subscript 'm' denotes the master system. The extended modes on each slice are determined as

$$\hat{\boldsymbol{\psi}}\,(\boldsymbol{x}, i) = \frac{\sum_{\tau=1}^{N_s} \hat{\boldsymbol{U}}_m\,(\boldsymbol{x}, \tau)\,\hat{\alpha}^{*}\,(i, \tau)}{\sum_{\tau=1}^{N_s} \hat{\alpha}\,(i, \tau)\,\hat{\alpha}^{*}\,(i, \tau)}. \tag{2.15}$$

According to the properties of the ESPOD modes, $\hat{\boldsymbol{\psi}}$ represents the only part of the flow in the master system that is correlated with the dynamic behaviours of the sparse observations. As the turbulence information contained in the observations varies with the variation of the observation density, the energy of the extended modes adjusts accordingly, directly reflecting the degree of partial synchronisation achieved in the sparse-data-driven slave system. The correlated part of the flow can be reconstructed using the same model coefficients $\hat{\alpha}$, i.e.

$$\hat{\boldsymbol{u}}_{mc}\,(\boldsymbol{x}, \tau) = \sum\nolimits_{i=1}^{N_s} \hat{\boldsymbol{\psi}}\,(\boldsymbol{x}, i)\,\hat{\alpha}\,(i, \tau), \tag{2.16}$$

while the following relation holds for each frequency:

$$\sum\nolimits_{\tau=1}^{N_s} \left[\hat{\boldsymbol{u}}_m\,(\boldsymbol{x}, \tau) - \hat{\boldsymbol{u}}_{mc}\,(\boldsymbol{x}, \tau)\right] \hat{\boldsymbol{u}}^{*}_{obs}\,(\boldsymbol{x}, \tau) = 0. \tag{2.17}$$

Equation (2.17) reveals that only the correlated flow component $\hat{\boldsymbol{u}}_{mc}$ in a turbulence system can be driven by sparse observations. Specifically, $\hat{\boldsymbol{u}}_{mc}$ is the only component that exhibits dynamical behaviour consistent with the observations, as reflected by the shared coefficients $\hat{\alpha}$ in DA. The remaining component, $\hat{\boldsymbol{u}}_m - \hat{\boldsymbol{u}}_{mc}$, arises from the turbulence production due to the nonlinearity of the Navier–Stokes equations. Notably, when the observations possess sufficient degree of freedom, $\hat{\boldsymbol{u}}_m = \hat{\boldsymbol{u}}_{mc}$ holds in the slave system. This implies that the entire energy cascade can be reconstructed from the observations once sufficient dynamics have been captured by the SPOD modes.

The motivation in the present study is to identify the synchronised modes in the slave system with DA implementation. It can be achieved by projecting a data slice in the slave onto the ESPOD modes $\hat{\boldsymbol{\psi}}$ derived earlier. This process is rather involved as the ESPOD

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

modes are not orthogonal. The Gram–Smith orthogonalisation is thus applied to obtain an orthogonal basis $\hat{\boldsymbol{\psi}}_{orth}$ that has the same span as $\hat{\boldsymbol{\psi}}$. The correlated flow in the slave system is computed as

$$\hat{\boldsymbol{u}}_{sc}(\boldsymbol{x}, \tau) = \sum_{i=1}^{N_s} \left[\hat{\boldsymbol{\psi}}(\boldsymbol{x}, i) \sum_{\boldsymbol{x}} \left(\hat{\boldsymbol{\psi}}^*_{orth}(\boldsymbol{x}, i) \hat{\boldsymbol{U}}_s(\boldsymbol{x}, \tau) \right) \right], \tag{2.18}$$

where subscript 's' denotes the slave system. Here $\hat{\boldsymbol{U}}_s$ is the DFT of the flow slice in the slave system constructed from $\hat{\boldsymbol{u}}_s$ in a manner analogous to (2.14) and $\hat{\boldsymbol{u}}_{sc}$ is the corresponding synchronised flow. This enables the computation of the synchronised mode coefficients $\hat{\beta}$ in the slave system satisfying the relation

$$\hat{\boldsymbol{u}}_{sc}(\boldsymbol{x}, \tau) = \sum_{i=1}^{N_s} \hat{\boldsymbol{\psi}}(\boldsymbol{x}, i) \hat{\beta}(i, \tau). \tag{2.19}$$

As the correlated flow $\hat{\boldsymbol{u}}_{sc}$ has the same span as modes $\hat{\boldsymbol{\psi}}$, $\hat{\beta}$ can be precisely determined by solving the linear equation (2.19). Upon full development of the assimilated flow, specific modes in both master and slave systems exhibit synchronised temporal evolution that remains invariant to changes in initial conditions. These ESPOD modes, hereafter referred to as synchronised modes, can be quantified through cross-correlation analysis of their mode coefficients at each frequency:

$$\mathcal{C}(i) = \frac{\left\langle \hat{\alpha}(i, \tau), \hat{\beta}(i, \tau) \right\rangle}{\left\| \hat{\alpha}(i, \tau) \right\| \cdot \left\| \hat{\beta}(i, \tau) \right\|}. \tag{2.20}$$

Here, the inner product $\langle\rangle$ and norm $\|\|$ denote computations performed separately for each mode i. Modal synchronisation is defined when the cross-correlation satisfies $\mathcal{C} > 0.9$. The energy ratio of the ith synchronised mode, representing the portion of energy that can be driven in the slave system by sparse observations, is calculated as

$$\varkappa(i) = \frac{\left\langle \hat{\beta}(i, \tau), \hat{\beta}(i, \tau) \right\rangle}{\left\langle \hat{\alpha}(i, \tau), \hat{\alpha}(i, \tau) \right\rangle}. \tag{2.21}$$

2.3. *Computational set-up*

The present study focuses on the fully developed turbulent channel flow with a computational domain of dimensions $L_x \times L_y (2H) \times L_z$. For each DA implementation in the slave system, a master system with an identical computational domain is precomputed using high-fidelity DNS or LES. The observations for DA are incorporated by imposing a mask function M to the global master flow as shown in (2.5), where M is defined as

$$M = \exp\left\{ -\sigma \left[\left(\frac{x - x_0}{\Delta x} \right)^2 + \left(\frac{y - y_0}{\Delta y} \right)^2 + \left(\frac{z - z_0}{\Delta z} \right)^2 \right] \right\}, \tag{2.22}$$

where (x_0, y_0, z_0) denote the centre coordinates of the mesh cells where the observations are located (marked by the yellow star in figure 2). These cells are randomly selected according to their indexes in the computational domain and remain stationary in time. As a result, observation points are more densely distributed in regions with finer mesh scales near the wall, ensuring that the ratio of turbulence scales captured by sparse observations to local mesh filter scales remains constant at different positions in the channel. The coordinates (x, y, z) represent the centres of mesh cells surrounding (x_0, y_0, z_0) within

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

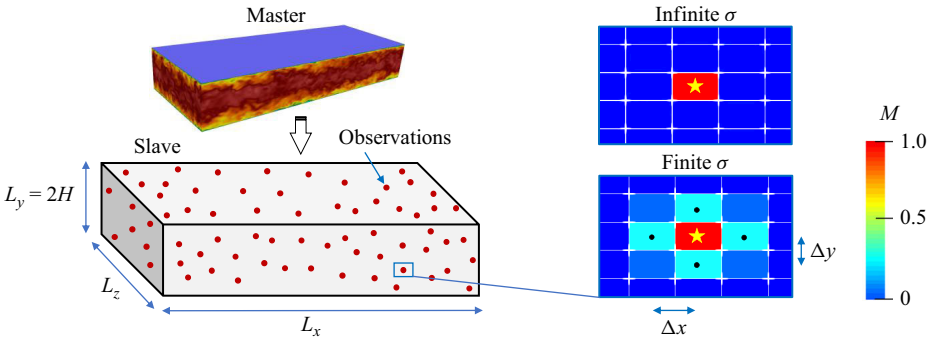


Figure 2. Schematic diagram of the computational set-up.

a $3 \times 3 \times 3$ cell box. The distances between these cell centres and (x_0, y_0, z_0) are denoted by Δx, Δy and Δz in each respective direction. The weighting relaxation coefficient σ governs the behaviour of the mask function: when σ approaches infinity, it results in single-cell 'nudging', while $\sigma = 0$ produces a 'solid-nudging' effect (Clark Di Leoni, Mazzino & Biferale 2020). For finite values of σ, M decreases radially way from (x_0, y_0, z_0). Mask function M is set to zero outside the $3 \times 3 \times 3$ cell box. The condition $M < 1$ $(\sigma > 0)$ in the surrounding cells accounts for scenarios where only the centre cells (marked with a star) provide exact observations, while the surrounding values are interpolated to enhance the DA performance. This situation commonly arises in practical applications, where experimental measurements typically capture flow information only at the centre cells, leaving data on the surrounding cells unavailable. Data assimilation that relies solely on velocity measurements from the centre cells would therefore require a relatively dense set of observations, as is demonstrated later. Whether the surrounding-cell velocities are obtained by directly copying the centre-cell values or by interpolating the centre-cell measurements, the resulting observations inevitably contain substantial errors. In this study, we introduce a parameter $\sigma > 0$ to down-weight the observations on the surrounding cells, which significantly reduces the influence of these observational errors and improves the agreement between the assimilated and true flow fields. Notably, the length scales in the mask function are normalised by the local mesh scales Δx, Δy and Δz rather than the commonly defined turbulence scales, since LES inherently involves explicit mesh truncation. Physically, such mesh scales represent the smallest resolved turbulent structures present in LES turbulent systems. The results will demonstrate that synchronisation primarily depends on mesh-based observation density, with the turbulence-scale-based criterion becoming relevant only when the mesh resolution meets DNS requirements. This does not violate the fundamental principle that turbulence synchronisation relies on small-scale information in turbulent systems. In fact, previous studies have predominantly employed DNS, in which the characteristic small-scale structures in the system correspond to the Kolmogorov (Lalescu *et al.* 2013) or Taylor (Wang & Zaki 2022) scales. However, in LES, the smallest scales are truncated by the LES mesh filter. The synchronisation criterion naturally shifts to the mesh-based density of the observational data. The observation density $\mathcal{S}$ is defined as

$$\mathcal{S} = \sqrt[3]{\frac{N_{obs}}{N_x \times N_y \times N_z}}, \tag{2.23}$$

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

Re_τ	$L_x \times L_y \times L_z$	Case no.	$N_x \times N_y \times N_z$	x^+	y^+	z^+	L_{mesh}	Type
		1	$450 \times 240 \times 360$	4.8	0.17	3.0	0.034	DNS
		2	$300 \times 160 \times 240$	7.0	0.25	4.4	0.051	DNS
180	$4\pi \times 2 \times 2\pi$	3	$192 \times 130 \times 160$	10.9	0.25	6.6	0.078	DNS
		4	$128 \times 90 \times 110$	16.2	0.33	9.4	0.116	q-DNS*
		5	$85 \times 60 \times 75$	24.8	0.42	14.1	0.173	q-DNS*
		6	$500 \times 150 \times 250$	18.4	0.55	13.8	0.034	LES
		7	$400 \times 130 \times 200$	23.0	1.1	17.3	0.042	LES
395	$4\pi \times 2 \times 1.5\pi$	8	$300 \times 120 \times 150$	30.7	1.1	23.0	0.055	LES
		9	$200 \times 80 \times 100$	46.2	1.47	34.6	0.082	LES
		10	$120 \times 60 \times 80$	76.7	2.93	43.1	0.125	q-LES*
		11	$80 \times 45 \times 50$	115.1	2.93	69.0	0.189	q-LES*
		12	$500 \times 150 \times 250$	30.2	0.9	22.6	0.034	LES
		13	$400 \times 130 \times 200$	37.7	1.8	28.3	0.042	LES
640	$4\pi \times 2 \times 1.5\pi$	14	$300 \times 120 \times 150$	50.3	1.8	37.7	0.055	LES
		15	$200 \times 80 \times 100$	75.4	2.4	56.5	0.082	q-LES*
		16	$120 \times 60 \times 80$	125.7	4.8	70.7	0.125	q-LES*
		17	$80 \times 45 \times 50$	188.5	4.8	113.1	0.189	q-LES*

Table 1. List of case configurations.

*For these cases, mesh resolutions are insufficient for DNS or LES.

where N_{obs} is the number of star-marked centre cells and N_x, N_y and N_z are the number of cells in x, y and z directions of the computational domain, respectively. In both the master and slave systems, a periodic boundary condition is applied to the front–back and inlet–outlet surfaces. A constant pressure gradient is imposed in the streamwise direction to achieve the target Reynolds numbers. Both DNS and LES employ the second-order central scheme for spatial discretisation and the second-order backward scheme for temporal discretisation. The time step is adjusted to maintain a maximum Courant number of approximately 0.5. The LES utilise the wall-adapting local eddy-viscosity model for subgrid stress modelling. All computations, including DNS, LES and DA, are performed on the OpenFOAM (www.openfoam.org) platform.

The cases investigated in the present study are summarised in table 1. Three Reynolds numbers are considered with only the case of $Re_\tau = 180$ simulated using DNS. Observations are denoted in different ways as described in table 2. Unless otherwise emphasised, each slave DA system uses the observations from the master system with the same mesh and set-up (observation I). In other words, within a given case number, the only difference between the master and slave systems is the initial condition. However, observations mapped from higher-resolution case (observation II) or interpolated from the down-sampled centre-cell values (observation III) are employed to support the analysis. The computation of each slave system is conducted for different values of $\mathcal{S}$ and σ. Notably, in full synchronisation study (Wang & Zaki 2022), DNS with identical computational mesh and discretisation scheme is required for both maser and slave to eliminate uncertainties arising from numerical methods. However, this study adopts a 'weak-constraint' approach to account for computational uncertainties such as subgrid model errors and discretisation errors, making it more applicable to realistic flow conditions. All master flows are validated against DNS data (Abe, Kawamura & Matsuo 2001), while the DNS or LES cases marked with an asterisk fail to accurately capture

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

Observation no.	Description
I	All the observations are derived from a master system using the same mesh and numerical set-up
II	All the observations are mapped from a master system using a higher-resolution mesh. It is exclusively applied to case 14 which uses a master system computed in case 12.
III	Only the observations in the centre cells (marked with a star in figure 2) are derived from the master system, while the data in the other cells are linearly interpolated from these centre-cell values

Table 2. List of observation configurations.

the mean velocity and fluctuation profiles due to insufficient mesh resolution. These are specially marked with 'q-' in the table. Nonetheless, these cases serve as a valuable test case to examine the dependence of modal synchronisation on master flow validity and the impact of significant model and discretisation errors. The assimilation procedure is executed every 10 time steps to minimise storage and computational costs, a sufficient temporal density to avoid affecting the DA results as demonstrated by Wang *et al.* (2022). For each case, 1000 instantaneous flow fields are saved, with modal synchronisation analysis conducted using the final 500 snapshots.

3. Results and discussion

3.1. *Synchronisation with sparse observations*

In the slave systems, the fluid in unobserved regions – those without direct measurements – also undergoes correlated motion as synchronisation develops. The flow dynamics in these regions remain strongly correlated with the available observational data. Our primary objective is to identify and extract these embedded flow components from the chaotic turbulent system. The master flow is directly integrated into the slave system, maintaining identical mesh resolution and numerical discretisation for each DA case. Observations from both the central and surrounding cells (as shown in figure 2) are sourced directly from the master flow. In the first DA scenario with $\sigma = +\infty$, a single-cell assimilation approach is applied, disregarding observations from surrounding cells. In all slave systems, the initial conditions are deliberately set to differ significantly from those of the master flow, as shown in the first row of figure 3 where case 3 as listed in table 1 is used for demonstration. At the onset of DA, spot defects appear on the velocity contours at the observation sites, reflecting the underdevelopment of global flow in the slave system. The motivation behind the SeVar approach is to achieve global synchronisation of the slave flow by leveraging inherent spatiotemporal correlations. This process unfolds over time, as evidenced in the second row of figure 3, where smooth velocity contours emerge at $t^+ = 1000$. Full synchronisation is defined as the state where pointwise errors (normalised by bulk velocity) are reduced to less than 1 %. The results show that this criterion is satisfied for the entire slave flow at an observation density of $\mathcal{S} = 0.33$. However, at lower densities such as $\mathcal{S} = 0.25$, while the assimilated flow maintains strong overall similarity to the master flow, discernible pointwise errors persist near the wall region. Further reduction

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

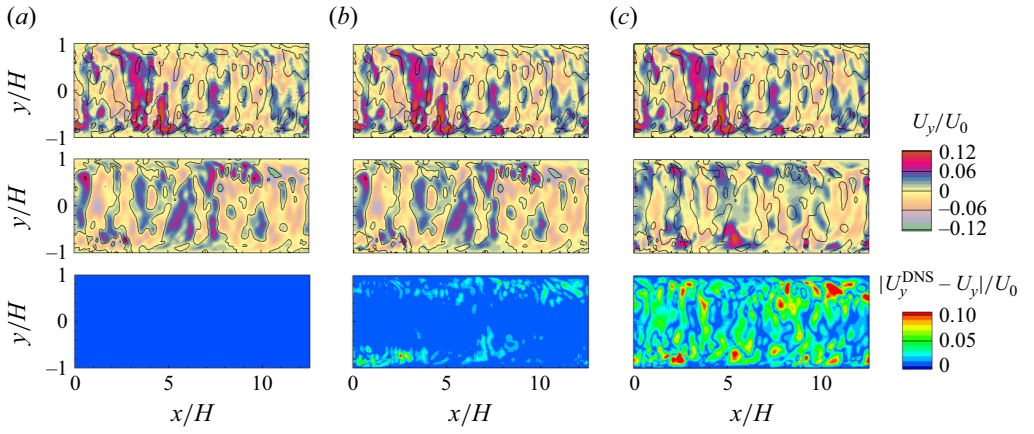


Figure 3. Flow fields and error distributions in the slave systems with different densities at $Re_\tau = 180$ in case 3 with $\sigma = +\infty$: (*a*) $\mathcal{S} = 0.33$, (*b*) $\mathcal{S} = 0.25$ and (*c*) $\mathcal{S} = 0.12$. First row: the initial time step; second row: long time; third row: errors at long time. Contours in the first and second rows denote the vertical velocity of the assimilated slave flows, while the black curves represent the contours of $U_y/U_0 = \pm 0.1$ in the DNS master flow.

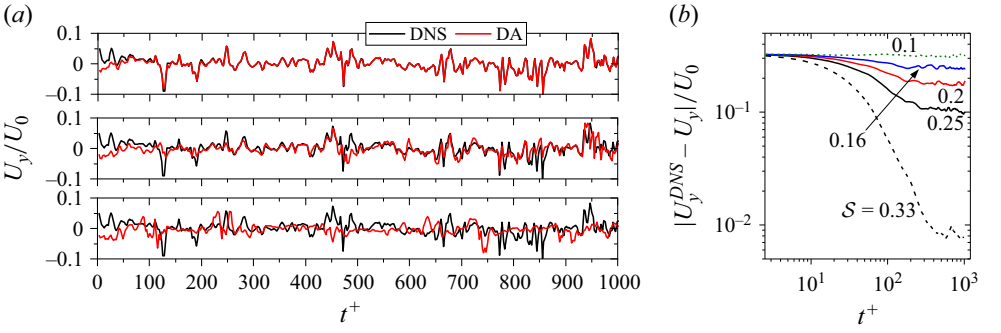


Figure 4. Asymptotic behaviours of the vertical velocity at $y^+ = 15$ in case 3 with $\sigma = +\infty$. (*a*) Temporal variation at a selected point. First row: $\mathcal{S} = 0.33$; second row: $\mathcal{S} = 0.25$; third row: $\mathcal{S} = 0.12$. (*b*) Normalised errors at various densities.

in observation density to $\mathcal{S} = 0.12$ results in significantly degraded correlation between the assimilated and master flows, with substantial errors distributed across the channel.

The asymptotic behaviours of the assimilated flows are shown in figure 4, where the temporal variation of vertical velocities in a free cell (without observation located) at $y^+ = 15$ and the normalised errors are presented. As no observations are available at the selected point, the assimilated slave flow initially exhibits significant deviation from the master DNS data as shown in figure 4(*a*). In the full-synchronisation case at $\mathcal{S} = 0.33$, the DA data rapidly converge to the DNS results, with the two curves overlapping for $t^+ > 200$. While full synchronisation also occurs for $\mathcal{S} > 0.33$, such cases are not discussed in this study. For the non-synchronisation case at $\mathcal{S} = 0.12$, the two curves evolve rather independently without any correlation that can be discerned in the temporal domain. Focusing on an intermediate density $\mathcal{S} = 0.25$, the DA data display strong temporal correlation with the DNS results, though they do not fully coincide. Figure 4(*b*) further clarifies the asymptotic trends by depicting the normalised errors across different observation densities. At $\mathcal{S} = 0.10$ and 0.12, the errors remain consistently high

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

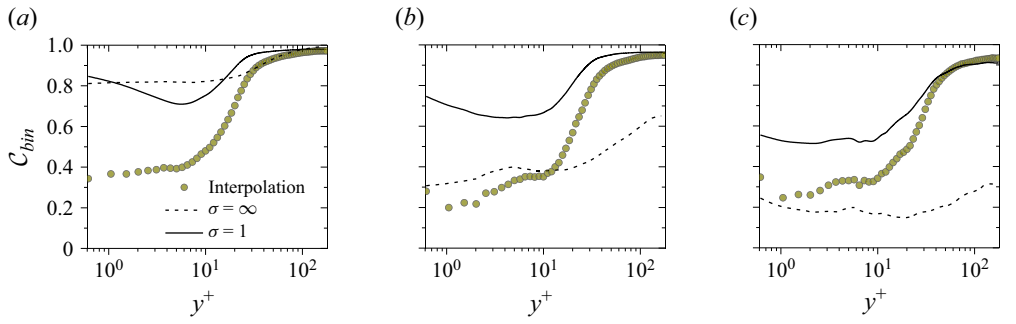


Figure 5. Bin-wise correlation of the flows with the master DNS at $Re_\tau = 180$ in case 3: (*a*) $\mathcal{S} = 0.25$, (*b*) $\mathcal{S} = 0.16$ and (*c*) $\mathcal{S} = 0.12$.

throughout the assimilation period. However, a notable reduction in errors emerges at later times for densities in the range $\mathcal{S} = 0.16$–0.25, albeit with residuals still exceeding those of the $\mathcal{S} = 0.33$ case. These findings suggest that partial flow recovery occurs in the slave systems for intermediate densities $\mathcal{S} = 0.16$–0.25, even though full synchronisation is unattainable due to insufficient observation density.

The adoption of a $3 \times 3 \times 3$ assimilation cell configuration, weighted by coefficient σ, is motivated by practical experimental conditions where observational data are typically limited to centre cells. The inclusion of surrounding cells with diminished weighting (reflecting lower confidence) serves to enhance DA, thereby improving flow reconstruction accuracy despite sparse direct measurements. While expanding the assimilation cube to larger dimensions (e.g. $5 \times 5 \times 5$ or beyond) might improve accuracy further, such benefits would be inherently constrained by a threshold determined by surrounding estimation errors. In the present study, the values in the surrounding cells are estimated through linear interpolation of the sparse observations (observation III). The confidence level of these interpolated values is reduced by increasing the coefficient σ from zero during the DA process. Here, we consider DNS cases at $Re_\tau = 180$ (case 3) with densities $\mathcal{S} = 0.12$, 0.16 and 0.25 below the full-synchronisation threshold. To evaluate the similarity between the assimilated flows and the reference DNS data within the assimilation cube, a bin-wise correlation metric is defined as

$$\mathcal{C}_{bin}(y) = \frac{\langle \boldsymbol{u}, \boldsymbol{u}_{DNS} \rangle_y}{\|\boldsymbol{u}\|_y \cdot \|\boldsymbol{u}_{DNS}\|_y}. \tag{3.1}$$

In this formulation, $\boldsymbol{u}_{DNS}$ represents the fluctuating component of the DNS data, with all vector components and cell values within the $3 \times 3 \times 3$ assimilation cube arranged as a single vector. The same vectorisation approach is applied to $\boldsymbol{u}$, which may represent either DA results or interpolated observations. Temporal averaging is performed for each wall-normal position at identical wall distances. Figure 5 presents the wall-normal distribution of $\mathcal{C}_{bin}$. The interpolation accuracy depends strongly on the characteristic length scales of vortical structures at a given observation density. Values of $\mathcal{C}_{bin}$ are higher in the channel centre region and decrease progressively towards the channel walls. A particularly rapid decline in $\mathcal{C}_{bin}$ is observed in the interpolated field (approximately equivalent to DA $\sigma = 0$) within the buffer layer ($y^+ = 10$–30) where intense small-scale vortical structures are generated. The case $\sigma = \infty$ corresponds to single-cell assimilation, where M has non-zero values only in the centre cells. This configuration leads to significant improvement in $\mathcal{C}_{bin}$ within the viscous sublayer and buffer layer for DA at $\mathcal{S} = 0.25$, as demonstrated in figure 5(*a*). However, this single-cell assimilation approach shows deteriorating correlation

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

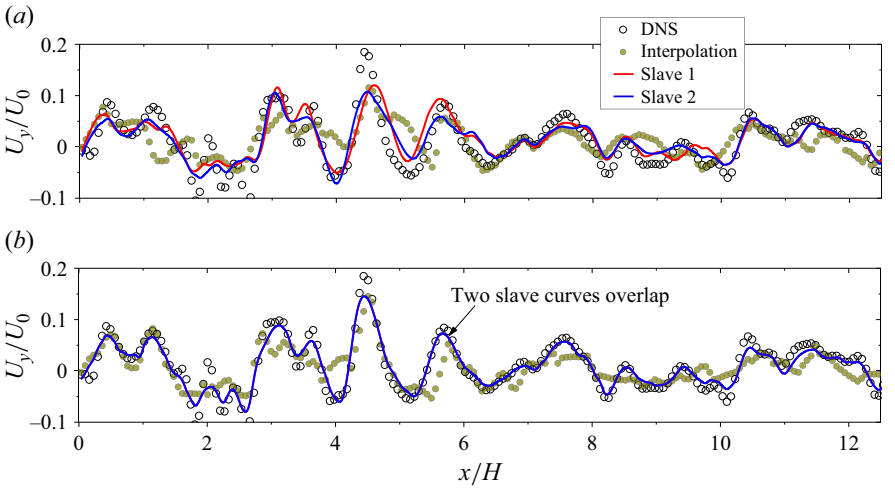


Figure 6. Streamwise variations of the wall-normal velocity component at $y^+ = 40$ and $Re_\tau = 180$ in case 3 for $\sigma = 1$: (*a*) $\mathcal{S} = 0.12$ and (*b*) $\mathcal{S} = 0.16$.

performance with sparser observations as evidenced in figures 5(*b*) and 5(*c*). The assimilation using $\sigma = 1$ significantly enhances correlation in both the sublayer and buffer layer across all considered densities. At $\mathcal{S} = 0.12$ and 0.16, this correlation substantially exceeds that observed with $\sigma = \infty$, while at $\mathcal{S} = 0.25$, it becomes marginally lower in the region beneath the buffer layer ($y^+ = 1$–20). These findings demonstrate that the degree of synchronisation exhibits a positive correlation with $\mathcal{C}_{bin}$. More importantly, DA employing interpolated observations with an appropriately selected σ value proves highly effective in enhancing the synchronisation level, thereby improving the accuracy of global flow reconstruction from sparse observations. Conversely, figures 5(*b*) and 5(*c*) reveal that DA implementation in non-synchronised regimes results in even lower $\mathcal{C}_{bin}$ values than direct interpolation, particularly within the logarithmic layer. This indicates that the current DA framework fails to achieve global flow reconstruction when non-synchronisation dominates at the present observation density. While using a finite σ value effectively increases the equivalent observation density $\mathcal{S}$ and may appear advantageous for flow reconstruction, it is crucial to recognise that an optimal σ exists for each specific density. This is evidenced by the counterexample of $\sigma = 0$, where the slave flow in the assimilating cube converges towards the direct interpolation result. Furthermore, the $\mathcal{C}_{bin}$ values in the $y^+ = 1$–20 region shown in figure 5(*a*) suggest that smaller σ values should be employed at higher observation densities.

Figure 6 presents the streamwise variation of wall-normal velocities at a selected point, with $\mathcal{S} = 0.12$ and 0.16 selected for demonstration. For each density, two assimilated slave systems with uncorrelated initial conditions are compared. The two slave systems exhibit nearly identical streamwise velocity variations at $\mathcal{S} = 0.12$, with only minor discrepancies apparent in figure 6(*a*). Notably, these assimilated velocity distributions capture the essential features of the real flow significantly better than the interpolation when compared with DNS data. When the density increases to $\mathcal{S} = 0.16$, the two slave flows overlap completely, though noticeable differences from the master DNS data emerge as shown in figure 6(*b*). This state does not precisely reproduce the interpolated field, but rather generates a new state that more closely resembles the DNS results. This finding demonstrates that global flow-field reconstruction remains possible even when the interpolated observations capture substantially different flow details. The primary

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

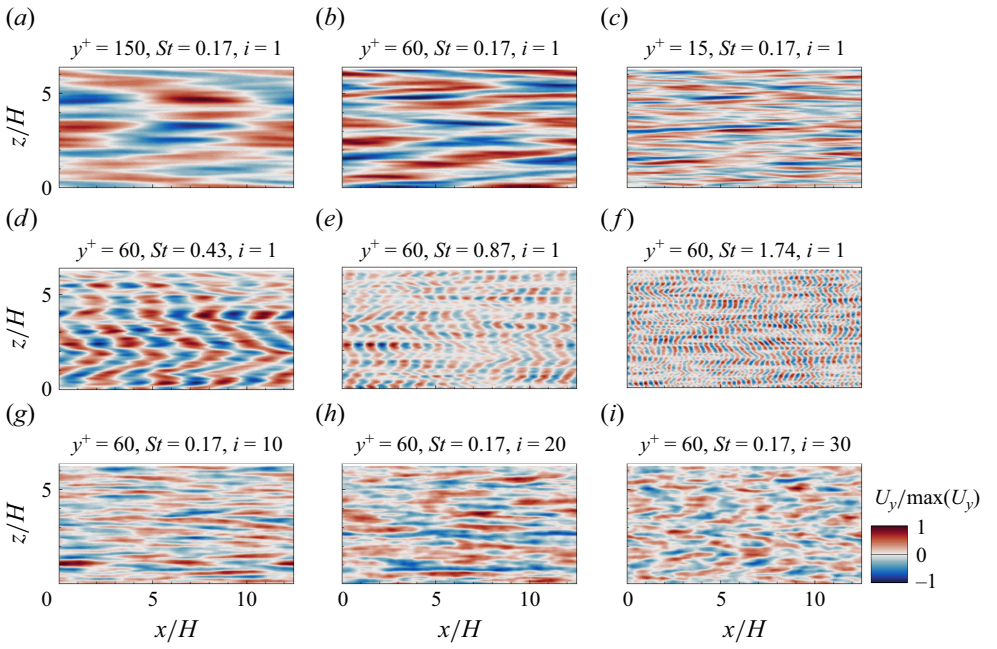


Figure 7. The ESPOD modes in the master system at $\mathcal{S} = 0.25$ with $\sigma = +\infty$ in case 3. Vertical component of the velocity normalised by the maximal value is shown (*a–c*) at different wall distances y^+, (*d–f*) at different Strouhal numbers St and (*e–i*) at different POD orders i.

discrepancy between the assimilated slave flow and master DNS occurs at the velocity peak location. This feature is particularly susceptible to being filtered out by sparse observations and cannot be adequately recovered through interpolation. While this issue also occurs in perfect-observation DA systems, its impact diminishes when using higher observation densities or smaller σ values.

3.2. *Modal synchronisation identification*

The synchronised components of the flow, which exhibit strong correlations with the sparse observations, are identified using the ESPOD method, as detailed in § 2.2. For each case, 500 snapshots are analysed. The data are divided into segments of 50 consecutive instantaneous fields for DFT computation, with a 40-field overlap between adjacent segments, resulting in a total of 46 segments for POD analysis. As an illustration, figure 7 presents the ESPOD modes in case 3 in the master system at an observation density of $\mathcal{S} = 0.25$. These modes were obtained by extending the SPOD modes of the observational data layer by layer to different wall-parallel planes. Notably, in other slave systems under the same case label and DA parameters, correlated flows are extracted by projecting onto these ESPOD modes as indicated by (2.18). Consequently, the spatial distributions of the correlated modes across different slave systems are fully consistent with those shown in figure 7. According to the ESPOD definition, the modes in figure 7 share identical temporal coefficients with the corresponding SPOD modes of the observational data as shown in (2.16). This implies that the global flow excited by the sparse observations in this study is extracted through temporal correlation, with no direct spatial-scale relationship. As seen in figure 7(*a–c*), the same observations excite distinctly different global flow structures at different wall distances, with the spatial scale decreasing as the wall distance decreases. This is consistent with the distribution characteristics of turbulent scales in

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

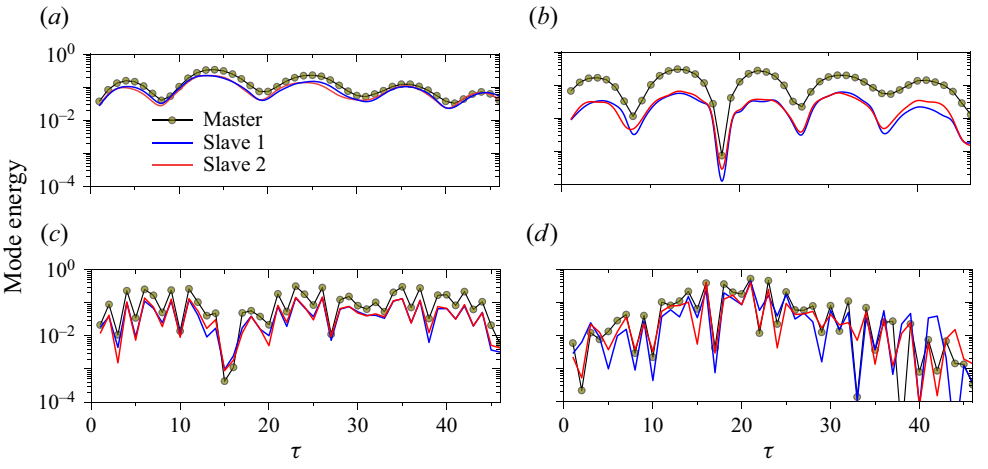


Figure 8. Energy evolution of the ith-order ESPOD modes $\hat{\boldsymbol{\psi}}$ evaluated at $y^+ = 2$ with $\mathcal{S} = 0.25$ and $\sigma = +\infty$ in case 3: (*a*) $St = 0.17$ and $i = 5$; (*b*) $St = 0.87$ and $i = 5$; (*c*) $St = 1.74$ and $i = 20$; (*d*) $St = 0.17$ and $i = 30$.

channel flows. At the same wall distance, the spatial scale of the modes decreases significantly with increasing frequency, as shown in figure 7(*d–f*). A key insight from these results is that although sparse observations contain only large-scale flow information, they can excite a variety of small-scale vortical structures during the DA process, with both sharing similar temporal evolution patterns. This phenomenon does not arise from direct spatial correlation between the observations and the global flow. Rather, it is initially excited by temporal correlation, which induces the formation of small-scale structures downstream of the observation locations, and subsequently propagates through synchronisation mechanisms to excite small-scale global modes. High-frequency sampling of observations is a crucial prerequisite, as also demonstrated by Wang *et al.* (2022). The role of POD in this context is mainly to filter the Fourier modes, helping to determine whether the excited flows in the master and slave systems exhibit temporal correlation. As shown in figure 7(*g–i*), different POD orders show minimal differences in the spatial scales of the modes but differ primarily in the level of noise. Therefore, using low-order POD modes is more appropriate for evaluating flow synchronisation characteristics across different frequencies.

The dynamic characteristics of synchronised modes can be illustrated through their kinetic energy evolution, as depicted in figure 8. The kinetic energy is normalised by the total kinetic energy at the corresponding wall distance in the master system. We examine two slave systems that differ solely in their initial conditions during the assimilation process. As demonstrated in figure 8, the energy evolution of slave-system modes closely matches that of the corresponding ESPOD modes in the master system, regardless of the Strouhal number and mode order. Notably, the slave modes exhibit remarkable insensitivity to initial conditions, as evidenced by the near-perfect overlap between the two slave-system curves, even at $St = 1.74$ and for higher-order modes (up to $i = 20$) as shown in figure 8(*a–c*). However, a gradual decline in curve similarity emerges with further increases in mode order, even at low Strouhal numbers. The observed similarity in modal energy evolution confirms successful modal synchronisation in the slave system, achieved via sparse-data-driven assimilation. This synchronisation typically develops alongside the slave system's dynamic evolution, which loses dependence on its initial state. Nevertheless, compared with the master-system reference, each synchronised

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

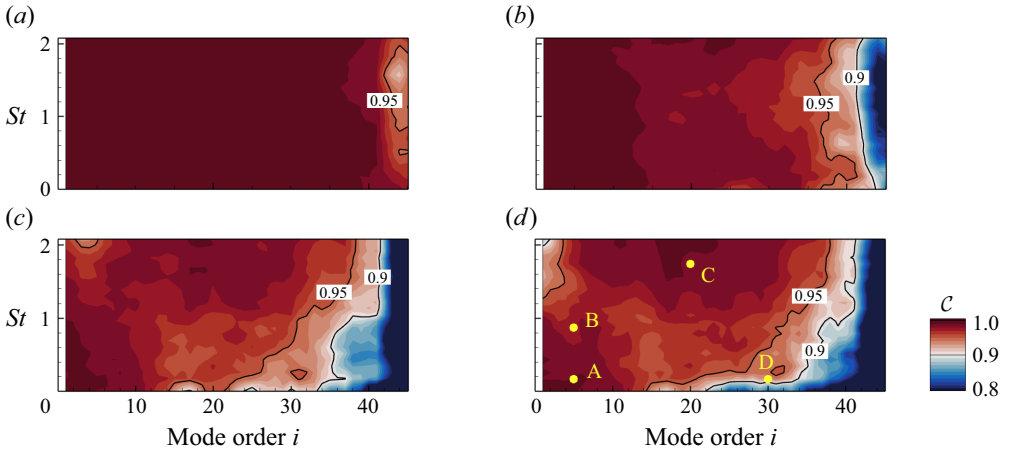


Figure 9. Cross-correlation of the mode coefficients in the slave and master systems with $\mathcal{S} = 0.25$ and $\sigma = +\infty$ in case 3: (*a*) $y^+ = 150$, (*b*) $y^+ = 60$, (*c*) $y^+ = 15$ and (*d*) $y^+ = 2$. Modal synchronisation is defined when $\mathcal{C} > 0.9$. The yellow circles A, B, C and D represent the (*St, i*) parameters extracted from Figures 8(*a*)–8(*d*), respectively.

mode exhibits reduced energy and becomes more pronounced at higher Strouhal numbers. This suggests that high-frequency, small-scale structures are more challenging to recover within the modal synchronisation framework.

The similarity in mode energy evolution can be quantified using the cross-correlation coefficient $\mathcal{C}$ of the ESPOD mode coefficients for each frequency and mode order, as defined in (2.20). This approach is justified because synchronisation involves the same modes in both the slave and master systems. Figure 9 presents the calculated values of $\mathcal{C}$ at various wall distances for $\mathcal{S} = 0.25$. Regions with high $\mathcal{C}$ values indicate strong similarity in mode energy evolution, and a threshold of $\mathcal{C} > 0.9$ is adopted to define modal synchronisation. The results reveal that modal synchronisation is more readily achieved in the channel centre than near the wall. Notably, while the Strouhal number influences mode energy, it does not play a decisive role in modal synchronisation. However, modal synchronisation breaks down at higher mode orders ($i > 40$), as these modes predominantly represent random noise and contain substantially less energy than the synchronised modes. Additionally, modal synchronisation tends to fail at moderate and high mode orders near $St = 0$, where the modes correspond to large-scale flow variations that extend across multiple data segments as captured by ESPOD using a finite-length DFT window. The data in figure 8 correspond to the (*St, i*) parameters marked by the yellow circles in figure 9(*d*), all of which satisfy the $\mathcal{C} > 0.9$ synchronisation criterion.

For the same DA case with $\mathcal{S} = 0.16$, the energy evolution and cross-correlation of mode coefficients are analysed, with focus restricted to the channel centre. Consistent with the earlier findings, modal synchronisation is achieved primarily near the channel centre ($y^+ = 150$). Although minor discrepancies in mode energy evolution are evident between the two slave systems as observed in figures 10(*a*) and 10(*b*), a substantial region with $\mathcal{C} > 0.9$ persists as shown in figure 10(*c*). The results indicate that the transition between modal-synchronisation and non-synchronisation states occurs gradually rather than abruptly. Notably, as demonstrated in figure 10(*d*), certain individual modes at $y^+ = 60$ still exhibit modal synchronisation characteristics. However, this study primarily focuses on clearly established cases of modal synchronisation, particularly those where the

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

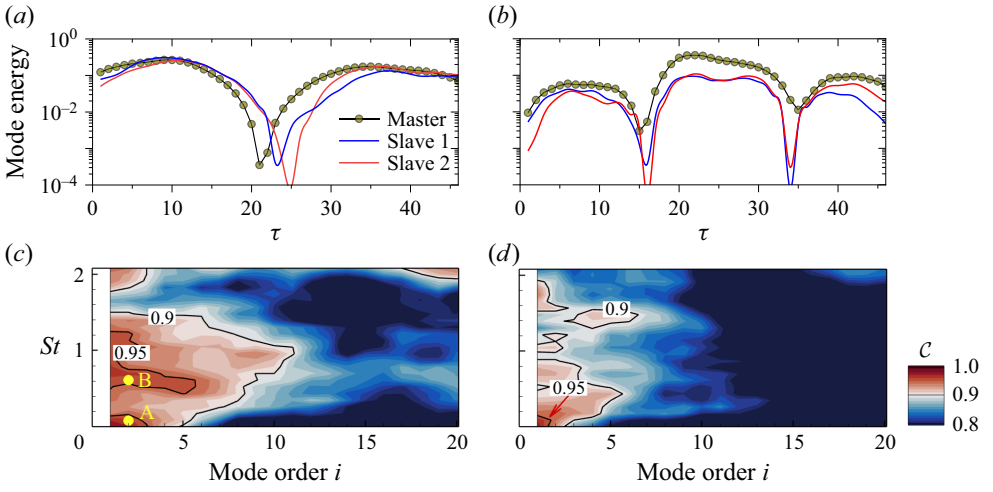


Figure 10. Modal synchronisation at $\mathcal{S} = 0.16$ and $\sigma = +\infty$ in case 3. Energy evolution of the ith-order ESPOD modes $\hat{\boldsymbol{\psi}}$ evaluated at $y^+ = 150$: (*a*) $St = 0.08$ and $i = 2$; (*b*) $St = 0.6$ and $i = 2$. Cross-correlation of the mode coefficients between the slave and master systems at (*c*) $y^+ = 150$ and (*d*) $y^+ = 60$.

leading-order ESPOD modes satisfy the 0.9 correlation criterion across most Strouhal numbers.

The failure of modal synchronisation in ESPOD modes under low observation density stems from incomplete information transfer between observations and the slave flow within the DA framework. This phenomenon primarily manifests as a velocity slip between observation points and surrounding fluid parcels. Under sparse observation conditions, the acquired information predominantly captures large-scale vortical structures, which must then be propagated to adjacent fluid through the existing nudging-like assimilation strategy. While the global Fourier mode injection method employed by Nikolaidis & Ioannou (2022) demonstrates theoretical advantages in this regard, the present study focuses on more realistic observation, mirroring practical experimental scenarios where data acquisition is limited to spatially sparse locations. A straightforward strategy for enhancing modal synchronisation involves minimising velocity slip by incorporating information from neighbouring regions of the master flow. This is achieved through finite values of the coefficient σ, which define $3 \times 3 \times 3$ assimilation cells with spatially decaying weights from the centre cell. It is noteworthy that although observations in both centre and surrounding cells are directly obtained from the master system, effectively increasing the observation density, this approach retains practical relevance for experimental applications. In such cases, surrounding values can be reasonably approximated through interpolation or duplication techniques.

Examining case 14 shown in table 1, figure 11 presents the cross-correlation maps of the mode coefficients with the observation density $\mathcal{S} = 0.14$. A LES model is employed for both the master and slave systems at $Re_\tau = 640$ to assess the influence of model imperfections on modal synchronisation. Unlike previous studies, which typically required the master and slave systems to employ identical meshes and computational settings in order to achieve turbulence synchronisation, the present work further investigates the effects of numerical discrepancies. Specifically, we use the mapped fields from a higher-resolution LES case (case 12) as the master flow, as shown in the second and fourth columns of figure 11. Initial results demonstrate that with this observation density and $\sigma = \infty$, the system fails to achieve modal synchronisation due to the insufficient

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

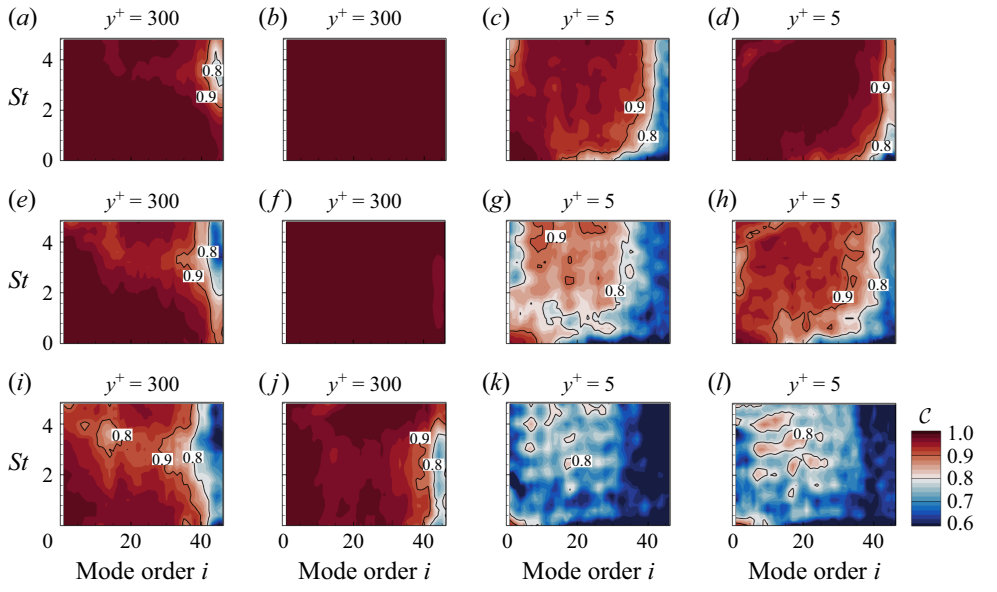


Figure 11. Cross-correlation of the mode coefficients in the slave and master systems with $\mathcal{S} = 0.14$ at $Re_\tau = 640$ in case 14: (*a–d*) $\sigma = 0$, (*e–h*) $\sigma = 1.5$ and (*i–l*) $\sigma = 2.5$. First and third columns: using master data computed from higher-resolution mesh of case 12; second and fourth columns: using master data computed from the same mesh.

observation density. From these results, it is also clearly evident that using finite values of σ can significantly enhance the degree of modal synchronisation. Reducing the σ value effectively increases the density of available observational data. However, it should be noted that this method of increasing data density does not directly reduce the scale of turbulent structures captured by the observations. Instead, it efficiently increases the assimilated volume fraction of the computational domain. As noted by Clark Di Leoni *et al.* (2020), the nudged volume fraction in physical space is an important parameter for turbulent synchronisation. Here, we interpret the influence of this parameter as the velocity slip between observation points and the surrounding fluid parcels. Another notable feature is the consistently lower correlation when using the mapped master flows from the higher-resolution mesh at all wall distances. Although the master flow mapped from the higher-resolution mesh occupies each mesh element in the slave system, minor interpolation errors are introduced during the mapping process. These errors reduce the energy proportion of synchronised modes during the ESPOD decomposition, leading to a decrease in $\mathcal{C}$ value shown here. However, as demonstrated later, this reduction in correlation does not fundamentally affect the flow synchronisation characteristics. In other words, the turbulent synchronisation behaviour remains unchanged while only the strictly correlated flow components are diminished within the current ESPOD framework.

Through comprehensive analysis of DA cases across varying Reynolds numbers, observation densities and σ values, a synchronisation map is constructed, as presented in figure 12. In this framework, full synchronisation is achieved only when the flow satisfies the synchronisation criterion on all wall-parallel planes. In contrast, modal synchronisation is identified when the criterion is met on at least one wall-parallel plane, reflecting partial yet coherent alignment of dominant flow structures. The $\mathcal{S}$–σ parameter space can be categorised into three distinct regimes representing full synchronisation, modal synchronisation and non-synchronisation states. A key finding is that the threshold

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

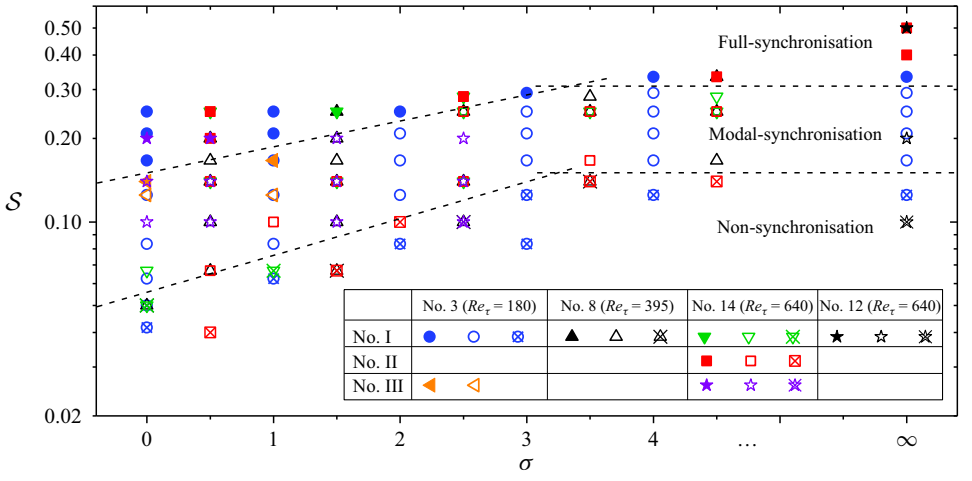


Figure 12. Synchronisation map constructed based on the cases 3, 8, 12 and 14. The filled symbols represent the full-synchronisation state, the open symbols denote the modal-synchronisation state, while the symbols with cross denote the non-synchronisation state. The black dashed lines indicate the approximate boundary between the three different regimes.

boundaries between these regimes are independent of the predictive model, Reynolds number and mesh resolution. This contrasts with prior studies, which demonstrated that full synchronisation primarily depends on the length scale of the observations and the Kolmogorov scale in the flow system. While the full-synchronisation criterion should be also applicable in the present study for DNS cases, its practical evaluation is complicated by the challenge of determining the observational length scales in these strongly anisotropic wall-bounded flows. Figure 12 reveals that observation density serves as the fundamental metric for assessing synchronisation within the current DA framework, given that the parameter σ also represents an equivalent density measure. This interpretation is intuitive because only a portion of the turbulence spectrum prevails in the master LES system. Consequently, the observation density governs the degrees of freedom in the slave flows that can be adjusted during the DA process, provided the smallest turbulence scale in the flow is finer than that which the mesh resolution can resolve. In other words, the Kolmogorov-scale or Taylor-scale criterion for full synchronisation, previously established for homogeneous or channel flows, represents the limiting case of the present density-based criterion when mesh resolution is progressively refined and the predictive model transitions to DNS. This assertion is supported by the results in figure 13, which compares the equivalent length scales resolved by sparse observations with twice the Taylor microscale, i.e. a relevant metric for assessing synchronisation in turbulent channel flows proposed by Wang & Zaki (2022). The Taylor microscale is computed using DNS data from Abe *et al.* (2001) and the anisotropic turbulence formulation by Pope (2000), while the equivalent observational scale is estimated as

$$\Lambda_{obs} \approx \frac{u_\tau}{\mathcal{S}\nu}\sqrt{\left(\frac{L_x}{N_x}\right)^2+\left(\frac{L_y}{N_y}\right)^2+\left(\frac{L_z}{N_z}\right)^2}=\frac{u_\tau}{\mathcal{S}\nu}L_{mesh}, \tag{3.2}$$

where u_τ represents the wall friction velocity and L_{mesh} defines the mesh scale. A full-synchronisation density of $\mathcal{S}=0.33$ is selected, yielding equivalent length scales $\Lambda_{obs}=$ 41.2, 63.8 and 103.4 for $Re_\tau=180$ (case 3), 395 (case 8) and 640 (case 14), respectively. Case 3 approximately satisfies the Taylor-scale criterion proposed by Wang & Zaki (2022),

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

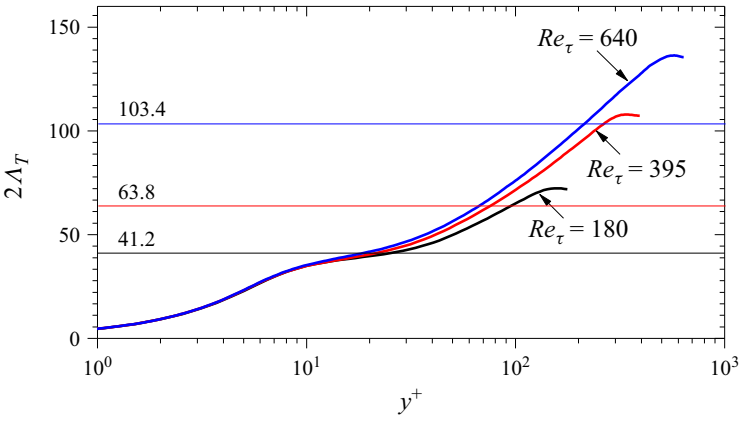


Figure 13. Taylor scales of the turbulence at different Reynolds numbers and wall distances. The horizontal lines represent the equivalent length scales captured at $\mathcal{S} = 0.33$ in cases 3, 8 and 14 for each Reynolds number.

while sparse observations are sufficient to synchronise all scales in the inertial sublayer. In contrast, the LES cases 8 and 14 resolve only a subset of the scales in the inertial sublayer and do not meet the Taylor-scale criterion. However, they still fall within the full-synchronisation regime in the present study. Notably, full-scale alignment in the full-synchronisation regime cannot be achieved with observation cases II and III, because the interpolation errors in the observations cannot be fully compensated by any momentum forcing within the current DA framework. A key reason for this limitation is that the slave flow inherently satisfies the continuity condition, whereas the mapped observations do not. Nevertheless, the final converged state of the slave system under this regime is independent of the initial conditions; that is, slave systems initialised differently ultimately align with each other across all scales, as also reported by Li, Tian & Li (2022). Another notable phenomenon is that when observation configuration III is employed, the flows tend to achieve full synchronisation more readily at smaller values of σ. This may be attributed to the fact that such down-sampling in physical space constrains the slave system into a state devoid of fine-scale vortex structures. As a result, the flow field possesses fewer degrees of freedom, making it more susceptible to full synchronisation. However, this tendency gradually weakens as σ increases.

To further examine the dependence of flow synchronisation on turbulent scales and mesh resolution, we analysed the synchronisation characteristics under different mesh scales for each Reynolds number, i.e. for all cases listed in table 1. To ensure consistent comparisons, all cases were configured with $\sigma = 0$ and with observation configuration I. The relative mesh scale, L_{mesh}, is defined in (3.2), and the corresponding computational cases with varying grid scales are provided in table 1. First, considering the DNS cases at $Re_\tau = 180$, the mesh employed in case 3 meets the Taylor-scale criterion established by Wang & Zaki (2022), meaning it approaches the resolution threshold to resolve all energy-containing flow structures and thereby enables the full-synchronisation regime. The corresponding mesh scale is $L_{mesh} = 0.078$. However, when the mesh resolution is further increased, the same observation density is able to capture smaller turbulent scales (i.e. Λ_{obs}), even though no smaller energy-containing turbulent structures are present in the flow. This implies that at higher mesh resolutions, a lower observation density is sufficient to attain the full-synchronisation regime. From figures 13 and 14, it can be estimated that $L_{mesh} = 0.078$ and $\mathcal{S} = 0.14$ are approximately sufficient to achieve full synchronisation. If the flow is

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

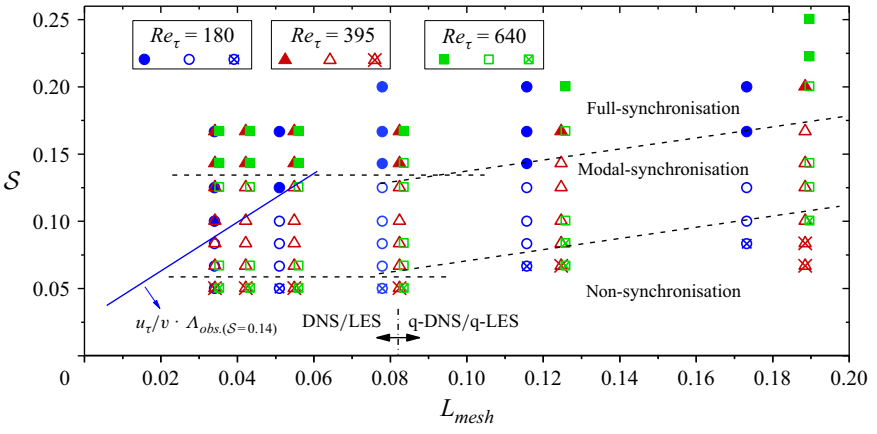


Figure 14. Synchronisation map using different meth resolutions with $\sigma = 0$. The filled symbols represent the full-synchronisation state, the open symbols denote the modal-synchronisation state, while the symbols with cross denote the non-synchronisation state. The black dashed lines indicate the approximate boundary between the three different regimes. The blue solid line represents the slope of (3.3). The DNS/LES and q-DNS/q-LES are separated by $L_{mesh} \approx 0.082$.

to remain fully synchronised at increased mesh resolutions, Λ_{obs} should remain constant. Therefore, according to (3.2), it follows that

$$\frac{\mathcal{S}}{L_{mesh}} \approx \frac{u_\tau}{\upsilon \cdot \Lambda_{obs(\mathcal{S}=0.14)}}. \tag{3.3}$$

This relationship indicates that a straight line with this slope separates the regimes of full synchronisation and modal synchronisation. The results for $Re_\tau = 180$ in figure 14 corroborate this observation: in cases 1 and 2, the minimum observation density required for the flow to achieve full synchronisation lies just above this sloped line, whereas below it, the system transitions into the modal synchronisation regime. It is noteworthy that, according to the analysis in § 3.2, achieving full synchronisation requires reaching a certain assimilated volume fraction. Consequently, at higher grid resolutions, the reduced observation density leads to a lower assimilated volume fraction, which shifts the threshold for full synchronisation away from the reference line. A similar trend is observed in figure 14. In this case, when the mesh resolution is coarse in the range $L_{mesh} < 0.09$, the threshold observation density for the three regimes remains relatively constant. This occurs because, for low-resolution grids, the DNS effectively behaves as an implicit LES, with the computational grid acting as an implicit filter. In such a scenario, the smallest energy-containing vortices in the turbulent system are no longer characterised by the Taylor scale, but rather by the mesh filter scale. At a fixed observation density, the ratio of the observational length scale to the scale of the smallest energy-containing vortices remains constant. It is this scale ratio that governs the synchronisation behaviour of the flow. This interpretation is further supported by the cases at $Re_\tau = 395$ and 640, where the threshold observation density similarly remains nearly constant for $L_{mesh} < 0.09$. However, when mesh resolution is increased beyond $L_{mesh} < 0.05$, the slope change observed in the DNS case does not appear. This is because the LES framework inherently includes smaller-scale turbulent structures; increasing mesh resolution must therefore be accompanied by a corresponding reduction in the turbulence scale resolved by the mesh, so as to preserve

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

the constancy of the ratio between the observational length scale and the smallest energy-containing turbulence scale in the system. According to the statistical data presented in figure 14, the boundaries of the three synchronisation regimes across different Reynolds numbers consistently lie near $S = 0.14$ and 0.06. However, when the mesh resolution becomes very coarse (i.e. $L_{mesh} > 0.09$), these boundaries begin to deviate from the constant values, indicating that a higher observational density is required to achieve either full or modal synchronisation. At such mesh resolutions, turbulent statistics obtained via DNS or LES also exhibit significant deviations from the accurate values.

These findings indicate that, beyond the previously established regime of full turbulence synchronisation, a state of modal synchronisation also exists in data-assimilated turbulent systems under sparse observations. In both DNS and LES frameworks, the criterion for achieving full synchronisation depends on the ratio between the observational scale and the smallest energy-containing scale in the slave system. When the computational mesh satisfies DNS requirements, the slave system in physical space achieves full synchronisation with the criterion for the observation data coinciding with the Taylor-scale criterion proposed by Wang & Zaki (2022). In contrast, when mesh resolution only meets LES standards, the smallest turbulent scale in the slave system becomes the mesh filter scale, and the synchronisation criterion shifts to observation density, i.e. effectively the ratio of the mesh filter scale to the observational scale. The emergence of modal synchronisation across all systems is predominantly governed by this observation density. Moreover, when mesh resolution falls below LES requirements, the synchronisation behaviour of the flow field deviates from the above criteria.

3.3. *Delay-coordinate assimilation*

Based on the preceding analysis, incorporating historical information from limited observations emerges as one of the most effective strategies for mitigating discrepancies between slave and master flows under low observation densities. While positioning observations at velocity peak locations would maximise assimilation accuracy, this approach is inherently impractical due to the absence of *a priori* knowledge of the instantaneous flow field. In practical applications such as LPT measurements, observation sites are typically randomly distributed across the flow domain. In such scenarios, leveraging the historical variation of velocity offers a viable alternative: through flow convection, it implicitly integrates information from adjacent regions and past velocity peaks that may not be captured at the current time step. This approach thereby enhances assimilation performance by spatially and temporally extending the influence of sparse measurements. The present study employs a delay-coordinate strategy, which incurs negligible additional computational cost compared with conventional DA methods. Unlike nudging techniques that directly introduce feedback terms into the momentum equations, the proposed SeVar framework requires modifications to the adjoint momentum equations (see Appendix A for details). Here, observations are considered at the current time step (κ) as well as the two preceding steps ($\kappa - 1, \kappa - 2$):

$$\frac{\boldsymbol{V}}{\delta t} + \nabla \boldsymbol{U} \cdot \boldsymbol{V} - (\boldsymbol{U} \cdot \nabla)\, \boldsymbol{V} - \nabla \cdot \left[\nu_{eff}\left(\nabla \boldsymbol{V} + (\nabla \boldsymbol{V})^T\right)\right] + \frac{1}{\rho}\nabla q + \frac{2\xi M}{U_0^2}\left[\left(\boldsymbol{U}'' - \boldsymbol{U}''_{obs}\right) + \left(\boldsymbol{U}' - \boldsymbol{U}'_{obs}\right) + \left(\boldsymbol{U} - \boldsymbol{U}_{obs}\right)\right] = 0, \tag{3.4}$$

where $\boldsymbol{U}'$ and $\boldsymbol{U}'_{obs}$ denote the assimilated flow and the observations, respectively, at the previous observational time $t_{obs}^{(\kappa-1)}$, while $\boldsymbol{U}''$ and $\boldsymbol{U}''_{obs}$ correspond to those at

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

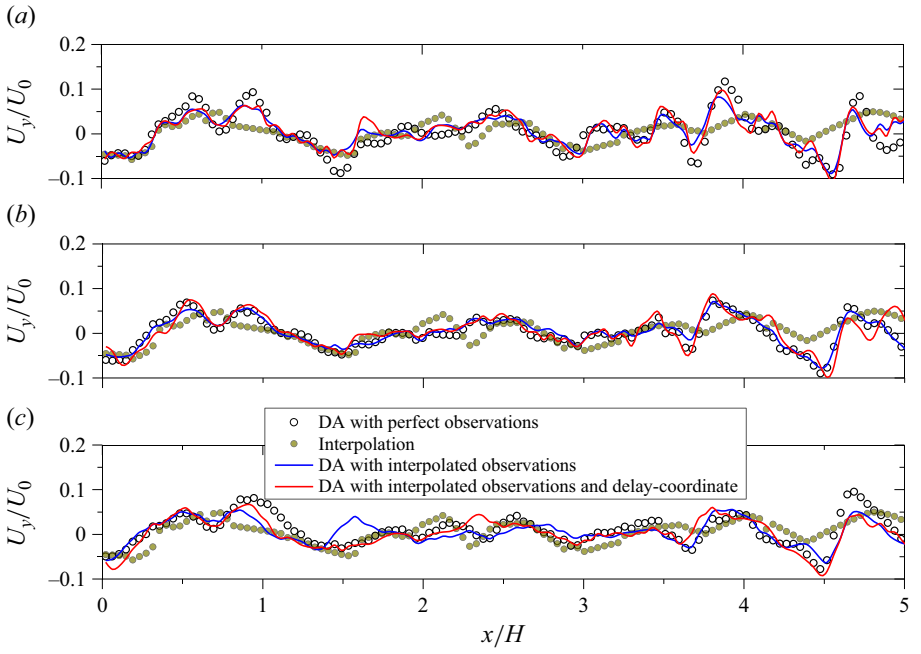


Figure 15. Streamwise variations of the wall-normal velocity component at $y^+ = 40$ in case 14 using observation III and $\mathcal{S} = 0.14$. The DA is performed with (*a*) $\sigma = 0.5$, (*b*) $\sigma = 1.5$ and (*c*) $\sigma = 2.5$. The dark yellow circles denote the direct interpolation, duplicated in each subfigure for comparison.

$t_{obs}^{(\kappa-2)}$. Notably, the delay-coordinate assimilation approach does not necessitate additional data or supplementary observation sites; instead, it leverages the same observational resources employed in conventional DA strategies. Furthermore, given the minimal objective residuals, the delay-coordinate strategy is only meaningful when $0 < \sigma < \infty$ in the surrounding cells of the assimilating cube. This approach integrates historical data into the current assimilation time, rather than considering a long evolution trajectory as done in 4DVar. Following the optimisation proposed by Pazo, Carrassi & Lopez (2016), only three time instants are incorporated, as excessive historical data have been shown to diminish the improvement in DA results.

All the following analysis is performed with case 14 using observation configuration III (table 2) at $Re_\tau = 640$. This case serves as a representative example of weak-constraint DA, where the master flow cannot be fully reproduced by any predictive model, being significantly different from those examined in previous studies (Lalescu *et al.* 2013; Wang & Zaki 2022; Li *et al.* 2024). The configuration also serves as a benchmark test, simulating real-world applications where measurement data are employed as observations. Figure 15 presents the streamwise variation of wall-normal velocities at a selected point for DA cases with and without the delay-coordinate strategy. The observation density is fixed at $\mathcal{S} = 0.14$, which lies within the modal-synchronisation regime for $\sigma = 0.5$, 1.5 and 2.5. Since the direct interpolation from the down-sampled observations fails to capture fundamental flow characteristics, we focus on DA approaches that demonstrate improved global feature recovery capabilities. Each subfigure compares DA results using perfect observations (configuration I) with interpolated observations (configuration III) at corresponding σ values. When σ approaches infinity, both perfect and interpolated observations produce identical DA results, as the assimilation converges to a single-cell scenario where the observations become equivalent. However, the agreement between

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

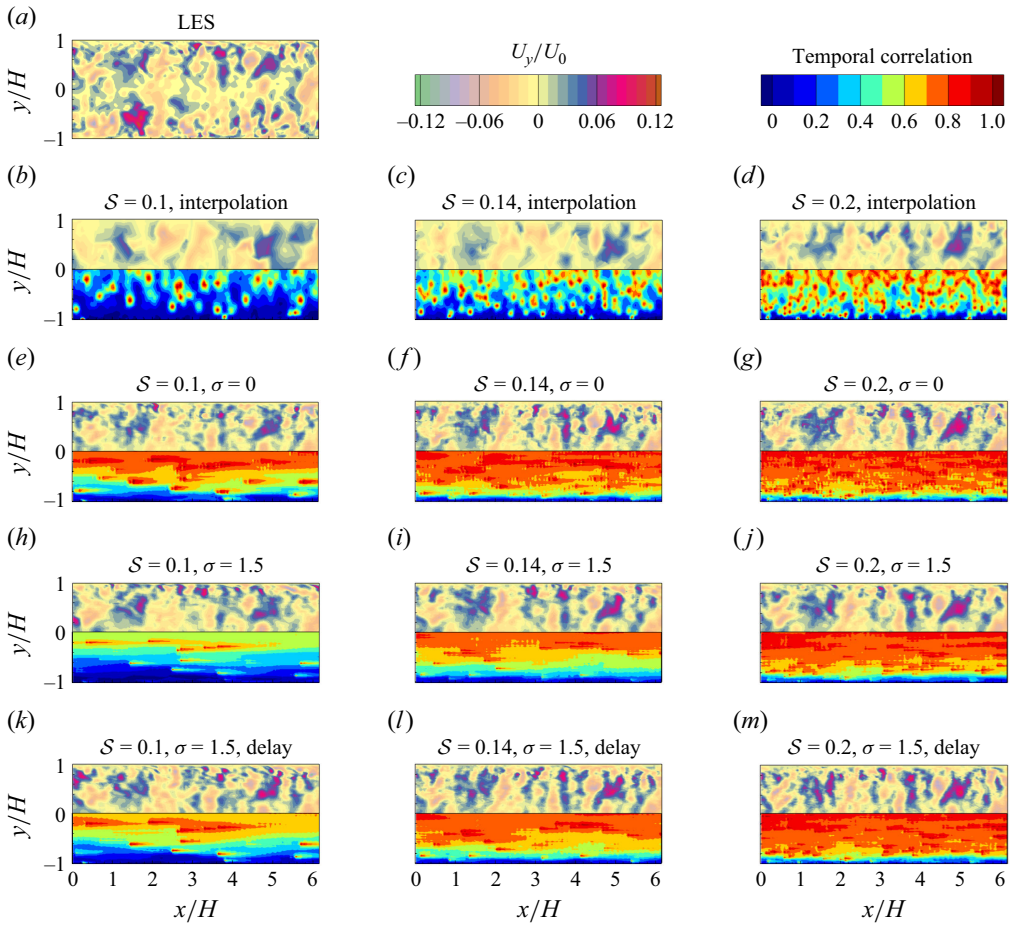


Figure 16. Velocity snapshots and temporal correlations of the interpolated and assimilated slave flows in case 14 using observation III. Parameter values are indicated above each subfigure for clarity.

these two cases does not monotonically improve as σ increases from 0.5 to 2.5. Notable deviations between the DA results with interpolated observations (blue curve) and those with perfect observations (open circles) are observed at $\sigma = 0.5$ (figure 15*a*) and $\sigma = 2.5$ (figure 15*c*), while optimal agreement occurs at $\sigma = 1.5$ (figure 15*b*). This behaviour suggests the existence of an optimal σ that minimises DA errors induced by observation interpolation. Although the delay-coordinate strategy significantly enhances velocity peak recovery, it also generates more spurious peaks as σ decreases.

To more comprehensively evaluate the enhanced performance of the proposed DA framework incorporating interpolated assimilation and delay-coordinate strategies, figure 16 presents both instantaneous snapshots and temporal correlations of the assimilated fields under varying observation densities and σ values. The temporal correlation is calculated as the cross-correlation between the target flow field and the LES reference data using a sliding-window approach with a window size of 50 instantaneous fields and 50 % overlap. Figure 16(*a*) displays the LES reference snapshot for comparison. At each observation density, sparse observations are interpolated onto the global mesh to generate the master interpolated flows, as shown in figure 16(*b–d*). The interpolation process primarily captures large-scale flow patterns (evident in the upper halves of each subfigure),

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

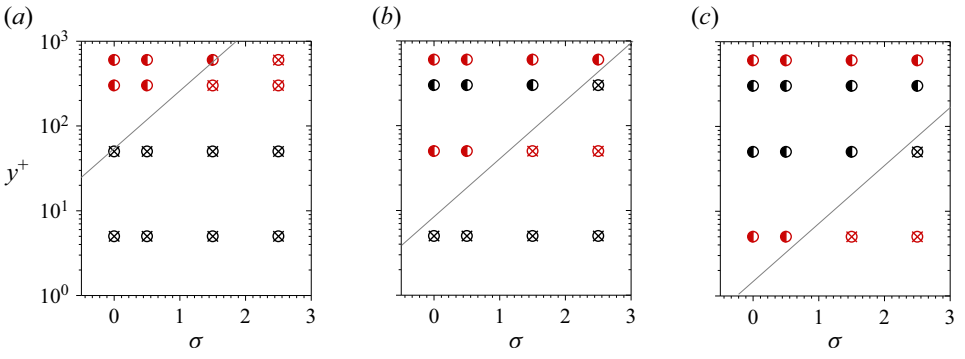


Figure 17. Synchronisation maps at various densities and wall distances in case 14 using observation III: (*a*) $\mathcal{S} = 0.1$, (*b*) $\mathcal{S} = 0.14$ and (*c*) $\mathcal{S} = 0.2$. Half-filled circles indicate full or modal synchronisation; circles with cross indicate non-synchronisation. Red circles highlight the parameter sets used for the energy comparison.

with temporal correlations remaining near unity at observation sites but decreasing sharply elsewhere. The DA substantially improves the flow fields by reconstructing vortical structures and enhancing temporal correlations. Notably, high correlation values are observed in slender regions downstream of observation sites. As demonstrated in figure 16(*e–g*), DA with $\sigma = 0$ yields the highest temporal correlation across all densities, though with noticeable noise (to be discussed later). Increasing σ to 1.5 adversely affects DA performance as shown in figure 16(*h–j*), but this effect diminishes rapidly with increasing observation density. Particular attention should be given to the delay-coordinate assimilation cases in figure 16(*k–m*), which show marked improvement at $\mathcal{S} = 0.1$ but only marginal noise reduction at $\mathcal{S} = 0.14$ and 0.2. Importantly, while $\mathcal{S} \geq 0.14$ with $\sigma = 1.5$ lies within the modal-synchronisation regime (figure 12), $\mathcal{S} = 0.1$ with $\sigma = 1.5$ approaches non-synchronised conditions. This suggests a preliminary conclusion that when the slave system falls in the modal-synchronisation regime, the delay-coordinate interpolated assimilation strategy with $\sigma > 0$ can effectively reconstruct global flow dynamics from sparse observations. However, $\sigma = 0$ performs inferiorly to $\sigma = \infty$ at high observation densities due to introduced interpolation errors.

To further substantiate the preceding argument, synchronisation maps derived from interpolated assimilation data are re-evaluated on each wall-parallel plane as illustrated in figure 17. This figure does not differentiate between full synchronisation and modal synchronisation across all planes, since even in the full-synchronisation regime the slave flow cannot perfectly match the master flow. Rather, it illustrates the second scenario of full synchronisation, where multiple slave systems converge to mutually identical states. This possibility warrants further investigation through additional computational examples. The comparison of its overall synchronisation map with other cases has been shown in case 14 and observation III in figure 12. It can be observed that the synchronisation characteristics under this observation configuration are generally consistent with other cases. Figure 18 presents the streamwise wave spectra computed from the streamwise velocity component at $y^+ = 50$ and 600. Since the master LES data are interpolated from a finer mesh, the spectra display an upward trend at high wavenumbers due to interpolation artefacts. In contrast, interpolation from sparse observations leads to severe spectral degradation, attributable to the filtering of numerous vortical structures. This manifests as a reduction in power spectral density across most wavenumbers, except at the high-wavenumber range where interpolation errors dominate. Notably, all spectral curves of the assimilated flow exhibit anomalous growth at high wavenumbers, consistent with

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

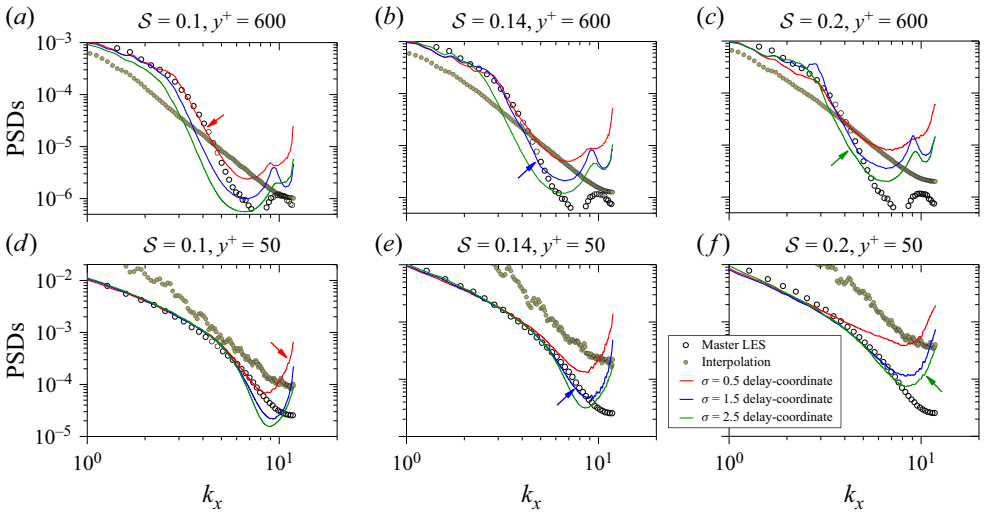


Figure 18. (*a–f*) Streamwise wave spectra computed from the streamwise velocity component in case 14 using observation III. Arrows indicate the optimal σ value in each subfigure. Parameter values are listed above each subfigure. PSD, power spectral density.

the nudging results reported by Clark Di Leoni *et al.* (2020). Our analysis focuses on the spectral range below $k_x < 8$, where the current DA framework operates effectively. As the σ value increases, each spectral curve demonstrates attenuated power spectral density at relatively high wavenumbers, with this attenuation being weakened progressively at lower wavenumbers. This behaviour clearly indicates that relaxed weighting for surrounding cells primarily influences small-scale structures. This is reasonable since large-scale vortices are already captured by the sparse observations. The key conclusion drawn from figure 18 is that the optimal σ for spectral reproduction within the present interpolated delay-coordinate DA framework lies near the modal synchronisation/non-synchronisation threshold, as demarcated by the oblique lines in figure 17. However, the precise positioning of σ within this transitional regime should not be overinterpreted, as the shift from modal synchronisation to non-synchronisation constitutes a gradual rather than abrupt transition. The above argument can be explained through the trade-off between error suppression effects and modal synchronisation strength. Given that modal synchronisation depends solely on equivalent density, it follows that neither delay-coordinate nor weighting-relaxation strategies prove effective for flow systems firmly within the non-synchronisation regime, as these approaches do not increase the number of observation sites involved in the DA process.

The underlying mechanism for these flow reconstruction improvements can be elucidated through examining energy modifications in the synchronised modes. As representative cases, we analyse the densities and wall-parallel planes, as indicated by filled and open red circles in figure 17, respectively. Figure 19 displays the corresponding energy ratios $\varkappa$ relative to the reference LES data as defined in (2.21). Only the leading modes for each frequency are shown, as the remaining synchronised modes exhibit similar trends. The cases with $\sigma = 0$ are selected as the benchmark for comparison since they exhibit the highest temporal correlation (figure 16), despite introducing undesirable interpolation errors across all observation densities. Notably, the delay-coordinate approach demonstrates limited effectiveness when $\sigma = 0$, as the contribution of historical data to the adjoint source in (3.4) becomes negligible, being distinct from

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

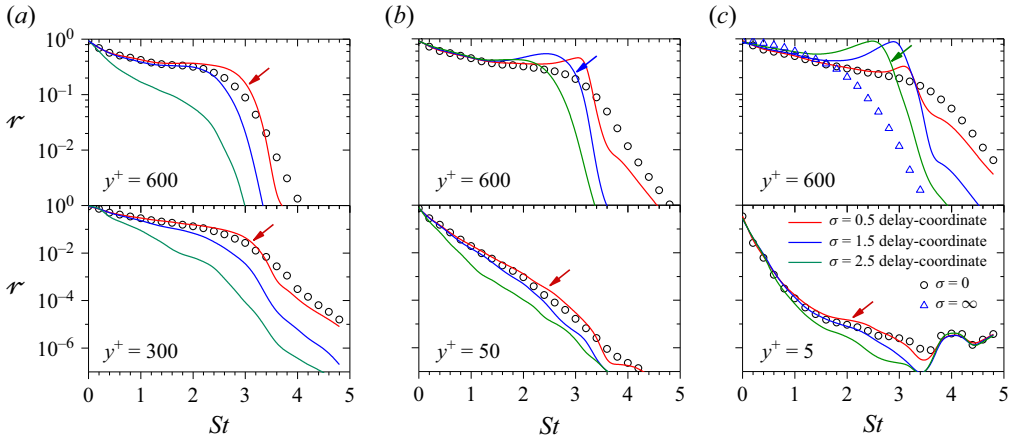


Figure 19. Energy ratios of the synchronised modes at different wall distances in case 14 using observation III, as indicated by red in figure 17: (*a*) $\mathcal{S} = 0.1$, (*b*) $\mathcal{S} = 0.14$ and (*c*) $\mathcal{S} = 0.2$. Arrows indicate the optimal σ value in each subfigure.

the nudging mechanism (Pazo *et al.* 2016). The assimilation results show rapid decay with increasing frequency, indicating that the current DA framework cannot adequately reproduce the energy levels of small-scale structures, even when modal synchronisation is achieved. Importantly, the observed high-frequency decay in correlation coefficient r does not correspond to a similar power spectral density reduction in the slave flows (figure 18), as interpolation errors introduce substantial flow components that remain uncorrelated with synchronised modes. For clarity, datasets showing faster attenuation with increasing σ without delay coordinate are omitted from figure 19 to avoid excessive curve clustering. The analysis focuses primarily on the delay-coordinate DA results, which demonstrate significantly improved correlation coefficients r in the Strouhal number range $St = 1.5$–3.5 across different σ values compared with the $\sigma = 0$ baseline. Remarkably, the optimal σ values (indicated by arrows) show excellent agreement with those identified in figure 18 at $y^+ = 600$, and consistently align near the modal synchronisation/non-synchronisation threshold at all wall distances. This finding strongly supports our earlier argument regarding flow reconstruction improvement within the modal-synchronisation regime. Based on the Taylor's frozen hypothesis, we estimate that flow structures with $St \approx 3$ correspond to a streamwise wavenumber $k_x \approx 3$ in the current flow configuration. Given that delay-coordinate assimilations show minimal discrepancies in the $k_x = 1$–3 range in figure 18, the observed r increase suggests kinetic energy transfer from uncorrelated components to synchronised modes, while maintaining nearly constant total energy in this frequency band. Although the weighting relaxation strategy accelerates r decay at $St > 3.5$, particularly at $y^+ = 600$, the DA process paradoxically benefits from this uncorrelated component attenuation. This effect ultimately yields optimal wave spectra in the slave flows when using the same σ values.

4. Conclusions

This study employs the SeVar technique to reconstruct global velocity fields in turbulent channel flow using sparse, randomly distributed observations within the fluid domain. Unlike 4DVar, which optimises over a long temporal window, SeVar operates within a localised temporal subdomain, focusing on individual time steps where observations are available. This sequential approach eliminates the need for backward integration

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

of the adjoint system, significantly reducing computational costs while preserving the advantages of weighted relaxation. In contrast to prior studies on full synchronisation in homogeneous turbulence relying on perfect predictive models, this work investigates partial synchronisation in wall-bounded turbulence under conditions of significantly sparser observations. The analysis addresses inherent uncertainties and identifies the specific flow components directly synchronised by the observations. The DA framework consists of a master system that provides sparse velocity observations from specified mesh locations, and a slave system whose global flow evolution strictly maintains velocity consistency with the master system at all observation sites.

In the first phase of this study, a comprehensive statistical analysis of turbulence synchronisation across various Reynolds numbers, mesh resolutions and observation sparsity parameters is presented. Regardless of whether the observations are derived from the same numerical case or obtained via mesh mapping or down-sampling interpolation, this study successfully identifies three distinct synchronisation regimes in turbulent flows: full synchronisation, modal synchronisation and non-synchronisation. The results indicate that the criterion for achieving full synchronisation depends on the ratio between the observational scale and the smallest energy-containing scale in the slave system. When the computational mesh satisfies DNS requirements, the slave system in physical space achieves full synchronisation with the criterion for the observation data coinciding with the Taylor-scale criterion. In contrast, when mesh resolution only meets LES standards, the smallest turbulent scale in the slave system becomes the mesh filter scale, and the synchronisation criterion shifts to observation density, i.e. effectively the ratio of the mesh filter scale to the observational scale. The emergence of modal synchronisation across all systems is predominantly governed by this observation density. Moreover, when mesh resolution falls below LES requirements, the synchronisation behaviour of the flow field deviates from the above criteria.

The second part of this study concentrates on master flows interpolated from sparse observations to simulate real experimental conditions. A delay-coordinate approach is implemented to incorporate historical observation data into the current time step, thereby refining synchronisation performance. Notably, the optimal weighting relaxation coefficient σ for each sparsity level is identified near the critical threshold between synchronised and non-synchronised states, ensuring synchronisation while permitting error correction through the predictive model. This finding further demonstrates that the previously observed phenomenon of modal synchronisation serves as a critical prerequisite for leveraging historical information in DA to improve the accuracy of turbulence reconstruction. In the non-synchronisation regime, however, the contribution of historical data to enhancing the reconstruction accuracy of instantaneous fields becomes almost negligible. The delay-coordinate DA framework not only enhances the statistical dynamics of vortical structures in slave flows, but also facilitates kinetic energy transfer from uncorrelated flow components to synchronised modes.

This study advances a theoretical foundation for instantaneous turbulence reconstruction, with direct applicability to tomo-PIV and 3D-LPT measurements. By extending established full-synchronisation theory, it offers novel insights into flow reconstruction under sparse observational constraints. Additionally, the findings, particularly those concerning historical data integration, deliver actionable principles for optimising 4DVar DA frameworks in future implementations.

Funding. This work was supported by the National Natural Science Foundation of China (Grant Nos. 12272231 and 12227803).

Declaration of interests. The authors report no conflict of interest.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

Appendix A. Derivation of the SeVar system

The objective of the present DA is to determine the spatiotemporal variation of the forcing $\boldsymbol{F}$ in (2.1) to minimise the cost function $\mathcal{J}$ subject to the governing equations:

$$\text{minimise } \mathcal{J}=\xi\int_{t_0}^{T}\int_{\Omega} M\left(\frac{\boldsymbol{U}-\boldsymbol{U}_{obs}}{U_0}\right)^2 \mathrm{d}\Omega\mathrm{d}t, \quad \text{subject to } \mathcal{R}\left(\boldsymbol{U},\, p,\, \boldsymbol{F}\right)=0, \tag{A1}$$

in which $[t_0,\, T]$ is the temporal window where the optimisation is performed and $\mathcal{R}=[\mathcal{R}_{Mom},\, \mathcal{R}_C]$ denotes the momentum and continuity equations as shown in (2.1) and (2.2). This constrained optimisation problem can be converted to the unconstrained problem by introducing the Lagrange function $\mathcal{L}$ to be minimised instead of (A1):

$$\mathcal{L}=\mathcal{J}+\int_{t_0}^{T}\int_{\Omega}\left(\boldsymbol{V},\, q\right)\mathcal{R}\,\mathrm{d}\Omega\mathrm{d}t, \tag{A2}$$

where $\boldsymbol{V}$ and q are introduced as the adjoint velocity and adjoint pressure, respectively. The total variation of $\mathcal{L}$ with respect to the state variables is calculated as

$$\delta\mathcal{L}=\delta_{\boldsymbol{U}}\mathcal{L}+\delta_p\mathcal{L}+\delta_{\boldsymbol{F}}\mathcal{L}. \tag{A3}$$

This can be simplified by choosing the appropriate $\boldsymbol{V}$ and q to deplete the variation with respect to the state variables:

$$\delta_{\boldsymbol{U}}\mathcal{L}+\delta_p\mathcal{L}=0. \tag{A4}$$

Substituting equations (A1) and (A2) into (A4) yields

$$\delta_{\boldsymbol{U}}\mathcal{J}+\delta_p\mathcal{J}+\int_{t_0}^{T}\int_{\Omega}\left(\boldsymbol{V},\, q\right)\delta_{\boldsymbol{U}}\mathcal{R}\mathrm{d}\Omega\mathrm{d}t+\int_{t_0}^{T}\int_{\Omega}\left(\boldsymbol{V},\, q\right)\delta_p\mathcal{R}\,\mathrm{d}\Omega\mathrm{d}t=0, \tag{A5}$$

with

$$\delta_{\boldsymbol{U}}\mathcal{J}=2\xi\int_{t_0}^{T}\int_{\Omega} M\frac{\boldsymbol{U}-\boldsymbol{U}_{obs}}{U_0^2}\delta\boldsymbol{U}\,\mathrm{d}\Omega\mathrm{d}t, \tag{A6}$$

$$\delta_p\mathcal{J}=0, \tag{A7}$$

$$\delta_{\boldsymbol{U}}\mathcal{R}_{Mom}=\frac{\partial\delta\boldsymbol{U}}{\partial t}+\left(\delta\boldsymbol{U}\cdot\nabla\right)\boldsymbol{U}+\left(\boldsymbol{U}\cdot\nabla\right)\delta\boldsymbol{U}-\nabla\cdot\left[\nu_{eff}\left(\nabla\delta\boldsymbol{U}+\left(\nabla\delta\boldsymbol{U}\right)^T\right)\right], \tag{A8}$$

$$\delta_{\boldsymbol{U}}\mathcal{R}_C=-\nabla\cdot\delta\boldsymbol{U}, \tag{A9}$$

$$\delta_p\mathcal{R}_{Mom}=\frac{1}{\rho}\nabla\delta p, \tag{A10}$$

$$\delta_p\mathcal{R}_C=0. \tag{A11}$$

Equation (A5) must hold for any value of $\delta\boldsymbol{U}$ and δp, from which the adjoint equations and the boundary and terminal conditions can be derived. The adjoint equations reads

$$-\frac{\partial\boldsymbol{V}}{\partial t}+\nabla\boldsymbol{U}\cdot\boldsymbol{V}-\left(\boldsymbol{U}\cdot\nabla\right)\boldsymbol{V}-\nabla\cdot\left[\nu_{eff}\left(\nabla\boldsymbol{V}+\left(\nabla\boldsymbol{V}\right)^T\right)\right]+\frac{1}{\rho}\nabla q+2\xi M\frac{\boldsymbol{U}-\boldsymbol{U}_{obs}}{U_0^2}=0, \tag{A12}$$

$$\nabla\cdot\boldsymbol{V}=0, \tag{A13}$$

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

with the boundary conditions $\boldsymbol{V}=0$ and $\boldsymbol{n}\cdot\nabla q=0$ on the wall where $\boldsymbol{n}$ is the unit normal vector, while a periodic condition is imposed for the other surfaces. In deriving the adjoint equations, the variation of the turbulent eddy viscosity with respect to the velocity is neglected. In practice, as long as the iterative coupling of the primal and adjoint equations converges, this simplification does not affect the final outcome. The terminal condition reads

$$\boldsymbol{V}\,(T)=0. \tag{A14}$$

Due to the terminal condition specified above, the adjoint equation must be integrated backward in time. As a result, the conventional 4DVar algorithm requires storing all primary and adjoint variables at every time step, leading to substantial data storage requirements and time-intensive read/write operations. The proposed SeVar reduces the temporal window to only one time step δt in which the optimisation is conduced from the instant when the observation exits, i.e. $[t_{obs},\, t_{obs}+\delta t]$. This yields a sequential DA scheme where the optimisation occurs exclusively at the observation instants, while maintaining conventional computation with $\boldsymbol{F}=0$ while the observation is not available. By employing the terminal condition $\boldsymbol{V}(t_{obs}+\delta t)=0$ and the first-order Euler approach for the time derivative, the adjoint equation is reduced as

$$\frac{\boldsymbol{V}}{\delta t}+\nabla\boldsymbol{U}\cdot\boldsymbol{V}-(\boldsymbol{U}\cdot\nabla)\,\boldsymbol{V}-\nabla\cdot\left[\nu_{eff}\left(\nabla\boldsymbol{V}+(\nabla\boldsymbol{V})^{T}\right)\right]+\frac{1}{\rho}\nabla q+2\xi M\frac{\boldsymbol{U}-\boldsymbol{U}_{obs}}{U_0^2}=0, \tag{A15}$$

$$\nabla\cdot\boldsymbol{V}=0. \tag{A16}$$

The sensitivities can thus be obtained using the variation of $\mathcal{L}$ with respect to $\boldsymbol{F}$:

$$\delta\mathcal{L}=\delta_{\boldsymbol{F}}\mathcal{J}+\int_{t_{obs}}^{t_{obs}+\delta t}\int_{\Omega}(\boldsymbol{V},\,q)\,\delta_{\boldsymbol{F}}\mathcal{R}\,\mathrm{d}\Omega\,\mathrm{d}t=-\int_{t_{obs}}^{t_{obs}+\delta t}\int_{\Omega}\boldsymbol{V}\,\mathrm{d}\Omega\,\mathrm{d}t. \tag{A17}$$

This yields the updating of $\boldsymbol{F}$ using the deepest descending algorithm shown in (2.7).

The delay-coordinate assimilation strategy formulated in (3.4) builds upon the methodological framework originally proposed by Pazo *et al.* (2016). This approach not only incorporates current prediction–observation discrepancies but also integrates a finite historical sequence of past observations. In contrast to conventional nudging techniques that directly incorporate discrepancy terms into the governing equations, the present SeVar system embeds these discrepancies within the adjoint equations to compute optimal forcing terms. Mathematically, the discrepancy term in (3.4) can be interpreted as equivalent to the definition of a cost function:

$$\mathcal{J}=\frac{\xi}{U_0^2}\int_{t_{obs}}^{t_{obs}+\delta t}\int_{\Omega}M\left[\alpha\,(\boldsymbol{U}-\boldsymbol{U}_{obs})+(\boldsymbol{U}-\boldsymbol{U}_{obs})^{\mathbf{2}}\right]\mathrm{d}\Omega\,\mathrm{d}t, \tag{A18}$$

where the slope

$$\alpha=2\left(\boldsymbol{U}''-\boldsymbol{U}''_{obs}+\boldsymbol{U}'-\boldsymbol{U}'_{obs}\right) \tag{A19}$$

accounts for the contribution of the historical information of the flow and observations. As shown in figure 20, the state variable $\boldsymbol{U}$ converges to the local observation $\boldsymbol{U}_{obs}$ when the historical information is excluded from the assimilation. When employing the delay-coordinate technique with $\alpha<0$ (indicating that prior estimates are lower than observations), the DA system adjusts the state variable $\boldsymbol{U}$ to a value marginally higher

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

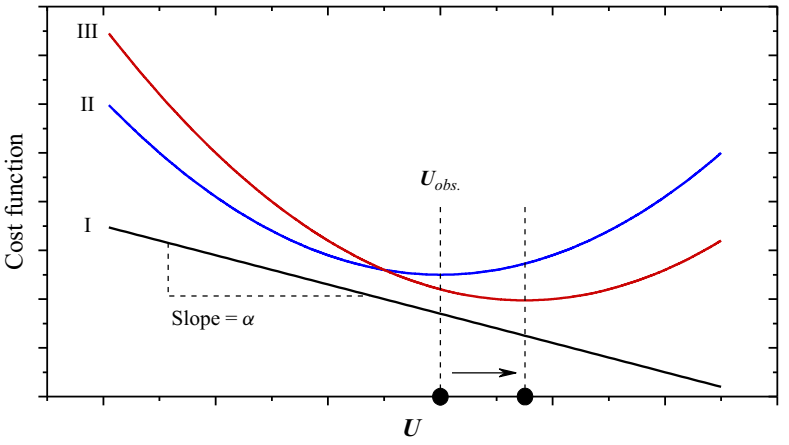


Figure 20. Schematic demonstration of the cost function: I, the discrepancy term containing the historical information; II, the cost function without the delay-coordinate technique; III, the cost function with the delay-coordinate technique.

than the local observation. This adjustment creates a negative feedback mechanism that effectively counterbalances discrepancies introduced through the historical data pathway. However, it is worth noting within the SeVar algorithm that the delay-coordinate technique demonstrates effectiveness only when $\sigma > 0$ in the surrounding cells of the assimilation cube, as the historical discrepancy approaches zero in centre cells or when $\sigma = 0$. Consequently, the current implementation of the delay-coordinate technique preserves accurate flow reconstruction in cells containing high-precision observations. It is worth noting that the delay-coordinate mechanism does not strictly guarantee that the reconstructed flow matches the observations, as it introduces past errors into current computations. In the historical data pathway, minor discrepancies between the reconstruction and observations do not vanish, which also means the feedback mechanism remains active at all times. However, from the perspective of DA using erroneous or noisy observations, this mechanism is statistically effective. Indeed, the approach that incorporates historical information via delay coordinates is less rigorous than the 4DVar method, yet it requires significantly less computational effort. While both mechanisms share certain functional similarities, in this study the delay-coordinate approach is employed to elucidate the relationship between historical information and the modal synchronisation, thereby providing valuable insights for subsequent 4DVar research.

Appendix B. Statistics of the master and slave flows

The flow fields in the master systems of cases 1–4 are validated against DNS data from Abe *et al.* (2001). Figure 21 demonstrates that all mean velocity profiles show reasonable agreement with the data of Abe *et al.*, except for case 14 which exhibits discrepancies due to its relatively low mesh resolution. In case 14, fluctuations are also overpredicted in the buffer layer. Mesh refinement is required to improve predictions using the current LES model, as clearly demonstrated by the improved results in case 12. Nevertheless, the present study extensively investigates the DA scenario using this low-resolution mesh. This case is designed to simulate an experimental situation where sparse measurements are accurate, while direct LES fails to capture the correct flow dynamics. The results demonstrate the potential of DA to reconstruct partial turbulence details even with reduced mesh resolutions, provided that precise sparse observations are available.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

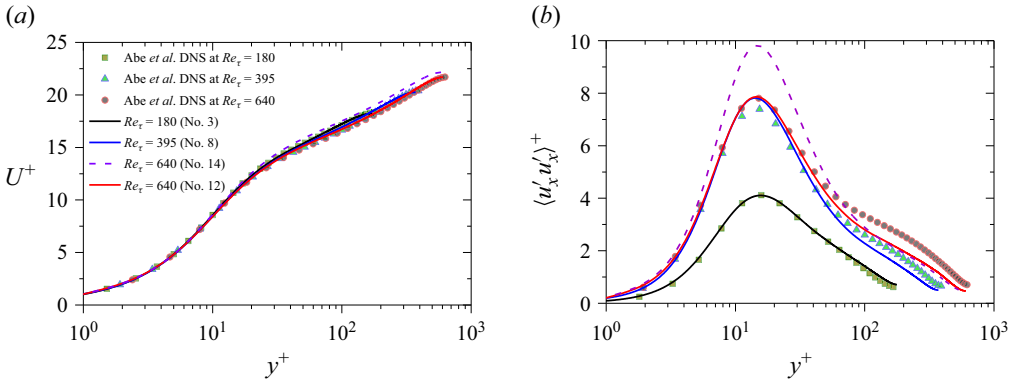


Figure 21. Validation of the (*a*) mean flows and (*b*) fluctuations in the master systems for cases 3, 8, 12 and 14.

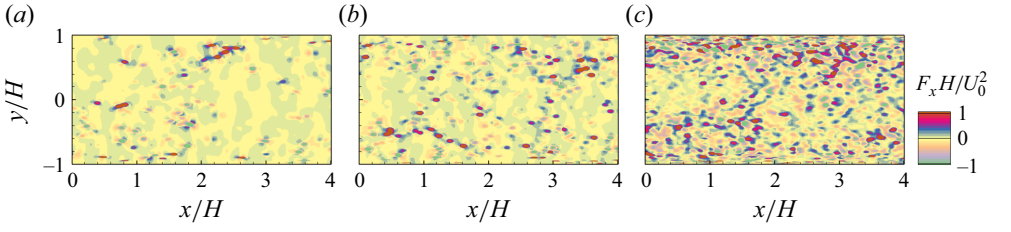


Figure 22. Streamwise component of the body force in SeVar with single-cell DA ($\sigma = \infty$). Demonstration is selected from case 14 using observation II: (*a*) $\mathcal{S} = 0.14$, (*b*) $\mathcal{S} = 0.2$ and (*c*) $\mathcal{S} = 0.4$.

In the SeVar systems, the optimised body force $\boldsymbol{F}$ serves as a compensation for model errors, exhibiting behaviour analogous to that of the feedback term in the nudging technique (Zauner *et al.* 2022). Specifically, both the body force and the feedback term diminish in a synchronised system when a perfect model, i.e. a DNS with identical numerical configurations, is employed for both the master and slave systems. Previous studies on laminar cylinder wakes have demonstrated that $\boldsymbol{F}$ becomes negligible at large times once the system achieves synchronisation with the master. However, in practice, model imperfections and numerical errors prevent $\boldsymbol{F}$ from vanishing entirely, a behaviour fundamentally consistent with the feedback mechanism in nudging. Figure 22 illustrates the normalised F_x for varying observation densities for case 14 using observation configuration II. The core principle of synchronisation lies in using the imposed body force at observation points to induce targeted local fluid adjustments, while exploiting the spatiotemporal correlations inherent to the flow to propagate these changes and achieve coherent motion in the surrounding fluid field.

The introduction of $\boldsymbol{F}$ modifies the flow dynamics relative to unforced conditions. In fully synchronised slave systems using observation configuration I, the flow characteristics precisely match those of the master case, including complete replication of both mean and fluctuating velocity components. However, for slave cases with sparser observations, $\boldsymbol{F}$ only partially contributes to model synchronisation while simultaneously influencing the remaining flow dynamics. This secondary influence results in minor deviations from the logarithmic law in the boundary layer. Although velocity fluctuations show substantial improvement in the slave systems (figure 23), precise reconstruction of the mean velocity profile remains challenging even with relatively dense observations. This finding suggests

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

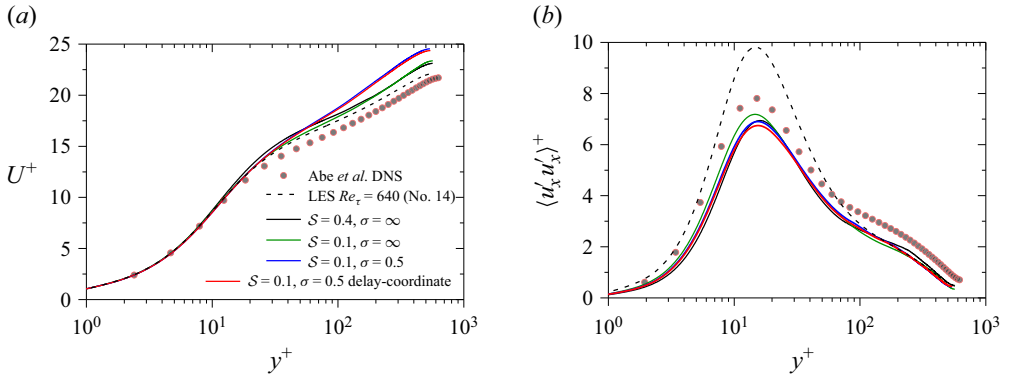


Figure 23. Validation of the assimilated (*a*) mean flows and (*b*) fluctuations in case 14 using the observation configuration III.

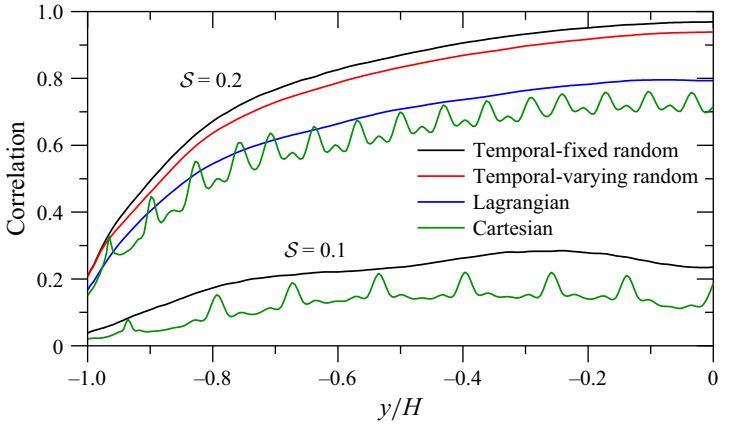


Figure 24. Correlation of the assimilated flow with that in the master system on various wall-parallel planes in case 12 using the observation configuration I. The upper four curves denote the results at $\mathcal{S} = 0.2$ and the lower two curves denote the results at $\mathcal{S} = 0.1$.

that preserving the statistical properties of turbulent boundary layers will represent a crucial research focus for future studies on modal synchronisation.

Appendix C. Comparison of different observation distributions

The effects of different observation distributions on the DA results are compared in terms of correlation of the assimilated flow with that in the master system on various wall-parallel planes. These include random distribution in the entire computational domain with all the observation sites remaining fixed with time (temporal-fixed random), random distribution with all the observation sites reinitialised at each instant (temporal-varying random), Lagrangian distribution with all the observation sites following the trajectories of the fluid parcels and Cartesian distribution with all the observations located on a uniform $\mathcal{S}N_x \times \mathcal{S}N_y \times \mathcal{S}N_z$ Cartesian grid. As evident from figure 24, the Cartesian distribution exhibits pronounced correlation peaks at observation locations, yet maintains significantly lower overall correlation compared with other distributions. The Lagrangian approach demonstrates particularly poor performance, consistent with the findings of Clark Di Leoni *et al.* (2020), who attributed this to insufficient time for flow synchronisation at each point

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

due to probe movement. An alternative explanation may lie in the limited sampling of flow events, as observations track specific fluid parcels. This interpretation is supported by the enhanced correlation achieved with the temporal-varying random strategy, which incorporates more diverse flow events from various fluid parcels. However, the hypothesis of Clark Di Leoni *et al.* gains further validation from the superior performance of the temporal-fixed random distribution, which shows the highest correlation values. These results indicate that the temporal-fixed random distribution represents the optimal strategy for sparse observations, and, consequently, this approach was adopted throughout the present study.

References

Abe, C., Kanda, N., Nakai, K. & Nonomura, T. 2025 Robust and multiresolution sparse processing particle image velocimetry for improvement in spatial resolution. *J. Vis.-Japan* **28**, 21–38.

Abe, H., Kawamura, H. & Matsuo, Y. 2001 Direct numerical simulation of a fully developed turbulent channel flow with respect to the Reynolds number dependence. *J. Fluids Engng* **123**, 382–393.

Boree, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35**, 188–192.

Casa, L.D.C. & Krueger, P.S. 2013 Radial basis function interpolation of unstructured, three-dimensional, volumetric particle tracking velocimetry data. *Meas. Sci. Technol.* **24**, 065304.

Deng, Z., He, C., Wen, X. & Liu, Y. 2018 Recovering turbulent flow field from local quantity measurement: turbulence modeling using ensemble-Kalman-filter-based data assimilation. *J. Vis.-Japan* **21**, 1043–1063.

Di Leoni, PClark, Mazzino, A. & Biferale, L. 2020 Synchronization to big data: nudging the Navier–Stokes equations for data assimilation of turbulent flows. *Phys. Rev. X* **10**, 011023.

Fiscaletti, D., Attili, A., Bisetti, F. & Elsinga, G.E. 2016 Scale interactions in a mixing layer–the role of the large-scale gradients. *J. Fluid Mech.* **791**, 154–173.

Fukami, K., Fukagata, K. & Taira, K. 2019 Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.

Fukami, K., Fukagata, K. & Taira, K. 2021 Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* **909**, A9.

Ganapathisubramani, B., Hutchins, N., Monty, J.P., Chung, D. & Marusic, I. 2012 Amplitude and frequency modulation in wall turbulence. *J. Fluid Mech.* **712**, 61–91.

Gesemann, S., Huhn, F., Schanz, D. & Schroder, A. 2016 From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. In *18th International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal*.

He, C., Liu, Y. & Gan, L. 2020 Instantaneous pressure determination from unsteady velocity fields using adjoint-based sequential data assimilation. *Phys. Fluids* **32**, 035101.

He, C., Wang, P., Liu, Y. & Gan, L. 2022*a* Flow enhancement of tomographic particle image velocimetry measurements using sequential data assimilation. *Phys. Fluids* **34**, 035101.

He, C., Wang, P. & Liu, Y. 2022*b* Sequential data assimilation of turbulent flow and pressure fields over aerofoil. *AIAA J.* **60**, 1091–1103.

He, C., Zeng, X., Wang, P., Wen, X. & Liu, Y. 2024 Four-dimensional variational data assimilation of a turbulent jet for super-temporal-resolution reconstruction. *J. Fluid Mech.* **978**, A14.

Huhn, F., Schanz, D., Manovski, P., Gesemann, S. & Schröder, A. 2018 Time-resolved large-scale volumetric pressure fields of an impinging jet from dense Lagrangian particle tracking. *Exp. Fluids* **59**, 81.

Jeon, Y.J., Müller, M. & Michaelis, D. 2022 Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+. *Exp. Fluids* **63**, 70.

Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S. & Yang, L. 2021 Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440.

Kim, H., Kim, J., Won, S. & Lee, C. 2021 Unsupervised deep learning for super-resolution reconstruction of turbulence. *J. Fluid Mech.* **910**, A29.

Lalescu, C.C., Meneveau, C. & Eyink, G.L. 2013 Synchronization of chaos in fully developed turbulence. *Phys. Rev. Lett.* **110**, 084102.

Li, S., He, C. & Liu, Y. 2023 Unsteady flow enhancement on an airfoil using sliding window weak-constraint four-dimensional variational data assimilation. *Phys. Fluids* **35**, 065122.

Li, T., He, C., Wen, X. & Liu, Y. 2023 Data assimilation of rotor flow at hovering state using ensemble Kalman filter. *J. Vis.-Japan* **26**, 815–839.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press

LI, L. & PAN, Z. 2024 Three-dimensional time-resolved Lagrangian flow field reconstruction based on constrained least squares and stable radial basis function. *Exp. Fluids* **65**, 57.
LI, J., TIAN, M. & LI, Y. 2022 Synchronizing large eddy simulations with direct numerical simulations via data assimilation. *Phys. Fluids* **34**, 065108.
LI, J., TIAN, M., LI, Y., SI, W. & MOHAMMED, H.K. 2024 The conditional Lyapunov exponents and synchronisation of rotating turbulent flows. *J. Fluid Mech.* **983**, A1.
LI, Y., ZHANG, J., DONG, G. & ABDULLAH, N.S. 2020 Small-scale reconstruction in three-dimensional Kolmogorov flows using four-dimensional variational data assimilation. *J. Fluid Mech.* **885**, A9.
LI, S., ZHANG, X., ZHOU, W., HE, C. & LIU, Y. 2024 Particle image velocimetry, delayed detached eddy simulation and data assimilation of inclined jet in crossflow. *J. Vis.-Japan* **27**, 307–322.
LIU, B., TANG, J., HUANG, H. & LU, X.-Y. 2020 Deep learning methods for super-resolution reconstruction of turbulent flows. *Phys. Fluids* **32**, 025105.
LUMLEY, J.L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Propagation, Nauka* (ed. A.M. Aglom & V.I. Tatarski), pp. 166–178.
MATHIS, R., HUTCHINS, N. & MARUSIC, I. 2009 Large-scale amplitude modulation of the small-scale structures in turbulent boundary layers. *J. Fluid Mech.* **628**, 311–337.
NIKOLAIDIS, M.-A. & IOANNOU, P.J. 2022 Synchronization of low Reynolds number plane Couette turbulence. *J. Fluid Mech.* **933**, A5.
PAZO, D., CARRASSI, A. & LOPEZ, J.M. 2016 Data assimilation by delay-coordinate nudging. *Q. J. R. Meteorol. Soc.* **142**, 1290–1299.
POPE, S.B. 2000 *Turbulent Flows*. Cambridge University Press.
SCARANO, F. & MOORE, P. 2012 An advection-based model to incease the temporal resolution of PIV time series. *Exp. Fluids* **50**, 919–933.
SCARANO, F., SCHNEIDERS, J.F., SAIZ, G.G. & SCIACCHITANO, A. 2022 Dense velocity reconstruction with VIC-based time-segment assimilation. *Exp. Fluids* **63**, 96.
SCHNEIDERS, J.F., DWIGHT, R.P. & SCARANO, F. 2014 Time-supersampling of 3D-PIV measurements with vortex-in-cell simulation. *Exp. Fluids* **55**, 1692.
SCHNEIDERS, J.F. & SCARANO, F. 2016 Dense velocity reconstruction from tomographic PTV with material derivatives. *Exp. Fluids* **57**, 139.
SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Maths.* **45**, 561–571.
SPEROTTO, P., PIERACCINI, S. & MENDEZ, M.A. 2022 A meshless method to compute pressure fields from image velocimetry. *Meas. Sci. Technol.* **33**, 094005.
TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.
VELA-MARTÍN, A. 2021 The synchronisation of intense vorticity in isotropic turbulence. *J. Fluid Mech.* **913**, R8.
WANG, Y., YUAN, Z., XIE, C. & WANG, J. 2022 Temporally sparse data assimilation for the small-scale reconstruction of turbulence. *Phys. Fluids* **34**, 065115.
WANG, M. & ZAKI, T.A. 2022 Synchronization of turbulence in channel flow. *J. Fluid Mech.* **943**, A4.
WELCH, P. 1967 The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* **15**, 70–73.
YOSHIDA, K., YAMAGUCHI, J. & KANEDA, Y. 2005 Regeneration of small eddies by data assimilation in turbulence. *Phys. Rev. Lett.* **94**, 014501.
ZAUNER, M., MONS, V., MARQUET, O. & LECLAIRE, B. 2022 Nudging-based data assimilation of the turbulent flow around a square cylinder. *J. Fluid Mech.* **937**, A38.
ZHU, L., JIANG, X., LEFAUVE, A., KERSWELL, R.R. & LINDEN, P. 2024 New insights into experimental stratified flows obtained through physics-informed neural networks. *J. Fluid Mech.* **981**, R1.

https://doi.org/10.1017/jfm.2026.11407 Published online by Cambridge University Press