

UCLA

UCLA Electronic Theses and Dissertations

Title

Flow estimation with point vortex models

Permalink

<https://escholarship.org/uc/item/5mb5g9x4>

Author

Le Provost, Mathieu Jacques Philippe

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Flow estimation with point vortex models

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mechanical Engineering

by

Mathieu Jacques Philippe Le Provost

2022

© Copyright by
Mathieu Jacques Philippe Le Provost
2022

ABSTRACT OF THE DISSERTATION

Flow estimation with point vortex models

by

Mathieu Jacques Philippe Le Provost

Doctor of Philosophy in Mechanical Engineering

University of California, Los Angeles, 2022

Professor Jeffrey D. Eldredge, Chair

In many applications, there is growing interest to use limited pressure observations to estimate the flow behavior. In this dissertation, we represent the flow field by the positions and strengths of a collection of point vortices. We perform the assimilation of the pressure observations with the ensemble Kalman filter (EnKF), which builds a Monte-Carlo approximation of the Kalman gain. As a result of the limited ensemble size, the estimated Kalman gain suffers from ill-conditioning issues, sampling errors, and spurious long-range correlations. We propose two strategies to resolve these problems in the context of flow estimation. First, we show that the flow estimator introduced by Darakananda et al. (Phys. Rev. Fluids 3, 124701 (2018)) is greatly improved by replacing the stochastic version of the EnKF with the ensemble transform Kalman filter (ETKF): a deterministic version of the EnKF that reduces sampling errors. We assess this improved flow estimator on two challenging flow configurations: a flat plate is subjected to strong and overlapping disturbances applied near the leading edge to mimic flow actuation, and a flat plate is placed in the wake of a cylinder. The ETKF significantly improves the estimation of the flow field. Second, predominant methods for regularizing the EnKF suppress correlations at long distances.

In incompressible fluid problems, the observations are given by elliptic partial differential equations, e.g. the pressure Poisson equation. Distance localization is not applicable here, as we cannot distinguish the slowly decaying physical correlations from the spurious long-range ones. In elliptic inverse problems, we observe that a low-dimensional projection of the observations is only informative of a low-dimensional subspace of the state space. We introduce the low-rank EnKF (LREnKF): a novel version of the EnKF that leverages this structure. We identify the most informative directions of the state and observation spaces as the leading eigenvectors of Gramian matrices based on the sensitivity of the observation operator. From the rapid spectral decay, we can estimate a lower-dimensional Kalman gain in the low-dimensional subspace spanned by the leading eigenvectors, hence reducing the variance of the estimator. The LREnKF avoids any *ad-hoc* tuning by adaptively identifying the dimensions of the informative subspace. We show the LREnKF significantly improves the estimate of the stochastic EnKF on two potential flow examples.

The dissertation of Mathieu Jacques Philippe Le Provost is approved.

Christopher R. Anderson

Kunihiko Taira

Xiaolin Zhong

Jeffrey D. Eldredge, Committee Chair

University of California, Los Angeles

2022

To my family...

TABLE OF CONTENTS

1	Introduction	1
2	Background on potential flow theory	9
2.1	The basics of potential flow theory	9
2.2	Pressure calculation in potential flows	11
2.3	Flow modeling over an airfoil with an aggregated vortex model	12
3	Background on the ensemble Kalman filter	16
3.1	A review of the filtering problem	16
3.2	Additional notation	22
3.3	The stochastic ensemble Kalman filter	23
3.4	The ensemble transform Kalman filter	26
3.5	Classical regularization techniques: covariance inflation and localization	29
4	Aerodynamic flow estimation with point vortex models	32
4.1	Flow estimation with an ensemble of aggregated vortex models	34
4.2	Results	36
4.2.1	Translating plate subjected to pulse actuation disturbances	40
4.2.2	Plate in the wake of a cylinder	50
5	Regularization of the ensemble Kalman filter in elliptic inverse problems	57
5.1	Low-rank ensemble Kalman filter	59
5.1.1	Ensemble Kalman filter	59

5.1.2	Assimilation in low-dimensional subspaces	59
5.1.3	Estimation of the leading directions from samples	69
5.2	Examples	70
5.2.1	Multipole expansion and leading directions of the state and observation Gramians	70
5.2.2	Inference of the properties of point vortices from pressure observations along a wall	74
5.2.3	Inference of a vortex patch advected along a wall from pressure obser- vations.	83
6	Conclusions	91
6.1	Summary	91
6.2	Future directions	93
A	Algorithms of the stochastic, deterministic, and low-rank ensemble Kalman filters	96
B	Mean transport of inertial particles in viscous streaming flows	101
B.1	Introduction	101
B.2	Basic transport for inertial particles	105
B.2.1	The Maxey–Riley equation with Saffman lift	107
B.2.2	The inertial particle velocity field	111
B.2.3	Asymptotic expansion of the Maxey–Riley equation in small Stokes number	112
B.3	Development of the equations for mean particle transport	114
B.3.1	Development of the Lagrangian mean field equations	116

B.3.2	Asymptotic reduction for small oscillation amplitude	127
B.3.3	Algorithm for Lagrangian averaged transport of particles	134
B.4	Results	137
B.4.1	Small Stokes expansion of inertial particle velocity field	141
B.4.2	Mean fluid particle trajectories	143
B.4.3	Mean inertial particle trajectories	145
B.4.4	The effect of transient behavior	149
B.5	Conclusions	153
References	157

LIST OF FIGURES

2.1	Schematic of the two-dimensional infinitely-thin plate.	13
4.1	One time step of the data-assimilated aggregated vortex model	33
4.2	(a): Schematic of the plate at 20° at $t^* = 0$ subject to actuation. Green dot depicts the location of the actuation. (b) Time history of the flow actuation. . .	40
4.3	Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° subjected to pulse actuation disturbance from (a) high-fidelity numerical simulation at Reynolds number 500, and mean over 100 realizations of an inviscid vortex model with (b) the sEnKF and (d) the ETKF. Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for (c) the sEnKF and (e) the ETKF.	42
4.4	Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° subject to actuation, from high-fidelity numerical simulation at Reynolds number 500 (—), mean over 100 realizations of the inviscid vortex model with sEnKF (—), mean over 100 realizations of the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Time history of the flow actuation (—). Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—) and ETKF (—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—) and ETKF (—) without application of the median filter.	43

4.5	Snapshots of the vorticity distribution at $t^* = 3.0$ (left column), $t^* = 4.0$ (middle column) and $t^* = 5.0$ (right column) for an impulsively translating plate at 20° subject to actuation, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.	45
4.6	Time history of the ensemble mean value of the LESPc of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.	46
4.7	Time history of the particle count for an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.	48
4.8	Time history of the ensemble variances of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations. Mean variances for (a) the x coordinate of the blobs, (b) the y coordinate of the blobs, (c) the circulation of the blobs, and (d) the LESPc from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Fainter dashed lines show the standard deviation of the different variances over the 100 realizations for the inviscid vortex model with sEnKF and ETKF.	49
4.9	Schematic of a plate at 20° behind a cylinder. The plate and cylinder translate to the right at uniform speed.	50

4.10	Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° in a cylinder wake (a) high-fidelity numerical simulation at Reynolds number 500, mean over 100 realizations of an inviscid vortex model with the sEnKF (b) and the ETKF (d). Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for the sEnKF(c) and ETKF (e)	51
4.11	Snapshots of the vorticity distribution at $t^* = 3.5$ (left column), $t^* = 4.2$ (middle column) and $t^* = 10.0$ (right column) for an impulsively translating plate at 20° in a cylinder wake, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.	52
4.12	Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° in a cylinder wake, from high-fidelity numerical simulation at Reynolds number 500 (—), mean over 100 realizations of the inviscid vortex model with sEnKF (—), mean over 100 realizations of the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—) and ETKF (—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—) and ETKF (—) without application of the median filter.	53

4.13	Time history of the ensemble mean value of the LESPc of an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.	54
4.14	Time history of the particle count for an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.	54
5.1	Schematic of the setup. Orange dots depict the location of the point sources. The yellow square (located at $(0,0)$) depicts the location of the centroid of the point sources. Grey diamonds depict the location of the evaluation points of the velocity potential ϕ	71
5.2	Top panel (a): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th state modes (\mathbf{v}_i s) obtained from the SVD of $\nabla_{\mathbf{x}}\mathbf{h}$ (solid turquoise and solid blue lines), and from orthogonalization of the columns of \mathbf{Q}_x (dashed red and dashed purple lines). Lower panel (b): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th observation modes (\mathbf{u}_i s) obtained from the SVD of $\nabla_{\mathbf{x}}\mathbf{h}$ (solid turquoise and solid blue lines), and from the columns of $\mathbf{Q}_{z'}$ (dashed red and dashed purple lines). The vectors have unit norm.	73

5.3	<p>(a): Schematic of the setup. True trajectories of the point vortices (for one realization of the system) are represented by colored lines. The location of the point vortices, sampled every three convective times, is depicted by colored dots. Fainter dashed lines and fainter dots refer to the image point vortices. Location of the pressure sensors are depicted by grey diamonds. (b): Estimation of the trajectories of the vortices with the LREnKF for $M = 40$ with the ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ set to capture 99% of the cumulative energy spectra. Solid lines depict the time history of the median posterior estimate for the position of the different point vortices. Fainted areas show the 5% and 95% quantiles of the posterior estimate for the position of the point vortices.</p>	75
5.4	<p>Median spectrum of the state Gramian $\mathbf{C}_{\mathbf{x}}$ (left panels (a)-(c)) and observation Gramian $\mathbf{C}_{\mathbf{y}}$ (right panels (d)-(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles. . . .</p>	76
5.5	<p>1, 2, 3 and 8th eigenvectors of the observation Gramian $\mathbf{C}_{\mathbf{y}}$, i.e. observation modes \mathbf{u}_is, at $t = 0.1$ (left column), $t = 7.0$ (middle column), and $t = 10.0$ (right column). (a)-(c): first mode, (d)-(f): second mode, and (g)-(i): third mode, (j)-(l): eighth mode. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. The grey horizontal line corresponds to 0 ordinate. Statistics are obtained from a run of the sEnKF with $M = 1000$. The eigenvectors have unit norm.</p>	77

- 5.6 Left column (a): Time-averaged evolution of the median RMSE with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The tracking performance of the sEnKF is unstable for $M < 40$. Dashed lines depict the 25% and 75% quantiles. Right column [(b)-(c)]: Time-history of the median ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF for $M = 30$ (computed over 50 realizations) for different ratios of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$, $\mathbf{C}_{\mathbf{y}}$: 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue. 79
- 5.7 Panel (a), (b), (c) depicts the magnitude of the empirical cross-covariance between the x , y coordinate and strength of the J point vortex with the k th pressure observation at $t = 1.0$, respectively. The curves are plotted against the distance $d_{J,k} = |z_J - z'_k|$, where $z_J = x_J + iy_J$ is the (complex) position of the J th point vortex, and z'_k is the (complex) location of the k th pressure sensor. The magnitude of the cross-covariances are normalized by the maximum cross-covariance (in magnitude) between the x , y coordinate or strength of the J point vortex with the different pressure observation, respectively. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. Red dashed and dotted curve depicts the algebraic decay $O(1/d_{J,k}^2)$. (d) Mean position of the different point vortices at $t = 1.0$. The same color is used to depict the properties of a point vortex on the panels (a)-(d). The results are obtained from a run of the sEnKF for $M = 1000$ 82

- 5.8 Snapshots of the vorticity distribution at $t = 0.5$ (left column), $t = 1.5$ (second column), $t = 6.0$ (third column), and $t = 12.0$, for a vortex patch and its image, predicted from [(a)-(d)] high-fidelity numerical simulation at Reynolds number 1000, [(e)-(h)] inviscid vortex model with LREnKF with $M = 30$, and [(i)-(l)] inviscid vortex model with sEnKF with $M = 30$. The ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF are set to capture 99% of the normalized cumulative energy of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$, respectively. Pressure sensors are depicted with grey diamonds. Orange denote positive vorticity, while blue vorticity denotes negative vorticity. 84
- 5.9 Median spectrum of the state Gramian $\mathbf{C}_{\mathbf{x}}$ (left panels (a)-(c)) and observation Gramian $\mathbf{C}_{\mathbf{y}}$ (right panels (d)-(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles. . . . 88
- 5.10 Left column (a): Time-averaged evolution of the median MSE of the posterior pressure with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The sEnKF is found unstable for $M < 30$. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles. Right column [(b)-(c)]: Time-history of the median ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF for $M = 20$ (computed over 50 realizations) for different ratios of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$, $\mathbf{C}_{\mathbf{y}}$: 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue. 89

5.11	Top line [(a), (b), (c)]: Spatiotemporal map of the pressure induced by a pair of two co-rotating vortex patches from (a) high-fidelity numerical simulation at Reynolds number 1000, and the mean posterior of an inviscid vortex model for one realization of the LREnKF (b), and the sEnKF (c) with $M = 30$. Bottom line [(d), (e)]: Spatiotemporal magnitude of the error between the truth pressure and the mean posterior pressure for one realization of the LREnKF (d), and the sEnKF (e) with $M = 30$. The ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF are set to capture 99% of the normalized cumulative energy of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$, respectively. The same color levels are used for all the top panels, and another set for the bottom panels.	90
B.1	Schematic of flow maps used in this work. \mathcal{X}_t and \mathcal{Y}_t represent slices of the configuration spaces $\mathcal{X}_t \times [0, \infty)$ and $\mathcal{Y}_t \times [0, \infty)$ at some instant t . Illustrations of particle trajectories are shown as colored curves (though strictly, these trajectories would proceed along the time axis of the respective space).	122
B.2	Mean (green) and actual (blue) particle trajectories with initial position \mathbf{x}_0 . The position \mathbf{x}^ξ is the <i>actual</i> position whose mean is \mathbf{x} . Adapted from Bühler et al. [1].	123
B.3	Overall algorithm for the fast Lagrangian averaged transport of particles.	135
B.4	Diagram of the oscillating cylinder (right), and a time sequence illustrating the repeated oscillation cycles.	138
B.5	Configuration and oscillation sequence with two cylinders adapted from the work of Chong et al. [2].	138
B.6	Time-resolved inertial particle trajectories from the Maxey–Riley equation (black) and the asymptotic expansion in small Stokes number (green) for particles initially located at $\mathbf{x}_0 = (2, 2)$ (left column) and $\mathbf{x}_0 = (1, 3)$ (right column). The top row (a,d) shows the trajectories over the first 20 oscillation periods, and the middle (b,e) and lower (c,f) rows depict the time histories of the x and y components, respectively, sampled once per period.	140

B.7	Trajectory of an inertial particle initially located at $\mathbf{x}_0 = (-2, 3)$ over 65000 periods of oscillation. Trajectories from the Maxey–Riley equation (black) and small Stokes number expansion (green) are both sampled once per period. . . .	143
B.8	Lagrangian mean trajectories (magenta circles, spaced by 30 periods) generated for fluid particles started at $\mathbf{x}_0 = (2, 2)$ and $\mathbf{x}_0 = (1, 3)$ (denoted by the larger grey circle), compared with the Lagrangian streamlines, depicted as black lines. (Other Lagrangian streamlines are shown in light grey.) The black region shows the mean position of the cylinder, and the lighter shaded region in the vicinity of the cylinder shows the range of displacement of the cylinder over one oscillation cycle.	144
B.9	(a) Lagrangian mean trajectory (magenta) for fluid particle started at $\mathbf{x}_0 = (2, 2)$, compared with the full time-resolved trajectory (blue) for the same particle. (b) Magnified view of fluid particle trajectory. (c,d) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle).	146
B.10	Left: Time resolved fluid particle trajectory, sampled once per cycle (blue), and Lagrangian mean trajectory (magenta circles, spaced by 30 periods), and Lagrangian mean streamlines (in light gray) for $\mathbf{x}_0 = (-1, 3)$, with mean cylinder configurations depicted in black. Right: Time histories of the x and y components of both trajectories shown in the left panel.	147
B.11	Lagrangian mean trajectories (magenta) generated for inertial particles started at (a) $\mathbf{x}_0 = (2, 2)$ and (b) $\mathbf{x}_0 = (1, 3)$ (denoted by the small circle), compared with the full time-resolved trajectories (blue) of the same particles, sampled once per cycle. The Lagrangian streamlines of the fluid are shown for reference as light grey lines. The black region shows the mean position of the cylinder, and the lighter shaded region in the vicinity of the cylinder shows the range of displacement of the cylinder over one oscillation cycle.	149

B.12	Left column: (a,b) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle) for particle $\mathbf{x}_0 = (2, 2)$. (c) Lagrangian mean trajectory (magenta) for inertial particle $\mathbf{x}_0 = (2, 2)$, compared with the full time-resolved trajectory (blue) for the same particle. Right column: (d,e) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle) for particle $\mathbf{x}_0 = (1, 3)$. (f) Lagrangian mean trajectory (magenta) for inertial particle started at $\mathbf{x}_0 = (1, 3)$, compared with the full time-resolved trajectory (blue) for the same particle.	150
B.13	Lagrangian mean trajectories generated for inertial particles with $\rho_p/\rho_f = 0.95$ (magenta) and 0.05 (blue) started at $\mathbf{x}_0 = (2, 2)$	151
B.14	(a) Full time resolved inertial particle trajectory, sampled once per cycle (blue), and Lagrangian mean trajectory (magenta) for a particle released from $\mathbf{x}_0 = (-2, 3)$ in the sequential oscillator configuration. (b,c) Time histories of the x and y trajectory components, respectively.	152
B.15	Time resolved inertial particles trajectories (blue) and mean transport algorithm (yellow) for $\mathbf{x}_0 = (-2.0, 3.0)$ over the transient regime of the left cylinder (left) and $\mathbf{x}_0 = (-1.9285, 1.0266)$ over the transient regime of the right cylinder (right).	153

LIST OF TABLES

4.1	Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject to actuation with the sEnKF and the ETKF, without application of the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.	49
4.2	Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject in a cylinder wake with the sEnKF and the ETKF without the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.	56
B.1	Relative errors on particle position, after 1 period and at the final time t_f , in two different inertial particles trajectories predicted by the small Stokes number expansion.	142

ACKNOWLEDGMENTS

First, I would like to express my sincere gratitude to my advisor, Professor Jeff D. Eldredge for your invaluable guidance and kindness. Thank you for all opportunities you have offered me. Finishing a M.S. in December, traveling to LA and starting a Ph.D. in January would not have been possible without you.

I would like to express my deepest appreciation to Professor Christopher R. Anderson, Professor Kunihiko Taira, and Professor Xiaolin Zhong for serving on my committee.

I cannot thank enough Ricardo Baptista and Professor Youssef Marzouk for their devoted support and kindness. It has been such an incredible experience to collaborate with you. Thank you Ricardo for always taking the time to meet with me. I cannot stress how much I have learnt from our discussions.

I gratefully acknowledge the support of this work by the National Science Foundation under Grant No. 1538824, the Air Force Office of Scientific Research under Grant No. FA9550-18-1-0440, and the Dissertation Year Fellowship offered by the Graduate Division of the University of California, Los Angeles.

Many thanks to my labmates of the Sofia lab: Dany, Darwin, Diederick, Rhuizi, Yi-Jui, Yuxuan and Wei. I really appreciate the time I have spent with you.

Chapter 3 and Chapter 4 are adapted from Le Provost, M., & Eldredge, J. D. (2021). Ensemble Kalman filter for vortex models of disturbed aerodynamic flows. *Physical Review Fluids*, 6(5), 050506 [3]. Chapter 2 and Chapter 5 are adapted from Le Provost, M., Baptista, R., Marzouk, Y., & Eldredge, J. D. Regularization of the ensemble Kalman filter for elliptic observations. arXiv, 2203.05120.

Many thanks to Antoine and Justine for being the coolest roommates in the West. I can't describe how much I learnt and discovered by living with you. To David, Dulaj, Meryem (and Chouquette), thank you for always making feel at home at your place. I also would like to thank my amazing art teacher Buena Johnson. To my other friends, here and elsewhere,

thank you for always being here.

Finally, I would never be thankful enough to my family, my sister, and my parents for their love and encouragement.

VITA

- 2018 M.S. in Mechanical and Aerospace Engineering, Illinois Institute of Technology, Chicago, Illinois.
- 2018 M.S. in Mechanical and Aeronautical Engineering, Ecole Nationale Supérieure de Mécanique et d'Aérotechnique, Poitiers, FRANCE

PUBLICATIONS

Le Provost, M., Baptista, R., Marzouk, Y., & Eldredge, J. D. (2022). A low-rank ensemble Kalman filter for elliptic observations. arXiv, 2203.05120.

Le Provost, M., & Eldredge, J. D. (2021). Ensemble Kalman filter for vortex models of disturbed aerodynamic flows. *Physical Review Fluids*, 6(5), 050506.

Le Provost, M., Baptista, R., Marzouk, Y., & Eldredge, J. D. (2021). A low-rank nonlinear ensemble filter for vortex models of aerodynamic flows. In *AIAA Scitech 2021 Forum* (p. 1937).

Eldredge, J. D., Le Provost, M., Baptista, R., & Marzouk, Y. (2021). Applications of ensemble Kalman filtered vortex modeling to gust–wing interactions. In *AIAA Scitech 2021 Forum* (p. 1936).

Le Provost, M., & Eldredge, J. D. (2020). Mean transport of inertial particles in viscous streaming flows. *Physical Review Fluids*, 5(5), 054302.

Le Provost, M., Hou, W., & Eldredge, J. D. (2020). Deep learning and data assimilation approaches to sensor reduction in estimation of disturbed separated flows. In AIAA Scitech 2020 Forum (p. 0799).

CHAPTER 1

Introduction

Light-weight and agile aerial vehicles can experience large amplitude and overlapping flow perturbations. Unfortunately, in this regime, the linear theory developed for airfoils at small angle of attack [4, 5, 6] is inadequate to estimate the nonlinear interactions between the disturbances and the vehicle. The development of unmanned aerial vehicles operating in this extended flight envelope requires an accurate model of the flow response to disturbances and flow actuation. This question is the common thread of this dissertation.

A major challenge in real world flow estimation problems is the limited (both spatially and temporally) observations of the surrounding flow field available onboard. Indeed, it is not possible to probe locally the fluid velocity field about the vehicle, as we cannot have floating Pitot sensors in the surrounding fluid region. Unfortunately, the flow estimation must be performed from a sparse set of surface pressure receptors. While pressure observations provide valuable information on the flow field, these observations are limited, noisy, and nonlinearly related to a velocity description of the surrounding flow field. In real flight conditions, the goal of flow estimation is to use these limited surface pressure observations to estimate the surrounding flow field in the presence of unknown disturbances and flow actuation. This inverse problem can be tackled with a purely regressional approach, where we seek to learn the mapping from the pressure observations to an estimate of the flow field (potentially returned as an estimation of the surrounding fluid velocity field) [7, 8]. This direction has gained a considerable interest over the last decade with the novel opportunities offered with deep learning [9, 10]. We refer readers to Fukami et al. [11] for a review of

supervised machine learning techniques for fluid flow estimation. In this dissertation, we prefer a hybrid solution, where the flow estimation is based on a simplified physics-based dynamical model, whose missing physical effects are reintroduced in the flow estimate via the assimilation of pressure observations.

Aerodynamic flow estimation problems are usually ill-posed: different flow fields can produce the same pressure distribution at limited locations on the surface of the airfoil. Individually taken, computational models, or noisy observations are often not sufficient to accurately estimate the flow field. The mathematical model for the flow field suffers from errors due to the choice of the discretization scheme, the initial and boundary conditions, the parameters of the model or simplifying assumptions. Moreover, the dynamical model cannot represent *a priori* the flow disturbances, which are by nature unknown. The observations of the state are usually limited, noisy and nonlinearly related to the state. Using the dynamical problem to generate a prior estimate for the flow state is essential to regularize the inverse problem. The field of *data assimilation (DA)* provides an elegant means to incorporate these observations in the dynamical model [12, 13]. The dynamical model gives an *a priori* prediction for the probability distribution of the state. Using Bayes theorem, the observations are used to improve the estimation of the state distribution and form the *a posteriori* state estimate. This dissertation focuses on the *filtering problem*, which is of primer interest for real-time flow estimation. We seek to estimate the conditional distribution of the state given the knowledge of all the observations up to that time. For high dimensional problems, we usually use a Monte-Carlo approximation of the filtering distribution: an ensemble of samples that approximate the filtering density is recursively updated. Within each assimilation cycle, we perform the following two steps: a *forecast step* and an *analysis step*. In the forecast step, the filtering ensemble is propagated through the dynamical model to generate samples from the *forecast* ensemble. The analysis step updates the forecast ensemble by assimilating the newly available observation from the true system. Unlike the forecast, the analysis step does not involve time propagation and can be treated as a *static* Bayesian

inference problem. It is enlightening to see this Bayesian inference step as the application of a prior-to-posterior transformation – the *analysis map* – to the forecast samples. Different ensemble filtering methods estimate this analysis map under different assumptions given limited samples from the forecast density. In particular, the *ensemble Kalman filter (EnKF)* (Evensen [14]) builds an empirical estimation of the Kalman gain — the linear mapping from observations discrepancies to state updates — derived for linear-Gaussian state-space models [15].

The dynamical model in a flow estimation problem can take diverse forms. da Silva and Colonius [16] used a discrete solver of the Navier-Stokes equation based on the immersed boundary projection method [17]. They successfully estimated the flow field about a two-dimensional airfoil subject to a time-varying freestream from limited surface pressure observations collected on the surface of an airfoil. To alleviate the computational cost of a fully-resolved Navier-Stokes simulation, they relied in a following work on a coarse-grid representation of the flow field [18]. To balance the bias error introduced by the truncation errors of the grid mesh, they developed a closure model for the bias error based on autoregressive algorithm. Deep learning has an important role in the development of closure models for truncation errors terms [19, 20]. In the current state of the art, it is not clear that these coarse-mesh strategies can be fast enough for real-time flow estimation without dramatic truncation errors. However, recent progresses in quantum computing could significantly reduce the computational cost of classical algorithms like the Poisson solver, which is the limiting component of any Navier-Stokes solver [21].

In this dissertation, we pursue an alternative direction, and use a low-dimensional Lagrangian representation of the flow field by tracking the positions and strengths of a limited collection of point vortices. Point vortex models constitute a versatile and highly expressive template to explain and model unsteady aerodynamic phenomena [22, 6, 23, 24, 25]. Indeed, changing the location of the vortex elements is a highly sensitive knob for flow approximation. Obviously, inviscid point vortex models do not explicitly account for viscosity. Fortunately,

for moderate to high Reynolds number aerodynamic flows, the effects of viscosity are limited to two main contributions: the creation of vorticity about the surface of the airfoil, and the viscous dissipation [6]. In an inviscid vortex model, the vorticity flux from the edges of the airfoil can be modeled by the shedding of new vortex elements. About the trailing edge, the Kutta condition is commonly used to determine the strength of the newly shed point vortices. However, the application of the Kutta condition is questionable about the leading edge, as it leads to a non-physical reverse flow. Ramesh et al. [26, 27] introduced the concept of leading edge suction parameter (LESP), as a non dimensional measure of the amount of suction (negative integrated pressure) about the leading edge of the airfoil. They showed from a combination of wind-tunnel and numerical experiments that the flow over the leading-edge of an airfoil can support a finite amount of suction before flow separation. The associated value of the LESP that triggers flow separation is called the critical LESP (LESP_c). One can formulate a simple vortex shedding criterion from this definition: a new vortex element is shed about the leading edge if the LESP exceeds the LESP_c, otherwise no vorticity is shed. The strength of the newly shed vortex is set by the amount by which the LESP exceeds the LESP_c. Eldredge [6] showed that this shedding criterion generalizes the Kutta condition, which is recovered for LESP_c = 0. Ramesh et al. [26, 27] hypothesized that the LESP_c is only a function of the shape of the airfoil. This hypothesis has been confirmed to hold for a wide variety of unsteady flows generated by an airfoil undergoing unsteady maneuvers [28]. However, in the presence of unknown flow disturbances, LESP_c is unknown and time-varying. With the inference of this single parameter, an inviscid vortex model can represent the vorticity flux generated by unknown flow perturbations. We should emphasize that the LESP_c is likely to be observable from the surface pressure observations, as it represents a measure of the integrated pressure load about the leading edge of the airfoil.

The continuous shedding of point vortices (two vortex elements are typically shed at each time step) constitutes a major limitation of classical point vortex models for simulations

past several convective times. To maintain a limited vortex population, Darakananda et al. [29, 25] developed an aggregation procedure for point vortices based on the impulse matching procedure. This aggregated vortex model outfitted with an estimate of the time-varying LESPC was used as a basis for flow estimation in [29]. The state variable was composed of the positions and strengths of a limited collection of point vortices, and the value of LESPC for a flat plate. The EnKF was used to assimilate surface pressure jumps across a flat plate in the state estimate. The truth pressure observations were generated from a high-fidelity simulation at $Re = 500$ [17, 30]. This framework successfully estimated the normal force of a flat plate at 20° with less than 50 vortex elements. This dissertation builds upon this flow estimation framework by following two main directions. First, one of their examples look at a flat plate subject to vortical gusts applied near the leading edge of the airfoil. Their framework accurately captured the flow response to these disturbances, with no representation of the perturbation in the dynamical model. The presence of the gusts was only accounted through the assimilation of the truth pressure observations. This is an appealing feature, as gusts are in essence unknown. In this dissertation, we show that this flow estimator remains accurate in the presence of stronger and overlapping perturbations, for which the individual contribution of the perturbations cannot be disentangled.

Second, Darakananda et al. [29] noted that the performance of the estimator can vary significantly from one realization to the next of the same flow experiment. We observe the creation of nonphysical spikes in the pressure and force estimates at random instants. These undesired events are caused by spurious updates of the vortex elements in the analysis step of the EnKF. The EnKF estimates the Kalman gain, a linear operator that maps pressure discrepancies to updates of the vortex elements, from limited samples of the forecast distribution. Thus, the empirical estimate of the Kalman gain suffers from sampling errors, rank-deficiency and possible spurious long-range correlations. To alleviate these effects, this dissertation explores several options to regularize the EnKF with point vortex models. The original algorithm of the EnKF, called the stochastic EnKF (sEnKF), artificially perturbs

the observations to account for the additive observation noise [14]. This introduces additional sampling errors in the empirical Kalman gain. The sEnKF does not reproduce the exact posterior covariance relation of the Kalman filter with a finite ensemble [13]. Since this seminal work, deterministic versions of the EnKF have been introduced, such as the ensemble transform Kalman filter (ETKF) (Bishop et al. [31]), to exactly reproduce the expected covariance relation with a finite ensemble. We have found that using the ETKF instead of the sEnKF significantly reduces the occurrence of the spurious updates in the flow estimation [3].

Localization is another important regularization technique based on the assumption that the observations are *local*, i.e. an observation only provides information about a subset of the state variables which are close-by in physical distance. This assumption is supported by the rapid decay of the correlations between state and observation variables. In this setting, distance localization removes long-range correlations [13]. However, in many inverse problems, the observations are given as *non-local* mappings of the state variables. Non-local observation operators such as integral of linear/nonlinear functions of the state are frequently encountered in fluid mechanics: mass flux through a surface, force exerted by a fluid on a surface, or the solution of an elliptic partial differential equation (PDE). For the Poisson equation, a canonical elliptic PDE, the solution can be obtained by convolution of the right hand-side with the Green's function associated with the Laplace operator. Elliptic PDEs are of primer importance in flow estimation. For the incompressible Navier-Stokes equations, the pressure field is given as a solution of the pressure Poisson equation: a Poisson equation whose right-hand side depends nonlinearly on the vorticity field (i.e. our state representation of the flow field). From the logarithmic or algebraic decay of the Green's function of the Laplace operator (depending on the dimension of the physical space), we expect long-range physical correlations between the state and observation variables. By systematically removing long-range interactions between the state and observation variables, distance localization would be harmful to regularize flow estimation problems. A major achievement of this dissertation is to

introduce a novel regularization of the EnKF for filtering problems with elliptic observation models. In fluid mechanics, fluid particles (particles of infinitesimally small radius) behave as coherent structures. Intuitively, limited pressure observations provide information about the large-scale structures of the flow, and not about the contribution of each fluid particle. Second, observations are corrupted by noise, making it difficult to extract features smaller in amplitude than the level noise. Based on these observations, incompressible flow estimation problems usually possess a low effective dimension: a low-dimensional projection of the observations variables is only informative along a limited number of directions of the state space. Spantini et al. [32] identified the informative directions of the state space, based on the spectrum of the observation matrix in the linear-Gaussian case (a linear observation model corrupted with an additive Gaussian noise). Cui et al. [33] generalizes this result for a nonlinear observation model, based on the Jacobian of the observation model with respect to the state variable. To the best of our knowledge, there is no procedure to identify the most informative directions in the observation space. By analogy, we propose a similar dimension reduction for the observation space. The methodology is applicable to any observation model, for which the Jacobian can be evaluated. To leverage the structure of the inference problem, we derive a factorization of the Kalman gain based on the identified modes in the linear and nonlinear cases. We observe that a few modes can well approximate the row and column spaces of the Kalman gain. We regularize the Monte-Carlo estimate of the Kalman gain of the EnKF, by performing the inference in the low-dimensional subspace spanned by these modes. For a limited ensemble size, our strategy significantly reduces the variance of the Kalman gain, while making little sacrifice on the bias error [34]. The proposed observation-informed low-rank EnKF (LREnKF) can be summarized as the following sequence of steps: identify and project the state and observation variables on the leading directions of the state and observation spaces, compute the Kalman gain in these reduced coordinates, and lift the result to the original variable space. We show that this novel regularization significantly improves the estimation of the positions and strengths of point vortices provided by the

sEnKF.

The rest of this manuscript is organized as follows. Chapter 2 presents a review of potential flow theory. Chapter 3 provides background information on the filtering problem, and the ensemble Kalman filter. Chapter 4 presents the results of flow estimation with the aggregated point vortex model. Chapter 5 introduces a novel regularization of the EnKF for elliptic inverse problems, and applies it to flow estimation problems with point vortices. Conclusions and directions for future work are presented in Chapter 6. For the sake of completeness and clarity, pseudo-codes for the sEnKF, the ETKF, and the LREnKF are presented in Appendix A. In the course of my doctoral degree, I developed computational tools to predict the mean transport of inertial particles in oscillating viscous flows generated by weakly oscillating surfaces (viscous streaming flow)[35]. This work was published as Le Provost, M. and Eldredge, J. D. (2020). Mean transport of inertial particles in viscous streaming flows. *Physical Review Fluids*, 5(5), 054302. [35]. An adapted version is presented in Appendix B.

The contributions of this dissertation can be summarized as follows:

- I improved the flow estimator based on a point vortex model — first introduced by Darakananda et al. [29] — to predict the flow response to large-amplitude and overlapping flow perturbations about an impulsively translating flat plate at an angle of attack of 20° .
- I introduced a novel regularization of the EnKF for elliptic inverse problems, and applied it to potential flow problems.
- I developed a closed set of equations that governs the mean motion of inertial particles in oscillatory viscous flows (viscous streaming flows). I showed that these equations allow for efficient and accurate predictions of particle transport with numerical steps that are $O(10^3)$ larger than the existing approach.

CHAPTER 2

Background on potential flow theory

This dissertation focuses on the problem of flow estimation about an impenetrable surface (i.e. the surface of an airfoil) with a point vortex model. This chapter provides a minimal review of potential flow theory (see Eldredge [6] for further details). We use this review to highlight the challenges of estimating the positions and strengths of point singularities from limited pressure observations. These challenges motivate the introduction of a novel regularization of the EnKF presented in Chapter 5. Section 2.1 introduces our notations and presents the Biot-Savart law: the dynamical model that advects the point singularities. Section 2.2 presents the pressure calculation in potential flows from the inversion of the pressure Poisson equation or from the unsteady Bernoulli equation. To simplify our discussion, the Biot-Savart law, the pressure Poisson equation, and the unsteady Bernoulli equation are presented for a collection of point singularities without body. Section 2.3 briefly describes how to account for the presence of an additional body in the previous results and presents the aggregated vortex model.

2.1 The basics of potential flow theory

A potential flow refers to a flow field which is both incompressible and irrotational almost everywhere [6]. We restrict ourselves to two-dimensional problems, but we remark that all of the material is easily extended to three-dimensional problems. Here we use “almost everywhere” to allow for non-zero dilatation rate $\theta = \nabla \cdot \mathbf{v}$ and/or non-zero vorticity $\boldsymbol{\omega} = \nabla \times \mathbf{v}$ at most on sets of zero volume of \mathbb{R}^2 , where $\mathbf{v}(\mathbf{r}, t)$ denotes the velocity field. With

these assumptions, the velocity field $\mathbf{v}(\mathbf{r}, t)$ can be derived from the velocity potential $\phi(\mathbf{r}, t)$ as $\mathbf{v} = \nabla\phi$, or alternatively, from the streamfunction $\psi(\mathbf{r}, t)$ as $\mathbf{v} = \nabla \times \psi \mathbf{e}_z$. In this study, we use both vector notation and complex notation. A point in space will either be referred in vector notation as $\mathbf{r} = (x, y) \in \mathbb{R}^2$, or in complex notation as $z = x + iy$. The conjugate of a complex number $z = \Re(z) + i\Im(z)$ is denoted as $\bar{z} \equiv \Re(z) - i\Im(z)$. We denote the canonical basis of \mathbb{R}^3 as $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$.

We consider a set of N singularities located at $\{z_1, \dots, z_N\}$ with complex strengths $\{S_1, \dots, S_N\}$, subject to a uniform flow with velocity $\mathbf{U}_\infty = (U_\infty, V_\infty)$ or in complex notation $W_\infty = U_\infty - iV_\infty$. The strengths are represented as $S_J = Q_J - i\Gamma_J$ where $Q_J, \Gamma_J \in \mathbb{R}$ denote the volume flux and the circulation of the J th singularity, respectively. A point source of volume flux Q_J corresponds to a singularity of real strength $S_J = Q_J$, while a point vortex of circulation Γ_J corresponds to a singularity of imaginary strength $S_J = -i\Gamma_J$. By linearity, the complex potential $F = \phi + i\psi$ induced by these elements at a location $z \in \mathbb{C}$ is

$$F(z) = \sum_{J=1}^N \frac{S_J}{2\pi} \log(z - z_J) + W_\infty z. \quad (2.1)$$

The derivative of the complex potential is the complex velocity w given by:

$$w(z, t) = \frac{dF}{dz} = \sum_{J=1}^N \frac{S_J}{2\pi} \frac{1}{z - z_J} + W_\infty. \quad (2.2)$$

The induced velocity can also be written in vector notation as:

$$\mathbf{v}(\mathbf{r}, t) = \sum_{J=1}^N -Q_J \mathbf{k}(\mathbf{r} - \mathbf{r}_J) + \sum_{J=1}^N \mathbf{k}(\mathbf{r} - \mathbf{r}_J) \times \Gamma_J \mathbf{e}_z + \mathbf{U}_\infty, \quad (2.3)$$

where \mathbf{k} is the velocity kernel given by $\mathbf{k}(\mathbf{r}) = -\mathbf{r}/(2\pi|\mathbf{r}|^2)$. For a velocity field given in vector notation by $\mathbf{v}(\mathbf{r}, t) = (u(\mathbf{r}, t), v(\mathbf{r}, t))$, we should stress the classical conjugation (i.e. minus sign) in the definition of the equivalent complex velocity $w(z, t) = u(z, t) - iv(z, t)$ for consistency with the Cauchy-Riemann equation [6]. Using the induced velocity (2.2), we can write the transport equation for the positions and strengths of the singularities. By simple inspection of (2.2), the velocity field is singular at the locations of the singularities.

Per Kirchhoff's law, a point singularity cannot induce velocity on itself [6]. Instead, a singularity is transported by the total induced velocity on itself minus its own contribution. The regularized velocity of the J th singularity is called the Kirchhoff velocity, denoted \mathbf{v}_{-J} in vector notation, or w_{-J} in complex notation. The dynamic of these singularities is governed by the following set of ordinary differential equations (ODEs):

$$\frac{dz_J}{dt} = \overline{w_{-J}(t)} = \sum_{K=1, K \neq J}^N \frac{\overline{S_J}}{2\pi} \frac{1}{\overline{z_J} - \overline{z_K}} + \overline{W_\infty}, \quad \frac{dS_J}{dt} = 0, \quad J = 1, \dots, N. \quad (2.4)$$

2.2 Pressure calculation in potential flows

For potential flows, the pressure can be computed by two means: from the inversion of the pressure Poisson equation or from the unsteady Bernoulli equation. The Euler equation and the associated divergence and curl conditions on the velocity field are given by:

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} = -\nabla p, \quad (2.5a)$$

$$\nabla \cdot \mathbf{v}(\mathbf{r}, t) = \sum_{J=1}^N Q_J \delta(\mathbf{r} - \mathbf{r}_J). \quad (2.5b)$$

$$\nabla \times \mathbf{v}(\mathbf{r}, t) = \sum_{J=1}^N \Gamma_J \delta(\mathbf{r} - \mathbf{r}_J). \quad (2.5c)$$

The pressure Poisson equation is formed by taking the divergence of the Euler equation (2.5a) and using the identity $\mathbf{v} \cdot \nabla \mathbf{v} = \nabla(\|\mathbf{v}\|^2/2) - \mathbf{v} \times \boldsymbol{\omega}$:

$$\nabla^2 \left(p + \frac{1}{2} \rho \|\mathbf{v}\|^2 \right) = \rho \nabla \cdot (\mathbf{v} \times \boldsymbol{\omega}) - \rho \frac{\partial \nabla \cdot \mathbf{v}}{\partial t}, \quad (2.6)$$

where $\boldsymbol{\omega}(\mathbf{r}, t)$ is the vorticity. Using (2.5b) and (2.5c), we obtain the pressure Poisson equation:

$$\nabla^2 \left(p + \frac{1}{2} \rho \|\mathbf{v}\|^2 \right) = \rho \sum_{J=1}^N [Q_J \mathbf{v}_{-J} \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_J) - \Gamma_J \mathbf{e}_z \cdot (\mathbf{v}_{-J} \times \nabla \delta(\mathbf{r} - \mathbf{r}_J))]. \quad (2.7)$$

From the elliptic nature of eq. (2.7), we expect long-range interactions between point singularities and pressure observations. This aspect is critical for inference in incompressible

fluid flows. In Chapter 1, we showed the limitations of the existing regularization methods of the EnKF to handle these long-range interactions. These limitations will motivate a novel regularization technique that exploits the structure of the inference problem introduced in Chapter 5. Equivalently to (2.7), the pressure induced by point singularities can be obtained in closed form from the unsteady Bernoulli equation. For a fixed evaluation point denoted \mathbf{r}' in vector notation, or z' in complex notation, the unsteady Bernoulli equation reads [6]:

$$p(\mathbf{r}', t) + \rho \frac{1}{2} \|\mathbf{v}(\mathbf{r}', t)\|^2 + \rho \frac{\partial \phi(\mathbf{r}', t)}{\partial t} = B(t), \quad (2.8)$$

where $B(t)$ is the Bernoulli constant. The second term of (2.8) is called the quadratic term, while the third term is called the unsteady term. Using the results derived in 2.1, we get

$$p(z', t) - B(t) = -\frac{1}{2} \rho \left| \sum_{J=1}^N \frac{S_J}{2\pi} \frac{1}{z' - z_J} + W_\infty \right|^2 + \rho \Re \left(\sum_{J=1}^N \frac{S_J}{2\pi} \frac{1}{z' - z_J} \overline{w_{-J}(t)} \right) \quad (2.9)$$

From this last equation, a singularity induces a *direct* and an *indirect* pressure field. First, a singularity directly induces velocity at the evaluation points, accounted in the quadratic term $-1/2\rho\|\mathbf{v}\|^2$. Second, a singularity is also responsible for an indirect contribution through triadic interactions with the other singularities and the target evaluation points. These indirect contributions are accounted in the unsteady term $-\rho\partial\phi/\partial t$. Physically, one singular point induces velocity on a second, which in turn induces pressure at the target location. Again, the interactions only have an algebraic decay, responsible for the strong interconnections between the singularities. A pressure observation nonlinearly encodes information about all the singularities. Thus, estimating the characteristics of the singularities from limited pressure observations present many challenges.

2.3 Flow modeling over an airfoil with an aggregated vortex model

This section is largely inspired by Le Provost et al. [3].

The two previous sections focus on the potential flow created by a collection of point singularities without body. To account for the presence of an additional body, we recall that the Poisson equations verified by the velocity potential ϕ and the streamfunction ψ are linear. Thus, the potential flow created by a collection of point singularities, a freestream, and a body can be constructed by linear superposition of these elements. From the linearity of the Helmholtz decomposition, the induced velocity $\mathbf{v}(\mathbf{r}, t)$ in (2.3) can also be decomposed into the contribution of the point vortices, the contribution of the freestream, and the contribution of the body. Note that in a potential flow, the induced velocity $\mathbf{v}(\mathbf{r}, t)$ only verifies the no-flow-through on the surface of the body. We refer readers to Eldredge [6] for a thorough treatment. This section succinctly describes the aggregated vortex model, which constitutes the dynamical model of our flow estimator, see [29] for more details.




Figure 2.1: Schematic of the two-dimensional infinitely-thin plate.

In this dissertation, we consider the response of a two-dimensional infinitely-thin plate of chord length c at an angle of attack α , translating impulsively from rest at a velocity $U\mathbf{e}_1$ in a fluid of density ρ , see Fig. 2.1. Throughout this study, positions will be reported in chord lengths and time will be measured in convective units $t^* = tU/c$. As in Darakananda et al. [29], the flow is modeled with a collection of N regularized point vortices. To regularize the interactions between nearby point vortices in the Biot-Savart law (2.4), we replace the singular Cauchy kernel $\mathbf{k}(\cdot)$, with the regularized algebraic blob kernel $\mathbf{k}_\varepsilon(\mathbf{r}) = \mathbf{k}_\varepsilon(\mathbf{r}) = (\varepsilon^2/\pi)(\|\mathbf{r}\|^2 + \varepsilon^2)^{-2}$, where ε is called the blob radius. The regularized point vortices are called blobs. We refer readers to [6] for a review of modeling strategies of two dimensional

inviscid flows.

As we have mentioned earlier, the main role of viscosity in moderate to high Reynolds number external flows, is to generate vorticity at the surface of a body. We model this vorticity flux by shedding new vortex elements from the edges of the plate. In particular, we apply the classical Kutta condition at the trailing edge. However, enforcing the Kutta condition at the leading edge at small to moderate angles of attack leads to a non-physical flow. The regularized flow will leave tangentially the leading-edge in the opposite direction to the freestream. Instead, Ramesh et al. [26] found experimentally that the flow at the leading edge can support a finite amount of suction (integrated pressure) before to trigger a flow separation. They defined the *leading edge suction parameter (LESP)* as a nondimensional measure of the amount of suction at the nose of an airfoil. An elegant vortex shedding criterion can be formulated from this definition: vortex elements are shed only if the LESP exceeds a critical value denoted $LESP_c$ [26, 27]. Eldredge [6] showed that the Kutta condition is recovered if $LESP_c$ is set to zero. The circulation of the new vortex element is proportional to the amount by which LESP exceeds $LESP_c$. We apply this shedding criterion at the leading edge. The $LESP_c$ is included with the vortex elements' positions and strengths as part of the system state vector, \mathbf{x} . Following Ramesh et al. [27], the $LESP_c$ is forecast to remain constant. Ramesh et al. [28] and Darakananda et al. [29] have shown that $LESP_c$ has a strong authority over the vortex dynamics near the leading edge. The true vortex dynamics are encoded in the pressure measurements of the plate, from which we can distill an estimate of the $LESP_c$ [7, 29]. In other words, though $LESP_c$ is predicted to remain constant, the assimilated pressure measurements will tend to cause it to vary, thereby triggering the release of weaker or stronger vorticity.

We use the same vortex element aggregation scheme developed by Darakananda et al. [29]. By aggregating vortex elements at every time step, the overall population remains modest (smaller than 60 in all cases explored) and the aggregated elements stay relatively farther from the plate than without such treatment, dramatically reducing the occurrence of spurious

pressure disturbances.

In our flow estimator, the state variable \mathbf{x}_k contains the positions and circulations of N blobs and the critical leading-edge suction parameter LESP_c ,

$$\mathbf{x}_k = \left[x_k^1 \quad y_k^1 \quad \Gamma_k^1 \quad \dots \quad x_k^N \quad y_k^N \quad \Gamma_k^N \quad \text{LESP}_{c_k} \right]^\top. \quad (2.10)$$

Thus, the state vector dimension is $n = 3N + 1$. The dynamical model applies the following operations:

1. Enforce the no-flow-through condition on the plate,
2. Introduce new vortices according to the Kutta condition at the trailing edge and the current estimate of LESP_c at the leading edge,
3. Advect vortices and plate,
4. Aggregate vortex elements, and nullify the strength of the source blobs.

More details on the dynamical model can be found in [29].

CHAPTER 3

Background on the ensemble Kalman filter

This chapter is adapted from [3].

In this section, we describe the filtering methodology that underpins our data-assimilated vortex model. We start by reviewing the basic filtering problem and the purpose of addressing this problem with an ensemble approach. We then present two algorithms of the ensemble Kalman filter (EnKF): the sEnKF as well as a deterministic variant called the ensemble transform Kalman filter (ETKF) (Bishop et al. [31]). Both will be evaluated in the vortex model applications that follow in Chapter 4. Finally, we review two classical regularization techniques of the EnKF: the *covariance inflation* and the *localization*.

3.1 A review of the filtering problem

In this section, we present a basic outline of the general filtering problem and its connection with Bayesian inference; details on this can be found in several references, including Asch et al. [13], Bishop et al. [31], Carrassi et al. [36], Evensen [37] and Vetra-Carvalho et al. [38]. Though the variants of the ensemble Kalman filter can be explained without this background, it is useful to discuss it in order to justify the use of the ensemble Kalman filter in a filtering problem and to identify some of its limitations. In this dissertation, we use the following font conventions. Serif fonts refer to random variables, e.g. \mathbf{Q} on \mathbb{R}^n , or Q on \mathbb{R} . Lowercase roman fonts refer to realization of random variables, e.g. \mathbf{q} on \mathbb{R}^n , or q on \mathbb{R} . $\pi_{\mathbf{Q}}$ denotes the probability density function for the random variable \mathbf{Q} , and $\mathbf{q} \sim \pi_{\mathbf{Q}}$

means that \mathbf{q} is a realization of \mathbf{Q} . $\pi_{\mathbf{Q}|\mathbf{R}=\mathbf{r}} := \pi_{\mathbf{Q}|\mathbf{R}}(\cdot|\mathbf{r})$, is the probability density function for the random variable \mathbf{Q} knowing that random variable \mathbf{R} takes the value \mathbf{r} . It is called the *conditional probability* of \mathbf{Q} given \mathbf{r} . $\mu_{\mathbf{Q}}, \Sigma_{\mathbf{Q}}$ denote the mean, the covariance matrix of the random variable \mathbf{Q} , respectively. $\Sigma_{\mathbf{Q},\mathbf{R}}$ denotes the covariance matrix of the random variables \mathbf{Q} and \mathbf{R} . Empirical quantities are differentiated by the carets.

For the sake of generality, we frame our discussion on a generic discrete nonlinear state-space model, but we will connect the reader throughout to the present case of a low-order vortex model. Formally, the model is described by the pair of random variables $(\mathbf{X}_k, \mathbf{Y}_k)$ for $k \geq 1$, where \mathbf{X}_k is the state variable of a Markov process in \mathbb{R}^n and $\mathbf{Y}_k \in \mathbb{R}^d$ is an observation of the state \mathbf{X}_k , assumed to be conditionally independent of the state (detailed later in this section). In our case, the state variable is composed of the positions and strengths of the entire set of point vortices, and potentially of the LESPC, while the observation vector is given by the vector of pressure jump coefficients along the plate. The dynamics of the state \mathbf{X}_k is described by the probability distribution for the initial condition $\pi_{\mathbf{x}_0}$ and a propagation equation, also called the dynamical equation:

$$\mathbf{X}_k = \mathbf{f}_k(\mathbf{X}_{k-1}) + \mathbf{W}_k, \quad (3.1)$$

where $\mathbf{f}_k : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the forward operator and \mathbf{W}_k is an additive process noise. In the present case, \mathbf{f}_k is the time-discretized vortex model. Propagating the dynamical equation for the state variable \mathbf{X}_{k-1} is equivalent to sampling from the transition distribution $\pi_{\mathbf{x}_k|\mathbf{x}_{k-1}}(\cdot|\mathbf{x}_{k-1})$ where \mathbf{x}_{k-1} is one realization of the random variable \mathbf{X}_{k-1} .

In general, the state \mathbf{X}_k is only indirectly observed in a noisy and nonlinear fashion through \mathbf{Y}_k :

$$\mathbf{Y}_k = \mathbf{h}_k(\mathbf{X}_k) + \mathcal{E}_k, \quad (3.2)$$

where $\mathbf{h}_k : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is the *observation operator* and \mathcal{E}_k is an additive *measurement noise*. In our case, \mathbf{h}_k applies the unsteady Bernoulli equation to the state vector to obtain a vector of pressure differences, at discrete locations, between the upper and lower surfaces of the

plate. The gradient of the observation operator \mathbf{h}_k evaluated at the current state \mathbf{x}_k is called the *tangent linear* of the observation operator, denoted \mathbf{H}_k , i.e., $\mathbf{H}_k = \nabla \mathbf{h}_k$ [13]. It is important to note that we do not restrict ourselves in this section to Gaussian distributions for the initial condition, process noise, or measurement noise. Furthermore, the forward and observation operators can be time-dependent. In general, the dynamical and observation models could include forcing terms. In our case, the system is not subject to known forcing terms, and we have omitted these terms for brevity.

In the filtering problem, we leverage the realizations of the observation process $(\mathbf{y})_{1:k}$ to infer the realization of the state variable \mathbf{x} at time step k . Our main objective is to estimate the posterior distribution of this state—providing us with complete information about the state and its uncertainty—given all of the observations made thus far, $\pi_{\mathbf{x}_k | \mathbf{y}_{1:k}}(\mathbf{x} | \mathbf{y}_{1:k}) = \pi_{\mathbf{x}_k | \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k}(\mathbf{x} | \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k)$. In the present case, we seek to estimate the positions and strengths of the vortex elements and the LESPC based on all the pressure observations that are available up to this time. Let us consider the estimation of this posterior distribution from the most recent observation, \mathbf{y}_k , which can be computed from Bayes’ theorem [38, 36, 13]:

$$\pi_{\mathbf{x}_k | \mathbf{y}_{1:k}}(\mathbf{x} | \mathbf{y}_{1:k}) = \frac{\pi_{\mathbf{y}_k | \mathbf{x}_k}(\mathbf{y}_k | \mathbf{x}) \pi_{\mathbf{x}_k | \mathbf{y}_{1:k-1}}(\mathbf{x} | \mathbf{y}_{1:k-1})}{\pi_{\mathbf{y}_k}(\mathbf{y}_k)}, \quad (3.3)$$

where $\pi_{\mathbf{y}_k | \mathbf{x}_k}(\mathbf{y}_k | \mathbf{x})$ is the likelihood distribution, i.e., the probability of the observation \mathbf{y}_k if we knew the state \mathbf{x} ; $\pi_{\mathbf{x}_k | \mathbf{y}_{1:k-1}}(\mathbf{x} | \mathbf{y}_{1:k-1})$ is the prior distribution, our estimate of the state before knowledge of the new observation \mathbf{y}_k ; and $\pi_{\mathbf{y}_k}(\mathbf{y}_k)$ is the distribution of the observation. In our case, Bayes’ theorem tells us how the probability densities of the positions and strengths of the vortex elements and the LESPC value get updated by the assimilation of a new pressure observation. Note that, in the likelihood, we have assumed that the observation errors are independent in time, so it is only conditioned on the current state and not on past observations. Indeed, the Bernoulli equation (2.8) only depends on the current collection of vortex elements and LESPC value, and not on its past values. Since

$\pi_{\mathbf{Y}_k}(\mathbf{y}_k)$ can be thought of as a normalizing constant, we can rewrite (3.3) as:

$$\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k}}(\mathbf{x} | \mathbf{y}_{1:k}) \propto \pi_{\mathbf{Y}_k | \mathbf{x}_k}(\mathbf{y}_k | \mathbf{x}) \pi_{\mathbf{x}_k | \mathbf{Y}_{1:k-1}}(\mathbf{x} | \mathbf{y}_{1:k-1}). \quad (3.4)$$

This equation provides us with a means of assimilating a new observation, \mathbf{y}_k , into our forecast probability distribution for the state [39]. The forecast itself comes from our dynamical model. Under the Markov assumption for the state dynamics, the distribution of the state at time step k depends only on the state at the previous time step $k - 1$, i.e., $\pi_{\mathbf{x}_k | \mathbf{x}_{0:k-1}} = \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}}$. The Markov assumption is appropriate in this problem as the time marching of the discrete vortex model only depends on the distribution of vortex elements and the LESPC value at the previous time step. The joint distribution $\pi_{\mathbf{x}_{k-1:k}}$ can thus be factorized into this state transition and the posterior distribution at the end of the previous step:

$$\pi_{\mathbf{x}_{k-1:k} | \mathbf{Y}_{1:k-1}} = \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}} \pi_{\mathbf{x}_{k-1} | \mathbf{Y}_{1:k-1}}. \quad (3.5)$$

By recursively applying this equation with (3.4) substituted for the posterior distribution, it is easy to show that

$$\pi_{\mathbf{x}_{0:k} | \mathbf{Y}_{1:k}} \propto \pi_{\mathbf{x}_0} \prod_{i=1}^k \pi_{\mathbf{Y}_i | \mathbf{x}_i} \pi_{\mathbf{x}_i | \mathbf{x}_{i-1}}. \quad (3.6)$$

This equation is central in Bayesian inference since it justifies the use of sequential methods to estimate the posterior distribution [39, 36]. We can update our previous estimate with new observations sequentially without having to restart the calculation at every time step.

The prior distribution $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k-1}}$ can be obtained by marginalizing (3.5) over all the possible realizations of the state at time $k - 1$:

$$\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k-1}} = \int_{\mathbf{x}_k} \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}} \pi_{\mathbf{x}_{k-1} | \mathbf{Y}_{1:k-1}} d\mathbf{x}_{k-1}. \quad (3.7)$$

This equation is called the Chapman–Kolmogorov equation and corresponds to a direct integration of the state transition kernel [36, 39].

Thus, Bayes’ theorem (3.4) coupled with equation (3.7) provides us with an elegant update rule for incorporating new measurements into the probability distribution, and thus, improving our estimate of the mean and uncertainty of the state. It requires a time-discretized

dynamical model (3.1) for the state transition, $\pi_{\mathbf{x}_k|\mathbf{x}_{k-1}}$, and an observational model (3.2) for the likelihood $\pi_{\mathbf{y}_k|\mathbf{x}_k}$.

However, there are two primary challenges with applying this in the contexts of interest here. First, we do not generally know the forms of distributions involved in any of these formulas. In what follows in this paper, we will make the typical assumption that all errors are drawn from Gaussian distributions with zero mean, but there are good reasons to doubt that this is reasonable in the present nonlinear context of a vortex model. Second, for high-dimensional problems, the integration of the Chapman–Kolmogorov equation (3.7) is intractable. Indeed, here, as in many physics contexts, the forecast step corresponds to the time advancement of a partial differential equation (the Euler equations), so it is infeasible that we could advance such an equation over all possible values of the state. In our case, this corresponds to one time step of the discrete vortex model for all the possible values of the strength and position of the vortex elements and the LESPc. Even if our discrete vortex model is lower dimensional than the discretized Navier–Stokes system, to be sure, it still contains tens to hundreds of degrees of freedom.

Both of these challenges motivate our use of ensemble filtering methods in this work to sequentially estimate the posterior distribution. We use a set of M *particles* $\{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M\}$ sampled from the state distribution $\pi_{\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1}}$ and we aim to construct a particle approximation of the posterior distribution $\pi_{\mathbf{x}_k|\mathbf{y}_{1:k}}$. It should be stressed that the term “particle” is used here in its usual sense in stochastic estimation as a member of the ensemble; it does not denote a vortex particle, a term that we avoid in this paper in favor of “vortex element” or “vortex blob”. In the present case, a particle \mathbf{x}^i is a sample of the distribution of interest for the positions and strengths of the vortex elements and the LESPc.

To construct this particle approximation of the posterior distribution, we perform the following two steps: a forecast step and an analysis step. In the forecast step, each particle \mathbf{x}^i is propagated through the dynamical equation (3.1) (the vortex model) and randomly perturbed with noise (called additive covariance inflation, discussed below) to form samples

from the prior distribution $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k-1}=\mathbf{y}_{1:k-1}}$. This ensemble forecast step constitutes a Monte Carlo approximation of the Chapman–Kolmogorov equation (3.7). Given this sampling $\{\mathbf{x}^i\}$ from the prior distribution, one can easily create a sampling $\{\mathbf{y}^i\}$ from the likelihood distribution $\pi_{\mathbf{y}_k | \mathbf{x}_k}$ by evaluating the observation equation (3.2) (the Bernoulli equation) for each ensemble member $\mathbf{x}^i, i = 1, \dots, M$.

The analysis step then updates the set of ensemble members by assimilating the new realization \mathbf{y}_k^* of the observation variable \mathbf{Y}_k ; in this paper, \mathbf{y}_k^* represents a new pressure measured from the truth system, a high-fidelity Navier–Stokes simulation. With the probability distributions unknown, this would require that we estimate the posterior distribution via Bayes’ theorem (3.4) given the finite set of samples $\{\mathbf{x}^i, \mathbf{y}^i\}$ from the joint distribution formed by the prior and likelihood. We note that the different ensemble filtering methods have a common forecast step but differ in the analysis step. The analysis step updates M samples $\{\mathbf{x}^i\}$ of the prior density with the new observation of the truth system \mathbf{y}_k^* to form samples of the posterior density. Our treatment of the analysis step is built on the idea that there is an underlying transformation \mathbf{T} that directly maps samples from the prior density $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k}=\mathbf{y}_{1:k-1}}$ to the posterior density $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k}=\mathbf{y}_{1:k}}$ [40, 39]. For a Gaussian-linear state-space model, Kalman [15] derived in closed form an exact linear prior-to-posterior transformation for the analysis step. Unfortunately, there is no closed-form otherwise. To bypass this issue, ensemble filtering methods estimate the analysis map \mathbf{T} from samples. For instance, Evensen introduced the *ensemble Kalman filter (EnKF)* that approximates the analysis map found by the Kalman filter from samples [14]. However, the task is made considerably simpler if we assume that the distributions are Gaussian. Under that assumption, and with knowledge of the prior mean and covariance matrix, then Bayes’ theorem leads naturally to the analysis step of the classical Kalman filter [15]. In the particle approximation, in which we estimate this mean and covariance from the finite ensemble statistics that emerge from the forecast, we obtain the ensemble Kalman filter [13]. It should be noted, however, that since the underlying distributions are likely non-Gaussian, *our EnKF framework is not expected*

to converge in the limit of large ensemble size to the true Bayesian solution (Mandel et al. [41]). Indeed, the EnKF can only produce a Gaussian approximation of the posterior density. In the following sections we will present two forms of this method, but first we present some important notation.

3.2 Additional notation

We will use the superscript f to denote prior, forecast quantities and the superscript a to denote analysis, posterior quantities. In particular, the exact prior mean is

$$\mathbf{x}_k^f = \mathbb{E} [\mathbf{X}_k | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}] = \mathbb{E} [\mathbf{X}_k | \mathbf{y}_{1:k-1}], \quad (3.8)$$

and exact prior covariance is $\mathbf{P}_k^f = \Sigma_{\mathbf{x}_k | \mathbf{y}_{1:k-1}}$. The posterior mean and covariance are given by $\mathbf{x}_k^a = \mathbb{E} [\mathbf{X}_k | \mathbf{y}_{1:k}]$ and $\mathbf{P}_k^a = \Sigma_{\mathbf{x}_k | \mathbf{y}_{1:k}}$, respectively. The notations used in this section are similar to those of Asch et al. [13] and Raanes [42].

Given an ensemble (\mathbf{x}^i) of size M , we define the ensemble matrix $\mathbf{X} \in \mathbb{R}^{n \times M}$ as [13, 42]:

$$\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M]. \quad (3.9)$$

We use an overbar to denote statistics obtained from the ensemble, such as the sample mean and covariance:

$$\bar{\mathbf{x}} = \frac{1}{M} \sum_{i=1}^M \mathbf{x}^i, \quad \bar{\mathbf{P}} = \frac{1}{M-1} \sum_{i=1}^M (\mathbf{x}^i - \bar{\mathbf{x}}) (\mathbf{x}^i - \bar{\mathbf{x}})^\top, \quad (3.10)$$

where \top denotes the transpose operator. We define the anomaly matrix $\mathbf{X}' \in \mathbb{R}^{n \times M}$ of an ensemble as

$$\mathbf{X}' = \frac{1}{\sqrt{M-1}} [\mathbf{x}^1 - \bar{\mathbf{x}}, \mathbf{x}^2 - \bar{\mathbf{x}}, \dots, \mathbf{x}^M - \bar{\mathbf{x}}]. \quad (3.11)$$

The anomaly matrix obviously has zero mean. Using $\mathbf{1}$ to denote the vector of ones of length M , we can conveniently define the sample mean, anomaly matrix, and sample covariance from the ensemble matrix:

$$\bar{\mathbf{x}} = \frac{1}{M} \mathbf{X} \mathbf{1}, \quad \mathbf{X}' = \frac{1}{\sqrt{M-1}} \mathbf{X} (\mathbf{I} - \mathbf{1} \mathbf{1}^\top / M), \quad \bar{\mathbf{P}} = \mathbf{X}' \mathbf{X}'^\top. \quad (3.12)$$

Finally, we have the relation between the ensemble matrix and the anomaly matrix:

$$\mathbf{X} = \overline{\mathbf{X}} + \sqrt{M-1}\mathbf{X}' \quad (3.13)$$

where $\overline{\mathbf{X}} = \overline{\mathbf{x}}\mathbf{1}^\top = [\overline{\mathbf{x}}, \overline{\mathbf{x}}, \dots, \overline{\mathbf{x}}] \in \mathbb{R}^{n \times M}$. It is easy to verify that $\mathbf{1}^\top \mathbf{1} = M$, $\mathbf{X}'\mathbf{1} = 0$, and $\overline{\mathbf{X}}\mathbf{1}/M = \overline{\mathbf{x}}$.

3.3 The stochastic ensemble Kalman filter

The application of the stochastic ensemble Kalman filter for aerodynamic flow estimation was already presented in [16, 29], but we review it here in order to identify an important deficiency that motivates our current treatment. The forecast step of the sEnKF—and, indeed, of all forms of the EnKF—is given by applying (3.1) (the vortex model) to each ensemble member [37]:

$$\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_k^i \quad \text{for } i = 1, \dots, M, \quad (3.14)$$

where the random process noise \mathbf{w}_k^i is drawn from the distribution $\mathbf{w}_k^i \sim \pi_{\mathbf{w}_k}$. It should be noted that this process noise is not inherently part of the dynamical model and must be explicitly introduced; this is discussed in Section 3.5.

In order to simplify notation, we drop the time dependence subscript of the variables in the rest of this section, since the analysis step does not involve time propagation. The analysis step of the sEnKF seeks a linear transformation of the form:

$$\mathbf{x}^{a,i} = \mathbf{x}^{f,i} + \mathbf{K}(\mathbf{y}^* - (\mathbf{h}(\mathbf{x}^{f,i}) + \boldsymbol{\epsilon}^i)), \quad \text{for } i = 1, \dots, M, \quad (3.15)$$

where $\mathbf{x}^{f,i}$ and $\mathbf{x}^{a,i}$ are the i th prior and posterior ensemble member, respectively; $\boldsymbol{\epsilon}^i$ is drawn from the measurement noise distribution $\pi_{\boldsymbol{\epsilon}}$; and \mathbf{y}^* is the realization of the observation variable \mathbf{Y} at the current assimilation time. The expression $\mathbf{y}^* - (\mathbf{h}(\mathbf{x}^{f,i}) + \boldsymbol{\epsilon}^i)$ is called the innovation for the i th ensemble member. In our case, (3.15) shows how discrepancies

in the pressure observations are mapped via the Kalman gain \mathbf{K} into linear updates of the positions and strengths of the vortex elements and the LESPc for each ensemble member.

The algorithm of the sEnKF is provided in Algorithm 1. The original analysis step derived by Evensen [14] did not include the measurement noise term, but this was corrected by Burgers et al. [43] to reflect that \mathbf{y}^* is drawn from a random distribution and to avoid some spurious correlations within the ensemble. The matrix $\mathbf{K} \in \mathbb{R}^{n \times d}$ is called the *Kalman gain* and is identical to its form in the standard Kalman filter, derived to minimize the trace of the posterior covariance matrix \mathbf{P}^a (as we will discuss further below):

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^\top (\mathbf{H} \mathbf{P}^f \mathbf{H}^\top + \mathbf{V})^{-1} \quad (3.16)$$

with \mathbf{V} the covariance matrix of the measurement noise and \mathbf{H} the tangent linear of the observation operator. The exact prior covariance \mathbf{P}^f is approximated by the sample prior covariance $\overline{\mathbf{P}}^f$ formed from the prior ensemble $\{\mathbf{x}^{f,i}\}$ with (3.12), $\overline{\mathbf{P}}^f = \mathbf{X}'^f \mathbf{X}'^{f\top}$; similarly, the posterior covariance is approximated by $\overline{\mathbf{P}}^a = \mathbf{X}'^a \mathbf{X}'^{a\top}$.

The Kalman gain suggests that we must calculate the tangent linear of the observation operator. Evensen [14] proposed a technique called implicit linearization that approximates $\mathbf{P}^f \mathbf{H}$ and $\mathbf{H} \mathbf{P}^f \mathbf{H}^\top$ given the prior ensemble. First, we construct $\overline{\mathbf{y}}^f$ the mean of the observations $\{\mathbf{h}(\mathbf{x}^{f,i})\}$ for $i = 1, \dots, M$:

$$\overline{\mathbf{y}}^f = \frac{1}{M} \sum_{i=1}^M \mathbf{h}(\mathbf{x}^{f,i}). \quad (3.17)$$

Let us then define the innovation anomaly matrix \mathbf{Z}'^f , with i th column given by

$$\mathbf{Z}'^{f,i} = \frac{\mathbf{h}(\mathbf{x}^{f,i}) - \overline{\mathbf{y}}^f - \boldsymbol{\epsilon}^i + \bar{\boldsymbol{\epsilon}}}{\sqrt{M-1}}, \text{ for } i = 1, \dots, M, \quad (3.18)$$

and with $\bar{\boldsymbol{\epsilon}}$ the sample mean of $\{\boldsymbol{\epsilon}^i\}$. It should be noted that this mean is itself a random number whose expected value is 0.

The key ingredient of the technique is to make the following approximation [13]:

$$\mathbf{H}(\mathbf{x}^{f,i} - \overline{\mathbf{x}}^f) \simeq \mathbf{h}(\mathbf{x}^{f,i}) - \overline{\mathbf{y}}^f. \quad (3.19)$$

Using this approximation, it is easy to show that $\mathbf{P}^f \mathbf{H}^\top$ and $\mathbf{H} \mathbf{P}^f \mathbf{H}^\top$ are approximated by $\mathbf{X}'^f \mathbf{Z}'^{f\top}$ and $\mathbf{Z}'^f \mathbf{Z}'^{f\top} - \mathbf{V}$, respectively [13], and thus, the Kalman gain (3.16) takes the simple form

$$\mathbf{K} = \mathbf{X}'^f \mathbf{Z}'^{f\top} \left(\mathbf{Z}'^f \mathbf{Z}'^{f\top} \right)^{-1}. \quad (3.20)$$

To gain more insights on this formulation of the filter, we can rewrite the analysis update (3.15) with ensemble matrix notations. From (3.13), the update of each posterior ensemble member is equivalent to a separate update of the posterior mean $\bar{\mathbf{x}}^a$ and the posterior anomaly matrix \mathbf{X}'^a . From the linear analysis step (3.15), the posterior anomaly matrix \mathbf{X}'^a is updated according to ([13]):

$$\mathbf{X}'^a = \mathbf{X}'^f - \mathbf{K} \mathbf{Z}'^f; \quad (3.21)$$

the update equation for the posterior mean state $\bar{\mathbf{x}}^a$ is obtained by taking the expectation of the analysis step (3.15):

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{K}(\mathbf{y}^* - \bar{\mathbf{y}}^f), \quad (3.22)$$

where $\bar{\mathbf{y}}^f$ is defined in (3.17). Equation (3.22) is the exact Kalman update equation.

Our derivations thus far have omitted the fact that our actual ensemble is finite sized. Because of the assumed linear form of the analysis step, equation (3.22) is reproduced almost exactly for finite-sized ensembles, except for the addition of a lingering sample mean of the measurement noise, $\bar{\boldsymbol{\epsilon}}$, in parentheses. However, a more problematic discrepancy lies in the relationship between the posterior and prior covariance matrices. For obtaining this relationship, let us first denote by \mathbf{W}' the anomaly matrix of the measurement noise $\mathbf{W}'^{i,j} = (\boldsymbol{\epsilon}^i - \bar{\boldsymbol{\epsilon}}) / \sqrt{M-1}$; the product $\mathbf{W}' \mathbf{W}'^\top$ approximates the measurement covariance \mathbf{V} . Then, from the update equation (3.21) and the definitions of the covariances and innovation anomaly matrix, we can construct the equation for the posterior covariance matrix [13]:

$$\begin{aligned} \bar{\mathbf{P}}^a &= (\mathbf{I}_n - \mathbf{K} \mathbf{H}) \bar{\mathbf{P}}^f (\mathbf{I}_n - \mathbf{K} \mathbf{H})^\top + \mathbf{K} \mathbf{W}' \mathbf{W}'^\top \mathbf{K}^\top \\ &+ (\mathbf{I}_n - \mathbf{K} \mathbf{H}) \mathbf{X}'^f \mathbf{W}'^\top \mathbf{K}^\top + \mathbf{K} \mathbf{W}' \mathbf{X}'^{f\top} (\mathbf{I}_n - \mathbf{K} \mathbf{H})^\top. \end{aligned} \quad (3.23)$$

By taking the expectation of this matrix over the measurement noise process and then minimizing its trace over \mathbf{K} , we recover the classical Kalman gain (3.16) and a simple form of the expected value of the posterior covariance:

$$\mathbb{E} [\overline{\mathbf{P}}^a] = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\overline{\mathbf{P}}^f. \quad (3.24)$$

Indeed, the classical Kalman filter adopts this expected value in order to propagate the covariance matrix, $\mathbf{P}^a = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{P}^f$. However, the finite sample covariance $\overline{\mathbf{P}}^a = \mathbf{X}'^a \mathbf{X}'^a{}^\top$ does not reproduce this ideal relation except in the limit of infinite ensemble sizes. Rather, the covariance of the finite ensemble only strictly satisfies (3.23), due to spurious correlations between the forecast anomalies and observation noise, $\mathbf{X}'^f \mathbf{W}'{}^\top$. Thus, the stochastic analysis step of the sEnKF introduces error that degrades the performance of the ensemble Kalman filter [31, 13]. A pseudo-code for the sEnKF is presented in Algorithm 1. An alternative form of the EnKF will be presented in the next section that addresses this issue.

3.4 The ensemble transform Kalman filter

In the previous section, it was shown that finite ensembles do not reproduce the expected value of the posterior covariance matrix $\overline{\mathbf{P}}^a$, and that this can cause the performance of the sEnKF to degrade. To circumvent the issue, Bishop et al. [31] developed the *ensemble transform Kalman filter (ETKF)* that exactly reproduces the ideal propagation equation for the covariance. This has been shown to give better performance in other applications [42, 13]. The ETKF belongs to a more general class of ensemble Kalman filters—called *deterministic ensemble Kalman filters*—that verify exactly the correct propagation equation through various analysis schemes.

To generate a posterior anomaly matrix \mathbf{X}'^a , one can start from the desired propagation

equation for the covariance and use the factorization of the sample covariance matrix (3.12):

$$\bar{\mathbf{P}}^a = \mathbf{X}'^a \mathbf{X}'^{a\top} = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\bar{\mathbf{P}}^f \quad (3.25)$$

The right-hand side of this equation can also be factorized when $\bar{\mathbf{P}}^f$ is replaced with its own factorization and the ensemble expression (3.20) of the Kalman gain is introduced. In the following, it should be noted the innovation anomaly matrix \mathbf{Z}'^f is defined as in (3.18), but now without the observation noise, so that $\mathbf{Z}'^f = \mathbf{H}\mathbf{X}'^f$:

$$\begin{aligned} \mathbf{X}'^a \mathbf{X}'^{a\top} &= (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{X}'^f \mathbf{X}'^{f\top} \\ &= (\mathbf{I}_n - \mathbf{X}'^f \mathbf{Z}'^{f\top} (\mathbf{Z}'^f \mathbf{Z}'^{f\top} + \mathbf{V})^{-1} \mathbf{H})\mathbf{X}'^f \mathbf{X}'^{f\top} \\ &= \mathbf{X}'^f (\mathbf{I}_M - \mathbf{Z}'^{f\top} (\mathbf{Z}'^f \mathbf{Z}'^{f\top} + \mathbf{V})^{-1} \mathbf{Z}'^f) \mathbf{X}'^{f\top} \\ &= \mathbf{X}'^f \mathbf{G} \mathbf{X}'^{f\top}. \end{aligned} \quad (3.26)$$

Matrix $\mathbf{G} \in \mathbb{R}^{M \times M}$ is symmetric positive definite. In order to produce an analysis equation for the posterior anomaly matrix, we seek a square-root factorization of \mathbf{G} [13]. In other words, we look for a matrix $\mathbf{G}^{1/2}$ such that $\mathbf{G} = \mathbf{G}^{1/2} \mathbf{G}^{1/2\top} = \mathbf{G}^{1/2\top} \mathbf{G}^{1/2}$. For a positive-definite matrix, the square-root decomposition exists but is not unique. Indeed for an arbitrary orthogonal matrix $\mathbf{U} \in \mathbf{O}(M)$ —i.e., any matrix such that $\mathbf{U}\mathbf{U}^\top = \mathbf{U}^\top \mathbf{U} = \mathbf{I}_M$ —then $\mathbf{G}^{1/2}\mathbf{U}$ is also a square-root of \mathbf{G} [13, 42].

Therefore, the analysis update of the anomaly can be written as a right linear transformation [13]:

$$\mathbf{X}'^a = \mathbf{X}'^f \mathbf{G}^{1/2} \mathbf{U} \quad (3.27)$$

with some choice of $\mathbf{U} \in \mathbf{O}(M)$, discussed below. Furthermore, the analysis update of the sample mean (3.21) can be written with the help of (3.20) as

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{X}'^f \mathbf{Z}'^{f\top} (\mathbf{Z}'^f \mathbf{Z}'^{f\top} + \mathbf{V})^{-1} \boldsymbol{\delta}, \quad (3.28)$$

where $\boldsymbol{\delta} = \mathbf{y}^* - \mathbf{h}(\bar{\mathbf{x}}^f)$ is the mean innovation. Equations (3.27) and (3.28) are used together, through (3.13), to update the ensemble.

It is notable that this analysis step assembles both the posterior anomaly and the update to the mean from the columns of \mathbf{X}'^f . In other words, the analysis is said to be performed in the *ensemble space*: the update to the state vector is a linear combination of the deviations among the forecast ensemble members from the mean. In our case, for example, the position of a certain vortex element will be updated with a linear combination of the ensemble deviations of that element’s position from the ensemble mean, after the ensemble has been advanced by the vortex model. This emphasizes the importance of such variance among ensemble members: without it, any discrepancy $\boldsymbol{\delta}$ between the new observation and the predicted observation is simply ignored.

The form of this update is an attractive property of the ETKF since the ensemble size M —which is on the order of 100—is typically similar to the size of the state n encountered in the vortex model, and certainly much smaller than the state vector associated with a CFD simulation [31, 13, 38]. Using Sherman–Morrison–Woodbury identities, the formula for $\mathbf{G}^{1/2}$ can be simplified [38]:

$$\mathbf{G}^{1/2} = (\mathbf{I}_M - \mathbf{Z}'^f{}^\top (\mathbf{Z}'^f \mathbf{Z}'^f{}^\top + \mathbf{V})^{-1} \mathbf{Z}'^f)^{1/2} = (\mathbf{I}_M + \mathbf{Z}'^f{}^\top \mathbf{V}^{-1} \mathbf{Z}'^f)^{-1/2}. \quad (3.29)$$

In this work, we choose the symmetric square-root factorization, $\mathbf{G} = \mathbf{G}^{1/2} \mathbf{G}^{1/2}$, among the different possibilities. Given an eigendecomposition $\mathbf{G} = \mathbf{R} \boldsymbol{\Sigma} \mathbf{R}^\top$, the symmetric square root $\mathbf{G}^{1/2}$ is given by $\mathbf{R} \boldsymbol{\Sigma}^{1/2} \mathbf{R}^\top$, with $\boldsymbol{\Sigma}^{1/2}$ the entry-wise positive square root of the diagonal matrix $\boldsymbol{\Sigma}$. Several studies [44, 37, 42] have shown that the symmetric root has useful properties, particularly that it does not introduce a spurious mean in the posterior anomaly matrix: $\mathbf{G} \mathbf{1} = \mathbf{1}$, so $\mathbf{1}$ is an eigenvector of \mathbf{G} with unit eigenvalue; thus, this is also true of $\mathbf{G}^{1/2}$, so $\mathbf{X}'^f \mathbf{G}^{1/2} \mathbf{1} = \mathbf{X}'^f \mathbf{1} = \mathbf{0}$. The overall construction of the size M square-root $\mathbf{G}^{1/2}$ is computationally inexpensive.

Sakov and Oke [45] have found that the choice of the symmetric square root leads consistently to better performance. However, for large ensemble, the symmetric square root can lead to the creation of outliers in the ensemble. One can prevent the appearance of these

spurious members by right multiplying $\mathbf{G}^{1/2}$ by a random orthogonal matrix \mathbf{U} from time to time. To avoid bias in the posterior anomaly matrix, we must also ensure that $\mathbf{X}^{/a}$ has zero mean. To do so, \mathbf{U} must preserve the mean, equivalent to requiring that $\mathbf{U}\mathbf{1} = \mathbf{1}$. Some authors [46, 47] have proposed a scheme based on sampling of the standard normal distribution and use of Householder reflections to construct such mean-preserving random rotations; this algorithm is presented in Algorithm 2. For the other assimilation steps, \mathbf{U} is simply the identity \mathbf{I}_M . A pseudo-code for the ETKF is presented in Algorithm 3. For further reading on the differences between the sEnKF and the ETKF, we refer readers to Katzfuss et al. [48].

3.5 Classical regularization techniques: covariance inflation and localization

In the different algorithms of the EnKF, we estimate the forecast covariance matrix $\Sigma_{\mathbf{x}_k | \mathbf{y}_{1:k-1}}$ and the covariance of the state and observation densities $\Sigma_{\mathbf{x}_k | \mathbf{y}_{1:k-1}, \mathbf{y}_k}$ from M prior samples $\{\mathbf{x}^{f,i}\}$, where the ensemble size $M \sim 100$ is typically smaller by several order of magnitudes than the dimensions of the state and observation spaces. Thus, the estimated Kalman gain suffers from rank-deficiency, sampling errors, and spurious long-range correlations. Another consequence of the limited ensemble size is the underestimation of the posterior covariance [42]. While the EnKF can be successfully applied in high-dimensional problems, its success is dependent to an adequate regularization of the EnKF. The two classical regularization techniques of the EnKF are the *covariance inflation* and the *localization*.

Over multiple assimilation cycles, the under-estimation of the covariance matrices drive the entries of the empirical Kalman gain to zero, leading to the filter divergence. Covariance inflation aims to increase the forecast state covariance and improve the conditioning of the estimated Kalman gain. Covariance inflation is typically applied after the forecast step and before the analysis step. A physical interpretation of the covariance inflation in the context

of our vortex modeling is presented in Chapter 4. The *multiplicative inflation* increases the spread of the ensemble about the sample mean, by rescaling the deviation $\mathbf{x}^{f,i} - \bar{\mathbf{x}}^f$ by the multiplicative factor $\beta > 1$ for each ensemble member:

$$\mathbf{x}^{f,i} \leftarrow \bar{\mathbf{x}}^f + \beta(\mathbf{x}^{f,i} - \bar{\mathbf{x}}^f), \text{ for } i = 1, \dots, M. \quad (3.30)$$

This form of inflation is equivalent to multiplying the sample prior covariance $\bar{\mathbf{P}}^f$ by β^2 . We recommend setting the value of β to 1.01 and 1.03 for the sEnKF and ETKF, respectively.

Additive inflation adds to each ensemble member a sample from a Gaussian distribution with zero mean and covariance \mathbf{Q} . Equivalently, additive inflation acts as a Tikhonov regularization of the prior covariance $\bar{\mathbf{P}}^f$, by adding the covariance matrix \mathbf{Q} . Thus, it enforces a lower bound on the spread of the prior covariance matrix and improves the conditioning of the analysis step. Additive inflation can also account for model errors in the dynamical model. Note that one can combine additive and multiplicative inflations as they account for different kind of errors [49]. The tuning of the multiplicative inflation only requires one parameter, while additive inflation requires $n(n+1)/2$ parameters for an arbitrary positive definite matrix \mathbf{Q} . In this work, we restrict to isotropy covariance matrices $\alpha \mathbf{I}_n$, where $\alpha > 0$. In general, the tuning of the additive inflation requires expert knowledge of the system.

The limited ensemble size is also responsible for spurious correlations over long distances between the state and observation variables. Here, the notion of “long distance” is intrinsically related to the definition of a metric d for the problem of interest. In our case, d can be the Eulerian distance between the position of a pressure sensor and the location of a point vortex. In many high-dimensional problems, such as atmospheric problems [13], the correlations between the state and observation variables are rapidly decaying with the distance. Distance localization regularizes the EnKF by systematically removing all the long-range correlations. However, in inverse problems of incompressible fluid mechanics, the observation models are given by elliptic partial differential equations (e.g. the pressure Poisson equation) for which there are physical long-range correlations between the state and

observation variables. Distance localization cannot be applied for these problems, as we cannot disentangle the true slowly decaying correlations from the spurious long-range ones. Chapter 5 will introduce a novel regularization of the EnKF for elliptic observation models.

CHAPTER 4

Aerodynamic flow estimation with point vortex models

This chapter is adapted from [3].

With the description of the aggregated vortex model in Chapter 2 and the review of the ensemble Kalman filter in Chapter 3, we are now in position to summarize one time step of the flow estimator developed in this work. A schematic of this method is presented in Fig. 4.1. In the forecast step, we propagate an ensemble of inviscid vortex models through the Biot-Savart law. At the end of this forecast, multiplicative and additive covariance inflation are applied to each member of the ensemble. In the analysis step, we update the positions and strengths of the different vortex models with the pressure observations obtained from the true system. In our case, the true observations will be obtained from high-fidelity simulations of the Navier-Stokes equations, see 4.2 for more details. The assimilation of the observations is either carried with the stochastic EnKF (sEnKF) [14] or the ensemble transform Kalman filter (ETKF) [31]. The rest of this chapter is organized as follows. Section 4.1 provides a physical interpretation of our method. Section 4.2 applies our flow estimator to two challenging scenarios. First, we consider an impulsively translating flat plate at 20° of angle of attack subjected to large-amplitude and overlapping perturbations applied near the leading edge of the airfoil to mimic flow actuation. Second, we look at an impulsively translating flat plate at 20° placed in the von Kármán street of a cylinder placed upstream.

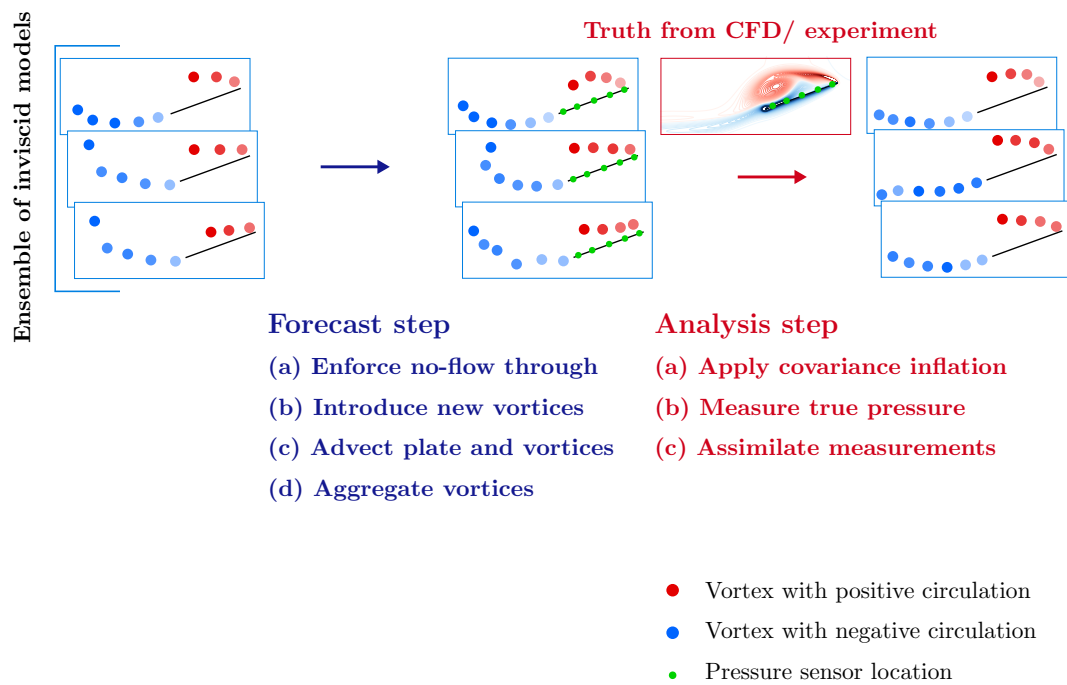


Figure 4.1: One time step of the data-assimilated aggregated vortex model

4.1 Flow estimation with an ensemble of aggregated vortex models

In this section, we interpret physically the role of each step of the filter. First, it is important to understand the effect of the ensemble and its propagation on the vorticity field. To demonstrate this effect in the clearest manner, let us assume that the N vortex elements in each ensemble member are singular: point vortices rather than blobs. Their blob form is only used to regularize the Biot–Savart interactions between them, and our discussion here focuses only on the interpretation of the vorticity field itself.

We can then write the vorticity field at location \mathbf{r} and time step k as a function of the random state vector \mathbf{X}_k , whose components (aside from LESPc) constitute the strengths and positions of the singular elements:

$$\omega(\mathbf{r}, \mathbf{X}_k) = \sum_{J=1}^N \Gamma_k^J \delta(\mathbf{r} - \mathbf{r}_k^J), \quad (4.1)$$

where δ is the Dirac delta function. The expected value of the vorticity field at the end of time step k is given (ideally, for infinite ensemble) in terms of the posterior distribution function

$$\mathbb{E} [\omega(\mathbf{r}, \mathbf{X}_k)] = \int \omega(\mathbf{r}, \mathbf{x}_k) \pi_{\mathbf{X}_k}(\mathbf{x}_k) d\mathbf{x}_k. \quad (4.2)$$

(For simplicity of notation, we have omitted the fact that the distribution $\pi_{\mathbf{X}_k}$ is conditioned on the observations made thus far.) Under our Gaussian assumption, it is particularly easy to evaluate these integrals, and we arrive at

$$\mathbb{E} [\omega(\mathbf{r}, \mathbf{X}_k)] = \sum_{J=1}^N \frac{\bar{\Gamma}_k^J}{2\pi |\mathbf{P}_k^{rJ}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{r} - \bar{\mathbf{r}}_k^J)^\top \mathbf{P}_k^{rJ}{}^{-1} (\mathbf{r} - \bar{\mathbf{r}}_k^J) \right), \quad (4.3)$$

where $|\cdot|$ denotes determinant; $\bar{\Gamma}_k^J$ and $\bar{\mathbf{r}}_k^J$ denote, respectively, the mean circulation and position of vortex element J at time step k ; and \mathbf{P}_k^{rJ} is the 2×2 covariance submatrix associated with the position of vortex element J at step k .

In other words, we can interpret each vortex element’s uncertainty as defining an elliptically-shaped region, centered at its mean location and endowed with its mean strength. The behavior of this elliptical region’s shape over time is determined by the combined influences of three processes in the filter: the forecast, the inflation, and the analysis.

The role of the forecast step is straightforward: it constitutes an inviscid (i.e., advective) advancement of the vortex elements by one step and the creation of new vorticity to satisfy modeled edge conditions. The ensemble of such inviscid models establishes a set of slightly different displacements of each vortex element. Since we interpret this ensemble as approximating a Gaussian distribution both before and after the forecast, the set of displacements of each vortex define a constrained transformation of the vortex’s ellipse: a net advection of the center and a stretching and rotation of its shape.

The inflation step imposes two influences on each region’s shape. The multiplicative inflation stretches the ellipse uniformly by a small fraction in every direction. The additive inflation, on the other hand, comprises a single step of a random walk. It is well known that a random walk, applied over a large number of steps, approaches a Wiener process, and the associated probability distribution satisfies a linear diffusion equation. This diffusion causes the elliptical region to spread, reminiscent of core spreading in vortex methods [50]. To simulate diffusion of viscosity ν , the random step is chosen from a normal distribution with standard deviation $\sqrt{2\nu\Delta t}$.

In the context here, the additive inflation occurs among other processes, so its interpretation is less clear. Chorin’s random vortex method [51], utilizing a large number of overlapping vortex blobs undergoing random walks, is known to converge to the solution of the Navier–Stokes equation as $N^{-1/2} \log N$, where N is the number of blobs [52]. In the EnKF context, it is possible to identify a stochastic differential equation that asymptotically describes the forecast and inflation steps in the limit of large ensemble [53]. We leave a rigorous interpretation of the ensemble of vortex models in this manner for future work. We merely observe that the variance of the distribution from which we select the additive

inflation parameter loosely specifies the viscosity of a diffusion process. However, we do not attempt to match this effective viscosity to the actual viscosity in the truth system; rather, we tune the additive inflation to balance the trust between the forecast and analysis steps.

Finally, in order to reconcile the new observations made in this step—namely, the pressures measured on the surface of the wing—with the pressures predicted via the Bernoulli equation, the vortex elements’ positions and strengths (and LESPC) need to be adjusted. The EnKF analysis step provides this adjustment by assembling a minimum least-squares solution to the problem $\mathbf{h}(\bar{\mathbf{x}}^a) = \mathbf{y}^*$, subject to measurement noise \mathbf{V} ; the minimization is regularized by the forecast $\bar{\mathbf{x}}^f$, with its associated covariance $\bar{\mathbf{P}}^f$, ensuring that the new state vector does not stray far from its forecast. Mathematically, this problem is expressed as [18, 12]:

$$\bar{\mathbf{x}}^a = \operatorname{argmin}_{\mathbf{x} \in \mathbf{R}^n} \frac{1}{2}(\mathbf{y}^* - \mathbf{h}(\mathbf{x}))^\top \mathbf{V}^{-1}(\mathbf{y}^* - \mathbf{h}(\mathbf{x})) + \frac{1}{2\beta}(\mathbf{x} - \bar{\mathbf{x}}^f)^\top \bar{\mathbf{P}}^{f^{-1}}(\mathbf{x} - \bar{\mathbf{x}}^f), \quad (4.4)$$

with β the multiplicative covariance inflation factor. The analysis step causes the covariance to shrink, according to (3.23). The elliptical region associated with each vortex element is translated, stretched, and rotated by the analysis step, but its area shrinks.

4.2 Results

In this section, we present the results of flow estimation for two strongly-perturbed flows about an infinitely thin plate at 20° . In the first case, a sequence of perturbations that mimic pulse actuation is applied near the leading edge of the plate. In the second case, the plate is subject to large scale and coherent perturbations created in the wake of an upstream cylinder. In the discussion of the results, positive and negative vortex elements will refer to vortex elements with positive (counter-clockwise) and negative (clockwise) circulation respectively.

We discretize the vortex model with the forward Euler time scheme with time step $\Delta t^* = 0.01$. The blob radius δ (normalized by c) is set to 5×10^{-3} and 9×10^{-3} for the sEnKF and the ETKF, respectively. The difference between these two blob sizes is of little physical

relevance: each blob radius was chosen concomitantly with the inflation and noise parameters to reduce the time-averaged root mean-square error of the normal force prediction and the appearance of spurious pressures on the surface for the sEnKF and the ETKF. We use the same blob radius for the sEnKF as the previous study by Darakananda et al. [29]. For the aggregated vortex model, the vortex elements are mostly isolated and their interactions require little regularization. The non-zero blob radius primarily regularizes the interactions of the vortex elements soon after their release from the edges. Throughout this study, we use an ensemble of size $M = 50$. At the initial time, no vortices are present in the state vector of the vortex model. The ensemble is initiated with random samples for the LESPC drawn from $\mathcal{N}(0.5, 0.1)$, i.e., a normal distribution with mean 0.5 and covariance 0.1. In this work, we rely on the assimilation of pressure jump coefficients— $\Delta C_p = 2(p^+ - p^-)/\rho U^2$ where $+$ and $-$ denote the upper and lower side of the plate—obtained from the truth system to correct our aggregated vortex model [29]. The role of the assimilation is twofold: to better account for the viscous effects, and to account for the presence of flow disturbances in the true system that are not modeled into the vortex model. The observation operator \mathbf{h}_k in (3.2) uses the unsteady Bernoulli equation to predict the pressure jump coefficients at d locations on the plate [54]. In this work, pressure measurements are calculated (and also provided by the truth system) at the following $d = 50$ sampled Chebyshev points:

$$\frac{c}{2} \cos\left(\frac{10i\pi}{512+1}\right) \quad \text{for } i = 1, \dots, d. \quad (4.5)$$

The true pressure jump measurements are generated from simulations of a flat plate at Reynolds number $Re = 500$, carried out with a high-fidelity Navier–Stokes solver, based on the immersed boundary projection method with lattice Green’s function [17, 30]. It is important to emphasize that the flow perturbations are only present in the truth system, and their effect is only made available to the vortex model through the assimilation of the true pressure measurements.

We use a different set of inflation parameters for the sEnKF and the ETKF:

- A multiplicative inflation $\beta = 1.01, 1.028$ for the sEnKF and ETKF, respectively.
- An additive inflation for the position of the vortices drawn from $\mathcal{N}(0, 1 \times 10^{-5})$ (normalized by c).
- An additive inflation for the strength of the vortices drawn from $\mathcal{N}(0, 1 \times 10^{-3} \Delta t^*)$ (normalized by Uc).
- An additive inflation for the LESPC drawn from $\mathcal{N}(0, 5 \times 10^{-5}), \mathcal{N}(0, 8.5 \times 10^{-5})$ for the sEnKF and ETKF, respectively.
- The measurement noise ϵ_k is drawn from $\mathcal{N}(0, 1 \times 10^{-8})$ (normalized by ρU^2).

These parameters and the blob radius have been manually determined, independently for the sEnKF and the ETKF, to minimize the time-averaged root-mean-squared error of the predicted normal force and the occurrence of spurious pressures on the surface. The sEnKF and ETKF share the same additive inflation parameters for the vortex properties. It is important to note that each of these parameters is chosen in order to balance the trust between the vortex model and the analysis step; lower values, for example, lead to a vortex model that is less responsive to measurement innovation. For example, the relatively larger additive inflation for the LESPC makes this parameter more responsive than the vortex parameters. The same parameters are used for all examples in the next section. These parameters are nondimensionalized by the undisturbed characteristic scales of the problem: the density of the fluid ρ , the chord length c and the translational velocity of the plate U . In our various numerical experiments, we have found that this choice of the parameters work well for perturbations of different strengths and shapes.

It should be noted that the dimension of the state vector, n , changes with each time step of this filtering process: it increases typically by six (three per newly released blob originating from each edge), and occasionally decreases as elements of zero strength are eliminated. There are no inherent restrictions in the EnKF on changes to the dimension of the state.

However, in order to keep a consistent state dimension, i.e., the same number of blobs across the different ensemble members, the vortex elements whose circulations are aggregated into another element are not removed, but rather, are simply assigned zero circulation. Similarly, even if the LESP does not exceed the current estimate of the LESP_c, a new vortex with zero circulation is still introduced. If the blob has zero circulation across all ensemble members, then it is removed from the state vector.

Due to the stochastic nature of the filtering problem, results may vary from one simulation to another for the same filter. We estimate the uncertainty in the results in order to draw consistent conclusions about the performance of each filter. For assessment purposes, the results presented in this work have been obtained by running an ensemble of 100 realizations of the same filter on each flow configuration, with each realization consisting of an application of the EnKF (using an $M = 50$ ensemble of vortex models). From this ensemble of realizations, we construct the sample mean and standard deviation of the quantity of interest for comparison with the truth. If we assume that the results obtained over the different runs follow a Gaussian distribution, the 95% confidence interval can be estimated by considering the plus and minus deviation of twice the standard deviation from the mean. It should be noted that the number of vortex elements—and hence, the dimensionality of the state vector—varies from one realization of the filter to the next, so it is not possible to define a mean state from the ensemble of realizations.

As mentioned in Section 2.3, the near encounters between the singular vortex elements and the plate can lead to spurious errors in the estimate of the pressure distribution, and hence, in the prediction of overall normal force on the plate. We use a median filter to remove these spurious spikes from our presented results. At each time step, the current value is replaced by the median value of the $l = 7$ previous time steps. The median filter is causal, so it is directly integrated into the estimation sequence.

Unfortunately, there is no simple and direct way to compare the discrete vorticity distribution generated by the state estimate with the true continuous vorticity distribution.




Figure 4.2: (a): Schematic of the plate at 20° at $t^* = 0$ subject to actuation. Green dot depicts the location of the actuation. (b) Time history of the flow actuation.

Instead, we compare the normal force coefficient C_n (i.e., the integral of this pressure distribution on the plate rescaled by $\rho c U^2/2$) from the high-fidelity simulation and from its evaluation for the analysis ensemble obtained from either the sEnKF or the ETKF. The performance of each filter is assessed with the root-mean-squared error (RMSE), the standard deviation of the ensemble [39], and the interquantile range defined by the 2.5% and the 97.5% quantiles. These metrics are time-averaged. We define the RMSE between the true observation C_n^* (from the high-fidelity simulation) and the mean posterior normal force coefficient \bar{C}_n^a (i.e., the mean normal force coefficient for the different realizations computed from the posterior ensemble) as $\text{RMSE} = \|C_n^* - \bar{C}_n^a\|_2$. We quantify the dispersion of the ensemble with the sample standard deviation of the normal force coefficient $\sigma_{C_n^a}$. While the RMSE and the standard deviation of the ensemble can be affected by the spurious force events, the interquantile range defined between the 2.5% and the 97.5% quantiles is insensitive to outliers, and is used as a companion to the standard deviation to quantify the spread of the ensemble.

4.2.1 Translating plate subjected to pulse actuation disturbances

In this section, we assess the data-assimilated vortex model on the response of an impulsively translating plate at 20° to disturbances applied near the leading edge. We apply the

perturbations at a point $0.475c$ from the centroid along the plate and $0.1c$ above it, as shown in Fig. 4.2(a). The perturbations are introduced as a superposition of vertical body force of nominal strength $F_n = 0.03\rho U^2 c$ at $t^* = 2.5, 2.9, 3.0$ and 3.2 , distributed in Gaussian form in time and space, with temporal standard deviation of $t_{std}^* = 0.1$ and spatial standard deviation $0.1c$. The vertical body force $f(t^*)$ is given by:

$$\frac{f(t^*)}{F_n} = \mathcal{N}(t^*; 2.5, t_{std}^*) + \mathcal{N}(t^*; 2.9, t_{std}^*) + \mathcal{N}(t^*; 3.0, t_{std}^*) + \mathcal{N}(t^*; 3.2, t_{std}^*), \quad (4.6)$$

where $\mathcal{N}(t^*; \mu, \sigma)$ denotes the temporal Gaussian kernel, with mean μ and standard deviation σ , evaluated at t^* . Fig. 4.2(b) shows the time history of the force actuation. For reference, the gusts considered in Darakananda et al. [29] were weaker and non-overlapping, applied further upstream of the plate at $t^* = 3, 4$ with nominal amplitude $F_n = 0.01\rho U^2 c$. The disturbance considered here is more challenging due to the presence of strong and overlapping perturbations.

Fig. 4.3 compares the history of the surface pressure response from the truth with the mean of 100 realizations of the sEnKF and the ETKF. During the first two convective times, we observe the development of the leading-edge vortex. The successive flow actuation disturbances are easily detected by the distinct regions of strongly negative pressure. The first suction region is due to the disturbance applied at $t^* = 2.5$, while the second one is due to the superposed response to the flow disturbances centered at $t^* = 2.9, 3.0$, and 3.2 . Both filters match well with the true pressure distribution. The uncertainty of each filter is characterized by the sample standard deviation of the pressure, computed over the 100 runs; these are shown in the right column of Fig. 4.3. In both filters, two narrow bands of high variance, particularly strong near the leading edge, can be identified at $t^* = 2.5$ and 3.0 . These correspond to the instants of local maxima in the disturbance force. However, the sEnKF has a higher level of variability in the pressure distribution from one run to another. In particular, there is a significant band of high dispersion for the sEnKF at around $t^* = 3.8$, and additional smaller bands at other times. The ETKF does not exhibit such bands, a direct consequence of its statistically consistent analysis update.




Figure 4.3: Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° subjected to pulse actuation disturbance from (a) high-fidelity numerical simulation at Reynolds number 500, and mean over 100 realizations of an inviscid vortex model with (b) the sEnKF and (d) the ETKF. Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for (c) the sEnKF and (e) the ETKF.

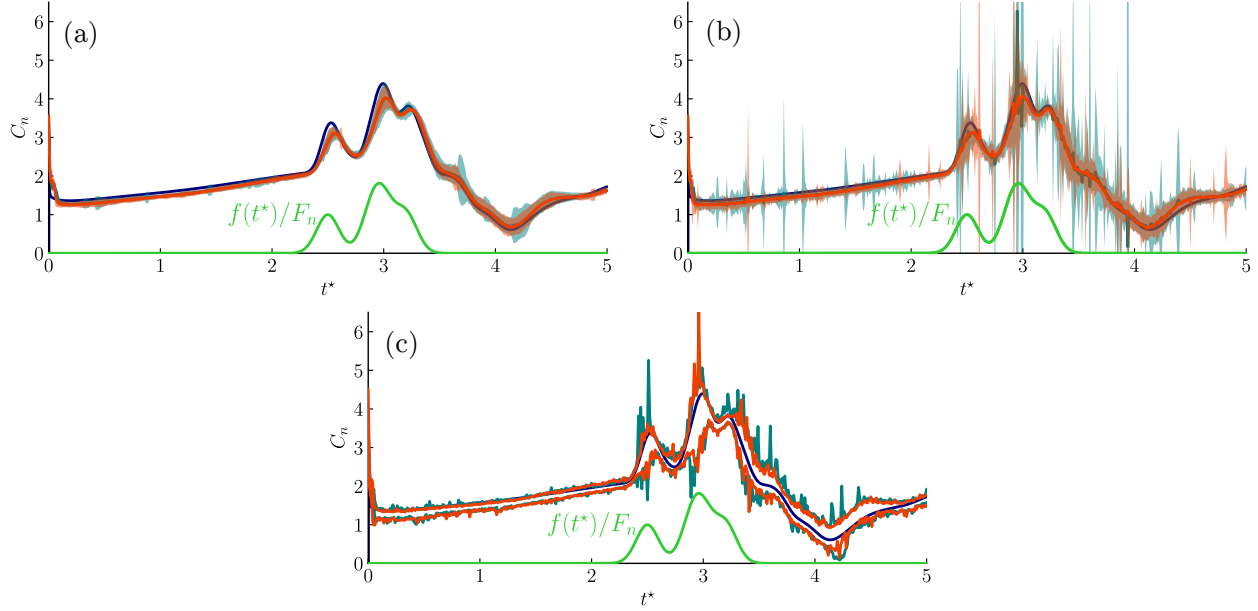


Figure 4.4: Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° subject to actuation, from high-fidelity numerical simulation at Reynolds number 500 (—), mean over 100 realizations of the inviscid vortex model with sEnKF (—), mean over 100 realizations of the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Time history of the flow actuation (—). Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—) and ETKF (—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—) and ETKF (—) without application of the median filter.

The normal force coefficient is the integral of this pressure distribution on the plate. The top left panel in Fig. 4.4 compares the force coefficient obtained from the truth system with the mean of the sEnKF and the ETKF applications over 100 realizations. The mean force predicted by each filter agrees very well with the true force response. The peaks created by the actuation and the subsequent drop of force around $t^* = 4$ are also well predicted. However, consistent with our observations of the surface pressure data, both filters show variability near the peak disturbances; Fig. 4.4 also shows that the peaks are slightly underpredicted.

The sEnKF exhibits more variability than the ETKF at all times, and this higher variability is apparent in the wider uncertainty envelope for the normal force coefficient, particularly after $t^* > 3.5$. To better appreciate the raw behavior of each filter, the top right panel in Fig. 4.4 show the same results, but without the use of the median filter. The sEnKF, in Fig. 4.4(b), exhibits frequent spikes throughout the simulation, and particularly so after $t^* > 3.5$. These spikes are not eliminated by the ETKF, but they are much weaker and less frequent. Due to the presence of spikes, the 95% confidence interval based on the standard deviation can be corrupted. To better appreciate the raw performance of the two filters, we have also plotted the 2.5% and 97.5% quantiles for each filter in the lower panel in Fig. 4.4. Table 4.1 reports the time-averaged RMSE, the standard deviation and the interquantile range of the normal force coefficient for the sEnKF and the ETKF, without the application of the median filter. The ETKF outperforms the sEnKF with an average reduction of 29% of the RMSE, 49% of the standard deviation and 17% of the interquantile range. We should note that there is no downside to apply the median filter.

The origins of this high uncertainty are clear when we examine the flow behavior. In Fig. 4.5 we show the vorticity distribution from the truth system at three instants, $t^* = 3.0$, 4.0, and 5.0. Since the vorticity in the fluid is modeled with limited number of regularized point vortices, it is not possible to make a direct comparison; instead, we plot the locations of the elements for one realization of each filter, with the elements' signed circulations

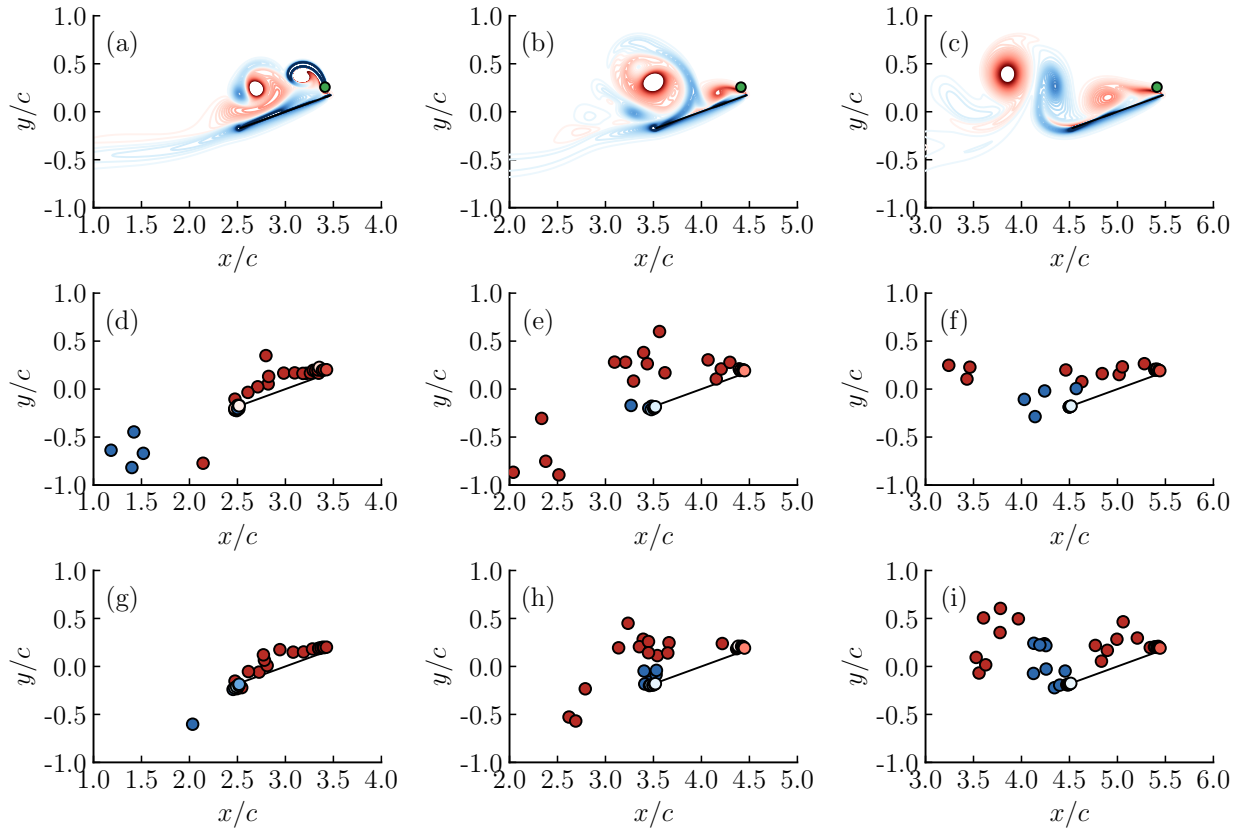


Figure 4.5: Snapshots of the vorticity distribution at $t^* = 3.0$ (left column), $t^* = 4.0$ (middle column) and $t^* = 5.0$ (right column) for an impulsively translating plate at 20° subject to actuation, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.




Figure 4.6: Time history of the ensemble mean value of the $LESP_c$ of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

represented in colors consistent with the sign of vorticity in the truth data. Most of the large-scale vortex dynamics are captured by the data-assimilated vortex model. Over the first two convective times, the impulsive translation of the plate creates a leading-edge vortex that grows in size and strength. This feature is captured by the continuous release of new vortex elements in both filters. The disturbance applied at $t^* = 2.5$ creates a coherent structure that is rapidly advected along the plate and merges with the initial leading-edge vortex. At around $t^* = 3.0$, the resulting vortex lies slightly above the mid chord of the plate, evident in the left column of Fig. 4.5. Both the sEnKF and ETKF capture this feature with positive vortex elements clustered at a similar position. It is also important to note that the vorticity associated with the disturbance itself, clearly evident at $t^* = 3.0$, is not observed in the vortex model results. The absence of this disturbance vorticity is by design; only its effect on vortex elements shed from the edge of the plate is captured by the assimilation of pressure data.

Over the time interval $t^* \in [3.0, 4.0]$, the vortex models predict a separation of a cluster of positive vortex elements from the plate with dynamics similar to the truth flow. The middle column of Fig. 4.5 compares the results at $t^* = 4.0$. The large coherent structure with positive vorticity is well captured by the inviscid model with both filters, but is more tightly clustered in the case of the ETKF. From $t^* = 4.0$ to 5.0 , the vortex is shed into the

wake and triggers a large flux of negative vorticity from the trailing edge. The right column of Fig. 4.5, for $t^* = 5.0$, shows a good visual agreement between the truth and the prediction of the ETKF. The spatial vortex distribution predicted by the sEnKF is less representative of the true vorticity distribution, however. Overall, the ETKF predicts a more structured and physically consistent spatial distribution of the vortices than the sEnKF. The coherence of these vortex element clusters is consistent with the narrower uncertainty in the normal force coefficient, and demonstrates a clear advantage of the ETKF over the sEnKF for modeling the flow response to unknown flow perturbations.

In Fig. 4.6 we compare the time histories of the LESPC estimated by the sEnKF and the ETKF. The LESPC is constrained to remain positive, and the constraint reverts to the Kutta condition if LESPC becomes zero. Before the flow is disturbed, the mean estimate of the LESPC stays on a plateau about 0.5, the mean value in the initial ensemble. This behavior supports the hypothesis of Ramesh et al. [26, 27]: the LESPC remains constant for a given Reynolds number and airfoil section. The time variation of the imposed disturbance is reflected in a similar variation of the LESPC. Indeed, the application of an actuation-like disturbance near the leading edge directly controls the vorticity flux about this edge. Large values of LESPC lead to weaker vorticity, temporarily suppressing the flux into the shear layer. The small decay of LESPC after the first disturbance increases the vorticity flux, triggering the creation of a new coherent structure that merges with the initial leading-edge vortex about the mid-chord. This leading-edge development is then halted after the next disturbance peak, and the leading-edge vortex is shed. The uncertainty envelopes of the two filters are very similar and tend to grow over time. The width of these envelopes is large, reflecting significant variation in the estimated values of LESPC from one realization to the next. This variation indicates a weak physical correlation between this threshold value and the pressure on the plate: this threshold's effect on pressure is only exerted indirectly, through the subsequent release of vorticity. (The LESPC itself, in contrast, is more strongly correlated, since shed vortex elements contribute to this value [27, 6].)




Figure 4.7: Time history of the particle count for an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

The population histories of vortex elements is depicted in Fig. 4.7. The histories are nearly identical for each filter. In the initial time steps, the number of vortex elements increases linearly from 0 to 20; the model prevents aggregation of elements during this interval. Subsequently, the population grows slowly up to 35 ± 7 at $t^* = 5$. It should be noted that, without aggregation, this population would be approximately 1000 vortex elements (500 time steps, with two elements shed per step). The variation in population among the different realizations of each filter is attributable to the variation in LESPC, which sets the initial strengths of the elements, which in turn affects their later aggregation. It is interesting to note that this variation of vortex element populations is proportionally larger than the variation in the pressure and normal force, indicating that there is some non-uniqueness in the mapping from surface pressures to vortex element dynamics.

Fig. 4.8 depicts the ensemble variances of the positions and strengths of the vortex elements and the LESPC estimate, averaged over the 100 realizations of each filter. Each variance is lower-bounded by the additive covariance inflation to avoid filter divergence. The variances of the strengths of the vortices and LESPC are fairly constant (and near the values set by the inflation parameter) while those of the x and y positions are more variable. These values are larger than the variance set by the inflation, likely due to additional error incurred by aggregation.

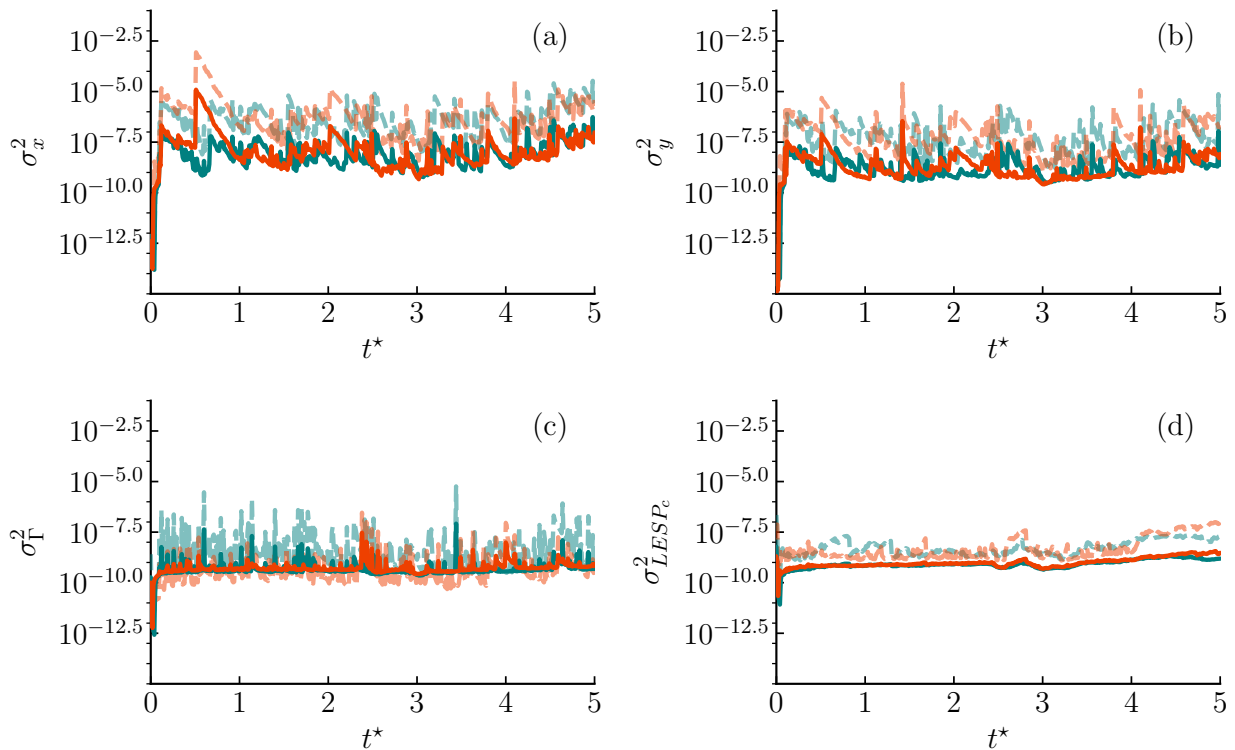


Figure 4.8: Time history of the ensemble variances of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations. Mean variances for (a) the x coordinate of the blobs, (b) the y coordinate of the blobs, (c) the circulation of the blobs, and (d) the LESPC from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Fainter dashed lines show the standard deviation of the different variances over the 100 realizations for the inviscid vortex model with sEnKF and ETKF.

	RMSE	Standard deviation	Interquantile range
sEnKF	3.67	3.06×10^{-1}	4.71×10^{-1}
ETKF	2.61	1.56×10^{-1}	3.92×10^{-1}

Table 4.1: Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject to actuation with the sEnKF and the ETKF, without application of the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.




Figure 4.9: Schematic of a plate at 20° behind a cylinder. The plate and cylinder translate to the right at uniform speed.

4.2.2 Plate in the wake of a cylinder

In this part, we assess our flow estimator when applied to a plate at 20° angle of attack in the wake of a cylinder. In the truth system, high fidelity simulations are conducted of a cylinder of diameter $0.16c$ is centered 2 chord lengths upstream and 0.3 chord lengths above the plate's centroid, as shown in Fig. 4.9. The plate and cylinder are both impulsively set in motion at $t^* = 0$ at speed U , so that their relative configuration remains fixed for all time. The Reynolds number based on the cylinder diameter is 80, sufficiently large that the cylinder's wake exhibits a von Kármán vortex street. The presence of the plate triggers the wake to break symmetry and achieve this vortex street. This flow configuration distills the main features of a vehicle flying through the wake of a structure, e.g., buildings in an urban environment or other flying vehicles.

It is important to recall that we do not include the cylinder and its wake in our vortex model flow estimator; its effects are only felt through the pressure jump measurements obtained from surface sensors along the plate. Fig. 4.10 depicts the history of the true surface pressure distribution and those estimated by the sEnKF and ETKF; the associated

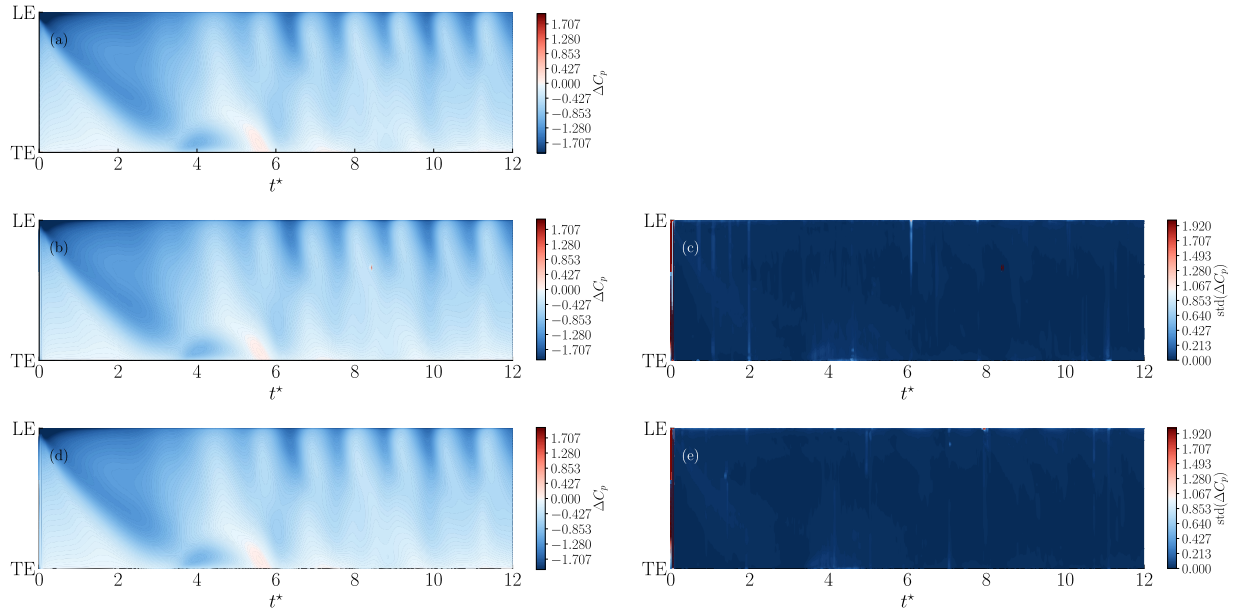


Figure 4.10: Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° in a cylinder wake (a) high-fidelity numerical simulation at Reynolds number 500, mean over 100 realizations of an inviscid vortex model with the sEnKF (b) and the ETKF (d). Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for the sEnKF(c) and ETKF (e)

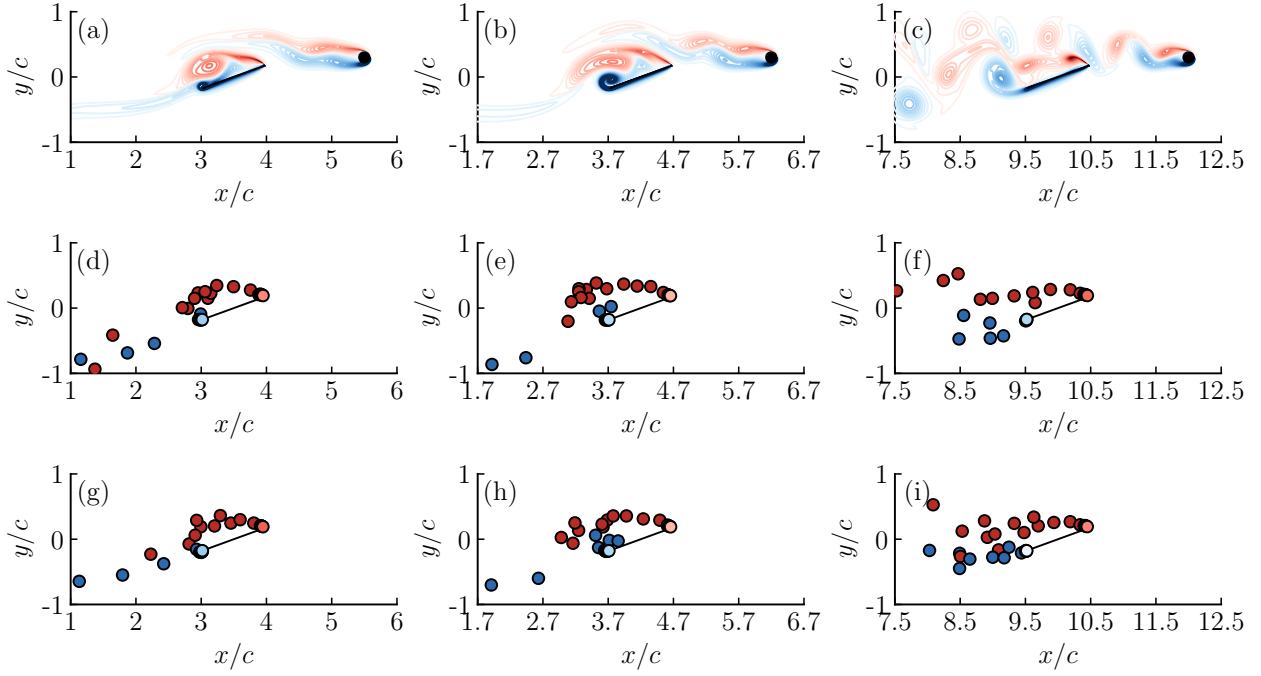


Figure 4.11: Snapshots of the vorticity distribution at $t^* = 3.5$ (left column), $t^* = 4.2$ (middle column) and $t^* = 10.0$ (right column) for an impulsively translating plate at 20° in a cylinder wake, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.

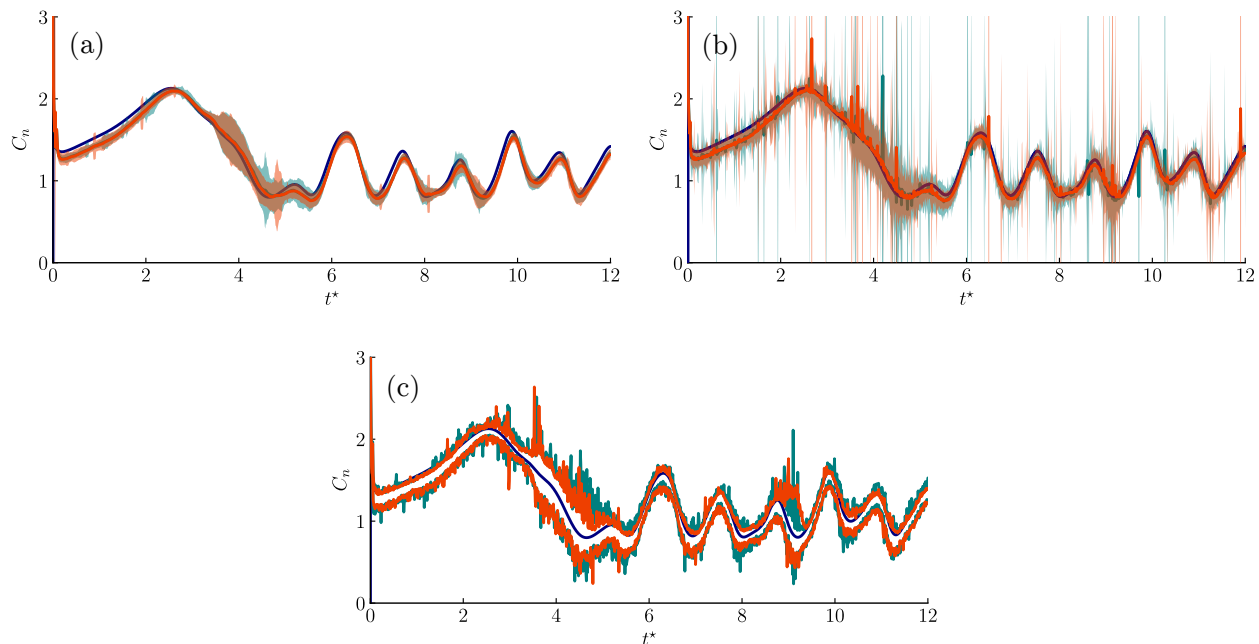


Figure 4.12: Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° in a cylinder wake, from high-fidelity numerical simulation at Reynolds number 500 (—), mean over 100 realizations of the inviscid vortex model with sEnKF (—), mean over 100 realizations of the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—) and ETKF (—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—) and ETKF (—) without application of the median filter.




Figure 4.13: Time history of the ensemble mean value of the $LESP_c$ of an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.




Figure 4.14: Time history of the particle count for an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

standard deviations are shown on the right. Both filters estimate pressure fields that agree very well with the true distribution. As expected, it takes about 2 convective times for the cylinder wake to reach and be sensed by the plate. Over this time window, the pressure field is essentially disturbance-free. The encounter of the vortex structures shed by the cylinder leaves successive short-lived pressure disturbances.

The true vorticity distribution and mean sets of vortex elements are compared in Fig. 4.11 at $t^* = 3.5$ (soon after the wake has reached the plate), at $t^* = 4.2$ (when the cylinder wake transitions to a von Kármán vortex street) and at $t^* = 10$ (long after the flow has achieved periodic vortex shedding). As expected, the vortex models do not attempt to represent the cylinder wake with vortex elements; rather, the filter accommodates the influence of the wake vorticity by modifying the behavior of the vortex elements shed from the plate. Large scale structures of the flow around the plate are captured and match visually with the true vorticity field.

Concomitantly, the normal force estimate agrees well with the truth, as seen in Fig. 4.12. As in the first example, we have plotted the normal force estimate of the two filters, without the application of the median filter, and the associated uncertainty between different realizations based on the standard deviation (Fig. 4.12 (b)) and the interquantile range (Fig. 4.12 (c)). The transition from a symmetric cylinder wake to a periodic vortex shedding causes some challenge to the discrete vortex models, as highlighted by the temporary growth of the uncertainty envelope around $t^* = 4.0$. The envelope remains small for both filters after the periodic wake behavior has been established. Table 4.2 reports the time-averaged RMSE and the standard deviation and the interquantile range of the normal force coefficient for the sEnKF and the ETKF. The ETKF reduces on average the standard deviation by 16% and the interquantile range by 16%. The RMSE performance of the two filters are almost identical. The history of the LESPc estimate in Fig. 4.13 reveals the same growth and decay of this critical value observed in the previous example; here, they represent a response to individual cylinder wake vortices passing the leading edge of the plate. These variations are essential

	RMSE	Standard deviation	Interquantile range
sEnKF	2.63	1.63×10^{-1}	3.18×10^{-1}
ETKF	2.66	1.41×10^{-1}	2.69×10^{-1}

Table 4.2: Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject in a cylinder wake with the sEnKF and the ETKF without the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.

to our inviscid framework to control the leading-edge vortex shedding in the presence of flow perturbations. As in the previous example, the weak correlation between the LESPC and measured pressure causes significant volatility in the estimated LESPC among the different realizations. The vortex element population remains small, $O(60)$, after 12 convective time units. Without aggregation, we would have to track 2400 vortex elements.

CHAPTER 5

Regularization of the ensemble Kalman filter in elliptic inverse problems

This chapter is adapted from [55].

In this chapter, we are interested in filtering problems where the assimilated observations correspond to *non-local* functions of the state, such as integrals of linear and nonlinear functions of the state [56]. As examples, we can cite the radiance measured by satellites, heat or mass fluxes through surfaces, forces measured on a body immersed in a fluid, or as we focus in this chapter, solutions of elliptic partial differential equations (PDEs). An elliptic PDE is given by

$$\mathcal{L}\mathbf{u}_k(\mathbf{r}) = \mathbf{q}_k(\mathbf{r}, \mathbf{x}_k), \quad \mathbf{r} \in \Omega, \quad (5.1)$$

where Ω is the physical space (a subspace of \mathbb{R}^2 or \mathbb{R}^3), $\mathbf{r} \in \Omega$ denotes the point of evaluation of (5.1), and \mathcal{L} is an elliptic linear operator (e.g. the Laplacian ∇^2 , in which case (5.1) is a Poisson equation). \mathbf{q}_k is a forcing term that depends nonlinearly on the state \mathbf{x}_k and comprises information about the spatial distribution and strength of the forcing. This chapter focuses on *elliptic inverse problems* where we seek to determine the state \mathbf{x}_k from noisy and spatially limited evaluations (i.e. for limited points $\mathbf{r} \in \Omega$) of the solution \mathbf{u}_k of the elliptic PDE (5.1). Up to an homogeneous part, the solution \mathbf{u}_k can be obtained by convolution of the Green's function of the elliptic operator \mathcal{L} with the forcing term \mathbf{q}_k over Ω . Elliptic inverse problems are particularly challenging as the observations inherit the non-locality and the nonlinearity of the Green's function and further nonlinear state dependence in the forcing term \mathbf{q}_k .

While in many high-dimensional problems the EnKF can successfully track the state with very limited samples, this success is predicated on an adequate regularization of the empirical Kalman gain. Predominant regularization techniques assume that the observations are *local*, i.e. an observation only provides information about a subset of the state variables which are close-by in physical distance. This assumption is supported by the rapid decay of the correlations between state and observation variables. In this setting, distance localization cuts off long-range correlations [13].

Distance localization is inappropriate in the current context, however. For a Poisson equation, the Green’s function of the Laplacian has a logarithmic or algebraic decay (based on the dimension of the physical space) as a function of distance, so we expect long-range physical interactions between the state and observation variables. Distance localization should therefore be avoided, as it will remove all long-range correlations, spurious or physical. In this chapter, we develop a novel regularization technique (discussed in the next paragraph) to assimilate non-local observations. While the methodology can be applied to a wide range of problems, we focus our discussion on the representative context of incompressible fluid mechanics. In these problems, the flow field is most compactly represented by the vorticity, or curl of the velocity field. We use a low-dimensional Lagrangian representation of the flow field by tracking the positions and strengths of a small collection of point vortices. Inviscid point vortex models have been a long-standing tool to model and explain incompressible fluid phenomena [57, 6, 3]. In this chapter, we focus on filtering problems where we seek to estimate the characteristics of point singularities over time from limited potential or pressure observations. Despite their relative low-dimensionality, these problems can be particularly challenging as the transport equation for point vortices (the Biot-Savart law) and the observation model (the pressure Poisson equation) involve the resolution of Poisson equations like (5.1) whose forcing terms nonlinearly couple all the singularities’ contributions.

The rest of this chapter is organized as follows. Section 5.1 presents our methodology. Example problems are treated in Section 5.2. A pseudo-code for the proposed low-rank

ensemble Kalman filter is provided in the Appendix A.

5.1 Low-rank ensemble Kalman filter

5.1.1 Ensemble Kalman filter

For the remainder of this section, we focus on a single analysis step of the EnKF. Our treatment of the analysis step is built on the idea that there is an underlying transformation \mathbf{T}_k , called the *prior-to-posterior transformation* or *analysis map*, that directly maps samples from the forecast (i.e. prior) density $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}}$ to the filtering (i.e. posterior) density $\pi_{\mathbf{x}_k | \mathbf{Y}_{1:k} = \mathbf{y}_{1:k}}$ [40, 39]. Further details on the filtering problem and the analysis map are given in Chapter 3. Kalman [15] showed that the best linear estimator for the prior-to-posterior transformation \mathbf{T}_k is given by:

$$\mathbf{T}_k(\mathbf{y}_k, \mathbf{x}_k) = \mathbf{x}_k - \Sigma_{\mathbf{x}_k, \mathbf{Y}_k} \Sigma_{\mathbf{Y}_k}^{-1} (\mathbf{y}_k - \mathbf{y}_k^*), \quad (5.2)$$

where \mathbf{y}_k^* is the observation to be assimilated, $\Sigma_{\mathbf{x}_k, \mathbf{Y}_k} \in \mathbb{R}^{n \times d}$ is the cross-covariance matrix of the state and observation, while $\Sigma_{\mathbf{Y}_k}^{-1} \in \mathbb{R}^{d \times d}$ is the precision matrix (inverse of the covariance matrix) of the observation's marginal distribution. The linear operator $\mathbf{K}_k = \Sigma_{\mathbf{x}_k, \mathbf{Y}_k} \Sigma_{\mathbf{Y}_k}^{-1} \in \mathbb{R}^{n \times d}$ is called the Kalman gain and maps observation discrepancies ($\mathbf{y}_k - \mathbf{y}_k^*$) to the state correction. The *stochastic ensemble Kalman filter (sEnKF)* introduced by Evensen [14] estimates the transformation (5.2) by replacing the covariances with empirical covariances that are computed from the joint samples of the observations and states $\{\mathbf{y}_k^i, \mathbf{x}_k^i\} \sim \pi_{\mathbf{x}_k, \mathbf{Y}_k | \mathbf{Y}_{1:k-1}}$, i.e.

$$\widehat{\mathbf{T}}_k(\mathbf{y}_k, \mathbf{x}_k) = \mathbf{x}_k - \widehat{\Sigma}_{\mathbf{x}_k, \mathbf{Y}_k} \widehat{\Sigma}_{\mathbf{Y}_k}^{-1} (\mathbf{y}_k - \mathbf{y}_k^*). \quad (5.3)$$

5.1.2 Assimilation in low-dimensional subspaces

For elliptic inverse problems, the conditional structure of the joint density of the observation and the state $\pi_{\mathbf{x}_k, \mathbf{Y}_k | \mathbf{Y}_{1:k-1}}$ is not localized, i.e. there is no rapid decay of the correlations as

a function of the distance between variables. Hence, distance based localization in this problem results in biased estimators for the Kalman gain. Instead, we will exploit another kind of low-dimensional structure in the joint distribution. Our treatment draws inspiration from the fast multipole method (FMM) [58], which uses a hierarchical clustering of “source” and “target” elements to accelerate the calculation of the potential field from a large set of point singularities [59]. These clusterings are typically based only on spatial distance, but here we will use information from the prior distribution and observation model (i.e. observation operator and observation noise) to infer these clusters automatically in a way that works independently of the spatial distribution of the singularities and observation locations. To regularize the EnKF, we identify important directions in the state and observation spaces, perform the assimilation in these low-dimensional subspaces, and finally, lift the result to the original state space. First, we present the treatment of low-rank structure in the case of a linear-Gaussian observation model in 5.1.2(5.1.2.1). The main result of this derivation is a factorization of the Kalman gain that exploits the existence of this low-dimensional subspace. Then, we extend this decomposition to the nonlinear-Gaussian setting in 5.1.2(5.1.2.2). An algorithm that summarizes the overall low-rank assimilation procedure is provided in Appendix A. For convenience, we drop the time dependence subscripts of the variables in the rest of this chapter, since the analysis step does not involve time propagation.

5.1.2.1 Low-rank assimilation for the linear-Gaussian case

In this section, we consider the inference problem for the linear-Gaussian observation model:

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathcal{E}, \tag{5.4}$$

where the state is given by $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{X}}, \boldsymbol{\Sigma}_{\mathbf{X}})$ and the observational error is given by $\mathcal{E} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathcal{E}})$ where \mathcal{E} is independent of \mathbf{X} . The matrix $\mathbf{H} \in \mathbb{R}^{d \times n}$ is called the observation matrix. In order to identify the important assimilation directions, we first define the whitened variables $\tilde{\mathbf{X}} = \boldsymbol{\Sigma}_{\mathbf{X}}^{-1/2} (\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}}) \in \mathbb{R}^n$, $\tilde{\mathcal{E}} = \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} \mathcal{E} \in \mathbb{R}^d$. We use $\mathbf{B}^{1/2}, \mathbf{B}^{-1/2}$ to denote a

square root of the matrix \mathbf{B} and its inverse, respectively. These whitened variables satisfy $\tilde{\boldsymbol{\epsilon}} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, $\tilde{\mathbf{X}} \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$. Applying the same whitening to the observation variable, the observation model becomes

$$\tilde{\mathbf{Y}} = \Sigma_{\boldsymbol{\epsilon}}^{-1/2} (\mathbf{Y} - \boldsymbol{\mu}_{\mathbf{Y}}) = \tilde{\mathbf{H}}\tilde{\mathbf{X}} + \tilde{\boldsymbol{\epsilon}}, \quad (5.5)$$

where $\boldsymbol{\mu}_{\mathbf{Y}} = \mathbf{H}\boldsymbol{\mu}_{\mathbf{X}}$ and $\tilde{\mathbf{H}} = \Sigma_{\boldsymbol{\epsilon}}^{-1/2}\mathbf{H}\Sigma_{\mathbf{X}}^{1/2} \in \mathbb{R}^{d \times n}$ is the whitened observation matrix. In this derivation we assume that $d \leq n$. We can now identify the important assimilation directions with a singular value decomposition (SVD) of the whitened observation matrix $\tilde{\mathbf{H}} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{d \times d}$ and $\mathbf{V} \in \mathbb{R}^{n \times d}$ are the left and right singular vectors, and $\boldsymbol{\Lambda} \in \mathbb{R}^{d \times d}$ is the diagonal matrix of singular values. From this decomposition, we can rotate and project the whitened state on the subspace spanned by the columns of \mathbf{V}^\top — $\check{\mathbf{X}} = \mathbf{V}^\top\tilde{\mathbf{X}} \in \mathbb{R}^d$ — and similarly rotate and project the whitened observational error and observation on the subspace spanned by the columns of \mathbf{U}^\top — $\check{\boldsymbol{\epsilon}} = \mathbf{U}^\top\tilde{\boldsymbol{\epsilon}} \in \mathbb{R}^d$, $\check{\mathbf{Y}} = \mathbf{U}^\top\tilde{\mathbf{Y}} \in \mathbb{R}^d$. The SVD of the whitened observation matrix simultaneously identifies a pair of orthogonal bases for the state and observation spaces and an ordering for these directions. In the rotated spaces, the observation model can be written as:

$$\check{\mathbf{Y}} = \boldsymbol{\Lambda}\check{\mathbf{X}} + \check{\boldsymbol{\epsilon}}, \quad (5.6)$$

where the rotated observation operator $\boldsymbol{\Lambda}$ and the observation error covariance are diagonal matrices. Hence, the inference in this rotated space can be performed in a fully decoupled manner for each state variable. Using the decoupled observation model (5.6), the Kalman gain in the rotated space is given by:

$$\check{\mathbf{K}} = \boldsymbol{\Lambda}(\boldsymbol{\Lambda}^2 + \mathbf{I}_d)^{-1}. \quad (5.7)$$

Thus, the linear analysis map in the rotated space $\check{\mathbf{T}}$ is given by:

$$\check{\mathbf{T}}(\check{\mathbf{y}}, \check{\mathbf{x}}) = \check{\mathbf{x}} - \check{\mathbf{K}}(\check{\mathbf{y}} - \check{\mathbf{y}}^*) = \check{\mathbf{x}} - \boldsymbol{\Lambda}(\boldsymbol{\Lambda}^2 + \mathbf{I}_d)^{-1}(\check{\mathbf{y}} - \check{\mathbf{y}}^*), \quad (5.8)$$

where $\check{\mathbf{y}}^* = \mathbf{U}^\top \Sigma_{\boldsymbol{\epsilon}}^{-1/2} \mathbf{y}^*$ denotes the realization of the assimilated data in the rotated observation space. In these rotated coordinates, the analysis map is local, i.e. each component of

the rotated observation only updates one associated component of the rotated state. The d columns of \mathbf{V} span the informative subspace of the whitened state space (i.e. where the random variable $\tilde{\mathbf{X}}$ lives). Thus, $\mathbf{P}_d = \mathbf{V}\mathbf{V}^\top \in \mathbb{R}^{n \times n}$ is an orthogonal projector from the whitened state space to this informative subspace. Note that \mathbf{P}_d has rank d as the subscript suggests. Therefore, the whitened state variable can be decomposed as:

$$\tilde{\mathbf{X}} = \mathbf{P}_d \tilde{\mathbf{X}} + (\mathbf{I}_n - \mathbf{P}_d) \tilde{\mathbf{X}} = \tilde{\mathbf{X}}^\parallel + \tilde{\mathbf{X}}^\perp, \quad (5.9)$$

where $\tilde{\mathbf{X}}^\parallel := \mathbf{P}_d \tilde{\mathbf{X}}$ and $\tilde{\mathbf{X}}^\perp := (\mathbf{I}_n - \mathbf{P}_d) \tilde{\mathbf{X}}$. Since \mathbf{P}_d is an orthogonal projector, the decomposition above is unique and $\tilde{\mathbf{X}}^\parallel \in \mathbb{R}^n$ is orthogonal to $\tilde{\mathbf{X}}^\perp \in \mathbb{R}^n$. We can always complete the orthogonal family of columns of \mathbf{V} to form a basis for \mathbb{R}^n with $n - d$ orthogonal columns \mathbf{V}^\perp . The projector on this complementary subspace is given by $\mathbf{P}_{n-d}^\perp = \mathbf{V}^\perp \mathbf{V}^{\perp\top} = \mathbf{I}_n - \mathbf{P}_d$. We can now connect the decomposition (5.9) with the projected state variables on the informative and complementary subspaces:

$$\tilde{\mathbf{X}} = \mathbf{V}\mathbf{V}^\top \tilde{\mathbf{X}} + \mathbf{V}^\perp \mathbf{V}^{\perp\top} \tilde{\mathbf{X}} = \mathbf{V}(\mathbf{V}^\top \tilde{\mathbf{X}}) + \mathbf{V}^\perp(\mathbf{V}^{\perp\top} \tilde{\mathbf{X}}) = \mathbf{V}\check{\mathbf{X}} + \mathbf{V}^\perp \mathbf{X}^\perp, \quad (5.10)$$

with $\check{\mathbf{X}} \in \mathbb{R}^d$ and $\mathbf{X}^\perp \in \mathbb{R}^{n-d}$. We emphasize that the inference performed with the analysis map in (5.8) only affects the rotated state variable $\check{\mathbf{X}} \in \mathbb{R}^d$, while the complementary component $\tilde{\mathbf{X}}^\perp$ of the whitened state is unaffected by the assimilation. Using eq. (5.10), the low-rank analysis map in the original space is given by:

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \Sigma_{\mathbf{x}}^{1/2} (\mathbf{V}\check{\mathbf{T}}(\check{\mathbf{y}}, \check{\mathbf{x}}) + (\mathbf{I}_d - \mathbf{V}\mathbf{V}^\top) \Sigma_{\mathbf{x}}^{-1/2} \mathbf{x}) = \mathbf{x} - \mathbf{K}(\mathbf{y} - \mathbf{y}^*), \quad (5.11)$$

where \mathbf{K} denotes the Kalman gain. We note that this recovers the canonical form of the linear analysis map given in eq. (5.3) with the Kalman gain $\mathbf{K} \in \mathbb{R}^{n \times d}$ factorized as

$$\mathbf{K} = \Sigma_{\mathbf{x}}^{1/2} \mathbf{V} \Lambda (\Lambda^2 + \mathbf{I}_d)^{-1} \mathbf{U}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1/2}. \quad (5.12)$$

A few important comments are in order about this factorization of the Kalman gain. The application of \mathbf{K} in (5.11) and (5.12) provides a concise summary of the inference process in the low-rank informative subspaces: the innovation term $(\mathbf{y} - \mathbf{y}^*)$ is whitened and rotated,

assimilated in this new set of coordinates, and finally lifted to the original state space. In the whitened space (where $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ live), it is easy to see that $\mathbf{V}\mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1}\mathbf{U}^\top$ constitutes the singular value decomposition of the Kalman gain. Thus, in the whitened space, exploiting the spectrum of the whitened observation matrix gives us the best low-rank approximation of the Kalman gain. Unfortunately, (5.12) is no longer the SVD of the Kalman gain in the original space as the columns of the matrices $\Sigma_{\mathbf{X}}^{1/2}\mathbf{V}$ and $\Sigma_{\mathcal{E}}^{-\top/2}\mathbf{U}$ are not necessary orthogonal. Nonetheless, the proposed factorization gives us a constructive means to form a low-rank approximation of the Kalman gain in the original space (even if not optimal in the Frobenius norm) without having to form the entire Kalman gain. This is clearly a desired feature for large inference problems where the Kalman gain cannot be practically formed. In the next section, we will generalize this factorization to perform the assimilation using low-dimensional subspaces in the nonlinear-Gaussian setting. From the decay of the singular values $\mathbf{\Lambda}$, we can instead use a truncated singular value decomposition of $\tilde{\mathbf{H}} \approx \mathbf{U}_r\mathbf{\Lambda}_r\mathbf{V}_r^\top$, where $r \leq d$, $\mathbf{U}_r \in \mathbb{R}^{r \times r}$, $\mathbf{V}_r \in \mathbb{R}^{n \times r}$ are the first r left and right singular vectors, and $\mathbf{\Lambda} \in \mathbb{R}^{r \times r}$. A rank- r approximation of the Kalman gain is then given by:

$$\mathbf{K}_r = \Sigma_{\mathbf{X}}^{1/2}\mathbf{V}_r\mathbf{\Lambda}_r(\mathbf{\Lambda}_r^2 + \mathbf{I}_r)^{-1}\mathbf{U}_r^\top\Sigma_{\mathcal{E}}^{-1/2}. \quad (5.13)$$

We should emphasize that, without the truncation of the SVD, eq. (5.12) is an exact factorization of the Kalman gain. With the truncated SVD, the state and observation variables are no longer just rotated but also projected to a subspace of dimension $r < \min(d, n)$.

5.1.2.2 Low-rank assimilation for the nonlinear-Gaussian case

In this section, we consider the nonlinear observation model of eq. (3.2), recalled for reference:

$$\mathbf{Y} = \mathbf{h}(\mathbf{X}) + \mathcal{E}, \quad (5.14)$$

where the state is $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{X}}, \Sigma_{\mathbf{X}})$ and the observational error is given by $\mathcal{E} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathcal{E}})$ where \mathcal{E} is independent of \mathbf{X} . To handle a nonlinear observation operator with a linear Gaussian filter, we usually use the Jacobian of the observation operator about the

prior mean as the observation matrix [53, 13]. Unfortunately, this treatment is only viable for functions which are well approximated by a linear function over the bulk of the prior distribution. Clearly, the proposed linearization about the prior mean, and consequently the results of the previous section, cannot be applied to the observation operator $\mathbf{h}: \mathbb{R}^n \rightarrow \mathbb{R}, \mathbf{x} \mapsto (1 + \|\mathbf{x}\|^2)^{-1}$ with a zero-mean prior distribution. This section resolves this issue in two steps. First, we present a methodology to identify the important subspaces for the state and observation variables for an arbitrary nonlinear observation model corrupted by an additive Gaussian noise. Second, we generalize the factorization of the Kalman gain (5.12) to the nonlinear-Gaussian case.

To understand the construction of the informative subspaces in the nonlinear case, it is enlightening to examine the directions that maximize the Rayleigh quotient of the posterior to prior precision in the linear-Gaussian case:

$$\mathcal{R}(\mathbf{w}) = \frac{\langle \mathbf{w}, \Sigma_{\mathbf{x}|\mathbf{y}}^{-1} \mathbf{w} \rangle}{\langle \mathbf{w}, \Sigma_{\mathbf{x}}^{-1} \mathbf{w} \rangle}, \text{ for } \mathbf{w} \in \mathbb{R}^n. \quad (5.15)$$

Given that $\Sigma_{\mathbf{x}|\mathbf{y}}^{-1} = \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} + \Sigma_{\mathbf{x}}^{-1}$, maximizing this Rayleigh quotient is equivalent to maximizing the Rayleigh ratio of the Hessian of the log-likelihood to the prior precision given by:

$$\mathcal{S}(\mathbf{w}) = \frac{\langle \mathbf{w}, \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} \mathbf{w} \rangle}{\langle \mathbf{w}, \Sigma_{\mathbf{x}}^{-1} \mathbf{w} \rangle}. \quad (5.16)$$

This equation shows the connection between the directions that maximize the Rayleigh quotient of the posterior to prior, and the directions (in the state space) where the observations are most informative with respect to the prior. With the change of variable $\mathbf{v} = \Sigma_{\mathbf{x}}^{-1/2} \mathbf{w} \in \mathbb{R}^n$, we also obtain

$$\tilde{\mathcal{S}}(\mathbf{v}) = \frac{\langle \mathbf{v}, \Sigma_{\mathbf{x}}^{1/2} \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} \Sigma_{\mathbf{x}}^{1/2} \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle}. \quad (5.17)$$

We denote the columns of \mathbf{V} that span the image space of the projector \mathbf{P}_d introduced in the previous section as $\{\mathbf{v}_1, \dots, \mathbf{v}_d\}$. Spantini et al. [32] showed that the vector \mathbf{v}_j maximizes the Rayleigh quotient $\tilde{\mathcal{S}}$ over the subspace $\mathbb{R}^n \setminus \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_{j-1}\}$, which is the null space of

the projector generated by the previous columns vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_{j-1}\}$. In other words, we are successively identifying the directions in the whitened state space where the observations are the most expressive relative to the prior and which are not in the same subspace as the previous directions. The matrix $\mathbf{C}_{\mathbf{x}} = \Sigma_{\mathbf{x}}^{1/2} \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} \Sigma_{\mathbf{x}}^{1/2} \in \mathbb{R}^{n \times n}$ can in fact be rewritten as the inner product of the whitened observation matrix introduced in 5.1.2.1, i.e. $\mathbf{C}_{\mathbf{x}} = \tilde{\mathbf{H}}^\top \tilde{\mathbf{H}}$. It is straightforward to show that the vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_d\}$ are the d eigenvectors associated with the d largest eigenvalues of the positive semi-definite matrix $\mathbf{C}_{\mathbf{x}}$. The following $n - d$ eigenvectors form an orthonormal basis for the non-informative subspace \mathbf{V}^\perp . The interpretation of the vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ as the eigenvectors of the matrix $\mathbf{C}_{\mathbf{x}}$ is an important step in generalizing to the setting with nonlinear observational models. In the nonlinear and non-Gaussian setting, Cui et al. [33] showed that the most important assimilation directions (for any realization of the observations) in the whitened state space can be identified by the eigenvectors of the *state space Gramian*:

$$\mathbf{C}_{\mathbf{x}} = \int \left(\Sigma_{\boldsymbol{\varepsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{x}}^{1/2} \right)^\top \left(\Sigma_{\boldsymbol{\varepsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{x}}^{1/2} \right) d\pi_{\mathbf{x}}(\mathbf{x}) \in \mathbb{R}^{n \times n}, \quad (5.18)$$

where the expectation is taken over the prior distribution. As expected, eq. (5.18) reverts to $\tilde{\mathbf{H}}^\top \tilde{\mathbf{H}}$ in the linear-Gaussian case. As before, $\mathbf{C}_{\mathbf{x}}$ is positive semi-definite and its eigendecomposition can be written as

$$\mathbf{C}_{\mathbf{x}} = \mathbf{V} \Lambda_{\mathbf{x}}^2 \mathbf{V}^\top, \quad (5.19)$$

where $\mathbf{V} \in \mathbb{R}^{n \times n}$ is an orthogonal basis for the whitened state space with associated eigenvalues $\Lambda_{\mathbf{x}}^2 \in \mathbb{R}^{n \times n}$.

To the best of our knowledge, there is no proved procedure to identify the most important directions of the observation space for a nonlinear observation model. We propose a heuristic inspired from the construction of eq. (5.18) for the state space that reverts to the columns of the orthogonal matrix \mathbf{U} in the linear-Gaussian case. It is reasonable to look for the directions in the observation space that maximize the relative ratio of the signal to the

observational noise. In the linear-Gaussian case, we can form the Rayleigh quotient \mathcal{T} that conveys this comparison as

$$\mathcal{T}(\mathbf{q}) = \frac{\langle \mathbf{q}, \boldsymbol{\Sigma}_{\mathbf{Y}} \mathbf{q} \rangle}{\langle \mathbf{q}, \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}} \mathbf{q} \rangle} - 1 = \frac{\langle \mathbf{q}, (\boldsymbol{\Sigma}_{\mathbf{Y}} - \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}) \mathbf{q} \rangle}{\langle \mathbf{q}, \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}} \mathbf{q} \rangle}, \text{ for } \mathbf{q} \in \mathbb{R}^d. \quad (5.20)$$

From eq. (5.4), we have $\boldsymbol{\Sigma}_{\mathbf{Y}} = \mathbf{H} \boldsymbol{\Sigma}_{\mathbf{X}} \mathbf{H}^\top + \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}$. With the change of variable $\mathbf{u} = \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \mathbf{q} \in \mathbb{R}^d$, we obtain:

$$\tilde{\mathcal{T}}(\mathbf{u}) = \frac{\langle \mathbf{u}, \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \mathbf{H} \boldsymbol{\Sigma}_{\mathbf{X}} \mathbf{H}^\top \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \mathbf{u} \rangle}{\langle \mathbf{u}, \mathbf{u} \rangle} \quad (5.21)$$

It is easy to show that the directions that maximize the ratio in (5.21) are the eigenvectors of the matrix $\mathbf{C}_{\mathbf{Y}} = \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \mathbf{H} \boldsymbol{\Sigma}_{\mathbf{X}} \mathbf{H}^\top \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} = \tilde{\mathbf{H}} \tilde{\mathbf{H}}^\top \in \mathbb{R}^{d \times d}$. The eigenvectors of $\mathbf{C}_{\mathbf{Y}}$ are also the column vectors of the matrix \mathbf{U} introduced in the previous section. Inspired by the treatment in the state space, we propose to use the eigenvectors of the *observation space Gramian* $\mathbf{C}_{\mathbf{Y}}$ to select the important assimilation directions in the whitened observation space:

$$\mathbf{C}_{\mathbf{Y}} = \int \left(\boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right) \left(\boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)^\top d\pi_{\mathbf{X}}(\mathbf{x}) \in \mathbb{R}^{d \times d}. \quad (5.22)$$

The matrix $\mathbf{C}_{\mathbf{Y}}$ is also PSD and has the eigendecomposition $\mathbf{C}_{\mathbf{Y}} = \mathbf{U} \boldsymbol{\Lambda}_{\mathbf{Y}}^2 \mathbf{U}^\top$, where $\mathbf{U} \in \mathbb{R}^{d \times d}$ is an orthonormal basis for the whitened observation space with associated eigenvalues $\boldsymbol{\Lambda}_{\mathbf{Y}}^2 \in \mathbb{R}^{d \times d}$. For convenience, we assume that the eigenvectors of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ are ordered by decreasing eigenvalues. For a nonlinear observation model, the eigenvalues of the state and observation Gramians can be different. In practice, we use Monte-Carlo approximations of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ that are estimated using the prior samples to identify the important subspaces. Depending on the inference problem, the Jacobian $\nabla \mathbf{h} \in \mathbb{R}^{d \times n}$ of the observation operator \mathbf{h} with respect to the state components can be computed either analytically, with automatic differentiation, with complex step differentiation, or, in the worst case, from finite differences. To perform the low-rank assimilation, we only retain the first $r_{\mathbf{X}} \leq n$ eigenmodes for the state space, and $r_{\mathbf{Y}} \leq d$ eigenmodes for the observation space. The ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ can be tuned independently based on the decay of $\boldsymbol{\Lambda}_{\mathbf{X}}^2$ and $\boldsymbol{\Lambda}_{\mathbf{Y}}^2$; typically we recommend setting these ranks to achieve a threshold $\alpha \in [0, 1]$ for the cumulative normalized energy of the

eigenvalue spectra, e.g. in the state space, we set $r_{\mathbf{X}} = \min \{r_{\alpha} \in \llbracket 1, n \rrbracket \mid E_{\mathbf{X}, r_{\alpha}} \geq \alpha\}$, where $E_{\mathbf{X}, r_{\alpha}} = \sum_{i=1}^{r_{\alpha}} \lambda_{\mathbf{X}, i}^2 / \sum_{i=1}^n \lambda_{\mathbf{X}, i}^2$. The new projected state and projected observation variables are defined as

$$\begin{aligned}\check{\mathbf{X}} &= \mathbf{V}_{r_{\mathbf{X}}}^{\top} \tilde{\mathbf{X}} = \mathbf{V}_{r_{\mathbf{X}}}^{\top} \Sigma_{\mathbf{X}}^{-1/2} (\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}}) \in \mathbb{R}^{r_{\mathbf{X}}} \\ \check{\mathbf{Y}} &= \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \tilde{\mathbf{Y}} = \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \Sigma_{\mathbf{Y}}^{-1/2} (\mathbf{Y} - \boldsymbol{\mu}_{\mathbf{Y}}) \in \mathbb{R}^{r_{\mathbf{Y}}},\end{aligned}\tag{5.23}$$

where $\mathbf{V}_{r_{\mathbf{X}}}, \mathbf{U}_{r_{\mathbf{Y}}}$ denotes the first $r_{\mathbf{X}}, r_{\mathbf{Y}}$ columns of \mathbf{V}, \mathbf{U} , respectively. We should emphasize that the different eigenvectors tell us the importance of each state and observation component in the inference problem. In particular, one can exploit the observation modes to determine the most impactful sensors, e.g. in an optimal sensor placement procedure [16, 8]. We will address this issue in the nonlinear setting in future work.

We now outline how to obtain the factorization of the Kalman gain in the nonlinear setting. In the linear-Gaussian case, the Gramians in (5.18) and (5.22) revert to $\mathbf{C}_{\mathbf{X}} = \tilde{\mathbf{H}}^{\top} \tilde{\mathbf{H}}$ and $\mathbf{C}_{\mathbf{Y}} = \tilde{\mathbf{H}} \tilde{\mathbf{H}}^{\top}$, respectively. From the definition of the SVD of the whitened observation matrix $\tilde{\mathbf{H}}$, we have $\tilde{\mathbf{H}} \mathbf{v}_i = \lambda_i \mathbf{u}_i$ for some left singular vector \mathbf{u}_i , singular value λ_i and right singular vector \mathbf{v}_i . This results in the eigendecompositions $\mathbf{C}_{\mathbf{X}} \mathbf{u}_i = \lambda_i^2 \mathbf{u}_i$ and $\mathbf{C}_{\mathbf{Y}} \mathbf{v}_i = \lambda_i^2 \mathbf{v}_i$. Therefore, if the triplet $\{\mathbf{u}_i, \lambda_i, \mathbf{v}_i\}$ is obtained from the eigendecomposition of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$, the sign of the singular vectors is lost, i.e. we only have $\tilde{\mathbf{H}} \mathbf{v}_i = \pm \lambda_i \mathbf{u}_i$. For a nonlinear observation model, the factorization of the Kalman gain given in eq. (5.12) is no longer applicable and needs to be generalized. In the rotated and whitened space, the analysis map is given by:

$$\check{\mathbf{T}}(\check{\mathbf{y}}, \check{\mathbf{x}}) = \check{\mathbf{x}} - \check{\mathbf{K}}(\check{\mathbf{y}} - \check{\mathbf{y}}^*),\tag{5.24}$$

where $\check{\mathbf{K}} = \Sigma_{\check{\mathbf{X}}\check{\mathbf{Y}}} \Sigma_{\check{\mathbf{Y}}}^{-1} \in \mathbb{R}^{r_{\mathbf{X}} \times r_{\mathbf{Y}}}$ is the Kalman gain in the informative space with $\Sigma_{\check{\mathbf{X}}\check{\mathbf{Y}}} \in \mathbb{R}^{r_{\mathbf{X}} \times r_{\mathbf{Y}}}$ and $\Sigma_{\check{\mathbf{Y}}}^{-1} \in \mathbb{R}^{r_{\mathbf{Y}} \times r_{\mathbf{Y}}}$. In the original space, we get

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \mathbf{x} - \Sigma_{\mathbf{X}}^{1/2} \mathbf{V} \check{\mathbf{K}} \mathbf{U}^{\top} \Sigma_{\boldsymbol{\varepsilon}}^{-1/2} (\mathbf{y} - \mathbf{y}^*).\tag{5.25}$$

This new factorization of the Kalman gain nicely generalizes eq. (5.12) to an inference problem with a nonlinear observation model. We emphasize that the sign issue presented above

is obviated by absorbing it into the definition of $\check{\mathbf{K}}$. Let us remark that the definition of $\check{\mathbf{K}}$ reduces to the diagonal matrix $\mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I})^{-1}$ in the linear-Gaussian case, where the matrices \mathbf{U}, \mathbf{V} are obtained from the SVD of $\tilde{\mathbf{H}}$. Indeed, eq. (5.25) constitutes *a change of coordinates for the Kalman gain* between the original space and the informative space. We denote the algorithm applying the linear update in (5.25) in each analysis step as the low-rank EnKF (LREnKF). A pseudo-code for the proposed low-rank EnKF is presented in A.

We should emphasize that one can consider different bases to project the state and observation variables. Solonen et al. [60] used the eigenvectors of the covariance matrices for the state and observation variables to reduce the dimension of the inference problem. This option can be significantly cheaper than our proposed methodology, as it does not require evaluating the Jacobian of the observation operator. We detail several remedies to this limitation in the conclusion 6. However, it is not clear that the covariance matrices will have a low-rank structure if the dynamical model does not have an intrinsic low dimension. Furthermore, for given ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$, these bases give a suboptimal low-rank approximation of the Kalman gain, compared to our proposed bases. Therefore, a larger number of modes will be required to obtain the same error on the Kalman gain. By estimating a larger number of components with a limited budget of samples, additional sampling errors are created in the posterior estimate.

In comparison, the sEnKF estimates the Kalman gain in the original space, a linear operator of dimensions $n \times d$, where n and d are the dimensions of the space and observation spaces. Leveraging the dimension reduction offered by the state and observation Gramians, the LREnKF only has to estimate the Kalman gain in the informative subspaces of \mathbf{X} and \mathbf{Y} , whose dimension is only $r_{\mathbf{X}} \times r_{\mathbf{Y}}$, where $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ can be significantly smaller than n and d , respectively. To appreciate the benefit of this dimension reduction in the inference problem, it is useful to recall the classical bias-variance trade-off in any machine learning problem [61]: Given an ensemble size M , the error in the estimation of the Kalman gain \mathbf{K} can be decomposed into a variance term and a bias term. For a limited ensemble size

(compared to n and d), the estimated Kalman gain from the vanilla sEnKF will have a large variance, potentially leading to a filter with diverging state estimation error, as is the case for small ensemble size in the example problems. We show later in this chapter that with the proposed dimension reduction, the Kalman gain in the informative subspaces $\check{\mathbf{K}}$ contains a much smaller number of entries than \mathbf{K} . This results in smaller variance. The bias in computing $\check{\mathbf{K}}$ can be controlled by setting the ranks of the state and observation projections. Due to the rapid decay of the spectrum of the state and observation Gramians, increasing the rank of the projected subspace beyond a certain ratio of the cumulative energy will only marginally reduce the bias, but greatly increases the variance. In other words, we choose the ranks $r_{\mathbf{X}}, r_{\mathbf{Y}}$ to capture the column/row space of the Kalman gain and make the bias small for any finite-sample-estimator of \mathbf{K} .

5.1.3 Estimation of the leading directions from samples

In practice, we only have access to limited samples $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ from the prior distribution to estimate the state and observation Gramians $\mathbf{C}_{\mathbf{X}}, \mathbf{C}_{\mathbf{Y}}$ and their associated eigenvectors. One can show that the empirical state and observation Gramians, namely $\hat{\mathbf{C}}_{\mathbf{X}}$ and $\hat{\mathbf{C}}_{\mathbf{Y}}$, converge to the true Gramians at the rate $O(1/\sqrt{M})$. However, this result is not sufficient to justify that we can accurately estimate the leading eigenvectors from these empirical Gramians. To justify this more rigorously, we recall this corollary of the Davis-Kahan theorem [62, 63]:

Theorem 5.1.1 (Corollary of the Davis-Kahan theorem [62, 63]¹) *Let $\mathbf{\Gamma}$ and $\hat{\mathbf{\Gamma}}$ be $m \times m$ positive semi-definite matrices with eigendecompositions $\mathbf{\Gamma}\mathbf{w}_i = \lambda_i^2\mathbf{w}_i$ and $\hat{\mathbf{\Gamma}}\hat{\mathbf{w}}_i = \lambda_i^2\hat{\mathbf{w}}_i$ where the eigenvalues are ordered by decreasing order: $\lambda_1^2 \geq \lambda_2^2 \geq \dots \geq \lambda_m^2 \geq 0$ and $\hat{\lambda}_1^2 \geq \hat{\lambda}_2^2 \geq \dots \geq \hat{\lambda}_m^2 \geq 0$. For $1 \leq r < m$ with $\lambda_r^2 - \lambda_{r+1}^2 > 0$, we define $\mathbf{W}_r = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r) \in \mathbb{R}^{m \times r}$ and $\hat{\mathbf{W}}_r = (\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2, \dots, \hat{\mathbf{w}}_r) \in \mathbb{R}^{m \times r}$. Then, the distance between the subspaces spanned by the columns of \mathbf{W}_r and $\hat{\mathbf{W}}_r$ denoted $d(\mathbf{W}_r, \hat{\mathbf{W}}_r)$ satisfies*

$$d(\mathbf{W}_r, \hat{\mathbf{W}}_r)^2 := \|\mathbf{W}_r\mathbf{W}_r^\top (\mathbf{I} - \hat{\mathbf{W}}_r\hat{\mathbf{W}}_r^\top)\|_F^2 \leq \frac{C}{\sqrt{M}} \frac{1}{\lambda_r^2 - \lambda_{r+1}^2}, \quad (5.26)$$

where C is a strictly positive constant that depends only on Γ .

Let us now apply the Davis-Kahan theorem to the pairs $(\Gamma, \hat{\Gamma}) = (\mathbf{C}_x, \hat{\mathbf{C}}_x)$, and $(\mathbf{C}_y, \hat{\mathbf{C}}_y)$. Two important conclusions can be made from this result. First, the error in the subspace spanned by the first r eigenvectors also converges at the rate $O(1/\sqrt{M})$. This is a critical result, as a slower convergence rate would have been a serious limitation to estimate the informative directions with limited samples. Second, the error in the estimated subspace spanned by the first r eigenvectors of the two matrices depends on the difference between two consecutive eigenvalues λ_r, λ_{r+1} , namely $\lambda_r - \lambda_{r+1}$, which is called the spectral gap (or eigengap). Therefore, the inverse of the spectral gap is a useful indicator of the estimation error of the informative subspaces. Intuitively, two eigenvectors with close eigenvalues are difficult to distinguish. For the fluid problems considered in this chapter (see examples in Section 5.2), the spectra of the state and observation Gramians rapidly decay. Thus, the leading eigenvalues are well separated. This suggests that we can accurately capture the row and column spaces of the Kalman gain with a few directions in the state and observation spaces.

5.2 Examples

5.2.1 Multipole expansion and leading directions of the state and observation Gramians

In this example, we develop intuition for the ideas presented in this chapter by connecting the leading eigenvectors of the state and observation space Gramians discussed in eqns. (5.18) and (5.22) to the multipole expansion. We consider a set of N point sources located at $\{z_1, \dots, z_N\}$ with strengths $\{Q_1, \dots, Q_N\}$, and a set of d evaluation points located at $\{z'_1, \dots, z'_d\}$. We assume that the strengths are perfectly known, but we seek to estimate the positions of the singularities from observations of the velocity potential ϕ at the evaluation




Figure 5.1: Schematic of the setup. Orange dots depict the location of the point sources. The yellow square (located at $(0, 0)$) depicts the location of the centroid of the point sources. Grey diamonds depict the location of the evaluation points of the velocity potential ϕ

points corrupted by additive Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}_d, \boldsymbol{\Sigma}_\epsilon)$. To ensure a convergent multipole expansion, we place all the evaluation points outside a circle of radius R where $R = \max_J |z_J|$. We use the following notations: $\mathbf{x} = [z_1, \dots, z_N]^\top \in \mathbb{C}^N$, $\mathbf{S} = [Q_1, \dots, Q_N]^\top \in \mathbb{R}^N$ and $\mathbf{z}' = [z'_1, \dots, z'_d]^\top \in \mathbb{C}^d$. The superscript \top denotes the transpose of a complex vector/matrix without conjugation, while the superscript \mathbf{H} denotes the transpose with conjugation of a complex vector/matrix.

The i th component of the observation vector $\mathbf{y} \in \mathbb{R}^d$ is given by $y_i = \phi(\xi_i; \mathbf{x}) + \epsilon_i, i = 1, \dots, d$, where the notation $\phi(z'_i; \mathbf{x})$ is used to highlight the dependence on the positions of the point sources. We seek to estimate the position vector \mathbf{x} from the noisy observations \mathbf{y} . The observation model reads in vector form as $\mathbf{y} = \mathbf{h}(\mathbf{x}) + \boldsymbol{\epsilon}$, where $\mathbf{h}: \mathbb{C}^N \rightarrow \mathbb{R}^d$, $\mathbf{x} \mapsto [\phi(z'_1; \mathbf{x}), \dots, \phi(z'_d; \mathbf{x})]^\top$. From section 5.1.2, the informative subspaces are identified from the Jacobian of the observation operator \mathbf{h} with respect to the state variable \mathbf{x} , denoted $\nabla_{\mathbf{x}} \mathbf{h}(\mathbf{x})$. If $|z_J|/|z'_k| < 1$ for all k, J , then we get from (2.1)

$$\nabla_{\mathbf{x}} \mathbf{h} = \frac{1}{4\pi} \begin{bmatrix} \frac{-Q_1}{z'_1 - z_1} & \cdots & \frac{-Q_N}{z'_1 - z_N} \\ \vdots & & \vdots \\ \frac{-Q_1}{z'_d - z_1} & \cdots & \frac{-Q_N}{z'_d - z_N} \end{bmatrix} \approx \underbrace{\begin{bmatrix} \frac{1}{z'_1} & \cdots & \frac{1}{z'_1{}^p} \\ \vdots & & \vdots \\ \frac{1}{z'_d} & \cdots & \frac{1}{z'_d{}^p} \end{bmatrix}}_{\mathbf{A}_{\mathbf{z}'}} \underbrace{\begin{bmatrix} 1 & \cdots & 1 \\ z_1 & \cdots & z_N \\ \vdots & & \vdots \\ z_1^p & \cdots & z_n^p \end{bmatrix}}_{\mathbf{B}_{\mathbf{x}}^\top} \text{diag} \left(-\frac{\mathbf{S}}{4\pi} \right), \quad (5.27)$$

where the approximation comes from truncating the multipole expansion at the p -th order. This factorization decouples the contribution of the position of the evaluation points ($\mathbf{A}_{\mathbf{z}'}$), the position of the point sources ($\mathbf{B}_{\mathbf{x}}^\top$), and the volume fluxes of the point sources ($\text{diag}(-\mathbf{S})$). Unfortunately, it is not possible to compute analytically the Gramians $\mathbf{C}_{\mathbf{x}}, \mathbf{C}_{\mathbf{y}}$ for this inference problem, even with a Gaussian distribution for the x and y coordinates of the point source locations. To highlight the connection between the eigenvectors of these matrices and the multipole expansion, we compare the singular vectors obtained from the SVD of $\nabla_{\mathbf{x}}\mathbf{h}$, with the orthonormal bases obtained by orthonormalization of the matrices $\mathbf{A}_{\mathbf{z}'}, \mathbf{B}_{\mathbf{x}}$ for a particular position vector \mathbf{x} . We recall that the left singular vectors \mathbf{u}_i of $\nabla_{\mathbf{x}}\mathbf{h}$ are also the eigenvectors of the Gram matrix $\nabla_{\mathbf{x}}\mathbf{h}^H\nabla_{\mathbf{x}}\mathbf{h}$, and similarly, the right singular vectors \mathbf{v}_i of $\nabla_{\mathbf{x}}\mathbf{h}$ are also the eigenvectors of the Gram matrix $\nabla_{\mathbf{x}}\mathbf{h}\nabla_{\mathbf{x}}\mathbf{h}^H$. To compute an orthonormal basis for the columns of $\mathbf{A}_{\mathbf{z}'}$ and $\mathbf{B}_{\mathbf{x}}$, we extract the \mathbf{Q} factor of the QR factorizations of these matrices, denoted by $\mathbf{Q}_{\mathbf{z}'} \in \mathbb{C}^{d \times p}$ and $\mathbf{Q}_{\mathbf{x}} \in \mathbb{C}^{n \times p}$, respectively.

We place a set of $N = 16$ point sources distributed according to $z_J = \rho \exp(2\pi i J/N)$ for $J = 1, \dots, N$, where $\rho = 0.3$. The volume flux of the point sources is set to 4π . The centroid of this set of point sources is at the origin. This is critical for a sensible comparison with the factorization of eq. (5.27). We use $d = 50$ evaluation points distributed along a horizontal line as $z'_k = -10 + (k-1)\Delta s + 4i$, for $k = 1, \dots, d$ with interspace $\Delta s = 20/(d-1)$. Fig. 5.1 depicts the configuration of the different elements. We truncate the multipole expansion at the $p = 30$ th order to get a machine precision approximation of $\nabla_{\mathbf{z}}\mathbf{h}$ in eq. (5.27).

Fig. 5.2 (a) compares the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th modes obtained from the right singular vectors of $\nabla_{\mathbf{x}}\mathbf{h}$, and from the corresponding columns of the unitary matrix $\mathbf{Q}_{\mathbf{x}}$. Fig. 5.2 (b) compares the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th modes obtained from the left singular vectors of $\nabla_{\mathbf{x}}\mathbf{h}$, and from the corresponding columns of the unitary matrix $\mathbf{Q}_{\mathbf{z}'}$. Overall, the state and observation modes obtained from these two procedures agree well, and particularly for the first three modes. The first state mode \mathbf{v}_1 is $\pm \mathbf{1}/\sqrt{N}$, where $\mathbf{1}$ denote a vector of ones of length N . Therefore, the leading

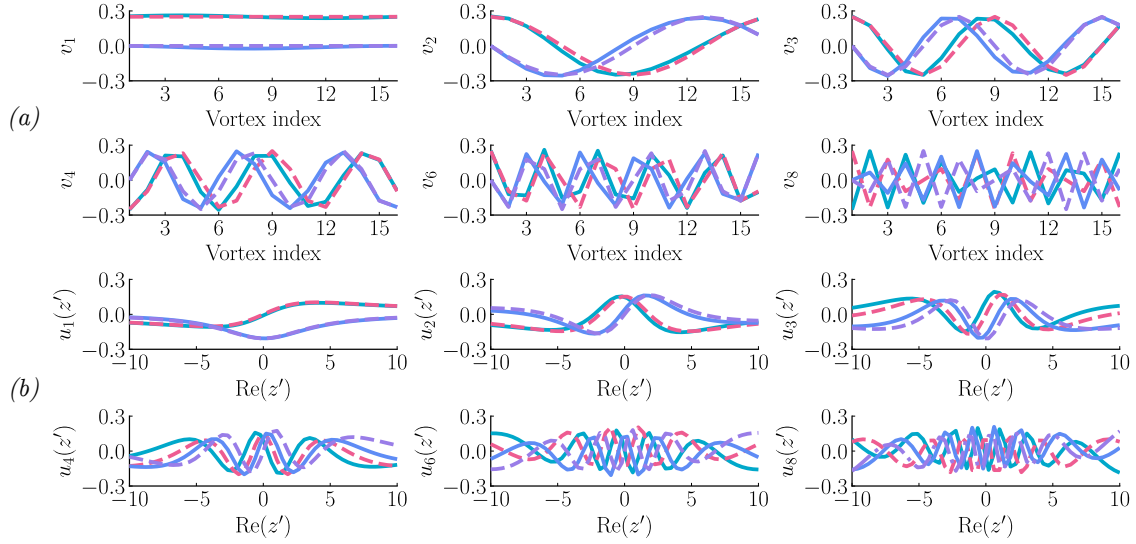


Figure 5.2: Top panel (a): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th state modes (\mathbf{v}_i s) obtained from the SVD of $\nabla_{\mathbf{x}}\mathbf{h}$ (solid turquoise and solid blue lines), and from orthogonalization of the columns of $\mathbf{Q}_{\mathbf{x}}$ (dashed red and dashed purple lines). Lower panel (b): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th observation modes (\mathbf{u}_i s) obtained from the SVD of $\nabla_{\mathbf{x}}\mathbf{h}$ (solid turquoise and solid blue lines), and from the columns of $\mathbf{Q}_{z'}$ (dashed red and dashed purple lines). The vectors have unit norm.

state correction corresponds to a rigid body translation of the point sources. Three modes capture more than 99.999999% of the cumulative energy of $\nabla_z \mathbf{h}^H \nabla_z \mathbf{h}$. The differences in the vectors obtained from these two approaches may be attributed to $\mathbf{Q}_{z'}$ only relying on the positions of the evaluation points, and \mathbf{Q}_x only relying on the positions of the point sources. In contrast, the SVD simultaneously has access to the locations of the evaluation points and the point sources to identify the left and right singular vectors. As a result, the SVD constructs an *optimized* “multipole expansion” of the Jacobian (see the Eckart-Young theorem [64]), specifically designed for the particular set of the evaluation points and the point sources.

5.2.2 Inference of the properties of point vortices from pressure observations along a wall

In this second example, we compare the LREnKF with the sEnKF for estimating the positions and strengths of a collection of point vortices advecting along a horizontal wall. We rely on pressure observations collected along the wall to estimate the state. These observations were generated from the same observation model used for inference, thereby making this a *twin experiment* [13]. We consider a set of $N = 5$ point vortices located at $\{z_1, \dots, z_N\}$ with circulations $\{\Gamma_1, \dots, \Gamma_N\}$, and a set of d evaluation points located at $\{z'_1, \dots, z'_d\}$, see Fig. 5.3 (a). The state variable \mathbf{x} contains the positions and circulations of the N point vortices: $\mathbf{x} = [x_1, y_1, \Gamma_1, \dots, x_N, y_N, \Gamma_N]^T \in \mathbb{R}^{3N}$. To avoid singular interactions between nearby vortices, we replace the singular Cauchy kernel $k(z) = 1/(2\pi z)$ used to compute the Kirchhoff velocities w_{-j} in (2.4) and (2.9), with the regularized algebraic blob kernel $k_\epsilon(z) = \bar{z}/(2\pi(|z|^2 + \epsilon^2))$, where ϵ is called the blob radius [6], set here to 5×10^{-2} . To enforce the no-flow-through condition along the x axis, we use the method of images and augment our collection of vortices with another set of N vortices at the conjugate positions $\{\bar{z}_1, \dots, \bar{z}_N\}$ with opposite circulation $\{-\Gamma_1, \dots, -\Gamma_N\}$. We emphasize that these mirrored vortices are only an artifice to enforce the no-flow-through in the forecast step, and play no role in the analysis step. We

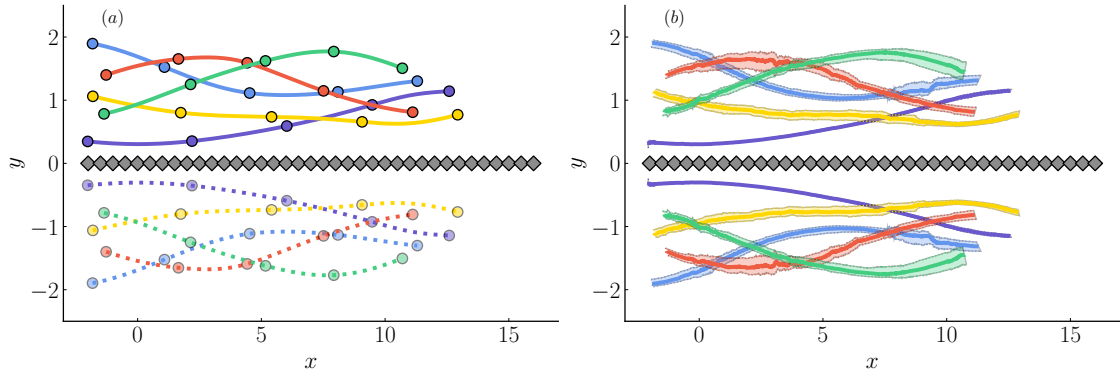


Figure 5.3: (a): Schematic of the setup. True trajectories of the point vortices (for one realization of the system) are represented by colored lines. The location of the point vortices, sampled every three convective times, is depicted by colored dots. Fainter dashed lines and fainter dots refer to the image point vortices. Location of the pressure sensors are depicted by grey diamonds. (b): Estimation of the trajectories of the vortices with the LREnKF for $M = 40$ with the ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ set to capture 99% of the cumulative energy spectra. Solid lines depict the time history of the median posterior estimate for the position of the different point vortices. Fainted areas show the 5% and 95% quantiles of the posterior estimate for the position of the point vortices.




Figure 5.4: Median spectrum of the state Gramian $\mathbf{C}_{\mathbf{X}}$ (left panels (a)-(c)) and observation Gramian $\mathbf{C}_{\mathbf{Y}}$ (right panels (d)-(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles.

add a freestream flow directed along increasing x values with velocity $\mathbf{U}_{\infty} = [1, 0]^{\top}$. The initial position of the J th point vortex is generated randomly with the form $z_J + \rho_r \exp(i\theta)$ where $\{z_1, \dots, z_N\} = \{-2.0 + 0.3i, -1.9 + 1.9i, -1.8 + 1.1i, -1.3 + 1.4i, -1.4 + 0.8i\}$, ρ_r is drawn from $\mathcal{N}(0.0, 0.1)$ and θ is drawn from the uniform distribution on $[0, \pi]$. The initial circulation of the point vortices is drawn from $\mathcal{N}(0.4, 0.1)$. The collection of point vortices is advanced using a forward Euler scheme with time step $\Delta t = 1 \times 10^{-3}$. The observation vector \mathbf{y} consists of pressure observations at $d = 37$ locations linearly distributed with an interspace of 0.5 along the segment $[-2, 16]$, see Fig. 5.3 (a). We use a Gaussian observation noise with zero mean and covariance 1×10^{-4} . We obtain the Jacobian of the observation operator used in the state and observation Gramians $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ by analytical differentiation of eq. (2.9).

We assess the performance of the two filters with a *twin experiment*. We draw a random initial condition for the positions and strengths of the different point vortices. Then we




Figure 5.5: 1, 2, 3 and 8th eigenvectors of the observation Gramian $\mathbf{C}_\mathbf{y}$, i.e. observation modes \mathbf{u}_i s, at $t = 0.1$ (left column), $t = 7.0$ (middle column), and $t = 10.0$ (right column). (a)-(c): first mode, (d)-(f): second mode, and (g)-(i): third mode, (j)-(l): eighth mode. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. The grey horizontal line corresponds to 0 ordinate. Statistics are obtained from a run of the sEnKF with $M = 1000$. The eigenvectors have unit norm.

simulate the dynamics of the point vortices over the time interval $[0, 12]$ (corresponding to 12000 assimilation cycles with the choice of time step Δt). The true state at the time step k is denoted by \mathbf{x}_k^* . At every time step, we generate the noisy pressure observation at the sample locations based on the true state \mathbf{x}_k^* . The realization of the true observation vector at the time step is denoted by \mathbf{y}_k^* . In the twin experiment, we seek to estimate the true state \mathbf{x}_k^* only from the knowledge of the noisy and indirect observations \mathbf{y}_k^* and the prior.

First, we assess the spectra and leading subspaces of the state and observation Gramians. Fig. 5.4 depicts the median eigenvalues of the state and observation Gramians over the time interval $[0, 12]$. We run the sEnKF with ensemble size $M = 1000$ to generate the samples used to compute empirical Gramians in each analysis step. We choose this larger ensemble size to reduce sampling errors. The median ranks $r_\mathbf{x}$ to capture 80%, 90%, 95%, and 99% of the cumulative energy of $\mathbf{C}_\mathbf{x}$ are 1, 2, 3, and 5, respectively. Similarly, the median

ranks $r_{\mathbf{Y}}$ to capture 80%, 90%, and 99% of the cumulative energy of $\mathbf{C}_{\mathbf{Y}}$ are 1, 2, and 4, respectively. The sharp decay of the spectra of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ supports our hypothesis of low-rank structure existing in the prior-to-posterior update. A low-dimensional subspace of the pressure observations is only informative along a limited number of directions in the state space. Fig. 5.5 shows the 1, 2, 3 and 8th eigenvectors of the observation Gramian at the three times: $t = 0.1, 7.0,$ and 10.0 . These modes clearly illustrate how information is extracted from the different pressure discrepancies during the analysis step. The first observation mode corresponds to a spatially weighted average of the pressure discrepancies about the mean x coordinate of the position of the vortices (depicted by the dashed vertical line on the panels of Fig. 5.5). The second mode captures differences between the pressure observations collected on the left and right side of the centroid location. The higher modes follow the same pattern and act as refined stencils with growing support to extract higher order features from the pressure observations. The leading eigenvectors of the state Gramian are more difficult to interpret, due to the manner in which vortices are coupled in the determination of pressure.

Next, we compare the performance of the LREnKF with the stochastic EnKF (sEnKF). The comparison of these two filters is natural, as the LREnKF without dimension reduction (i.e. for $r_{\mathbf{X}} = n = 3N$ and $r_{\mathbf{Y}} = d$) reverts to the sEnKF. For a given ensemble size M , one could perform a parametric study to determine the ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ which give the best performance of the LREnKF. We found, however, that using fixed ranks for all assimilation cycles is suboptimal as we face one of two possible scenarios: either the ranks are too large, leading to additional variance, or the ranks are too small, leading to additional bias. Furthermore, we expect the dimension of the informative subspaces may vary over time as the flow field evolves. Instead, we select the ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ adaptively at each time step based on a predetermined ratio of the cumulative energy of the Gramians $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. The performance of each filter is assessed using the *root mean squared error (RMSE)* for tracking the true state \mathbf{x}_k^* . The best estimate of the true state at a given assimilation time is given by the mean of the analysis ensemble at this step, that we denote by $\bar{\mathbf{x}}_k^a$. We define the RMSE




Figure 5.6: Left column (a): Time-averaged evolution of the median RMSE with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The tracking performance of the sEnKF is unstable for $M < 40$. Dashed lines depict the 25% and 75% quantiles. Right column [(b)-(c)]: Time-history of the median ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF for $M = 30$ (computed over 50 realizations) for different ratios of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$, $\mathbf{C}_{\mathbf{y}}$: 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue.

at one time between the true state \mathbf{x}_k^* and the estimate $\bar{\mathbf{x}}_k^a$ as $\text{RMSE}_k = \|\mathbf{x}_k^* - \bar{\mathbf{x}}_k^a\|_2 / \sqrt{3N}$, where N is the number of point vortices. The mean RMSE of each filter is computed over the time interval $[8, 12]$ (i.e. the last 4000 assimilation steps), to remove any influence of the initial conditions and to ensure empirical stationarity of the filters. The first 8000 assimilation steps are called the spin-up phase and are discarded. The uncertainty in the RMSE is quantified by the 5%, 25%, 75% and 95% quantiles of the time-averaged RMSE over 50 realizations of the same experiment for different initial ensembles.

Fig. 5.6(a) shows the time-averaged RMSE for the LREnKF and the sEnKF. We assess the performance of the LREnKF for 85%, 95%, and 99% of the normalized cumulative energy of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$. For $M < 40$, the sEnKF is unstable, while the LREnKF accurately estimates the true state, even for $M = 10$. For $M \in [40, 60]$, the RMSE of the sEnKF significantly increases as M decreases. For $M = 40$, the median RMSE of the sEnKF is 0.9, while the median RMSEs of the LREnKF with the different energy ratios are smaller than 0.16. For large M , the RMSE of the LREnKF decreases with the energy ratio, but remains larger than the RMSE of the sEnKF. The proposed low-rank approximation of the Kalman gain can lead to a small bias for large ensemble sizes. The regularization technique proposed in this work, however, is designed for the small ensemble size regime; one can always increase the energy ratio (beyond 99%) to recover the performance of the sEnKF for large ensemble sizes. The RMSE of the LREnKF with the different energy ratios decreases for $M \in [10, 40]$, then plateaus. In the small ensemble size regime, adding samples improves the estimate of the dominant directions (see the Davis-Kahan theorem 5.1.1), leading to a reduction in variance. As we continue to increase the number of samples for a given number of dimensions to estimate, the variance becomes smaller than the bias from using a truncated basis to approximate the Kalman gain. For $M = 10$, the LREnKF with 85% energy ratio performs better than the LREnKF with 99% energy ratio. For $M = 10$, the ranks needed to achieve 99% is large, leading to additional variance. On the other hand, 85% energy requires a much smaller number of entries to estimate, leading to less variance compared to

the bias of this approximation. We recommend choosing the target energy ratio based on the available ensemble size.

Fig. 5.6(b)-(c) show the history of the median ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ (computed over 50 realizations) for the different ratios of the cumulative energy for $M = 30$. The ranks plateau over the time window and remain small compared to the dimension of the state and observation spaces. For the LREnKF with 99% of the cumulative energy, $r_{\mathbf{X}}, r_{\mathbf{Y}}$ are smaller than 8, 6, respectively. Fig. 5.3 (b) shows one estimate of the trajectories of the vortices over the time interval $[0, 12]$. The results are obtained from the LREnKF with $M = 50$ with the ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ set to capture 99% of the energy spectra. We quantify the uncertainty with the 5% and 95% quantiles of the posterior for the point vortex positions. We observe an excellent agreement with the true trajectories with a time-averaged RMSE of 0.038.

In Section 5.1.2, we argue that distance localization schemes will be harmful to regularize inference problems with elliptic observation models, as studied herein. We conclude this section by providing an *a posteriori* justification based on the results of the sEnKF run with a large ensemble ($M = 1000$). Fig. 5.7 (a)-(c) shows the magnitude of the cross-covariance between the positions and strengths of the different point vortices and the pressure observations at $t = 1.0$. We notice that the cross-covariance entries decay very quickly on a small support about each point vortex. Past a certain radius, however, the cross-covariance entries quickly increase, and beyond this point, they only decay algebraically. More precisely, we observe a decay as the inverse square of the distance between the point vortex and the pressure observation. This decay rate is expected from the pressure field derived in (2.9). In Fig. 5.7(c), the covariance between the circulation of three point vortices and pressure observations at a distance of 5 and 10 units away is still about 20% and 10% of its maximal value, respectively. This algebraic decay clearly violates the assumption of rapidly decaying correlations of distance localization schemes. The cross-covariance of the components of the position and the strength for the different point vortices can have different algebraic decay rates as a function of the distance. Even for a particular vortex, the variations of the cross-

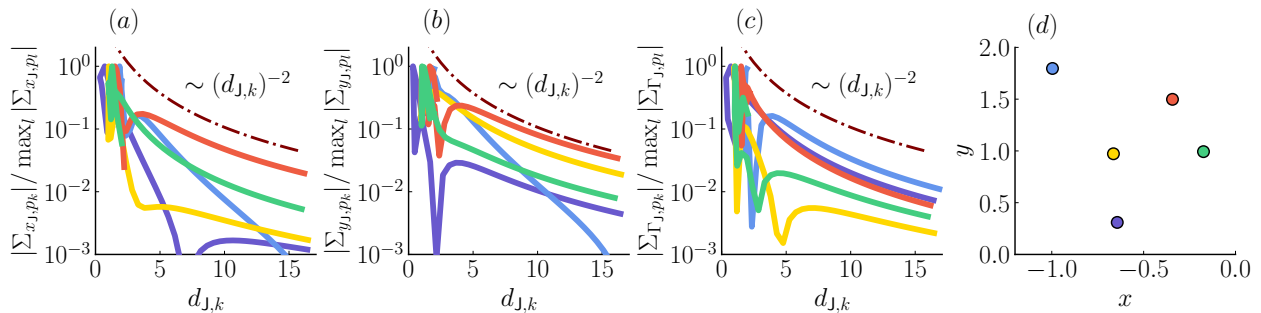


Figure 5.7: Panel (a), (b), (c) depicts the magnitude of the empirical cross-covariance between the x , y coordinate and strength of the J point vortex with the k th pressure observation at $t = 1.0$, respectively. The curves are plotted against the distance $d_{J,k} = |z_J - z'_k|$, where $z_J = x_J + iy_J$ is the (complex) position of the J th point vortex, and z'_k is the (complex) location of the k th pressure sensor. The magnitude of the cross-covariances are normalized by the maximum cross-covariance (in magnitude) between the x , y coordinate or strength of the J point vortex with the different pressure observation, respectively. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. Red dashed and dotted curve depicts the algebraic decay $O(1/d_{J,k}^2)$. (d) Mean position of the different point vortices at $t = 1.0$. The same color is used to depict the properties of a point vortex on the panels (a)-(d). The results are obtained from a run of the sEnKF for $M = 1000$.

covariance for different components of the position and the strength are drastically different. In unreported results, we observe that the decay rate of the correlations also varies over time as the vortices evolve. In the best case, this suggests that three different localization radius are needed and would have to be tuned independently: an impractical procedure. These different facts support our *a priori* hypothesis that distance localization is not suited to regularize filtering problems with elliptic observation operators.

5.2.3 Inference of a vortex patch advected along a wall from pressure observations.

In this third example, we look at the inference of a circular vortex patch advected along the x axis. A circular vortex patch \mathcal{V} is a disk of uniform vorticity ω_0 . As in 5.2.2, the vortex patch is constrained to advect along the x axis by the method of images [6]. The total circulation of the vortex patch \mathcal{V} , denoted $\Gamma_{\mathcal{V}}$, is given by $\Gamma_{\mathcal{V}} = \int_{\mathbf{x} \in \mathcal{V}} \omega_0 \, dS(\mathbf{x})$ and the vorticity centroid, denoted $\mathbf{x}_{\mathcal{V}}$, is given by $\mathbf{x}_{\mathcal{V}} = \Gamma_{\mathcal{V}}^{-1} \int_{\mathbf{x} \in \mathcal{V}} \mathbf{x} \omega_0 \, dS(\mathbf{x})$. The problem is parameterized by the ratio of the radius $r_{\mathcal{V}}$ of the vortex patch to the distance $d_{\mathcal{V}}$ between the centroid of the vortex patch and its image centroid.

We discretize each vortex patch with a collection of regularized point vortices; further details are given in [6]. In this study, the vortex patch is discretized by $N_r = 4$ concentric rings of blobs, leading to a subdivision of the vortex patch into $N_{pv} = 1 + 4N_r(N_r - 1) = 49$. The initial position of the J th point vortex is generated randomly with the form $z_J + \rho_r \exp(i\theta)$, where z_J denotes the nominal position of the J th vortex. ρ_r is drawn from $\mathcal{N}(0.0, \sigma_r^2)$, where σ_r corresponds to 10% of the radius between two concentric rings of point vortices, and θ is drawn from the uniform distribution on $[0, \pi]$. The initial circulation of the point vortices is drawn from $\mathcal{N}(\Gamma_{\mathcal{V}}/N_{pv}, 10^{-4})$. For a blob at z_J with circulation Γ_J in the vortex patch \mathcal{V} , there is an image blob located at \bar{z}_J with circulation $-\Gamma_J$. The state is defined, as in the previous example, by the positions and strengths of the point vortices in the upper half-plane. For N_{pv} blobs, the state dimension is $n = 3N_{pv} = 147$. We use a forward Euler model scheme with time step $\Delta t = 5.0 \times 10^{-3}$. The vortex patches are evolved over the time interval $[0, 12]$. The observation vector \mathbf{y} consists of pressure observations collected at $d = 24$ locations with interspace 0.5 along the segment $[-1.5, 10.0]$ of the x axis. The observation noise is Gaussian with zero mean and covariance 4×10^{-2} (corresponding to a ratio of peak pressure amplitude to standard deviation of the noise equal to 10).

In this example, the forward model that generated the observations is given by solving

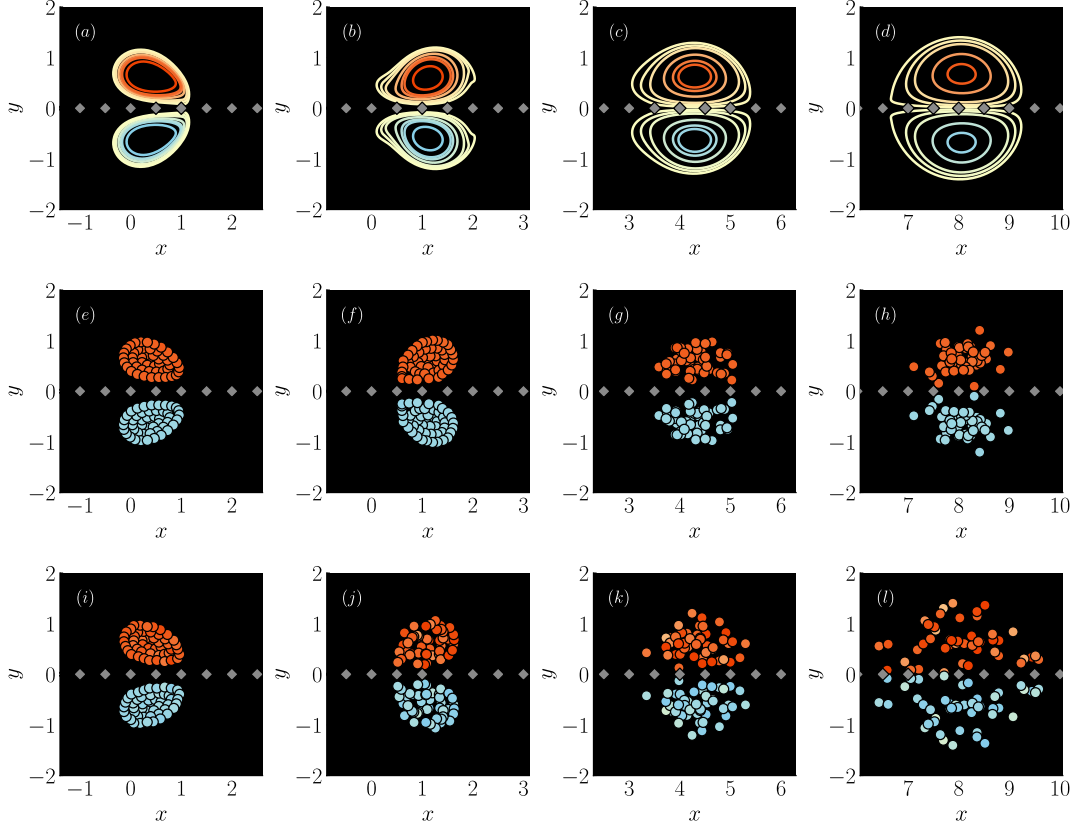


Figure 5.8: Snapshots of the vorticity distribution at $t = 0.5$ (left column), $t = 1.5$ (second column), $t = 6.0$ (third column), and $t = 12.0$, for a vortex patch and its image, predicted from [(a)-(d)] high-fidelity numerical simulation at Reynolds number 1000, [(e)-(h)] inviscid vortex model with LREnKF with $M = 30$, and [(i)-(l)] inviscid vortex model with sEnKF with $M = 30$. The ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF are set to capture 99% of the normalized cumulative energy of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$, respectively. Pressure sensors are depicted with grey diamonds. Orange denote positive vorticity, while blue vorticity denotes negative vorticity.

the incompressible Navier-Stokes equations at a Reynolds number $\Gamma_{\mathcal{V}}/\nu = 1000$, where ν is the kinematic viscosity of the fluid. The true solution contains viscous effects, which are not explicitly included in the forecast model we use for inference. The initial condition is constructed by discretizing the two vortex patches on a uniform Cartesian grid with grid spacing $\Delta x = 0.01$ for the same geometry (i.e. $r_{\mathcal{V}}$, $\mathbf{x}_{\mathcal{V}}$, and $d_{\mathcal{V}}$ are identical) and the same total circulation $\Gamma_{\mathcal{V}}$. We use a high-fidelity Navier Stokes solver with lattice Green function (LGF) [17, 30]. The true pressure observations are generated by inverting the pressure Poisson equation (2.7) with LGF. Panels (a)-(d) of Fig. 5.8 depict the truth vorticity field at $t = 0.5, 1.5, 6.0$ and 12.0 .

To the best of our knowledge, there is no straightforward way to compare the discrete vorticity distribution generated by the ensemble filter with the true continuous vorticity distribution obtained from the Navier-Stokes solver. Instead, we compare the pressure distribution along the x axis from the high-fidelity simulation with the evaluation of the observation model (3.2) at the posterior ensemble generated by the LREnKF, or the sEnKF (i.e. posterior predictive samples for the pressure distribution). The performance of the LREnKF is assessed for 85%, 95% and 99% of the normalized cumulative energy for $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. We assess the performance of the filter on the segment $[-1.5, 11.0]$ with a finer interspace distance of 0.01. We define the median-squared-error (MSE) at time k between the true pressure distribution \mathbf{y}_k^* , and the median posterior pressure distribution $\text{median}(\mathbf{y}_k^a)$ (i.e. the median pressure distribution for the posterior ensemble at time k) as $\text{MSE}_k = \|\mathbf{y}_k^* - \text{median}(\mathbf{y}_k^a)\|_2 / \text{dim}(\mathbf{y}_k^*)$, where $\text{dim}(\mathbf{v})$ denotes the dimension of a vector \mathbf{v} . Due to the singular nature of the pressure field, we found that the MSE is a more robust estimator of the performance of each filter than the RMSE due to the presence of outliers in the ensemble estimate of the sEnKF. The spin-up phase of the filters is $[0.0, 8.0]$. The reported MSE is time-averaged over the remaining time interval $[8.0, 12.0]$. The uncertainty in the MSE is quantified by the 5%, 25%, 75% and 95% quantiles of the MSE over 50 realizations of the same experiment with different samples for the initial condition.

Fig. 5.9 shows the median spectrum of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$ over the entire time interval. The statistics are obtained from a run of the sEnKF for a large ensemble ($M = 1000$). As in the previous example, the inference problem possesses the low-rank informative structure. The median ranks $r_{\mathbf{x}}$ needed to capture 80%, 90%, 95%, and 99% of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$ are 2, 3, 4, and 16, respectively. Similarly, the median ranks $r_{\mathbf{y}}$ to capture 80%, 90%, and 99% of the cumulative energy of $\mathbf{C}_{\mathbf{y}}$ are 2, 3, and 5, respectively. These ranks are small compared to the state dimension $n = 147$, and the observation dimension $d = 24$.

Fig. 5.10 (a) reports the evolution of the median MSE (computed over 50 realizations) for the sEnKF and the LREnKF with the ensemble size M . We assess the performance of the LREnKF for 85%, 95%, and 99% of the normalized cumulative energy of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$. We recall that the LREnKF reverts to the sEnKF when the dimensions are not reduced. For $M \geq 60$, there is no significant difference in the MSE of the sEnKF and the LREnKF (for the different energy ratios). For $M \in [30, 50]$, the MSE of the sEnKF significantly increases as M decreases. Over this interval, the MSE of the LREnKF shows no major variation. Overall, capturing 85% of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$ is sufficient to yield stable inference results. For $M < 30$, the MSE of the sEnKF diverges, while the LREnKF leads to a reasonable estimate of the pressure distribution for ensemble size as small as 10. This clearly demonstrates the benefit of our regularization for the EnKF. Fig. 5.10(b)-(c) shows the time history of the median ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ required to capture at least 85, 95, 99% of the cumulative energy (computed over 50 realizations) for $M = 20$. The rank $r_{\mathbf{x}}$ initially increases with time, until it reaches a plateau at about $t = 3.0 - 7.0$ depending on the energy ratio. The increase of the rank $r_{\mathbf{x}}$ can be related to the growing role of viscosity in the true pressure response. The rank $r_{\mathbf{y}}$ remains close to 5 over time for the different energy ratios.

Fig. 5.8 compares the vorticity distribution at four times, $t = 0.5, 1.5, 6.0$, and 12.0 , from the truth with the distribution of vortex elements of the posterior mean of the LREnKF and the sEnKF. The ensemble estimates are obtained from one realization of each filter for $M = 30$. The ranks of the LREnKF are set to capture at least 99% of the normalized cumu-

lative energy. The large-scale vortex features are well captured by the inviscid vortex model, despite the obvious absence of the viscous effects in our flow representation. The state estimate is informed of the viscous effects through the assimilation of the pressure observations [3]. Overall, the LREnKF provides a more physically consistent estimate of the vorticity distribution than the sEnKF. With a limited ensemble size, the sEnKF cannot distinguish between physical and spurious long-range correlations between the pressure observations and the vortices. As a result, the sEnKF inconsistently displaces the vortices. At $t = 12.0$, the vorticity distribution of the LREnKF contains a dense core of vortices, with a few vortex satellites to capture the viscous diffusion of the vortex patch. However, the sEnKF poorly estimates the truth vorticity distribution for $t > 6.0$. The spatial structure of the vortices is lost, and they occupy a much larger support than the true vorticity distribution.

Fig. 5.11 compares the history of the pressure from the truth with the posterior mean estimate of the LREnKF and the sEnKF for the same realization of the filters. Over time, the pair of vortex patches is advected along the x axis. Concomitant to the diffusion of the vortex patches, the peak pressure (in magnitude) decreases over time. The true pressure is globally well approximated by the two filters. As we pointed in the previous paragraph, the sEnKF creates spurious displacements of the vortices over time, leading to a growing error in the posterior predictive pressure at later times (the red regions in Fig. 5.11(e) for $t > 6.0$).




Figure 5.9: Median spectrum of the state Gramian $\mathbf{C}_{\mathbf{X}}$ (left panels (a)-(c)) and observation Gramian $\mathbf{C}_{\mathbf{Y}}$ (right panels (d)-(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$. The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles.




Figure 5.10: Left column (a): Time-averaged evolution of the median MSE of the posterior pressure with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The sEnKF is found unstable for $M < 30$. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles. Right column [(b)-(c)]: Time-history of the median ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ of the LREnKF for $M = 20$ (computed over 50 realizations) for different ratios of the cumulative energy of $\mathbf{C}_{\mathbf{x}}$, $\mathbf{C}_{\mathbf{y}}$: 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue.




Figure 5.11: Top line [(a), (b), (c)]: Spatiotemporal map of the pressure induced by a pair of two co-rotating vortex patches from (a) high-fidelity numerical simulation at Reynolds number 1000, and the mean posterior of an inviscid vortex model for one realization of the LREnKF (b), and the sEnKF (c) with $M = 30$. Bottom line [(d), (e)]: Spatiotemporal magnitude of the error between the truth pressure and the mean posterior pressure for one realization of the LREnKF (d), and the sEnKF (e) with $M = 30$. The ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ of the LREnKF are set to capture 99% of the normalized cumulative energy of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$, respectively. The same color levels are used for all the top panels, and another set for the bottom panels.

CHAPTER 6

Conclusions

6.1 Summary

This dissertation studies the regularization of flow estimation problems, where the state is composed of the positions and strengths of a collection of point vortices and the observations are given by pressure readings. We focus our attention on the regularization of the ensemble Kalman filter, which builds an empirical estimate of the Kalman gain from limited samples of the forecast distribution. Unfortunately, for high dimensional and nonlinear problems, the finite ensemble size leads to rank-deficiency, sampling errors and spurious long-range correlations in the estimated Kalman gain. For point vortex models, spurious updates can have dramatic consequences on the flow estimation, with the creation of nonphysical spikes in the pressure and normal force response. This work proposes two strategies to resolve these issues.

Stochastic versions of the EnKF, such as the sEnKF, only verify the propagation equation of the posterior covariance matrix of the Kalman filter with an infinite ensemble. Deterministic versions of the EnKF have been introduced to verify the desired propagation equation with a finite ensemble size. Chapter 4 examines the influence of using a deterministic version of the EnKF, called the ETKF, instead of the sEnKF. We show that the flow estimator introduced by Darakananda et al. [29] is greatly improved with the ETKF. We assess this flow estimator on two challenging scenarios: a flat plate subject to strong and overlapping disturbances applied near the leading edge of the airfoil, and a flat plate placed in the wake

of a circular cylinder over an extended time window of 12 convective times. On these two examples, we demonstrate the superiority of the ETKF over the sEnKF on the estimation of the coherent structures of the flow field, and the normal force response. In both cases, we show that the proposed flow estimator accurately estimates the flow field, with no *a priori* knowledge of the flow perturbations. The perturbations are only accounted by the assimilation of the truth pressure observations.

Chapter 5 proposes a novel regularization of the EnKF for *non-local* observation models, such as models given by elliptic PDEs. Elliptic PDEs are ubiquitous in fluid mechanics: the relation between vorticity and pressure, the relation between streamfunction and vorticity, and the relation between the scalar potential and dilation rate of the velocity field are all given by Poisson equations. Classical methods for regularizing the EnKF are based on the *locality* of the observations: an observation only provides information about the state variables which are close-by in physical distance. From the logarithm or algebraic decay of the Green’s function of the Laplace operator, we expect physical long-range correlations between the state and observation variables. Distance localization is not suited for elliptic inverse problems, as it would blindly remove spurious and physical long-range correlations. We show that elliptic inverse problems usually possess a low effective dimension: a low-dimensional projection of the observations is mostly informative along a limited number of directions of the state space. Chapter 5 leverages this structure, to propose a novel regularization of the EnKF. To do so, we use a factorization of the Kalman gain based on the singular value decomposition of the whitened observation matrix. This factorization illuminates the inference process: the innovation term is first whitened, projected onto the leading modes of the observation space (the right singular vectors of the whitened observation matrix), assimilated in the rotated informative subspace, and finally lifted to the original space. We generalize this decomposition to a nonlinear-Gaussian observation model with two key steps. First, we identify the informative directions of the state and observation spaces as the leading eigenvectors of the state Gramian $\mathbf{C}_{\mathbf{x}}$ and the observation Gramian $\mathbf{C}_{\mathbf{y}}$, based on the

sensitivity of the observation operator. We show that the leading eigenvectors of the state and observation Gramians are reminiscent of the source and target modes of the multipole expansion. Nonetheless, our methodology only requires evaluations of the Jacobian of the observation operator. Second, we interpret the factorization of the Kalman gain as a change of coordinates between the original space and the informative space. We introduce a novel algorithm for the EnKF, the low-rank EnKF (LREnKF), which regularizes the inference by computing the Kalman gain in the span of the informative subspace (i.e. in the span of the leading eigenvectors of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$), and then lifting the result to the original space. For elliptic inverse problems, the spectrum of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$ is rapidly decaying. This is the central property exploited by the LREnKF. Our methodology concomitantly reduces the variance error due to the smaller number of coefficients to identify, while making the bias error very small for any budget of samples from the prior distribution. To avoid any *ad hoc* tuning of the ranks $r_\mathbf{x}$ and $r_\mathbf{y}$, the ranks are adaptively selected based on a given ratio of the cumulative energy of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$. We assess the sEnKF and the LREnKF on two problems where we seek to estimate the positions and strengths of point vortices from pressure observations. On both problems, the LREnKF significantly improves the estimate of the sEnKF, and remains stable with limited samples ($M = 10$).

6.2 Future directions

To conclude this dissertation, we propose several directions to further improve the assimilation of pressure observations into a point vortex model.

For frequent assimilation cycles and Gaussian observation noise, the Gaussian assumption is valid. Indeed, over a short time window any nonlinear system with Gaussian noise can be well approximated by a linear-Gaussian system (using first order Taylor series approximations). However, real-world pressure observations collected on the surface of an airfoil contain many outliers, and are not well represented by light-tailed Gaussian densities. To

perform robust flow estimation with experimental observations, one could model the observation noise with a Laplace distribution [34]. A Laplace distribution has heavy tails that will assign higher probability to outliers, and be more representative of the statistics of real-world pressure data. From the nonlinearity of the dynamical and observation operators (the Biot-Savart law and the pressure Poisson equation), one can expect the state distribution to be non-Gaussian. Therefore, the linear analysis map of the EnKF can be inadequate for estimating the true update of a pressure discrepancy on a vortex element. Indeed, the analysis step of the EnKF uses an implicit linearization of the observation operator about the mean of the forecast ensemble [53]. In future work, it could be useful to explore nonlinear filters that embrace the nonlinear and non-Gaussian properties of the flow estimation problem. Spantini et al. [39] introduces a nonlinear generalization of the sEnKF, the *stochastic map filter* (SMF), based on parsimonious and interpretable nonlinear prior-to-posterior transformations to reduce the bias of the EnKF but maintain its robustness. Le Provost et al. [34] shows encouraging results with the SMF on flow estimation problems with only two pressure sensors placed at the edges of a flat plate undergoing strong perturbations. This work introduced a low-rank version of the SMF by estimating the nonlinear analysis map in the informative subspace spanned by the leading eigenvectors of the state and observation Gramians.

The major limitation of the proposed joint dimension reduction is the computational cost to form the state and observation Gramians. Fortunately, several techniques can be used in tandem to reduce this cost. First, we do not need the entire set of eigenpairs of these Gramians. The algorithm of the LREnKF works for any basis of the low-dimensional subspace spanned by the leading eigenvectors of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$. From the Davis-Kahan theorem 5.1.1, and the rapid spectral decay of the eigenvalues of $\mathbf{C}_\mathbf{x}$ and $\mathbf{C}_\mathbf{y}$, the top eigenvectors can be identified with limited samples. Instead, one can evaluate the Jacobian for a subset of the samples, either selected by random subsampling without replacement, or by a clustering algorithm (e.g. k-means algorithm). Moreover, it is not necessary to compute the entire

Jacobian matrix for all the samples. Let assume that one has scalable routines to evaluate the matrix-vector products $(\nabla \mathbf{h})\mathbf{v}$, and $(\nabla \mathbf{h}^\top)\mathbf{u}$ where $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{u} \in \mathbb{R}^d$. The product $(\nabla \mathbf{h})\mathbf{v}$ corresponds to a single directional derivative of \mathbf{h} along \mathbf{v} . The product $(\nabla \mathbf{h}^\top)\mathbf{u}$ can be computed with a single call of an adjoint solver for the observation model. By leveraging ideas from randomized linear algebra [65, 66, 67], we can construct accurate approximation of the informative subspaces from limited matrix-vector products $(\nabla \mathbf{h})\mathbf{v}$ and $(\nabla \mathbf{h}^\top)\mathbf{u}$. Another interesting direction is to recognize that the evaluation of the observation model at the prior samples provide unstaggered information about its sensitivity. One can seek to construct a Jacobian-free estimate for the state and observation Gramians from these evaluations [66].

In many problems of engineering, it is sometimes unclear how the state and observation variables are related. The spectrum of the state and observation Gramians provides a principled way to discover and interpret these interactions. The algorithm of the LREnKF based on the idea of a low-rank prior-to-posterior transformation is not limited to the flow estimation problems studied in this dissertation. Indeed, elliptic observation models are ubiquitous in science: geoscience, electromagnetism, magneto-hydrodynamics to name a few. For all these problems, the regularization of the EnKF is usually *ad-hoc* [68, 69]. The LREnKF can be highly beneficial to regularize these problems. We firmly believe that the methodology of the LREnKF can pave the way for the development of principled and scalable algorithms to assimilate non-local observations.

APPENDIX A

Algorithms of the stochastic, deterministic, and low-rank ensemble Kalman filters

For the sake of completeness, we provide pseudo-codes for the stochastic EnKF (sEnKF) [14, 13] in Algorithm 1, the ensemble transform Kalman filter (ETKF) [31, 13] in Algorithm 3, and the low-rank EnKF (LREnKF) in Algorithm 4. Algorithm 2 describes an algorithm to construct a mean-preserving random rotations based on [47, 46] used in the ETKF's algorithm.

Algorithm 1: $\text{senkf}(\pi_{\mathbf{x}_0}, \mathbf{f}, \mathbf{h}, K, \{\mathbf{y}_k^*\}_{1:K})$ recursively assimilates the observations $\{\mathbf{y}_k^*\}_{1:K}$ with the ETKF, where K is the number of assimilation steps and $\pi_{\mathbf{x}_0}$ is the initial state distribution. Adapted from Le Provost et al. [3].

Input: $\pi_{\mathbf{x}_0}, \mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n, \mathbf{h} : \mathbb{R}^n \rightarrow \mathbb{R}^d$, the number of time steps $K \in \mathbb{N}$, and the observations to assimilate $\{\mathbf{y}_k^*\}_{1:K}$.

Output: $\{\mathbf{x}_k^{a,i}\}_{k=1:K, i=1:M}$

```

1 for  $k = 1 : K$  do
    /* Compute the ensemble forecast: */
2    $\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_{k-1}^i$ , for  $i = 1, \dots, M$ 
    /* Sample from the measurement noise distribution  $\pi_{\mathcal{E}_k}$ : */
3    $\boldsymbol{\epsilon}_k^i \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_k)$ , for  $i = 1, \dots, M$ 
    /* Compute the sample means: */
4    $\bar{\mathbf{x}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_k^{f,i}$ ,  $\bar{\boldsymbol{\epsilon}}_k = \frac{1}{M} \sum_{i=1}^M \boldsymbol{\epsilon}_k^i$ ,  $\bar{\mathbf{y}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{h}(\mathbf{x}_k^{f,i})$ ,
    /* Compute the anomalies: */
5    $\mathbf{X}_k'^{f,i} = (\mathbf{x}_k^{f,i} - \bar{\mathbf{x}}_k^f) / \sqrt{M-1}$ ,
      $\mathbf{Z}_k'^{f,i} = (\mathbf{h}(\mathbf{x}_k^{f,i}) - \boldsymbol{\epsilon}_k^i - \bar{\mathbf{y}}_k^f + \bar{\boldsymbol{\epsilon}}_k) / \sqrt{M-1}$ , for  $i = 1, \dots, M$ 
    /* Solve the linear systems for  $(\mathbf{b}_k^i)$ : */
6    $\mathbf{Z}^{\prime f} \mathbf{Z}^{\prime f \top} \mathbf{b}_k^i = \mathbf{y}_k^* - \mathbf{h}(\mathbf{x}_k^{f,i}) + \boldsymbol{\epsilon}_k^i$ , for  $i = 1, \dots, M$ 
    /* Update the ensemble: */
7    $\mathbf{x}_k^{a,i} = \mathbf{x}_k^{f,i} + \mathbf{X}_k^{\prime f} \mathbf{Z}_k^{\prime f \top} \mathbf{b}_k^i$ 
8 return  $\{\mathbf{x}_k^{a,i}\}_{k=1:K, i=1:M}$ 

```

Algorithm 2: Algorithm to generate mean-preserving random rotations, adapted from [47, 46, 3].

Input: $M \in \mathbb{N}$

Output: $U \in \mathbb{R}^{M \times M}$

```

/* Draw  $(M - 1) \times (M - 1)$  samples from a standard normal distribution and store them
   in  $\Omega \in \mathbb{R}^{(M-1) \times (M-1)}$ : */
1  $\Omega = \text{randn}(M - 1, M - 1)$ 
   /* SVD decomposition of  $\Omega$ : */
2  $\Omega = U\Sigma V^\top$ 
3  $\Omega = V^\top$ 
4  $\mathbf{b} = \mathbf{1}/\sqrt{M}$ 
   /* Compute  $\mathbf{b}^{sign}$ : */
5  $\mathbf{b}^{sign} = \mathbf{b}$ 
   /* Add  $\text{sign}(\mathbf{b}[M])$  to the last entry: */
6  $\mathbf{b}^{sign}[M] += \text{sign}(\mathbf{b}[M])$ 
7 Initialize  $\mathbf{B} \in \mathbb{R}^{M \times M}$ 
   /* Fill columns of  $\mathbf{B}$ : */
8  $\mathbf{B}[:, 1] = \mathbf{b}$ 
9  $\mathbf{B}[:, 2 : M] = -(\mathbf{I}_M + \mathbf{b}^{sign}\mathbf{b}^{sign^\top}/(\|\mathbf{b}\| + 1))[:, 1 : M - 1]$ 
10 Initialize  $\Lambda_{\mathbf{b}} \in \mathbb{R}^{M \times M}$ 
   /* Fill  $\Lambda_{\mathbf{b}}$ : */
11  $\Lambda_{\mathbf{b}}[1, 1] = 1$ 
12  $\Lambda_{\mathbf{b}}[1 : M - 1, 1 : M - 1] = \Omega$ 
13 return  $U = \mathbf{B}\Lambda_{\mathbf{b}}\mathbf{B}^\top$ 

```

Algorithm 3: $\text{etkf}(\pi_{\mathbf{x}_0}, \mathbf{U}, \mathbf{f}, \mathbf{h}, K, \{\mathbf{y}_k^*\}_{1:K})$ recursively assimilates the observations $\{\mathbf{y}_k^*\}_{1:K}$ with the ETKF, where K is the number of assimilation steps, $\pi_{\mathbf{x}_0}$ is the initial state distribution, and \mathbf{U} is a mean-preserving rotation generated from time to time, see Algorithm 2, otherwise set to \mathbf{I}_M . Adapted from Le Provost et al. [3].

Input: $\pi_{\mathbf{x}_0}$, $\mathbf{U} \in \mathbb{R}^{M \times M}$, $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $\mathbf{h} : \mathbb{R}^n \rightarrow \mathbb{R}^d$, the number of time steps

$K \in \mathbb{N}$, and the observations to assimilate $\{\mathbf{y}_k^*\}_{1:K}$.

Output: $\{\mathbf{X}_k^a\}_{1:K}$

```

/* Generate M samples from  $\pi_{\mathbf{x}_0}$ : */
1  $\mathbf{x}_0^{a,i} \sim \pi_{\mathbf{x}_0}$  for  $i = 1, \dots, M$ 
2 for  $k = 1 : K$  do
    /* Compute the ensemble forecast: */
3      $\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_{k-1}^i$  for  $i = 1, \dots, M$ 
    /* Compute the prior mean and anomaly matrix: */
4      $\bar{\mathbf{x}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_k^{f,i}$ ,  $\mathbf{X}_k^{f,i} = (\mathbf{x}_k^{f,i} - \bar{\mathbf{x}}_k^f) / \sqrt{M-1}$ , for  $i = 1, \dots, M$ 
5      $\mathbf{Z}_k^{f,i} = \mathbf{h}(\mathbf{x}_k^{f,i})$ , for  $i = 1, \dots, M$  (compute the observations)
    /* Compute mean of the observations  $\bar{\mathbf{y}}_k^f$ : */
6      $\bar{\mathbf{y}}_k^f = \frac{1}{M} \mathbf{Z}_k^f \mathbf{1}$ 
    /* Construct the matrices  $\mathbf{S}$  and  $\mathbf{G}$ : */
7      $\mathbf{S} = \mathbf{V}_k^{-1/2} (\mathbf{Z}_k^f - \bar{\mathbf{y}}_k^f \mathbf{1}^\top) / \sqrt{M-1}$ 
8      $\mathbf{G} = (\mathbf{I}_M + \mathbf{S} \mathbf{S}^\top)^{-1}$ 
    /* Construct the vectors  $\delta$  and  $\mathbf{c}$ : */
9      $\delta = \mathbf{V}_k^{-1/2} (\mathbf{y}_k^* - \bar{\mathbf{y}}_k^f)$ ,  $\mathbf{c} = \mathbf{G} \mathbf{S}^\top \delta$ 
    /* Update the ensemble: */
10     $\mathbf{X}_k^a = \bar{\mathbf{x}}_k^f \mathbf{1}^\top + \mathbf{X}_k^{f,i} (\mathbf{c} \mathbf{1}^\top + \sqrt{M-1} \mathbf{G}^{1/2} \mathbf{U})$ 
11 return  $\{\mathbf{X}_k^a\}_{1:K}$ 

```

Algorithm 4: $\text{lowrankenkf}(\mathbf{y}^*, \mathbf{h}, \{\mathbf{x}^i\}, r_{\mathbf{X}}, r_{\mathbf{Y}})$ assimilates the data \mathbf{y}^* in the prior samples $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ with the low-rank EnKF. The integers $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ determine the rank of the projected subspace in the state space and in the observation space, respectively.

Input: $\mathbf{y}^* \in \mathbb{R}^d$, $\mathbf{h} : \mathbb{R}^n \rightarrow \mathbb{R}^d$, M samples $\{\mathbf{x}^i\}$ from $\pi_{\mathbf{X}}$, ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$

Output: M samples $\{\mathbf{x}_a^i\}$ from $\pi_{\mathbf{X}|\mathbf{Y}}(\cdot | \mathbf{y}^*)$

```

1 for  $i = 1 : M$  do
2    $\mathbf{z}^i \leftarrow \mathbf{h}(\mathbf{x}^i)$ ,  $\boldsymbol{\epsilon}^i \sim \mathcal{N}(\mathbf{0}_d, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}})$ 
   /* Compute a Monte-Carlo approximation of  $\mathbf{C}_{\mathbf{X}}$  and  $\mathbf{C}_{\mathbf{Y}}$ . */
3    $\widehat{\mathbf{C}}_{\mathbf{X}} \leftarrow \frac{1}{M-1} \sum_{i=1}^M \left( \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)^{\top} \left( \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)$ 
4    $\widehat{\mathbf{C}}_{\mathbf{Y}} \leftarrow \frac{1}{M-1} \sum_{i=1}^M \left( \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right) \left( \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)^{\top}$ 
   /* Perform a low-rank eigendecomposition of  $\widehat{\mathbf{C}}_{\mathbf{X}}$  and  $\widehat{\mathbf{C}}_{\mathbf{Y}}$ : */
5    $\mathbf{V}_{r_{\mathbf{X}}} \leftarrow \text{eigenvector}(\widehat{\mathbf{C}}_{\mathbf{X}}, r_{\mathbf{X}})$ ,  $\mathbf{U}_{r_{\mathbf{Y}}} \leftarrow \text{eigenvector}(\widehat{\mathbf{C}}_{\mathbf{Y}}, r_{\mathbf{Y}})$ 
6 for  $i = 1 : M$  do
7    $\check{\mathbf{x}}^i \leftarrow \mathbf{V}_{r_{\mathbf{X}}}^{\top} \boldsymbol{\Sigma}_{\mathbf{X}}^{-1/2} (\mathbf{x}^i - \widehat{\boldsymbol{\mu}}_{\mathbf{X}})$ ,  $\check{\mathbf{z}}^i \leftarrow \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} (\mathbf{z}^i - \widehat{\boldsymbol{\mu}}_{\mathbf{Z}})$ ,  $\check{\boldsymbol{\epsilon}}^i \leftarrow \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} (\boldsymbol{\epsilon}^i - \widehat{\boldsymbol{\mu}}_{\boldsymbol{\epsilon}})$ 
8 for  $i = 1 : M$  do
9    $\mathbf{A}_{\check{\mathbf{X}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\mathbf{x}}^i - \widehat{\boldsymbol{\mu}}_{\check{\mathbf{X}}})$ ,  $\mathbf{A}_{\check{\mathbf{Z}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\mathbf{z}}^i - \widehat{\boldsymbol{\mu}}_{\check{\mathbf{Z}}})$ 
10   $\mathbf{A}_{\check{\boldsymbol{\epsilon}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\boldsymbol{\epsilon}}^i - \widehat{\boldsymbol{\mu}}_{\check{\boldsymbol{\epsilon}}})$ 
   /* Solve the linear system for  $\check{\mathbf{b}} \in \mathbb{R}^{r_{\mathbf{Y}} \times M}$ : */
11  $(\mathbf{A}_{\check{\mathbf{Z}}} \mathbf{A}_{\check{\mathbf{Z}}}^{\top} + \mathbf{A}_{\check{\boldsymbol{\epsilon}}} \mathbf{A}_{\check{\boldsymbol{\epsilon}}}^{\top}) \check{\mathbf{b}} = \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}^{-1/2} (\mathbf{y}^* \mathbf{1}_M^{\top} - (\mathbf{E}_{\mathbf{Z}} + \mathbf{E}_{\boldsymbol{\epsilon}}))$ ,
   /* where  $\mathbf{E}_{\mathbf{Z}} \in \mathbb{R}^{d \times M}$  and  $\mathbf{E}_{\boldsymbol{\epsilon}} \in \mathbb{R}^{d \times M}$  are the ensemble matrices of the samples  $\{\mathbf{z}^i\}$ 
   and  $\{\boldsymbol{\epsilon}^i\}$ . Lift result to the original space: */
12 for  $i = 1 : M$  do
13    $\mathbf{x}_a^i \leftarrow \mathbf{x}^i + \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \mathbf{V}_{r_{\mathbf{X}}} (\mathbf{A}_{\check{\mathbf{X}}} \mathbf{A}_{\check{\mathbf{Z}}}^{\top}) \check{\mathbf{b}}[:, i]$ 
14 return  $\{\mathbf{x}_a^i\}$ 

```

APPENDIX B

Mean transport of inertial particles in viscous streaming flows

B.1 Introduction

Recent developments in the fields of biomedical diagnosis, pollutant treatment, drug delivery, and microfluidics—to name a few—have motivated the need for efficient and fast methods to transport, cluster or trap inertial particles (small finite-sized particles) in a fluid environment. The particles transported, such as drugs or biological cells, are fragile, and any direct contact creates undesirable stresses on the particles that may cause irreversible damage. Though no method of transport can avoid applying stress, non-contact methods can provide opportunities to distribute the stresses over the particle more uniformly, reducing the possibility for damage. Techniques using ultrasound [70, 71], lasers [72, 73], magnetic effects [74], dielectrophoresis [75] or inertial hydrodynamics effects [76, 77] have emerged as some of the most effective methods to manipulate inertial particles.

Another attractive possibility for non-contact particle transport is based on the notion of viscous streaming. A streaming flow is a weak but large-scale steady response of the fluid to oscillatory forcing, brought about through the Reynolds stresses imparted on the fluid. Numerous studies have shown the promises of viscous streaming to transport and trap inertial particles. Classical works have focused on viscous streaming created by a cylinder oscillating weakly in rectilinear motion [78, 79, 80, 81]. Lutz et al. [82] have been able to trap particles in steady streaming eddies arranged in a clover-shaped pattern around a cylindrical

post fixed in a micro-channel through which fluid was forced in oscillatory fashion. Chong et al. have identified the mechanisms that underlie this trapping [81], and have shown that an arrangement of multiple cylinders forced in sequence with oscillatory motions can be used to construct desired inertial particle trajectories [2]. Abadi et al. [83] have recently designed a closed-loop controller for the position and velocity of inertial particles inside a two-dimensional square box using steady streaming mechanisms. The control actuation is made through four vibrating piezoelectric beams inclined at 45° at each corner of the square box. Very good performance was reported: inertial particles were successfully forced to carry out a variety of prescribed motions, such as eight-branch star trajectories or the transport of a constellation of inertial particles without changing the distance between them. Parthasarathy et al. [84] have recently shown that, in an arrangement of two cylinders in a fluid, in which one is actively moved and the other is passively transported by the resulting flow, the passive cylinder's transport is enhanced by adding oscillations to the active cylinder's motion. The additional oscillations generate a streaming flow.

Rigid walls are not the only means for introducing oscillatory motion into the fluid to enact a streaming flow. For example, recent works have shown the potential of using streaming flows created by bubbles undergoing oscillatory volume and shape changes [85, 86, 87, 88]. Two reasons motivate the use of bubbles for actuation of viscous streaming: Larger amplitude motions can be created compared to rigid bodies, resulting in quadratic increase of the streaming speed [86]. Also, the bubble–fluid interface allows non-zero tangential velocity, leading to less deceleration of inertial particles in the vicinity of bubbles' surface compared to rigid surfaces.

All of these works have shown the potential for particle manipulation using viscous streaming. In order to devise means of strategically exploiting this mechanism for transport, it is important to have a mathematical model for predicting the particle trajectories effected by a given geometry and forcing modality. Chong et al. [81] showed that the dynamics of isolated inertial particles in viscous streaming flows are well captured by the Maxey–Riley

(MR) equation [89] with the addition of a Saffman lift force [90]. (We refer readers to Michaelides [91] for a review and history of various transport equations for inertial particles.) In particular, this approach accounts for the local velocity of the fluid in the particle transport, but the particle’s influence on the fluid motion can be reasonably neglected.

This prediction of particle transport is challenged by the underlying mechanisms responsible for the transport. A viscous streaming flow has two well-separated time scales: a fast oscillatory scale t_f and a slow one associated with the steady streaming t_s [81, 92]. Consider a weakly oscillating cylinder with angular frequency Ω , amplitude of oscillation A and radius R such that, $\epsilon = A/R \ll 1$, as is typical in streaming applications. The fast time scale is set by the period of oscillation, $t_f = T \sim 1/\Omega$. Inertial particles are transported at the characteristic (drift-corrected) speed of the streaming flow, $V_s = \epsilon\Omega A$, in a streaming cell of characteristic size δ_{DC} . Chong et al. [81] have shown that this streaming cell size remains $\delta_{DC} = O(R)$ over a wide range of Reynolds number $Re = \Omega R^2/\nu$. Hence, the characteristic convective time of an inertial particle around a streaming cell is of order $t_s \sim \delta_{DC}/V_s \sim R/(\epsilon\Omega A) = 1/(\epsilon^2\Omega)$. Therefore the slow time scale t_s is a factor $1/\epsilon^2$ larger than the fast scale, i.e., $t_s = t_f/\epsilon^2$ with $\epsilon \ll 1$. An inertial particle of small radius $a \ll R$ has a third, even faster time scale: the time of viscous response when its velocity deviates from that of the fluid. The ratio of this scale to that of the oscillations is measured by the Stokes number, denoted in this paper by τ , which is proportional to Rea^2/R^2 .

Indeed, in a viscous streaming flow, a particle is continuously wiggling around its mean trajectory at the fast time scale and translating at the slow time scale. In order to discern long-time behaviors of inertial particles in viscous streaming flows, simulations must be carried out over several transport time scales, i.e., several multiples of t_s , and consequently, several thousands of oscillation cycles. Since the oscillatory motion of the fluid—occurring on the fast time scale—has non-negligible influence on the particle transport, the simulations of both the particle transport and the governing equations in the fluid must ostensibly be well-resolved temporally at the scale t_f . Even with the one-way coupling described above, a

full simulation of a single particle trajectory is computationally expensive when carried out in this manner.

Clearly, the primary desire is to predict the slow-timescale (i.e, ‘mean’) trajectory of the particle and to seek only the averaged influence of the fast-timescale on this trajectory. In this work, we seek to provide a framework for accelerated prediction of inertial particle trajectories in this fashion. We are not the first to pursue such a strategy. It should be noted that Thameem et al. [92] and Agarwal et al. [93] have proposed a time-scale separation of the Maxey–Riley equation to derive equations resolved at the slow time scale. However, their approach is restricted to a purely radial velocity field, periodic in time at leading order (in powers of ϵ) and steady at second order.

The approach we propose in this work is also focused on oscillatory fluid velocity fields, though it allows for slow transient changes of such fields, and places no restrictions on the field’s spatial structure. We will use the framework of the Generalized Lagrangian Mean (GLM) theory of Andrews and McIntyre [94] to form an expression for the Lagrangian mean velocity field, $\overline{\boldsymbol{w}}^L$, associated with an underlying oscillatory velocity field \boldsymbol{w} . This velocity field, defined as the time-average velocity of the particle passing through any field point, is a cornerstone of our method: once we have it, we can directly compute the mean particle trajectories, using numerical time steps equal to several oscillation periods. The field $\overline{\boldsymbol{w}}^L$ explicitly filters the fast fluctuating component from the mean motion of particles; we will show that the fast component’s effect is confined to the Stokes drift. To the best of our knowledge, no previous work has used these tools to efficiently solve the Maxey–Riley equation in a setting of disparate time scales.

We will form Lagrangian mean velocity fields for fluid particles as well as inertial particles. The underlying velocity field for inertial particles will be obtained by an asymptotic expansion of the Maxey–Riley equation in small Stokes number for the deviation of the particle’s velocity from that of the fluid. In so doing, we will extend an approach used previously by Maxey [95] and Ferry and Balachandar [96]; here, we will add the important Faxén cor-

rection to these earlier treatments. This approach allows us to obtain the inertial particle’s velocity from that of the fluid at little extra cost. The problem will be further simplified by applying a separate asymptotic expansion in small oscillation amplitude ϵ . Truncating this expansion at second order, we will arrive at a compact and self-consistent form of equations for the fluid velocity field and the subsequent mean transport of fluid and inertial particles.

The remainder of the paper is organized as follows. The description of the Maxey–Riley equation and their expansion in small Stokes number will be presented in Section B.2. The equations for the Lagrangian mean velocity field, their simplification in small-amplitude viscous streaming flow, and the algorithm for the fast Lagrangian-averaged transport of particles will be discussed in Section B.3. Applications of our algorithm to the cases of one and two weakly oscillating cylinders will be presented in Section B.4. Concluding remarks will follow in Section B.5. We would also like to note that all of the computational tools developed for this paper are available at <https://github.com/jdeldre/ViscousStreaming.jl>.

B.2 Basic transport for inertial particles

In this work, we are interested in computing the trajectory an inertial particle immersed in an incompressible flow. As discussed in Section B.1, we assume that the coupling between the fluid motion and the particle trajectory is one way: the particle’s motion is determined by the local fluid velocity and generates a disturbance field that is unmodified by the proximity to oscillating bodies. Thus, the fluid’s time-varying velocity field, $\mathbf{u}(\mathbf{x}, t)$ can be assumed known—for example, through analytical or computational means—without regard for the particle’s presence, and our focus is only on obtaining the particle trajectory in this field. The goal of this section is to obtain the general transport equations for inertial particles in a form conducive for the next section, in which we distill this transport into the mean motion.

The fluid’s density and kinematic viscosity are denoted by ρ_f and ν , respectively. The

fluid velocity field, \mathbf{u} , is governed by the incompressible Navier–Stokes equations,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{\text{Re}} \nabla^2 \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0, \quad (\text{B.1})$$

in which all quantities (including pressure, p) have been non-dimensionalized by the uniform fluid density ρ_f and the characteristic length and time scales of the flow. These scales are established by the driving mechanism: The fluid is bounded on the interior by impenetrable surfaces that are either stationary or oscillating with angular frequency Ω —generically, we will refer to these surfaces as ‘oscillators’. Thus, the characteristic time scale is taken as $1/\Omega$ and, in the case of a cylindrically-shaped oscillator, the characteristic length taken as the cylinder’s radius, R . The flow Reynolds number, Re , is thus defined as

$$\text{Re} = \frac{\Omega R^2}{\nu}. \quad (\text{B.2})$$

In viscous streaming applications, we anticipate $\text{Re} = O(10)$. We assume that the fluid is initially quiescent and that the flow is generated in an infinite domain in which the fluid remains at rest at infinity,

$$\mathbf{u}(\mathbf{x}, 0) = 0, \quad \mathbf{u} \rightarrow 0, \quad |\mathbf{x}| \rightarrow \infty, \quad (\text{B.3})$$

though this condition at infinity can easily be replaced with, e.g., a steady uniform flow or stationary enclosing walls. The form of the boundary conditions on the oscillators will be discussed later in the paper. For now, we simply note that the displacement amplitude of the oscillations, A , will be assumed small compared with the size of the oscillator. The ratio of these scales is denoted by ϵ , so we are assuming that

$$\epsilon \equiv A/R \ll 1. \quad (\text{B.4})$$

Inertial particles are assumed to be rigid spheres with density ρ_p and radius much smaller than the oscillating object, e.g., $a \ll R$. The particle’s mass is denoted by $m_p = 4\pi\rho_p a^3/3$, and the displaced fluid mass by $m_f = 4\pi\rho_f a^3/3 = m_p\rho_f/\rho_p$. We will denote the particle

trajectory by $\mathbf{x}_p(t)$ and associated velocity by $\mathbf{V}_p(t)$:

$$\frac{d\mathbf{x}_p}{dt} = \mathbf{V}_p(t). \quad (\text{B.5})$$

It will be assumed that the particle starts each trajectory at the same velocity as the surrounding fluid,

$$\mathbf{V}_p(t_0) = \mathbf{u}(\mathbf{x}_p(t_0), t_0). \quad (\text{B.6})$$

It is also useful to define the particle ‘slip’ velocity, $\mathbf{V}_p(t) - \mathbf{u}(\mathbf{x}_p(t), t)$, the particle’s velocity relative to the surrounding fluid, which is initially zero by virtue of the initial condition (B.6).

B.2.1 The Maxey–Riley equation with Saffman lift

We will assume throughout this work that the Reynolds number redefined on the particle radius is small,

$$\text{Re}(a/R)^2 \ll 1. \quad (\text{B.7})$$

In this work, for the transport of an inertial particle, we use a form of the Maxey–Riley (MR) equation [89] that includes the Saffman lift [90], as was done by Chong et al. [81] or Ferry and Balachandar [96]. If we neglect gravity, the trajectory of an inertial particle is governed by

$$\begin{aligned} m_p \frac{d\mathbf{V}_p}{dt} = & 6\pi\rho_f\nu a \left(\mathbf{u}(\mathbf{x}_p(t), t) + \frac{1}{6}a^2\nabla^2\mathbf{u}(\mathbf{x}_p(t), t) - \mathbf{V}_p(t) \right) \\ & + m_f \left. \frac{D\mathbf{u}}{Dt} \right|_{\mathbf{x}_p(t)} \\ & - \frac{1}{2}m_f \left[\frac{d\mathbf{V}_p}{dt} - \left. \frac{D\mathbf{u}}{Dt} \right|_{\mathbf{x}_p(t)} - \frac{d}{dt} \left(\frac{1}{10}a^2\nabla^2\mathbf{u}(\mathbf{x}_p(t), t) \right) \right] \\ & + 2\sqrt{3}\pi\nu^{1/2}a^2\rho_f\mathcal{L}_B \left[\mathbf{u}(\mathbf{x}_p(t), t) + \frac{1}{6}a^2\nabla^2\mathbf{u}(\mathbf{x}_p(t), t) - \mathbf{V}_p(t) \right] \\ & + 2\sqrt{3}\pi\nu^{1/2}a^2\rho_f\mathcal{L}_S [\mathbf{u}(\mathbf{x}_p(t), t) - \mathbf{V}_p(t)]. \end{aligned} \quad (\text{B.8})$$

Two different time derivatives act on field quantities in equation (B.8): by d/dt and D/Dt we denote, respectively, the time derivative operators following the particle and the fluid:

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{V}_p(t) \cdot \nabla \quad (\text{B.9})$$

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla. \quad (\text{B.10})$$

The set of terms on each line of the right-hand side of equation (B.8) represent, respectively, the Stokes drag, the fluid acceleration force, the added mass effects, and finally, the Basset history force and the Saffman lift, with linear operators respectively defined as

$$\mathcal{L}_B[\mathbf{f}](t) = \sqrt{\frac{3}{\pi}} \int_{-\infty}^t \frac{d\mathbf{f}/d\tau}{\sqrt{t-\tau}} d\tau, \quad (\text{B.11})$$

$$\mathcal{L}_S[\mathbf{f}](t) = \frac{3\sqrt{3}J_\infty}{2\pi^2\sqrt{|\boldsymbol{\omega}(\mathbf{x}_p(t), t)|}} \mathbf{f}(t) \times \boldsymbol{\omega}(\mathbf{x}_p(t), t), \quad (\text{B.12})$$

where $\boldsymbol{\omega} = \nabla \times \mathbf{u}$ denotes the associated vorticity of the fluid flow at the location of the particle. For the coefficient J_∞ , we use $J_\infty = 2.255$: the limit of the lift coefficient function $J(\eta)$ as the ratio $\eta = \text{Re}_G^{1/2}/\text{Re}_p$ goes to infinity. (For details on this function J , see [97, 98].)

The Basset history force is a memory term due to the unsteady diffusion of vorticity from the particle during its traveling history. Several studies [99, 100] have shown that it can be of significant importance. We retain the term for now for the sake of generality and comparison with previous works, and our scaling analysis below will not reveal it to be clearly smaller than other terms. However, in the context of particle transport in viscous streaming, Chong et al. [81] have shown empirically that this term is of negligible importance in the current parameter regime and can safely be ignored. We will do the same later in the paper.

Equation (B.8) also contains the Faxén corrections (the Laplacian of the fluid velocity field), which were shown by Chong et al. [81] to be crucial in regions of high shear to cause the particle's trajectory to deviate from that of the fluid. That study also demonstrated the important role of the Saffman lift in ultimately trapping the inertial particle at the center of a viscous streaming cell, observed in previous experiments by, for example, Lutz et al. [82].

Several comments are in order regarding our inclusion of the Saffman lift. This term represents an inertial influence of the fluid when the particle moves relative to the fluid in a region of shear, generating a force on the particle perpendicular to the motion. As such, it is non-linearly dependent on the fluid velocity field. The parameter regime of viscous streaming described in this paper justifies the inclusion of such lift, as we will discuss below. However, it does not strictly meet all of the conditions under which Saffman derived the expression for lift [90]. That derivation relies on the particle lying in a region of shear that is nearly uniform well beyond a region of length $L_S = (\nu/G)^{1/2}$ (the so-called ‘Saffman length’), where G is the norm of the local velocity gradient; such shear uniformity enables Saffman’s rigorous singular perturbation treatment [90].

In streaming flows, the particle encounters significant shear within the Stokes boundary layer generated around the oscillating body, a region of thickness $\delta_s = (\nu/\Omega)^{1/2}$. The ratio of the Saffman length to the Stokes boundary thickness should be small to justify the singular perturbation treatment. Here, that ratio is $(\Omega/|\boldsymbol{\omega}|)^{1/2}$, where $|\boldsymbol{\omega}|$ is representative of the instantaneous vorticity in the Stokes layer. In the limit of vanishing oscillation amplitude ϵ , this ratio increasingly fails to abide by the required separation of scales. But this limit is of little practical relevance, as the flow itself vanishes in this limit. In the scenarios described later in this paper, the ratio is of order 1: still not quite sufficient for the strict separation of scales. As Saffman’s own analysis of a particle in steady Poiseuille flow showed [90], this separation of scales is difficult to meet even in many simpler flows. Thus, we interpret the mathematical form here as a representative model of the phenomenon, albeit not fully justified mathematically. We believe that the Saffman lift has served as a useful model in this capacity for many other studies.

Aside from our relaxation of proper scale separation, the other conditions of Saffman’s derivation [90]—placed on the various Reynolds numbers—are satisfied in the parameter regime considered in this paper. The shear Reynolds number, $\text{Re}_G = Ga^2/\nu$, describes the squared ratio of the particle size to the Saffman length, and should be much smaller than

unity. This Reynolds number is approximately $\text{Re}_G \sim \text{Re}(a/R)^2 |\boldsymbol{\omega}|/\Omega$, and thus satisfies its condition by virtue of (B.7). The ‘slip’ Reynolds number, $\text{Re}_p = a|\mathbf{V}_p - \mathbf{u}|/\nu$, is more difficult to ascertain a priori, but based on the analysis that follows in this paper (demonstrated in equation (B.62)), the slip velocity is dominated by the Faxén correction and approximately $a^2|\nabla^2\mathbf{u}|$. In the Stokes layer, where this correction is most active, the Laplacian of the fluid velocity scales like $|\boldsymbol{\omega}|/\delta_s$, and thus $\text{Re}_p \sim \text{Re}^{3/2}(a/R)^3 |\boldsymbol{\omega}|/\Omega$, also much less than unity. Using these scalings, the requirement that $\text{Re}_s/\text{Re}_G^{1/2} \ll 1$ is also met.

For arbitrary shear flows, different generalizations of the Saffman lift can be found in the literature: Tio et al. [101] used expressions involving coordinate-independent fluid shear rate and the norm of the particle slip velocity. However, the form used here, due to Ferry and Balachandar [96], is written in a manner that is linearly dependent on the particle velocity. These two formulations are not equivalent for arbitrary shear flows but reduce to the same formula with Saffman’s assumptions. In unreported tests, only minor differences in the transport of inertial particles were observed between these two formulas. We retain the second formulation whose linearity in the slip velocity will be helpful for deriving the asymptotic expansion of the Maxey–Riley equation.

It should also be emphasized that the particle transport model omits other effects—namely, hydrodynamic interactions with the wall of the oscillator—that are undoubtedly significant in some parts of the trajectory. In particular, both the tangential and normal components of the particle’s motion would be slowed relative to the fluid during encounters with the wall, providing an additional mechanism for the particle to be pushed toward the center of the streaming cell. Rather than compute the full hydrodynamics of these encounters, however, we rely instead on capturing similar effects in a manner that does not require a full coupling with the fluid flow field: the Faxén correction to generate a relative slip velocity, the Saffman lift to effect transverse motions, and kinematic constraints to prevent penetration.

Overall, these simplifying assumptions, while omitting some of the physics of the particle

transport, are made in order to obtain a model that can predict long-range trajectories in extended arrays of oscillators while preserving the basic mechanisms of particle trapping.

B.2.2 The inertial particle velocity field

The velocity $\mathbf{V}_p(t)$ is clearly a quantity associated with a particle-centered (i.e., Lagrangian) perspective. However, our treatment in this paper benefits greatly from changing our view of particle motion into an Eulerian perspective: the velocity and other quantities observed at a fixed position \mathbf{x} are those attributable to the inertial particle currently occupying that position. That is, we define the inertial particle velocity field $\mathbf{v}(\mathbf{x}, t)$ such that

$$\frac{d\mathbf{x}_p}{dt} = \mathbf{v}(\mathbf{x}_p(t), t) = \mathbf{V}_p(t), \quad (\text{B.13})$$

By differentiating this expression, it is clear that the time derivative of \mathbf{v} following the particle trajectory (B.9) is identical to $d\mathbf{V}_p/dt$.

This definition $\mathbf{v}(\mathbf{x}, t)$, and much of the remainder of this section, draw closely from the work of Ferry and Balachandar [96]. We briefly review the treatment here, and adapt it to account for the Faxén corrections, which were neglected by Ferry and Balachandar [96] but which we expect are non-negligible in the current context. It is important to note that the integral curves (i.e., pathlines) of this time-varying inertial particle velocity field describe every possible inertial particle trajectory. The flow map of this field, once obtained, provides a comprehensive solution for inertial particle transport, an extremely valuable result.

However, it should also be noted that the definition of \mathbf{v} depends on our choice of initial condition for the particle: with a different choice, a different particle would generally occupy a position \mathbf{x} at time t . Ferry and Balachandar [96] reason that the trajectories for two different choices of initial condition converge toward each other over time, losing memory of their different initial velocities.

From hereon, we will presume that independent and dependent variables have been non-dimensionalized by the characteristic time and length scales of the flow. In our viscous

streaming context, it is reasonable to take these, respectively, as the fast flow time scale—the inverse of the oscillation frequency, $1/\Omega$ —and the radius R of an oscillating cylinder. Thus, for example, \mathbf{u} and \mathbf{v} will henceforth denote the fluid and inertial particle velocity fields scaled by ΩR , \mathbf{x} will be the position scaled by R , and time t will represent the dimensional time multiplied by Ω .

Using the definitions presented above and some simple manipulation, we can rewrite (B.8) in a more compact dimensionless form:

$$\frac{d\mathbf{v}}{dt} = \frac{1}{\tau} (\mathbf{u} + \mathbf{q}_F - \mathbf{v}) + \beta \frac{D\mathbf{u}}{Dt} + \frac{\beta}{5} \frac{d\mathbf{q}_F}{dt} + \sqrt{\frac{\beta}{\tau}} (\mathcal{L}_B [\mathbf{u} + \mathbf{q}_F - \mathbf{v}] + \mathcal{L}_S [\mathbf{u} - \mathbf{v}]), \quad (\text{B.14})$$

in which, for convenience, we have defined a Faxén correction velocity,

$$\mathbf{q}_F = \frac{1}{6} (a/R)^2 \nabla^2 \mathbf{u}. \quad (\text{B.15})$$

We have also defined two dimensionless parameters: a density ratio parameter, β , and a particle Stokes number, τ , respectively, as

$$\beta \equiv \frac{3}{2\rho_p/\rho_f + 1}, \quad \tau \equiv \frac{\Omega a^2}{3\beta\nu}. \quad (\text{B.16})$$

This latter parameter represents the ratio of the characteristic response time of the Stokes drag on the particle to the fastest characteristic flow time scale.

B.2.3 Asymptotic expansion of the Maxey–Riley equation in small Stokes number

The two parameters introduced in equation (B.16), β and τ , each play an important role in dictating the behavior of the inertial particle motion. For example, $\beta = 1$ represents a neutrally buoyant particle, which, in the absence of the Faxén corrections, remains on the same trajectory as a fluid particle. The Stokes number, τ , is proportional to $\text{Re}(a/R)^2$, which we have already assumed to be small in (B.7), and thus make the same assumption for the Stokes number: $\tau \ll 1$. In the absence of Faxén corrections, we would expect this small

Stokes number to quickly penalize deviations of the inertial particle's velocity from that of the surrounding fluid (and to render the governing equation (B.14) numerically stiff). Maxey [95] used this argument to develop an asymptotic expansion in τ for the inertial particle velocity's departure from that of the fluid. Later, Ferry and Balachandar [96] extended this expansion to include all of the terms that we have included in the MR equation (B.14) except for the Faxén correction.

In regions of shear, the Faxén correction cannot be neglected a priori. For example, in the Stokes boundary layer formed by an oscillating cylinder, whose thickness scales like $1/\sqrt{\text{Re}}$, one expects $\mathbf{q}_F \sim \text{Re}(a/R)^2 \mathbf{u}$, and thus, $\mathbf{q}_F \sim \tau \mathbf{u}$. For the sake of keeping our analysis in this section general, we will not yet explicitly invoke this scaling of \mathbf{q}_F . It is important simply to note that the Faxén correction is at least comparable to the other dominant terms in the analysis. We will include the Faxén correction only as \mathbf{q}_F , unadorned with scaling; once its scaling in τ is determined, its placement in the asymptotic expansion can be adjusted accordingly. We will do so for the case of an oscillating cylinder.

Note that this inclusion changes the apparent target velocity at vanishing τ from \mathbf{u} to $\mathbf{u} + \mathbf{q}_F$, and also changes the lowest power of τ in the expansion from τ to $\tau^{1/2}$, as we shall see. Adapting the approach of Maxey [95] and Ferry and Balachandar [96], we write the inertial particle velocity field in terms of the fluid velocity field as

$$\mathbf{v} = \mathbf{u} + \mathbf{q}_F + \tau^{1/2} \mathbf{q}. \quad (\text{B.17})$$

The deviation of the particle velocity from the target is now borne by the third term on the right-hand side of (B.17). A derivation of the resulting equation for \mathbf{q} is presented in the appendix of Le Provost et al. [35]. In the course of that derivation, they use the expected scaling of \mathbf{q}_F in the Stokes layer surrounding the oscillating cylinder, and furthermore, neglect the Basset memory term as consistent with the analysis of Chong et al. [81]. The resulting expression is

$$\mathbf{v} = \mathbf{u} + \tau \mathbf{a} - \tau^{3/2} \beta^{1/2} \mathcal{L}_S [\mathbf{a}] + O(\tau^2), \quad (\text{B.18})$$

in which the fluid acceleration force, \mathbf{a} , has been defined as

$$\mathbf{a} \equiv (\beta - 1) \frac{D\mathbf{u}}{Dt} + \frac{1}{2} \frac{\beta}{\text{Re}} \nabla^2 \mathbf{u}. \quad (\text{B.19})$$

Equation (B.18) forms one of the cornerstones of our proposed method for accelerated simulation of inertial particle transport, since it provides a velocity field that describes this transport everywhere, entirely in terms of the local fluid velocity and its derivatives. However, before we proceed to the distillation of this equation into fast and slow time-scales, we make a few observations. First, it is important to note that the equation reduces to that of Ferry and Balachandar [96] and Haller et al. [102] when the Faxén correction velocity is omitted. In that situation, the acceleration force (B.19) reduces to

$$\mathbf{a} = (\beta - 1) \frac{D\mathbf{u}}{Dt}. \quad (\text{B.20})$$

We can observe from this reduced form that, without the Faxén corrections, the inertial particle’s motion can only depart from that of the fluid if the particle is not neutrally buoyant (i.e., if $\beta \neq 1$). But the retention of these Faxén correction terms into the expanded field emphasizes the observations made by Chong et al. [81]: When the particle is neutrally buoyant or nearly so, deviation of the particle’s motion from that of the fluid is solely brought about by the Faxén correction velocity, and the particle’s subsequent dynamics are dominated by the Saffman lift. In equation (B.18), these observations are confirmed to be the two dominant disturbances from the fluid velocity. For lighter or heavier particles, the dynamics are effected by a mixture of this influence with that from fluid acceleration.

B.3 Development of the equations for mean particle transport

In the previous section, we obtained a velocity field for inertial particle transport that derives from the velocity field of the fluid. The trajectories of both fluid and inertial particles are described by the general transport equation

$$\frac{d\mathbf{x}_p}{dt} = \mathbf{w}(\mathbf{x}_p(t), t), \quad \mathbf{x}_p(0) = \mathbf{x}_0, \quad (\text{B.21})$$

where the generic velocity field \boldsymbol{w} can be interpreted as either the fluid velocity field \boldsymbol{u} or the inertial particle velocity field \boldsymbol{v} . In the current discussion it is not important to make the distinction, and we will use the generic term ‘particle’ to describe either a fluid particle or an inertial particle.

In the flows associated with \boldsymbol{u} or \boldsymbol{v} , the trajectories contain a mixture of fast (fluctuating) and slow (mean) scales. The main objective of this work is to seek the mean trajectories of such particles directly, skipping over the integration of the fast scales to the extent possible. This challenge to derive equations resolved only at the slower (or larger) scales exists in many realms of physics, e.g., turbulence or climate modeling. The classical Reynolds decomposition [103] of a fluid quantity (such as velocity or pressure) into a mean and a fluctuating component provides the basic machinery for developing such equations from an Eulerian perspective, i.e., for fluctuations observed from a fixed spatial location. Because we are interested in this paper in seeking the slow-scale transport of individual particles, we cannot perform a standard Eulerian average of the Maxey–Riley equation; rather, we need to average it in a Lagrangian sense, i.e., for a fixed particle label. Following the work of Andrews and McIntyre [94, 104], we will introduce a Reynolds-like decomposition of the motion of a particle into a slow (mean) Lagrangian component and a fast (fluctuating) component. In the context of Lagrangian fluid stability, Bernstein [105] first argued that the fluctuating component of the motion of a particle can be derived from an Eulerian disturbed displacement field, $\boldsymbol{\xi}(\boldsymbol{x}, t)$, evaluated at the mean Lagrangian position of the particle. This led to Generalized Lagrangian Mean (GLM) theory, developed by Andrews and McIntyre [94, 104], who successfully applied the theory to wave problems in the contexts of stratified and rotating fluid flows. Holm analyzed the GLM theory from a geometric point of view and derived the Lagrangian averaged Navier-Stokes-alpha (LANS- α) model for turbulent flows [106, 107, 108].

We will use GLM theory to provide a framework in which to analyze particle transport into fast and slow components. The basic aspects of the theory’s application are outlined

in Section B.3.1. Like the Reynolds-averaged Navier–Stokes equations, GLM theory still leaves the treated equations with a closure problem, analogous to the one encountered in the Reynolds-averaged Navier–Stokes equations. However, rather than seek to replace the influence of the fastest scales with a model, we instead account for their influence by explicitly computing the disturbed displacement field, $\boldsymbol{\xi}$. Following the work of Holm [108], we formulate a simplified form of the equation for $\boldsymbol{\xi}$ for small disturbances. In Section B.3.2, we will clarify this equation via an asymptotic expansion in the small oscillation amplitude that underpins viscous streaming theory.

B.3.1 Development of the Lagrangian mean field equations

Given a time-varying Eulerian field $f(\boldsymbol{x}, t)$ of arbitrary tensor rank, we can define the following averaging operator

$$\bar{f}(\boldsymbol{x}, t) = \frac{1}{T} \int_{t-T}^t f(\boldsymbol{x}, t') dt'. \quad (\text{B.22})$$

In anticipation of the emergence of two timescales, it is important to note that this operator is intended to average the fast scales and leave the slow scales unaffected. Following Holm [108], we could formally define f with distinct dependencies on time in these two separate scales. By expanding the averaging operator in the ratio $t_f/t_s \ll 1$, it is straightforward to show that the operator’s leading-order behavior preserves the slow behavior of f , provided that $t_s \ll T \ll t_f$. In the case of strictly periodic fast scales, it is sufficient for T to be the period of oscillation or some integer multiple thereof.

We will refer to the field defined in this fashion as the Eulerian mean field. The averaging operator has the following properties:

- Linearity: For constant scalars a, b and two Eulerian fields f and g , $\overline{af + bg} = a\bar{f} + b\bar{g}$
- Idempotence: $\overline{\bar{f}} = \bar{f}$

It also commutes with spatial and time derivatives, but importantly (and famously in the

Navier–Stokes equations), does not commute with the advection operator: That is,

$$\overline{\mathbf{u} \cdot \nabla f} \neq \bar{\mathbf{u}} \cdot \nabla \bar{f}. \quad (\text{B.23})$$

In this paper, we seek the mean paths of fluid or inertial particles, and the mean variation of quantities along those paths. The Eulerian mean, assessed at a fixed location, is not the appropriate measure of average in this context. However, it can be used to ‘induce’ a definition of a Lagrangian mean, $\overline{(\cdot)}^L$: that is, an average of a field taken along the trajectory of a particle, for fixed Lagrangian label. This definition, and several useful tools associated with it, are provided by the Generalized Lagrangian Mean (GLM) theory of Andrews and McIntyre [94]. In this section, we present a basic outline of the Generalized Lagrangian Mean theory of Andrews and McIntyre [94].

Consider a flow map, $\mathbf{X}(\mathbf{x}_0, t)$, illustrated in Figure B.1, from material coordinates (i.e., Lagrangian label) \mathbf{x}_0 in a reference space \mathcal{X}_0 to a location in a space-time configuration space $\mathcal{X}_t \times [0, \infty)$. In this latter space, each slice \mathcal{X}_t represents an evolved form of \mathcal{X}_0 at some later time t ; we assume that $\mathcal{X}_t = \mathcal{X}_0$ at $t = 0$. We denote the velocity field associated with this map as \mathbf{w} , and will refer to the infinitesimal bit of material associated with \mathbf{x}_0 as a ‘particle’. We will use the notation $\mathbf{w}_0(\mathbf{x}_0, t)$ to denote the velocity \mathbf{w} in its Lagrangian form, i.e.,

$$\mathbf{w}_0(\mathbf{x}_0, t) \equiv \frac{\partial \mathbf{X}(\mathbf{x}_0, t)}{\partial t}. \quad (\text{B.24})$$

As is typical, we regard the flow map as invertible, so that we can uniquely associate a Lagrangian label with any fixed location. To distinguish from later terminology, we will refer to the location $\mathbf{X}(\mathbf{x}_0, t)$ as the actual location of the particle \mathbf{x}_0 at time t , and the velocity $\mathbf{w}_0(\mathbf{x}_0, t)$ as the particle’s actual velocity.

We can always think of $\mathbf{X}(\mathbf{x}_0, t)$ as a composition of two maps,

$$\mathbf{X}(\mathbf{x}_0, t) = \mathbf{X}^\xi(\mathbf{Y}(\mathbf{x}_0, t), t), \quad (\text{B.25})$$

as illustrated in Figure B.1. The first, $\mathbf{Y}(\mathbf{x}_0, t)$, maps \mathbf{x}_0 in \mathcal{X}_0 to some location in another

configuration space $\mathcal{Y}_t \times [0, \infty)$ at time t ; and the second,

$$\mathbf{X}^\xi(\mathbf{x}, t) \equiv \mathbf{x} + \boldsymbol{\xi}(\mathbf{x}, t), \quad (\text{B.26})$$

maps from that location in the slice \mathcal{Y}_t at t to the actual location in the corresponding slice \mathcal{X}_t at time t .

We will require that \mathcal{Y}_t coincides with \mathcal{X}_0 at $t = 0$, just as \mathcal{X}_t does. In other words, the two spaces are identical at the initial instant. At all times, we assume that the mapping $\mathbf{Y}(\cdot, t)$ —like $\mathbf{X}(\cdot, t)$ —is invertible. Also, just as $\mathbf{X}(\cdot, t)$ is associated with the velocity field \mathbf{w} , the flow map $\mathbf{Y}(\cdot, t)$ is generated by its own velocity field, \mathbf{W} , i.e.,

$$\mathbf{W}_0(\mathbf{x}_0, t) \equiv \frac{\partial \mathbf{Y}(\mathbf{x}_0, t)}{\partial t}. \quad (\text{B.27})$$

From the definition (B.26) and its use in the composition (B.25), it is clear that $\boldsymbol{\xi}(\mathbf{x}, t)$ provides an additive correction from the location provided by mapping $\mathbf{Y}(\mathbf{x}_0, t)$ to the actual location of the particle, provided by $\mathbf{X}(\mathbf{x}_0, t)$:

$$\boldsymbol{\xi}(\mathbf{Y}(\mathbf{x}_0, t), t) = \mathbf{X}(\mathbf{x}_0, t) - \mathbf{Y}(\mathbf{x}_0, t). \quad (\text{B.28})$$

If we differentiate this with respect to time (keeping the Lagrangian label fixed), then by the chain rule we obtain

$$\frac{\partial}{\partial t} \boldsymbol{\xi}(\mathbf{Y}(\mathbf{x}_0, t), t) + \mathbf{W}_0(\mathbf{x}_0, t) \cdot \nabla \boldsymbol{\xi}(\mathbf{Y}(\mathbf{x}_0, t), t) = \mathbf{w}_0(\mathbf{x}_0, t) - \mathbf{W}_0(\mathbf{x}_0, t), \quad (\text{B.29})$$

which relates the velocity fields, \mathbf{w} and \mathbf{W} , associated with each flow map.

For GLM theory and our later applications, it will be very useful to regard $\boldsymbol{\xi}$ as an Eulerian vector field, $\boldsymbol{\xi}(\mathbf{x}, t)$. To make sense of this interpretation, let us make use of the inverse mapping $\mathbf{x}_0 = \mathbf{Y}^{-1}(\mathbf{x}, t)$, uniquely associating any fixed location \mathbf{x} in \mathcal{Y}_t to the particle currently residing there at time t via the mapping $\mathbf{Y}(\cdot, t)$. Then relation (B.28) can be written as

$$\boldsymbol{\xi}(\mathbf{x}, t) = \mathbf{X}(\mathbf{Y}^{-1}(\mathbf{x}, t), t) - \mathbf{x}, \quad (\text{B.30})$$

showing that, at each point \mathbf{x} , $\boldsymbol{\xi}(\mathbf{x}, t)$ provides the actual location of the particle relative to \mathbf{x} itself. In fact, it is clear from (B.25) that $\mathbf{X}^\xi(\mathbf{x}, t) = \mathbf{X}(\mathbf{Y}^{-1}(\mathbf{x}, t), t)$: the two sides of the equality just represent two different routes to the same map, as can be observed in Figure B.1. Indeed, through this map formalism, any field quantity ϕ can be viewed from one of three perspectives: as an Eulerian field quantity in the configuration space \mathcal{Y}_t , $\phi^\xi(\cdot, t)$; as the ‘actual’ Eulerian field in \mathcal{X}_t , $\phi(\cdot, t)$; or as a Lagrangian (i.e, particle-centered) field, $\phi_0(\cdot, t)$, associating ϕ to particles in \mathcal{X}_0 . They are related by

$$\phi^\xi(\mathbf{x}, t) \equiv \phi(\mathbf{X}^\xi(\mathbf{x}, t), t) \equiv \phi_0(\mathbf{Y}^{-1}(\mathbf{x}, t), t). \quad (\text{B.31})$$

Similarly, we can rewrite the velocity relation (B.29) as

$$\frac{\partial}{\partial t} \boldsymbol{\xi}(\mathbf{x}, t) + \mathbf{W}(\mathbf{x}, t) \cdot \nabla \boldsymbol{\xi}(\mathbf{x}, t) = \mathbf{w}^\xi(\mathbf{x}, t) - \mathbf{W}(\mathbf{x}, t), \quad (\text{B.32})$$

where we have defined the Eulerian velocity field, $\mathbf{W}(\mathbf{x}, t) \equiv \mathbf{W}_0(\mathbf{Y}^{-1}(\mathbf{x}, t), t)$, associated with the flow map $\mathbf{Y}(\cdot, t)$. By $\mathbf{w}^\xi(\mathbf{x}, t)$, we denote the actual velocity of the particle currently mapped to \mathbf{x} by $\mathbf{Y}(\cdot, t)$:

$$\mathbf{w}^\xi(\mathbf{x}, t) \equiv \mathbf{w}_0(\mathbf{Y}^{-1}(\mathbf{x}, t), t). \quad (\text{B.33})$$

Equation (B.32) shows that this velocity differs from $\mathbf{W}(\mathbf{x}, t)$ by a correction described by the rate of change of the $\boldsymbol{\xi}$ field measured along the $\mathbf{Y}(\cdot, t)$ trajectory passing through \mathbf{x} at time t .

With this formalism in place, following Andrews and McIntyre [94], we can define the Lagrangian mean $\overline{(\cdot)}^L$ of any field $\phi^\xi(\mathbf{x}, t)$ as equal to the Eulerian mean along the trajectory followed by \mathbf{x} in the configuration space \mathcal{X}_t under the map $\mathbf{X}^\xi(\mathbf{x}, t) \equiv \mathbf{x} + \boldsymbol{\xi}(\mathbf{x}, t)$:

$$\overline{\phi(\mathbf{x}, t)}^L \equiv \overline{\phi(\mathbf{x} + \boldsymbol{\xi}(\mathbf{x}, t), t)} \equiv \overline{\phi^\xi(\mathbf{x}, t)}. \quad (\text{B.34})$$

Thus far, we have not specified anything about the map $\mathbf{Y}(\cdot, t)$ and its associated field $\boldsymbol{\xi}(\mathbf{x}, t)$. GLM theory assigns $\boldsymbol{\xi}(\mathbf{x}, t)$ the role of a fluctuation field, and furthermore, asserts that it has zero Eulerian mean and that the velocity field $\mathbf{W}(\mathbf{x}, t)$ is its own mean.

$$\overline{\boldsymbol{\xi}(\mathbf{x}, t)} = 0, \quad \overline{\mathbf{W}(\mathbf{x}, t)} = \mathbf{W}(\mathbf{x}, t). \quad (\text{B.35})$$

We can immediately note that, by taking the mean of equation (B.26) and applying the first of these axioms, we get

$$\overline{\mathbf{X}^\xi(\mathbf{x}, t)} = \mathbf{x}. \quad (\text{B.36})$$

That is, the location \mathbf{x} in $\mathcal{Y}_t \times [0, \infty)$ maps on average to the location \mathbf{x} in $\mathcal{X}_t \times [0, \infty)$. Furthermore, by taking the Eulerian mean of the relationships (B.30) and (B.32), the axioms (B.35) immediately imply other important relationships between the configuration spaces $\mathcal{Y}_t \times [0, \infty)$ and $\mathcal{X}_t \times [0, \infty)$: The trajectory in $\mathcal{Y}_t \times [0, \infty)$ described by the flow map $\mathbf{Y}(\mathbf{x}_0, t)$ is the mean of the actual trajectory $\mathbf{X}(\mathbf{x}_0, t)$ in $\mathcal{X}_t \times [0, \infty)$, so that (B.28), rewritten trivially as

$$\mathbf{X}(\mathbf{x}_0, t) = \mathbf{Y}(\mathbf{x}_0, t) + \boldsymbol{\xi}(\mathbf{Y}(\mathbf{x}_0, t), t), \quad (\text{B.37})$$

represents a Reynolds decomposition of the trajectory, and $\boldsymbol{\xi}$ can be called the disturbed displacement field; and the velocity \mathbf{W} represents the Lagrangian mean of the actual velocity \mathbf{w} ,

$$\mathbf{W}(\mathbf{x}, t) = \overline{\mathbf{w}(\mathbf{x}, t)}^L, \quad (\text{B.38})$$

or simply, $\mathbf{W} = \overline{\mathbf{w}}^L$. From hereon, we will refer to $\mathcal{Y}_t \times [0, \infty)$ as the mean configuration space.

If we define the Lagrangian mean material derivative as the rate of change while moving along a mean trajectory,

$$\frac{D^L}{Dt} \equiv \frac{\partial}{\partial t} + \overline{\mathbf{w}}^L \cdot \nabla, \quad (\text{B.39})$$

and the Lagrangian disturbance velocity as the difference between the actual velocity and the Lagrangian mean velocity,

$$\mathbf{w}^l \equiv \mathbf{w}^\xi - \overline{\mathbf{w}}^L, \quad (\text{B.40})$$

then we can rewrite the velocity relationship (B.32) as

$$\frac{D^L \boldsymbol{\xi}}{Dt} = \mathbf{w}^l. \quad (\text{B.41})$$

This equation provides the basis for generating the actual trajectory of the particle while following the particle's mean trajectory. It is useful to note that the Lagrangian mean of this equation is identically zero.

The analysis of other fields follows from the definitions thus far. For a general field ϕ , if we differentiate the relationship (B.31) with respect to time and apply the chain rule, then, with the help of (B.32), it can be shown that

$$\frac{D^L \phi^\xi}{Dt}(\mathbf{x}, t) = \left(\frac{D\phi}{Dt} \right)^\xi(\mathbf{x}, t) \equiv \left(\frac{\partial \phi}{\partial t} + \mathbf{w} \cdot \nabla \phi \right)(\mathbf{X}^\xi(\mathbf{x}, t), t). \quad (\text{B.42})$$

In other words, the rate of change of ϕ^ξ measured while moving along the mean trajectory is identical to the rate of change measured while moving along the actual trajectory. That is, no information has been lost while following a different trajectory. The Reynolds decomposition of this field follows naturally from the Lagrangian mean,

$$\phi^\xi(\mathbf{x}, t) = \overline{\phi}^L(\mathbf{x}, t) + \phi^l(\mathbf{x}, t). \quad (\text{B.43})$$

In particular, the Lagrangian disturbance velocity is defined by $\mathbf{v}^l(\mathbf{x}, t) \equiv \mathbf{v}^\xi(\mathbf{x}, t) - \overline{\mathbf{v}}^L(\mathbf{x}, t)$ [35].

B.3.1.1 The basic equations for mean particle transport

Using the notation for GLM theory defined in the previous section and illustrated in Figure B.1, our objective is to seek the mean flow map $\mathbf{Y}(\cdot, t)$ for particular values of the particle label \mathbf{x}_0 . The equation generating this trajectory for a specific particle \mathbf{x}_0 follows directly from equation (B.27). When we substitute the velocity with the Lagrangian mean velocity using relation (B.38), we obtain the kinematic equation for a mean particle trajectory:

$$\frac{d\overline{\mathbf{x}}^L}{dt} = \overline{\mathbf{w}}^L(\overline{\mathbf{x}}^L(t), t), \quad \overline{\mathbf{x}}^L(0) = \mathbf{x}_0, \quad (\text{B.44})$$

where we have used the shorthand notation $\overline{\mathbf{x}}^L(t) \equiv \mathbf{Y}(\mathbf{x}_0, t)$ to denote the mean trajectory of a single particle, \mathbf{x}_0 , and explicitly included its initial condition. By definition, the




Figure B.1: Schematic of flow maps used in this work. \mathcal{X}_t and \mathcal{Y}_t represent slices of the configuration spaces $\mathcal{X}_t \times [0, \infty)$ and $\mathcal{Y}_t \times [0, \infty)$ at some instant t . Illustrations of particle trajectories are shown as colored curves (though strictly, these trajectories would proceed along the time axis of the respective space).

Lagrangian mean velocity field requires averaging while following the actual trajectory of the particle in $\mathcal{X}_t \times [0, \infty)$, obtained by adding the local disturbed displacement, $\boldsymbol{\xi}(\bar{\mathbf{x}}^L(t), t)$ to the trajectory in $\mathcal{Y}_t \times [0, \infty)$. The disturbed displacement field can be generated from the velocity field via its transport equation (B.41), which we rewrite here with relevant definitions for the purpose of elucidating the underlying (and thus far, exact) computational problem:

$$\begin{aligned} \frac{\partial \boldsymbol{\xi}}{\partial t} &= -\bar{\mathbf{w}}^L \cdot \nabla \boldsymbol{\xi} + \mathbf{w}^\xi - \bar{\mathbf{w}}^L, \\ \boldsymbol{\xi}(\mathbf{x}, 0) &= 0, \\ \mathbf{w}^\xi(\mathbf{x}, t) &= \mathbf{w}(\mathbf{x} + \boldsymbol{\xi}(\mathbf{x}, t), t), \\ \bar{\mathbf{w}}^L(\mathbf{x}, t) &= \frac{1}{T} \int_{t-T}^t \mathbf{w}^\xi(\mathbf{x}, t') dt'. \end{aligned} \tag{B.45}$$

We have included here the initial condition on $\boldsymbol{\xi}$, which was established by requiring that $\mathcal{Y}_t = \mathcal{X}_t = \mathcal{X}_0$ at $t = 0$. Figure B.2 illustrates the relationships between the mean and actual trajectories.




Figure B.2: Mean (green) and actual (blue) particle trajectories with initial position \mathbf{x}_0 . The position \mathbf{x}^ξ is the *actual* position whose mean is \mathbf{x} . Adapted from Bühler et al. [1].

The set of equations (B.44) and (B.45) does not obviously achieve our goal of ‘skipping over’ the fast (oscillatory) time scales of the flow to accelerate the solution for mean trajectories. However, it is important to observe that the coupled equations (B.45) are Eulerian in the mean configuration space $\mathcal{Y}_t \times [0, \infty)$. Furthermore, when they are supplied with the actual velocity field, \mathbf{w} , they can be solved a priori to generate the (slowly-varying) Lagrangian mean field $\overline{\mathbf{w}}^L(\mathbf{x}, t)$, either simultaneously with \mathbf{w} or in a subsequent procedure. With this Lagrangian mean velocity field in hand, the slow particle trajectories are easily obtained with no further regard for the fast (oscillatory) timescale by integrating equation (B.44).

In Section B.3.2.2, we will clarify features of the equations (B.45) that can be used to simplify our task. In particular, we will make the assumption that the disturbed displacement field $\boldsymbol{\xi}$ is small everywhere (compared with the characteristic length scale of the flow, e.g., R) and exploit this assumption to reduce the equations to a more computationally tractable form.

B.3.1.2 A note on the disturbed displacement field on moving surfaces

In the context of this paper, the problems we wish to solve involve flows whose fastest time-scales are generated by impenetrable surfaces undergoing motion consisting entirely of rapid disturbances. The partial differential equation in (B.45) for the disturbed displacement field, $\boldsymbol{\xi}(\boldsymbol{x}, t)$, is already closed—it purely describes transport and has no need of boundary conditions. However, it is nonetheless useful to reconcile the surface motion (which we will assume is prescribed) with our definition of the mean configuration space $\mathcal{Y}_t \times [0, \infty)$ and the field $\boldsymbol{\xi}(\boldsymbol{x}, t)$; intuitively, we should expect that the surface motion is somewhat simpler in this space. In fact, we will restrict our attention in this work to surfaces that are at rest in the mean configuration space, and consider the implications of this restriction. After that, we will discuss its rationale.

Let us consider a subset of the particles in \mathcal{X}_0 to comprise a reference surface, S_{b_0} . This surface represents the interface between the body and the surrounding fluid at $t = 0$. We can also interpret S_{b_0} as comprising multiple disconnected surfaces in case there are many bodies. The (actual) subsequent motion of any surface particle $\boldsymbol{x}_0 \in S_{b_0}$ is described by a map $\boldsymbol{X}_b(\boldsymbol{x}_0, t)$, and the collection of all such points is denoted by $S_b(t)$. For the velocity field \boldsymbol{w} , whether it represents the fluid or inertial particle motion, we will require that the no-slip and no-penetration conditions are both enforced at the surface. (It should be apparent from the inertial particle velocity field (B.18) that this field does not satisfy these conditions even if the fluid velocity does. We will augment the field with a constraint, to be discussed in Section B.3.3.) We thus insist that $\boldsymbol{X}(\boldsymbol{x}_0, t) = \boldsymbol{X}_b(\boldsymbol{x}_0, t)$ for all $\boldsymbol{x}_0 \in S_{b_0}$, which ensures that

$$\boldsymbol{w}_0(\boldsymbol{x}_0, t) = \frac{\partial \boldsymbol{X}_b(\boldsymbol{x}_0, t)}{\partial t} \quad (\text{B.46})$$

for all such particles.

We also define the map $\boldsymbol{Y}_b(\boldsymbol{x}_0, t)$ and insist that it, too, agrees with the overall map to $\mathcal{Y}_t \times [0, \infty)$ for all particles in S_{b_0} . But, in line with our assumption that the surface is at

rest in the mean configuration space, we require that this map is the identity:

$$\mathbf{Y}(\mathbf{x}_0, t) = \mathbf{Y}_b(\mathbf{x}_0, t) \equiv \mathbf{x}_0, \quad \mathbf{x}_0 \in S_{b0}. \quad (\text{B.47})$$

Applying our Lagrangian form of the Reynolds decomposition (B.37) to particles in S_{b0} , we can relate these surface maps to the disturbed displacement field:

$$\boldsymbol{\xi}(\mathbf{x}_0, t) = \boldsymbol{\xi}_b(\mathbf{x}_0, t) \equiv \mathbf{X}_b(\mathbf{x}_0, t) - \mathbf{x}_0, \quad \mathbf{x}_0 \in S_{b0}. \quad (\text{B.48})$$

where we have defined a surface displacement field, $\boldsymbol{\xi}_b(\mathbf{x}, t)$ for particles in the surface. Note that this relationship also ensures that $\boldsymbol{\xi}(\mathbf{x}_0, 0) = 0$, as desired.

To express this relationship in the usual Eulerian form of the field $\boldsymbol{\xi}$, we will define \bar{S}_b as the fixed position of the surface in the mean configuration space. By assumption, $\bar{S}_b = S_{b0}$, and clearly, the coordinates of any fixed location on this surface $\mathbf{x} \in \bar{S}_b$ are equal to the particle label there, $\mathbf{x} = \mathbf{x}_0$. Thus, we obtain the following simple expression for the disturbed displacement field at points on the surface:

$$\boldsymbol{\xi}(\mathbf{x}, t) = \boldsymbol{\xi}_b(\mathbf{x}, t) \equiv \mathbf{X}_b(\mathbf{x}, t) - \mathbf{x}, \quad \mathbf{x} \in \bar{S}_b. \quad (\text{B.49})$$

Applying the GLM axiom that the mean of $\boldsymbol{\xi}$ vanishes identically, it is clear that requiring the surface to remain at rest in the mean configuration space is equivalent to requiring that $\overline{\boldsymbol{\xi}_b} = 0$, or equivalently,

$$\overline{\mathbf{X}_b(\mathbf{x}, t)} = \mathbf{x}. \quad (\text{B.50})$$

That is, the location $\mathbf{x} \in \bar{S}_b$, as intuitively expected, is identically the mean position of this moving surface point on $S_b(t)$ in the full configuration space $\mathcal{X}_t \times [0, \infty)$, and $\boldsymbol{\xi}(\mathbf{x}, t)$ entirely describes its motion in that space. It should also be clear that, being stationary in the mean configuration space implies that the Lagrangian mean velocity, $\overline{\mathbf{w}^L}(\mathbf{x}, t)$, is exactly zero at all points on the surface:

$$\overline{\mathbf{w}^L}(\mathbf{x}, t) = 0, \quad \mathbf{x} \in \bar{S}_b. \quad (\text{B.51})$$

In the specific case of the fluid velocity field, $\mathbf{w} \rightarrow \mathbf{u}$, the no-slip boundary condition (B.46) can be rewritten in terms of the surface displacement field $\boldsymbol{\xi}_b$ as

$$\mathbf{u}(\mathbf{x}_0 + \boldsymbol{\xi}_b(\mathbf{x}_0, t), t) = \frac{\partial \boldsymbol{\xi}_b(\mathbf{x}_0, t)}{\partial t}, \quad \mathbf{x}_0 \in S_{b0}, \quad (\text{B.52})$$

where $\overline{\boldsymbol{\xi}_b} = 0$ and \mathbf{x}_0 coincides with the mean location of the surface point on $S_b(t)$.

As an example of an admissible surface motion, relevant for the results we will show later in this work, let us consider a rigid body in oscillatory translational motion. Then we can write the full surface map for any surface label $\mathbf{x}_0 \in S_{b0}$ as

$$\mathbf{X}_b(\mathbf{x}_0, t) = \mathbf{x}_c(t) + (\mathbf{x}_0 - \mathbf{x}_{c0}), \quad (\text{B.53})$$

where $\mathbf{x}_c(t)$ is the time-varying position of the centroid of the body, and $\mathbf{x}_c(0) = \mathbf{x}_{c0}$ is that centroid's initial location. Applying the restriction to this motion, the surface point must remain fixed at \mathbf{x}_0 in the mean configuration space, and, to ensure that this truly is the mean space,

$$\overline{\mathbf{x}_c^L} = \mathbf{x}_{c0}. \quad (\text{B.54})$$

That is, the centroid must start at its mean position. The resulting disturbed displacement field on the mean surface is described by

$$\boldsymbol{\xi}(\mathbf{x}, t) = \boldsymbol{\xi}_b(\mathbf{x}, t) \equiv \mathbf{X}_b(\mathbf{x}, t) - \mathbf{x} = \mathbf{x}_c(t) - \mathbf{x}_{c0}, \quad \mathbf{x} \in \overline{S}_b. \quad (\text{B.55})$$

In other words, when the surface is in rigid translation, the disturbed displacement field is uniform. It is straightforward to conceive of other admissible surface motions that would generate streaming flows, including oscillatory rigid rotation or time-varying deformations about a stationary mean surface.

We should observe that our restriction (B.50) precludes combinations of faster (fluctuating) motions with slower motions that have non-zero mean. We make this restriction to avoid ambiguity involving our definition of the Eulerian mean (B.22) when applied in the mean configuration space at fixed positions on or near the surface. If such a surface were

moving in this space, then it would move relative to this fixed averaging position during the averaging interval, obscuring the decomposition of the surface’s motion. Of course, by construction, any motion in the mean space is presumed to occur on a much slower time-scale than the averaging interval. In fact, if we rely on two independent measures of time, slow and fast—as, for example, by Holm [106] and others—then the surface can be treated as stationary with respect to the time averaging over the fast scales. However, we have chosen to use only a single measure of time, primarily because the time-scale associated with the slow motion effected by viscous streaming only presents itself a posteriori. Nevertheless, the viscous streaming flows we focus on in this work do not contain such slow surface motions.

In passing, we note that there is nothing in the analysis of this section that prevents us from applying it to surfaces formed from bubble (liquid–gas) interfaces. However, in that case, one would have to allow for mean transport within the mean surface. That is, although \bar{S}_b would still be stationary, $\mathbf{Y}_b(\cdot, t)$ would no longer be the identity, but would allow for mean transport on the surface.

B.3.2 Asymptotic reduction for small oscillation amplitude

The full governing equations describing the mean transport of fluid and inertial particles have now been specified. These include the equations for the velocity fields themselves—the Navier–Stokes equations (B.1) for the fluid velocity field \mathbf{u} and equation (B.18) for the inertial particle velocity field \mathbf{v} induced by this fluid motion, and the initial condition and boundary conditions on the fluid velocity (B.3) and (B.52). They also include the equations for mean transport in these velocity fields, including (B.45) for the Lagrangian mean $\bar{\mathbf{w}}^L$ of either of the velocity fields, and equation (B.44) for the mean transport $\bar{\mathbf{x}}^L(t)$ of any particle, fluid or inertial. Collectively, they represent a map from a given surface displacement field of the oscillator(s), $\boldsymbol{\xi}_b$, to the resulting mean transport of the particles in the fluid.

In this section, we aim to simplify the calculation of this mean transport by exploiting the small amplitude, $\epsilon = A/R \ll 1$, of the oscillations described by $\boldsymbol{\xi}_b$. We will help our

cause by exposing the oscillation amplitude with a unit form, Ξ_b , of the surface displacement field:

$$\boldsymbol{\xi}_b(\mathbf{x}, t) = \epsilon \Xi_b(\mathbf{x}, t), \quad (\text{B.56})$$

where $\Xi_b(\mathbf{x}, t) = O(1)$. With the driving mechanism proportional to ϵ , we expect that all other quantities, including either velocity field embodied by \mathbf{w} , and its associated disturbed displacement field, $\boldsymbol{\xi}$, are themselves proportional to ϵ at leading order (and we do not expect flow instabilities to emerge in this parameter regime that might change this fact). Thus, we will expand all such quantities in powers of ϵ ,

$$\mathbf{w} = \epsilon \mathbf{w}_1 + \epsilon^2 \mathbf{w}_2 + O(\epsilon^3), \quad \boldsymbol{\xi} = \epsilon \boldsymbol{\xi}_1 + \epsilon^2 \boldsymbol{\xi}_2 + O(\epsilon^3), \quad (\text{B.57})$$

where \mathbf{w} , as usual, could be either the fluid velocity \mathbf{u} or the inertial particle velocity \mathbf{v} . We will also use the same expansion for the pressure, p (which, more precisely, represents the pressure disturbance from ambient).

B.3.2.1 Reduction of the Navier–Stokes equations

Let us first introduce the asymptotic expansions for \mathbf{u} and p into the Navier–Stokes equations (B.1). We can also do the same for the boundary condition (B.52) and expand the velocity in a Taylor series about $\epsilon = 0$. Equating powers of ϵ , it is easy to show that we get

$$\frac{\partial \mathbf{u}_1}{\partial t} - \frac{1}{\text{Re}} \nabla^2 \mathbf{u}_1 + \nabla p_1 = 0, \quad \nabla \cdot \mathbf{u}_1 = 0, \quad (\text{B.58})$$

and the associated initial condition $\mathbf{u}_1(\mathbf{x}, 0) = 0$ and boundary conditions

$$\mathbf{u}_1(\mathbf{x}, t) = \frac{\partial \Xi_b(\mathbf{x}, t)}{\partial t}, \quad \mathbf{x} \in \bar{S}_b, \quad \mathbf{u}_1 \rightarrow 0, \quad |\mathbf{x}| \rightarrow \infty. \quad (\text{B.59})$$

At the next level, we get

$$\frac{\partial \mathbf{u}_2}{\partial t} - \frac{1}{\text{Re}} \nabla^2 \mathbf{u}_2 + \nabla p_2 = -\mathbf{u}_1 \cdot \nabla \mathbf{u}_1, \quad \nabla \cdot \mathbf{u}_2 = 0, \quad (\text{B.60})$$

with initial condition $\mathbf{u}_2(\mathbf{x}, 0) = 0$ and boundary condition

$$\mathbf{u}_2(\mathbf{x}, t) = -\Xi_b(\mathbf{x}, t) \cdot \nabla \mathbf{u}_1(\mathbf{x}, t), \quad \mathbf{x} \in \bar{S}_b, \quad \mathbf{u}_2 \rightarrow 0, \quad |\mathbf{x}| \rightarrow \infty. \quad (\text{B.61})$$

Equations (B.58)–(B.61) describe the dominant fluid behavior in a viscous streaming problem for a given unit surface motion, $\Xi_b(\mathbf{x}, t)$. It is important to note that the equations for \mathbf{u}_1 and \mathbf{u}_2 are both unsteady Stokes equations and linear. The non-linear effects enter the second-order equation, as a forcing term involving only the first-order flow field. The boundary conditions are applied at the mean location of the cylinder surface. Because of this, the second-order boundary condition contains a correction to account for the application of the first-order boundary condition at this mean location rather than its actual location. In fact, as will be revealed in equations (B.67) and (B.73) below, the boundary conditions on \mathbf{u}_1 and \mathbf{u}_2 ensure that fluid particles initially on the surface remain on the surface and that the Lagrangian mean fluid velocity $\bar{\mathbf{u}}^L$ will remain zero on the mean surface to $O(\epsilon^3)$.

There are two significant advantages gained by this asymptotic treatment of the governing equation. First, the geometry of the problem, including that of the oscillators, is fixed, which greatly improves the computational efficiency of the solution procedure; and second, the flow field generated at each asymptotic level has unit order of magnitude, which reduces the effects of numerical error on the solution.

With this expansion of the fluid velocity in small amplitude, the inertial particle velocity field, \mathbf{v} , provided by equation (B.18), can be asymptotically expanded in the same manner, though with the inclusion of an intermediate term at $\epsilon^{3/2}$ due to the Saffman lift, $\mathbf{v} = \epsilon\mathbf{v}_1 + \epsilon^{3/2}\mathbf{v}_{3/2} + \epsilon^2\mathbf{v}_2$. It is straightforward to show that the leading contributions are given by

$$\mathbf{v}_1 = \mathbf{u}_1 + \tau\mathbf{a}_1, \quad (\text{B.62})$$

and

$$\mathbf{v}_{3/2} = -\tau^{3/2}\beta^{1/2}\frac{3\sqrt{3}J_\infty}{2\pi^2\sqrt{|\boldsymbol{\omega}_1|}}\mathbf{a}_1 \times \boldsymbol{\omega}_1, \quad \mathbf{v}_2 = \mathbf{u}_2 + \tau\mathbf{a}_2, \quad (\text{B.63})$$

where $\boldsymbol{\omega}_1 = \nabla \times \mathbf{u}_1$ and where the leading-order acceleration forces \mathbf{a}_1 and \mathbf{a}_2 are given by

$$\mathbf{a}_1 = (\beta - 1)\frac{\partial\mathbf{u}_1}{\partial t} + \frac{1}{2}\frac{\beta}{\text{Re}}\nabla^2\mathbf{u}_1 \quad (\text{B.64})$$

and

$$\mathbf{a}_2 = (\beta - 1) \left(\frac{\partial \mathbf{u}_2}{\partial t} + \mathbf{u}_1 \cdot \nabla \mathbf{u}_1 \right) + \frac{1}{2} \frac{\beta}{\text{Re}} \nabla^2 \mathbf{u}_2. \quad (\text{B.65})$$

In order to unify our later discussions on averaging the trajectories of fluid and inertial particles under the generic velocity \mathbf{w} , we take some liberty in asymptotic notation by lumping $\epsilon^{3/2} \mathbf{v}_{3/2} + \epsilon^2 \mathbf{v}_2$ into a single second-order term, $\epsilon^2 \mathbf{v}_2$, where

$$\mathbf{v}_2 = -\epsilon^{-1/2} \tau^{3/2} \beta^{1/2} \frac{3\sqrt{3} J_\infty}{2\pi^2 \sqrt{|\boldsymbol{\omega}_1|}} \mathbf{a}_1 \times \boldsymbol{\omega}_1 + \mathbf{u}_2 + \tau \mathbf{a}_2. \quad (\text{B.66})$$

Though there is some awkwardness in this notation with a negative power of ϵ , there should be no ambiguity, and the $O(\epsilon^{3/2})$ term is easily recovered in every expression that follows.

With the advantages presented by the mean configuration space, it is worth wondering whether we might formulate and solve governing equations directly for \mathbf{u}^ξ in this space to provide a more direct path to the Lagrangian mean velocity field. Such equations have been derived, for example, by Andrews and McIntyre [94, 104]. These equations introduce new quantities, such as the pseudo-momentum density field, $-\mathbf{u}^l \cdot \nabla \boldsymbol{\xi}$, that couple the disturbed displacement field into the equations. We have chosen instead to solve for \mathbf{u} in the actual configuration space and then follow the procedure described in Section B.3.2.2 to relate this velocity (or the inertial particle velocity) to its corresponding Lagrangian mean field. Ultimately, after asymptotic expansions in ϵ have been used to reduce the equations, one can show that both procedures reduce to the same result.

B.3.2.2 Reduction of the Lagrangian mean velocity field

Now let us apply our asymptotic expansions (B.57) to the particle transport equations. We start with the equation for \mathbf{w}^ξ in (B.45) and carry out a Taylor expansion about $\epsilon = 0$. We get

$$\mathbf{w}^\xi(\mathbf{x}, t) = \epsilon \mathbf{w}_1(\mathbf{x}, t) + \epsilon^2 (\mathbf{w}_2(\mathbf{x}, t) + \boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1(\mathbf{x}, t)) + O(\epsilon^3). \quad (\text{B.67})$$

It should be noted that this expanded form of \mathbf{w}^ξ relates the Eulerian forms of the actual velocity field in our two spaces: between the value at fixed \mathbf{x} in the mean configuration space

and its value at the mean location to which \mathbf{x} is mapped in the actual configuration space. We can then substitute this expanded form of \mathbf{w}^ξ into the definition of $\overline{\mathbf{w}}^L$ in (B.45) and easily get an expanded form of this mean velocity:

$$\overline{\mathbf{w}}^L = \epsilon \overline{\mathbf{w}}_1 + \epsilon^2 (\overline{\mathbf{w}}_2 + \overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1}) + O(\epsilon^3). \quad (\text{B.68})$$

This shows that, at leading order, the Lagrangian mean velocity at some location \mathbf{x} in \mathcal{Y}_t is equal to the Eulerian mean of the leading-order velocity at the same location in \mathcal{X}_t . At the next order, however, an additional term appears: the Eulerian mean velocity is modified by the Stokes drift velocity, $\overline{\mathbf{w}}_d \equiv \overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1}$. When the fluid velocity field is purely oscillatory (i.e., without transient behavior), it can be easily verified that the leading velocity has zero Eulerian mean, $\overline{\mathbf{w}}_1 = 0$, for both types of particles. Thus, the Stokes drift has an essential role in determining the mean trajectories of particles. For inertial particles, where \mathbf{w} is taken to be \mathbf{v} , the Saffman lift in (B.18) generates non-zero Eulerian mean at smaller order ($\epsilon^{3/2}$) than for a fluid, but the Stokes drift still cannot be neglected in such a case.

Now let us complete the asymptotic analysis by substituting the expansion of $\boldsymbol{\xi}$ and both expanded forms of the velocities into the equation for $\boldsymbol{\xi}$ in (B.45) and equating like powers of ϵ . At the leading two asymptotic levels, we get

$$\frac{\partial \boldsymbol{\xi}_1}{\partial t} = \mathbf{w}_1 - \overline{\mathbf{w}}_1, \quad (\text{B.69})$$

$$\frac{\partial \boldsymbol{\xi}_2}{\partial t} = \mathbf{w}_2 - \overline{\mathbf{w}}_2 - \overline{\mathbf{w}}_1 \cdot \nabla \boldsymbol{\xi}_1 + \boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1 - \overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1}. \quad (\text{B.70})$$

First, let us note that at each asymptotic level these equations preserve the zero mean of $\boldsymbol{\xi}$, thereby ensuring that we remain within the constraints of GLM theory. Second, we observe that we can obtain a completely self-consistent algorithm for generating the Lagrangian mean velocity field by retaining only the equation for $\boldsymbol{\xi}_1$. In other words, for a given pair of

asymptotic velocity fields \mathbf{w}_1 and \mathbf{w}_2 , we can solve

$$\frac{\partial \boldsymbol{\xi}_1}{\partial t} = \mathbf{w}_1 - \overline{\mathbf{w}}_1, \quad (\text{B.71})$$

$$\overline{\mathbf{w}}_d \equiv \overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1}, \quad (\text{B.72})$$

$$\overline{\mathbf{w}}^L(\mathbf{x}, t) = \epsilon \overline{\mathbf{w}}_1(\mathbf{x}, t) + \epsilon^2 (\overline{\mathbf{w}}_2(\mathbf{x}, t) + \overline{\mathbf{w}}_d(\mathbf{x}, t)), \quad (\text{B.73})$$

to generate a Lagrangian mean field, $\overline{\mathbf{w}}^L$, valid to $O(\epsilon^{5/2})$. This reduced form only requires the leading Eulerian velocity fluctuation, $\mathbf{w}_1 - \overline{\mathbf{w}}_1$ to obtain the required disturbed displacement field, and thence, the Stokes drift's important contribution, $\overline{\mathbf{w}}_d$, to the Lagrangian mean velocity field for particle transport.

It is worth making a few other notes on the Lagrangian velocity before we close this section. To support the first two notes, let us first develop an alternative form of the Stokes drift (B.72). If we substitute (B.71) for \mathbf{w}_1 and remember that $\overline{\boldsymbol{\xi}}_1 = 0$, then this Stokes drift can be written as $\overline{\boldsymbol{\xi}_1 \cdot \nabla \partial \boldsymbol{\xi}_{1,t}}$, where $\boldsymbol{\xi}_{1,t}$ denotes $\partial \boldsymbol{\xi}_1 / \partial t$. Using our definition of the Eulerian mean in (B.22), we can integrate this form of the Stokes drift by parts:

$$\overline{\mathbf{w}}_d = \frac{1}{T} \int_{t-T}^t \boldsymbol{\xi}_1 \cdot \nabla \boldsymbol{\xi}_{1,t'} dt' = \frac{1}{T} [\boldsymbol{\xi}_1 \cdot \nabla \boldsymbol{\xi}_1]_{t-T}^t - \frac{1}{T} \int_{t-T}^t \boldsymbol{\xi}_{1,t'} \cdot \nabla \boldsymbol{\xi}_1 dt' \quad (\text{B.74})$$

with T the averaging interval. The first term on the right-hand side of this equation is identically zero when the field $\boldsymbol{\xi}_1$ is periodic and T is an integer multiple of the period. For transient problems, in which T is taken to be much larger than the oscillatory time-scale, the term does not strictly vanish. However, it should be noted that, had we formally defined separate fast and slow measures of time, then T would be much longer than this fast time scale (while the slow time is effectively held fixed). The term would vanish in this limit. For this reason, we argue that it can be neglected in general cases without consequence. Thus, the Stokes drift can also be written as $\overline{\mathbf{w}}_d = -\overline{\boldsymbol{\xi}_{1,t} \cdot \nabla \boldsymbol{\xi}_1}$, or, alternatively, as a combination of the two forms, with the velocity substituted back in,

$$\overline{\mathbf{w}}_d = \frac{1}{2} (\overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1 - \mathbf{w}_1 \cdot \nabla \boldsymbol{\xi}_1}). \quad (\text{B.75})$$

This latter form of the Stokes drift has several merits, as we discuss below.

Mean trajectories in an incompressible velocity field. In the case when the velocity field \mathbf{w} is incompressible, $\nabla \cdot \mathbf{w} = 0$, and the Lagrangian mean of this field is steady, $\partial \overline{\mathbf{w}}^L / \partial t = 0$, then we can obtain mean particle trajectories directly from the contours of a Lagrangian streamfunction field associated with $\overline{\mathbf{w}}^L$ [109], as we will show here. When \mathbf{w} is divergence-free, then \mathbf{w}_1 and \mathbf{w}_2 and their means are, as well. Each of these means can thus be derived from a streamfunction field, which we will denote by $\overline{\psi}_1$ and $\overline{\psi}_2$, respectively (where, e.g., $\overline{\mathbf{w}}_1 = \nabla \times \overline{\psi}_1$). It is less obvious that the Stokes drift term, $\overline{\mathbf{w}}_d = \overline{\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1}$ can be derived from a streamfunction, as well. However, when the velocity field \mathbf{w}_1 is divergence-free, then by (B.71), so are $\boldsymbol{\xi}_1$ and its time derivative, and the right-hand side of (B.75) can thus be written as $\overline{\nabla \cdot (\boldsymbol{\xi}_1 \mathbf{w}_1) / 2 - \nabla \cdot (\mathbf{w}_1 \boldsymbol{\xi}_1) / 2}$. After applying a standard vector identity on this latter form, the Lagrangian mean of an incompressible periodic velocity field can be written as

$$\overline{\mathbf{w}}^L = \nabla \times \overline{\psi}^L, \quad (\text{B.76})$$

where the Lagrangian streamfunction, $\overline{\psi}^L$, is defined as

$$\overline{\psi}^L = \epsilon \overline{\psi}_1 + \epsilon^2 (\overline{\psi}_2 + \overline{\psi}_d), \quad (\text{B.77})$$

and we have defined a Stokes drift streamfunction, $\overline{\psi}_d$, as

$$\overline{\psi}_d = \frac{1}{2} \overline{\mathbf{w}_1 \times \boldsymbol{\xi}_1}. \quad (\text{B.78})$$

Numerical computation of the Lagrangian mean velocity field. In our present viscous streaming context, the assumptions necessary to derive the trajectories from a Lagrangian streamfunction are only fulfilled by the fluid velocity field. The inertial particle velocity field is not divergence-free due to the Saffman lift, and we must compute trajectories by numerically integrating the mean transport equation (B.44). For most problems, this numerical integration must make use of a $\overline{\mathbf{w}}^L$ field numerically computed from grid velocity data. This grid-based approximation inevitably introduces error, and the specific form of this error can be highly deleterious. The error is most clearly revealed in the cases in which

we have the Lagrangian streamfunction available to verify our result—that is, cases in which $\overline{\boldsymbol{w}}^L$ is divergence-free. Most contours of the Lagrangian streamfunction form closed loops, whereas numerically computed trajectories generally fail to close unless we adopt a numerical approximation that has certain key properties. In particular, if, for two vector fields \boldsymbol{a} and \boldsymbol{b} , the underlying approximation satisfies the product rule $\boldsymbol{a} \cdot \nabla \boldsymbol{b} = \nabla \cdot (\boldsymbol{a}\boldsymbol{b}) - (\nabla \cdot \boldsymbol{a})\boldsymbol{b}$ in a discrete sense, and if the discrete divergence vanishes when the continuous divergence does, then the form of the Stokes drift given by equation (B.75) greatly mitigates the mismatch between the numerically-computed particle trajectories and the Lagrangian streamlines. Even when the mean particle trajectories cannot be otherwise obtained from a Lagrangian streamfunction, form (B.75) retains many of its benefits for reducing the accumulated error in the trajectories.

Reconciliation with surface motion. Finally, it is useful to reconcile equations (B.71)–(B.73) with the expected behavior of these quantities on the surface \overline{S}_b at this order of approximation. Most obviously, the first-order disturbed displacement field $\boldsymbol{\xi}_1$ on this surface is described by the fluctuating velocity of the moving surface $S_b(t)$ evaluated at its mean location in the actual space. Furthermore, by (B.51), we restrict the Lagrangian mean velocity to be exactly zero on \overline{S}_b . Equation (B.73) shows that this restriction requires that $\overline{\boldsymbol{w}}_1 = 0$ at the mean location of the surface in actual space; at second order, it is further required that

$$\overline{\boldsymbol{w}}_2 = -\overline{\boldsymbol{w}}_d, \tag{B.79}$$

which reflects that the Eulerian mean velocity at this mean location is not exactly zero, but must vary with the displacement of the surface from this location.

B.3.3 Algorithm for Lagrangian averaged transport of particles

With the governing equations now developed and reduced for $O(\epsilon)$ oscillations, we can now summarize the proposed algorithm for computing the fast Lagrangian averaged transport of

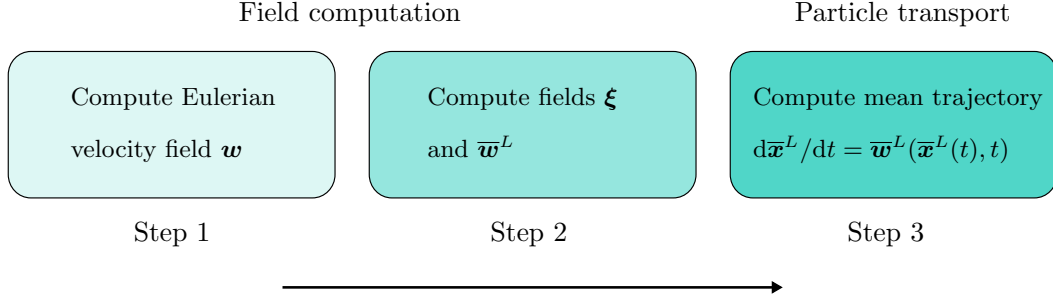


Figure B.3: Overall algorithm for the fast Lagrangian averaged transport of particles.

fluid or inertial particles in viscous streaming flows. The algorithm, illustrated in Figure B.3, involves three steps. The first step consists of computing the underlying velocity field, \mathbf{w} , for a given surface motion $\epsilon\boldsymbol{\Xi}_b(\mathbf{x}, t)$. If our interest is in fluid particle trajectories, \mathbf{w} represents the fluid velocity field \mathbf{u} . This velocity is assembled from $\mathbf{u} = \epsilon\mathbf{u}_1 + \epsilon^2\mathbf{u}_2$, using the solutions of the sequence of unsteady Stokes equations (B.58)–(B.61). In this work, these equations are solved numerically with the Immersed Boundary Projection Method [17, 30]. For inertial particles, \mathbf{w} corresponds to the particle velocity field \mathbf{v} , induced by the fluid velocity \mathbf{u} via the expansion in small Stokes number (B.18). This velocity field’s own asymptotic expansion in ϵ , where needed, is provided through $O(\epsilon^2)$ by equations (B.62) and (B.66).

In the second step of the algorithm, the leading-order disturbed displacement field, $\boldsymbol{\xi}_1$, is computed by integrating equation (B.71), rewritten here for reference:

$$\frac{\partial \boldsymbol{\xi}_1}{\partial t} = \mathbf{w}_1 - \overline{\mathbf{w}}_1. \quad (\text{B.80})$$

In the case of fluid particles, we simply integrate this equation simultaneously with the Navier–Stokes equations in the first step. For inertial particles, we integrate with a third-order Runge-Kutta method, using cubic splines to interpolate the underlying time-discretized velocity field. From this field, the Lagrangian mean velocity field $\overline{\mathbf{w}}^L$ is computed from

equation (B.73) with the Stokes drift calculated with equation (B.75), i.e.,

$$\overline{\mathbf{w}}^L = \epsilon \overline{\mathbf{w}}_1 + \epsilon^2 \left(\overline{\mathbf{w}}_2 + \frac{1}{2} \overline{(\boldsymbol{\xi}_1 \cdot \nabla \mathbf{w}_1 - \mathbf{w}_1 \cdot \nabla \boldsymbol{\xi}_1)} \right). \quad (\text{B.81})$$

Finally, in the third step, the Lagrangian mean trajectory of each particle is computed by integrating the mean transport equation (B.44) from some initial location \mathbf{x}_0 , with a fifth-order Adams–Bashforth method. We generally use a time step that is 10 times larger than the period of oscillation. As we observed in Section B.3.1.2, the inertial particle velocity field (B.18) does not inherently satisfy the no-slip or no-flow-through conditions, even if the fluid velocity field does. This is also true of the solution from the full Maxey–Riley equation, and Chong et al. [2] handled the issue by adding a penalty force inspired from lubrication theory to prevent penetration of inertial particles through the surfaces of oscillators. Here, we use an alternative approach wherein we augment the transport equation with a constraint that the particle remain in the region external to the oscillators. The constraint is posed as

$$H_\delta(d(\mathbf{x})) = 0, \quad (\text{B.82})$$

where $d(\mathbf{x})$ is a signed distance field with respect to the oscillator boundaries—positive in the interior of the oscillator and negative in the exterior—and H_δ is a smooth Heaviside function used previously by Li et al. [110], defined as

$$H_\delta(z) = \begin{cases} 0 & z < -\delta, \\ \frac{1}{2} (1 + z/\delta + \pi^{-1} \sin(\pi z/\delta)) & |z| \leq \delta, \\ 1 & z > \delta, \end{cases} \quad (\text{B.83})$$

where δ is a smoothing distance set equal to the particle radius a for all our simulations. The constrained system of ordinary differential equations is then solved with the manifold projection method described by Hairer et al. [111].

For both types of particles, all Eulerian means in the algorithm are computed with the time average defined in equation (B.22). The averaging interval T is taken to be 10 periods of oscillation. This ensures that, during transient phases, the interval is long compared

with the fast time scale but short compared with the trajectory. When the flow reaches a stationary periodic state, the averaging interval can be reduced without consequence to a single period of oscillation.

B.4 Results

In this section, we present the results from the application of the particle transport algorithm to two representative viscous streaming flows. First, in Section B.4.1, we verify that our asymptotic expansion (B.18) of the inertial particle velocity field in small Stokes number produces time-resolved trajectories that are accurate when compared with the solution of these trajectories from the full Maxey–Riley equation (B.8). Then, in Sections B.4.2 and B.4.3 we investigate the performance of the mean transport algorithm when applied to fluid and inertial particles, respectively, compared with the fully time-resolved integration of these trajectories, and, in the case of fluid particles, with the Lagrangian streamlines.

The two viscous streaming flows we consider in this work each consist of a flow generated by a rigid cylinder of radius R in weak oscillatory translation with angular frequency Ω . An isolated cylinder in such motion creates four streaming cells arranged along 45-degree rays [112]; in arrays of multiple cylinders, the cells are still present, though somewhat deformed by the presence of other cylinders. Inertial particles tend to become trapped in these streaming cells, as evident from previous work [82, 81, 2]. Our focus in this paper is primarily to confirm the various aspects of the proposed transport algorithm. The first case consists of a single cylinder, while the second case consists of two cylinders that oscillate in sequence: one cylinder oscillates while the other remains stationary, then, after a certain interval, they exchange their roles. In this second case, we are particularly interested in the second cylinder’s ability to draw an inertial particle originally trapped near the first cylinder towards one of its own streaming cells.

For both cases, the surface displacement ξ_b of any oscillator is described entirely by the




Figure B.4: Diagram of the oscillating cylinder (right), and a time sequence illustrating the repeated oscillation cycles.




Figure B.5: Configuration and oscillation sequence with two cylinders adapted from the work of Chong et al. [2].

motion of the cylinder's centroid, as expressed in equation (B.55). In our investigations of this section, the centroid motion is purely sinusoidal, $\mathbf{x}_c(t) = \mathbf{x}_{c0} + \epsilon \sin t \mathbf{e}_x$, where $\epsilon = A/R$. (We continue to non-dimensionalize all quantities in this section with Ω and R , as discussed just before equation (B.14).) In terms of the unit form of this surface displacement, $\boldsymbol{\xi}_b = \epsilon \boldsymbol{\Xi}_b$, we can express the motion as

$$\boldsymbol{\Xi}_b(\mathbf{x}, t) = \sin t \mathbf{e}_x. \quad (\text{B.84})$$

Throughout, the Reynolds number, Re , is held fixed at 40, and the amplitude, ϵ , for any oscillator is 0.1. The Stokes number of the inertial particles is set at $\tau = 0.1$ and the particle density ratio at $\rho_p/\rho_f = 0.95$, which correspond to $\beta = 1.034$ and a particle of radius $a/R = 0.088$.

A schematic of the unit surface motion is depicted in Figures B.4 and B.5. The blue dot represents the position of the centroid of the cylinder, and the blue dashed line depicts the unit surface displacement applied at points on the fixed cylinder surface. The right diagram in each figure depicts the manner in which we generate flow fields over long time horizons. After each change in the oscillator motion, the flow does not become statistically stationary until viscous diffusion has had sufficient time to act. Once this transient phase has ended and the flow's mean has become stationary, the solution over the last oscillation cycle is re-used as many times as necessary to generate the flow field's history. In the conditions specified above, we find that the flow becomes statistically stationary in the region within 6 radii of the oscillator after around $n_{s1} = 20$ periods of oscillation. We have two such transient phases in the case of two oscillators. As Figure B.5 shows, these oscillators are arranged 6 radii apart along the same axis on which they oscillate, and the periodic flow solution developed by each oscillator is repeatedly recycled as needed in the respective intervening periods. For particle trapping purposes in this conditions, we find that these recycling intervals require $n_{f1} - n_{s1} = 25000$ periods and $n_{f2} - n_{s2} = 40000$ periods, respectively.

Throughout this investigation we will rely on fluid velocity fields obtained by numerical solution of the unsteady Stokes equations with a procedure based on the immersed boundary projection method with lattice Green's function [17, 30]. The validation of this procedure, including its convergence to the analytical solution in the case of a single cylinder in oscillatory translation, has been confirmed but is omitted from this work for brevity. We note that the simulations reported here are carried out on a Cartesian grid with spacing $\Delta x/R = 0.02$ and a time-step size $\Omega\Delta t/(2\pi) = 0.004$, or 250 time steps per period, to satisfy the viscous stability constraint. The computational domain in both cases, $[-6R, 6R] \times [-6R, 6R]$, is relatively more compact than required by other numerical methods due to the use of the lattice Green's function and associated viscous integrating factor [30].

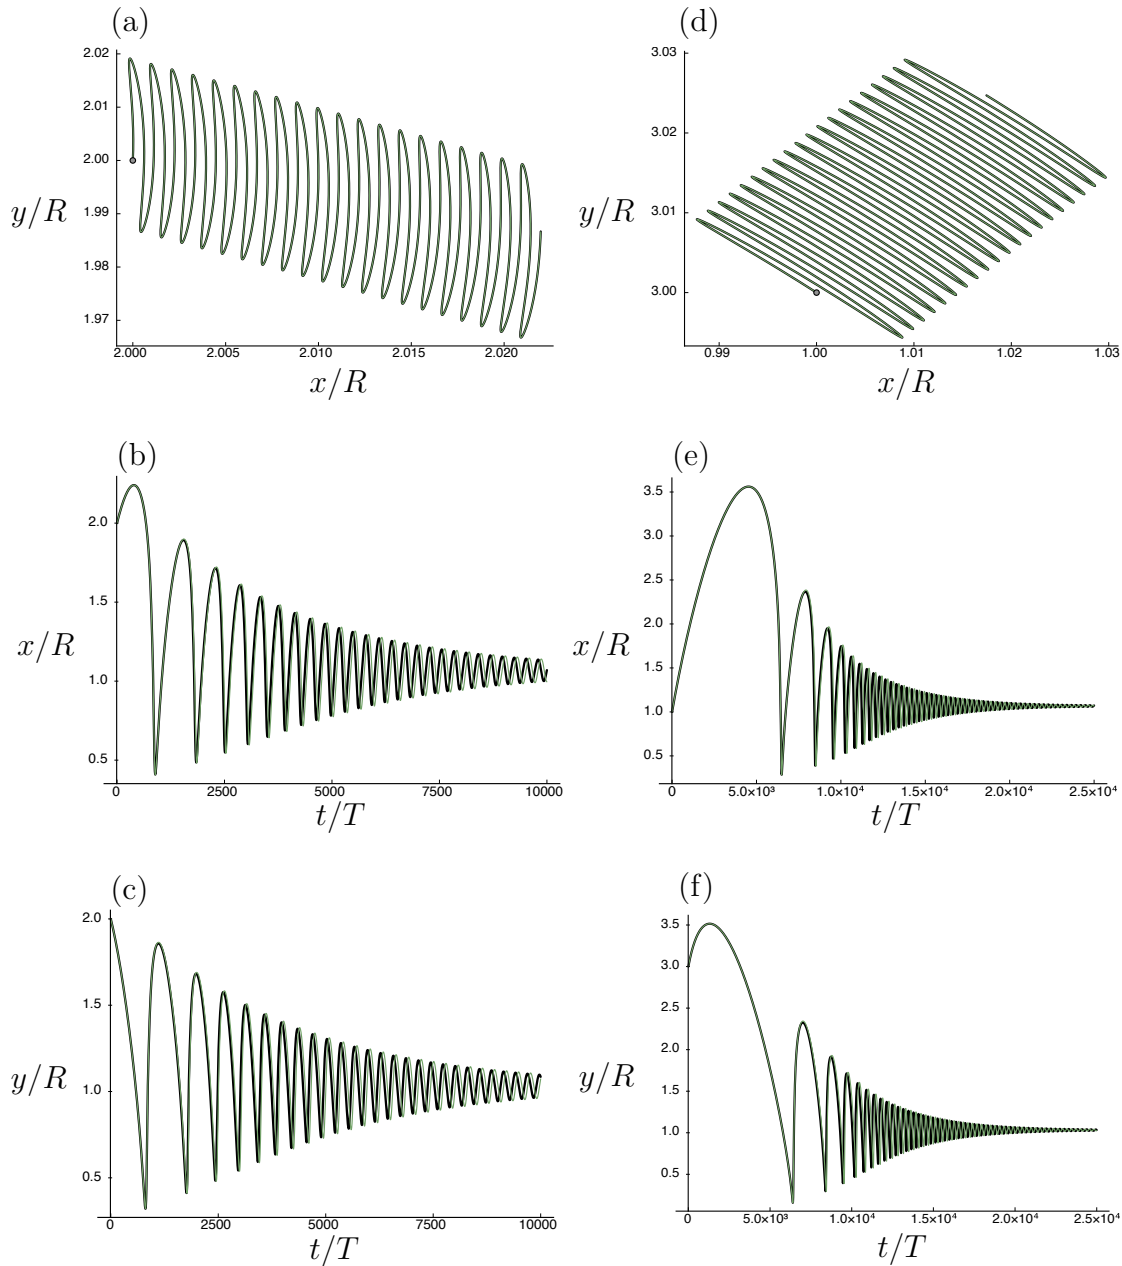


Figure B.6: Time-resolved inertial particle trajectories from the Maxey–Riley equation (black) and the asymptotic expansion in small Stokes number (green) for particles initially located at $\mathbf{x}_0 = (2, 2)$ (left column) and $\mathbf{x}_0 = (1, 3)$ (right column). The top row (a,d) shows the trajectories over the first 20 oscillation periods, and the middle (b,e) and lower (c,f) rows depict the time histories of the x and y components, respectively, sampled once per period.

B.4.1 Small Stokes expansion of inertial particle velocity field

In this section, we assess the accuracy of the asymptotic expansion of the inertial particle velocity field in small Stokes number, developed in equation (B.18). For evaluation purposes, we compare the trajectories of inertial particles transported by this velocity field, with the trajectories of the same particles with velocity obtained from the full Maxey–Riley equation (B.8). It should be noted that we are not yet assessing the mean transport algorithm in this section, so our comparison is made of the full time-resolved trajectory, computed from (B.13) for both velocities. As discussed earlier in the paper, the Basset term is neglected in both forms of velocity. For the numerical integration of these trajectories, we use a fifth order Adams–Bashforth method with time step $\Omega\Delta t/(2\pi) = 0.004$.

For the case of a single oscillating cylinder, we simulate two trajectories: one for a particle starting from $\mathbf{x}_0 = (2, 2)$ and tracked for 10000 periods of oscillation, and another for a particle released from $\mathbf{x}_0 = (1, 3)$ and tracked for 25000 periods. Both particles are released after the flow has reached its stationary periodic state. The comparisons of these trajectories are shown in Figure B.6. The small portions of the full trajectories shown in the top row exhibit the characteristic fluctuations of these trajectories about a mean. To reveal this mean behavior more clearly, we sample these trajectories only once per cycle in the middle and lower rows, with the history of each component depicted in a separate plot. These plots exhibit the trapping behavior: both particles converge toward a fixed point inside the streaming cell along the 45-degree ray.

The plots in Figure B.6 show that the asymptotic expansion in small Stokes number has very accurately preserved the behavior of the Maxey–Riley equation. Though small errors accumulate over time, the trajectories apparently agree well even after 25000 periods. Table B.1 reports a quantitative measure of this comparison, with error defined as the difference of the asymptotically-approximated trajectory components from those of the Maxey–Riley trajectory at the same instant, normalized by the current radial distance from. The error

\mathbf{x}_0	error after 1 period(%)	t_f/T	error at t_f (%)
(2, 2)	$(-1.19 \times 10^{-5}, 8.77 \times 10^{-5})$	10000	(-6.91, -1.11)
(1, 3)	$(1.29 \times 10^{-4}, 2.04 \times 10^{-5})$	25000	(-0.693, -0.736)

Table B.1: Relative errors on particle position, after 1 period and at the final time t_f , in two different inertial particles trajectories predicted by the small Stokes number expansion.

remains small throughout, and the final trapping location is predicted with less than one percent error.

Now, let us validate our small Stokes number expansion on the transport of inertial particles in the two-cylinder case. Here, we release particles just after the initiation of motion of the left cylinder. This case is potentially more challenging due to the transient behavior of the flow after each oscillator’s motion is initiated. The results in Figure B.7, which depicts the full trajectory sampled once per period for a particle released from $(-2, 3)$, show that the particle is first trapped near the center of a streaming cell near the left cylinder at $(x, y) = (-1.98, 1.03)$; and after the right cylinder starts its own motion, the particle is eventually drawn to a new trapping location at $(x, y) = (1.98, 1.03)$. It should be noted that this problem requires the no-penetration constraint described in equation (B.82) when the inertial particle approaches the right oscillator. As observed in Figure B.7, the particle is drawn toward this oscillator along the axis of symmetry. Without the explicit enforcement of this constraint the particle would spuriously pass across the oscillator surface. Instead, the particle remains offset from the oscillator by a small distance set by the smoothing parameter δ in this constraint and is quickly drawn into an orbit that spirals toward the trapping point. Throughout this sequence, the asymptotically-approximated trajectory agrees well with the Maxey–Riley trajectory, with error less than 0.01 percent after the first transient phase. The final percentage error in the trapping location, after 65000 cycles, is $(1.76 \times 10^{-2}, 1.19 \times 10^{-2})$.




Figure B.7: Trajectory of an inertial particle initially located at $\mathbf{x}_0 = (-2, 3)$ over 65000 periods of oscillation. Trajectories from the Maxey–Riley equation (black) and small Stokes number expansion (green) are both sampled once per period.

B.4.2 Mean fluid particle trajectories

In this section, we verify our algorithm for computing mean particle trajectories by applying it to fluid particles. In the first case, we compute trajectories from the fluid velocity field generated by the single oscillating cylinder after it has achieved periodic behavior. As we discussed in Section B.3.2.2, the mean trajectories of fluid particles are equivalently derived from the contours of a Lagrangian streamfunction field, $\overline{\psi}^L$, in equation (B.77). This alternative approach provides a natural target for verification. Examples of this comparison are shown in Figure B.8 for two different particles, both of whose mean trajectories have been integrated with a time step size of 10 periods. The agreement is very good, and importantly, the trajectories generated by the algorithm are closed after each orbit to within small numerical error.

In Figure B.9 we compare the mean trajectory of the particle $\mathbf{x}_0 = (2, 2)$ with the full time-resolved trajectory. This latter trajectory is obtained from the same (numerically-




Figure B.8: Lagrangian mean trajectories (magenta circles, spaced by 30 periods) generated for fluid particles started at $\mathbf{x}_0 = (2, 2)$ and $\mathbf{x}_0 = (1, 3)$ (denoted by the larger grey circle), compared with the Lagrangian streamlines, depicted as black lines. (Other Lagrangian streamlines are shown in light grey.) The black region shows the mean position of the cylinder, and the lighter shaded region in the vicinity of the cylinder shows the range of displacement of the cylinder over one oscillation cycle.

computed) fluid velocity field, but by integrating the unsteady fluid velocity with 250 time steps per period and cubic spline interpolation of the Cartesian grid values. The full trajectory reveals the oscillations incurred by the particle as it orbits about the streaming cell. The mean trajectory from the algorithm displays the expected behavior, passing through the first point in each cycle, as shown in the small section of trajectory in Figure B.9(b). Indeed, when the full trajectory is sampled once per cycle (starting with its initial position), as shown in Figure B.9(c,d), the mean trajectory agrees well with it even after 10000 oscillation periods, corresponding to nearly 6 orbits.

In the second case, we use the configuration of two cylinders. As a target of comparison, we release a particle from $\mathbf{x}_0 = (-1, 3)$, near the left cylinder, after statistically stationary behavior has been achieved from the right cylinder’s motion. The comparison with the full time-resolved trajectory is shown in Figure B.10 and exhibits very good agreement. The particle is initially drawn toward the right cylinder and achieves a closed orbit about the streaming cell; each orbit requires approximately 13000 periods of oscillation.

B.4.3 Mean inertial particle trajectories

The previous section demonstrated that our proposed algorithm can successfully predict the mean trajectories of fluid particles. In this section, we apply the algorithm to inertial particle transport. Each trajectory computed from this algorithm is compared with the full time-resolved trajectory obtained from the same inertial particle velocity field, \mathbf{v} , derived from the fluid velocity field after it has reached a periodic state. Thus, the differences between these trajectories is due entirely to errors in truncating the asymptotic expansion in ϵ in the construction of the Lagrangian mean velocity in (B.71)–(B.73).

The panels in Figure B.11 depict the predicted mean trajectories of inertial particles initially released from $(2, 2)$ and $(1, 3)$, respectively, in the single cylinder configuration. Both exhibit good agreement with the full trajectory, though small errors incurred during the nearest approach to the cylinder (due to the aforementioned truncation of the expansion

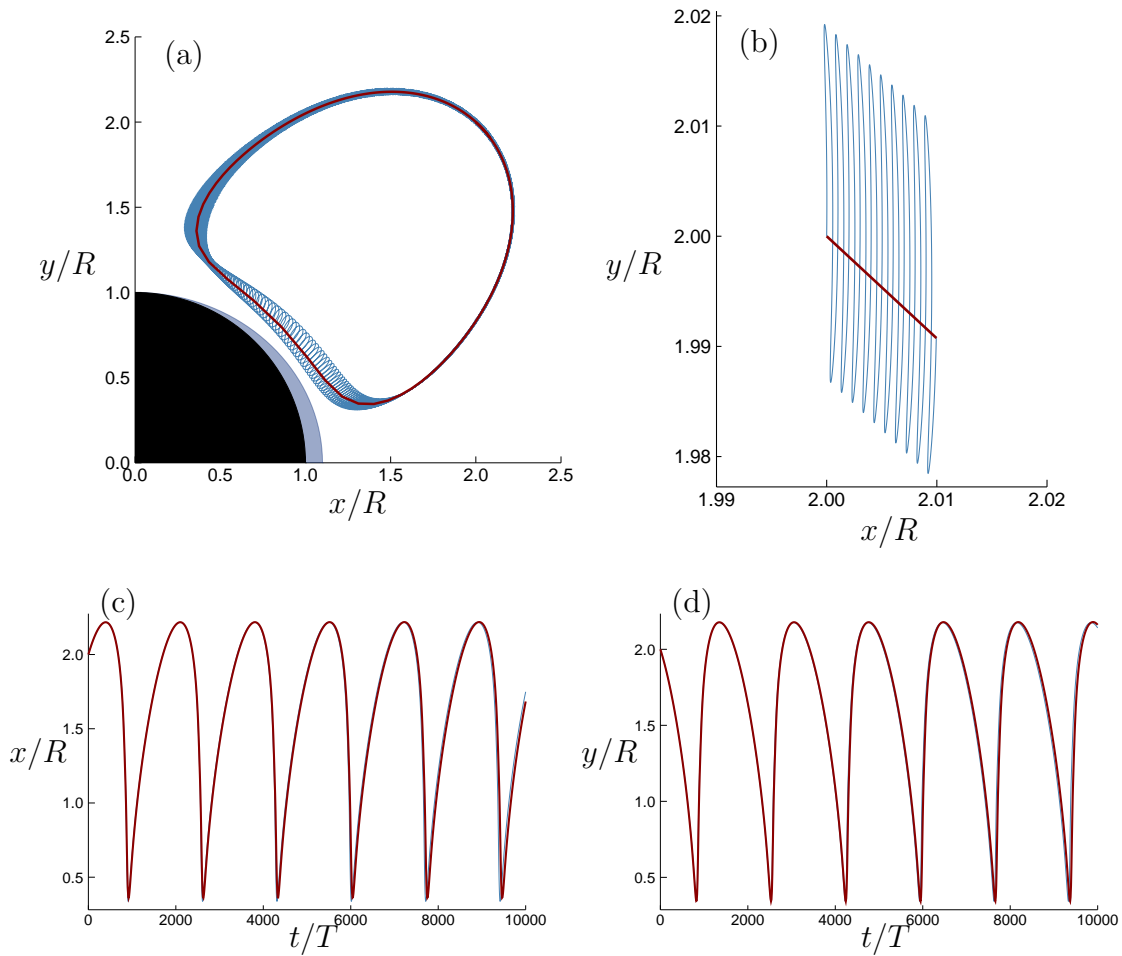


Figure B.9: (a) Lagrangian mean trajectory (magenta) for fluid particle started at $\mathbf{x}_0 = (2, 2)$, compared with the full time-resolved trajectory (blue) for the same particle. (b) Magnified view of fluid particle trajectory. (c,d) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle).

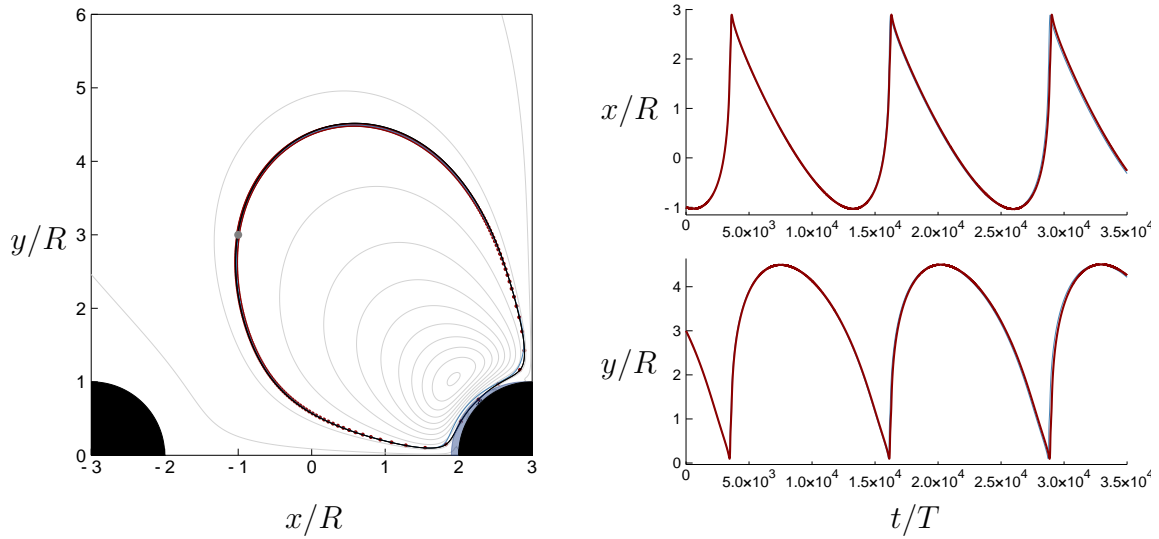


Figure B.10: Left: Time resolved fluid particle trajectory, sampled once per cycle (blue), and Lagrangian mean trajectory (magenta circles, spaced by 30 periods), and Lagrangian mean streamlines (in light gray) for $\mathbf{x}_0 = (-1, 3)$, with mean cylinder configurations depicted in black. Right: Time histories of the x and y components of both trajectories shown in the left panel.

in ϵ) tend to push the particle onto a slightly different orbit in each encirclement of the streaming cell. However, these errors are largely irrelevant, as both the mean trajectory and the full time-resolved trajectory converge on the same trapping point. This is more clearly revealed in panels Figure B.12(a,b) and (d,e), which show the respective comparisons of each coordinate’s history along the trajectory (plotted once per cycle in the case of the full trajectory). The orbits of the streaming cell predicted by the mean transport algorithm are slightly faster than those of the full trajectories due to the slightly larger push toward the center experienced by the particle in the mean algorithm as it passes closest to the cylinder. To better illustrate the relationship between these two approaches to computing the trajectories, in Figure B.12(c) and (f) we plot every point along a portion of the full time-resolved trajectory and overlay the same portion of trajectory predicted by the mean transport algorithm. The algorithm visually tracks the center of oscillations, but skews slightly inward as it moves away from the cylinder. It can be observed that, for $\mathbf{x}_0 = (1, 3)$, the oscillations along the full trajectory overlap with the right-most configuration of the cylinder; however, those portions of the trajectory that overlap correspond to the phase in the cycle when the cylinder is in its left-most configuration.

Most of our results in this paper are focused on a single type of particle, nearly neutrally buoyant with $\rho_p/\rho_f = 0.95$ (or $\beta = 1.034$). To demonstrate the effect of density ratio, we compare this particle’s mean trajectory in Figure B.13 with that of a very light particle, with density ratio $\rho_p/\rho_f = 0.05$ (so that $\beta = 2.73$). The light particle is affected much more by the buoyancy term—the first term in (B.19)—which applies a centrifugal motion directed toward the center of the streaming cell. Interestingly, this contribution is most active during the intervals when the Saffman lift is not, on the outermost parts of the orbit when the particle’s trajectory is most curved.

Figure B.14 shows the results for a particle released from $(-2, 3)$ in the two-cylinder configuration. Over the first 25000 periods, the particle is drawn toward the streaming cell of the left cylinder during that cylinder’s motion. Then, when the left cylinder stops its

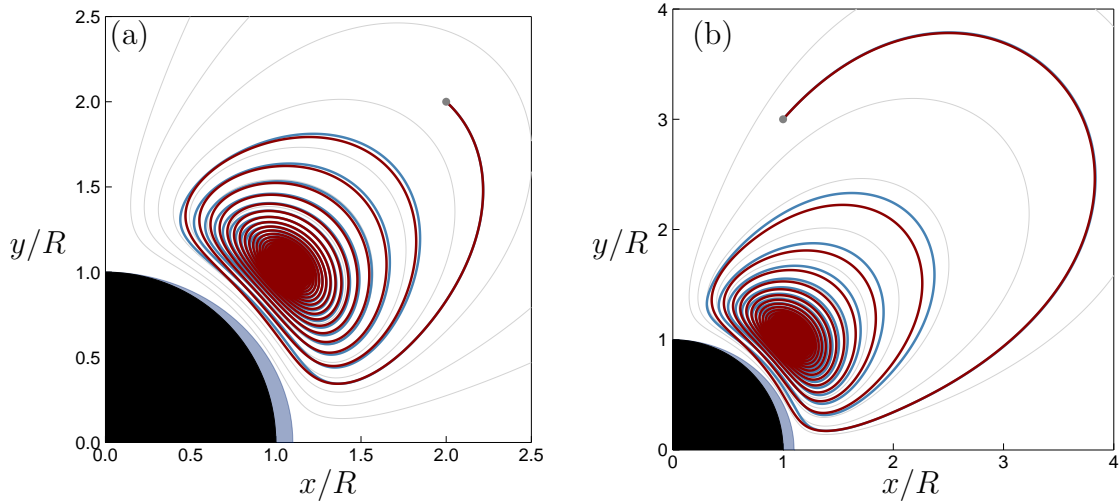


Figure B.11: Lagrangian mean trajectories (magenta) generated for inertial particles started at (a) $\mathbf{x}_0 = (2, 2)$ and (b) $\mathbf{x}_0 = (1, 3)$ (denoted by the small circle), compared with the full time-resolved trajectories (blue) of the same particles, sampled once per cycle. The Lagrangian streamlines of the fluid are shown for reference as light grey lines. The black region shows the mean position of the cylinder, and the lighter shaded region in the vicinity of the cylinder shows the range of displacement of the cylinder over one oscillation cycle.

motion and the right cylinder starts to oscillate, the particle is entrained into the cell nearest to the right cylinder over the ensuing 40000 periods. Both trajectories are predicted well by the mean transport algorithm. It should be noted that the no-penetration constraint is active for both the mean and the full trajectory predictions during the interval in which the particle reaches the right-most cylinder and is drawn along its boundary (at around 30000 periods).

B.4.4 The effect of transient behavior

An important question that overlies the transport of particles in viscous streaming is the effect of transient behavior in the fluid during changes of oscillator motion. In the examples we have profiled in this paper, these changes in motion occur suddenly: each cylinder stops

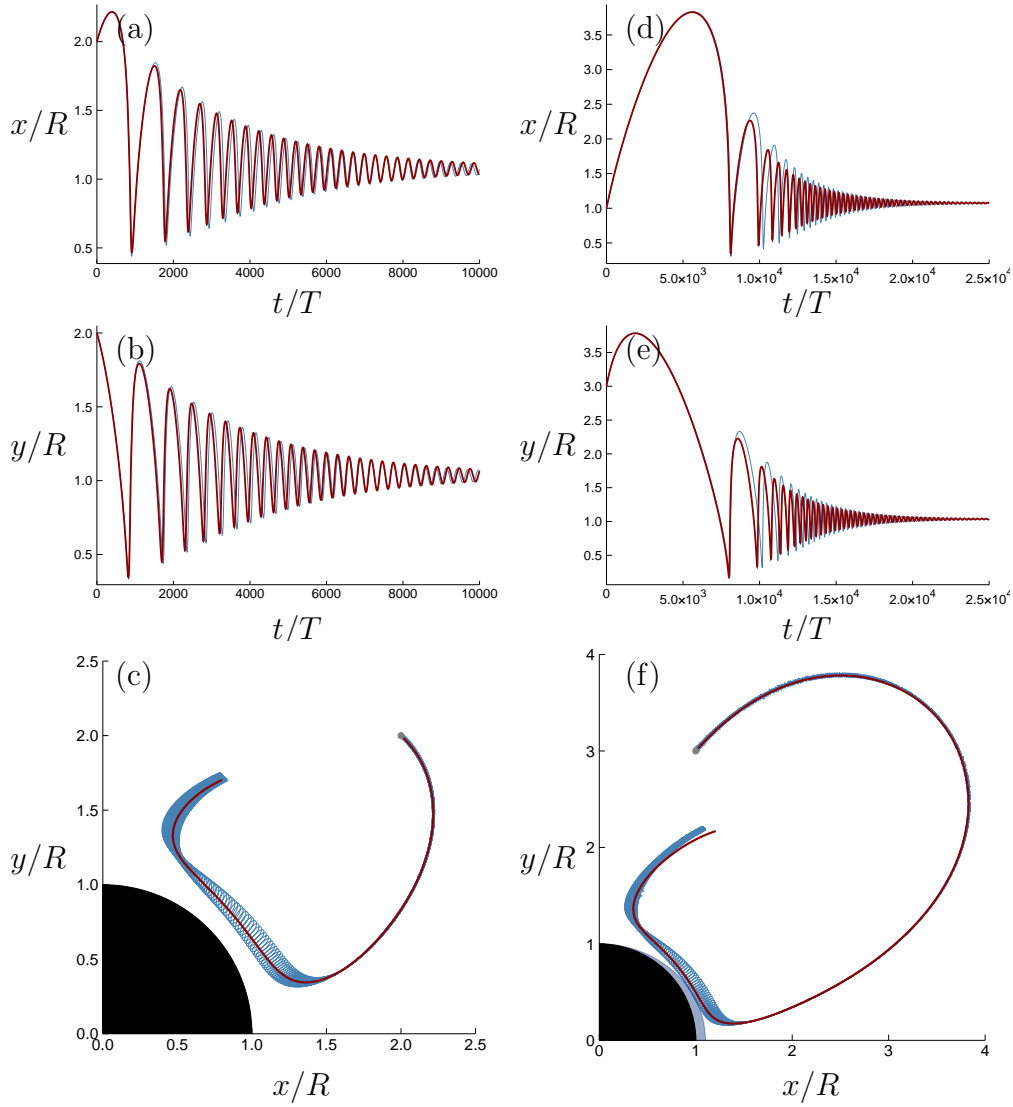


Figure B.12: Left column: (a,b) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle) for particle $\mathbf{x}_0 = (2, 2)$. (c) Lagrangian mean trajectory (magenta) for inertial particle $\mathbf{x}_0 = (2, 2)$, compared with the full time-resolved trajectory (blue) for the same particle. Right column: (d,e) Comparison of x and y components, respectively, of the Lagrangian mean trajectory and the full time-resolved trajectory (sampled once per cycle) for particle $\mathbf{x}_0 = (1, 3)$. (f) Lagrangian mean trajectory (magenta) for inertial particle started at $\mathbf{x}_0 = (1, 3)$, compared with the full time-resolved trajectory (blue) for the same particle.




Figure B.13: Lagrangian mean trajectories generated for inertial particles with $\rho_p/\rho_f = 0.95$ (magenta) and 0.05 (blue) started at $\mathbf{x}_0 = (2, 2)$.

or starts instantaneously. Such sudden changes provide an unambiguous context in which to assess the transient effects since the flow is necessarily approaching a well-defined stationary state when such transient effects will vanish. Some of the sudden change of motion is communicated everywhere instantaneously through pressure. The transient effects are due to viscous diffusion, which plays a particularly important role along directions transverse to the motion. As mentioned earlier, at $\text{Re} = 40$, this diffusion requires (empirically) around 20 oscillation cycles to spread the information about motion changes across the entire region of interest and thus establish stationary periodic behavior.

In Figure B.15 we examine the effects of transient behavior for two inertial particles over the first 14 oscillation periods. Each case depicts the inertial particle's trajectory predicted by the mean transport algorithm over one time step (in this example only, taken to be 14 periods) during a transient interval of the two-cylinder array—in the first case during the initial motion of the left cylinder, and in the second case during the newly-initiated motion of the right cylinder (after the left cylinder has stopped). The full time-resolved trajectory




Figure B.14: (a) Full time resolved inertial particle trajectory, sampled once per cycle (blue), and Lagrangian mean trajectory (magenta) for a particle released from $\mathbf{x}_0 = (-2, 3)$ in the sequential oscillator configuration. (b,c) Time histories of the x and y trajectory components, respectively.

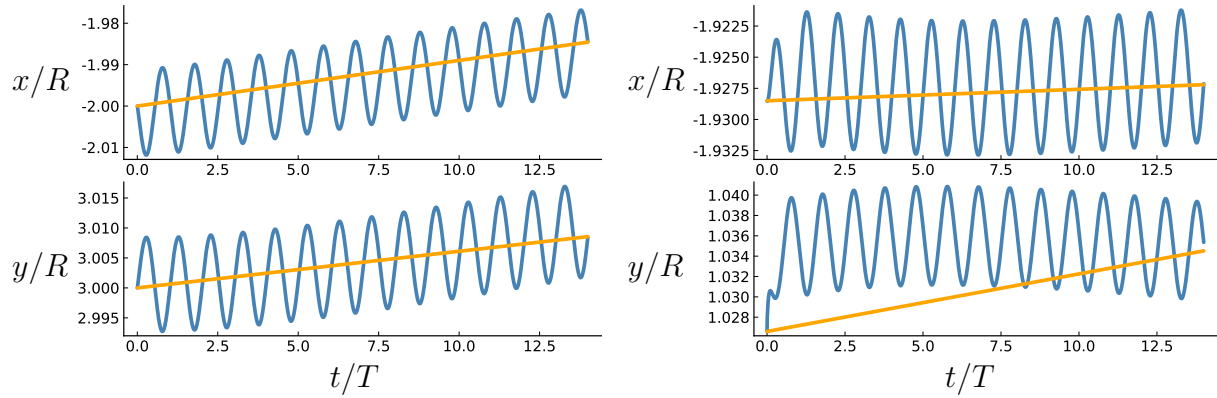


Figure B.15: Time resolved inertial particles trajectories (blue) and mean transport algorithm (yellow) for $\mathbf{x}_0 = (-2.0, 3.0)$ over the transient regime of the left cylinder (left) and $\mathbf{x}_0 = (-1.9285, 1.0266)$ over the transient regime of the right cylinder (right).

is depicted for reference in both cases. It is important to observe first that, as a result of the pressure-driven part of the flow, each particle's trajectory achieves approximately periodic behavior very quickly; the viscous adjustment takes longer. However, the mean transport algorithm predicts the behavior very well during such intervals: the particle's final location at the end of the step agrees well with where it is expected to be along the full trajectory. Furthermore, the plots show that the influence of transient behavior in the flow is likely negligible. In 14 periods, neither particle has moved more than $0.02R$ from its initial location. Even if we were to ignore the mean transport during this transient interval, our error would be equivalent to assuming that the particle had started at a negligibly different location and then been subject to a truly periodic flow.

B.5 Conclusions

In this paper, we have developed simplified equations governing the mean transport of inertial particles in viscous streaming flows. In flows generated by weakly oscillating rigid surfaces, the motions of fluid and inertial particles exhibit two distinct time scales: a fast

scale associated with the particle’s oscillatory motion and a slow scale associated with its mean translation. Previous work by Chong et al. [81] has shown that the mean motion of small inertial particles in streaming flows is well described by the Maxey–Riley equation augmented with Saffman lift and with Faxén corrections retained, but with the Basset memory term neglected. Collectively, the Faxén correction and the Saffman lift effect the trapping of inertial particles in streaming cells generated near the oscillating object.

In this work, we have analyzed such transport with the help of three key tools. First, we conceived an Eulerian field for inertial particle velocity by asymptotically expanding the Maxey–Riley equation in the small Stokes number associated with small particles in moderate Reynolds number flows. This approach follows the earlier works of Maxey [95] and Ferry and Balachandar [96], but importantly here, retains the essential Faxén term. This expansion has confirmed the observations made by Chong et al. [81]: A small neutrally-buoyant particle moves at leading order like a fluid particle, but in regions of shear near the oscillating body, the Faxén correction alters the particle’s velocity from that of the fluid and the Saffman lift then causes it to move transversely to the shear, ultimately causing it to spiral toward a trapping point in the center of a streaming cell.

The second tool has been the Generalized Lagrangian Mean theory of Andrews and McIntyre [94]. This theory’s exact distinction between the mean and fluctuating parts of a trajectory has allowed us to construct the Lagrangian mean velocity field, which is ultimately responsible for a particle’s mean transport. This mean field receives an essential contribution from Stokes drift, based on an Eulerian disturbed displacement field that accompanies the time-varying velocity field.

The third important tool has been an expansion in the small oscillation amplitude. This expansion’s effect on the fluid velocity field was already known from early work on streaming (e.g., Holtmark et al. [112]). However, with the availability now of the inertial particle velocity field and its subsequent decomposition into mean and fluctuating parts, we have been able to identify the dominant effects of small amplitude oscillation on mean inertial

particle transport. For fluid particles, the particle trajectories are directly obtained from the contours of a mean Lagrangian streamfunction field.

By applying the resulting algorithm to two basic oscillator flows, we have demonstrated that the approximations we have made by truncating these expansions have generally preserved the accuracy of the original treatment. Furthermore, the application of these tools has made a tremendous impact on the efficiency of computing mean particle trajectories. The previous approach to predicting such trajectories involved two straightforward steps: first, to compute the fluid velocity field until it reached stationary periodic behavior; and second, to advance the particle in this oscillatory flow field with the Maxey–Riley equation, with time steps that sufficiently resolve the fast scales. Each such time step requires the evaluation of the instantaneous forces on the particle, generally obtained by interpolating the velocity field and its derivatives from the computational grid. A full trajectory generally requires $O(10^6)$ such time steps and is quite slow to compute. The new approach presented here, which constructs the aforementioned Eulerian fields to develop the Lagrangian mean velocity field, allows time steps that are $O(1000)$ times larger than the previous approach.

It is also important to stress that, in spite of our examples, no aspect of our treatment of this problem is limited to two-dimensional oscillators. Indeed, the cost reduction would be proportionally greater in three-dimensional problems, where the calculations of forces on full trajectories require more taxing interrogations and calculations of the flowfield data.

We have also shown in this work that the viscous transients in the fluid that arise after changes of oscillator motion have insignificant effect on mean particle transport. This observation depends on the distinction of time scales in this regime of Reynolds number and oscillator amplitude: a particle moves very little in the time it takes for viscous diffusion to communicate the oscillator’s change of motion, so it is safe to assume that the fluid has already reached a stationary periodic state. We have not exploited this feature in this paper, but in a paper currently in development we will demonstrate that, by treating the underlying flow field as strictly periodic, we can solve for this flow field and then construct the

Lagrangian mean velocity of either type of particle entirely in the frequency domain. This leads to a further substantial gain in computational efficiency.

REFERENCES

- [1] Oliver Bühler and Michael E. McIntyre. On non-dissipative wave–mean interactions in the atmosphere or oceans. J. Fluid Mech., 354:301–343, 1998.
- [2] Kwitae Chong, Scott D. Kelly, Stuart T. Smith, and Jeff D. Eldredge. Transport of inertial particles by viscous streaming in arrays of oscillating probes. Phys. Rev. E, 93(1):013109, 2016.
- [3] Mathieu Le Provost and Jeff D Eldredge. Ensemble kalman filter for vortex models of disturbed aerodynamic flows. Physical Review Fluids, 6(5):050506, 2021.
- [4] Herbert G Küssner. Zusammenfassender bericht über den instationären auftrieb von flügeln. Luftfahrtforschung, 13(12):410–424, 1936.
- [5] William R Sears. Some aspects of non-stationary airfoil theory and its practical application. Journal of the Aeronautical Sciences, 8(3):104–108, 1941.
- [6] Jeff D. Eldredge. Mathematical Modeling of Unsteady Inviscid Flows. Springer International Publishing, 2019.
- [7] Wei Hou, Darwin Darakananda, and Jeff Eldredge. Machine learning based detection of flow disturbances using surface pressure measurements. In AIAA Scitech 2019 Forum, page 1148, 2019.
- [8] Mathieu Le Provost, Wei Hou, and Jeff Eldredge. Deep learning and data assimilation approaches to sensor reduction in estimation of disturbed separated flows. In AIAA Scitech 2020 Forum, page 0799, 2020.
- [9] J Nathan Kutz. Deep learning in fluid dynamics. Journal of Fluid Mechanics, 814:1–4, 2017.
- [10] Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. Annual Review of Fluid Mechanics, 52:477–508, 2020.
- [11] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Assessment of supervised machine learning methods for fluid flows. Theoretical and Computational Fluid Dynamics, pages 1–23, 2020.
- [12] Kody Law, Andrew Stuart, and Konstantinos Zygalakis. Discrete Time: Filtering Algorithms, pages 79–114. Springer International Publishing, Cham, 2015.
- [13] Mark Asch, Marc Bocquet, and Maëlle Nodet. Data assimilation: methods, algorithms, and applications, volume 11. SIAM, 2016.

- [14] Geir Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. Journal of Geophysical Research: Oceans, 99(C5):10143–10162, 1994.
- [15] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. Journal of Basic Engineering, 82(1):35–45, 03 1960.
- [16] Andre FC da Silva and Tim Colonius. Ensemble-based state estimator for aerodynamic flows. AIAA Journal, 56(7):2568–2578, 2018.
- [17] Kunihiko Taira and Tim Colonius. The immersed boundary method: a projection approach. J. Comput. Phys., 225(2):2118–2137, 2007.
- [18] Andre FC da Silva and Tim Colonius. Flow state estimation in the presence of discretization errors. Journal of Fluid Mechanics, 890, 2020.
- [19] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. Annual Review of Fluid Mechanics, 51:357–377, 2019.
- [20] Jaideep Pathak, Mustafa Mustafa, Karthik Kashinath, Emmanuel Motheau, Thorsten Kurth, and Marcus Day. Using machine learning to augment coarse-grid computational fluid dynamics simulations. arXiv preprint arXiv:2010.00072, 2020.
- [21] René Steijl and George N Barakos. Parallel evaluation of quantum algorithms for computational fluid dynamics. Computers & Fluids, 173:22–28, 2018.
- [22] Peter J Baddoo and Lorna J Ayton. An analytic solution for gust–cascade interaction noise including effects of realistic aerofoil geometry. Journal of Fluid Mechanics, 886, 2020.
- [23] Sébastien Michelin and Stefan G Llewellyn Smith. An unsteady point vortex method for coupled fluid–solid problems. Theoretical and Computational Fluid Dynamics, 23(2):127–153, 2009.
- [24] Luca Cortelezzi and A Leonard. Point vortex model of the unsteady separated flow past a semi-infinite plate with transverse motion. Fluid dynamics research, 11(6):263, 1993.
- [25] Darwin Darakananda and Jeff D Eldredge. A versatile taxonomy of low-dimensional vortex models for unsteady aerodynamics. Journal of Fluid Mechanics, 858:917–948, 2019.
- [26] Kiran Ramesh, Ashok Gopalarathnam, Jack R Edwards, Michael V Ol, and Kenneth Granlund. An unsteady airfoil theory applied to pitching motions validated against experiment and computation. Theoretical and Computational Fluid Dynamics, 27(6):843–864, 2013.

- [27] Kiran Ramesh, Ashok Gopalarathnam, Kenneth Granlund, Michael V Ol, and Jack R Edwards. Discrete-vortex method with novel shedding criterion for unsteady aerofoil flows with intermittent leading-edge vortex shedding. Journal of Fluid Mechanics, 751:500–538, 2014.
- [28] Kiran Ramesh, Kenneth Granlund, Michael V Ol, Ashok Gopalarathnam, and Jack R Edwards. Leading-edge flow criticality as a governing factor in leading-edge vortex initiation in unsteady airfoil flows. Theoretical and Computational Fluid Dynamics, 32(2):109–136, 2018.
- [29] Darwin Darakananda, André Fernando de Castro da Silva, Tim Colonius, and Jeff D Eldredge. Data-assimilated low-order vortex modeling of separated flows. Physical Review Fluids, 3(12):124701, 2018.
- [30] Sebastian Liska and Tim Colonius. A fast immersed boundary method for external incompressible viscous flows using lattice Green’s functions. J. Comput. Phys., 331:257–279, 2017.
- [31] Craig H Bishop, Brian J Etherton, and Sharanya J Majumdar. Adaptive sampling with the ensemble transform kalman filter. part i: Theoretical aspects. Monthly weather review, 129(3):420–436, 2001.
- [32] Alessio Spantini, Antti Solonen, Tiangang Cui, James Martin, Luis Tenorio, and Youssef Marzouk. Optimal low-rank approximations of bayesian linear inverse problems. SIAM Journal on Scientific Computing, 37(6):A2451–A2487, 2015.
- [33] Tiangang Cui and Olivier Zahm. Data-free likelihood-informed dimension reduction of bayesian inverse problems. Inverse Problems, 37(4):045009, 2021.
- [34] Mathieu Le Provost, Ricardo Baptista, Youssef Marzouk, and Jeff Eldredge. A low-rank nonlinear ensemble filter for vortex models of aerodynamic flows. In AIAA Scitech 2021 Forum, page 1937, 2021.
- [35] Mathieu Le Provost and Jeff D Eldredge. Mean transport of inertial particles in viscous streaming flows. Physical Review Fluids, 5(5):054302, 2020.
- [36] Alberto Carrassi, Marc Bocquet, Laurent Bertino, and Geir Evensen. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. Wiley Interdisciplinary Reviews: Climate Change, 9(5):e535, 2018.
- [37] Geir Evensen. The ensemble kalman filter for combined state and parameter estimation. IEEE Control Systems Magazine, 29(3):83–104, 2009.
- [38] Sanita Vetra-Carvalho, Peter Jan Van Leeuwen, Lars Nerger, Alexander Barth, M Umer Altaf, Pierre Brasseur, Paul Kirchgessner, and Jean-Marie Beckers. State-of-the-art stochastic data assimilation methods for high-dimensional non-gaussian problems. Tellus A: Dynamic Meteorology and Oceanography, 70(1):1–43, 2018.

- [39] Alessio Spantini, Ricardo Baptista, and Youssef Marzouk. Coupling techniques for nonlinear ensemble filtering. arXiv preprint arXiv:1907.00389, 2019.
- [40] Youssef Marzouk, Tarek Moselhy, Matthew Parno, and Alessio Spantini. Sampling via measure transport: An introduction. Handbook of Uncertainty Quantification, pages 1–41, 2016.
- [41] Jan Mandel, Loren Cobb, and Jonathan D Beezley. On the convergence of the ensemble kalman filter. Applications of Mathematics, 56(6):533–541, 2011.
- [42] Patrick N Raanes. Improvements to ensemble methods for data assimilation in the geosciences. PhD thesis, University of Oxford, 2015.
- [43] Gerrit Burgers, Peter Jan van Leeuwen, and Geir Evensen. Analysis scheme in the ensemble kalman filter. Monthly weather review, 126(6):1719–1724, 1998.
- [44] David M Livings, Sarah L Dance, and Nancy K Nichols. Unbiased ensemble square root filters. Physica D: Nonlinear Phenomena, 237(8):1021–1028, 2008.
- [45] Pavel Sakov and Peter R Oke. A deterministic formulation of the ensemble kalman filter: an alternative to ensemble square root filters. Tellus A: Dynamic Meteorology and Oceanography, 60(2):361–371, 2008.
- [46] Lars Nerger, Tijana Janjić, Jens Schröter, and Wolfgang Hiller. A unification of ensemble square root kalman filters. Monthly Weather Review, 140(7):2335–2345, 2012.
- [47] Julian Tödter and Bodo Ahrens. A second-order exact ensemble square root filter for nonlinear data assimilation. Monthly Weather Review, 143(4):1347–1367, 2015.
- [48] Matthias Katzfuss, Jonathan R Stroud, and Christopher K Wikle. Understanding the ensemble kalman filter. The American Statistician, 70(4):350–357, 2016.
- [49] Jeffrey S Whitaker and Thomas M Hamill. Evaluating methods to account for system errors in ensemble data assimilation. Monthly Weather Review, 140(9):3078–3089, 2012.
- [50] A. Leonard. Vortex methods for flow simulation. J. Comput. Phys., 37(3):289–335, 1980.
- [51] Alexandre Joel Chorin. Numerical study of slightly viscous flow. Journal of fluid mechanics, 57(4):785–796, 1973.
- [52] Ding-Gwo Long. Convergence of the random vortex method in two dimensions. Journal of the American Mathematical Society, 1(4):779–804, 1988.
- [53] Geir Evensen. The Ensemble Kalman Filter: Theoretical formulation and practical implementation. Ocean Dynamics, 53(4):343–367, 2003.

- [54] Darwin Darakananda. Vortex Models for Data Assimilation. PhD thesis, UCLA, 2017.
- [55] Mathieu Le Provost, Ricardo Baptista, Youssef Marzouk, and Jeff D. Eldredge. A low-rank ensemble kalman filter for elliptic observations. arXiv preprint arXiv:2203.05120, 2022.
- [56] Peter Jan Van Leeuwen. Non-local observations and information transfer in data assimilation. Frontiers in Applied Mathematics and Statistics, 5:48, 2019.
- [57] Turgut Sarpkaya. An inviscid model of two-dimensional vortex shedding for transient and asymptotically steady separated flow over an inclined plate. Journal of Fluid Mechanics, 68(1):109–128, 1975.
- [58] Leslie Greengard and Vladimir Rokhlin. A fast algorithm for particle simulations. Journal of computational physics, 73(2):325–348, 1987.
- [59] Frank Ethridge and Leslie Greengard. A new fast-multipole accelerated poisson solver in two dimensions. SIAM Journal on Scientific Computing, 23(3):741–760, 2001.
- [60] Antti Solonen, Tiangang Cui, Janne Hakkarainen, and Youssef Marzouk. On dimension reduction in gaussian filters. Inverse Problems, 32(4):045003, 2016.
- [61] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. The elements of statistical learning, volume 1. Springer series in statistics New York, 2001.
- [62] Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. SIAM Journal on Numerical Analysis, 7(1):1–46, 1970.
- [63] Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the davis–kahan theorem for statisticians. Biometrika, 102(2):315–323, 2015.
- [64] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. Psychometrika, 1(3):211–218, 1936.
- [65] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM review, 53(2):217–288, 2011.
- [66] Armin Eftekhari, Michael B Wakin, Ping Li, and Paul G Constantine. Randomized learning of the second-moment matrix of a smooth function. arXiv preprint arXiv:1612.06339, 2016.
- [67] Paul G Constantine. Active subspaces: Emerging ideas for dimension reduction in parameter studies. SIAM, 2015.

- [68] Marie Bocher, Alexandre Fournier, and Nicolas Coltice. Ensemble kalman filter for the reconstruction of the earth’s mantle circulation. Nonlinear Processes in Geophysics, 25(1):99–123, 2018.
- [69] Matthew Lang, Philip Browne, Peter Jan Van Leeuwen, and Mathew Owens. Data assimilation in the solar wind: Challenges and first results. Space Weather, 15(11):1490–1510, 2017.
- [70] H. M. Hertz. Standing-wave acoustic trap for nonintrusive positioning of microparticles. J. Appl. Phys., 78(8):4845–4849, 1995.
- [71] Albrecht Haake, Adrian Neild, Gerald Radziwill, and Jurg Dual. Positioning, displacement, and localization of cells using ultrasonic forces. Biotechnol. Bioeng., 92(1):8–14, 2005.
- [72] Arthur Ashkin. Acceleration and trapping of particles by radiation pressure. Phys. Rev. Lett., 24(4):156, 1970.
- [73] Justin E. Molloy and Miles J. Padgett. Lights, action: Optical tweezers. Contemp. Phys., 43(4):241–258, 2002.
- [74] Alexander Van Reenen, Arthur M. de Jong, Jaap M. J. den Toonder, and Menno W. J. Prins. Integrated lab-on-chip biosensing systems based on magnetic particle actuation—a comprehensive review. Lab. Chip, 14(12):1966–1986, 2014.
- [75] Manfred Dürr, Jörg Kentsch, Torsten Müller, Thomas Schnelle, and Martin Stelzle. Microdevices for manipulation and accumulation of micro-and nanoparticles by dielectrophoresis. Electrophoresis, 24(4):722–731, 2003.
- [76] Emil Chmela, Robert Tijssen, Marko T. Blom, Han J. G. E. Gardeniers, and Albert van den Berg. A chip system for size separation of macromolecules and particles by hydrodynamic chromatography. Anal. Chem., 74(14):3470–3475, 2002.
- [77] Marko T. Blom, Emil Chmela, R. Edwin Oosterbroek, Rob Tijssen, and Albert Van Den Berg. On-chip hydrodynamic chromatography separation and detection of nanoparticles and biomolecules. Anal. Chem., 75(24):6761–6768, 2003.
- [78] W. P. Raney, J. C. Corelli, and P. J. Westervelt. Acoustical streaming in the vicinity of a cylinder. The Journal Of The Acoustical Society of America, 26(6):1006–1014, 1954.
- [79] N. Riley. Oscillating viscous flows. Mathematika, 12(2):161–175, 1965.
- [80] N. Riley. On a sphere oscillating in a viscous fluid. The Quarterly Journal of Mechanics and Applied Mathematics, 19(4):461–472, 1966.

- [81] Kwitae Chong, Scott D Kelly, Stuart Smith, and Jeff D Eldredge. Inertial particle trapping in viscous streaming. Physics of Fluids, 25(3):033602, 2013.
- [82] Barry R. Lutz, Jian Chen, and Daniel T. Schwartz. Hydrodynamic tweezers: 1. Non-contact trapping of single cells using steady streaming microeddies. Anal. Chem., 78(15):5429–5435, 2006.
- [83] Avi Abadi and Gabor Kosa. Closed loop control of microscopic particles incorporating steady streaming and visual feedback. Biomed. Microdevices, 20(2):28, 2018.
- [84] T. Parthasarathy, F. K. Chat, and M. Gazzola. Streaming-enhanced flow-mediated transport. J. Fluid Mech., 878:647–662, 2019.
- [85] Hoang V. Phan, Muhsincan Şeşen, Tuncay Alan, and Adrian Neild. Single line particle focusing using a vibrating bubble. Appl. Phys. Lett., 105(19):193507, 2014.
- [86] Bhargav Rallabandi, Alvaro Marin, Massimiliano Rossi, Christian J. Kähler, and Sascha Hilgenfeldt. Three-dimensional streaming flow in confined geometries. J. Fluid Mech., 777:408–429, 2015.
- [87] M. Costalonga, Philippe Brunet, and H. Peerhossaini. Low frequency vibration induced streaming in a Hele-Shaw cell. Phys. Fluids, 27(1):013101, 2015.
- [88] Rocío Bolaños-Jiménez, Massimiliano Rossi, David Fernandez Rivas, Christian J. Kähler, and Alvaro Marin. Streaming flow by oscillating bubbles: quantitative diagnostics via particle tracking velocimetry. J. Fluid Mech., 820:529–548, 2017.
- [89] Martin R Maxey and James J Riley. Equation of motion for a small rigid sphere in a nonuniform flow. The Physics of Fluids, 26(4):883–889, 1983.
- [90] P. G. T. Saffman. The lift on a small sphere in a slow shear flow. J. Fluid Mech., 22(2):385–400, 1965.
- [91] Efsthios E. Michaelides. The transient equation of motion for particles, bubbles, and droplets. J. Fluids Eng., 119(2):233–247, 1997.
- [92] Raqeeb Thameem, Bhargav Rallabandi, and Sascha Hilgenfeldt. Fast inertial particle manipulation in oscillating flows. Phys. Rev. Fluids, 2(5):052001, 2017.
- [93] Siddhansh Agarwal, Bhargav Rallabandi, and Sascha Hilgenfeldt. Inertial forces for particle manipulation near oscillating interfaces. Phys. Rev. Fluids, 3(10):104201, 2018.
- [94] D. G. Andrews and M. E. McIntyre. An exact theory of nonlinear waves on a Lagrangian-mean flow. J. Fluid Mech., 89(4):609–646, 1978.

- [95] Martin R. Maxey. The gravitational settling of aerosol particles in homogeneous turbulence and random flow fields. J. Fluid Mech., 174:441–465, 1987.
- [96] Jim Ferry and S. Balachandar. A fast Eulerian method for disperse two-phase flow. Int. J. Multiphase Flow, 27(7):1199–1226, 2001.
- [97] Renwei Mei. An approximate expression for the shear lift force on a spherical particle at finite Reynolds number. Int. J. Multiphase Flow, 18(1):145–147, 1992.
- [98] Q. Wang, K. D. Squires, M. Chen, and J. B. McLaughlin. On the role of the lift force in turbulence simulations of particle deposition. Int. J. Multiphase Flow, 23(4):749–763, 1997.
- [99] Nicolas Mordant and J.-F. Pinton. Velocity measurement of a settling sphere. Eur. Phys. J. B, 18(2):343–352, 2000.
- [100] Stefano Olivieri, Francesco Picano, Gaetano Sardina, Daniele Iudicone, and Luca Brandt. The effect of the Basset history force on particle clustering in homogeneous and isotropic turbulence. Phys. Fluids, 26(4):041704, 2014.
- [101] Kek-Kiong Tio, Amable Liñán, Juan C Lasheras, and Alfonso M Gañán-Calvo. On the dynamics of buoyant and heavy particles in a periodic stuart vortex flow. Journal of Fluid Mechanics, 254:671–699, 1993.
- [102] George Haller and Themistoklis Sapsis. Where do inertial particles go in fluid flows? Physica D, 237(5):573–583, 2008.
- [103] Osborne Reynolds. On the dynamical theory of incompressible viscous fluids and the determination of the criterion. Philos. Trans. Roy. Soc. London, Ser. A., 186:123–164, 1895.
- [104] D. G. Andrews and M. E. McIntyre. On wave-action and its relatives. J. Fluid Mech., 89(4):647–664, 1978.
- [105] Ira B. Bernstein, E. A. Frieman, Martin David Kruskal, and R. M. Kulsrud. An energy principle for hydromagnetic stability problems. Proc. R. Soc. Lond. A, 244(1236):17–40, 1958.
- [106] Darryl D. Holm. Fluctuation effects on 3D Lagrangian mean and Eulerian mean fluid motion. Physica D, 133(1-4):215–269, 1999.
- [107] Shiyi Chen, Darryl D. Holm, Len G. Margolin, and Raoyang Zhang. Direct numerical simulations of the Navier–Stokes alpha model. Physica D, 133(1-4):66–83, 1999.
- [108] Darryl D. Holm. Averaged Lagrangians and the mean effects of fluctuations in ideal fluid dynamics. Physica D, 170(3-4):253–286, 2002.

- [109] Michael Selwyn Longuet-Higgins. Mass transport in water waves. Phil. Trans. R. Soc. Lond. A, 245(903):535–581, 1953.
- [110] Zhilin Li. An overview of the immersed interface method and its applications. Taiwanese Journal of Mathematics, 7(1):1–49, 2003.
- [111] Ernst Hairer, Christian Lubich, and Gerhard Wanner. Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations, volume 31. Springer Science & Business Media, 2006.
- [112] J. Holtmark, I. Johnsen, T. Sikkeland, and S. Skavlem. Boundary layer flow near a cylindrical obstacle in an oscillating, incompressible fluid. The Journal of the Acoustical Society of America, 26(1):26–39, 1954.