INFLOW AND INITIAL CONDITIONS FOR DIRECT NUMERICAL SIMULATION BASED ON ADJOINT DATA ASSIMILATION

Alejandro Gronskis
INRIA Rennes, Bretagne Atlantique
Campus universitaire de Beaulieu
F-35042 Rennes, France
agronskis@gmail.com

Dominique Heitz
CEMAGREF
17 avenue de Cucillé, CS64427
F-35044 Rennes, France
dominique.heitz@cemagref.fr

Etienne Mémin
INRIA Rennes, Bretagne Atlantique
Campus universitaire de Beaulieu
F-35042 Rennes, France
etienne.memin@inria.fr

ABSTRACT

A method for generating inflow conditions for direct numerical simulations (DNS) of spatially-developing flows is presented. The proposed method is based on variational data assimilation and adjoint-based optimization. The estimation is conducted through an iterative process involving a forward integration of a given dynamical model followed by a backward integration of an adjoint system defined by the adjoint of the discrete scheme associated to the dynamical system. The approach's robustness is evaluated on two synthetic velocity field sequences provided by numerical simulation of a mixing layer and a wake flow behind a cylinder. The performance of the technique is also illustrated in a real world application by using PIV measurements to acquire the database. This method allows to denoise experimental velocity fields and to reconstruct a continuous trajectory of motion fields from discrete and unstable measurements.

INTRODUCTION

In spite of significant developments in computational methods over the past few decades, a number of flows are still too expensive to simulate because of their strong unsteady character that must be taken into account to reproduce the main flow features. To overcome these issues and limit the computational cost, new strategies that aim at simulating only the region of interest of the flow have recently been developed. A critical issue that arises is the imperative need for correct specification of all the boundary conditions of the spatial domain. Generating proper inlet conditions for unsteady simulations of spatially developing flows requires the generation of bidimensional vector fields evolving in time in agreement with the spatiotemporal dynamics of the flow.

Usually, turbulent mean velocity profiles can be used and some random noise superimposed in order to start some artificial perturbation supposed to mimic the real instantaneous behaviour of the turbulent flow. In this case, the lack of realistic turbulent structures induces a transient region near the inlet. This development region has no physical significance in general and is of no practical interest. So far, the proposed stategies are divided into two separate parts. In a first step, the inflow condition is built from the data and then the simulation of the downstream development of the flow is conducted from these information. We could address both steps by modifying the inflow condition in a way to get the simulation close to the data. That means that something useful might be learnt from optimal control theory.

This paper presents a new approach based on variational data assimilation (VDA), generating simultaneous transitional initial and inflow boundary conditions and reproducing the spatiotemporal dynamics of the targeted flow. VDA (Le-Dimet and Talagrand, 1986) is a technique derived from optimal control theory (Lions, 1971). It is expressed as the minimization with respect to a control variable of an objective function that measures a discrepancy between a state variable and noisy measurements, subject to a constraint given by the state variable dynamics. The control variable may be for instance a parameter of the dynamics or the initial condition (Papadakis and Memin, 2008). Assuming that both the model and the objective function are differentiable, VDA proposes to solve this inverse problem looking for a control that cancels out the gradient of this cost function through the use of adjoint minimization techniques. Such techniques enable computing the functional gradient by means of the adjoint of the tangent linear dynamics. In practice, the tangent linear dynamics and its adjoint are provided by automatic differentiation (AD) tools. In the present study, the gradient descent minimization is coupled with a limited memory BFGS deterministic gradient based optimization algorithm (Liu and Nocedal, 1989).

DATA ASSIMILATION

We rely on this technique to perform the estimation over time of state variables representing a system of interest. This framework allows to handle high-dimensional state spaces and is thus intensively used in environmental sciences (Le-Dimet and Talagrand, 1986) for atmospheric flows or oceanic simulations. The problem we are dealing with consists in recovering a system's state $\mathbf{X}(\mathbf{x},t)$ obeying a dynamical law, given some noisy and possibly incomplete measurements $\mathbf{Y}$ of the state. The measurements, in this context also called observations, are assumed to be available only at discrete instants t^* separated by a given latency Δt_{obs}. This is formalized, for any location, $\mathbf{x}$, at time $t \in [t_0, t_f]$, by the system

$$\partial_t \mathbf{X}(\mathbf{x},t) + \mathbb{M}(\mathbf{X}(\mathbf{x},t), \eta(t)) = 0 \qquad (1)$$

$$\mathbf{X}(\mathbf{x},t_0) = \mathbf{X_0}(\mathbf{x}) + \varepsilon(\mathbf{x}) \qquad (2)$$

where $\mathbb{M}$ is a nonlinear dynamical operator depending on a control parameter η. The term $\mathbf{X_0}$ is the initial vector at time t_0, and ε is an (unknown) additive control variable on the initial condition.

Continuous Model

The dynamical model that relates the state function to the unknown of our inverse problem is based on *pressure-velocity* formulation of *Navier-Stokes* equations and is briefly described below.

Governing Equations
Mass and momentum conservation principles are represented by the Navier-Stokes equations, which have the following form for an incompressible fluid

$$\nabla \cdot \mathbf{u} = 0 \qquad (3)$$

$$\frac{\partial \mathbf{u}}{\partial t} = -\nabla p - \omega \times \mathbf{u} + \nu \nabla^2 \mathbf{u} \qquad (4)$$

where ν is the kinematic viscosity, $p(\mathbf{x},t)$ the dynamic pressure field $(P + \frac{1}{2}\rho|\mathbf{u}|^2)$, $\mathbf{u}(\mathbf{x},t)$ the velocity field, and $\omega(\mathbf{x},t)$ the vorticity field $(\nabla \times \mathbf{u})$.

Numerical Method
A numerical code fully based on sixth-order compact finite difference schemes and a Cartesian grid is used to solve the incompressible Navier-Stokes equations. The incompressibility condition is ensured via a fractional step method introducing a Poisson equation for the pressure. An original characteristic of the present code, called *Incompact3d*, is that this equation is directly solved in the framework of the modified spectral formalism (Laizet and Lamballais, 2009). The time advancement is performed using a second-order Adams-Bashforth scheme. Free-slip boundary conditions are applied at $y = \pm L_y/2$, whereas outflow boundary conditions at $x = L_x$ are determined through the resolution of a simplified convection equation.

Cost Functional

Basically, we are considering a minimization problem under a particular additional constraint: the state dynamics (1-2), which expresses the dependence of the flow field $\mathbf{X} \equiv \mathbf{u}(\gamma)$

on the control variable $\gamma = \{\varepsilon(\mathbf{x}), \eta(t)\} \equiv \{\mathbf{u}(\mathbf{x},t_0), \mathbf{u}(\mathbf{x}_{in},t)\}$. Here, $\mathbf{x}_{in}$ stands for the inlet section of the domain, whereas $\mathbf{u}(\mathbf{x},t_0)$ is the initial condition. The objective consists then to find an optimal control of low energy γ^* that leads to the lowest discrepancy between the measurements and the state variable. This leads to the minimization problem

$$J(\gamma) = \int_{t_1}^{t_N} \int_{\Omega_A} \|\overline{\omega}(\mathbf{x},t^*) - \omega^{obs}(\mathbf{x},t^*)\|_R^2 \, d\mathbf{x} \, dt^*$$
$$+ \int_{t_0}^{t_N} \int_{\Omega_C} \|\gamma(\mathbf{x},t) - \overline{\gamma}(\mathbf{x},t)\|_Q^2 \, d\mathbf{x} \, dt \qquad (5)$$

where $t_1 \in [t_0, t_f]$ is the initial time of the assimilation window, $\overline{\omega}$ is a spatial average of the vorticity field generated from the model output, ω^{obs} the vorticity field coming from the observations, and $\overline{\gamma} = \{\overline{\mathbf{u}}(\mathbf{x},t_0), \overline{\mathbf{u}}(\mathbf{x}_{in},t)\}$ a smoothed value of the control variable. Here, $\overline{\mathbf{u}}(\mathbf{x},t_0)$ is a spatial average on the initial condition whereas $\overline{\mathbf{u}}(\mathbf{x}_{in},t)$ is a temporal average on the inflow condition. The norms $\|\cdot\|_R$ and $\|\cdot\|_Q$ are induced norms of the inner products $< R\cdot, \cdot >$ and $< Q\cdot, \cdot >$; R and Q are covariance matrices of the assimilation space (spatial domain Ω_A) and control space (spatial domain Ω_C). Few remarks can be done here. For numerical stability reasons, the spatial resolution (dx) used in our simulation code is generally higher than that of the observations (dx_g) coming from Particle Image Velocimetry (PIV) measurements. Thus, the spatial averaging applied to the fine grid vorticity field obtained by DNS allows us comparing the state variable with the coarse grid vorticity field corresponding to the observations. This spatial averaging is made by means of a Gaussian weighting function over the region Λ, where $\Lambda = \{(x,y) \in \Omega_C \mid (x^2 + y^2) \leq (dx_g/dx)^2\}$. In this way, this procedure exhibits an analogy with the weighting over the span of the interrogation window in conventional correlation PIV implementations. Nevertheless, the first term of the functional causes the optimization system to give a peak shape response in the model trajectory at instants t^*. One possible approximation to handle this problem is to perform a temporal average on the inflow condition over the interval $[t^* - \Delta t_{obs}, t^* + \Delta t_{obs}]$. This strategy allows the inflow condition to change over the whole sequence time range $[t_0, t_f]$. Thus, the second term of the functional enforces dynamical coherence and a temporal continuous trajectory of the solution.

Adjoint Model

Regarding the minimization of the objective function, a direct numerical evaluation of the functional gradient is computationally infeasible, because this would require to compute perturbations of the state variables along all the components of the control variables $(\delta\varepsilon, \delta\eta)$ - i.e. to integrate the dynamical model for all pertubed components of the control variable, which is obviously not possible in practice. A solution to this problem consists to rely on an adjoint formulation (Le-Dimet and Talagrand, 1986). Within this formalism, the gradient functional is obtained by a forward integration of the dynamical system (1-2) followed by a backward integration of an adjoint dynamical model given by

$$-\partial_t \lambda(\mathbf{x},t) + (\partial_{\mathbf{X}}\mathbb{M})^* \lambda(\mathbf{x},t) = R(\mathbf{Y} - \mathbf{X})(\mathbf{x},t) \qquad (6)$$

$$\lambda(\mathbf{x},t_f) = 0 \qquad (7)$$

where the adjoint variable λ and the adjoint of the tangent linear operator $(\partial_X \mathbb{M})^*$ have been introduced. This adjoint model is defined by the adjoint of the discrete scheme associated to the dynamical system. For our purpose, AD adjoint is generated by a tool, called *Tapenade* (Hascoet, 2003). This systematic approach leads however to massive use of storage, requiring code transformation by hand to reduce memory usage as explained in the next section.

Building Adjoint Algorithms through AD

Automatic Differentiation (AD) is a technique to evaluate derivatives of a function $F : X \in \mathbb{R}^m \mapsto Y \in \mathbb{R}^n$ defined by a computer program P. In AD, the original program is automatically transformed or extended to a new program P' that computes the derivatives analytically. In our case, we set the input $X \equiv \gamma$ and function F has a single output (the cost).

Application of AD An AD tool uses the source of the program that computes the state dynamics (1-2), and identifies this program with a composition of mathematical functions, one per run-time instruction. Calling $\{I_k\}_{k=1 \to p}$ the sequence of instructions executed at run-time, each of them implementing an elementary function f_k, the function F computed by P is:

$$F = f_p \circ f_{p-1} \circ \cdots \circ f_2 \circ f_1 \tag{8}$$

Setting $W_0 = X$ and $W_k = f_k(W_{k-1})$, the chain rule gives us the Jacobian of F:

$$F'(X) = f'_p(W_{p-1}) \cdot f'_{p-1}(W_{p-2}) \cdot \cdots \cdot f'_1(W_0) \tag{9}$$

The so-called *adjoint* (or *reverse*) mode of AD aims at computing the product of the transposition of F' by a given weight vector $\overline{Y}$ to get:

$$\overline{X} = F'^t(X) \cdot \overline{Y} = f'^t_1(W_0) \cdot f'^t_2(W_1) \cdot \cdots \cdot f'^t_p(W_{p-1}) \cdot \overline{Y} \tag{10}$$

In this way, the adjoint mode of AD builds a new code $\overline{\text{P}}$ that computes $\{\overline{W}_{k-1} = f'^t_k(W_{k-1}) \cdot \overline{W}_k\}_{k=p \to 1}$, by using values from P in the reverse of their computation order. The stack needed to store W_k is the bottleneck of reverse AD. In order to keep it small enough, we applied a storage/recomputation strategy described as follows:
(i) During the computation of the cost function J, we store in memory only W_k at instants t^*.
(ii) During the computation of the adjoint variable $\overline{X} \equiv \lambda$, we restart the program on snapshot W_k until W_{k+1} in the forward sweep of P, storing in memory all the intermediate values W_{itime} at each DNS time step. Later, in the backward sweep, each W_{itime} is restored from the stack to be used by $\overline{\text{P}}$.

In our case, by setting the weight vector $\overline{Y} \equiv Y$, the code $\overline{\text{P}}$ enables to get the gradient functional $\overline{X} = F'^t(X) \cdot Y \equiv \nabla Y$.

Optimization with Gradient Descent

In our optimization problem, Hessian matrix of the cost function is too dense if we take into account the large number of variables involved, i.e. $(2 \cdot n_x \cdot n_y + 2 \cdot n_y \cdot \Delta\text{itime})$, where n_x, n_y is the number of grid points in control space domain Ω_C, and Δitime the number of DNS time steps in $[t_0, t_f]$. So, in order to reduce the cost of storing and manipulating it, we have chosen a limited-memory quasi-Newton method which maintains simple and compact approximations of Hessian matrices (Liu and Nocedal, 1989). We have used an algorithm known as L-BFGS, which is based on the BFGS updating formula

$$\gamma_{k+1} = \gamma_k + \alpha_k p_k \tag{11}$$

where α_k is the step length, $p_k = -H_k \nabla Y_k$ is the search direction, and H_k denotes an approximation of the Hessian which is updated at every iteration. The step length is computed from a line search procedure to satisfy the Wolfe conditions.

RESULTS

At first, we generated tangent (forward mode AD) and adjoint (reverse mode AD) codes for the present DNS code. We performed validation experiments (dot-product test) to check for correctness of both AD generated codes.

Control on the Initial Condition

In order to assess the benefits of our technique, we carried out a VDA experiment applied to the identification of the initial condition from a spatially evolving 2D mixing layer flow by using numerical data. The velocity ratio between low and high speed is 0.66 and the Reynolds number based on the velocity difference of two streams and the inflow vorticity thickness $\delta_{\omega 0}$ is $Re = 400$. The governing equations are directly solved using 401×257 grid points in x and y directions respectively. The assimilation domain size was $L_x = 200\delta_{\omega 0}$, $L_y = 512\delta_{\omega 0}$, and the grid resolution in x direction was $\Delta x = 0.5\delta_{\omega 0}$. The stretching of the grid in y direction leads to a minimal mesh size of $\Delta y_{min} = 0.15\delta_{\omega 0}$.

As a first step, a precursor simulation was run to provide artificial inflow conditions for the assimilation experiment. The spanwise section retained which will correspond to the inlet section of the assimilation domain is located at the beginning of the convection region of vortical structures. Subsequently, synthetic observations generated by the direct model trajectory are perturbed using zero mean Gaussian noise and an additional spatial filtering is applied to these vector velocity fields in order to provide smoother motion flow.

In this case the functional $J(\gamma)$ depends only on the initial condition, and comes to an initial value control problem. Figure 1 indicates that the true state at the end of the assimilation window (a) is recovered with a good approximation (d) from the perturbed one (b), showing the quality of the derivatives obtained.

Control on the Initial and Inflow Condition

In the previous assimilation experiment, an artificial inlet condition obtained from the reference DNS trajectory was used as a fixed parameter in our optimization system, and then $\gamma = \{\mathbf{u}(\mathbf{x}, t_0)\}$. From now on, we will try to incorporate inflow condition as a control parameter. In order to make it possible, we need an initial estimate for the inflow condition, i.e. $\mathbf{u}_{k=0}(\mathbf{x}_{in}, t)$. At this point we put emphasize on the initial estimate chosen for the initial condition

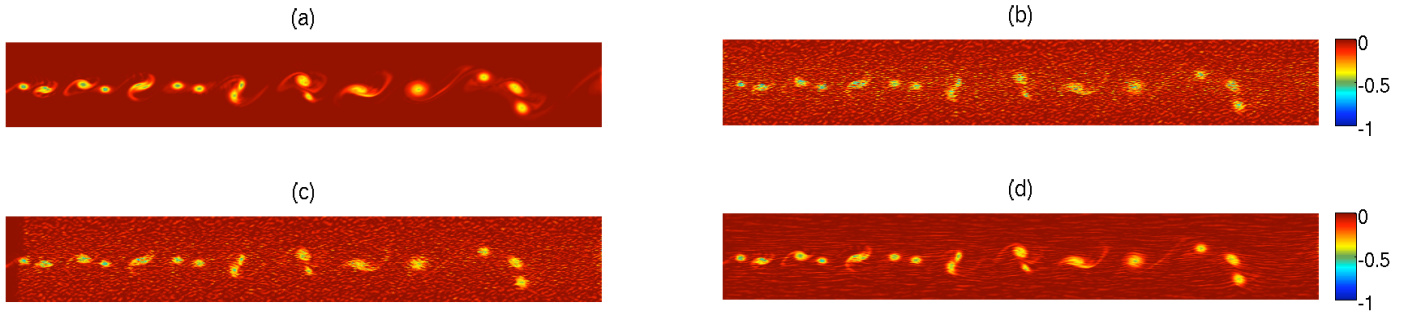$$\mathbf{u}_{k=0}(\mathbf{x}, t_0) = \mathbf{u}^{obs}(\mathbf{x}, t_0) \tag{12}$$

Figure 1. Vorticity fields at the end of the assimilation window. (a) True field; (b) Noisy observation; (c) Field obtained by starting from noisy observation at the begining of the assimilation window; (d) Field obtained by starting from the retrieved optimal initial conditions.

A natural choice consists to use the complete sequence of observations to construct $\mathbf{u}_{k=0}(\mathbf{x}_{in}, t)$. From Taylor's hypothesis we can find the spatial change caused by advection between two contiguous observations by $\Delta x_{obs} = U_m \Delta t_{obs}$, where U_m is the mean convection velocity at the inlet section. Then we constructed the initial estimate for the inflow condition as

$$\mathbf{u}_{k=0}(\mathbf{x}_{in}, t_\sigma) = \mathbf{u}^{obs}(\mathbf{x}_\sigma, t^*) \qquad (13)$$

where $\mathbf{x}_\sigma = \mathbf{x}_\sigma(\Delta x_{obs}, dx)$, $t_\sigma = t_\sigma(t^*, \Delta t_{obs}, dt)$.

In order to assess the performance of the VDA method for the specification of inflow condition we have constituted a benchmark composed of DNS results and experimental PIV data. The numerical simulation and the experimental data concerns both a wake behind a circular cylinder at Reynolds 125 and 200 respectively (Reynolds number based on the free stream velocity U_∞, the kinematic viscosity ν and the diameter of the circular cylinder D). For these Reynolds numbers, the transition to turbulence takes place in the wake. This regime is identified with the Bénard-von Karman vortex street and the largest scales remain bidimensional.

Numerical Data Assimilation The VDA approach is now validated with an identical twin experiment that generates the observation data using a numerical simulation of reference. Our computational domain size is $L_x \times L_y = 20D \times 20D$ and the corresponding number of points is $n_x \times n_y = 1801 \times 721$. A constant flow is imposed at the entrance of this reference (or true state) domain, and the center of the cylinder is located at $x_{cyl} = 8D$ downstream of the inflow. The simulation was carried out with a constant time step size of $\Delta t = 0.012 D/U_\infty$. A sequence of 50 velocity fields with a $10\Delta t$ time steps between them has been kept for building a snapshots sequence of the flow motion. The spatial domain size has been reduced from this database to $9D \times 9D$. The inlet section of the assimilation domain has been chosen at $22D$ from the center of the cylinder, i.e. $(x_{in} - x_{cyl}) = 14D$, in a way to satisfy Taylor's hypothesis by verifying $\langle u_y' u_y' \rangle_{max}/U_m \approx 0.1$ at the inlet. To mimic a typical experimental situation, the spatial resolution has been reduced in a factor 5, yielding 65×64 points for the assimilation domain in the streamwise and normal directions.

Figure 2 presents the observation and model dynamics trajectories at two different stages of the assimilation process.
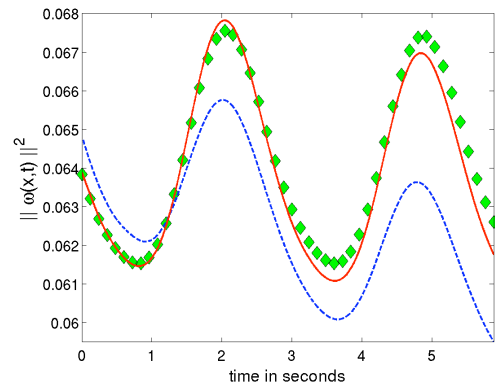

Figure 2. VDA cylinder wake twin test. Temporal evolution of the squared norm of the vorticity field; diamonds: observations; dashed line: initial approximation; solid line: assimilated solution.

The initial trajectory is generated from the model output by

1. starting from the spatially interpolated observation at $t = t_0$, and
2. using a temporally smoothed inflow condition which has been constructed from the observations by applying Taylor's hypothesis.

The assimilated trajectory is obtained from the best estimate for the initial and inflow condition found by L-BFGS algorithm. As it can be observed from these results, this trajectory fits almost perfectly the observation data. The inflow control variable enables to complete the missing elements of the initial DNS dynamics trajectory to explain the data. Moreover, as described below, the analysis of the velocity spectra corroborates the fact that the assimilated solution reproduces the correct turbulent scales. Spectra of the longitudinal velocity component corresponding to the assimilated and true trajectories obtained on the $y = 0$ axis are presented in Figure 3 and are compared to spectrum obtained at the same location coming from observations.

Spectrum of the assimilated data shows a wide-range spectral content in agreement with the reference. As expected, the assimilation method allows to recover the low frequency range missing in the observations. Moreover, vortex shed-
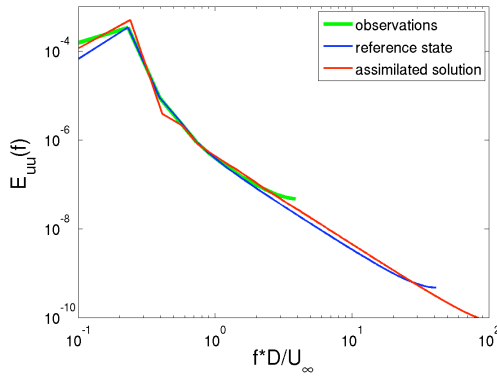
Figure 3. VDA cylinder wake twin test. Mean energy spectrum of the longitudinal velocity component on the $y = 0$ axis.

Figure 4. Temporal evolution of the squared norm of the vorticity field for (top) VDA cylinder wake experimental test B and (bottom) test A; diamonds: observations; dashed line: initial approximation; solid line: assimilated solution.

ding frequencies characteristic of von Karman street, corresponding to a peak in the spectra around frequencies of order $fD/U_\infty \approx 0.2$, are also well reproduced for the assimilated data.

Clearly, this twin configuration provides enough information (e.g. $\Delta t_{obs} = 0.12D/U_\infty$) to yield meaningful results. But in real world applications it may not be feasible to have noise-free velocity fields or high temporal resolution sequence, since observations are provided by PIV.

Experimental Data Assimilation Time-resolved 2D PIV measures of the wake behind a circular cylinder at Reynolds 200 has been carried out at the laboratory of Fluid Dynamics (Buenos Aires University, Argentina). A sequence of 10 velocity fields with a temporal resolution of $\Delta t_{obs} = 1.1D/U_\infty$ and with a temporal length of approximatively 4 vortex shedding has been kept for building the snapshots sequence of the flow motion. Two assimilation tests were performed by changing the inlet location of the assimilation domain. Table 1 summarizes the main spatial assimilation domain characteristics of these velocity fields.

Figure 4 presents the observation and DNS trajectories for both VDA tests at two different stages of the optimization cycle.

As it can be observed from these curves assimilation technique enables to modify inflow condition to recover with a good accuracy the trajectory corresponding to the observations. When the inlet of the assimilation domain is located downstream of the vortex formation region (VDA test A), initial guess stay close to the observation data and the optimization algorithm requires 80 iterations to reach the assimilated trajectory. On the other hand, if the inlet belongs to the vortex formation region (VDA test B), Taylor's hypothesis does not hold (the relative turbulence intensity is high) and the initial approximation is far from the data. This leads to an increase of the number of iterations required (150) to get the assimilated solution.

To characterize the assimilation results further and more quantitatively, mean flow characteristics are compared. Figure 5 presents the downstream evolution of the profiles of both the mean and fluctuating vorticity for VDA test A.

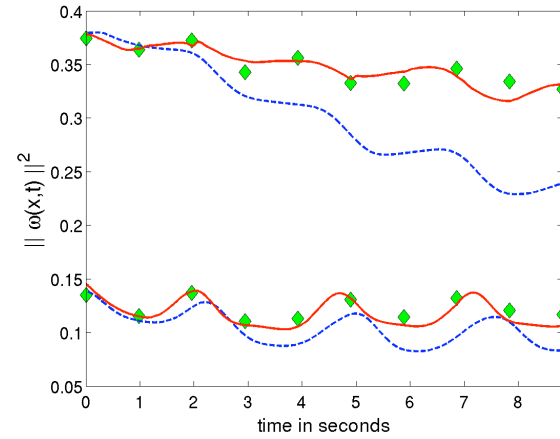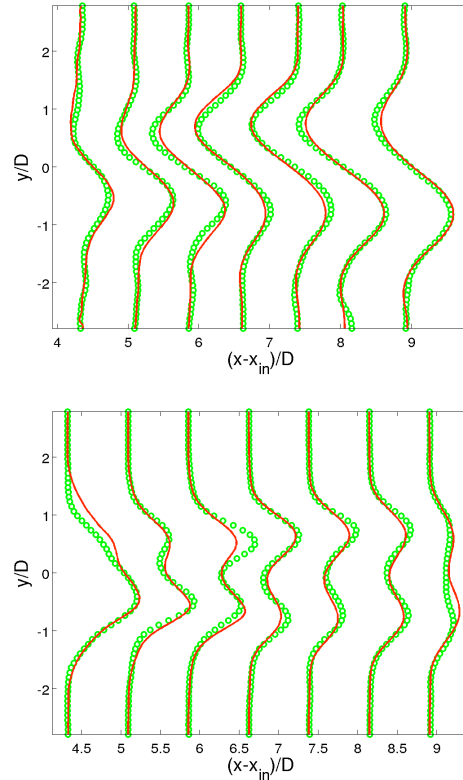It reveals that temporal mean properties of assimilated vor-

Figure 5. VDA cylinder wake experimental test A. Longitudinal evolution of (top) the mean vorticity and (bottom) the vorticity fluctuation $\langle \omega' \omega' \rangle$; solid lines: assimilated solution; symbols: observations.

ticity are in good agreement with experimental ones. To illustrate the vorticity fields estimated through the assimilation procedure, we plot in Figure 6 two pairs of consecutive snapshots of the vorticity corresponding to the DNS and the PIV

5

Table 1. VDA cylinder wake experimental tests from PIV velocity fields. Characteristics of the spatial assimilation domain.

VDA Test	$L_x \times L_y$	$n_x \times n_y$	$(x_{in} - x_{cyl})$
A	$5D \times 6D$	55×67	$4.5D$
B	$9D \times 6D$	99×67	D

$t^* = 6.8$ $t^* = 7.8$
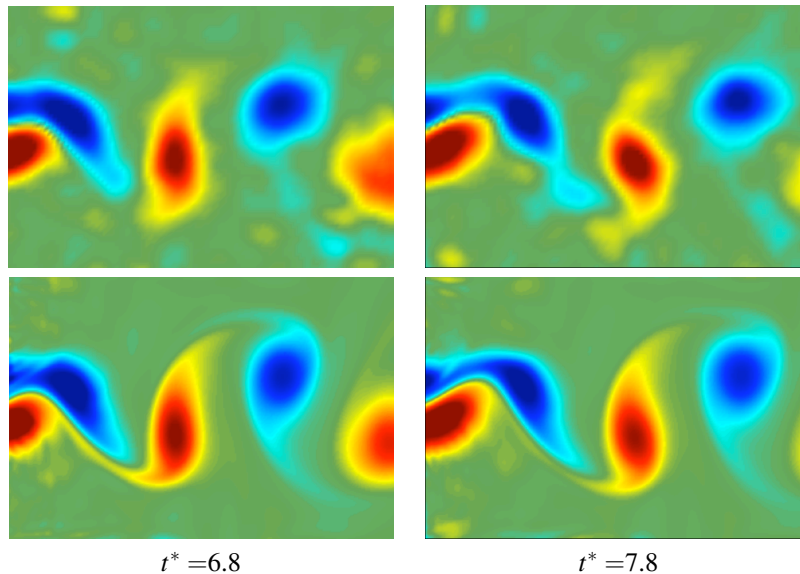
Figure 6. VDA cylinder wake experimental test B. Vorticity fields at two consecutive discrete instants in the assimilation window. Experimental observations (*Top*); Assimilated state (*Bottom*).

sequence for VDA test B.

Results indicate that despite low spatial and temporal resolution observations the assimilated state exhibits fine scale details of vorticity like vortex filaments. Furthermore, it should be noted that the proposed method provides a means to simulate a wake flow without simulating the flow around the obstacle.

code is now being considered. Future work will also include extending the adjoint code capabilities to include three-dimensional effects in data assimilation system. Furthermore, considering the outflow condition, the actual configuration using a purely convective flow assumption could be improved by introducing the outflow boundary condition as a control parameter of the optimization problem.

CONCLUSIONS

In this work, a new method for generating inflow boundary conditions for DNS has been introduced. This approach relies on variational data assimilation principles and adjoint-based optimization. By modifying the initial and inflow condition of the system, the proposed method allows us to recover the state of an unknown function on the basis of a DNS model and noisy measurements. In order to test this new approach, DNS of a 2D mixing layer flow and a wake flow behind a circular cylinder have been performed to provide a synthetic database. Both twin experiments allowed the validation of the solution methodology in a controlled scenario and demonstrated the feasibility and the reliability of the proposed method. The potential of the assimilation technique were also illustrated in a real world application. For this purpose, PIV has been chosen to acquire the database.

To go further, it could be interesting to introduce dynamic laws related to the observed phenomenon at higher Reynolds numbers; in this sense, an attempt at combining an experimental database to a large-eddy simulation (LES)

REFERENCES

Hascoet, L., Greborio, R., and Pascual, V., 2003, "Computing adjoints by automatic differentiation with tapenade", Research report, INRIA.

Laizet, S., and Lamballais, E., 2009, "High-order compact schemes for incompressible flows: a simple and efficient method with the quasi-spectral accuracy", J. Comp. Phys., 228(16), 5989-6015.

Le-Dimet, F. X., and Talagrand, O., 1986, "Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects", Tellus 38 (A), 97-110.

Lions, J.L., 1971, "Optimal control of systems governed by PDEs", Springer-Verlag, New York.

Liu, D., and Nocedal, J., 1989, "On the limited memory BFGS method for large scale optimization", Mathematical Programming, Series B, 45(3), 503-528.

Papadakis, N., and Memin, E., 2008, "Variational assimilation of fluid motion from image sequences", SIAM Journal on Imaging Science, 1(4), 343-363.