10TH INTERNATIONAL SYMPOSIUM ON PARTICLE IMAGE VELOCIMETRY - PIV13
Delft, The Netherlands, July 1-3, 2013

A matching pursuit approach to solenoidal filtering of three-dimensional velocity measurements

D. Schiavazzi[1], F. Coletti[2], G. Iaccarino[2] and J.K. Eaton[2]

[1] Mechanical and Aerospace Engineering Department, University of California, San Diego, US
[2] Department of Mechanical Engineering, Stanford University, Stanford, US
email: dschiavazzi@ucsd.edu, colettif@stanford.edu

ABSTRACT

Methodologies to acquire three-dimensional velocity fields are becoming increasingly available, generating large datasets of steady state and transient flows of engineering and/or biomedical interest. This paper presents a novel linear filter for three-dimensional velocity acquisitions, which eliminates the spurious velocity divergence due to measurement errors. The noise reduction properties of the associated linear operator are discussed together with the treatment of boundary conditions. Examples show the application of the technique to real velocity fields acquired through volumetric Particle Image Velocimetry. The effectiveness of the filter is demonstrated by application to synthetic velocity fields obtained from analytical solutions and computations. The filter eliminates about half of the noise, without artificial smoothing of the original data, and conserves localized flow features.

1. Introduction

Fluid flows in engineering and biomedical applications are typically complex, unsteady and highly three-dimensional. Recent advances have led to several approaches capable of measuring the three velocity components over a three-dimensional volume (often referred to as 3D-3C velocity measurements). This opens exciting new possibilities for a deeper understanding of complex flow phenomena, as well as new avenues to extract information from experiments.

Among the various three-dimensional velocimetry techniques available today, Particle Tracking Velocimetry (PTV) was probably the first to be successfully demonstrated, and it has been employed extensively for the study of Lagrangian particle motion in turbulence [10]. Defocusing digital particle image velocimetry has been applied to single-phase and bubbly flow [8]. Scanning-light-sheet PIV has been able to describe complex flow patterns in relatively low velocity flows [5]. Holography is three-dimensional in nature and can be very effective in small measurement volumes [4]. Tomographic PIV [1] has found the largest success due to its broad applicability and has triggered widespread interest in the fluid mechanics community.

Noise and intrinsic errors inevitably affect three-dimensional flow measurements generating a spurious, finite divergence. Many flows of interest are effectively incompressible, which translates into a zero-divergence constraint for the velocity field as dictated by mass-conservation. The development of a filter to extract a divergence-free field of minimum distance (in the sense of the closest approximation) from the measurements is therefore of primary interest. Besides improving accuracy, the resulting divergence-free velocities are most suitable to extract physical quantities such as pressure and vorticity.

We propose a technique that generates a least-squares zero-divergence approximation of the acquired velocity field in a way compatible with the measurement resolution. For this purpose, we use tools originally developed for signal processing applications [7] to decompose a given velocity field in a linear combination of local solenoidal waveforms. The proposed strategy shows promising noise reduction capabilities, is amenable to parallel implementation and allows seamless treatment of boundary conditions and solid walls.

The paper is structured as follows. The problem of extracting a divergence-free field from measured velocities is analyzed in Section 2, where a Sequential Matching Pursuit (SMP) algorithm is proposed and its properties analyzed. The proposed algorithm is applied to two- and three-dimensional velocity fields in Sections 3. Section 4 is devoted to the analysis of the de-noising properties of the proposed approach. Conclusions and future perspectives are given in Section 5.

2. A signal processing interpretation of divergence free filtering

Consider a compact volume $\Omega \subset \mathbb{R}^3$ of fluid, with a piecewise smooth boundary $\partial\Omega$. If $\mathbf{v}$ is an incompressible, continuous and differentiable vector field defined on a neighborhood of Ω of normal $\mathbf{n}$, the following integral constraint holds

$$\int_{\Omega} (\nabla \cdot \mathbf{v})\, d\Omega = \int_{\partial\Omega} (\mathbf{v} \cdot \mathbf{n})\, dS = 0, \tag{1}$$

i.e., for every closed surface fully contained in the domain of interest, the divergence-free condition corresponds to zero net flux across the boundaries. The linear nature of the above constraint suggests a possible strategy to extract the divergence-free component from an

arbitrary vector field, i.e., by projection onto the space generated by linear combination of velocity patterns having the above integral property.

2.1 Projections onto Divergence-free vector spaces

Assume we acquire a finite dimensional representation of a signal f (next section will clarify the relationship between f and the actual measured velocity field) at a discrete number of locations $x^{(i)}$, $i = 1,\ldots,N$, where N is the total number of available samples. A representation of f is sought in terms of a set of *waveforms* (also referred to as *atoms*), as follows:

$$f|_{x^{(i)}} = f_i = \sum_{j=1}^{P} \alpha_j \, w_j(x^{(i)}) \quad i = 1,\ldots,N \quad \text{or, in matrix form} \quad \mathbf{f} = \mathbf{W}\boldsymbol{\alpha}, \tag{2}$$

where vector $\mathbf{f} \in \mathbb{R}^N$ contains the measurements of signal f, $\mathbf{W} \in \mathbb{R}^{N\times P}$ is the waveform dictionary of cardinality P, and $\boldsymbol{\alpha} \in \mathbb{R}^P$ is the unknown coefficient vector. The columns of $\mathbf{W}$ are discrete realizations of single waveforms. The vector $\boldsymbol{\alpha}$ can be computed using different techniques, depending on the choice of N, P and using prior information on the nature of the resulting approximation. For the case where $N \gg P$, an approximate solution $\boldsymbol{\alpha}^*$ can be computed using a least-squares minimization of the residual,

$$\boldsymbol{\alpha}^* = \min_{\boldsymbol{\alpha}} \|\mathbf{f} - \mathbf{W}\boldsymbol{\alpha}\|_2. \tag{3}$$

This will lead to a projection in the space spanned by the atoms in $\mathbf{W}$ of minimal distance (in ℓ_2 sense) from the original signal. We are particularly interested in the case $N < P$ where the above approximation is sought according to a redundant dictionary of waveforms.

2.2 Sequential matching pursuit algorithm (SMP)

The vector f is projected onto span$\{\mathbf{w}_i,\, i = 1,\ldots,P\}$ by iterative refinements of the coefficients α_i. These coefficients are computed by correlating the i-th waveform with the residual, as follows:

$$\alpha_i = \mathbf{r}_{i-1}^T \mathbf{w}_i, \text{ where} \quad \mathbf{r}_i = \mathbf{r}_{i-1} - \alpha_i \mathbf{w}_i \quad \text{and} \quad \mathbf{r}_0 = \mathbf{f}. \tag{4}$$

Note that we always assume a normalized waveform, i.e., $\|\mathbf{w}_i\|_2 = 1$, $i = 1,\ldots,P$ in our development. The reconstructed solution $\mathbf{f}_i^*$ at the i-th iteration is computed as

$$\mathbf{f}_i^* = \mathbf{f}_{i-1}^* + \alpha_i \mathbf{w}_i, \text{ where} \quad \mathbf{f}_0^* = \mathbf{0}. \tag{5}$$

Note the minimal storage requirements of the algorithm where, at every iteration, only two terms of the decomposition $\mathbf{f}_i^* + \mathbf{r}_i = \mathbf{f}$ need to be stored, i.e., the reconstructed solution and the residual vector. With simple arguments it is also possible to show that the residual monotonically decreases and the rate of convergence at every iteration depends on the correlation angle θ_i between the residual at the current iteration and the selected waveform. Finally, note that the SMP heuristic shown above allows efficient correlation of waveforms that are sparse in $\mathbb{R}^N$ or associated with a pre-defined structure.

2.3 Velocity encoding-decoding

Our approach implements a normal face flux (normal edge flux in two-dimensional applications) encoding of the measured velocities. In practice, we choose N as the total number of faces in the dual mesh associated with the measurement grid, while f_i contains the normal flux across face i, i.e., before starting the evaluation of the representation coefficients through SMP, the acquired velocities are transformed to face fluxes (encoding). At the end of the projection step, filtered velocities are transformed back from reconstructed fluxes (decoding).

Assume N_p to be the total number of grid points for the acquired velocity field, we define the three vectors $\boldsymbol{v}_x, \boldsymbol{v}_y, \boldsymbol{v}_z \in \mathbb{R}^{N_p}$ containing the three spatial components for all the acquired grid points. We order sequentially these components in the vector $\boldsymbol{v} \in \mathbb{R}^{3N_p}$. Encoding and decoding operations can be written as:

$$\boldsymbol{f} = \boldsymbol{E}\boldsymbol{v} \quad \text{and} \quad \boldsymbol{v} = \boldsymbol{D}\boldsymbol{f} \tag{6}$$

2.4 Solenoidal atoms for discretized domains

We note that solenoidal atoms should be selected consistently with the resolution of the acquired data, and need to be generators for all possible divergence-free flow patterns for the selected discrete grid. A representation of the flow in terms of normal face fluxes is naturally suited to be decomposed using discrete divergence-free atoms. Examples of atoms are illustrated in Figure 5a and 5b, for 2D unstructured and 3D structured acquisition grids, respectively. The patterns are analogous to vortices rotating around every vertex (in two-dimensional layouts) or around the $\{X,Y,Z\}$ axes (in 3D structured layouts). The above concept can be easily generalized to arbitrary discretizations in 3D.

To increase the convergence rate of the method, every other divergence-free waveform can be added to the dictionary $\mathbf{W}$. For example, a constant flow component throughout the domain, correlates slowly with local vortex-like diverge-free atoms illustrated in Figure 5a and 5b. Global atoms are therefore added to $\mathbf{W}$ for situations characterized by significant advective flow.

2.5 Sparsity and parallel computations

A divergence-free atom $\mathbf{w}_i$ is said to be *sparse* if it can be represented in terms of a small number of normal face fluxes. In particular, we define an atom to be s-sparse, if it can be expressed exactly using s normal fluxes. The graphical representation provided in Figure 5a and 5b, helps in quantifying such a sparsity. In the two-dimensional, unstructured case where velocity acquisitions are located at

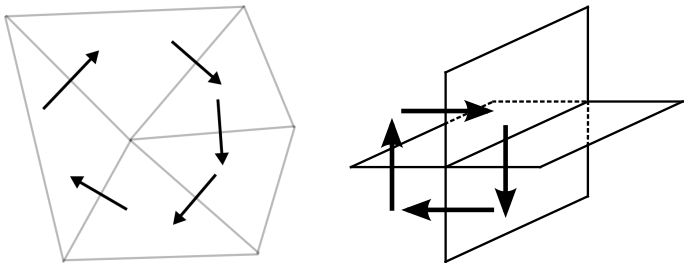

(a) (b)

Figure 1: Examples of divergence-free flux patterns (atoms) for unstructured 2D discretization (a) and structured 3D mesh (b).

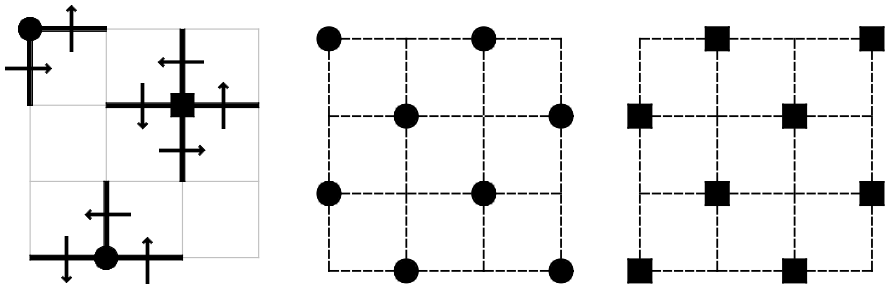

(a) (b) (c)

Figure 2: Two-dimensional visualization of faces forming divergence-free atoms (a). Waveforms with disjoint support whose correlation can be performed concurrently in the first stage (b) and second stage (c) of execution.

triangle centroids, s is the number of edges shared by the same vertex. For three-dimensional configurations, 2, 3 and 4-sparse atoms are possible depending on the location of the edge (segment shared by neighbor faces). This means that, in the latter case, up to 4 multiplications and 3 additions are needed to find the incremental coefficient correlating the components of the residual at the current iteration with the non-zero components of $\mathbf{w}_i$. We also define the support of atom $\mathbf{w}_i$ as the set $\Gamma_i = \{j : w_{i,j} \neq 0, i = 1,\ldots,P\}$ and say that the two atoms $\mathbf{w}_i$ and $\mathbf{w}_j$ have disjoint support if $\Gamma_i \cap \Gamma_j = \{0\}$. A key feature of the proposed algorithm is that correlation and update with waveforms of disjoint support can be performed asynchronously and therefore is parallelizable. This can be thought as an instance of the standard graph coloring problem (see, e.g., [6]). For every vortex orientation, practically half of the atoms have disjoint support, allowing the above correlations to be performed in only two steps, assuming a sufficient number of available cores (see Figure 2).

2.6 Formalism through Linear Transformations

In SMP, the ℓ_2 norm of the residual is progressively reduced, leading to a least-squares approximation for a sufficiently large number of iterations. In the following, we use subscript U to indicate the *unfiltered* data and DF meaning *divergence-free*. Following standard least-squares, we have:

$$\boldsymbol{\alpha}_{DF} = (\mathbf{W}^T\mathbf{W})^{-1}\mathbf{W}^T\mathbf{f}_U \quad \text{and} \quad \mathbf{f}_{DF} = \mathbf{W}\boldsymbol{\alpha}_{DF} = \mathbf{W}\left(\mathbf{W}^T\mathbf{W}\right)^{-1}\mathbf{W}^T\mathbf{f}_U \tag{7}$$

Existence of the inverse of $\mathbf{W}^T\mathbf{W}$ is a key factor for the solution of the least-squares problem above and translates in the conditions that $\mathbf{W}$ needs to be full column rank. A closer look at the structure of $\mathbf{W}$ reveals that it is a rectangular matrix with more columns (the number of vortexes for edges along the three coordinate axis) than rows (the total number of faces). Moreover, as a result of the expected redundancy of divergence-free waveforms, $\mathbf{W}$ is neither full column rank nor full row rank.

Complications resulting from the singularity of $\mathbf{W}^T\mathbf{W}$ can be overcome by introducing the Moore-Penrose pseudo-inverse (see, e.g., [2]) $\mathbf{W}^+$ for the matrix $\mathbf{W}$. In this case, the expression for our linear filter becomes:

$$\mathbf{f}_{DF} = \mathbf{W}\mathbf{W}^+\mathbf{f}_U \tag{8}$$

2.7 Treatment of boundary conditions

The SMP filter developed in the previous Sections, is not aware of boundary conditions or the presence of solid walls. As a result, non-zero velocities might be found at locations where no flux is expected. Note that information related to location of solid walls or boundaries is generally available for experiments. A simple but generally inaccurate approach would consist in simply thresholding the filtered velocities at those locations, with the result of generating a spurious divergence at the interface between regions with and without flow. A better way to proceed is suggested by the linear nature of the filter algorithm. If we define Ω as the full acquisition grid, we can identify as Ω_S the subset of grid points located within solid walls. The vector $\mathbf{v}_S$ thus defines the restriction of vector $\mathbf{v}$ to Ω_S. We now apply our filter to $\mathbf{v}_S$, resulting in the face vector $\mathbf{f}_{DF,S}$. A filtered normal face flux vector compliant with the boundary conditions is therefore:

$$\mathbf{f}^*_{DF} = \mathbf{f}_{DF} - \mathbf{f}_{DF,S} \tag{9}$$

Note that $\mathbf{f}_{DF}$ is globally divergence-free and so it is also within Ω_S. This means that face fluxes inside solid walls will be closely approximated by the SMP filter restricted to Ω_S and they will therefore vanish after subtraction in (9). Figure 3 shows the typical steps when applying a boundary conditions-aware filter. In this case the initial flow has the non-physical feature of being directed straight towards the wall. The filtered field instead goes around the obstacle, consistent with the continuity constraint.

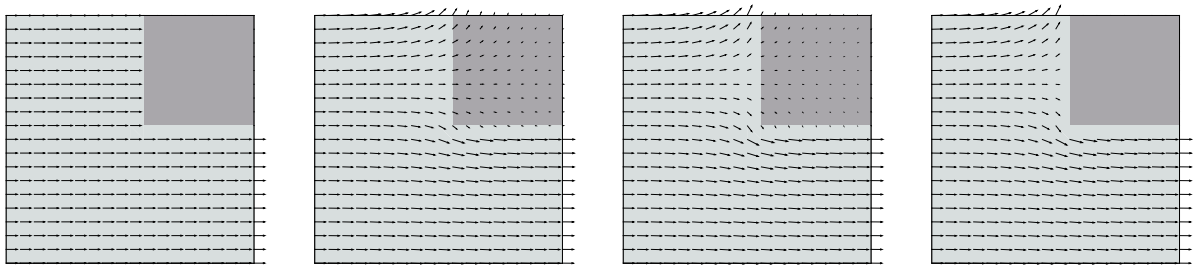

Figure 3: Typical filter operation. The original measurements (a) are processed using the SMP filter (b). A new application of the same procedure restricted to Ω_S (here represented by the darker region) remove most of the velocity component in the solid region (c). The final threshold operation (d) minimizes the amount of generated divergence.

2.8 Comparison with existing methods

For the sake of performing comparisons and better understanding the properties of the proposed methodology, we introduce a reference approach which is easy to implement and does not perform any smoothing in the acquired data. We refer to it as *isotropic divergence approach* (IDA). The IDA algorithm is illustrated in Figure 4 for the case of a regular measurement grid.

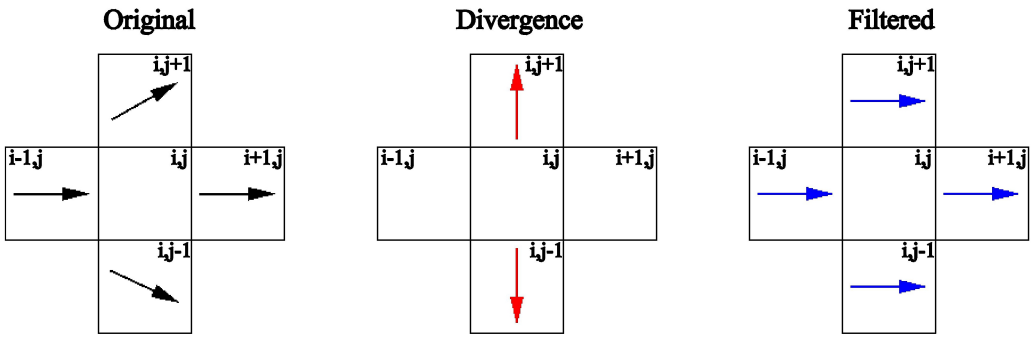

Figure 4: IDA approach applied to a simple two-dimensional case. Given a cell-centered velocity field (left) the divergence is evaluated by subtraction of the X,Y and Z components of the velocity, respectively (center), and re-assigned isotropically (right). The procedure is then iterated for all internal cells.

The algorithm iteratively *smears* the divergence of internal cells by isotropic alterations of the neighbor velocities.

An upper bound in noise reduction for the IDA filter can be estimated as follows. Let us assume that only one velocity component is modified to restore the solenoidal field. The remaining two components will retain their noise, while the third will be completely determined by satisfying the divergence-free requirement. As a result, the total variance will be diminished by one-third, leading to a reduction of $\sqrt{2/3} = 0.8165$ in terms of standard deviation. In other words, only about 18% of the noise will be removed. We will show later that the SMP filter has a superior performance.

3. Application to 2D and 3D fields

The proposed procedure is applied to fluxes computed by solving the diffusion equation with a Galerkin residual formulation on linear triangles for a 2D flow through a set of baffles. Results are shown in terms of streamlines generated by integrating trajectories from the inflow using the computed node-velocities with a fixed time-step (see Figure 5). The streamlines start from the left edge and, given the imposed pressure difference, should all exit from the right edge. Only a few streamlines are able to follow the correct velocity path in the Galerkin solution, as the velocity field is not locally conservative. Application of the SMP filter significantly improves streamline patterns, allowing all of them to reach the exit boundary.

The SMP filter is tested on volumetric PIV measurements of the flow in the wake of a swimming moon jellyfish (Aurelia aurita), see Figure 6. The data were collected using the V3V^{TM} system from TSI (Pereira et al.). It is a 3-aperture imaging technique which measures the 3D-3C velocity field instantaneously. The resolution was about 1.7 mm in all directions. Measurements were kindly provided by Brad Gemmell, (Marine Biological Laboratory, Woods Hole, MA). The spurious divergence in the original data set (see Figure 6) is removed completely by the application of the filter, and the noise level is reduced, accordingly.

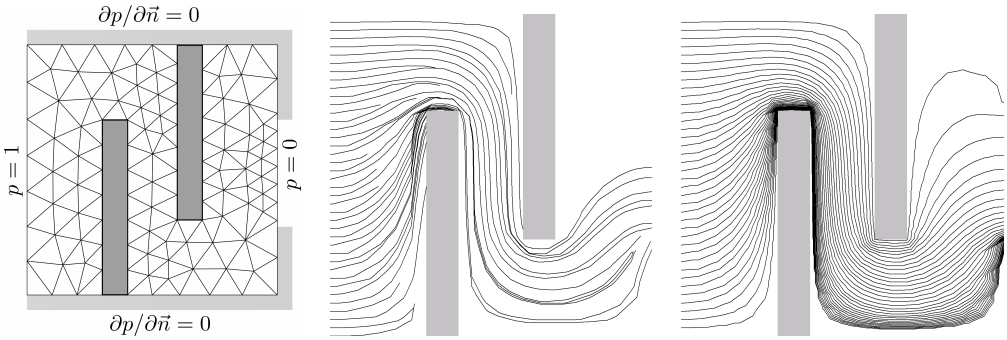

Figure 5: The velocity field generated using a Galerkin residual functional on a discretization with linear triangular finite elements has been filtered with the proposed procedure. (a) Mesh and boundary conditions are shown, together with (b) a visualization of the streamlines before and (c) after application of the SMP filter.

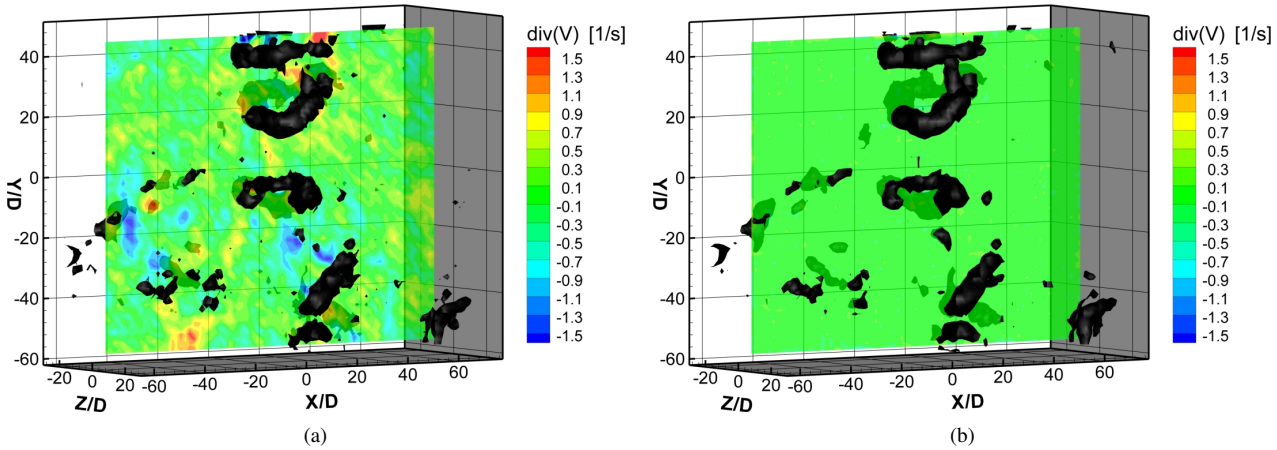

Figure 6: Volumetric PIV measurements of the flow behind a swimming jellyfish. The divergence of velocity is shown together with iso-surfaces in black identifying vortical structures based on the λ_2 criterion: original data set (a) and filtered data (b). The largest flow structures, e.g. the vortex rings, remain well represented on the filtered field and many of the small features are removed.

In order to characterize the modification to the above experimental velocity fields imposed by the SMP filter, Figure 7 displays the PDF of the difference between filtered and initial fields: the modifications are typically within 5% of the characteristic velocity. This shows that the applied correction are well within the uncertainty of the original measurements.

4. Noise reduction assessment for the SMP filter

The representation of the filter in terms of linear transformation derived in Sections 2.3 and 2.6 allows us to study its expected noise reduction properties. This estimate is verified by the application on known fields obtained from CFD simulations.

4.1 Analysis of denoising and numerical verification

Consider a random vector $\mathbf{y} \in \mathbb{R}^d$, where each y_i, $i = 1,\ldots,d$ is a random variable with distribution $\mathcal{N}(0,\sigma^2)$. The components of $\mathbf{y}$ are uncorrelated, i.e., with diagonal covariance matrix $cov(y_1,\ldots,y_d) = \sigma^2 I_d$, where I_d is the identity matrix of order d. If a new random vector $\mathbf{x}$ is derived from $\mathbf{y}$ by means of a linear transformation, it follows that:

$$\mathbf{x} = \mathbf{a}_0 + \mathbf{A}\mathbf{y} \quad \text{and} \quad cov(x_1,\ldots,x_d) = \mathbf{A}\, cov(y_1,\ldots,y_d)\,\mathbf{A}^T = \sigma^2 \mathbf{A}\mathbf{A}^T \tag{10}$$

The proposed filter can be considered as a sequence of encoding, SMP projection and decoding operations (see Section 2.3 and 2.6) and expressed as follows:

$$\mathbf{v}_{DF} = \mathbf{D}\mathbf{W}\mathbf{W}^{+}\mathbf{E}\mathbf{v}_U \tag{11}$$

If we therefore assume all components of $\mathbf{v}_U$ to be distributed as a standard Gaussian (zero mean, unit variance), the covariances resulting from encoding-decoding, projection and the complete process, respectively, are:

$$cov(\mathbf{v}_{ED}) = \mathbf{D}\mathbf{E}\mathbf{E}^T\mathbf{D}^T, \quad cov(\mathbf{f}_{DF}) = \mathbf{W}\mathbf{W}^{+}\mathbf{W}^{+,T}\mathbf{W}^T, \quad cov(\mathbf{v}_{DF}) = \mathbf{D}\mathbf{W}\mathbf{W}^{+}\mathbf{E}\mathbf{E}^T\mathbf{W}^{+,T}\mathbf{W}^T\mathbf{D}^T. \tag{12}$$

As a result of the linear transformation, the final velocity (or face flux) vectors contain correlated components; the square roots of the eigenvalues for the transformation matrices in (12) therefore provide the amplification/reduction factors for a noise pattern associated to the latter modes. Figure 8 shows the ordered square-roots of the eigenvalues obtained from the filter associated with a $5 \times 5 \times 5$

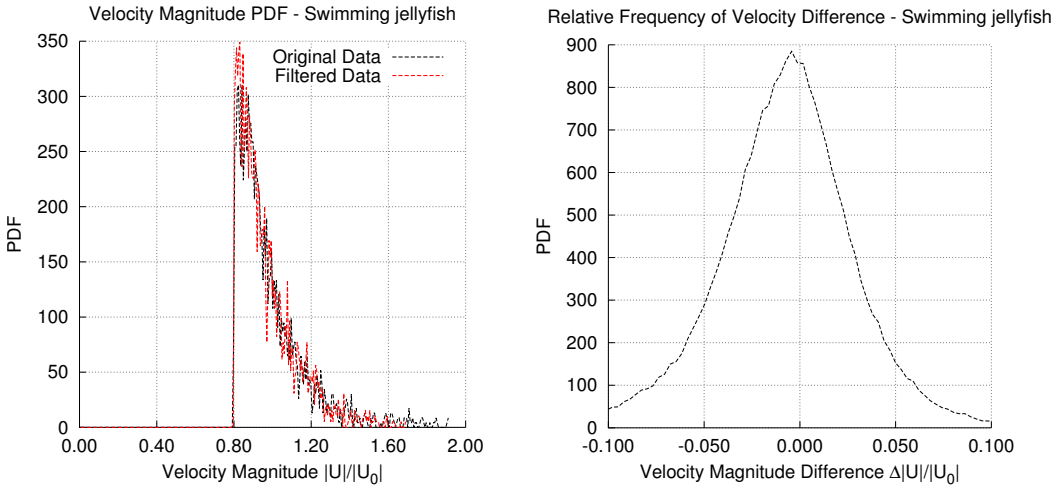

(a) (b)

Figure 7: PDF of velocity magnitudes (a) and velocity magnitude differences (b) for the three-dimensional swimming jellyfish dataset.

cubic region. The total number of velocity measurements for this case is 125×3, resulting in 375 possible modes and 450 faces. Two different geometries were used: cubic cells (isotropic case) or prismatic cells of edge lengths proportional to $\{1,2,3\}$, respectively (anisotropic case) as illustrated in Figure 9a and 9b. The graph on the left of Figure 8 shows modal noise reductions referred only to the encoding-decoding operations. As a reduction at the boundaries is always produced by encoding constant noise patterns, all the ordinates in the graph are strictly less than 1.0. No mode is completely removed (zero eigenvalue) as cancellation never occurs at boundary faces. The SMP projection operator is also analyzed in the center graph in Figure 8; it can be seen how noise patterns are either left unaltered or completely removed. In other words, if a noise face flux pattern is divergence-free, then its energy is not modified by the SMP projection, as expected. The modal noise reduction factors for the full filter (combination of encoding-decoding and SMP projection) are also illustrated on the right. This graph provides information on the stability of the implemented filter, meaning that no significant amplification can be noticed for noise patterns, both for the isotropic and anisotropic configurations. Another important indicator is the average noise reduction; this value can be used to provide estimates of the expected noise reduction, under the assumption that a broad range of modes are contained in the flow of interest. A value of approximatively 0.5077 is found; we use this estimate as a reference for the numerical experiments performed in the following sections.

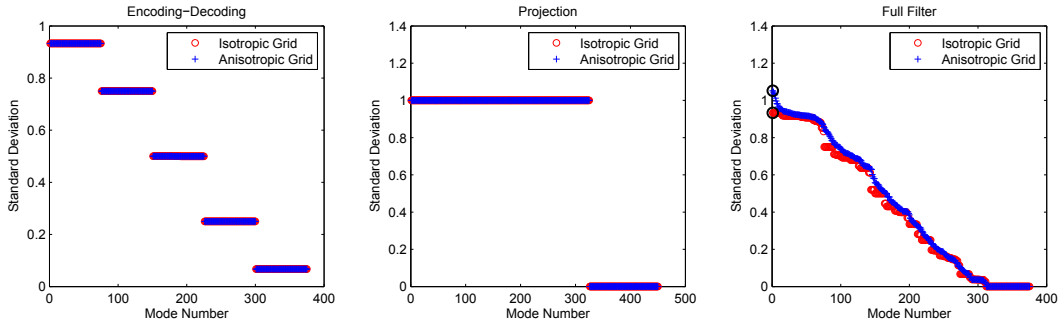

Figure 8: Standard Deviations of uncorrelated random variables as resulting from encoding-decoding operations (a), SMP projection (b) and application of the full filter, from measured velocities at grid points to divergence-free velocities at the same locations.

4.2 Noise reduction assessment

Practical noise reduction is assessed using the velocity field obtained by direct numerical simulation (DNS) of the flow around a wall-mounted cube [9]. The incompressible Navier-Stokes equations are integrated using an unstructured finite-volume scheme. The Reynolds number based on the free-stream velocity and the cube height h is 5000. A domain of size $19h \times 10h \times 10h$ is discretized with a grid of 12.6 millions hexahedral elements. The numerical set-up is described in detail in [9]. Here one instantaneous realization of the time-varying turbulent flow is considered, and is visualized in Figure 9.

We consider three cases in which we add Gaussian noise with zero average and increasing standard deviation equal to the 15%, 30%, 45% of the maximum velocity magnitude. We then apply the SMP filter and compare with results generated by the IDA filter. The noise reduction estimates provided in the above Sections are also reported, and found in good agreement with the actual filtered data for both flow configurations. Overall, the SMP filter removes about 50% of the noise, while the IDA filter removes about 17% of the noise, as shown in Figure 10.

It also is important to test the ability of the SMP approach to resolve localized features without alterations or distortions. The Hill spherical vortex [3] is chosen for this purpose. Figure 11 shows a negligible amount of smoothing (e.g., at the boundaries), which

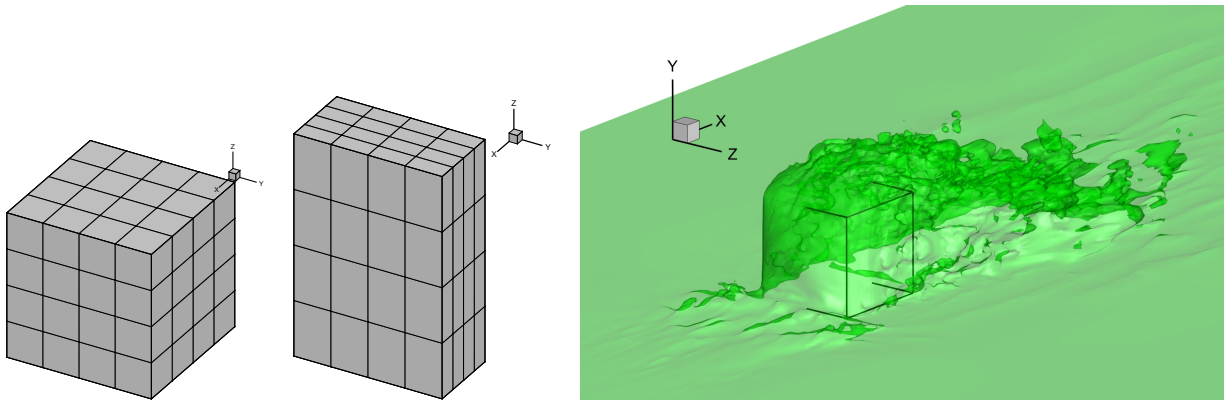

(a) (b) (c)

Figure 9: Isotropic (a) and anisotropic (b) measurement grids employed for de-noising assessment. Direct numerical simulation of the flow around a wall-mounted cube (c). A streamwise velocity isosurface is shown which corresponds to 30% of the free stream value.

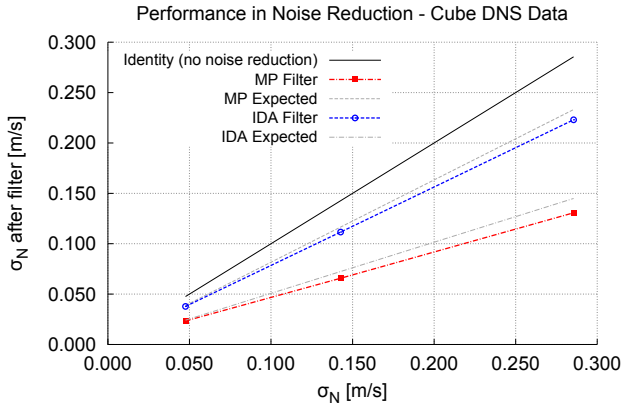

Figure 10: Noise reduction results from DNS of the flow around a wall mounted cube. The values of noise reductions obtained in the numerical experiments agree well with their expected values.

occurs as a result of the finite grid representation of the spherical vortex.

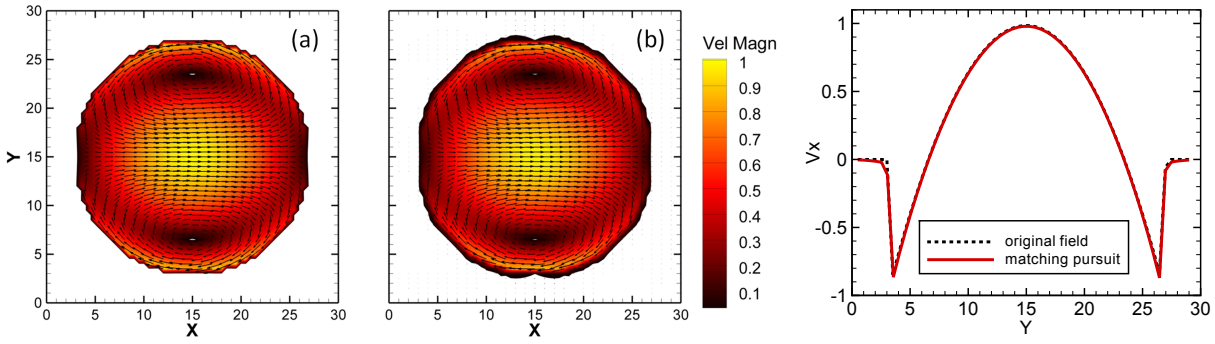

(a) (b)

Figure 11: Capabilities of the proposed filter in resolving localized flow features. The analytical flow field generated by a spherical vortex (see [3]) is mapped to an acquisition grid of reasonable size (a). The application of the SMP filter (b) provides almost no alteration of the localized vortex structure and very limited changes can also be observed at the boundary.

5. Conclusion and future work

This study proposes a methodology based on waveform matching to iteratively project measured velocity fields onto a divergence-free space. The algorithm, which can be represented as a linear operator, acts as a filter in the sense that it removes the non-solenoidal part of the noise. The method proves effective in removing the spurious divergence from 3D volumetric velocity fields obtained by Particle

Imaging Velocimetry. It is demonstrated analytically (and confirmed by numerical experiments) that the present approach removes about 50% of the noise. By contrast, an iterative redistribution of the velocity to the neighboring cells removes only about 17% of the noise. The application of the filter does not result in artificial smoothing, allowing it to reproduce localized features of the flow with high accuracy. Further research efforts will focus on using this type of filter as a basis for computing the relative pressure from instantaneous and time-averaged velocity acquisitions by integration of the Navier-Stokes equations.

Acknowledgments

We are indebted to Dr. Bernd Wieneke for suggesting the IDA method as a reference approach and providing insights on its de-noising properties. We wish to thank Riccardo Rossi at Università di Bologna for providing the direct simulation of the flow around the wall-mounted cube. PIV 3D measurements of the flow around a swimming jellyfish were kindly provided by Brad Gemmell, Marine Biological Laboratory, Woods Hole, MA.

G.I.'s and D.S.'s work is supported under Subcontract No. B597952, with Lawrence Livermore National Security under Prime Contract No. DE-AC52-07NA27344 from the Department of Energy National Nuclear Security Administration for the management and operation of the Lawrence Livermore National Laboratory. F.C.'s work is supported by a grant from the Honeywell Corporation.

REFERENCES

[1] G. E. Elsinga, F. Scarano, B. Wieneke, and BW Van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41:933–947, 2006.

[2] G. Golub and W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *Journal of the Society for Industrial & Applied Mathematics, Series B: Numerical Analysis*, 2(2):205–224, 1965.

[3] M.J.M. Hill. On a Spherical Vortex. *Proceedings of the Royal Society of London*, 55(331-335):219–224, 1894.

[4] K.D. Hinsch and S.F. Herrmann. Holographic particle image velocimetry. *Measurement Science and Technology*, 15, 2004.

[5] T. Hori and J. Sakakibara. High-speed scanning stereoscopic PIV for 3D vorticity measurement in liquids. *Measurement Science and Technology*, 15:1067, 2004.

[6] Enrico Malaguti and Paolo Toth. A survey on vertex coloring problems. *International Transactions in Operational Research*, 17(1):1–34, 2010.

[7] S.G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41:3397–3415, 1993.

[8] F. Pereira, M. Gharib, D. Dabiri, and D. Modarress. Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. Application to bubbly flows. *Experiments in Fluids*, 29:78–84, 2000.

[9] R. Rossi, D.A. Philips, and G. Iaccarino. A numerical study of scalar dispersion downstream of a wall-mounted cube using direct simulations and algebraic flux models. *International Journal of Heat and Fluid Flow*, 31(5):805–819, 2010.

[10] M. Virant and T. Dracos. 3D PTV and its application on lagrangian motion. *Measurement Science and Technology*, 8:1539, 1999.