



Compact representation of wall-bounded turbulence using compressive sampling

J.-L. Bourguignon, J. A. Tropp, A. S. Sharma, and B. J. McKeon

Citation: *Physics of Fluids* (1994-present) **26**, 015109 (2014); doi: 10.1063/1.4862303

View online: <http://dx.doi.org/10.1063/1.4862303>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/pof2/26/1?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Power-law versus log-law in wall-bounded turbulence: A large-eddy simulation perspective](#)

Phys. Fluids **26**, 011703 (2014); 10.1063/1.4862919

[How linear is wall-bounded turbulence?](#)

Phys. Fluids **25**, 110814 (2013); 10.1063/1.4819081

[Constrained large-eddy simulation of wall-bounded compressible turbulent flows](#)

Phys. Fluids **25**, 106102 (2013); 10.1063/1.4824393

[Stokes number effects on particle slip velocity in wall-bounded turbulence and implications for dispersion models](#)

Phys. Fluids **24**, 021705 (2012); 10.1063/1.3690071

[Generation of synthetic turbulent inflow data for large eddy simulation of spatially evolving wall-bounded flows](#)

Phys. Fluids **21**, 045103 (2009); 10.1063/1.3103881



Compact representation of wall-bounded turbulence using compressive sampling

J.-L. Bourguignon,¹ J. A. Tropp,¹ A. S. Sharma,² and B. J. McKeon¹

¹*Division of Engineering and Applied Science, California Institute of Technology, Pasadena, California 91125, USA*

²*Aerodynamics and Flight Mechanics, Faculty of Engineering and the Environment, University of Southampton, Southampton SO17 1BJ, United Kingdom*

(Received 6 March 2013; accepted 20 December 2013; published online 22 January 2014)

Compressive sampling is well-known to be a useful tool used to resolve the energetic content of signals that admit a sparse representation. The broadband temporal spectrum acquired from point measurements in wall-bounded turbulence has precluded the prior use of compressive sampling in this kind of flow, however it is shown here that the frequency content of flow fields that have been Fourier transformed in the homogeneous spatial (wall-parallel) directions is approximately sparse, giving rise to a compact representation of the velocity field. As such, compressive sampling is an ideal tool for reducing the amount of information required to approximate the velocity field. Further, success of the compressive sampling approach provides strong evidence that this representation is both physically meaningful and indicative of special properties of wall turbulence. Another advantage of compressive sampling over periodic sampling becomes evident at high Reynolds numbers, since the number of samples required to resolve a given bandwidth with compressive sampling scales as the logarithm of the dynamically significant bandwidth instead of linearly for periodic sampling. The combination of the Fourier decomposition in the wall-parallel directions, the approximate sparsity in frequency, and empirical bounds on the convection velocity leads to a compact representation of an otherwise broadband distribution of energy in the space defined by streamwise and spanwise wavenumber, frequency, and wall-normal location. The data storage requirements for reconstruction of the full field using compressive sampling are shown to be significantly less than for periodic sampling, in which the Nyquist criterion limits the maximum frequency that can be resolved. Conversely, compressive sampling maximizes the frequency range that can be recovered if the number of samples is limited, resolving frequencies up to several times higher than the mean sampling rate. It is proposed that the approximate sparsity in frequency and the corresponding structure in the spatial domain can be exploited to design simulation schemes for canonical wall turbulence with significantly reduced computational expense compared with current techniques. © 2014 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4862303>]

I. INTRODUCTION

Compressive sampling, which employs optimization-based algorithms to reconstruct a signal from a set of measurements, presents a means of circumventing Nyquist limitations on sample size and thus reducing the amount of data required for signal reconstruction. This approach is valid only if the measurement signals are sparse with respect to a known orthogonal basis. The technique has been applied, for example, in magnetic resonance imaging (MRI),¹ also for MRI applied to the interrogation of bubble-flow interactions in multiphase flows, Tayler *et al.*² Other authors have also successfully used compressive sampling techniques in fluid dynamics, e.g., Tu and Rowley,³ Jovanovic, Schmid, and Nichols,⁴ Jovanovic and Schmid,⁵ however the application of compressive sampling in turbulence modeling has been thwarted because the broadband spectrum that is observed in

wall-bounded turbulence based on spatially localized spectral measurements lacks the requisite signal sparsity.

Frequency domain (temporal) analysis of data is not often attempted or explored in either experiment or simulation because of the associated requirements for the storage of large time-resolved datasets. The frequency spectrum of wall-bounded turbulence is known to be broadband when measured locally in space, with the local bandwidth dependent on the wall normal location and the Reynolds number. For this reason, it is necessary to acquire many samples to resolve the frequency content of the flow so as to satisfy the Nyquist criterion. Most of the samples are needed just to resolve the high-frequency, low-energy modes of the flow so they do not alias onto the lower frequency modes.

When a sinusoidal signal of frequency f_0 is sampled periodically at a rate f_s , each sinusoid with frequency $|f_0 - Nf_s|$, where N is an integer, also interpolates the samples. These spurious modes are called aliases of the fundamental signal. Bilinski⁶ suggests that aliasing effects can be eliminated by sampling randomly in time because the aliases do not interpolate the randomized samples. As a consequence, randomized sampling seems to offer a way to reduce the number of samples required for perfect reconstruction of the signal below the rate predicted by the Nyquist criterion.

To extract frequency content from randomized samples, additional care becomes necessary. If a Fast Fourier Transform (FFT) is used, then fuzzy aliasing will occur (i.e., aliasing that is more diffuse than in the case of periodic sampling) because the Fourier basis functions are not orthogonal on the set of randomized samples. Bilinski⁶ describes several techniques to remove, or at least attenuate, fuzzy aliasing when a FFT is applied to nonperiodic samples.

In recent work,⁷⁻⁹ we have demonstrated that key features of the fluctuating velocity in fully developed wall turbulence can be captured by approximating the transfer function (or resolvent) at each wavenumber/frequency combination, i.e., under a triple Fourier decomposition in the homogeneous directions, between forcing arising from the interaction between different scales, and the velocity response. A basis for the wall-normal direction can be obtained by singular value decomposition of the resolvent, where the singular functions (or resolvent modes) represent the most amplified velocity response to the “most dangerous” input forcing and are ordered by the magnitude of the amplification arising from the forcing. As such, a low-rank approximation can be obtained by retaining a limited number of singular functions; recognizable statistical and structural features of wall turbulence can be identified in even the rank-1 approximation, in which only the principal singular functions are investigated.⁹

For canonical turbulent pipe and channel flows, the output of this analysis at each (k, n, ω) , where k and n are the streamwise and spanwise wavenumbers, respectively, and ω is the frequency, consists of a three-dimensional, three velocity component propagating wave which we hypothesize can be considered to represent a building block of wall turbulence. These are deemed *resolvent modes*; they represent the most amplified disturbance. This model provides a natural order reduction by emphasizing only the amplified resolvent modes with wall normal coherence that allows matching of the fluctuation field and neglecting those that contribute a negligible amount.¹⁰ It also suggests storing only (complex) coefficients for the wall-normal variation of the Fourier mode amplitudes over a range of wavenumber/frequency combinations instead of a long series of instantaneous direct numerical simulations (DNS) velocity fields. In theory, given the correct amplitudes and phases for the coefficients, both are complete representations of the velocity field, but the wavenumber-frequency combination is extremely compact: the velocity field at any instance in space and time, \mathbf{x}_0 and t_0 , can be obtained by simply evaluating the linear sum of Fourier modes at that \mathbf{x}_0 and t_0 , i.e., advancing a set of waves that propagate relative to each other, as determined by their relative streamwise wavespeeds, \mathbf{k} . While a powerful reduction of information already, this investigation of the application of compressive sampling reveals that not all coefficients are active for a given wavenumber combinations, such that the array of coefficients is not fully populated.

Theory and experiment motivated the scientific hypothesis that the velocity field admits a sparse representation after an appropriate transformation into the frequency domain, namely, that the number of energetic frequencies among all resolvent modes with a given (k, n) is limited. Under this hypothesis, compressive sampling is an ideal tool for reducing the amount of information required to approximate the velocity field. Indeed, one of the early papers¹¹ on compressive sampling shows

that random samples in time suffice to acquire a frequency-sparse signal, making this a potentially useful tool to formulate a compact representation of the velocity field.

However, we can infer a stronger result from the effectiveness of compressed sensing in this application: the success is itself serves as a validation of the hypothesis that the velocity field admits a sparse representation in frequency. The mathematical signal processing literature establishes rigorously that sparsity does not occur by accident. First of all, a generic signal never has a sparse representation.¹² Furthermore, when a signal does admit a representation as a sum of a small number of frequency modes, every other representation requires a large number of frequency modes.^{13–16}

In this paper, it is shown using DNS of turbulent pipe flow that when wall-bounded turbulence is Fourier decomposed in the homogeneous spatial (wall-parallel) directions, and in time, as explored by McKeon and Sharma,⁷ a compact representation of the velocity field can be found by applying compressive sampling to identify the energetically dominant frequencies of the most energetic spatial Fourier modes, because their frequency content turns out to be approximately sparse,¹⁷ in the sense that the signal is close (in the L_2 sense) to an exactly sparse signal. The fact that compressed sensing works, i.e., we can approximate the velocity field with a small number of frequencies, provides strong evidence that this representation is physically meaningful and indicative of special properties of wall turbulence. The use of compressive sampling in wall-bounded turbulence is demonstrated by a series of tests, based on synthetic velocity fields with known frequency content, aimed at verifying the output from the compressive sampling routine. The data storage requirements for reconstruction of the full field using compressive sampling are shown to be significantly less than for periodic sampling, in which the Nyquist criterion limits the maximum frequency that can be resolved; conversely, compressive sampling maximizes the frequency range that can be recovered if the number of samples is limited, resolving frequencies up to several times higher than the mean sampling rate. The paper is organized as follows. An introduction to our specific implementation of compressive sampling and the datasets to which it was applied is given in Sec. II, results are summarized in Sec. III, before conclusions are drawn in Sec. IV.

II. METHODOLOGY

Two different types of data sets were used to investigate the application of compressive sampling to wall-bounded turbulence, specifically for the pipe flow geometry shown in Figure 1. Frequency analysis was performed to confirm that the frequency content is approximately sparse and to validate the use of both periodically and randomly sampled DNS data. The details of the compressive sampling approach and the data sets used in this study are described in Secs. II A–II C.

A. Problem formulation

In all cases, compressive sampling was applied to the Fourier-transformed three-dimensional (spatial) velocity fields in order to reconstruct their temporal variation. It is well-known^{18,19} that the Fourier series decomposition in the homogeneous (spatial and temporal) directions is optimal in the sense that it maximizes the turbulent kinetic energy captured for a given number of basis




FIG. 1. Schematic of pipe geometry and nomenclature.

functions; the Fourier modes are known to be the (orthogonal) eigenmodes of the linear Navier-Stokes (NS) equations in the homogeneous directions. We focus on the efficiency, or compactness, of the reconstruction of the streamwise velocity component, i.e., the reduction in storage requirements, here; subsequent work will investigate the three-dimensional velocity field and offer an interpretation into the characterization of wall turbulence that can be obtained using the compressive sampling technique. As such, we consider here only the streamwise velocity, u , by way of example and denote the coefficients of the 2D Fourier series representation in the streamwise and azimuthal directions by $c_{k,n}(r, t)$, henceforth referred to as the 2D spatial Fourier modes and indexed by (k, n) . The full streamwise velocity field, u , is then given by

$$u(x, r, \theta, t) = \sum_{k,n} c_{k,n}(r, t) e^{i(kx+n\theta)}, \quad (1)$$

normalized such that

$$\sum_k \sum_n c_{k,n}(r, t) c_{k,n}^*(r, t) = u^2(r, t). \quad (2)$$

Here, $*$ denotes the complex conjugate. Only the positive k half plane is retained since the spectrum of a real-valued signal is symmetric. The 2D Fourier spectra were integrated in the wall-normal direction and averaged in time to identify the 2D Fourier modes that contribute most to the streamwise turbulence intensity, effectively a crude Proper Orthogonal Decomposition (POD) over a limited resolvable parameter space.

After Fourier decomposition in the homogeneous directions (the optimal choice in the L_2 sense as described above with respect to Eq. (1)), the full vector velocity field $\mathbf{u}(x, r, \theta, t)$ can be written in terms of streamwise and azimuthal wavenumbers (k, n) and angular frequencies $\omega = 2\pi f$ as follows:

$$\mathbf{u}(x, r, \theta, t) = \sum_{k,n,\omega} \mathbf{c}_{k,n,\omega}(r) e^{i(kx+n\theta-\omega t)}, \quad (3)$$

which can be interpreted⁷ as a summation of propagating waves with relative magnitude and phase given by the complex-valued coefficients $c_{k,n,\omega}$. The radial distribution of momentum associated with each (k, n, f) combination is given by $\mathbf{c}_{k,n,\omega}(r)$. Here, the wavenumbers and frequencies are non-dimensionalized with the pipe radius, R , and centerline velocity, U_{cl} .

B. Compressive sampling

The theory of compressive sensing (Candés *et al.*¹¹) suggests that for sparse signals that are sampled randomly in time, the solution of a constrained convex optimization problem to recover sparse energetic content is a computationally efficient problem. For the Fourier coefficient representation given in Eq. (3), the formal optimization problem for the recovery of sparse energetic content, namely, a few large Fourier coefficients for a given (k, n) , is

Minimize over all $c_{k,n,2\pi jdf}$

$$\sum_{j=1}^{N_{opti}} |c_{k,n,2\pi jdf}(r)|, \quad (4)$$

for each r , under the constraint

$$\sum_{j=1}^{N_{opti}} c_{k,n,2\pi jdf}(r) e^{i2\pi jdf t_s} = c_{k,n}(r, t_s), \quad \forall t_s, \quad (5)$$

where N_{opti} and df are the number of samples and the frequency resolution, respectively. Since this corresponds to optimization at each radial grid point, we term this a ‘‘local optimization.’’

Alternatively, a ‘‘global optimization’’ problem may be formulated based on the minimization of the ℓ_2 -(energy) norm of the coefficients of the basis functions (the temporal Fourier

coefficients), i.e.,

Minimize

$$\sum_{j=1}^{N_{opti}} \int_0^1 c_{k,n,2\pi jdf}(r) c_{k,n,2\pi jdf}^*(r) r dr, \quad (6)$$

under the constraint

$$\sum_{j=1}^{N_{opti}} c_{k,n,2\pi jdf}(r) e^{i2\pi jdf t_s} = c_{k,n}(r, t_s), \quad \forall t_s. \quad (7)$$

The latter method only needs to be applied once, instead of at each wall-normal location separately, and tends to select basis functions (frequencies) that are energetic over a wide range of wall-normal locations. However, the optimization algorithm converges significantly slower due to the significant increase in the number of constraints from N_s (the number of samples) to $N_s \times N_r$ (the number of samples times the number of grid points in the radial direction).

Constrained optimization problems such as those in Eqs. (4)–(6) will identify a sparse representation of the signal, i.e., a few energetic temporal Fourier modes, provided that such a representation exists and that sufficient samples have been acquired. The number N_s of samples required for perfect reconstruction can be estimated using an empirical relationship obtained in Tropp *et al.*,²⁰ henceforth called the sparsity relationship,

$$N_s \geq 1.7 K \log \left(\frac{\Omega}{K} + 1 \right), \quad (8)$$

where K is the number of energetic frequencies in the approximately sparse signal and Ω is the bandwidth, or ratio of the highest to lower frequency of interest. Evidently, the number of samples scales as the logarithm of the input signal bandwidth, in contrast to the linear scaling associated with periodic sampling. Alternatively, the lower the number of energetic frequencies present in the input signal the broader the frequency range that can be resolved with compressive sampling. The sparsity relationship has been justified theoretically by other studies, including Candés *et al.*,¹¹ Donoho²¹ and the analysis in Tropp *et al.*,²⁰ however, the numerical constant 1.7 results from a linear regression on experimental data in Ref. 20, and appears to be problem specific.

In practice, application of the compressive sampling technique to a randomly sampled signal requires selection of the input frequency range and resolution for the constrained minimization problem, $[df, N_{opti} \times df]$, such that the sparsity relationship is satisfied and the frequency resolution appropriate to capturing the energetic content in the signal is obtained. Either the number and bandwidth of the energetic frequencies must be known, a conservative estimate must be made or iteration must be performed before applying compressive sampling. Extensive studies of synthetic velocity fields,¹⁷ not described here in the interests of brevity, were used to investigate the combinations of sampling and optimization parameters in the compressive sampling minimization problems of Eqs. (4) and (6) that lead to the correct recovery of the input frequencies for modes spanning the frequency content of the full turbulent velocity field. For simplicity, the input frequencies were chosen to be equi-spaced within the optimization frequency range, i.e., $df = const$ (but this is not a requirement of the technique).

In all cases, for correctly selected sampling and optimization parameters, the frequency range that could be resolved by compressive sampling is significantly broader than the range that could be resolved using periodic sampling with the same number of samples and sampling duration. The highest recoverable frequency was found to be about 6 times higher than the mean sampling frequency ($f_s = 0.3$). The maximum error in the recovered signal was always 1% or less.

However, in cases when either the number of input frequencies was increased or the sampling duration (which is related to the resolvable bandwidth Ω of the signal in Eq. (8)) was decreased such that the sparsity relationship (Eq. (8)) was violated, the success of the convex optimization degraded significantly, there was a “leaking” of energy into adjacent frequencies, and several erroneous frequency peaks were identified.

Note that the minimum and maximum frequencies used to reconstruct the signal can, respectively, be lower than the inverse of the sampling duration and higher than the mean sampling rate. In contrast, in the case of periodic sampling, the minimum and maximum frequencies of the Fourier series representing the input signal are directly related to the sampling parameters: the minimum frequency corresponds to the inverse of the sampling duration, and the maximum frequency equals the sampling rate for complex-valued signals and half the sampling rate for real-valued signals.

A conclusion from our studies reported here and in Bourguignon¹⁷ is that the compressive sampling technique seemed to perform slightly better than predicted by the sparsity relationship of Eq. (8), i.e., the constant takes a value lower than 1.7. One possible explanation for this improvement is differing requirements for a successful optimization: the amplitude of the coefficients was required to contain at least 10% of the peak frequency energy in our studies, whereas matching of the input signal up to machine precision was used by Tropp *et al.*²⁰ to determine a multiplicative constant of 1.7.

In this work, the minimization problems of Eqs. (4) and (6) were solved with Matlab using the CVX toolbox for convex optimization;^{22,23} an example of the code used is shown in the Appendix. Results from each of the two optimization approaches are compared in Sec. III B. Sparse frequencies were identified as those for which the amplitude of the corresponding coefficients contained at least 10% of the peak frequency energy. Other choices could be made, for example, requiring matching of the input signal up to machine precision.²⁰ The frequency spectrum was subsequently integrated in the wall-normal direction to identify the dominant frequencies, defined as the frequencies containing not less than 10% of the energy in the peak frequency.

C. DNS velocity fields

The DNS velocity fields were kindly provided by Wu²⁴ who used the second-order finite difference code described in Wu and Moin²⁴ at $Re = 24\,580$, or $Re_\tau = Ru_\tau/\nu = 685$ where $u_\tau = \sqrt{\tau_w/\rho}$, τ_w is the mean wall shear stress, ρ is the fluid density, and ν is the kinematic viscosity, and with a domain length of $30R$ ($k_{min} = 0.21$). The number of grid points in the streamwise, wall-normal, and azimuthal directions was, respectively, $2048 \times 256 \times 1024$. The DNS data were subsampled by a factor of 4 in the streamwise and azimuthal directions to decrease the size of the data files. The flow domain is large enough for the spatially-averaged velocity profile and streamwise turbulence intensity profile to be converged, even with the subsampling. Three sets of velocity fields were studied: one sampled under a temporally periodic scheme, one with randomized sampling and one containing equivalent length records under both types of sampling.

For the first dataset, the flow was periodically sampled in time at a rate of 1 sample every 7.2 dimensionless time units, based on the pipe radius and bulk velocity. A total of 21 samples was taken over 150 dimensionless time units. The number of samples and sampling duration for the first data set was not chosen by the authors and constrained the number of 2D spatial Fourier modes whose frequency content could be resolved with the available data. The 21 periodically sampled DNS velocity fields were used to check whether the frequency content of the 2D spatial Fourier modes was indeed approximately sparse, as described in Sec. III A.

The range of correctly resolved 2D Fourier modes in the periodically sampled DNS, i.e., those modes with (k, n) resolved within a single velocity field and with frequency content free from aliasing effects across the temporal record, was estimated by comparing the frequency range that can be resolved from specific sampling parameters with empirical bounds on the frequency content of the flow (Figure 2). The dynamically significant frequency content in wall-bounded turbulence can be estimated as a function of the streamwise wavenumber under the assumption that the highest and lowest streamwise convection velocities, namely, $2\pi f/k$, correspond, respectively, to the centerline velocity and 10 times the friction velocity, u_τ . The upper bound comes from an extrapolation to turbulent flow of the analysis of the Orr-Sommerfeld equation by Joseph²⁵ which showed that the real parts of the eigenvalues, which can be interpreted as disturbance convection velocities, are restricted to the range set by the laminar base profile. The lower bound was estimated based on the near-wall experimental measurements of Morrison *et al.*,²⁶ who determined a minimum convection velocity of approximately ten times the friction velocity, $u = 10u_\tau$, or $0.44\bar{U}$ at $Re = 24\,580$. This

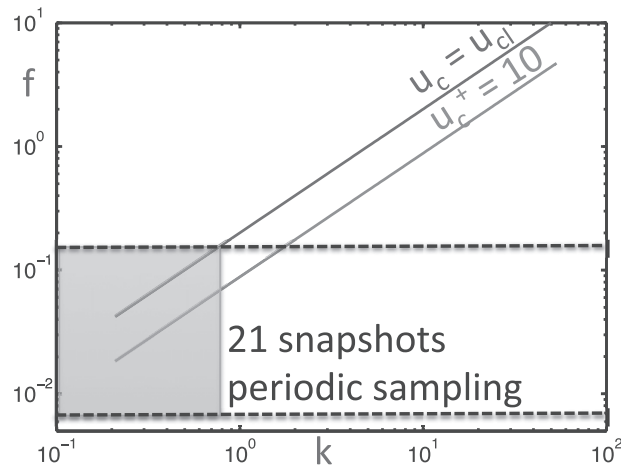


FIG. 2. Frequency range resolved by the periodically sampled DNS data (delimited by the two horizontal dashed lines) compared to the empirical upper and lower bounds on the DNS frequency content (solid lines) as a function of the streamwise wavenumber k . The shaded area shows the time-resolved streamwise wavenumber range for the available data. Notice that while the area of the graph (resolution) goes up as the product of resolution of f and k , the area between the two lines grows linearly in k . The consequence of this is that resolving only the area between the two lines of constant c is more efficient as Reynolds number increases.

value for the lowest convection velocity is broadly supported in the literature, as summarized by LeHew *et al.*²⁷

These upper and lower bounds on the frequency content as a function of the streamwise wavenumber are depicted as straight lines in Figure 2. The shaded area shows the range of wavenumber/frequency space that can be considered fully resolved. The highest frequency that could be resolved with the available DNS data is given by the inverse of the sampling rate $f_{max} = \frac{1}{7.2} = 0.14$ (recall that the frequency normalization is performed with bulk velocity); at $k = 0.69$, the modal streamwise convection is equal to the centerline velocity. Part of the frequency content of all 2D spatial Fourier modes with $k \geq 0.69$ is therefore aliased to lower frequencies, meaning that such modes cannot be considered correctly resolved within the constraints of our framework. Therefore, only the frequency content of the lowest three streamwise wavenumbers $k = 0.21, 0.42, \text{ and } 0.63$, for various azimuthal wavenumbers, can be accurately extracted from the available DNS data by applying a FFT in time, because the Nyquist criterion is satisfied for these three modes. Note that the energetic range of k can be estimated to exceed $k = 5$ by consideration of the k location corresponding to the near-wall activity with a dominant streamwise wavelength of 1000 viscous wall-units.

The second set was obtained by random temporal sampling in order to enable application of compressive sampling. A new run of the turbulent pipe flow DNS was performed by Wu,²⁴ and the velocity fields were recorded at the 50 sampling time instants shown on Figure 3, randomly distributed over 100 dimensionless time units. The number of samples, N_s and the total sample length were selected by consideration of synthetic velocity fields.¹⁷

In order to enable a fair comparison between the periodic and randomized sampling approaches, a third dataset was constructed from the first 14 periodically sampled fields from the first dataset and




FIG. 3. The 50 DNS sampling time instants randomly distributed over 100 dimensionless time units based on the radius and bulk velocity. The last sampling time instant is at $\tau = 96.57$.

14 velocity fields randomly chosen from the 50 fields available (total non-dimensional sample length of 100 dimensionless time units in each case). In this case, the minimum and maximum frequencies for the FFT are given by $f_{min} = 0.01$ and $f_{max} = 0.13$; consequently, only the 2D Fourier modes corresponding to the lowest three streamwise wavenumbers could be fully resolved (the other modes had higher frequencies aliased to the range $[0.01, 0.13]$). Similarly, for the randomly sampled series, the full frequency content of the most energetic 2D Fourier modes could be resolved for $k \leq 1.47$ instead of $k \leq 4.3$ when the 50 samples were used.

III. RESULTS AND DISCUSSION

In this section, we first confirm the approximate sparsity of the 2D Fourier modes in the frequency domain. Then we demonstrate the success of compressive sensing in randomly sampled DNS velocity fields, describe the associated compact representation of wall-bounded turbulence and provide a brief comparison of the results with those from periodically sampled DNS fields.

A. Sparsity check based on periodically sampled DNS data

An explicit check for approximate sparsity of the velocity field in the frequency domain under the 2D Fourier mode decomposition, an essential component for success of compressive sampling, was performed using the periodically sampled DNS data before proceeding to apply compressive sampling. The frequency spectrum of the most energetic 2D Fourier modes was computed as a function of the wall-normal distance by performing an FFT of the 21 samples available, in time, at each wall-normal location. Only the modes with $k \leq 0.63$ which are free from aliasing effects were considered.

The time-averaged profiles in the wall-normal direction, $y = 1 - r$, of the three most energetic 2D Fourier modes are plotted on Figure 4. The modes are tall in the wall-normal direction and extend to the centerline. A representative frequency spectrum corresponding to the mode $(0.42, 5)$, which contains significant energy but is not one of the three most energetic 2D Fourier modes and thus potentially represents a tougher test of sparsity, is plotted on Figure 5 as a function of the wall-normal distance, together with the empirical bounds on the frequency content corresponding to convection velocities equal to the centerline velocity and 10 times the friction velocity. The frequency spectrum exhibits vertical contours of constant energy, highlighting the wall-normal coherence of the energetic 2D Fourier modes, which suggests constant phase speed in the propagating wave interpretation of the modes. The frequency content of all the modes studied was always approximately sparse, exhibiting between one and three dominant frequencies; all other frequencies contain at least an order of magnitude less energy than the dominant frequency(ies). The dominant frequencies fall in between




FIG. 4. Wall-normal profile ($y = 1 - r$) of the magnitude (a) and phase (b) for the first three most energetic modes $(k, n) = (0.21, -2)$ (dotted, "x"), $(k, n) = (0.42, 3)$ (dashed, diamonds), and $(k, n) = (0.21, 2)$ (solid) of the periodically sampled DNS flow field.




FIG. 5. Representative frequency spectrum as a function of the wall-normal distance for the 2D Fourier mode $(k, n) = (0.42, 5)$ extracted from the periodically sampled DNS data using a temporal FFT, showing approximate sparsity in the frequency domain. The contour levels span an order of magnitude. The empirical upper and lower bounds on frequency corresponding to convection velocities equal to the centerline velocity and 10 times the friction velocity are demarcated by the two vertical lines.

the two empirical bounds on the frequency content. Hence, even though turbulent pipe flow exhibits a continuous frequency spectrum when measured locally in space, a sparse representation in the frequency domain is obtained under this particular decomposition of the velocity field. As such, we now exploit the reduction in the number of samples required for success of compressive sampling, thus validating our original hypothesis, and explore key features of this representation.

B. Application of compressive sampling to frequency analysis of randomly sampled DNS velocity fields

The application of compressive sampling to the turbulent velocity field was investigated using the second, randomly sampled DNS dataset (Figure 3). The optimization parameters N_{opti} and df resulting in the highest frequency resolution that still satisfies the sparsity relationship were determined by analysis of the frequency content of several 2D Fourier modes. An optimization frequency range from -1 to 1 allows for both upstream and downstream propagating waves, and with a frequency increment of 0.005 results in the highest frequency resolution that can be obtained with the given number of samples. Only the downstream propagating waves carry a significant fraction of the streamwise turbulence intensity, but the noise is spread over both upstream and downstream propagating waves, leading to a sparser solution. (The presence of negative optimization frequencies is required for the optimization to converge, and it was found that more sparse solutions are obtained using both positive and negative optimization frequencies, instead of increasing the upper bound on the optimization frequencies.) The fundamental frequency is then $f_{min} = 0.005$ (half the inverse of the sampling duration) and the number of optimization frequencies is 400 . By way of comparison, the flow would have to be sampled for 200 dimensionless time units (instead of 100) in order to obtain a fundamental frequency of 0.005 using periodic sampling, and 400 samples would be required as opposed to 50 in the case of compressive sampling.

Figure 6 shows an example of the radially integrated power spectral density corresponding to correctly extracted frequency content of the 2D Fourier mode $(k, n) = (0.21, 2)$. Most of the energy is concentrated in the range $[0, 0.1]$, while the basis functions corresponding to higher frequencies, as well as negative frequencies corresponding to upstream propagating modes, only capture the noise. Note, however, that the optimization does not converge if the frequency range is restricted to $[0, 0.1]$ because in that case the high frequency noise above the noise floor cannot be resolved. The noise floor is about 5 orders of magnitude lower than the peak. The distribution of energy with respect to distance from the wall and the corresponding detailed radially integrated power spectral density are




FIG. 6. Power spectral density for the 2D Fourier mode $(k, n) = (0.21, 2)$ from the randomly sampled DNS obtained via optimal compressive sampling.

shown in Figures 7(a) and 7(b). The frequency spectrum in the range $[0, 0.1]$ contains two dominant frequencies $f = 0.035$ and 0.040 , with an energy content of not less than 10% of the peak value.

The solution to the convex optimization problem was shown to be robust with respect to variations of the fundamental and maximum frequencies, as long as the sparsity relationship is satisfied, and the maximum frequency is high enough to resolve the noise. When the resolution is increased by lowering the fundamental frequency, the width of the peak in the power spectral density stays constant, but the number of frequencies within the peak increases, resulting in a violation of the sparsity relationship (Eq. (8)). Figures 7(b) and 7(d) show that when the resolution is doubled in a local minimization in y (Eq. (4)) the number of dominant frequencies within the peak increases from 2 to 4, and the frequency spectrum as a function of the wall-normal distance exhibits distorted contours (Figures 7(a) and 7(c)), indicative of a violation of the sparsity relationship, Eq. (8). The minimum number of samples required to resolve the spectrum correctly is 32, whereas for the spectrum on Figure 7(d) 57 samples are needed, but only 50 are available. This issue is an artefact of imposing a discrete frequency resolution onto a velocity series obtained in the temporal rather than frequency domain.

The 2D Fourier modes can be reassembled by summation over the basis functions (temporal Fourier modes). The reconstructions of the 2D Fourier mode amplitude variation based on summation over the sparse frequencies identified in Figure 7 are compared with the true amplitude from the DNS in Figures 8(a) and 8(c). The frequency $f = 0.030$ is not considered to be a dominant frequency, based on the criterion that dominant frequencies need to contain at least 10% of the peak energy, but is retained here such that the superposition of the three most energetic frequencies captures over 90% of the energy in the 2D Fourier mode $(k, n) = (0.21, 2)$. Relaxation of the criterion distinguishing sparse/dominant frequencies to those containing, say, 1% of the energy associated with the most dominant mode would also admit $f = 0.065$, with an associated increase in the percentage of energy captured.

When the sparsity relationship is violated by increasing the number of optimization frequencies to 800 and therefore halving the frequency resolution to $f_{min} = 0.0025$, the results degrade differently depending on whether a global (Eq. (6)) or local (Eq. (4)) minimization (in the wall-normal direction) is used. In the case of local minimization, the dominant frequencies can vary slightly from one




FIG. 7. Power spectral density over the frequency range $f \in [0, 0.1]$ as a function of the wall-normal distance (a) and (c) and integrated in the wall-normal direction (b) and (d) for the 2D Fourier mode $(k, n) = (0.21, 2)$ from the randomly sampled DNS. The top and bottom rows correspond to a local minimization with, respectively, 400 and 800 optimization frequencies. The dots in (b) and (d) indicate the dominant frequencies.

wall-normal location to the next, leading to distorted wall-normal amplitude profiles that span only a portion of the radius, as can be seen by comparison of Figures 7(a) and 7(c) and Figures 8(a) and 8(c). In the case of global minimization, corresponding to the results in Figures 8(b) and 8(d), the energy is spread over more frequencies and the contribution of each dominant frequency to the streamwise turbulence intensity decreases when the resolution is doubled and the sparsity relationship is violated. However, the wall-normal profiles of the temporal Fourier modes are still smooth and span the whole radius, as shown in Figure 8(d).

Figure 8 shows that the amplitude variation of the 2D Fourier modes reconstructed by summation of the amplitudes of the dominant frequency components extracted via compressive sampling do not change when the sparsity relationship is no longer satisfied, implying that the same amount of energy is captured by the dominant frequencies, but distributed differently among them. Similarly, the dominant frequencies obtained using the local or global minimizations were shown to be identical for the 2D Fourier modes studied, as were the reconstructed 2D Fourier modes, illustrating the robustness of the method. However, the wall-normal profiles of the temporal Fourier modes were smoother in the case of global minimization (see Figure 8), justifying the choice of the global minimization method to analyze the 2D Fourier modes. Henceforth, we present the frequency content of a broad range of energetic 2D Fourier modes analyzed via compressive sampling using the global optimization method of Eq. (6) and the optimization parameters $df = 0.005$ and $f \in [-1, 1]$.

C. Features of the sparse representation

Having determined the correct sampling parameters for this application of compressive sampling and determined the effects of deviating too far from them, we now explore key features of the Fourier

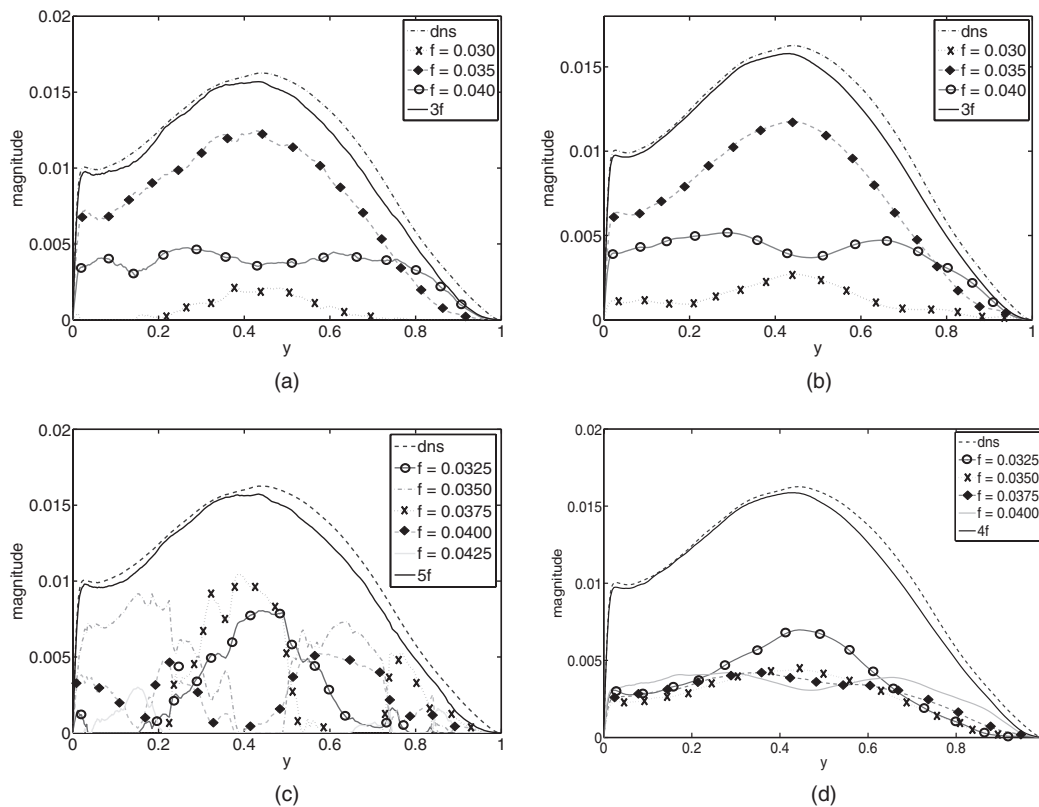


FIG. 8. Wall-normal profiles of the Fourier modes $(k, n, \omega) = (0.21, 2, 2\pi f)$ from the randomly sampled DNS, for the different dominant frequencies, compared to the time average of the original and reconstructed signals computed using the local (a) and global (b) minimizations in the wall-normal direction with 400 optimization frequencies, and the local (c) and global (d) minimizations with 800 optimization frequencies. The maximum frequency for the optimization is 1 in both cases.

modes recovered by the technique, which have some physical significance associated with the sparsity of the representation, as discussed in Sec. I above.

The most energetic 2D Fourier mode $(k, n) = (0.21, 2)$ described in Figure 7 was explored further as a representative example of the low streamwise wavenumber modes. The temporal Fourier modes corresponding to the three most energetic frequencies $f = 0.030, 0.035,$ and 0.040 of this 2D Fourier mode shown in Figure 8(b) together contain 91% of the streamwise turbulence intensity compared to 52% for the peak frequency alone. Isocontours of the real part of the full 2D Fourier mode and the reconstruction from the three dominant frequencies are shown in Figure 9 (with correct account taken of the wall-normal phase variations). The reconstructed signal clearly captures the wall-normal extent of the full mode, but lacks some of the small-scale variability. However, the presence of dominant frequencies adjacent to the peak frequency, which can be interpreted as a (dispersive) range of wavespeeds, does capture the apparent amplitude modulation in time of the full 2D Fourier mode at a fixed height (see, for example, the increase in magnitude of the reconstructed and full 2D Fourier modes from the first to the second period at $y \approx 0.4$).

Another two 2D Fourier modes $(k, n) = (1.05, 2)$ and $(k, n) = (3.14, 4)$ were selected for illustration purposes and represent the range of streamwise wavenumbers that can be analyzed with the 50 velocity fields available. The isocontours of the real part of the two 2D Fourier modes in a wall-normal-temporal plane are plotted on Figures 10(a) and 10(c) and the wall-normal profile of their dominant temporal Fourier modes is shown on Figures 10(b) and 10(d). A comparison between Figures 9 and 10 shows that the energetic radial extent of the reconstructed 2D Fourier mode decreases as the streamwise wavenumber increases, and the reconstructed modes contain less and less energy near the wall. In addition, the sparse frequency reconstruction captures less of the total energy in the 2D Fourier modes than in the $(k, n) = (0.21, 2)$ case. A reconstruction of the




FIG. 9. Isocontours of the real part of the 2D Fourier mode $(k, n) = (0.21, 2)$ from the randomly sampled DNS, as a function of the wall-normal distance and time. (a) All frequencies included. (b) The three most energetic frequencies recovered from compressive sampling recover the large scale features of the full signal.

$(k, n) = (1.05, 2)$ mode using the three sparse frequencies is shown in Figure 11(a). The corresponding full field was shown in Figure 10(a); there are still common key features, but the agreement between the reconstruction and the full mode is less satisfactory than for the $(k, n) = (0.21, 2)$ case.

Upon further inspection, two different zones of approximately uniform momentum in the wall-normal direction can be identified for the higher wavenumber cases, as observed in Figure 10(c)

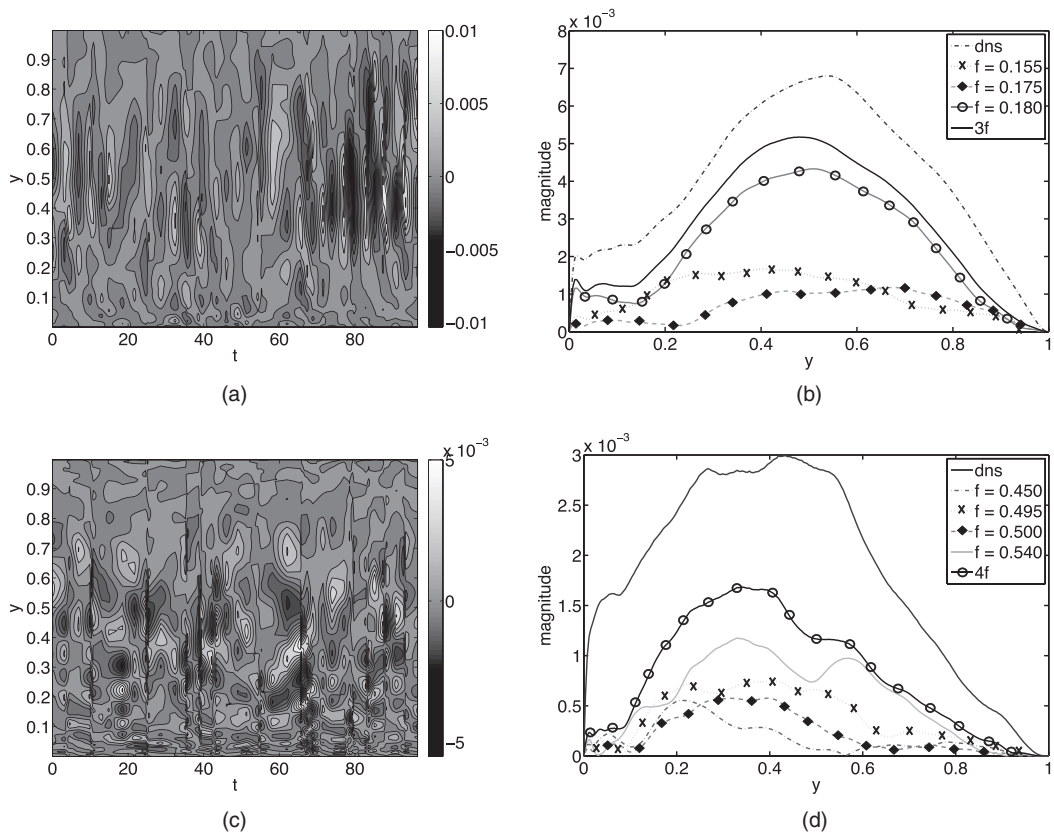


FIG. 10. Isocontours of the real part of the 2D Fourier modes $(k, n) = (1.05, 2)$ (a) and $(k, n) = (3.14, 4)$ (c) from the randomly sampled DNS as a function of the wall-normal distance and time. Wall-normal profiles of the Fourier modes $(k, n, \omega) = (1.05, 2, 2\pi f)$ (b) and $(k, n, \omega) = (3.14, 4, 2\pi f)$ (d) for the different dominant frequencies compared to the time average of the original and reconstructed signals. The reconstructions with three or four frequencies capture a significant amount of the energy in the original signal.




FIG. 11. (a) Isocontours of the real part of the 2D Fourier mode $(k, n) = (1.05, 2)$ reconstructed with only the three most energetic frequencies recovered from compressive sampling. The corresponding full field was shown in Figure 10(a). (b) Time-averaged power spectral density in the wall-normal direction for the 2D Fourier mode $(k, n) = (3.14, 4)$ from the randomly sampled DNS. Here, q is the quantum number, related to the number of zero crossings of the radial amplitude profile.

above and below $y = 0.1R$. Close to the wall, these zones are short in the wall-normal direction whereas further from the wall, the uniform momentum zones have a radial wavelength of about a quarter radius inferred from the sharp peak in the radial power spectrum shown in Figure 11(b). This spectrum was obtained by applying a FFT in the radial direction to the 2D Fourier mode amplitude (without windowing because the profile vanishes at the boundaries, the wall, and the centerline), and can therefore be assumed periodic with a fundamental wavelength equal to R . The spectrum is decomposed by quantum number, q , which in this case corresponds to the number of zero crossings of the mode amplitude plus one. Using a sharp Fourier filter (here with a cutoff at the 14th mode), the “near-wall type uniform momentum zones” can be separated from the “core uniform momentum zones” in the wall-normal direction, as shown in Figures 12(a) and 12(b) for $(k, n) = (3.14, 4)$. The filter successfully separates the two different types of uniform momentum zones that are integrated in Figure 10(c).

In terms of the compressive sampling approach, the contributions of these two types of zones to the full 2D Fourier modes, and in particular any variation in sparse frequency content, cannot be distinguished. However, separate minimization problems for each zone can be solved after use of the sharp radial filter. The frequency content of the near-wall uniform momentum zones is broadband and cannot be captured using compressive sampling with the chosen sampling parameters, reflecting the decrease in the energy captured near the wall as k increases observed on Figures 10(b) and 10(d). However, the sparse frequency representation in the core improves when only the core uniform




FIG. 12. (a) Isocontours of the real part of the 2D Fourier mode $(k, n) = (3.14, 4)$ low-pass filtered and (b) high-pass filtered showing the two different zones present in the wall-normal direction.

TABLE I. Frequency analysis of 5 representative modes using only 14 snapshots acquired over 100 dimensionless time units from two different runs of the turbulent pipe flow DNS. The dominant frequencies obtained via compressive sampling (CS) analysis of the randomly sampled DNS and their streamwise energy content (in % of the 2D Fourier mode energy content) are reported in the second and third columns, respectively. The fourth and fifth columns correspond to the frequencies and respective energy content obtained by FFT of the periodically sampled DNS.

(k,n)	f (CS)	% u^2 (CS)	f (FFT)	% u^2 (FFT)
(0.21,2)	0.035	82	0.03	51
	0.040	14	0.04	20
(0.42,-3)	0.065	95	0.06	21
			0.07	58
(0.63,-2)	0.09	23	0.08	50
	0.105	29	0.110	27
	0.110	23
(1.05,2)	0.180	59
(1.47,2)	0.255	15
	0.260	25
	0.265	8

momentum zones are considered. Compressive sampling thus can be used to give insight into the compactness of representation that can be achieved in different regions of the flow using the (k, n, ω) decomposition, and, more fundamentally, variations in sparsity in the frequency domain with wall-normal location and spatial wavenumber (assuming that the sampling parameters have been correctly chosen to resolve the corresponding bandwidth).

Besides the three representative (k, n) values described above, compressive sampling was applied to a broad range of 2D Fourier modes. The technique successfully captured sparse frequency content in the core of the pipe for energetic 2D Fourier modes ranging in size from the largest modes




FIG. 13. Comparison between the time-averaged wall-normal profile of the 2D Fourier mode $(k, n) = (0.21, 2)$ from the randomly sampled DNS and its two dominant frequencies obtained using compressive sampling. The mode shapes obtained with 14 and 50 snapshots show robust similarity.




FIG. 14. Power spectral density of a superposition of three Fourier modes with unit magnitude and frequencies $f = 0.035$, 0.052 , 0.087 sampled 50 times during 100 time units computed using (a) periodic sampling with FFT, and (b) compressive sampling with a frequency increment of $df = 0.01$ and (c) $df = 0.001$, showing the improvement in resolution obtained with compressive sampling.

($k = 0.21$) all the way down to the near-wall type modes of size $\lambda_x^+ = 1000$ ($k = 4.3$). Implications of this finding will be discussed in Sec. IV below.

D. Comparison with periodic sampling

The compressive sampling approach was briefly compared with periodic sampling using the third dataset (containing equal numbers of samples over identical total sample lengths) to emphasize the superiority of the former method in resolving the frequency content of the most energetic 2D Fourier modes when only a limited amount of data are available.

The dominant frequencies and their energy content for five representative 2D Fourier modes, as obtained from the FFT and compressive sampling, are reported in Table I. The frequencies obtained from the two techniques match relatively well, especially for the two modes with the lowest streamwise wavenumber, even though the samples were generated by two different runs of the DNS and the statistics would not be expected to be converged in either case. The dominant frequencies obtained by compressive sampling applied to either 14 or 50 samples are identical, but the wall-normal profiles of the temporal Fourier modes vary slightly, as shown in Figure 13 and the energy content of the dominant frequencies is higher when 50 samples are used instead of 14. This energy content is also higher than for the frequencies obtained via FFT, with, respectively, 96% vs. 71% of the energy in the dominant frequencies for the first mode, and 95% vs. 79% for the second. In other words, the compressive sampling solution requires less basis functions (frequencies) to capture a given percentage of streamwise turbulence intensity than the solution obtained with periodic sampling.

Conversely, compressive sampling also provides a way to increase the frequency resolution for a given number of samples above the Nyquist limit, since many more frequencies can be chosen to perform the optimization than the number of samples available. Consider a toy temporal signal composed of three Fourier modes $e^{i2\pi ft}$ with frequencies $f = 0.035, 0.052, 0.087$ and unit magnitude, sampled 50 times over 100 time units. The PSD found by FFT of the periodically sampled signal is shown in Figure 14(a). The minimum and maximum frequencies that can be resolved are given by $f_{min} = 0.01$ and $f_{max} = 0.49$ (note that the maximum frequency nearly equals the sampling frequency rather than half the sampling frequency because the input signal is complex-valued), and the three input frequencies are not clearly recovered for these sample parameters.

The same signal, now randomly sampled but with the same total number of samples over the same total sample length, was also interrogated using compressive sampling and two different values of frequency resolution. The first one matched the periodic sampling case, $df = f_{min} = 0.01$ and the second one was ten times higher $df = 0.001$. The PSDs are plotted on Figures 14(b) and 14(c). For the lower resolution case, frequencies away from the input frequencies have significantly lower energy content than in the periodically sampled case, and the input frequencies are more clearly, but not definitively, identified. The PSD exhibits a lower level of noise, or equivalently, a sparser solution than the PSD obtained via FFT of periodic samples. For the higher-resolution case, the three input frequencies are recovered exactly, because the set of temporal Fourier modes used for the optimization includes the three input frequencies, and the noise floor drops even further.

IV. CONCLUSIONS

The success of compressive sampling, or conversely the presence of a sparse representation of the energetic streamwise velocity field in the frequency domain under this formulation, has strong implications for our understanding of wall turbulence, and for future requirements on both data sampling and simulation strategies.

First and foremost, the sparse representation is likely to have physical significance as a property of wall turbulence. Consider the frequency content of all the dominant modes identified in this study, shown in Figure 15(a) as a function of the streamwise wavenumber. Within a resolution error of $df = \pm 0.0025$, all the frequencies fall in between the lower and upper empirical bounds corresponding to convection velocities equal to, respectively, 10 times the friction velocity and the centerline velocity. Note that a range of azimuthal wavenumbers were considered for each k (an effective integration in the azimuthal direction); in the full (k, n) domain there are significantly less frequencies associated




FIG. 15. Frequency content of the 2D Fourier modes as a function of the streamwise wavenumber, k (a range of azimuthal wavenumbers, n , was considered at each k). The two solid lines indicate the upper and lower empirical bounds on the frequency content from consideration of the convection velocities and the dashed line shows the streamwise wavenumber corresponding to the near-wall type modes. The shaded areas highlight the dynamically significant bandwidth and streamwise wavenumber range at (a) $Re = 24\,580$ and (b) extrapolated to $Re = 300\,000$. Dots represent 2D Fourier modes considered in this study.

with each wavenumber pair. Nonetheless, a visual identification of sparsity in the frequency domain in Figure 15 is still possible.

Clearly, under the Fourier decomposition described here, the energetic turbulent activity is confined to a much smaller domain than the range of frequencies and streamwise wavenumbers resolved by the simulation, which are denoted by the gray shaded area. Using the empirically defined, but physically and observationally justified, convection velocity limits of ten times the friction velocity and the centerline velocity together with a compressive sampling approach suggests that a truly compact representation of the energetic content of wall turbulence is possible. In the three-dimensional frequency/streamwise/spanwise wavenumber space, (k, n, ω) , this corresponds to a shrinking of a cuboid, defined by minimum and maximum energetic streamwise and spanwise wavenumbers and frequencies across the whole wall-normal domain, to a “pancake” demarked by the convection velocity limits and the radially local range of n . This pancake is not fully populated in the frequency domain, and as such this representation leads to a significant reduction of complexity.

McKeon and Sharma⁷ and subsequent works⁸ have described the resolvent framework for determining the form of the nonlinear interactions giving rise to a specific resolvent mode shape. The fact that only a sparse set of frequencies are active, or energetic, for a given wavenumber pair imposes additional constraints beyond triadic and physical compatibility of interacting modes. Further development of the sparsity result in the context of the resolvent framework is reserved for future work, but the discovery of a sparse representation is believed to be important to ongoing resolvent analysis.

The advantage of compressive sampling over periodic sampling in the case of approximately sparse data will become even more evident at higher Reynolds number, because the dynamically significant frequency range increases with Reynolds number. A simple extrapolation of Figure 15(a), namely, retaining the upper and lower empirical bounds on the convection velocity and recalculating the maximum k required to capture the near-wall activity for a Reynolds number that is an order of magnitude higher leads to the results of Figure 15(b). Extending the k range by an order of magnitude leads to linearly proportional (with coefficient one) broadening of the frequency range, implying that ten times more periodically sampled velocity fields would be needed in order to resolve the frequency content via FFT. The number of samples required for compressive sampling scales as the logarithm of the bandwidth, leading to an exponential gain in terms of number of samples required to fully resolve the frequency domain compared to the corresponding requirement for periodic sampling.

The reduction of complexity identified above suggests strategies for new, efficient simulation techniques for canonical wall-bounded turbulent flows, such as the pipe flow explored here. The expense of direction numerical simulation is proportional to the range of scales to be resolved; a simulation strategy that exploits our compact representation to reduce the dimensionality of the problem will limit that range and thus reduce the stiffness of the problem. Further, we propose that the correspondence of high wavenumber modes with high frequencies indicated in Figure 15 suggests the possible efficiency of scale-dependent time-stepping schemes.

The present study has used solely streamwise velocity data from DNS as the input to the compressive sampling optimization, in an effort to recover wall-normal coherence within a particular decomposition of the field. Clearly, obtaining experimental data capable of resolving a range of streamwise and spanwise wavenumbers and frequencies over sizable wall-normal extents, at turbulent Reynolds numbers, represents a significant challenge, even to such techniques as tomographic Particle Image Velocimetry (PIV). As such, progress in this regard has come via phase-locking to synthetic modes rather than decomposing unperturbed wall turbulence.^{8,28} However, efforts to extract a sufficient characterization of the full turbulent velocity field to allow the use of compressive sampling to extract resolvent modes in an experimental turbulent boundary layer are under way. Future work also includes analysis of the potential of the technique to recover the full three-component velocity field.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) under CAREER Award No. CBET-0747672. We are extremely grateful to Dr. Xiaohua Wu for his generosity in performing

new runs of the pipe flow DNS and providing us with the data used in this study. We thank the anonymous reviewers for critiques that led to significant improvements in this paper.

APPENDIX: MATLAB CODE USED TO SOLVE THE CONVEX OPTIMIZATION PROBLEM

```

% Nr is the number of grid points in the radial direction
% Ns is the number of samples
% 2D_Fourier(Nr,Ns) is the 2D Fourier mode as a function of the
    wall-normal
% distance r discretized over Nr points and the Ns number of samples
% (time instant ts)
% F_out(Nr,Nopti) is the output of the minimization corresponding to
% the frequency content as a function of the wall-normal distance.

Nopti=10*Ns; % number of frequencies for the
              optimization

f=(-Nopti/2+1:1:Nopti/2)/(Nopti/2); % equispaced frequencies for the
              optimization

C=exp(-1i*2*pi*ts*f); % constraint matrix: output must
                     match input
                     % at the sampling time instants

for j=1:Nr
    cvx_begin
    cvx_quiet(true)
        variable x(Nopti) complex;
        minimize(norm(x,1)); % minimize the sum of the absolute
                             value
                             % of the temporal Fourier
                             coefficients

        subject to
            C * x==2D_Fourier(j,:);
    cvx_end
    F_out(j,:)=x;
end.

```

¹M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.* **58**, 1182–1195 (2007).

²A. B. Tayler, D. J. Holland, A. J. Sederman, and L. F. Gladden, "Exploring the origins of turbulence in multiphase flow using compressed sensing MRI," *Phys. Rev. Lett.* **108**, 264505 (2012).

³J. H. Tu and C. W. Rowley, "Dynamic mode decomposition with sub-Nyquist-rate data samples," in *Proceedings of the SIAM Conference on Applications of Dynamical Systems*, Snowbird, UT, May 2013.

⁴M. R. Jovanovic, P. J. Schmid, and J. W. Nichols, "Low-rank and sparse dynamic mode decomposition," *Center for Turbulence Research Annual Research Briefs* (Center for Turbulence Research, Stanford University, 2012).

⁵M. R. Jovanovic and P. J. Schmid, "Sparsity-promoting dynamic mode decomposition," in *Bulletin of the American Physical Society*, San Diego, CA, November 2012.

⁶I. Bilinskis, *Digital Alias-Free Signal Processing* (John Wiley, 2007).

⁷B. J. McKeon and A. Sharma, "A critical layer framework for turbulent pipe flow," *J. Fluid Mech.* **658**, 336–382 (2010).

⁸B. J. McKeon, A. S. Sharma, and I. Jacobi, "Experimental manipulation of wall turbulence: a systems approach," *Phys. Fluids* **25**, 031301 (2013).

⁹A. S. Sharma and B. J. McKeon, "On coherent structure in wall turbulence," *J. Fluid Mech.* **728**, 196–238 (2013).

¹⁰R. Moarref, A. S. Sharma, J. A. Tropp, and B. J. McKeon, "Reynolds number scaling of the low rank approximation to turbulent channel flow," *J. Fluid Mech.* **734**, 275–316 (2013).

¹¹E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory* **52**, 489–509 (2006).

¹²J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory* **50**, 2231–2242 (2004).

- ¹³D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM J. Appl. Math.* **49**, 906–931 (1989).
- ¹⁴E. J. Candès and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Found. Comput. Math.* **6**, 227–254 (2006).
- ¹⁵J. A. Tropp, "On the linear independence of spikes and sines," *J. Fourier Anal. Appl.* **14**, 838–858 (2008).
- ¹⁶J. A. Tropp, "The sparsity gap: Uncertainty principles proportional to dimension," in *Proceedings of the 44th IEEE Conference on Information Sciences and Systems (CISS), Princeton, NJ, 2010* (IEEE, 2010), pp. 1–6.
- ¹⁷J.-L. Bourguignon, "Models of turbulent pipe flow," Ph.D. thesis, California Institute of Technology, 2012.
- ¹⁸G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annu. Rev. Fluid Mech.* **25**, 539–575 (1993).
- ¹⁹Z. Liu, R. J. Adrian, and T. J. Hanratty, "Large-scale modes of turbulent channel flow: Transport and structure," *J. Fluid Mech.* **448**, 53–80 (2001).
- ²⁰J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Trans. Inf. Theory* **56**, 520–544 (2010).
- ²¹D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006).
- ²²M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.0 beta," 2012, see <http://cvxr.com/cvx>.
- ²³M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, edited by V. Blondel, S. Boyd, and H. Kimura (Springer-Verlag Limited, 2008), pp. 95–110, see http://stanford.edu/~boyd/graph_dcp.html.
- ²⁴X. Wu and P. Moin, "A direct numerical simulation study on the mean velocity characteristics in turbulent pipe flow," *J. Fluid Mech.* **608**, 81–112 (2008).
- ²⁵D. D. Joseph, "Eigenvalue bounds for the Orr-Sommerfeld equation," *J. Fluid Mech.* **33**, 617–621 (1968).
- ²⁶W. R. B. Morrison, K. J. Bullock, and R. E. Kronauer, "Experimental evidence of waves in the sublayer," *J. Fluid Mech.* **47**, 639–656 (1971).
- ²⁷J. LeHew, M. Guala, and B. J. McKeon, "A study of the three-dimensional spectral energy distribution in a zero pressure gradient turbulent boundary layer," *Exp. Fluids* **51**, 997–1012 (2011).
- ²⁸I. Jacobi and B. McKeon, "Dynamic roughness-perturbation of a turbulent boundary layer," *J. Fluid Mech.* **688**, 258–296 (2011).