Multi-fidelity uncertainty quantification using RANS and DNS

By R. Ahlfeld, S. Laizet, F. Montomoli†, G. Geraci AND G. Iaccarino

Multi-fidelity uncertainty quantification using RANS and DNS for turbulent flow is explored. The objective is to achieve statistical results of higher quality than obtainable through RANS simulation but with significantly lower cost than by using DNS. The novelty of this work is that previous research on multi-fidelity modeling for uncertainty quantification in CFD has only investigated the use of partially converged or coarser mesh solutions as lower-fidelity models. This ensured a high correlation and low bias between the two models. A multi-fidelity solution combining RANS and DNS is potentially more beneficial but also more difficult to achieve. The use of two models with such large conceptual differences makes the correlation between the models more challenging. We therefore investigate two recent methods for model combination: a multi-grid approach of nested optimal Gaussian sparse grids and weighted regression favoring high-fidelity collocation points. The results show that a beneficial connection between the two models is indeed difficult, because the correlation between the responses is overall low. However, there are quantities for which the correlation is high enough and others for which a sufficiently high correlation can be achieved using trust-region correction methods.

1. Introduction

Real-world engineering systems are always affected by a range of uncertainties. In this work, we focus on geometrical variation that is caused by manufacturing errors, assembling inaccuracy, or in-service degradation. Various authors have shown that the impact of small geometrical profile variations can become disproportionately large under turbulent conditions. For example, Schnell *et al.* (2014) found that manufacturing errors reduce compressor efficiency on average by 2%. The effect of random geometric variations on the final product can be quantified before manufacturing using CFD. Such robust design methods make possible the production of more efficient blades and airfoils. One problem with this approach is, however, that the reliability of CFD predictions for geometrically changing geometries is difficult to validate. While various authors have performed uncertainty quantification (UQ) for geometric uncertainty using RANS simulations already, it remains difficult to evaluate to what degree their statistical predictions were affected by modeling errors. The problem is that even the most efficient propagation of aleatory uncertainty (parameter uncertainty, like geometric profile variations) becomes meaningless if epistemic uncertainty (model-form uncertainty) is too large. Dow & Wang (2011) therefore calibrate the RANS simulation using an adjoint method to efficiently solve an inverse problem from DNS data.

In this research, the RANS model is not calibrated. Instead, we compare statistics propagated using RANS simulations with statistics from nine DNS simulations. Then, we investigate how DNS and RANS can be combined into a multi-fidelity framework. In

† Imperial College London, UK

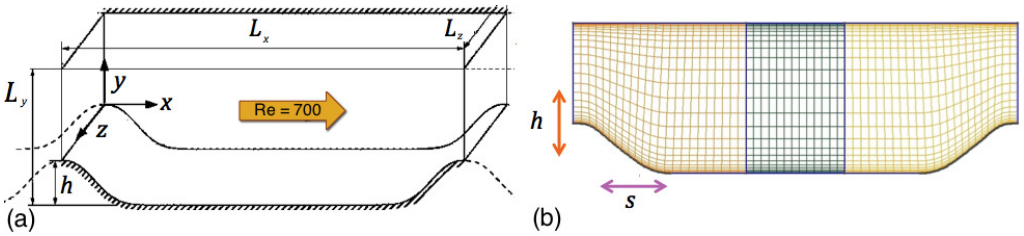

FIGURE 1. (a) Computational domain. (b) Parametric random variation of domain.

this way, we provide the first insight into how model-form errors of RANS simulations affect the shape of probability distributions under turbulent flow conditions, although be it at a Reynolds number much too low to be of industrial importance. The low Reynolds number is unavoidable, however, to obtain a sufficiently high number of DNS simulations to perform uncertainty quantification.

As test case we consider flow over periodically arranged hills at a Reynolds number of 700 (Figure 1(a)), as described for example in Marquillie *et al.* (2008). This scenario is chosen because a highly reliable experimental and computational database is available, see, for instance, Breuer *et al.* (2009). The flow separates from the curved surface of the hill, recirculates in its downstream side and reattaches at the flat channel bottom. Although this is a simple scneario, lower-fidelity CFD methods still struggle with predicting the point of reattachment. Important for this work is that the shape of the curved surface strongly affects the reattachment position. Hence, large model-form uncertainty can be expected from RANS simulations if geometrical uncertainty is present. We introduce geometric uncertainty by parametrizing the hill shape with two parameters: the hill height h and the hill skewness s as shown in Figure 1(b). The mean height is chosen at 75% of the reference case and the hill shape is skewed by asymmetrically stretching and compressing the upstream and downstream sides of the hill. For example, $s = 0.3$ means that the upstream part of the hill is stretched by a factor of 1.3, whereas the downstream side is compressed by a factor of 0.7. The two normal distributions $N(\mu, \sigma^2)$ for the height and skewness are defined as

$$h \sim N\left(0.75, 0.25^2\right) \qquad s \sim N\left(0, 0.3^2\right). \tag{1.1}$$

Much of the research done in UQ focuses on more efficient uncertainty propagation methods. Due to the use of DNS, we also require an efficient uncertainty propagation framework to keep the number of model runs at a minimum. We therefore use polynomial chaos methods to expand the system response into a series of orthogonal polynomials, which are optimally determined from the input probability density functions (PDFs) of the stochastic geometry parameters. For a small number of random parameters and a smooth model response this results into a convergence speedup of the statistical quantities that is several orders of magnitudes faster than using Monte Carlo sampling. To avoid interfering with the DNS and RANS codes, we use only non-intrusive methods, in which the projection of the model response is approximated through numerical integration (quadrature), as described by Ng & Eldred (2012), or determined through linear regression (point collocation), as detailed by Palar *et al.* (2016).

Within the efficient framework of polynomial chaos expansions we investigate multi-fidelity approaches combining RANS with DNS to compute the output statistics with the accuracy of DNS but at lower computational cost. DNS is chosen because it provides the highest fidelity possible in CFD. RANS is chosen because it is much cheaper and the

current standard in the industry. The concept of exploiting multiple levels of fidelity has only recently entered the field of UQ. Giles (2008) proposed to speedup the solution of stochastic partial differential equations using multi-level Monte Carlo sampling. Following this work, Ng & Willcox (2014) implemented a multi-level Monte Carlo method for engineering models. Within the field of polynomial chaos, two authors recently developed new methods: Ng & Eldred (2012) developed an approach using different levels of sparse grids and Palar *et al.* (2016) developed an approach using a multi-level point collocation. Slight variations of the methods by Ng & Eldred and Palar *et al.* are employed in this work for multi-level RANS and DNS combinations. Unlike Ng & Willcox (2014) we use sparse grids based on optimal Gaussian quadrature points, which are generally not nested. However, a nested approach was developed by the authors in a previous work, see Ahlfeld *et al.* (2016). We reuse the optimal Gaussian collocation points for the second approach of weighted linear regression, where we assign higher priority to them. The weights are derived from the model correlation coefficient R^2 introduced by Palar *et al.* (2016).

2. Numerical modeling and flow parameters

The governing equations are the forced incompressible Navier-Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} = -\nabla p - \frac{1}{2}\left[\nabla\left(\mathbf{u}\otimes\mathbf{u}\right) + (\mathbf{u}\cdot\nabla)\mathbf{u}\right] + \nu\nabla^2\mathbf{u} + \mathbf{f} \tag{2.1}$$

$$\nabla\cdot\mathbf{u} = 0, \tag{2.2}$$

where $p(\mathbf{x},t)$ is the pressure field (for a fluid with a constant density $\rho = 1$), $\mathbf{u}$ the velocity field and ν the kinematic viscosity of the fluid. The high-fidelity 3D simulations presented in this study and shown in Figure 2 were performed for a computational domain $9h \times 3.036h \times 4.5h$, where h is the height of the hill, discretized on a Cartesian mesh with $512 \times 257 \times 128$ mesh nodes. The mesh is slightly refined in the y direction near the walls. The spatial resolution is as good as the ones found in previous numerical studies for the same Reynolds number. The time step is equal to $0.0005h/U_b$, where U_b is the bulk velocity. For all calculations, turbulent statistical data have been collected on a time period $T = 150h/U_b$, with a spatial average in the span-wise direction in order to reach a statistical convergence. The coordinate system is orthonormal, with coordinate x in the streamwise direction, y in the vertical direction and z in the span-wise direction. The boundary conditions are of periodic type in the streamwise and span-wise directions and zero-velocity type in the vertical direction. The initial condition for the velocity field is a Poiseuille profile $(1-(y/h)^2)$, where h is half the height of the channel. First, the simulations are launched without any hill in order to get a fully developed flow. In order to drive the flow, a pressure gradient is adjusted dynamically at each time step to maintain a constant mass flow rate through the channel. In our setup, the reference mass flow rate is computed from the Poiseuille profile and is equal to

$$\dot{m} = \rho\frac{1}{2h}\int_{-h}^{h} 1 - \left(\frac{y}{h}\right)^2 dy. \tag{2.3}$$

In our setup we get $\dot{m} = 2/3$ as $\rho = 1$ and $h = 1$. This parameter is kept constant for all the simulations. Then the hill is introduced into the computational domain and the data are collected only after a transitional period when the flow is fully developed again.

ANSYS Fluent is used for the 2D RANS simulations, as suggested by Iaccarino (2001). The SST k-ω turbulence model developed by Menter *et al.* (2003) is chosen because of

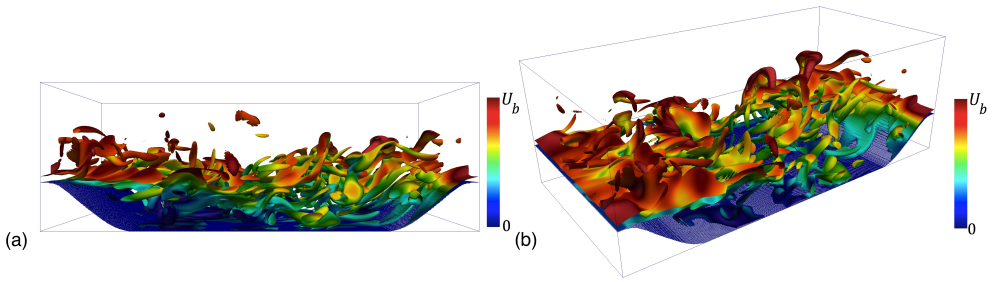

FIGURE 2. Instantaneous DNS visualization of the vorticity modulus (for $||W|| = 5Uc/h$) colored by the streamwise velocity. (a) side view (b) angles view of domain.

its suitability for separated flows. The mesh is constructed from curvilinear blocks with 192×120 nodes as illustrated in Figure 1(b) . This high resolution is necessary for an accurate comparison of the bubble height and length scale with DNS; otherwise, the same configurations are used as in DNS.

2.1. *Numerical method for DNS*

In this work, we use a customized immersed boundary method, which is based on a forcing field $\mathbf{f}(\mathbf{x}, t)$ in order to take into account the 2D hill inside the computational domain. Note that the convective terms are written in skew-symmetric form. This form reduces aliasing errors while remaining energy conserving for the type of spatial discretization considered here. To solve the incompressible Navier-Stokes equations, we use the high-order flow solver Incompact3D, which is based on sixth-order compact schemes for spatial discretization on a Cartesian mesh and a third-order Adams-Bashforth scheme for time advancement. To treat the incompressibility condition, a fractional step method is used to solve a Poisson equation. For efficiency reasons, this equation is solved in spectral space using appropriate 3D fast fourier transforms (FFT). In order to have a strict equivalence between finite-difference operators in physical space and spectral operators, we use the concept of modified wave number introduced by Lele (1992), which reduces the accuracy of the spectral operators to sixth-order accuracy. Note that the divergence free condition is ensured up to machine accuracy. The pressure mesh is staggered from the velocity mesh by half a mesh to avoid spurious pressure oscillations. The modeling of the hill is performed with a customized immersed boundary Method based on a direct forcing approach that ensures a zero-velocity boundary condition at the wall of the hill. More details about the present code and its validation, especially the original treatment of the pressure in spectral space, can be found in Laizet & Lamballais (2009). Because of the size of the simulations, the parallel version of Incompact3d has been used for this numerical work. Based on a highly scalable 2D decomposition library and a distributed FFT interface, it is possible to use it on thousands of computational cores. More details about this efficient parallel strategy can be found in Laizet & Li (2011).

3. Polynomial chaos expansions

The concept of polynomial chaos is to expand a stochastic model output $f(\bar{\xi})$, dependent on a vector of d independent stochastic input random variables $\bar{\xi} = \xi_1, \xi_2, ..., \xi_d$ with event space Ω as a linear combination of N_P stochastic multivariate orthogonal

polynomials $\Psi_i(\bar{\xi})$ and deterministic coefficients α_i

$$f(\xi) \approx \sum_{i=0}^{N_P} \alpha_i(x)\,\psi_i(\xi). \tag{3.1}$$

In case a full tensor product is used, the number of linear combination terms N_P for a polynomial expansion of order p is truncated to $N_P = \frac{(d+p)!}{d!p!}$. The multivariate orthogonal polynomials Ψ_k are calculated as products of univariate orthogonal or orthonormal polynomials ψ_k

$$\Psi_k(\bar{\xi}) = \prod_{i=1}^{d} \psi^{(i)}(\xi_i),\ k \in \overline{1, N_P}, \tag{3.2}$$

and for them holds

$$\int_{\xi\in\Omega} \psi_k(\xi)\,\psi_l(\xi)\,w(\xi)\,d\xi = \delta_{kl} \quad \forall k,l = \overline{0, N}. \tag{3.3}$$

To achieve exponential convergence, the ψ_k should be chosen as orthogonal polynomials belonging to the optimal Gaussian quadrature rule corresponding to the individual input random variables. For example, the Hermite polynomials are optimal orthogonal for the normal distribution and are used in this work. The polynomial chaos coefficients can then be obtained by computing

$$\alpha_k = \frac{\langle f(\xi), \psi_k(\xi)\rangle_w}{\langle \psi_k(\xi), \psi_k(\xi)\rangle_w}. \tag{3.4}$$

In this work, we use optimal Gaussian numerical quadrature based on sparse grids and weighted linear regression in the sparse grid collocation points. The latter choice of collocation points is referred to as probabilistic collocation and was succesfully applied by Loeven *et al.* (2008) to investigate geometric airfoil uncertainty.

3.1. *Nested optimal Gaussian sparse grids*

The details of the nested Gaussian sparse quadrature formulation can be found in Ahlfeld *et al.* (2016), and the method to combine multiple grids to obtain polynomial chaos coefficients is described in Ng & Eldred (2012). Unlike Ng & Eldred (2012), we use sparse grids based on optimal Gaussian quadrature points, which are generally not nested. However, a nested approach was developed by Ahlfeld *et al.* (2016). Following this work, the sparse quadrature is calculated using an adaptive and anisotropic version of Smolyak's rule

$$A(d+l,d) = \sum_{\substack{l+1 \\ \leq |i| \leq \\ l+d}} (l-1)^{l+d-|i|} \begin{pmatrix} d-1 \\ l+d-|i| \end{pmatrix} \otimes_{k=1}^{d} U^{i_k}, \tag{3.5}$$

where $\{U^{i_j}\}_{j=1,\ldots,d}$ are d sequences of one-dimensional quadrature rules, l is the level and the term $|i|$ is the sum of a row of the index matrix $I_{d+l,d}$. The operator $S_{q,d}$ is introduced to describe a sparse quadrature approximation of the d-dimensional polynomial chaos expansion, as described in Eq. (3.5), where $q = d + l$ is used to shorten the notation. The multi-grid combination between RANS and DNS is set up by evaluating the RANS model response f_{RANS} more frequently on a higher-order grid of level q, whereas the DNS response f_{DNS} is evaluated less frequently on a lower-level grid of reduced level $q - r$. The sparse quadratures are defined as

$$f_{RANS}(\bar{\xi}) \approx S_{q,d}\left[f_{RANS}(\bar{\xi})\right] \qquad f_{DNS}(\bar{\xi}) \approx S_{q-r,d}\left[f_{DNS}(\bar{\xi})\right], \tag{3.6}$$

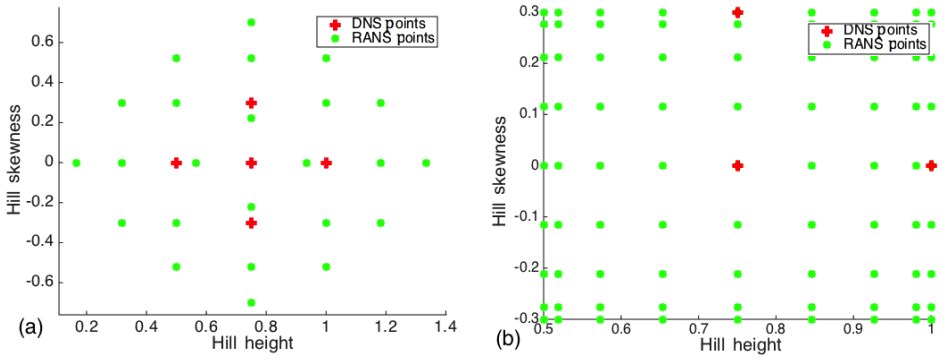

FIGURE 3. (a) Multi-fidelity nested Gaussian sparse grids. (b) Clenshaw-Curtis RANS points and three DNS points used for weighted regression.

where the lower-order grid is a subset of the higher-level grid, that is, $S_{q-r,d} \subset S_{q,d}$. In this framework, the DNS response can be approximated through the multi-fidelity approximation

$$\tilde{f}_{DNS}\left(\bar{\xi}\right) \approx S_{q,d}\left[f_{RANS}\left(\bar{\xi}\right)\right] + S_{q-r,d}\underbrace{\left[f_{DNS}\left(\bar{\xi}\right) - f_{RANS}\left(\bar{\xi}\right)\right]}_{C(\xi)}. \tag{3.7}$$

Figure 3(a) illustrates this concept, where the DNS points are obtained from a level 1 grid and the RANS points from a level three grid. The DNS points are also contained in the RANS grid, but only influence the first three polynomial chaos coefficients. In the form of polynomial chaos expansions, the multi-fidelity correction becomes

$$S_{q,d}\left[f_{RANS}\left(\bar{\xi}\right)\right] = \sum_{i \in I_{q,d}} \alpha_{RANS,i}\Psi_i\left(\bar{\xi}\right) \qquad S_{q-r,d}\left[C\left(\bar{\xi}\right)\right] = \sum_{i \in I_{q-r,d}} \alpha_{C,i}\Psi_i\left(\bar{\xi}\right), \tag{3.8}$$

where C represents the correction function introduced in Eq. (3.7). The polynomial chaos multi-fidelity approximation of f_{DNS} then becomes

$$\tilde{f}_{DNS}\left(\bar{\xi}\right) \approx \sum_{i \in I_{q,d}} \left(\alpha_{RANS,i} + \alpha_{C,i}\right)\Psi_i\left(\bar{\xi}\right) + \sum_{i \in I_{q,d} \setminus I_{q-r,d}} \alpha_{RANS,i}\Psi_i\left(\bar{\xi}\right). \tag{3.9}$$

3.2. *Weighted regression (point collocation) polynomial chaos*

The concept of point-collocation non-intrusive polynomial chaos is to obtain the polynomial chaos coefficients by solving a linear system of equations

$$\begin{bmatrix} \psi_1\left(\xi_1\right) & \cdots & \psi_1\left(\xi_d\right) \\ \vdots & \ddots & \vdots \\ \psi_{N_P}\left(\xi_1\right) & & \psi_{N_P}\left(\xi_d\right) \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \vdots \\ \alpha_{N_P} \end{bmatrix} = \begin{bmatrix} f\left(\bar{\xi}_0\right) \\ \vdots \\ f\left(\bar{\xi}_{N_p}\right) \end{bmatrix}. \tag{3.10}$$

The full details of this point collocation approach are described in Palar *et al.* (2016). Reasonable accuracy can be achieved if this system is oversampled and a linear regression approach is used to find $\hat{\alpha} = \arg\min\left[\bar{\alpha}^T\Psi\left(\bar{\xi}\right) - f\left(\bar{\xi}\right)\right]^2$. Unlike Palar *et al.* (2016) we do not use quasi-random samples. To keep the computational costs low, we reuse the DNS collocation points. More details about such a probabilistic collocation approach can be found in Loeven *et al.* (2008). In addition, we introduce a weighted linear regression to

QoI	R^2	$\|M\|_{rel}$
Reattachment location	0.8504	0.6414
Height of the separation bubble at $x/h = 2$	0.9570	0.1049
Height of the separation bubble at $x/h = 4$	0.7840	2.0988
Maximum horizontal backflow velocity at $x/h = 2$	0.6088	0.6197

TABLE 1. Correlation test of DNS and RANS for separation bubble parameters .

solve for the polynomial coefficients

$$\hat{\alpha} = \left(\Psi^T w \Psi\right)^{-1} \Psi^T w \bar{Y}. \tag{3.11}$$

The weights are derived from the model correlation coefficient R^2 shown in Eq. (4.1) that Palar *et al.* (2016) introduce. DNS measurements receive weight 1 and RANS measurements receive the calculated correlation coefficient as weight. Hence, higher priority is assigned to RANS if it captures the DNS trend well. In the limiting case that it captures the trend perfectly, both models receive the same weight.

4. Results of RANS/DNS multi-fidelity modeling

In order to achieve computational cost reductions using a multi-fidelity approach, the lower-fidelity model has to capture the general trends observed in the higher-fidelity model. To measure the correlation between RANS and DNS for quantities of interest, we use the R^2 correlation, as suggested by Palar *et al.* (2016), to quantify the trend compatibility

$$R^2 = \left[\frac{\sum_{i=1}^{m} \left(y_i^{[H]} - \bar{y}^{[H]}\right)\left(y_i^{[L]} - \bar{y}^{[L]}\right)}{\left(\sum_{i=1}^{m} \left(y_i^{[H]} - \bar{y}^{[H]}\right)^2\right)\left(\sum_{i=1}^{m} \left(y_i^{[L]} - \bar{y}^{[L]}\right)^2\right)} \right]^2 \tag{4.1}$$

and the mean absolute relative error $|M|_{rel}$

$$|M|_{rel} = \frac{1}{m} \frac{\sum_{i=1}^{m} \left| y_i^{[H]} - y_i^{[L]} \right|}{\bar{y}^{[H]}} \tag{4.2}$$

to quantify the offset between the two models. The results for various quantities of interest with regard to the size and length of the separation bubble are summarized in Table 1. In the comparison of probability distributions that follows, the DNS distributions are always computed by performing a Monte Carlo simulation on a linear interpolation between the DNS points. This is done to keep response surface shape assumptions at a minimum.

Since the bubble height shows a very high correlation between the two models, we apply the sparse multi-grid approach to it. The high correlation allows us to treat the RANS samples like DNS solutions within the trust region spanned by the DNS points. Therefore, the additional samples can be used to create a more accurate probability distribution. Figure 4 shows that the multi-fidelity grid combines the smoothness of the RANS PDF with the slightly more accurate trends contributed by the DNS. The maximum horizontal backflow velocity is found inside the bubble at the specified probe location x/h=2. It

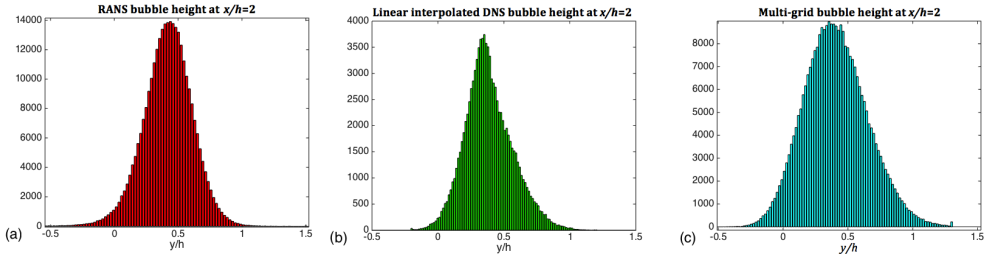

FIGURE 4. (a) 81 RANS points. (b) Sampled linear interpolation of 9 DNS points. (c) Combined multi-grid of 24 RANS and 5 DNS points.

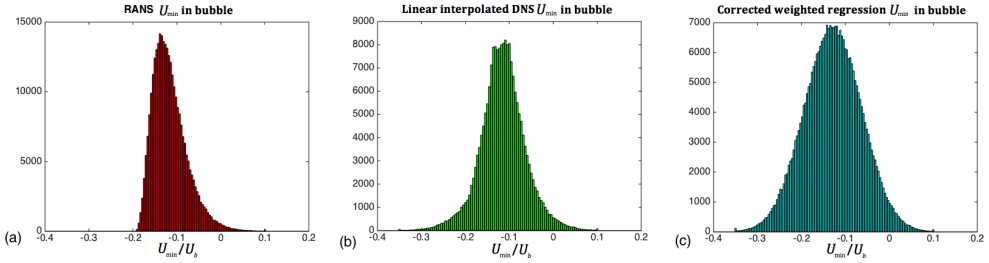

FIGURE 5. (a) 81 RANS points. (b) Sampled linear interpolation of 9 DNS points (c) Weighted regression of 81 RANS and 3 DNS points

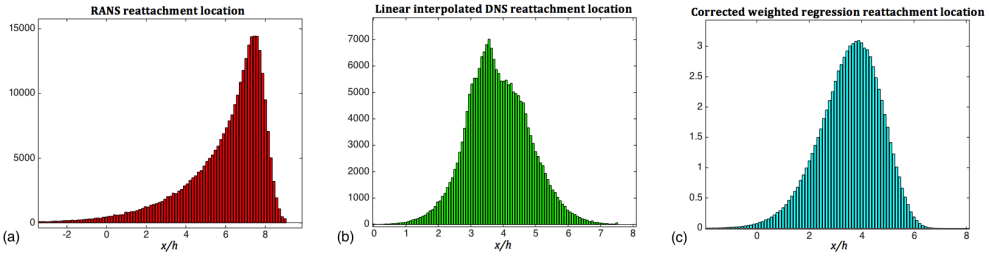

FIGURE 6. (a) 81 RANS points. (b) Sampled linear interpolation of 9 DNS points. (c) Weighted regression of 81 RANS and 3 DNS points.

is a good indicator of the overall strength of the separation bubble. Unfortunately, the correlation between the two model responses is fairly low at approximately 0.61. Nevertheless, a comparison of the response surfaces in all nine points in Figure 7(a) shows that the RANS follows the DNS trend reasonably well. To combine the models we use weighted regression, as it is the more efficient and more general approach. In addition, we use a linear plane as a correction function

$$C\left(\xi_1,\xi_2\right) = f_{DNS}\left(\bar{\xi}\right) - f_{RANS}\left(\bar{\xi}\right) = C_0 + C_1\xi_1 + C_2\xi_2 \tag{4.3}$$

to correct the RANS response surface. The idea is that three DNS are the absolute minimum number to define a correction plane. When they are chosen at the corners, they span the full stochastic space and create a large trust region. Thus, only three DNS simulations are needed to significantly correct the distribution shape. The RANS weights for the regression are chosen equal to 0.61 and the DNS weights as 1. Figure 5 shows that the RANS PDF is incorrectly skewed to the left and its variance is too low. The

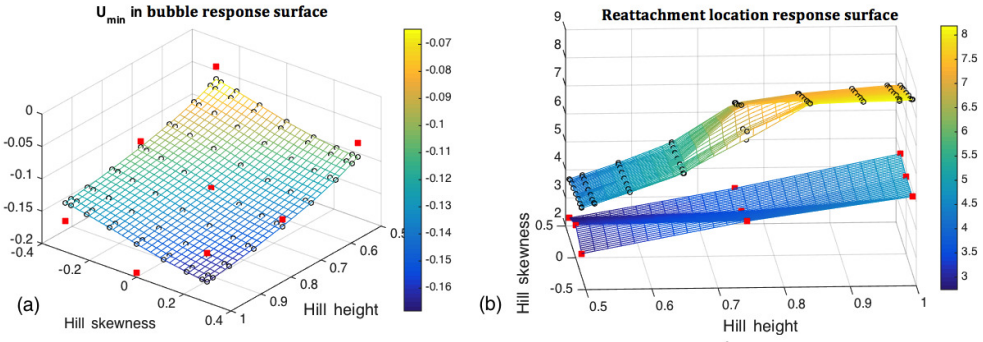

FIGURE 7. Response surfaces for (a) maximum backflow velocity in bubble and (b) reattachment point location.

multi-fidelity PDF resembles the pure DNS PDF more closely and is also smoother due to the additional RANS collocation points.

For the reattachment point, the models show a reasonably high correlation of approximately 0.85. However, the RANS response surface in Figure 7(b) shows a sharp bend in the response surface. This is caused by the fact that the kω-SST model overpredicts the position of the reattachment location. As the investigated geometric uncertainty is large with respect to the domain size, the bubble growth is impeded by the next hill. This creates a in the RANS response surface a sharp bend that is not present in the DNS simulations. This defect materializes itself in the PDF of the multi-fidelity solution in Figure 6(c) in form of an incorrect skewness, even though mean and variance values are improved.

5. Conclusions

We investigated the possibility of using multi-fidelity uncertainty quantification combining RANS and DNS for turbulent flows. Due to the large conceptual differences of these models, we tested two methods for combination: a multi-grid approach of nested optimal Gaussian sparse grids and a weighted regression favoring high-fidelity collocation points. As a test case, we investigated flow over periodic hills with geometrically uncertain hill height and skewness. RANS simulations are known to struggle with predicting the size and length of the separation bubble for this case. Hence, we focused the uncertainty quantification on proxies of the size and position of the bubble. We found that, in general, the connection between RANS and DNS is low and a direct connection between the models using multi-fidelity approaches is difficult. Moreover, pure RANS methods can lead to significantly incorrect PDF shapes. This was shown using the example of the distribution of the reattachment point location. However, a beneficial connection between the two models can be achieved for specific parameters by using correlation and correction functions. Further investigation of correction and combination, which should be conducted at a higher Reynolds number, is a promising opportunity for further research.

Acknowledgments

The authors acknowledge use of computational resources from the Certainty cluster awarded by the National Science Foundation to CTR.

REFERENCES

Ahlfeld, R., Belkouchi, B. & Montomoli, F. 2016 SAMBA: sparse approximation of moment-based arbitrary polynomial chaos. *J. Comput. Phys.* **319**, 1–29.

Breuer, M., Peller, N., Rapp, C. & Manhart, M. 2009 Flow over periodic hills–numerical and experimental study in a wide range of Reynolds numbers. *Comput. & Fluids.* **38**, 433–457.

Dow, E. & Wang, Q. 2011 Uncertainty quantification of structural uncertainties in RANS simulations of complex flows. *AIAA Paper* 2011-3865.

Giles, M. B. 2008 Multilevel Monte Carlo path simulation. *Oper. Res.* **56**, 607–617.

Iaccarino, G. 2001 Predictions of a turbulent separated flow using commercial CFD codes. *J. Fluids Eng.* **123**, 819–828.

Laizet, S. & Lamballais, E. 2009 High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *J. Comput. Phys.* **228**, 5989–6015.

Laizet, S. & Li, N. 2011 Incompact3d: A powerful tool to tackle turbulence problems with up to O (105) computational cores. *Int. J. Numer. Meth. Fl.* **67**, 1735–1757.

Lele, S. K. 1992 Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* **103**, 16–42.

Loeven, G., Witteven, J. & Bijl, H. 2008 Airfoil analysis with uncertain geometry using the probabilistic collocation method. *AIAA Paper* 2008-2070.

Marquillie, M., Laval, J.-P. & Dolganov, R. 2008 Direct numerical simulation of a separated channel flow with a smooth profile. *J. Turbul.* p. N1.

Menter, F. R., Kuntz, M. & Langtry, R. 2003 Ten years of industrial experience with the SST turbulence model. *Int. J. Heat Mass Tran.* **4**, 625–632.

Ng, L. W.-T. & Eldred, M. S. 2012 Multifidelity uncertainty quantification using nonintrusive polynomial chaos and stochastic collocation. In *AIAA Paper* 2012-1852.

Ng, L. W. T. & Willcox, K. E. 2014 Multifidelity approaches for optimization under uncertainty. *Int. J. Numer. Meth. Eng.* **100**, 746–772.

Palar, P. S., Tsuchiya, T. & Parks, G. T. 2016 Multi-fidelity non-intrusive polynomial chaos based on regression. *Comput. Method. Appl. M.* **305**, 579–606.

Schnell, R., Ashcroft, G. & Hoehe, L. 2014 Geometric variability and its influence on steady and unsteady fan aerodynamic performance. In *29th Congress of the International Council of the Aeronautical Sciences*, pp. 1–9.