



High-order computational fluid dynamics simulations of a spinning golf ball

Jacob Crabill¹ · Freddie Witherden¹ · Antony Jameson¹

Published online: 20 February 2019
© International Sports Engineering Association 2019

Abstract

This paper presents the first high-order computational fluid dynamics (CFD) simulations of static and spinning golf balls at realistic flow conditions. The present results are shown to capture the complex fluid dynamics inside the dimples which lead to drag reduction versus a smooth sphere, and compare well to previous experimental and computational studies. The high-order flux reconstruction method has been paired with the artificial boundary overset method to enable simplified mesh generation and grid motion. The compressible Navier–Stokes equations are modeled using a scale-resolving large eddy simulation (LES) approach with no sub-grid models. The codes implementing these methods have been implemented for NVIDIA graphical processing units (GPUs), enabling large speedups over traditional computer hardware. The new method allows for the simulation of golf balls, and other objects at moderate Reynolds numbers, to be simulated in a matter of days on large computing clusters.

Keywords Computational fluid dynamics · Large eddy simulation · Finite element methods · Golf ball · Sports aerodynamics

1 Introduction

The flow physics behind the phenomenon of drag reduction of dimpled spheres—golf balls—has been investigated since at least the 1970s [3, 22]. Early studies relied primarily on wind tunnel experiments, typically collecting force data, and occasionally also performing flow visualizations with, for example, oil streaks. It is only recently, with the advent of large-scale large eddy simulation (LES) and direct numerical simulation (DNS) simulations of modest Reynolds numbers (Re), that the accurate, predictive computational fluid dynamics (CFD) simulation of a golf ball has allowed deeper insight into the effects of dimples.

The typical goal of a golf ball design for a tee shot is to maximize the range it can be driven in a straight line. This

primarily leads to the desire to reduce its drag as much as possible, with secondary goals of minimizing variation in side forces to maintain straight-line flight, and of maximizing the lift force produced by backspin. Putting backspin on the ball produces lift via the Magnus effect, which extends the flight time and distance of the ball. Drag reduction is mostly due to the dimples, which are sized to create a series of separation bubbles that will lead to early transition in the unstable shear layer above the bubbles. The exact size, depth, and arrangement of dimples all contribute to the final aerodynamic properties of a golf ball under various conditions, and the full resolution of the flow details within each dimple is required to fully characterize these properties. However, collecting such data is arduous, requiring either fine-scale wind tunnel measurements or extremely large CFD simulations. The objective of this work was to show that even though large-scale CFD simulations are required to capture the underlying physics involved in golf ball aerodynamics, their use is now feasible on modest computing clusters. Further, such simulations are shown to provide a level of accuracy and detail suitable to better understand golf ball aerodynamics and develop new designs.

✉ Jacob Crabill
jcrabill@stanford.edu

Freddie Witherden
fdw@stanford.edu

Antony Jameson
ajameson@stanford.edu

¹ Stanford University, Stanford, USA

1.1 Previous studies

The first noteworthy experimental investigation of the aerodynamics of golf balls under a variety of flow conditions is by Bearman and Harvey in 1976 [3]. Their study used wind tunnel testing of scaled golf ball models to compare the characteristics of round vs. hexagonal dimples, with a smooth sphere used to assess the validity of their experimental setup. The hexagonally dimpled ball also had far fewer dimples than the “conventional” ball (240 vs. 330 or 336). They found that the hexagonally dimpled ball had a lower drag coefficient (C_D) and higher lift coefficient (C_L) over most of the Re and spin rate ranges of interest, hypothesizing that the hexagonal dimples led to more discrete vortices due to the straight edges of the dimples. The effect of the dimple edge radius was not studied. For both dimple types, however, they showed that the dimples serve to reduce the critical Re at which a drag reduction occurs and that the drag coefficient remains nearly constant for a large range of Re after this point.

Another detailed wind tunnel study was performed by Choi et al. [7]. They studied both fully dimpled and half-dimpled spheres without rotation. In comparison to Bearman and Harvey, they used only round dimples with a much smaller depth ($k/d = 4 \times 10^{-3}$ for Choi et al. vs. $k/d = 9 \times 10^{-3}$ for Bearman and Harvey, where k is the dimple depth and d is the sphere diameter) and also with a larger number of dimples (392 vs. approximately 330). Their results showed a slightly higher critical Re ($\sim 80,000$ vs. $\sim 50,000$) with a slightly lower C_D (~ 0.21 vs. ~ 0.25) afterwards, with a much more noticeable rise in C_D after the initial drop. Velocity data collected with a hot-wire anemometer was used to confirm that the turbulence generated by the free shear layer over the dimples led to an increase in momentum near the surface of the golf ball after reattachment and that the separation angle remained at a constant 110° after the critical Re .

Further studies have been performed using a combination of Reynolds-averaged Navier–Stokes (RANS) [28, 29], LES [1, 20, 21], DNS [4, 26], and wind tunnel experiments [1, 8, 27]. Li et al. proposed a link between small-scale vortices created at the golf ball dimples and a reduction in side-force variations at supercritical Reynolds numbers. Also, several studies have discussed an apparent positive correlation between dimple depth the supercritical drag coefficient, and a negative correlation between dimple depth and critical Reynolds number [4, 8, 29]. Although no specific mechanism has been proposed to explain the correlation with supercritical drag coefficient, the correlation with the critical Reynolds number is likely due to the instability of the free shear layer over the dimples becoming more unstable as the dimple (and hence separation

bubble) becomes deeper, leading to quicker transition at lower Reynolds numbers. Numerous wind tunnel and computational experiments have also shown a strong positive correlation between lift force and spin rate (though smaller in magnitude than the negative lift force generated by a similar smooth sphere) and a slight correlation between drag and spin rate as well [1, 3, 4, 23].

2 Simulation overview

2.1 Simulation method

The flow of air around a golf ball is governed by the Navier–Stokes equations. For the present study, we have utilized the implicit LES (ILES) method, sometimes referred to as under-resolved DNS. The full viscous compressible Navier–Stokes equations are solved, with no additional models used to model the dissipation of kinetic energy at the smallest length scales present in the flow. Instead, since the Reynolds number of the golf ball we simulate here—150,000—is low enough to capture most of the scales involved, we size the numerical grid to capture all but the smallest structures in the fluid flow, and allow the dissipation inherent to our spatial discretization method to ‘model’ the dissipation that would occur at the smallest scales. Such an approach is currently practical on modern computing clusters up to Reynolds numbers of approximately 500,000. Furthermore, since we are going to the effort of capturing the vast majority of features in the flow, we use explicit time-stepping to capture the time evolution of these features. Implicit methods, in addition to being memory intensive and complicated to implement for high-order schemes, would not greatly increase the maximum allowable time step given the small time scales that must be captured to preserve accuracy.

The spatial discretization method used here is the flux reconstruction (FR) approach, a unifying framework encompassing a variety of high-order methods, including the discontinuous Galerkin (DG) and spectral difference (SD) methods [16]. The method uses Lagrange polynomials to represent the solution within each element of the grid. However, in contrast to typical finite element methods, these polynomials are permitted to be *discontinuous* between elements. Within FR elements are coupled through the (approximate) solution of a Riemann problem. This yields a common numerical flux at element interfaces which is then lifted into the interior of elements via specially defined ‘correction functions’. The method is covered in more detail in other works [2, 6, 31, 32, 34], with an excellent summary being given in [38].

High-order methods such as FR are particularly attractive for use in a DNS or ILES setting, and in fact have been shown to give remarkably good results in this context [11,

30, 33]. Not only are they less dissipative, enabling the simulation of vortex-dominated flows with fewer degrees of freedom than a lower-order method, they are also far better suited for utilizing modern hardware than traditional second-order CFD methods. This is due to the number of floating-point operations (FLOPs) performed per byte of memory accessed for each algorithm: an optimized second-order finite-volume solver will achieve less than 3% of peak performance on graphical processing units (GPUs) (which have an available FLOPs-to-bytes ratio of over 5) [19], while a high-order discontinuous finite-element method (DFEM) is able to achieve over 50% of peak performance on the same hardware [33].

The novelty of our current approach lies in the use of these high-order methods on overset grids. In the overset approach, a numerical grid is created around each body of interest, then the grids are patched together into a coherent whole through a process termed ‘domain connectivity’. The key benefits of the approach are simplified mesh generation and the ability to easily handle multiple objects in relative motion, such as the main rotor blades, tail rotor, and fuselage of a helicopter [35, 36]. While traditional overset methods lose much accuracy at the boundaries between grids through the use of simple, low-order interpolation methods, the present approach maintains high-order accuracy across overset boundaries [9–11, 13], and is hence better at preserving complex, vortex-dominated flows such as those around rotorcraft, high-lift wing systems, and in the present case, spinning golf balls.

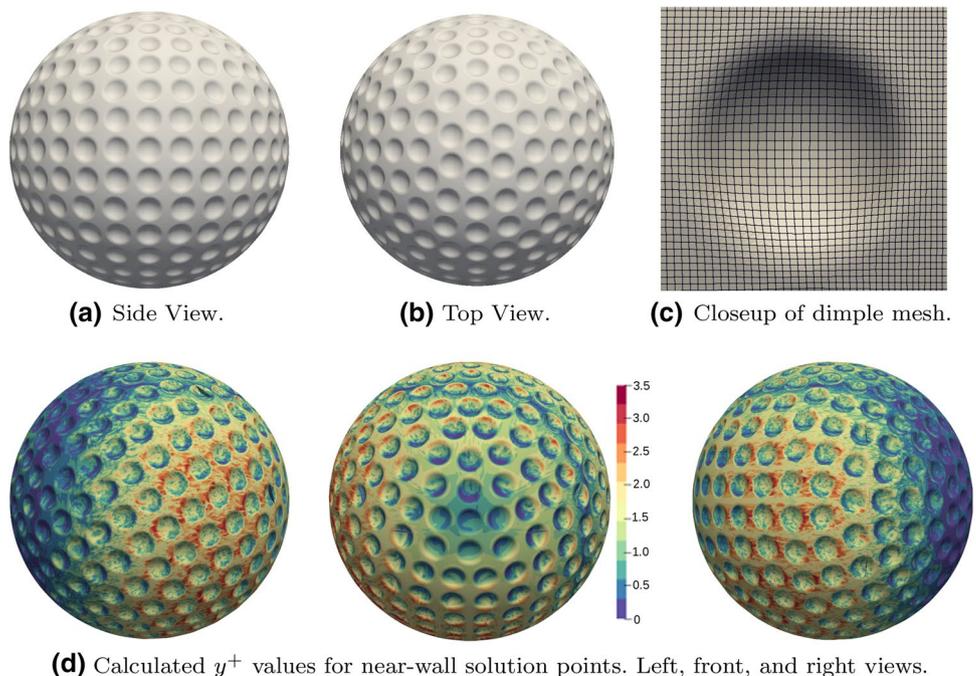
Additional consideration was given to efficient execution on modern, highly parallel computing platforms such

as NVIDIA graphical processing units (GPUs). Six out of the top ten fastest supercomputers in the world have either Intel Xeon Phi or NVIDIA Tesla accelerators as of November 2017, and the use of accelerators is continuing to grow. Since accelerators use different architectures than traditional CPUs, new approaches and algorithms are required to make efficient use of their available computing power. To enable a degree of performance good enough to solve large cases on accelerators then, a new method for domain connectivity was developed and implemented in the overset connectivity library TIOGA [25], and combined with our in-house FR solver, ZEFR [24]. The performance of the combined system on static and moving grids has previously been shown to be high enough to solve large-scale flow physics problems on multi-grid systems in a reasonable amount of time [10, 11].

2.2 Golf ball geometry

In this study, the golf ball surface geometry was created as a parameterized CAD model with 19 rows of circular dimples (9 rows per hemisphere + 34 dimples around centerline), for a total of 388 dimples, as shown in Fig. 1. The golf ball diameter was 42.7 mm, the dimple depth was 6.41×10^{-4} m ($k/D = 0.015$), and the dimple diameter was a constant 2.99 mm ($c/D = 7.0 \times 10^{-2}$). The dimple edges were filleted with a radius of 0.75 mm. The surface was exported in the STL format and used within the multiblock structured mesh generator GridPro [12] to create a spherical grid with a boundary layer. The surface of the golf ball was divided into 24 roughly square regions, each with a resolution of 144×144 quadrilaterals, with 60 layers in the radial

Fig. 1 Golf ball surface and grid resolution. **c** The pre-agglomerated element sizes; the actual surface resolution of the golf ball is 4/3 times greater than that shown. The color mapping in **d** ranges from $y^+ = 0$ to $y^+ = 3.5$



direction, for a total of 29,859,840 linear hexahedra, or 1,105,920 cubically curved hexahedra after agglomeration. The first cell height was chosen to be at an estimated y^+ value of 6.667 (3.4×10^{-5} m), the first 18 layers were held to a constant thickness, and the remaining 42 layers were allowed to grow out to a final outer diameter of 31.82 mm. The first cell height was chosen such that after agglomeration into cubically curved hexahedra and run with 4th order tensor-product solution polynomials, the first solution point inside the element would lie at a y^+ of approximately 1. The surface mesh resolution was chosen to match the recommendations of Li et al. [20], which were based upon recommendations from Muto et al. [23]. The golf ball grid of Li used a surface resolution of less than $\delta_B/2$, where $\delta_B = 3\sqrt{D\nu/(2V)}$ is the estimated laminar boundary layer thickness 90° from the stagnation point [23] (D is the diameter, ν is the kinematic viscosity, and V is the freestream velocity). Here, with $\delta_B \approx 2 \times 10^{-4}$ m, our surface mesh resolution at the level of the linear grid was slightly more than $\delta_B/2$, with the final resolution being slightly less than $\delta_B/2$ once the high-order polynomials are introduced into the agglomerated hexahedra. Fig. 1d shows the actual distribution of y^+ for the solution points nearest the wall during the simulation, computed using the first cell height of 0.12 mm and $p = 3$ Gauss–Legendre solution point locations. Although the grid was originally designed to be run using 4th order instead of 3rd order polynomials, the maximum y^+ value was 3.5, and the average value over the entire surface was 1.

The mesh was output in the CGNS structured multiblock format and imported into HOPR (High-Order Pre-Processor) [15], a utility which can agglomerate the cells of a structured mesh into high-order curved hexahedra. The new high-order mesh, in an HDF5-based HOPR-specific format, was then converted into the PyFR mesh format [37], which ZEFR has the capability to read. Finally, all grids were partitioned for parallel computation using METIS [17].

The pseudo-structured golf ball grid was then combined with a mostly Cartesian background grid created in Gmsh [14] to fill the desired extents of the full computational domain. The box has a width and height of 0.6832 m (16 times the golf ball diameter D), and length 1.0248 m ($-12D$ to $12D$). A refined region was created in the area to be occupied by the golf ball, with a refined wake region stretching out to the rear of the domain for a total of 715,750 linear hexahedra elements.

3 Static golf ball

The simulation was advanced in time using an explicit 5-stage 4th-order Runge–Kutta time integration method. Specifically, a low-storage method with an embedded 3rd order method and error estimate (RK4(3)5[2R+J]) [18] was

Table 1 Simulation conditions for all golf ball test cases

Reynolds number	150,000	ρ	1.0 kg/m ³
Mach	0.2	V	1.0 m/s
Prandtl	0.72	P	17.85714286 Pa
γ	1.4	R	17.85714286 J/(Kg K)
L	0.0427 m	T	1 K

The freestream quantities have been scaled such that no further non-dimensionalization is required for accuracy within the solver

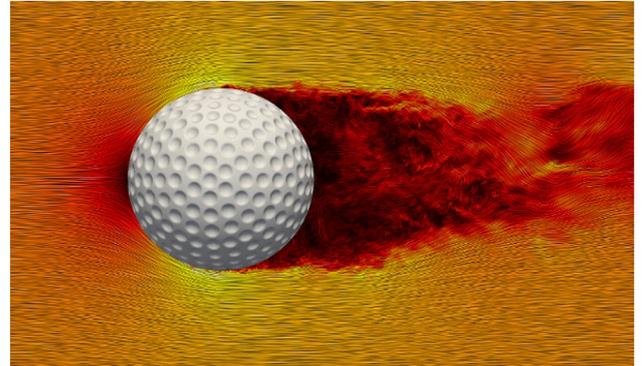
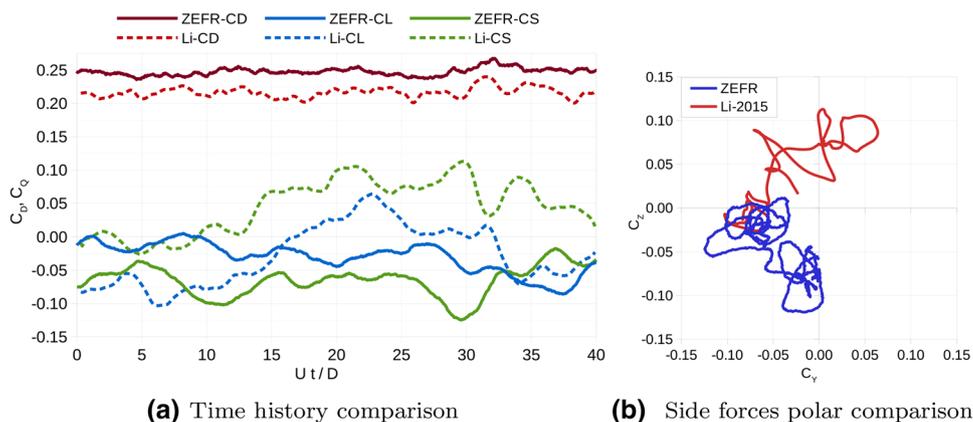


Fig. 2 View of approximate streamlines and velocity magnitude field through the $y = 0$ plane (golf ball centerline)

used to enable adaptive time stepping with a PI-type controller [5]. Third-order solution polynomials were utilized, as 4th order polynomials resulted in too restrictive a time step on this grid to generate results in a reasonable amount of time. The flow is along the x axis, with a Reynolds number of 150,000 based upon the golf ball diameter of 0.0427 m and a Mach number of 0.2. The full physical freestream conditions used (scaled such that the freestream velocity is 1) are shown in Table 1. An instantaneous view of velocity contours and approximate streamlines in the mid plane of the ball are shown in Fig. 2. The time histories of the drag and both side forces are shown in Fig. 3a, and a polar plot of the two side forces C_y and C_z are shown in Fig. 3b.

As a verification that our results are correct, Fig. 3 also plots our force coefficient histories against those generated by Li et al. for a very similar case. The conditions for their study were $Re = 110,000$ incompressible flow: a lower Reynolds number than that used here, but still corresponding to the supercritical regime where the drag coefficient should remain nearly constant. A second-order finite-volume LES solver was utilized for their simulation, using implicit time-stepping. Their golf ball had 392 dimples with a dimensionless diameter $c/D = 9.0 \times 10^{-2}$ and depth $k/D = 0.005$. An unstructured prism/tetrahedron grid was used with a total of approximately 1.45×10^6 elements in the domain, with overall extents $-13D \leq x \leq 13D$ and $-5.6D \leq y, z \leq 5.6D$.

Fig. 3 Force coefficients for the static golf ball simulation, compared with those produced by Li et al. under similar conditions



Since the dimples primarily change the drag, it would be expected that the side forces should be quite similar between the two cases. Indeed, that is the case as shown in Fig. 3a; the average and standard deviation of the side force histories are nearly identical. The side-force polar plot in Fig. 3b shows this as well; the two studies show similar trajectories, simply offset by a rotation about the axis of the flow. The present study was run for much longer (100 passes vs. 40), leading to a more visibly bimodal polar plot, but the trends remain the same. The drag histories are also in agreement; the offset between the two is to be expected, as the dimple depth used here is far greater than the dimple depth used by Li et al. Results from a variety of studies have shown a direct correlation between dimple depth and supercritical drag coefficient, along with an inverse correlation to the critical Reynolds number [8, 29].

In Figs. 4 and 5, we confirm the results of others in showing that the mechanism of transition is the growth of instabilities in the shear layer which forms over the recirculation regions inside the dimples not far from the stagnation point. Figure 4 shows contours of instantaneous vorticity magnitude through the $y = 0$ plane (mid plane of the golf ball). Separation and reattachment can be seen in dimples near the stagnation point, then near the top of the image, the shear layer breaks down and becomes turbulent. Figure 5 shows the same from a view above the centerline of the golf ball, with the three lines of dimples around $y = 0$ shown looking down at the stagnation point. Figure 5a shows the strip for $z > 0$ (the flow is roughly radially symmetric from the stagnation point). The dimples near the stagnation point have well-defined reattachment areas; but when the flow reaches the next dimple, instabilities are visible in the shear layer over the dimple, clearly seen in Fig. 5b. Over the 4th row of dimples (beginning at 31° from the stagnation point), the

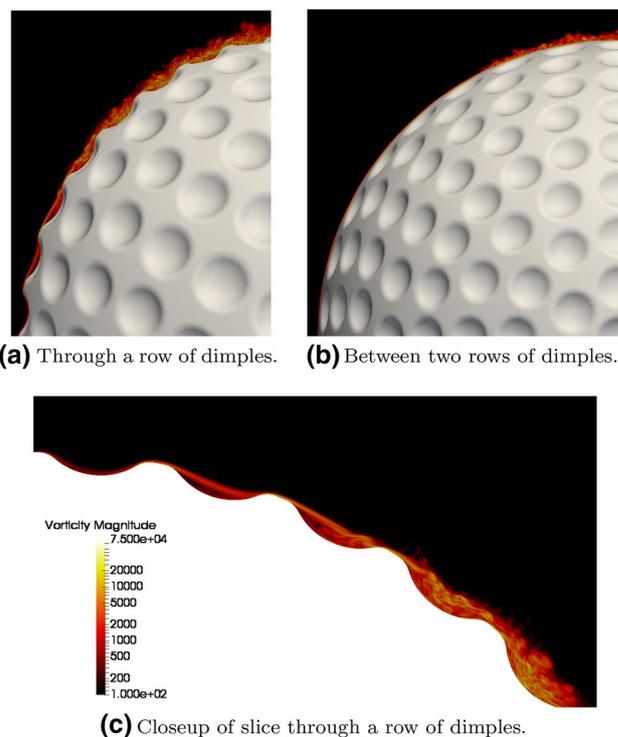
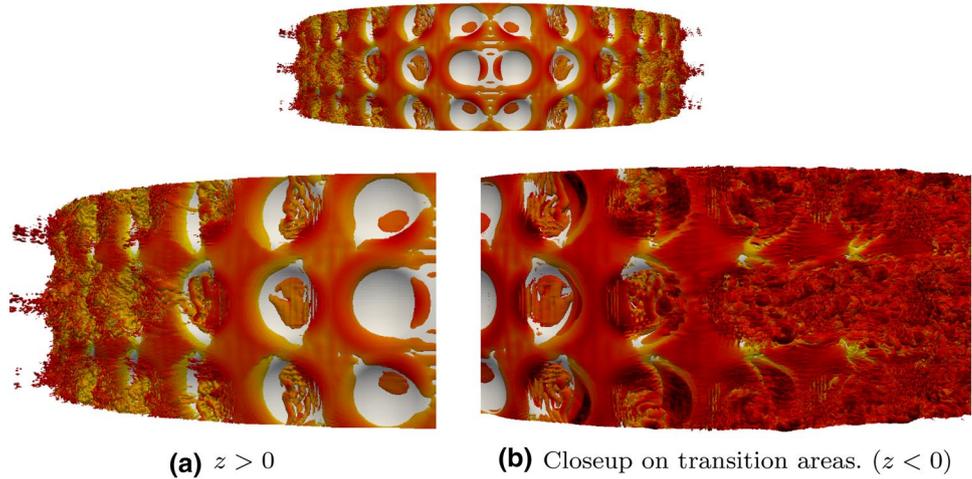


Fig. 4 Closeup view of log of vorticity magnitude through several slices near or at the $y = 0$ plane, showing boundary layer transition occurring in the shear layer above a dimple. In **b**, the flow remains laminar until the point of transition is clearly visible after the 6th dimple from the stagnation point. In **c**, the stagnation point is at the top left of the image

shear layer has broken down and become fully turbulent. The turbulent region then spreads out downstream from each dimple until the entire flow becomes turbulent by $\sim 63^\circ$. Figure 4b shows the point at which the smooth line between two rows of dimples becomes turbulent.

Fig. 5 Isosurfaces of Q -criterion colored by log of vorticity magnitude in a strip along the $y = 0$ centerline of the golf ball. Note that the location of transition is clearly visible as starting from the instabilities in the free shear layer over the dimples not far from the stagnation point



4 Spinning golf ball

We next move on to the case of a spinning golf ball. We keep the golf ball fixed at the origin, but apply a constant rotation rate around the z axis; to fall in line with other studies, we choose a non-dimensional spin rate $\Gamma = \omega r / U_{\text{inf}} = 0.15$. All other physical flow parameters are left the same. To handle the movement of the golf ball grid, we use the arbitrary Lagrangian–Eulerian (ALE) form of the Navier–Stokes equations, which involves adding an extra term to the convective fluxes due to the grid velocity. Only a constant rotation rate around one axis is used, but future work may involve using our 6 degree of freedom capability to simulate the entire trajectory of the golf ball under the influence of its surface forces and moments.

Our average C_D and C_L values are compared against the results from a number of other studies, both experimental and computational, in Fig. 6. As expected from previous literature, the spin induces a slightly higher drag coefficient than the static case but imparts a more regular variation in side forces upon the golf ball; the averages

for all force coefficients (with standard deviations) are summarized in Table 2 and the time history is shown in Fig. 7. While the out-of-plane side force (C_Z) hovers near zero, the lift (C_L or C_Y) hovers around a value of 0.16, with relatively large low-frequency oscillations. However, looking at a polar plot of the side forces, the oscillations are far more constrained than in the static case, where the symmetric nature of the flow allows the wake to oscillate randomly with no preferred direction. In addition to providing a sizable lift force, the spin has the effect of imposing some structure and a more preferred direction to the oscillations of the wake.

As was done with the case of the static golf ball, we may compare against the prior results of Li et al. [21], who have also performed detailed LES calculations of golf balls at a similar Reynolds number (110,000 vs. 150,000) and spin rate (0.1 vs. 0.15). The time histories and polar plot of the force coefficient results from Li are plotted alongside the present results in Fig. 7. The polar plots from the same cases without spinning are also included for comparison. As in the static case, the present drag value is larger, as should be expected; our present value is about 5% larger than that of the static golf ball. The average lift value is also larger,

Fig. 6 Comparison of C_D and C_L values from several experimental and computational studies alongside the present results produced with ZEFR

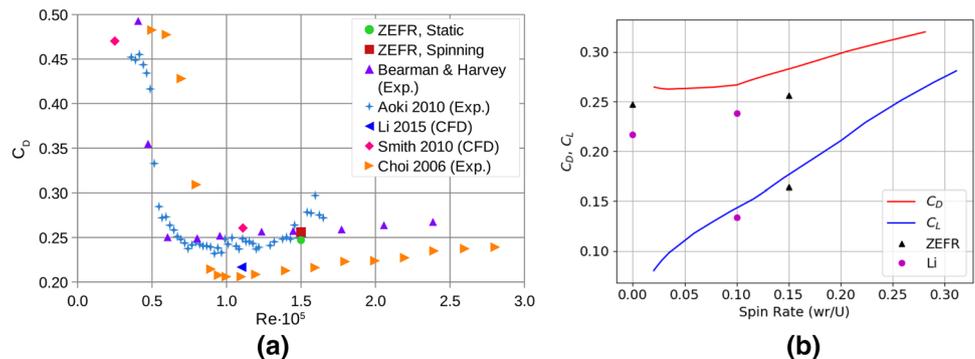


Fig. 7 Comparison of force time histories for our spinning golf ball vs. the results from Li et al. [21]. The dashed curves in **b** re-plot the same polar data from the static golf ball cases

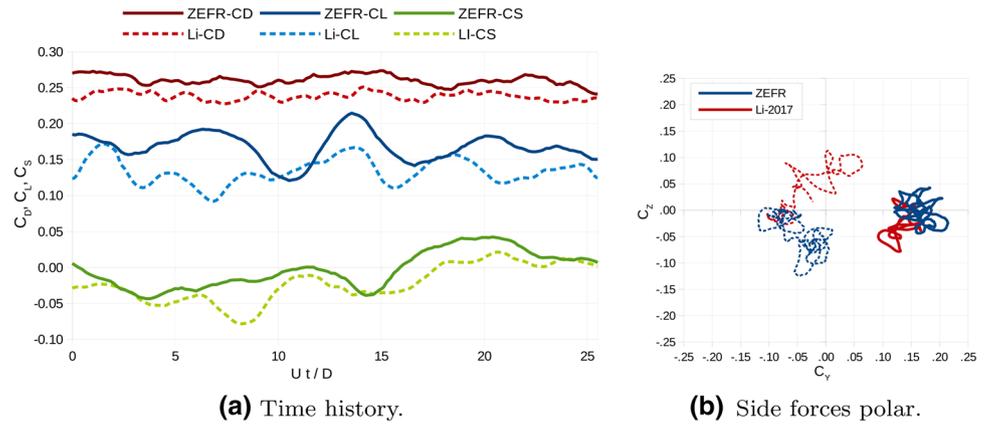


Fig. 8 Comparisons of power spectrum density for our spinning golf ball vs. the results from Li et al. [21]

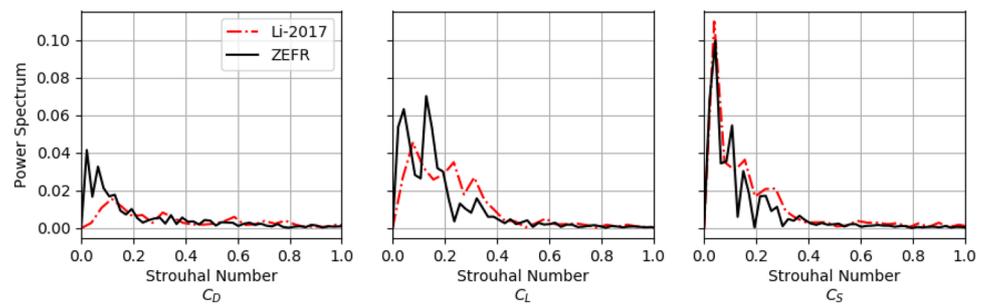


Table 2 Summary of average force coefficients for the static and spinning golf balls

	Static	Spinning ($\Gamma = 0.15$)	Li et al. [20] (static)	Li et al. [21] ($\Gamma = 0.1$)
C_D	0.2469 ± 0.005	0.256 ± 0.010	0.217 ± 0.008	0.238 ± 0.0057
C_Q	0.076 ± 0.020	0.165 ± 0.021	0.079 ± 0.019	0.190 ± 0.025
C_Y	-0.047 ± 0.032	0.164 ± 0.021	-0.029 ± 0.045	0.134 ± 0.018
C_Z	-0.044 ± 0.032	0.002 ± 0.022	0.046 ± 0.040	-0.022 ± 0.026

C_Q refers to the combined magnitude of the lift and side forces C_Y and C_Z . Present results compared to the similar study from Li et al.

as is also expected due to the higher spin rate used in the present study.

The present values of lift and drag coefficients also agree well with those of Bearman and Harvery. Using the data shown in Fig. 6b, the estimated C_D for a conventional golf ball at a nondimensional spin rate of 0.15 would be about 0.28, or 8% higher than that of a static golf ball, with a lift coefficient of about 0.18. Our results align more with their results at a spin rate of 0.13, with a lift coefficient of 0.16 and a spinning-to-static C_D increase of 5%. The results of Li et al., meanwhile, predict a much higher rise in C_D with respect to spin rate, with a change of 11% at a spin rate of 0.1.

We can also more quantitatively compare our results to those of Li by using the power spectrum of the golf ball forces, shown in Fig. 8. While the two side force spectra appear nearly identical, the lift and drag coefficients show

slight differences. In particular, the drag spectra show a difference in peak, with the present results showing more low-frequency components. Similarly, the present lift coefficient spectra also show a slight shift to lower frequencies for the two primary frequencies which appear, although the two smaller peaks are in the same locations as those of Li.

While the overall flow features are quite similar from the static to the spinning golf ball, a few differences can clearly be seen. Figure 9 shows how the locations of transition shift; Fig. 10 likewise shows a direct comparison between flow features near the stagnation points of the static and spinning cases. On the advancing side of the stagnation point, both figures show that the increased relative velocity over the surface of the golf ball leads to earlier transition; conversely, on the retreating side, the transition is slightly delayed relative to the static case.

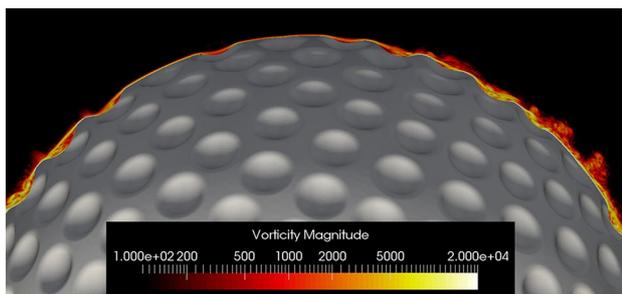


Fig. 9 Closeup view of $z = 0$ centerplane showing log of vorticity magnitude. The stagnation point is at the top of the image and the ball is spinning counter-clockwise

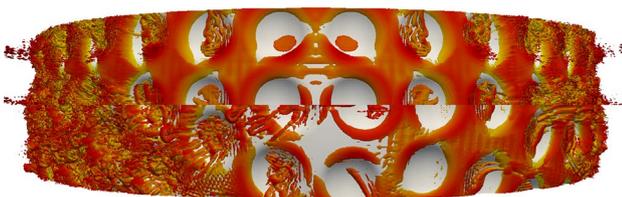


Fig. 10 Isosurfaces of Q -criterion colored by log of vorticity magnitude along the centerline of the golf ball. Top: static golf ball, bottom: spinning golf ball at $T = 49$ s

5 Conclusions

In this work, the fluid dynamics of static and spinning golf balls were simulated using a high-order numerical scheme on overset grids. These are the first simulations of static and spinning golf balls using high-order numerical methods and represent an advance in the state of the art in both scale-resolving CFD and overset grid calculations. New algorithms were developed to leverage the modern hardware accelerators becoming increasingly common on large-scale computing clusters. These new algorithms allow the application of moving overset grids to high-order methods to solve large-scale fluid physics problems in a reasonable amount of time. By comparing to a variety of previous experimental and computational studies, we have confidence that our present approach is able to accurately predict the complex, turbulent flow fields around golf balls and other sports balls, which operate at modest Reynolds numbers of less than 500,000. Other sports applications are also within the reach of high-fidelity simulation using our methods, such as hockey pucks, small or slow-speed sailboats, and bicycles or cyclists at modest speeds.

Acknowledgements The authors would like to acknowledge the Army Aviation Development Directorate (AMRDEC) for providing funding for this research under the oversight of Roger Strawn, the Air Force Office of Scientific Research for their support under Grant FA9550-14-1-0186 under the oversight of Jean-Luc Cambier, and Margot Gerritsen for access to the XStream GPU computing cluster, which is supported

by the National Science Foundation Major Research Instrumentation program (ACI-1429830). We would also like to thank Dr. Peter Eiseinan for providing academic licensing to the GridPro meshing software and assisting with the creation of several golf ball grids. Last, we would like to thank Dr. Jay Sitaraman for his expertise and help on overset connectivity methods, and his help in ensuring our numerical methods were robust enough for broad applicability.

References

1. Aoki K, Muto K, Okanaga H (2010) Aerodynamic characteristics and flow pattern of a golf ball with rotation. *Procedia Eng* 2:2431–2436
2. Asthana K, Jameson A (2014) High-order flux reconstruction schemes with minimal dispersion and dissipation. *J Sci Comput* 62:913–944
3. Bearman PW, Harvey JK (1976) Golf ball aerodynamics. *Aeronaut Q* 27:112–122
4. Beratlis N, Squires K, Belaraz E (2012) Numerical investigation of Magnus effect on dimpled spheres. *J Turbul* 13(15):1–15. <https://doi.org/10.1080/14685248.2012.676182>
5. Butcher JC (2016) Numerical methods for ordinary differential equations, 3rd edn. Wiley, New York
6. Castonguay P (2012) High-order energy stable flux reconstruction schemes for fluid flow simulations on unstructured grids. Ph.D. thesis, Stanford University
7. Choi J, Jeon WP, Choi H (2006) Mechanism of drag reduction by dimples on a sphere. *Phys Fluids* 18:041702
8. Chowdhury H, Loganathan B, Wang Y, Mustary I, Alam F (2016) A study of dimple characteristics on golf ball drag. *Procedia Eng* 147:87–91
9. Crabill J, Jameson A, Sitaraman J (2016) A high-order overset method on moving and deforming grids. In: AIAA modeling and simulation technologies conference (AIAA2016-3225). <https://doi.org/10.2514/6.2016-3225>
10. Crabill JA (2018) Towards industry-ready high-order overset methods on modern hardware. Ph.D. thesis, Stanford University
11. Crabill JA, Witherden FD, Jameson A (2018) A parallel direct cut algorithm for high-order overset methods with application to a spinning golf ball. *J Comput Phys* 374:692–723
12. Eiseinan PR, Ebenezer SJ, Sudharshanam V, Anbumani V (2016) GridPro manuals, version 2.0. <http://www.gridpro.com>. Accessed 2 Sept 2016
13. Galbraith MC (2013) A discontinuous galerkin overset solver. Ph.D. thesis, University of Cincinnati
14. Geuzaine C, Remacle JF (2009) Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *Int J Numer Methods Eng* 79:1309–1331 (**Mesh-Generation Software**)
15. Hindenlang F, Bolemann T, Munz CD (2015) Mesh curving techniques for high order discontinuous Galerkin simulations. In: Kroll N, Hirsch C, Bassi F, Johnston C, Hillewaert K (eds) IDIHOM: industrialization of high-order methods—a top-down approach. Springer, Berlin, pp 133–152
16. Huynh HT (2007) A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods. In: 47th AIAA aerospace sciences meeting
17. Karypis G, Kumar V (1998) A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J Sci Comput* 20(1):359–392
18. Kennedy CA, Carpenter MH, Lewis RM (2000) Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations. *Appl Numer Math* 35(3):177–219

19. Langguth J, Wu N, Cai JCX (2013) On the GPU performance of cell-centered finite volume method over unstructured tetrahedral meshes. In: S. SC13 International conference for high performance computing networking, analysis (eds.) *IA³: Workshop on irregular applications: architectures and algorithms*. ACM, Denver, CO, USA. <https://doi.org/10.1145/2535753.2535765>
20. Li J, Tsubokura M, Tsunoda M (2015) Numerical investigation of the flow around a golf ball at around the critical Reynolds number and its comparison with a smooth sphere. *Flow Turbul Combust* 95:415–436. <https://doi.org/10.1007/s10494-015-9630-4>
21. Li J, Tsubokura M, Tsunoda M (2017) Numerical investigation of the flow past a rotating golf ball and its comparison with a rotating smooth sphere. *Flow Turbul Combust* 99:837–864. <https://doi.org/10.1007/s10494-017-9859-1>
22. Mehta RD (1985) Aerodynamics of sports balls. *Ann Rev Fluid Mech* 17:151–189
23. Muto M, Tsubokura M, Oshima N (2012) Negative Magnus lift on a rotating sphere at around the critical Reynolds number. *Phys Fluids* 24:014102
24. Romero JD (2017) On the development of the direct flux reconstruction scheme for high-order fluid flow simulations. Ph.D. thesis, Stanford University
25. Sitaraman J (2015) TIOGA: topology independent overset grid assembly library. <https://github.com/jsitaraman/tioga> (**Overset Connectivity Library for CFD**). Accessed 26 June 2018
26. Smith CE, Beratlis N, Balaras E, Squires K, Tsunoda M (2010) Numerical investigation of the flow over a golf ball in the subcritical and supercritical regimes. *Int J Heat Fluid Flow* 31:262–273. <https://doi.org/10.1016/j.ijheatfluidflow.2010.01.002>
27. Smits AJ, Smith DR (1994) A new aerodynamic model of a golf ball in flight. In: Cochran AJ, Farrally MR (eds) *Science and golf II: proceedings of the world scientific congress of golf*. E & FN Spon, London, pp 340–347
28. Ting LL (2002) Application of CFD technology analyzing the three-dimensional aerodynamic behavior of dimpled golf balls. In: ASME international mechanical engineering congress & exposition
29. Ting LL (2003) Effects of dimple size and depth on golf ball aerodynamics. In: 4th ASME/JSME joint fluids engineering conference
30. Vermeire BC, Witherden FD, Vincent PE (2017) On the utility of GPU accelerated high-order methods for unsteady flow simulations: a comparison with industry-standard tools. *J Comput Phys* 334:497–521
31. Vincent P, Castonguay P, Jameson A (2010) A new class of high-order energy stable flux reconstruction schemes. *J Sci Comput* 47:50–72
32. Vincent P, Castonguay P, Jameson A (2011) Insights from von Neumann analysis of high-order flux reconstruction schemes. *J Comput Phys* 230:8134–8154
33. Vincent P, Witherden F, Vermeire B, Park JS, Iyer A (2016) Towards green aviation with Python at petascale. In: *Proceedings of the international conference for high performance computing, networking, storage and analysis, SC '16*. IEEE Press, Piscataway, NJ, USA. <http://dl.acm.org/citation.cfm?id=3014904.3014906>
34. Williams DM, Castonguay P, Vincent PE, Jameson A (2013) Energy stable flux reconstruction schemes for advection–diffusion problems on triangles. *J Comput Phys* 230:8134–8154
35. Wissink A (2012) Helios solver developments including strand meshes. In: Oral presentation, 11th symposium on overset composite grids and solution technology
36. Wissink A (2012) An overset dual-mesh solver for computational fluid dynamics. In: 7th International conference on computational fluid dynamics (ICCFD7)
37. Witherden FD, Farrington AM, Vincent PE (2014) PyFR: an open source framework for solving advection-diffusion type problems on streaming architectures using the flux reconstruction approach. *Comput Phys Commun* 185:3028–3040
38. Witherden FD, Vincent PE, Jameson A (2016) High-order flux reconstruction schemes. In: *Handbook of numerical analysis*, vol 17, pp 227–2633. Elsevier, Boca Raton. <https://doi.org/10.1016/bs.hna.2016.09.010> (**chap. 10**)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.