Journal of Computational Science 1 (2010) 13–23

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

Unstable periodic orbits in weak turbulence

L. Fazendeiro [a], B.M. Boghosian [b], P.V. Coveney [a,*], J. Lätt [c]

[a] Centre for Computational Science, University College of London, 20 Gordon Street, London WC1H 0AJ, UK
[b] Department of Mathematics, Tufts University, Bromfield-Pearson Hall, Medford, MA 02155, USA
[c] Institute of Mechanical Engineering, École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

ARTICLE INFO

Article history:
Received 15 December 2009
Received in revised form 30 March 2010
Accepted 30 March 2010

PACS:
47.27.E−
05.45.−a
47.11.−j

Keywords:
Fluid turbulence
Periodic theory
Petascale computing
Lattice-Boltzmann

ABSTRACT

We report on a sophisticated numerical study of a parallel space–time algorithm for the computation of periodic solutions of the driven, incompressible Navier–Stokes equations in the turbulent regime. Efforts to apply the machinery of dynamical systems theory to fluid turbulence depend on the ability to accurately and reliably compute such unstable periodic orbits (UPOs). For example, the UPOs may be used to construct the dynamical zeta function of the system, from which very accurate turbulent averages of observables may be extracted.

Though a number of algorithms for computing such orbits have been proposed and tested, in this paper we focus on a space–time variational principle introduced by Lan and Cvitanović in 2004 [15]. This method has not, to our knowledge, been tested on dynamical systems of high dimension because of the formidable storage and computation required. In this paper, we use petascale computation to apply this algorithm to weak hydrodynamic turbulence.

We begin with a brief description and reformulation of the space–time algorithm of Lan and Cvitanović. We then describe how to apply this algorithm to the lattice-Boltzmann method for the solution of the Navier–Stokes equations. In particular, we describe the fully parallel implementation of this algorithm using the Message Passing Interface. This implementation, called HYPO4D, has been successfully deployed on a large variety of platforms both in the UK and the USA and has shown very good scalability to tens of thousands of computing cores.

Finally, we describe the application of this implementation to the problem of weak homogeneous turbulence driven by an Arnold–Beltrami–Childress force field in three spatial dimensions, at a Reynolds number of 371. We commence by systematically searching for nearly periodic orbits as candidate solutions from which to begin the relaxation; we then apply the variational algorithm until convergence is obtained. Because the algorithm requires storage of the space–time lattice, even the smallest orbits require resources on the order of tens of thousands of computing cores. Using this approach, two UPOs have been identified and some of their properties have been analysed.

1. Introduction

Turbulence has been called the last unsolved problem of classical physics. Turbulent fluids are well known to exhibit large non-normal fluctuations from equilibrium and macroscopic spatio-temporal structure. Although the Navier–Stokes equations [1], (NSE), which describe the behaviour of a turbulent viscous fluid, have been known for more than a century, very few exact results about turbulence have followed from them [2–5]. In particular, we still lack the ability to make *a priori* predictions of the turbulent averages of observables.

The state space of a driven, viscous, incompressible fluid in the turbulent regime is the infinite-dimensional function space of all divergenceless vector fields. The dynamics of the Navier–Stokes equations define a trajectory in this space. In the long-time limit, this trajectory settles onto a finite-dimensional attracting set [6], and the turbulent average of an observable may be thought of as its integral over a *natural measure* on this attracting set.

Strange attracting sets with highly structured natural measures are understood from low-dimensional examples, such as the Lorenz attractor [7]. In such systems, the attracting sets are replete with unstable periodic orbits (UPOs). These comprise a set of measure zero that is nonetheless dense in the attracting set. That is, the attracting set can be thought of as the closure of the set of all UPOs. The countable sequence of UPOs provide a useful characterization of the structure and dynamics of the attractor [8].

* Corresponding author. Tel.: +44 20 7679 4560.
E-mail addresses: l.fazendeiro@ucl.ac.uk (L. Fazendeiro), bruce.boghosian@tufts.edu (B.M. Boghosian), p.v.coveney@ucl.ac.uk (P.V. Coveney), jonas.latt@epfl.ch (J. Lätt).

1877-7503/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jocs.2010.03.004

The natural measure of the attractor is an eigenfunction of the Frobenius–Perron operator, $\mathcal{L}$, of the dynamics under consideration. The characteristic equation of this operator is one way of defining the dynamical zeta function,

$$\zeta(z) = \frac{1}{\det(\mathcal{I} - z\mathcal{L})},$$

where $\mathcal{I}$ is the identity. The location of the poles of $\zeta(z)$, as well as closely related functions, may be used to extract turbulent averages of observables and their correlations [9]. As is the case for other zeta functions, most notably the Riemann zeta function, $\zeta(z)$ can be expressed as an infinite product over the prime periodic orbits of the corresponding dynamical system. For many systems, a good approximation to the dynamical zeta function (DZF) can be constructed by using the lowest-period UPOs of the flow, determined numerically [8]. As described in detail by Cvitanović et al., knowledge of a finite set of low-period UPOs is often sufficient to estimate statistical averages over the natural measure of the attracting set. More recently, Kawasaki and Sasa [10] have argued that even a single UPO, with a large period, might suffice to characterize high-dimensional chaotic dynamical systems.

There are a number of potential advantages to this approach. The averages obtained by using the DZF formalism can be very accurate, particularly if the symbolic dynamics of the system in question are known [11]. Averages thus obtained are not stochastic in nature, as would be those obtained from a numerical time average; their statistical errors do not decay as the inverse square root of the number of decorrelation times. Moreover, by maintaining a library of low-period UPOs for a specified dynamical system, turbulent averages of observables may be computed without the need to perform additional simulations.

The main challenge then lies in developing efficient algorithms to compute the prime periodic orbits of the dynamical system under study. These cannot be captured by a simple forward-time integration due to their instability. Over the years, many methods have been proposed for this. Some are "shooting methods" that begin with orbits that do not quite close and use a Newton–Raphson method to close them. In recent years, such methods have been employed to facilitate the numerical discovery of time-periodic solutions in plane Couette flow [12], in isotropic turbulence [13], and in shell models of turbulence [14].

In 2004, a new variational principle for finding UPOs was proposed by Lan and Cvitanović [15]. Whereas the shooting method begins with an orbit that satisfies the dynamical equations but does not close, the variational procedure begins with an orbit that is closed but does not satisfy the dynamical equations. It then relaxes this orbit to a periodic solution by minimizing the norm of the residual of the dynamical equations themselves. Because this method calls for the storage of the entire orbit in space–time, it requires high performance computing resources. It is perhaps for this reason that it has received less attention than shooting methods.

In the present paper we reformulate the variational principle approach for finding UPOs, and apply it to the lattice-Boltzmann method [16] (LBM) for simulating the NSE. This is a quasi-compressible numerical method (although it is often not characterized as such) that can be used to simulate the incompressible NSE when it is applied in the limit of low Mach number. The LBM is a highly scalable numerical algorithm, particularly suitable for deployment on very large computational resources, and can be straightforwardly folded into the variational method for finding UPOs.

The paper is organised as follows: In Section 2 we present our reformulation of the variational principle which will be applied to the LBM. Section 3 briefly describes the particular LBM model considered, its implementation and performance scaling on a range of supercomputing platforms. In this section we also report on the use of this model to simulate weakly-turbulent fluids, and demonstrate good agreement with results in the research literature. Section 4 describes the methodology that was used to identify suitable starting estimates for the space–time variational relaxation procedure. Finally, Section 5 presents the main results obtained, including a description of the UPOs discovered by this method. We conclude with a discussion of desirable future developments.

2. Variational principle

In this section we reformulate and discuss the space–time variational principle for the determination of periodic orbits in a specified dynamical system.

As noted in Section 1, most prior work on finding periodic orbits in continuous-time dynamical systems has employed shooting methods. To describe this procedure, let S be a codimension-one surface in the system's state space, Ω, through which the orbit is known to pass frequently. A point in S is chosen as an initial condition, $\mathbf{R}(0)$, and its time evolution is followed until it returns to the surface S at some later time, T. If we define the displacement $\delta = \mathbf{R}(T) - \mathbf{R}(0)$ then the idea is to vary the initial starting point, $\mathbf{R}(0)$, so as to make $\delta = 0$. This variation can be achieved through a multidimensional Newton–Raphson iteration method. Because the procedure is unlikely to converge unless the initial guess is already exceptionally close to a closed orbit [17], recent work by Viswanath has introduced the use of hook step methods to improve the robustness of the method [18].

While the storage requirements of the shooting method are not extraordinary, it does require very long computation times, since the NSE must be integrated forward in time at each step of the algorithm. Lan and Cvitanović's variational method requires instead to store an entire space–time trajectory and subsequently relax it towards a solution of the dynamical equations. This dramatically increases the storage requirements but offers the possibility of reduced computation time, since it allows one to reap the benefits of parallelism in time as well as space. In recent work [19], this method has been tested on the Lorenz equations [7] and other low-dimensional dynamical systems, and found to yield robust convergence to UPOs.

To understand the variational approach, suppose that our goal is to find a solution $\mathbf{R}(t)$ for a differential equation of the form

$$\dot{\mathbf{R}} = \mathbf{f}(\mathbf{R}), \tag{1}$$

where $\mathbf{R}$ is a point in the state space, $\mathbf{f}(\mathbf{R})$ is the vector field describing the dynamics, and the dot denotes differentiation with respect to time t. Writing T for the unknown period, the approach is to minimize

$$\mathcal{F}([\mathbf{R}], T) \equiv \frac{1}{2}\int_0^T dt\, |\dot{\mathbf{R}}(t) - \mathbf{f}(\mathbf{R}(t))|^2. \tag{2}$$

This is a function of the orbit $\mathbf{R}(t)$, and a function of the unknown period T. Note that $\mathcal{F} \geq 0$ from the definition, and $\mathcal{F} = 0$ only for solutions of Eq. (1). We propose to find periodic orbits of Eq. (1) by minimizing $\mathcal{F}$ with respect to both $\mathbf{R}(t)$ and T. A straightforward calculation yields the Fréchet derivative

$$\frac{\delta \mathcal{F}}{\delta \mathbf{R}(t)} = -\ddot{\mathbf{R}} - \left[\nabla \mathbf{f} - (\nabla \mathbf{f})^{\top}\right] \cdot \dot{\mathbf{R}} + \frac{1}{2}\nabla \left|\mathbf{f}\right|^2,$$

and the partial derivative

$$\frac{\partial \mathcal{F}}{\partial T} = \frac{1}{2}\left|\dot{\mathbf{R}}(T) - \mathbf{f}(\mathbf{R}(T))\right|^2 - \frac{1}{2}\left|\dot{\mathbf{R}}(0) - \mathbf{f}(\mathbf{R}(0))\right|^2$$

and both of these derivatives must vanish for a periodic orbit. Here we have used ∇ to denote derivatives with respect to the components of $\mathbf{R}$. Also, the superscript $\top$ denotes the transpose operator, not be confused with the period T.