Joint Contour Net analysis of lattice QCD data

Dean P. Thomas, Rita Borgo, Hamish Carr and Simon Hands

Abstract Lattice Quantum Chromodynamics (QCD) is an approach used by theoretical physicists to model the strong nuclear force. This works at the sub-nuclear scale to bind quarks together into hadrons including the proton and neutron. One of the long term goals in lattice QCD is to produce a phase diagram of QCD matter as thermodynamic control parameters temperature and baryon chemical potential are varied. The ability to predict critical points in the phase diagram, known as phase transitions, is one of the on-going challenges faced by domain scientists. In this work we consider how multivariate topological visualisation techniques can be applied to simulation data to help domain scientists predict the location of phase transitions. In the process it is intended that applying these techniques to lattice QCD will strengthen the interpretation of output from multivariate topological algorithms, including the joint contour net. Lattice QCD presents an interesting opportunity for using these techniques as it offers a rich array of interacting scalar fields for analysis; however, it also presents unique challenges due to its reliance on quantum mechanics to interpret the data.

Dean P. Thomas
Departments of Computer Science and Physics, College of Science, Swansea University, Swansea, UK e-mail: 798295@swansea.ac.uk

Rita Borgo
Informatics Department, Kings College London, London, UK e-mail: rita.borgo@kcl.ac.uk

Hamish Carr
School of Computing, University of Leeds, Leeds, UK e-mail: h.carr@leeds.ac.uk

Simon Hands
Department of Physics, College of Science, Swansea University, Swansea, UK e-mail: s.j.hands@swansea.ac.uk

1 Introduction

Multivariate topology, in contrast to univariate topology, is a relatively recent addition to the set of visualisation and analysis tools available to scientists. Many fields of science present scalar data obtained either from simulation or observation upon common sampling points. Whilst univariate topology can evaluate topological features in each field as unique entities, multivariate topology can further understanding of correlations between scalar fields.

Several different approaches exist for computing features in multi-fields including Jacobi Sets, Reeb graph comparison, Reeb spaces, and range tessellation. We concentrate on the Joint Contour Net [1], which builds upon existing theoretical and practical aspects of the Reeb graph to represent the Reeb space in a discrete graph based format. The Reeb space [2] addresses the relation between multiple sampled fields by contracting multivariate contours to singular points.

In this work we make the following contributions:

- Extend the use of the multivariate topological techniques to a new scientific domain, lattice Quantum Chromodynamics (QCD)
- Show how analysis of the Joint Contour Net using non-visual techniques can allow it to be used on large, complex data sets that are beyond the scope of visual inspection
- Investigate the use of multivariate persistence for predicting properties of lattice QCD data

2 Lattice Quantum Chromodynamics

Quantum Chromodynamics is the theory used to describe interactions between subnuclear particles that bind the hadron group of particles together. In its most basic form it describes quark particles and how they interact with one another via the exchange of massless gluon particles (Fig. 1). Hadrons exist in two states; those consisting of three quarks are known as baryons, and quark-antiquark configurations make up the meson group. Protons and neutrons are part of a subset of the baryon group called nucleons that combine with the electron particle to create the atoms that form the periodic table of elements.

The way in which quarks and gluons behave across a range of temperatures and chemical potentials can be modelled using a phase diagram (Fig. 2). The design is similar to that of water, and can be used to understand how states of matter change as chemical potential and temperature vary. A proposed model given in [3] shows that situations at neutron star cores are most similar to a system with a non-zero chemical potential. In addition, the diagram shows output from real-world experiments, allowing a sense of how work in this area fits in with experiments in nuclear and particle physics.

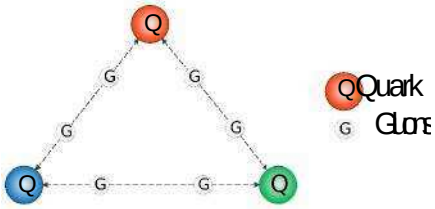

Fig. 1 Colour neutrality in a 3-quarksystem (Baryon).

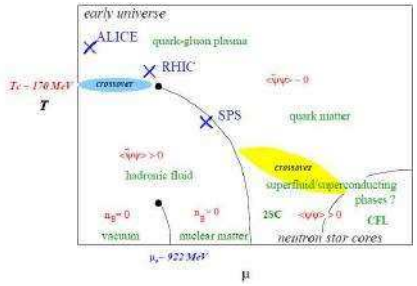

Fig. 2 A proposed phase diagram for QCD [3].

Quarks and gluons appear impossible to observe in an unbound state. A single quark has never been observed and they are always found in colour neutral bound states [4], held together by the exchange of gluons. This phenomenon is known as colour confinement and unlike its electromagnetic force equivalent the energy needed to separate two or more bound particles increases indefinitely with distance. Hence, in order to test theories of the quark model computer simulations are employed to model quark-gluon interactions. The process of computing Quantum Chromodynamics as discrete models is known as lattice QCD.

2.1 Lattice structure

Kenneth Wilson was the first physicist to suggest that QCD could be approximated on a discrete lattice to model properties of quark and gluon fields [5]. The structure of the lattice is a hyper-torus in Euclidean space-time, meaning that the three spatial dimensions and the time dimension are treated as equal and translationally-invariant boundary conditions $f(x) = f(x - L_x)$ are used so that all sites in the lattice are equivalent. As the lattice exists in four dimensions it can be convenient to think of the lattice of being made up of hyper-cubic cells, as shown in Figure 3.

Quarks are placed onto the lattice at positions with integer indices, referred to as sites. In the remainder of this work we will label sites on the lattice using the notation $U(X)$ where $X \in Z^4$. In Lattice QCD data is not placed on these sites, instead they represent the starting point for computations with a given origin. From each lattice site four link variables are used to model the gluon potential in the x, y, z and t directions between neighbouring sites. Link variables are represented as $U_\mu(X)$, i.e. the variable defined on the link emerging from the site X in the direction μ. When traversing the lattice, the relation $U_\mu(X) \equiv U_\mu(X - \hat{\mu})^\dagger$ allows us to define movements in the reverse direction using the adjoint (2) form of link variables.

Each link variable is a member of the special unitary group of matrices (1), identified using the notation $SU(n)$. The value of n represents the number of charge colours used in the gauge theory, with true QCD defined with $n = 3$. However, in

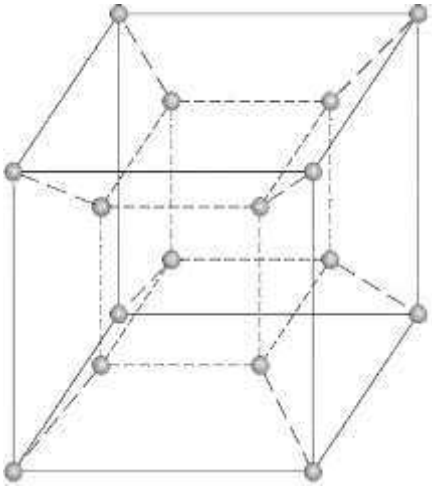

Fig. 3 Arrangement of lattice sites for one 4D cell (dashed line represents t dimension).

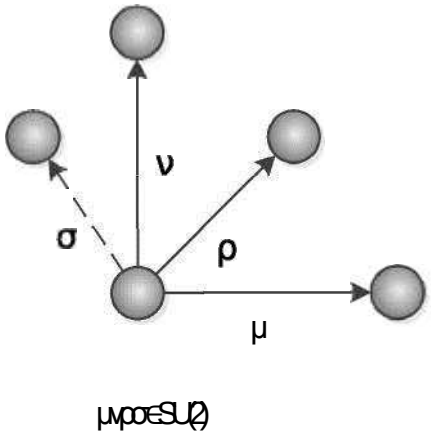

μμpoeSU(2)

Fig. 4 Placement of data is on the four lattice links from each site (U(X)).

this work we use a simplified two colour model of the theory using $SU(2)$ matrices. Colour is used in this context to parametrise the concept of colour neutrality, required to explain confinement (see Figure 1), rather than a parameter of appearance. One of the primary reasons for using a simplified model is that it allows us the freedom to vary the chemical potential of the system.

$$SU(2) = \left\{ \begin{bmatrix} a & -b \\ b & a \end{bmatrix} : a,b \in \mathbb{C}, |a|^2 + |b|^2 = 1, \det = 1 \right\} \quad (1)$$

where x represents the complex conjugate.

$$U^\dagger = \begin{bmatrix} a - ib & e - if \\ c - id & g - ih \end{bmatrix} \text{ with } U = \begin{bmatrix} a + ib & c + id \\ e + if & g + ih \end{bmatrix} \quad (2)$$

3 Topological Analysis

Topological analysis, in particular for univariate data, is an established technique of forming models of data for indirect volume rendering. For visualisation purposes it allows scalar fields to be presented in the context of connected regions, which can also be used to optimise the rendering process. Information regarding internal structure can also be relayed back to the user allowing further insights to be made about the data. It is this use of topology that is of most interest in the context of

forming insights into large data sets. This section introduces the most relevant topo-

logical structures to this work; first in the setting of univariate data, then extended to multivariate inputs.

3.1 Univariate Topology

The Reeb graph and closely associated contour tree are methods for reducing the topology of a scalar field to a graph structure. Topological events are mapped to vertices of a directed graph, with edges representing a continual deformation of a manifold between two critical vertices. Following the path of the graph allows tracking of splits and joins in the topology as the function value, or isovalue, is varied. In [6] the contour tree is used as the underlying data structure to store and generate object meshes using path seeds.

Whilst being quicker to compute than the Reeb graph, a limitation of the contour tree algorithm [7] is that it can only be applied to data without periodic boundaries. This limitation can be avoided by making symbolic cuts in the topology for the purposes of computing the contour tree. These cuts can then be stitched together to form the Reeb graph of a periodic domain using a technique used in [8] to accelerate Reeb graph computations.

Persistence

Univariate persistence is typically used in two scenarios; the first is an aid to simplify the Reeb graph (and contour tree) using geometric measures to determine noise from features [9]. The second is to provide additional information about distinct topological objects in a scalar volume. In this work we focus on the second approach as domain scientists have their own noise removal methods (cooling).

The contour spectrum [10] was introduced as a method of relaying quantitative information about individual contours in scalar data. Attributes that could be computed included surface area and volume of contours. In [11] isosurface statistics were directly compared against raw histograms of scalar data for a number of data sets. Measurements evaluated included the cell intersection count, triangle count and isosurface area. It was found that using these measures a truer distribution of the scalar field could be computed. An improvement was given in [12] using concepts from geometric measure theory that minimised the effect of noise on the observed distributions. The key to this improvement was introducing a normalisation of the individual contour statistics to the domain average.

3.2 Multivariate Topology

The Reeb space is a generalisation of the Reeb graph for multivariate or temporal data. The first discussion of using the Reeb space to compute topological structure

of multiple functions appears in [2]. Here it is suggested that the Reeb space can be modelled mathematically in the form $f : \mathbf{M} \mapsto \mathbf{R}^k$, where $\mathbf{M}$ represents the domain and f the output of k scalar functions. For the simple case, where $k = 1$, this is directly comparable to the Reeb graph. The Reeb space extends this formulation to situations where $k \geq 2$.

Joint Contour Net

Carr et al. [1] presented the first discrete representation of the Reeb space using the Joint Contour Net (JCN). The algorithm computes the Reeb Space as a number of multi-variate contours, named slabs, representing connected regions of the domain with respect to multiple functions. In comparison to the contour tree, the JCN lends itself to parallelisation significantly more easily as each slab is constructed from smaller discrete regions labelled fragments [13].

In nuclear physics the Joint Contour Net has previously been used to visualise and analyse scission datasets, where it was used to identify the splitting of an atomic nucleus into multiple parts [14]. It was found that the Joint Contour Net was well suited to capturing this divergent behaviour, using proton and neutron density fields as inputs. This experiment was initially performed at a single temperature [15], but later repeated at multiple temperatures [16] due to its ability to capture the splitting of the compound nucleus as a forking in the multi-field topology. Whilst performing the analysis a number of other events were captured and linked to the scission theory. It was found that starburst-like features in the graph structure of the JCN likely equated to well formed fragments [15]. A zippering effect was noted at large temperatures that was determined to represent a spatial connectivity between distinct regions in the multi-field. In terms of the underlying physical phenomena, domain experts interpreted it as a de-localisation of quasi-particles at high temperatures.

More recently the joint contour net was used to visually analyse data from hurricane Isabel [17], allowing vertices in the Joint Contour Net to be mapped to their spatial co-ordinates. An interactive environment was developed that allowed users to relate interactions in the temperature, pressure and precipitation fields to physical phenomena such as rain bands and the eye of the hurricane. The ability to relate properties of the joint contour net to known physical features helped to increase understanding of how the joint contour net is able to capture multi-field interactions.

Persistence

Persistence in multivariate data sets is a more difficult problem to define than in the univariate case. Simplification and persistence metrics can be defined on a number of structures present within the multi-field topology. It is also possible to extend the concept of isosurface statistics [11], [12] to multi-field inputs through the use of Continuous Scatterplots [18]. These can be defined to show relations between n dimensional inputs with m scalar fields; in the case where $n = 3$ and $m = 1$ the output approximates to the output of [12].

The Jacobi set defines the set of points where the gradient of multiple functions align or have a gradient of zero [19]. When the multi-field is used to represent time dependent data this can be used to augment the univariate notion of persistence with a lifetime parameter. This approach was used in [20] to compute persistence in the context of the Morse-Smale complex. However, when generalised to non-temporal functions,defining persistence as a feature of the Jacobi set becomes a non-trivial task [21]. The Reeb skeleton relies on an extended Jacobi set, the Jacobi Structure primarily to aid multi-dimensional simplification [22]. This is a simplified graph structure that allows measures of persistence to be assigned to its arcs. Lip pruning techniques, similar to the leaf pruning method of simplification found in univariate topological structures [9] can then be applied to progressively remove noisy features in the multi-field.

4 Lattice QCD data analysis

Lattice QCD simulations, despite their relatively small space-time size, are extremely complex due to the ability to compute a rich set of observables on the lattice. Furthermore, the method of noise reduction used in this study, known as cooling, generates a unique lattice for each iteration. Due to the quantum nature of lattice QCD each experiment must be computed multiple times as physical observables are defined as averages or expectation values. This results in an ensemble of suitably weighted configurations that are representative of the Feynman path integral over every possible state of the system. The ensemble emerges from a Markov chain with updates generated by an algorithm, such as hybrid Monte Carlo [23], that respects a physical condition known as detailed balance.

Each of these experiments are conducted as part of multi-ensemble studies [24]; in our particular work each ensemble is generated with a differing level of chemical potential (μ). This takes the form of computing around fifty separate joint contour nets for each ensemble; measures from the joint contour nets are then computed and averaged to produce an ensemble average. Typically domain scientists would then present this information with regard to chemical potential as a histogram or line graph to look for signals of phase changes in the simulated quark-gluon matter.

The experiments were carried out on a Dell cluster made up of four homogenous compute nodes each with access to 16GB RAM and a separate dedicated front-end node. Each compute node contains four AMD Opteron 6376 CPUs, each with eight physical cores (16 logical cores) running 64-bit Debian 8.6. The software is built using VTK 6.1 [25] with extensions from the Multi-field Extension of Topological Analysis (META) project [26] in C++11. Quantisation parameters of the Joint Contour Net were chosen to best suit the available resources in terms of memory and number of parallel processes.

4.1 Joint Contour Net comparison of data under cooling

The cooling algorithm is an established method of iteratively removing noise from lattice QCD configurations whilst respecting the underlying physical theory. However, the process is not without pitfalls; in particular, it is possible for over-cooling to remove the actual intended observables. As a physical noise reduction technique is available we are restricted in our use of using persistence as a noise reduction technique. Instead we use it to evaluate the effect of the physical process on the data.

Overview of input scalar fields

The Polyakov loop, otherwise known as the Wilson line operator, can be used as a method for computing the symmetry of a lattice. Breaking of symmetry (Fig. 5) is one signal that can indicate a transition to a de-confined state, achieved by varying ensemble parameters such as temperature or chemical potential. This method of locating critical temperatures is well established in SU(2) and SU(3) lattice gauge theories [27, 28, 29, 30].

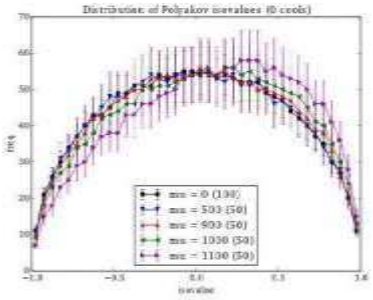

Fig. 5 By graphing the scalar values making up the Polyakov loop observable it is possible to see the breaking of symmetry at high chemical potentials (μ) in un-cooled data.

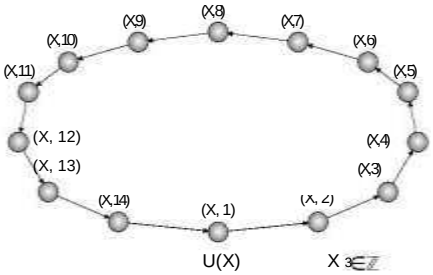

Fig. 6 The Polyakov loop is a scalar field $f : \mathbf{R}^4 \mapsto \mathbf{R}^3$ computed by visiting each site in a given space-time direction. It most cases we loop over the time axis.

From a computational point-of-view the Polyakov loop presents a convenient method for reducing the $\mathbf{R}^4$ lattice to a $\mathbf{R}^3$ scalar field. In order to compute the Polyakov loop, we take the product of all time-like link variables from a given lattice site in three dimensions (see Eq. 3). On a lattice defined with a periodic time axis, as is the case in lattice QCD, the effect is a closed straight line, as visualised in Fig. 6. Thus the Polyakov loop represents an attractive method for identifying de-confinement in a form that can easily be visualised and analysed using existing

methods.

$$f(x) = \frac{1}{2^9}\left(\mathrm{Tr}\left(\prod_{n=1}^{L_t} U_i(x+n\hat{i})\right)\right) \quad (3)$$

$$\text{where } i = 0,\ x \in \mathbb{Z}_3$$

Previous studies using univariate topological measures calculated using the Reeb graph [31] suggest that the Polyakov loop observable may lose its ability to identify de-confinement under cooling. We are interested to see if this could phenomena can be identified using multivariate topology. For this study we calculate the Polyakov loop lattice observable for each configuration in an ensemble at cooling iterations zero to twenty. Joint contour nets are computed for each neighbouring pair of cooling iterations and their structure analysed.

4.2 Results and analysis

Below we present a selection of the most promising results gathered by analysing different aspects of the Joint Contour Net.

Joint contour net vertex connectivity

We present the output from multiple ensembles by considering the connectivity of vertices in the JCN, focusing upon highly connected vertices. These are signatures of the starburst effect present in the JCN study on nuclear scission [15]. From the perspective of the scalar data, this corresponds to two regions of highly correlated geometric features. A high correlation is expected for cooled data as each slice should be a smoothed version of its proceeding configuration.

When considering the distributions for multiple ensembles we found interesting behaviour as the level of cooling was increased. For uncooled data (Fig. 7) there is very little difference between lattice sizes and levels of chemical potential (i), meaning all ensembles are roughly similar using this measurement. However, as cooling takes effect we notice a distinct banding in the distributions (Fig. 8). This appears to not only make a distinction between differing chemical potentials, but also the hot and cold lattice. We are currently in the process of evaluating the output visually with the domain experts, to better understand what type of event this phenomena might be linked to at the level of the lattice QCD simulation.

Multivariate persistence

Results from analysis of the graph structure of the JCN suggested that the algorithm was able to distinguish between different thermodynamic control parameters. It also revealed a banding effect in the closely correlated regions of the multi-field geome-

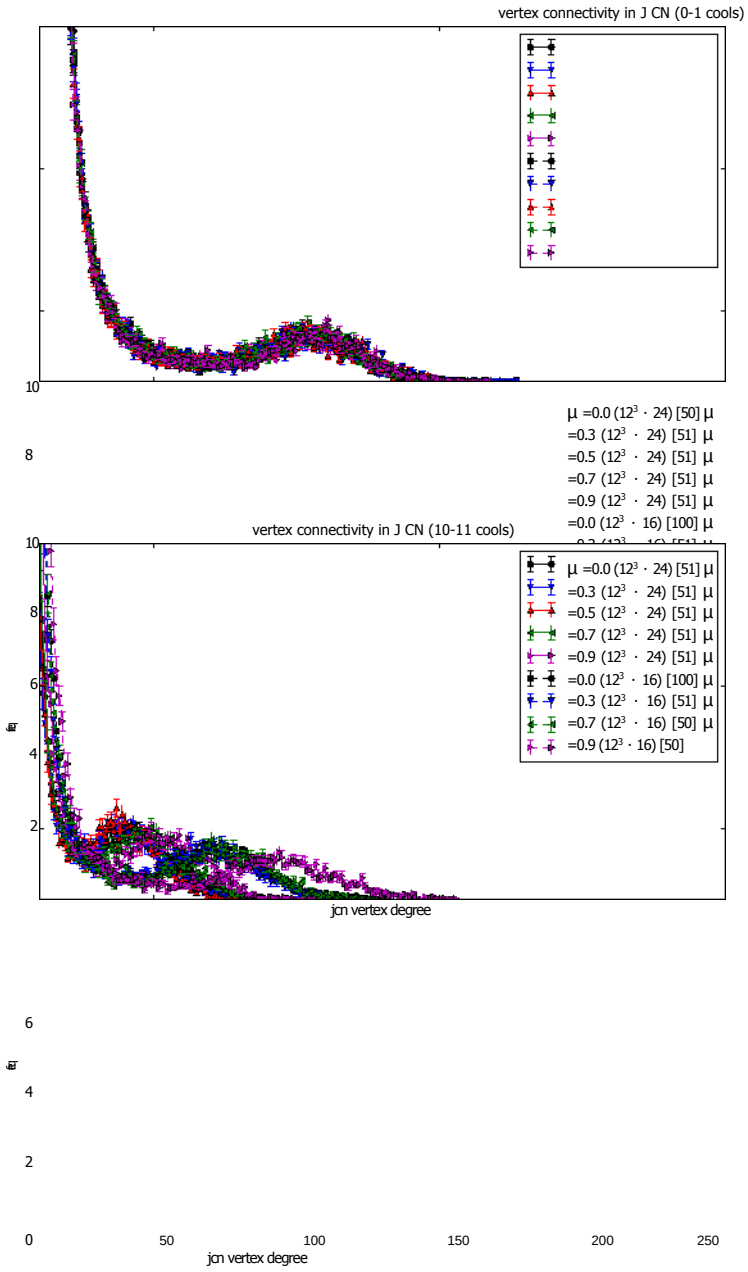

Fig. 7 Degree of connectivity of vertices in the Joint Contour Net as an average distribution for uncooled data. All ensembles, when uncooled, seem to exhibit the same connectivity distribution.

Fig. 8 Degree of connectivity of vertices in the Joint Contour Net as an average distribution for data cooled for 10 iterations. A distinct banding effect can be observed, separating the hot and cold lattices and differing chemical potentials.

try. In order to further examine this we next considered measures of slab persistence for multiple ensembles.

The Joint Contour Net allows the Reeb space to be segmented into regions or Joint Contour Slabs. Each slab is associated with an n-tuple of isovalue, where n represents the number of fields, with each slab corresponding to a vertex in the JCN. Using this segmentation of the Reeb space we compute measures of topological persistence by querying properties of the slabs. In Figure 9 we have visualised this within the data domain by rendering the slabs using colour to highlight persistence.

In order to understand the effect of cooling we created surface plots of the persistence measures against the chemical potential (μ) and number of cooling iterations.

Fig. 9 Left: Quantised contours generated from Polyakov loop at 19 cooling iterations. Centre: Joint contour slabs are a union of the quantised contours at 19 and 20 cooling iterations. Right: Quantised contours generated from Polyakov loop at 20 cooling iterations. Red regions represent most persistent slabs, measured by counting triangles.

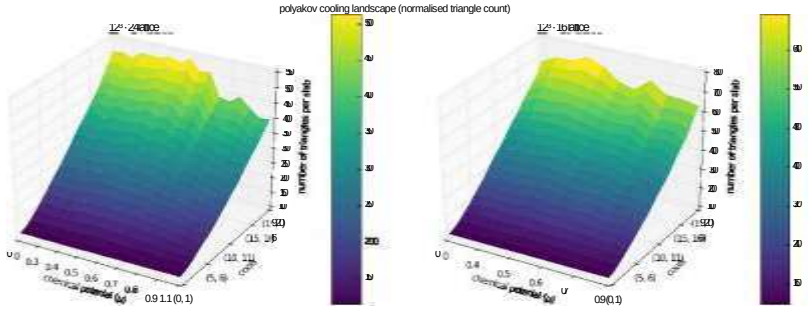

Fig. 10 Bivariate persistence, measured as the average number of triangles per slab, computed using the Polyakov loop across the cooling range.

Persistence measures have been averaged in each case by dividing the totals by the number of slabs. In Fig. 10 we present a cooling landscape generated by counting the average number of triangles per slab on the hot and cold lattices. An upward trend is shown in the averages as the number of cooling iterations is increased and number of distinct topological objects decreases, representing the simplification effect the algorithm has on the input field and associated Reeb space. The same trend is visible when comparing with other persistence measures including the average surface area.

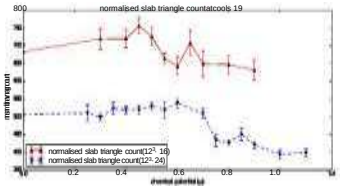

Fig. 11 Comparing the average number of triangles per slab on the two lattices at cools 19 and 20.

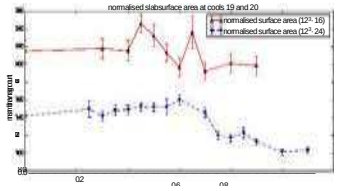

Fig. 12 Comparing the average surface area per slab on the two lattices at cools 19 and 20.

In Fig. 11 we show a direct comparison of the persistence as a measure of the number of triangles on the two lattices at the maximum number of cools. Both lattices feature similar trends in the data; an initial largely flat region at low values of i, followed by a global peak. A second peak follows higher in the chemical potential (i) range following a region with a downward trend in persistence before both distributions begin to plateau at high values of i. Initial examination of the two distributions suggest that on the hotter lattice ($n_t = 16$) the trend is shifted towards the lower end of chemical potentials. This result seems to agree with recent evidence produced from statistical physics [32] that see a shift in de-confinement on hotter lattices. While this is an encouraging result, it should be noted that due to a limited number of configurations available for the analysis data exhibits large error bars. Consultation with physicists suggested that availability of a larger number of configurations will reduce the level of uncertainty.

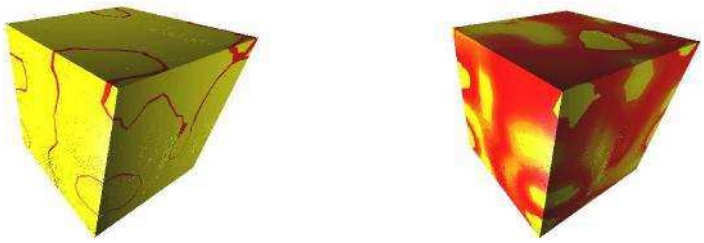

Fig. 13 A direct comparison of the multivariate persistence measures. Left: measuring persistence as the number of triangles per slab. Right: measuring persistence as the surface area of each slab.

We were also able to compute a more sophisticated measure of multivariate persistence by computing the surface meshes for each slab (Fig. 12). This is a slower process to compute, requiring the computation of the area of million of triangles using Herons formula but encouragingly converges to a similar distribution. There are some subtle differences in the two persistence measures; most evident in the hotter ($n_t = 16$) lattice. In particular there is a difference in the size of the peaks at $i = 0.45$ and $i = 0.65$. Figure 13 demonstrates the difference between the two persistence measures in a visual form. The triangle count measure is able to pick-out boundaries between regions in the multi-field, whereas the surface area highlights is able to pick out additional structure. Collection of further samples would allow us to confirm which of the two measures is more expressive with respect to the analysis of the Reeb space structure.

Other computed observables

Besides the two topological measures we have discussed here we also analysed other aspects of the JCN graph structure. We found that computing the ratio of Jacobi Nodes to JCN vertices gave an interesting overview of the effect of the cooling

algorithm, with the ratio tending towards 1.0 as the level of cooling reached it's peak. This confirmed expectations as it would indicate an overall convergence to a stable configuration with minor changes between cooling iterations. The effect wasn't exactly constant across the chemical potential range however, with ensembles at higher chemical potentials seemingly taking longer to converge on a optimal level of cooling.

5 Conclusions and future work

We have presented a number of uses of the Joint Contour Net for analysing data from lattice QCD ensemble data sets. Due to the quantity of data analysis required because of the quantum mechanics involved, visual inspection is not a feasible method for analysing the data. Instead we have proposed a number of measures taken directly from the multi-field topology that can then be displayed according to a number of domain specific parameters. Through the use of ensemble averages it is then possible to understand if patterns present in the multi-field topology share a correlation with existing statistical physics predictions. Existing results suggest that some measurements, in particular the multi-field persistence, could correlate well with physical observations.

In this work we have concentrated on a single lattice observable, the Polyakov loop. However, lattice QCD presents many other observables that can be analysed for hints of de-confinement, many of which are defined on four dimensional space-time fields. We are currently in the process of analysing this data to look for correlations between different lattice fields.

Acknowledgements This work used the resources of the DiRAC Facility jointly funded by STFC, the Large Facilities Capital Fund of BIS and Swansea University, and the DEISA Consortium (www.deisa.eu), funded through the EU FP7 project RI- 222919, for support within the DEISA Extreme Computing Initiative.

References

1. H. Carr and D. Duke, "Joint Contour Nets," IEEE Transactions on Visualization and Computer Graphics, vol. 20, no. 8, pp. 1100 – 1113, 2013.
2. H. Edelsbrunner, J. Harer, and A. K. Patel, "Reeb spaces of piecewise linear mappings," in Proceedings of the twenty-fourth annual symposium on Computational geometry. ACM, 2008, pp. 242–250.
3. S. Hands, "The phase diagram of QCD," Contemporary Physics, vol. 42, no. 4, pp. 209–225, 2001.
4. M. Creutz, Quarks, gluons and lattices, L. P.V., D. Weinberg, W. Sciama, and M. W.H., Eds. Cambridge: Cambridge University Press, 1983.
5. K. G. Wilson, "Confinement of quarks," Physical Review D, vol. 10, no. 8, p. 2445, 1974.

6. H. Carr and J. Snoeyink, “Path seeds and flexible isosurfaces using topology for exploratory visualization,” in Proceedings of the symposium on Data visualisation 2003. Eurographics Association, 2003, pp. 49–58.
7. H. Carr, J. Snoeyink, and U. Axen, “Computing contour trees in all dimensions,” Computational Geometry, vol. 24, pp. 75–94, 2003.
8. J. Tierny, A. Gyulassy, E. Simon, and V. Pascucci, “Loop surgery for volumetric meshes: Reeb graphs reduced to contour trees,” Visualization and Computer Graphics, IEEE Transactions on, vol. 15, no. 6, pp. 1177–1184, 2009.
9. H. Carr, J. Snoeyink, and M. van de Panne, “Simplifying flexible isosurfaces using local geometric measures,” in *Proceedings of the conference on Visualization’04*. IEEE Computer Society, 2004, pp. 497–504.
10. C. L. Bajaj, V. Pascucci, and D. R. Schikore, “The contour spectrum,” in Proceedings of the *8th conference on Visualization’97*. IEEE Computer Society Press, 1997, pp. 167—-173.
11. H. Carr, D. Brian, and D. Brian, “On histograms and isosurface statistics,” IEEE Transactions on Visualization and Computer Graphics, vol. 12, no. 5, pp. 1259–1266, 2006.
12. M. Meyer, C. E. Scheidegger, J. M. Schreiner, B. Duffy, H. Carr, and C. T. Silva, “Revisiting histograms and isosurface statistics,” IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 6, pp. 1659–1666, 2008.
13. D. J. Duke, F. Hosseini, and H. Carr, “Parallel computation of multifield topology: experience of Haskell in a computational science application,” in Proceedings of the 3rd ACM SIGPLAN workshop on Functional high-performance computing. ACM, 2014, pp. 11–21.
14. D. Duke, H. Carr, A. Knoll, N. Schunck, H. A. Nam, and A. Staszczak, “Visualizing nuclear scission through a multifield extension of topological analysis,” Visualization and Computer Graphics, IEEE Transactions on, vol. 18, no. 12, pp. 2033–2040, 2012.
15. N. Schunck, D. Duke, H. Carr, and A. Knoll, “Description of induced nuclear fission with Skyrme energy functionals: Static potential energy surfaces and fission fragment properties,” Physical Review C, vol. 90, no. 5, p. 54305, 2014.
16. N. Schunck, D. Duke, and H. Carr, “Description of induced nuclear fission with Skyrme energy functionals. II. Finite temperature effects,” Physical Review C, vol. 91, no. 3, p. 34327, 2015.
17. Z. Geng, D. Duke, H. Carr, and A. Chattopadhyay, “Visual analysis of hurricane data using joint contour net,” in Computer Graphics and Visual Computing (CGVC), W. Tang and R. Borgo, Eds. The Eurographics Association, 2014.
18. S. Bachthaler and D. Weiskopf, “Continuous scatterplots,” IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 6, pp. 1428–1435, 2008.
19. H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci, “Local and global comparison of continuous functions,” in Visualization, 2004. IEEE. IEEE, 2004, pp. 275–280.
20. P. T. Bremer, E. M. Bringa, M. A. Duchaineau, A. G. Gyulassy, D. Laney, A. Mascarenhas, and V. Pascucci, “Topological feature extraction and tracking,” in Journal of Physics: Conference Series, vol. 78, no. 1. IOP Publishing, 2007, p. 12007.
21. N. Suthambhara and V. Natarajan, “Simplification of jacobi sets,” in Topological Methods in Data Analysis and Visualization. Springer, 2011, pp. 91–102.
22. A. Chattopadhyay, H. Carr, D. Duke, Z. Geng, and O. Saeki, “Multivariate topology simplification,” arXiv preprint arXiv:1509.04465, 2015.
23. H. J. Rothe, Lattice gauge theories: an introduction, 4th ed. Singapore: World Scientific, 2012, vol. 82.
24. S. Cotter, P. Giudice, S. Hands, and J.-I. Skullerud, “Towards the phase diagram of dense two-color matter,” Physical Review D, vol. 87, no. 3, p. 34507, 2013.
25. W. Schroeder, K. Martin, and B. Lorensen, The visualization toolkit: an object oriented approach to 3D graphics. New York: Kitware, 4th ed. Kitware Inc., 2006.
26. H. Carr and D. Duke, “Multifield Extension of Topological Analysis (META) project.” [Online]. Available: http://www.engineering.leeds.ac.uk/computing/research/cse/meta/
27. L. D. McLerran and B. Svetitsky, “A Monte Carlo study of SU (2) Yang-Mills theory at finite temperature,” Physics Letters B, vol. 98, no. 3, pp. 195–198, 1981.

28. J. Kuti, J. Polonyi, and K. Szlachanyi, "Monte Carlo study of SU (2) gauge theory at finite temperature," Physics Letters B, vol. 98, no. 3, pp. 199–204, 1981.
29. J. Engels, F. Karsch, H. Satz, and I. Montvay, "High temperature SU (2) gluon matter of the lattice," Physics Letters B, vol. 101, no. 1, pp. 89–94, 1981.
30. K. Kajantie, C. Montonen, and E. Pietarinen, "Phase transition of SU (3) gauge theory at finite temperature," Zeitschrift für Physik C Particles and Fields, vol. 9, no. 2, pp. 253–254, 1981.
31. D. P. Thomas, R. Borgo, and S. Hands, "Topological Visualisation techniques for the understanding of Lattice Quantum Chromodynamics (LQCD) simulations," in EGUK Computer Graphics & Visual Computing Conference. Eurographics UK, 2016.
32. T. Boz, S. Cotter, L. Fister, D. Mehta, and J.-I. Skullerud, "Phase transitions and gluodynamics in 2-colour matter at high density," The European Physical Journal A, vol. 49, no. 7, pp. 1–13, 2013.