

Development of an Efficient and Flexible Pipeline for Lagrangian Coherent Structure Computation

Siavash Ameli, Yogin Desai, and Shawn C. Shadden

Abstract The computation of Lagrangian coherent structures (LCS) has become a standard tool for the analysis of advective transport in unsteady flow applications. LCS identification is primarily accomplished by evaluating measures based on the finite-time Cauchy Green (CG) strain tensor over the fluid domain. Sampling the CG tensor requires the advection of large numbers of fluid tracers, which can be computationally intensive, but presents a large degree of data parallelism. Processing can be specialized to parallel computing architectures, but on the other hand, there is compelling need for robust and flexible implementations for end users. Specifically, code that can accommodate analysis of wide-ranging fluid mechanics applications, while using a modular structure that is easily extended or modified, and facilitates visualization is desirable. We discuss the use of Visualization Toolkit (VTK) libraries as a foundation for object-oriented LCS computation, and how this framework can facilitate integration of LCS computation into flow visualization software such as ParaView. We also discuss the development of CUDA GPU kernels for efficient parallel spatial sampling of the flow map, including optimizing these kernels for better utilization.

1 Introduction

The computation of Lagrangian coherent structures (LCS) originated as a means to compute stable and unstable manifold type structures in vector fields with aperiodic time dependence. This was motivated by knowledge that the interaction

S. Ameli • S.C. Shadden (✉)
University of California, Berkeley, CA, USA
e-mail: sameli@berkeley.edu; shadden@berkeley.edu

Y. Desai
Georgia Tech, Atlanta, GA, USA
e-mail: yogindesai02@gmail.com

of such manifolds gives rise to chaotic dynamics, and hence understanding these interactions helped bring an ordered understanding to chaotic advection in fluid flow. This approach was originally applied to time periodic systems, especially using Poincaré mappings that make the dynamics autonomous. The applicability of traditional invariant manifold concepts breaks down for time aperiodic vector fields for practical and conceptual reasons. Notably, asymptotic notions associated with invariant manifold theory are neither applicable nor desirable for understanding inherently transient phenomena associated with unsteady computational or experimental fluid flow data.

LCS computational techniques do not typically solve for the stable and unstable manifolds of explicit trajectories. A global approach was developed from observations that material points straddling stable and unstable manifolds typically separate faster in forward and backward time than pairs of points not straddling such manifolds. That is, the manifold geometry may be inferred by considering the stretching associated with the hyperbolicity of these structures [14]. This led to an alternative characterization of organizing structures in fluid flows as invariant manifolds (material surfaces) that satisfy certain locally attracting or repelling properties, which became the basis for formalizing the concepts of LCS. LCS computations have been applied to diverse applications, as reviewed in [26], demonstrating wide-ranging utility in analysis of unsteady fluid advection.

Scalar measures are often used to identify LCS, such as the (largest) finite time Lyapunov exponent FTLE field [29]. This is accomplished by plotting the field and visually identifying such structures, or using an algorithmic approach to extract features such as ridges in the field [19, 21, 24]. The FTLE is derived from the largest eigenvalue of the Cauchy Green (CG) strain tensor

$$\mathbf{C}(\mathbf{x}_0, t_0, t_f) = \nabla \mathbf{F}_{t_0}^{t_f}(\mathbf{x}_0)^T \cdot \nabla \mathbf{F}_{t_0}^{t_f}(\mathbf{x}_0), \quad (1)$$

where $\mathbf{F}_{t_0}^{t_f} : \mathbf{x}(t_0) \mapsto \mathbf{x}(t_f)$ denotes the flow map, and $\mathbf{x}(t)$ is a fluid element trajectory with $\mathbf{x}_0 = \mathbf{x}(t_0)$. The full CG tensor encodes direction-dependent stretching information that can be leveraged, for example in defining normally hyperbolic LCS, i.e. material surfaces that are locally the most normally repelling over a chosen time interval [9]. Therefore, the computation of the CG tensor over the flow domain can be thought of as a common, or at least representative, target in LCS identification strategies.

The global approach to LCS identification requires a highly resolved sampling of the CG tensor over the fluid domain to locate potential LCS. One typically starts with fluid velocity field data $\mathbf{u}(\mathbf{x}, t)$, obtained from computation or measurement, and the flow map is computed by seeding the fluid domain with tracers and integrating the advection equation

$$\dot{\mathbf{x}}(\mathbf{x}_0, t) = \mathbf{u}(\mathbf{x}, t), \quad (2)$$

over a finite time interval (t_0, t_f) for a grid of seed points \mathbf{x}_0 . This is typically the most computationally intensive aspect of LCS identification since the seed grid may be composed of millions of material points especially in 3D flows. Furthermore, the integration of Eq. (2) for each seed point typically requires thousands of integration steps or more, and each integration step can require several space-time interpolations of $\mathbf{u}(\mathbf{x}, t)$. Furthermore, this process may be repeated for a time series of seed grids at different t_0 to obtain the time evolution of LCS. We note that the high resolution of seed points is needed in part for finite differencing to compute $\nabla \mathbf{F}_{t_0}^{t_f}(\mathbf{x}_0)$. Alternatively, integration of the variational equations, as performed in [15], may also be used to obtain the linearized flow map directly. However, sufficient resolution is still needed for LCS detection, especially when ridge extraction is performed [24].

It has been our experience that two major deterrents for wider adoption of LCS computations for flow post-processing are (1) computational time and (2) ease of use. With regards to (1), we will discuss the use of general purpose computing on graphics processing units (GPGPU) for accelerating flow map computations. This work builds on previous works of [7, 8, 13]. With regards to (2), we will discuss the use of the Visualization Toolkit (VTK) for creating an LCS computational pipeline that on the “front-end” is capable of handling various input data, and on the “back-end” facilitates visualization of LCS with standard flow analysis tools. In Sect. 2 we describe the main component of the LCS computational pipeline. In Sect. 3 we describe the implementation of the flow map computation on the GPU and results for various applications, ranging from 2D Cartesian to 3D unstructured velocity data processing. In Sect. 4 we will discuss some results and relationship with previous work.

2 LCS Pipeline

Here we describe a pipeline developed to process velocity field data for the purpose of computing LCS. This pipeline was written in C++ and will be made openly available on GitHub (www.github.com/FlowPhysics). This pipeline was developed using an object-oriented approach to create filters (classes that modify data) that can be extended or modified with minimal effort. Specifically, this pipeline makes extensive use of the popular Visualization ToolKit (www.vtk.org). VTK is a broad collection of open-source libraries for 3D computer graphics, image processing and visualization. The filters that we developed adhere to VTK coding standards and conventions for modularity, portability and debugging. We also chose to develop our pipeline using VTK for two additional reasons. First, VTK has a number of classes for reading standard file formats and data grid types commonly used by the fluid mechanics community. We could therefore leverage existing functionality to better support input from a variety of applications. Second, open-source programs for visualizing scientific data such as ParaView and VisIt are built on VTK.

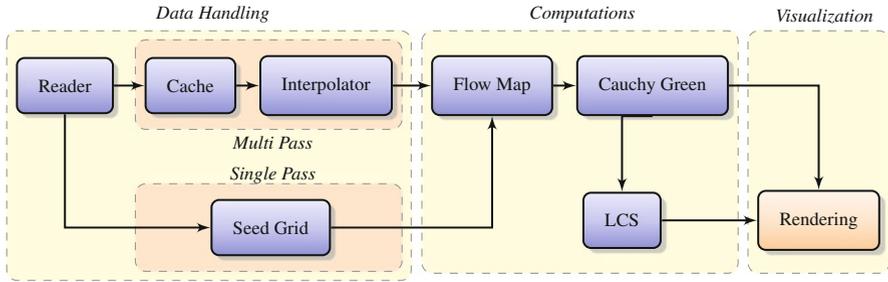


Fig. 1 Pipeline architecture

Developing an LCS pipeline in VTK facilitates integration with these programs, and alternatively development of custom rendering for a stand-alone software.

VTK is highly object-oriented and supports a pipeline architecture that generalizes the connection and execution of algorithms. The benefits of this methodology are (1) a universal interface between different filters in the pipeline, (2) greater control over information flow to better manage cache. The pipeline enables streaming of data, which is important for the application here because unsteady velocity field data are typically stored by a series of files that are often altogether too large to be loaded in memory. Streaming requires the use of *time-aware* filters and for this we developed time-related request keys, similar to the approach described by Biddiscombe et al. [4]. This pipeline does not support spatial streaming of velocity data for applications where individual time frames are too large for memory.

Our pipeline architecture includes the following custom filters, which are connected as shown in Fig. 1. We note that the pipeline does not represent data flow, but rather connectivity of filters and the management of requests.

Reader. A reader filter was developed to accept a variety of input data. VTK contains a number of classes to load various types of data (e.g. structured and unstructured grid data) and file formats. There are few readers in VTK that are time-aware, and those may only support one or a few file formats. Our filter implements a reader class that casts to the appropriate readers of both legacy and XML file types to support 11 file formats commonly used in VTK. This filter provides additional functionality by providing two different types of output based on requests of the pipeline. Output can be either a single data object or a data object that encapsulates multiple data objects depending on if the request comes from the Seed or the Cache filter. This filter also adds metadata on output objects such as times and indices that are needed throughout the pipeline to facilitate temporal streaming and interpolation. Lastly, this filter provides a framework for implementing necessary pre-processing of velocity data, such as tetrahedralization as motivated below.

Cache. This filter is connected to the output of Reader. The Cache filter is essentially a wrapper that manages VTK's `vtkTemporalDataSetCache` class.

It stores a window of the last requested data time frames. The amount of data cached depends on the upstream request. This filter is used to avoid repetitive streaming processes in the pipeline and deletes unnecessary data on memory.

Interpolator. This filter is used to interpolate the velocity field data in space and time. Data in the pipeline before this filter are provided discretely and all filters after `Interpolator` treat velocity data as continuous. Tracers are updated altogether each time step. This ensures that tracers are requesting the same flow field information each update step to best handle data streaming and memory utilization. This is also consistent with our GPU implementation whereby tracers are updated altogether to keep threads short-lived. We note that both space and time interpolation is performed on the GPU. Specifically, a window of velocity data frames (nominally two velocity files for linear interpolation) are loaded on the GPU's memory, and interpolation is performed on the GPU as necessary as integration proceeds between these time frames.

Seed. This filter is used to initialize the seed points where the CG tensor is to be sampled. Currently these locations are defined as structured grid data. This filter defines their initial conditions, other attributes, and for unstructured velocity data computes which cell in the velocity field mesh each seed point is located in. While the seed grid are nominally the locations where the CG tensor is evaluated, highly localized auxiliary points may also be defined about each seed point for improving flow map gradient computation as discussed below.

Flow Map. This filter maps the seed points forward in time. Therefore it is connected to both `Seed` and `Interpolator`. It is capable of implementing various types of integration routines, e.g. single step (Runge Kutta) and multi step (e.g. Adams Bashforth) methods. The same points passed from the input (`Seed`) are passed to the output with the final seed point locations added as a vector attribute to the point data. If a trajectory leaves the spatial domain of the velocity data, two options are possible. The CG tensor evaluation can be performed early at all locations relying on this trajectory. Alternatively, velocity extrapolation can be performed inside the interpolation filter, however one must take care is this regard [26, 30].

Cauchy-Green. This filter calculates the finite time Cauchy-Green strain tensor over the seed points using the initial and final point locations output from `Flow Map`. Currently, this is performed using standard central difference formula for the flow map gradient matrix entries. This differencing can be applied directly to the evolution of the seed point grid, or by differencing auxiliary points attributed to each seed location as described above. The auxiliary point method is more computationally expensive, but can provide more accurate CG tensor computation at less computational expense than an equivalent increase of seed grid resolution when seed spacing is reduce by more than 1/2. However, an equivalent increase in seed grid resolution also improves resolution of the CG tensor field, which is not accomplished through the auxiliary points method. Additionally, the auxiliary point

method can result in missed structures, when LCS are not properly straddled by points used in finite differencing [29, 31]. This filter also computes the eigenvalues and eigenvectors of the CG tensor at each seed point, which are commonly used for LCS detection. The seed points are passed to the output with the eigenvalue/vector fields added as point data attributes.

LCS. This filter processes the output from the Cauchy Green filter to apply criterion for LCS detection, for example, further processing of FTLE field data, such as Gaussian smoothing to remove computational noise as well as C-Ridge computation [23, 24]. While not currently implemented, strainline and shearline computation as described in [10] could be implemented in this filter.

Visualization. This is not a separate filter in our pipeline, but we have separated this conceptually. One important benefit of using the new pipeline functionality of VTK is that we can compile each filter to a shared library that can be loaded by the ParaView plugin manager, so an end user can simply import the filters as plugins in ParaView and then use these to create a pipeline for computing and visualization of fields defined from the CG filter (e.g. FTLE), or visualize features from the LCS filter. We note that support for FTLE computation was recently added to VisIt [2] as used by Özgökmen et al. [18].

3 Flow Map Computation on the GPU

The computation of the CG tensor field requires the advection of a large set of tracer trajectories. Since each tracer moves independently, these calculations present a parallel workload. GPGPU was used to perform acceleration/parallelization of this computation. On the hardware side, GPUs are widely available, cheap and scalable and on the software side GPU programming has become widely accessible, portable and supported by various compilers.

We used NVIDIA's compute unified device architecture (CUDA) platform [1]. Optimization of a CUDA program depends somewhat on the GPU architecture used. We report results run on a consumer-level graphics card with the NVIDIA GeForce GTX 670 GPU and the higher-end NVIDIA Tesla K20 GPU. Both GPUs are based on the Kepler architecture. GeForce GTX 670 has peak theoretical double floating point performance of 0.19 TFlops whereas Tesla K20 has peak theoretical double floating point performance of 1.17 TFlops. While GPU threads are plentiful and can be launched and terminated with minimal overhead, performance of parallelized flow map (trajectory) computation is primarily limited by memory bandwidth.

For the GPU implementation only registers, global memory and constant memory were utilized. Registers were used by individual threads to store local variables in specific interpolation and integration kernels discussed below, and were reused in each kernel to the greatest extent possible. Since constant memory is read only and limited to 64 kb but can be accessed by all threads, we used this for mesh parameters

and other constants. Tracer and velocity field arrays were always stored on global memory.

Each streaming multiprocessor (SMX) can run up to 16 resident blocks or 2,048 resident threads concurrently for the architecture we used. A common target in optimizing the code was to make sure every streaming processor had roughly 2,048 resident threads running at any given time. The number of registers available per SMX limits the number of threads that can reside on a SMX. Since the number of registers were limited to 64 K per multiprocessor, this implied that every thread needed to use only 31 registers for 100 % occupancy. We developed our kernels, as discussed below, to keep register usage nominally within this bound.

3.1 Implementation

For the purpose of explaining the implementation and performance results we focus discussion on the 4th order explicit Runge-Kutta integration method (RK4). To decrease register usage and hence improve GPU utilization, we broke up the integration kernel into smaller kernels. For example, RK4 requires the vector field to be evaluated at four different locations in space-time for a single update step. These interpolations are done as separate kernels. Furthermore, for unstructured grid data each interpolation requires a cell-search, i.e. determination of which velocity grid cell the interpolation point is located. Therefore, for each integration step, eight kernels are run; four kernels to evaluate the velocity field, and each of these requires a preceding cell-search kernel call. This strategy involves the use of a large number of threads that are short-lived, which are what GPUs are designed to handle. Using this strategy also enables the CG tensor to be evaluated at any time step during the integration interval (t_0, t_f) with minimal effort. This can be advantageous if the CG tensor is to be evaluated early for trajectories that leave the fluid domain before the full integration interval is reached. This is common with velocity data coming from modeling or measurement over a truncated domain. But moreover, some LCS detection strategies require certain criteria to be evaluated on the fly as integration proceeds. This need also motivated updating tracers synchronously for each integration step, which is consistent with maintaining efficient data flow in our pipeline, Sect. 2. That is, the pipeline is most efficient if all upstream requests require the same velocity data loaded in memory. We note however that velocity data is loaded into GPU memory only every time a new velocity data file is needed, not every integration time step. Figure 2 shows the runtime comparison of the single-kernel code and split-kernel code. It can be seen that split-kernel strategy executes in roughly 30 % less time for all by very small seed grid sizes.

We use an efficient local cell search strategy for unstructured grids [27]. We have demonstrated that this method outperforms other cell search methods we have tested such as the structured auxiliary mesh approach, and VTK's Kd-tree and Oct-tree methods. This method is designed for tetrahedral grids (or triangular grids in 2D). Therefore, other grid topologies of velocity data that might need to be

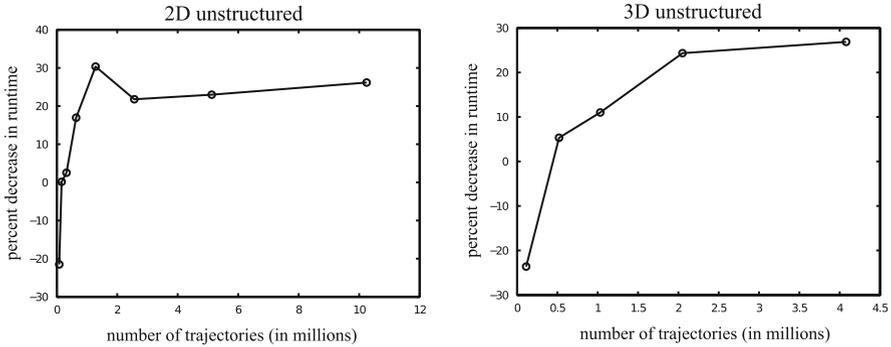


Fig. 2 Runtime comparison between a single-kernel RK4 time step implementation and a split-kernel RK4 time step implementation. Graphs plot the percent decrease in runtime of the split kernel implementation from that of the single kernel implementation. Runs were performed on both the 2D and 3D unstructured data examples described below

processed are tetrahedralized using built in VTK functionality as a preprocessing step in the `Reader` filter. This enables a single, efficient cell search algorithm to be used, which alleviates the need to develop and optimize different interpolation and integration kernels for different velocity grid topologies. However the tradeoff is that piece-wise linear representation using a tetrahedral grid may degrade interpolation accuracy when the native CFD grid has higher order elements.

3.2 Performance

Most of the processing time for the CG tensor field computation is spent inside the velocity field interpolation function. This function gets called one or more times each integration step for each trajectory update. Efficiency of this function is a main determinant of computational time. The interpolation function is mainly dependent on the grid topology of the velocity data, e.g. interpolation on Cartesian grids differs in strategy and performance than interpolation on unstructured grids. Different kernels were developed to handle interpolation on different grid types. The efficient cell search algorithm for tetrahedralized unstructured grid data that is used readily facilitates spatial interpolation as described in [27]. We present results for four different grid types considered: 2D Cartesian, 3D Cartesian, 2D unstructured, 3D unstructured. Specifically, the example applications include the double-gyre flow [29], which has become a standard test case for LCS computations (2D Cartesian); the 3D Rayleigh-Bénard convection cell [17] (3D Cartesian); a coronary stenosis model [28] (2D unstructured); and an abdominal aortic aneurysm (AAA) model [3] (3D unstructured). These examples are shown in Fig. 3.

Figure 4 plots the performance results from the GeForce and Tesla GPUs. The values plotted are the runtimes on the respective GPU normalized by the serial CPU

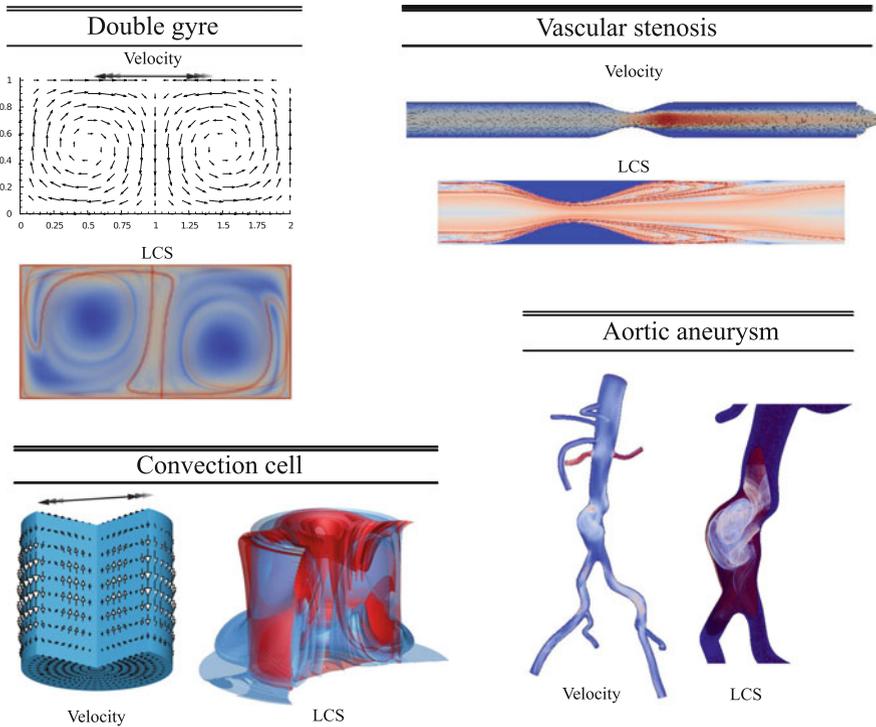


Fig. 3 Snapshots of the flow fields and LCS from the four applications used for performance testing, clockwise from *upper-left*: double gyre [29]; vascular stenosis [28]; abdominal aortic aneurysm [3]; Rayleigh-Bénard convection [17]

runtime, which gives the speedup over serial CPU processing. The CPU used was an Intel Core i7-3770 Ivy Bridge 3.5 GHz processor and the implementation was FlowVC [25], which was written in C. We observed up to roughly 70× speedup for computations on the 2D Cartesian data down to roughly 12× speed up for 3D unstructured grid data. As expected, the Tesla K20 yields higher speedup due to a higher number of streaming multiprocessors. The AAA velocity data was specified on an unstructured tetrahedral grid of 1.01 million elements. Unstructured velocity grids up to several million elements were tested and yielded similar speedup results as the AAA model.

We note that we achieved significant performance improvement by properly coalescing tracer grid arrays, which is common practice, whereby nearby threads processed tracer data coalesced in memory. We did not notice any improvement of performance by sorting velocity field data arrays. This is because as the trajectories of the seed grid evolve, trajectories with nearby indices in memory (and hence locally processed on the GPU) do not necessarily require velocity data that is localized in space. A recent solution to this problem was proposed by Chen and Hart

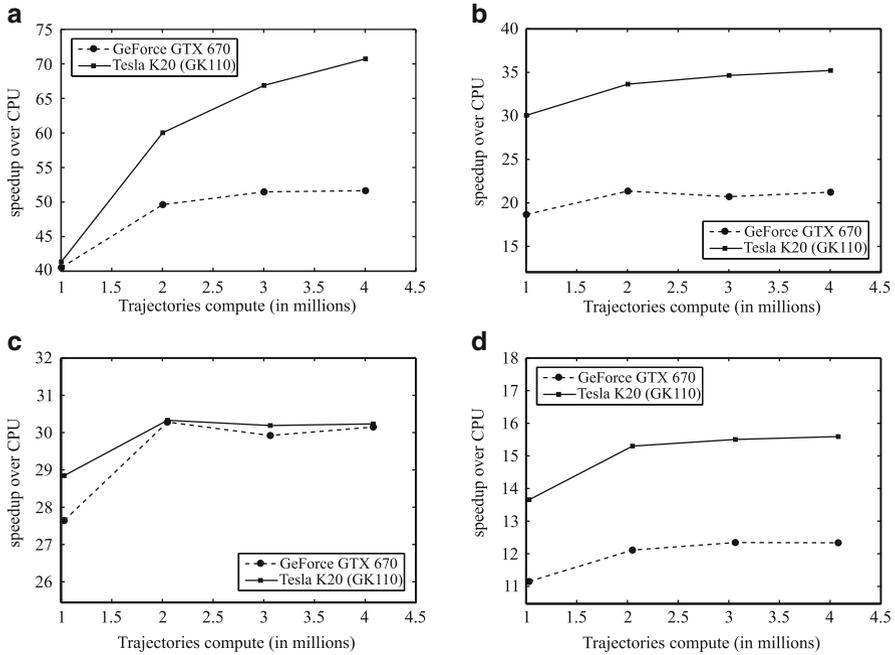


Fig. 4 Speedup of particle trajectory computation on a GeForce GTX 670 GPU and a Tesla K20 GPU compared to serial processing on an Intel Core i7-3770 CPU. All computations were performed using double precision for floating point variables. (a) 2D Cartesian. (b) 3D Cartesian. (c) 2D unstructured. (d) 3D unstructured

[6], which reorganizes particles into spatially coherent bundles as they are advected to improve memory coherence and shared-memory GPU performance.

3.3 Verification

Computations performed on the GPU were verified against existing CPU code that has been extensively used and tested, FlowVC. We present verification data from the 2D Cartesian and 3D unstructured grid applications, as these cases represent the extrema in results from the four examples considered. For each application we released several thousand tracers in the domain and integrated their trajectories for the nominal minimum time needed for LCS computation. Specifically, we chose this time based on the AAA model, and scaled the other integration times by the nominal edge size x_l divided by the mean velocity magnitude $\langle \mathbf{u}(\mathbf{x}, t) \rangle$ at peak flow. This defines a characteristic time scale similar to use of the CFL number. This ensured that all cases were integrated over roughly the same number of elements in their respective domains. We performed these computations using the GPU and CPU

codes on the GeForce GTX 670 and Intel Core i7-3770. The average error was defined by computing of the L_2 norm of the differences in tracer locations over time between the two runs, and averaging over all tracers as follows

$$e_D(t) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i^{DG}(t) - \mathbf{x}_i^{DC}(t)\| ,$$

where N is the number of trajectories, $\mathbf{x}_i^{DG}(t)$ is the trajectory of tracer i computed on the GPU using double precision floating point numbers, $\mathbf{x}_i^{DC}(t)$ is the trajectory of tracer i computed on the CPU using double precision floating point numbers, and $\mathbf{x}_i^{DG}(0) = \mathbf{x}_i^{DC}(0)$. In all applications the error stayed below 1×10^{-10} , indicating that both GPU and CPU implementations were performing equivalent tracer trajectory computation.

Single precision calculation can be several times faster than the double precision calculation, however accuracy may be unacceptable. The results generated using a single precision floating point GPU implementation were compared against the double precision CPU results. As described above, the average L_2 norm of the error in trajectories of several thousand tracers was computed from a single precision GPU run and a double precision CPU run as

$$e_S(t) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i^{SG}(t) - \mathbf{x}_i^{DC}(t)\| ,$$

where $\mathbf{x}_i^{SG}(t)$ is the trajectory of tracer i computed on the GPU using single precision for floating point variables. For the 2D Cartesian data, we noticed fairly acceptable errors in tracer trajectories and subsequent LCS computation. However, for the more complex 3D unstructured grid data, we notice unacceptable degradation in accuracy, e.g. to a point where noticeable degradation of the FTLE field occurred. The errors $e_d(t)$ and $e_s(t)$ are plotted against integration time for these two cases in Fig. 5 using the NVIDIA GeForce GTX 670 GPU.

Because the plots in Fig. 5 represent averages over many tracers, some tracer trajectories may deviate to far greater extent than the mean values shown. Indeed, in the AAA flow, some tracers were advected to different arterial branches based on these errors, which has important consequences for that application. Not surprisingly, errors can be worse near LCS due to inherent sensitivity to initial conditions at these locations, which can be problematic for accurate LCS detection. In addition, while we consider the double precision CPU results as a baseline for comparison, this does not imply that these results represent the “true” trajectories. The double precision CPU computations are subject to normal truncation and round-off errors. However, since double precision computation represents the de facto method for minimizing numerical error, and since no application considered nor of practical importance has a closed-form analytic solution, this was deemed an appropriate baseline for comparison.

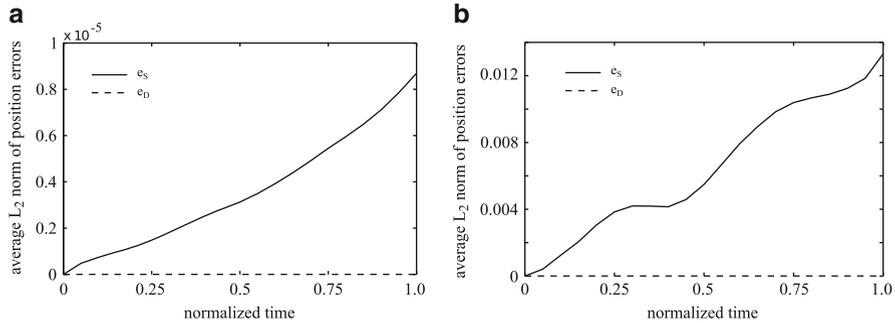


Fig. 5 Error between single precision (*solid line*) and double precision (*dashed line*) GPU computations compared to a double precision CPU computations versus integration time for the double gyre flow (*left*) and AAA flow (*right*). Integration times have been normalized and the value of 1.0 approximately represents the nominal minimum integration time needed to compute LCS for each application. (a) 2D Cartesian. (b) 3D unstructured

4 Discussion

We have developed a modular pipeline for LCS computation that is capable of loading a wide variety of fluid mechanics data, and that can be easily interfaced with ParaView for visualization of LCS computational results alongside other flow visualization tools. This pipeline was designed to be modular and flexible so that modifications and additions can be made with minimal effort.

As the LCS method has gained considerable popularity in the past few year, there have been several developments to improve computation and visualization of LCS. Strategies have been devised so that tracer seeding can be adapted to cluster sampling of the flow map near LCS for improved detection [8,16,21,22]. Also, since vector field interpolation is computationally intensive, strategies have also been developed to locally approximate the flow map [5], which can be beneficial when time series of tracer grids are considered. As well, trajectory computations obtained in a hierarchical manner have been considered for efficient FTLE computation [11]. We note that sampling the flow map (gradient) needs to be highly resolved in space but not necessarily time for LCS computations. That is, sampling in space is driven by LCS identification methods; mainly computation of the flow map gradient or subsequently the CG tensor field. Sampling time is done primarily for visualizing the evolution of the structures. Regardless of the strategy of sampling or interpolating the flow map or its derivative over the fluid domain, this process can be parallelized since advection of tracers can be performed for the most part independently. Because of this high degree of parallelism, and the fact the LCS are typically used for desktop postprocessing and flow visualization, many-core implementation on a single workstation is desirable, as opposed to a visualization cluster, which are less accessible.

Jimenez and Vankerschaver [13] discussed the computation and visualization of FTLE by GPGPU using CUDA. They also made their source code publicly available [12]. Their implementation was, by their own admission, naive since it only could handle FTLE computation for analytically defined vector fields. This greatly reduced memory accesses that, as discussed above, are the bottleneck in computing FTLE fields using GPU processing in practical applications. Garth et al. [8] and Hlawatsch et al. [11] leveraged GPU processing for FTLE computation as well, though these papers did not discuss details of their implementation or performance results as this was not a main focus. The recent paper by Conti et al. [7] described FTLE computation for bluff body flows using OpenCL, which allows implementation on mixed architectures, e.g. AMD's accelerated processing units. Their implementation was specialized to remeshed vortex methods for bluff body flows that involve mesh-particle interpolations previously tailored for GPU implementation [20]. They were able to achieve around one order magnitude speedup compared to serial CPU implementation, similar to what we observed here. However, their application was specialized to a particular flow problem. Most LCS computation is performed as a post processing step on fluid velocity field data. Since this is the most common and general scenario, it was the one which we developed our framework around, consistent with the design specification of our LCS pipeline.

We considered application to 3D flow and flow on unstructured grids, as [7, 13] reported performance results for 2D flows on structured grids. As shown and discussed above, we have run this implementation to integrate millions of tracers on velocity grids with several million element. This represents a reasonable limit for most LCS applications. For significantly larger velocity field grids, or tracer grids, one will overflow the global memory for the GPU. We generally noticed peak performance when kernels were kept within register memory bounds. This required each trajectory update to be divided into a series of kernels. Therefore, kernels were very short. Because of the way data is processed through our pipeline, we perform synchronous integration between the tracers. We believe this has advantages for LCS detection as well. For example, in truncated domains, which represent the vast majority of fluid mechanics data, tracers leave the domain before the finite time interval being considered for flow map computation. But perhaps more importantly, one may want to have access to the CG tensor at various times. Most LCS criteria are a one parameter family dependent on the chosen integration window. Using one integration window may identify a manifold as an LCS, but another integration window may not. Similarly, a manifold may satisfy one LCS criterion but not the other for a particular integration time. Therefore, having access to this information enables implementation of methods that may depend on how strain rate or direction changes over the integration time parameter. Or alternatively, schemes that adaptively sample the flow map based on CG tensor information may need sequential access to this information.

While this pipeline can function as a stand alone program, it can also be compiled as libraries to be used with ParaView. Such integration provides a natural platform for visualizing not only LCS computations, but also integrating these results with other existing flow visualization tools provided by this programs, and

also GPU-based visualization techniques available. Indeed, LCS as a collection of codimension one surfaces are rarely useful in understanding the flow. Knowledge of these manifolds must be integrated with other knowledge of the flow to gain fundamental insight into the flow topology, and these platforms provide significant capability in this regard.

Acknowledgements This work was supported by the National Science Foundation, award number 1047963.

References

1. NVIDIA CUDA C Programming Guide, Version 4.2 (2012)
2. VISIT – Software that delivers Parallel, Interactive Visualization. <http://visit.llnl.gov/>
3. A. Arzani, S.C. Shadden, Characterization of the transport topology in patient-specific abdominal aortic aneurysm models. *Phys. Fluids* **24**(8), 081,901-1-16 (2012)
4. J. Biddiscombe, B. Geveci, K. Martin, K. Moreland, D. Thompson, Time dependent processing in a parallel pipeline architecture. *IEEE Trans. Vis. Comput. Graph.* **13**(6), 1376–1383 (2007)
5. S.L. Brunton, C.W. Rowley, Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *Chaos* **20**(1) (2010)
6. M. Chen, J.C. Hart, Fast coherent particle advection through time-varying unstructured flow datasets (2013, Preprint). <http://graphics.cs.illinois.edu/papers/fastvection>
7. C. Conti, D. Rossinelli, P. Rossinelli, GPU and APU computations of finite time Lyapunov exponent fields. *J. Comput. Phys.* **231**, 2229–2244 (2012)
8. C. Garth, F. Gerhardt, X. Tricoche, H. Hagen, Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Trans. Vis. Comput. Graph.* **13**(6), 1464–1471 (2007)
9. G. Haller, A variational theory of hyperbolic Lagrangian coherent structures. *Physica D* **240**(7), 574–598 (2011)
10. G. Haller, F.J. Beron-Vera, Geodesic theory of transport barriers in two-dimensional flows. *Physica D* **241**, 1680–1702 (2012)
11. M. Hlawatsch, F. Sadlo, D. Weiskopf, Hierarchical line integration. *IEEE Trans. Vis. Comput. Graph.* **17**(8), 1148–1163 (2011)
12. R. Jimenez, J. Vankerschaver, CUDA_FTLE. http://www.its.caltech.edu/~raymondj/LCS/cuda_ftle-0.9.tar.bz2
13. R. Jimenez, J. Vankerschaver, Optimization of FTLE calculations using nVidia’s CUDA. Technical report, California Institute of Technology, 2009. http://www.its.caltech.edu/~raymondj/LCS/FTLE_on_GPU.pdf
14. C.K.R.T. Jones, S. Winkler, Invariant manifolds and Lagrangian dynamics in the ocean and atmosphere, in *Handbook of Dynamical Systems*, vol. 2, ed. by B. Fiedler (Elsevier, Amsterdam/Boston, 2002), pp. 55–92
15. J. Kasten, C. Petz, I. Hotz, B. Noack, H.C. Hege, Localized finite-time lyapunov exponent for unsteady flow analysis, in *Vision, Modeling and Visualization*, ed. by M. Magnor, B. Rosenhahn, H. Theisel (2009), pp. 265–274
16. F. Lekien, S.D. Ross, The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds. *Chaos* **20**(1) (2010)
17. F. Lekien, S.C. Shadden, J.E. Marsden, Lagrangian coherent structures in n-dimensional systems. *J. Math. Phys.* **48**, 065,404-1-19 (2007)
18. T.M. Özgökmen, A.C. Poje, P.F. Fischer, H. Childs, H. Krishnan, C. Garth, A.C. Haza, E. Ryan, On multi-scale dispersion under the influence of surface mixed layer instabilities and deep flows. *Ocean Model.* **56**, 16–30 (2012)

19. R. Peikert, F. Sadlo, Height ridge computation and filtering for visualization, in *Pacific Visualization Symposium*, Kyoto, ed. by I. Fujishiro, H. Li, K.L. Ma (2008), pp. 119–126
20. D. Rossinelli, C. Conti, P. Koumoutsakos, Mesh-particle interpolations on graphics processing units and multicore central processing units. *Philos. Trans. R. Soc. A – Math. Phys. Eng. Sci.* **369**(1944), 2164–2175 (2011)
21. F. Sadlo, R. Peikert, Efficient visualization of Lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Trans. Vis. Comput. Graph.* **13**(5), 1456–1463 (2007)
22. F. Sadlo, A. Rigazzi, R. Peikert, Time-dependent visualization of lagrangian coherent structures by grid advection, in *Topological Methods in Data Analysis and Visualization*, ed. by V. Pascucci, X. Tricoche, H. Hagen, J. Tierny (Springer, Dordrecht, 2010), pp. 151–165
23. B. Schindler, R. Fuchs, S. Barp, J. Waser, A. Pobitzer, R. Carnecky, K. Matkovic, R. Peikert, Lagrangian coherent structures for design analysis of revolving doors. *IEEE Trans. Vis. Comput. Graph.* **18**(12), 2159–2168 (2012)
24. B. Schindler, R. Peikert, R. Fuchs, H. Theisel, Ridge concepts for the visualization of lagrangian coherent structures, in *Topological Methods in Data Analysis and Visualization II*, ed. by R. Peikert, H. Hauser, H. Carr, R. Fuchs (Springer, Heidelberg/New York, 2012), pp. 221–236
25. S.C. Shadden, FlowVC (Version 1) [Computer software]. Retrieved from <http://shaddenlab.berkeley.edu/software/>
26. S.C. Shadden, Lagrangian coherent structures, in *Transport and Mixing in Laminar Flows: From Microfluidics to Oceanic Currents*, chap. 3, ed. by R. Grigoriev (Wiley, Weinheim, 2012)
27. S.C. Shadden, M. Astorino, J.F. Gerbeau, Computational analysis of an aortic valve jet with Lagrangian coherent structures. *Chaos* **20**, 017,512-1-10 (2010)
28. S.C. Shadden, S. Hendabadi, Potential fluid mechanic pathways of platelet activation. *Biomech. Model. Mechanobiol.* **12**(3), 467–474 (2013)
29. S.C. Shadden, F. Lekien, J.E. Marsden, Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Phys. D: Nonlinear Phenom.* **212**(3–4), 271–304 (2005)
30. W. Tang, P.W. Chan, G. Haller, Accurate extraction of lcs over finite domains, with applications to flight data analyses over Hong Kong International Airport. *Chaos* **20**(1), 017,502-1-8 (2010)
31. M. Üffinger, F. Sadlo, M. Kirby, C. Hansen, T. Ertl, FTLE computation beyond first-order approximation, in *Short Paper Proceedings of Eurographics 2012*, Cagliari, pp. 61–64