Fireflies: New Software for Interactively Exploring Dynamical Systems Using GPU Computing

Robert Merrison-Hort

*School of Computing, Electronics and Mathematics,
University of Plymouth, Drake Circus,
Plymouth, Devon, PL4 8AA, UK
robert.merrison@plymouth.ac.uk*

Received May 12, 2015

In nonlinear systems, where explicit analytic solutions usually cannot be found, visualization is a powerful approach which can give insights into the dynamical behavior of models; it is also crucial for teaching this area of mathematics. In this paper, we present new software, Fireflies, which exploits the power of graphical processing unit (GPU) computing to produce spectacular interactive visualizations of arbitrary systems of ordinary differential equations. In contrast to typical phase portraits, Fireflies draws the current position of trajectories (projected onto 2D or 3D space) as single points of light, which move as the system is simulated. Due to the massively parallel nature of GPU hardware, Fireflies is able to simulate millions of trajectories in parallel (even on standard desktop computer hardware), producing “swarms” of particles that move around the screen in real-time according to the equations of the system. Particles that move forwards in time reveal stable attractors (e.g. fixed points and limit cycles), while the option of integrating another group of trajectories backwards in time can reveal unstable objects (repellers). Fireflies allows the user to change the parameters of the system as it is running, in order to see the effect that they have on the dynamics and to observe bifurcations. We demonstrate the capabilities of the software with three examples: a 2D “mean field” model of neuronal activity, the classical Lorenz system, and a 15D model of three interacting biologically realistic neurons.

Keywords: Visualization; software; GPU computing; neuroscience.

1. Introduction

Many mathematical models are described by nonlinear ordinary differential equations (ODEs). It is therefore important to be able to visualize solutions to these equations, and study how their dynamic behavior changes when parameter values are modified. The theory of bifurcation analysis provides a rigorous mathematical framework for understanding how the stability of fixed points and limit cycles change as the parameters of the system are varied. Bifurcation theory can be combined with

numerical continuation in order to trace the boundaries in parameter space that separate different dynamical regimes, and several software packages exist for this purpose, such as AUTO [Doedel & Oldeman, 2012] and MATCONT [Dhooge *et al.*, 2003]. Before applying numerical continuation tools, in many cases it is useful to first get a rough understanding of how the system behaves for particular parameter values, since the initial points for continuation (such as approximate fixed point and limit cycle positions) must be known. This is where

more qualitative investigation (such as visualization of the system) can be useful.

In this paper, we introduce new software, “Fireflies”¹ for the dynamic visualization of ODE solutions. Instead of the traditional method of showing complete trajectories in phase space, Fireflies presents the user with a two- or three-dimensional view of a cloud of moving particles. Each particle represents the position in state space of one trajectory at the current point in time, and as the simulation runs, the particles move around the screen according to the equations of the system. Different colored groups of particles can be given different ranges of initial conditions, and each group can be either integrated forwards or backwards in time. Forwards moving particles illuminate attractors (e.g. stable fixed points and limit cycles) of the system, while backward moving ones illuminate repelling (unstable) objects. By watching the motion of the particles, it is also possible to see features such as saddle separatrices, and how the speed of movement varies around limit cycles. Finally, Fireflies can also generate “dynamic” bifurcation diagrams, in which one of the visualized dimensions corresponds to a chosen parameter and each particle is given a different value for this parameter.

When only the current position of each trajectory is shown, a very large number of particles must be used in order for the structure of the system to be visible — we find that several million are required for good results. Numerically integrating this many equations quickly enough to display an interactive animation is not typically possible using the limited parallelism of traditional central processing units (CPU), but is an ideal problem for so-called “massively parallel” graphical processing units (GPUs). Unlike CPUs, which typically contain a small number (up to 18, at time of writing) of largely independent processing cores, GPUs consist of hundreds or thousands of cores, all of which execute the same code at the same time. Nowadays GPUs, which were originally developed to accelerate rendering in 3D graphics applications, sit alongside the central processing unit (CPU) in almost all modern computers, including tablets and smartphones. Despite their origins in the video games industry, the use of GPUs is rapidly becoming an essential part of many scientific computing applications, due to their

low cost, widespread availability, and potential for large efficiency gains in certain types of computations. Fireflies consists of a user friendly graphical interface which can be used to produce simulations of N -dimensional systems of ODEs. Thanks to the power of GPU computing, these simulations can contain millions of particles while still running quickly enough to be interactive, allowing the user to change parameter values and immediately observe the effect of this on the particles’ motion.

In this paper, we will demonstrate the capabilities of Fireflies by showing visualizations of three example systems of ODEs. We will show how using Fireflies to study these systems can give insights into their behavior and how their dynamics depends on their parameters. Note that the figures show static screenshots from moving simulations; in order to properly appreciate the capabilities of the software we recommend watching videos of Fireflies in motion. A collection of videos showing the systems described in this paper is available at <https://vimeo.com/album/3391568>.

2. Example Systems

2.1. A two-dimensional model of the basal ganglia

In this section, we show an example of a two-dimensional system of nonlinear ODEs, corresponding to a simple model of neuronal activity. We show how Fireflies makes the dynamics of this system clear, and how the bifurcations of the system can be observed in an exciting new way. Although the system is two-dimensional, we also show how Fireflies can be used to create an interactive three-dimensional bifurcation diagram, by extending the visualization into parameter space.

Based on earlier modeling work [Gillies *et al.*, 2002; Pavlides *et al.*, 2012] we developed a mathematical model of activity in two connected regions in the brain: the subthalamic nucleus (STN) and external globus pallidus (GPe) [Merrison-Hort *et al.*, 2013]. These regions are both part of the “basal ganglia”, a group of nuclei that are known to be involved in movement control and that are severely affected by Parkinson’s Disease. A current subject of much interest is the fact that the Parkinsonian basal ganglia show a much larger degree

¹Named after the classic dynamical systems example of firefly phase synchronization, and also because the visualizations produced (slightly) resemble swarms of fireflies.

of rhythmically modulated (oscillatory) activity [Boraud *et al.*, 2005]. We therefore use our model to study the conditions under which oscillations can appear in the STN-GPe network.

The equations of our model are:

$$\tau_s \dot{x} = -x + Z_s(w_{ss}x - w_{gs}y + I). \quad (1)$$

$$\tau_g \dot{y} = -y + Z_g(-w_{gg}y + w_{sg}x). \quad (2)$$

In these equations, which are based on the Wilson–Cowan formulation [Wilson & Cowan, 1972], the variables x and y correspond to the average level of neuronal activity in the STN and GPe populations, respectively. The populations are coupled to each other and themselves through chemical connections (synapses), and the strength of these connections are given by the parameters w_{gg} , w_{gs} , w_{ss} and w_{sg} (so, for example, w_{sg} is the strength of the connection from STN to GPe). The STN population is excitatory, meaning it acts to increase the activity of its synaptic targets, while the GPe is inhibitory, meaning it acts to decrease the activity of its synaptic targets. The functions Z_s and Z_g are monotonically increasing sigmoid curves that describe how each population responds to synaptic input, and the parameters τ_s and τ_g are time constants that determine how quickly the activity in each population changes.

Here we will show how Fireflies can be used to investigate how the behavior of the system changes with one of its parameters. The parameter we will vary is w_{ss} , which corresponds to the strength of self-excitation in the STN. Gillies *et al.* [2002] demonstrated that if the neurons within the STN are able to excite each other (i.e. $w_{ss} > 0$) then the STN-GPe circuit is able to generate oscillations, although currently there is no known biological mechanism for such self-excitation. The simulations shown all contain two “particle groups”, one of which is integrated forwards in time and other backwards. Green and pink particles are used to render the forwards and backwards particle groups, respectively. Currently Fireflies supports only one solver method, fourth order Runge–Kutta, with a constant time step that can be adjusted by the user during the simulation. All particles are given random initial conditions, with values of the system variables drawn from a uniform random distribution on the interval $(0, 1)$; the values 0 and 1 correspond to the minimum and maximum levels of population activity respectively. Any trajectories

that leave this square region, or which have not been reset for more than time $T_{\max}$, are reset to a new random set of initial conditions.

Figure 1(a) shows a screenshot of the simulation when $w_{ss} = 0$, corresponding to the situation with no STN self-excitation. Under these conditions there is a single stable fixed point, towards which all of the green (forwards time) particles spiral in. The pink (backwards time) particles are difficult to see in this image, as they move very quickly out of the bounds of the system (to infinity) and are constantly being reset to new random initial conditions. It is interesting to note from this image that the density of particles approaching the spiral is not uniform, and has several discontinuities where the density suddenly changes. We suspected that this phenomena is a result of the boundaries of the phase space, and confirmed this with a traditional phase portrait where all the trajectories began on the edges of phase space [Fig. 1(e)]. The borders where particle density changes quickly correspond to the trajectories that begin in the four corners of phase space [red lines in Fig. 1(e)]; these four trajectories divide the phase plane into regions that “funnel” trajectories into the spiral.

The slider controls [visible in the bottom left of Fig. 1(a)] can be used to explore the effects of changing the parameters of the system. When the value of a parameter is changed using its slider, the dynamics of the system change immediately and the difference can be clearly seen in the movement of the particles. Figure 1(b) shows the STN/GPe system after the strength of STN self-excitation (w_{ss}) has been increased to 4.9, taking the system through a fold of limit cycles bifurcation. For this parameter value, a pair of limit cycles (stable and unstable) now encircle the original stable fixed point. This illustrates the purpose of the group of particles that are integrated backwards (pink): these particles are attracted to unstable objects, here they reveal the location of the unstable limit cycle. As w_{ss} is increased further, the unstable limit cycle shrinks around the fixed point and eventually disappears in a subcritical Andronov–Hopf bifurcation, at which point the fixed point becomes unstable. After this bifurcation the limit cycle is globally stable, but as w_{ss} is increased further the period of oscillation increases. This can be seen on the visualization in the form of particles “bunching up” and moving very slowly around one part of the cycle [Fig. 1(c)]. Eventually, a saddle-node on invariant

(a)

(b)

(c)

(d)

(e)

Fig. 1. Exploring the 2D STN/GPe model in Fireflies under variation of w_{ss} . (a)–(d) Show screenshots from Fireflies for different values of w_{ss} , with 700 000 particles, half of which are integrated forwards in time (green) and half of which are integrated backwards (pink). (a) Full screenshot of the Fireflies window with $w_{ss} = 0$. All forward-moving particles approach the globally stable spiral (in this still picture it is very difficult to see the backwards particles, which diverge quickly to infinity). (b) $w_{ss} = 4.9$: A stable spiral, unstable limit cycle, and stable limit cycle. (c) $w_{ss} = 7.8$: An unstable spiral and a stable limit cycle. (d) $w_{ss} = 11$: An unstable node, saddle (not visible), and stable node. (e) Phase portrait (not generated by Fireflies) of the STN/GPe system with $w_{ss} = 0$. All trajectories start from the borders of the phase plane and spiral in to the fixed point. The trajectories that start in the four corners of the phase plane are shown in pink.

circle (SNIC) bifurcation occurs, after which all trajectories approach a new, globally stable, fixed point.

Fireflies can also be used to generate three-dimensional animated bifurcation diagrams of two-dimensional systems. To see this, we redefine the system so that the bifurcation parameter (w_{ss}) is a new state variable, subject to $dw_{ss}/dt = 0$. The initial condition range for this new state variable is set to be the range of parameter values that are of interest. With this configuration, each particle has its own value for w_{ss} , and the space consists of a “continuum” of phase portraits, beautifully illustrating the dynamics of the system (Fig. 2). The user can further investigate the effects of the other parameters on the system’s dynamics by varying them interactively and observing how this changes the 3D bifurcation diagram.

2.2. Lorenz equations

In this example, we use Fireflies to visualize the dynamics of the 3D Lorenz system. We describe in detail the results of various visualizations obtained for different values of a parameter of the system, and then produce an animated bifurcation diagram that summarizes these results.

The classical Lorenz system consists of three coupled nonlinear differential equations:

$$\dot{x} = \sigma(y - x), \quad (3)$$

$$\dot{y} = x(r - z) - y, \quad (4)$$

$$\dot{z} = xy - \beta z, \quad (5)$$

where σ , r and β are parameters of the system ($\sigma, r, \beta > 0$). These equations were originally studied as a simplified model of convection in the atmosphere [Lorenz, 1963], and a straightforward description of the different dynamical regimes that they can produce, including chaos, can be found in any textbook on nonlinear dynamics (e.g. [Strogatz, 1994, pp. 311–320]). In this section, we will demonstrate the ability of Fireflies to visualize three-dimensional systems by exploring the behavior of the Lorenz equations in the case where $\sigma = 10$ and $\beta = \frac{8}{3}$, while the parameter r is gradually increased from zero.

For values of r less than 1, the origin is the only stable fixed point. Figure 3(a) shows Fireflies when $r = 0.5$, shortly after the simulation has started and all the particles are quickly approaching the origin. As was also seen in the 2D system described above, Fireflies shows how the space occupied by the cube of initial conditions is deformed around the fixed point as the particles move. Note that we do not show any backwards-moving particles in this section, as they are not generally useful in visualizations of the Lorenz equations. This is because any unstable fixed point or cycle are of saddle type, which means that neither forward nor backward moving particles tend to them as $t \rightarrow \infty$.

At $r = 1$, a supercritical pitchfork bifurcation occurs, and two new stable fixed points emerge from the origin. Figure 3(b) shows a screenshot from Fireflies with $r = 5$. With the new parameter value the origin becomes a saddle, and particles begin to move away from it in one of two directions, spiraling

Fig. 2. 3D bifurcation diagram of the STN/GPe system under variation of the parameter w_{ss} . The simulation contains eight million particles, divided into forwards in time group (green) and backwards in time group (pink). The three visible bifurcations are marked with asterisks, from left to right: fold of limit cycles, subcritical Andronov–Hopf, saddle-node on invariant circle. Note how the particles on the stable limit cycle become increasingly “bunched” at the top as the cycle approaches the SNIC bifurcation. The asterisks and direction arrows were added to the image manually and were not generated by Fireflies.

Fig. 3. Exploring the Lorenz equations by varying r . All particles are integrated forwards in time, different colors use different initial condition ranges. Asterisks approximately indicate the origin. (a) $r = 0.5$: All particles approach the stable fixed point at the origin. The complex shape is due to the boundaries of the initial conditions ($x \in (-10, 10), y \in (-30, 30), z \in (0, 50)$) and the shape of the fixed point's manifolds. (b) $r = 5$: The origin is now a saddle and two new stable fixed points have emerged. Both groups of particles have initial conditions near to the origin, but on opposite sides of its incoming separatrix (green: $x, y \in (0, 0.01), z \in (0, 0.01)$, blue: $x, y \in (-0.01, 0), z \in (0, 0.01)$). Trajectories approach one of the spirals, based on which side of the separatrix they start on. (c)–(d) $r = 15$: Trajectories now switch between the two spirals for some time, before settling down and approaching one of them (transient chaos). (c) Initial conditions near the saddle, as in (b); these trajectories loop once around one spiral before approaching the other. (d) Initial conditions further away from the saddle; these trajectories can switch sides repeatedly. (e) $r = 25$: A strange attractor. Trajectories starting near the saddle (green and blue) begin by approaching one of the spirals, before slowly spiraling away and swapping sides chaotically. Another group of particles (bronze) with initial conditions spanning a large area show the shape of the attractor.

in to one of the new stable fixed points that were created in the bifurcation. To clearly show how trajectories spiral in to the fixed point, this simulation has two groups of particles (green and blue) with initial conditions that are drawn from separate small volumes in state space, both very close

to the saddle at the origin, but on opposite sides of its incoming separatrix. The paths that the particles in these two groups take away from the saddle are very close to the saddle's outgoing separatrices. Trajectories starting at other points in state space all approach the spiral that is on the same

side of the saddle's incoming separatrix as their initial position.

At $r \approx 13.926$ a homoclinic bifurcation occurs, at which point the saddle's outgoing separatrices join up with its incoming ones, resulting in the creation of a pair of unstable limit cycles. For values of r beyond the bifurcation the separatrices have "crossed over", and trajectories that begin near the saddle make one cycle around their nearest spiral before returning back towards the saddle, and then spiraling in to the stable fixed point on the opposite side of the incoming separatrix, as shown in Fig. 3(c). Trajectories that start a bit further away from the saddle, however, rotate around their closest spiral at a much greater distance. When this rotation brings them close to the saddle's incoming separatrix, some of them split off and begin rotating around the opposite spiral. This unpredictable swapping, which corresponds to transient chaos, can happen many times before a particle finally spirals fully in to one of the two stable fixed points. Figure 3(d) shows one group of particles which all start closer to the top spiral than the bottom one, but a little further away from the origin than those in Fig. 3(c). Although most spiral in to the top fixed point, with each rotation a mass of particles splits off and switches to rotate around the bottom spiral.

Finally, at $r \approx 24.06$ a strange attractor appears. Now, even though the two spirals remain stable until $r \approx 24.74$, almost all trajectories flip backwards and forwards between the two spirals infinitely and chaotically. In Fig. 3(e), the two groups of particles that start near to the origin (green and blue) initially behave as in Fig. 3(c), performing one large loop around their closest spiral before approaching quite close to the opposite spiral. The spirals are only very weakly stable now, however, and the trajectories spiral slowly away from them, instead of into them. After some time [as is just beginning to happen in Fig. 3(e)], the particles come far enough away from the spiral that some of them switch sides, beginning an endless series of such seemingly random side swappings on the strange attractor. The figure also contains a third set of particles (bronze colored) which start from a much wider set of initial conditions; these show the general shape of the strange attractor's surface.

By applying the same technique as in the previous section, we can also use Fireflies to generate an animated bifurcation diagram for the Lorenz equations. To achieve this, we make parameter r

Fig. 4. A "live" bifurcation diagram for the Lorenz equations. The diagram was created by setting up a simulation with eight million particles, each with a random fixed value for r , projected onto the (r, y) plane. The supercritical pitchfork at $r = 1$ is clearly visible, as is the appearance of a strange attractor at $r \approx 13.9$ and various transient windows of nonchaotic behavior such as at $r = 92$. Axes added manually.

a state variable with $dr/dt = 0$ and set up a two-dimensional projection with axes r, y . Note that a 3D projection could also be used, but due to the perspective transformation the results are not shown as clearly in this case. Figure 4 shows a screenshot from a simulation using this configuration, where each particle's initial value of r is chosen from the interval $(0 \cdots 110)$. The supercritical pitchfork bifurcation is clearly visible at $r = 1$, and the appearance of a strange attractor at $r \approx 13.9$ is marked by particles beginning to form a cloud that resembles noise. This cloud clearly has some very detailed structure, however, as can be seen more and more clearly as r increases. Several small windows of parameter values where the dynamics become regular are also visible, for example, at $r = 92$. Again, it should be kept in mind that Fig. 4 is a static screenshot of a moving animation of eight million particles. The other two parameters of the system can be changed interactively, and the effect of this variation on the bifurcation diagram is seen immediately.

3. Coupled Hodgkin–Huxley Neurons

In this section, we demonstrate the use of Fireflies to visualize a system that is considerably more complex than those presented above, consisting of 15 coupled ODEs that describe the activity in three synaptically connected neurons.

In 1963, Alan Hodgkin and Andrew Huxley received a Nobel prize for their discovery of the mechanism that allows neurons to generate action potentials (the electrical impulses, or “spikes”, that neurons use to communicate with each other). More than 50 years later, the general structure of the equations that Hodgkin and Huxley used to describe this mechanism still forms the basis of an enormous number of biologically realistic computational models of neuronal activity. Here, however, we will use the original equations and parameters, which specifically relate to biophysical activity in the giant axon of the squid [Hodgkin & Huxley, 1952]. In this model, the electrical potential across a part of the neuron’s membrane evolves according to the flow of sodium (Na) and potassium (K) ions through the membrane. The permeability of the membrane to these ions is controlled by gates, which open and close according to the membrane potential. Equations (6)–(9) describe the “classical” Hodgkin–Huxley model of neuronal dynamics in a population of N neurons.

$$C\dot{V}_i = \bar{g}_{lk}(e_{lk} - V_i) + h_i m_i^3 \bar{g}_{na}(e_{na} - V_i) + n_i^4 \bar{g}_k(e_k - V_i) + I_{\text{syn}}^i(t) + I_i. \quad (6)$$

$$\dot{h}_i = \alpha_h(V_i)(1 - h_i) - \beta_h(V_i)h_i. \quad (7)$$

$$\dot{m}_i = \alpha_m(V_i)(1 - m_i) - \beta_m(V_i)m_i. \quad (8)$$

$$\dot{n}_i = \alpha_n(V_i)(1 - n_i) - \beta_n(V_i)n_i. \quad (9)$$

$$i = 1, 2, \dots, N$$

Here V_i is the membrane potential of the i th neuron and h_i , m_i and n_i represent the average state of the gates on its ion channels ($0 \leq h_i, m_i, n_i \leq 1$). The parameters C , $\bar{g}_{lk}$, $\bar{g}_{na}$, $\bar{g}_k$, e_{lk} , e_{na} and e_k , along with the functions $\alpha_X(\cdot)$ and $\beta_X(\cdot)$ ($X \in \{h, m, n\}$), are described in [Hodgkin & Huxley, 1952] or any computational neuroscience textbook. All parameters mentioned so far take the values given in [Hodgkin & Huxley, 1952]. Parameter I_i controls how much external current is injected into the cell — increasing this parameter causes a transition from quiescence to single action potential firing to repetitive firing. Finally, I_{syn}^i represents the total synaptic current that arises as a result of inputs from the other neurons in the model. If $I_i = I_{\text{syn}}^i = 0$, then the membrane potential of neuron i approaches a fixed point at the “resting” potential of 0 (mV). In this example, we use a network where the neurons are arranged in a loop, with

each receiving synaptic input from only the previous neuron:

$$I_{\text{syn}}^i(t) = \bar{g}_{\text{syn}}(e_{\text{syn}} - V_i)s_{(i-1) \bmod N}. \quad (10)$$

$$\dot{s}_i = \tau_r^{-1}(1 + \exp(-\sigma(V_i - \theta)))^{-1} \times (1 - s_i) - \tau_d^{-1}s_i. \quad (11)$$

The parameter $\bar{g}_{\text{syn}}$ is the maximum conductance (mS) of a synaptic connection (its “strength”); each synaptic connection has the same strength in this model. e_{syn} is the synaptic equilibrium potential: if $e_{\text{syn}} > 0$ (mV) then synaptic input is excitatory and acts to raise the membrane potential of the post-synaptic neuron above the resting state; in this example we set $e_{\text{syn}} = 10$. For each neuron we consider a variable s_i , where $0 < s_i < 1$, which acts as an indicator for when the neuron is spiking. This variable increases to 1 very rapidly (with time constant τ_r) when the neuron is firing an action potential, and then decays slowly (time constant τ_d) after a spike, as described by Eq. (11). Parameters σ and θ are the slope and shift respectively of the sigmoid function which is used to increase s_i .

We begin by briefly showing a visualization of the 4D phase space of a single independent neuron under varying strengths of current injection. For low values of I_1 all trajectories approach the fixed point which corresponds to the resting state. As I_1 is increased past $I_1 \approx 6.25$ the resting state undergoes an Andronov–Hopf bifurcation and a stable limit cycle appears. This limit cycle corresponds to periodic spiking with a frequency that increases with I_1 . Figure 5 shows a screenshot of a simulation with a single neuron and $I_1 = 10$. The three dimensions of the projection are the three gating variables: h_1 , m_1 and n_1 . All the particles in this simulation eventually approach the stable limit cycle. However, there is also a spiral-shaped cloud of particles (marked with an asterisk), which is made up of trajectories that start close to the formerly stable fixed point. For parameter value $I_1 = 10$ this fixed point is only weakly unstable, so trajectories nearby spend a long time oscillating with a low (gradually increasing) amplitude before they approach the limit cycle.

Next, we consider a ring of $N = 3$ neurons [Fig. 6(a)], where each neuron receives an identical current injection that is strong enough to cause repetitive firing ($I_j = 10, j = 1, 2, 3$). Such networks of coupled regular spiking neurons are typically able

Fig. 5. A screenshot of Fireflies running the single Hodgkin-Huxley neuron model with 500 000 particles. The position of each particle is projected into the 3D space (h_1, m_1, n_1) . Particles are colored according to their position on each axis. Asterisk indicates the approximate position of the (weakly) unstable fixed point, with a cloud of particles moving slowly away from it. Axes added manually.

to synchronize their firing in either in-phase or anti-phase, depending on the strength and nature of synaptic connections [Mirollo & Strogatz, 1990]. We set $e_{\text{syn}} = 10$ (mV) and $\bar{g}_{\text{syn}} = 0.5$ (mS) to simulate the case of weak excitatory coupling. From running a few simulations of the system from random initial conditions in XPPAUT [Ermentrout, 2002] it was clear that with these parameter values there was a very stable synchronous state, where all three neurons fired regularly with the same frequency and phase shift. Figure 6(b) shows a screenshot of Fireflies after this system has been simulated for long enough that all particles have settled down onto stable attractors. Each particle's position projected onto the space (V_1, V_2, V_3) . The color of each particle varies with its position, using a projection from the 3D coordinates to an RGB color code. The prominent solid white diagonal line is the synchronous limit cycle; most particles are attracted to this and continually move along the diagonal as the three neurons spike in unison. Figure 6(c) shows the synchronous state as produced by XPPAUT.

In addition to the synchronous state, however, the visualization also reveals four other stable limit

cycles, three of which are easily visible with 500 000 particles. The three limit cycles appear as six colored prongs in Fig. 6(b). By interactively moving around 3D space we were able to see that the particles were organized into three closed orbits, each of which included two differently colored prongs. Each of these closed orbits corresponds to a limit cycle where two of the neurons are spiking and the third is silent. Simulations in XPPAUT revealed that each of these limit cycles corresponds to the state where the presynaptic neuron fires first, followed by the postsynaptic neuron, followed by a pause before the cycle repeats; this is shown in Fig. 6(cii). This detail about the order in which the neurons fire is not immediately visible in the Fireflies visualization.

The other stable limit cycle that we have found has a relatively small basin of attraction, meaning that not many of the 500 000 particles approach it. It is also quite difficult to see in still screenshots, although some particles on this cycle are marked with an asterisk in Fig. 6(b). By following the path of these particles around the visualization space, we can see that their trajectories correspond approaching a stable limit cycle corresponding to the state where all three neurons spike in sequence. Interestingly, this order of this sequence is opposite to that of the direction of synaptic coupling (i.e. the firing sequence is neuron 1, neuron 3, neuron 2, ...). In order to verify that this was indeed a real limit cycle and not a numerical error in Fireflies, we wrote a short script to repeatedly (and sequentially) run simulations from random initial conditions in XPPAUT. The script recorded any set of initial conditions that lead to solutions where all the neurons fired repetitively but nonsynchronously. After performing a very large number of simulations this script had found only two sets of initial conditions that lead to the limit cycle that we had identified in Fireflies, and the resulting plots of $V(t)$ are shown in Fig. 6(ciii).

The example in this section demonstrates a case where Fireflies could reveal several stable dynamical regimes that were not immediately obvious from running individual simulations of the system from random initial conditions. Due to the fact that the system under study is composed of three identical elements, it is not surprising to find that multiple regimes corresponding to different symmetries are possible. Furthermore, it is likely that other limit cycles exist — for example, there may be one in which all three neurons fire in sequence in order of

Fig. 6. A Hodgkin-Huxley model with three neurons spiking repetitively in response to current injection, connected by weak excitatory synapses. (a) An overview of the model's structure. (b) A screenshot of Fireflies running the model with 500 000 particles. The position of each particle is projected into the 3D space (V_1, V_2, V_3) . Particles are colored according to their position on each axis. Asterisk indicates a small group of particles that are close to the limit cycle where each neuron spikes in turn. Axes added manually. (c) Individual plots (from XPPAUT) of trajectories starting from three different sets of initial conditions, leading to different limit cycles. (i) Synchronous spiking; (ii) two neurons spiking, one suppressed; (iii) neurons spiking in turn.

synaptic coupling, rather than in the opposite order as we observed here. However, because none of the hundreds of thousands of particles in our simulations approached such attractors, it is likely that they are either unstable or have extremely small basins of attraction.

4. Discussion

To facilitate the qualitative study of systems of ODEs, software such as XPPAUT [Ermentrout, 2002], often includes the ability to plot trajectories in phase space that start from multiple sets of initial conditions, chosen either from a regular grid in phase space or stochastically. However, in systems with high-dimensional phase space (many ODEs) or with finely structured dynamical regimes, a very large number of initial conditions must be used in order to have a high level of confidence that all stable attractors have been found. This can make the qualitative investigation a very slow process, especially if one wishes to also examine how the picture

varies with parameter values. It is also very difficult to visualize many different trajectories on the same phase portrait, as large numbers of them will completely fill the phase space. We believe that Fireflies offers a useful and exciting new way to investigate dynamical systems. This approach is more intuitive than traditional phase portraits, as it presents dynamical systems as they are: *dynamic*.

In this paper we have demonstrated, through several examples, the power of Fireflies for qualitatively exploring the behavior of dynamical systems. The visualization of systems in explorable 2D or 3D space provide a new perspective on such systems that can help with intuitively understanding them. In addition to this, since Fireflies can simulate millions of trajectories in parallel, stable attractors can be discovered that might not have been found with other methods of qualitative investigation — as in Sec. 3 (although clearly as the dimensionality of systems increases, the number of particles required to uniformly sample phase space with a given resolution increases exponentially).

Table 1. The time taken to advance each of the systems shown in this paper by a single step (averaged over 1000 steps), when executed on a single core CPU and a GPU. The best case performance for a quad-core CPU was estimated by dividing the single core figure by four. The CPU used was an Intel Core i5-2500k (3.3 GHz) and the GPU was an Nvidia Geforce GTX 460. The C code for the CPU benchmark (available at www.bitbucket.org/rmerrison/fireflies) was adapted from the corresponding OpenCL kernel for each system, and was compiled using GCC 4.8.2 on Ubuntu 14.04 using the `-Ofast` option.

System	Num. Particles	Single Core CPU	Quad Core CPU (est.)	GPU
STN-GPe	700 k	182 ms	46 ms	1 ms
Lorenz	3 m	109 ms	27 ms	3 ms
Hodgkin–Huxley	500 k	942 ms	236 ms	22 ms

Scientific computing is increasingly taking advantage of the power and availability of GPUs, and a range of techniques have emerged for using GPUs to solve problems in domains such as computational neuroscience [Brette & Goodman, 2012] and fluid dynamics [Harris, 2004]. To demonstrate the necessity of “massively parallel” GPU computing for producing Fireflies’ visualizations, we compared the average time taken to compute a single integration step using a traditional CPU (single core and estimated quad core) to the time taken when integration is performed on the GPU. Table 1 shows the results of this comparison. For interactive graphical applications, it is normally considered necessary to render at least 30 frames per second in order to give the user a smooth experience, which means that each frame must be generated in 33 ms or less. Since the times shown in Table 1 only represent the time taken to advance the simulation and do not include additional time needed to render particles to the screen, it would seem to be very difficult or impossible to run visualizations shown in this paper at an acceptable frame rate without using the GPU. Additionally, running the simulation on the CPU would also incur the significant additional overhead of passing the particles’ current positions to the computer’s graphical hardware for rendering; by performing integration on the GPU, Fireflies minimizes the number of CPU–GPU memory transfers.

We have found that our new visualization technique can be particularly useful in a teaching context, as it very easily demonstrates the behavior that a particular set of equations produces. Visualization of simple systems, such as the 2D model of neuronal activity illustrated in Sec. 2.1, can be used

to show different bifurcations in a very intuitive way. For example, the parameter w_{ss} is increased, Fireflies shows how particles move more and more slowly around one part of the limit cycle, becoming increasingly bunched together until finally being attracted into a new stable fixed point that appears in a SNIC bifurcation.

The ability of Fireflies to quickly simulate millions of trajectories at the same time is due to the fact that almost all processing occurs exclusively within the GPU, but this approach carries a number of disadvantages. In particular, due to the time taken to transfer data between GPU and CPU it is not possible to permanently record trajectories (e.g. for further analysis), however in future we plan to add the ability to record either a subset of trajectories, or all trajectories but with a recording interval that is greater than the time step. It should be noted, however, that this is not the main intended purpose of Fireflies, and other libraries for exploiting GPUs to solve ODEs, such as the Odeint library for C++ [Ahnert & Mulansky, 2011], may be more appropriate. Another limitation that arises from the use of the GPU is that Fireflies does not allow the equations of the system to contain branching (e.g. changing variable values in response to threshold crossing). This is because GPUs are based on a “single instruction multiple data” (SIMD) architecture, whereby each processing core executes the same instruction at the same time; due to this architecture branching causes a significant reduction in computation speed.

Fireflies is written in Python and utilizes a number of cross-platform and open source libraries, namely NumPy [van der Walt *et al.*, 2011], PyOpenGL,² PyOpenCL [Klöckner *et al.*, 2012],

²<http://pyopengl.sourceforge.net/>

PySide³ and Mako.⁴ The software has been tested in Linux and Windows, and should also work in Apple OSX. The full source code for Fireflies, along with instructions for its use, can be found at <https://bitbucket.org/rmerrison/fireflies>.

Acknowledgments

The author is extremely grateful to Roman Borisyuk for the valuable input that he gave during the development of Fireflies and the preparation of this manuscript. Funding was also received from BBSRC Grant BB/L000814/1.

References

- Ahnert, K. & Mulansky, M. [2011] “Odeint — Solving ordinary differential equations in C++,” <http://arXiv.org/abs/1110.3397v1>.
- Boraud, T., Brown, P., Goldberg, J., Graybiel, A. & Magill, P. [2005] “Oscillations in the Basal Ganglia: The good, the bad, and the unexpected,” *The Basal Ganglia VIII*, eds. Bolam, J., Ingham, C. & Magill, P., Section 1 (Springer Science and Business Media), pp. 1–24.
- Brette, R. & Goodman, D. F. M. [2012] “Simulating spiking neural networks on GPU,” *Network: Comput. Neur. Syst.* **23**, 167–182.
- Dhooge, A., Govaerts, W. & Kuznetsov, Y. A. [2003] “Matcont: A Matlab package for numerical bifurcation analysis of ODEs,” *ACM Trans. Math. Softw.* **29**, 141–164.
- Doedel, E. J. & Oldeman, B. E. [2012] “Auto-07p: Continuation and bifurcation software for ordinary differential equations,” Tech. rep., Concordia University, Montreal.
- Ermentrout, B. [2002] *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to XPPAUT for Researchers and Students* (SIAM, Philadelphia, PA, USA).
- Gillies, A., Willshaw, D. & Li, Z. [2002] “Subthalamic-pallidal interactions are critical in determining normal and abnormal functioning of the basal ganglia,” *Proc. Roy. Soc. Lond. Ser. B: Biol. Sci.* **269**, 545–551.
- Harris, M. J. [2004] “Fast fluid dynamics simulation on the GPU,” *GPU Gems* (Addison-Wesley), pp. 637–665.
- Hodgkin, A. L. & Huxley, A. F. [1952] “A quantitative description of membrane current and its application to conduction and excitation in nerve,” *J. Physiol.* **117**, 500.
- Klößner, A., Pinto, N., Lee, Y., Catanzaro, B., Ivanov, P. & Fasih, A. [2012] “PyCUDA and PyOpenCL: A

- scripting-based approach to GPU run-time code generation,” *Parall. Comput.* **38**, 157–174.
- Lorenz, E. N. [1963] “Deterministic nonperiodic flow,” *J. Atmosph. Sci.* **20**, 130–141.
- Merrison-Hort, R. J., Yousif, N., Njap, F., Hofmann, U. G., Burylko, O. & Borisyuk, R. [2013] “An interactive channel model of the basal ganglia: Bifurcation analysis under healthy and Parkinsonian conditions,” *J. Math. Neurosci.* **3**.
- Mirollo, R. E. & Strogatz, S. H. [1990] “Synchronization of pulse-coupled biological oscillators,” *SIAM J. Appl. Math.* **50**, 1645–1662.
- Pavlidis, A., John Hogan, S. & Bogacz, R. [2012] “Improved conditions for the generation of beta oscillations in the subthalamic nucleus-globus pallidus network,” *European J. Neurosci.* **36**, 2229–2239.
- Strogatz, S. H. [1994] *Nonlinear Dynamics and Chaos* (Westview Press).
- van der Walt, S., Colbert, S. & Varoquaux, G. [2011] “The Numpy Array: A structure for efficient numerical computation,” *Comput. Sci. Engin.* **13**, 22–30.
- Wilson, H. & Cowan, J. [1972] “Excitatory and inhibitory interactions in localized populations of model neurons,” *Biophys. J.* **12**, 1–24.

Appendix A

Implementation Details

In order to simulate a system of ODEs in Fireflies, the user must specify the *N State Variables*, along with the *Parameters* of the system, and one or more *Render Techniques* and *Particle Groups*:

- Each of the *N State Variables* corresponds to one of the system’s ODEs and comprises a mathematical expression describing the variable’s time derivative (right-hand side), along with the bounds of the variable. If a trajectory passes outside of bounds of any state variable during simulation it is reset to new random initial conditions.
- For each *Parameter* the user specifies a default value and the range of allowable values for the parameter.
- A *Render Technique* specifies a method for transforming the current position of each trajectory in phase space into a position of a particle on the screen. At present, particles can either be drawn using a simple 2D projection or a perspective 3D projection, and when specifying a projection the user must select which state variable

³<http://www.pyside.org>

⁴<http://www.makotemplates.org>

corresponds to each of the (two or three) projection dimensions. The color of rendered particles is also specified as part of a render technique; this can either be manually specified or varied automatically based on the position of the particle in 2D or 3D space.

- Each *Particle Group* consists of P particles that share the same render technique and range of initial conditions. The user specifies: the number of particles; the render technique that should be used to draw them; whether the particles in the group should be simulated forwards or backwards in time; and, for each state variable, a range of valid initial conditions. These initial condition ranges define a cube in phase space from which a position for each particle in the group is chosen when the particle is first initialized (or reset as a result of going out of bounds).

These objects that describe a system of equations are specified using the software's graphical user interface (Fig. 7). Once this has been done, Fireflies initializes the system before entering its main execution loop, which consists of repeatedly updating the particles' positions and rendering them to the screen. An overview of how the software is organized is shown in Fig. 8. The following sections describe each stage of processing in more detail.

A.1. Initialization

When a system is first created, or subsequently modified, the GPU must be configured to simulate it. This involves populating the GPU's memory with the particles' initial positions and compiling and uploading programs to update and render the

Fig. 7. A screenshot of the full Fireflies window showing the panel where the system to be simulated (the Lorenz equations in this case) is defined. Here, a state variable has been selected from the list of objects on the left-hand side, allowing the properties of that state variable to be set using the panel on the right-hand side.

Fig. 8. A schematic overview of Fireflies. The processes at the top of the figure occur during initialization, and involve compiling an OpenCL kernel to advance the simulation, and generating random initial conditions. The main (lower) part of the figure shows the program's main loop, during which particles are repeatedly moved and then rendered to the screen. Note that almost all of the processing in the main loop occurs entirely within the GPU, avoiding the overhead of transferring data between the GPU and CPU. The exception to this is the resetting procedure, which only processes a subset of particles each frame.

system onto the GPU. Setting up the particles' initial positions is straightforward: if the system is N -dimensional and consists of a total of P particles, then an array with space for $N \cdot P$ single-precision floating point values is created and filled with random initial conditions (subject to the initial condition ranges specified by each particle's *Group*) — this array is then transferred into the GPU's memory. Note that there are two possible ways of organizing this memory: the first involves storing the N values associated with the first particle, followed by those for the second particle, and so on (we call this “row-major” order), while the second method involves storing the P values corresponding to the particles' positions in the first dimension, followed by the values for the second dimension, etc. (we call this “column-major” order). At present Fireflies always uses row-major order, however it is likely that due to the way in which GPU memory access operates, column-major may

gives significant performance benefits; this is something we plan to investigate in a future version of the software.

In order to simulate the system on the GPU, a program (or “kernel”) must be compiled that updates the position of a single particle. The general form of this kernel is independent of the particular system of equations being integrated, and involves reading the particle's current position from memory, determining its new position using a fourth order Runge–Kutta routine, and then writing the new position back into memory. The only part of the kernel that changes for different systems of equations is that which calculates the time derivative for a given point in state space (the right-hand side of each ODE). This function is generated automatically using the definitions for the state variables given by the user. In order to convert the specified system into kernel source code (in the “OpenCL C” language), Fireflies uses a template source code file which is processed by the Mako templating library⁵ to produce the final kernel source. This source code is then compiled and uploaded onto the GPU using OpenCL.

In addition to the kernel, three other programs (or “shaders”) must be compiled onto the GPU which performs the transformations necessary to render the particles to the screen. These shaders form a pipeline that converts from particle positions in state space into a series of squares (pairs of triangles) in 2D screen space that can be drawn by the graphics hardware. The different stages of this pipeline are the vertex shader, the geometry shader, and the fragment shader.

- The *vertex shader* takes as input the two or three components (depending on whether a 2D or 3D projection is required) of the particle's current position that correspond to the X , Y and Z directions of the scene to be rendered and outputs a 2- or 3D vector containing the position of the point transformed to be relative to an imaginary camera. This transformation is a simple linear transformation obtained by multiplying the vertex position by the “Model-View” matrix. In the case of a 2D view, the camera always points directly at the plane containing the particles but it can be translated around and zoomed in and out. With a three-dimensional view, the user can

⁵<http://www.makotemplates.org>

interactively adjust the camera's position and orientation in space in order to "fly around" the system. As the user moves around the system, the Model-View matrix is updated accordingly.

- The *geometry shader* takes as input the camera-relative particle position from the vertex shader and first transforms this into 2D position, using the "View-Projection" matrix. If the particle is to be rendered as a 2D projection then this transformation is straightforward, while for a 3D projection a standard perspective transformation matrix is used (for details of this see any 3D graphics programming resource). After the particle position has been projected into 2D, the geometry shader outputs four 2D vertices that correspond to a square that is centered on the particle position and facing the camera.
- The *fragment shader* is the final stage of the shader pipeline. When the graphics hardware renders the squares generated by the geometry shader on to the screen, it determines the color of each pixel within the square by calling the fragment shader. The fragment shader generates an intensity value based on the distance from the center of the square, producing a smooth circular particle shape. The actual color output by the fragment shader is either specified explicitly by the particle's render technique, or varies linearly as a function of the particle's position in state space. The color of the pixel on the screen is set to its current color plus the color contributed by the particle being drawn.⁶

A.2. Running the simulation

Once the system has been specified and the GPU initialized with the kernel, shaders and initial particle positions, the simulation can be started by making the step size nonzero (both positive and

negative step sizes are possible, to simulate the system either forwards or backwards in time).

When the simulation is running, the main simulation loop repeatedly instructs the GPU to alternately execute the kernel and then the shader pipeline, causing the particles to be updated and drawn to the screen. The user can see, in real time, how the particles are moving around the state space and over time the system's attractors become clear. The user can explore the state space by changing the position of the camera in 2D or 3D space (effectively changing the Model-View matrix in GPU memory) using the mouse and keyboard. The system's parameters can also be changed during simulation by dragging sliders on the user interface, which updates the location in GPU memory where the corresponding parameter value is stored.

Some operations, such as checking for and resetting particles that have moved outside the bounds of the system to new random positions are not well-suited for GPU processing; Fireflies performs these operations on the CPU. Since copying data between the CPU and GPU is relatively slow, it is not practical to copy the positions of all of the particles back and forth at every time step. Instead, only a fraction of the particle position data is copied and processed by the CPU for each frame. The CPU processes this batch of particles by checking if any of them have left the bounds of the system and chooses new initial conditions where necessary, before copying the affected particles' positions back to the GPU. This method ensures that simulations run at a smooth frame rate, although it means that particles are typically reset several time steps after their reset conditions are first satisfied. The user can adjust the size of CPU processing batches while the simulation is running, in order to find the setting that maximizes performance.

⁶This corresponds to `glBlendFunc(GL_SRC_ALPHA, GL_ONE)`.