

Improving the FLAME GPU User Experience through Examples

PI: Paul Richmond (Professor of Research Software Engineering, University of Sheffield)

dRTPs: Peter Heywood (RSE in RSE Sheffield group), **Robert Chisholm** (Senior RSE in RSE Sheffield group)

FLAME GPU is a powerful open-source software framework that helps scientists and engineers simulate how large numbers of individual “agents” such as people in a city, cars in traffic, or cells in the body interact with one another. What makes FLAME GPU unique is its ability to execute these simulations on graphics processing units (GPUs), the same hardware used for gaming and artificial intelligence, allowing it to model millions of agents quickly and efficiently. The use of GPU high-performance computing enables researchers to explore complex systems, test scenarios, and better understand how simple individual behaviours can lead to large-scale patterns in the real world.

The overarching goal of the project is to increase the user base of the FLAME GPU software by improving the accessibility and hence user experience through the development of new example models. The project will implement a number of example models which are foundational within the agent-based approach. These will be adopted from examples in the NetLogo software and implemented within FLAME GPU as examples and starting points for users to develop their own models. Emphasis will be placed on Social Science and Humanities modelling to increase the user uptake of the software within these respective communities. Models will be executed and benchmarked on the Bede supercomputer, in particular the Grace-Hopper nodes with results expected to be published in the NVIDIA Technical Blog series (through our NVIDIA Applied Research Accelerator Award).

The intern will develop the models using our new Python based bindings for FLAME GPU which allow modellers to develop models using pure Python and avoiding the need for C++. Increasing the number of Python examples is expected to facilitate uptake of the software by inexperienced research developers creating a lower barrier of entry for HPC usage. The intern will have access to a self-hosted GPU development machine within the RSE lab, will be provided a laptop if necessary and will also be provided access to Sheffield Stange HPC system as well as Bede for benchmarking. The intern will join weekly FLAME GPU development meetings and will working closely with dRTPs developing the core FLAME GPU codebase (Robert Chisholm and Peter Heywood).

Schedule of Work

- Week 1-2: FLAME GPU familiarisation (through tutorials and mentoring), NetLogo model selection for porting.
- Week 2-6: Porting of NetLogo models (2-4 in total) to pyFLAMEGPU. Support and mentoring.
- Week 6-7: Development of a social science tutorial model using Python notebook (following existing tutorial templates).
- Week 7-8: Benchmarking of the models and comparisons with NetLogo, reporting and drafting of NVIDIA Development Blog article.
- Post Internship: The PI and dRTPs will finalise and agree the blog post content and benchmark results for publication.