Exascale agent-based modelling for policy evaluation in real-time (ExAMPLER)*

Gary Polhill[1][0000−0002−8596−0590], Alison Heppenstall[2][0000−0002−0663−3437], Mike Batty[3][0000−0002−9931−1305], Matt Hare[1][0000−0002−8508−0884], Doug Salt[1][0000−0001−5186−9388], Richard Milton[3][0000−0003−4135−094X], and Ricardo Colasanti[2][0000−0001−7718−5136]

[1] The James Hutton Institute, Aberdeen AB15 8QH, United Kingdom
{gary.polhill,matt.hare,doug.salt}@hutton.ac.uk
https://www.hutton.ac.uk/social-simulation/
[2] University of Glasgow, Glasgow G12 8RS, University
{alison.heppenstall,ricardo.colasanti}@glasgow.ac.uk
https://www.gla.ac.uk/schools/socialpolitical/staff/alisonheppenstall/
[3] Bartlett Centre for Advanced Spatial Analysis, University College London, London WC1E 6BT, United Kingdom
{m.batty,richard.milton}@ucl.ac.uk
https://www.ucl.ac.uk/bartlett/casa/

Keywords: Exascale computing · Agent-Based Modelling · Policy evaluation.

Abstract

Though there are a handful of exascale computers already, exascale computing is expected to become more widely available towards the latter half of this decade. We are gathering requirements for exascale computing for agent-based modelling to support policy evaluation. Exascale computing has a transformative potential in the area as it entails such a dramatic improvement in computing time that even work requiring days or weeks on a contemporary high-performance computing cluster would take a second or two on an exascale computer. This creates a context in which it is possible to conceive of creatively interacting with an existing agent-based model as part of a one-day transdisciplinary workshop. However, the reduction in time to experiment with a model then puts emphasis on the time taken to assemble and prepare data, and to build the model in the first place. Full exploitation of the potential of exascale computing therefore entails a more in-depth look at the software, data and institutional environment in which the agent-based models of the future will be prepared.

Exascale computing is defined with respect to the speed with which numerical calculations can be completed as a billion-billion (10^{18}) floating-point operations per second. Though instructions executed per second is of greater relevance to

* Supported by the Engineering and Physical Sciences Research Council (project reference EP/Y008839/1) and the Scottish Government Rural and Environment Science and Analytical Services Division (project reference JHI-C5-1)

agent-based modelling than purely floating-point operations, the two are approximately the same. Exascale computers are about a billion times faster than a personal computer, and a thousand to a million times faster than a university's computing cluster or industry cloud computing services such as Microsoft Azure, Amazon AWS or Google Cloud – each generally referred to as 'high performance computing' (HPC) facilities.

Use of HPC is rare in the agent-based modelling community, for various reasons. First is the steep learning curve that authors such as Alessa et al. [1] and An et al. [2] have pointed out is faced by social scientists simply to build a working model in the first place. HPC environments typically require a great deal more computing skills to access; a technical barrier that is somewhat needless. Second is the problem raised by Polhill [3] that agent-based models have computational features (such as varying population sizes) that make their memory and CPU-time unpredictable. These two points combine to exacerbate a third, cultural issue around HPC around the suitability of the code and the scientific problem it is investigating for the advanced computing machinery on which it will be run. These are points that we can expect to be more pronounced in exascale environments that cost hundreds of millions of euros to build.

In this poster, we will describe the ExAMPLER project, which is seeking input from the empirical agent-based modelling community on how high-performance and exascale computing could be designed to support their needs, as well as how it might affect future directions for agent-based modelling research.

References

1. Alessa, L.N., Laituri, M., Barton, M.: An "all hands" call to the social science community: Establishing a community framework for complexity modeling using agent based models and cyberinfrastructure. Journal of Artificial Societies and Social Simulation **9**(4), 6 (2006), https://www.jasss.org/9/4/6.html
2. An, L., Grimm, V., Turner II, B.L.: Editorial: Meeting grand challenges in agent-based models. Journal of Artificial Societies and Social Simulation **23**(1), 13 (2020). https://doi.org/10.18564/jasss.4012
3. Polhill, G.: Antisocial simulation: using shared high-performance computing clusters to run agent-based models. Review of Artificial Societies and Social Simulation **14 December 2022** (2022), https://rofasss.org/2022/12/14/antisoc-sim/