Development of an HPC-enhanced Code for Agent-based Simulation of Large Economies in 1:1 Scale

Maddegedara Lalith[1], Joshua Panganiban[2], Amit Gill[3], Sebastian Poledna[4], Yoshiki Ogawa[5], Kohei Fujita[1], and Tsuyoshi Ichimura[1]

[1]Earthquake Research Institute, The University of Tokyo, Japan
[2]Dept. of Civil Engineering, The University of Tokyo, Japan
[3]Sumitomo Mitsui Construction Company, Ltd., Tokyo, Japan
[4]International Institute for Applied Systems Analysis, Laxenburg, Austria
[5]Center for Spatial Information Science, The University of Tokyo, Japan

Abstract

We developed a hybrid parallel computing extension that combines distributed and shared memory (`MPI+OpenMP`) for agent-based economic models, allowing for the simulation of large economies on a 1:1 scale—that is, including all economic entities. The model was calibrated to the Japanese economy and validated through simulations covering the period from 2015 to 2018. Large-scale agent-based models are increasingly sought after in various applications, particularly in scenarios where the heterogeneous characteristics of economic entities are significant. While our primary focus is on geographically localized, highly heterogeneous conditions created by natural disasters like large earthquakes, the developed simulator can be adapted to a wide range of scenarios. By partitioning the employer-employee graph, we distribute agents across multiple CPUs. To mitigate the serious communication bottleneck stemming from billions of random interactions in the goods market, we implemented sales outlets on each CPU. The use of cache-friendly data structures and algorithms significantly improved computational efficiency; for instance, a single period of a model encompassing 331 million agents was simulated on 64 Intel Xeon CPU cores in just 153 seconds. Furthermore, we calibrated the model to simulate the Japanese economy using publicly available data from various data portals and institutions of the Japanese government. By initializing the model with data at three resolutions—economy-wide, industry-level, and firm-level—we validated it at all three levels and demonstrated that finer-grained data reduces uncertainty in the simulation results.