Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

TOWARDS PERFORMANCE-AWARE PARTITIONING FOR LARGE-SCALE AGENT-BASED MICROSCOPIC DISTRIBUTED TRAFFIC SIMULATION

Anibal Siguenza-Torres

Intelligent Cloud Technologies Laboratory
Technical University of Munich
Riesstr. 25
Munich, 80992, GERMANY

ABSTRACT

To scale up the performance in large-scale agent-based microscopic traffic simulations, parallel distribution is one of the ways to achieve it. One of the most determinant factors for good performance is the partitioning of the road network. To achieve high performance, the partitions need to have a good load balancing and minimize the communication cost. Many approaches use the number of agents as proxy to estimate the computational and communication costs, assuming a direct relation. However, depending on the simulation logic and the runtime environment, this may not hold true. This work instead proposes to directly measure the communication and computational cost, and to use this information to generate performance-aware partitions. We believe that this would exploit better the system capabilities resulting in higher performance.

1 INTRODUCTION AND MOTIVATION

Parallel distribution is one of the ways to scale up the performance of large-scale agent-based microscopic traffic simulations. Common approaches to distribute the workload into different Logical Processes (LPs) exist. One of these approaches is to partition the road network into sub-regions (one per LP), and each LP processes the agents within its sub-region and communicates the relevant information to the LPs of neighboring sub-regions at the borders (Xu et al. 2014).

The creation of these partitions is non-trivial. To have a high performance, the partitions need to have a good load balancing and minimize the communication cost simultaneously. A common approach to create the partitions is to take the road network graph and encode the computational cost in the node weights and the communication cost in the edge weights. Then it can be interpreted as a multi-constraint graph partitioning problem, and available solvers can be employed to create balanced partitions while trying to minimize the edge cut. The number of agents is commonly used as proxy for both the computational cost and the communication cost (Xu et al. 2014). However these approaches come with some limitations, for example it relies on the assumption that the computational cost is homogeneous for the agents in the simulation, which may not hold in the case of diverse and per-agent models, or if the simulation is distributed in a heterogeneous computing infrastructure (e.g., virtual machines on the cloud with diverse compute power). Similar to the computational cost, the *effective* communication cost is not necessarily proportional to the actual volume of communication. For instance, if the communication is overlapped with the computation, then part of the latency is hidden, and only the non-overlapped communication will contribute to the *effective* communication cost. Additionally, the communication cost itself may depend on the mapping of LPs to the hardware, since different communication channels may be used e.g., communication via shared memory between LPs on the same machine, otherwise via network.

In this work, we aim to develop a partitioning algorithm addressing the challenges outlined above to achieve high performance in the distributed simulation. Instead of using agent counts as proxy, we propose to directly measure the communication and computational cost on the system, and generate a performance-aware partitioning algorithm to better exploit the system capabilities.

2 DISTRIBUTED CITYMOS

This work explores the partitioning problem with distributed CityMoS (Zehe et al. 2017), a large-scale agent-based microscopic traffic simulator. CityMoS employs thread-based parallelism in shared-memory using OpenMP, while distributed CityMoS additionally spreads load among multiple LPs communicating via MPI. Load is distributed by partitioning the road network into sub-regions, and each LP is responsible to compute the agents state within its sub-region. Communication between LPs is required in two cases: agent *migration*, and *remote-sensing*. The *migration* is the process to exchange an agent from one sub-region to another as it crosses the boundary between them. The *remote-sensing* is the process of exchanging information between sub-regions required to compute the local agent state (i.e. to obtain the information of the driving agent in front of an agent). The simulation engine makes use of latency hiding techniques to reduce the impact of communication on overall simulation performance.

3 PRELIMINARY RESULTS AND FUTURE WORK

As a first step, we used a well established static partitioning method in distributed CityMoS, namely the k-way METIS partitioning with weights determined by precomputed traffic information as in (Xu et al. 2014). We used a synthetic benchmark simulation scenario based on the Shenzhen metropolitan area which provides a realistic topology with a mix of road types and networks patterns. Therefore it represents a good case study for load balancing partitioning and scalability studies. We generated traffic to have a single traffic wave of $\sim$4 million itineraries and a maximal load of $\sim$1.1 million concurrent agents simulated throughout 24 hours. The experiment was ran on a single machine with 16 Intel(R) Xeon(R) CPUs E7-8890 v4. The results for the strong scaling are shown in Table 1.

Table 1: Speed up against the single partition execution for the Shenzhen scenario using the k-way graph partitioning. The speedup at the maximum concurrent agent count and the overall speed up are presented.

	4 partitions	8 partitions	12 partitions	16 partitions
Speedup at $\sim$1.1 million agents	2.35	3.99	5.51	6.68
Overall speedup	2.49	4.30	5.58	6.37

By exploring the parameters of the METIS partitioning algorithm, we found partitions that were more balanced in terms of number of agents, but yielded lower performance, which substantiates our hypothesis that a trade-off between communication and computation according to the simulator logic is fundamental to achieve high performance. At the same time, even for the best partitions, profiling results indicate that communication is still a limiting factor for performance. Hence, we believe that a performance-aware partitioning approach explicitly accounting for the computation and effective communication cost can contribute to higher performance. We will therefore engage in a methodical study on a variety of hardware infrastructure using synthetic scenarios in which we can control communication and computation, and therefore work towards a novel partitioning algorithm to be validated using the Shenzhen scenario.

REFERENCES

Xu, Y., W. Cai, H. Aydt, and M. Lees. 2014. “Efficient Graph-Based Dynamic Load-Balancing for Parallel Large-Scale Agent-Based Traffic Simulation”. In *Proceedings of the Winter Simulation Conference 2014*, edited by A. Tolks, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3483–3494. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Zehe, D., S. Nair, A. Knoll, and D. Eckhoff. 2017. “Towards CityMoS: A Coupled City-Scale Mobility Simulation Framework”. In *Proceedings of the 5th GI/ITG KuVS Fachgespräch Inter-Vehicle Communication*, April 6th–7th, Erlangen, Germany, 26–28.