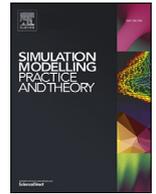


Contents lists available at [ScienceDirect](#)

Simulation Modelling Practice and Theory

journal homepage: www.elsevier.com/locate/simpat

Data-parallel agent-based microscopic road network simulation using graphics processing units

Peter Heywood^{a,*}, Steve Maddock^a, Jordi Casas^b, David Garcia^b,
Mark Brackstone^b, Paul Richmond^a

^a Department of Computer Science, The University of Sheffield, United Kingdom

^b Transport Simulation Systems Ltd (TSS), Spain

ARTICLE INFO

Article history:

Available online xxx

Keywords:

Agent-based simulation
GPU
Simulation framework
Transport microsimulation

ABSTRACT

Road network microsimulation is computationally expensive, and existing state of the art commercial tools use task parallelism and coarse-grained data-parallelism for multi-core processors to achieve improved levels of performance. An alternative is to use Graphics Processing Units (GPUs) and fine-grained data parallelism. This paper describes a GPU accelerated agent based microsimulation model of a road network transport system. The performance for a procedurally generated grid network is evaluated against that of an equivalent multi-core CPU simulation. In order to utilise GPU architectures effectively the paper describes an approach for graph traversal of neighbouring information which is vital to providing high levels of computational performance. The graph traversal approach has been integrated within a GPU agent based simulation framework as a generalised message traversal technique for graph-based communication. Speed-ups of up to $43\times$ are demonstrated with increased performance scaling behaviour. Simulation of over half a million vehicles and nearly two million detectors at a rate of $25\times$ faster than real-time is obtained on a single GPU.

© 2017 The Authors. Published by Elsevier B.V.
This is an open access article under the CC BY license.
(<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Simulations of road networks are used during the development and management of transport networks around the globe. Microscopic road network simulations are fine-grained simulations which simulate individual vehicles within the system, capturing low level behaviours. Agent Based Modelling (ABM) is one microscopic approach, where relatively simple individual behaviours are defined, which combined with interactions between agents and the environment, allows the emergence of complex behaviours. However, microscopic simulations are much more computationally expensive than the more traditional higher-level macroscopic simulations, which use a higher level of abstraction consisting of network flows rather than individual vehicles. Nonetheless, the level of detail captured by microscopic simulations is much greater than that of macroscopic and mesoscopic simulations, including the emergent behaviours enabled by the use of ABM.

* Corresponding author.

E-mail addresses: p.heywood@sheffield.ac.uk (P. Heywood), p.richmond@sheffield.ac.uk (P. Richmond).

<https://doi.org/10.1016/j.simpat.2017.11.002>

1569–190X/© 2017 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license.

(<http://creativecommons.org/licenses/by/4.0/>)

As microscopic simulations are computationally expensive, current commercial and open-source tools such as Aimsun [1], Vissim [2], MATsim [3] and SUMO [4] can have considerable execution run-times, especially for large scale simulations. This limits the overall effectiveness and uptake of microscopic simulation within the transport sector [5]. To increase simulator performance, existing simulation tools use parallel processing applied to multi-core processors to reduce the run-time of simulations. Task-parallel and coarse-grained data-parallel approaches are typically applied within existing tools. Task-parallelism distributes independent processing tasks to separate processing threads. Coarse-grained data-parallelism applies the same algorithms to different units of data, where the individual units of data are relatively large. These approaches are well-suited to multi-core Central Processing Units (CPUs), but can result in poor performance scalability.

Many-core architectures such as Graphics Processing Units (GPUs) offer significantly greater levels of parallelism than multi-core architectures, and significantly greater levels of raw compute performance. To access the high levels of performance, algorithms and data structures must expose high levels of parallelism and enable good memory-access patterns. Typically fine-grained data-parallelism is used, where the same algorithms are applied to relatively small individual units of data.

To demonstrate the simulation performance improvements which could be achieved by using many-core GPUs for microscopic road network simulations, a single-lane transport model with stop-sign-based yellow-box junctions has been implemented using FLAME GPU (Flexible Large Scale Agent Modelling Environment for Graphics Processing Units). FLAME GPU is the only general-purpose GPU-accelerated ABM framework [6]. The performance of the simulation is benchmarked using a procedurally generated artificial grid-based road network, and the simulation performance is compared to a high-performance multi-core CPU microscopic road network simulation tool, Aimsun 8.1, which implements the same models.

This paper presents two contributions: (i) a GPU accelerated data-parallel agent-based road network microsimulation model is evaluated against an equivalent model in a commercial multi-core CPU software tool, demonstrating considerable improvements to simulation performance and performance scalability; and (ii) a general-purpose graph-based communication strategy is presented for high performance agent communication for fine-grained data-parallel agent based simulations, implemented for the FLAME GPU ABM framework which enables high performance agent based simulations of transport networks on GPUs.

Section 2 provides a summary of related work. Section 3 details the implemented model, the benchmark network used and the FLAME GPU implementation. Section 4 describes the graph-based communication strategy to enable the high levels of performance in this model. Section 5 provides the results of a set of application benchmarks used to assess the performance impact of GPUs on microscopic road network simulation using ABM. Section 6 concludes the paper.

2. Related work

Microscopic simulation of transport simulation involves the simulation of individual vehicles and pieces of road network infrastructure to predict the effects of changes in vehicle behaviour or changes in the road network infrastructure. This is used as a tool in the planning and management of transport networks to improve the effectiveness of infrastructure and minimise the impact of changes in conditions on the transportation network. Microscopic simulations are computationally expensive, which has limited the adoption of microsimulation compared to higher level mesoscopic and macroscopic simulations [5]. A micro-scale simulation must include behavioural models for vehicles to follow as well as models for any dynamic infrastructure such as traffic signals and vehicle detectors, which attempt to accurately capture the behaviour of the real world. ABM is a powerful approach to defining microscopic models, which provides a natural method of describing individual behavioural models and then simulating the interaction between individuals in the simulation [7]. Some of the most important vehicle behavioural models are: (i) Car Following Models [8,9]; (ii) Lane Changing Behaviour [10,11]; and (iii) Gap Acceptance Modelling [1,12].

These behavioural models make use of several sets of input data, including: the attributes and structure of the transport network; transport demand which is used to populate the simulation and allow alternate scenarios to be simulated for predictive use; and simulation parameters which are used to modify the models within the simulator. Parameters can include distributions from which vehicle properties can be sampled, such as vehicle length, rate of acceleration or even properties such as pollution emission rates. These parameters can be manipulated to replicate observed behaviour [13,14].

Leading commercial software packages such as Aimsun and Vissim use multi-core CPU architectures to increase simulation performance through task-parallelism and coarse-grained data-parallelism [15,16], reducing the time required for simulations to execute. The work-load of the simulation is distributed as individual tasks or coarse-grained units of data across the available processing hardware, but, as with many task-parallel and coarse-grained data-parallel multi-core software applications, the performance improvement from each additional processing core reduces as the number of cores and threads is increased. Fig. 1 shows the application run-time of Aimsun 8.1 for a simulation containing approximately 25,000 concurrent vehicles as the simulation thread count is increased for two multi-core CPU systems. The diminishing returns of additional processor threads are shown, with no significant increases in performance observed beyond six threads. This limits the performance scalability of the application, with large simulations requiring considerable amounts of time to execute even using CPUs with high numbers of processing threads.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات