Solving Large-Scale Markov Decision Processes on Low-Power Heterogeneous Platforms

Denisa-Andreea Constantinescu*, Angeles Navarro*, Francisco Corbera*,
Juan-Antonio Fernández-Madrigal†, Rafael Asenjo*

* Dept. of Computer Architecture. Universidad de Málaga, Spain.
E-mail: dencon@uma.es, {angeles,corbera,asenjo}@ac.uma.es
† Dept. of Systems Engineering and Automation, Spain.
E-mail: jafma@ctima.uma.es

Abstract

Markov Decision Processes (MDPs) provide a framework for a machine to act autonomously and intelligently in environments where the effects of its actions are not deterministic. MDPs have numerous applications in artificial intelligence, robotics, logistics and maintenance, to name just a few. We focus on practical applications for decision making, such as autonomous driving and service robotics, that have to run on mobile platforms with scarce computing and power resources. In our study, we use Value Iteration to solve MDPs, a core method of the paradigm to find optimal sequences of actions, which is well known for its high computational cost.

In order to solve these computationally complex problems efficiently in platforms with stringent power consumption constraints, high-performance accelerator hardware and parallelised software come to the rescue. We introduce a generalisable approach to implement practical applications for decision making, such as autonomous driving on mobile and embedded low-power heterogeneous SoC platforms that integrate an accelerator (GPU) with a multicore. However, heterogeneous platforms are difficult to exploit efficiently when the application is irregular, as in our case, because bottlenecks due to load imbalance or contention for shared resources between the computing devices (the accelerator and multicore) might hurt performance. Another challenge is for our implementation to be able to provide a solution for real-time use while reducing the power footprint for large scale real-world use cases. The main novelty in this work is solving large-scale MDPs on low-power heterogeneous CPU-GPU platforms and demonstrate the feasibility of our proposal.

To solve this problem, we evaluate three scheduling strategies that enable concurrent execution and efficient use of resources on a variety of SoCs embedding a multicore CPU and integrated GPU, namely Oracle, Dynamic, and LogFit. The Oracle scheduler uses a static partition of the CPU and GPU resources for execution, it has little overhead and is used as baseline. The Dynamic scheduler adjusts the size of the CPU chunk (block of iterations) dynamically to optimise overall CPU-GPU throughput, and uses as input a static chunk size for the GPU provided by the user. In contrast to the first two schedulers that need user input for partitioning the workload, LogFit is specially designed for irregular applications on heterogeneous CPU-GPU chips, and it has an adaptive partitioning strategy that computes the optimal workload (chunk size) at runtime both for CPU and GPU, without user intervention. We compare these strategies for solving an MDP modelling the use-case of autonomous robot navigation in indoor environments on four representative platforms for mobile decision-making applications with a power use ranging from 4 to 65 Watts. We provide a rigorous analysis of the results to better understand their behaviour depending on the MDP size and the computing platform.

Our experimental results show that by using CPU-GPU heterogeneous strategies, the computation time and energy required are considerably reduced with respect to multicore implementation, regardless of the computational platform. We reduce the execution time by up to 64%, in most of the cases and, most importantly, the overall energy consumption, by up to 72%, when solving MDPs by efficiently using the computing resources on the heterogeneous system. We note that Oracle and Dynamic schedulers require offline training while LogFit provides similar performance with no previous training.

Keywords

Decision Making Under Uncertainty, Markov Decision Processes, Value Iteration, Mobile Robot Navigation, Low-power Heterogeneous Computing, Energy Reduction, Irregular Application