Massively Parallel Robot Simulations with the HBP Neurorobotics Platform

Florian Walter, Mahmoud Akl, Fabrice O. Morin, Alois Knoll
Chair of Robotics, Artificial Intelligence and Real-Time Systems, Technical University of Munich

CVPR VIRTUAL JUNE 19-25

Introduction

- The success of deep learning in robotics hinges on the availability of physically accurate virtual training environments and simulation tools that accelerate learning by scaling to many parallel instances.
- Common virtual training environments for deep learning in robotics are typically lightweight and support parallelization but often miss key features of traditional robot simulators that, by contrast, usually offer only limited support for parallelization and machine learning.
- We have implemented a new cloud-based simulation framework that combines the ease of use and scalability of virtual training environments with the full feature set of a robot simulator.

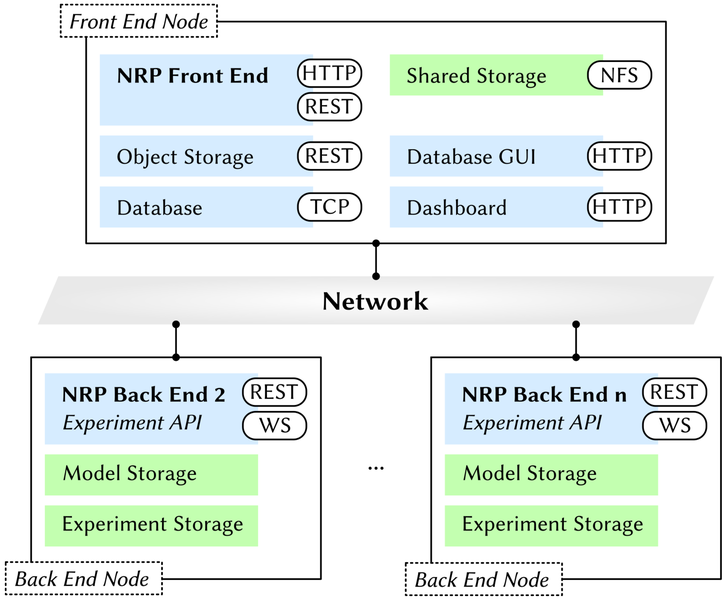

The HBP Neurorobotics Platform (NRP)

- The NRP is a cloud-based simulation framework for neuroscience, neuromorphic engineering, robotics and embodied AI developed in the European Human Brain Project (HBP).
- It is based on the robot simulator Gazebo and adds a *Closed-Loop Engine* that makes it easy to connect robots to neural models from both AI and neuroscience (including support for spiking neural networks and neuromorphic processors).
- Deployment is based on Docker to support flexible scaling from single machines to high performance computing environments.

Massively Parallel NRP Experiments

- By default, the NRP supports launching many simulations in parallel but has so far not allowed any interactions between them.
- To enable the accelerated collection of data sets on many concurrent simulations or distributed training with possibly hundreds of workers, we have added components for shared storage and data analysis to the original system architecture.

A Task Environment for Robot Grasping

- To evaluate our extension, we have created an experiment setup with a KUKA LBR iiwa robot that is placed in front of a tray with objects. An RGB camera behind it captures the workspace.
- Robot control is implemented using common tools from the ROS software stack (ros_control, MoveIt) and the objects in the tray can be spawned randomly to ensure diversity across parallel simulation instances.
- Like in other virtual training environments, users can connect their machine learning models through a new high-level *Experiment API*.

Evaluation

- We have successfully applied the new NRP extension to parallelize the sampling of random robot grasp attempts on an Open Stack cluster with more than 100 instances.
- With all simulations running asynchronously, the speedup can be considered linear.

Outlook

- In the future, we plan to add services for distributed reinforcement learning algorithms and to include support for other environment simulators in addition to Gazebo.

This research has received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3).