

Guest Editorial

Special Issue on Cloud Robotics and Automation

Javier Civera ^{*1}, Matei Ciocarlie ^{†2}, Alper Aydemir ^{‡3}, Kostas Bekris ^{§4} and Sanjay Sarma ^{¶5}

¹Department of Computer Science and Systems Engineering,
Universidad de Zaragoza, Spain

²Department of Mechanical Engineering, Columbia University, NY,
USA

³Center for Autonomous Systems, Royal Institute of Technology,
Stockholm, Sweden

⁴Department of Computer Science, Rutgers University, NJ, USA

⁵Department of Mechanical Engineering, MIT, MA, USA

The Internet and the availability of vast computational resources, ever-growing data and storage capacity have the potential to define a new paradigm for robotics and automation. An intelligent system connected to the Internet can expand its onboard local data, computation and sensors with huge data repositories from similar and very different domains, massive parallel computation from server farms and sensor/actuator streams from other robots and automata. It could even use crowdsourced human supervision or have access to human-made data sources. With the appropriate tools, such resources will be able not only to enhance the performance and robustness of intelligent systems; but also to make them operate beyond traditional limits and in new and more complex domains.

The use of the Cloud has several potential benefits and challenges for robotics and automation. First, computation outsourcing saves energy and physical space, which is highly constrained in mobile robots. And it can reduce the

*jcivera@unizar.es

†matei.ciocarlie@columbia.edu

‡aydemir@kth.se

§kostas.bekris@cs.rutgers.edu

¶sesarma@mit.edu

processing delay, particularly for parallelizable algorithms. A key aspect is to meet the hard real-time constraints of control systems with the variable network delays.

The same energy and space saving applies to data storage. But more importantly, the centralized data gathering opens the doors to the use of learning and AI techniques. The biggest gains of Cloud Robotics and Automation might be found there. The extraction of patterns and knowledge from the data and their transfer to potentially different machines, environments, and tasks could facilitate the programming interfaces and increase the flexibility and capabilities of current robots and automata.

The latest, though, also accounts for the most challenging questions. How to extract relevant patterns from the data? How to exploit unsupervised and weakly supervised data from different scenarios and machines in a similar or even different task? How to formulate abstractions for concepts and tasks from data; and how to ground them again as low-level commands for a possibly different system? How to use the human-made images, video, and text currently in the Internet to enhance the machine capabilities?

It is the potential and also the research challenges of the field what lead us to edit this Special Issue on Cloud Robotics and Automation. Our aims are to group together and to show the state-of-the-art of this newly emerged field, identify the relevant advances and topics, point out the current lines of research and potential applications, and discuss the main research challenges and future work directions. We received 18 submissions in March 2014, spanning a range of topics, and after careful review, 11 were accepted. The work began with the IROS 2013 Workshop on Cloud Robotics and its publication coincides with the CASE 2015 Workshop on Cloud Robotics and Automation.

In the first paper of the issue, Kehoe et al. present a survey that introduces the field, identifies the most relevant work and organizes it around the potential benefits of the Cloud for robotics and automation. Specifically, the main benefits identified by the authors are: 1) Access to Big Data; 2) Cloud Computing; 3) Collective Robot Learning; and 4) Human crowdsourcing. They also describe how Cloud R&A can facilitate open-source software, public databases and benchmarks, and open competitions. The paper also identifies the main challenges and future directions, as for example, privacy and security and the need for new algorithms that are scalable and robust to network latencies.

The next three papers explore the benefits of the Cloud for robotic sensing. The first two exploit Cloud computing to estimate a map of the environment from sensor data. The third one uses a Cloud knowledge-base, built from other robots data, to reason about the environment and to improve exploration and navigation tasks.

Portugal et al. evaluates the performance of two parallel architectures for particle-filter laser SLAM. The computation is distributed among a team of mobile robots, up to five in the experiments. In the first architecture, called stateless, the robots provide computation without sharing the SLAM state. In the second one, called stateful, the robots share their state. From the paper experiments, it can be seen how distributing the computation decreases the

SLAM processing time. Different types of connections are analyzed—cable and several wireless standards—being all of them suitable for the application. The stateful approach is shown to be more efficient due to the need of the stateless one of sending the whole state each time some computation is outsourced.

Similarly, Mohanarajah et al. presents a Cloud-based collaborative visual SLAM that enhances the capabilities of a team of low-cost robots with Cloud server computation. This paper shows how two robots, equipped with smartphone-like processors and RGB-D sensors, can use the remote Amazon servers via a wireless connection to jointly estimate a 3D map of an unknown environment. The experiments show that: 1) parallelizing the heavy map estimation—that includes a pose graph optimization and map merging—on several cluster nodes allows the real-time estimation of the map, something that would be out of reach for the robot local resources and 2) the bandwidth required is low enough for a standard wireless connection.

Riazuelo et al. presents a mapping approach that uses the robot sensing and the information in the Cloud knowledge-base RoboEarth—built from other robot’s data—to improve on semantic mapping and object search. For the semantic mapping task, the knowledge-base infers the set of landmark objects to be found and the robot builds a map containing them. For the object search task, the knowledge-base provides a partial map of the room built by other robots that is used by the current robot for navigation. The knowledge-base also infers likely object locations from previous data for active object search, producing time savings with respect to exhaustive search.

The next three papers show how not only robotic sensing but also robot actuation can benefit from the Cloud. Specifically, the study cases are collision-free navigation, grasping, and water quality monitoring.

Salmeron-Garcia et al. addresses the use of the Cloud for robotic navigation from stereo vision. This paper analyzes four different computation distributions; namely, all processing performed locally, all processing performed in the Cloud and the two intermediate solutions; local stereo vision plus Cloud navigation and Cloud stereo vision plus local navigation. The experimental results show that the most critical aspect is the computation-communication tradeoff, due to the big size of a stereo stream. In any case, a carefully chosen computation offloading can outperform an unconnected robot.

Kehoe et al. presents an algorithm for grasp planning in the presence of shape uncertainty that leverages the parallel computing power of the Cloud. The authors’ define a set of candidate grasps and perturbations of the uncertain 2D geometry. Each candidate grasp and perturbation is evaluated and a lower bound on the probability of force closure is estimated. This sampling-based algorithm is particularly well-suited for Cloud computing; the experiments demonstrate up to 515x speed-up with 500 computation nodes. The authors also propose a procedure for the early removal of low-quality grasps to reduce the grasp evaluations up to 91.5

Pandey et al. presents a Cloud-based framework for the collaboration of a team of robots for a water quality monitoring application in an underwater environment, where the resources are highly constrained. Their framework uses

the resources in the in-network composed by the team of robots and boats, and also external servers. The authors develop a resource-provisioning algorithm that assigns the computation tasks to the components of the net. The experimental results demonstrate that the collaboration of an internal and external computation network has benefits with respect to the purely local and purely Cloud computation.

Mohanarajah et al. present Rapyuta, an open source Platform-as-a-Service (PaaS) framework developed by the RoboEarth project team. Rapyuta lets robots perform computation remotely by providing secure customizable computing in the cloud and access to the RoboEarth knowledge repository. The system architecture is designed to facilitate multiprocess high-bandwidth robotics applications and provides a well documented open source implementation.

Agüero et al. presents a framework to facilitate cloud-hosted robot simulations for the Virtual Robotics Challenge, an open competition aimed to advance the state-of-the-art in robotics for disaster response. The parallel computation capabilities of the Cloud are essential here in order to simulate accurately and in real-time the physics of complex scenarios for a possibly high number of participants in different parts of the globe. The framework is based on the robot simulator Gazebo and the software CloudSim to host the simulations in the Cloud. This paper reports the performance of the system and the lessons learned during the challenge.

The distribution of the computation, sensing and communication load is one of the most important features of a Cloud system. Wang et al. presents a framework for the case of realtime multisensor data retrieval on an architecture composed of a data center, cloud cluster hosts, and robot clients. The scheduling is formulated as a Stackelberg game, where the utility of the robot clients and the revenue of the leader are maximized. The authors demonstrate that their proposal improves over the state-of-the-art in terms of time of response, bandwidth usage, and CPU load.

Finally, Stenmark et al. identifies as key demands for future manufacturing systems their flexibility and the simplification of the programming and usage. Online knowledge bases are essential for both demands, in which the high-level robot capabilities are modeled, shared and reused and the low-level parameters are learnt from data. This paper summarizes the authors' work and lessons learned on four different projects on the topic—SIARAS, ROSETTA, PRACE, and SMERobotics.

The Guest Editors would like to thank all the authors and the anonymous reviewers for their outstanding work. We would also like to thank Markus Waibel for the initial proposal of the issue, the Editor-in-Chief Ken Goldberg for his expert guidance and advice, the Editorial Assistant and Journal Administrator Samantha Jacobs for her excellent work and patient support, and the IEEE T-ASE Associate Editors M. Ani Hsieh, Mariagrazia Dotoli, Weiming Shen, and Bengt Lennartson for their valuable help with papers that fell into their areas of expertise.