The Use of Cloud Hosted Simulation to Support the Design of Habitats for Space

Sekou L. Remy, Ph.D.
Institute for African American Mentoring
Clemson University
Clemson SC, USA
sekouremy@gmail.com

Abstract - *NASA already has a very visible presence in the field of cloud computing. Projects like "Be A Martian", and the Jet Propulsion Lab's use of the Amazon Cloud, not just share the "7 minutes of Terror", but also to share much of the other data from the Curiosity Rover project, are just some of the ways that NASA is leveraging cloud computing technologies to successfully deliver engaging experiences of Science, Technology, Education and Math (STEM) to the public. The proposed concept performs science on the process of using the cloud to deliver robotic Simulation Platforms as a Service. These services leverage virtualization of computing resources over the network to provide access to simulated robots. Such an offering provides interesting opportunities for education and training since it can deliver on-demand access computing with staggering resource portfolios and provide new pathways to develop user interfaces even when equipment may not be available.*

Keywords: Robotic simulation.

1 Introduction

In 2012, a team of volunteers traveled to a workshop in Colorado to work on a telerobotics project. Most of the members of the team were new to the project but their collective backgrounds were all relevant. When work was started at the project site, the team faced a set of challenges that are unfortunately very common in robotics. These challenges included software/operating system compatibility, system integration, hardware failures, and more broadly a need for major components of the project to be properly functioning before meaningful progress could be made on the project.

The nature of these challenges also appear in research, class projects, competitions and even for hobby/DIY applications. In this work we consider the use of remotely hosted simulation resources. Such simulation resources provide a meaningful complementary pathway that permits progress to be made even when the hardware is not complete.

2 Background

Robots are well suited to the 3Ds "dull, dirty and dangerous", and in space exploration these three conditions abound. As such, the design of future human space exploration missions depends on robotic assets. Understanding how astronauts will interact with their robotic team members is critical as it influences both the design of the needed vehicles/habitats and also the process of allocating the required crew to perform missions.

Of specific focus in this work is the question of which teleoperation modes are most effective in supporting astronaut interaction. Earlier work [3] proposed the study of four teleoperation modes (EVA Direct Observation, Support Vehicle Direct Observation, Third-person View, and an Onboard Camera View) and in this work we seek to determine whether simulated robots can be used to assess of these modes.

The use of remotely hosted computing resources in this way is not new to the research community, NASA, or to space exploration. For nearly two decades there has been active research and development in this vein. Specifically, control of robots over the World Wide Web has been prominent during this time period, with especially early high profile efforts such as the Mercury Project [1] and the Web Interface for Telescience [2], which was initially evaluated during the Mars Pathfinder mission.

To implement the simulation environment, the Robot Operating System (ROS) [4] is used as the core infrastructure for message/services and abstractions of other useful robotic resources. ROS is, in large part, an example of a Service Oriented Architecture (SOA) [5], however it deviates in some key ways to provide functionality natural in robotics. The Open Source Robotics Foundation (OSRF) supports the development of ROS, and the foundation also supports the development of another complimentary software called Gazebo [6]. Gazebo is a 3D multi-robot simulator that provides "physically plausible" simulations. With ROS, Gazebo can be used to provide a compelling simulation of a rover or a lander. At present however, Gazebo is not well suited to simulate a large number of robots, yet there is another open source environment, Stage [7], which has successfully been used to simulate hundreds of mobile robots. Both Gazebo and Stage fill needed niches in these NASA relevant simulation missions. In this work, as a precursor to ROS+Gazebo, we implement scenarios with ROS+Stage. The software is run on a machine hosted in an Infrastructure as a Service cloud framework, and as such is accessed in an on-demand manner.

3 Infrastructure

Multiple Stage world files were created to capture the teleoperation modes of interest. EVA Direct Observation, Support Vehicle Direct Observation, and the third-person View teleoperation modes were each implemented by creating a simulation with two robots, one to provide the needed external view of the environment. These three cases differed only in the resolution and the fields of view provided. The Onboard Camera View

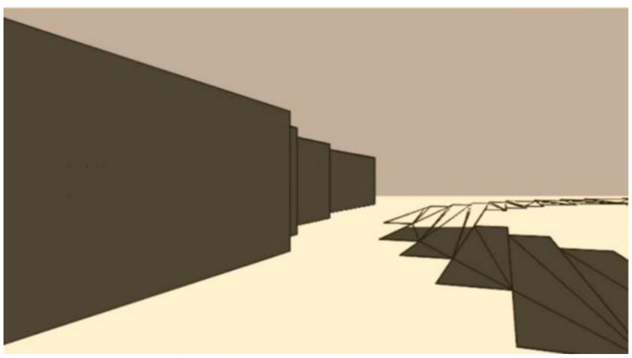

Onboard Camera View of environment with varying types of obstacles

3.1 Why Simulate Robots in the Cloud?

Simulating robots in the cloud has the potential to deliver the specific benefit of increased impact. This impact is associated with the promise of standard ways to develop, deploy, manage and share simulations and simulation infrastructure. It is also associated with the prospect of cost savings that result from power consumption, cooling, system (hardware) maintenance, space, and (physical) security.

Cloud computing provides another set of opportunities, ones that are really engaging in education and training. Since cloud computing is elastic and flexible, the resources for a class can be modified dynamically to adjust to need, within a few minutes. The resources can also be provisioned and decommissioned quickly making the start-up and shut-down time associated with the use of the robotic simulation much less daunting.

The multi tenant nature of cloud computing also makes it quite easy for computing in the margins, that is it provides the avenue for enterprises to donate spare capacity in a structured manner without disrupting the normal enterprise operation. In truth, this is the reason that many of cloud vendors provide their offerings, as they aim to more efficiently utilize their infrastructure. The same cloud infrastructure can be use to deploy and manage simulations in different geographic zones. This could be leveraged to provide standard support more readily, or even to centralize the delivery of localized resources. (e.g. the simulation could be located in Virginia for east coast users, or in California for west coast users). At the core, the cloud provides solutions that can readily be scaled.

This is all contributing to the impact that can be realized through its use. Cloud computing is scalable and flexible, but as a set of technologies, it also provides the opportunity to become a gateway to learning because it is extensible. For example, in the process of using a simulation of a robot in the cloud, it is possible to do many different types of "deep dives" depending on the level of expertise of the audience. Educators could use the simulation opportunity to discuss networking, programming, applications of STEM, etc.

Finally simulation in the cloud may not be as much as a desire, but a requirement as many point to the future of computation being in the "cloud", and researchers and academics may not have much choice in the future in whether they can have access to bare metal hardware. Corporate policy may dictate that this is the way that computing is done, so we must be prepared for this possibility.

3.2 Why Consider Not Simulating in the Cloud?

In addition to the challenges that relate in general to standards, there are (surprise, surprise) several that are specific to this application that should be considered. Many of these challenges exist for the use of technology of any kind and include: effectiveness of the tool, reliability of the resources, safety (or security) of using the resource, (cost is viewed as a favorable one, especially if the resource is only used part of the time), integration with existing (legacy) infrastructure and workflow. Typically listed as a challenge, but not the case with simulation in the cloud would be the challenge of how well things scale (how expandable is the solution).

Further, because of the intrinsically distributed nature of cloud based simulation, the network matters greatly, and at best you can get some type of service level agreement (SLA) that promises some performance. These agreements are based on technology-based predictions, and can only predict on average, not predict what will happen in any specific case.

In “the cloud” you do not truly have a means of measuring the true impact of neighboring simulations, as in other tenants in the resource, or clients simultaneously operating in multiple zones. For example is there a competitive advantage afforded to a client in one zone over another? Also, it is not (yet) clear whether simulation in the cloud is equivalent to remote simulation over the network and the reasons for such a concern include the use of virtualized hardware that are significant to the operation of the simulation (GPU, CPU, networks, etc.).

Cloud computing is also likely to enjoy wider deployment of software defined networking and as has been previously stated, the network matters greatly, so this can be a been a boon or a threat depending on configuration. Finally, time and timing variability in the cloud are concepts that should give developers interested in robotic simulation pause as time (and space) are critical to embodied simulation of robots. In this domain, it is not just input and output, but it is input (sensor data) and output (actuation data) from a particulate place and time, and these are used to determine what happens next. Computer simulations are afforded the luxury of approximating time (simulation time), so things can be simulated faster or slower than real time, but in all cases, confidence in the value of time, or the elapsed time is critical to the performance of simulations that are repeatable, verifiably, and realistic. As a specific example, many in the controls world expect control loops to be able to be controlled in 1ms for meaningful control to be performed (and thus for meaningful simulations to be performed). The timing resolutions in many of the cloud computing offerings cannot maintain such a request, and vendors recommend that time sensitive operations are not performed in the cloud.

4 Conclusions

In addition to experimental testing with human subjects, there are many technical questions that are still to be resolved. For example, what are the constraints on the number and kind of robots that can be used in a given experiment, and on a given type of computing infrastructure? What are the timing challenges that are associated with the use of the cloud for relevant teleoperation tasks? What are the impacts of network latency for teleoperation over the Internet, and also for teleoperation with heads-up displays as an alternative to the well-vetted keypad/joystick/mouse user interface devices?

Acknowledgment

This work would not have been possible without the support of Mr Kylan Race and the Institute for African American Mentoring in Computer Science (iAAMCS).

References

[1] K. Y. Goldberg, M. Mascha, S. Gentner, N. Rothenberg, C. Sutter, and J. Wiegley, “Destop teleoperation via the world wide web,” in ICRA, 1995, pp. 654–659.

[2] P. Backes, G. K. Tharp, and K. S. Tso, “The web interface for telescience (wits),” in in Proc. IEEE Int. Conf. Robot. Automat, 1997, pp. 411–417.

[3] R. Howard Jr., “Robotic Teleoperations Experiment for the Mars Desert Research Station,” in NSBE Aerospace Systems Conference, 2012.

[4] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “ROS: an open source Robot Operating System,” in ICRA Workshop on Open Source Software, 2009.

[5] M. Huhns and M. Singh, “Service-oriented computing: key concepts and principles,” Internet Computing, IEEE, vol. 9, no. 1, pp. 75 – 81, jan-feb 2005.

[6] N. Koenig and A. Howard, “Design and use paradigms for gazebo, an open-source multi-robot simulator,” in IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004, pp. 2149–2154.

[7] R. Vaughan, “Massively multi-robot simulation in stage,” Swarm Intelligence, vol. 2, no. 2, pp. 189–208, Dec. 2008.[Online].Available: http://dx.doi.org/10.1007/s11721-008-0014-4