

Introduction to Research Methodology - Homework

Tommi Nylander

22 June 2016

1 Introduction

In this short essay, I will begin by briefly describing my research field and continue with some reflections on the research methods that I currently use and perhaps will use in the future.

Less than six months ago, I began my PhD studies at the Department of Automatic Control at Lund University. As a fairly new PhD student, I am yet to complete my first paper and have only just begun with my first research project. The field of my research is cloud control. This is an area that covers control theoretical approaches for improving today's cloud infrastructure. Examples of problems that can be formulated within this area include how to improve the resilience of the cloud as well as how to optimize the user-perceived response times. The cloud infrastructure of today is almost entirely made up of servers residing in far-away, large-scale data centers. The goal with my future research is to investigate the opportunities and challenges that arise when the cloud becomes more distributed, that is, when more nodes in the cloud (e.g. base stations) participate in the computations.

2 My current research methods

The current research project that I am participating in together with my supervisors is on auto-scaling of brownout-aware replicas. The brownout concept was introduced to the cloud computing field by the Cloud Control research group in Umeå and Lund [3], and is inspired by brownouts in power grids, i.e. where the power is reduced to avoid complete blackouts. This idea is utilized in the cloud as allowing applications to only respond with some optional content if they are not under heavy load. My current project will expand this idea to also handle when the load changes for longer periods of time, by automatically adding and removing replicas that run the application.

The research methods that we have utilized thus far in this project can be

related to some parts of the scientific method. The initial research question was stated as "How can new auto-scaling algorithms be developed for replicas with brownout?". We then formulated a hypothesis we believe that the auto-scaler has to act on the proportion of optional content served by the replicas, rather than just their response times. Until now almost all of the time spent has been on gathering data to test and evaluate this hypothesis. This has been done by extending an existing Python-based simulator, and then by running tests in order to develop a model of the system. At the time being, this is how far we have come in this project. Hopefully this work will lead to us being able to draw some conclusions of our data and to develop some new theories of how auto-scaling of brownout-aware replicas should be performed.

The motivation behind this research approach lies in the fact that this project extends previous work where the same methods have been utilized. A clear example on how following a similar procedure simplifies the research, is that the simulator previously developed by our predecessors can be re-used with some slight modifications.

3 Possible future research methods

An interesting research method that I would like to try out is to perform case studies on existing cloud management systems that could benefit from control theoretical auto-scaling. Interesting examples include the cluster management systems Mesos [1] and Kubernetes [2], which both have been released during the last couple of years. Since these systems still have not matured, their solutions for handling load changes over long time intervals are just on a proof of concept level. As they are both open source software, I believe that it would definitely be possible to implement my future auto-scaler and evaluate it using case studies.

Another interesting approach for my future research could be to, on a higher extent than now, examine the actual need of certain cloud solutions in industry. This could be done by performing field studies or surveys, and the results of these investigations could then be helpful for defining and motivating the upcoming research projects. Since my research is part of the WASP program, with tight connections to industry, these studies could probably be launched without much external effort. The hard part would rather be to formulate the questions in these studies in a suitable way, such that meaningful and relevant answers are gained.

As described previously, my current research project is based on a concept borrowed from power systems. It is therefore not impossible that we will continue to look into existing methods in other fields, when trying to come up with new solutions that improve the cloud.

References

- [1] Hindman Benjamin, Konwinski Andy, Zaharia Matei, Ghodsi Ali, Joseph Anthony, Shenker Scott, and Stoica Ion. Mesos: a platform for fine-grained resource sharing in the data center. In *In Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, 2011.
- [2] Brendan Burns, Brian Grant, David Oppenheimer, Eric Brewer, and John Wilkes. Borg, omega, and kubernetes. *Commun. ACM*, 59(5):50–57, April 2016.
- [3] Cristian Klein, Martina Maggio, Karl-Erik Årzén, and Francisco Hernández-Rodríguez. Brownout: Building more robust cloud applications. In *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pages 700–711, New York, NY, USA, 2014. ACM.