The Journal of Systems and Software xxx (2012) xxx–xxx

Contents lists available at SciVerse ScienceDirect

The Journal of Systems and Software

journal homepage: www.elsevier.com/locate/jss

# A goal-oriented simulation approach for obtaining good private cloud-based system architectures

Lawrence Chung[a], Tom Hill[a], Owolabi Legunsen[a,*], Zhenzhou Sun[a], Adip Dsouza[a], Sam Supakkul[b]

[a] The University of Texas at Dallas, 800 West Campbell Road, Richardson, TX 75080, USA
[b] Sabre Inc., USA

ARTICLE INFO

*Article history:*
Received 7 April 2012
Received in revised form
16 September 2012
Accepted 15 October 2012
Available online xxx

*Keywords:*
Cloud computing
System architecture
Goal-oriented
NFR framework
Simulation model
CloudSim
Little's Law
Requirements engineering

ABSTRACT

The fast-growing Cloud Computing paradigm makes it possible to use unprecedented amounts of computing resources at lower costs, among other benefits such as fast provisioning and reliability. In designing a good architecture – the numbers, types and layouts of devices – for a cloud-based system, which meets the goals of all stakeholders, such goals need to be factored in from the earliest stages. However, there seems to be a lack of methodologies for incorporating stakeholder goals into the design process for such systems, and for assuring with higher confidence that the designs are likely to be good enough for the stated goals. In this paper, we propose a goal-oriented simulation approach for cloud-based system design whereby stakeholder goals are captured, together with such domain characteristics as workflows, and used in creating a simulation model as a proxy for the cloud-based system architecture. Simulations are then run, in an interleaving manner, against various configurations of the model as a way of rationally exploring, evaluating and selecting among incrementally better architectural alternatives. We illustrate important aspects of this approach for the private cloud deployment model and report on our experiments, using a smartcard-based public transportation system.

Published by Elsevier Inc.

## 1. Introduction

Would it be possible to predict whether a cloud-based service will indeed be good enough from the perspectives of multiple stakeholder goals, such as profitability, performance and scalability? Can we do this in the early stages of development, before committing to potentially costly and time-consuming implementation and testing? Answering these questions affirmatively is important to people who want to take advantage of the many benefits, such as cost reduction, fast service provision, reliability, and the like, promised by Cloud Computing (Armbrust et al., 2009; Mell and Grance, 2009). In particular, designers of proposed cloud-based systems will want to investigate how to come up with designs that meet the goals of all stakeholders before it is too late, disruptive or expensive to make changes. This paper proposes one approach for doing this by using goal-orientation together with cloud computing simulations that are cheap and quick to set up, since goal-oriented techniques allow for capturing stakeholder goals and using them in exploring, analyzing and selecting among architectural design alternatives. In the proposed approach, codenamed Silverlining,[1] goal orientation and simulation are used in an interleaving, iterative manner – goal-oriented aspects can be carried out simultaneously, and in any order with, simulation modeling and analysis.

By nature, early stage design in cloud-based systems is multi-stakeholder, multi-objective, multi-dimensional and large scale. It is "multi-stakeholder" in the sense that there are different types of stakeholders (e.g., cloud vendors, service providers, and end users) and "multi-objective" in that the goals of one group of stakeholders differ from those of another group and goals are oftentimes competing and conflicting, even within the same group. The "multi-dimensional" nature concerns the fact that stakeholder goals need to be simultaneously investigated for different stakeholder groups at different levels of abstraction (hardware level, Virtual Machine (VM) level, database level, etc.). It is also "large scale" because very large numbers of individuals in each stakeholder group need to be considered while designing such systems, necessitating the use of different kinds of workloads in assessing the quality of system design. These characteristics make cloud-based system design quite challenging. Without a systematic method, such as the one proposed in this paper, it would be a daunting challenge to understand, develop, and successfully operate cloud based systems. This

* Corresponding author.
*E-mail addresses:* chung@utdallas.edu (L. Chung), tom.hill.fellow@gmail.com (T. Hill), owolabi.legunsen@utdallas.edu (O. Legunsen), zxs101020@utdallas.edu (Z. Sun), amd061000@utdallas.edu (A. Dsouza), sam.supakkul@sabre.com (S. Supakkul).

1 http://utdallas.edu/~axg118830/silverlining.html.

0164-1212/$ – see front matter. Published by Elsevier Inc.
http://dx.doi.org/10.1016/j.jss.2012.10.028

goal-oriented, simulation-based approach provides a fast way with little financial and manpower resources to tackle the challenge.

Users' requirements for cloud computing architecture, as well as the role of goal-oriented requirements engineering in Cloud Computing adoption have been described (Clarke, 2010; Zardari and Bahsoon, 2011). However, there still needs to be a more systematic approach for integrating these into the architecture design process for cloud-based systems, toward reaching the level of rigor and success attained by goal-oriented techniques in traditional (not cloud-based) Software Engineering. Perhaps, the lack of such a rational approach in cloud-based system design is the main culprit in what has been described as "solutions that are looking for a problem" in the cloud (Lehman and Vajpayee, 2011)? Our main aim in this paper is to describe a systematic approach which may be used to design cloud-based systems based on stakeholder needs. We borrow heavily from best-of-class Goal-Oriented Software Engineering (van Lamsweerde, 2001) techniques, while showing how such might be used in a realistic system design scenario. We also focus on the so-called private cloud deployment model, in which one entity (the Cloud Service Creator) owns and operates the datacenter but gives shared access to other entities who subscribe on a pay-as-you-go basis.

The use of simulations for investigating complex system behavior is not new. However, in the Software Engineering community, there has been a recent increase in the recognition of the role that simulations can play in the early stages of system design, especially for evaluating and validating design decisions. This is perhaps due to growing realization among researchers that rising levels of systems complexity need to be matched by more reliable ways of investigating designs in the early stages. The marriage of goal-orientation (techniques for exploration, selection among, and evaluation of architectural alternatives based on stakeholders' goals) and simulation (which is fast and easy to set up) is expected to be highly beneficial but challenging. Some work has recently emerged in this regard, for using simulation to confirm and reconfirm architectural decisions (Hill et al., 2009), for runtime monitoring and system maintenance (Hill, 2011; Hill et al., 2010) and for simulating and optimizing design decisions in quantitative goal models (Heaven and Letier, 2011). The Cloud Computing field has also been quick to recognize the utility of simulations in investigating infrastructural level concerns (Kertész et al., 2011; Jeyarani et al., 2010). Our use of simulation is for investigating the impact of Cloud-Computing infrastructural design choices on the proposed system's ability to meet the needs of its stakeholders.

At a high level of description, our approach starts with the capture of stakeholder goals plus some domain characteristics such as workflows. The goals are then refined and extended quantitatively with numbers obtained from the domain, using estimation methods that we propose. All these are subsequently used to create a simulation model as a proxy for the cloud-based system architecture. Lastly, simulations are run iteratively, in an interleaving manner, against various configurations of the model as a way of rationally exploring, evaluating and selecting among incrementally better architectural alternatives that have a better chance of meeting the stakeholder goals.

The main contributions of this paper are as follows:

1. A method of combining goal-oriented requirements engineering techniques with simulation for cloud computing.
2. A technique for complementing the qualitative goal models in the NFR Framework (Chung et al., 2000) with a quantitative approach.
3. The development of an interactive, iterative and interleaving simulation technique based on this quantitative approach as well as stakeholder needs.
4. A new visual notation, along with heuristics and guidelines for reasoning about the goal models in the presence of the quantitative additions.

In the rest of this paper, Section 2 provides a background to our approach, including an overview of the real world system used for illustration and some key Goal-Oriented Requirements Engineering concepts to be used in the rest of the paper. Section 3 gives step-by-step details of our proposed approach as applied to the system under study. We discuss our experimental outcomes in Section 4

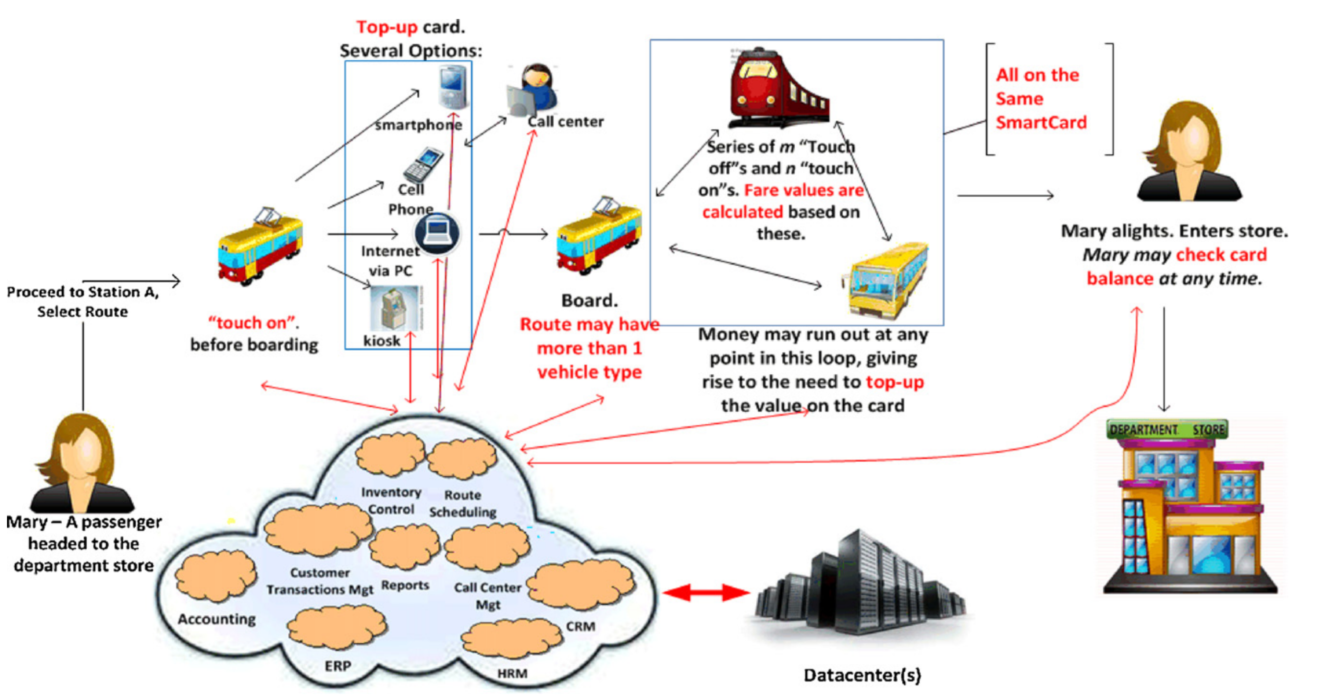


**Fig. 1.** In the "myki", very many transactions must be processed in real time. How do we come up with a cloud-based design to satisfy this and other design constraints for all stakeholders?

Please cite this article in press as: Chung, L., et al., A goal-oriented simulation approach for obtaining good private cloud-based system architectures. J. Syst. Software (2012), http://dx.doi.org/10.1016/j.jss.2012.10.028