

Simple Pricing Schemes for the Cloud

Extended Abstract*

Ian A. Kash
Microsoft Research
iankash@microsoft.com

Peter Key
Microsoft Research
Peter.Key@microsoft.com

Warut Suksompong
Stanford University
warut@cs.stanford.edu

ABSTRACT

The problem of pricing the cloud has attracted much recent attention due to the widespread use of cloud computing and cloud services. From a theoretical perspective, several mechanisms that provide strong efficiency or fairness guarantees and desirable incentive properties have been designed. However, these mechanisms often rely on a rigid model, with several parameters needing to be precisely known in order for the guarantees to hold. In this paper, we consider a stochastic model and show that it is possible to obtain good welfare and revenue guarantees with simple mechanisms that do not make use of the information on some of these parameters.

CCS CONCEPTS

•Networks → Cloud computing; •Theory of computation → Algorithmic mechanism design; Computational pricing and auctions;

KEYWORDS

cloud computing, pricing, mechanism design

ACM Reference format:

Ian A. Kash, Peter Key, and Warut Suksompong. 2017. Simple Pricing Schemes for the Cloud. In *Proceedings of NetEcon'17, Cambridge, MA, USA, June 27, 2017*, 1 pages.
DOI: 10.1145/3106723.3106727

1 INTRODUCTION

With cloud computing generating billions of dollars per year and forming a significant portion of the revenue of large software companies, the problem of how to price cloud resources and services is of great importance. On the one hand, for a pricing scheme to be used, it is necessary that the scheme provide strong welfare and revenue guarantees. On the other hand, it is also often desirable that the scheme be simple. We combine the two objectives in this paper and show that simple pricing schemes perform almost as well as more complicated ones with respect to welfare and revenue guarantees.

For our initial results we assume that there is a single server, which receives jobs of various lengths whose value per time step is drawn from the same probability distribution regardless of length.

*The full version is available at <http://arxiv.org/pdf/1705.08563.pdf>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

NetEcon'17, Cambridge, MA, USA

© 2017 Copyright held by the owner/author(s). 978-1-4503-5089-1/17/06...\$15.00
DOI: 10.1145/3106723.3106727

We compare the welfare and revenue that can be obtained by setting a price per time step that is independent of the job length against the corresponding objective obtained by setting an individual price for each job length. When we are allowed the freedom of setting different prices for different job lengths, intuitively we want to set a higher price per time step for longer jobs as a premium for reserving the server for a longer period of time. However, as we show, we do not lose more than 50% of the welfare or revenue if we are only allowed to set one price. We would like to emphasize that this is a worst-case bound over a wide range of parameters, including the number of job lengths, the distribution over job lengths, and the distribution over job values. Indeed, as we show, we can obtain improved bounds if we know the value of some of these parameters. The price that we use in the single-price setting can be chosen from one of the prices used in the multi-price setting, meaning that we do not have to calculate a price from scratch. Moreover, all of our approximation guarantees hold generally for arbitrary prices, meaning that for any prices that we may set in a multi-price setting (i.e., not necessarily optimal ones), we can obtain an approximation of the welfare or revenue by setting one of those prices alone. Finally, we emphasize that these results put no restrictions on the form of the distribution; it can be discrete, continuous, or mixed. The only substantive constraint is that jobs of all lengths share the same distribution of value per time step. However, in an extension we show that a version of our results continues to hold even if this constraint is relaxed.

We then generalize our results to a setting where there are multiple servers, each of which receives jobs of various lengths. The distribution over job lengths can be different for different servers. This is conceivable, for instance, if the servers are in various geographic locations or are utilized by various groups of users. We compare the welfare and revenue obtained by a simple pricing scheme that sets the same price for all servers against the corresponding objective achieved by a scheme that can set a different (single) price for each server. Roughly speaking, we show that as long as the parameters are not too extreme, e.g., the number of servers or the job lengths are not too large, then we do not lose too much of the welfare or revenue by setting a single price. Combining this with our initial results, we obtain an approximation of a very restricted pricing scheme where we must set the same price for all servers and all job lengths against one where we can set an individual price for each job length of each server. These results require an assumption that all servers have the same probability of not receiving a job at a time step. Using similar techniques, we also obtain approximation bounds when this assumption does not hold but there is only one job length across all servers.