Profitable cloud broker with two phases optimisation

Chalee Boonprasop
School of Computer Science
University of St Andrews
Fife KY16 9SX
cb330@st-andrews.ac.uk

Abstract

Cloud brokers act as an intermediary between cloud service providers (e.g., Amazon and Google)and cloud service end users. To operate at a profitable level, many cloud brokers employ a buy and sell strategy. User requests are placed into a pool of reserved instances, which the broker has prepurchased. By buying low and reselling the cloud instances, the broker can potentially make a profit from the price difference between reserved instances and on-demand instances. However, the spikey and unpredictable nature of user requests (both initialisation and termination) poses a challenge for maximising the usage of any prepurchased reserved instances, which sit idle if not fully utilised.

In this talk, I will explain how to design a cloud broker which utilises partial data optimisation in two phases, in order to optimise utilisation and reduce risk when prepurchasing reserved instances. The first phase involves, a moving window clustering algorithm, which is applied to adjust the available reserved instances based on the incoming demand. The second phase manages the overall risk of the broker using an anomaly detection algorithm. Our initial results show that the overall profit of the broker is increased when compared to a non-optimised strategy. In addition, the broker has less reliance on historical user data. The broker is still able to make a profit regardless of data distribution.

Keywords: Cloud broker, Unsupervised learning, Optimisation