

# Learning-based Cost Management for Cloud Databases

Olga Papaemmanouil  
Brandeis University

2/12/2018 at 10:30 am  
CoRE 301

## Abstract

Cloud computing has become one of the most active areas of computer science research, in large part because it allows computing to behave like a general utility that is always available on demand. While existing cloud infrastructures and services reduce significantly the application development time, significant effort is still required by cloud data management applications to manage their monetary cost, for often this cost depends on a number of decisions including but not limited to performance goals, resource provisioning and workload allocation. These tasks depend on the application-specific workload characteristics and performance objectives and today their implementation burden is left on application developers. We argue for a substantial shift away from human-crafted solutions and towards leveraging machine learning algorithms to address the above challenges. These algorithms can be trained on application-specific properties and customized performance goals to automatically learn how to provision resources as well as schedule the execution of incoming query workloads with low cost. Towards this vision, we have developed WiSeDB, a learning-based cost management service for cloud-deployed data management applications. In this talk, I will discuss how WiSeDB leverages (a) supervised learning to automatically learn cost-effective models for guiding query placement, scheduling, and resource provisioning decisions for batch processing, and (b) reinforcement learning to offer low cost online processing solutions, while being adaptive to resource availability and decoupled from notoriously inaccurate performance prediction models.

## Bio

Olga Papaemmanouil is an Associate Professor in the Department of Computer Science at Brandeis University since January 2009. Her research interest lies in the area of data management with a recent focus on cloud databases, data exploration, query optimization and query performance prediction. She is the recipient of an NSF Career Award (2013), an ACM SIGMOD Best

Demonstration Award (2005) and a Paris Kanellakis Fellowship from Brown University (2002). She received a undergraduate degree in Computer Engineering and Informatics from the University of Patras, Greece, in 1999, a Sc.M. in Information Systems from the University of Economics and Business, in Athens, Greece, in 2002, and a Ph.D in Computer Science from Brown University, in 2008.

Faculty Host: Shan Muthukrishnan