# Automated Modeling of Complex Data Center Applications

[Keynote Talk]

Prashant Shenoy School of Computer Science University of Massachusetts Amherst MA 01003 shenoy@cs.umass.edu

## ABSTRACT

As online services become increasingly common, the complexity of backend distributed server applications in data centers has also continued to grow. At the same time, there is an increasing need to enhance the manageability of these large applications by automating common management tasks, which requires a good understanding of the run-time behavior of the application under different scenarios. However, the rising complexity of these applications makes the tasks of manually modeling and analyzing their run-time behavior increasingly difficult. In this talk, I will argue for the need to automate the modeling of the run-time performance of distributed data center applications. I will present techniques for automatically deriving application models using statistical methods from machine learning. I will describe how we have put these ideas into practice into two systems that we have built, Modellus and Predico, and will present case studies of using these systems for management tasks such as capacity planning and what-if analysis of data center applications.

### **Categories and Subject Descriptors**

D.4.8 [**Operating Systems**]: Performance—Modeling and prediction, measurements, queuing theory

#### **General Terms**

Performance

#### Keywords

Performance modeling, data centers, learning

Copyright is held by the author/owner(s). *ICPE'13*, April 21–24, 2013, Prague, Czech Republic. ACM 978-1-4503-1636-1/13/04.