

Recent Trends in Performance Modeling of Big Data Systems

Varsha Apte
Computer Science and Engineering Department
Indian Institute of Technology - Bombay
Powai, Mumbai 400 076, India

ABSTRACT

With the advent of big data through social media and continuous creation of digital footprints through various mobile devices, special-purpose programming models were developed that would make it easy to write programs to process such data. MapReduce and its Hadoop implementation is one of the most popular platforms for writing such programs. The MapReduce framework involves a "map" phase where various tasks work in parallel for intermediate processing of data and a "reduce" phase where again various tasks work in parallel to extract information from this processed data. Performance modeling of such systems will need different approaches than are used for traditional multi-threaded multi-core systems supporting Web applications, primarily because the dependencies and synchronization required between various tasks is not easily expressible using standard queuing network models. In this talk we will review work done by researchers to address this modeling problem. The work done encompasses first-principles calculations of execution time completion, queuing network models, and finally, simulation. We will review these efforts as well as highlight opportunities for further work in this area.

BIO

Varsha Apte is Professor in the Department of Computer Science and Engineering, at IIT Bombay. She has been at IIT Bombay since 2002. Prior to that she was in the Performance Analysis department of AT&T Labs in New Jersey, US. She finished her PhD in Computer Science from Duke University in 1994 and prior to that has an M.Sc in Computer Science from Pune University. Her core interest and expertise is in solving performance modeling and management problems of multi-tier applications using Queuing Systems and other Stochastic Models. She has built tools that facilitate and simplify the process of performance measurement and modeling such applications. She is still interested in the same fundamental problem, but now in the context of virtualization and cloud computing.

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

ICPE'17, April 22-26, 2017, L'Aquila, Italy

© 2017 ACM. ISBN 978-1-4503-4404-3/17/04.

DOI: <http://dx.doi.org/10.1145/3030207.3030222>