COMPUTER SCIENCE
AND
SOFTWARE ENGINEERING

AUBURN UNIVERSITY

SAMUEL GINN COLLEGE
OF ENGINEERING

Data Center Specific Thermal and Energy Saving Techniques

Xiao Qin
Department of Computer Science and Software Engineering
Auburn University, Auburn, AL 36849-5347
xqin@auburn.edu
http://www.eng.auburn.edu/~xqin

Abstract: Data centers are ever increasing as we become more reliant of web based transactions. The benefits of such massive computing are obvious by the speed and ease we can get most media or information. A challenge is that new large data centers introduce a level of energy consumption that the world has not seen before. The obvious energy cost of running the computers is a billion dollar problem, but there are hidden costs like running cooling systems as well. To help combat the problems of large data centers, we aim at developing solutions that can work for each type of data center. This could entail creating tools that are generic enough to work for all data centers, or focusing on specific tools the type of software running in the data center. In this talk, we present a thermal model that is flexible enough to be applicable for all data centers; we show how our model can be used to save energy. We also discuss new energy saving techniques for Hadoop clusters specifically, where we focus on very data centric implementations of Hadoop to gain a significant energy savings. Finally, we propose a Spark specific process that takes what we have learned from our Hadoop and thermal research and try to develop techniques that offer large energy and thermal savings within Spark clusters.

Bio: Xiao Qin is an Associate Professor in the Department of Computer Science and Software Engineering at Auburn University. He received the B.S. and M.S. degrees in Computer Science from Huazhong University of Science and Technology, China, in 1996 and 1999, respectively. He received the Ph.D. in Computer Science from the University of Nebraska-Lincoln in 2004. Prior to joining Auburn University in 2007, he had been an assistant professor with New Mexico Institute of Mining and Technology (New Mexico Tech) for three years. He won an NSF CAREER award in 2009 to conduct research on multicore-based parallel disk systems. His research interests include parallel and distributed systems, real-time computing, storage systems, fault tolerance, and performance evaluation. His research is supported by the U.S. National Science Foundation, Auburn University, and Intel Corporation.