



# Improving Big Data Analytics with Software Defined Networking

**Theophilus Benson**

**Duke University**

**[tbenson@cs.duke.edu](mailto:tbenson@cs.duke.edu)**

Although there is tremendous interest in designing improved networks for big data analytics data centers, very little is known about the network-level traffic characteristics of current data centers. In this talk, I will present an empirical study of the network traffic in 10 data centers, describing the flow-level and packet-level transmission properties of the applications within these data centers, and their impact on network utilization, link utilization, congestion, and packet drops. Then, I will briefly discuss the implications of the observed traffic patterns for data center internal traffic engineering as well as for recently-proposed architectures for data center networks.

Next, I will shift to proactively controlling data center traffic patterns by proposing techniques that improve performance within data centers by leveraging SDN. I will present an overview of MicroTE, a system that adapts to traffic variations by using SDN to exploit the short-term predictability of data center traffic. I will show that MicroTE performs close to the optimal solution. Finally, I will conclude by presenting initial results from ongoing work, NaPS, that uses SDN to provide applications, such as Hadoop, with visibility into the network and thus allowing them to make more efficient decisions in a network aware manner.

## Biography

Theophilus Benson is passionate about eliminating the complexity of managing networks and tackling performance oriented problems both within data centers and clouds. His research focuses on Software Defined Networking, infrastructures for big data, and abstractions for managing various workloads in the cloud. This work has earned him IBM fellowships, a best paper award at IMC 2010, and, more recently, his cloud computing platform was acquired by a large cloud provider. He is an Assistant Professor at Duke University. He received his PhD from the University of Wisconsin – Madison in 2012 after which he spent a year at Princeton University as a Post-Doc with Jennifer Rexford. Prior to that, he received his B.S. at Tufts University and worked as a software engineer at an MIT based startup in Waltham, MA.