


# Computer Science Seminar


Tuesday, February 20<sup>th</sup>, 2018  
1:00 – 2:00 pm  
Atwood Chemistry Building Room 240



Dr. Nosayba El-Sayed

## Reading Between the Lines of Datacenter Logs

Designing datacenters that are reliable, energy-efficient, and capable of delivering high performance and high utilization is a nontrivial problem facing scientists, businesses, and governments alike. In this talk, I will demonstrate how analyzing large datasets from different organizations helped us uncover interesting (and often surprising) patterns in the behavior of systems and applications in these large-scale platforms. I will show how real-world data helped us tackle critical questions such as how does temperature impact server reliability in places like Google, or how well do users configure the computing jobs they submit to shared clusters (spoiler alert: not very well!). Finally, I will demonstrate how simple machine learning techniques can be leveraged to accurately predict job failures in datacenters, while using data that is easily collected in current platforms.



**Dr. Nosayba El-Sayed** is currently a Post-doc Research Associate at the Massachusetts Institute of Technology (MIT) and the Qatar Computing Research Institute (QCRI). She obtained her Ph.D. in Computer Science from the University of Toronto in 2016. Her main research area is High Performance Computing (HPC), with focus on improving the energy efficiency and reliability of large scale systems using data analysis and simulation techniques.