

# Identifying trending stories on Twitter and optimal temperature for data center computers

July 24 2012

---

Papers on how best to identify trending stories on Twitter and on just how cool computers in data centers need to be kept has won accolades for faculty and students from UTSC's Department of Computer & Mathematical Sciences.

Nick Koudas, UTSC computer science professor, and his graduate students Albert Angel and Nikos Sarkas, won the Best Paper award for the Very Large Databases Conference, which they will present at the meeting in Istanbul in August.

The paper presents a new way of identifying stories of interest on Twitter and other streaming microblogging sites. It looks at ways to identify stories of interest by looking at "entities" that occur together in the same tweet.

"When these events happen, there are certain entities – events, people, locations – that are used together in a tweet," Koudas says.

For instance, tweets about the US raid that killed Osama bin Laden would likely contain "entities" including bin Laden, the CIA, Barack Obama, Abbottabad, and possibly others. By keeping track of these entities and how they are grouped together, the method Koudas and his students invented can determine which tweets are likely to be about a particular story.

What's more, it can do so in an efficient way, making it useful in real-time. That's important when you consider that [Twitter](#) adds 4,000 tweets a second to its stream, Koudas says.

In another paper, Bianca Schroeder, assistant professor of computer science, and students Nosayba El-Sayed, Ioan Stefanovici, George Amvrosiadis and Andy A. Hwang were awarded best paper at the ACM Sigmetrics conference in London in June.

Their paper examines the issue of temperature management in data centres, and suggests that allowing warmer temperatures than are normally recommended might be justifiable.

As data centres have proliferated they have required more energy, accounting now for about 1 percent of global electricity useage. A sizeable fraction of that is the cooling necessary to keep the machinery functioning properly.

But Schroeder's paper found that warmer temperatures than are normally recommended might be able to save energy without negatively impacting equipment reliability and longevity.

Data centres typically operate at temperatures from 20C to 22C. Estimates show that just 1 degree increase in temperature could save 2 to 5 percent of energy consumed.

By collecting data from a large number of centres, and also running tests in which they measured the effect of temperature on performance, the UTSC researchers suggested higher temperatures might be okay. Their data showed that higher temperatures either weren't associated with negative effects on the equipment, or else the negative effects were smaller than predicted.

"We see our results as strong evidence that most organizations could run their data centers hotter than they currently are without making significant sacrifices in system reliability," the authors say.

Provided by University of Toronto Scarborough

Citation: Identifying trending stories on Twitter and optimal temperature for data center computers (2012, July 24) retrieved 1 July 2024 from <https://phys.org/news/2012-07-trending-stories-twitter-optimal-temperature.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.