

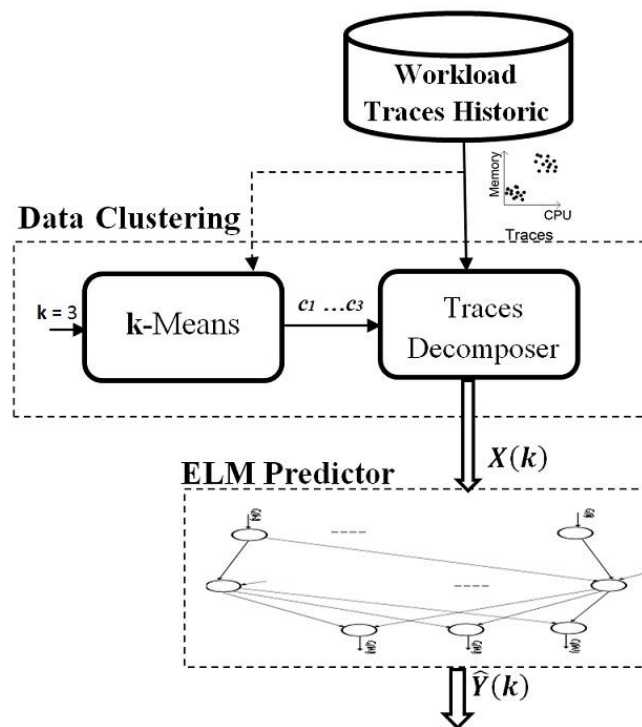
The role of workload prediction in reducing datacentre energy costs

It's not a stretch to say information technology services have become near-essential in North America today. On average we now connect to more than 22 applications daily and generate two GB of data while we're at it, says Salam Ismaeel, a PhD candidate at Ryerson in the [Department of Computer Science](#). "More than 77 percent of today's businesses," he adds, "rely on the Internet and data center resources."

But all that expansion poses a huge challenge. IT services require the pervasiveness of the Internet and thousands of computing resources in data centers, and the ability to manage costs – energy costs in particular – is crucial.

"While data-center computing platforms are now largely software driven, the primary resource remains power," says Ismaeel. "Datacenter power is among the largest commodity expenditures in IT services for most organizations." With 88 percent of this power going to powering and cooling the IT equipment, any reduction has a significant impact on the bottom line, he explains.

And along with such heavy reliance on electricity comes extreme vulnerability to market shifts. In Ontario for example, over the last seven years the electrical power rate more than doubled, according to Ontario Hydro figures, and by 2018 it's projected to increase by a further 33 percent.



Under the supervision of [Dr. Ali Miri](#), professor and associate chair of the computer science department, Ismaeel and fellow researchers have been conducting experiments in workload prediction. Why put the focus there? As Ismaeel

explains, if you can predict workload and switch servers to lower power states when not in use, that's an important way to reduce energy costs, not to mention conserve energy. The main focus of his team's research has been on predicting future Virtual Machine (VM) requests based on past usage. "This has been done by combining k-Means clustering and Extreme Learning Machines (ELM)," he says.

Most notably, Ryerson researchers found their model was associated with improved performance. As Ismaeel writes: "We have tested our model on real Google traces that feature over 25 million tasks collected over a 29-day time period. Experimental results presented show that our proposed system outperforms other models reported in the literature."

Ismaeel and Dr. Miri summarized their findings in a paper called ["Using ELM Techniques to Predict Data Centre VM Requests"](#) and presented them last November, in New York at the [2nd IEEE International Conference on Cyber Security and Cloud Computing](#).

Of his studies at Ryerson, Ismaeel says, "Working with my supervisor's group in the same lab and regular group meetings every week to discuss different research topics motivated me to overcome the challenges in my work." He also notes group meetings plus the opportunities to network with other researchers in the field as an invaluable part of his graduate work.

Ismaeel's research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC).