

An Indian-Australian research partnership



Industry sponsored project by Infosys

Project title Green Data-Centers: Optimizing Data-Center Layouts and Intra-Cloud Migration of Virtual Machines.

Project number: IMURA0177

Monash University supervisors: Professor Mohan Krishnamoorthy, Prof David Abramson

Monash University contact: Professor and Associate Dean of Research (Engineering) Email:

Mohan.Krishnamoorthy@adm.monash.edu.au

IITB supervisors: Prof Krithi Ramamritham, Prof Pushpak Bhattacharyya

IITB contact: Chair Professor Department of Computer Science and Engineering.; Email: krithi@iitb.ac.in,
pb@cse.iitb.ac.in

Infosys supervisors: Dr. Sumit Kumar Bose, Mr. Dheepak R. A

Research Academy theme/s

List only the research academy theme/s that is relevant to the project

1. **Advanced computational engineering, simulation and manufacture**
2. Infrastructure engineering
3. Clean energy
4. Water
5. Nanotechnology
6. Biotechnology and stem cell research

The research problem

There is a growing interest in sustainable design and management of IT systems and infrastructure. Data centers of today house potentially hundreds and thousands of systems with widely varying power and cooling requirements. It is estimated that one data-center housing a few hundred thousand racks requires a dedicated power station to cater to its power and cooling requirements. In essence, data-centers of today are huge behemoths gurgling enormous amounts of electrical power for generating huge amounts of computing cycles and has phenomenal cooling requirements. A typical 1U server can draw up-to 500W of power. When 40 of them are packed into a standard 42U rack the power consumption can be as high as 20KW and produces nearly 70,000 BTU. Further dissipating this heat requires almost 6 tons of cooling per rack. Add to this, the advent of new age computing paradigms – virtualization and cloud computing for example, and adoption of multi-core and multi-processor servers, the amount of power used per square foot is only going to double than what it was 5 years ago. It is therefore important to optimize the power and cooling requirements by having appropriate floor plan layouts, designing appropriate cabling and strategizing on the placement of air-ducts. This way virtualization technologies and policy based management systems though brings in autonomic and self-healing capabilities within the IT infrastructure, also adds to the existing ‘power’ and ‘cooling’ challenges within the data-center. A plain and unintelligent way of managing virtual machines can lead to an increase in power and cooling requirements of the ‘virtualized data center’. The challenge then is to improve the energy efficiency of the data-centers – both existing and new, even as we embrace new age computing paradigms

like cloud computing, virtualization and multi-core machines. Solution to the problem involves deciphering the complex relationship amongst multiple dimensions (e.g. rack space, temperature, humidity) to the problem. The advent of cloud technologies only adds one more dimension and hence complexity to the problem. The aim of this research work is to explore sensor technologies, multi-agent systems and operations research methods to make the data-center operation more 'Green' even as we adopt cloud technologies.

Project aims

The research work will explore advanced analytic techniques (such as support vector machines etc) to establish relationship between the different variables that affect the power and cooling requirements within a data-center. This in essence denotes the temperature profile of the data-center. This relationship will then act as the basis for the subsequent design of an efficient data-center. Further the analytical equation derived from the earlier step would also act as the basis for migrating VMs across a data-center to minimize the power and cooling requirements arising out of SLA breaches. The work will explore establishing an analytical relationship as a first step. This requires an in-depth knowledge of fluid dynamics and SVM. Subsequent to this, the work will explore the use of this relationship for design of the data-center layout optimization model and also as a basis for VM migration within a DC.