



Arizona State University
Tempe, AZ 85287
480-965-3806
<http://impact.asu.edu/>

“Do Cool”: Energy-Efficient and Autonomic Resource Management Technology for Data Centers

*By Sandeep Gupta, Professor
Georgios Varsamopoulos, Research Faculty*

September 2008

Data centers are increasingly using a larger portion of the total national energy supply. Currently, Data centers in US collectively use about 1.5 percent of nation's energy supply. Governments at all levels are providing incentives to companies for reducing their power consumption by using green technologies. Despite the increasing need for energy savings, in a recent survey of Uptime Institute¹, about 46% of the data center operators don't use power-save features and 38% don't turn off unused or under-used servers. The main reason was unfamiliarity with the impacts of energy-savings techniques on performance.

Moreover, conventional green techniques cannot yield a lot of savings, because they are *reactive* in nature and do not address the phenomenon of *heat recirculation*, that is, heat from the computing equipment that doesn't enter the cooling intakes but rather re-circulates back to the room (Figure 1), which is the most dominant factor of energy waste in a data center. Research at the IMPACT Laboratory has allowed for the development of green technology, "Do Cool", to enable energy-efficient management of data centers.

Do Cool Technology benefits

The Do Cool Technology is designed to address energy-inefficiencies in data centers, with focus on heat recirculation. Research results show that the *proactive thermal-aware scheduling schemes* of Do Cool Technology can yield more than 20% savings (Figure 2). Specifically, the research has allowed exploring the limits of such technique and also reducing

¹ Kenneth G. Brill, "Special Report: Energy Efficiency Strategies Survey Results." White Paper. Uptime Institute. April 2008

the computational complexity of the management software. The energy savings of a scheme that performs thermal and power management can exceed 30% and 60% for a typical and the worse case, respectively.

Thus, for a data center of 1MW power (i.e. about 400-500 servers), the savings can be about 400kW on average, which translates into 3.5MWh \approx \$315,000 (at 9¢/kWh) in a year. The savings can be greater, depending on the physical layout and workload of the data center. The reduction in carbon dioxide footprint is proportional to the energy savings. Moreover, our algorithms make the data centers avoid thermally critical situations, thus greatly reducing the risk of downtime or throttled-down performance.

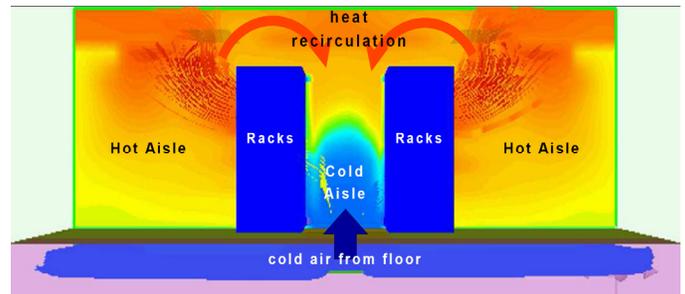


Figure 1: the heat recirculation effect.

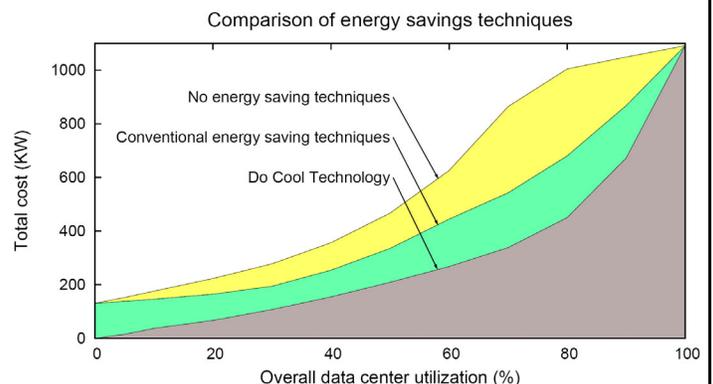


Figure 2: contribution of various energy-saving techniques

Do Cool Technology insights

Our technology integrates management of jobs, cooling and power to achieve synergistic effects and benefits. It is a **runtime** solution, that is, it does not require any changes in the physical layout, although it can point out its inefficiencies. It is also an **add-on** to the installed scheduling software of a data center (e.g. Moab or SLURM), thus it requires minimal changes to the installed software. Lastly, it is **autonomic**, which means that it requires minimal configuration and intervention from administrators.

The Do Cool technology relies on deriving a **thermal model** of the data center and using it to perform thermal-aware assignment of incoming jobs to the servers in the data center. This is achieved by using a **distributed software architecture** (Figure 3), which is used to collect sensor readings and statistics as well as to disseminate control decisions.

The architecture provides the Do Cool technology with the power consumption profiles or specifications of the equipment and the utilization profile of jobs to

translate this thermal model into a **cost model**. The cost model is then used to predict the cost associated with particular job placement. The thermal-aware scheduler then assigns jobs to the servers with the minimal cost.

In addition to the traditional scheduling output (i.e. order and start times), the thermal-aware extensions enhance the scheduler to produce:

- an enhanced energy-efficient job schedule, i.e. start times of the jobs, as well as server assignment;
- a cooling schedule, i.e. a time-based schedule of *CRAC* (*Computer Room Air Conditioner*) set points;
- a power schedule, i.e. a schedule of turning off and turning on machines.

Although energy savings can be achieved even with item (a) alone, the synergistic effects of applying all three schedules together multiply the savings.

Do Cool technical requirements

In order to enable its autonomic and

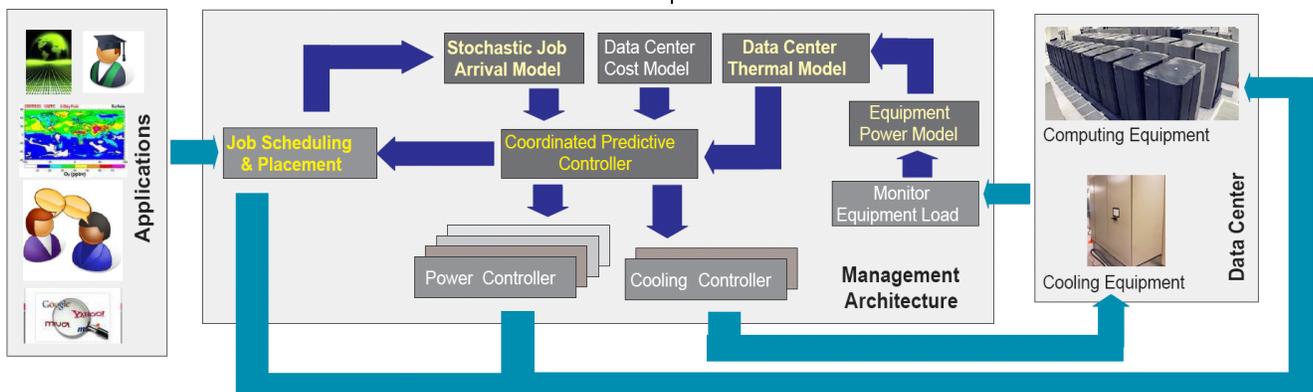


Figure 3: Software architecture of the Do Cool Technology.



energy-saving capabilities, Do Cool technology requires from the data center to provide:

- 1) inlet fan and outlet fan temperatures for each machine (blade or chassis)
- 2) CPU/core utilization for each machine
- 3) CRAC air inlet and outlet temperatures
- 4) all fan speeds (CRAC and computing units), either in RPM or air flow (cubic feet per minute)
- 5) power specs of each computing machine type
- 6) power specs of each CRAC unit
- 7) Submission/execution logs of i) job name/id, ii) arrival time, iii) start time, iv) deadline, v) placement (node assignment), and vi) finish time.

The availability of these data and sensor readings is common in modern (2005 or newer) computing equipment.

The IMPACT Lab (<http://impact.asu.edu>) has a proven record in developing thermal-aware management technologies with several on-going projects supported by funds from Science Foundation of Arizona (SFAz), Consortium of Embedded Systems (CES), National Science Foundation (NSF) and Intel Corp. For more information, visit <http://impact.asu.edu/~mcs/DataCenter.htm> or contact Prof. Sandeep K. S. Gupta, email: Sandeep.gupta@asu.edu phone: 480-759-4468.