# Center for Energy-Smart Electronic Systems (ES2)

State University of New York at Binghamton, Bahgat Sammakia, 607.777.6880, bahgat@binghamton.edu

State University of New York at Binghamton, Kanad Ghose, 607.777.4803, ghose@cs.binghamton.edu

Villanova University, Alfonso Ortega, 610.519.7440, alfonso.ortega@villanova.edu

University of Texas at Arlington, Dereje Agonafer, 817.272.7377, agonafer@uta.edu

Georgia Tech, Yogendra Joshi, 404.385.2810, yogendra.joshi@me.gatech.edu

Vishwakarma Institute of Technology (VIT), Siddharth Jabade +91.20.24202.198 siddharth.jabade@gmail.com

State University of New York at Binghamton, Andrea Palmeri, 607.777.3629, apalmeri@binghamton.edu

## Putting Data Centers on a Diet: Dynamic, Load-Dependent Rightsizing of Server Capacity

The past decade has witnessed unprecedented growth in the number, size and capacity of data centers that support our daily lives domestically, in commerce, and in governance. Data centers represent the backbone of cloud computing. This is the core infrastructure that supports social networking, provide web services that we access daily from computing devices at all scales, from cell phones to desktops. Data center deployments are growing worldwide at unprecedented rates and scales. U.S. data centers consume about 100 billion kilowatt-hours of electricity annually. Data centers use electricity not just to run the IT equipment but also to power cooling systems that take away the heat dissipated by the servers. They stress out our already taxed power generation units. Unfortunately, current data centers are not very energy-efficient. Significant amounts of energy are wasted in operating data centers.

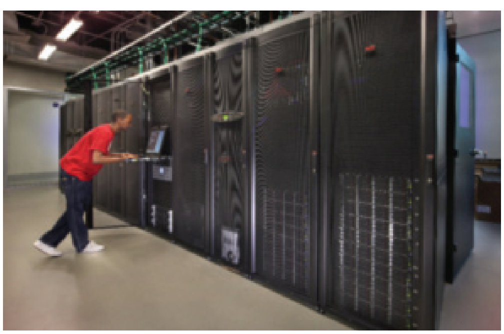

Data centers are usually designed to target the peak demands but such demands rarely occur. However, data center operators tend to keep all servers online, so as not to miss any requests should the request volume go up suddenly. This practice results in idling servers or servers running at low utilization. These practices lead to low energy efficiency because idling servers dissipate a significant amount of power. An industry-wide study done by the Gartner Group indicates that, on the average, server utilizations are at less than 15% of their maximum capacity.

In the data center industry, a great deal of effort has been spent to reduce inefficiencies in cooling systems. As a result, the energy spent on cooling in today's newer, relatively well-designed data centers is less than 30% of the total power drawn by the data centers. The bulk of the remaining energy is spent to support the IT equipment. Significant energy savings are possible by better managing server capacities in ways that track demand as it fluctuates from one instant to the next.

This breakthrough ES2 technology automatically provides just the right amount of server capacity needed at any time to handle current offered loads and activates additional servers when the demand grows. This saves power by shutting down servers when demand drops. The improvements in energy efficiency comes from two main practices. First, active servers are operated at high utilization levels; this improves the overall energy efficiency of the IT equipment in the data centers. Second, unused servers are shut off, avoiding any power wastage from idling servers.

The challenge in doing such automatic server capacity provisioning has to do with the time it takes to activate turned-off servers when the load grows. Since it takes a few minutes to turn on a server, a reactive solution that reacts to increased request volume by turning on servers will not work, because increases occur suddenly and can exceed the capacity of the currently turned-on servers well before the servers being turned on are ready to accept requests. The ES2's breakthrough technology uses a proactive server activation/deactivation strategy that uses the recent history of the actual and offered load to predict the expected load to turn additional servers on in advance, avoiding any service degradation. Additional features permit the degree of cooling provided to dynamically match the capacity of the servers, thus avoiding wastage due to overcooling or damage due to undercooling. Demonstrations of a prototype implementation shows that the new technology permits over 25% reductions in data center IT equipment power draws in realistic scenarios of operation with almost negligible impact on performance.

**Economic impact:** Automatic and dynamic server capacity provisioning in data centers, as enabled by this breakthrough technology, permit data center operating expenditures for the IT equipment to be reduced by as much as 25% without appreciable impacts on performance. Additional cascaded energy savings are possible in cooling system and in power distribution and conversion networks. The technology is particularly well suited to data centers that provide online services such as social networking, email, news, shopping, content searching etc. where the demands can fluctuate rapidly in volume. The resulting drop in electricity draw also lowers the dependence on fossil fuels and the carbon footprint.

For more information, contact Kanad Ghose, ghose@cs.binghamton.edu, Bio www.cs.binghamton.edu/~ghose/, 607.777.4608.

## More Efficient Data Centers: Maximizing Airside Cooling

Around the world, unprecedented and exceedingly large volumes digital data are stored in data centers. The data come from online transactions, social networking websites, banks, healthcare facilities, schools, government bodies, and industry, to name a few. It is essential that these data are reliably processed by and securely transmitted to a variety of user/customers. Within the last twenty years, the amount of digital data being generated has greatly increased resulting in ever increasing numbers of data centers.

Data centers house many servers with typically multiple thousands of IT equipment units. The centers require a considerable amount of cooling infrastructure to enable IT equipment to function properly. One way to reduce energy consumption of data centers is to use air-side economization (ASE) and indirect and direct evaporative cooling (I/DEC). These cooling methods do not use compressors, which is good because compressors consume large amount of energy to convert vaporized cooling fluid to liquid cooling fluid. Although not all data centers can be cooled 100% of the time using ASE and I/DEC methods, maximizing the use of ASE and I/DEC (minimizing compressor-based cooling system) can result in significant reductions in energy consumption of the cooling infrastructure. The annual number of hours ASE and I/DEC will be used in a year is expected to increase since IT equipment manufacturers have begun to make IT equipment that can be operated at wider temperature and humidity ranges than was possible few years ago.

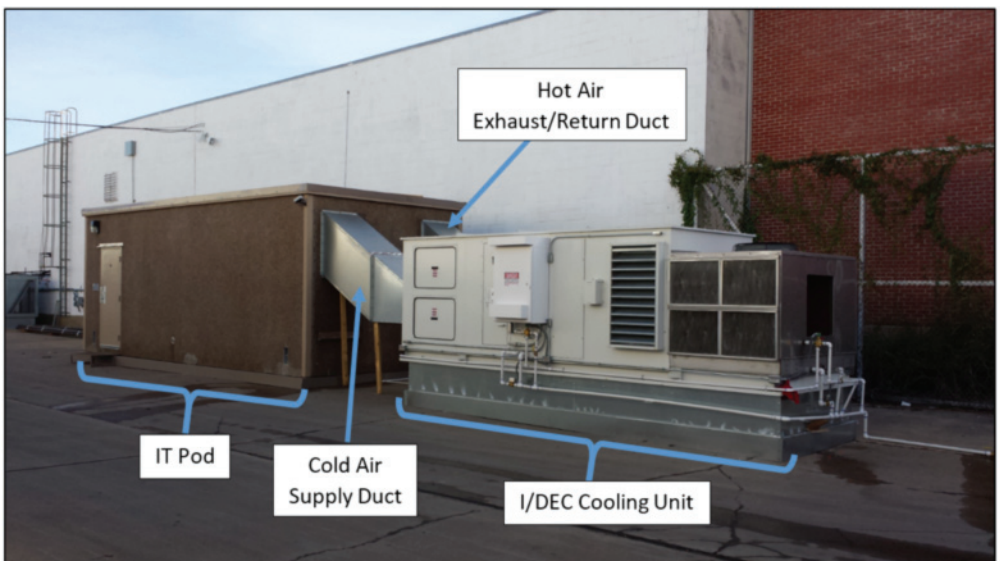


*Modular data center*

Research on ASE and I/DEC for data center cooling has been conducted on a testbed modular data center. Although these methods of cooling are not new to data centers, maximizing their use over compressor based cooling technologies has the advantage of reducing overall cooling costs and carbon footprints. In the specific I/DEC unit that has been studied, a cooling tower is used to cool water that runs through a water-to-air heat exchanger (indirect evaporative cooling). For this system, the cooling tower can be placed at a distance from the main cooling unit. This allows greater ease of maintenance and provides freedom to place the unit at a suitable place in the data center.

> **Economic impact**: Electricity consumption by data centers is expected to increase from 91 billion kWh in 2013 to 140 billion kWh by 2020. Reliable operation in data centers of IT equipment requires that the heat generated needs to be continuously and effectively removed; otherwise, the IT equipment may fail to operate resulting in data center downtime. A recent study indicated that about 30% of total data center energy is consumed by the cooling infrastructure of data centers. Reducing energy consumption of the cooling infrastructure of data centers by few percentage points will translate to many thousands of dollars of savings per year in savings for data center owners and significant reductions in carbon footprints of the data centers.

This breakthrough cooling approach avoids the use of compressors or chillers. The approach s 70% less energy than traditional air conditioning systems. Depending on location and environmental conditions, water usage can also be significantly reduced in these systems. This is because water is recirculated and used several times. Continuous R&D testing has helped replace conventional propeller type fans with VFD to a more efficient EC axial fan in the cooling tower. Data centers that implement these cooling technologies will benefit from the electric cost saving. Overall, the cost of purchasing, installing and servicing an EC fan/motor assembly can save significant dollars as opposed to conventional fan/motor selections tied to a VFD.

The nation will benefit from more sustainable growth of energy efficient data centers that maintain smaller carbon footprints that conventional cooling systems.

For more information, contact Dereje Agonafer at the University of Texas, Arlington,
agonafer@exchange.uta.edu, Bio emnspc.uta.edu/Director.html, 817-272-7377.