

Reducing Datacenter Energy Consumption Using Coordinated ICT-cooling Control Technology of Datacenter Energy Management System

*Masayuki Nakamura, Akira Takeuchi, and
Hiroyuki Enomoto*

Abstract

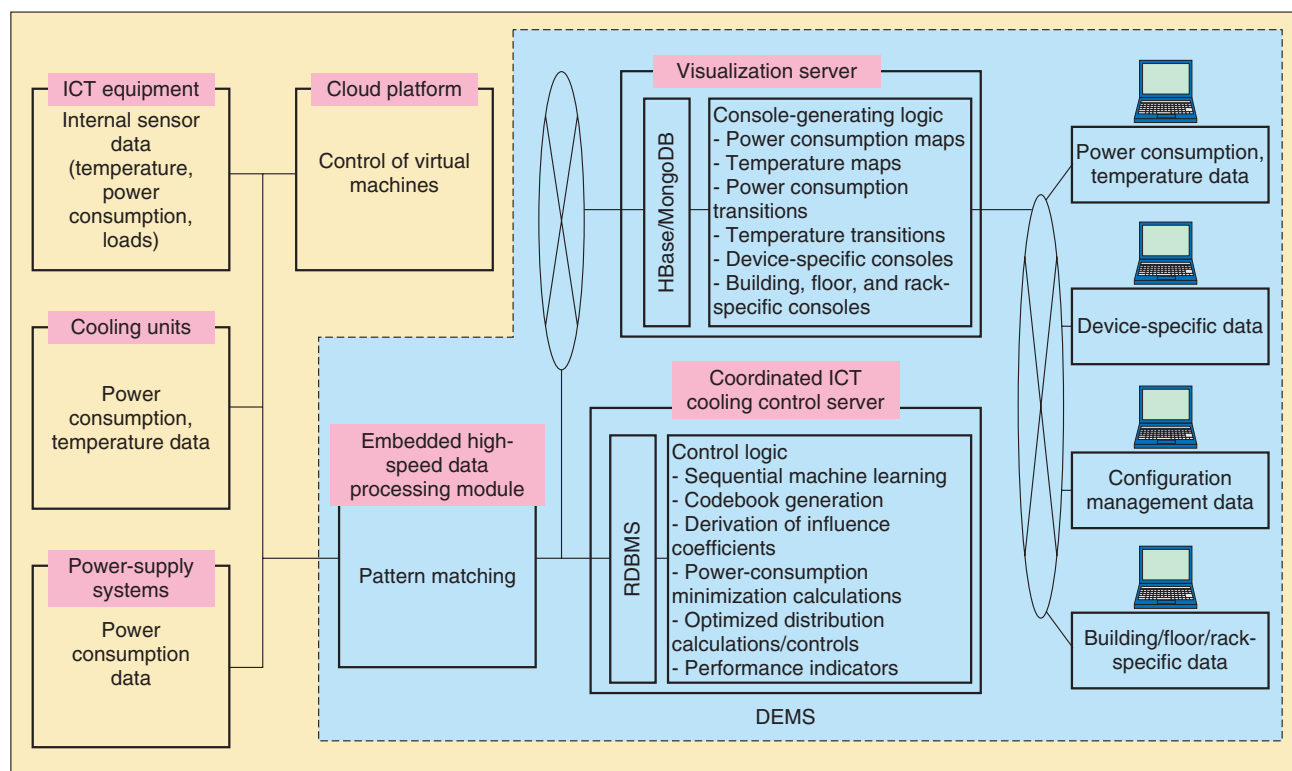
This article describes research efforts underway in NTT Energy and Environment Systems Laboratories in the area of datacenter energy management systems (DEMSs), which are systems designed to reduce energy consumption in datacenters and other telecommunication buildings. We present a new energy-saving system that comprises coordinated control of multiple cooling systems to minimize cooling power consumption, integration with information and communications technology (ICT) equipment to extend the benefits of the cooling-system controls, and high-speed data processing that can be adapted to large-scale datacenters.

1. Introduction

Energy consumption in datacenters has been increasing annually in recent years, driven by the rapid growth of cloud services. Information and communications technology (ICT) equipment, cooling units, and power supply systems account for the majority of energy consumed in datacenters. The energy use breakdown in one example revealed that 45% of the total energy was consumed by ICT equipment, 30% by cooling units, and 18% by power supply systems [1]. NTT has been researching and developing datacenter energy management systems (DEMSs) as a technology to cut energy consumption in datacenters and telecommunication buildings [2]. Our research has led to advances in coordinated ICT-cooling control technology that links multiple cooling units and regulates load distributions among ICT equipment to lower the total energy costs of ICT equipment, power supply systems, and cooling units.

We have also examined high-speed data processing technology designed to process the large amounts of device information within a datacenter.

A block diagram of the DEMS-coordinated ICT-cooling control architecture is shown in **Fig. 1**. The DEMS consists of a visualization server that enables operators to view temperatures and power consumption within the datacenter, a coordinated ICT-cooling control server, and an embedded high-speed data processing module. The visualization server collects temperature, power consumption, and load data using sensors inside ICT equipment rather than external sensors. It then provides fine distributions of the temperatures and power consumption of all the devices within the datacenter. The coordinated ICT-cooling control server takes the visualization server's sensor data and calculates the optimal cooling settings and ICT equipment load distributions to minimize the power consumption of cooling units and ICT equipment. The server then accordingly adjusts the cooling



RDBMS: relational database management system [3]

Fig. 1. Block diagram of DEMS-coordinated ICT-cooling control architecture.

controls and controls the ICT equipment's virtual machines and power supply systems via a cloud platform. The embedded high-speed data processing module quickly implements the optimized cooling settings and load distributions. It does this not by recalculating optimized values but by implementing pattern matching of sensor data using a codebook of optimal cooling settings and optimal load distributions calculated by the coordinated ICT-cooling control server.

2. Coordinated ICT-cooling control technology

Cooling unit assets located in datacenters and telecommunication buildings have conventionally been operated independently of other control systems. NTT Energy and Environment Systems Laboratories in partnership with NTT Facilities, Inc., has been working to raise cooling efficiency by unifying the controls that manage ICT equipment and multiple cooling units. An optimization problem can be formulated to find the combination of cooling settings that 1) ensure ICT equipment does not exceed its

upper temperature limits and 2) minimize the overall power consumption. A flowchart governing the coordinated control of multiple cooling units is shown in Fig. 2.

The first step necessary in implementing such controls is to develop a technique to estimate ICT equipment temperatures in response to changes in cooling unit settings. A temperature estimation equation can be found by analyzing data obtained during actual datacenter operation for the correlation between cooling-unit temperature changes and ICT equipment temperature changes.

A model of power consumption by cooling units is also necessary. The problem is the difficulty and calculation cost of deriving an accurate model, since cooling-unit power consumption is affected by multiple parameters.

To test and demonstrate the savings in cooling energy using the control flowchart mentioned above, we decided to use as our test case a linear model of cooling temperature settings for temperature estimates and cooling-unit power consumption. An automatic control system using prototype software

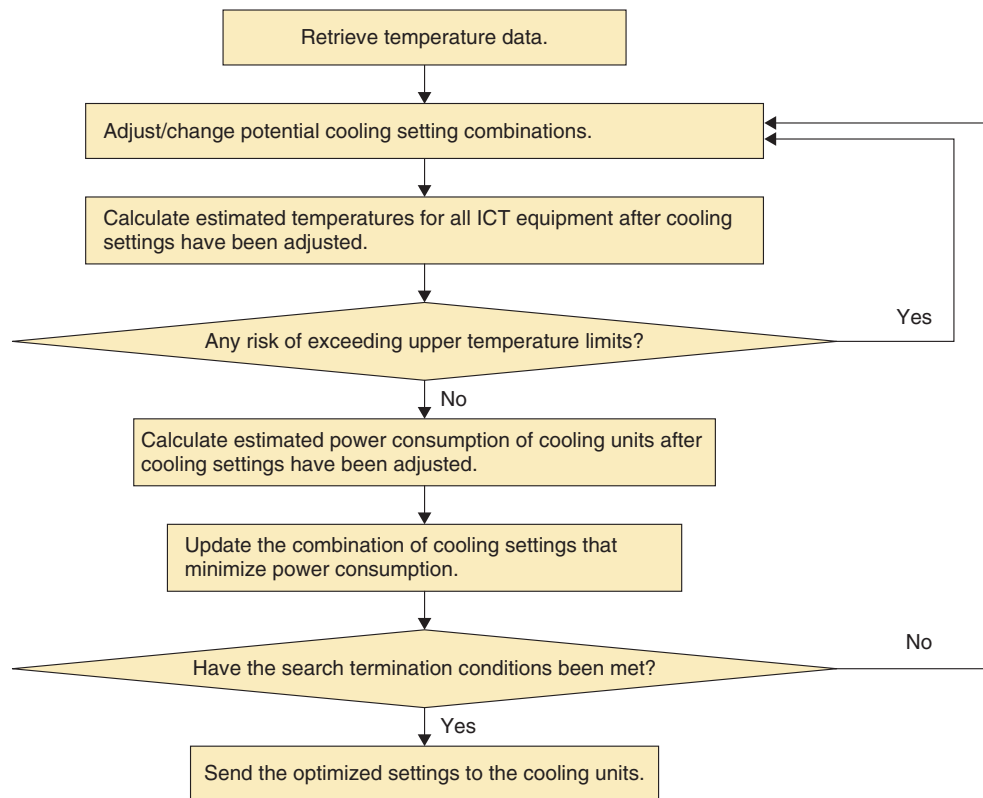


Fig. 2. Flowchart of coordinated control over multiple cooling units.




Fig. 3. Hourly correlations between PUE and outside temperature.

was constructed in a server room containing 9 cooling units and over 1000 operating servers. The test results are shown in **Figs. 3** and **4**. The test system obtained an average power usage effectiveness (PUE) [4] of

1.225 while maintaining ICT equipment temperatures at proper levels (at or under 27°C). Although the trials were done in an environment conducive to favorable PUE values, as aisle capping* was




Fig. 4. Distribution of measured temperatures.

implemented in a season with low external temperatures, the system did demonstrate that practical results could be achieved even with a linear model for temperature estimates and power consumption.

We believe that in addition to balanced control of multiple cooling units, coordinated ICT-cooling controls that regulate ICT equipment loads can further reduce datacenter energy consumption. Controlling ICT equipment via cloud platforms can reduce the number of pieces of equipment in operation and consequently reduce the load on the cooling system. Additionally, further controls of the loads placed on working ICT equipment can be added to raise the cooling system's efficiency. In this way, coordinated ICT-cooling control technology can cut the overall power consumption of ICT equipment while simultaneously reducing the cooling system's energy costs. This leads to the problem of determining which ICT equipment to run and how to distribute the load. At a typical datacenter, warm exhaust air of ICT equipment is recirculated with the cold supply air to other ICT equipment, and this lowers the cooling system's efficiency. Thus, governing loads in such a way as to minimize the recirculation of exhaust air is an effective means of reducing cooling power consumption.

The DEMS uses a temperature estimation model that includes the temperature relationship between

cooling and ICT equipment to forecast ICT equipment temperatures for a given load situation. Next, the DEMS sets limiting conditions on load distributions to implement the required cloud services with an ICT information model. This model, provided from a cloud platform, describes the relationship between cloud services and the ICT equipment necessary to deliver the cloud services. The DEMS finds a power consumption model for the ICT equipment, power supply systems, and cooling units and calculates cooling settings and load distributions that will minimize power consumption while meeting temperature conditions and load distribution constraints derived from the temperature estimation model and the ICT information model. The calculated load distribution is sent to a cloud platform, which adjusts the ICT equipment and the power supply systems. Cooling units are controlled following the calculated cooling settings (**Fig. 5**).

In datacenters that have areas with high and low volumes of recirculated heat from ICT equipment, hot spots—where ICT equipment near high recirculation areas runs hotter than that near low recirculation areas—will form if loads are distributed uniformly. In these cases, the cooling system works to eliminate the hot spots so that no ICT equipment operates above their upper temperature limits. Unfortunately, this results in overcooling of ICT equipment near low-recirculation areas. Without proper load distribution controls, the existence of excessively cooled ICT equipment means the cooling system is

* Aisle capping: an airflow control method that divides cool supply air and warm exhaust air by compartmentalizing the aisles between racks with walls or ceilings.

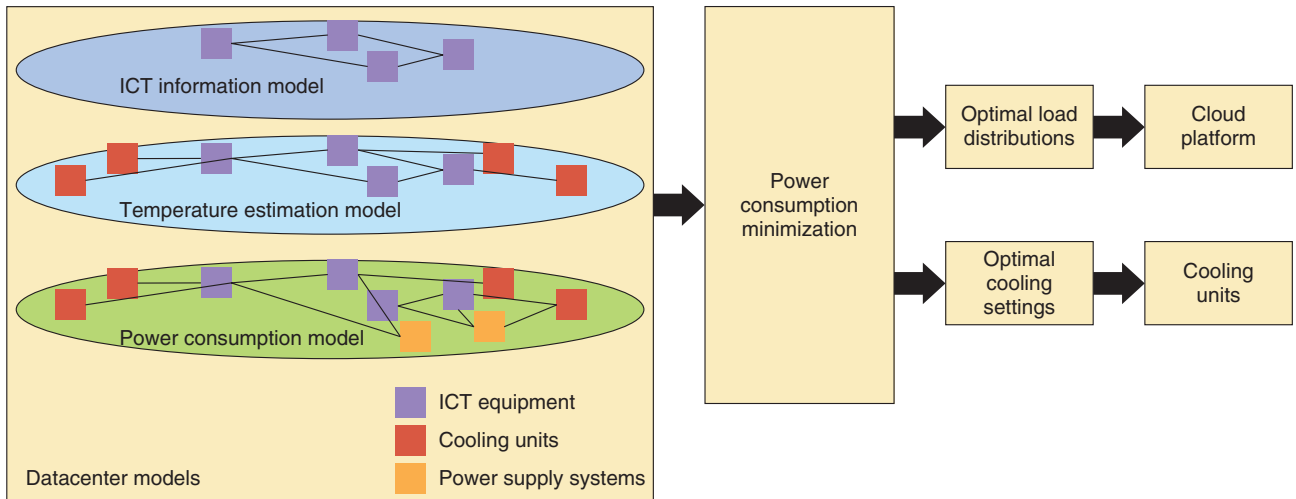


Fig. 5. Optimization of cooling settings and load distributions using coordinated ICT-cooling controls.


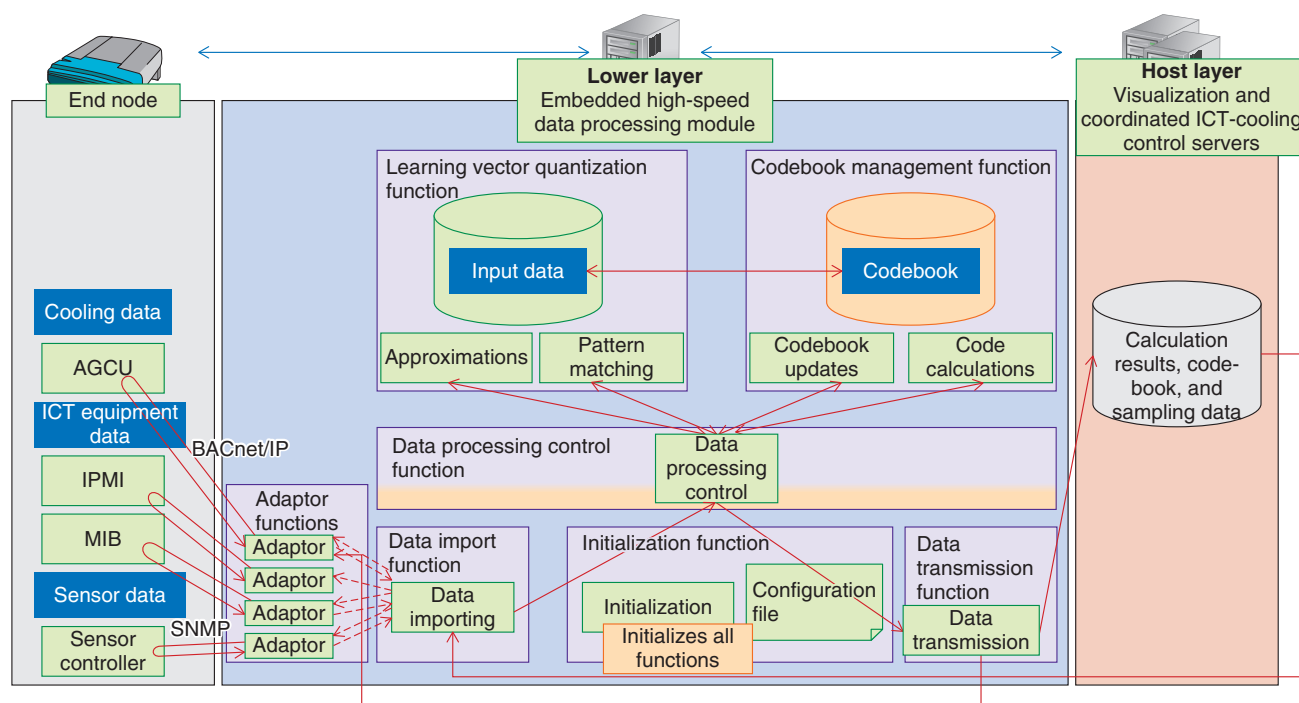


Fig. 6. Effect of load distribution control (overhead view of datacenter room).

using more energy than necessary. The DEMS, however, implements load distribution controls to allocate larger loads to ICT equipment near low-recirculation areas and lighter loads to that near high-recirculation areas and, thus, equalizes temperatures between devices. The DEMS combines load controls with coordinated cooling unit controls to calculate optimal cooling settings and eliminate overcooling conditions (**Fig. 6**). In the actual system, the relative amount of heat recirculation is accounted for in the temperature

estimation model that describes the temperature relationship between cooling and ICT equipment. In the model, the ICT equipment temperature is given as a function of the cooling settings and the load distribution. The DEMS then calculates the combination of cooling settings and load distributions that both meets ICT equipment maximum temperature limits and minimizes cooling power consumption. This optimization uses mathematical programming and other techniques.



BACnet/IP: building automation and control network, Internet protocol

Fig. 7. Block diagram of functions from high-speed processing technology for time-series numeric data.

3. High-speed processing technology for time-series numerical data

When the DEMS must collect data from thousands of pieces of ICT equipment in a datacenter or telecommunication building, and thousands of external sensors are installed to collect temperature and power consumption data, the sensor maintenance operations become exceedingly complex, and the cost is prohibitive. To overcome these problems, we investigated the possibility of collecting temperature, power consumption, and other data of ICT equipment using standard protocols, for example, the simple network management protocol (SNMP) or the Intelligent Platform Management Interface (IPMI), to construct a highly reliable, inexpensive, sensor-less system. Another issue we faced was how to process data inexpensively and energy-efficiently when calculating the optimal settings for multiple cooling units and the optimal load distribution among hundreds of pieces of ICT equipment from large data sets. Ordinary hardware solutions that scale up to meet increased calculation costs (i.e., distributing the processing load over additional servers) result in greater power con-

sumption and higher costs.

Our solution was to focus on a scalable architecture that contributes to reducing power consumption. In the DEMS data collection layer (lower layer), we implemented a high-speed processing module specialized for time-series numeric data that has interfaces to collect data from SNMP management information bases (MIBs) and from IPMIs. This module makes rapid adjustments to cooling unit controls and ICT equipment load distributions based on the collected data to minimize overall power consumption (**Fig. 7**).

The module collects cooling data, ICT equipment data, and other external sensor data through interface adaptors with air-conditioner group control units (AGCUs), MIBs, and IPMIs. The module manages a codebook generated from optimization calculations run by the host server and from input data in order to effect cooling controls and ICT equipment load distributions. Learning vector quantization is used for codebook pattern matching with sensor data as the input. In this way, quasi-optimal cooling settings and load distributions are found without redoing the optimization calculations, which would require significant

calculation costs. The quasi-optimal settings are sent to the cooling system and cloud platforms via interface adaptors to control the actual cooling units and ICT equipment. The module is capable of very high-speed processing because it is done in memory without any database access. We also believe the module can be implemented with inexpensive hardware because the approach uses quasi-optimal solutions instead of combinatorial computing and other costly optimization calculations.

4. Future research directions

NTT is moving ahead with energy-efficient data-center trials in which the cooling and ICT control

technology discussed in this article will be linked with cloud platforms.

References

- [1] "IT Trends Survey Report," the Japan Electronics and Information Technology Industries Association (JEITA), 2010 (in Japanese).
- [2] K. Matsuda, T. Tominaga, and M. Yanagi, "Energy Management of Telecommunication Buildings and Datacenters," NTT Technical Review, Vol. 10, No. 1, 2012.
<https://www.ntt-review.jp/archive/ntttechnical.php?contents=ntr201201fa2.html>
- [3] RDBMS.
http://en.wikipedia.org/wiki/Relational_database_management_system
- [4] The Green Grid.
http://www.thegreengrid.org/sitecore/Content/Global/Content/whitepapers/Green-Grid-Metrics?sc_lang=en



Masayuki Nakamura

Senior Research Engineer, Energy Optimization Technology Group, Energy System Project, NTT Energy and Environment Systems Laboratories.

He received the B.E., M.E., and Ph.D. degrees in mathematical engineering and information physics from the University of Tokyo in 1988, 1990, and 1998, respectively. Since joining NTT Applied Electronics Laboratories in 1990, he has studied information processing in sensing and control systems. He moved to NTT Energy and Environment Systems Laboratories in 2002. He is a member of the Society of Instrument and Control Engineers and the Institute of Electrical Engineers of Japan (IEEJ). He received a presentation award from IEEJ in 1998 and a research promotion award from the Japan Society for Environmental Chemistry in 1999.



Akira Takeuchi

Senior Research Engineer, Energy Optimization Technology Group, Energy System Project, NTT Energy and Environment Systems Laboratories.

He received the B.E. and M.E. degrees in electronics engineering from Kyushu University, Fukuoka, in 1990 and 1992, respectively. He joined NTT Interdisciplinary Research Laboratories in 1992. He moved to NTT Energy and Environment Systems Laboratories in 2002. His research interests are power converters, energy control technologies, and optimization techniques. He is a member of the Institute of Electronics, Information and Communication Engineers and IEEE.



Hiroyuki Enomoto

Senior Research Engineer, Energy Optimization Technology Group, Energy System Project, NTT Energy and Environment Systems Laboratories.

He received the M.E. degree in management of technology (MOT) from Shibaura Institute of Technology, Tokyo, in 2007. He joined NTT Communications in 2003 and NTT Energy and Environment Systems Laboratories in 2006. His research interests include power distribution control, optimization of power transmission, high-speed processing technology of numeric data, and analysis of economic effects of technology. He was one of the authors of the world scenarios series "Engineering & Construction: Scenarios to 2020" at Davos World Economic Forum 2008, and his name is listed in the Special Acknowledgement for Extraordinary Contributions of the report. He is a member of the Japan MOT Society.