ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

Control of Computer Room Air Conditioning using IT Equipment Sensors

Geoffrey C. Bell

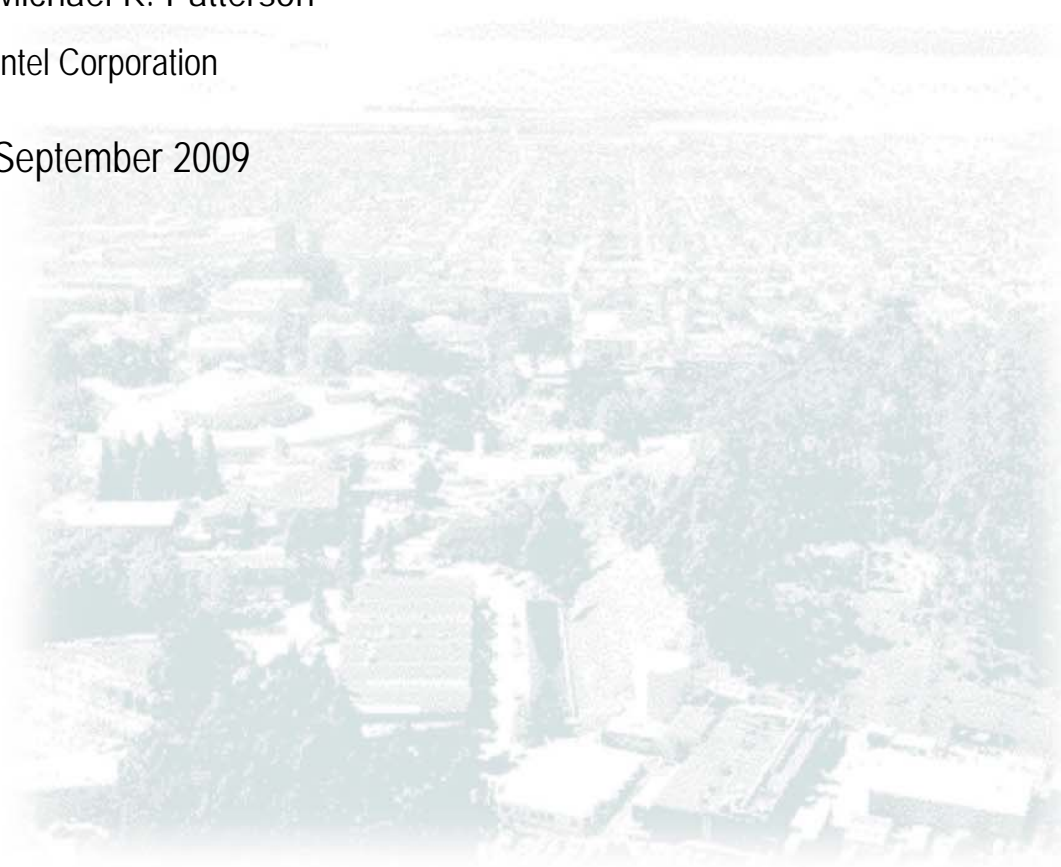Lawrence Berkeley National Laboratory

Bill Storey

Wunderlich-Malec Controls

Michael K. Patterson

Intel Corporation

September 2009

DISCLAIMER

Acknowledgments

Please cite this report as follows:

Bell, Geoffrey C., Bill Storey, and Michael K. Patterson. 2009. *Control of Computer Room Air Conditioning using IT equipment sensors.* California Energy Commission.

Table of Contents

List of Figures

Abstract

The goal of this demonstration was to show how sensors in IT equipment could be accessed and used to directly control computer room air conditioning. The data provided from the sensors is available on the IT network and the challenge for this project was to connect this information to the computer room air handler's control system. A control strategy was developed to enable separate control of the chilled water flow and the fans in the computer room air handlers. By using these existing sensors in the IT equipment, an additional control system is eliminated (or could be redundant) and optimal cooling can be provided saving significant energy. Using onboard server temperature sensors will yield significant energy reductions in data centers. Both chiller-plant and pumping energy can be reduced by 30 to 40 percent. Intel hosted the demonstration in its Santa Clara, CA data center. Intel collaborated with IBM, HP, Emerson, Wunderlich-Malec Engineers, FieldServer Technologies, and LBNL to install the necessary components and develop the new control scheme. LBNL also validated the results of the demonstration.

Key Words: control, data center, HVAC, IT sensors.

Executive Summary

The most common method used for controlling data center cooling devices employs the return air temperature from the center. This method provides a blended air temperature regime from the center to the cooling devices that masks the actual needs of the servers. Consequently, the cooling devices will operate inefficiently because they do not accurately anticipate cooling requirements.

However, by retrieving server inlet air temperature from the server's manageability network and linking this information to the building management system to control the cooling system in the data center, energy reductions can be achieved by more precisely controlling the amount of cooling required.

Accordingly, the primary goals for this project's proof-of-concept were to:

1. Demonstrate ability to provide operating information from the server's IT manageability network to a facility management system (FMS).

2. Demonstrate ability of the FMS to provide set point changes to data center conditioning systems in response to temperature information from the servers.

This project demonstrated and validated the ability of the latest generation computer servers to provide temperature information to a facility management system and to subsequently have the FMS determine and provide operating setpoint(s) for cooling system operations. Effective communications and closed loop control from the Intel SC-11 Data Center servers to the Intel Facility Management System (IFMS) were developed without significant interruption or reconfiguration of the devices.

Using onboard server temperature sensors will yield significant energy reductions in data centers. In this data center demonstration project, the method achieved a 43 percent reduction of airflow volume that reduced fan energy-use by more than 80 percent. Additional energy savings will be realized from higher inlet air temperatures to the data center servers allowing higher-temperature chilled water production. Thus, both chiller-plant and pumping energy can be reduced by 30 to 40 percent.

Introduction

Data center cooling is usually provided with computer room air conditioner (CRAC) devices, with direct-expansion refrigeration cooling coils, or computer room air handler (CRAH) devices, with chilled water cooling coils. Typically, these devices use return air temperature sensors as the primary control-variable to adjust the air temperature supplied to the data center. This control approach significantly limits energy efficiency because it is the worst location for maintaining temperatures at the inlet to IT equipment. Importantly, server manufacturers have agreed that their main operational parameter is the air temperature provided at the inlet of the server itself, not the proxy temperature returning to the cooling device. Therefore, a much higher degree of monitoring and efficient control would be achieved by using front-panel, inlet temperature sensor data.

The majority of servers have a host of platform information available from their Information and Communications Technology (ICT) management network. Server front-panel, i.e., server inlet, air temperature is monitored and available through each server's ICT manageability network connection that supports simple network management protocol (SNMP) or Intelligent Platform Management Interface (IPMI)[1].

By using the front-panel temperature sensor data from the ICT network as input information to the facility management system (FMS), the control of the cooling devices in a data center can be more effective. Energy waste would be reduced by networking the two management systems, ICT and FMS. Therefore, significant gains in energy efficiency can be realized while improving the visibility of data center operations for both the ICT and facilities staff.

Project Approach

The project's scope was developed by Lawrence Berkeley National Laboratory (LBNL) in response to a California Energy Commission (CEC) Memorandum Agreement NO. MUC-08-03 with funding from the Public Interest Energy Research (PIER) program. LBNL subcontracted with Wunderlich-Malec Engineering (WME) to support of the demonstration and proof-of-concept project. Services provided by WME included: instrumentation and controls design; equipment installation; controls programming and integration, and project commissioning and coordination. The demonstration project's proof-of-concept was performed at Intel Corporation's headquarters located at 2200 Mission College Ave., Santa Clara, CA. Intel provided a demonstration area in its active SC-11 Data Center. This demonstration area contained both HP C-class blade, IBM HS-20 blade chassis with onboard server temperature data used for server fan control. Participants in the project included Intel, Liebert, IBM, HP, and Fieldserver Technologies.

To demonstrate and validate the ability of the latest generation servers to provide two-way communication of operating information between the ICT manageability network and the FMS, two primary goals needed to be achieved:

1. Prove operating information can be uploaded/downloaded from/to the server's ICT manageability network to/from the building management system.

---

[1] The **Intelligent Platform Management Interface** (IPMI) specification defines a set of common interfaces to a computer system which system administrators can use to monitor system health and manage the system. Several dozen companies support IPMI. IPMI operates independently of the operating system and allows administrators to manage a system remotely and offers enhanced features when used with system management software.

2. Establish that the FMS can provide set point changes to data center conditioning systems in response to temperature information from the servers and receive verification of these changes.

Goal 1:

The center's servers were connected to Intel's Facility Management System (FMS). The existing Intel FMS, which included a data repository system known as SPNet, was used for acquiring server temperatures and power data. The SPNet data application system queries the servers through an Intelligent Platform Management Interface (IPMI). Therefore, the needed control points were being collected that allows operator monitoring, alarms, and trending. In addition, the SPNet was used for moving large amounts of trended data to the FMS server in real-time. Data were presented using Intel's FMS Cimplicity Human-Machine-Interface (HMI) software and stored within the system for input to the data center cooling system devices.

Goal 2:

The technology used to communicate to the data center's cooling system employed an existing GE Fanuc Series 90-70 Programmable Logic Controller (PLC) system. This PLC system provides process control through the GE Fanuc Cimplicity HMI, noted above. The interface between the existing GE 90-70 PLC and the Liebert CRAH unit was accomplished by installing a Fieldserver Technologies FS-B2010 bridge that provides direct communications from the PLC Ethernet port to the Fieldserver Modbus™ port which then communicated from the Modbus™ port to the Liebert CRAH cooling device. The Cimplicity FMS local script capability was used to develop graphics that allowed facility operators to choose which server temperature data points were used for control.

All data collected from facility monitoring points and additional monitoring points used to adjust facility setpoints were recorded for trending and analysis.

Operational Sequence

Since the demonstration project was performed in an active, on-line data center, a fall-back operational sequence was provided. The CRAH unit's setpoints inputs were selectable between one of two modes: room thermostat mode or server control mode, which used front-panel temperature sensors.

Room Thermostat Mode: When using the room thermostat for cooling, the control algorithm was identical to the current thermostat/CRAC Sequence of Operation.

Server Control Mode: The user was able to choose a rack or set of racks from the primary rack group to associate to a the CRAH unit. This selection allowed the association of the "worst case" rack (hottest) with the control of the CRAH. Within each rack, the user chose a server to assign for each process variable; 1) temperature input for the physically highest server, 2) temperature input for the physically lowest server in the rack.

Chilled water valve control was implemented in phases for demonstrating functionality while in server control mode. The process consisted of three steps control evolution:

Step 1: Chilled water (CW) valve modulation was controlled locally by the CRAH unit's built-in controls, as defined in their current and existing sequence of operations under the Room Thermostat Mode.

Step 2: Under the server control mode, chilled water valve was modulated locally by the Liebert control panel and served as a sub-master proportional-integral-derivative (PID) control loop to the room level Master PID. The room thermostat control point is substituted with a process variable from the physically highest server in the associated rack from the Master PID.

Step 3: The physically lowest server or blade chassis in the selected rack was used to control the CW valve and overriding the Liebert control over the CW valve. The temperature signal from the selected server became the process variable for the PID control loop for maintaining the correct airflow temperature. This will happen at the PLC versus the CRAC. The airflow temperature setpoint was user-adjustable and opened the associated CRAH 3-way CW control valve. Note that prior to implementing PLC control over the CRAC CW valve (Step 2) in the server mode, functionality was proven. Any faults that would have occured while in the server mode, "failed" to thermostat mode, i.e., localized CRAH controls.

Startup

Start-up services verified that all ICT values were correctly monitored, all alarms were functional, and all control loops worked and were well tuned. Prior to startup, one of the key aspects of good data center control is to implement some level of air management control by determining which perforated floor tiles affect which local server racks. Evaluation of the CRAH airflow and recommendations for tile adjustments for specific racks was part of the project's scope.

Implementation Details

The demonstration project area of work is shown in Figure 1. CRU 44A and CRU43A are Liebert Model FH740CVAI00, 460 Volts, 60Hz, 3 Phase Deluxe System/3 Environmental Control Systems. Racks A 01 through A 13 contained HP and IBM servers from which server ICT operation information was collected and used as inputs to control the CRAH operating setpoints.
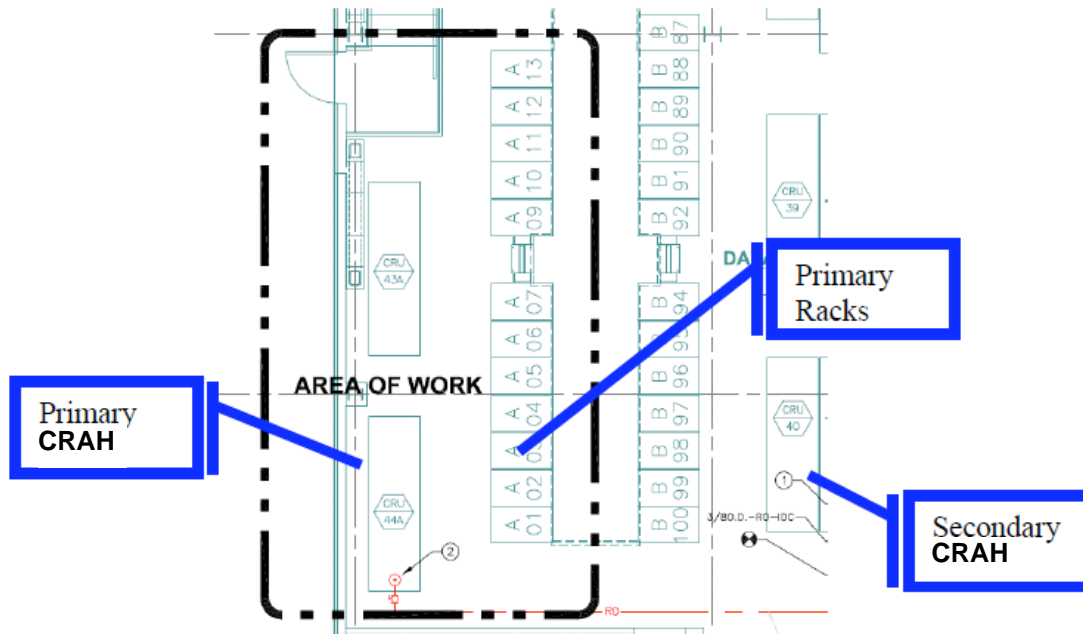
Figure 1: SC-11 Demonstration Project Work Area

Figure 2 depicts the data center HVAC mechanical layout of the demonstration area. Air temperatures surrounding the demonstration area were monitored by eight (8) independent Rotronics temperature sensors, providing 0.5 °F accuracy on temperatures in order to verify equipment values and process performance. Additional sensors were installed to provide monitoring of chilled water flow to the local CRAH unit within the demonstration area. Sensor points were connected to the existing GE Fanuc FMS using available input/output (I/O) points. Additionally, WME provided project
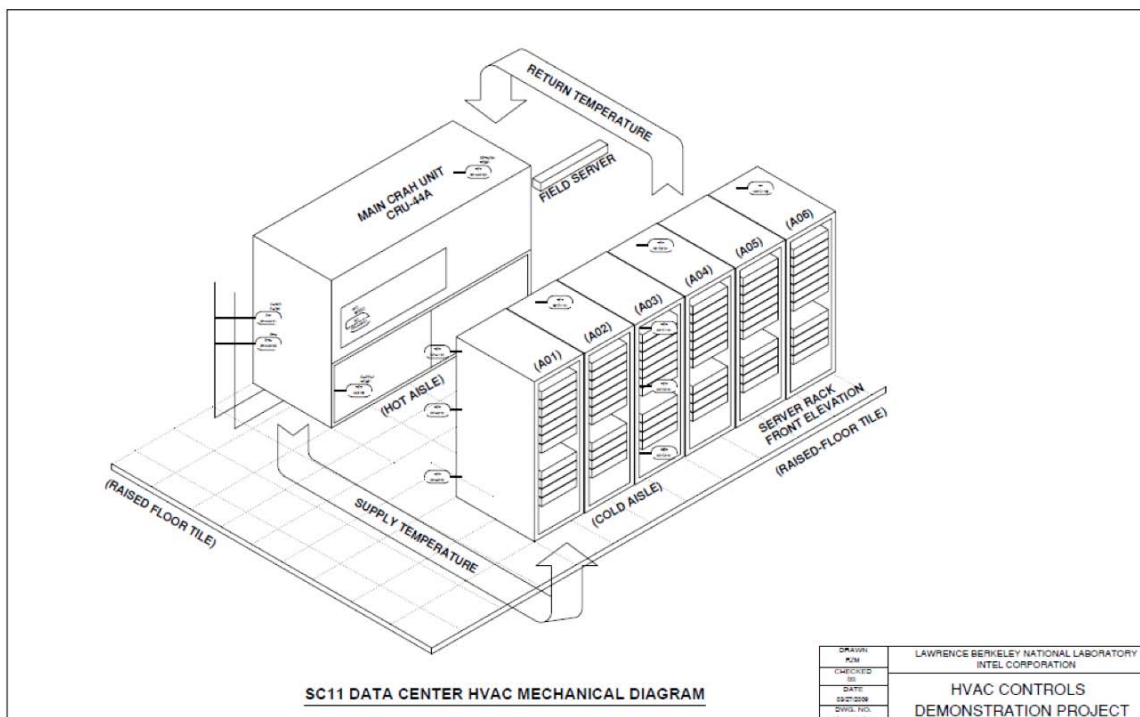

Figure 2: SC-11Data Center HVAC Mechanical Diagram

management, site coordination, meeting notes, scheduling, and installed network and interface equipment.

Control system architecture block diagrams, wiring diagrams, and equipment layout records were developed to support in field temperature, pressure and flow monitoring in the demonstration area. Refer to Figure 3 for the network communication configuration arrangement within the demonstration area.
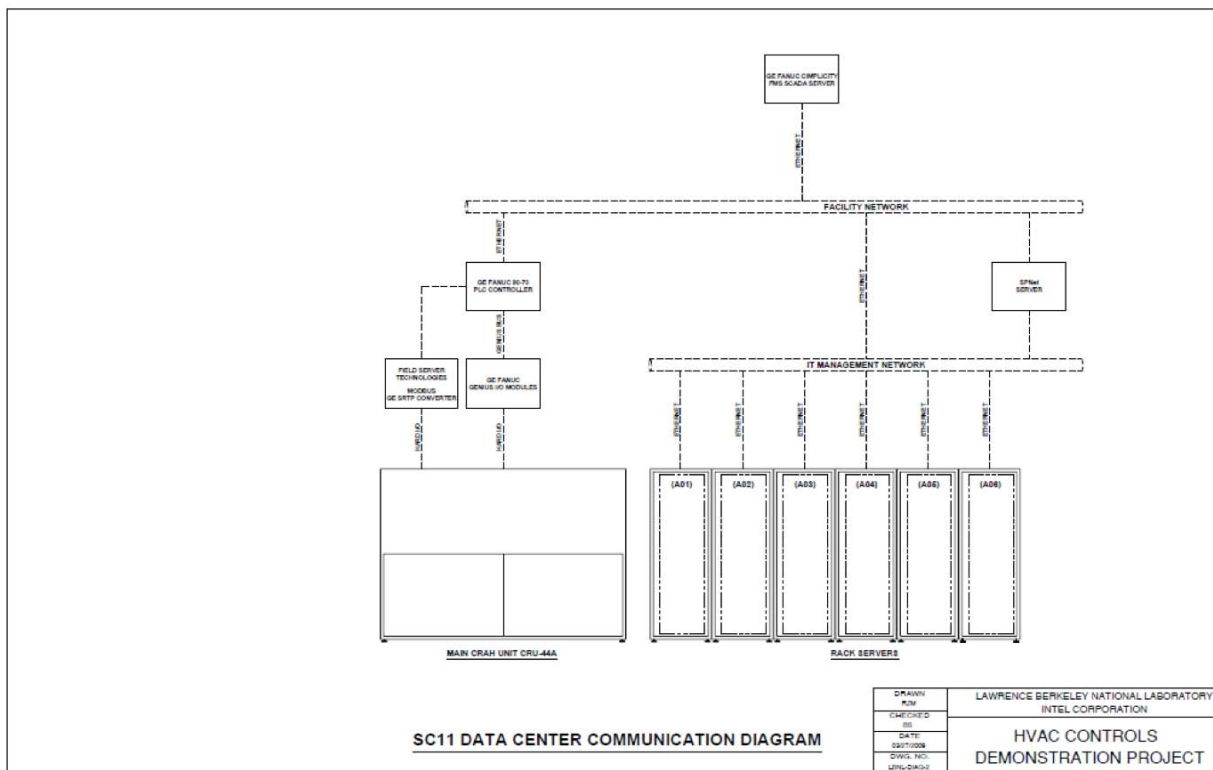

Figure 3: SC-11Communication Diagram

Project Outcomes

The primary goal for the demonstration project's proof-of-concept was achieved. Effective communications and closed-loop control were developed from the Intel Data Center servers to the Intel Facility Management System (FMS) to provide control through a programmable logic controller (PLC) with a proportional-integral-derivative (PID) control routine to adjust setpoints for supply air temperature and fan volume flow in a Emerson Electric (Liebert) CRAH unit without significant interruption or reconfiguration of the devices.

The existing Intel data repository system, known as SPNet, was used for acquiring server temperatures and power data. The SPNet data application system queried the servers through the Intelligent Platform Management Interface (IPMI). Therefore, the needed control points were being collected and the SPNet could be used for moving large amounts of data to the FMS server in real-time. Data were displayed using Intel's FMS Cimplicity Human-Machine-Interface (HMI) software and stored within the system

for input to the data center cooling system devices. See Figure 4; Legend, Server - Result (green plot line).
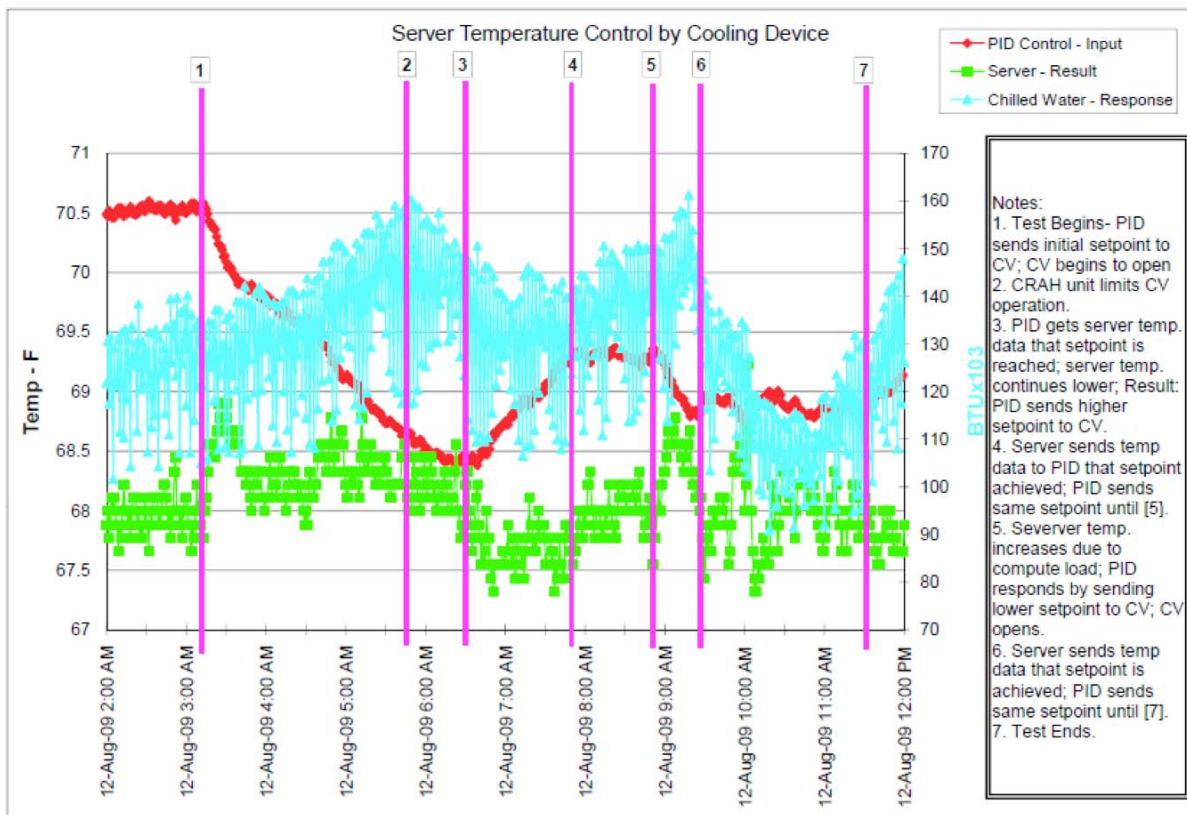

Figure 4: Server Temperature Control by Cooling Device

The technology used to communicate with the data center's cooling system (CRAH units) employed the existing FMS, which is a GE Fanuc Series 90-70 PLC system. It provides process control through a GE Fanuc Cimplicity HMI system that was used for operator monitoring, alarms, and trending. The interface between the existing GE 90-70 PLC and the Liebert CRAH unit was accomplished using a FieldServer Technologies FS-B2010 bridge, which provided direct communications from the PLC Ethernet port to the Fieldserver Modbus™ port to the Liebert CRAH cooling device. The Cimplicity FMS local script capability was used to develop graphics that allowed facility operators to choose which server temperature data points were used for control. See Figure 4: Notes.

Figure 5 provides the data from temperature sensors mounted at the inlet of the server rack. This plot verifies that the server's request for an inlet air temperature of 70°F was being achieved. Average temperature input to the racks was 69.5 °F +/- 0.1 °F. Little variation in the cold aisle temperature was observed. Figure 6 shows the change monitored in the leaving air temperature from the rack as the sever compute cycles increased within this rack. Figure 7 provides temperature information gathered above three racks. The average temperature of 70°F may indicate that the supply airflow volume was more than actually required to condition the racks. With additional study

and observation, this temperature can be used to modulate the CRAH fan speed. See Next Steps discussion below.
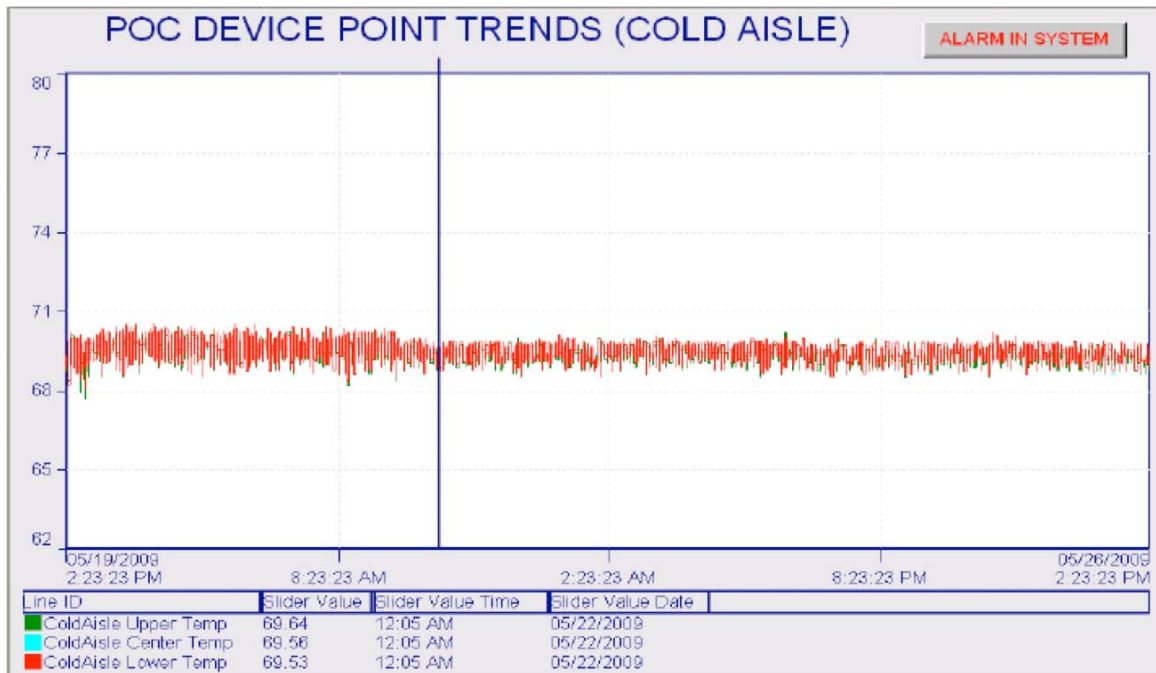
Figure 5: Server rack Inlet Air Temperature (SIAT)
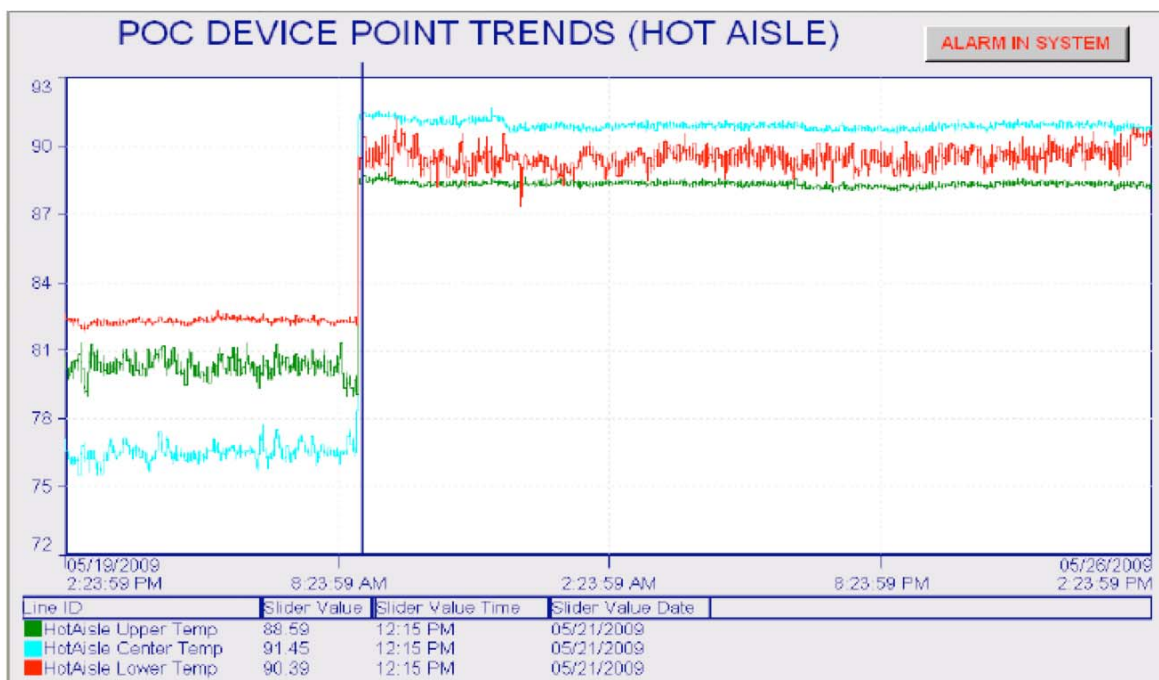
Figure 6: Server-rack Leaving Air Temperature (SLAT)
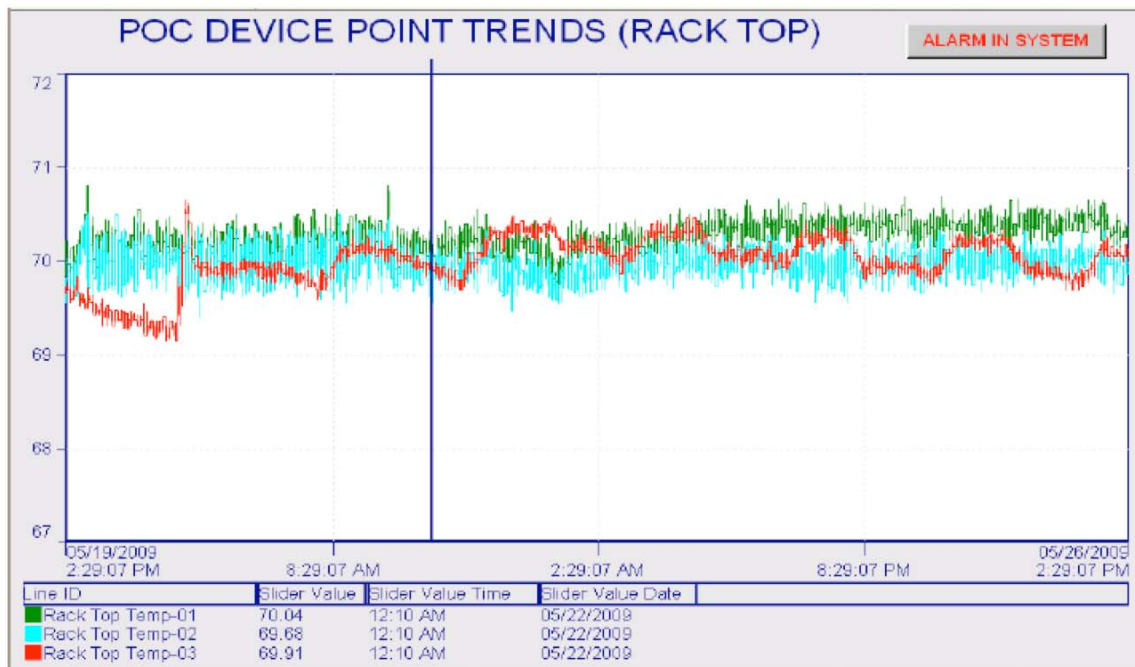
Figure 7: Air Temperature at Top of Rack

HMI screens were also developed during the demonstration project. Web access was available within the Intel facility for authorized personnel to monitor data and operations. The equipment/rack screens included 2D rack and individual server temperature, status, and power. See Figure 8 and Figure 9.

Figure 8: Main Menu Screen

Figure 9: Overall System Status Screen

An HMI screen was developed with Cimplicity graphics that allowed facility operators to choose which server temperature data points were used for controlling the CRAH device. In Figure 10, note (2) server locations outlined in color as opposed to standard gray. This simple convention allowed facilities to know the current location for temperature inputs in the system.

Figure 10: Control Input Server Location Screen

Conclusions

The project's primary goal to demonstrate a proof-of-concept that onboard server temperature sensors can provide usable data for controlling a data center's cooling devices was achieved. Effective communi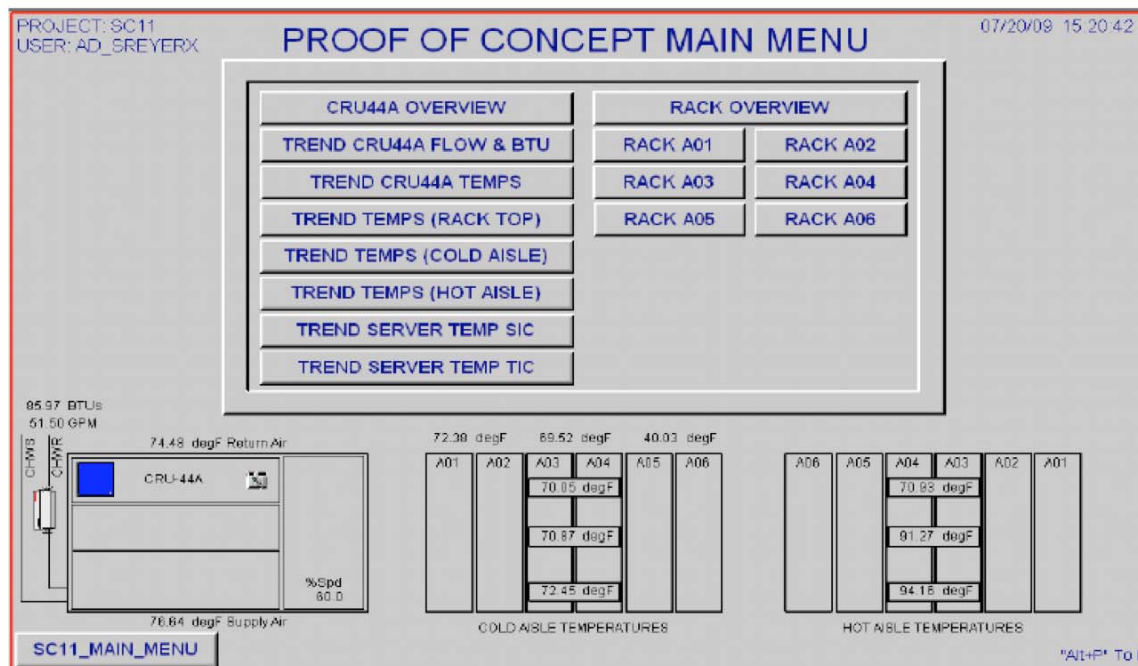cations and closed-loop control were developed from the Intel data center servers to the Intel Facility Management System (FMS) that provided control through a programmable logic controller (PLC) with a proportional-integral-derivative (PID) control routine to adjust setpoints for supply air temperature and fan volume flow in a Emerson/Liebert Electric CRAH unit without significant interruption or reconfiguration of the devices. It is clear that energy efficiency improvements can be realized by using the actual server temperature sensors to modulate CRAH device operation rather than a surrogate temperature.

Using onboard server temperature sensors will yield significant energy reductions in data centers. In this data center demonstration project, the method achieved a 43 percent reduction of airflow volume that reduced fan energy-use by more than 80 percent. Note that the internal CRAH fan control system limited achieving maximum energy reduction. Additional energy savings will be realized from higher inlet air temperatures to the data center servers allowing higher-temperature chilled water production. Thus, both chiller-plant and pumping energy can be reduced by an estimated 30 to 40 percent.

Next steps

An optimized control scheme was not achievable. The control of the CRAH unit's cooling valve was implemented by "fooling" the valve into modulating flow by providing it with a new set point any time a new flow value was calculated by the external PID control scheme. The control logic scheme within the CRAH unit could not be bypassed so it continually interrupted this external control scheme.

Fan speed in the CRAH unit was limited to no lower than 60 percent of full speed. The external PID control scheme tried to lower the fan's speed without result, again due to the onboard CRAH unit's control scheme.

Finally, due to the configuration of the data center area, isolation of a CRAH unit to provide dedicated cooling to only the servers used for control input was not feasible.

Recommendations

Control CRAH fan speed

The CRAH fan speed could be controlled by a user-defined differential between the inlet temperature the physically lowest and physically highest server in a selected rack by considering their respective front panel temperature readings. The expectation is that the upper racks typically suffer from recirculation, depending on airflow volumes. A high fan speed will cause more cool air to move out of the raised floor plenum and over-supply the rack resulting in excess airflow over the top of the racks. If the fan speed is too low, the upper servers will suffer recirculation from the hot aisle air that comes around over the top of the rack to feed the inlet to the upper server(s). The ideal airflow

will be just enough to carry the cooling airflow to the top of the rack, preventing recirculation, without over- or under-supplying the needed cooling air.

Coordinate data center control

An alternative FMS control solution that gives the data center operator a much higher degree of monitoring and control should be explored. The control of the cooling system would combine the CW valve and fan speed control loops to maximize efficiency. The CW control loop modulates the cooling fluid temperature, which is driven by the server inlet air temperature data from the physically-lowest server in the rack. This would ensure that the correct supply air temperature is provided. The fan-speed control loop modulates the airflow volume from the cooling system. As noted above, a variable speed fan would control the differential temperature between the lowest server in the rack and the highest server at the top of the rack. Coordinating both of these FMS control loops with the server's data would minimize of the data center's energy use. When an optimized control schemes are implemented in data centers, energy reductions in the range of 30 to 40% can be realized.

Integrate ICT and Facilities efforts

One of the persistent problems yet to be overcome in the data center space has been the separate worlds of the ICT and Facilities design and operations. Therefore, challenges with the integration of the ICT management data with the facilities control system continue to exist. It is expected that this will be a most productive area for improving data center efficiency over the next several years.

References

Citation

1. Patterson, M., M. Meakins, D. Nasont, P. Pusuluri, W. Tschudi, G. C. Bell, R. Schmidt, K. Schneebeli, T. Brey, M. McGraw, W. Vinson, J. Gloeckner, Energy-Efficiency Through the Integration of Information and Communications Technology Management and Facilities Controls, ASME, Proceedings of IPACK2009, InterPack'09, 2009.

Project Product Information

Technology used in the demonstration area included the existing Facility Management System, which was a GE Fanuc Series 90-70 PLC system used for process control and a GE Fanuc Cimplicity HMI system used for operator monitoring, alarms, and trending. Please refer to http://www.gefanuc.com/products/family/automation for details on these GE Fanuc products. A product data sheet is attached as well.

The interface between the existing GE 90-70 PLC and the Liebert CRAH unit was accomplished using a Fieldserver Technologies FS-B2010 bridge, allowing direct communications from the PLC Ethernet port to the Fieldserver Modbus port. Please refer to http://www.fieldserver.com/products/gateways/Ethernet-Serial_Gateway.php for details on the Fieldserver product. This product was chosen due to the flexibility to adapt to alternate protocols on both the controller and CRAH, providing a flexible

solution for future units/sites for demonstration. A product data sheet is attached as well. The CRAH unit contained a serial interface unit by Liebert, known as a Intellisolt 485. Please refer to http://www.liebert.com/product_pages/Product.aspx?id=185&hz=60. A product data sheet is attached as well.

Glossary

AHU: air-handling unit
ASHRAE: American Society of Heating, Refrigerating, and Air-Conditioning Engineers
CRAC: computer room air conditioner with DX coil and compressor
CRAH: computer room air-handler with chilled water coil
CW: chilled water
FMS: facility management system
HMI: Human-Machine-Interface
HVAC: heating, ventilating, and air-conditioning
IT: Information Technology
ICT: Information and Communications Technology
IPMI: Intelligent Platform Management Interface
PID: proportional-integral-derivative
PLC: Programmable Logic Controller
SNMP: simple network management protocol
UPS: uninterruptible power supply
VFD: variable frequency drive

Logos