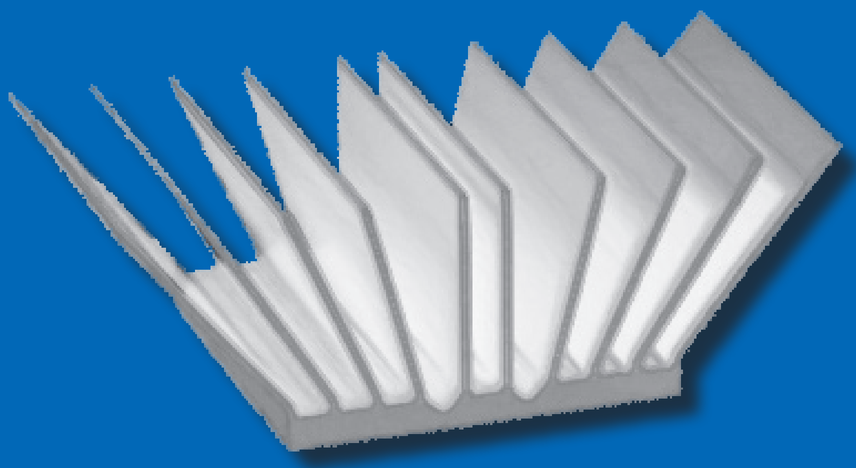

ATS WHITE PAPER

Adaptive Cooling in Datacenters

Advanced Thermal Solutions, Inc.
89-27 Access Road | Norwood, MA 02062 | USA
www.qats.com | T: 781.769.2800 | F: 781.769.9979

ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®

Adaptive Cooling

in Data Centers

With servers increasing in both power and cost, data centers are challenged not only to keep the components cool to prevent their failure, but also to control and ideally decrease the cost of cooling of the cabinets. One promising technique is adaptive cooling. As an example, consider a service provider housing 100 racks of 40 KW each in a data center. The total power required to run the servers is 4 MW. The fan power needed to remove the air for cooling purposes is about 400 KW. Another 4 MW is required to cool the data center assuming cooling requires the same amount of energy as power consumption. Any improvement in cooling efficiency brings down the required power consumption, cost of operation and maintenance needs.

Current application of heat removal technologies based on maximum heat load results in over provisioning at all levels of chips, systems and data centers [1]. This over provisioning stems from the inability to track the dissipated power and to allocate cooling resources accordingly due to insufficient sensing and control. A smart power cooling system must be developed, from chips to data centers, to improve the state of the art. Such a flexible system should be applied to chip level, system level and finally the data center, in order to optimize the cooling required. In general, a flexible cooling system will have four distinct areas [1]:

1. Flexible Actuators. Actuators are generally step motors connected to valves for controlling the amount of coolant. They receive commands and respond accordingly.
2. Distributed Sensing. A network of sensing elements for measuring component temperatures, inlet and outlet air temperatures of the air flow at the rack level, the temperature of the coolant and air temperature at various locations in the data center.
3. Control Policies. A set of rules that define the performance levels of the processors, operating costs, acoustic levels, etc.
4. Evaluations. A set of instructions in the control system that collect all the data from the sensing elements and send commands based on policies.

The smart system should be flexible at all levels to achieve the best performance. For example, the new system should be able to control the amount of liquid cooling going into the chip, if applicable, or controlling the voltage of thermoelectric coolers to match the junction temperature to the desired level. At the heat sink level, if it is an active one, the fan speed can be adjusted to regulate the performance of the sink. At the system level the fan tray speed can be changed or the amount of coolant flow can be adjusted according to the requirements of the system. At the data center level the flexible system must be able to adjust the air conditioner blower speed or change the CRAC temperature by adjusting the valves to the chilled water or changing the opening in the tiles to distribute the flow accordingly.

In an article [2], Turner Construction studied the cost of building and cooling a 4,000 KW data center. They analyzed three different sized buildings with different heat densities. In the first option, they designed a building to accommodate 450 W/m^2 (50 W/ft^2). This requires a 9000 m^2 (80,000 ft^2) space to accommodate the 4000 KW. This facility would cost approximately $17 million. By increasing the cooling capacity

to 3600 W/m² (400 W/ft²), the facility size could be reduced to 1000 m² (10,000 ft²) with the construction cost of about $10 million with a saving of approximately $7 million. In a bigger data center, e.g. 33,750 m² (300,000 ft²), the savings in power consumption for cooling would be
around $20 million. The numerical results are shown in Figure 1. This shows how important it is to optimize the cooling capacity, because real estate costs vastly outweigh any premium required for higher density cooling.

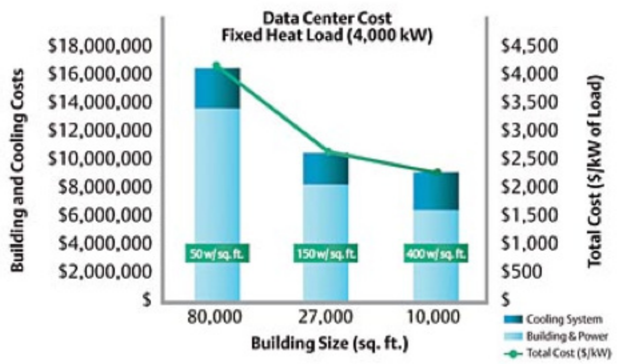

Figure 1. Cost Effectiveness of Higher Density Cooling [2].

But this is just the power consumption. The cost of operation can be significantly higher. The down time as a result of failures in a cooling mechanism system can add to the above costs and to the customer dissatisfaction. Ease of maintenance and reliability of the cooling system should have even more priority than the power consumption.
In [3], they indicate that increasing the efficiency of the infrastructure from 41 percent to 51 percent, could allow the addition of 25 percent more servers to the data centers. This is a very useful insight for IT organizations to improve the efficiency of cooling in a space-constrained data center. Data center managers generally focus on increasing the rack density to get more work out of their existing space, but the problem is typically not the lack of rack space, but how to cool the data center efficiently.

Figure 2 shows a conceptual schematic of an adaptive control system. The ultimate goal in any cooling system is to keep the junction temperature of the device under control.

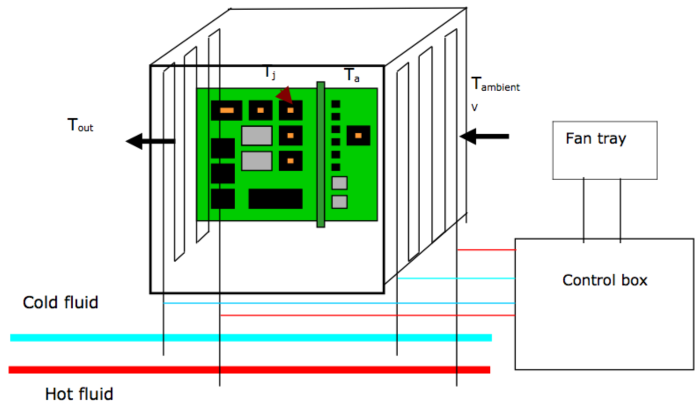

Figure 2. Hypothetical Schematic of an Adaptive System for a Server.

Assuming there is a heat sink on top of the device:

$$T_j = T_a + P \times R$$

Where,
T_j = junction temperature of the device (oC)
T_a = local air temperature in the server (oC)
P = device power (W)
R = thermal resistance of the heat sink (oC/W)

The thermal resistance of the heat sink is a function of air velocity:

$$R = a + b \times V^c$$

Where,
a, b and c are constants for a given heat sink
V = approach air velocity (m/s)

The local ambient temperature T_a is a function of data center ambient temperature $T_{ambient}$, volumetric flow rate of coolant going through the heat exchanger and speed of the flow generated by the fan tray.

$$T_a = f(\dot{m}, V, T_{ambient})$$

Where,
$\dot{m}$ = volumetric flow rate of coolant in the heat exchanger
$T_{ambient}$ = ambient air temperature outside the cabinet

If we substitute all the above variables in the equation for junction temperature:

$$T_j = f(\dot{m}, V, T_{ambient}) + P \times (a + b \times V^{c})$$

The control box can manage the valves for both the inlet and outlet heat exchangers, the voltage to the fan tray for adjusting the speed and communicating with server using I^2c protocol, wireless communication or any other protocol. At any point in time, the control system should examine the value of T_j. If it is high, there are two choices: either increase the mass flow rate through the heat exchanger to bring the T_a down, or increase the fan RPM to increase the V and subsequently decrease the T_a. What is the implication on the exit air temperature and global data center temperature? Is it more economical to change the flow in the heat exchanger or the RPM of the fans?

The adaptable data center not only has to answer the above questions, but it has to be scalable to address any increase in load or addition of servers with minimal engineering work.

Reference:

1. Chandrakant, D., Cullen, E., Sharma R., Beitelmal, A. and Malon, C., Smart Chip, System and Data Center Enabled by Advanced Flexible Cooling Resources, IEEE, 2005.

2. Madara, S., Adaptive Cooling for the Next Generation of IT Systems, Sustainablefacility.com.

3.Dell.com/downloads/global/power/ps1q07-20070210-CoverStory.pdf