EE249 Lab 2: Data Center Temperature Control (Part 2)

Xuening Sun
xuening@eec.berkeley.edu

September 17, 2009

Last week, we built a simple continuous-time model of the server room and its temperature control. This week, we continue with this example.

Tasks: (Work in groups of 3)

1. **Imagine now that you have 4 server racks, each with a temperature monitor attached to it. The goal now is keep all for monitors within the 20~25˚C range. Create a model to account for this, and find a good position of the server racks to satisfy the temperature constraint. How did the model or control change? (i.e. how scalable is your model)**
2. **Lets attach a power consumption number for the temperature control system, such that every time cold air is pumped into the room, it consumes 200kW of power, and when idle (i.e. no inflow air), it consumes 50kW of power. What is the hourly power consumption just for cooling the server room? How can you minimize this based on the model that you have created?**
3. **If you were to completely redesign the server room, e.g. room volume, number of servers, number of vents, position of vents, number of sensors, etc., what would you suggest? How do you model it?**

Addendum:

A. To simplify the problem, we can assume that the effect of the heat generated by each server rack on the rest of the room is negligible.

B. Each server rack should generate different heat due to different workloads. We'll use: **q1=1000W, q2=800W, q3=1500W, and q4=500W**.

C. We can assume the inflow vent to be located in the exact center of the room (at top), and the outflow vent to be in the exact center of the room (at bottom). We can also view them as just single point in which air enters and exits the room.

D. It maybe helpful to look at the problem in 2D, where server racks and vents are views as just points in the 2D space. We are looking for the best locations for the server racks, with respect to the vents, and a way to see what are the temperatures read from each sensor at any given time.

E. You should assume a simple linear relationship between the temperature change at any point in the room and the distance of that point from the inflow vent. Again, you can assume the distance is just with respect to the 2D plane.

For example: $\Delta T = (1/d) * f(\Delta t)$, where $f(\Delta t)$ is derived from the equations we outlined in Lab 1, and d is the distance away from the vents. Notice that now your wall models would also change.

F. In this lab, you have the freedom to change the inflow air pressure as you wish. So the controller will probably no longer be just a simple bang-bang controller.

G. **For the ambitious ones:** you can try to see what it'll be like if the negligible server heat assumption is removed. Also, typically, there is a minimum and maximum distance of separation constraint between the servers due to maintenance or connectivity needs, so it may be interesting to explore how that would change the model behavior.