ThermoCast: A Cyber-Physical Forecasting Model for Data Centers

Carnegie Mellon

Lei Li, Chieh-Jan Mike Liang, Jie Liu, Suman Nath, Andreas Terzis, Christos Faloutsos

JOHNS HOPKINS UNIVERSITY

Microsoft Research

Motivation

- US Data centers consume 12GW power (=$7.4 Billion)
- Traditional data centers are over provisioned, with ≈40% of energy spent for cooling

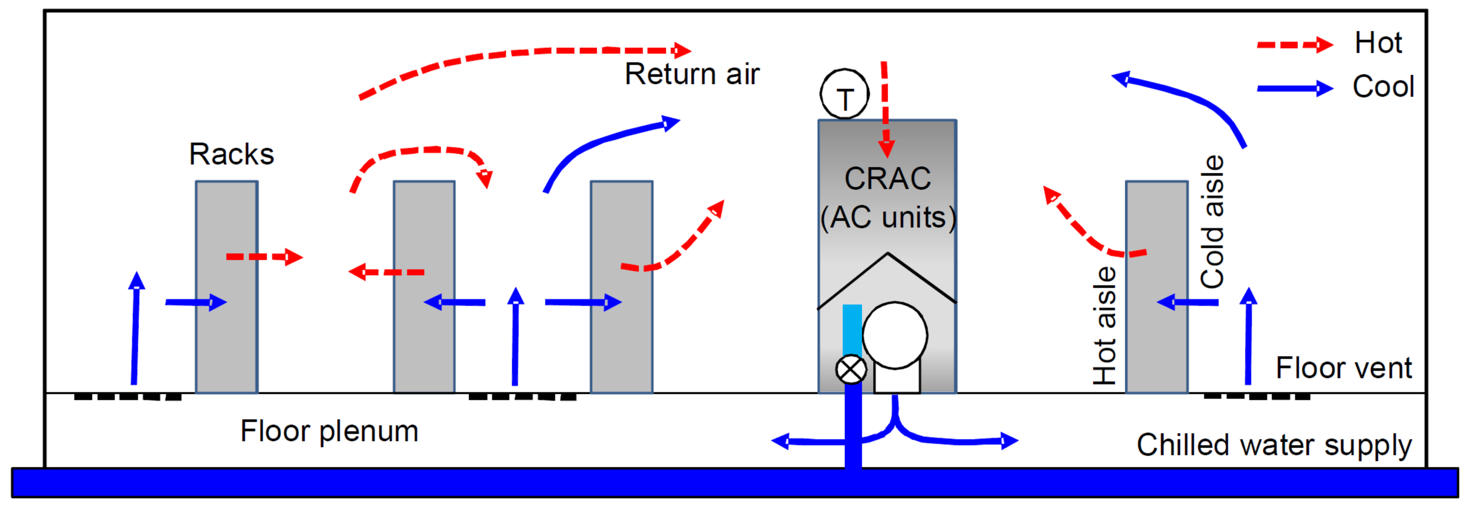

Airflow in a typical data center

Reactive energy saving:

- slow down cooling fan in CRAC
- raise AC temperature set points

Proactive data center management:

- predicting temperature distribution and thermal aware placement of workload

Setup & Observation

A university DC with 171 1U server nodes (8 cores).

A network of 80 sensors placed to monitor intake/outtake temperatures, and air flow speed.

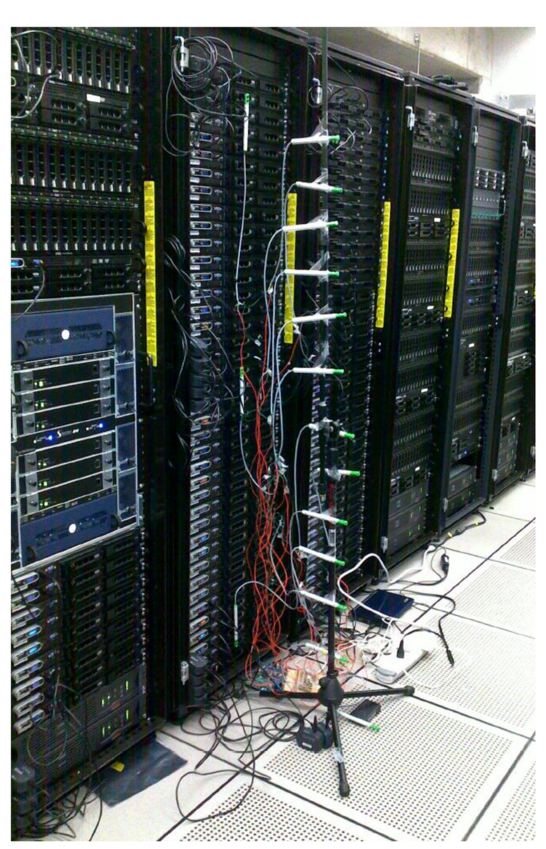

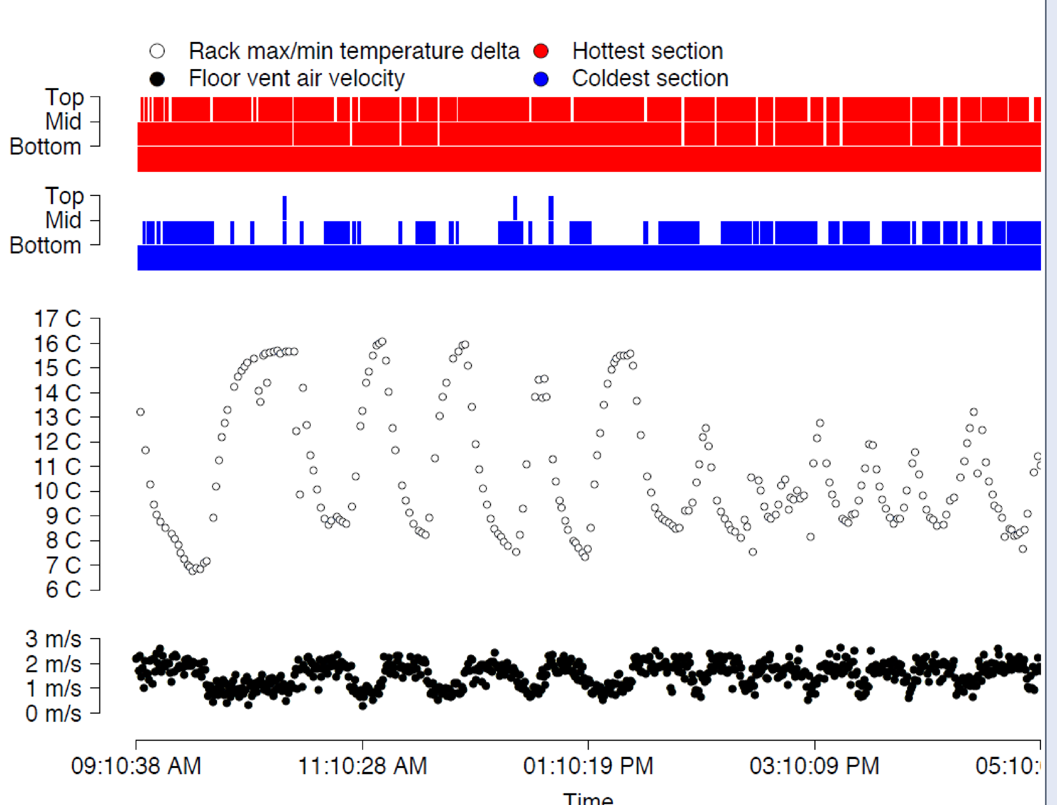

- Temperature difference cycle (max/min temp. on the same rack) is in antiphase with air velocity cycle.
- Middle and bottom sections are coldest; Top is hottest
- Shutting down under-utilized servers could reduce energy consumption.

Problem Definition

Given: intake temperatures, outtake temperatures, workload for each server , and floor air speed

To build a model that can forecast the intake temperature

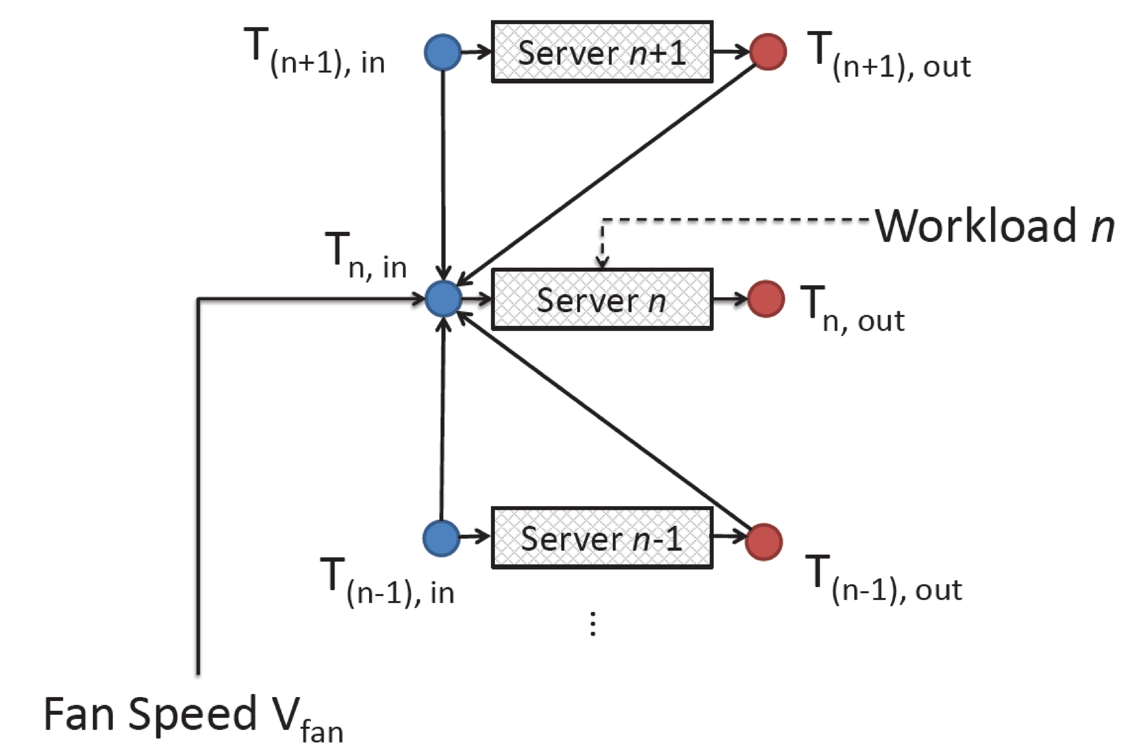

Proposed Approach: ThermoCast

Zonal model: divide the machine room into zones, and each rack into sections.

Assumptions:

A0: incompressible air

A1: environmental temperature is constant

A2: supply air temperature is constant within a period

A3: constant server fan speed

A4: vertical air flow at the outtake is negligible

A5: vertical air flow at the intake is linear to height

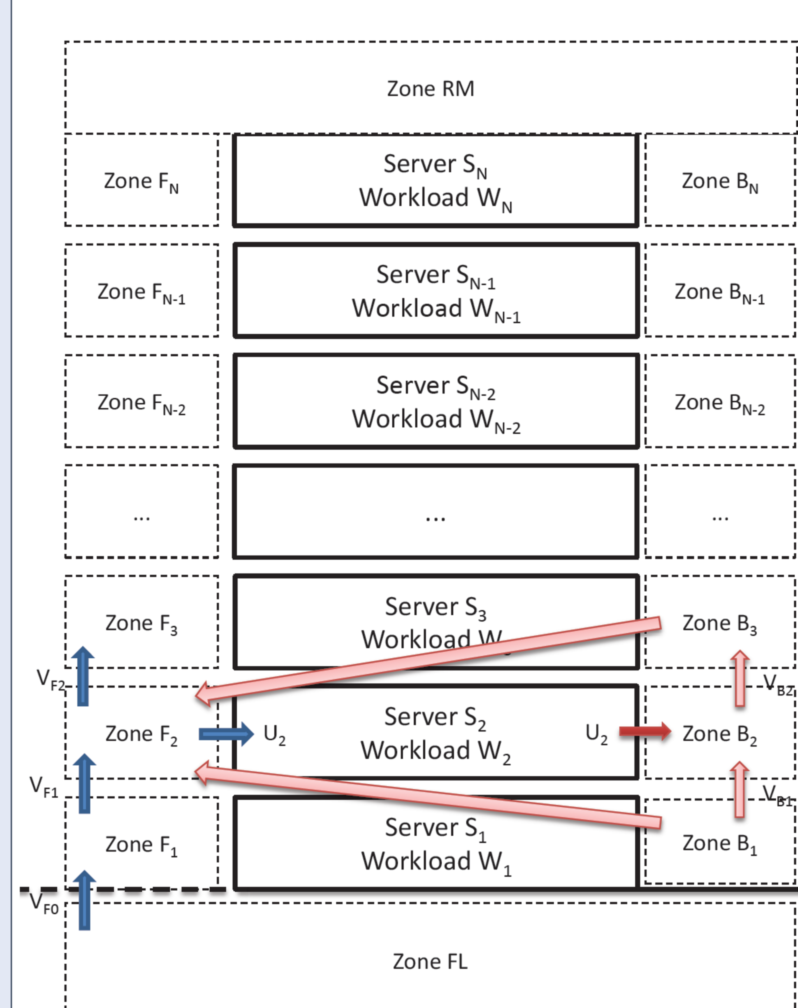

ThermoCast Model

$$
\begin{aligned}
T_{B_s}(t+1) &= f_1 \cdot T_{B_s}(t) + f_2 \cdot T_{F_s}(t) + \\
&\quad f_3 \cdot U_s(t) \cdot W_s(t) + f_4 \cdot T_{B_{s-1}}(t)
\end{aligned}
$$

$$
\begin{aligned}
T_z(t+1) = a \cdot T_{F_i}(t) &+ b_1 \cdot U_i(t) \cdot T_{B_i}(t) \\
&+ b_2 \cdot (1 - U_i(t)) \cdot T_{B_i}(t) + b_3 \cdot V_{FL}(t) \cdot T_{F_{i-1}}(t) \\
&+ b_4 \cdot V_{FL}(t) \cdot T_{F_{i+1}}(t)
\end{aligned}
$$

$$V_{FL}(t+1) = \eta_0 \cdot V_{FL}(t) + \eta_1 \cdot V_{FL}(t-1)$$

Parameter estimation

$$\hat{\theta}^{(i)} \leftarrow \arg\min f_\lambda(\theta^{(i)}) = \sum_{t=1}^{t_{max}-1} \exp(\lambda t) g(\theta^{(i)}, t)$$

s.t.

$$
\begin{aligned}
g(\theta^{(i)}, t) = \big(T_{F_i}(t+1) &- a \cdot T_{F_i}(t) - b_1 \cdot U_i(t) \cdot T_{B_i}(t) \\
&- b_2 \cdot (1 - U_i(t)) \cdot T_{B_i}(t) - b_3 \cdot V_{FL}(t) \cdot T_{F_{i-1}}(t) \\
&- b_4 \cdot V_{FL}(t) \cdot T_{F_{i+1}}(t)\big)^2 \\
&+ \big(T_{B_i}(t+1) - f_1 \cdot T_{B_i}(t) - f_2 \cdot T_{F_i}(t) \\
&- f_3 \cdot U_i(t) \cdot W_i(t) - f_4 \cdot T_{B_{i-1}}(t)\big)^2
\end{aligned}
$$

Experiments and Results

We collected real workload and sensor monitoring data at a university data center.

Q1: How accurately can a server learn its local thermal dynamics for prediction?

2x better

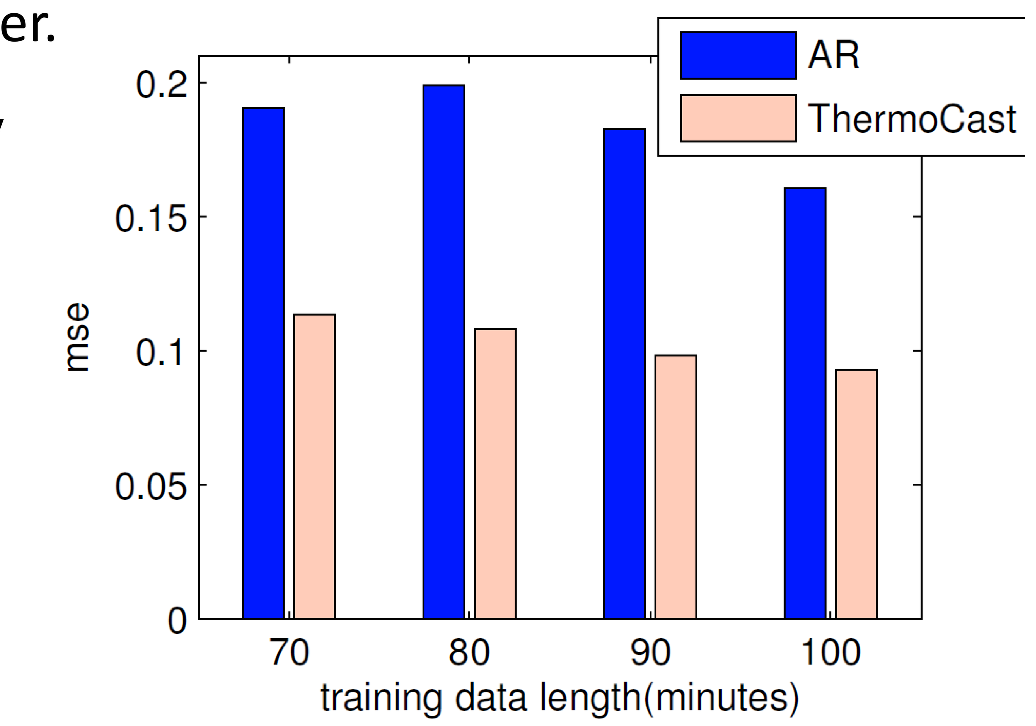

Q2: How long ahead can ThermoCast forecast thermal alarms?

2x faster

	Baseline	ThermoCast
Recall	62.8%	71.4%
FAR	45%	43.1%
MAT	2.3min	**4.2 min**

MAT=mean look-ahead time

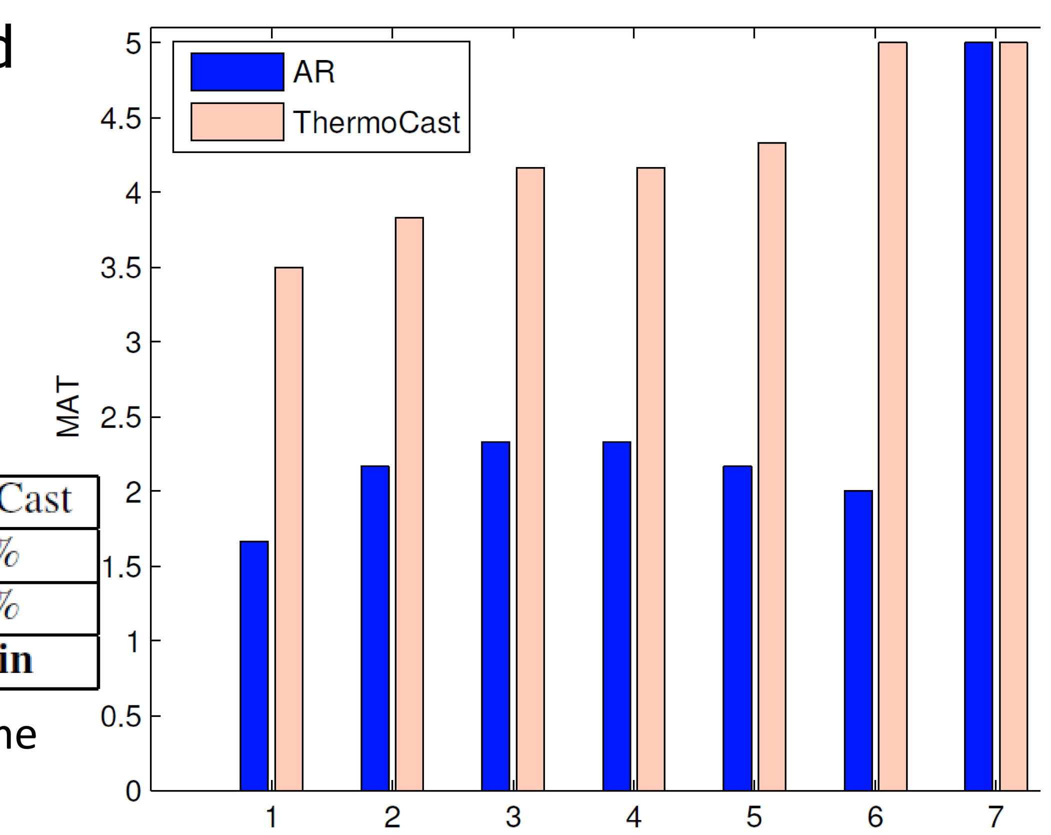

Prediction using 90 minutes of data trace; all predictions are made at 5 minutes away from training.

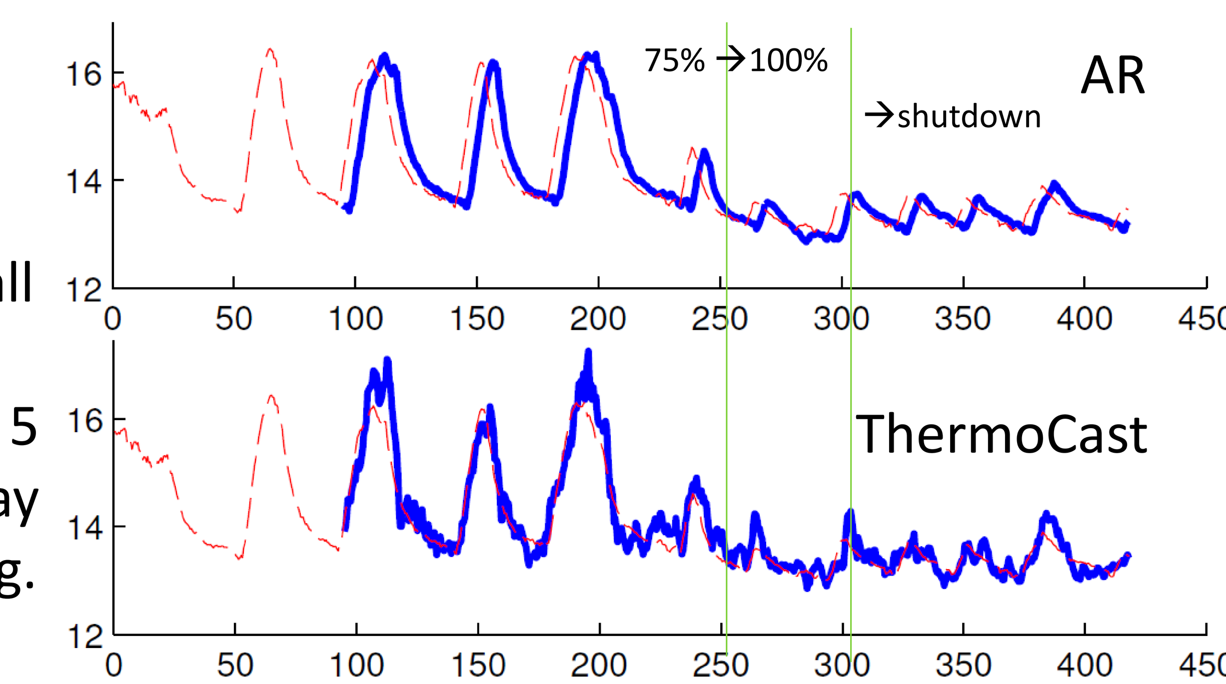

Contribution & Conclusion

- We provide a systematic approach to integrate the physical laws and sensor monitoring in a data center.
- We provide an algorithm to learn from sensor data for such cyber-physical system, reducing full fluid models.
- We instrument in a practical data center and evaluate the models against real workload and measured data.

KDD, 2011