Data Mining Issues for the Thermal Management of Data Centers

Myong K.(MK) Jeong[1,2], W. Art Chaovalitwongse[2], Hoang Pham[2] , and Norman Kim[1]

[1]RUTCOR (Rutgers Center for Operations Research)

[2]Department of Industrial and Systems Engineering

Thermal Management of a DC

- To provide ideal thermal environment to lower energy costs, reduce system failure rates, and consequently, optimize computing resources and maximize business outcome (Moore et al. 2006)

Thermal Mapping of Data Centers for Uniform Workload Assignment

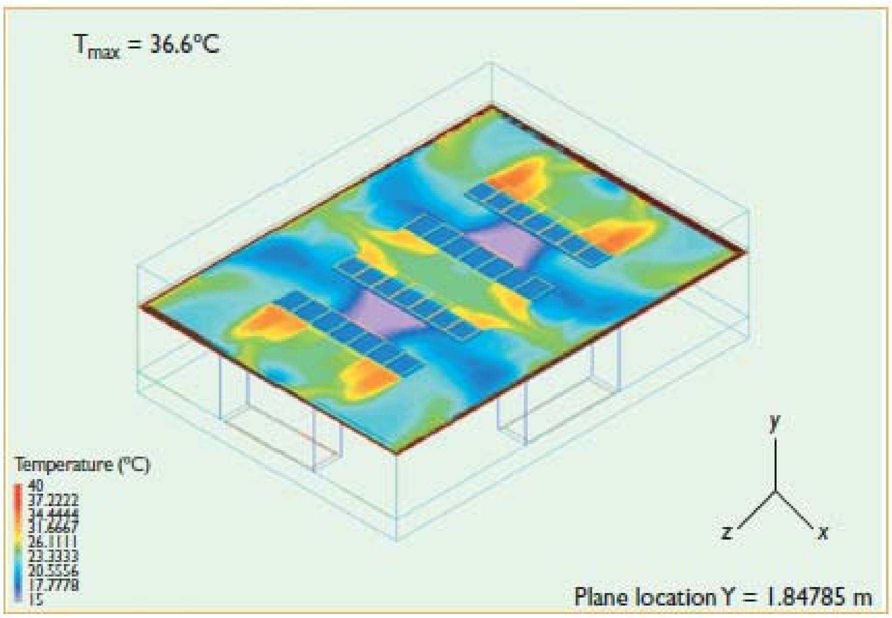

- Figure: Simulated temperature contour plot for a programmable data center where uniformly loaded servers dissipating heat at 75% of their rated capacity (252kW) using computational fluid dynamics model (source: Sharma et al. 2005)

Goals of Our Research

- Development of advanced data mining techniques for on-line thermal prediction of data centers

- Eventually, promote uniform temperature distribution of data centers, e.g., by developing workload placement model considering the thermal mapping of data centers

- This may result in
  - Decreasing cooling costs of a DC
  - Increasing the reliability and performance of hardware equipments
  - Improving operational efficiencies of a DC

Thermal-Aware Workload Placement (Tang 2008)

- A workload placement policies attempt to maximize cooling efficiency by allowing us to lower inlet temperature of the racks.

- Workload Placement Algorithm

$$\begin{aligned} &minimize\ (\max_i T_{in}^i) \\ &subject\ to \sum_{j=1}^{q} c_{ji} \le m_i, \forall i = 1,\dots,n \\ &\qquad\qquad \sum_{i=1}^{n} c_{ji} = c(j), \forall j = 1,\dots,q \end{aligned}$$

where

T_{in} : *inlet temerature of air into rack* i

c_{ji} : *amount of task* j *assigned to rack* i

$c(j)$: *total amount of task* j

m_i : *capacity of rack* i

n : *# of racks*

q : *# of tasks*

Workload Placement Models

- Workload Placement Algorithm that maximizes the uniformity of the temperatures of a DC (Jeong et al. 2009)

$$\begin{aligned} &\textit{minimize } (\mathrm{var}(T_{in})) \\ &\textit{subject to} \sum_{j=1}^{q} c_{ji} \le m_i, \forall i = 1,\ldots,n \\ &\qquad\qquad \sum_{i=1}^{n} c_{ji} = c(j), \forall j = 1,\ldots,q \end{aligned}$$

where

T_{in} : *inlet temerature of air into racks*

We need to develop the model for the on-line prediction of thermal mapping of a DC.

Development of data mining models for the prediction of thermal mapping

- Models (Gupta et al. 2008;Moore et al. 2006)

$$\mathbf{Y} = \boldsymbol{\alpha}\mathbf{W} + \boldsymbol{\beta}\mathbf{C} + \boldsymbol{\gamma}\mathbf{P} + \varepsilon$$

$$\begin{aligned}
&where \\
&\mathbf{W} = [\mathbf{W}_{11}, \mathbf{W}_{12}, \cdots, \mathbf{W}_{1p}, \mathbf{W}_{21}, \mathbf{W}_{22}, \cdots, \mathbf{W}_{2p}, \cdots, \mathbf{W}_{n1}, \mathbf{W}_{n2}, \cdots, \mathbf{W}_{np}] \\
&\mathbf{Y} = [\mathbf{Y}_1, \mathbf{Y}_2, \cdots, \mathbf{Y}_n] \\
&cf. \\
&n = number\ of\ blocks \\
&p = number\ of\ workload\ distribution\ variables
\end{aligned}$$

- **Workload distribution (W)**: utilization of hardware units (e.g., CPU utilization, disk I/O rates, rotate speed, memory I/O rates, and network activity)
- **Cooling configuration (C) :** number of CRACs, air velocity, the temperature of the air CRACs supply, fan speeds of the servers, etc.
- **Physical topology (P) :** locations of server racks, doors, slotted floor tiles, etc.

Challenging Issues in On-line Thermal Mapping Prediction (1/3)

- Thermal mapping over time are auto-correlated
- Complex nonlinear relationship between the thermal map and input variables
- Proposed Lagged Dependent Variables (LDVs) Model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Y}\boldsymbol{\alpha} + \boldsymbol{\varepsilon},$$

$$\mathbf{y} = \begin{pmatrix} y_{p+1} \\ y_{p+2} \\ \vdots \\ y_T \end{pmatrix}, \mathbf{X} = \begin{pmatrix} x_{p+11} & x_{p+12} & \cdots & x_{p+1k} \\ x_{p+21} & x_{p+22} & \cdots & x_{p+2k} \\ \vdots & \vdots & \vdots & \vdots \\ x_{T1} & x_{T2} & \cdots & x_{Tk} \end{pmatrix}, \boldsymbol{\beta} = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_k \end{pmatrix}, \mathbf{Y} = \begin{pmatrix} y_p & y_{p-1} & \cdots & y_1 \\ y_{p+1} & y_p & \cdots & y_2 \\ \vdots & \vdots & \vdots & \vdots \\ y_{T-1} & y_{T-2} & \cdots & y_{T-p} \end{pmatrix}, \boldsymbol{\alpha} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_p \end{pmatrix}, \boldsymbol{\varepsilon} \sim N(0, \sigma^2)$$

Prediction Model Using KRR-LDVs

- Kernel Ridge Regression (KRR) Model Considering Lagged Dependent Variables (LDVs)

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Y}\boldsymbol{\alpha} + \boldsymbol{\varepsilon},$$

$$\hat{\boldsymbol{\beta}}_{\gamma} = \Phi(\mathbf{X})^{T}((\mathbf{I}-\mathbf{K}_2(\mathbf{K}_2+\gamma\mathbf{I})^{-1})\mathbf{K}_1+\lambda\mathbf{I})^{-1}(\mathbf{I}-\mathbf{K}_2(\mathbf{K}_2+\gamma\mathbf{I})^{-1})\mathbf{y}$$
$$\hat{\boldsymbol{\alpha}}_{\gamma} = \Phi(\mathbf{Y})^{T}(\mathbf{K}_2+\gamma\mathbf{I})^{-1}(y-\Phi(\mathbf{X})\hat{\boldsymbol{\beta}}_{\gamma})$$

where

$$\mathbf{K}_1 = \Phi(\mathbf{X})\Phi(\mathbf{X})^{T} \textit{ and } \mathbf{K}_2 = \Phi(\mathbf{Y})\Phi(\mathbf{Y})^{T}$$

Challenging Issues in On-line Thermal Mapping Prediction (2/3)

- Huge number of input variables and small number of available samples (small n large p problem)

  Ex: (K1 servers) x (K2 number of W variables) + K3 Cooling configuration variables + K4 physical topology variables > 10,000 when K1=1000, K2=10, …

  - Overfitted prediction model can be produced and eventually produce the unstable prediction model

  => Dimensionality reduction and feature selection techniques should be developed

Feature Selection for TM

- Feature Selection
  - the technique of selecting a subset of relevant features for building robust learning models.
  - Helps alleviating the effect of the curse of dimensionality.
  - Enhances generalization.

- Most of existing works are developed for variable selection for single output variable or a few (2 or 3) output variables

- Development of new feature selection techniques when there are huge number of output variables (= # of servers/block size)

Challenging Issues in On-line Thermal Mapping Prediction (3/3)

- Robust Support Vector Regression (r-SVR)
  - Current N/N-based thermal mapping model could be very sensitive to outliers.
  - Also, popular SVR is not effective in dealing with outliers in thermal applications. So, few outliers result in a poor prediction performance.

  => Robust regression model with the robustness against outliers is required.
- Online Updating Model
  - When new observations are available, we need to update our current model in a fast way.
- Spatial Regression Model
  - Temperatures of servers are spatially correlated.

Conclusion

- Develop a new on-line prediction model for thermal map.
- Several challenging issues of data mining for the on-line prediction of thermal mapping are proposed.
- We can formulate different workload assignment algorithms considering the thermal mapping and compare their performance.

References

1. Filani, D., He, J., Gao, S., Rajappa, M, Kumar, A., Shah, R. and Nagappan, R, "Dynamic Data Center Power Management: Trends, Issues, and Solutions", white paper, Intel, 2008.
2. HP Labs, “Making Scheduling Cool: Temperature-Aware Workload Placement in Data Centers”.
3. Moore, J., Chase, J., Ranganathan, P. and Sharma, R., "Making scheduling "cool": temperature-aware workload placement in data centers", in Proc. of the annual conference on USENIX Annual Technical, 2005.
4. Moore, J., Chase, J. and Ranganathan, P., "ConSil: Low-Cost Thermal Mapping of Data Centers", In proc of SysML, 2006.
5. Moore, J., Chase, J. S. and Ranganathan, P., “Weatherman: Automated, Online, and Predictive Thermal Mapping and Management for Data Centers”, 2006 IEEE International Conference on Autonomic Computing, pp.155-164, 2006.
6. Ramos, L. and Bianchini, R., "C-Oracle: Predictive thermal management for data centers", in Proc. of IEEE 14th International Symposium on High Performance Computer Architecture, pp.111-122, 2008.
7. Sharma, R. K., Bash, C. E., Patel, C. D., Friedrich, R. J. and Chase, J. S., "Balance of Power: Dynamic Thermal Management for Internet Data Centers", IEEE Internet Computing archive, 9(1), pp.42-49, 2005.
8. Tang, Q., “Thermal-aware Scheduling in Environmentally Coupled Cyber-Physical Distributed Systems, Dissertation, Arizona State Univ, 2008.
9. Tang, Q., Gupta, S. K. S. and Varsamopoulos, G, "Energy-Efficient, Thermal-Aware Task Scheduling for Homogeneous, High Performance Computing Data Centers: A Cyber-Physical Approach", Transactions on Parallel and Distributed Systems, Special Issue on Power-Aware Parallel and Distributed Systems (TPDS PAPADS), 19(11), pp.1458–1472, 2008.
10. Tang, Q., Gupta, S. K. S. and Varsamopoulos, G, "Thermal Aware Task Scheduling for Datacenters through Minimizing Heat Recirculation", in Proc. of Cluster 2007, Austin, USA, 2007.