AUBURN

UNIVERSITY

SAMUEL GINN

COLLEGE OF ENGINEERING

TIGER : Thermal-Aware File Assignment in Storage Centers

Ajit Chavan

Outline

- Motivation
- Models
- TIGER
- Experimental Results
- Conclusion
- Future Work

Motivation

Power and Cooling Expenses as a Percentage of New Server spend

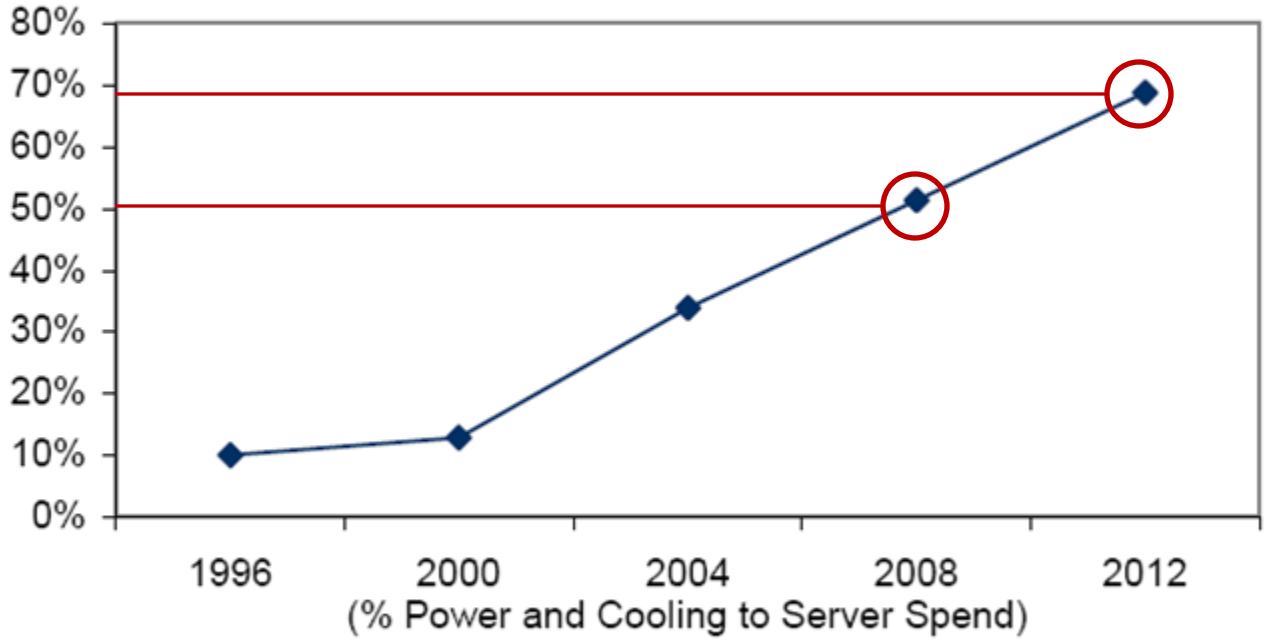

Source : IDC, 2009

Motivation (Contd.)

- Power Distribution:

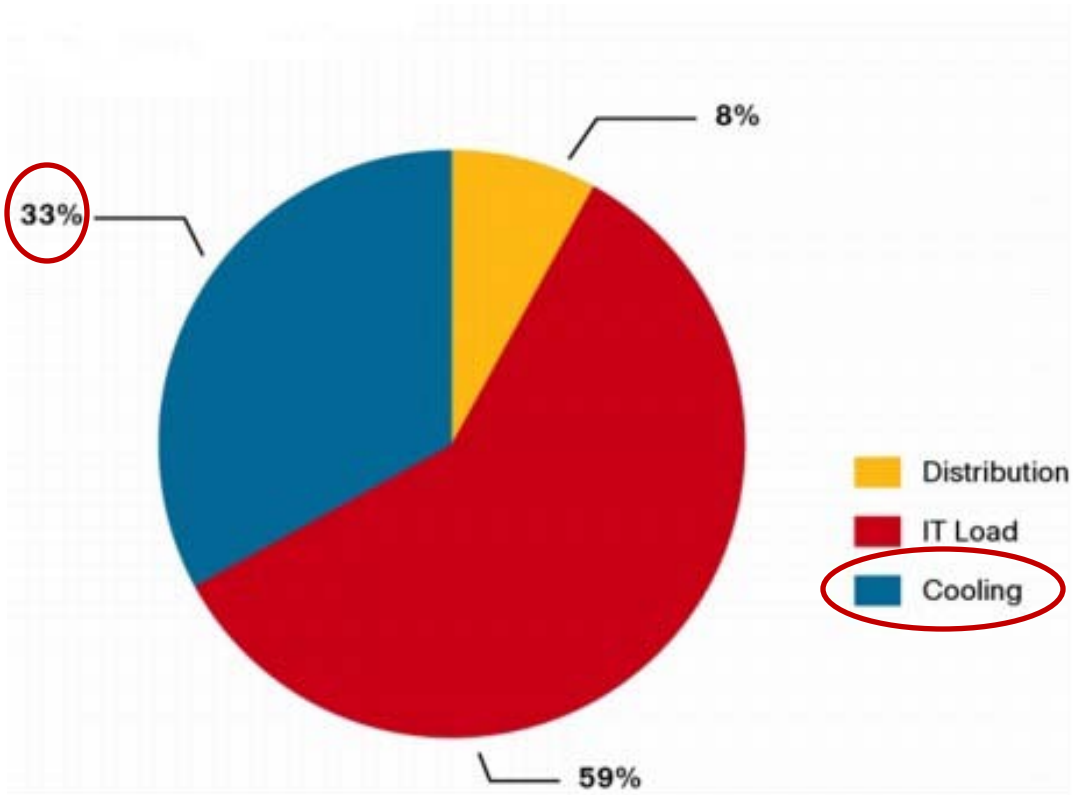

Source: Power Management in the CISCO Unified Computing System.

Models

- Heat Recirculation Model
  - $A_{N\times N}$: Cross-interference matrix
  - $\alpha_{i,j}$: Cross-interference coefficient

Diagonal matrix $K_i = \rho a_i c_p$

$$t_{in} = t_{sup} + [(K - A^T K)^{-1} - K^{-1}]p$$

Inlet temperature (T_i^{in}) vector

Supply temperature (T_i^{sup}) vector

Power (P_i^{Node}) vector

Models

- Cooling Cost Model

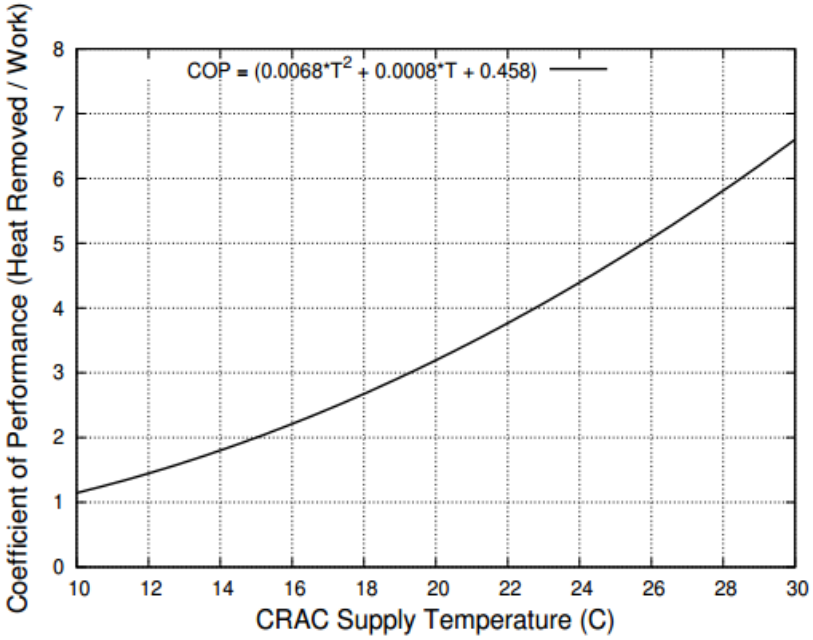

$$COP(T_{sup}) = 0.0068\,T_{sup}^2 + 0.0008\,T_{sup} + 0.458$$

TIGER

- Problem statement

Assigning *m* files to *D* disks residing in *N* nodes in order *to minimize the cooling cost* of the storage clusters by reducing heat recirculation in the data center, where file attributes like service time and arrival rates are known a priory.

TIGER

- Sort files according to service time

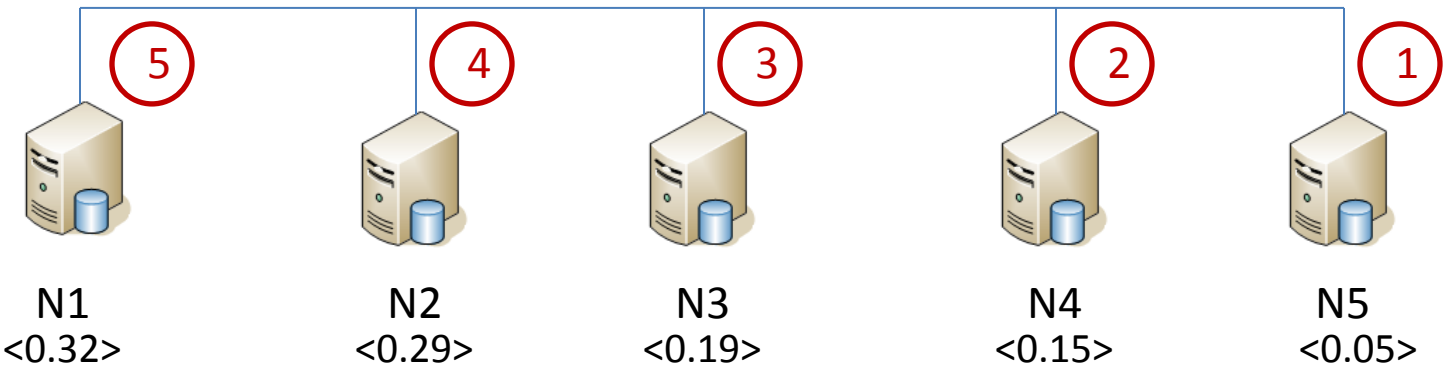

f1 <20,5> f2 <30,4> f3 <15,3> f4 <40,5> f5 <120,2> f6 <20,4> f7 <10,5> f8 <25,4> f9 <50,4> f10 <55,3>

$$U_{avg}^{Th} = \frac{1}{D}\sum_{k=1}^{m} u_k$$

f≤s,λ>,
s = service time in ms
λ= arrival rate in requests/s
u = s× λ (workload of file)

$$S_{avg} = \frac{1}{N} = 0.20$$

TIGER

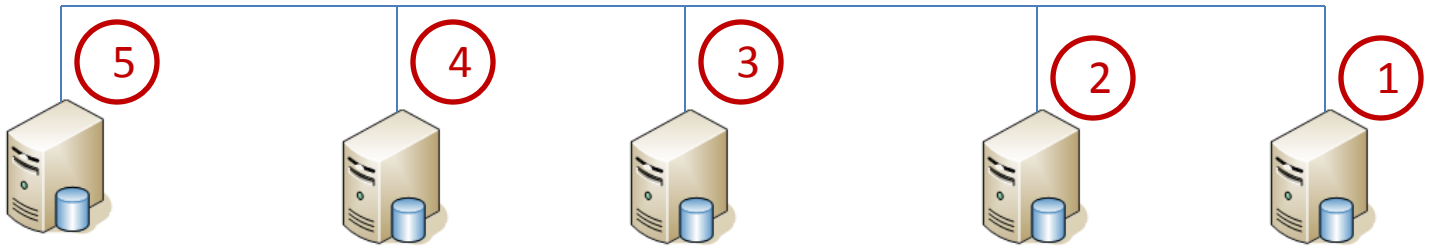

N1 <0.32> N2 <0.29> N3 <0.19> N4 <0.15> N5 <0.05>

f5	f10	f9	f4	f2	f8	f1	f6	f3	f7
u = 0.24	0.17	0.20	0.20	0.12	0.10	0.10	0.08	0.05	0.05

U_{avg}^{Th} = 0.27 S_{avg} = 0.20 U_2^{Th} = 0.34

Experimental Results

- Test bed

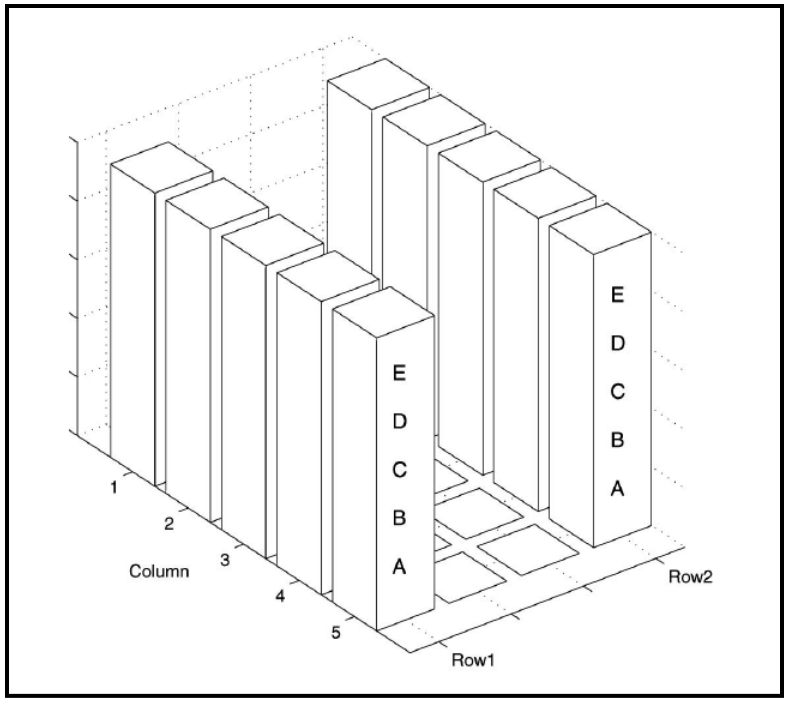

Number of nodes *(N)*	50
Number of disks in each node	24
Total number of disks (*D*)	1200

- Zipf distribution of file service time and arrival rate.

Experimental Results

Case 1: Energy conservation algorithm (like PRE-BUD) is used to turn down the idle disks.

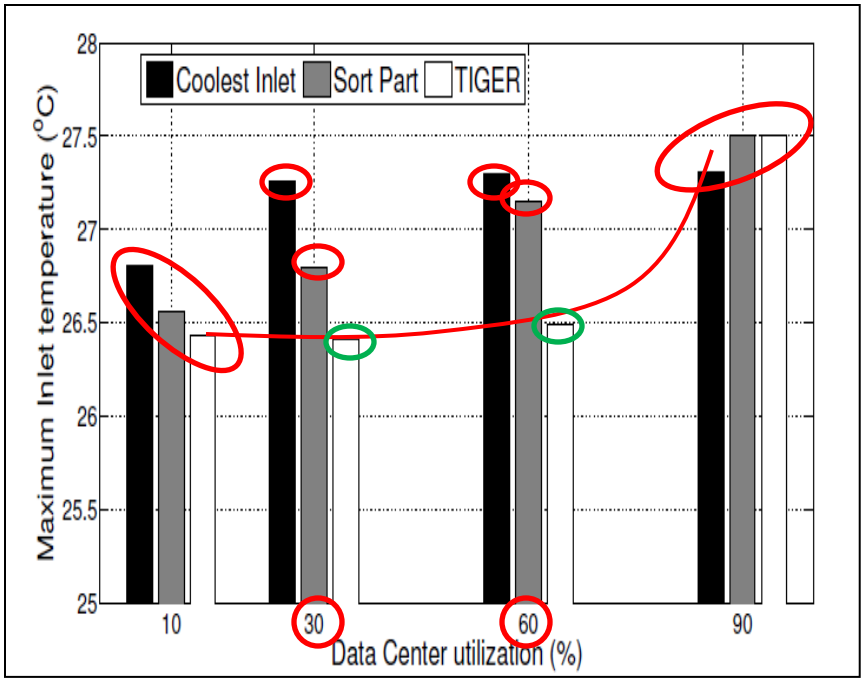

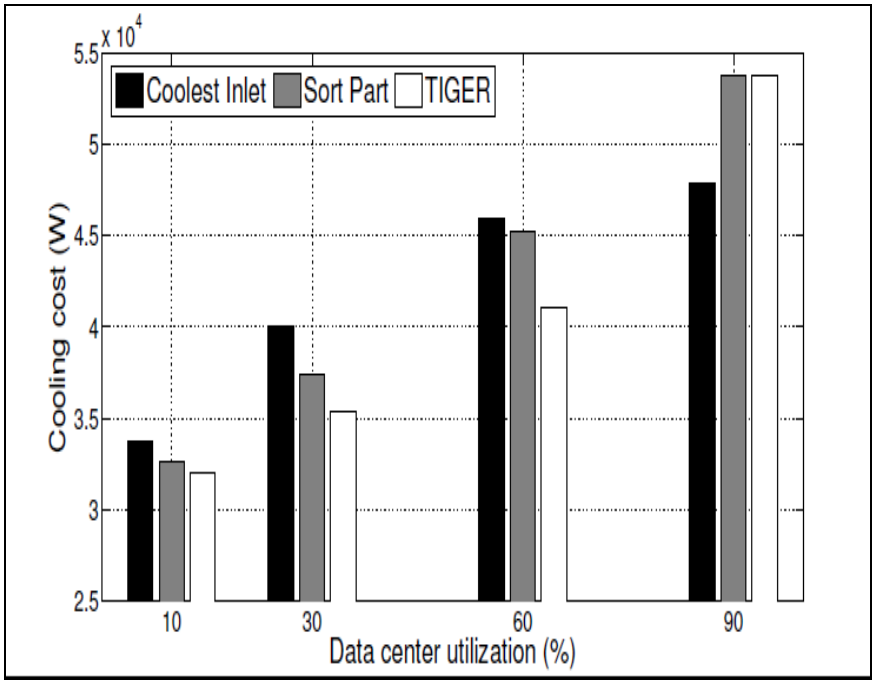

Conclusion

- We proposed file assignment algorithm to reduce cooling cost of the data centers.
- We use cross-interference matrix to characterize the heat recirculation in data center.
- TIGER calculates threshold on disk utilization based on the contribution of the node in the heat recirculation.
- Based on the threshold, TIGER assigns the files to the disks.
- TIGER offers about 10 to 15 percent cooling energy saving.

THANK YOU !

QUESTIONS ?