Keeping the Clouds Cool and Energy Efficient

Yogendra Joshi
Consortium for Energy Efficient Thermal Management (CEETHERM)
G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A.

Outline

- The energy and thermal management challenges in data centers
- Ways to address these for an existing facility
  - Dynamic load-re-allocation
- Ways to address these for a new facility
  - Air-side economizers
  - The S-Pod cabinet layout
- Conclusions

Infrastructure for Cloud Computing: Data Centers

- Ranging in size up to 11 acres
- Data processing equipment stacked vertically in 2-m tall enclosures
- In 1990 a typically rack dissipated ~ 1 kW, while today's racks can dissipate up to 30 kW, with no change in size
- Facility net power dissipation can be on the order of several MW
- In future the *data center will be the computer*

National Energy Research Scientific Computing Center

Data Center Cooling Requirements

- Much higher and spatially localized heat loads (over 200 W/ft^2 in some cases)
- Tighter relative humidity requirements (recommended 40%-50% during operation)
- Recommended air inlet dry bulb temperatures 20 ^oC – 25 ^oC during operation
- Very high reliability of cooling systems and power availability (99.9999% or 32 s/yr)

Chips are more demanding than people

Multiscale Nature of Data Center Thermal Problems

- Macro-scale covers $O(10^5)$ length scales

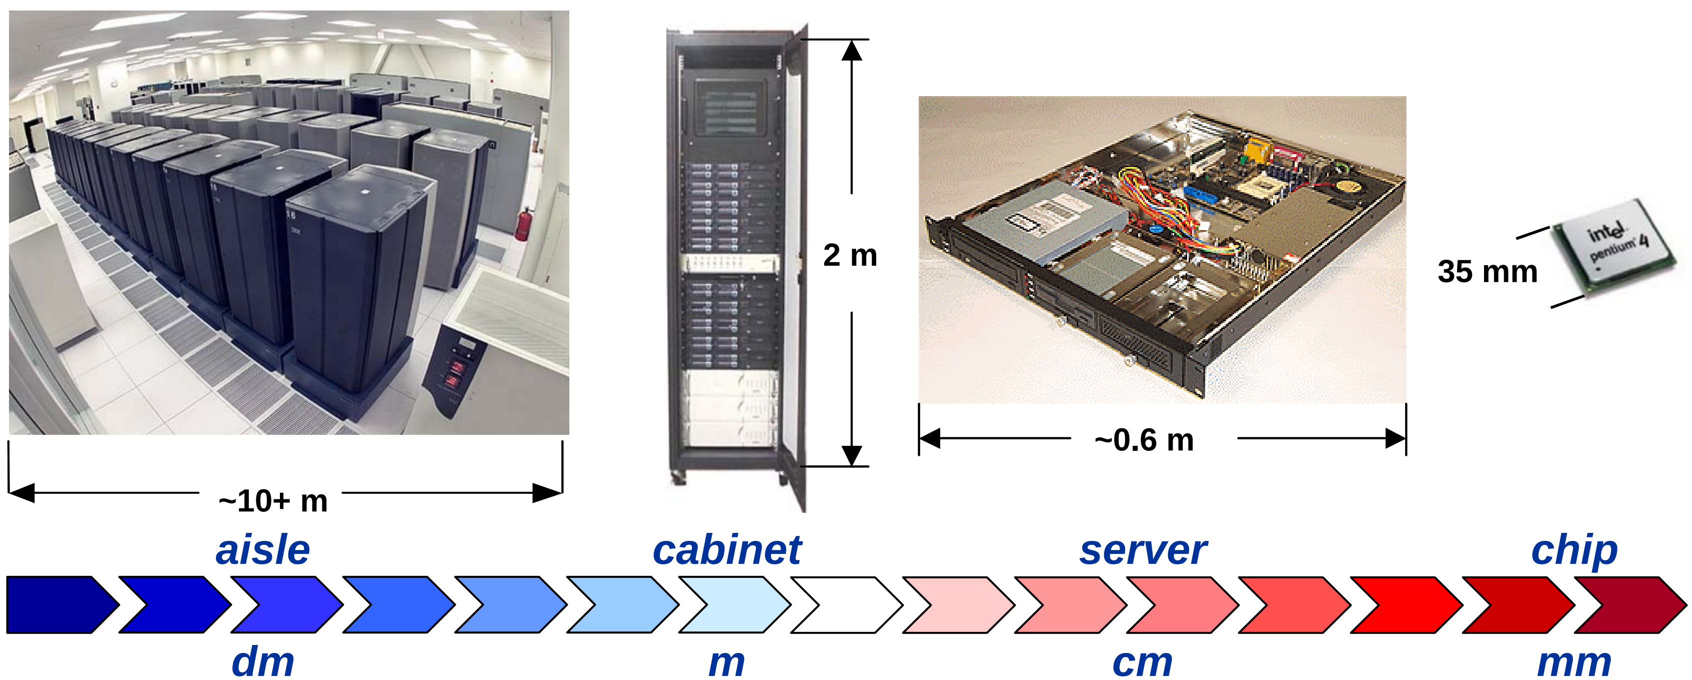

GT Research on Ultra-Compact High Performance Chip Cooling Devices

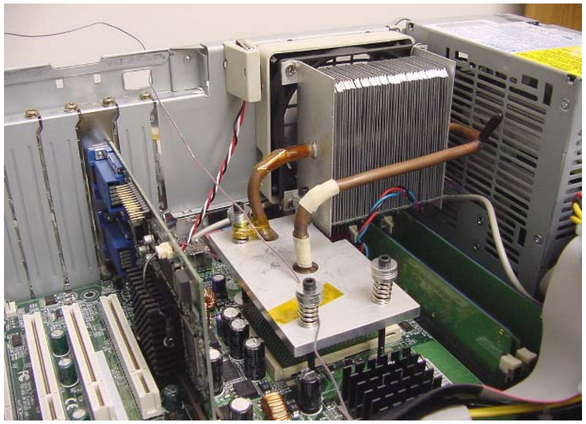

Compact Thermosyphon

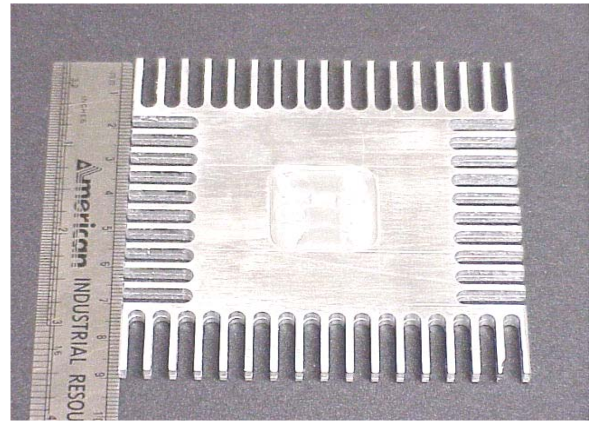

Two Phase Heat Spreader

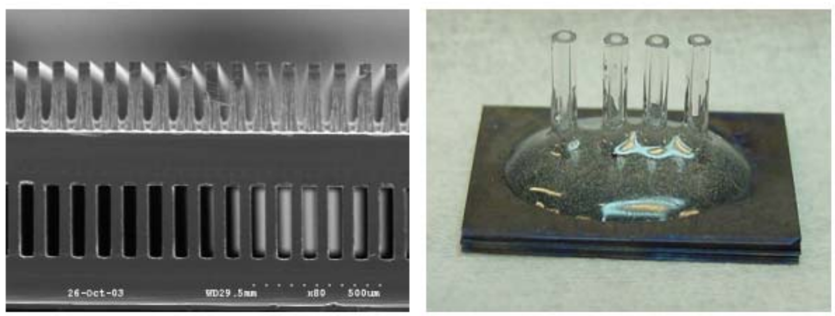

Stacked Microchannel Heat Sink

Electricity Flows in Data Centers

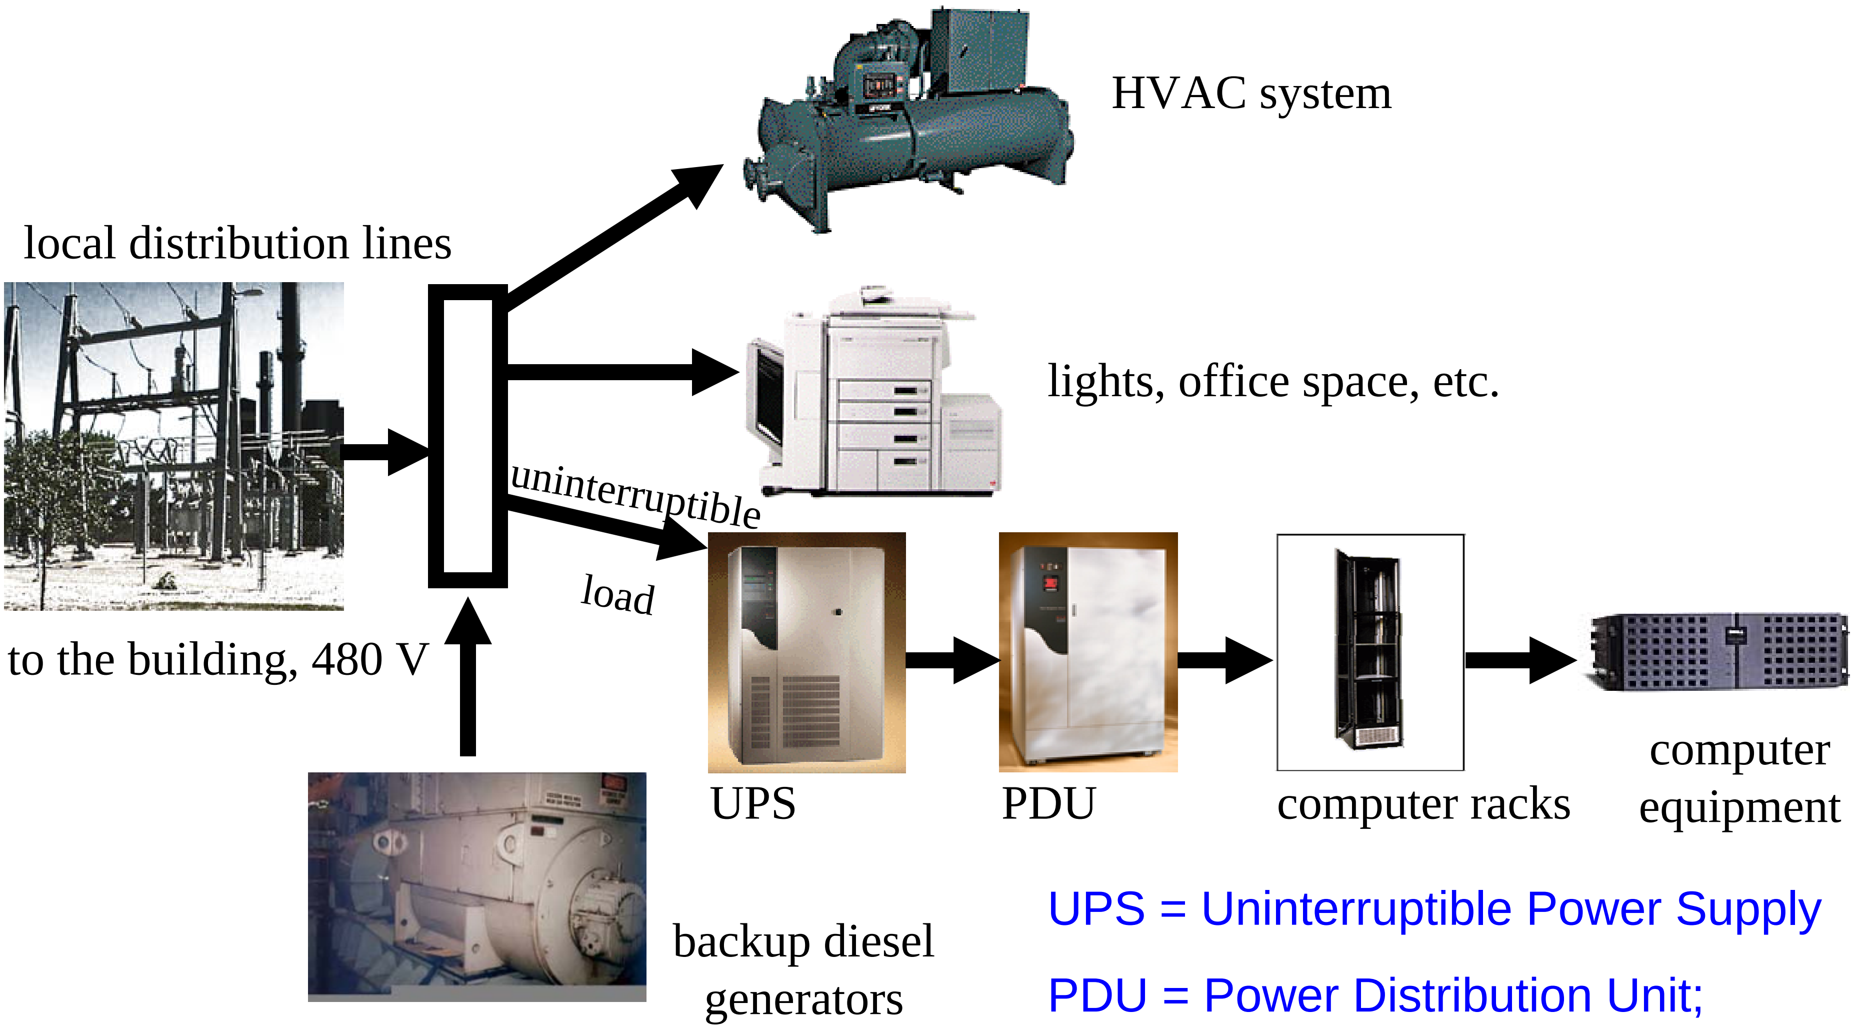

UPS = Uninterruptible Power Supply

PDU = Power Distribution Unit;

Source: W. Tschudi, Lawrence Berkeley Laboratories

Trend of Projected Electricity Use

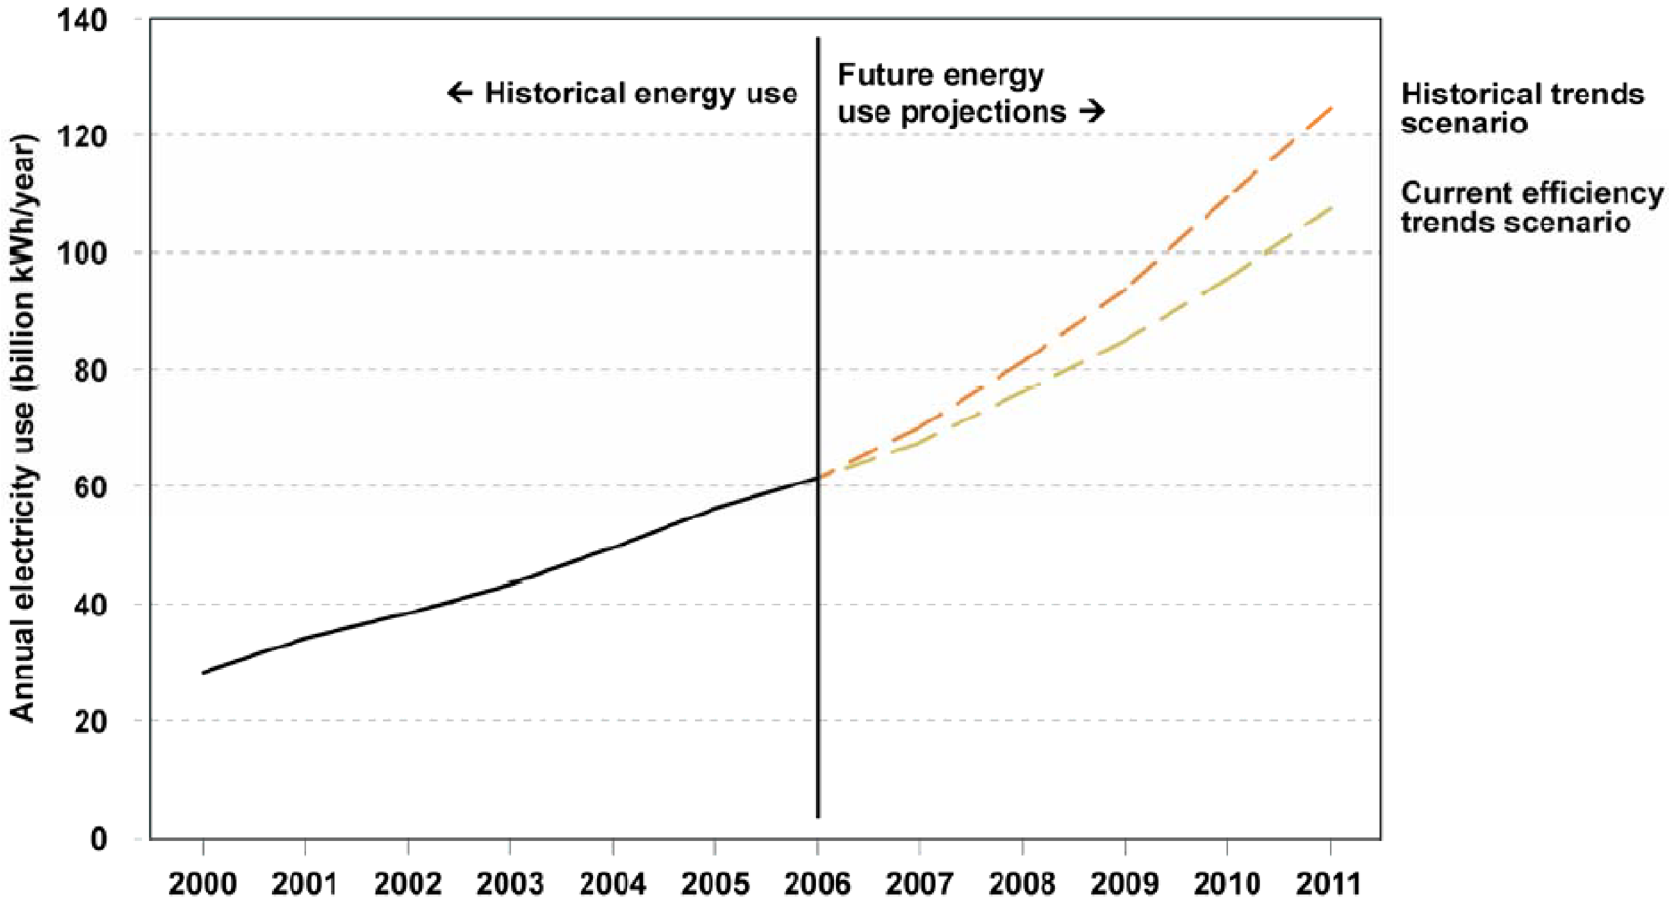

- ***Energy efficient design slows down the increasing historical trend of electricity use in data centers***
- ***How can we have a more energy efficient design in data centers?***

Source: US EPA Report to Congress on Server and Data Center Energy Efficiency; Public Law 109-431, pp. 52, August 2007.

Energy Consumption for HVAC

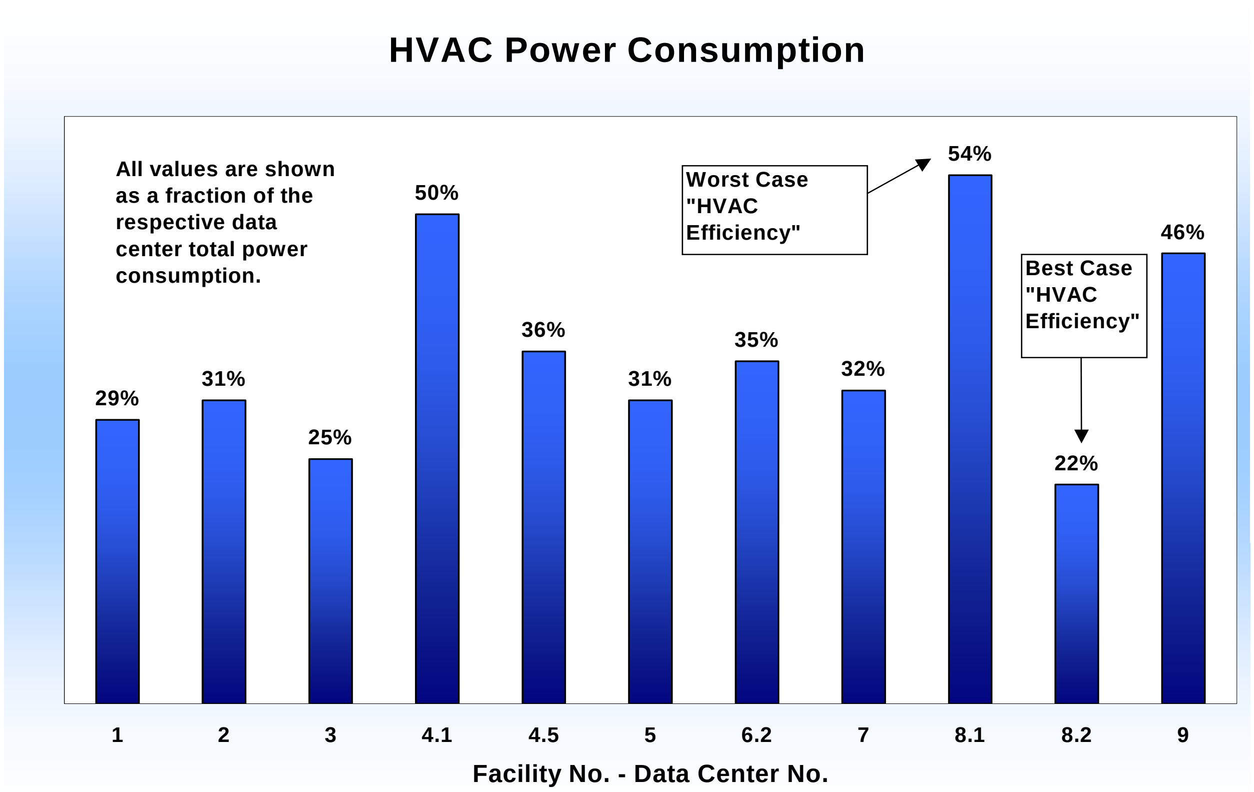

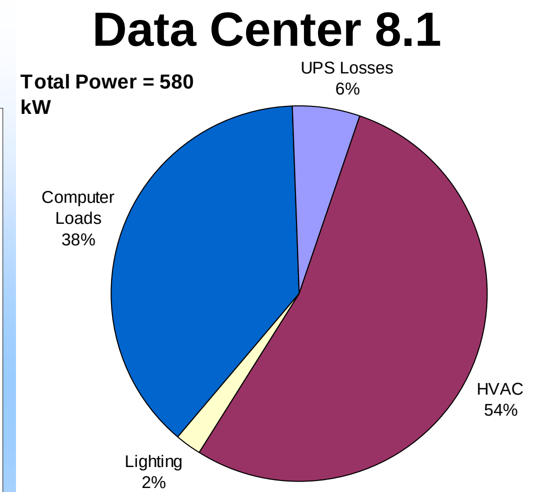

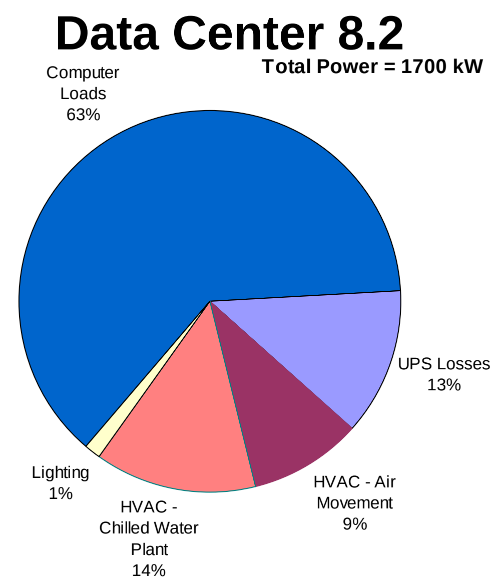

Message: Energy efficiency is key to reducing operation costs and saving environment

Source: W. Tschudi, Lawrence Berkeley Laboratories

CEETHERM Data Center Lab

- Moveable perforated tiles in raised floor plenum
- Flexible ductwork and movable ceiling grilles overhead
- Possible to backfill plenum to vary depth
- Chilled water taps in plenum for liquid cooled racks

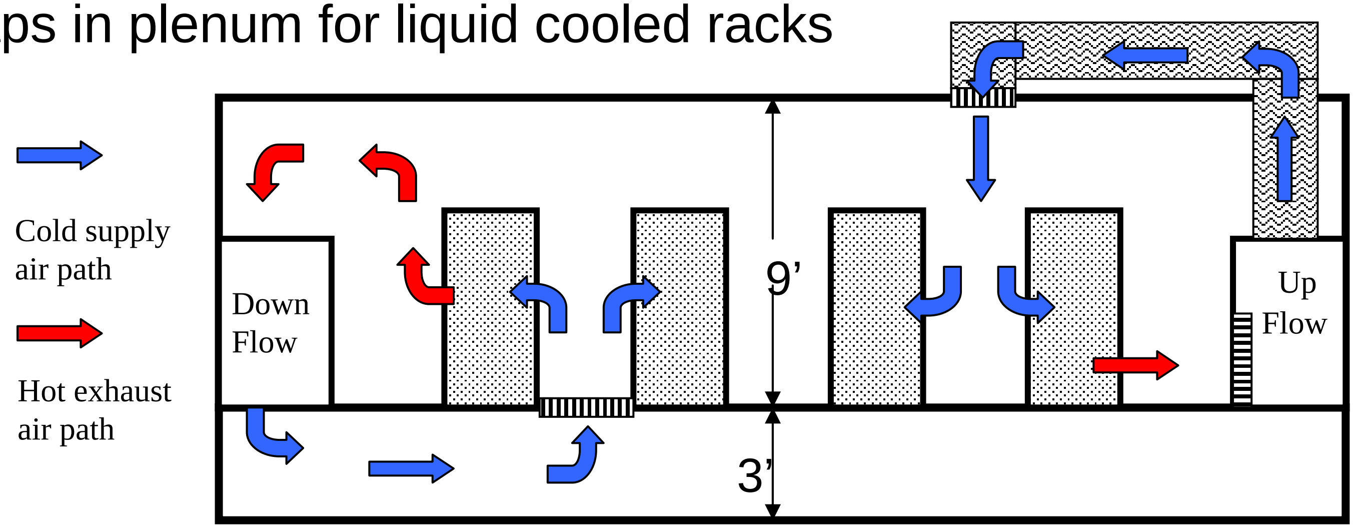

Model Based Load Reallocation

- 5 control variables: inlet velocity of CRAC unit, Vin; heat loads of Rack A1&B1, QRack1; Rack A2&B2, QRack2; Rack A3&B3, QRack3; Rack A4&B4, QRack4
- CRAC inlet velocity from 1 m/s - 3.2 m/s
- Rack heat load range: 10 kW to 20 kW
- Total center heat load: 120 kW
- Heat load allocation between 8 racks in the most efficient way

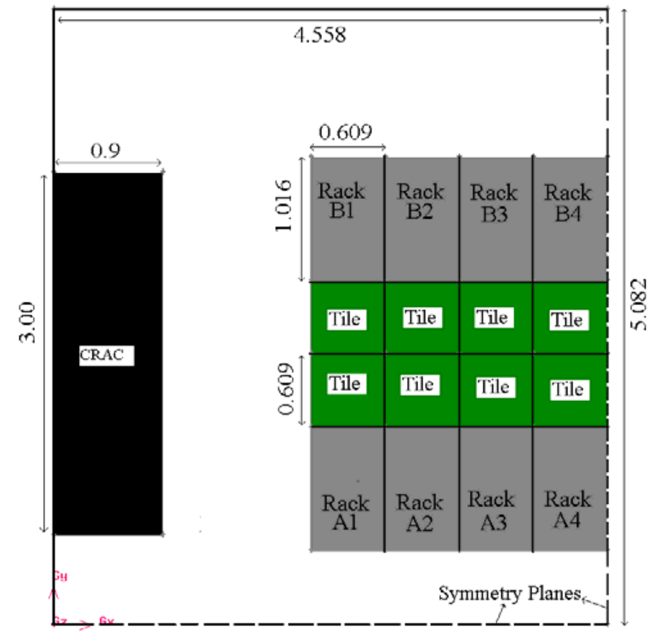

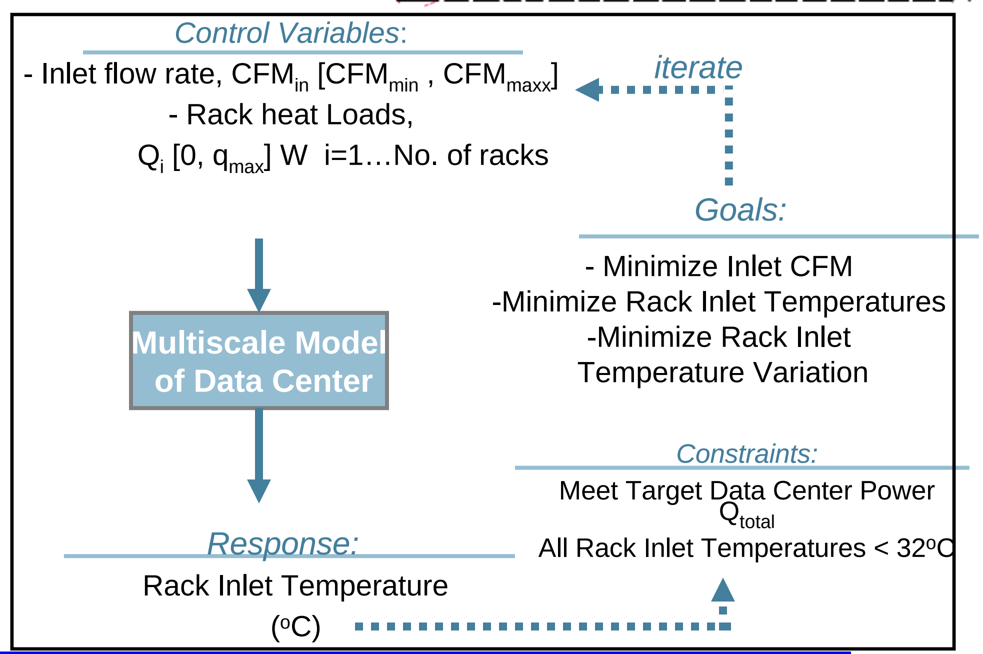

Energy Saving by Dynamic Load Distribution

Increasing the range of changes in the rack heat load

- Heat load distribution of [30 kW, 5 kW, 5 kW, 20 kW] in the case study only needs 1.7 m/s (9,726 CFM) cooling air flow
- It is 19% less than the uniform distribution needs
- This could save ~$189,000 annually in typical real world data centers

Temperature Contours Around Racks:

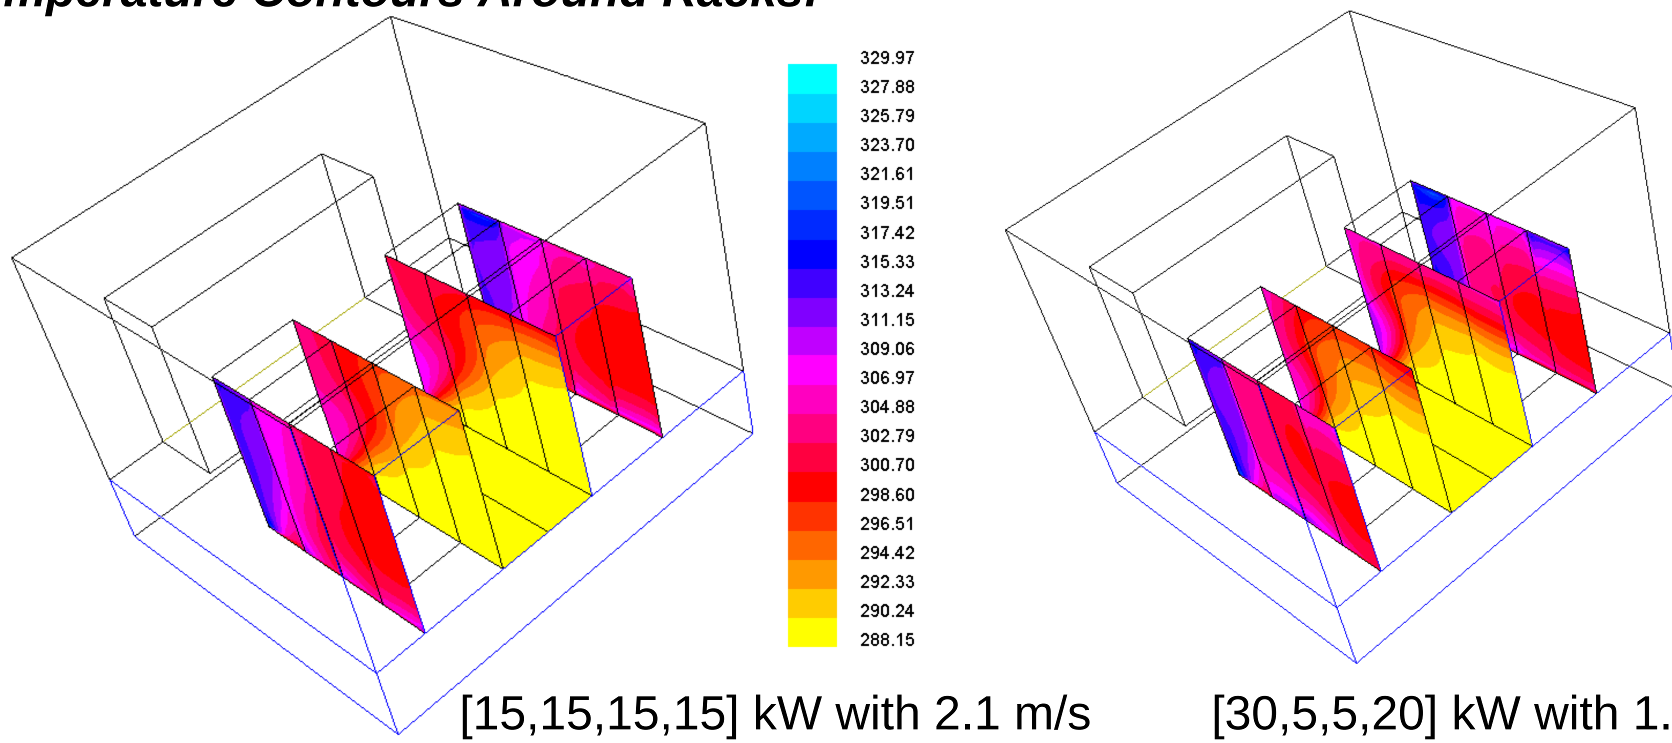

[15,15,15,15] kW with 2.1 m/s [30,5,5,20] kW with 1.7 m/s

The AILM Approach – *Think Globally, Act Locally*

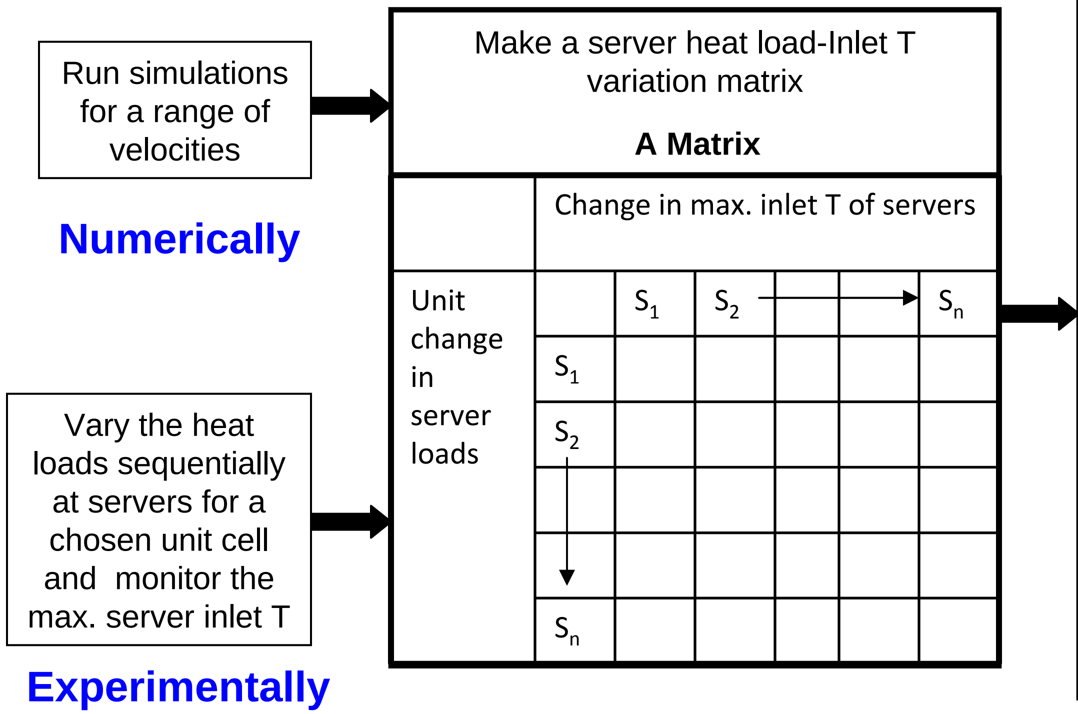

$$\max \sum_{i=1}^{n} l_i$$

$$s.t. \quad A^T l \leq T_{cr}$$

$$l_{\min} \leq l \leq l_{\max}$$

Where,

l_i *server I load*

$l_{\min}$ *Minimum load (startup)*

$l_{\max}$ *Max. load (full utilization)*

T_{cr} *Max. inlet T allowed by ASHRAE*

Advantage:

The simulations run for different velocities are not required for the experimental approach.

Modifications:

Blocks of servers can be identified with same effect or no effect on the inlet T.

- This will give insights on the sparsity of this matrix.
- Reduce the computational work.

An AILM Example

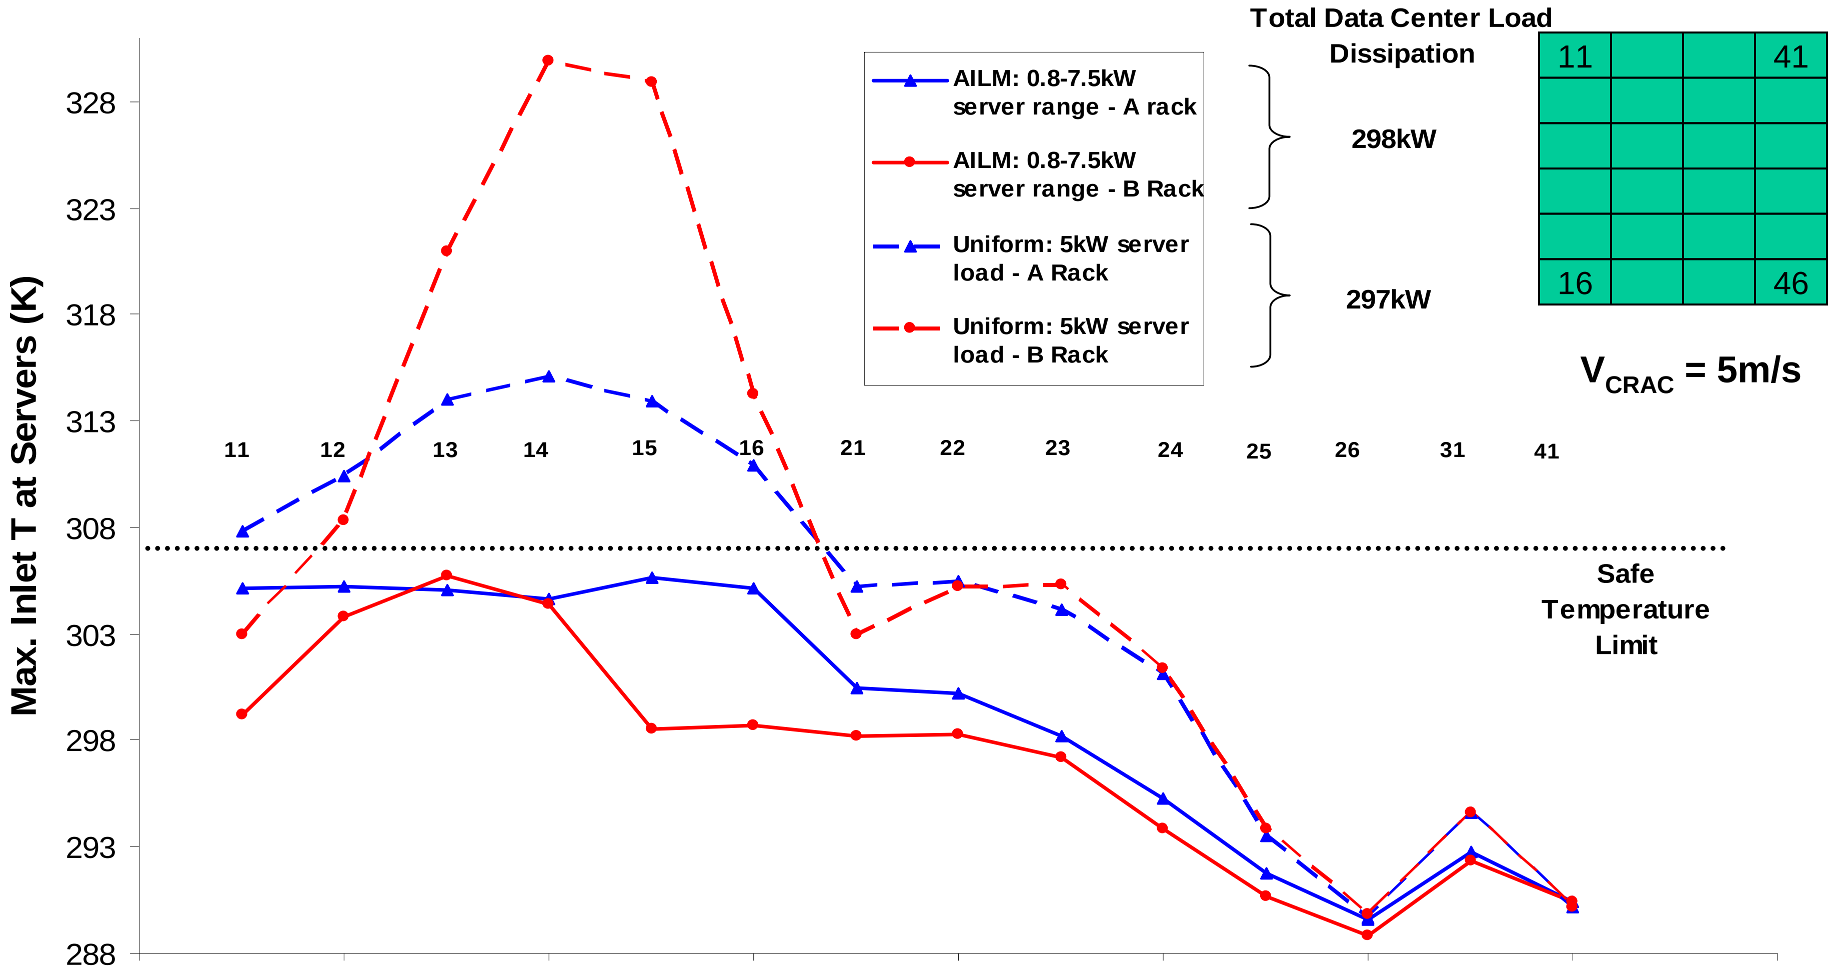

68% increase in allowed heat dissipation (For the same CRAC velocity)

37.5% decrease in Facilities Energy Consumption (For the same heat dissipation)

GT Invention Disclosure

Load Variation with AILM

The load distribution can also be intuitively understood by the recirculation patterns shown in the animation

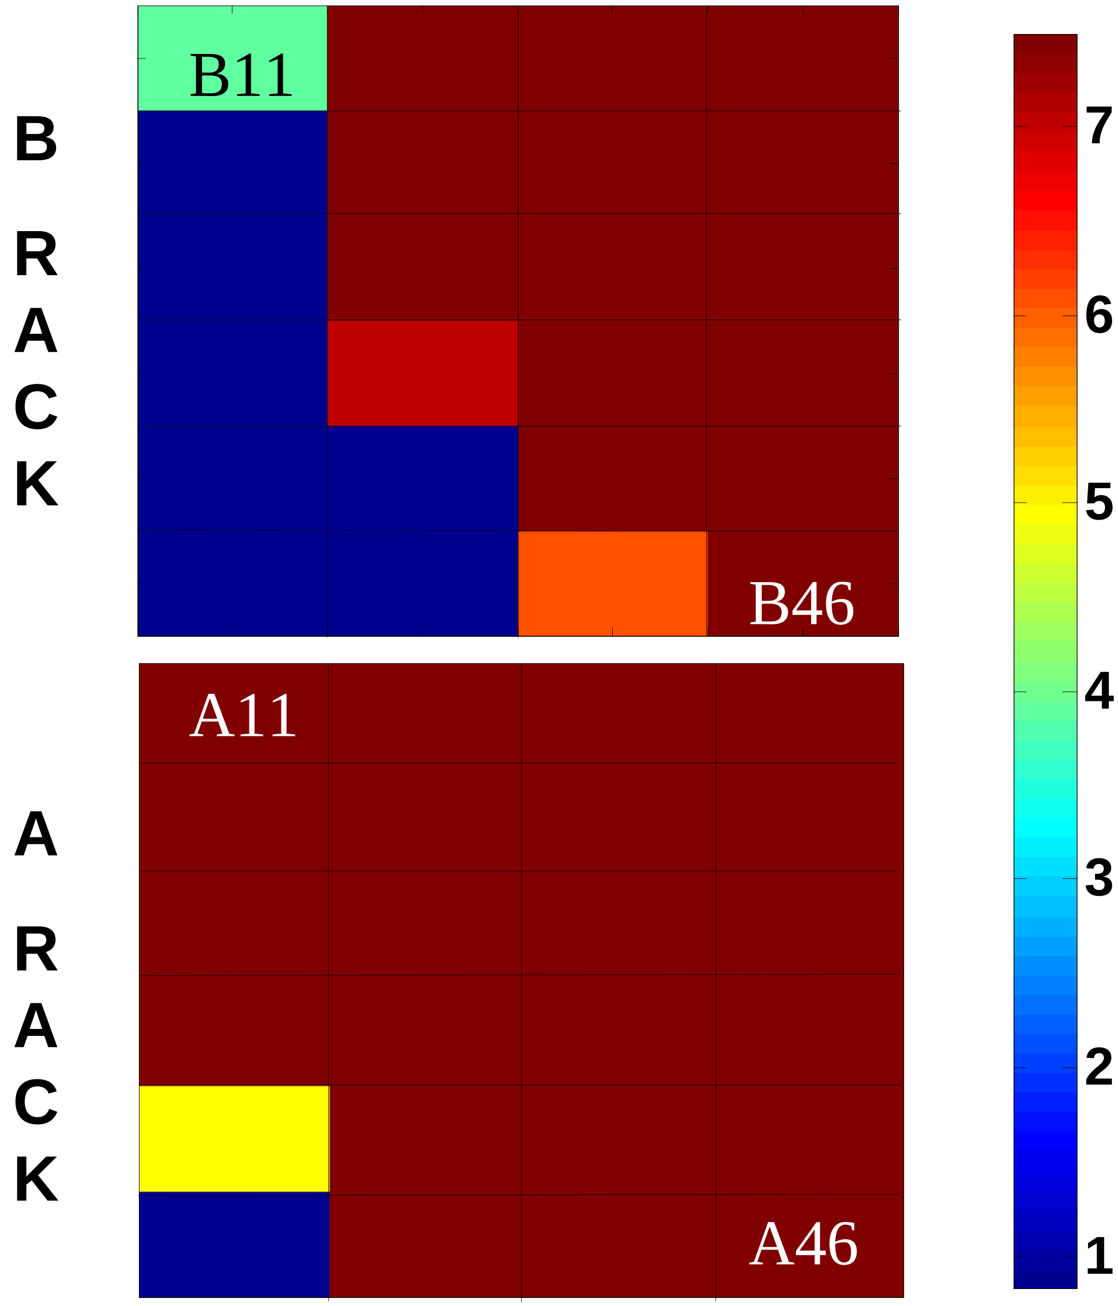

Airside Economizer

- When the outside air is cooler than the return air, hot return air is exhausted and replaced with cooler air
- Utilizes a system of sensors, ducts and dampers to allow entry of the appropriate volume of outside air to satisfy cooling demand
- When economizer system is operating, the use of air conditioner system's compressor and related components is reduced or eliminated

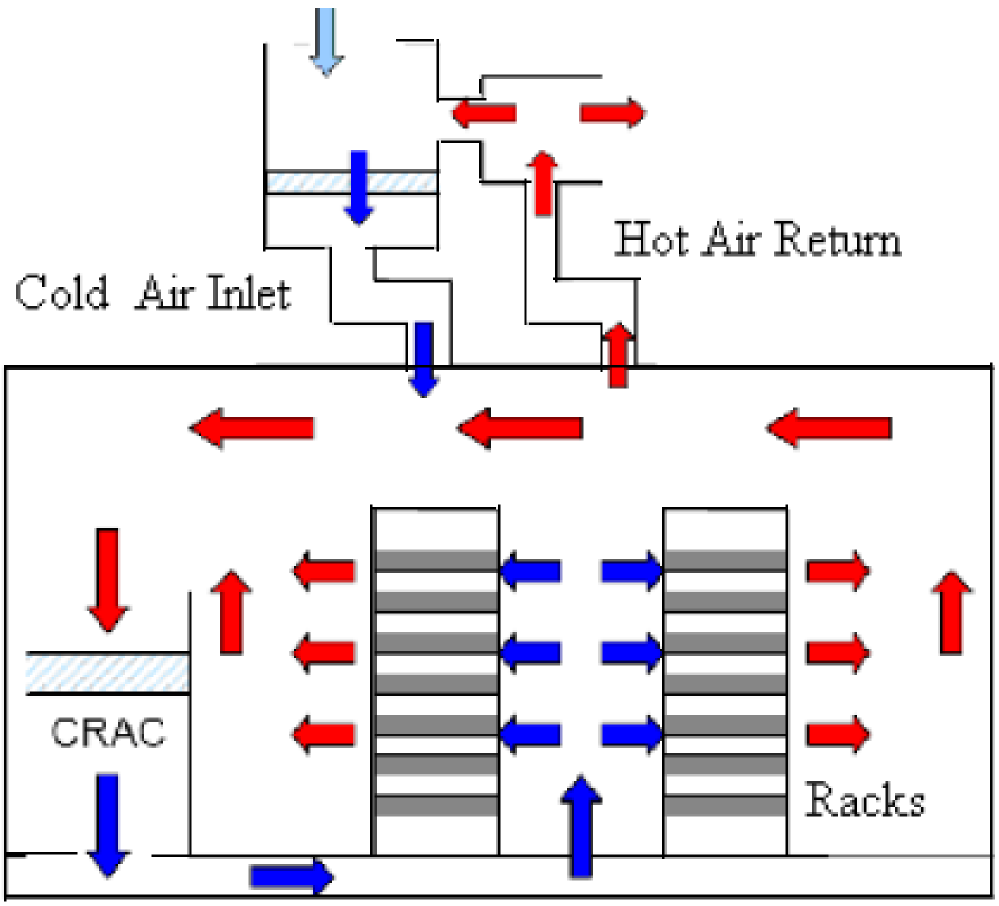

Model Data Center with Economizer – Top View

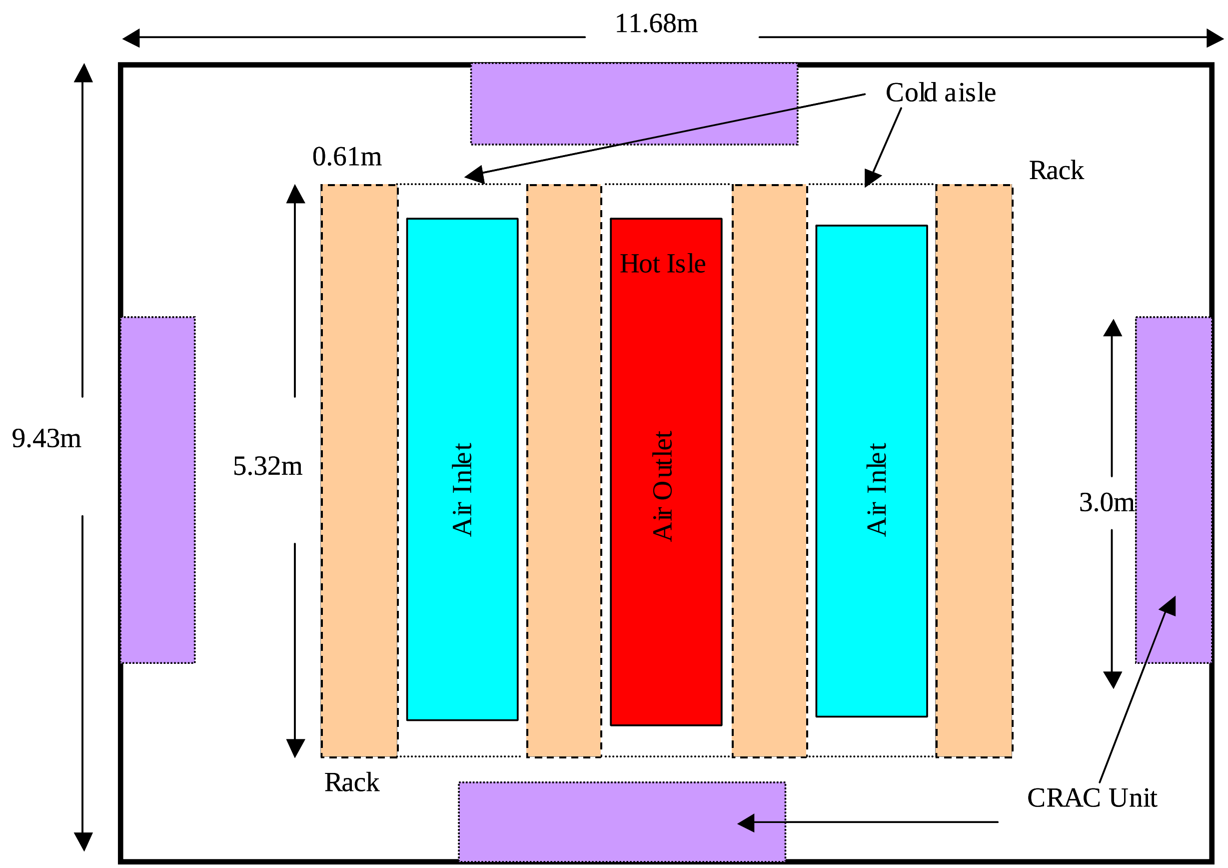

- Cold air inlet (2) directly above the cold aisles (5.0 m x 0.8m)
- Hot air outlet is directly above the hot aisle (5.0 m x 0.8m)
- This helps minimizing hot air recirculation
- Outside air does not interfere with the inlet air to the server

Air Flow With Economizer

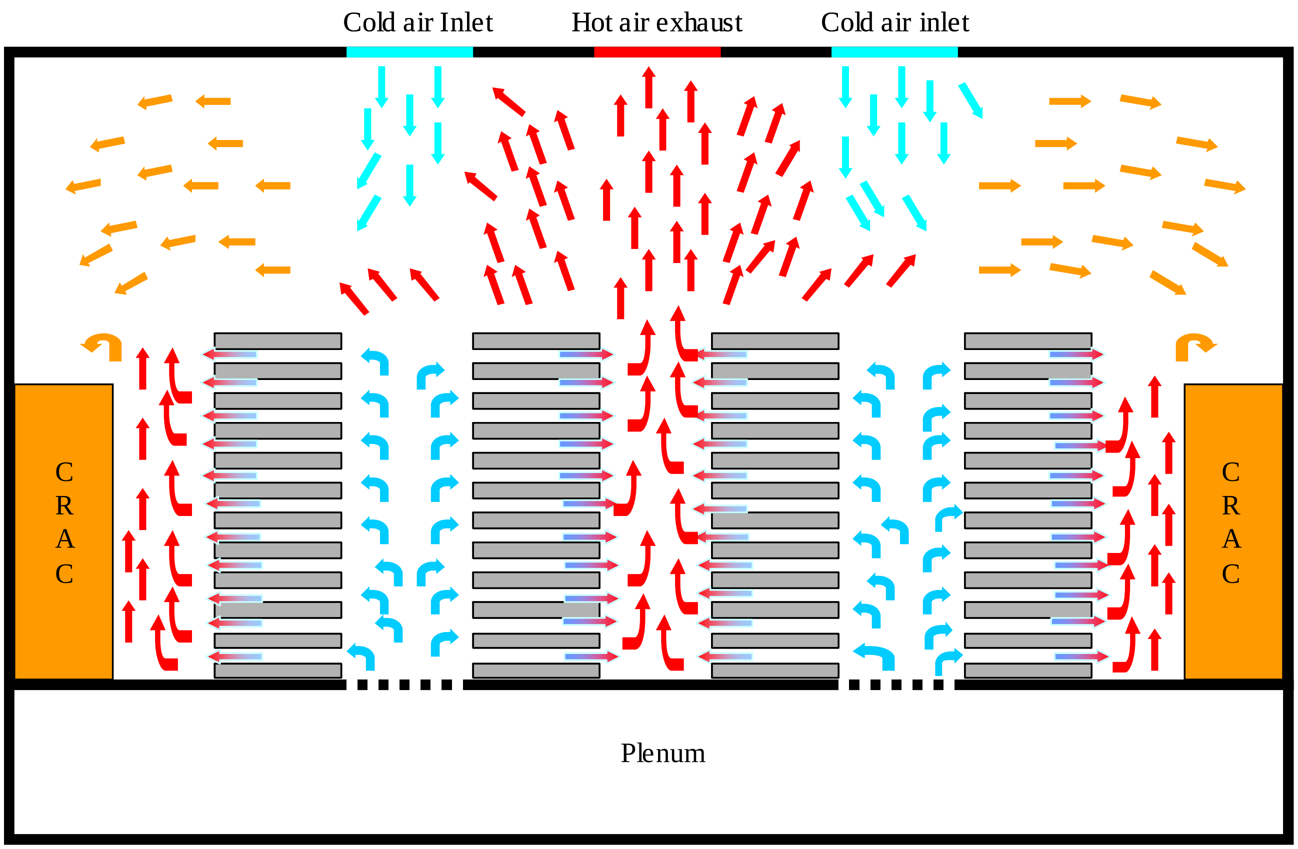

Energy Saving for Atlanta

Month	Average Outside Temerature (^{0}C)	Average Outside Humidity (RH %)
January	5.9	79
February	8.2	77
March	12.4	77
April	16.4	79
May	21.0	82
June	24.9	84
July	26.7	88
August	26.1	89
September	22.9	88
October	17.1	84
November	11.9	82
December	7.4	80

Month	Flow rate into the server room (kg/s)	Baseline Energy usage by CRAC units (kW)	Energy usage by CRAC using economizer (kW)	Saving (%)
January	7.9	118.4	40.9	65.4
February	7.9	118.4	60.5	48.9
March	7.9	118.4	93.5	21.0
April	0.0	118.4	118.4	0.0
May	0.0	118.4	118.4	0.0
June	0.0	118.4	118.4	0.0
July	0.0	118.4	118.4	0.0
August	0.0	118.4	118.4	0.0
September	0.0	118.4	118.4	0.0
October	0.0	118.4	118.4	0.0
November	7.9	118.4	93.1	21.4
December	7.9	118.4	57.1	51.8

- Economizer usable for 5 months, as the temperature is below data center operating temperature (15 ^{0}C)
- CRAC power usage reduction up to 65% achieved in January
- With increase in outside temperature the savings keep falling
- Mass flow rate has been fixed as not to affect the RH near the inlet of the servers

Can Changing the Layout Help ?
Hot Aisle Cold Aisle (HACA) Layout

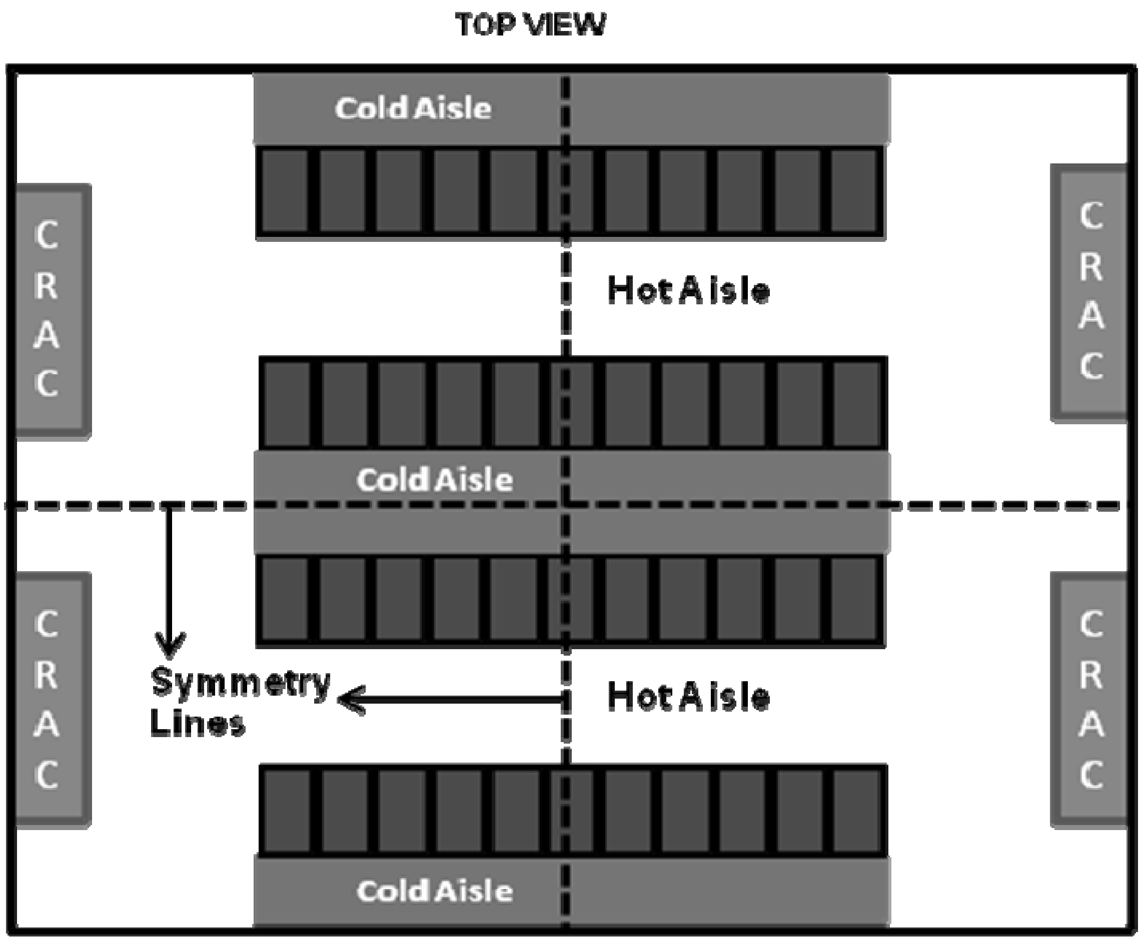

Issues

- No other ordered rack configuration studied.
- Highest footprint possible with HACA is around 30%.
- Crippled with mixing of cold inlet air with hot aisle air and thus reducing its cooling potential
- Short circuiting of cold air directly to the CRAC return

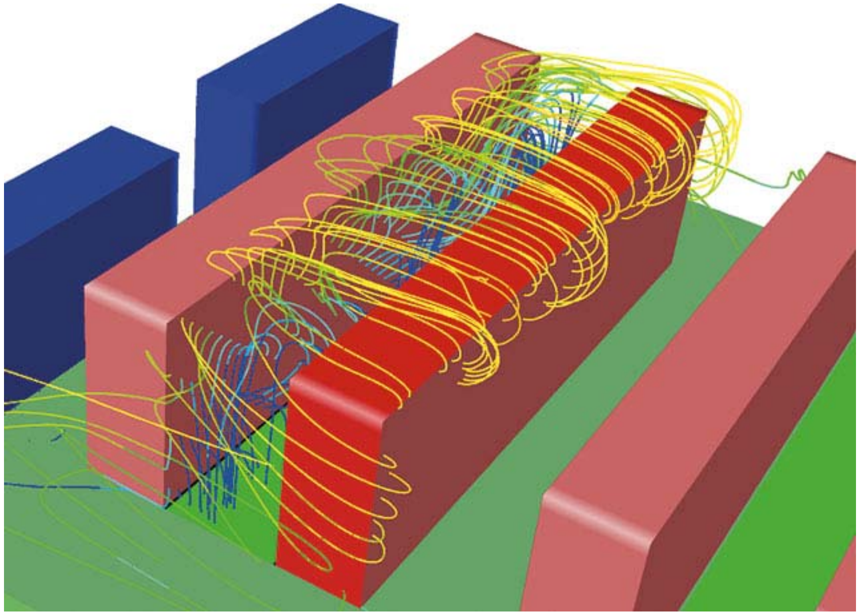

Top-side Recirculation*

** Source - http://www.fluent.com/about/news/newsletters/04v13i2/img/a19i1.jpg*

A New Concept : S-Pod Layout

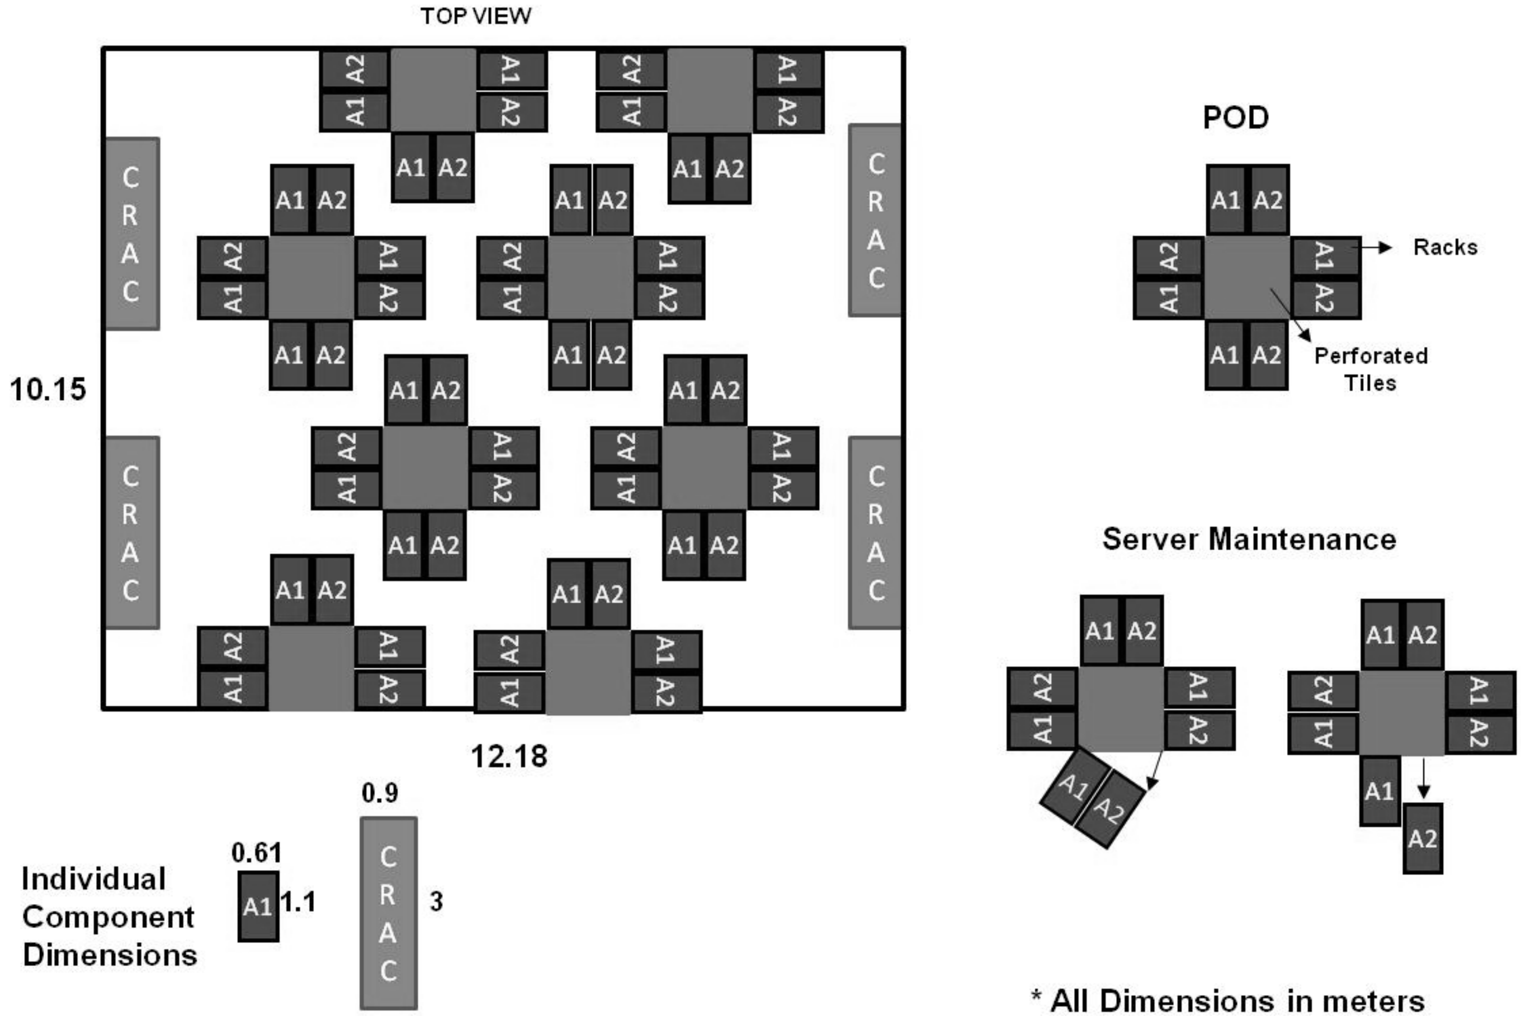

GT Invention Disclosure

Modeling

- ➢ Racks treated as black boxes with constant volumetric heat generation
- ➢ Standard k-ϵ model
- ➢ CRAC inlet temperature to plenum is 288.15K
- ➢ Modeling – GAMBIT 2.3, Simulation – FLUENT 6.3

HACA:
- Min. grid dimension – 0.12m
- Total Grid Size – 0.4 million
- Simulating one-fourth of the facility due to symmetry

S-Pod:
- Min. grid dimension – 0.14m
- Total Grid Size – 1.2 million

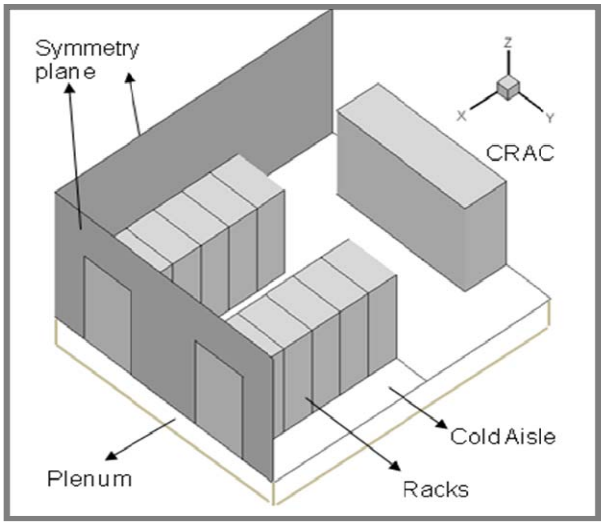

HACA

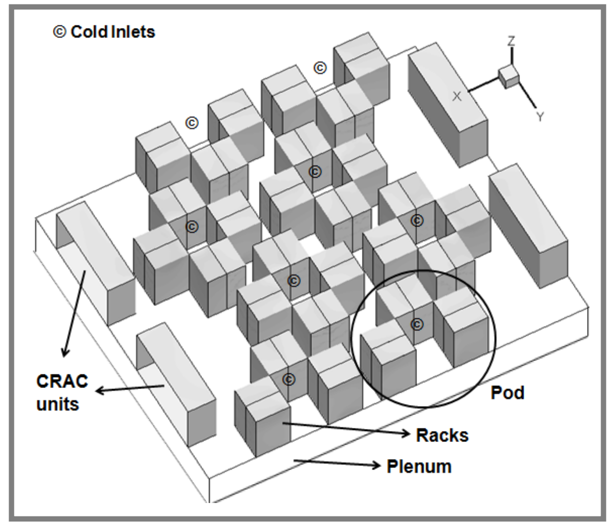

S-Pod

Results for V_{CRAC} = 7 m/s

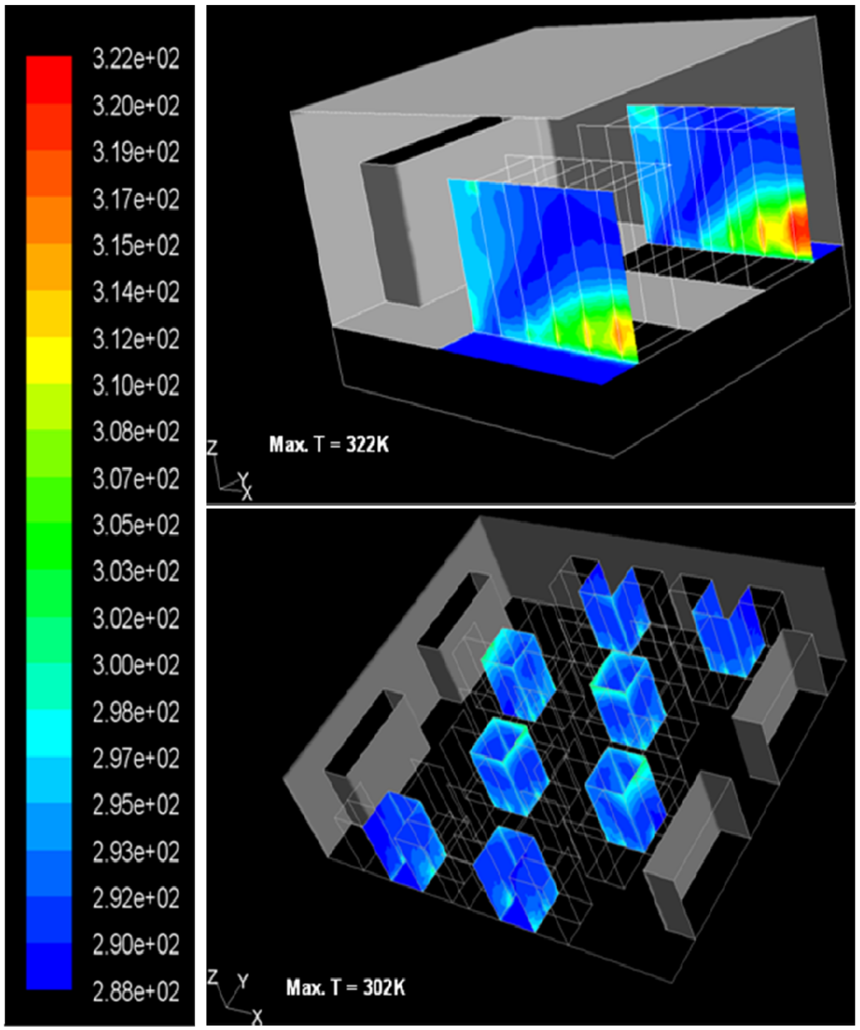

Temperature Contours

- Rack heat generation **16161.61 W/m³ (21.7 kW racks)**
- HACA layout: Hot spots at lower end of rack
- S-Pods: 3 column pods have lower temperatures than 4 column pods.
- Maximum allowed rack inlet temperature is 305.35K (32.2°C) [ASHRAE]
- Maximum allowed per rack uniform heat generation limit: HACA: 15 kW; S-Pod: 22 kW.
- S-Pod case has 27.3% more racks
- **Net heat load capacity of the Data Center with S-Pod layout increases by 95%.**

Results for V_{CRAC} = 4 m/s

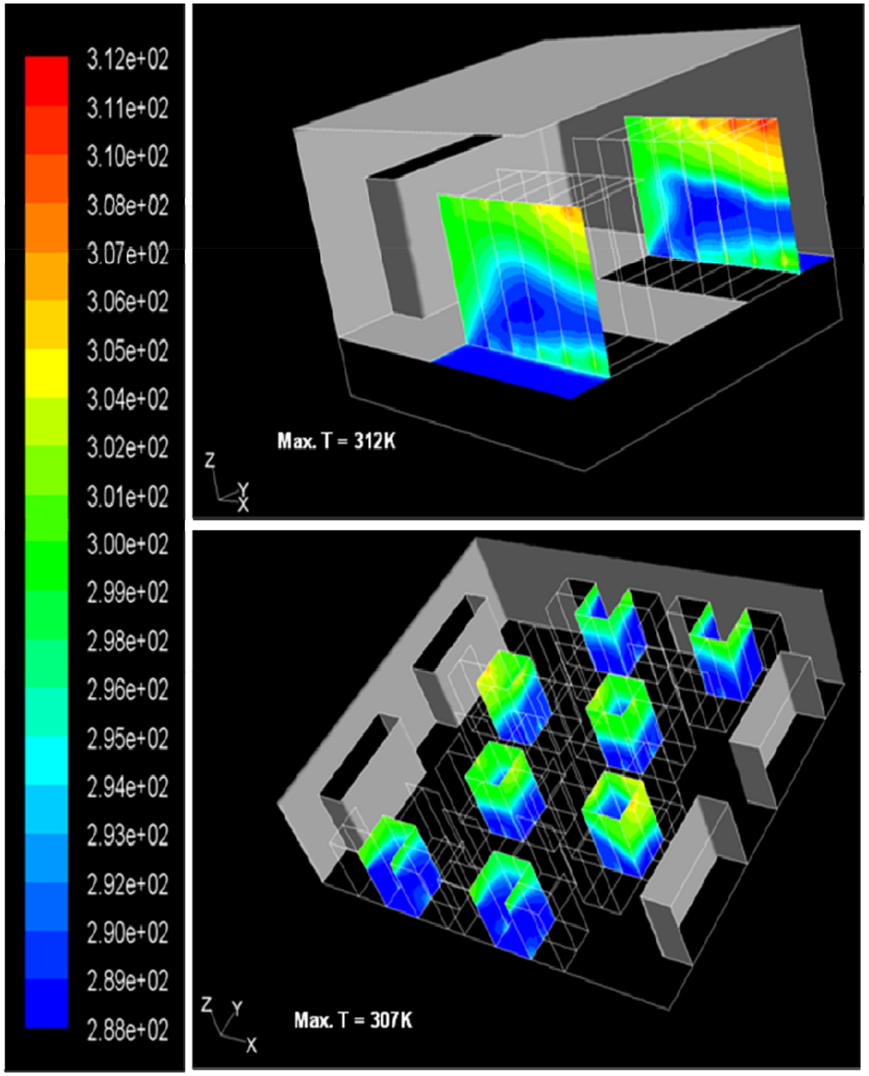

Temperature Contours

- Rack heat generation: **8888.89 W/m³ (11.93 kW racks)**
- HACA layout: Hot spots occur at the **higher** end of the rack
- Pods closer to CRAC unit have higher inlet temperatures
- Topside recirculation is prominent in racks closer to the CRAC.
- Maximum allowed per rack uniform heat generation limit - HACA: 8.6 kW; S-Pod: 11 kW.
- S-Pod case has 27.3% more racks
- **Net heat load capacity of the Data Center with S-Pod layout increases by 63.6%.**

Recirculation in Pod

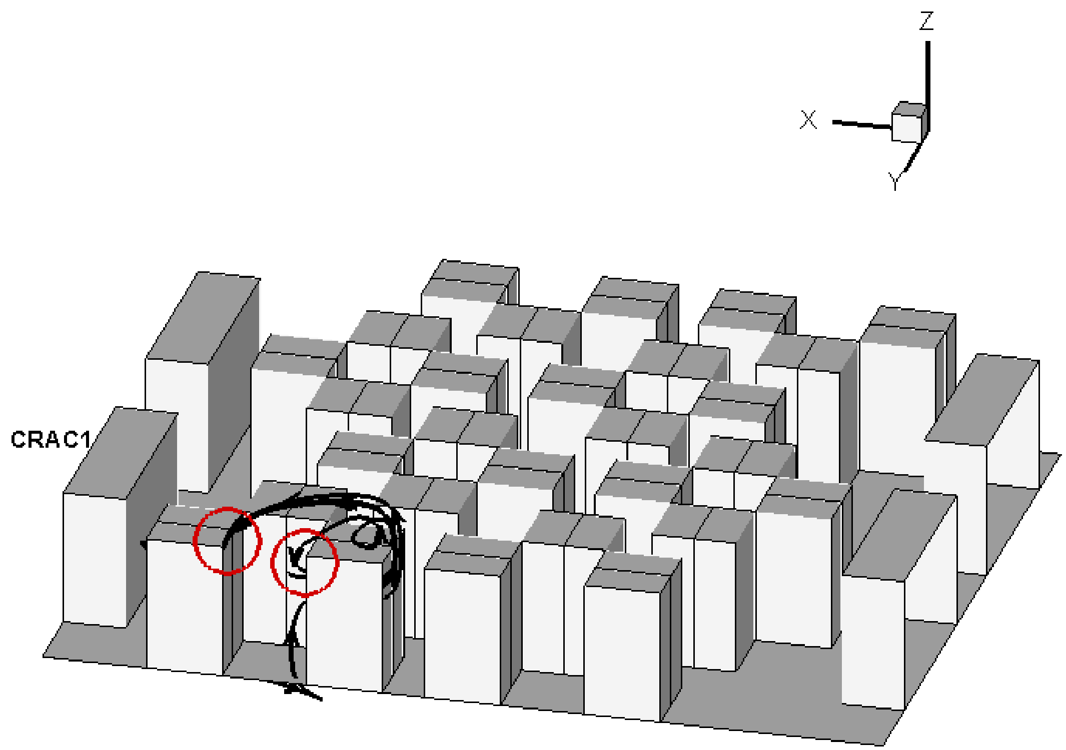

- Path of streamline (black curve): CRAC → Plenum → Cold Inlet → Outside Pod
- Hot air from the exhausts being recirculated in the cold inlet (Red Circles).
- No side recirculation; Topside recirculation has to be thwarted.

Introduction of Barriers

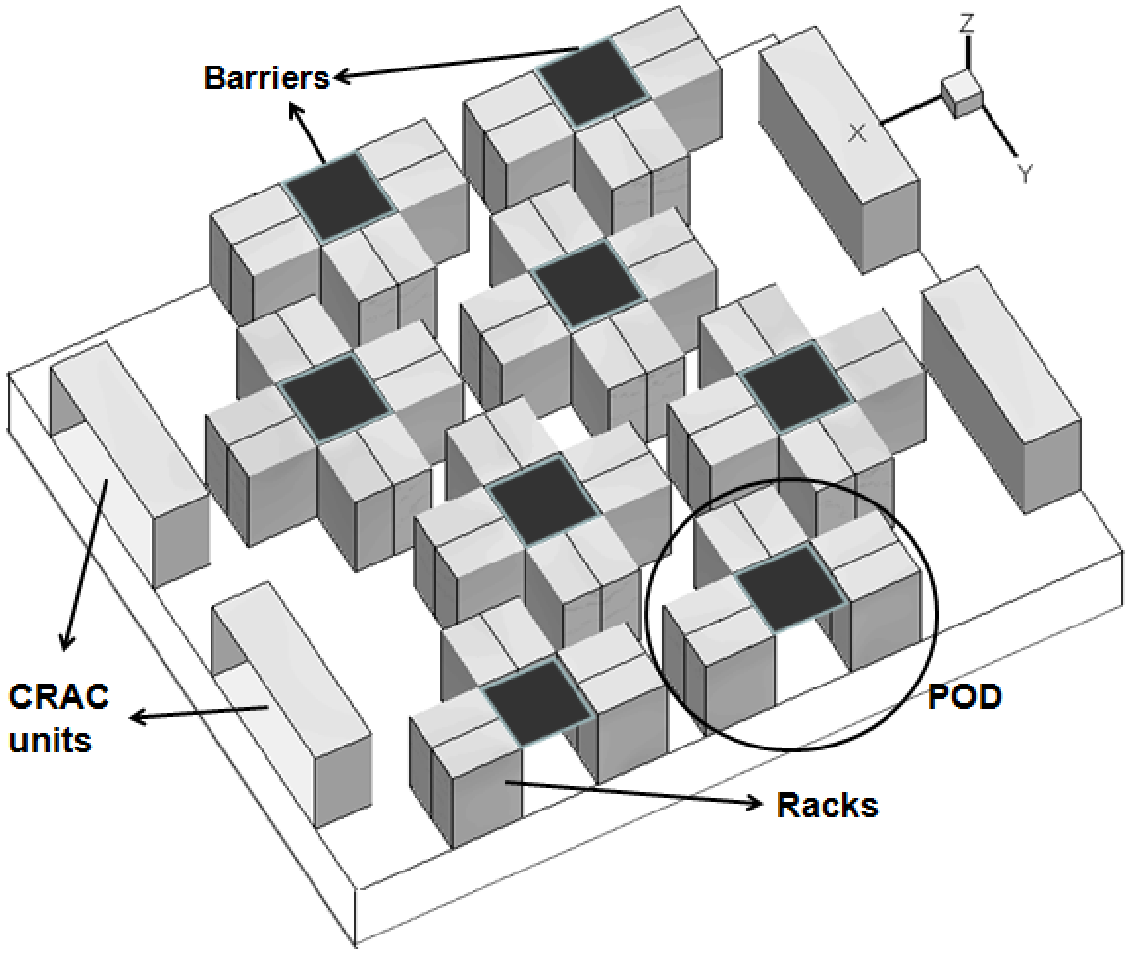

Reason

1. To Reduce topside recirculation
2. To stop short circuiting of cold air to the hot aisle directly.

Caution

Enough mass flow rate should be provided by plenum to **avoid "starving" of servers**

Conclusions

1. Large and increasing cooling costs of data centers require urgent attention
2. IT and facilities staff need to collaborate for holistic solutions
3. On-demand allocation of cooling resources for existing facilities is a promising approach
4. Air-side economizers present an attractive energy saving option for new facilities
5. Air-delivery layouts can be further optimized