Sustainable Information Technology Ecosystem

supply and demand side management

Chandrakant D. Patel

HP Fellow

chandrakant.patel@hp.com

HP Laboratories, Hewlett-Packard Company

Drivers of Next Generation of IT Services

Up and coming Small and Medium Enterprises

<u>IT Advantage in the Emerging Economies</u>

- Enabling Business Transformation

Retail Example:

<u>Subhiksha Trading Services</u>

"IT will be the key enabler and the key differentiator between operations that are or aren't well run. Whether it is managing the front end or logistics support or in stock and inventory management".

- R. Subramaniam, MD and Founder of Subhiksha Trading Services

- *The Economic Times, March 13, 2007*

Subhiksha Store

Vadodara, Gujarat, India

invent

Drivers of Next Generation of IT Services

Micro businesses in emerging economies

IT services as a means for micro-businesses to improve quality of life if provided at a given price point

Sustainable IT Ecosystem

Deconstruct conventional supply chain and replace with Sustainable IT ecosystem

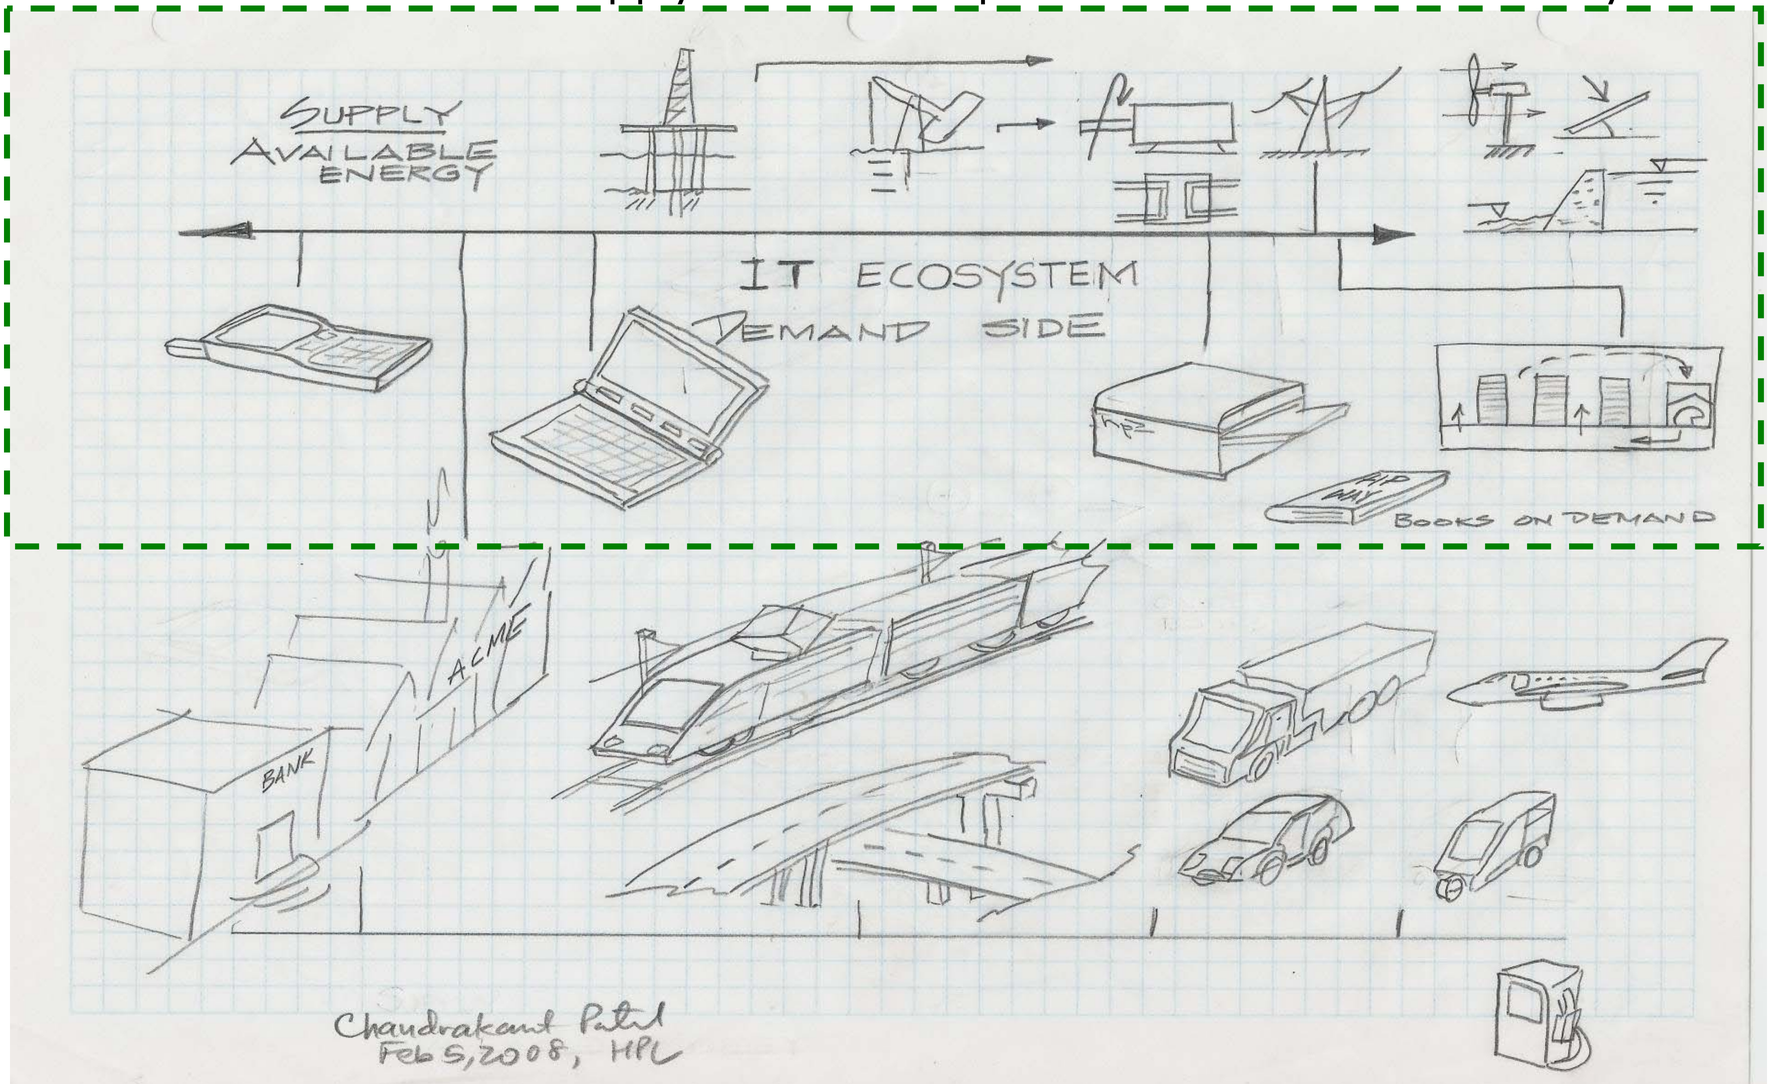

The Road Ahead

Next Generation IT Services

- IT and telecommunications is a means to overcome physical infrastructure issues such as transportation, increase productivity and improve the standard of living.
- Necessitates a total cost of ownership (TCO) that enables growth in new areas e.g. 800 million in India who want to use IT as a means to improve quality of life and business competitiveness.
  - Achieving the TCO necessitates management of available energy in the IT ecosystem as a key resource.

Devising a Sustainable IT Ecosystem

  - Evolve an "end to end" metric and process to address sustainability in IT
  - Design, manufacture, operate and manage end of life of products that minimize consumption of available energy – a "cradle to cradle" perspective

hp invent

Devising a Sustainable IT Ecosystem

supply side + demand side management

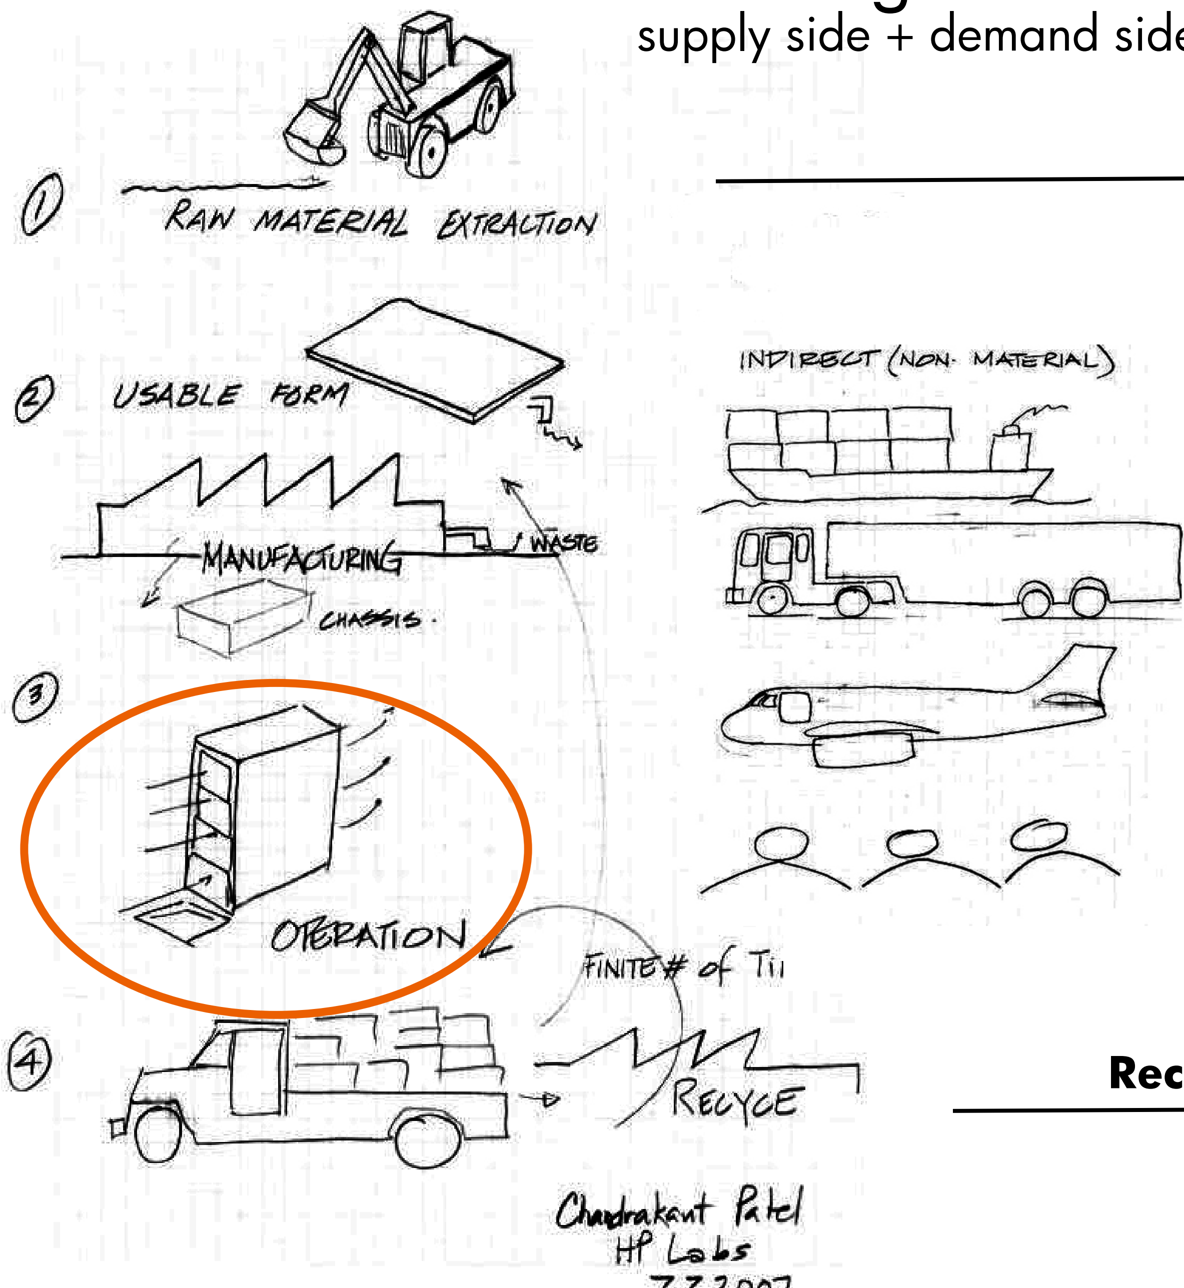

Ground State

irreversibility

Framework needed to quantify the consumption of available energy **(exergy)**

Reclaim

Exergy or Available Energy

- "Exergy" is a measure of the quality of energy
  - Or, alternatively, the amount of work available from a given amount of energy ('available energy')

- From the 2nd Law of Thermodynamics:
  - Irreversibilities in real processes continuously degrade the quality of energy, simultaneously destroying exergy
  - For example, the conversion of coal to electricity
    - Or electricity to heat
    - Or rejection of heat to ambient

Approach

- Can a measure of the total exergy destroyed across a product's lifetime ("lifetime exergy") be a measure of the environmental sustainability?

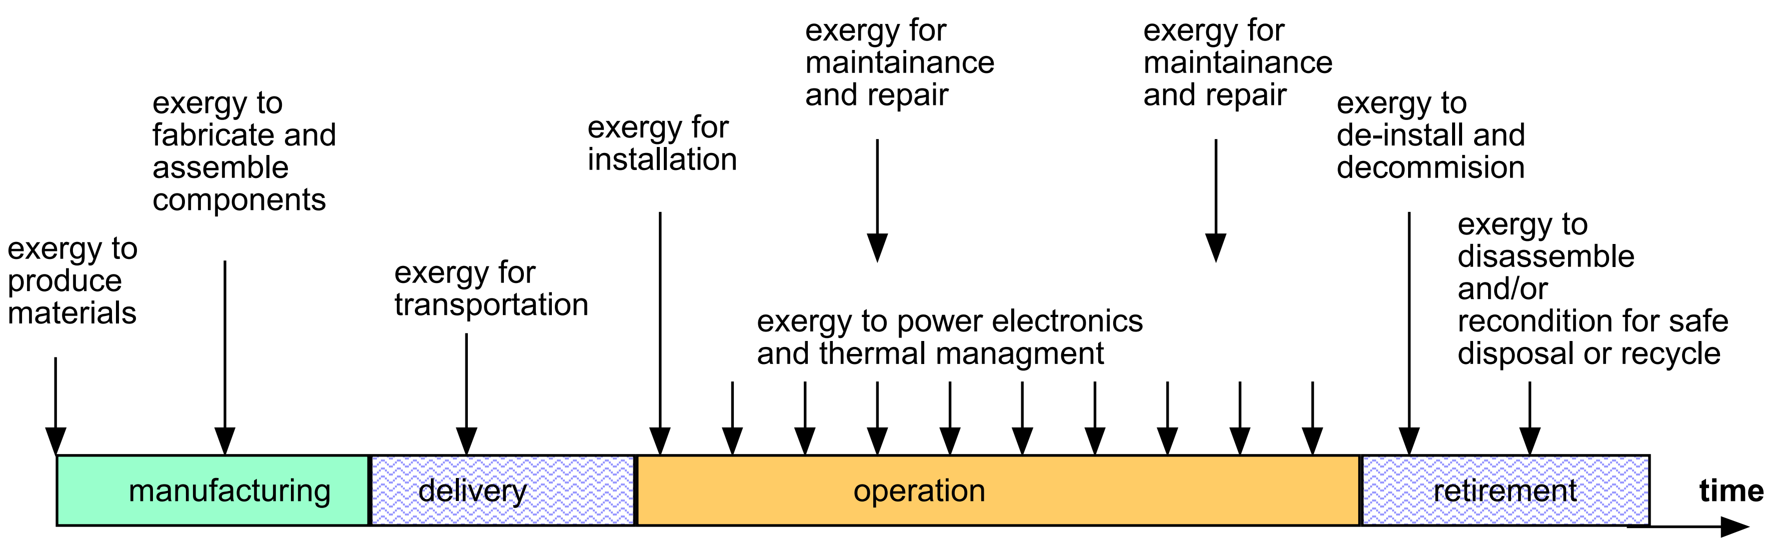

Operation

Powering the electronics and thermal Management

Operation

Exergy consumed:

- Electronics and electro-mechanical systems
- Work required to remove heat
  - Flow Work
    - *volume flow, m^3/s*
    - *pressure drop, Pa*
  - Thermodynamic Work
    - *Temperature, °C*

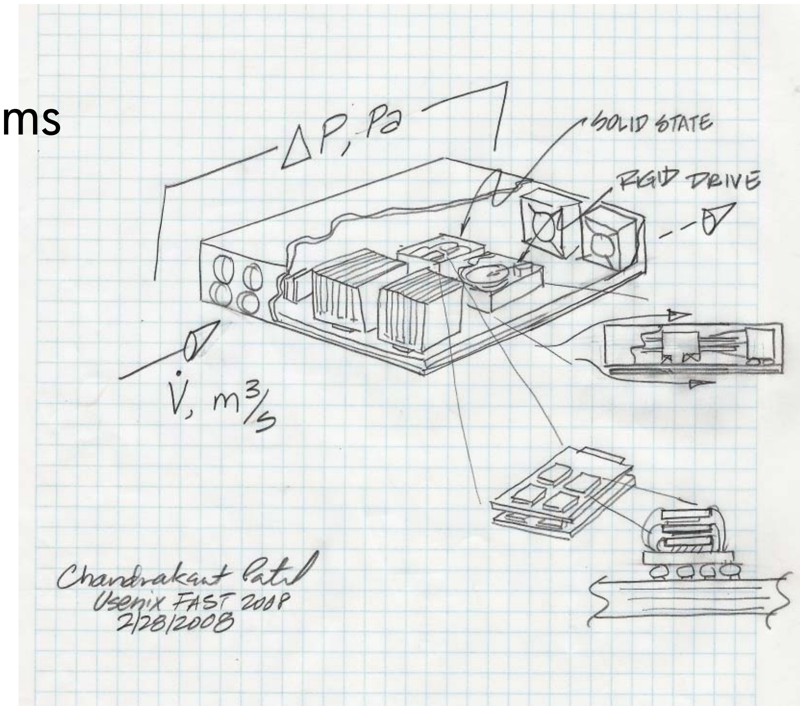

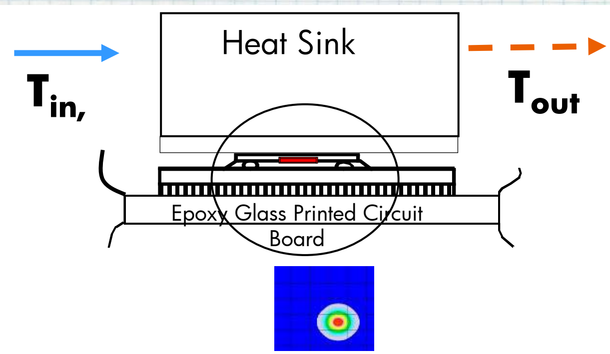

Technology Trends

chip packaging, integrated photonics, semiconductor technology

Source: Chandrakant Patel projections, non industry data

- High Power Density
- Temperature Control due to optical devices

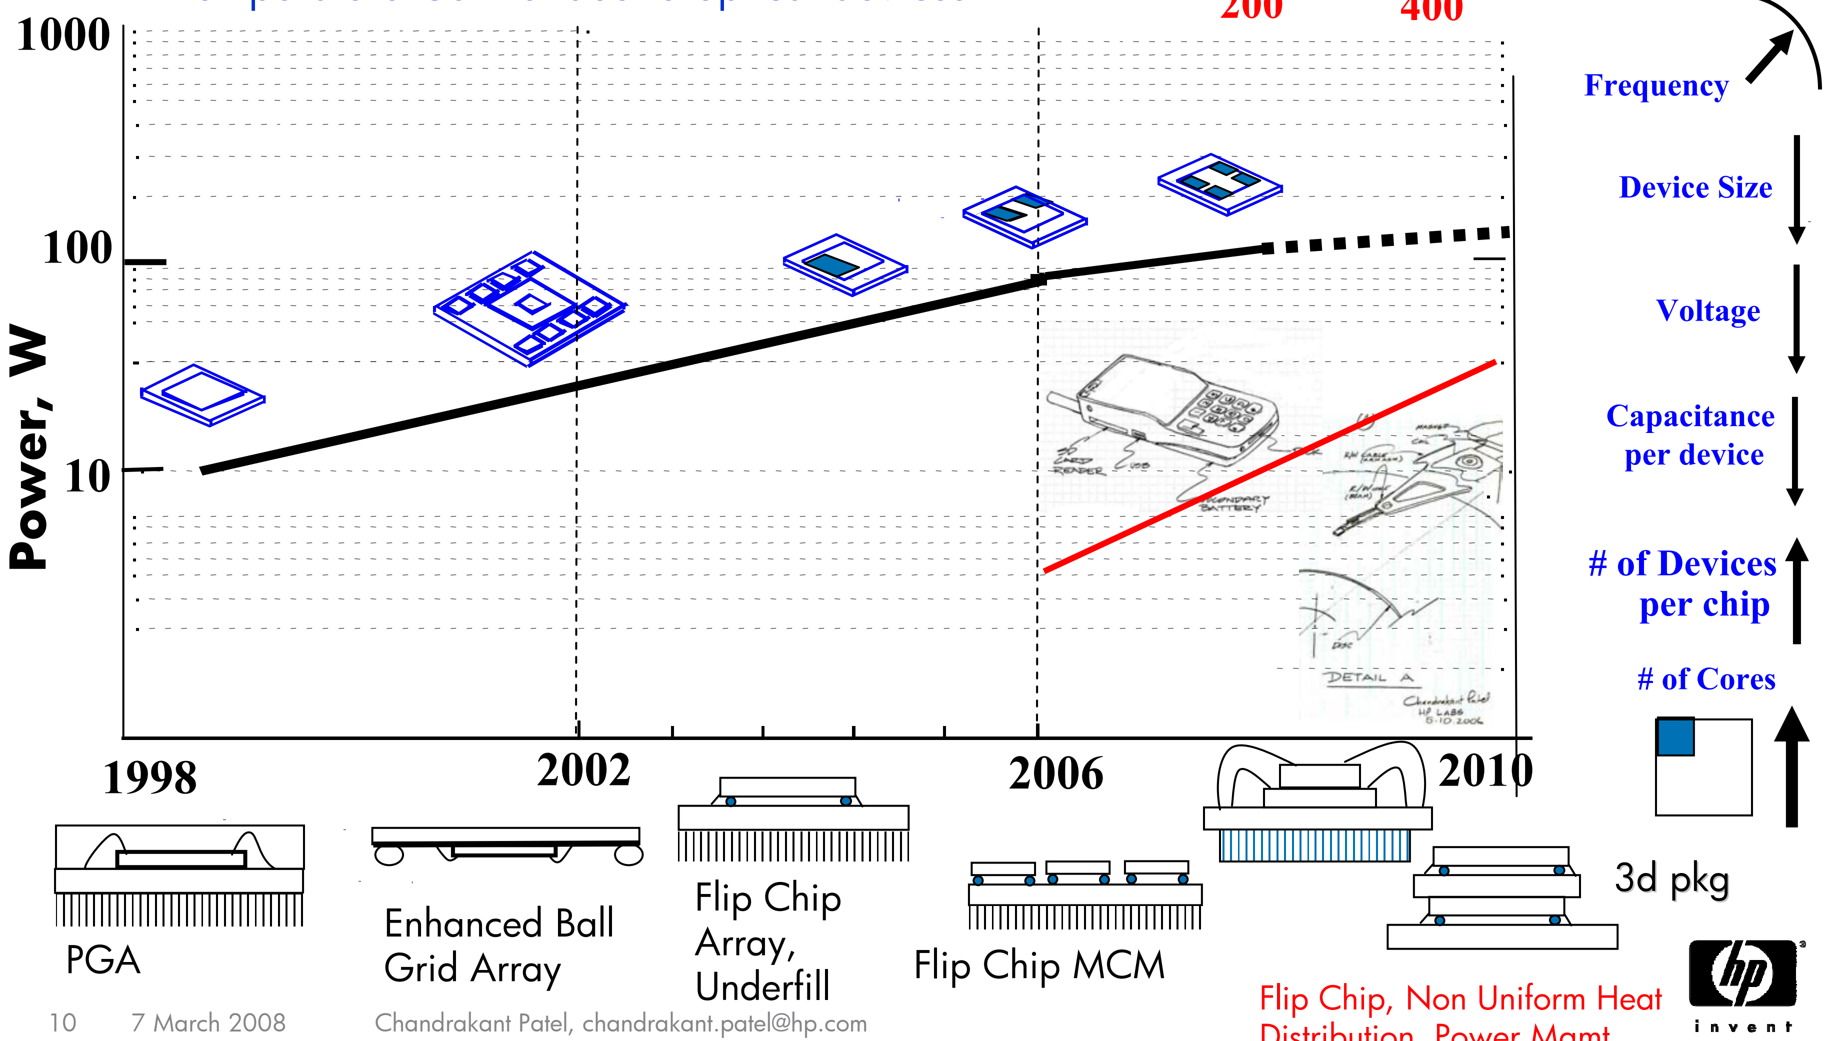

Thermal Management Challenge

stacked devices: chip and package scale

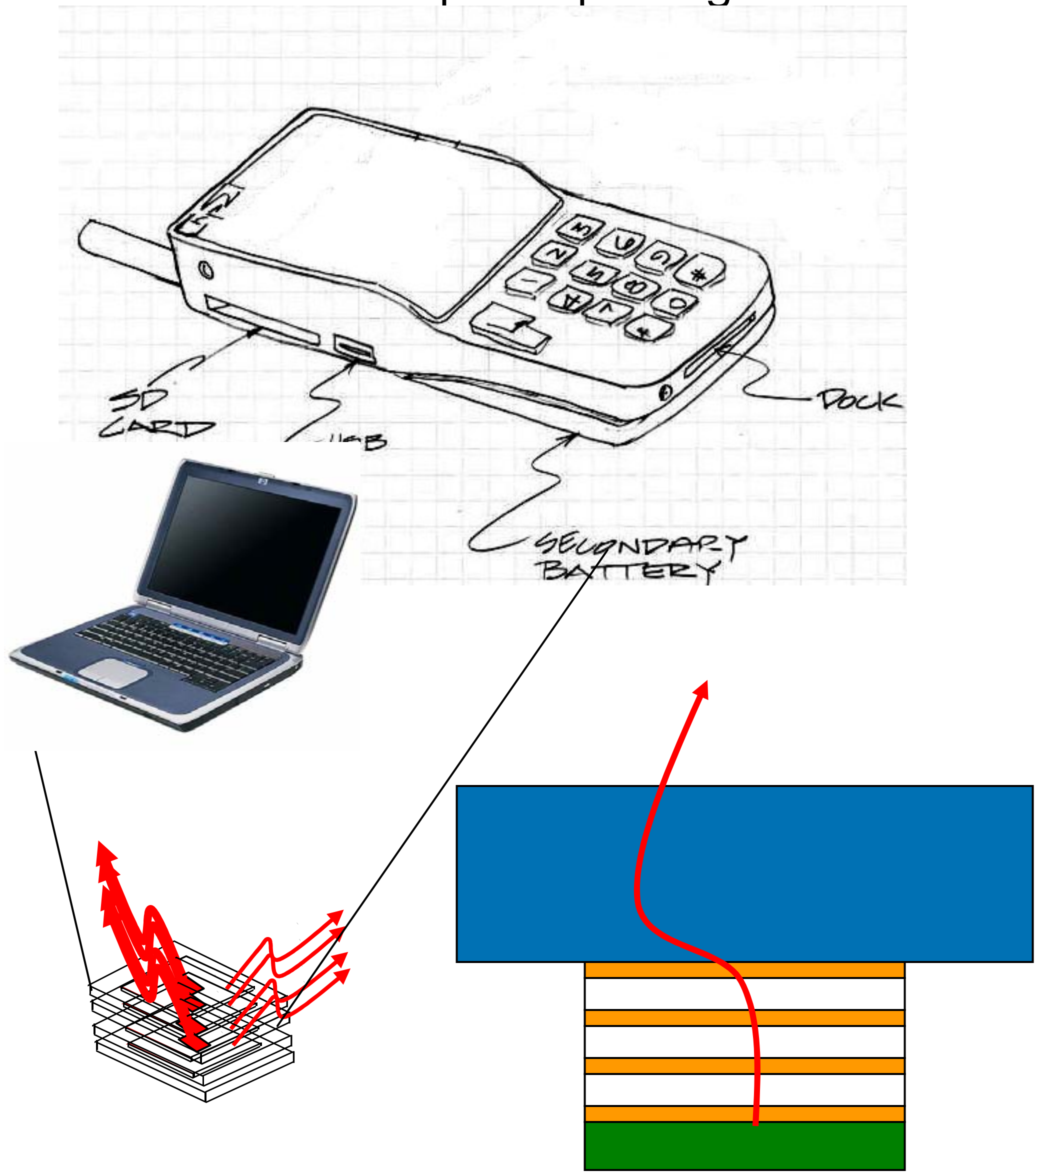

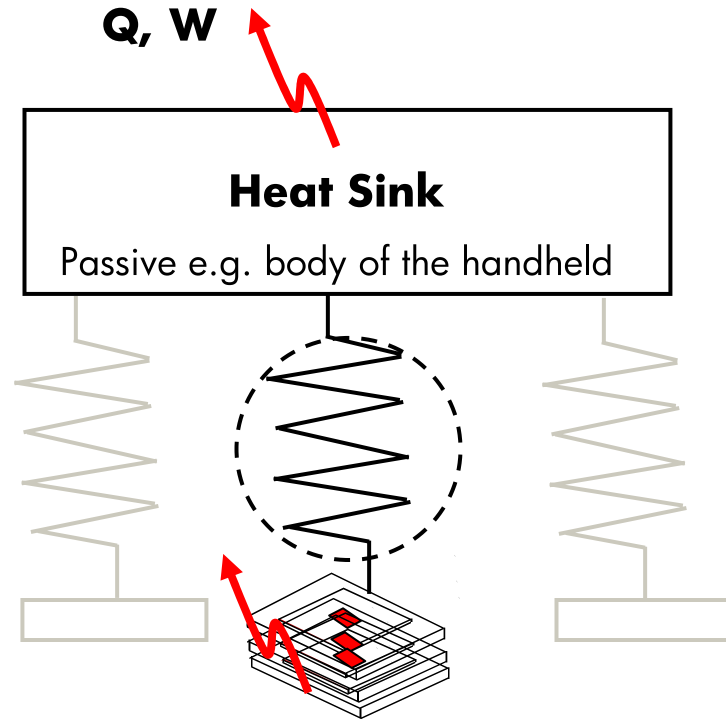

High power density microprocessor with stacked devices

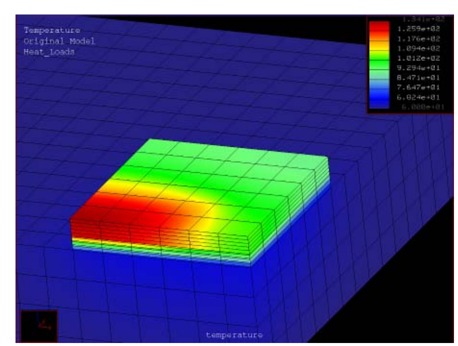

Chip, Package and System Scale

Work required to remove the heat

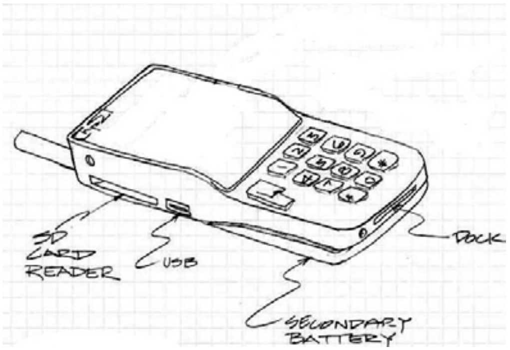

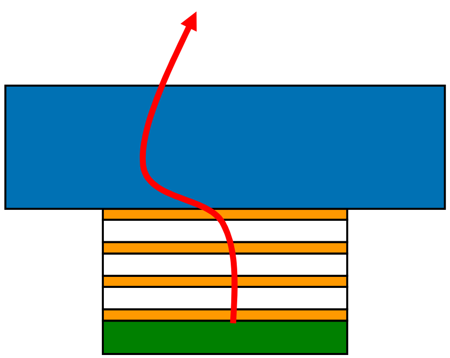

Demise of Passive Only Cooling Solution

- *Work Required at the chip interface level due to high power density and advent of new packaging*

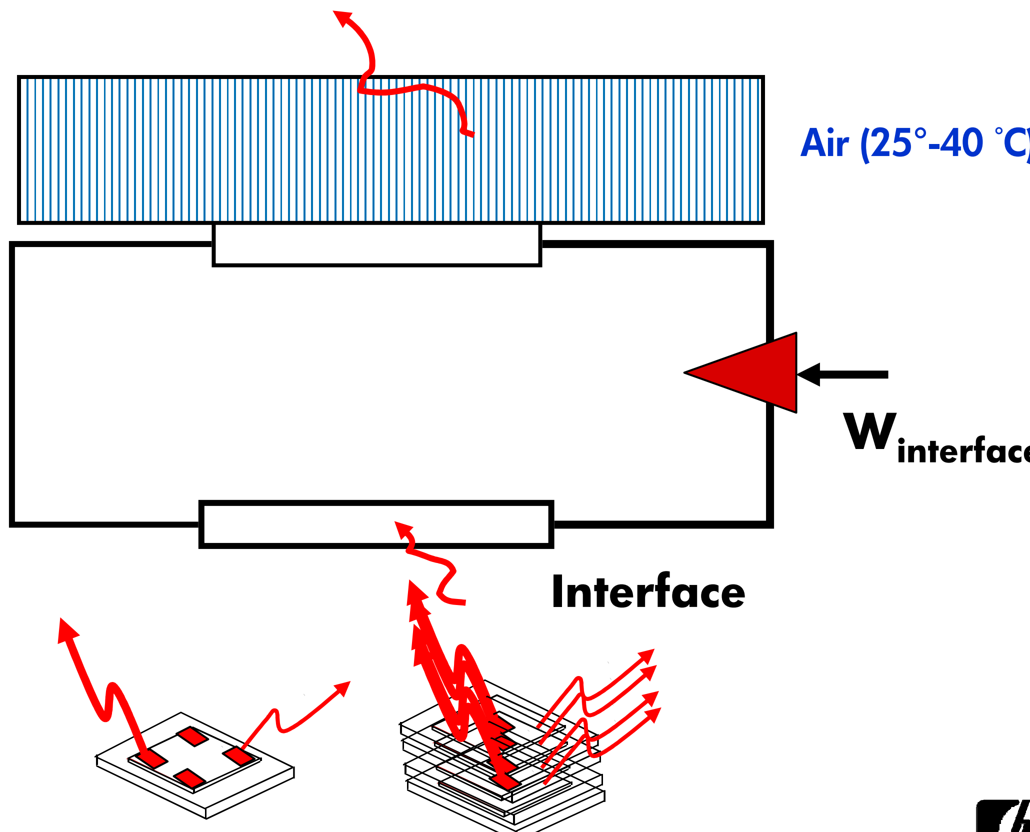

Data Center Scale

Work required to remove heat

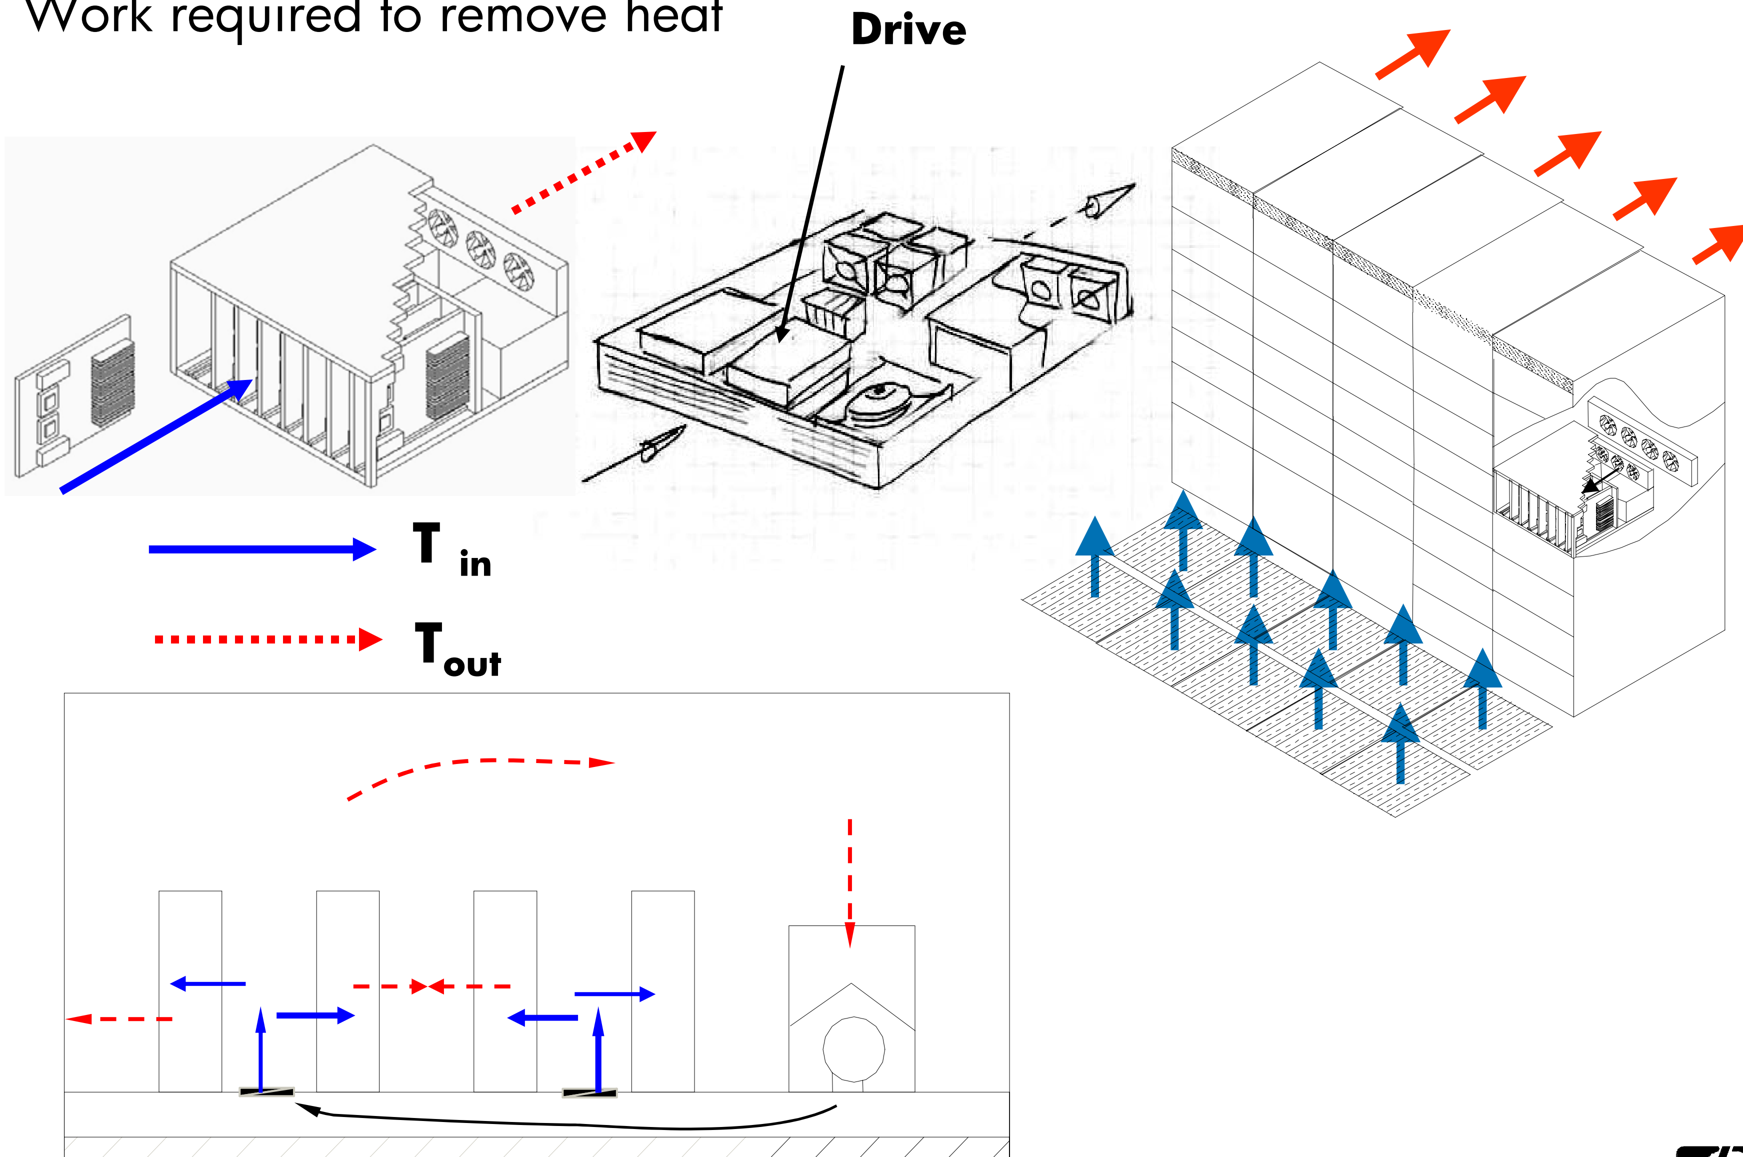

Data Center Cooling

Conventional approach in controlling temperature in the data center

- Single point temperature measurement at the return of the CRAC (Computer Room Air Conditioning Unit), typically set at 20 °C

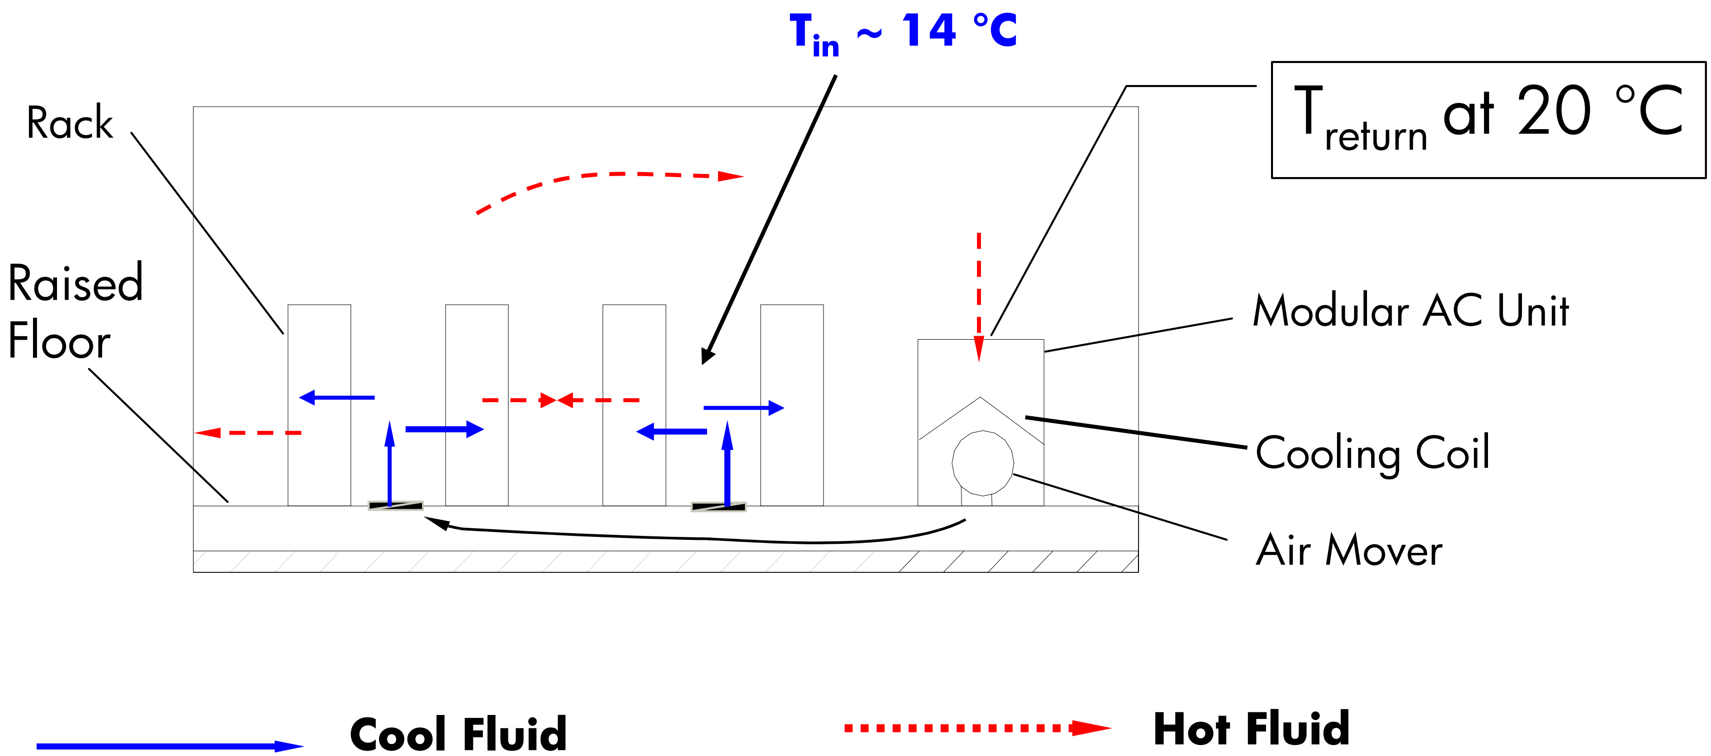

IT Ecosystem

Billions of Users and 1000s of Data Centers

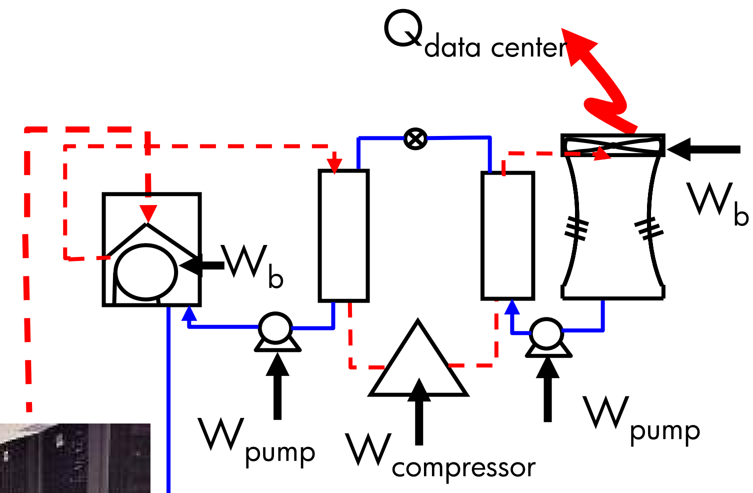

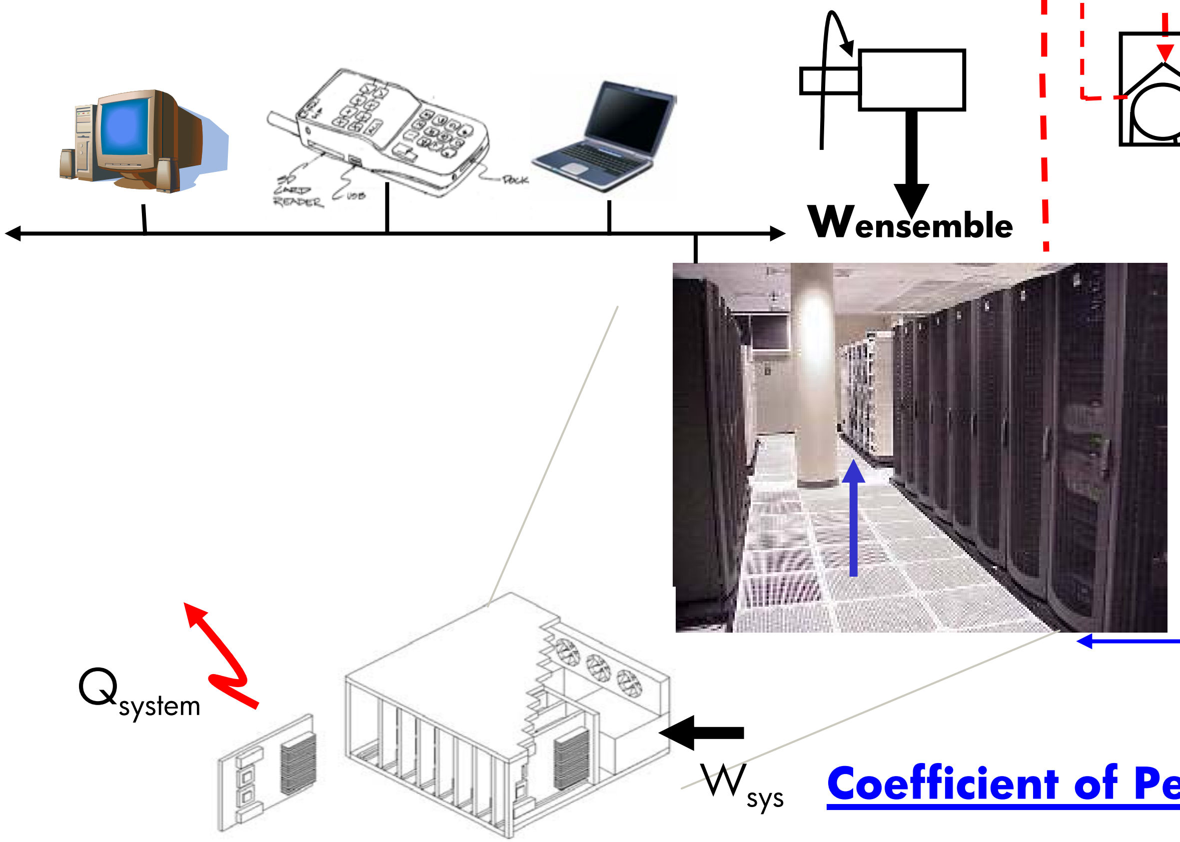

Coefficient of Performance of the Ensemble

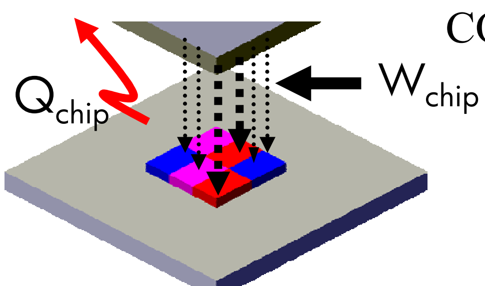

$$COP_{ensemble} = \frac{\text{Total Heat Dissipation}}{(\text{Flow Work} + \text{Thermodynamic Work})\ \text{of Cooling system}}$$

Patel, C.D., Sharma, R.K., Bash, C.E., Beitelmal, M, "Energy Flow in the Information Technology Stack: Introducing the Coefficient of Performance of the Ensemble", ASME International Mechanical Engineering Congress & Exposition, November 5-10, 2006, Chicago, Illinois

Coefficient of Performance of the Ensemble

$$COP_G = \frac{Q_{dc}}{\sum_k \left(\sum_j \left(\sum_i \left(W_{cp} + W_{\sup-dev} \right) \right) + W_r \right) + \sum_l W_{b-cr} + \sum_m W_p + W_{comp} + W_{ct}}$$

$$COP_{ct} = Q_{ct} / W_{ct} = \frac{Q_{dc} + W_{comp}}{W_{ct}}$$

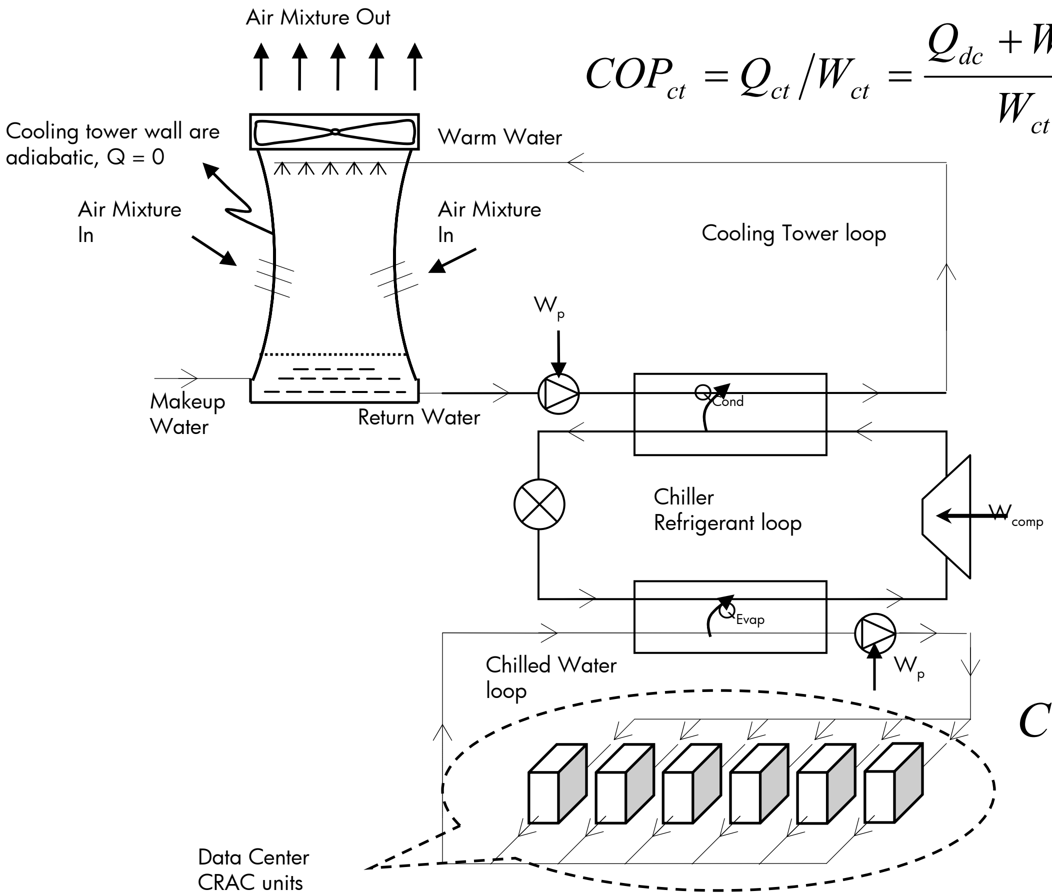

$$COP_{ch} = \frac{Q_{ch}}{W_{comp}}$$

$$W_{comp} = \frac{\dot{m}_{ref} n P_2 v_2}{\eta_p \eta_{motor} (n-1)} \left[\left(\frac{P_3}{P_2} \right)^{(n-1)/n} - 1 \right]$$

$$COP_{hydronics} = Q_{dc} \Big/ \sum \left(W_p \right)_{secondary}$$

Patel, C.D., Sharma, R.K., Bash, C.E., Beitelmal, M, "Energy Flow in the Information Technology Stack: Introducing the Coefficient of Performance of the Ensemble", ASME International Mechanical Engineering Congress & Exposition, November 5-10, 2006, Chicago, Illinois

Impact

Operation

Flow Work + Thermodynamic Work

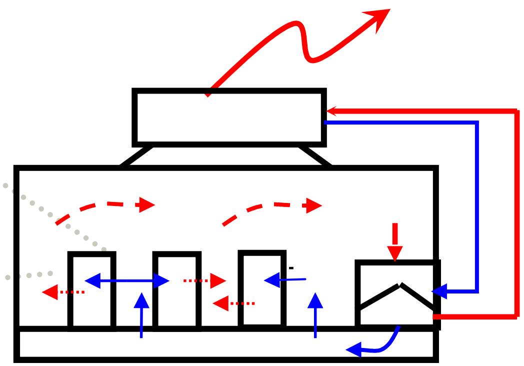

Services →

Heat Generated

10 to 100 W … **2 MW**

Energy to Remove Heat

2-40 W … **2 MW**

Total

48 GW Assume: 4 billion handhelds at 12 W

Global Power Consumption

Assume: 5000 data centers **20 GW**

200 Million Metric Tons of coal

Annualized Destruction of Available Resources – Coal

85 Million Metric Tons

175 Million Metric Tons CO_2 emission

Data Center Total Cost of Ownership

$$Cost_{total} = \left(\frac{\$}{ft^2}\right)\left(A_{critical}, ft^2\right) + \left(1 + K_1 + L_1 + K_2 L_1\right) U_{\$,grid} P_{\substack{consumed\\hardware}} + R\left(M_{total} S_{avg} + IT_{dep} + \sigma_1\right)$$

Real Estate

Burdened power consumption

Personnel, equipment, SW per rack

J_1 : capacity utilization factor, i.e. ratio of maximum design (rated) power consumption to the actual data center power consumption

$K_1 = F(J_1)$: burdened power delivery factor, i.e. ratio of amortization and maintenance costs of the power delivery systems to the cost of grid power

$K_2 = F(J_1)$: burdened cooling cost factor, i.e. ratio of amortization and maintenance costs of the cooling equipment to the cost of grid power

Depreciation factors

L_1: cooling load factor, i.e. ratio of power consumed by cooling equipment to the power consumed by compute, storage and networking hardware (inverse of $COP_{ensemble}$)

Factors that impact the "Burdened Power Cost"

Lack of vital information from single-input single-output environmental control systems

- Responsiveness → $\downarrow J_1$
- Energy Consumption → $\downarrow L_1$

Patel and Shah, Cost Model for Planning, Development and Operation of a Data Center
http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.html

invent

Smart IT

IT-Facility End to End Management

compute, power, cooling resources provisioned based on the need

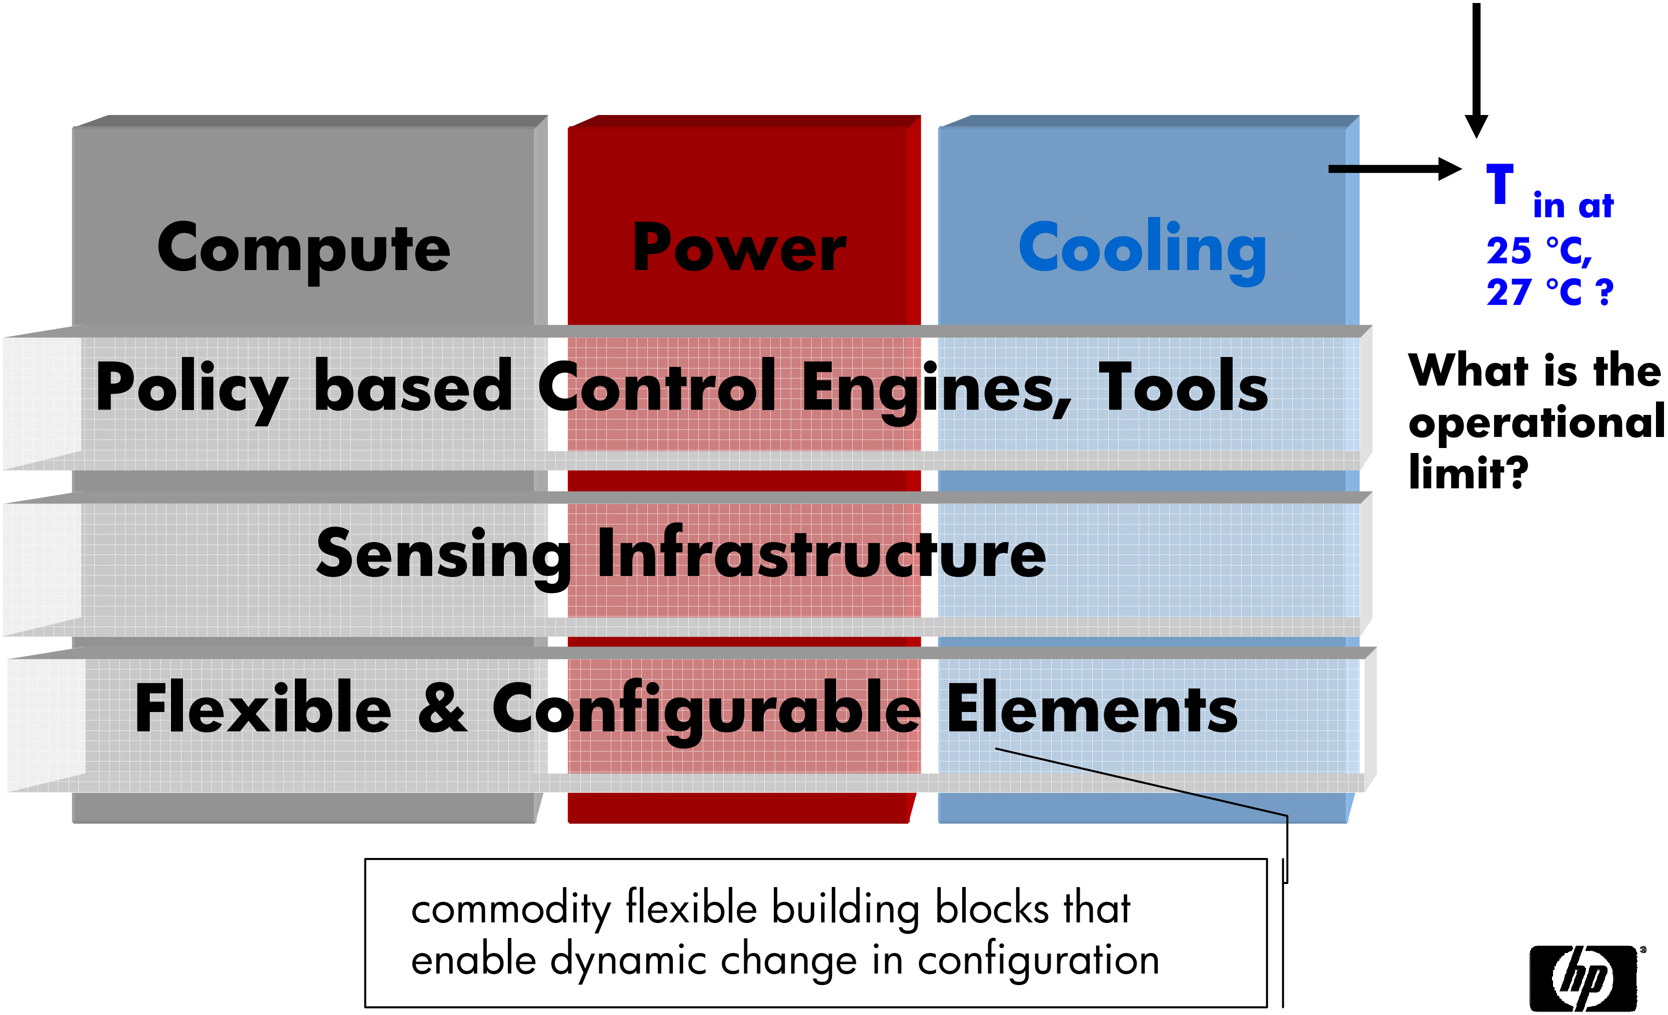

Dynamic Smart Cooling Architecture

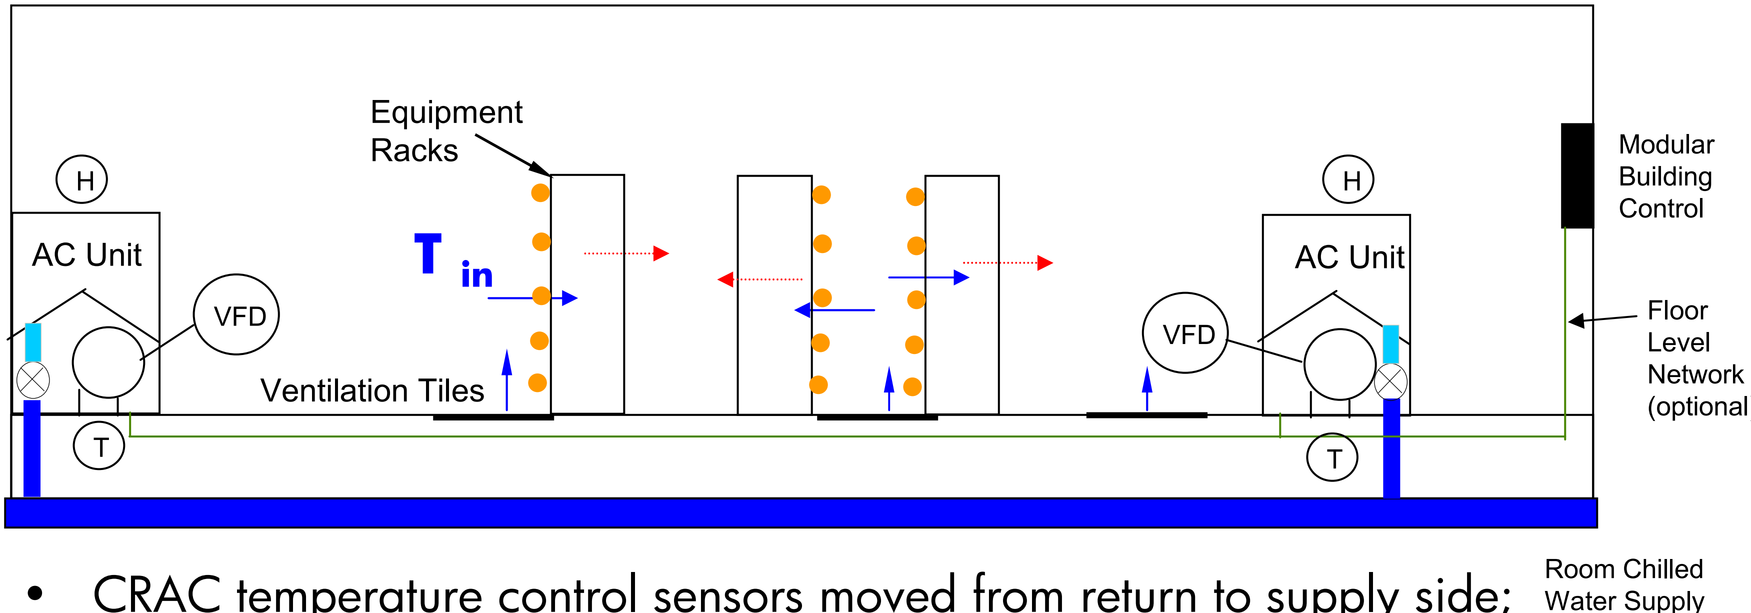

- CRAC temperature control sensors moved from return to supply side;
- Variable Frequency Drives added to AC units;
- Temperature sensors added to rack inlets;
- Advanced algorithms control operation of AC units.

HP Labs Smart Data Center

Palo Alto, California

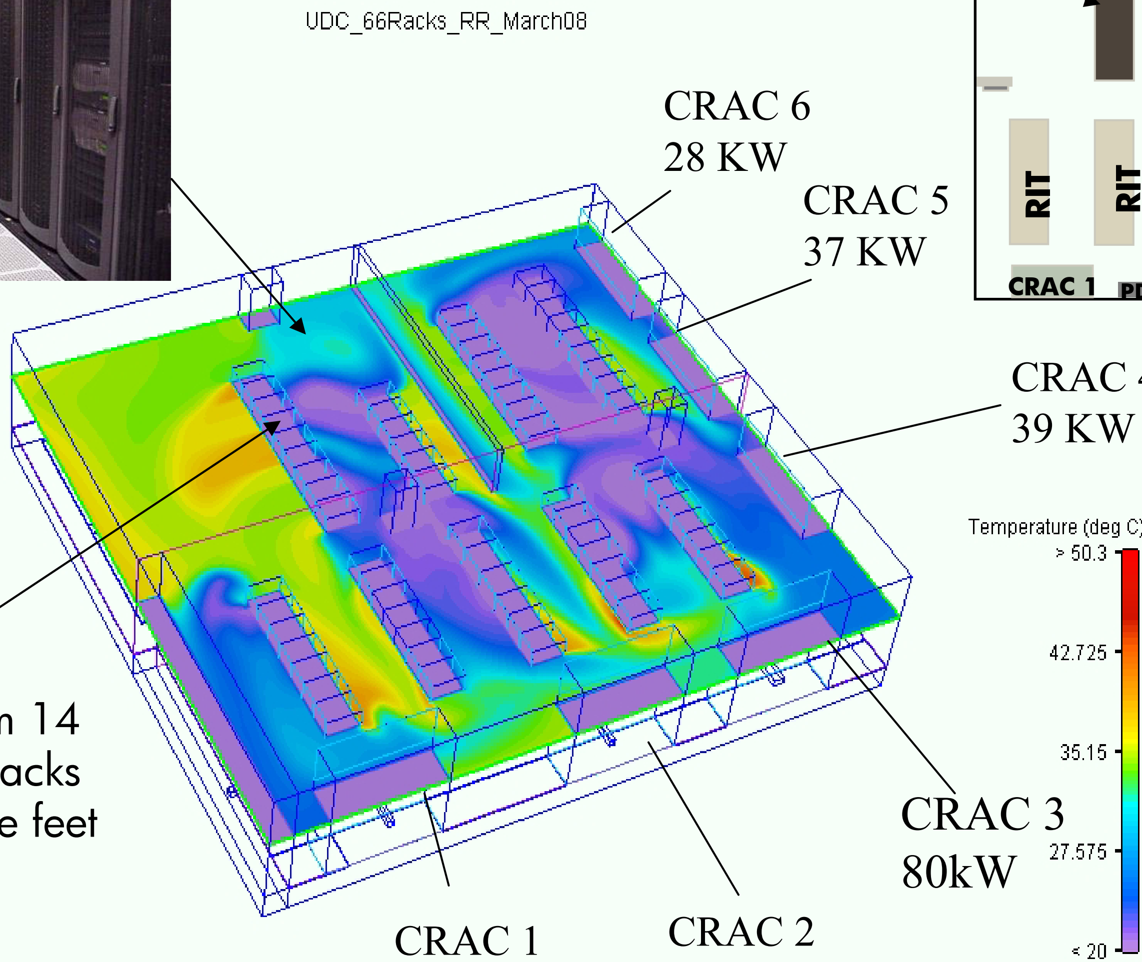

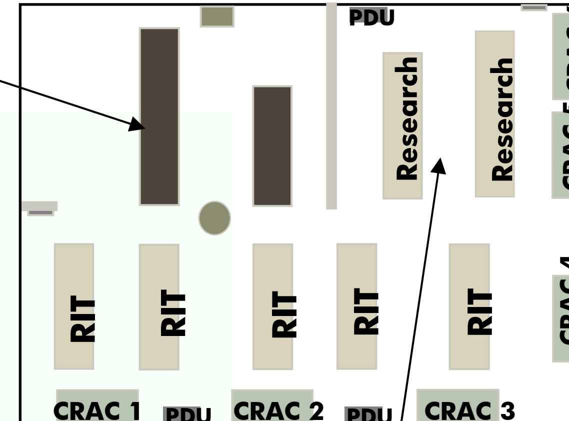

HD Area

CRAC 6
28 KW

CRAC 5
37 KW

CRAC 4
39 KW

High Performance Cluster – low load in this state

HD Area:

140 KW from 14 fully loaded racks in 400 square feet

CRAC 1
94KW

CRAC 2
92 KW

CRAC 3
80kW

HP Labs Data Center

Palo Alto, CA

Conventional Mode

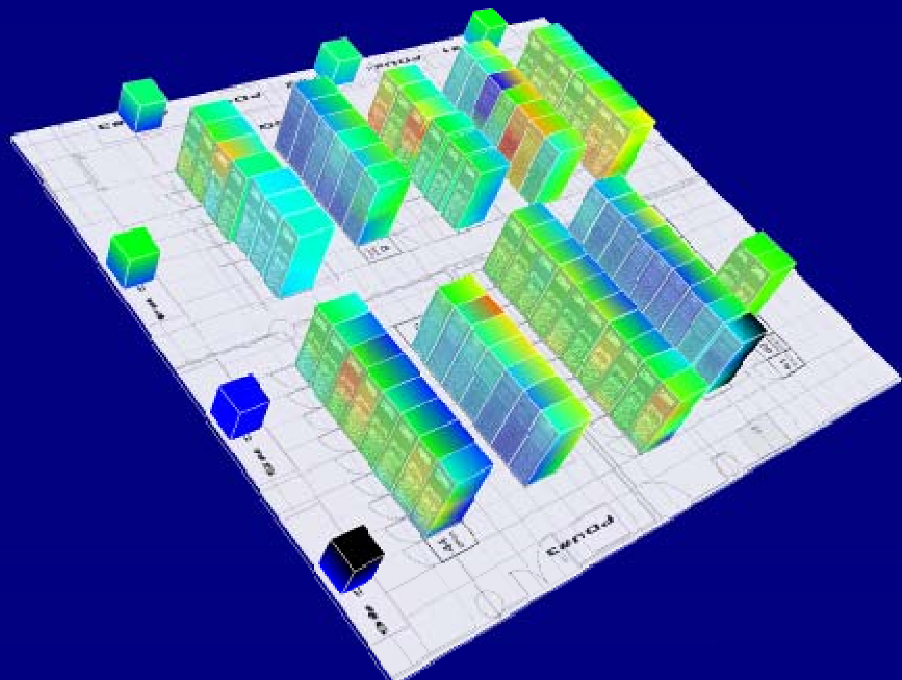

Dynamic Smart Cooling Mode

- **35% Energy Savings**
- **Improved reliability**

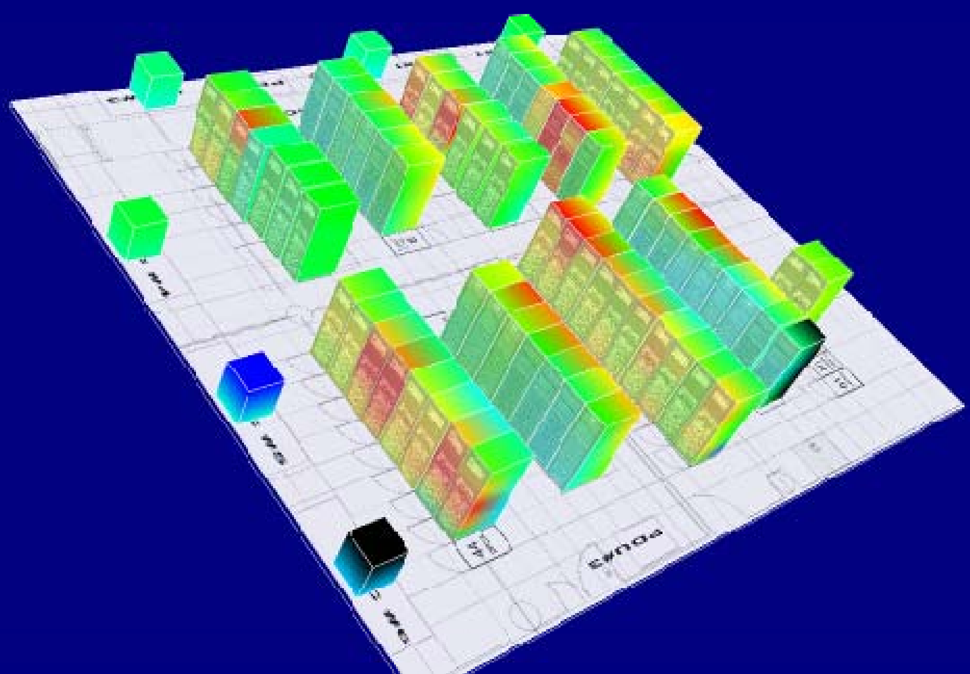

Local Disturbances – Vent Tile Blockage

Scenario: IT department performed maintenance on a rack and covered nearby vent tiles for 4 hours with bag. DSC responds and keeps impacted racks below maximum.

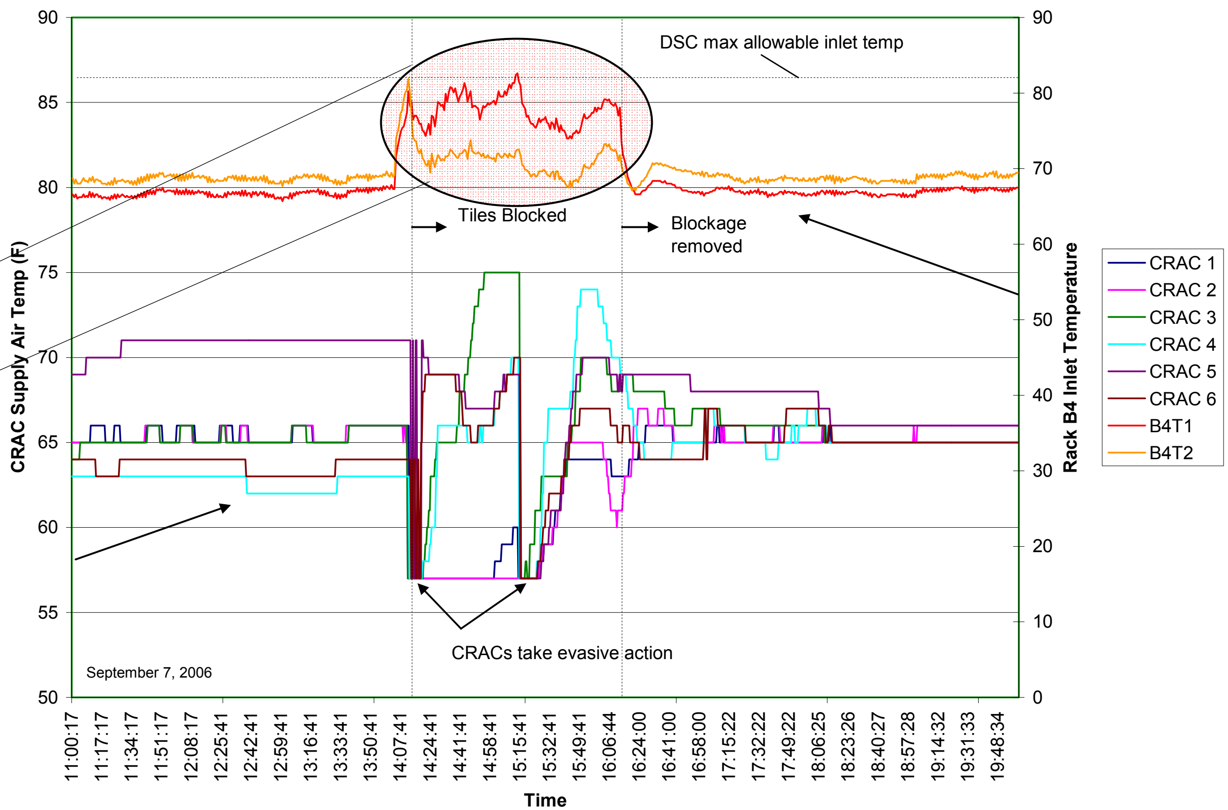

Global Disturbance – Chiller Failure

Scenario: Groundskeeper blows leaves into primary cooling tower during routine cleanup. Leaves block filter in cooling tower retarding condenser water flow that ultimately results in chiller failure. DSC delayed the impact of the failure allowing for repair.

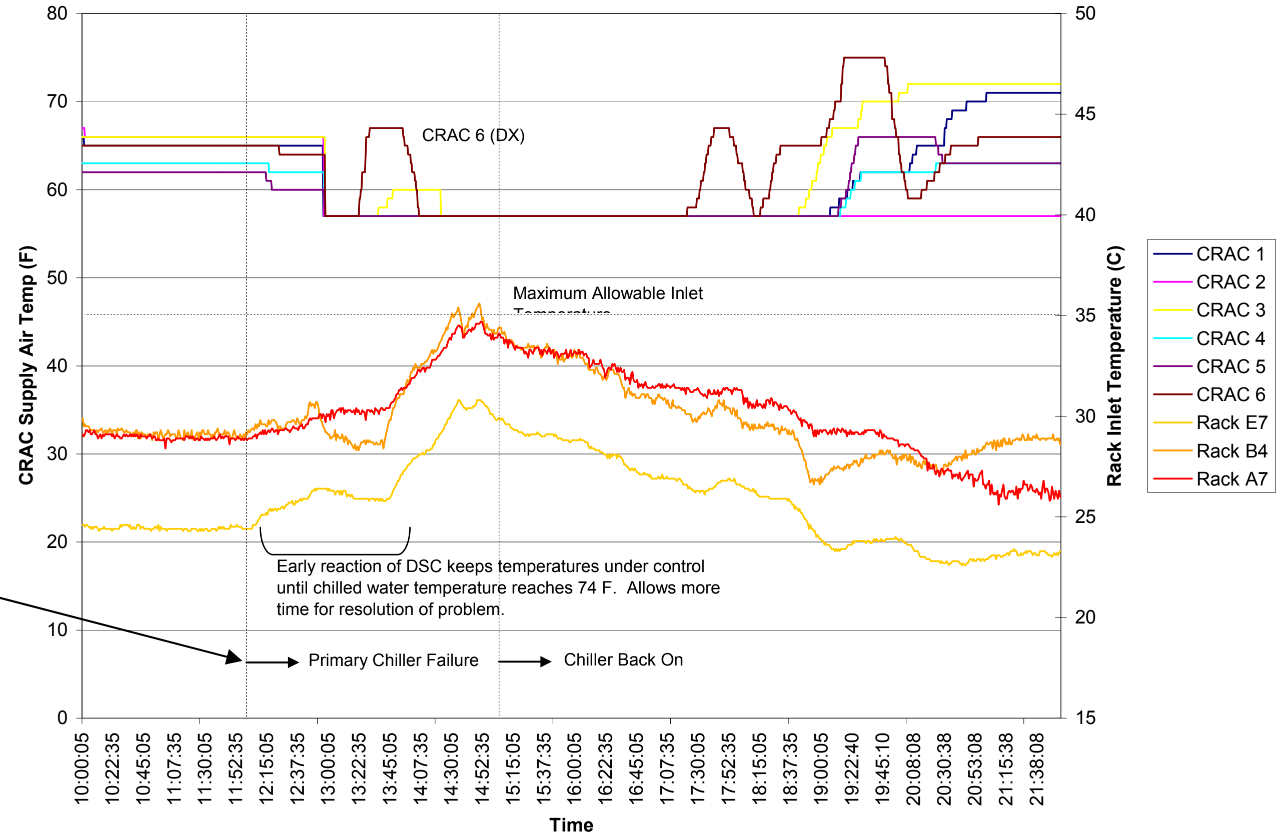

HP R&D Lab-Data Center

Dynamic Smart Cooling Implementation in Bangalore

Facility Building Blocks




- **Chillers**
  - 3 air-cooled
  - 2 water-cooled
- **Pumps**
  - 7 Primary
  - 5 Secondary
- **CRAC units**
  - 55 units
- **Diesel Generators**

- Scope/Sizing
- Layout/Topology
- Comm. Fabric/Integration
- Control and Monitoring
- Management

IT Building Blocks

- Servers
  - Non-Stop servers
  - Proliant servers
  - Blade servers
  - Custom Enclosures
- Storage (XP/EVA)
- Multiple Network topologies
- Sensor Network
  - 7500 sensors

5 floors @14k sq. ft.
900kW cooling per floor

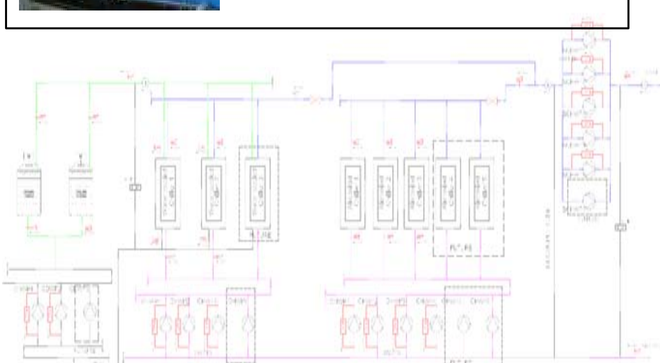

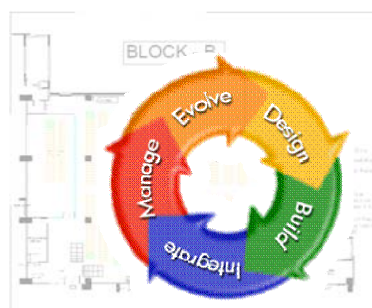

Supply Side and Demand Side Management

Demand Side (Consumption of Resources):

- Library of *flexible* and *scalable* building blocks overlaid with *pervasive sensing* and *control* to provision resources based on the need

Supply Side (Availability of Resources)

- Devise a 2nd law based tool that enables a "cradle to cradle" approach in quantifying and designing the IT ecosystem
- Can it be used to:
  - Dematerialize the ecosystem i.e. least materials design
  - Seek out appropriate materials
  - Seek out appropriate distributed energy sources
  - Seek out analytical techniques to address operations and "end of life"
    - **leverage** the rich instrumentation used for demand side management to detect and understand anomalies

Dematerializing the Ecosystem

Data Center Management: modeling, measurement and inference

manufacturing	operation	retirement

Reduce redundancy, manage "end of life":

- Minimize redundancy in the data center facility infrastructure
  - Example: empirical data and inference techniques to eliminate excessive redundancy e.g. standby air conditioners
- Hardware "Damage Boundary*" [1] and minimizing hardware redundancy
  - IT-Facility measurement and inference techniques in place, can we push the limits of operation
    - Example: Understanding the impact inlet temperature T_{in} to life of components by having thermo-mechanical models in place
    - Understanding thermo-mechanical behavior to determine root cause of failure and manage "end of life"
- Drive reliability studies have been presented before at USENIX using large samples, SMASH data – opportunity to extend the work [2][3]

*Used for fragility assessment – can this be extended to other areas to determine operational limts

[1] Lee Hedtke, Chandrakant Patel, Damage Boundary Testing of Hard Disk Drives, Proceedings of the ASME Winter Annual Meeting, 1990

[2] Pinheiro, Wolf-Dietrich Weber, Luiz André Barroso. Failure Trends in Large Disk Drive Population. *In Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST '08),* February 2007.

[3] Weihang Jiang, Chongfeng Hu, Yuanyuan Zhou and Arkady Kanevsky, Are Disks the Dominant Contributor for Storage Failures, 6th USENIX-FAST, February 2008

Rigid Drive Mechanism (Thermal)

manufacturing	operation	retirement

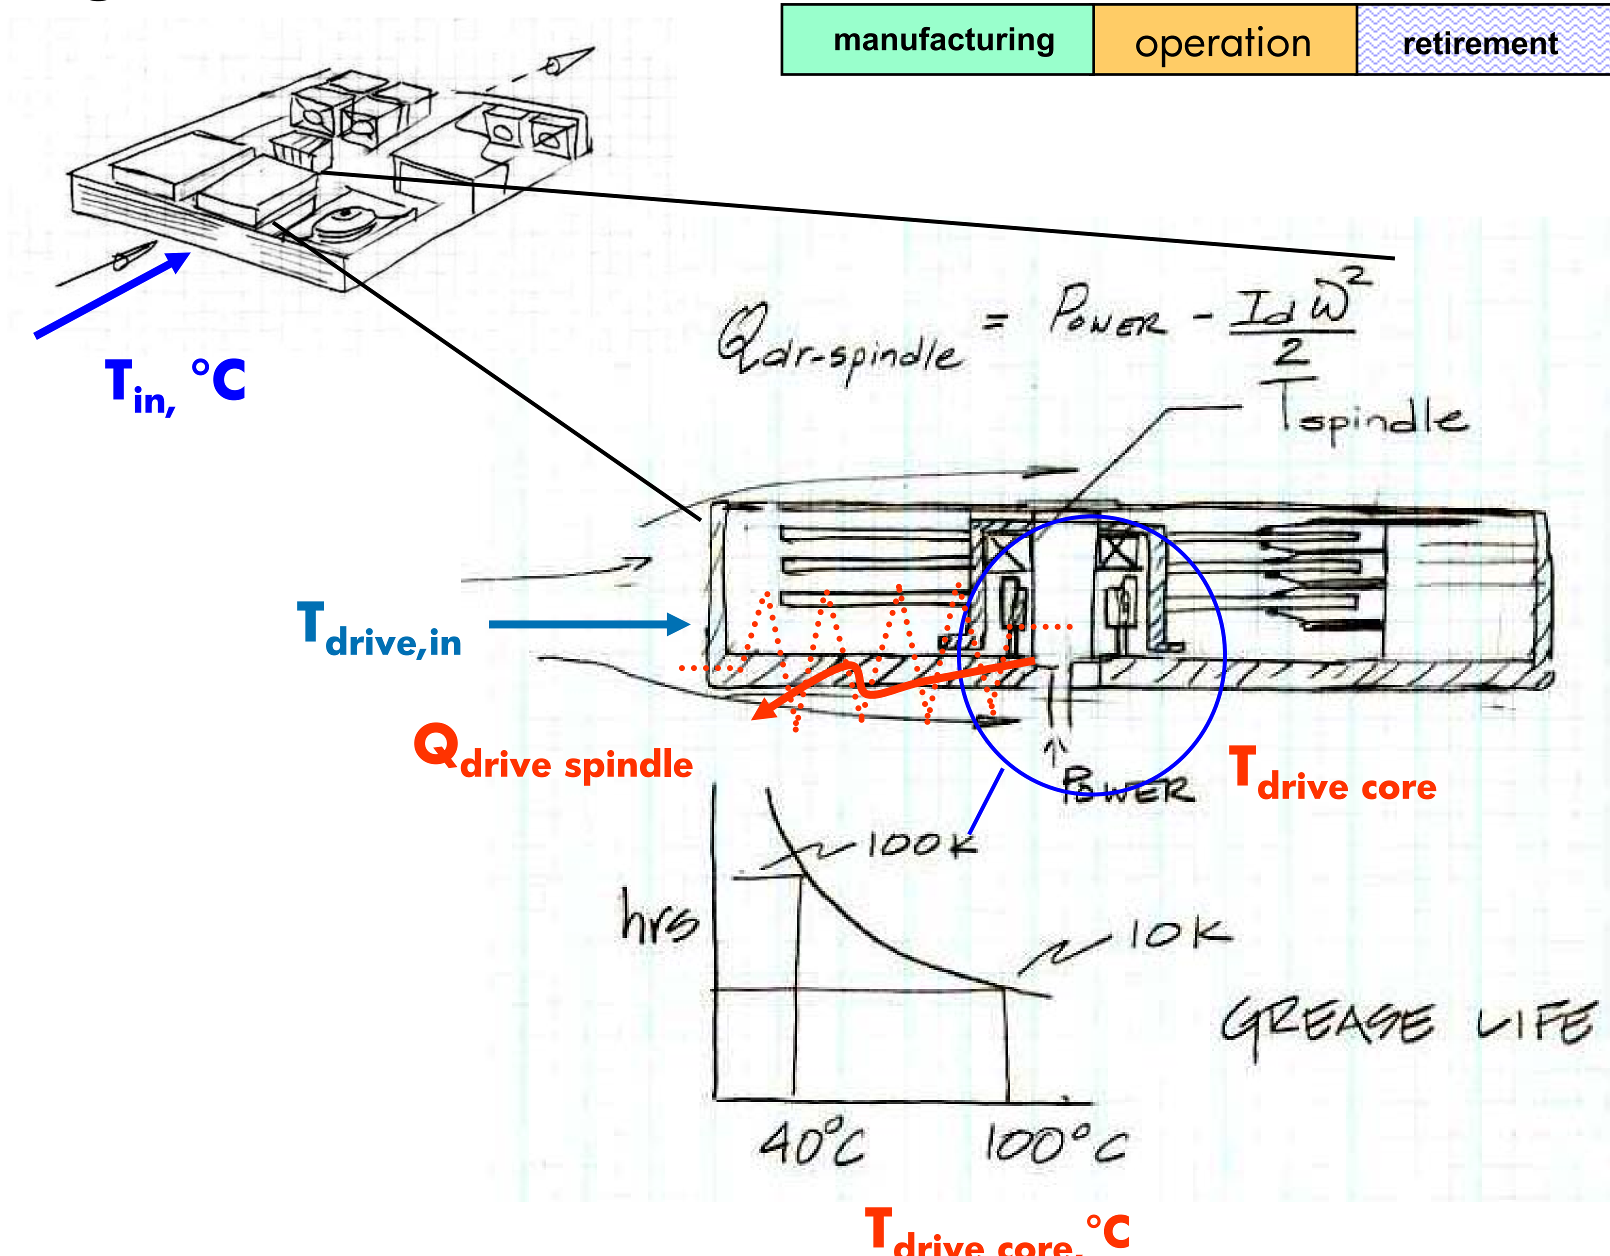

 Chandrakant Patel, chandrakant.patel@hp.com

Thermo-Volume Resistance

Expedient Thermo-Fluids Modeling of Systems

Can we represent sub-systems as "thermo-volumes" with thermo-mechanical attributes?

"Thermo-Volume" [Ref]
Resistance

System Blower (s)

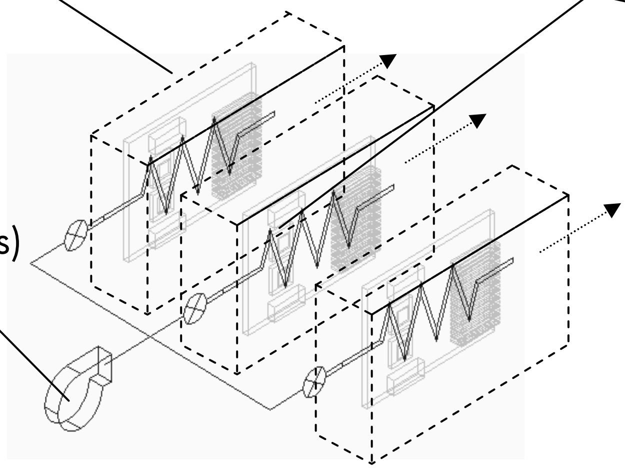

T_{chip}

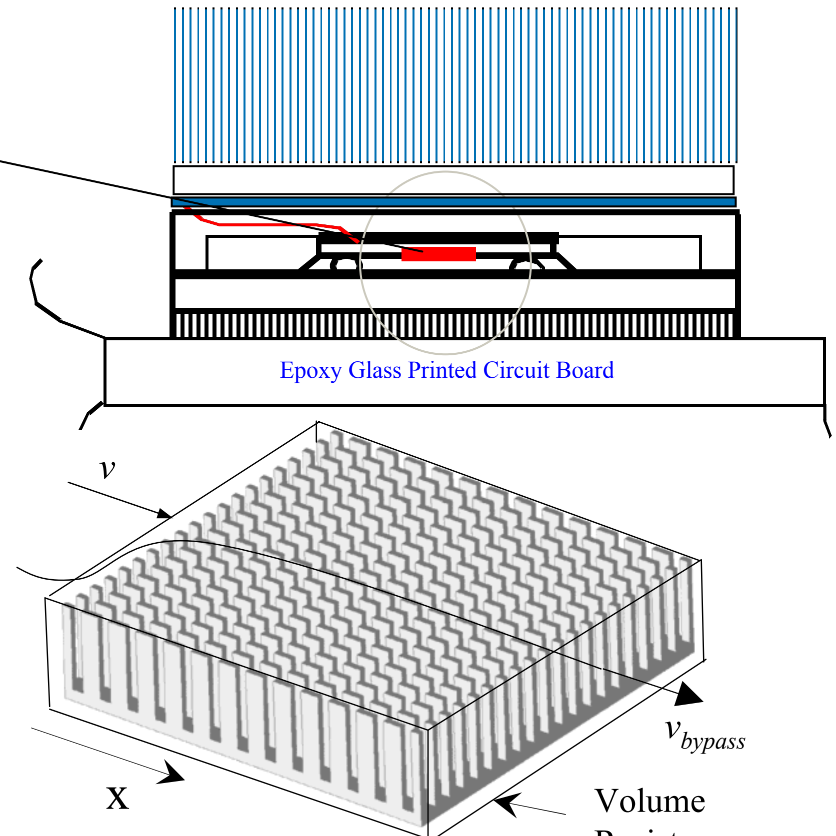

$$\frac{\partial P}{\partial x} = \frac{1}{2}\frac{f_l}{d_{duct}}\rho v_{duct} + \frac{1}{2}\frac{f_t}{d_{duct}}\rho v^2_{duct}$$

Extract the temperature at a given location

e.g. chip, hard drive temperature

- Patel and Belady, ECTC1997
- Bash and Patel, ISPS1999
- Patel and Bash, Thermal Management Course Notes, 2007

invent

Rigid Drive Mechanism (Mechanical)

retirement

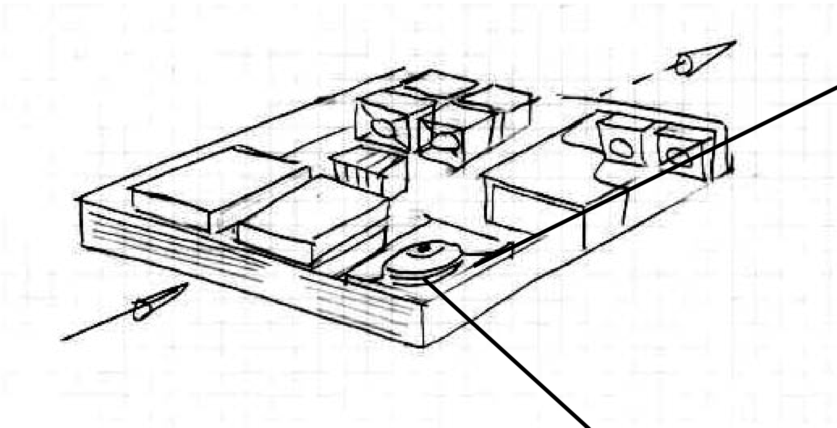

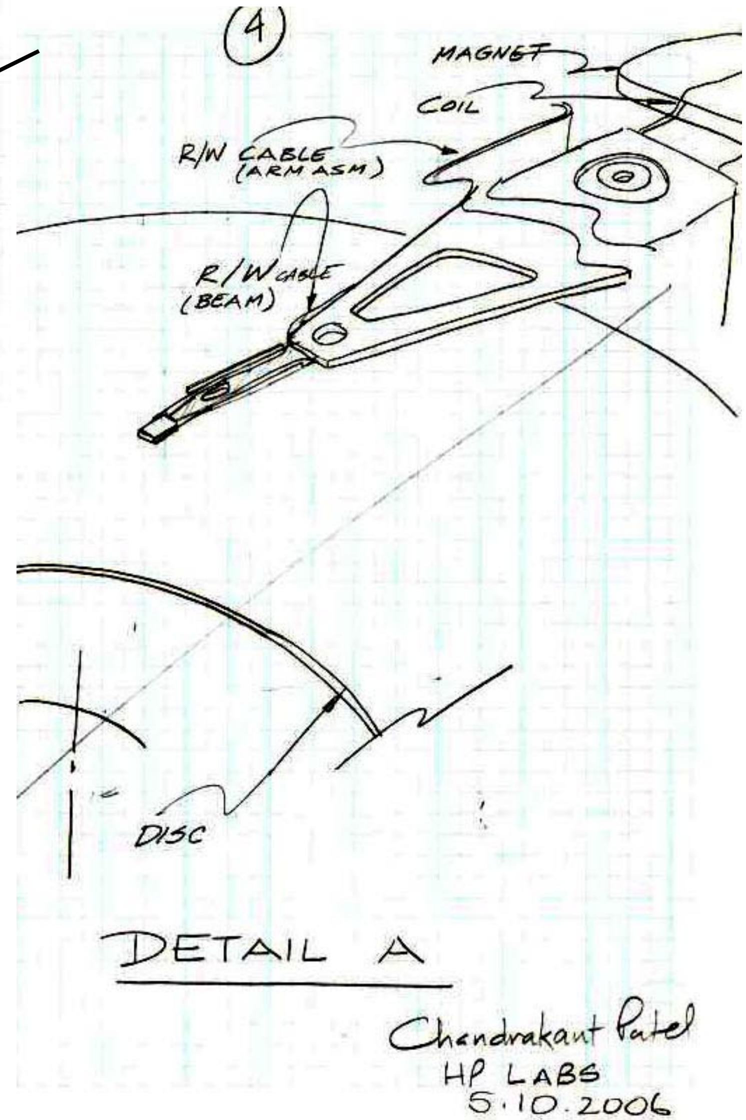

Susceptibility to vibration:

- over the years, stiffness to mass ratio has gone up due to miniaturization
- however, in-situ disturbances in data centers can coincide with natural frequency of the arm

Dematerializing the Ecosystem

Lifetime Exergy Advisor - Case Study of Sample PC

manufacturing

Extraction Data

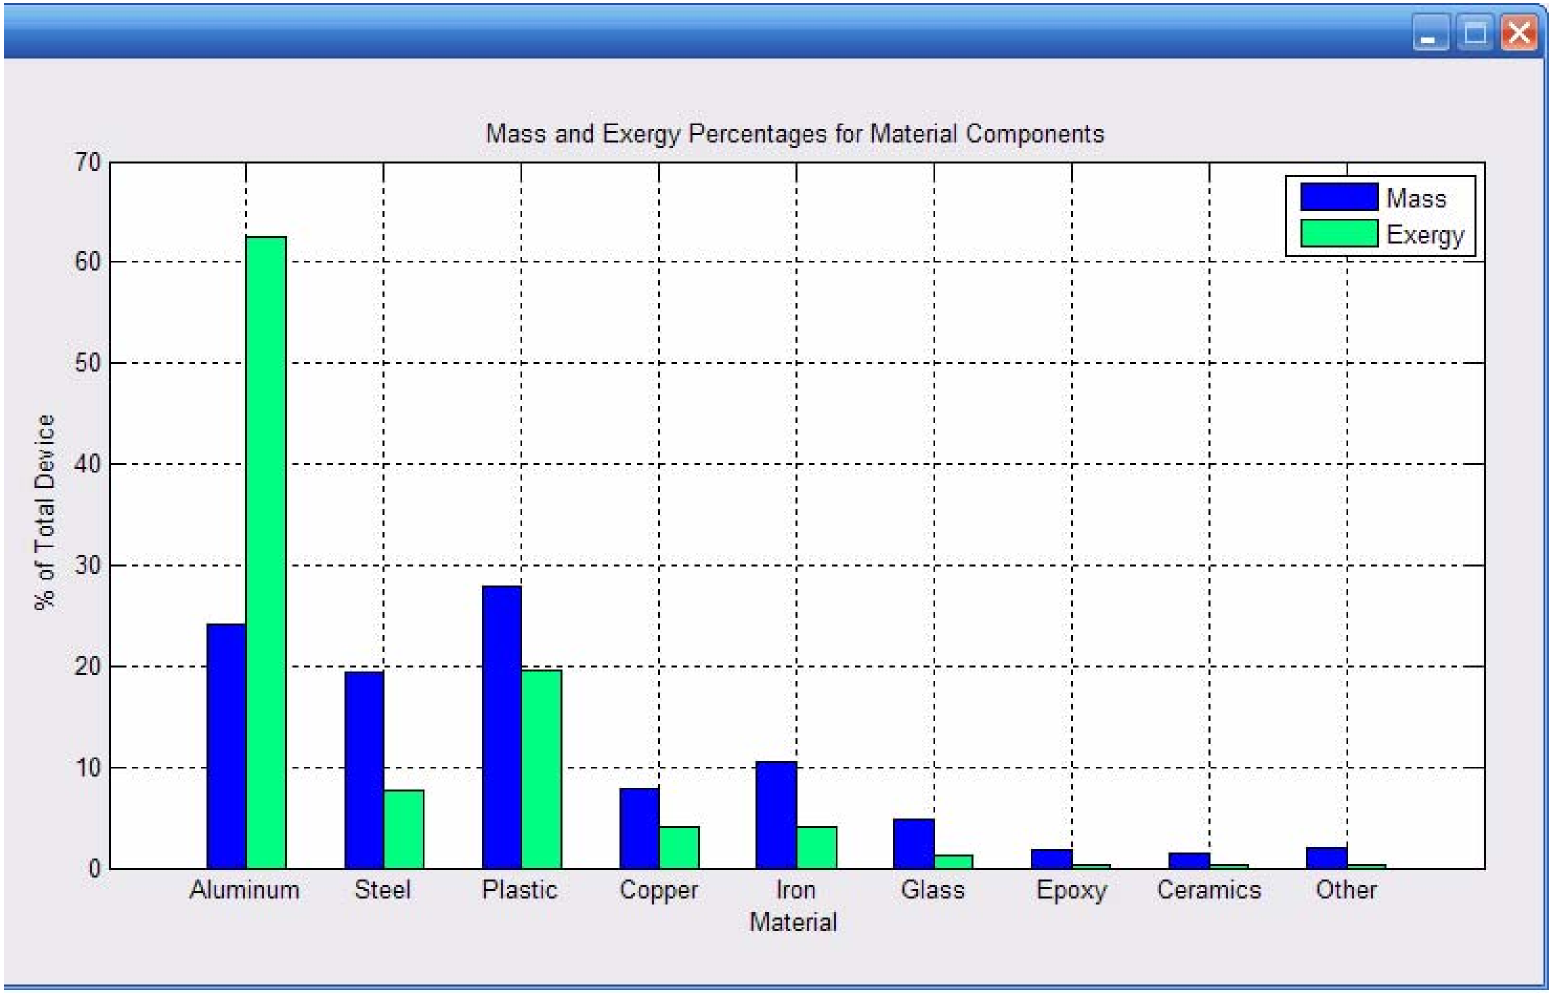

Ref: HP Labs – UC Berkeley data

Joules: Currency of the Flat World

ecosystem: billions of handhelds and printers, thousands of data centers and print factories

Cradle to Cradle Design: Least Material, Appropriate materials based on the 2nd Law of Thermodynamics

Least Energy through need based provisioning of resources

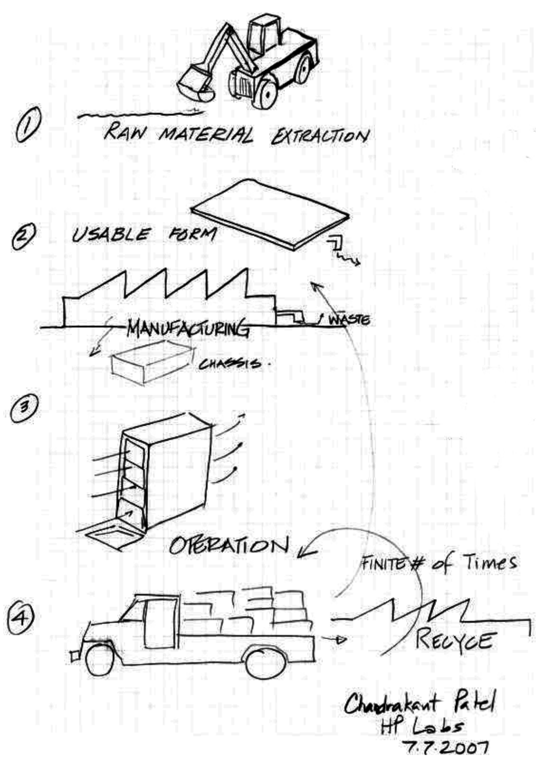

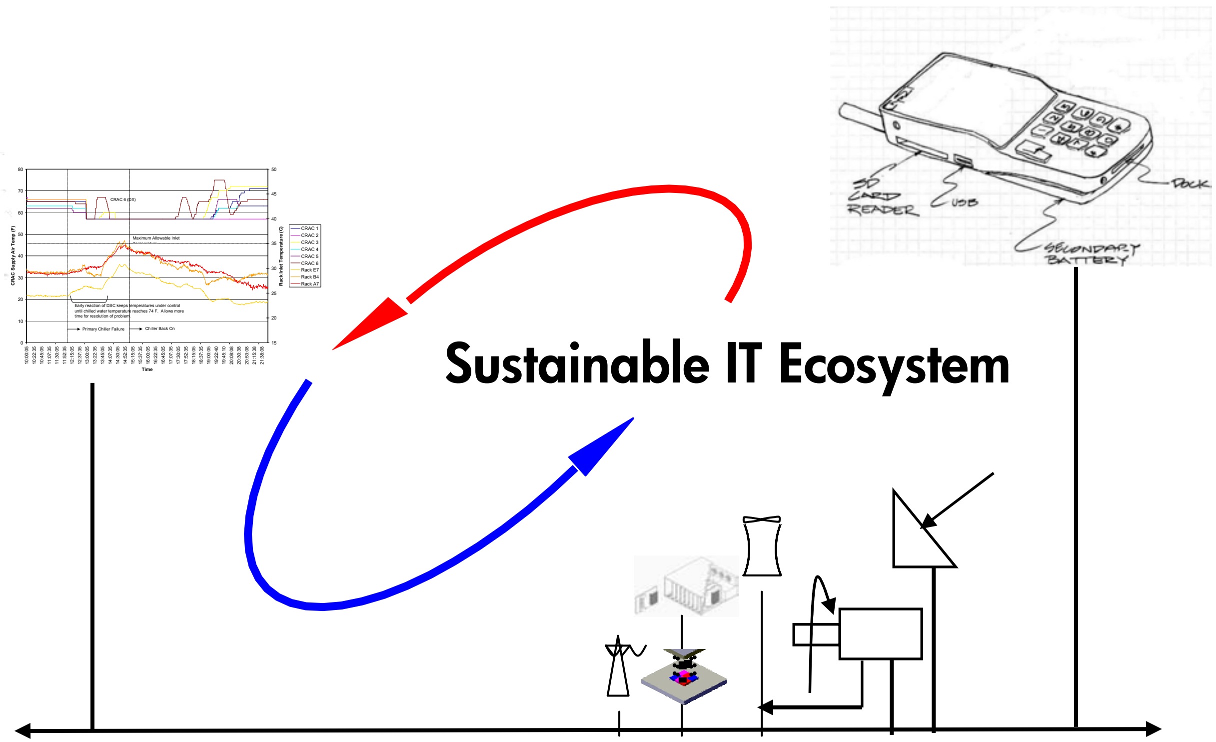

Opportunity for collaboration CS-ME-EE

Thank You

- I thank the USENIX-FAST organizers for giving me this opportunity
- I thank my HPL colleagues

&

- I thank you for your time