Energy Efficient Enterprise Computing and Green Datacenters

Massoud Pedram
University of Southern California
Dept. of Electrical Engineering
Oct. 8, 2010
Workshop on IT and Future Society
Jeju Island, South Korea

Outline

- ICT Ecosystem and ICT Sustainability
- Lifecycle Analysis and Environmental Footprint of ICT
- Cloud Computing as Today's Driving Platform
- Energy Use and Peak Power Crisis
- Power Minimization and Management Strategies and Solutions
- Emergence of Smart Grid and Dynamic Energy Pricing
- Conclusion

Internet Usage Trend

- People, organizations, and governments are becoming increasingly dependent on the Internet and information and communications technologies (ICT)
- As ICT becomes more deeply integrated with a myriad of societal infrastructure, ICT failures and inefficiencies cascade through many other critical infrastructures, making such problems untenably costly

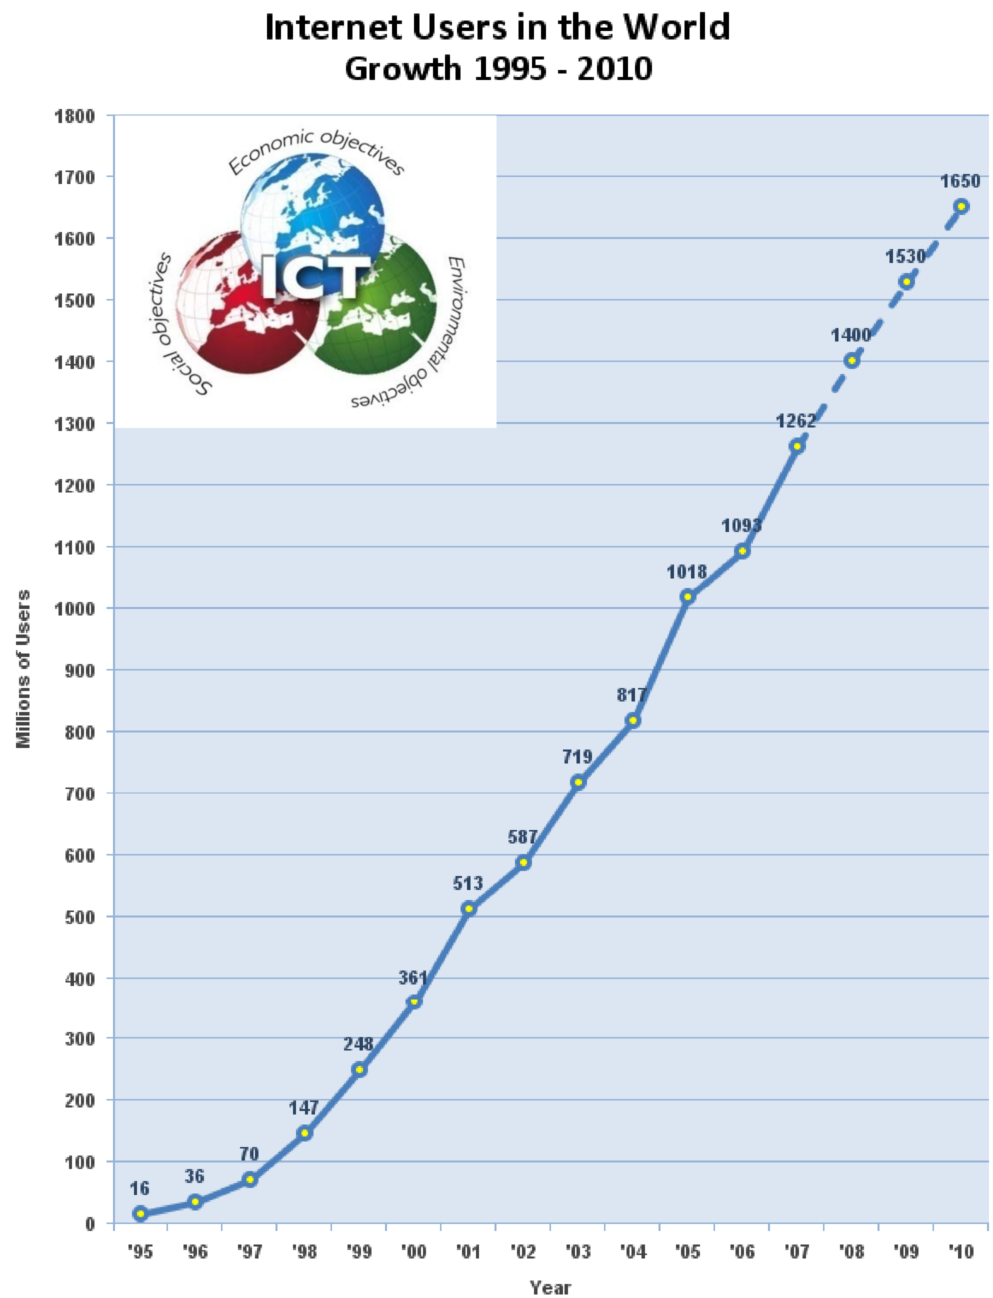

Source: www.internetworldstats.com - January, 2008
Copyright © 2008, Miniwatts Marketing Group

The Internet of Things

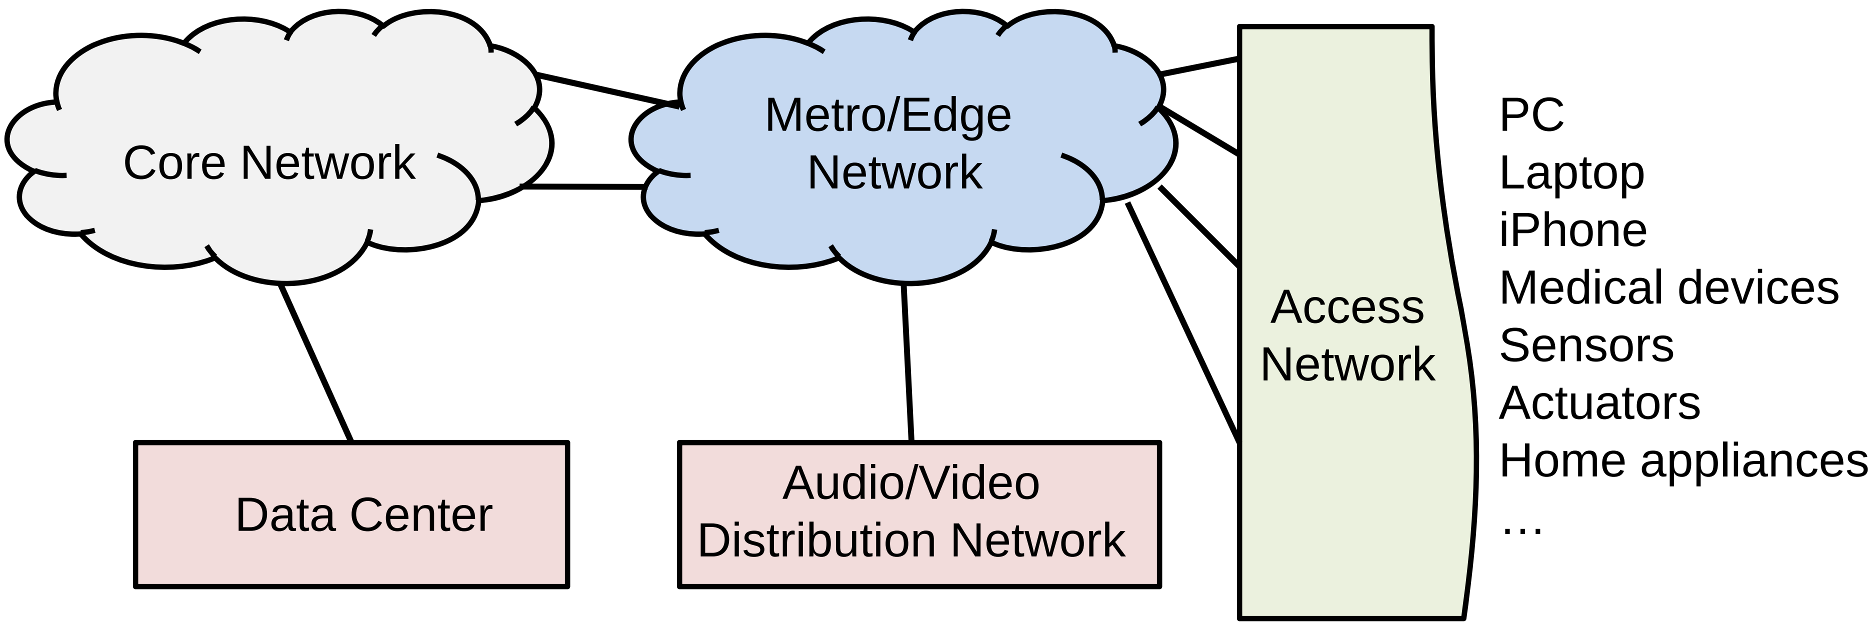

The ICT Ecosystem

Servers and thin mobile clients

Cloud computing and switching centers

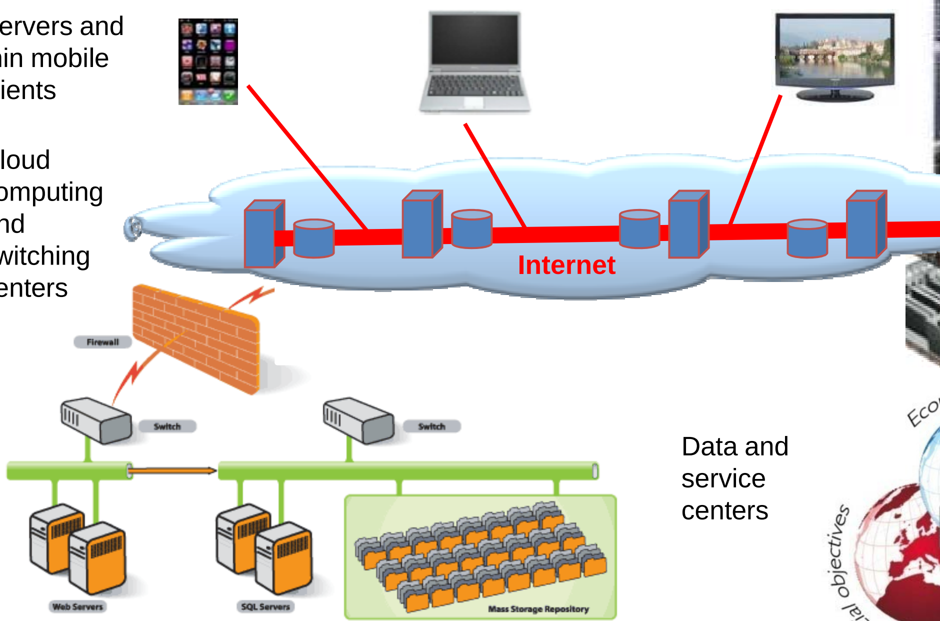

Data and service centers

The ICT ecosystem encompasses the policies, strategies, processes, information, technologies, applications and stakeholders that together make up a technology environment for a country, government or an enterprise.

--Source: Harvard Law

ICT Sustainability

- ICT sustainability goes beyond gains in energy efficiency; Instead, the ICT must offset its own (growing) environmental impacts (e.g., carbon footprint) on an absolute basis, by reducing its total pollution and consumption of resources
  - The ICT can play an important role in offsetting environmental impacts created by other much larger portions of overall electricity demand in the nation
  - Note however a “rebound effect” in which several indirect effects stemming from increased efficiencies increase consumption, thus overwhelming savings
  - Watch for the “free rider” problem

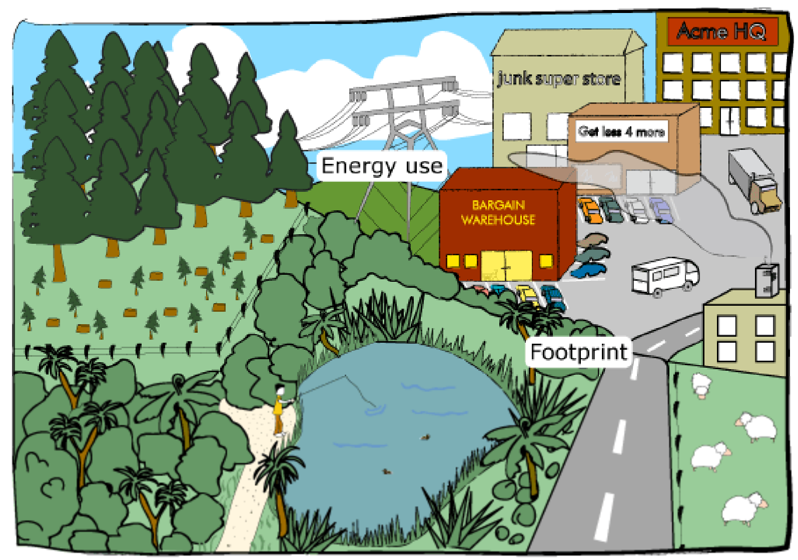

Life Cycle (Cradle-to-Grave) Analysis

- The analysis of the full range of environmental impacts of a product requires the assessment of raw material production, manufacture, distribution, use, and disposal
- Life cycle energy analysis (LCEA) is an approach in which all energy inputs to a product are accounted for, not only direct energy inputs during manufacture, but also all energy inputs needed to produce components, materials and services needed for the manufacturing process
- Net energy content is the energy content of the product minus energy input used during extraction and conversion, directly or indirectly.
  - Different energy forms (heat, electricity, chemical energy etc.) have different quality and value
  - A joule of electricity is 2.6 times more valuable than a joule of heat or fuel

> We must find a way to meet the increasing demand for energy without adding catastrophically to greenhouse gases.
> —Ray Orbach, Undersecretary for Science, U.S. Department of Energy

CO_2 Footprint of Some Consumer Devices

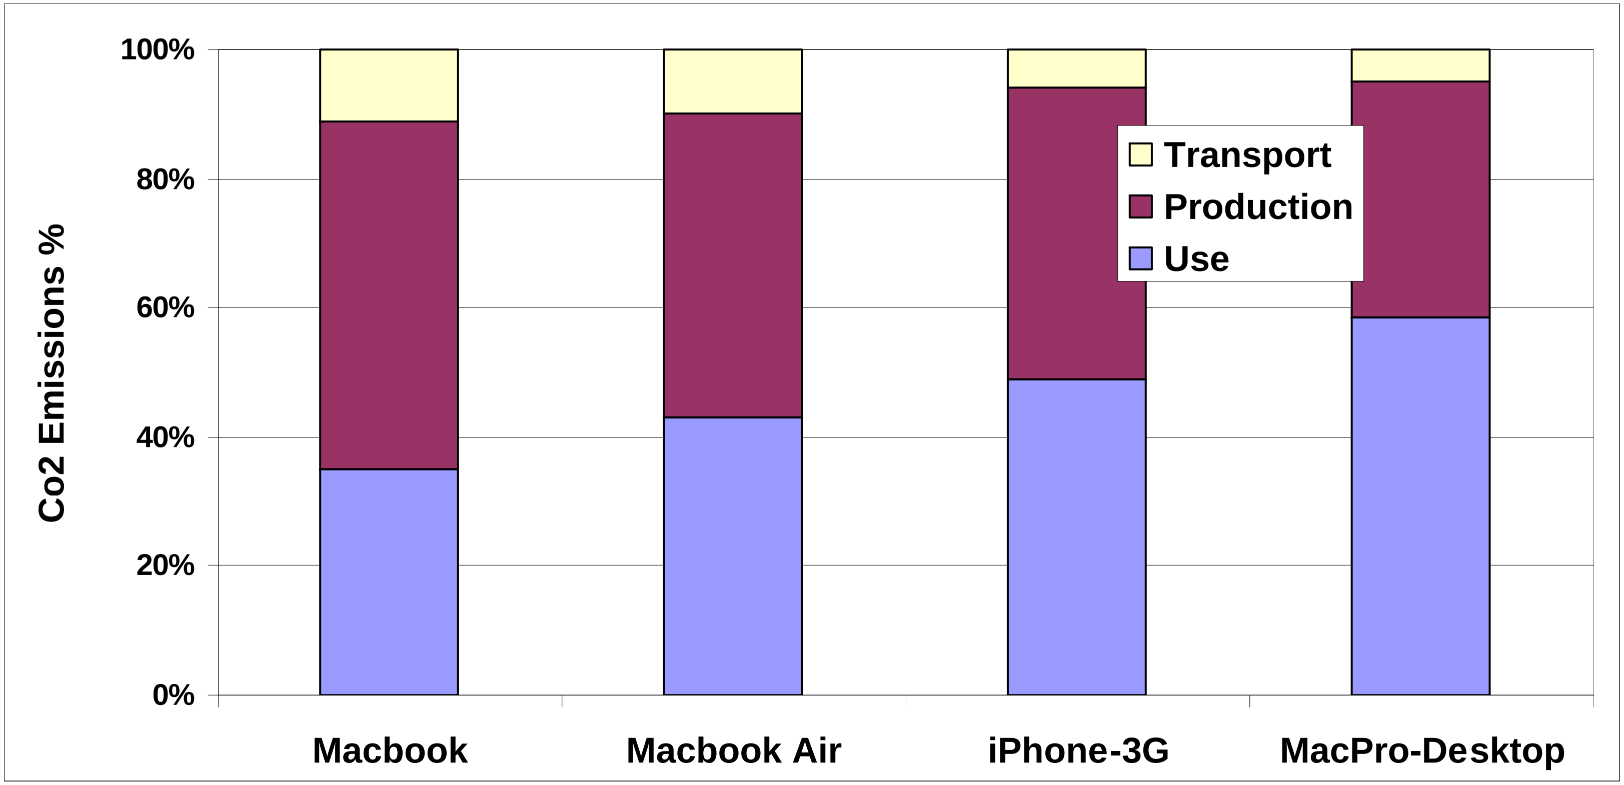

- CO_2 emissions are a good metric for comparison
- Need different strategies for each market segment

ICT Industry's Carbon Emissions in the U.S.

The ICT sector must ensure its own carbon footprint grows slower than the demand for its products

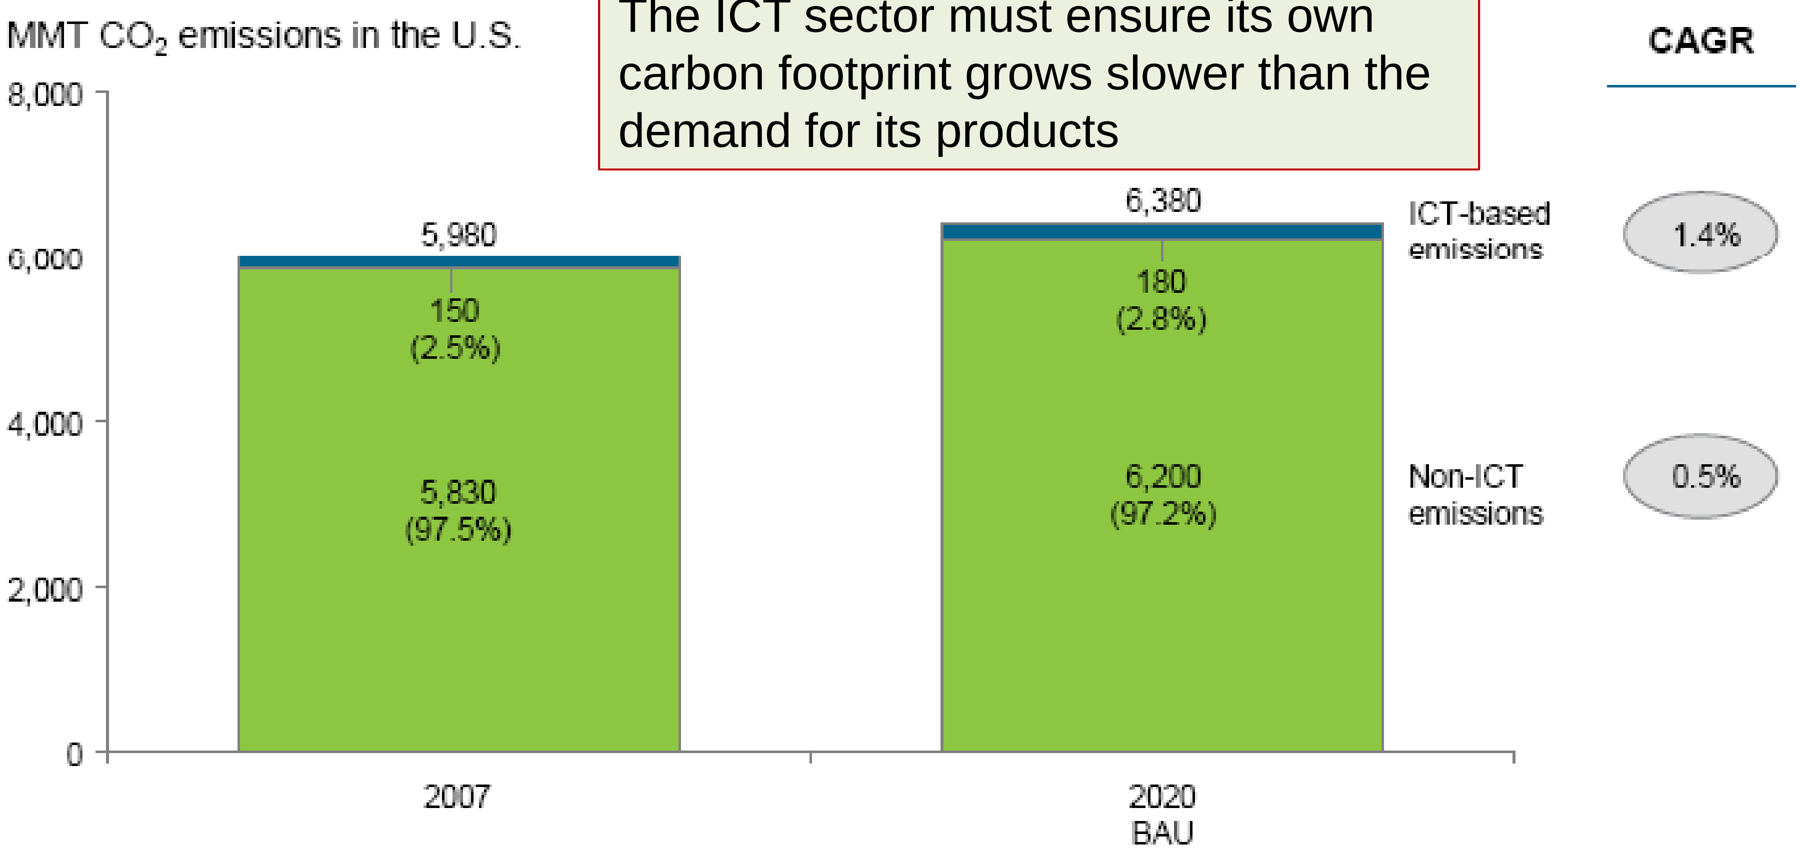

Sources: GeSI and the Climate Group, "Global Smart 2020 report", 2008; EIA Annual Energy Outlook 2008.

Target 2020

California's Global Warming Solutions Act of 2006
– Requires Reduction of GHG by 2020 to 1990 Levels
– 10% Reduction from 2008 Levels; 30% from BAU 2020 Levels
The European Union Requires Reduction of GHG by 2020 to 20% Below 1990 Levels (12/12/2008)

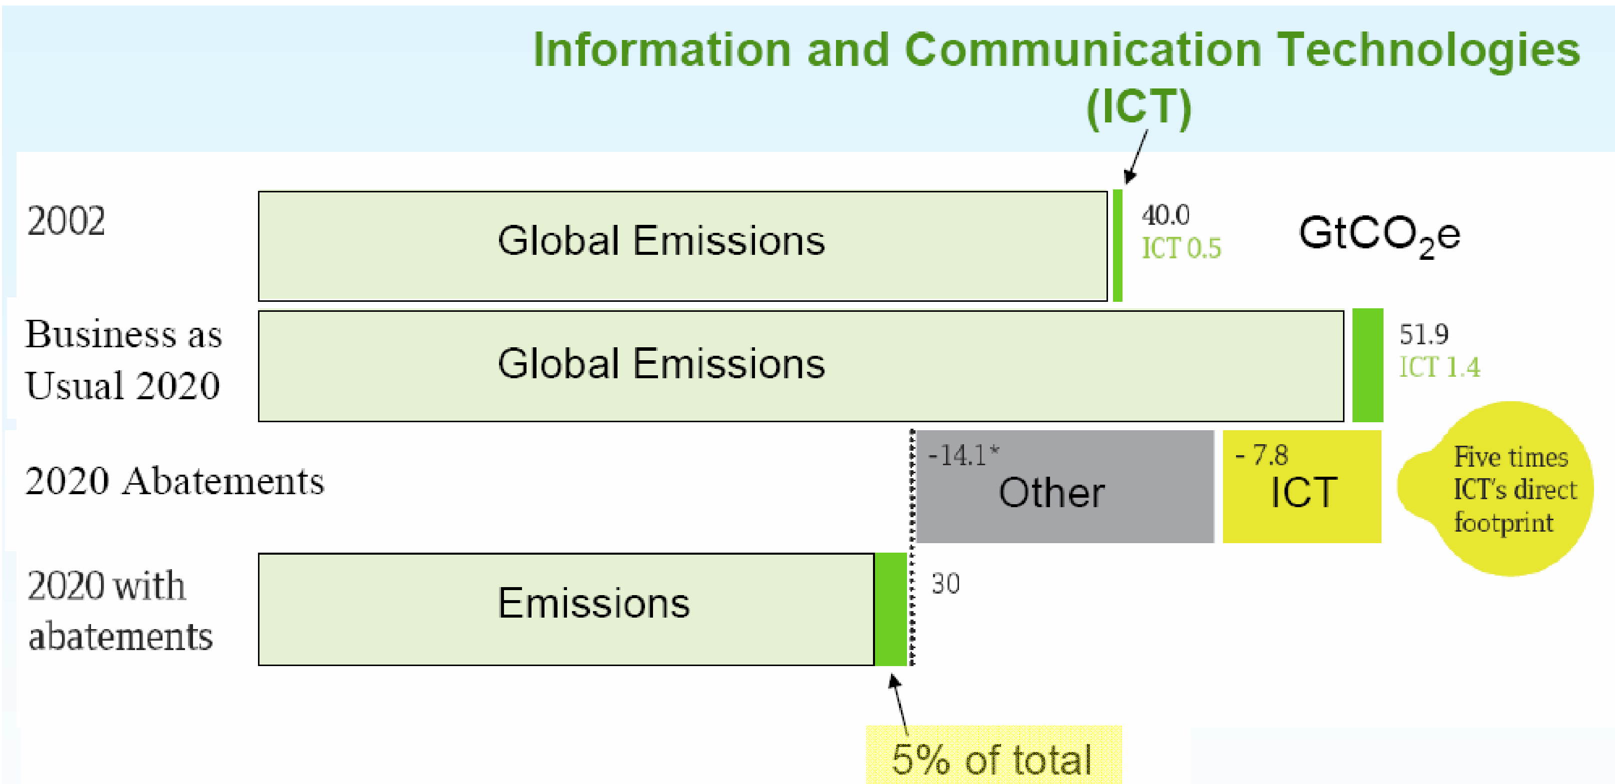

Green House Gases: Reduction Potential with ICT

	MMTCO2	Reduction With ICT
Transporta	1892.2	459.8
Industrial	1553.4	455.7
Residentia	1198	395.9
Commerci:	1041.4	344.1
Total	5685	1655.4

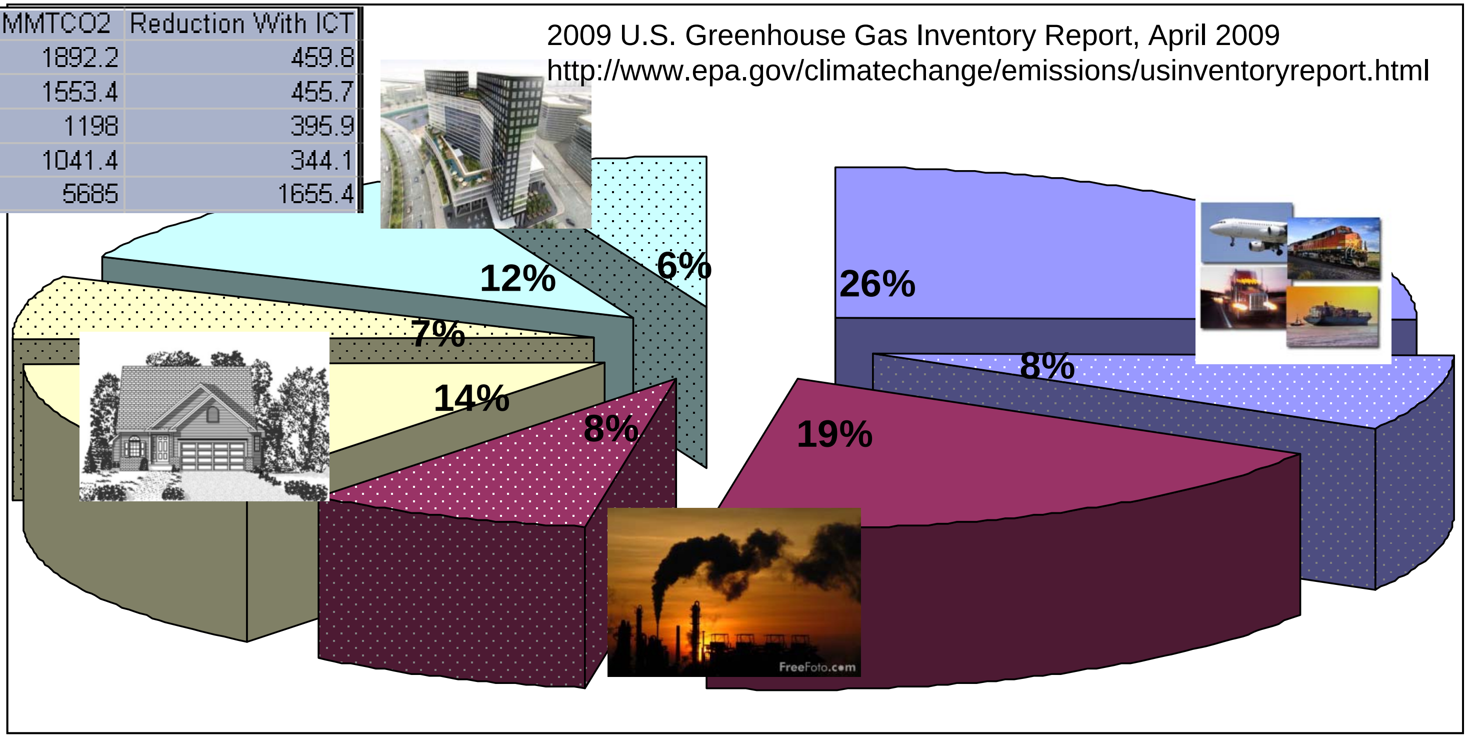

- Rising energy consumption and growing concern about global warming
  - World wide GHG emissions in 2006: 27,225 MMT of CO_2 Equivalent
- Major contributors: Transport, Industrial, Residential, Commercial
- Improved efficiency with Information Technology (IT) usage
  - 29% expected reduction in GHG
  - Equal to gross energy and fuel savings of $315 billion dollars

Cloud Computing as Today's Driving Platform

- Cloud computing provides a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service-provider interaction

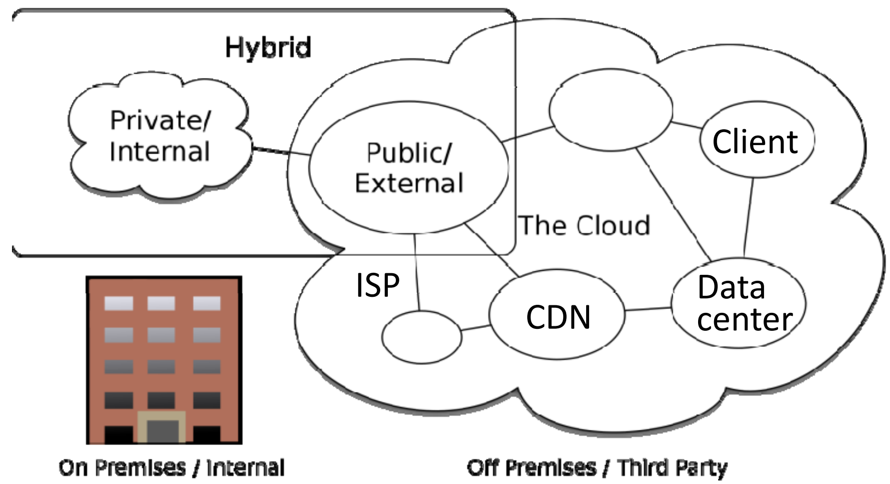

- Widespread acceptance and rapid growth of cloud computing are due to many factors including
  – Moore's Law in computing; storage becoming less expensive; increase in wide-area network bandwidth; web-based application delivery frameworks; and the high cost of IT insourcing to ensure security, scalability, failover, and software updates
- We expect an explosion in the number of cloud services in the coming decades

Cloud Infrastructure: Service Plane

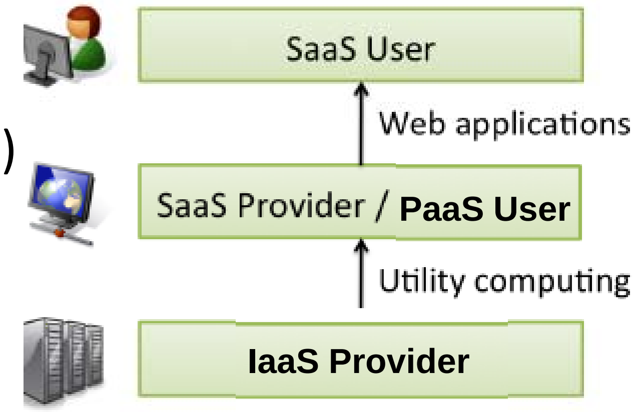

- A *service plane* comprised of the following:
  - Infrastructure (known as Infrastructure as a Service)
    - The provider only supplies the hardware infrastructure while the consumers provision processing, storage, networks, etc., and run arbitrary software and applications
  - Platform (known as Platform as a Service)
    - A set of software and product development tools hosted on the provider's hardware infrastructure. Developers create customized applications on this infrastructure over the Internet
    - Application development frameworks e.g., Google App Engine and Azure Services Platform
  - Software (known as Software as a Service)
    - The provider supplies the hardware infrastructure and the software product, and interacts with the user through a front-end portal
    - Basic services are provided e.g., Authentication, Billing & metering, Search, Web-based email, inventory control, database processing e.g., Amazon EC2, Microsoft Azure

Cloud Stack

- From a physical organizational view point the cloud can be treated as a stack containing five layers as depicted here
- The three types of service planes can be obtained by appropriate combination of these five layers
  - For example, an IaaS-type service plane mainly comprises the cyberphysical systems and computing/storage layers, whereas a SaaS-type service plane primarily consists of the applications/software and middleware/SOA layers
  - In both cases, the networks/distributed systems layer is involved since it provides load balancing and scheduling as well as data transfer either within or across datacenters

Examples of Key Challenges

- How does one guarantee a high degree of reliability and availability of critical Information and Communications Infrastructure (ICI) resources that relies on heterogeneity and replication at multiple levels, ranging from hardware to service providers?
- How does one equip the ICI to handle orders of magnitude more digital data and service requests securely and cost effectively?
- How should ICT components be architected so as to achieve true energy proportionality in the amount of work delivered?
- What are the economic and public policy implications of the dramatic changes in societal access to resources brought about by the ICI?

Internet Growth is Driving Data Center Usage

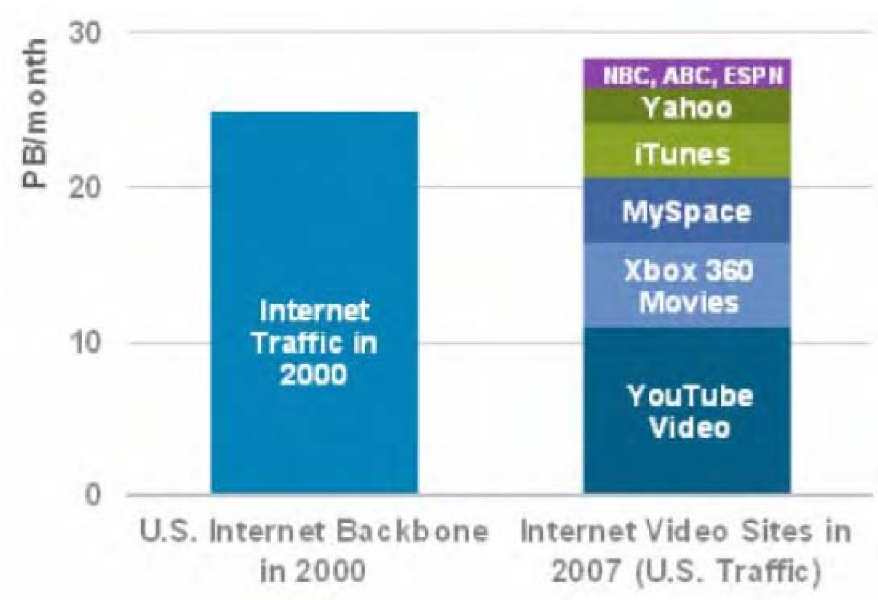

Source: ADC

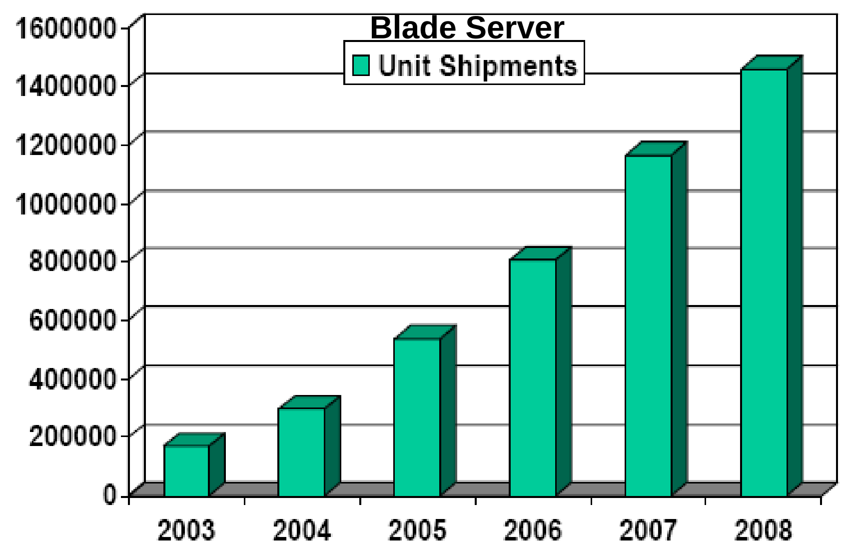

Increase from 500K to 1.45mn in 3 years

Source: Gartner

> Corporate data growing 50 fold in three years.
>
> —2007 Computerworld

Emergence of Mega Data Centers

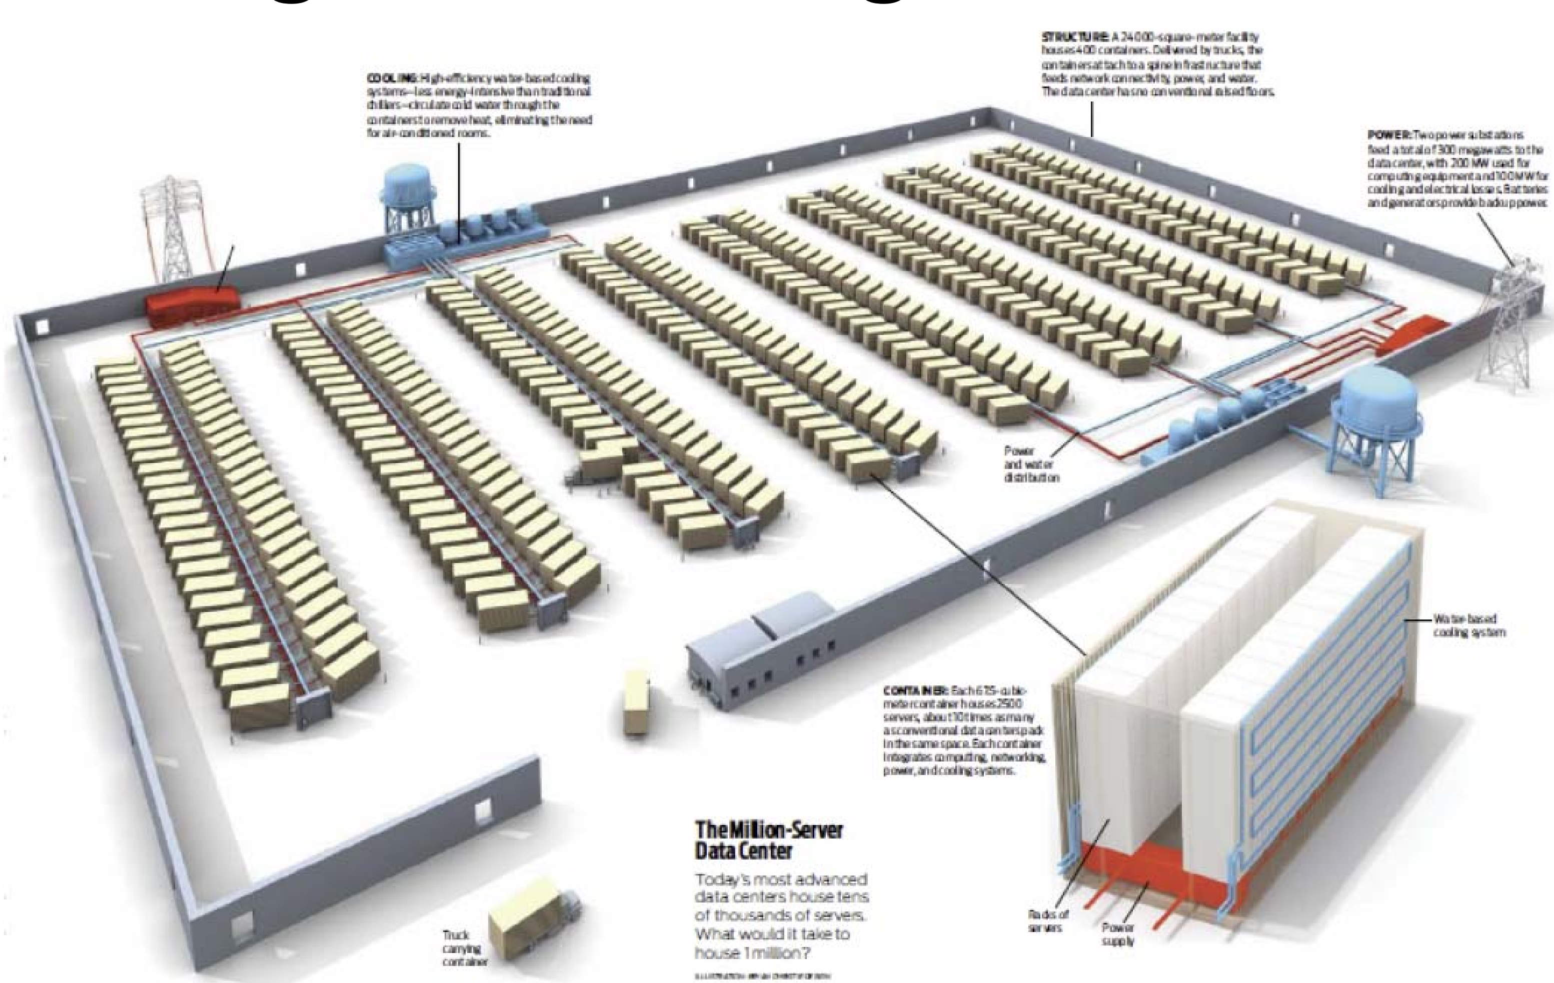

Microsoft's Chicago Datacenter: Entire first floor is full of containers; each container houses 1,000 to 2,000 systems; 150 - 220 containers on the first floor.

Green Data Centers: A Strategic Necessity

- Three drivers have led a "datacenter crisis"
    - Demand for digital services
    - Increase in energy dissipation of IT
    - Environmental concerns
- Datacenters consume ~ 2% of all US electricity
- *Datacenter annual growth* of 15% is unsustainable
- **Datacenter power projected to be > 8% of US power by 2020**
    - Datacenter carbon emissions are projected to exceed those of the airlines by 2020
- Need a <u>paradigm shift</u> in data center computing to put us on a more sustainable and scalable ICT energy efficiency curve

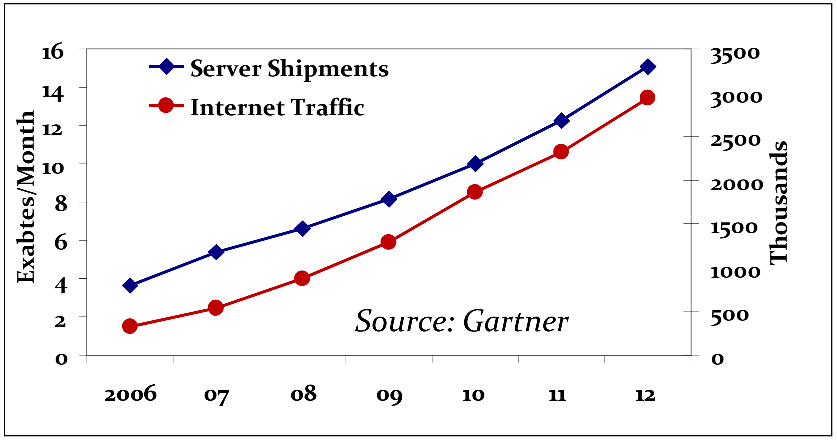

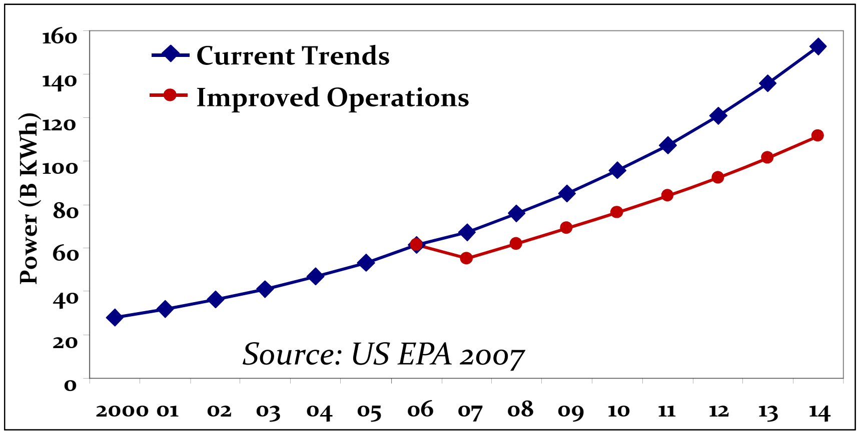

Peak Power Demands of Data Centers

- Apart from the total energy consumption, another critical component is the peak power; the peak load on the power grid from data centers is currently estimated to be approximately 7 gigawatts (GW), equivalent to the output of about 15 baseload power plants
  - This load is increasing as shipments of high-end servers used in data centers (e.g., blade servers) are increasing at a 20-30 percent CAGR.
  - If current trends continue, power demands are expected to rise to 12 GW by 2011
  - Indeed, according to a 2008 Gartner report, 50 percent of data centers will soon have insufficient power and cooling capacity to meet the demands of high-density equipment

Ecological Footprint of Data Centers

- Although carbon emission due to every kWh of electrical energy consumed in the U.S. varies depending on the type of power generator used to supply power into the grid, an **average conversion rate of 0.433 kg CO2 emission per kWh of electrical energy** can be assumed
- Data centers currently generate 23% of all emissions produced by the ICT industry, a figure that continues to trend upward (> 79 MMT of CO2 in 2011)
  - The Intergovernmental Panel on Climate Change (IPCC) has called for an overall greenhouse gas emissions reduction of 60-80% below 2000 levels by 2050

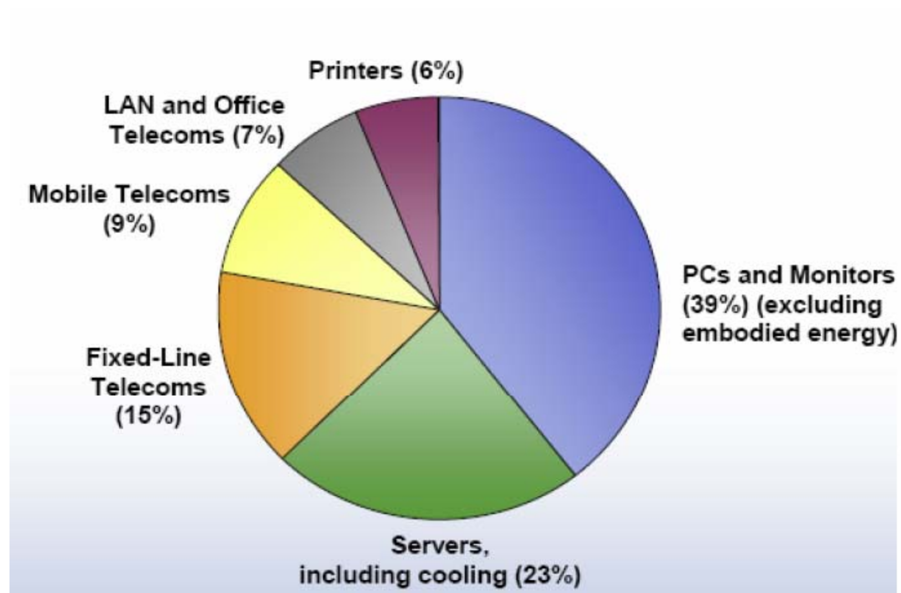

Gartner

Power Profile

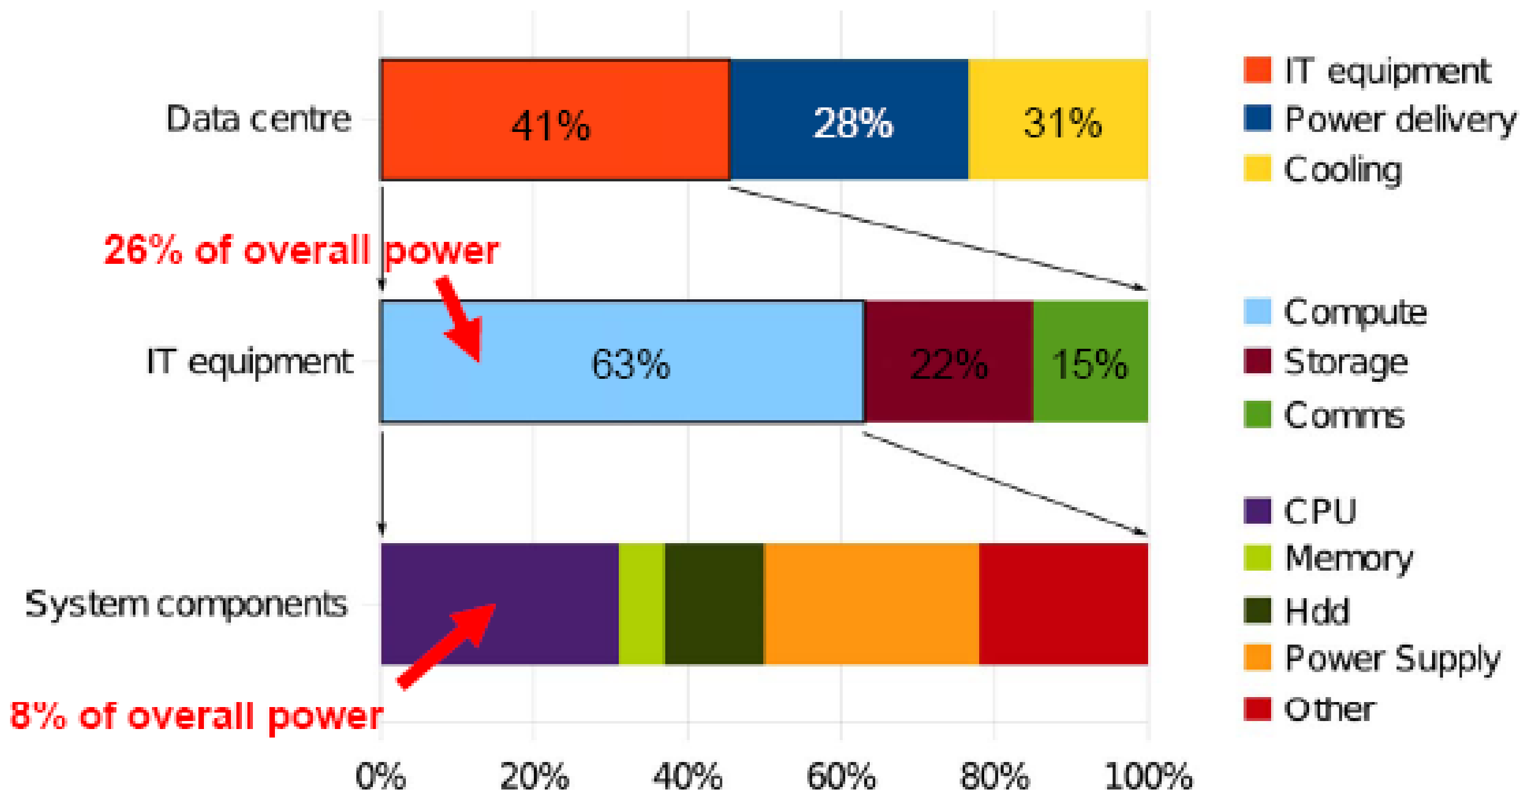

Source: Data Center Efficiency in the Scalable Enterprise,
Dell Power Solutions, Feb 2007

Idleness, State Transitions, and Power Savings

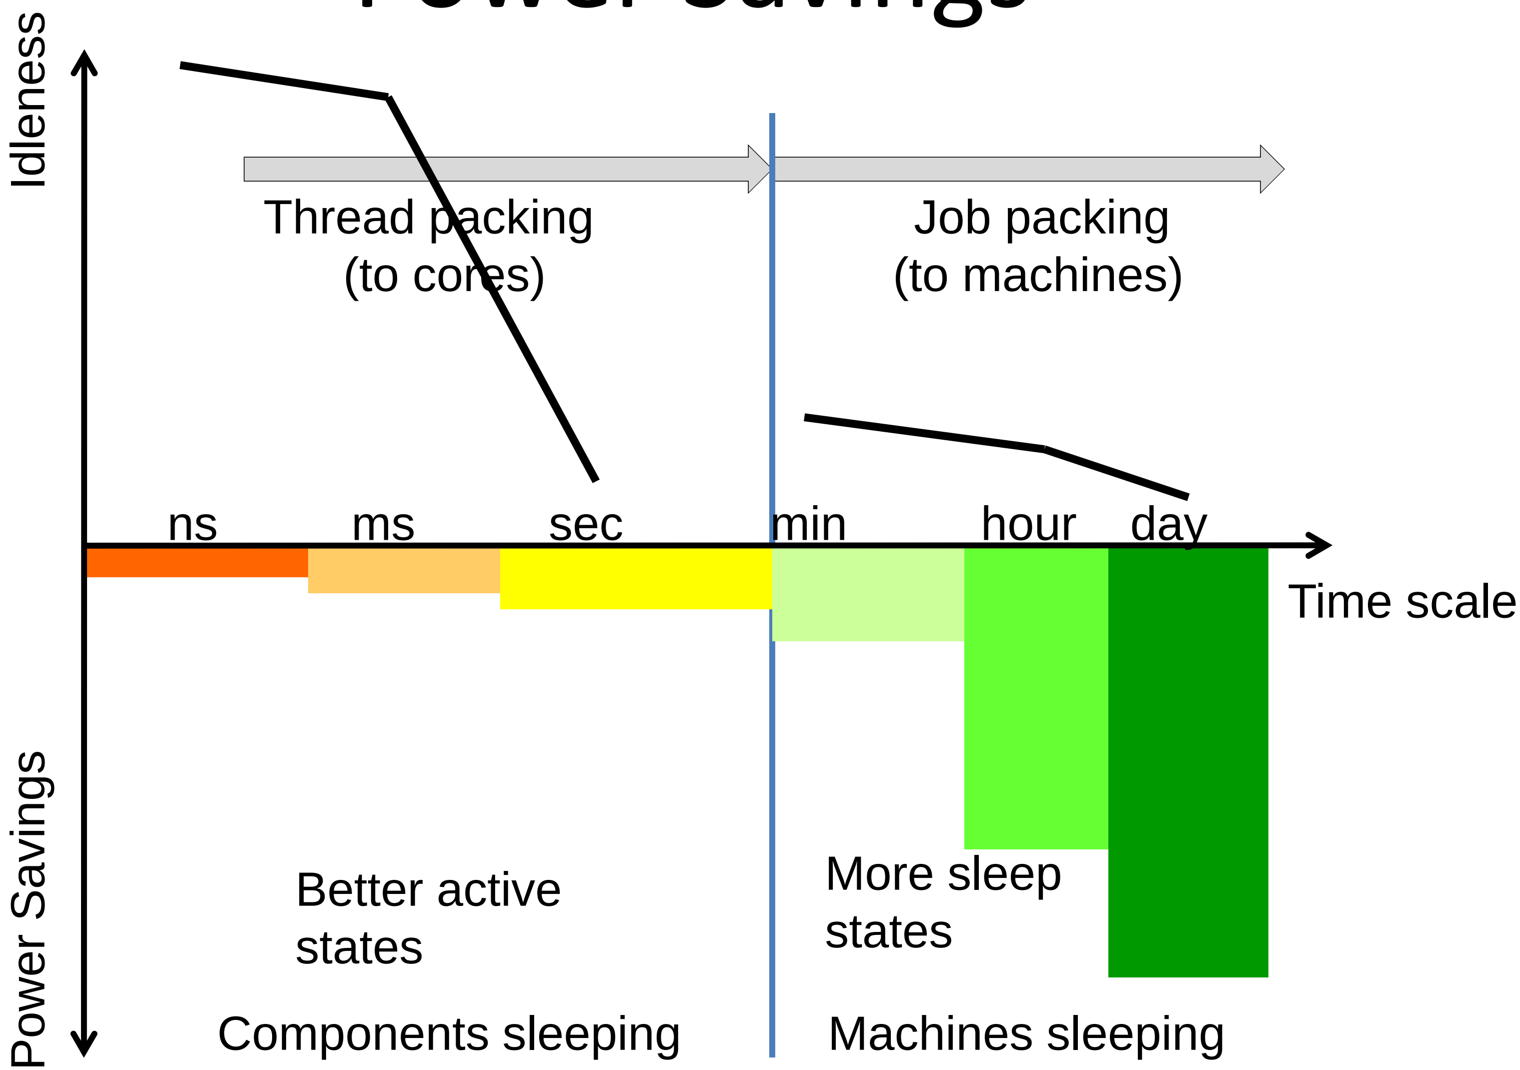

Source: W-D. Weber, Google

Minimum Energy CMP Design Through Consolidation and DVFS

- Chip Multiprocessor system with *M* identical cores
  - Per-core DVFS with *N* (*voltage-frequency*) configurations
  - Local Queue
- Power Management Unit
- Global Queue
- Task Dispatcher
- Tasks
  - Expected *execution time*, τ
  - Expected *instructions per cycle*

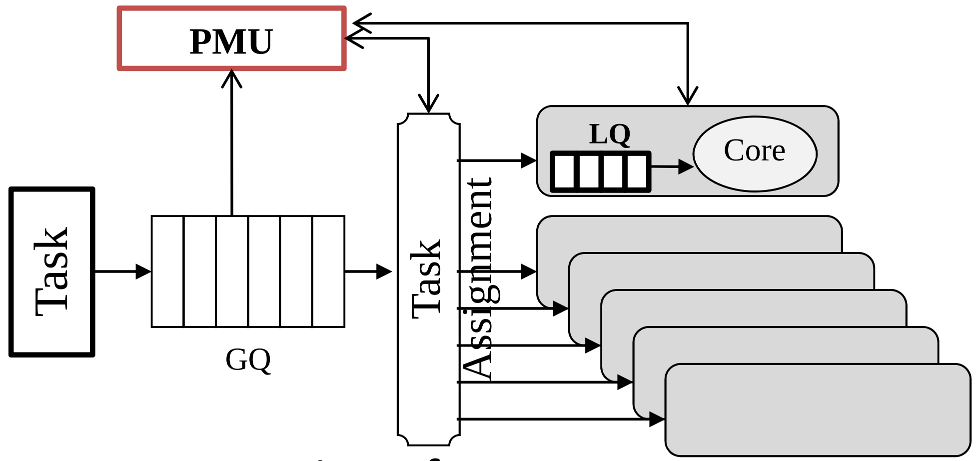

- Objective: Minimize the total Energy consumption of cores in a multicore system
- Constraints: Minimum required throughput, IPSreq
- Variables:
- Number of active cores –total number of cores: M
- v-f settings of cores –total settings: J
- Task distribution– total number of tasks: K

A Hierarchical Solution

- Determine the number of ON cores
  - ON cores are in C0 when active or C2 (halt or sleep) when idle

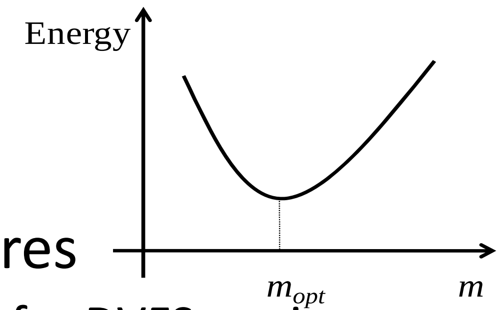

- Determine operating frequency of ON cores
  - A feedback-based control method is adopted for DVFS setting
    - This is needed due to inherent uncertainty and variability of task characteristics
    - PI Controller: controller adjusts the v-f setting to match the required throughput based on the observed error
    - Feedback control loop determines a single optimum frequency for all cores and then the Quantizer would translate f_{opt} to available DVFS steps
- Find a feasible assignment of tasks to the ON cores

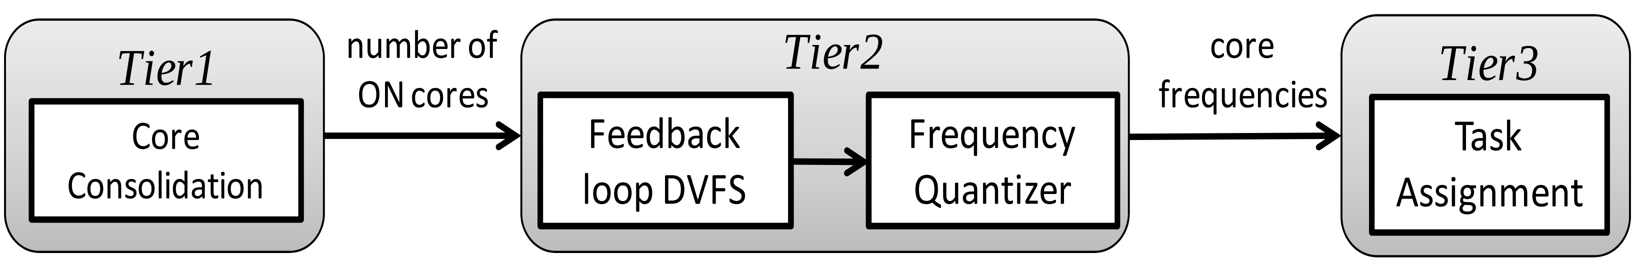

Evaluation Results

- Comparison to a relatively energy-efficient baseline PM
  - It implements the same power reduction techniques as our method
  - It utilizes open loop DVFS, and does not support core consolidation
- The figure compares the frequency setting and throughput
  - The baseline PM always runs with all cores ON
  - The closed loop frequency is always lower for same throughput constraint
  - In the second 100ms, the baseline chooses lower frequency with m=4 cores running while our proposed method uses m=3 cores
- The proposed PM gains an overall energy consumption of **17%**

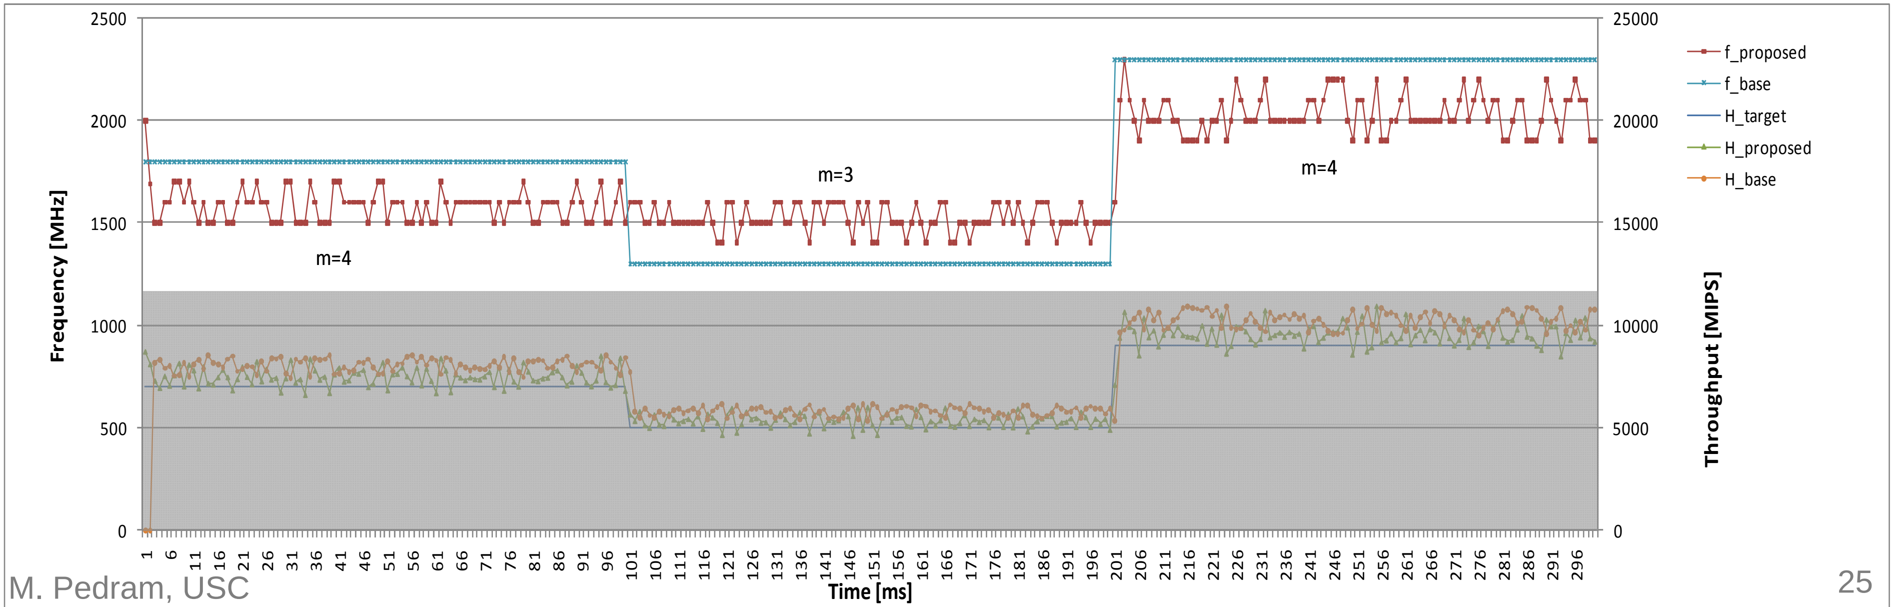

Server Utilization

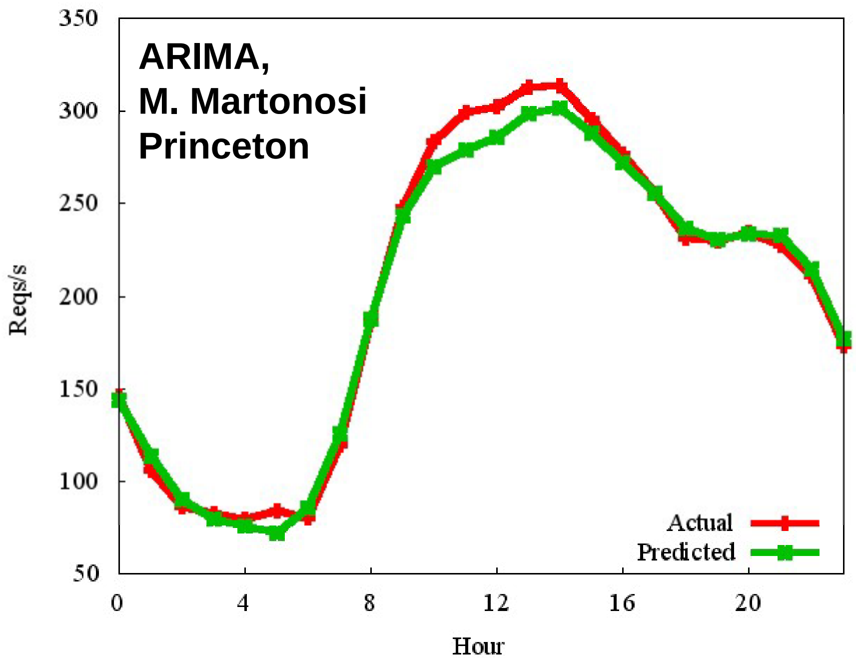

Workload intensity profile for a typical week day

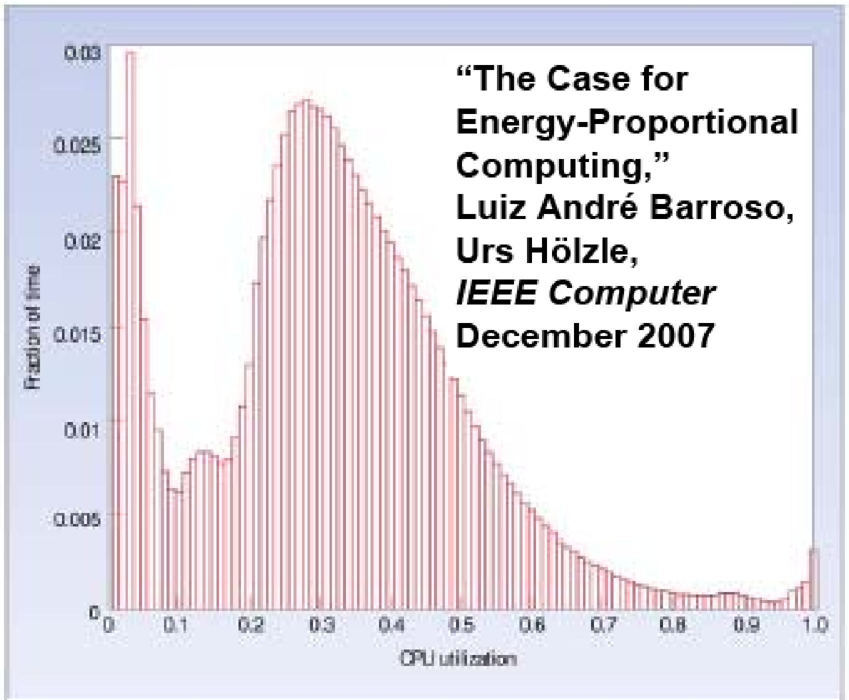

PDF of CPU utilization measured over a 6-month period

Servers are rarely completely idle and seldom operate near their maximum utilization, instead operating most of the time at between 10 and 50 percent of their maximum

Reality: Non-Energy-Proportional Servers

Power dissipation of Intel Xeon 5400 series server at high (2.3GHz) and low (2GHz) settings

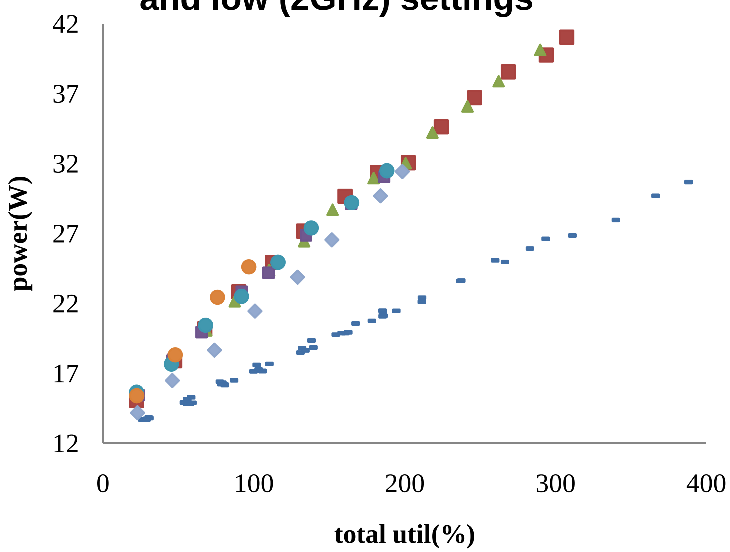

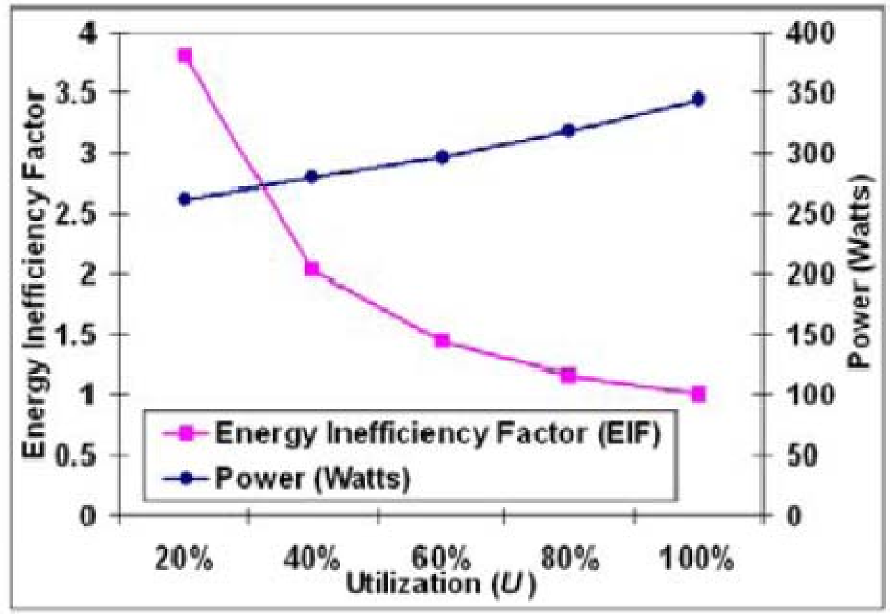

$$EIF_U = \frac{P_U}{P_1 U}$$

Here P_1 denotes server power dissipation at 100% utilization, whereas P_U is power at utilization level of U

An energy proportional system will have an Energy Inefficiency Factor (EIF) of one at all utilization levels

Virtualization

- Virtualization is disassociating the tight bond between software and hardware by introducing a hypervisor between the OS and hardware
  - One can then use the same hardware to serve up the needs of the different software servers: Oracle, MS SQL Server, Exchange, Dynamics CRM, etc.
  - It is also possible to run different operating systems so one could run MS SQL Server 2008 on Windows 2008 Server and run Oracle on Linux all running on the same hardware
- By doing this, the resources will be better utilized since we can easily add/migrate virtual machines
  - Examples include Microsoft Hyper-V, VMware ESX Server 3.5, Linux Xen hypervisor

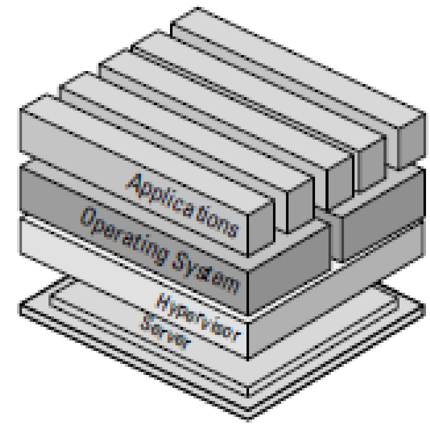

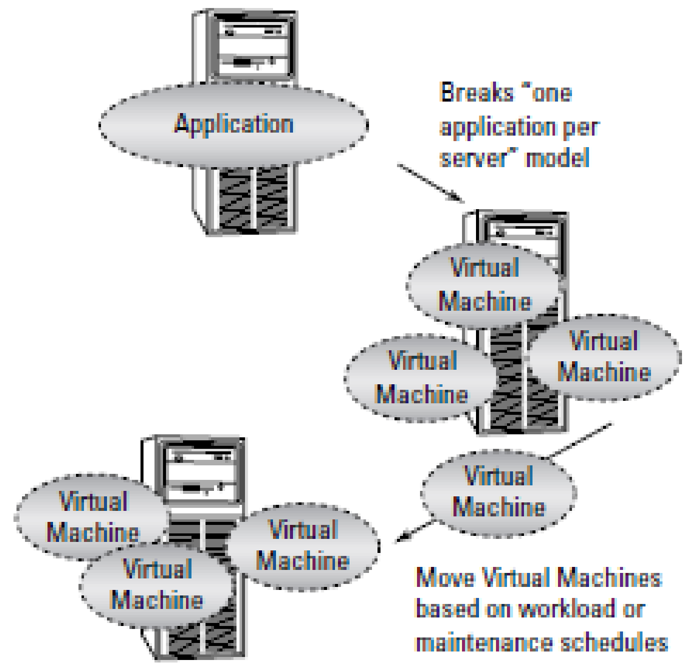

Consolidation: Power Analysis

- Core consolidation is ineffective from a power saving perspective
  - Processor power dissipation is nearly independent of which subset of CPUs is used by the running domain
- Processor consolidation is quite helpful in reducing the total power dissipation
  - Reported as a function of total utilization of all CPUs in the system

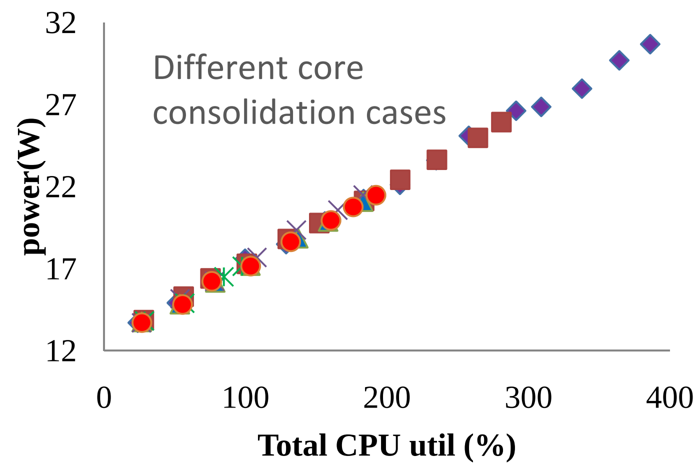

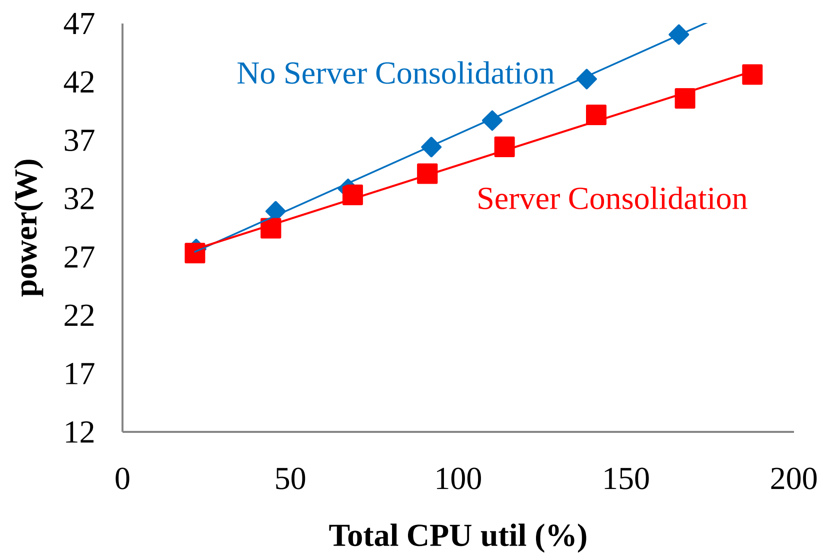

Consolidation: Performance Analysis

- For a given number of virtual CPUs, the response time decreases as the number of active CPUs increases
  – Normalized CPU util = Total CPU util / (# of active CPUs)
  – True when Normalized CPU utilization is unchanged
- For a given number of physical CPUs, the response time increases as the number of virtual CPU’s increases
  – True even when the total CPU utilization is unchanged

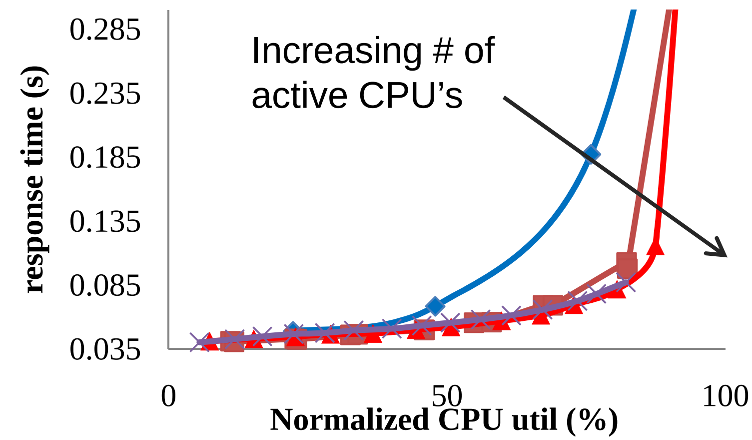

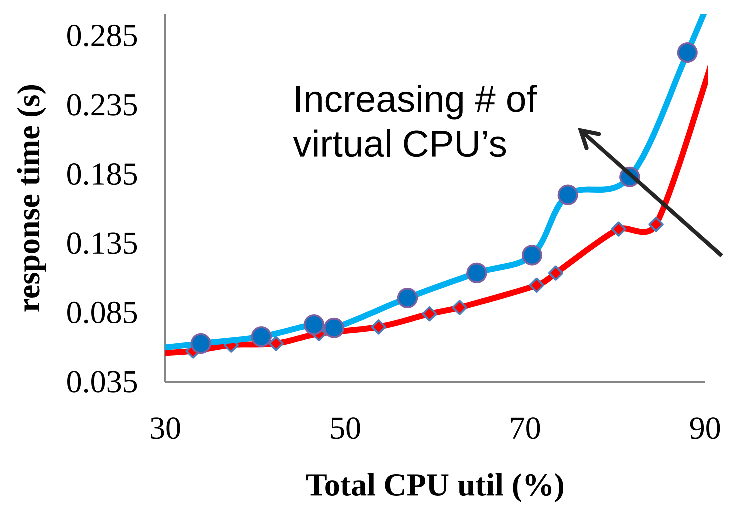

Power-Performance Tradeoff

- Pareto-optimal surface of Energy-Delay Product vs. the request arrival rate
  - Proper selection of the virtualization configuration (# of virtual CPUs vs. physical CPUs and their mappings) of an enterprise computing system can improve energy efficiency subject to given SLAs

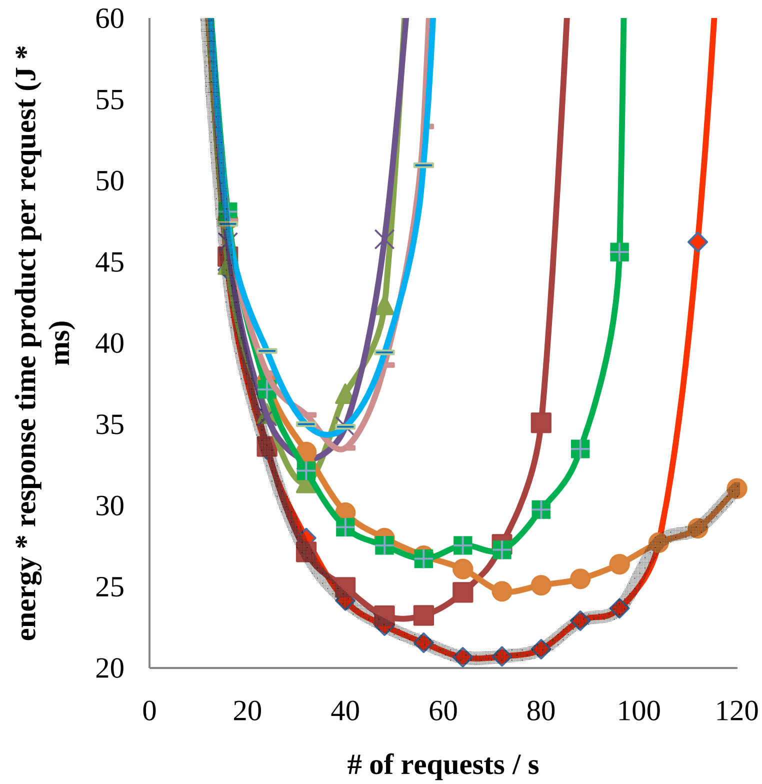

Energy * response time product per request as a function of virtualization configuration

Scheduling and Resource Assignment

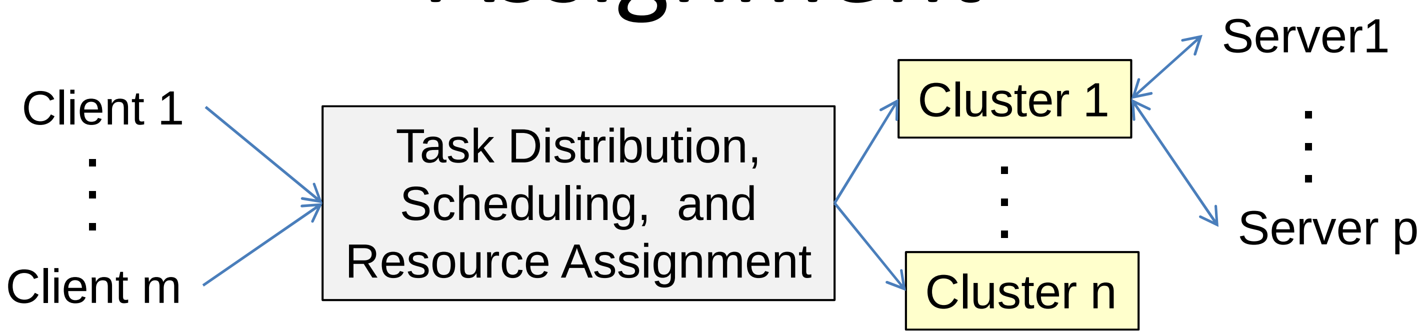

- A cloud computing system is an important example of a distributed computing system
  - Clients typically impose some Service Level Agreements (SLAs) while the cloud owner is trying to maximize its total profit
  - Server clusters are usually heterogeneous in the number and type of resources they control
  - Assume a step wise utility function for the clients based on their response time requirement
  - Must use a three dimensional model of the resources in the clusters, i.e., computational, storage and networking capabilities, and finally

A Distributed Algorithm

- Generate an initial solution for assignment of clients to clusters and allocation of processing and data storage resources to them
  - For each client, the best possible cluster to execute the application is found by finding the highest approximated profit for the client on each cluster with respect to the current state of the clusters
  - This process continues until all clients are assigned to the clusters
- Allocate communication resources to assigned clients
  - An instance of Multi Choice Knapsack Problem, solved by a pseudo polynomial dynamic programming (DP) method
- Maximize the total utility by improving the resource allocation within a cluster without changing the assigned set of clients to it
  - An instance of Multi dimensional Knapsack problem , solved by DP
- Iteratively Improve client assignment to clusters
  - In each step a client is picked randomly and the best cluster is chosen with respect to the current state of each cluster and the characteristics of the client

Some Results

- The number of clusters, the number of different server classes, and the number of different utility classes are set to 5, 10 and 5, respectively
- For each utility class, means of the service rates for the clients are generated with (uniformly distributed) random variables between 0.4 and 1
- The value of request arrival rate for each client is set with a random variable between 0.5 and 4.5
- The utility class of each client is set with a random assignment form the existing utility classes

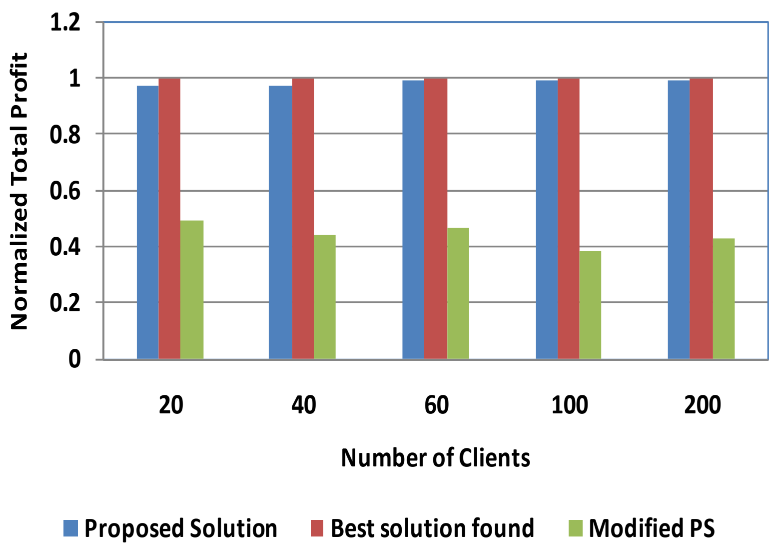

PS: Proportional Share Scheduling

Agent-Based Resource Management

Employ a hierarchical control architecture so as to enable tradeoffs at all tiers

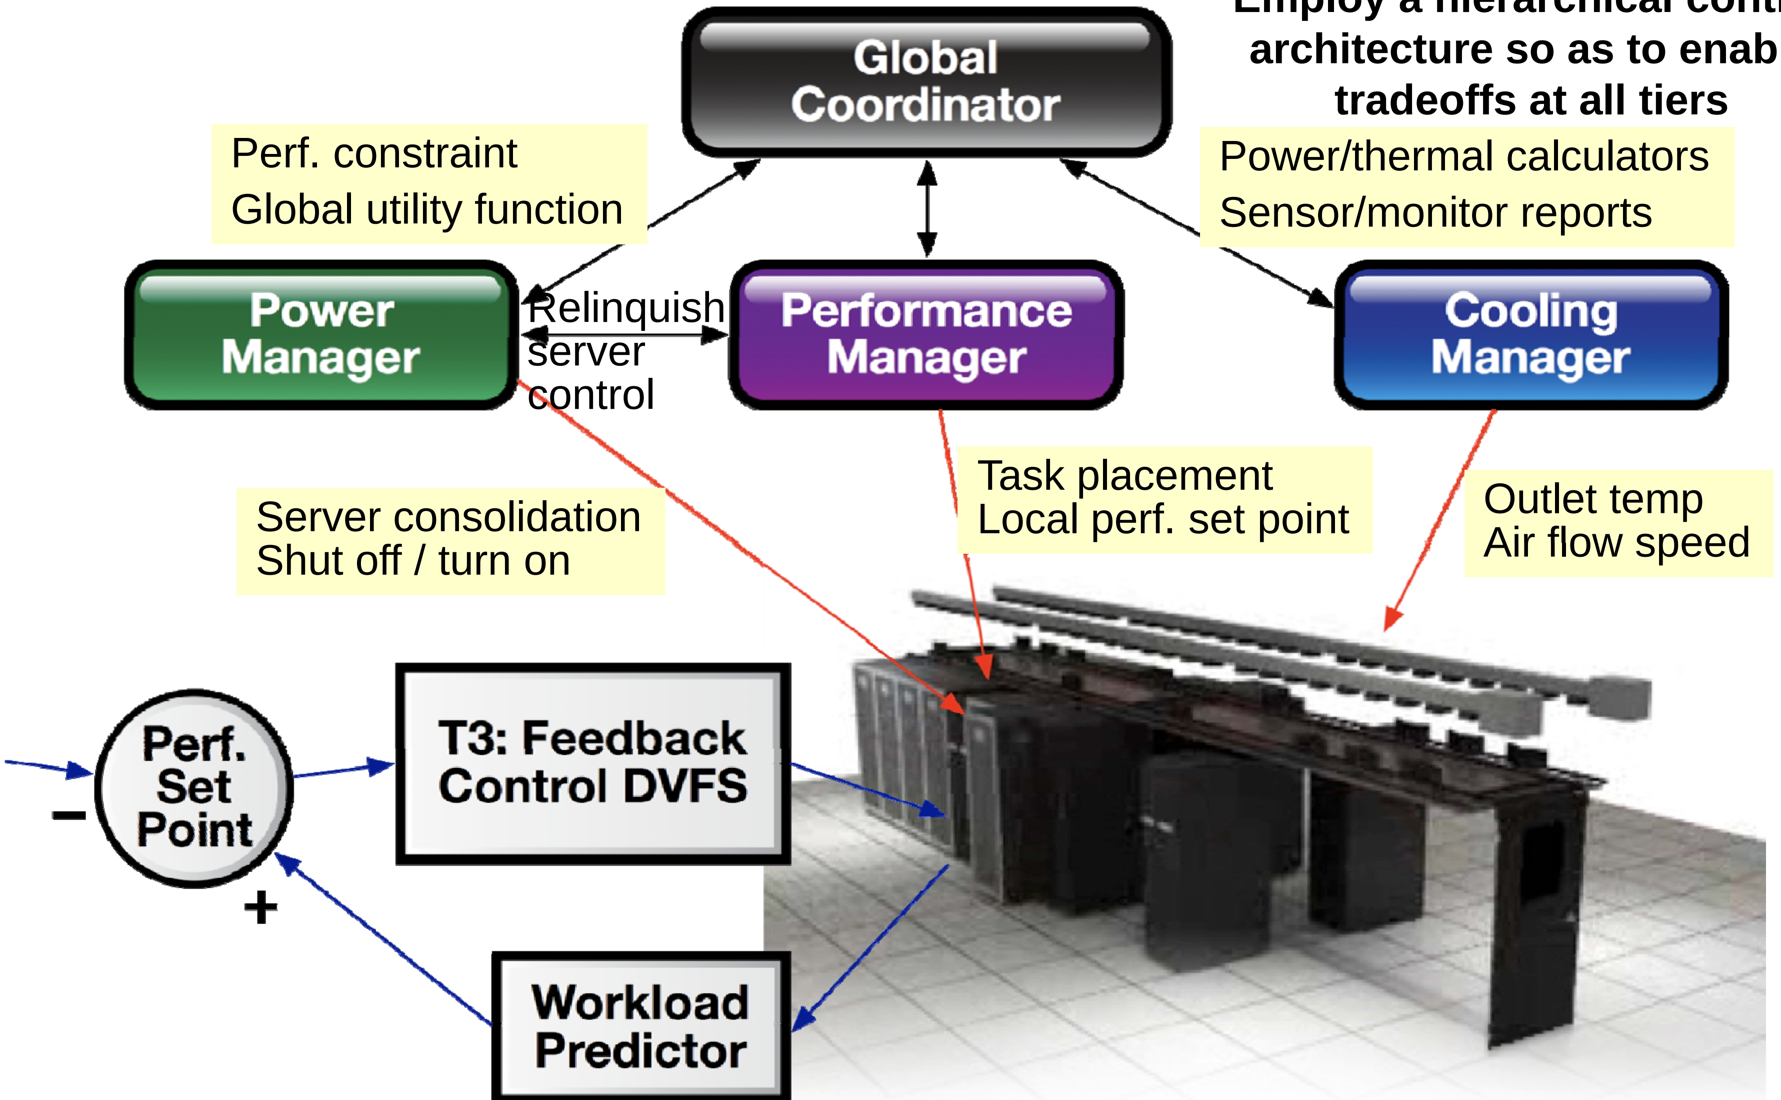

Dynamic Workload Placement based on Cooling Efficiency

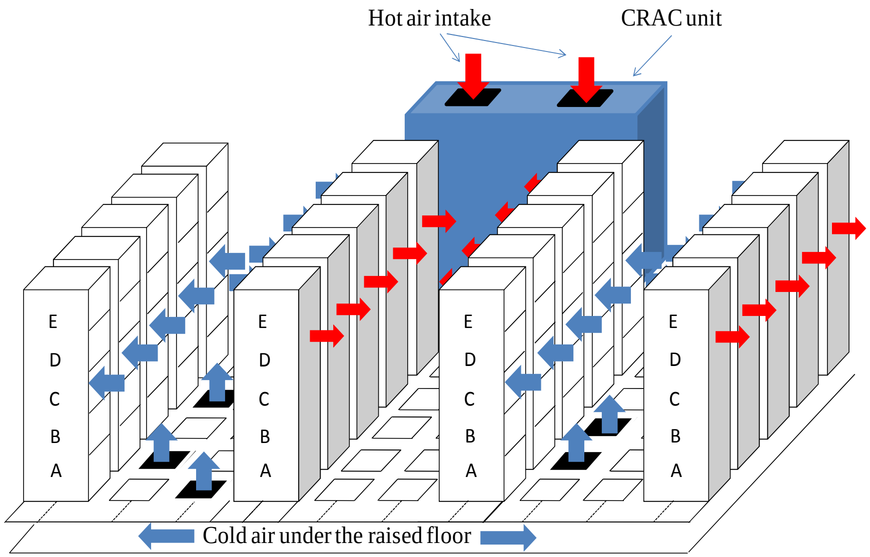

Premise:

- Hotspots exist that impact HVAC efficiency and energy consumption

Approach:

- Place workload in more power-efficient locations
- Perform server consolidation and DVFS setting

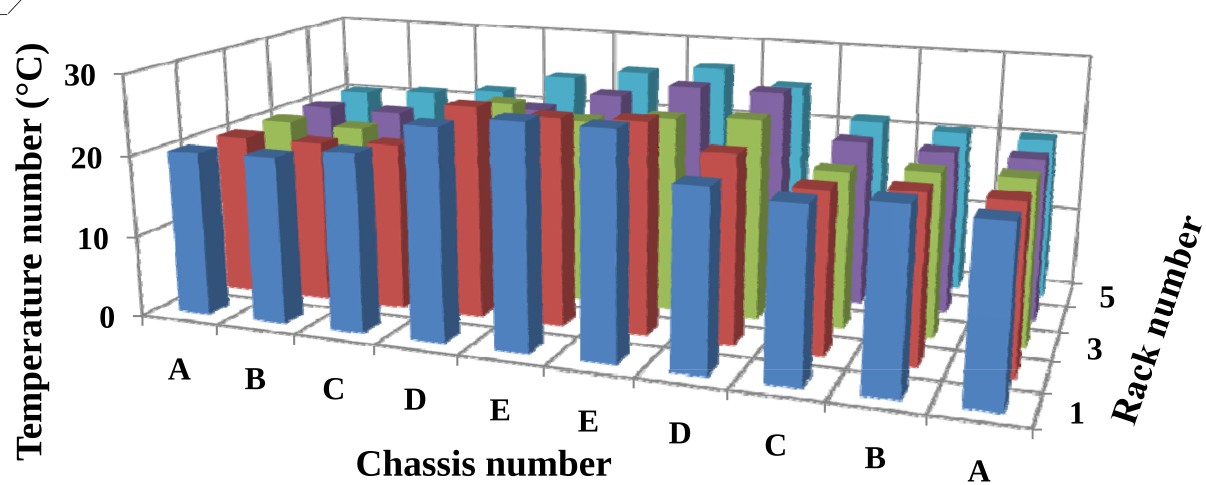

Smart Grid and Dynamic Energy Pricing

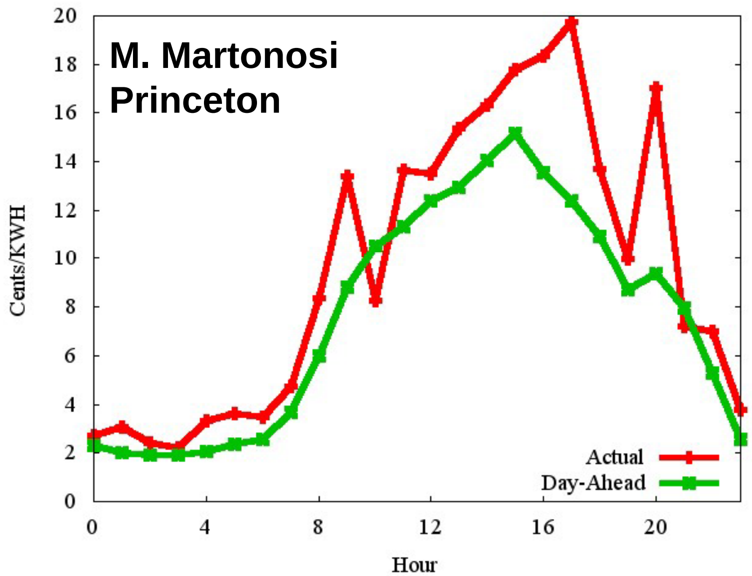

- Day-ahead price prediction is usefully accurate, although some deviations occur mid-day
- Information can be gathered on both dynamic electricity usage and prices with the aim of reducing energy consumption of both individual households and businesses

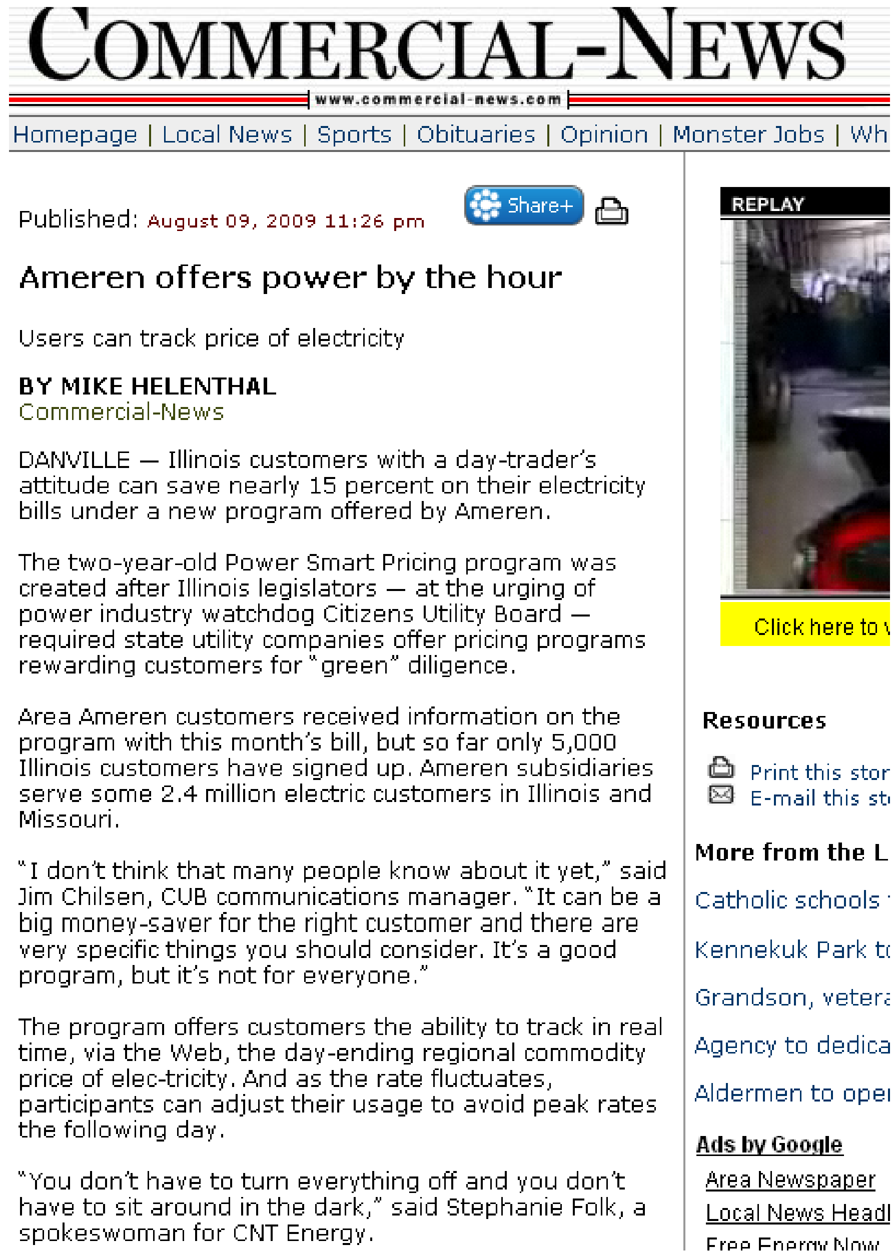

COMMERCIAL-NEWS

www.commercial-news.com

Homepage | Local News | Sports | Obituaries | Opinion | Monster Jobs | Wh

Published: August 09, 2009 11:26 pm

Share+

Ameren offers power by the hour

Users can track price of electricity

BY MIKE HELENTHAL
Commercial-News

DANVILLE — Illinois customers with a day-trader's attitude can save nearly 15 percent on their electricity bills under a new program offered by Ameren.

The two-year-old Power Smart Pricing program was created after Illinois legislators — at the urging of power industry watchdog Citizens Utility Board — required state utility companies offer pricing programs rewarding customers for "green" diligence.

Area Ameren customers received information on the program with this month's bill, but so far only 5,000 Illinois customers have signed up. Ameren subsidiaries serve some 2.4 million electric customers in Illinois and Missouri.

"I don't think that many people know about it yet," said Jim Chilsen, CUB communications manager. "It can be a big money-saver for the right customer and there are very specific things you should consider. It's a good program, but it's not for everyone."

The program offers customers the ability to track in real time, via the Web, the day-ending regional commodity price of elec-tricity. And as the rate fluctuates, participants can adjust their usage to avoid peak rates the following day.

"You don't have to turn everything off and you don't have to sit around in the dark," said Stephanie Folk, a spokeswoman for CNT Energy.

REPLAY

Click here to v

Resources

Print this stor
E-mail this st

More from the L

Catholic schools
Kennekuk Park t
Grandson, veter
Agency to dedica
Aldermen to ope

Ads by Google

Area Newspaper
Local News Headl
Free Energy Now

Summary

- These are exciting times with many new opportunities and challenges due to planned upgrades to the Power Grid, introduction of renewable sources of energy, smart metering and dynamic energy pricing, people's awareness of environmental issues, advent and popularity of social networking
- Must be able to characterize the life cycle environmental impacts of the ICT on carbon emission, water footprint, and air pollution
- A holistic , cross-layer approach to ICT energy efficiency and greenness is needed, which spans
  - Application efficiency and energy management, micro-architecture and system design, storage and networks, resource management and scheduling
- Must contribute to understanding the basic issues and options available to policy makers and guide them on how to make crucial decisions about proposed changes of the ICI and the "inter-cloud" by incentivizing commercial and individual users, through market design, or by regulation