Paying to Save: Reducing Cost of Colocation Data Center via Rewards

Mohammad A. Islam, Hasan Mahmud, **Shaolei Ren,** Xiaorui Wang*

Florida International University

*The Ohio State University

Supported in part by NSF CNS-1423137 and CNS-1143607

Data center

Google's data center in Mayes County, Oklahoma

Data centers are power-hungry

U.S. data centers in 2013

91 billion kWh; 50% increase by 2020

Power entire Washington; $10+ billion

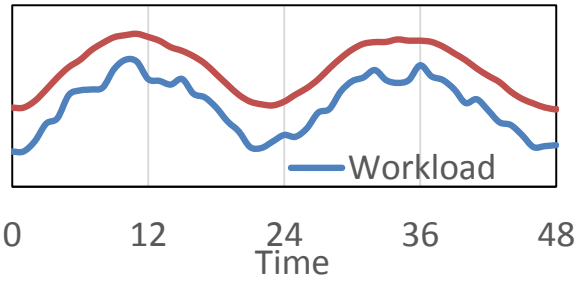

Many energy-saving techniques

Source: Natural Resources Defense Council

Bad news…

Most of the existing approaches ***cannot*** be applied in many large data centers…

What are data centers?

Google

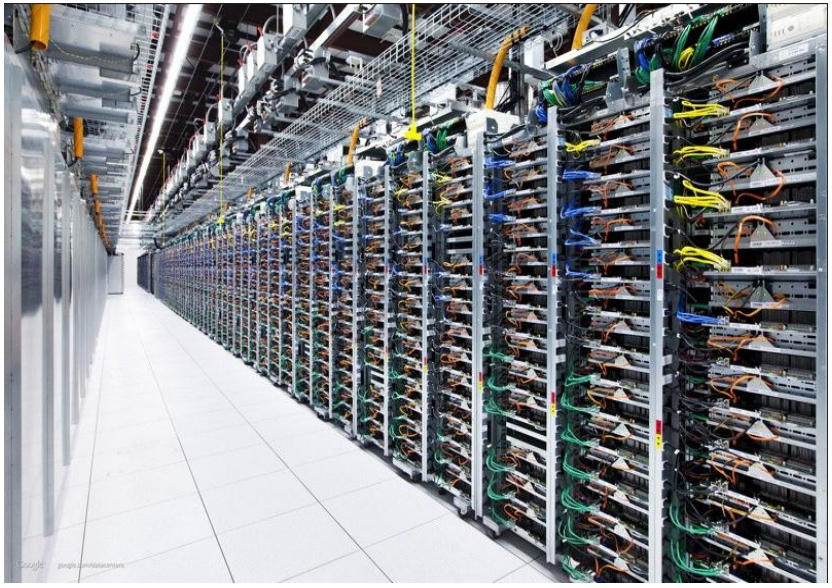

What are data centers?

Data centers

Google

facebook

Microsoft

EQUINIX

verizon terremark

TelecityGroup

Multi-tenant colocation data center

- Multiple tenants house their own servers in one ***shared*** space and manage their equipment independently
- Data center operator is mainly responsible for facility management (e.g., power distribution, cooling)

CoreSite's "One Wilshire" (Photo: CoreSite)

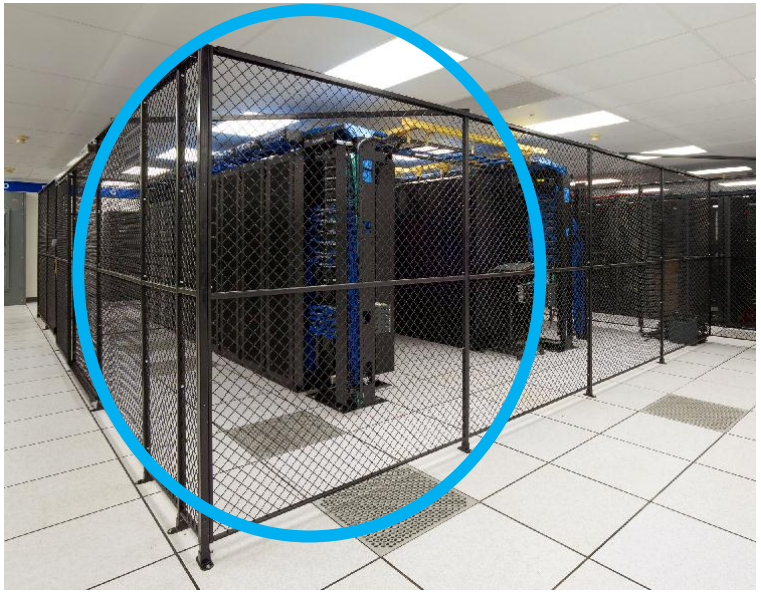

Who are using colocations?

- Almost all industry sectors
  - Including top-brand websites, e.g., Wikipedia, Twitter

- Many clouds

- Our Internet
  - According to Cisco, **55% Internet traffic** will be handled by CDN providers, e.g., Akamai, by 2018 (up from 33% in 2013)

37.3%

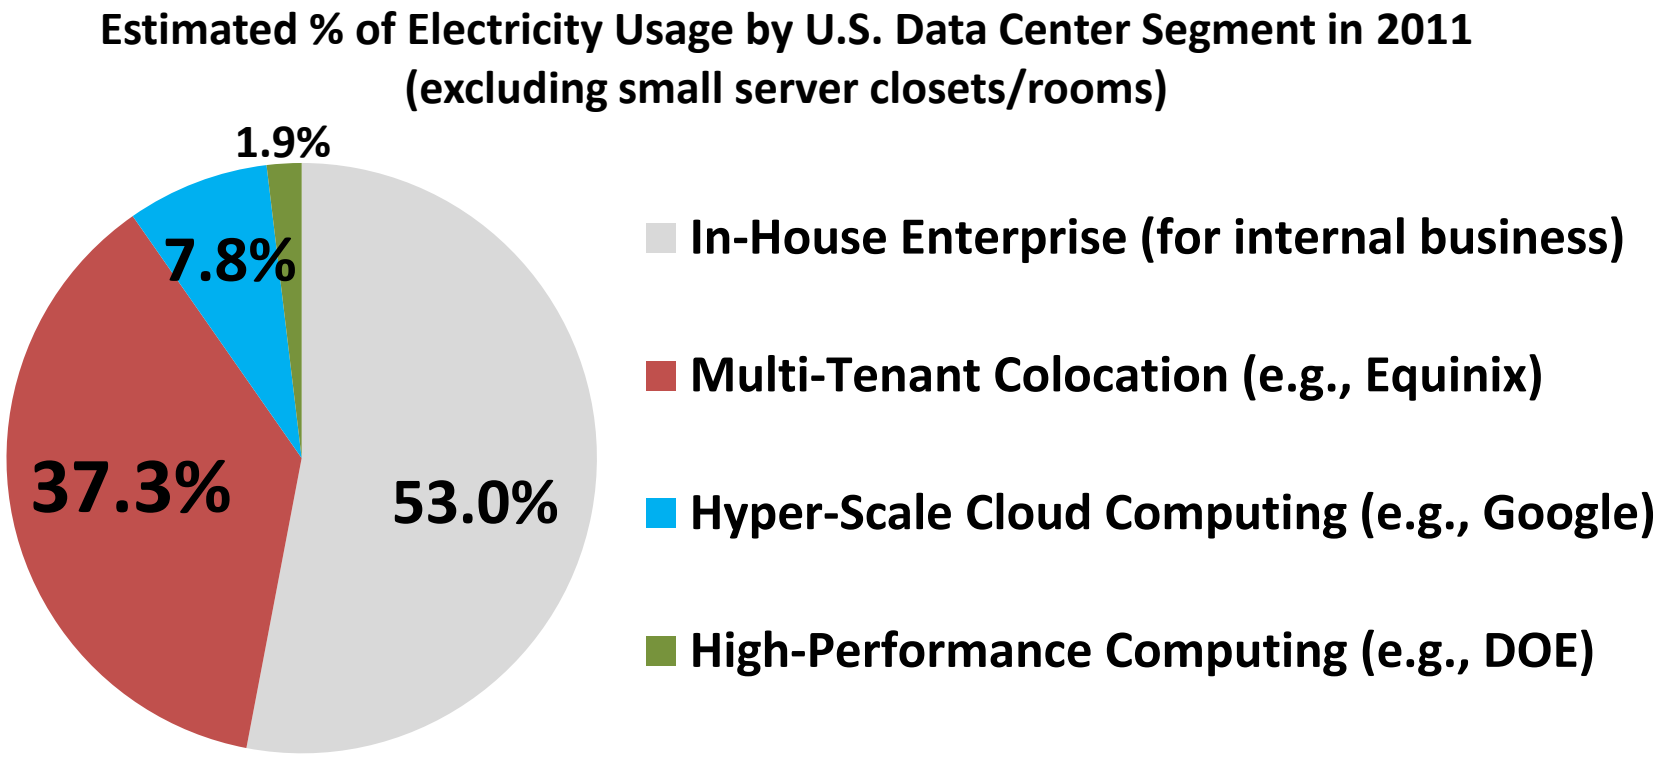

There're over **1,200** colocation data centers in the U.S., and according to IDC'14, many enterprise data centers are migrating to colocations!

Source: Natural Resources Defense Council, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," Issue Paper, Aug. 2014.

What's the problem with colocation data centers?

Cost, cost, and cost!

- Electricity cost:

Amortized Cost	Component	Sub-Components
~45%	Servers	CPU, memory, storage systems
~25%	Infrastructure	Power distribution and cooling
~15%	Power draw	Electrical utility costs
~15%	Network	Links, transit, equipment

Table: Guide to where costs go in an owner-operated data center.

Source: A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. 2008. The cost of a cloud: research problems in data center networks. SIGCOMM Comput. Commun. Rev.

Cost, cost, and cost!

- Electricity cost: ≈**40%** of TCO

Amortized Cost	Component	Sub-Components
~25%	Infrastructure	Power distribution and cooling
~15%	Power draw	Electrical utility costs

Table: Guide to where costs go in an owner-operated data center.

Let's reduce the electricity cost (**OpEx**)!

Source: A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. 2008. The cost of a cloud: research problems in data center networks. SIGCOMM Comput. Commun. Rev.

A major **barrier** for reducing OpEx

- A commonly used pricing model
  - **Power Subscriptions**
- **"Split incentive"**
  - Colocation operator desires energy saving
  - Tenants manage servers but have **no incentive** for energy saving

CIRCUITS	KW	RATE
20/120A	2.4	$ 350
30/120A	3.6	$ 525
20/208A	4.2	$ 610
30/208A	6.2	$ 910
50/208A	10.4	$ 1,520

Promotional pricing in Verizon Terremark's Miami data center

Pricing tenants' energy usage is **not** good (enough)

- **Uncoordinated** power management
  - High peak power demand charge
    - up to **40%** of total electricity bill

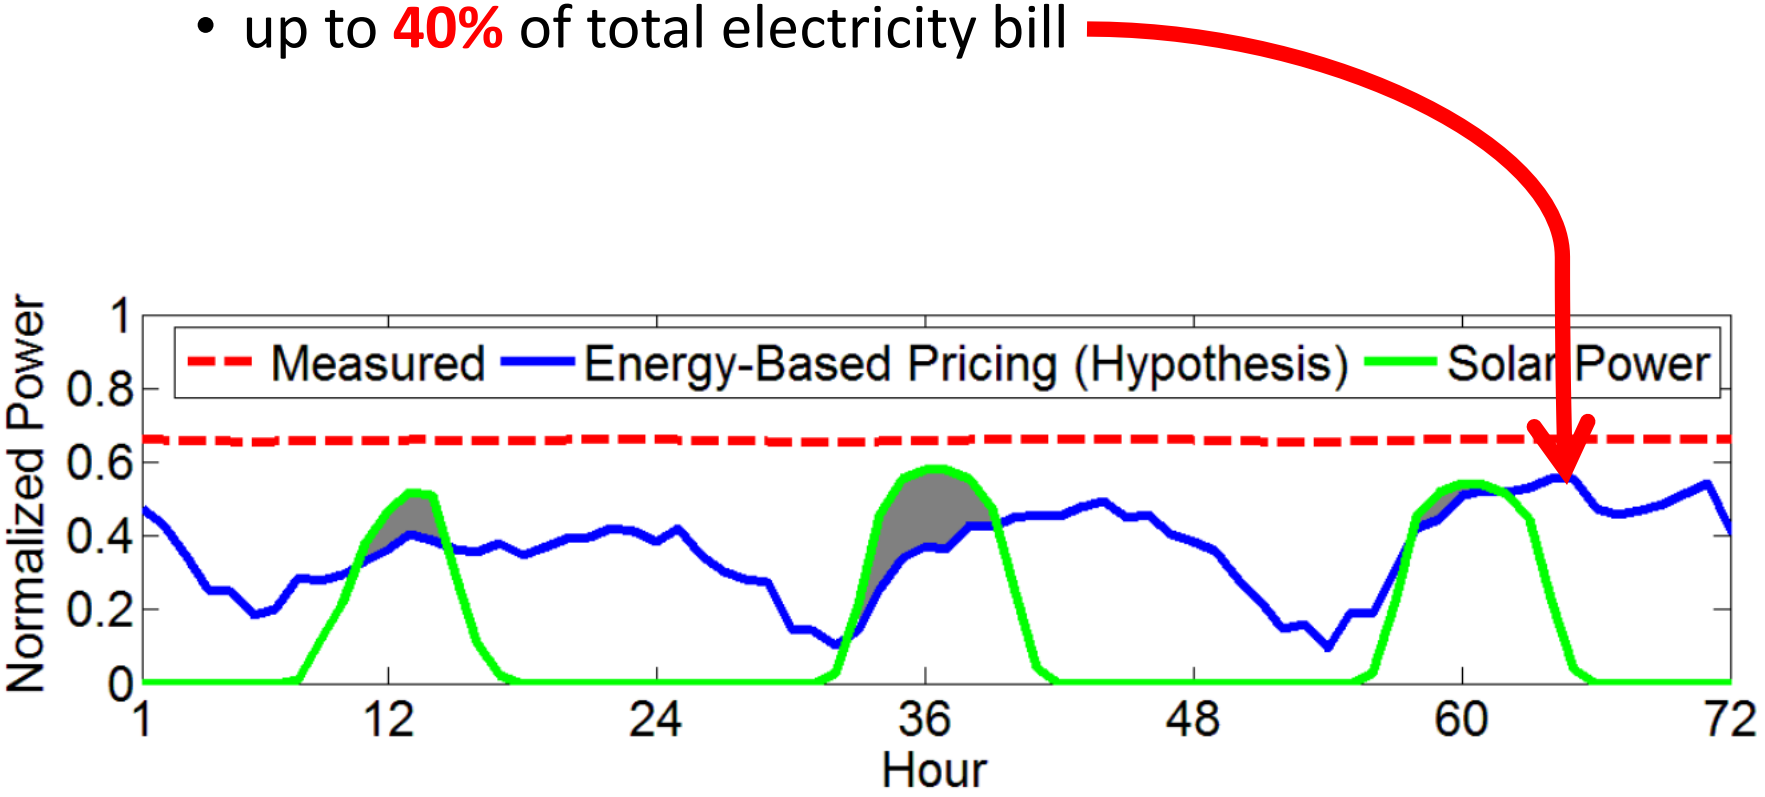

Pricing tenants' energy usage is **not** good (enough)

- **Uncoordinated** power management
  - High peak power demand charge
    - up to **40%** of total electricity bill
  - Fail to "follow the renewables"
    - Saving energy at an **appropriate** time!

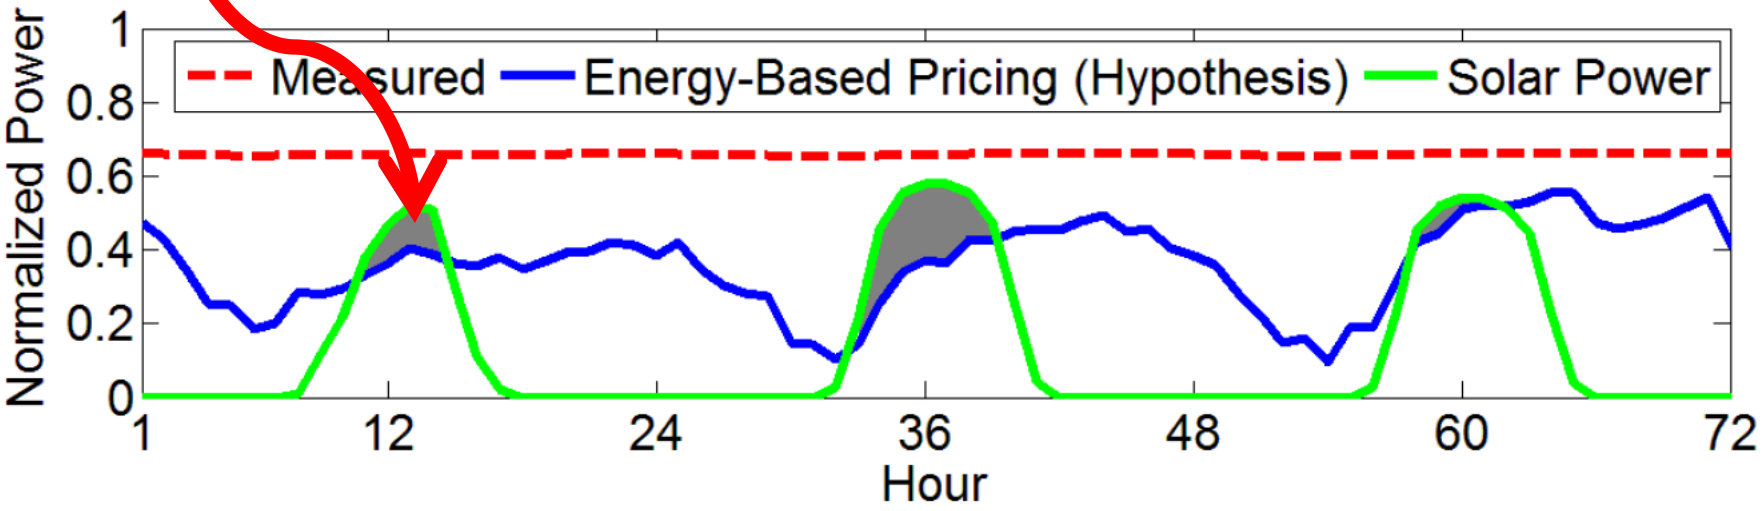

RECO

Pay tenants for **not** using energy at appropriate times

Challenges

- We need to dynamically set rewards ***online***
  - Time-varying cooling efficiency
  - Tenants' unknown responses

Prediction

  - **Peak demand charge**, often **$10+ per kW**, is typically determined as the maximum power (e.g., over 15-min interval) over a billing cycle
    - Set higher rewards during peak demand periods, but **when is peak demand period?**

Online feedback

RECO

- Track the “peak” power demand online
  - Set a higher reward only when the expected power usage exceeds the tracked peak

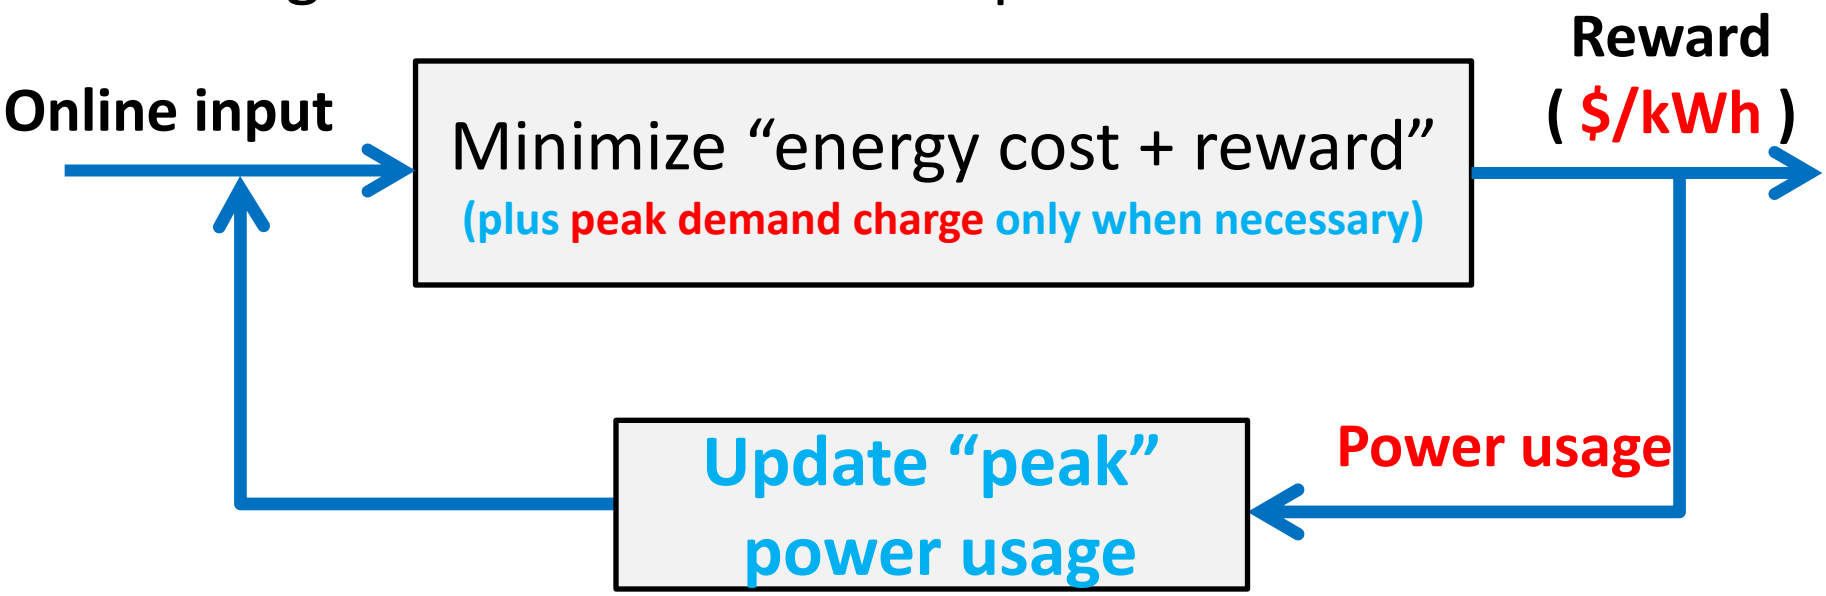

Case study

How to evaluate RECO?

- Scaled-down prototype system
  - 5 Dell PowerEdge servers, each having 6 VMs
    - **Tenant #1**: Web workloads based on key-value stores (e.g., Facebook), with a SLA of 95% delay not exceeding 500ms
    - **Tenant #2**: Hadoop workloads (e.g., data analytics), with a maximum deadline of 15 minutes
    - Workload trace: Gmail & MSR
  - Power management
    - **AutoScale**, which is being used in Facebook's production system

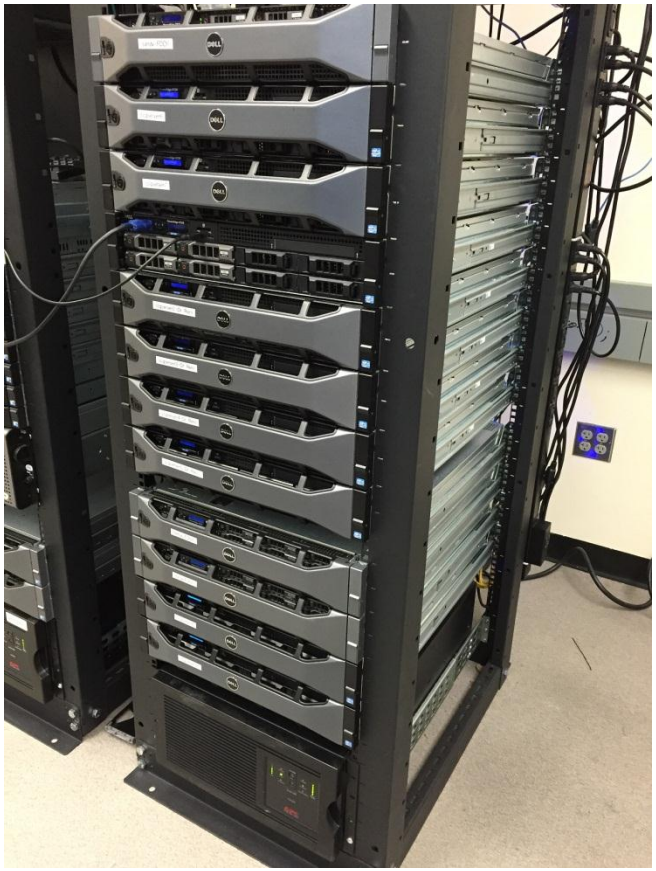

Our prototype system housed in FIU-SCIS data center

Settings

- 192 time slots, each 15 minutes
- Located in San Francisco, CA, cost pro-rated based on PG&E rate schedule E-20 for industry customers
- Benchmarks
  - **BASELINE**: power-based pricing without rewards
  - **EPR**: energy-based pricing (i.e., reward tenants based on electricity price)

RECO saves operator's OpEx

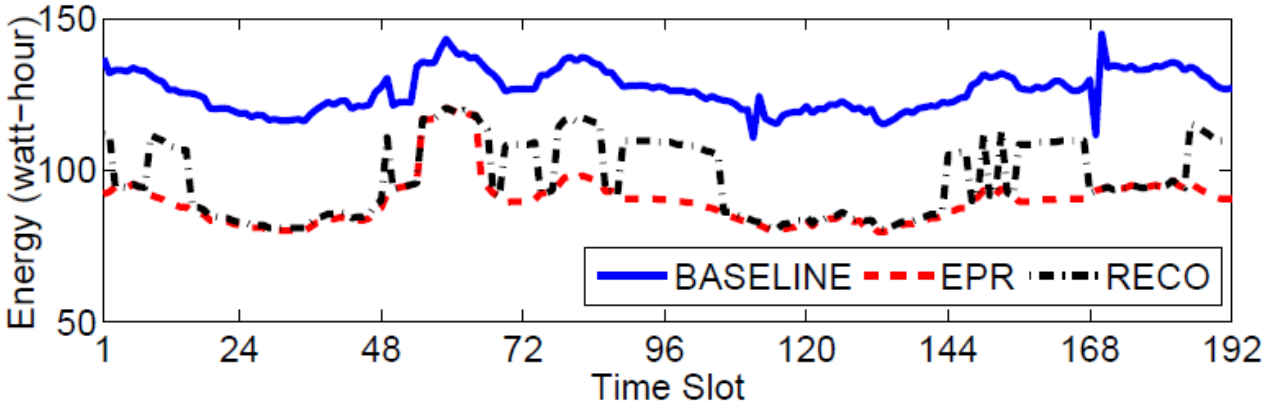

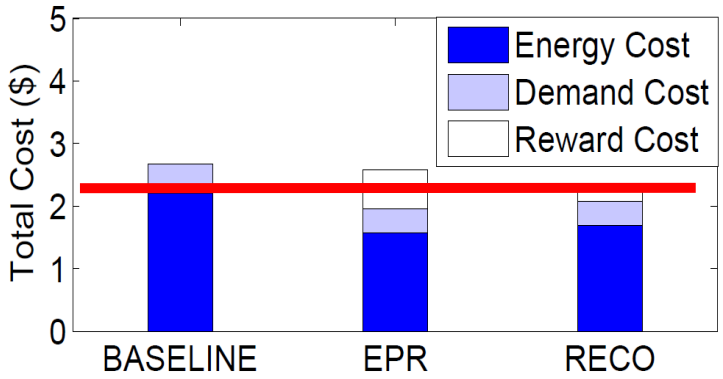

10+%
cost saving!

Tenants receive rewards for "free"

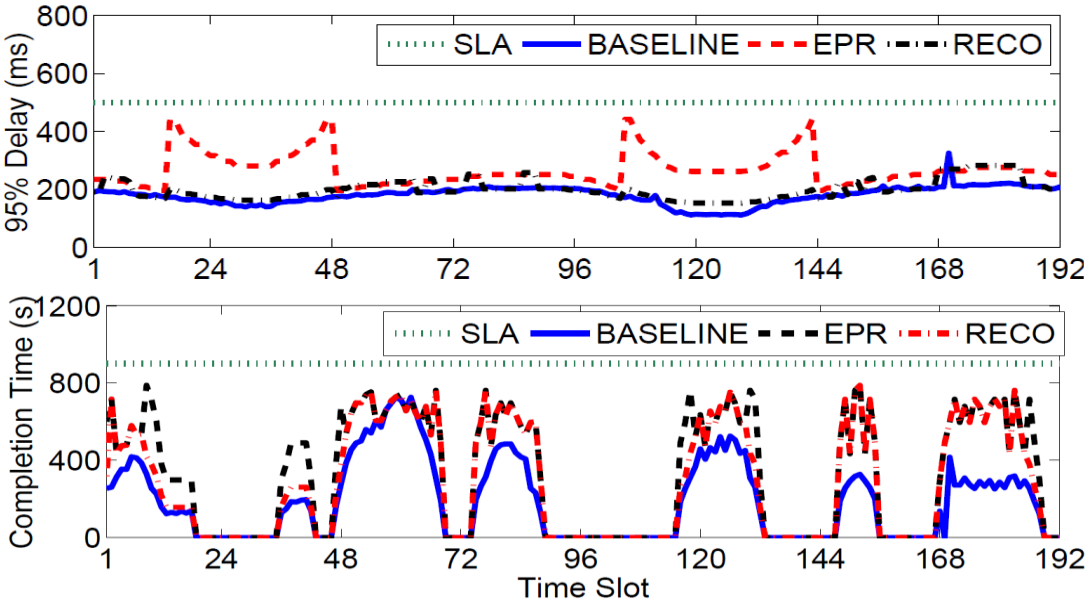

- Tenants 1 & 2 save 3.5% and 6.5%, respectively
  - ***Without*** violating SLA
- Results further confirmed by simulations

RECO is "win-win"!

Some messages

- Multi-tenant colocation data center is a very common and critical infrastructure for our Internet
  - **37.3%** energy of all data centers
- Electricity cost is nearly **40%** of operator's TCO
- **RECO**, an early step to reduce OpEx…

Our work v.s. others

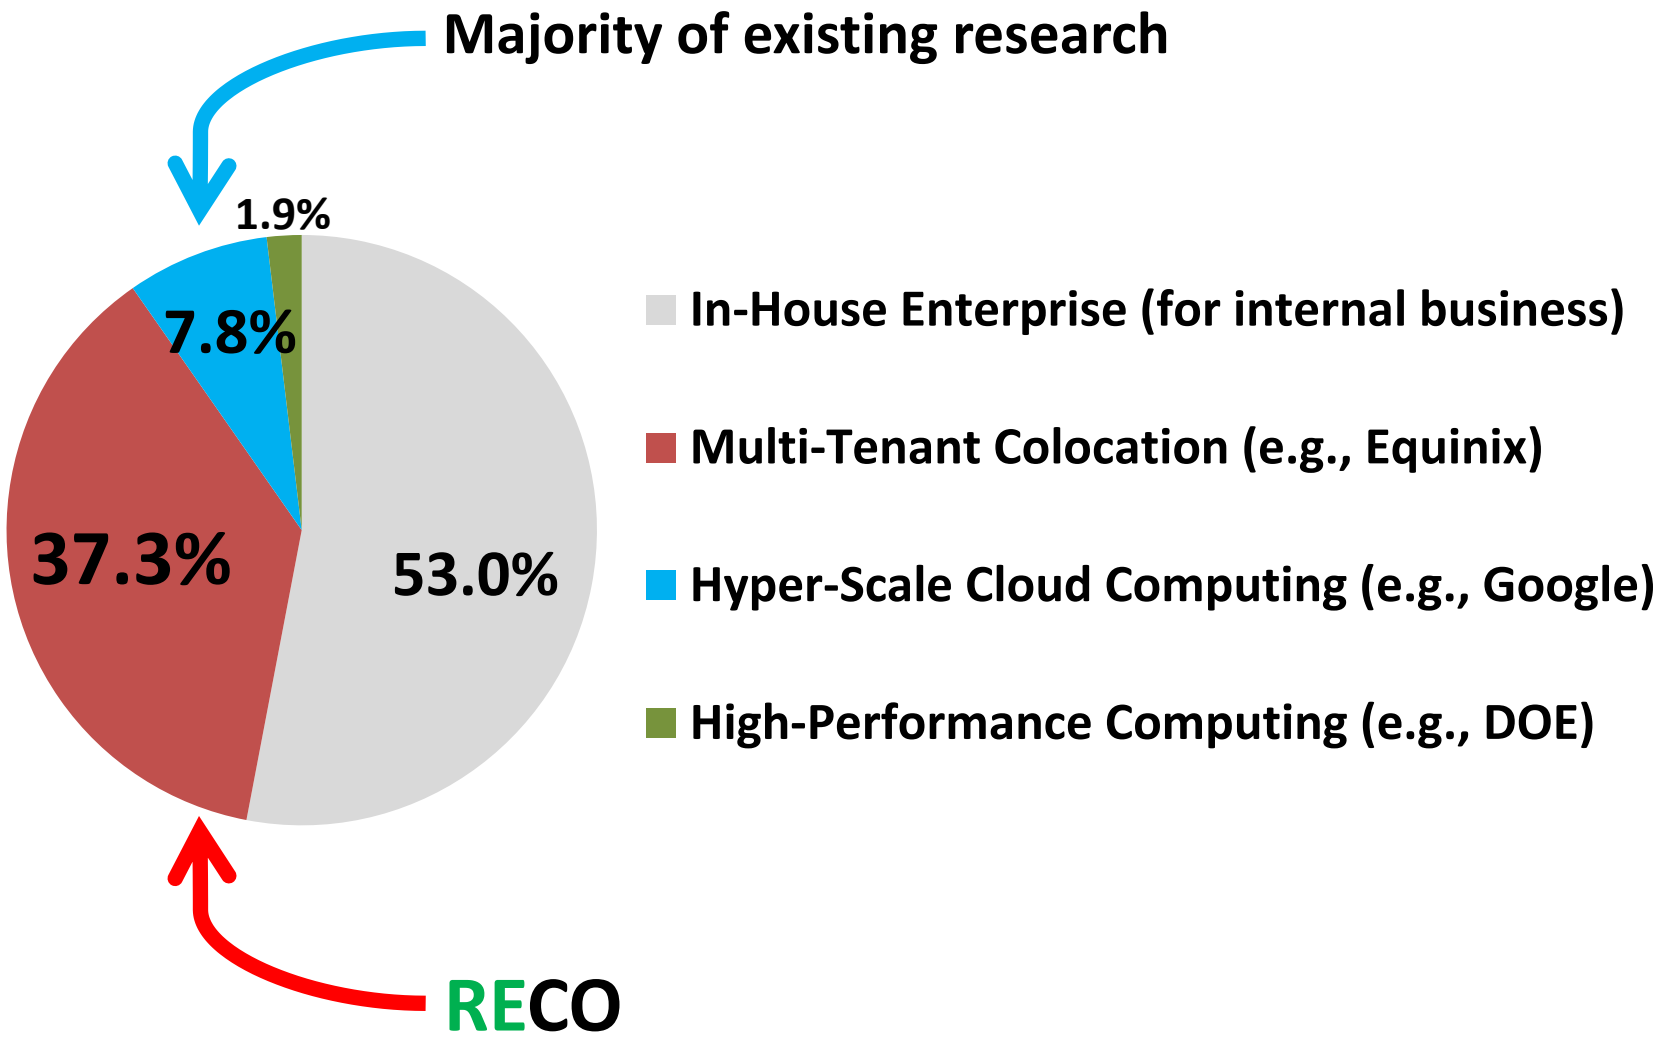

We're **not** alone…

colocation, Green Data Centers

Colocation Providers, Customers Trade Tips on Energy Savings

BY JORDAN NOVET ON NOVEMBER 1, 2013 ADD YOUR COMMENTS

21 Like | 22 Tweet | 21 | 1 +1

Environmental Issues > Energy Main Page > All Energy Documents

America's Data Centers Consuming and Wasting Growing Amounts of Energy

Critical Action Needed to Save Money and Cut Pollution

- *Data centers are one of the largest and fastest growing consumers of electricity in the United States. In 2013, U.S. data centers consumed an estimated 91 billion kilowatt-hours of electricity -- enough electricity to power all the households in New York City twice over -- and are on-track to reach 140 billion kilowatt-hours by 2020.*
- *Some large server farms operated by well-known Internet brands provide shining example of ultra-efficient data centers. Yet small, medium, and corporate data centers are responsible for the vast majority of data center energy consumption and are generally much less efficient.*
- *The largest issues and opportunities for energy savings include the under-utilization of data center equipment and the misalignment of incentives, including in the fast growing multi-tenant data center market segment.*
- *To move forward, systemic measures such as the public disclosure of efficiency metrics are necessary to create the conditions for best-practice efficiency behaviors across the data center industry.*

MORE FROM AUTHOR PIERRE DELFORGE

PUBLICATIONS

- America's Data Centers Consuming and Wasting Growing Amounts of Energy

RECENT BLOG POSTS

Why the majority of data centers are failing at energy efficiency

By **Heather Clancy**
Published August 26, 2014
Tags: Energy Conservation, Energy Efficiency

Email | Print | Single Page View

Hear more about data centers at the VERGE Salon NYC on Sept. 16.

Featured, Green Data Centers

NRDC: Multi-Tenant Data Centers Need To Play Bigger Energy Efficiency Role

BY JASON VERGE ON AUGUST 26, 2014 ADD YOUR COMMENTS

11 Like | 49 Tweet | 40 | 8 +1

Data centers are among the fastest growing users of electricity in the U.S. consuming an estimated 91 billion kilowatt-hours of electricity in 2013. Annual consumption is projected to increase by roughly 47 billion kilowatt-hours by 2020.

And while the industry has made progress in cutting energy waste, it still holds many wasteful practices, according to a Natural Resources Defense Council (NRDC) report.

Photo by Shaolei Ren at Chicago O'Hare International Airport in May, 2014