Cutting the Electric Bill for Internet-Scale Systems

Asfandyar Qureshi
MIT CSAIL
asfandyar@mit.edu

Rick Weber
Akamai Technologies
riweber@akamai.com

Hari Balakrishnan
MIT CSAIL
hari@mit.edu

John Guttag
MIT CSAIL
guttag@mit.edu

Bruce Maggs
Carnegie Mellon University
bmm@cs.cmu.edu

Motivation

- Increasing trends towards massive geographically distributed systems
- Large organizations have energy expenses in millions of dollars.
- Energy in – $Wh \approx n \cdot (Pidle + (Ppeak - Pidle) \cdot U + (PUE - 1) \cdot Ppeak) \cdot 365 \cdot 24$

Company	Servers	Electricity	Cost
eBay	16K	~0.6×10^5 MWh	~$3.7M
Akamai	40K	~1.7×10^5 MWh	~$10M
Rackspace	50K	~2×10^5 MWh	~$12M
Microsoft	>200K	>6×10^5 MWh	~$36M
Google	>500K	>6.3×10^5 MWh	~$38M
USA (2006)	10.9M	610×10^5 MWh	$4.5B

Observations and Idea

- Electricity prices vary between different locations at different time
- Mechanisms for request routing and replication already exist
- Develop a cost-aware request routing policy that maps requests to locations where energy is cheaper

Main Goals

- Identify relevance of electricity price differentials to large distributed systems
- Estimate cost savings if the system is deployed
  - Analysis concerned with reducing costs and not energy

Electricity Market

- Regional Transmission Organization (RTO)
  - Bilateral contracts (majority)
  - Wholesale electricity (40% of total electricity)
    - Auctioning mechanism: producers (supply offers), consumers (demand bids)
- Market Types
  - Day-ahead markets (futures)
    - Provide hourly prices for delivery during following day based on expected load
  - Real-time markets (spots)
    - Prices calculated every 5 minutes based on actual condition rather than expectations (small fraction of total energy transactions)

Price Variations

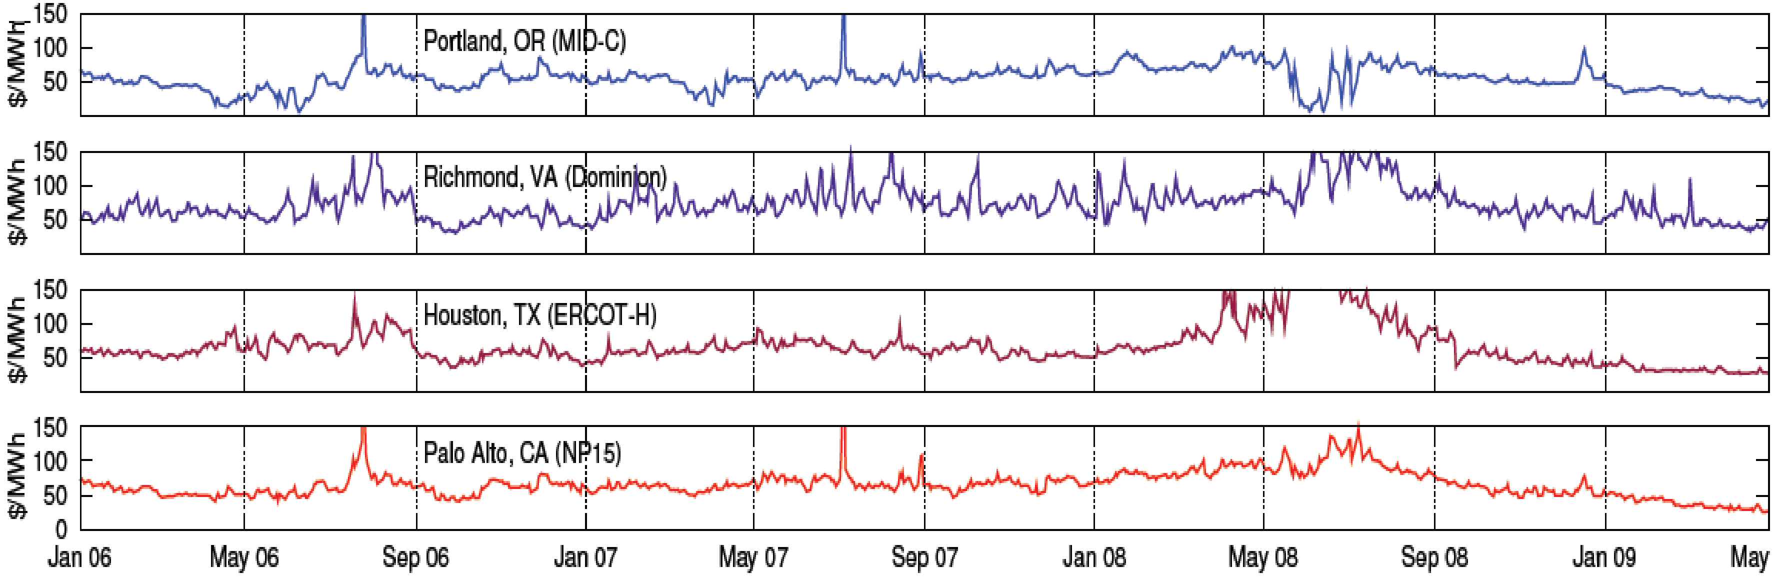

- Figure: Daily averages of day-ahead peak prices at different hubs
- The elevation in 2008 correlates with record high natural gas prices, and does not affect the hydroelectric dominated Northwest (seasonal rainfall).
- Last downward trend due to global economic downturn

Geographic Correlation

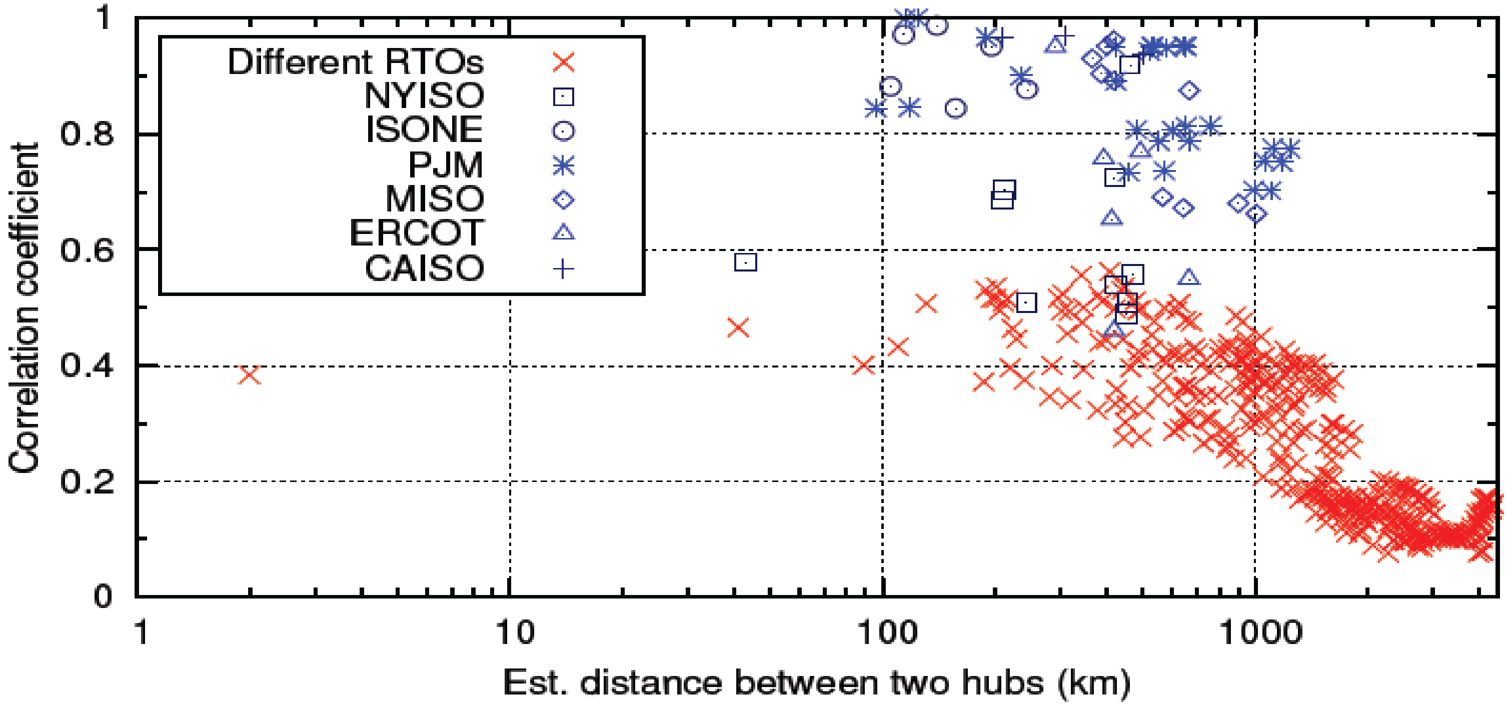

- Figure: Relationship between price correlation, distance, and parent RTO.
- Each point represents a pair of hubs (29 hubs, 406 pairs), and the correlation coefficient of their 2006-2009 hourly prices (> 28k samples each).
- Red points represent paired hubs from different RTOs.
- Blue points are labeled with the RTO of both.

Price Differentials

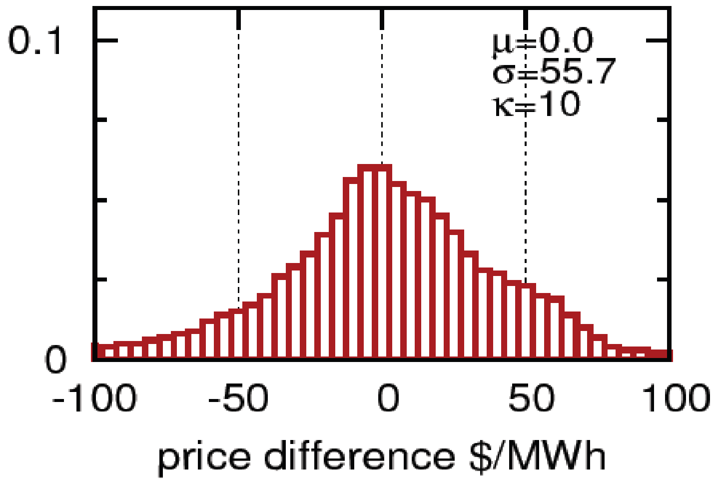

PaloAlto - Virginia

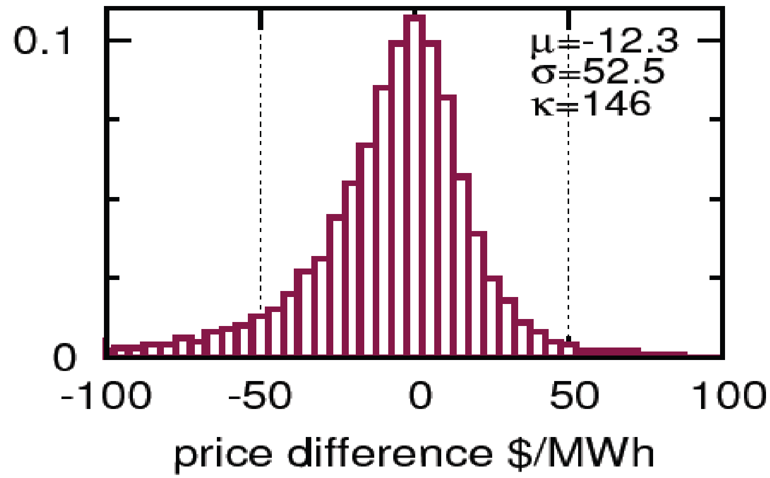

Boston - NYC

- Figure: Price differential histograms for five location pairs and 39 months of hourly prices.

Price Differentials

- Differential Distribution
- Evolution in Time
- Time-of-Day
- Differential Duration

Akamai – Traffic and Bandwidth

- Data set detailing traffic on Akamai's infrastructure
  - Covers 24 days with Peak rate of over 2 million hits/sec
  - 9-region traffic for which electricity price available
- Traffic Data
  - Data collected at 5 minute intervals on servers housed in Akamai's public cluster
  - Number of hits, bytes, clients' geographic location, load in each cluster
- Bandwidth Costs
  - Significant for Akamai (system is optimized to reduced bandwidth)
  - Cannot be ignored (But treatment of bandwidth costs is abstract)
- Client-Server Distances

Cluster Energy Consumption

- Adapted from Google's empirical study of data center
- P – Power usage of a cluster:

$$P_{cluster}(u_t) = F(n) + V(u_t, n) + \epsilon$$

- F – Fixed power
- V – Variable power

$$\begin{aligned} F(n) &= n \cdot \left(P_{idle} + (PUE - 1) \cdot P_{peak}\right) \\ V(u_t, n) &= n \cdot (P_{peak} - P_{idle}) \cdot (2u_t - u_t^r) \end{aligned}$$

- Critical in determining savings: $\frac{P_{cluster}(0)}{P_{cluster}(1)}$

Simulation

- Discrete time simulation
- Electricity Prices
  - Hourly real-time prices 29 locations
- Traffic and Server Data
  - 5 minute samples for 9 clusters (US traffic)
- Routing schemes- Original, Distance Constrained
- Energy Model
- Client-Server Distances

Results – 24 Days of Traffic

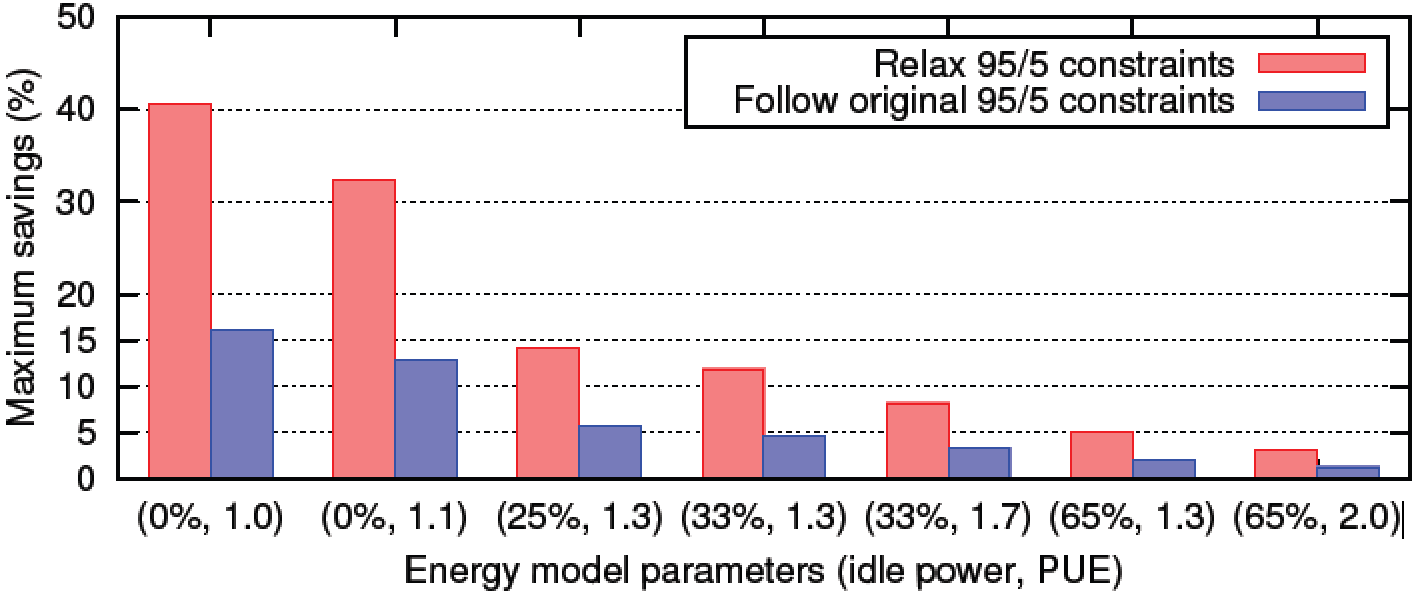

- 24-day savings for a different PUE and Pidle values with a 1500km distance threshold.
- Key : energy elasticity
- Obeying existing 95/5 bandwidth constraints reduces savings.

Results – 24 Days of Traffic

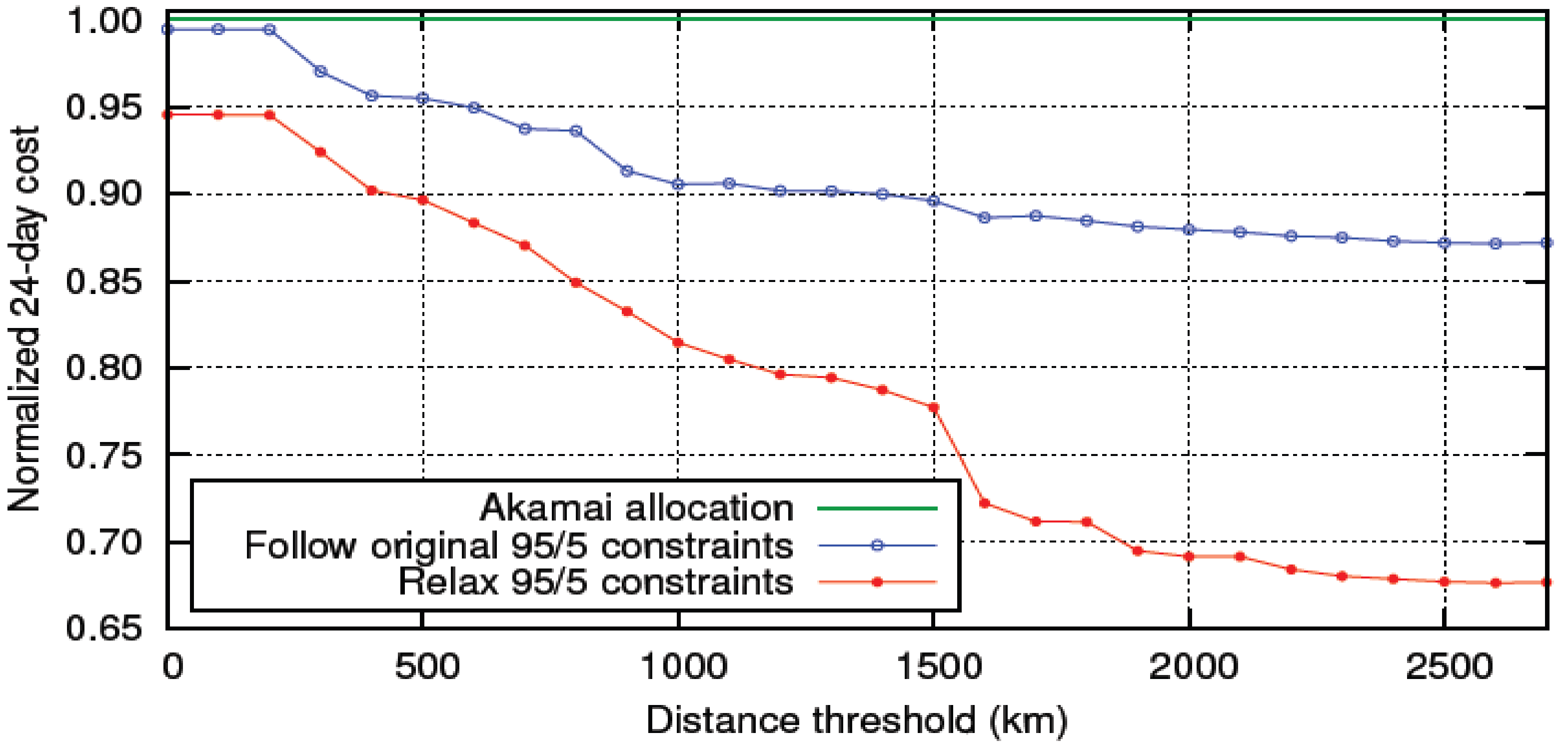

- Increasing distance threshold reduces electricity costs

Results – Synthetic Workload

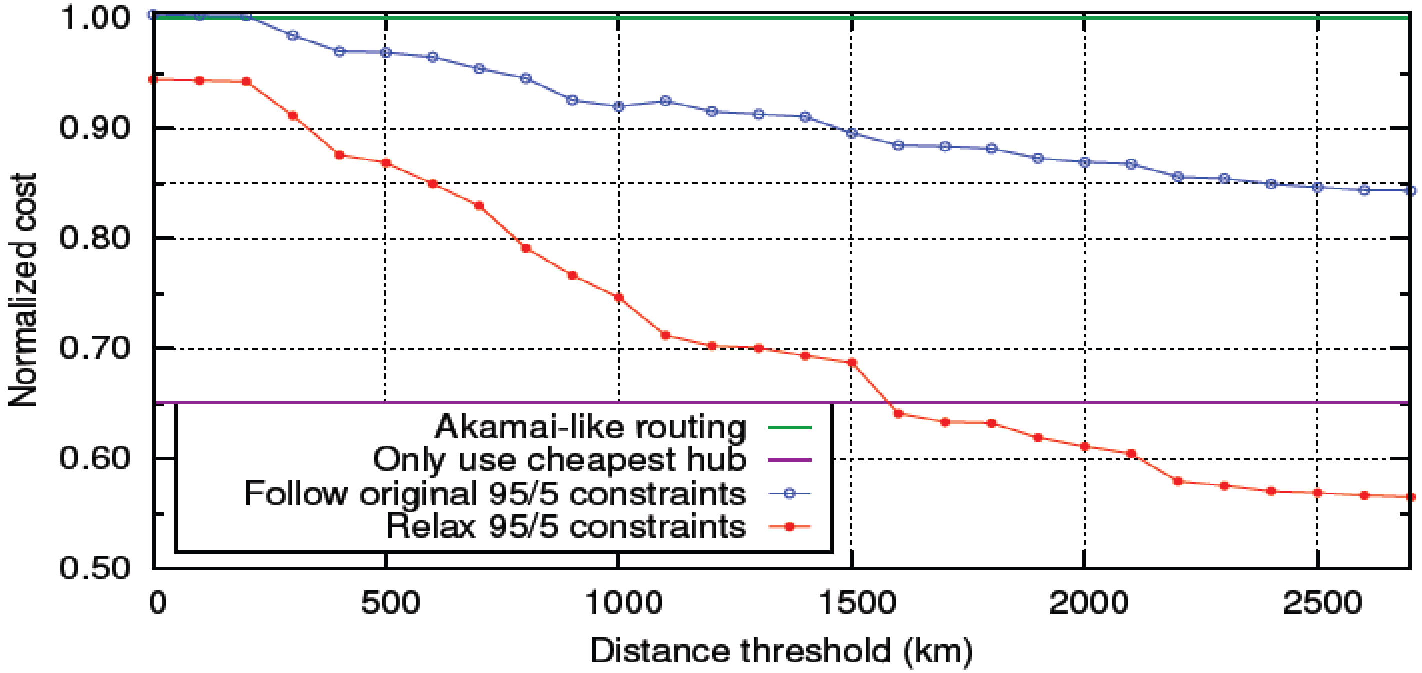

- Distance savings similar to 24 days simulation
- Dynamic beats Static (95/5 constraints ignored)

Reaction delay

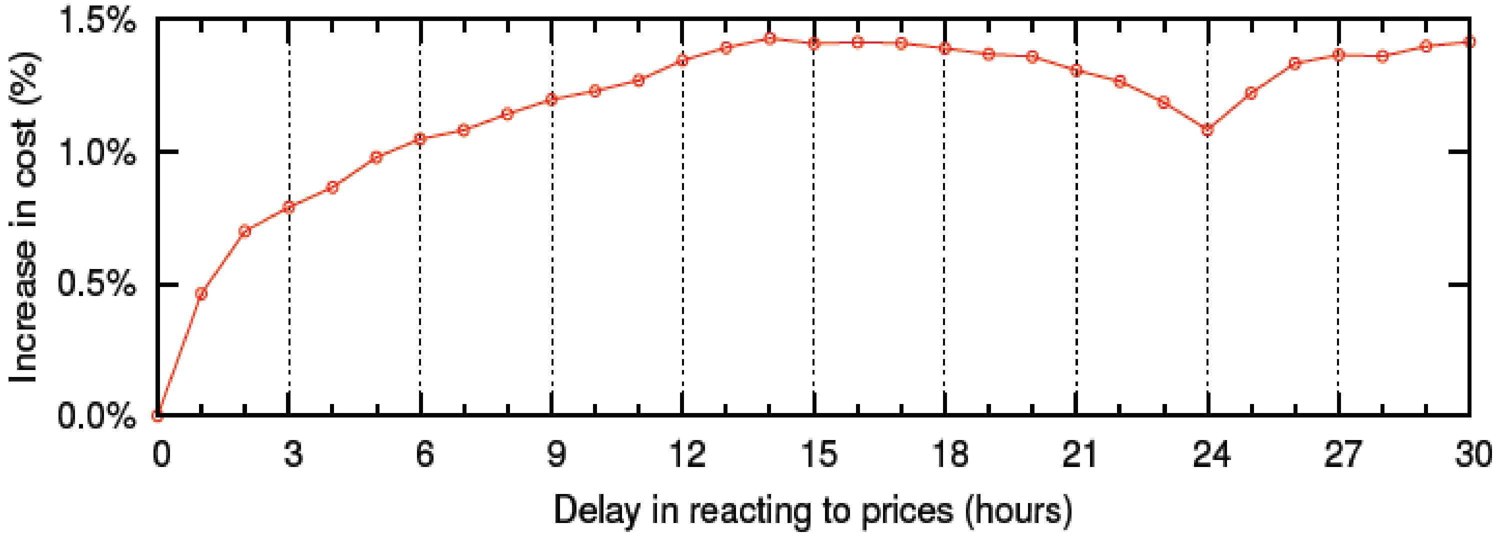

- Savings reduced when reaction to price change is slow

Future Work

- Implementing joint optimization
- RTO interaction
- Weather differentials
- Environmental costs