plugging into energy market diversity

asfandyar qureshi | **MIT-CSAIL**

photograph by .bullish @ flickr

electricity market deregulation

diversification of energy sources

distributed systems

escalating energy expenses

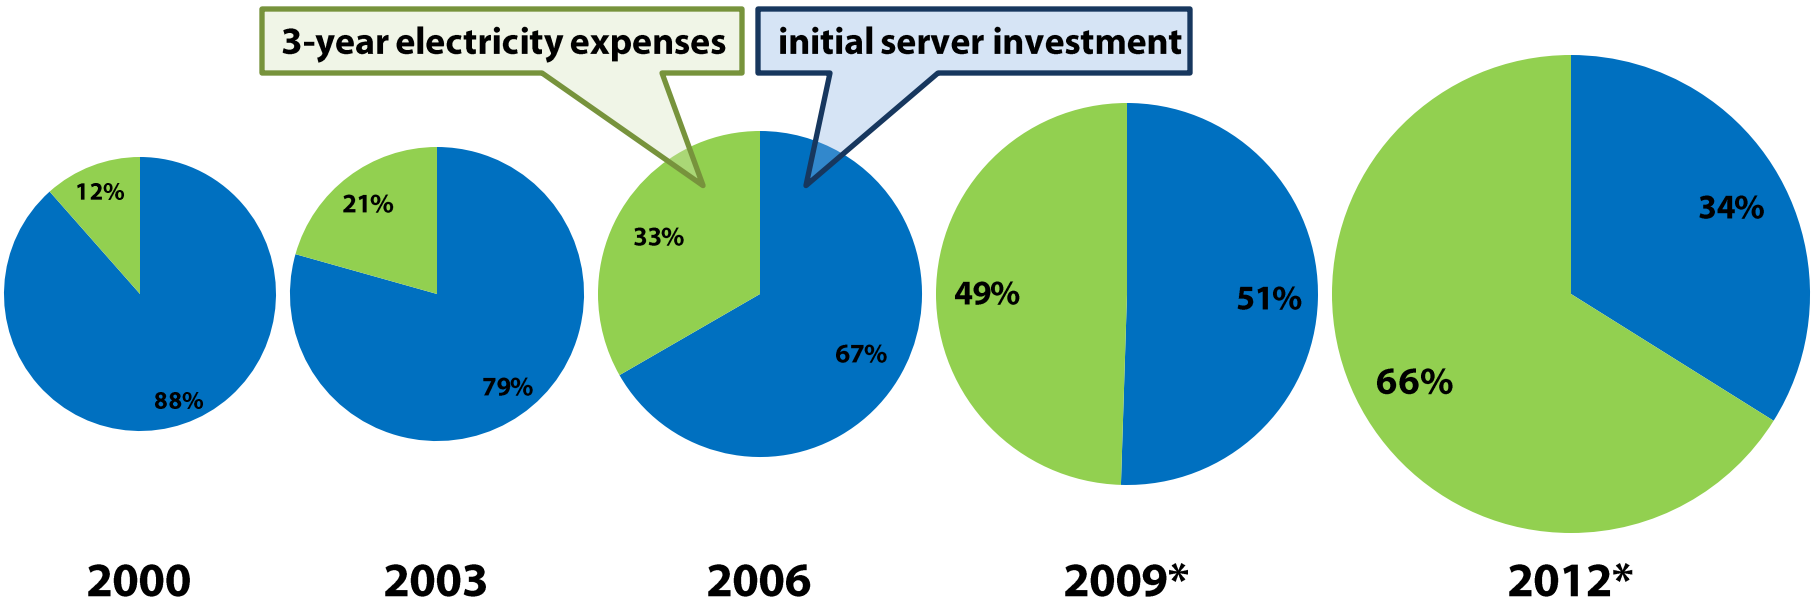

data centers consume many, many megawatts

- electricity is large fraction of total operating costs
- USA: around five gigawatts ~ $2.7 billion in 2005
- consumption projected to triple between 2005 and 2010

sources: koomey 2007, uptime institute

lowering electricity bills

more efficient data centers

- hardware, virtualization, cooling, DC power, …
- lots of potential here, but not what this talk is about

consume cheaper electricity

- data center planners already consider geographic variation
  - **based on annual price averages**
- e.g. build data centers near hydro-electric power plants
  - **Google in the Dalles, Microsoft in Northlake**

what else can we do?

observations

- electricity prices are not correlated at different places
- we are building distributed replicated systems
- computation can be moved in ways electricity cannot

pay attention to locational computation cost

- e.g. what does it cost to serve a single web request?
- what is the impact of electricity prices on this?

compute in cheap energy markets

- if a computation can be performed in multiple places, use up-to-date market prices to pick the cheaper replica.

could this work?

- infrastructure owners: 2%-10% savings are plausible

the fluidity of computation

hard: electricity transport

- inefficient grid, physics, markets seams issues, etc.
- price disparities arise

relatively easy: relocating computation

- we've proven we're good at this…
- we don't always need to move far away
- cheaper energy can be nearby (network distance)

when to relocate

what is the cost of running a web service?

- network, electricity, amortized infrastructure cost etc.
- some fixed / some depends on demand

comparing replicas A and B

- should a client's request be handled at A or B?
- what is the impact on client QoS?
  - **we don't analyze this rigorously, leaving it to future work…**
- what is the difference in electricity prices?
  - **for simplicity we assume network and fixed costs are similar**
- relocate when price differences overshadow QoS impact

observed electricity prices

how exactly do prices vary?

- anything we can exploit?
- interested in price differentials

purely observational approach

- not a theoretical/model-driven analysis of markets

historical market data

- over two years worth of weekday prices
  - **wholesale day-ahead market prices**
  - **seven different locations, six different markets**
  - **compiled from archives of daily market news publication**
- annual averages for three years
  - **publically available from US federal government**

historical market prices

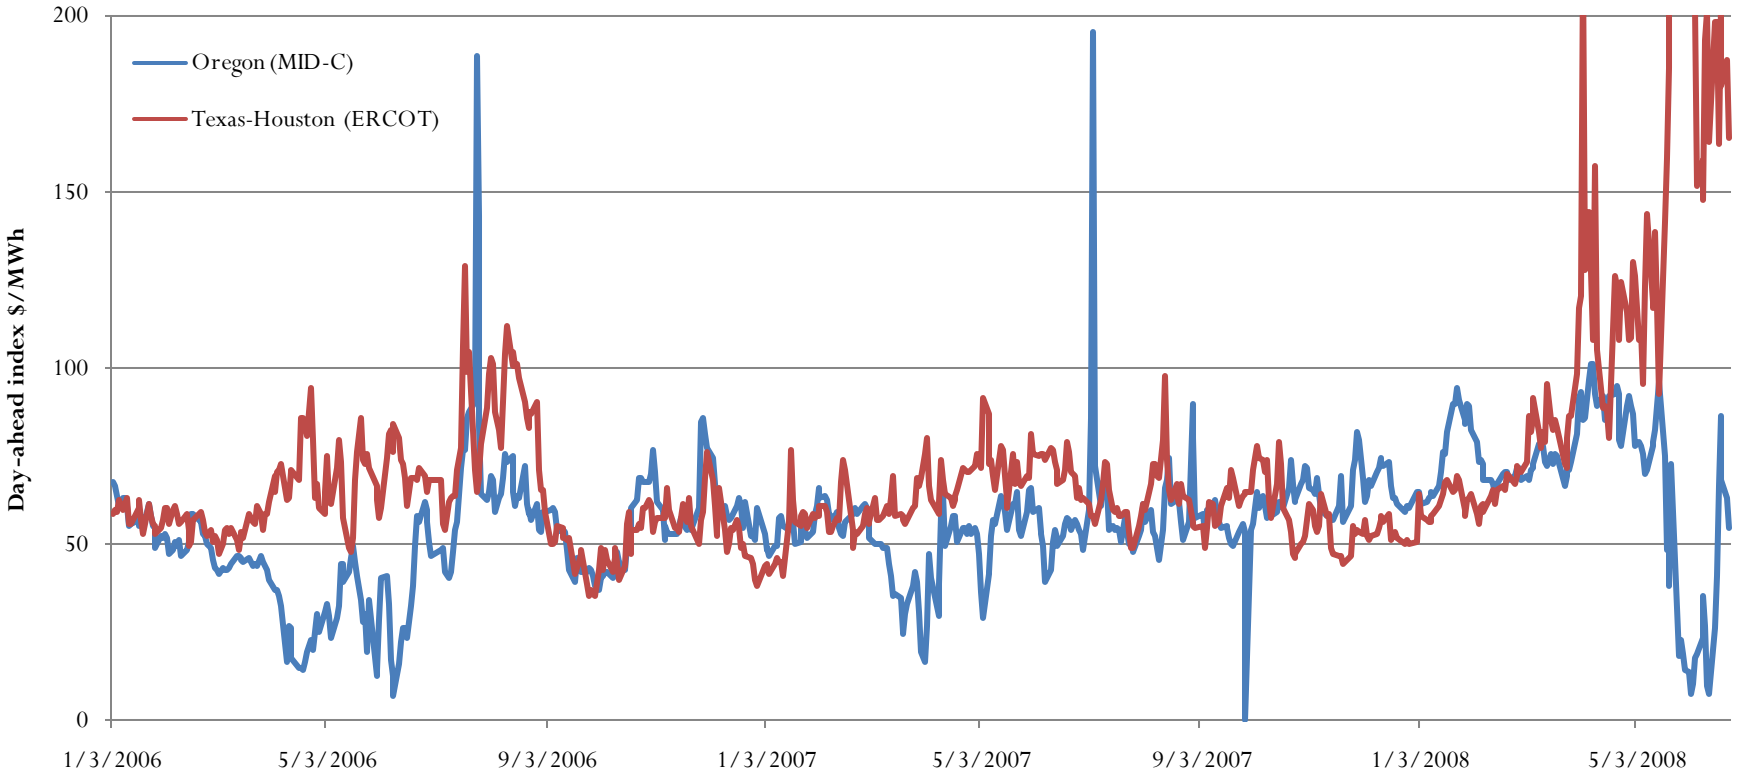

lots of day-to-day variation

- seasonal effects, order of magnitude differences, …
- significant differences exist; not perfectly correlated

prices: things change

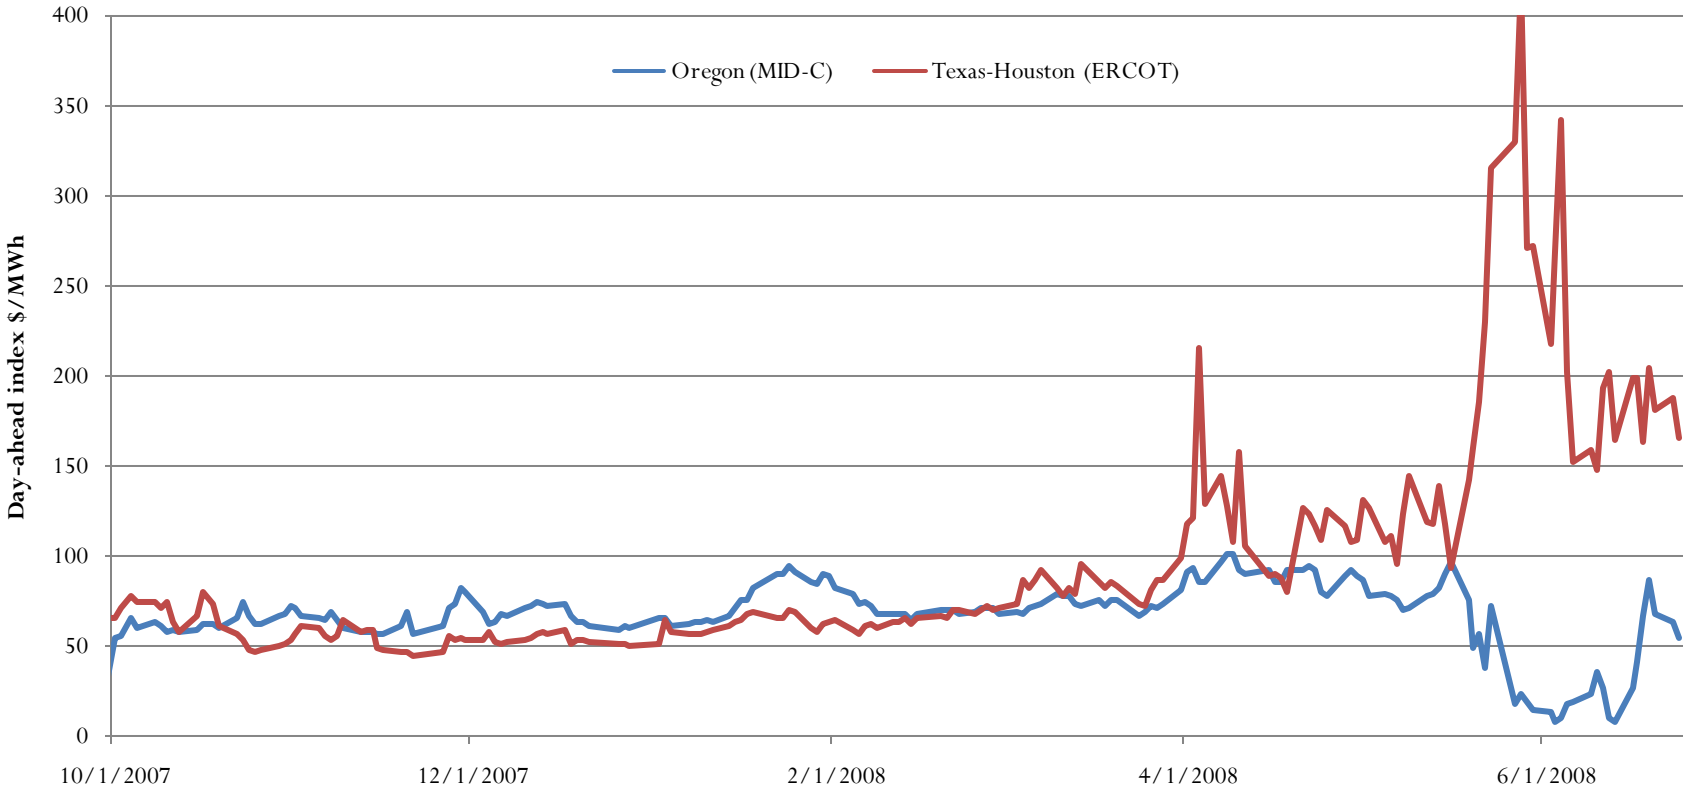

we can't predict the future…

- best choice today is not necessarily best tomorrow
- market deregulation, record high gas prices, …

reasons for price differences

cannot transport electricity efficiently

- inefficient electric grid; not all points are connected

cannot store electricity efficiently

- demand and supply must be continuously balanced

different sources of energy

- different dependencies
  - **weather (rainfall, wind, …)**
  - **fuel costs (natural gas, coal …)**

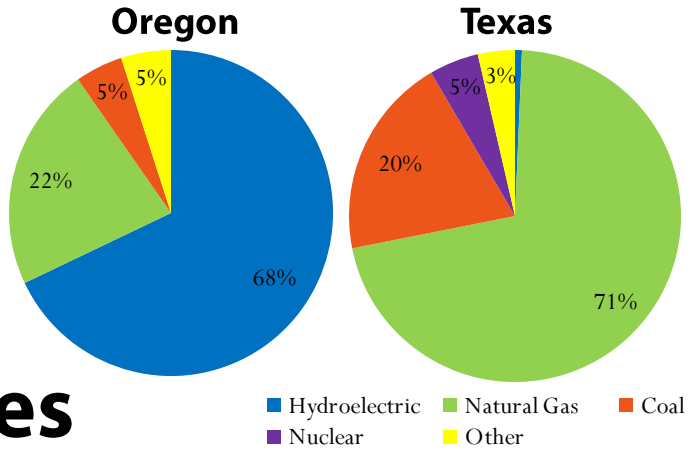

regional demand differences

- level of local demand relative to supply affects prices

other factors

- malfunctions, storms , deregulation, manipulation…

exploiting price differences

how do we go about building a mechanism?

- shuffle computation between energy markets
- shutdown (some) resources in high cost markets

distributed system model

- **N** nodes placed in different energy markets
- any **R<N** nodes form a complete replica of the system
- already build this way for performance and reliability
- nodes can be data centers or rented racks
  - **so this applies to both Google and Akamai**

electricity market model

- day-ahead model: buy today for delivery tomorrow
- wholesale markets: day-to-day flexibility

selective blackouts

dynamically disable expensive nodes

- optimize once a day, using day-ahead electricity prices
- rank nodes by operating expense
- ensure **R** least expensive are always running
- only disable others if price differences exceed threshold

useful properties

- parameters allow expense/performance trade-offs
- only update routing once a day (so DNS works…)
- enough warning time for provider to re-sell electricity
- max active nodes, unless price differentials are large

analyzing effectiveness

seven locations

- market diversity
- IXP proximity

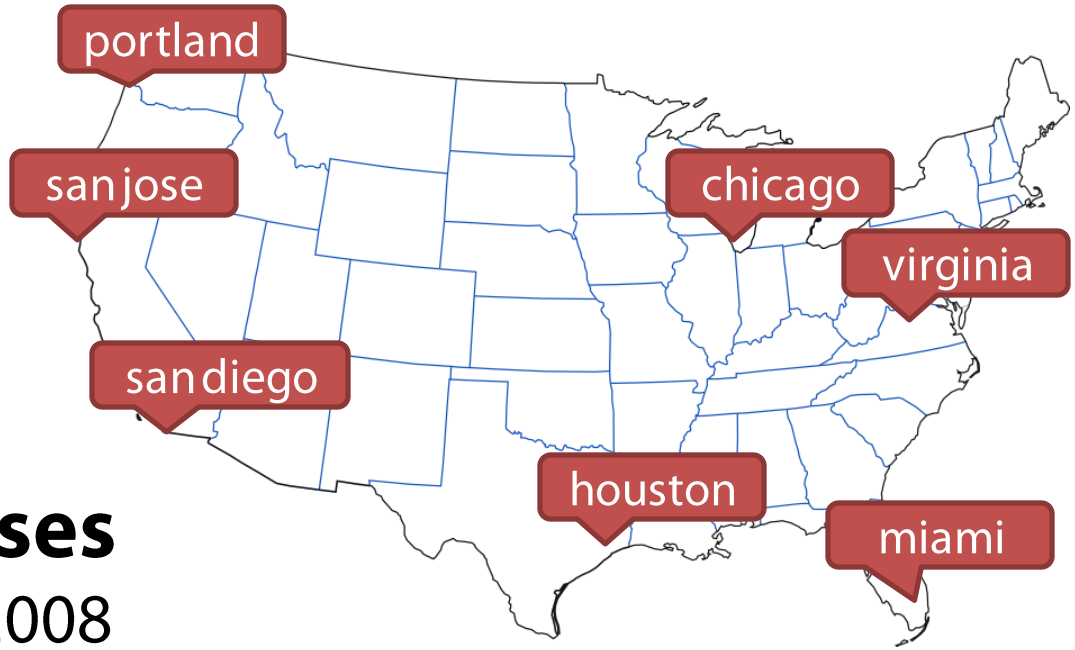

total energy expenses

- jan 2006 through july 2008
- observed wholesale market prices
- what would selective blackouts have done?
- assume constant demand, spread evenly over active nodes
- our results are upper bounds on savings

static solution is not optimal

best R nodes change with time

- conventional: pick R nodes to minimize average expenses
- blackouts: every day minimize next-day expenses

significant savings are possible

- 2%-30% reduction in total electricity expenses
- depending on performance targets and excess capacity

varying degrees of flexibility

R = N – 1 (one redundant replica)

- **2.5%** reduction in electricity expenses
- all seven online **96%** of the time (only 27 days with six)
- transient events cause significant expenses

R = $\lceil N/2 \rceil$ (double capacity)

- **7.2%** reduction in electricity expenses
- six or more online **74%** of the time
- can match cheapest market, with better performance

R = 1 (six redundant replicas)

- **26.5%** reduction in electricity expenses
- only one replica online **76%** of the time

who really saves?

contracts…

- data centers may be locked into fixed prices
- constant charges for rented rack space

the public good argument: it doesn't matter

- high prices can be equated with resource scarcity
- avoid putting strain on local resources if we don't have to
- won't save money, but others will benefit…

incentives exist for rational selfish actors

- right now, elevated prices impact someone
  - **bankruptcy of Pacific Gas and Electric in California '01**
- if we avoid consuming, there is a tangible benefit
- should be possible to negotiate savings-sharing deals

large or small?

monolithic data centers

- find cheapest average electricity
- cost hundreds of millions of dollars
- economies of scale

many smaller nodes

- improved performance and reliability
- dynamically minimize electricity cost
  - **does this challenge earlier economies of scale?**

smaller still…

- data-center-in-a-container
- ~250 machines

looking ahead…

we already do a lot of dynamic optimization

- electricity prices can be another input

greener computing: minimize carbon usage

- just a different cost function…

pricing in clouds

- differentiate pricing based on energy costs
- auctions based on price/performance

hour-ahead markets

- price variation is more pronounced

increasing electricity price variation

- already here: locational marginal pricing
- future: passing price variation on to consumers

questions?

photograph by
robertjosiah @ flickr