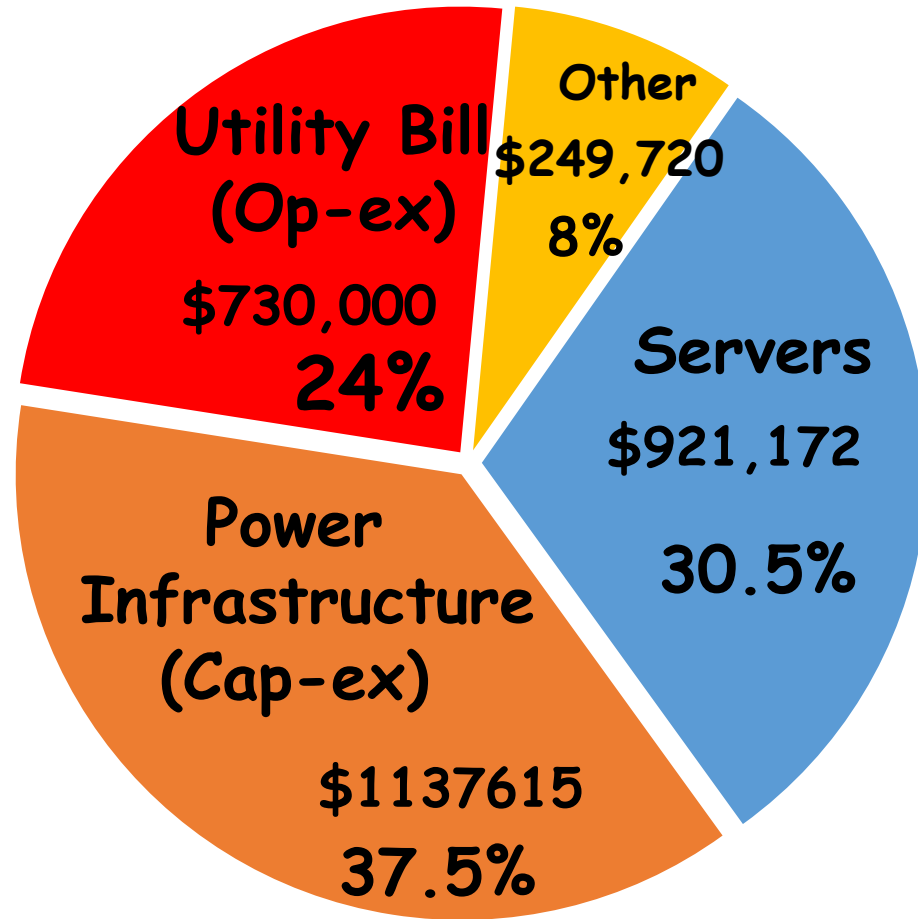


A Case for Virtualizing the Electric Utility in Cloud Data Centers

Cheng Wang, Bhuvan Urgaonkar, George Kesidis,
Uday Shanbhag, Qian Wang
The Pennsylvania State University



Powering Data Centers is Expensive!



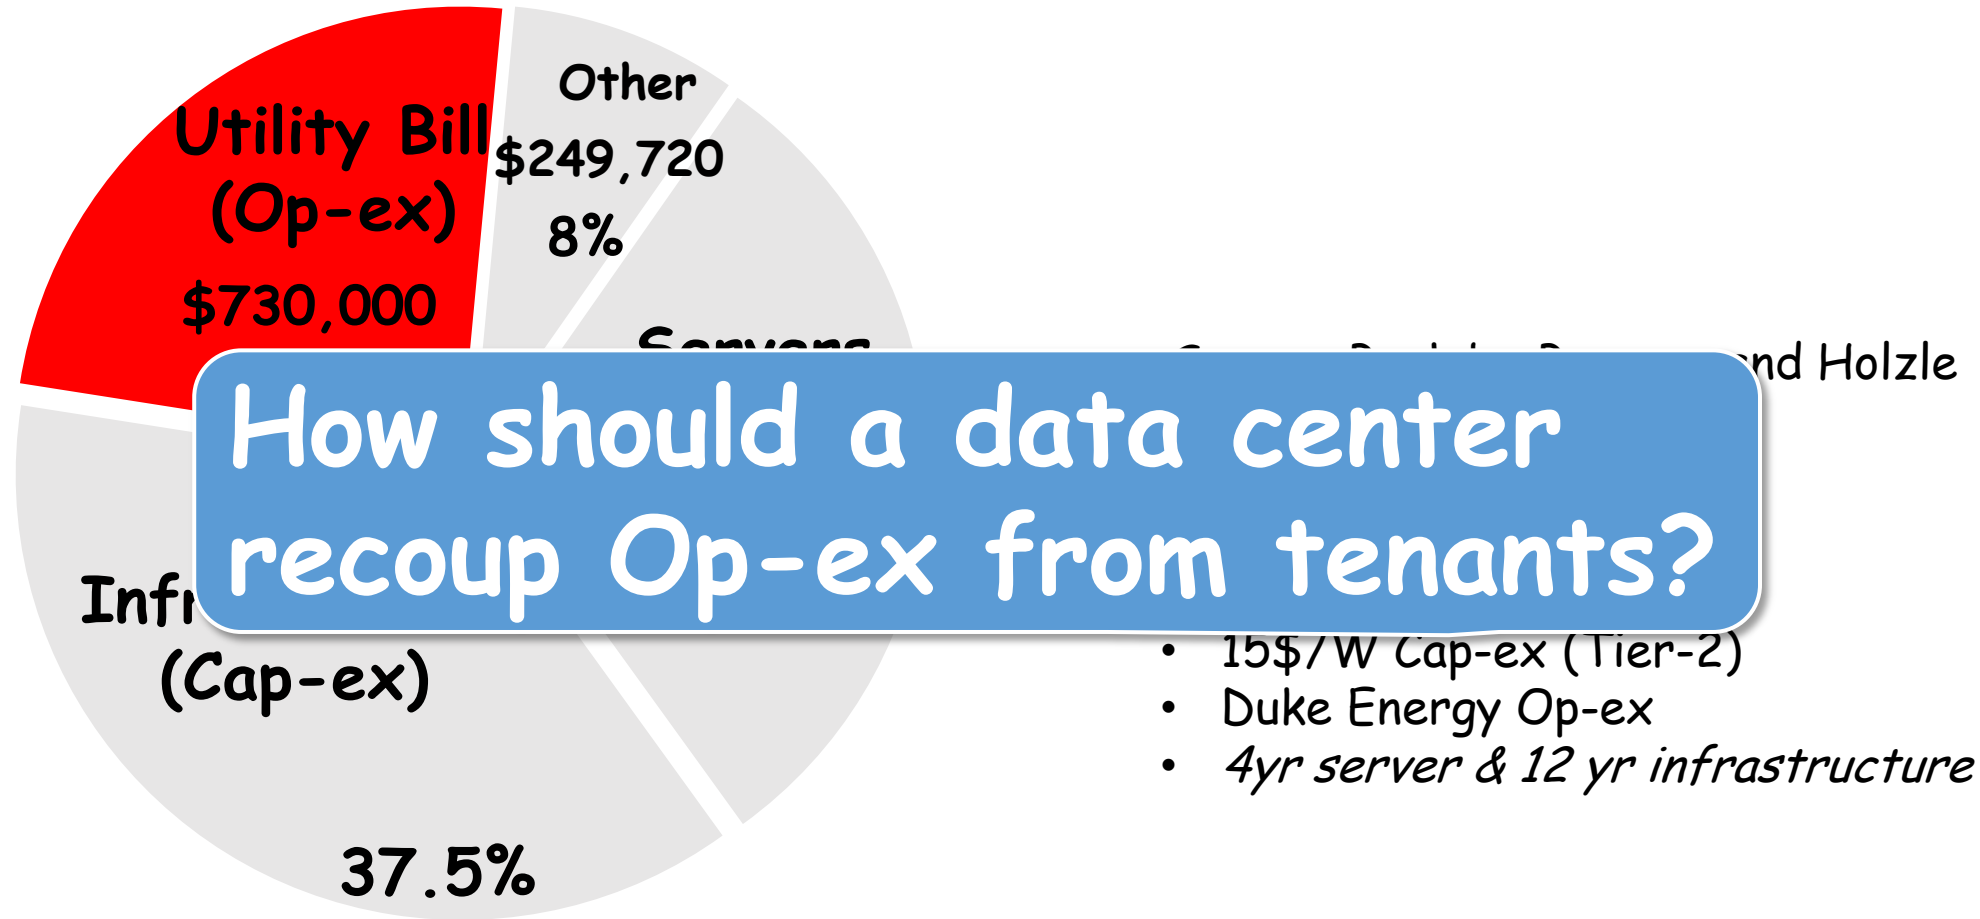
*All cost are amortized
at a monthly granularity*

Source: Book by Barroso and Holzle

Assumptions:

- 20,000 servers
- 1.5 PUE
- 15\$/W Cap-ex (Tier-2)
- Duke Energy Op-ex
- 4yr server & 12 yr infrastructure

Powering Data Centers is Expensive!



*All cost are amortized
at a monthly granularity*

How is Op-ex Recouped Today?

- Cloud resource interface is purely IT-based
 - E.g., IaaS: VMs, bytes stored, bytes transferred, ...; SaaS: queries, sessions, ...

Cloud provider	Compute		Storage (\$/GB·Month)	Data transfer out (\$/GB)
	Mem (GB)	Price (\$/Hour)		
Amazon EC2	7.5	0.14	0.05	0.12*
Windows Azure	3.5	0.12	0.068*	0.12*
Rackspace	2	0.08	0.12	0.12
Google Cloud	7.5	0.14	0.026	0.12*
SoftLayer	8	0.19	0.10	0.10
ProfitBricks	7.5	0.14	0.04	0.06
CloudSigma	7.5	0.15	0.14	0.05

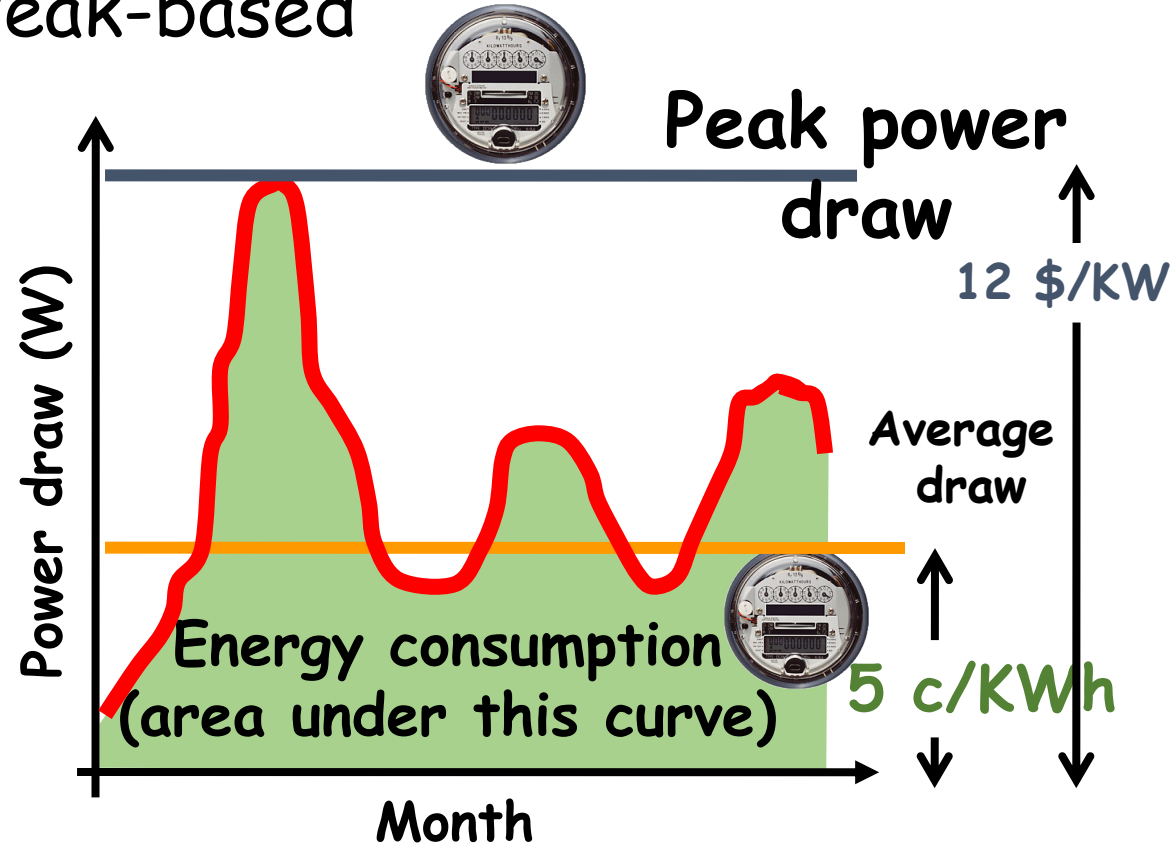
Compute prices are based on a Linux virtual machine with 2 vCPUs.

*: Tiered pricing. Only show price for first 1TB.

- Energy-related costs **bundled into** cloud resource prices without considering idiosyncrasies of electric utility pricing

Electric Utility Pricing

- Peak-based 15-min



Peak to
Average
ratio

3:1

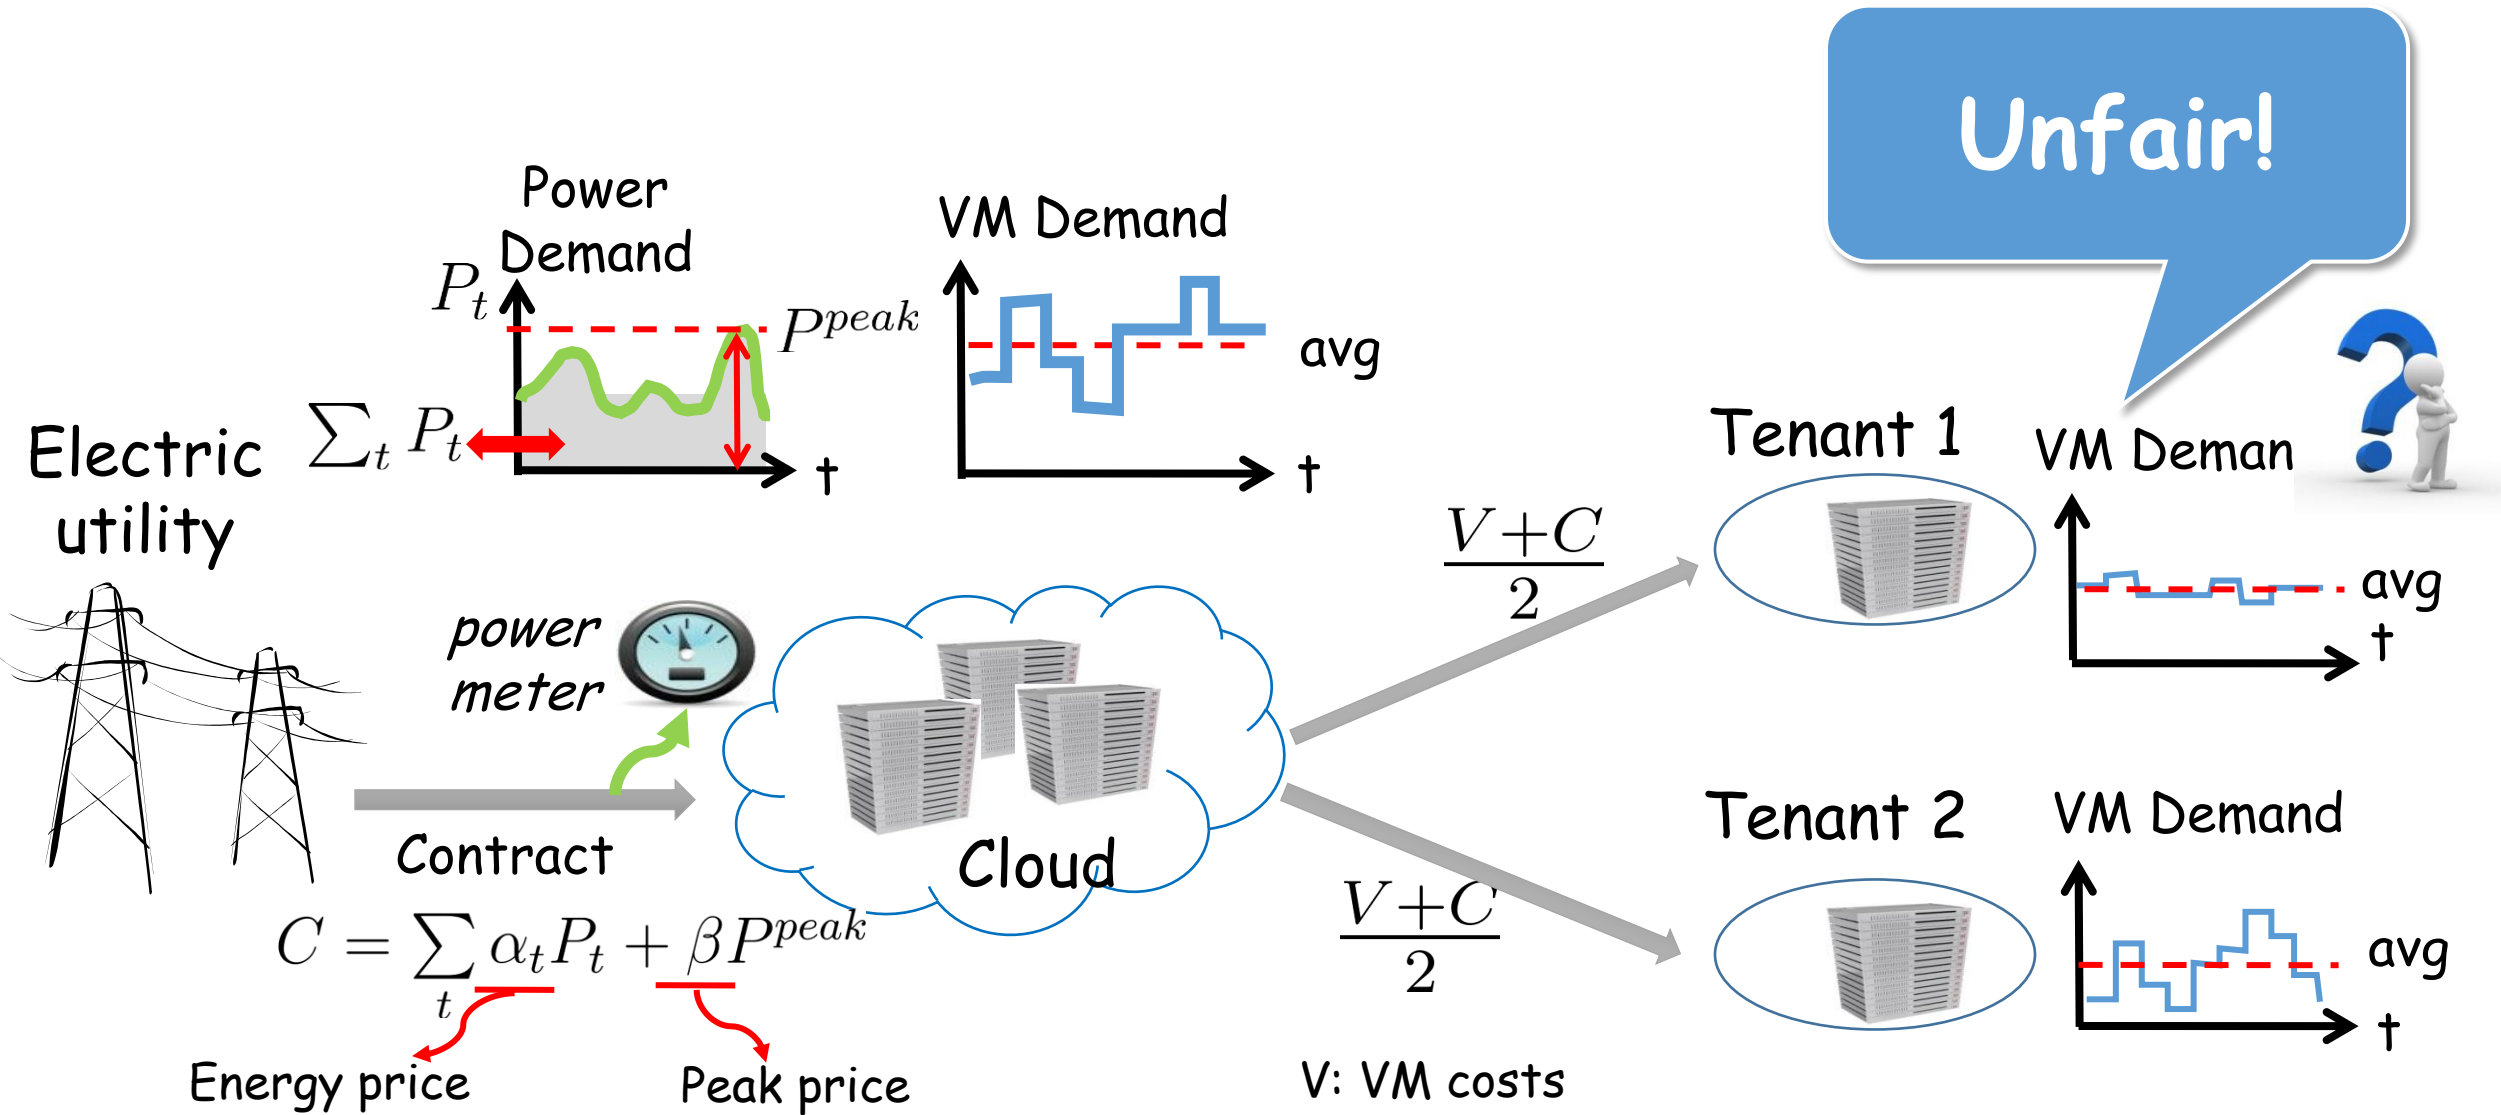
Duke Utility Tariffs
(12 \$/KW, 5 c/KWh)

Note: Tariff rates collected from Duke Energy Utility.

Gap: Utility vs. Cloud Resource Pricing

- Shortcomings due to this gap:
 - Lack of fairness (in how tenants are charged)
 - Loss of cost-efficacy
- Fixing these shortcomings:
 - Why is it non-trivial?
 - Our proposal: virtualize the electric utility

A Thought Experiment



Fixing the Gap: Key Idea

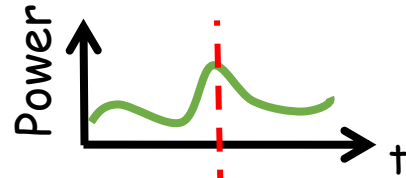
Pass on utility pricing structure to tenants

- Consider a strawman approach that passes on utility pricing "as is" to tenants

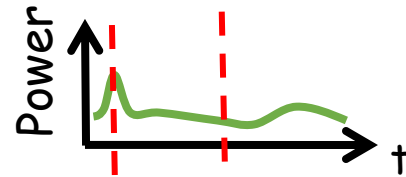
Strawman #1

- Base each tenant's peak charge on its local peak
- What if a tenant's individual peak does not coincide w/ aggregate peak?

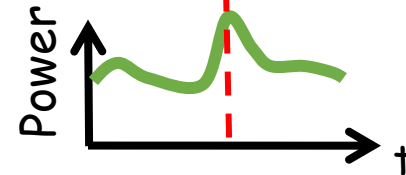
Tenant 1



Tenant 2



Cloud



Still unfair!



Fixing the Gap: Key Idea, Take 2

Pass on utility pricing structure to tenants such that
tenants are charged for their contribution to Op-ex

- What about charging tenants for their contribution (if any) to the aggregate peak?

Strawman #2

- Base each tenant's peak charge on its contribution to aggregate peak
- Tenants' demands may react to price differently
- Consider a single time slot optimization for cloud

Assume

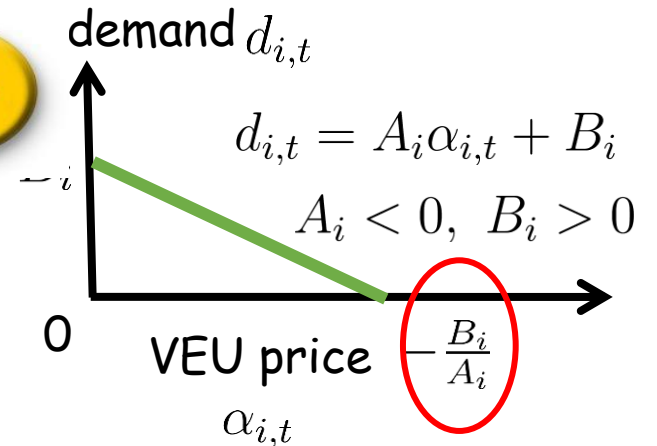
VM price

Instance cost

VM revenue $\theta k \sum_i \alpha_{i,t}$

VEU revenue $\sum_i d_{i,t} \alpha_{i,t}$

Not cost-effective!



Cloud profit $V_t = \sum_i d_{i,t}(\alpha_{i,t} + \theta k) - (\alpha_t + \lambda k) \sum_i d_{i,t}$

$$\max_{\alpha_{i,t}} V_t \longrightarrow \frac{\partial V_t}{\partial \alpha_{i,t}} = 0 \longrightarrow \alpha_{i,t}^* = \frac{1}{2}(\alpha_t + \lambda k - \theta k - \frac{B_i}{A_i})$$

Fixing the Gap: Key Idea, Take 3

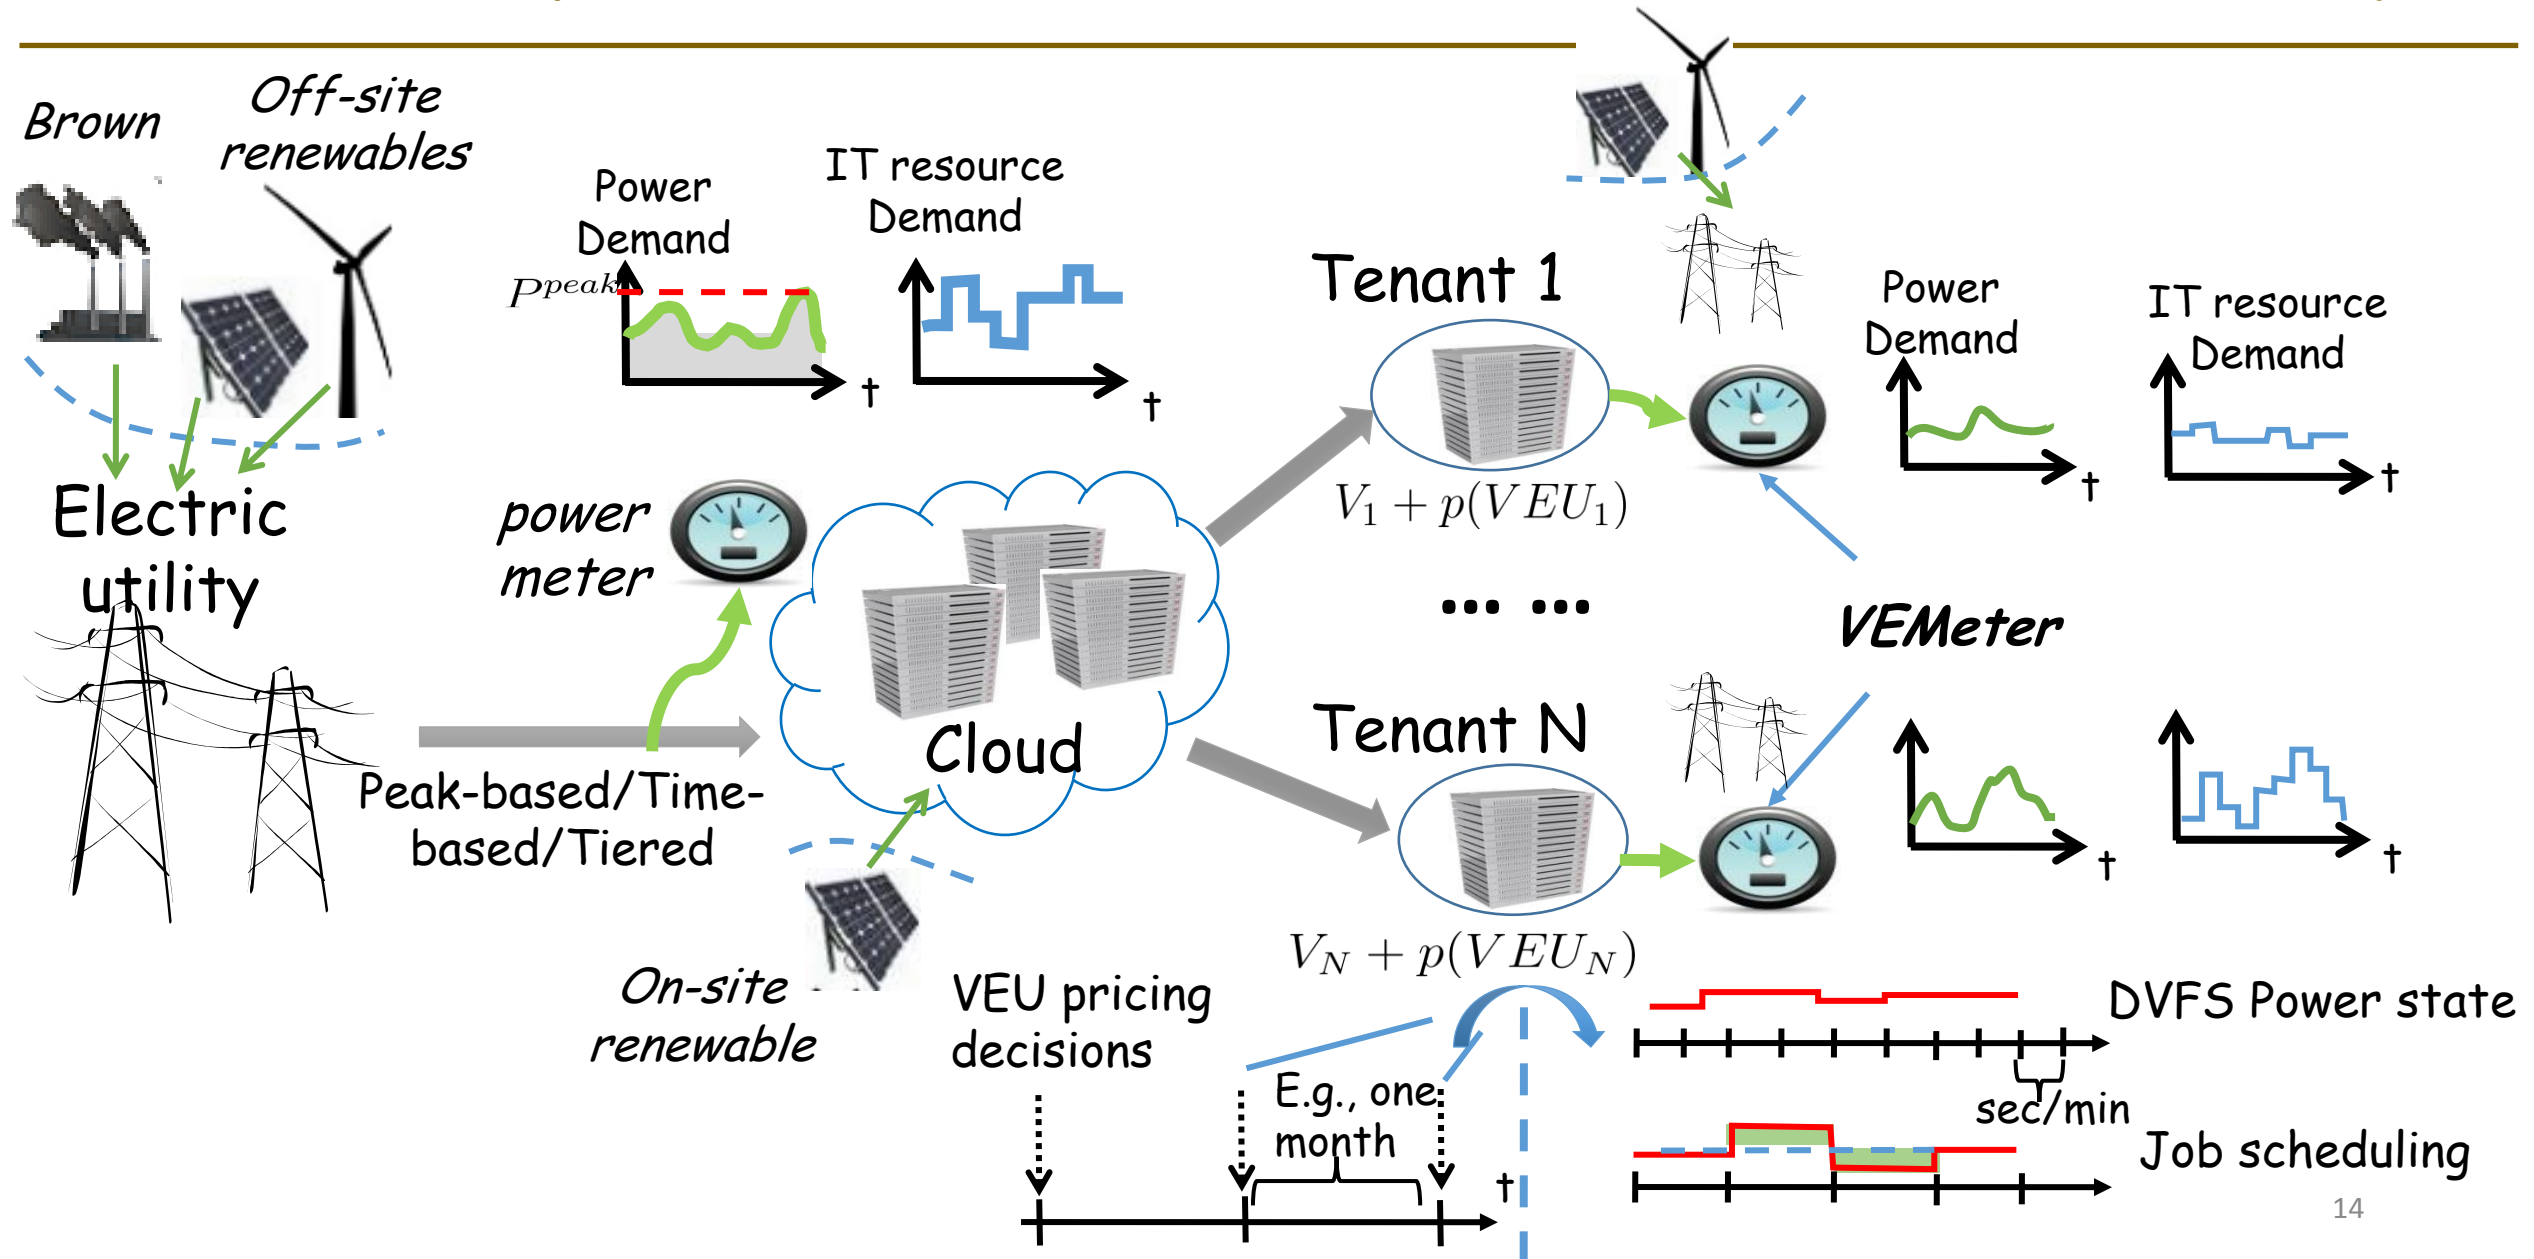
Pass on utility pricing structure to tenants such that tenants are charged for their contribution to Op-ex, and these prices reflect the (idiosyncratic) “value” the tenants derive from power

- Arguably, a tenant is the best judge of what “value” it derives from power and its demand response
 - Similar in spirit to e2e arguments, exo-kernels, ...

Our Proposal: Virtualize the Electric Utility

- Focus on “big”/long-lasting tenants
- Cloud exposes both IT resource and VEU pricing interfaces

Our Proposal: Virtualize the Electric Utility



Design: how to design and/or negotiate prices?

Control: How should cloud and tenants operate?



Design Effective VEU pricing

- Desirable features of VEU pricing design
 - Certain "fairness"
 - E.g., a tenant with higher demand variance might be charged at a higher price than another tenant w/ same total demand but lower variance
 - Cost-effective
 - Volume discount:
 - Price per unit demand decreases in consumption
 - Incentize tenants to participate into DR
 - Revenue neutral:
 - Revenue matches energy cost
- Not all are necessary
- Some might be conflicting

A Preliminary Game-based VEU Pricing

- Utility pricing: peak-based α, β
- Cloud pricing: time-varying
- Consider a *given interval of time*

κ_i Mean of tenant i 's energy consumption $\underline{\kappa} = (\kappa_1, \kappa_2, \dots, \kappa_N)$

s_i Standard deviation of tenant i 's energy consumption

$S^2 = \sum_i s_i^2$ Variance of aggregate energy consumption of data center

(If demands are uncorrelated)

Approx. aggregate peak by mean + 2*standard deviation

By revenue neutral: $\sum_i p_i(\underline{\kappa}) \kappa_i = \alpha \sum_i \kappa_i + \beta (\sum_i \kappa_i + 2S)$

Per-unit price
for tenant i

Total energy




Peak approximation

Total power cost of the data center

A Preliminary Game-based VEU Pricing

Design VEU price for tenant i :

$$p_i(\underline{\kappa}) = \alpha + \beta + \beta \frac{2S}{\kappa_i} \cdot \frac{g(\kappa_i, s_i)}{\sum_j g(\kappa_j, s_j)} \quad \text{where } g(\kappa, s) = \left(\frac{\kappa}{\kappa_{max}} \right) \gamma^{(s-s_{max})}$$

- Certain "fairness"
 - Price increases in demand variance 
- Volume discount:
 - Price decreases in average consumption 
- Revenue neutral:
 - Revenue matches energy cost 
- Cost-effective:
 - To be verified through numerical simulation

A Preliminary Game-based VEU Pricing

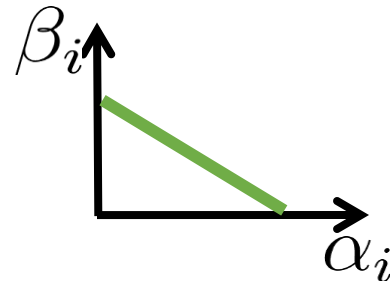
In this game, tenant i will optimize net profit by load shedding:

$$\max_{\kappa_i} v_i(\underline{\kappa}) = \underbrace{\mu_i(\kappa_i)}_{\text{Revenue}} - \underbrace{p_i(\underline{\kappa})\kappa_i}_{\text{Cost}} \quad \text{where } \mu_i \text{ is increasing, concave and bounded}$$

Theorem 1. If for $s_1 = s_2$ there is a symmetric Nash Equilibrium (NE) $\kappa_1^* = \kappa_2^*$ at which $\kappa(\mu'(\kappa) - \alpha - \beta)$ is decreasing in κ^* , then by perturbing s_1 so that $s_1 > s_2$ and $s_1 \approx s_2$, this symmetric NE changes such that $\kappa_1^* < \kappa_2^*$.

VEU Pricing Design: Challenges & Ideas

- Simplification/limitation of the pricing design
 - Decision-making during a single time interval.
 - Possible extension: Demand modulation over multiple time intervals.
- How about VEU prices resembling the actual electric utility?
 - The data center offers choices of energy charge and peak charge for tenants to choose from.
 - Tenant w/ lower variance might be willing to pay higher charge for peak if that allowed it to choose lower energy charge.

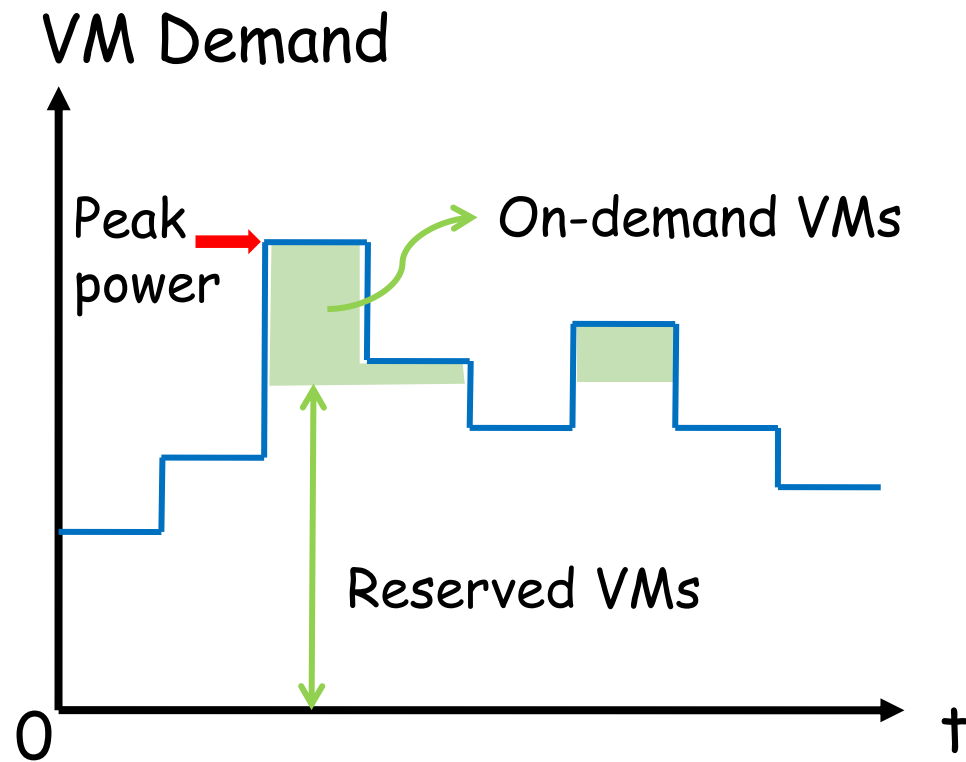


VEU Design: System S/W & Tools

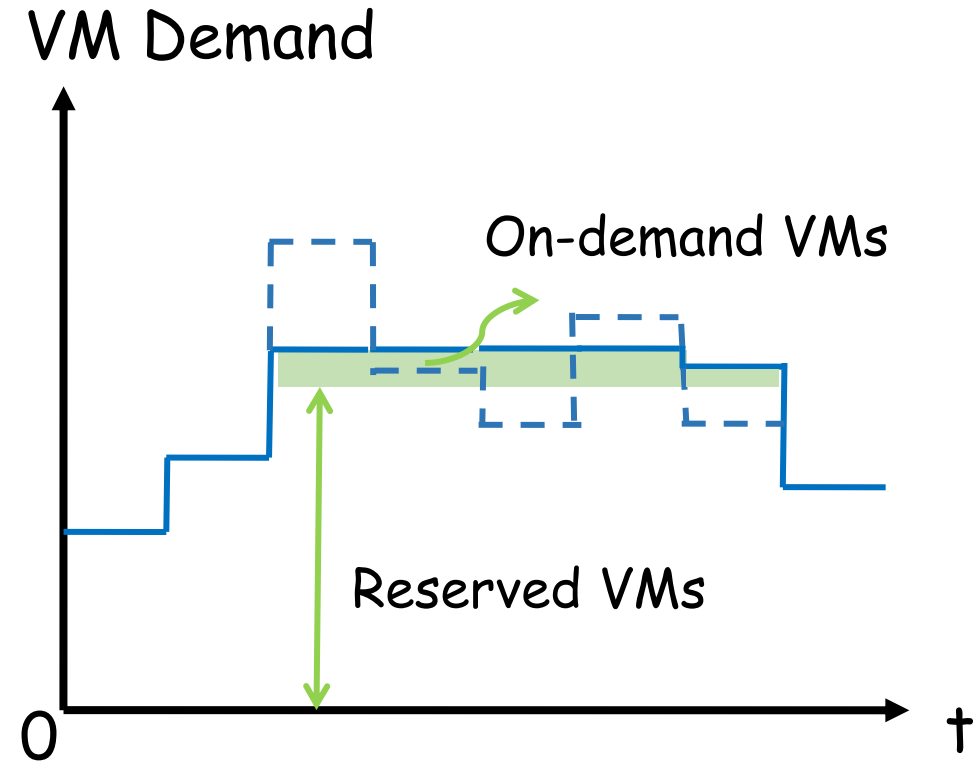
- VEMeter:
 - Infer tenant's local power consumption and estimate individual VEU
 - Challenge:
 - Attributing energy consumed by shard components
 - Tools in literature
 - Energy accounting
 - vPath

Tenant Operation w/ VEU

- Novel resource procurement problems:



(a) Decision making w/o VEU



(b) Decision making w/ VEU

Implementation Considerations

- Additional complexity for tenant's operation
 - Solving more complex stochastic optimization problems
 - Demand modulation via abstract knobs: Dropping and/or delaying.
 - Predicting/converting IT resource procurement into power consumption
 - Leveraging existing work: vPower, energy container, palloc, ...

Related Work

- Reducing energy-related costs
 - Reducing raw energy consumption
 - Improving PUE, cooling system, IT capacity modulation/shutdown (e.g., CPU, mem, disk, entire server, etc) ...

VEU serves as a mech. for propagating the energy/cost benefits offered by these techniques to tenants fairly.

- Alternative approaches that “virtualize” power
 - Treating energy as a first-class resource

We propose to virtualize NOT just power but the electric utility itself.

Conclusions

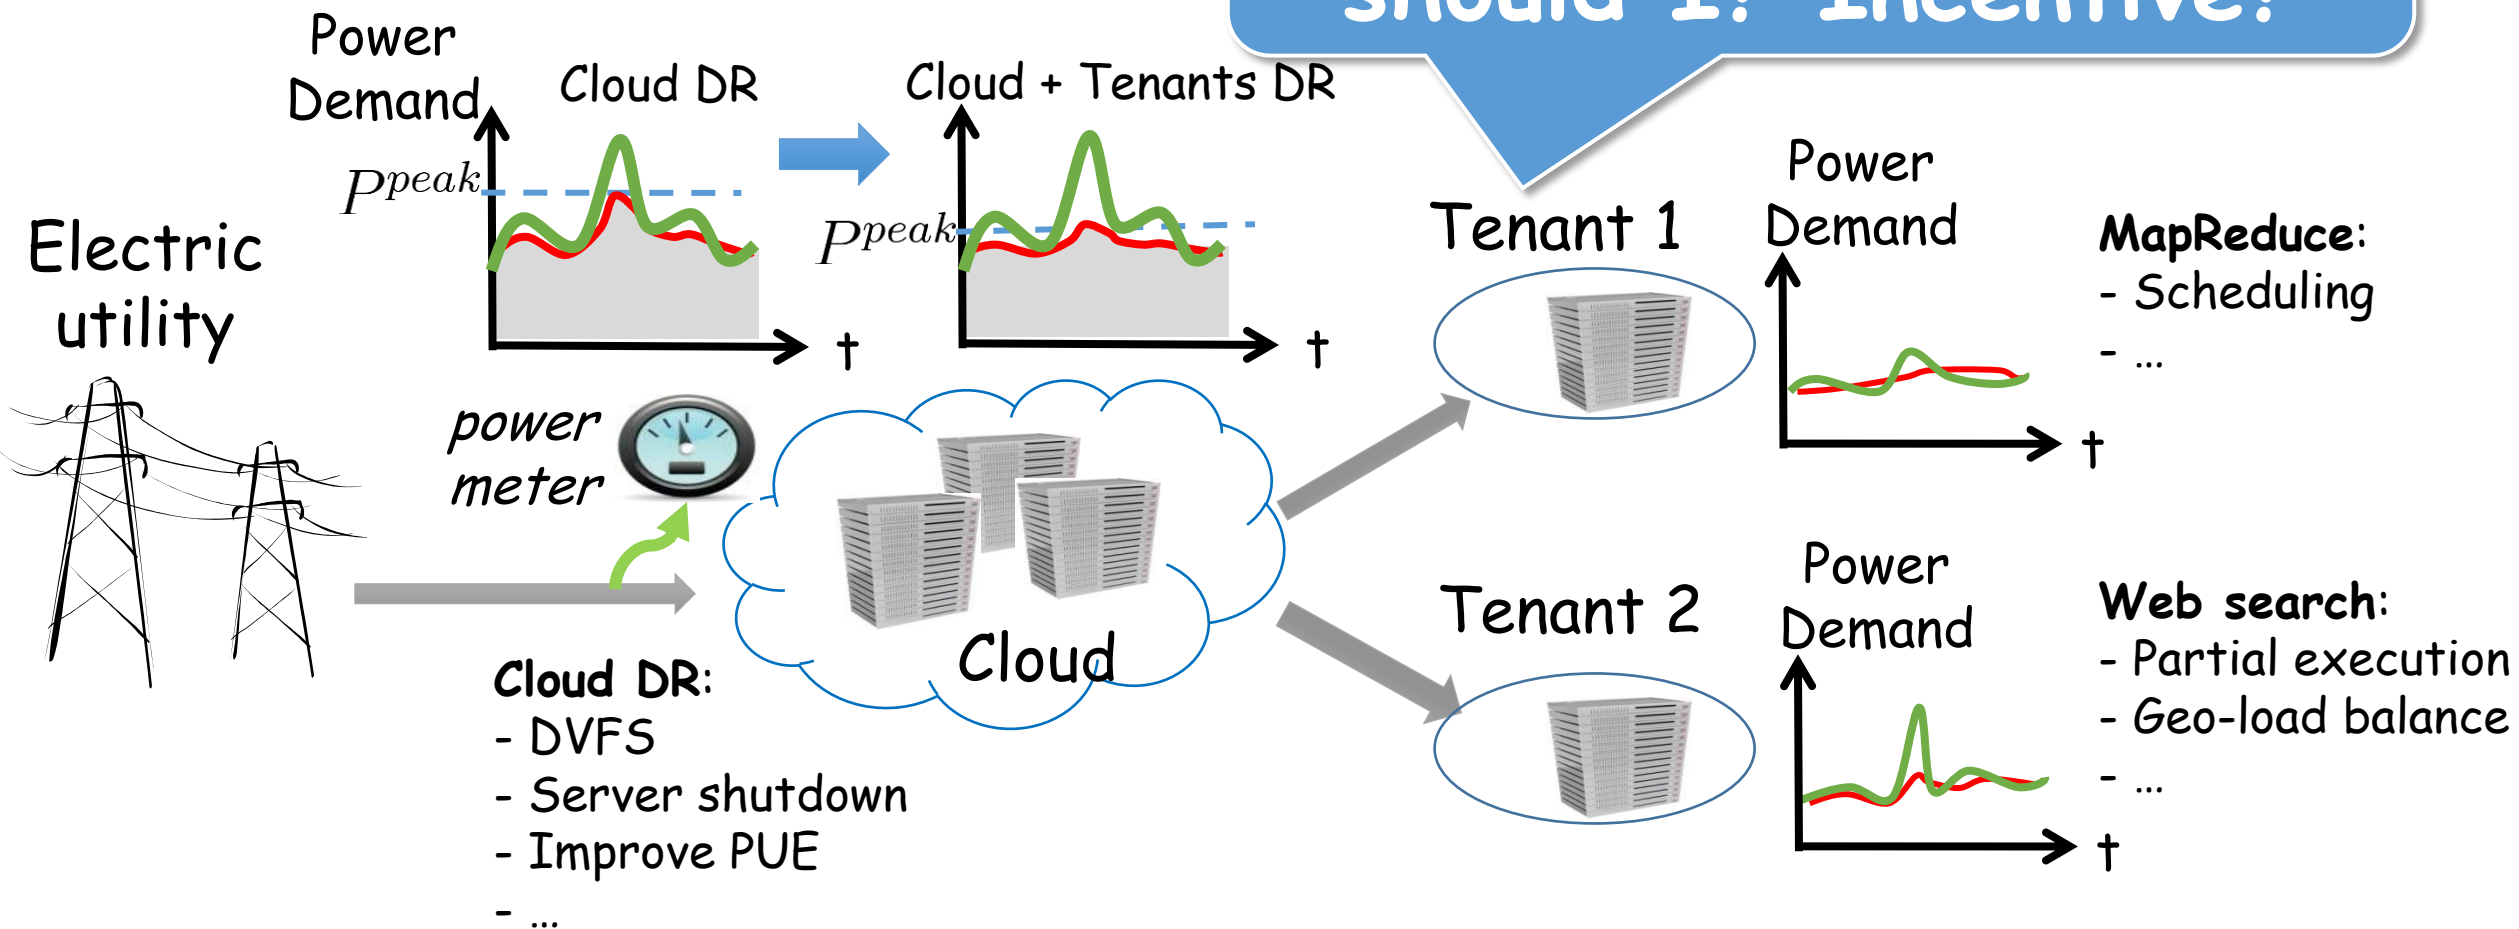
- Gap between electric utility pricing and cloud pricing can results in shortcomings
 - Fairness
 - Cost efficacy
- Virtualizing the electric utility might offer a remedy to these problems

An Example of Loss of Cost-Efficacy

- Cost-efficacy (From cloud's

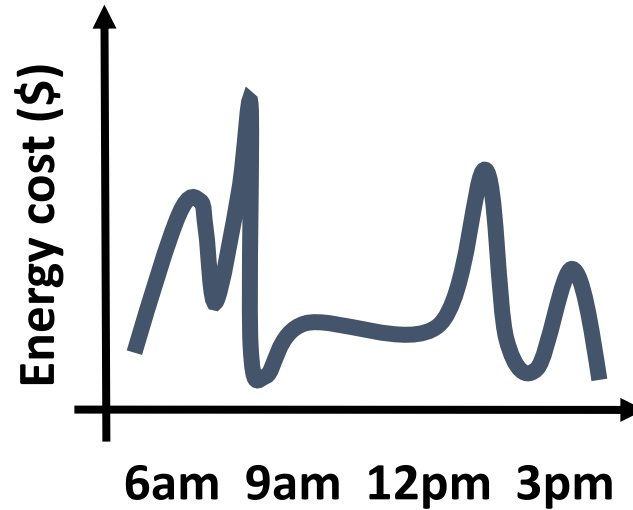
Good for cloud! But why should I? Incentive?

DR

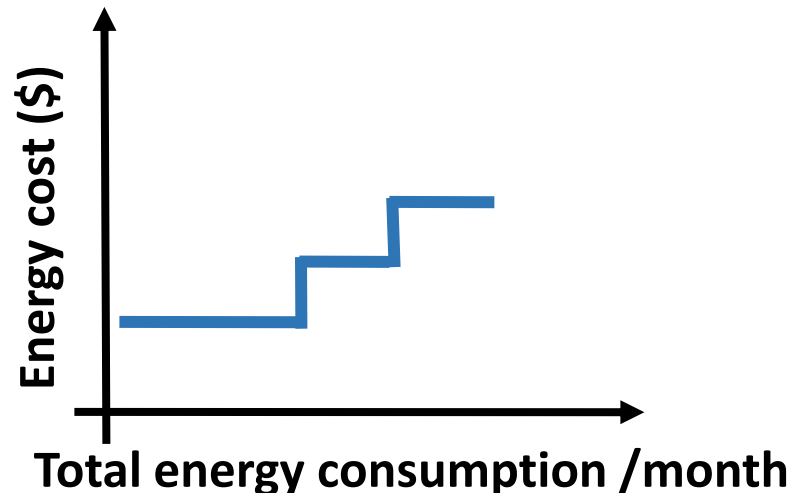


Electric Utility Pricing (contd.)

- Real-time pricing with high “coincident” peak charges



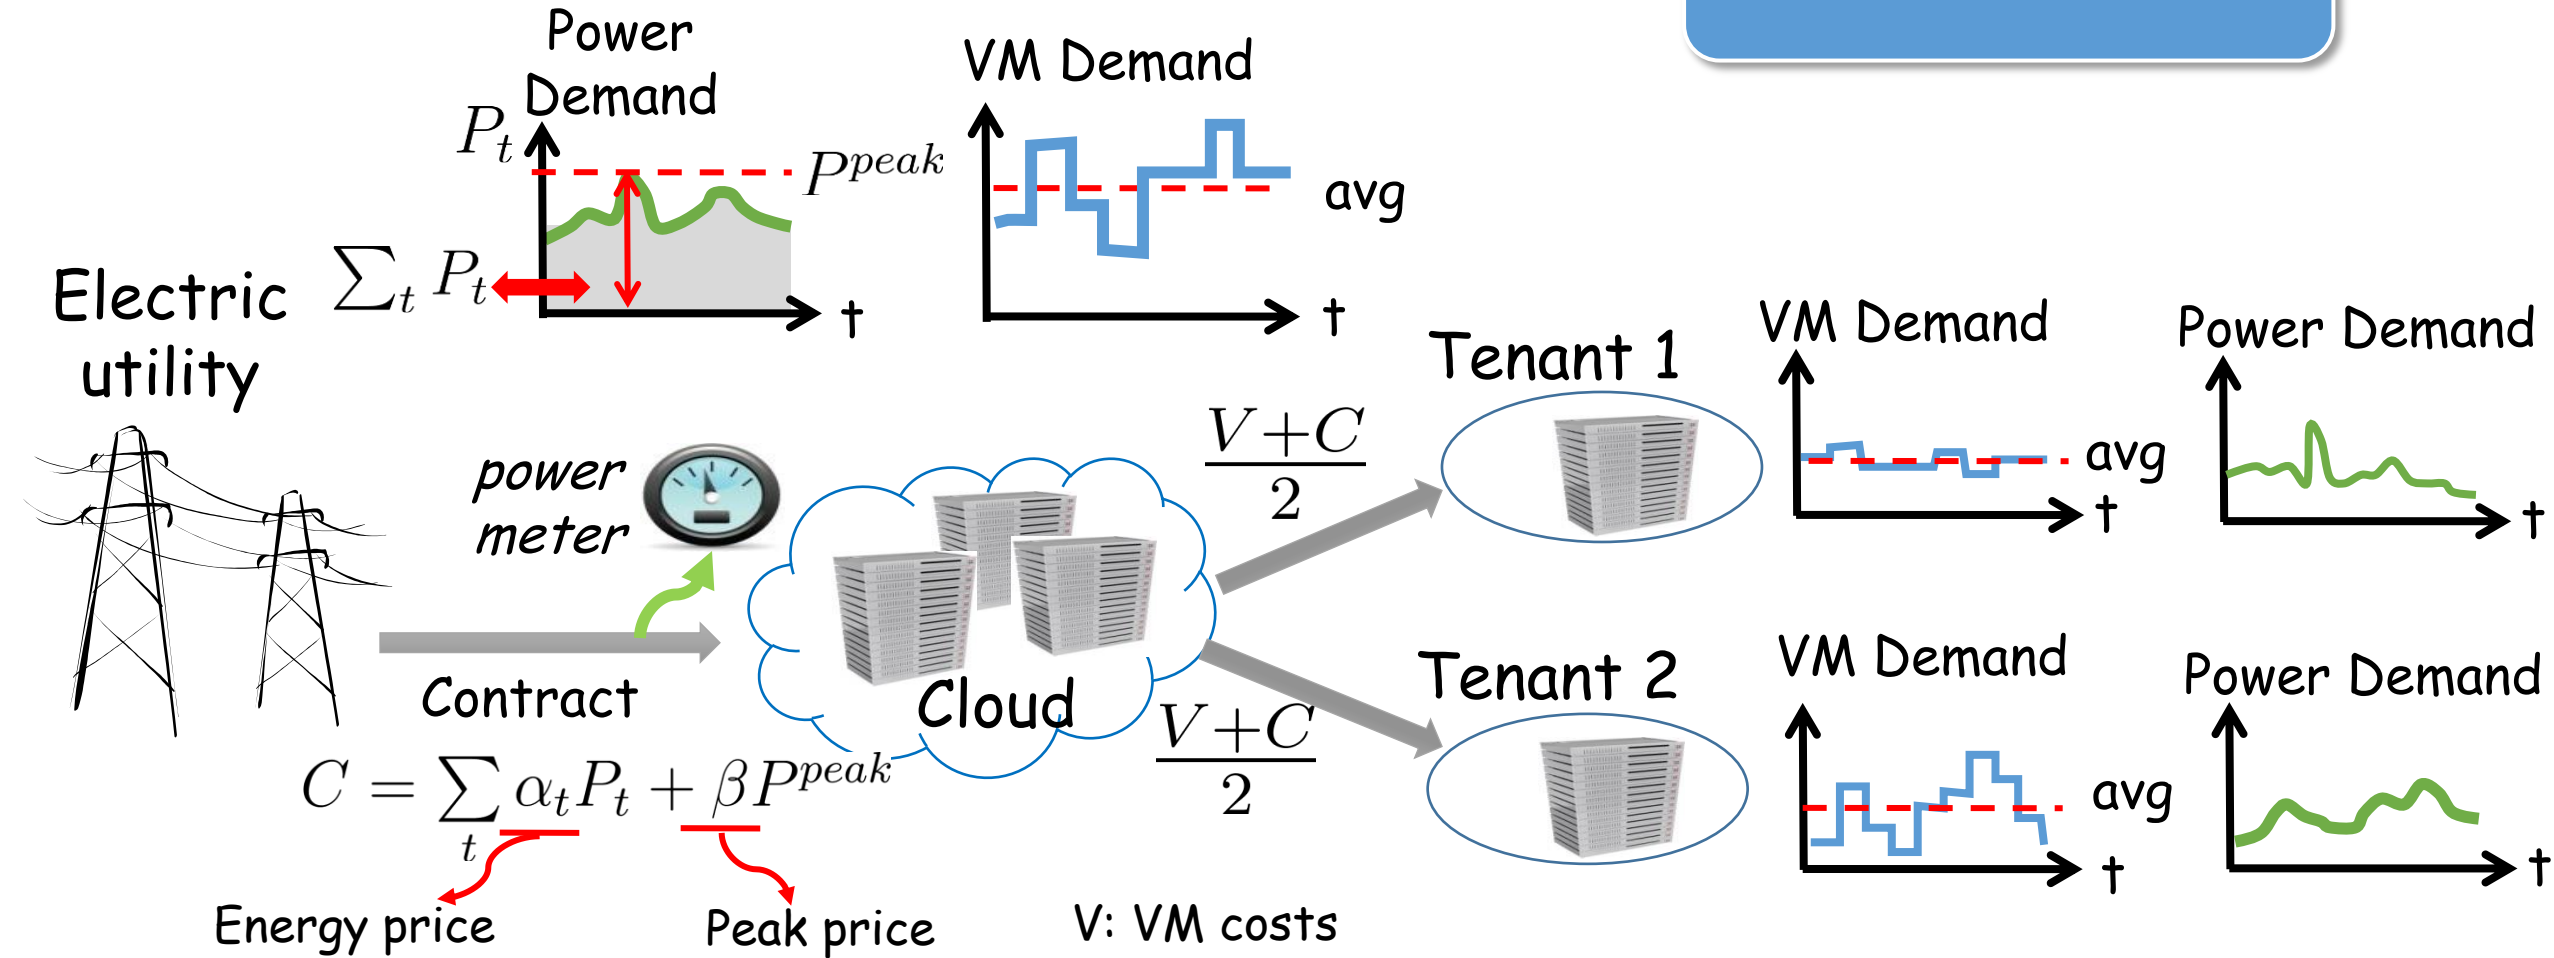
- Tiered



A Thought Experiment

- More complex VM v.s. power relation

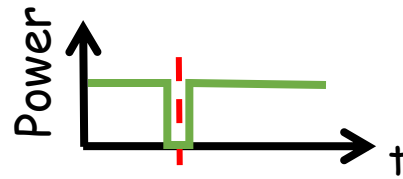
Unfair



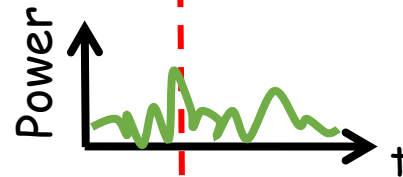
Strawman #2

- Base each tenant's peak charge on its contribution to the aggregate peak
- What if tenant 1 carries out demand response as shown?

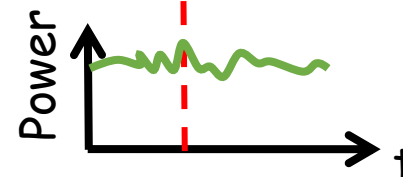
Tenant 1



Tenant 2



Cloud



Cloud's revenue
affected by tenant
demand response!