Exploiting a Thermal Side Channel for Power Attacks in Multi-Tenant Data Centers

Mohammad A. Islam, Shaolei Ren, and Adam Wierman

Acknowledgement: NSF under grants CNS-1551661, CNS-1565474, and ECCS-1610471, AitF-1637598, CNS-1518941, and CNS-1319820.

Multi-tenant data centers

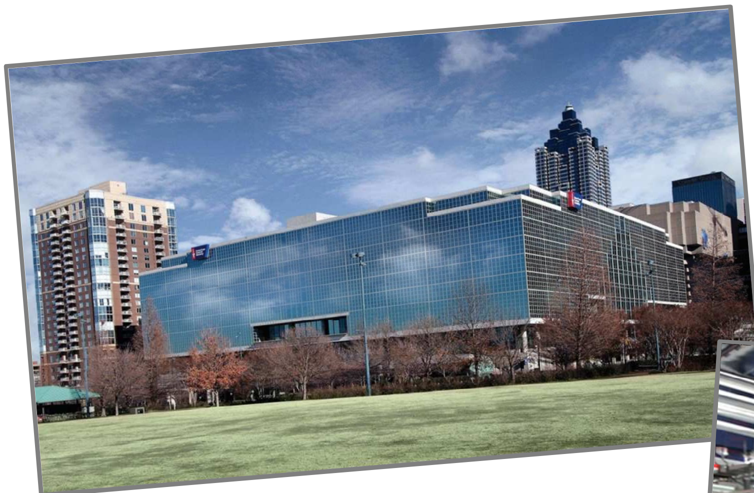

- Mission-critical infrastructure
- Backbone of digital economy
- 50% growth by 2020
- ……

Multi-tenant data centers

- Mission-critical infrastructure
- Backbone of digital economy
- 50% growth by 2020
- ……

A multi-tenant data center is a shared facility that houses multiple tenants, each managing its own servers…

Multi-tenant data centers are everywhere...

2,000+ in U.S.

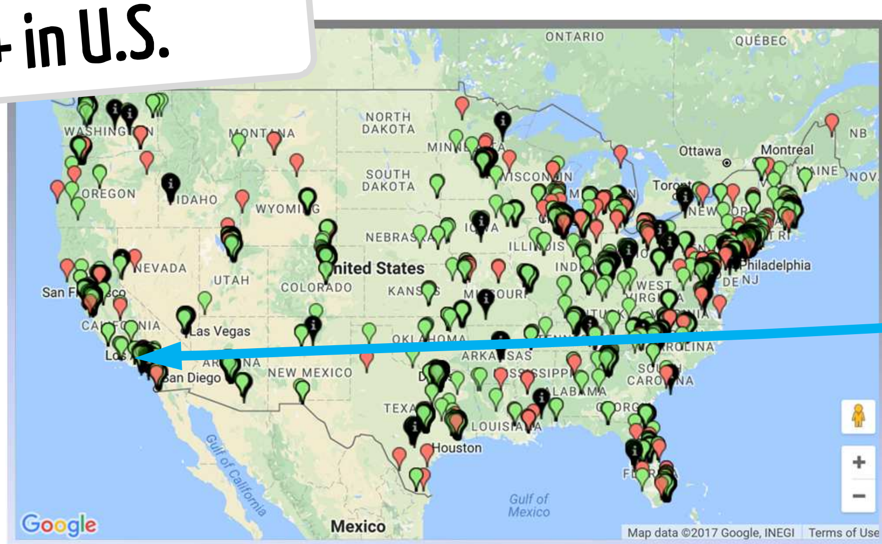

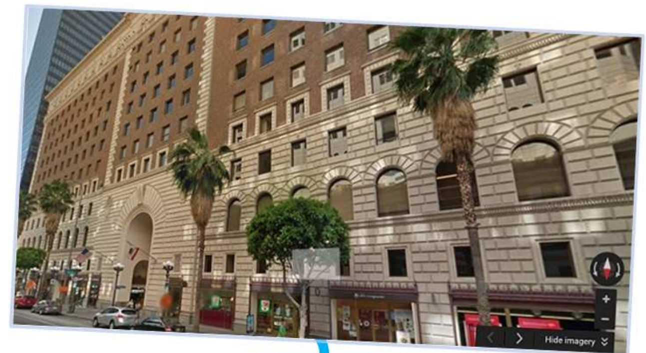

Multi-tenant data centers are everywhere...

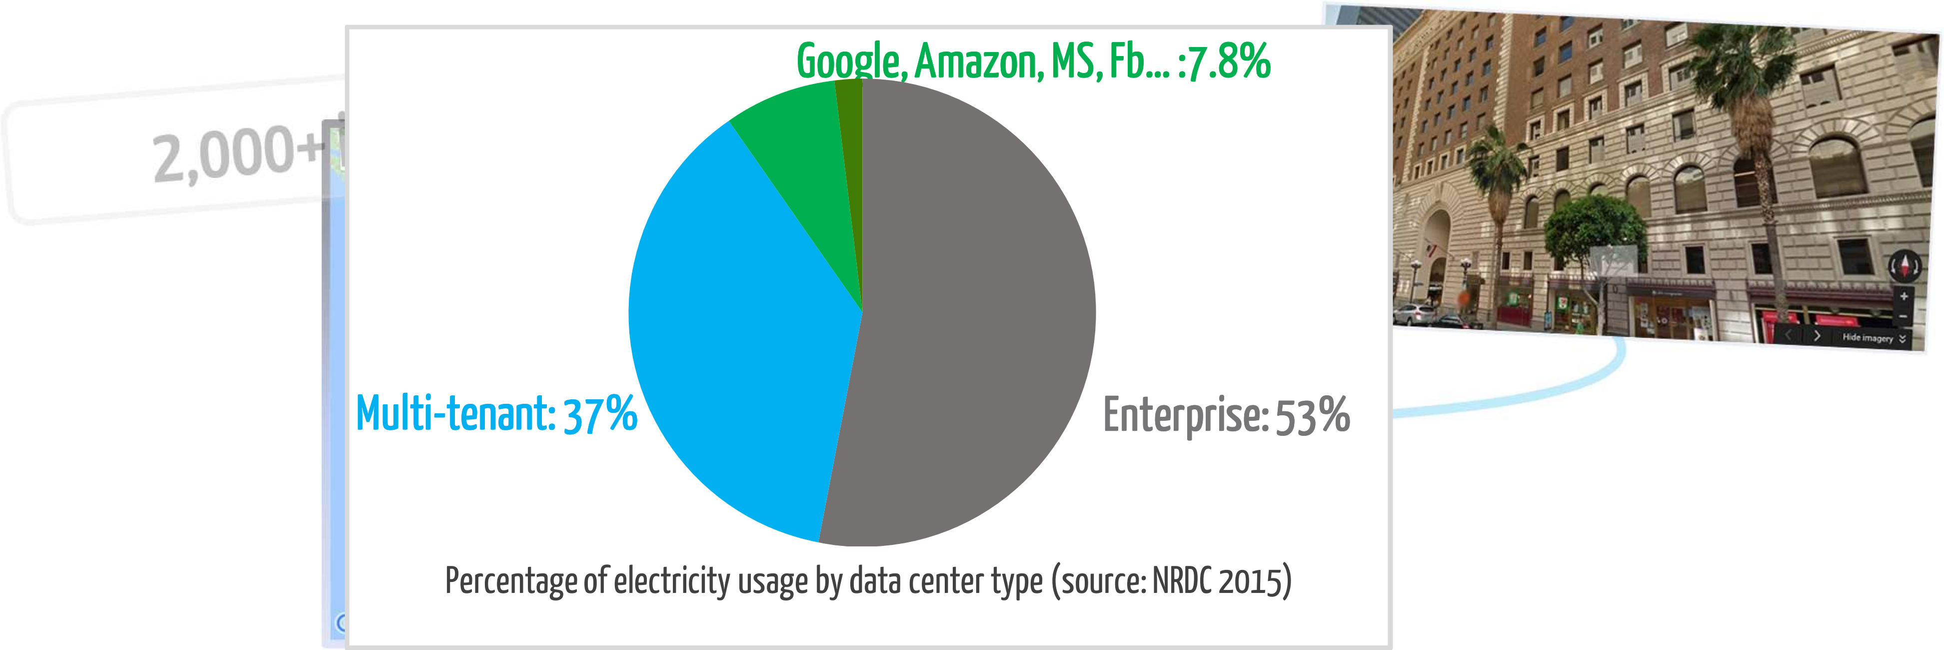

Using multi-tenant data centers for …

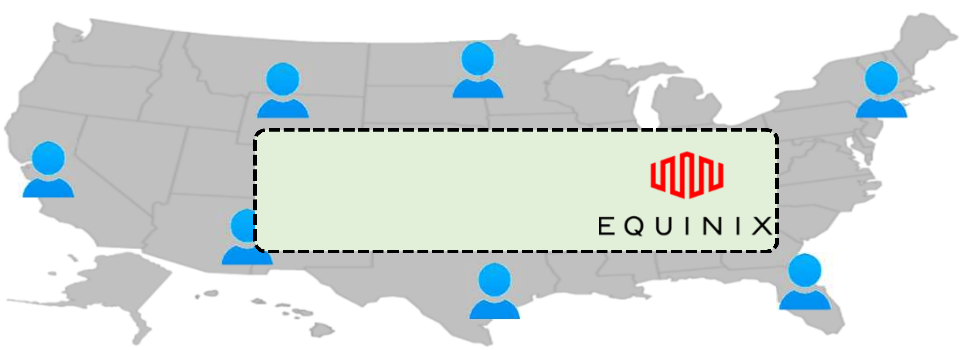

Using multi-tenant data centers for ...

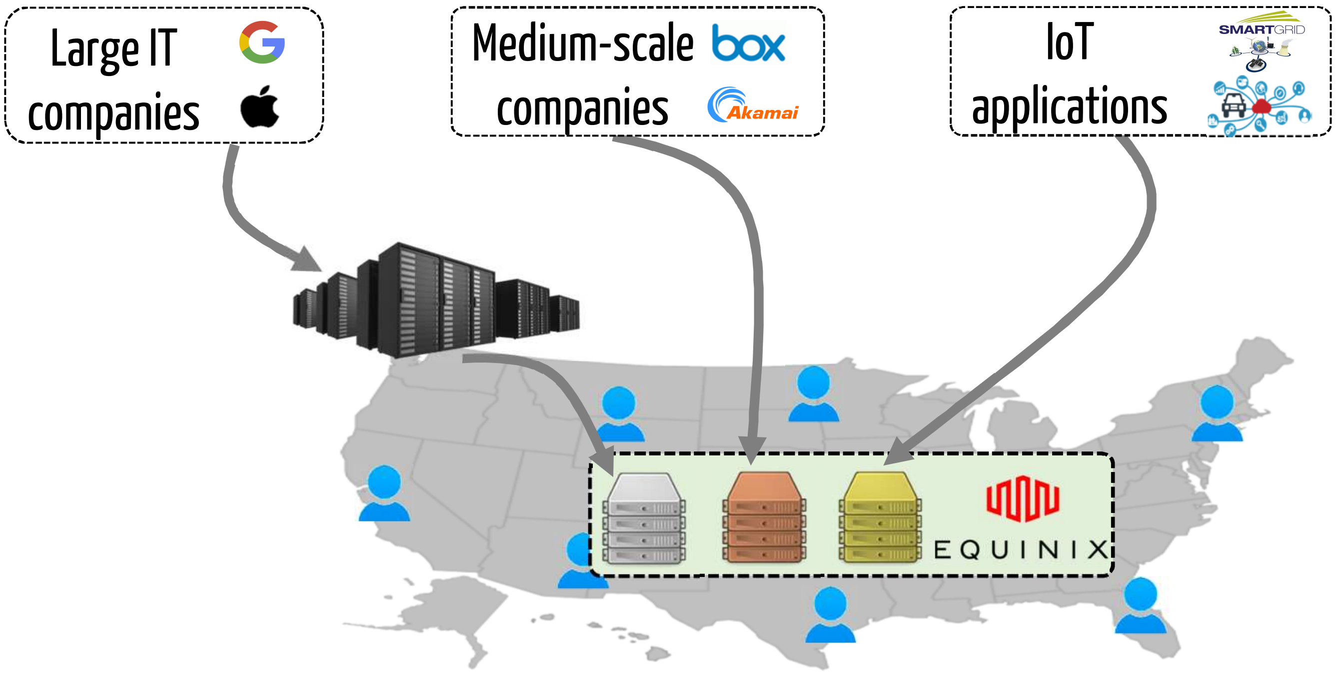

- Lower latency
- Lower CapEx & OpEx
- Better privacy
- Higher scalability
- ...

Apple houses 25% of its servers in multi-tenant data centers...

Securing multi-tenant data centers is extremely important!

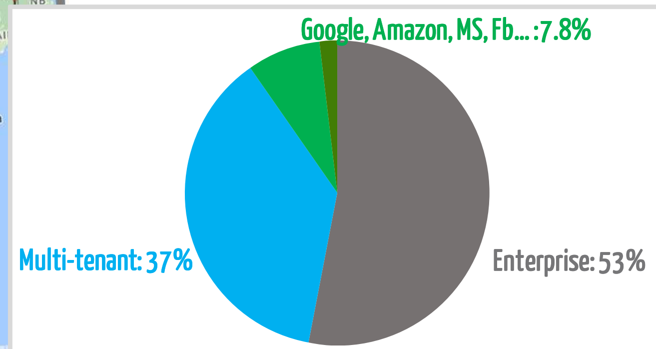

Percentage of electricity usage by data center type (source: NRDC 2015)

A cyber-physical view…

Utility

A cyber-physical view...

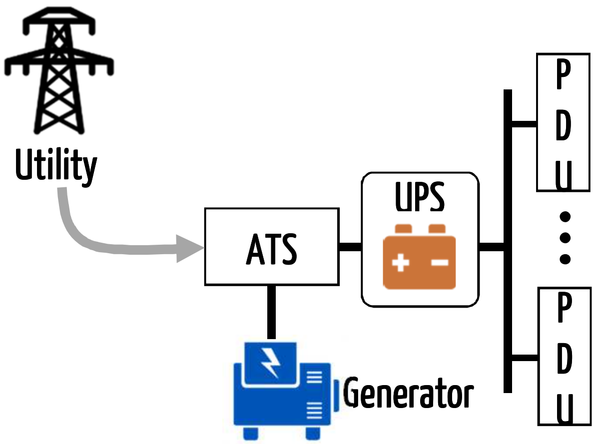

Managed by operator

A cyber-physical view...

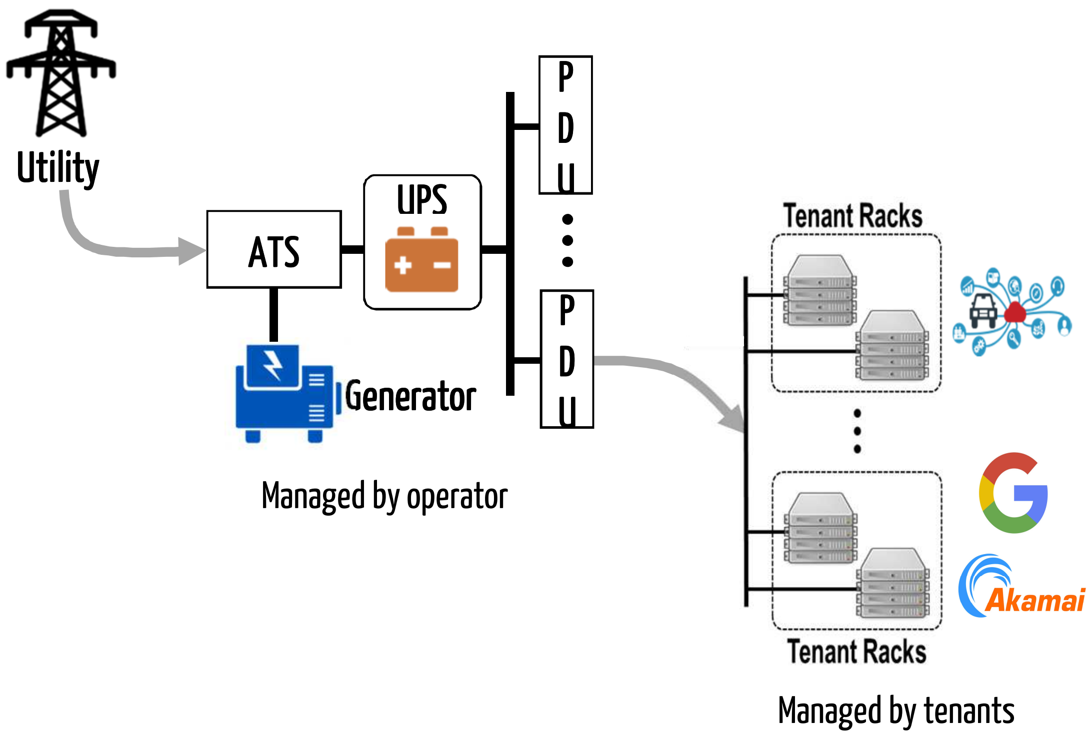

A cyber-physical view…

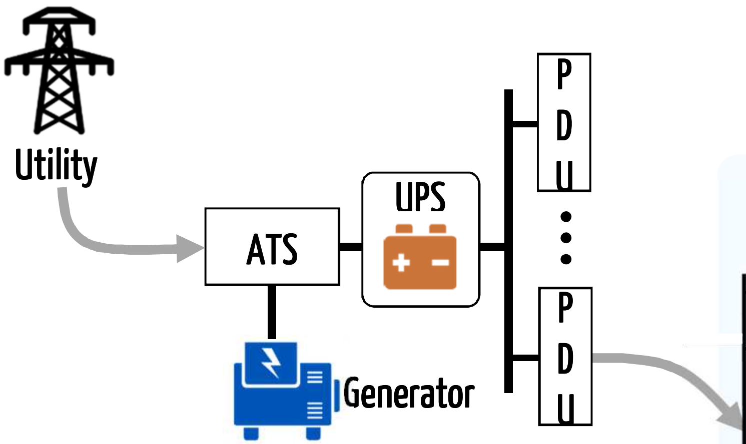

Managed by operator

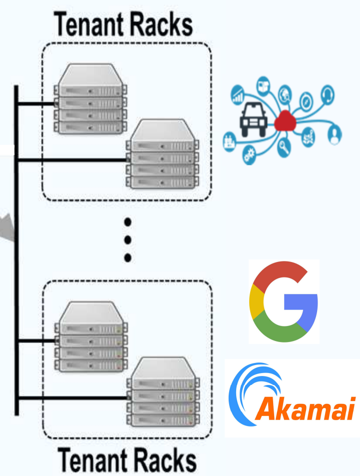

Managed by tenants

Securing the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

  [Mirkovic, Sigcomm'04] [Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]…

A cyber-physical view...

How about physical security?

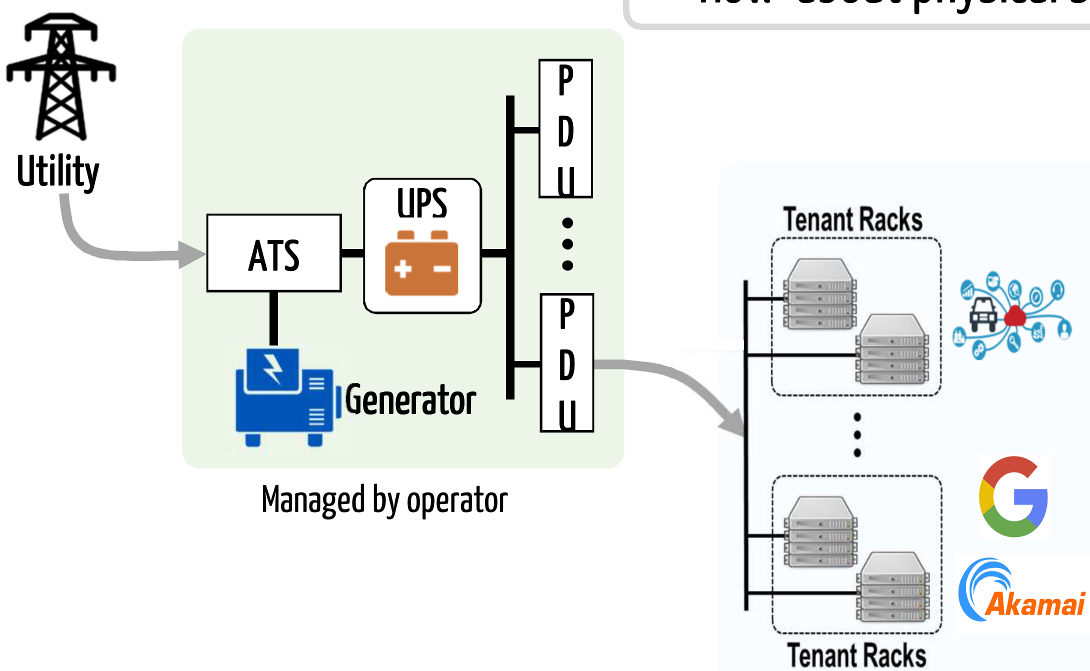

Managed by operator

Managed by tenants

Securing the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

[Mirkovic, Sigcomm'04] [Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]...

A cyber-physical view…

How about physical security?

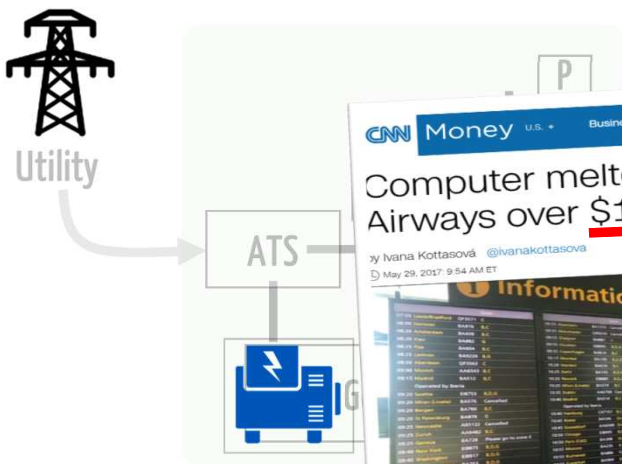

CNN Money

Computer meltdown may cost British Airways over $100 million

by Ivana Kottasová @ivanakottasova

May 29, 2017: 9:54 AM ET

BA MELTDOWN COULD COST OVER $100 MN

British Airways' costly computer system crash

ecuring the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

[Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]…

PCWorld

NEWS

Power surge at British Airways data center causes flight chaos

A cyber-physical view...

How about physical security?

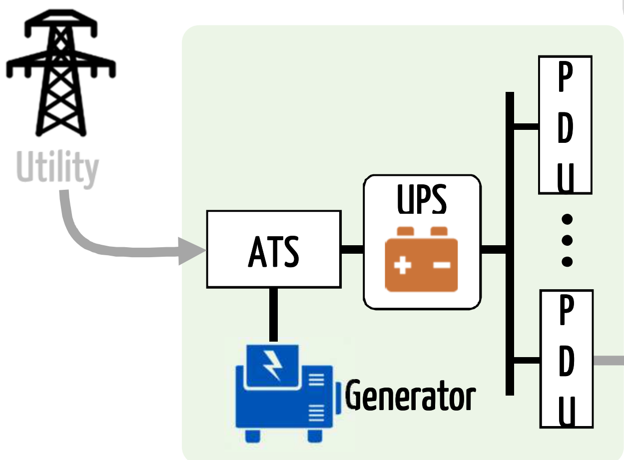

99.995+% availability!

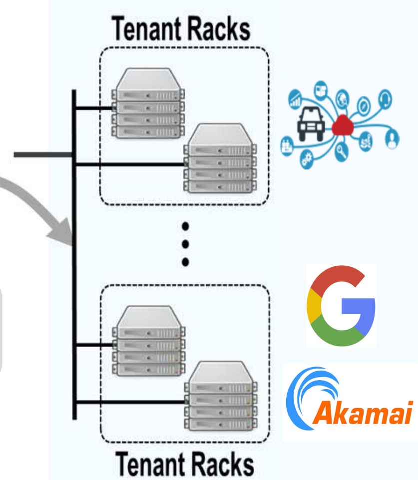

Securing the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

[Mirkovic, Sigcomm'04] [Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]...

We revisit the conventional wisdom and find…

We revisit the conventional wisdom and find...

Multi-tenant data centers are highly vulnerable to well-timed **power attacks!**

- Why are multi-tenant data centers vulnerable to power attacks?
- What is the potential impact of power attacks?
- How could an attacker mount a power attack?
- How to defend a data center against power attacks?

- Why are multi-tenant data centers vulnerable to power attacks?
- What is the potential impact of power attacks?
- How could an attacker mount a power attack?
- How to defend a data center against power attacks?

When building a data center…

US$ 10-25 per watt of data center capacity
CapEx is 60+% of the total cost of ownership

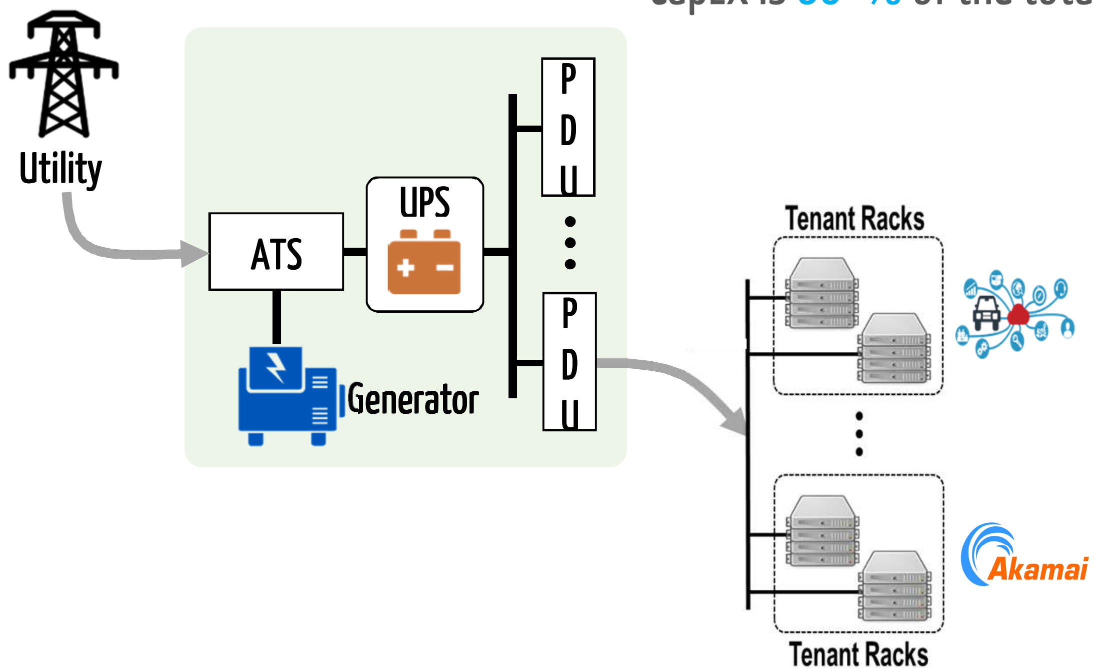

When building a data center…

US$ 10-25 per watt of data center capacity
CapEx is 60+% of the total cost of ownership

Oversubscribing the data center capacity is common!

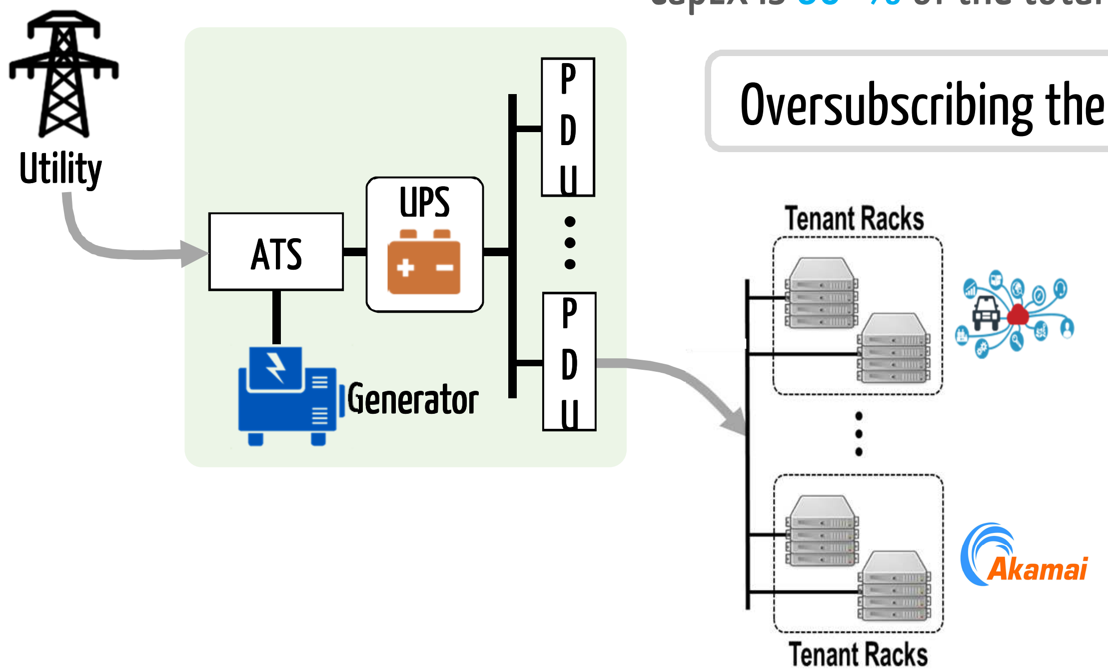

When building a data center...

US$ 10-25 per watt of data center capacity
CapEx is 60+% of the total cost of ownership

Oversubscribing the data center capacity is common!

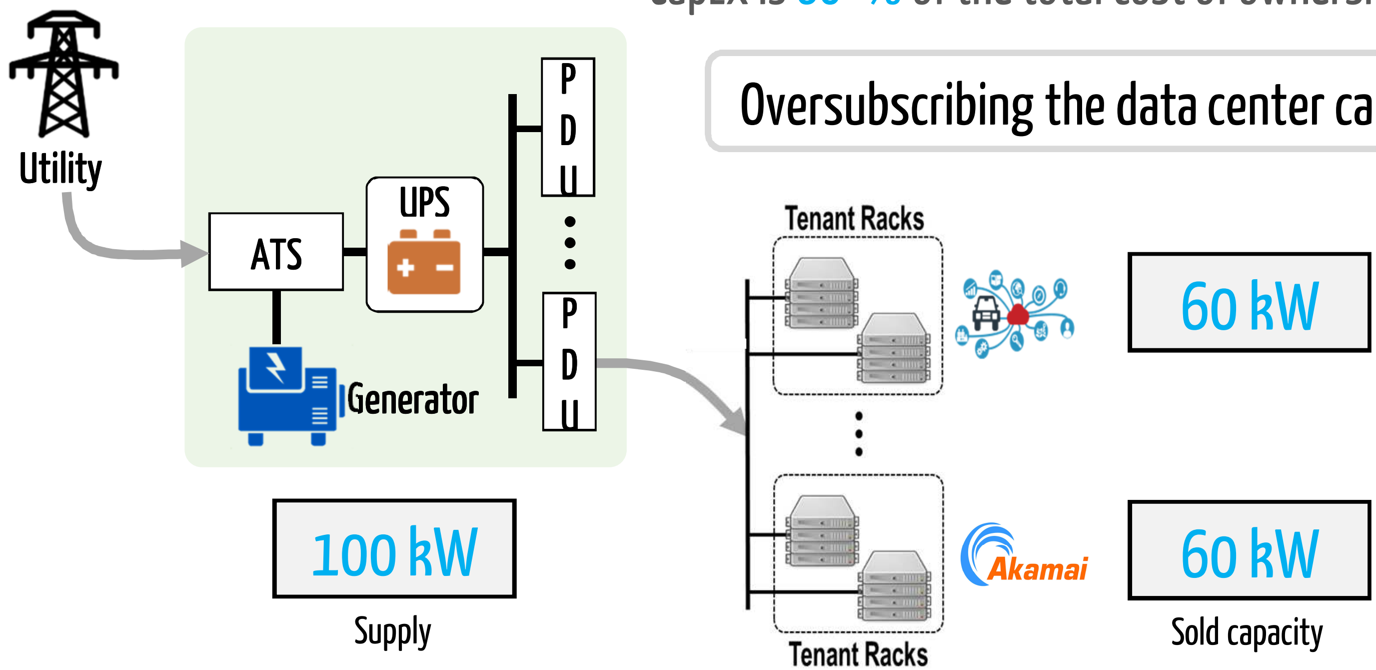

When building a data center...

US$ 10-25 per watt of data center capacity

CapEx is 60+% of the total cost of ownership

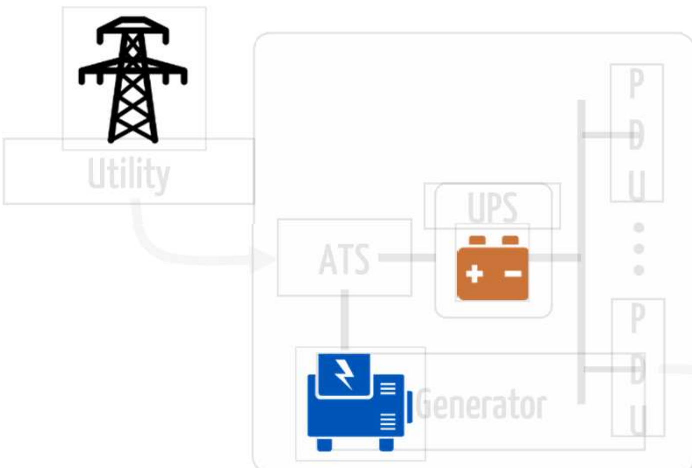

Oversubscribing the data center capacity is common!

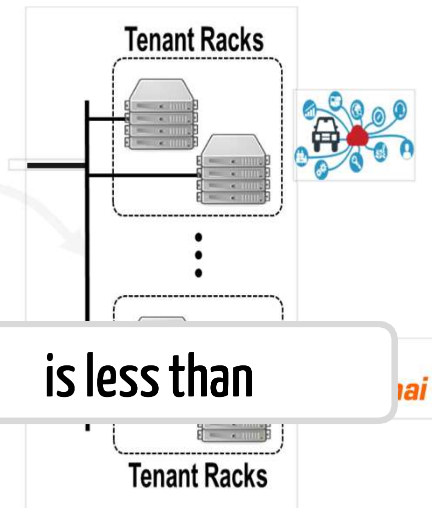

100 kW

Supply

is less than

60 kW

+

60 kW

Sold capacity

Rationale & safeguards

- Multiplex tenants' power demand
- Limit on tenants' power usage
- Infrastructure robustness and redundancy

Rationale & safeguards

- Multiplex tenants' power demand
  - Simultaneous peaks are very rare!
- Limit on tenants' power usage
  - Normal usage limited to 80% of tenant's subscribed capacity
  - Only occasional peak usage is allowed
- Infrastructure robustness and redundancy
  - Transient spikes are harmless

Rationale & safeguards

- Multiplex tenants' power demand
  - Simultaneous peaks are very rare!
- Limit on tenants' power usage
  - Normal usage limited to 80% of tenant's subscribed capacity
  - Only occasional peak usage is allowed
- Infrastructure robustness and redundancy
  - Transient spikes are harmless

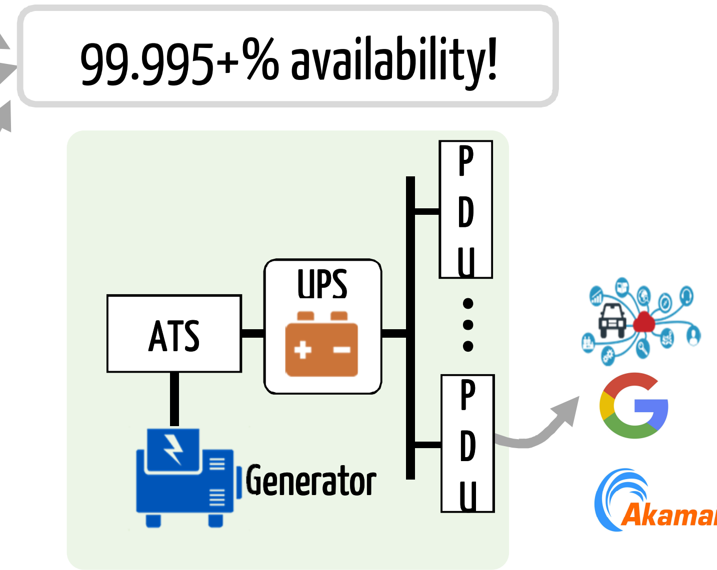

Rationale & safeguards

- Multiplex tenants' power demand
  - Simultaneous peaks are very rare!
- Limit on tenants' power usage
  - Normal usage limited to 80% of tenant's subscribed capacity
  - Only occasional peak usage is allowed
- Infrastructure robustness and redundancy
  - Transient spikes are harmless

99.995+% availability??

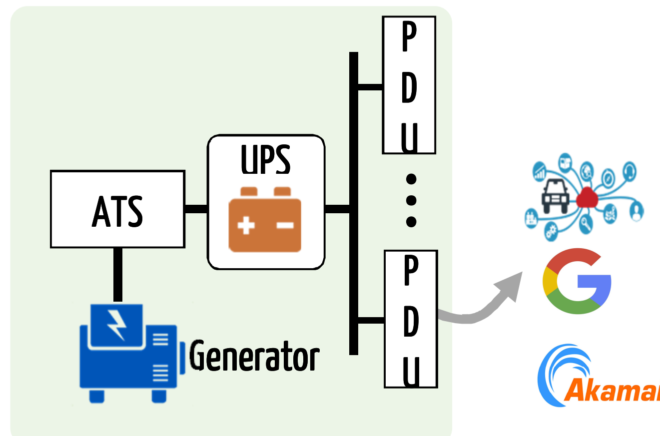

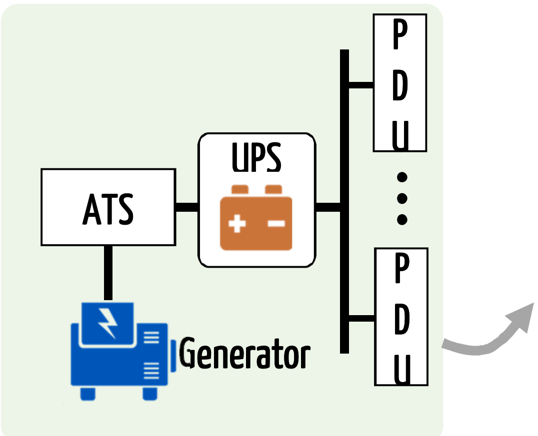
ATS
UPS
P
D
U
P
D
U
Generator

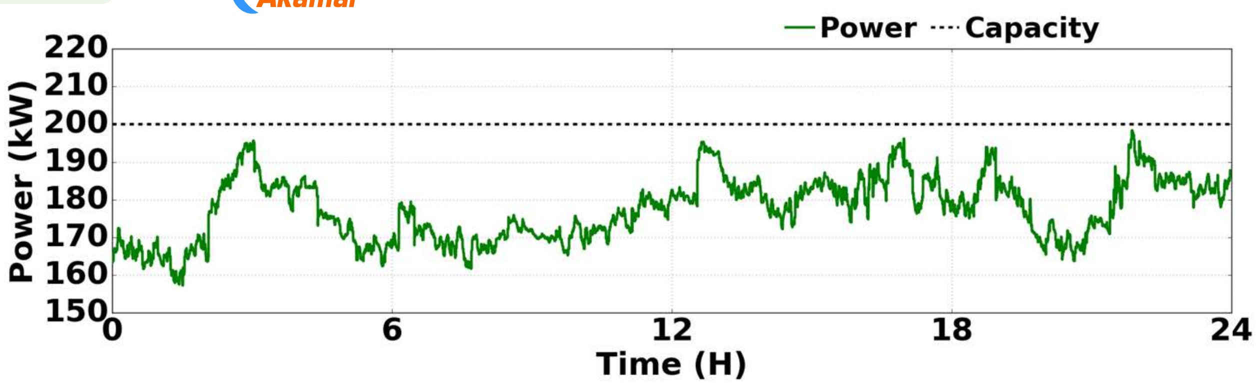
Power
Capacity
220
210
200
190
180
170
160
150
Power (kW)
0
6
12
18
24
Time (H)

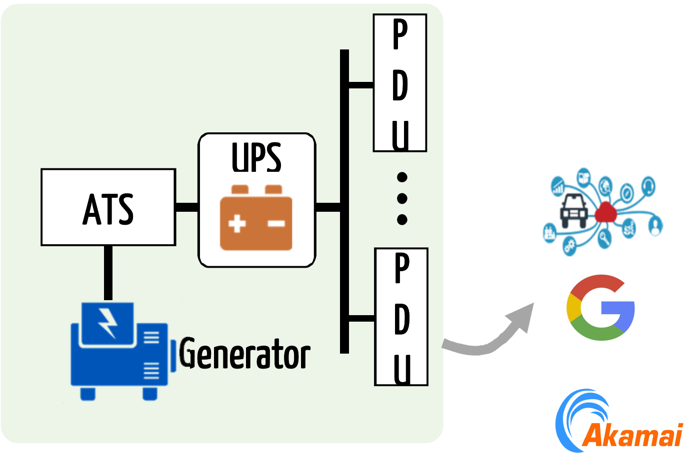
ATS
UPS
P
D
U
P
D
U
Generator
Akamai

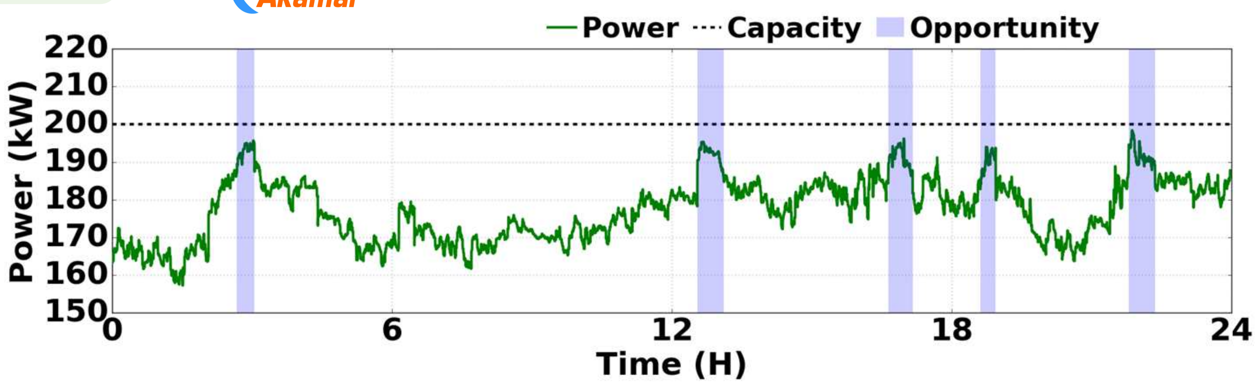
Power
Capacity
Opportunity
Power (kW)
Time (H)
220
210
200
190
180
170
160
150
0
6
12
18
24

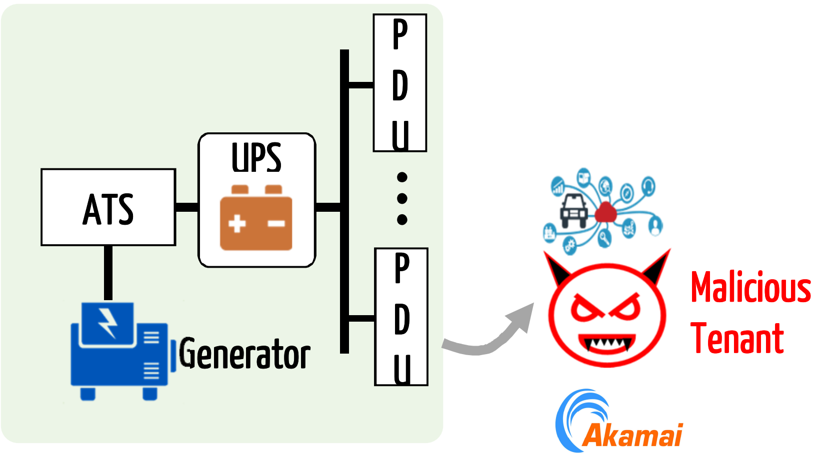

Frequent capacity overloads…

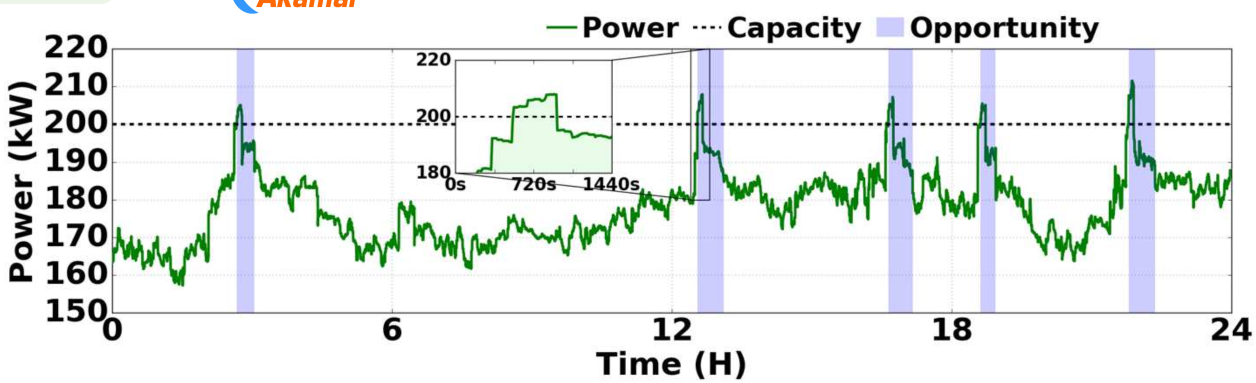

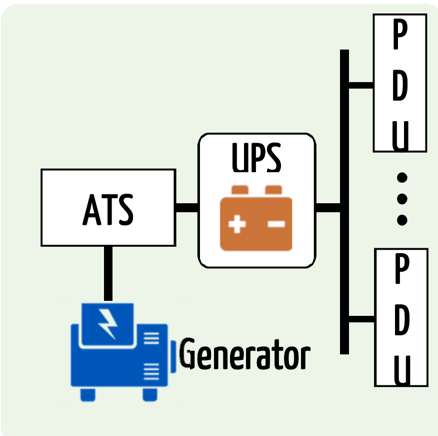

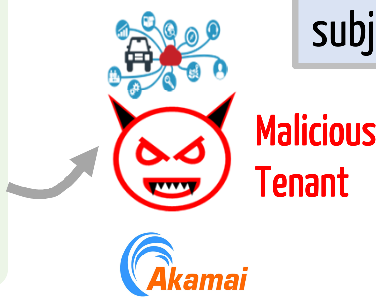

New threat model: Power attack

Well-timed power injection to overload the shared data center capacity, subject to all applicable usage constraints set by the operator

Frequent capacity overloads…

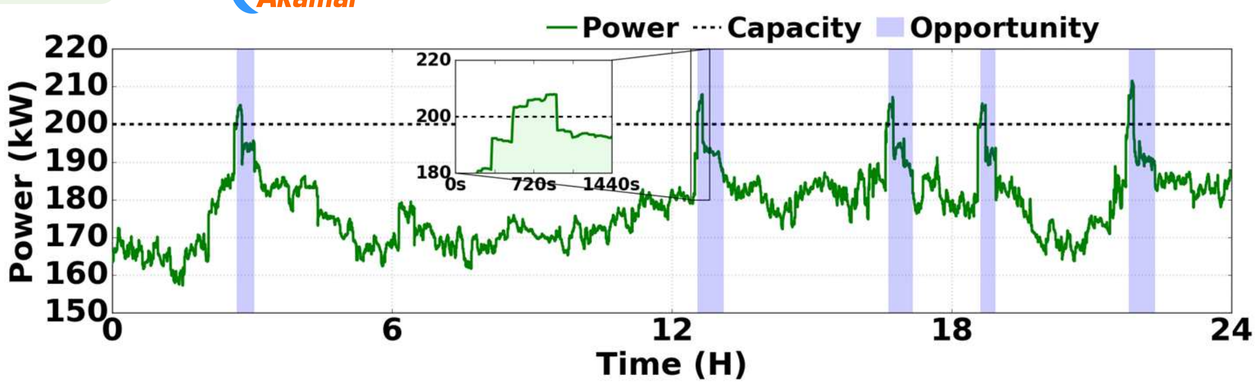

- Why are multi-tenant data centers vulnerable to power attacks?
    - Current safeguards are ineffective for well-timed power attacks

- How could an attacker mount a power attack?

- How to defend a data center against power attacks?

- Why are multi-tenant data centers vulnerable to power attacks?
  - Current safeguards are ineffective for well-timed power attacks
- **What is the potential impact of power attacks?**
- How could an attacker mount a power attack?
- How to defend a data center against power attacks?

Compromising data center availability…

- The outage risk is 280+ times higher during a capacity overload than otherwise

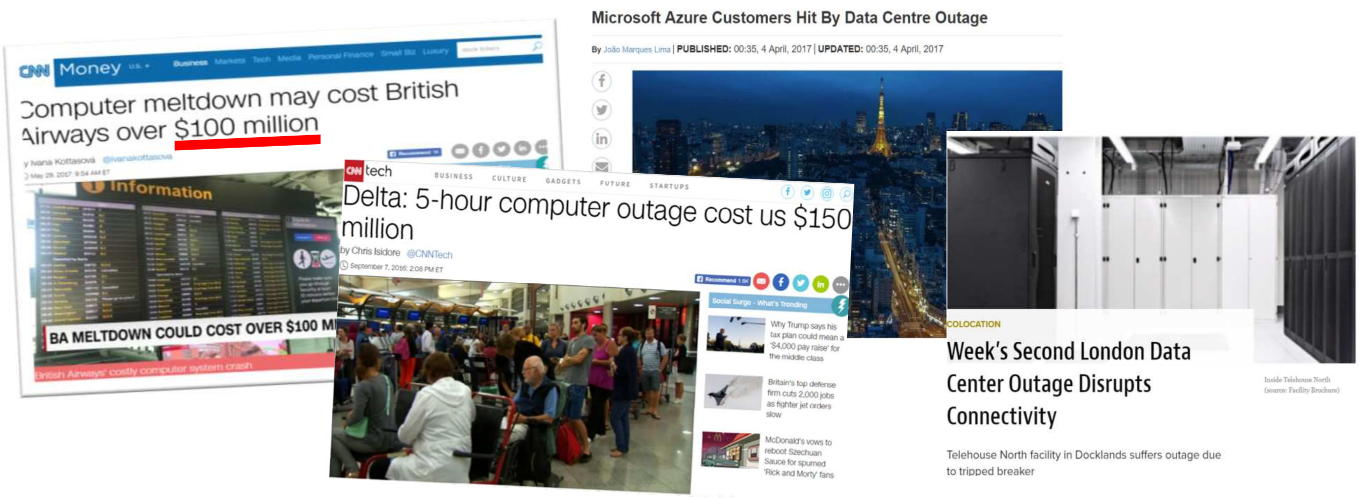

Rather than rare events, data center outages could be much more frequent

Cost analysis

- Estimated impact of capacity overloads (5% of the time) on a 1MW-10,000 sqft data center

Type	Redundancy	Downtime w/o Attack (hours/Yr)	Downtime w/ Attack (hours/Yr)	Increased Downtime Cost (mill. $/Yr)	Amortized Capital Loss (mill. $/Yr)	Total Cost (mill. $/Yr)
Tier-II	N+1 (generator/UPS/chiller)	22.69	366	22.12	0.1 (9+%↓)	22.22
Tier-III	N+1 (all non-IT equipment)	1.58	25.46	11.15	1.0 (50%↓)	12.15
Tier-IV	2N (all non-IT equipment)	0.44	6.59	3.42	1.1 (50%↓)	4.52

Cost analysis

- Estimated impact of capacity overloads (5% of the time) on a 1MW-10,000 sqft data center

Type	Redundancy	Downtime w/o Attack (hours/Yr)	Downtime w/ Attack (hours/Yr)	Increased Downtime Cost (mill. $/Yr)	Amortized Capital Loss (mill. $/Yr)	Total Cost (mill. $/Yr)
Tier-II	N+1 (generator/UPS/chiller)	22.69	366	22.12	0.1 (9+%↓)	22.22
Tier-III	N+1 (all non-IT equipment)	1.58	25.46	11.15	1.0 (50%↓)	12.15
Tier-IV	2N (all non-IT equipment)	0.44	6.59	3.42	1.1 (50%↓)	4.52

Million $ loss!

Cost analysis

- Estimated impact of capacity overloads (5% of the time) on a 1MW-10,000 sqft data center

Type	Redundancy	Downtime w/o Attack (hours/Yr)	Downtime w/ Attack (hours/Yr)	Increased Downtime Cost (mill. $/Yr)	Amortized Capital Loss (mill. $/Yr)	Total Cost (mill. $/Yr)
Tier-II	N+1 (generator/UPS/chiller)	22.69	366	22.12	0.1 (9+%↓)	22.22
Tier-III	N+1 (all non-IT equipment)	1.58	25.46	11.15	1.0 (50%↓)	12.15
Tier-IV	2N (all non-IT equipment)	0.44	6.59	3.42	1.1 (50%↓)	4.52

Million $ loss!

- Strong incentives: The attacker only spends US$ <500k (1-15% of the resulting loss)!
  - Data center operator's competitor
  - Against certain tenants to cause service disruptions
  - Creating chaos...

- Why are multi-tenant data centers vulnerable to power attacks?
  - Current safeguards are ineffective for well-timed power attacks
- What is the potential impact of power attacks?
  - Million dollar loss and service disruption
- How could an attacker mount a power attack?
- How to defend a data center against power attacks?

- Why are multi-tenant data centers vulnerable to power attacks?
  - Current safeguards are ineffective for well-timed power attacks
- What is the potential impact of power attacks?
  - Million dollar loss and service disruption
- How could an attacker mount a power attack?
- How to defend a data center against power attacks?

Attack opportunities are intermittent…

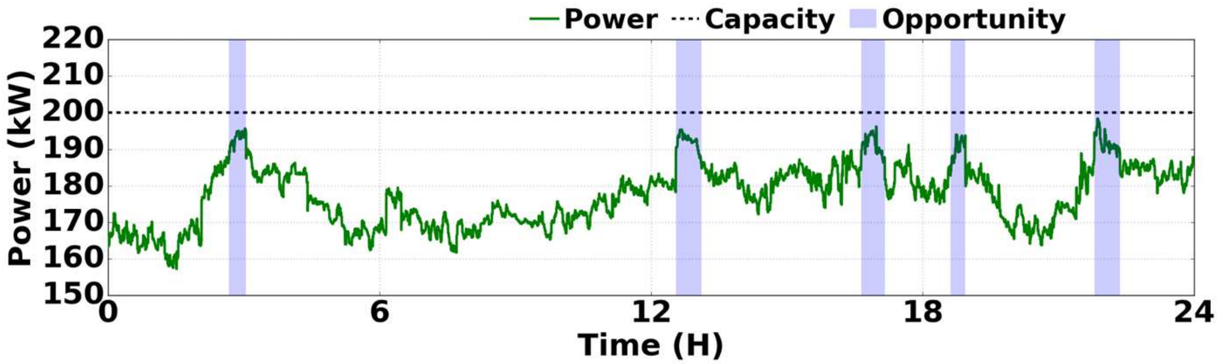

- Random attacks are unlikely to be successful, while constant full power is prohibited

Attack opportunities are intermittent…

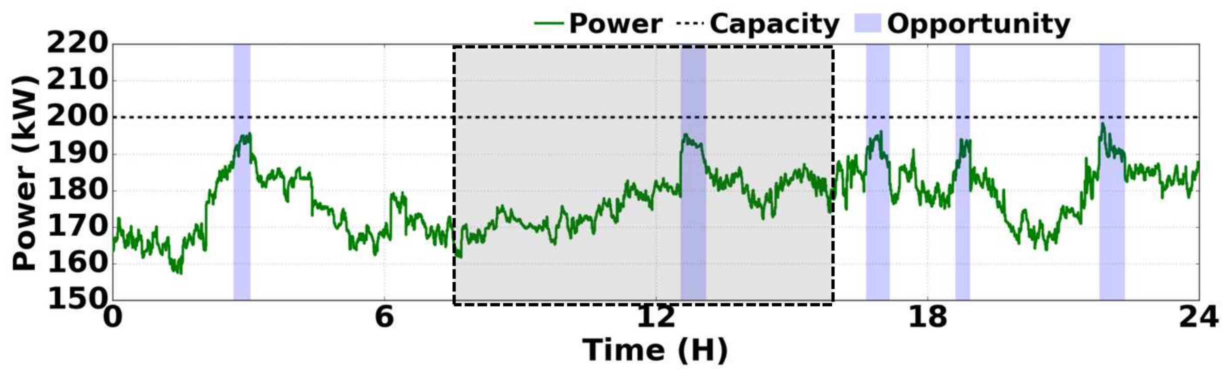

- Random attacks are unlikely to be successful, while constant full power is prohibited
- Coarse timing (e.g., based on "peak" hours) is ineffective

Attack opportunities are intermittent…

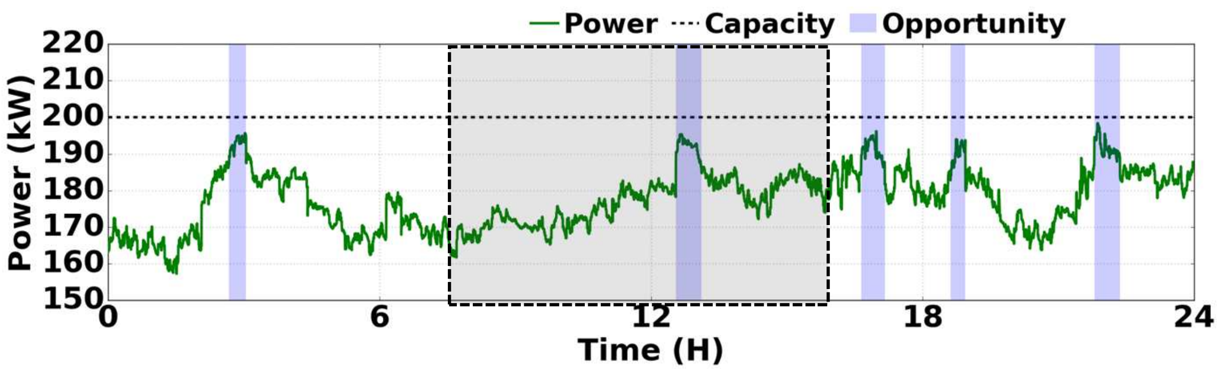

- Random attacks are unlikely to be successful, while constant full power is prohibited
- Coarse timing (e.g., based on "peak" hours) is ineffective

How to achieve a precise timing for successful power attacks?

In a multi-tenant data center…

Tenants co-locate their servers in a shared data center space

In a multi-tenant data center…

Tenants co-locate their servers in a shared data center space

Interconnected through physical processes that may leak power usage information

Power ⟶ Heat

A thermal side channel

- Hot air can travel to nearby racks, affecting their inlet temperatures

Demo of heat recirculation --- 5x speed viewing in Autodesk CFD

A thermal side channel

- Hot air can travel to nearby racks, affecting their inlet temperatures

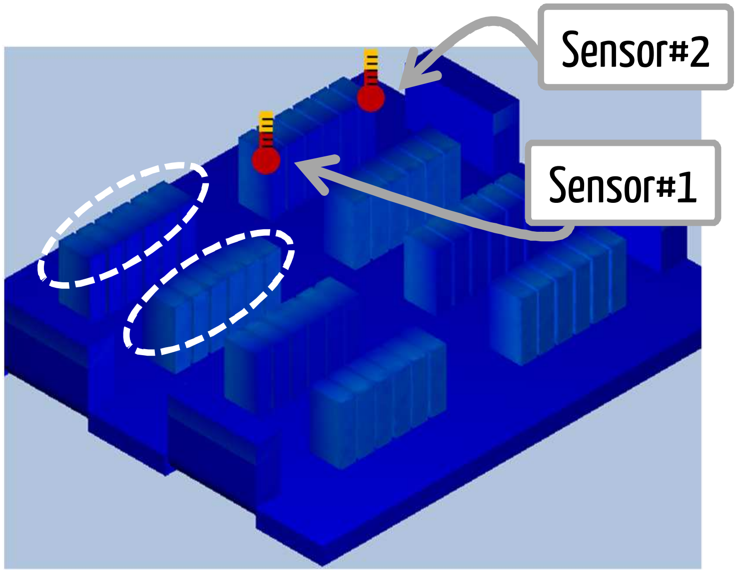

Demo of heat recirculation --- 5x speed viewing in Autodesk CFD

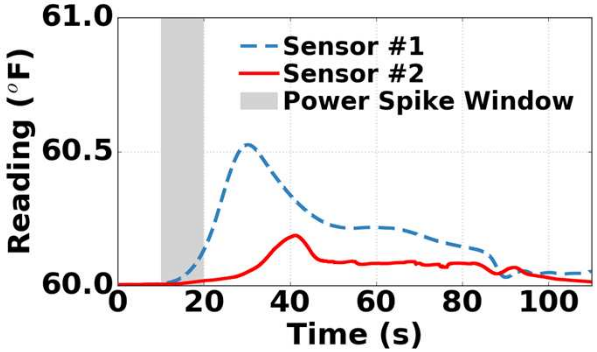

Temperature trace at select sensors

A high temperature doesn’t necessarily mean a high aggregate power usage…

A high temperature doesn’t necessarily mean a high aggregate power usage…

Heat recirculation is spatially non-uniform --- more significant among nearby racks!

A closer look at thermal network

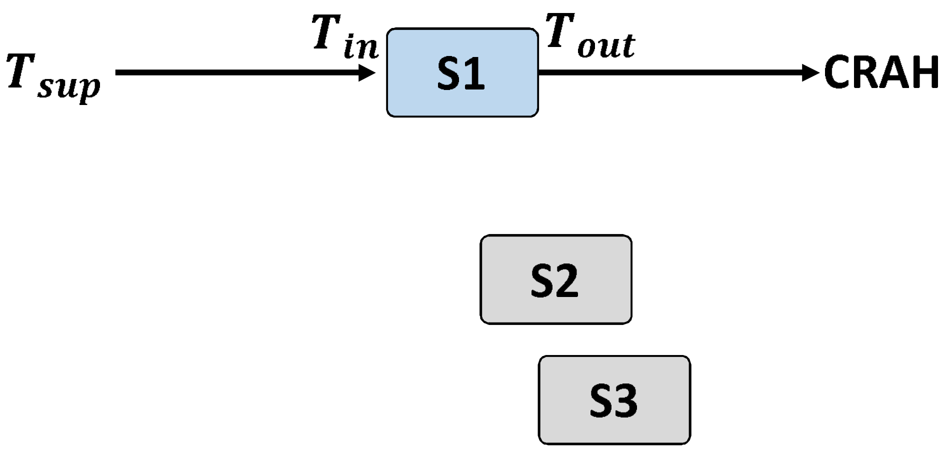

A closer look at thermal network

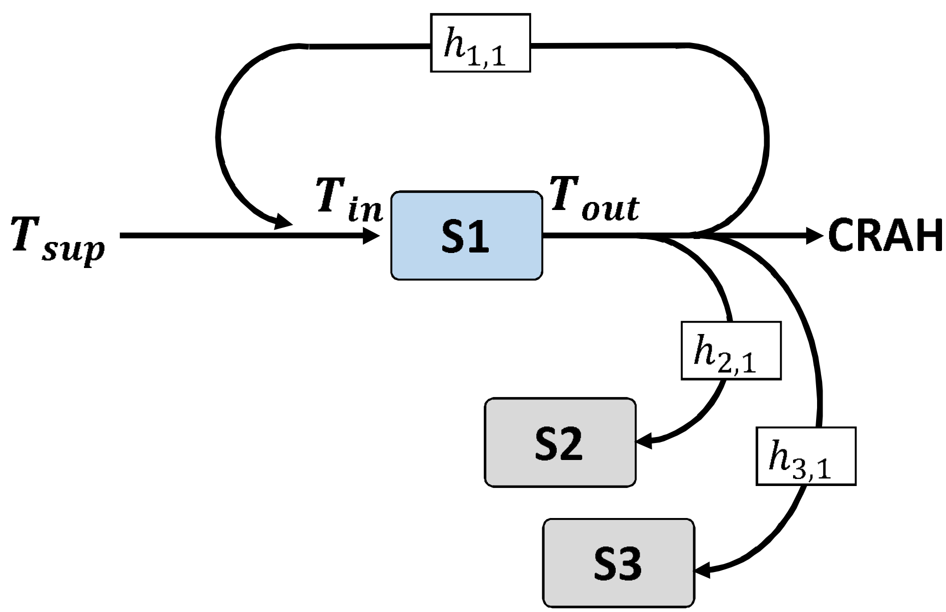

A closer look at thermal network

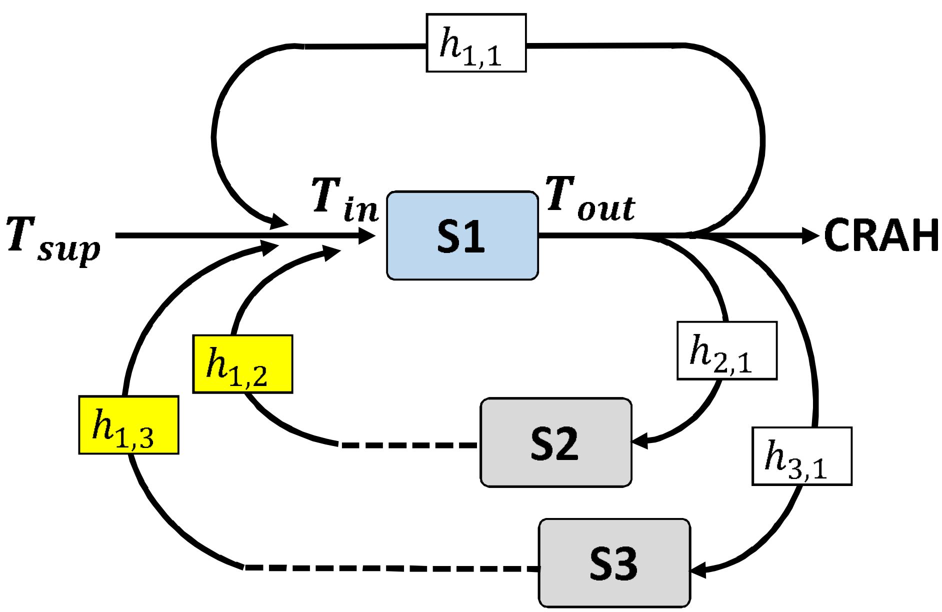

A closer look at thermal network

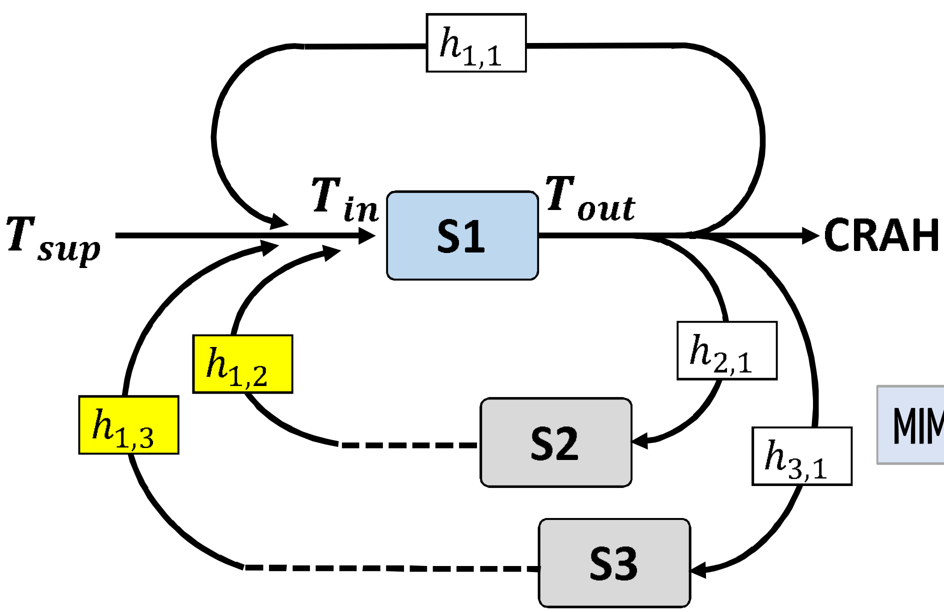

$$T_m(t) = T_{sup}(t) + \sum_{n=1}^{N+J}\sum_{\tau=0}^{K-1} p_n(t-\tau)\cdot h_{m,n}(\tau) + r_m(t)$$

Sensor reading

Heat recirculation impact

MIMO model: N benign servers, J attacker servers, and M sensors

A closer look at thermal network

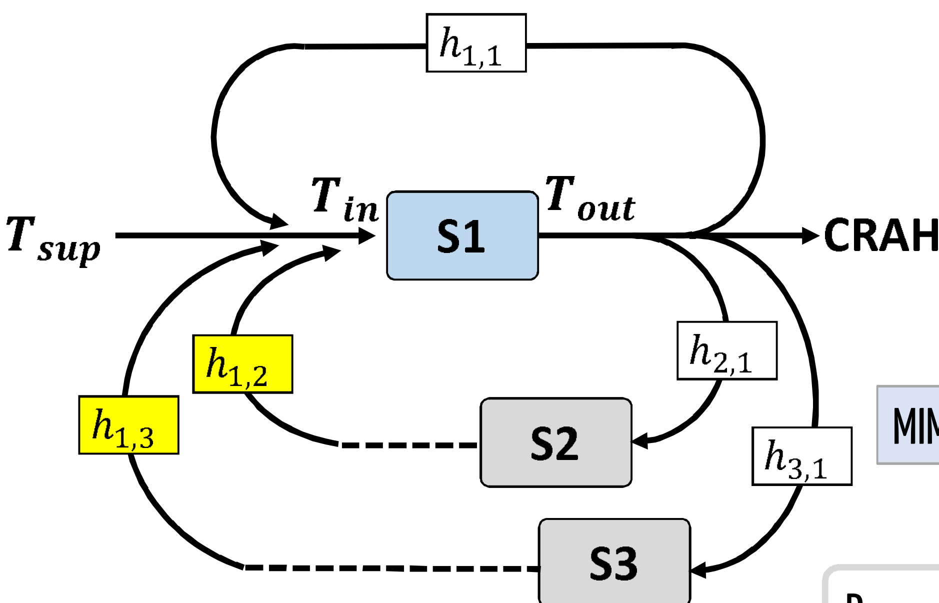

$$T_m(t) = T_{sup}(t) + \sum_{n=1}^{N+J} \sum_{\tau=0}^{K-1} p_n(t-\tau) \cdot h_{m,n}(\tau) + r_m(t)$$

Sensor reading

Heat recirculation impact

MIMO model: N benign servers, J attacker servers, and M sensors

Power information is leaked and embedded into temperature readings...

Rewriting the attacker's observation model

$$z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$$

Rewriting the attacker's observation model

Temperature increase due to benign tenants

Noise

$$z_t = \mathbf{H}_b x_t + r_t$$

Impact from previous slots

$$\mathbf{H}_b = \begin{bmatrix} h_{1,1}(t) & \cdots & h_{1,N}(t) & \cdots & h_{1,1}(t-K) & \dots & h_{1,N}(t-K) \\ & \vdots & & \ddots & & \vdots & \\ h_{M,1}(t) & \cdots & h_{M,N}(t) & \cdots & h_{M,1}(t-K) & \cdots & h_{M,N}(t-K) \end{bmatrix}$$

$$x_t = [p_1(t) \quad \cdots \quad p_N(t) \quad \cdots \quad p_1(t-K) \quad \cdots \quad p_N(t-K)]^{\mathrm{T}}$$

Rewriting the attacker's observation model

$$z_t = \mathbf{H}_b x_t + r_t$$

$$\mathbf{H}_b = \begin{bmatrix} h_{1,1}(t) & \cdots & h_{1,N}(t) & \cdots & h_{1,1}(t-K) & \dots & h_{1,N}(t-K) \\ & \vdots & & \ddots & & \vdots & \\ h_{M,1}(t) & \cdots & h_{M,N}(t) & \cdots & h_{M,1}(t-K) & \cdots & h_{M,N}(t-K) \end{bmatrix}$$

Challenges:

- $\mathbf{H}_b$ has a size of M by N·K, very large for $N \in [500,1000]$ servers
- Difficult to obtain accurately, and high computational complexity

Rewriting the attacker's observation model

A signal estimation problem with imperfect channel state information

$$z_t = \mathbf{H}_b x_t + r_t$$

$$\mathbf{H}_b = \begin{bmatrix} h_{1,1}(t) & \cdots & h_{1,N}(t) & \cdots & h_{1,1}(t-K) & \dots & h_{1,N}(t-K) \\ & \vdots & & \ddots & & \vdots & \\ h_{M,1}(t) & \cdots & h_{M,N}(t) & \cdots & h_{M,1}(t-K) & \cdots & h_{M,N}(t-K) \end{bmatrix}$$

Challenges:

- $\mathbf{H}_b$ has a size of M by N·K, very large for $N \in [500{,}1000]$ servers
- Difficult to obtain accurately, and high computational complexity

Approximate zone-level thermal network--- Divide data center into zones

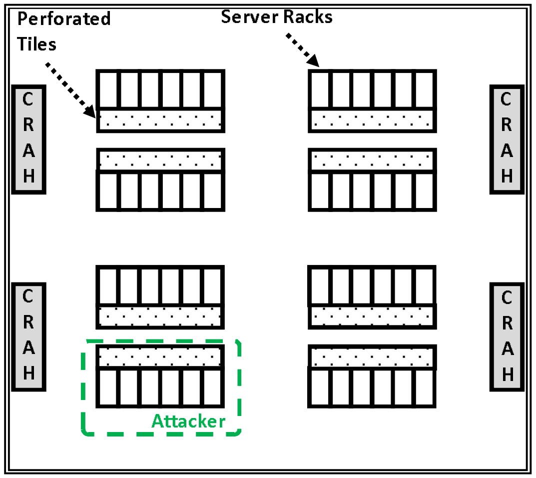

$$\mathbf{H}_b = \begin{bmatrix} h_{1,1} & \cdots & h_{1,50} & h_{1,51} & \cdots & h_{1,100} & h_{1,101} & \cdots & h_{1,N} & \cdots & \cdots \\ & & & & \vdots & & & & & \ddots & \vdots \\ h_{M,1} & \cdots & h_{M,50} & h_{M,51} & \cdots & h_{M,100} & h_{M,101} & \cdots & h_{M,N} & \cdots & \cdots \end{bmatrix}$$

Approximate zone-level thermal network--- Divide data center into zones

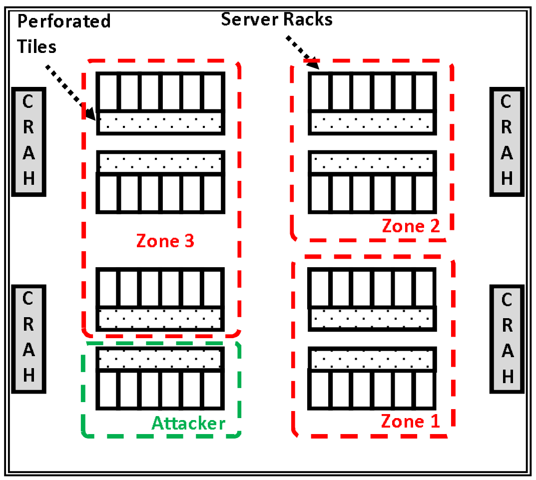

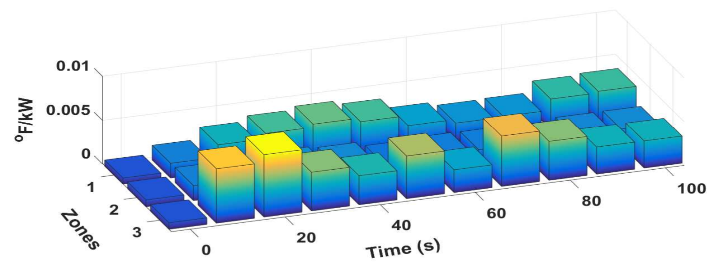

Zone-level heat recirculation matrix

Obtained offline (say, through CFD)

Replaced by one value

$$\mathbf{H}_b = \begin{bmatrix} h_{1,1} & \cdots & h_{1,50} & h_{1,51} & \cdots & h_{1,100} & h_{1,101} & \cdots & h_{1,N} & \cdots & \cdots \\ & & & & \vdots & & & & & \ddots & \vdots \\ h_{M,1} & \cdots & h_{M,50} & h_{M,51} & \cdots & h_{M,100} & h_{M,101} & \cdots & h_{M,N} & \cdots & \cdots \end{bmatrix}$$

Approximate zone-level thermal network--- Divide data center into zones

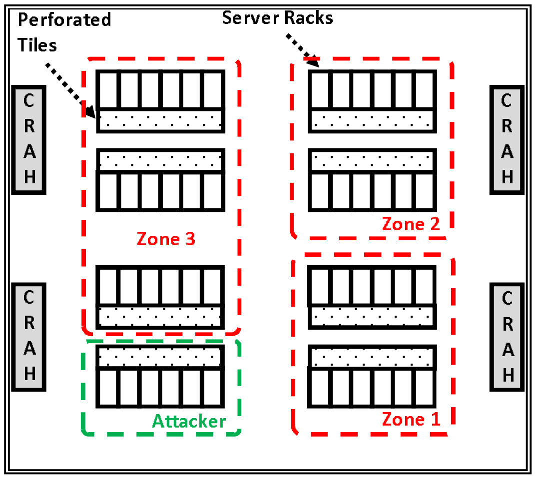

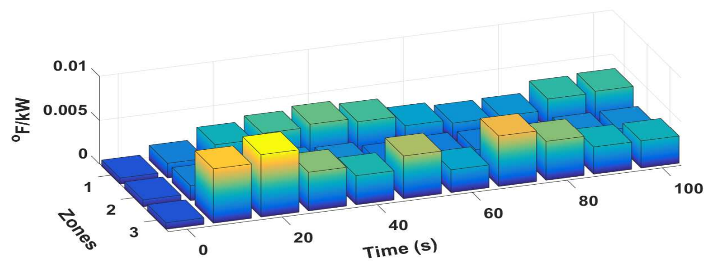

Zone-level heat recirculation matrix

Obtained offline (say, through CFD)

Error in $\mathbf{H}_b$

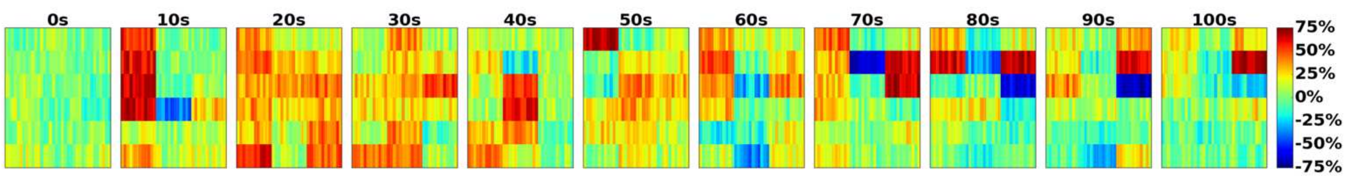

Estimating x_t from $z_t = \mathbf{H}_b x_t + r_t$

Solution: State-augmented robust Kalman filter

Estimating x_t from $z_t = \mathbf{H}_b x_t + r_t$

x_t is the augmented state, z_t is the observation

Assumed state transition model: $x_{t+1} = \mathbf{F}x_t + q_t$

Predict:
$$\begin{cases} \hat{x}_{t|t-1} = \mathbf{F}\hat{x}_{t-1|t-1} \\ \mathbf{P}_{t|t-1} = \mathbf{F}\mathbf{P}_{t-1|t-1} + \mathbf{Q} \end{cases}$$

Update:
$$\begin{cases} u_t = z_t - \mathbf{H}_b \hat{x}_{t|t-1} \\ \mathbf{S}_t = \mathbf{H}_b \mathbf{P}_{t|t-1} \mathbf{H}_b^{\mathrm{T}} + \mathbf{R} \\ \mathbf{G}_t = \mathbf{P}_{t|t-1} \mathbf{H}_b^{\mathrm{T}} \mathbf{S}_t^{-1} \\ \hat{x}_{t|t} = \hat{x}_{t|t-1} + \mathbf{G}_t u_t \\ \mathbf{P}_{t|\mathrm{t}} = (\mathbf{I} - \mathbf{G}_t \mathbf{H}_b) \mathbf{P}_{t|\mathrm{t}-1} \end{cases}$$

Solution: State-augmented robust Kalman filter

Estimating x_t from $z_t = \mathbf{H}_b x_t + r_t$

x_t is the augmented state, z_t is the observation

Assumed state transition model: $x_{t+1} = \mathbf{F} x_t + q_t$

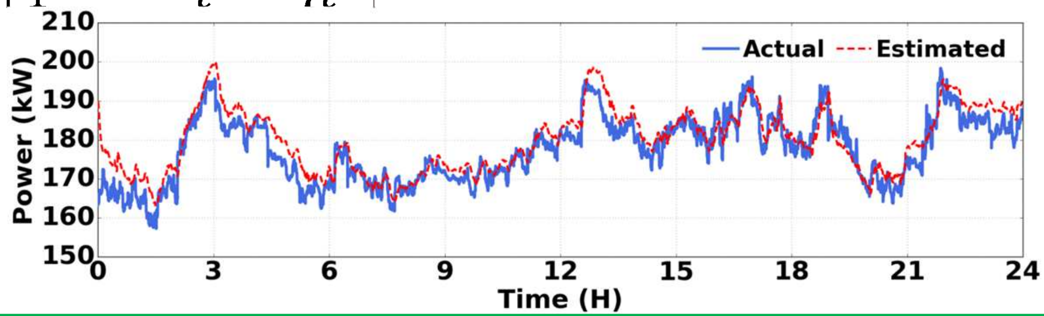

Predict: $\begin{cases} \hat{x}_{t|t-1} = \mathbf{F}\hat{x}_{t-1|t-1} \\ \mathbf{P}_{t|t-1} = \mathbf{F}\mathbf{P}_{t-1|t-1} + \end{cases}$

Update: $\begin{cases} u_t = z_t - \mathbf{H}_b \hat{x}_{t|t-1} \\ \mathbf{S}_t = \mathbf{H}_b \mathbf{P}_{t|t-1} \mathbf{H}_b^{\mathrm{T}} + \\ \mathbf{G}_t = \mathbf{P}_{t|t-1} \mathbf{H}_b^{\mathrm{T}} \mathbf{S}_t^{-1} \\ \hat{x}_{t|t} = \hat{x}_{t|t-1} + \mathbf{G}_t u_t \\ \mathbf{P}_{t|\mathrm{t}} = (\mathbf{I} - \mathbf{G}_t \mathbf{H}_b) \mathbf{P}_{t|\mathrm{t}-1} \end{cases}$

Avg. error < 3% for estimating benign tenants' aggregate power usage

An attack strategy…

- Attack when the estimate of benign tenants' power usage is sufficiently high
- Wait for some time before attacks
- Each attack lasts no more than T_{hold}, and no consecutive attacks

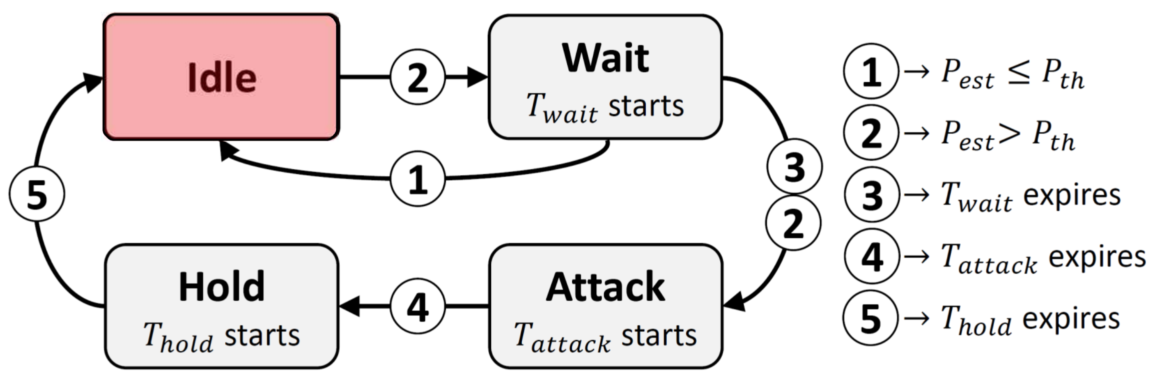

An attack strategy…

- Attack when the estimate of benign tenants' power usage is sufficiently high
- Wait for some time before attacks
- Each attack lasts no more than T_{hold}, and no consecutive attacks

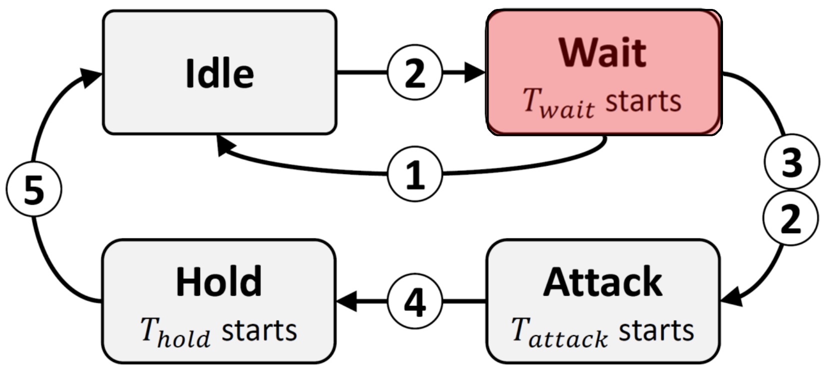

1 → $P_{est} \leq P_{th}$
2 → $P_{est} > P_{th}$
3 → T_{wait} expires
4 → T_{attack} expires
5 → T_{hold} expires

An attack strategy...

- Attack when the estimate of benign tenants' power usage is sufficiently high
- Wait for some time before attacks
- Each attack lasts no more than T_{hold}, and no consecutive attacks

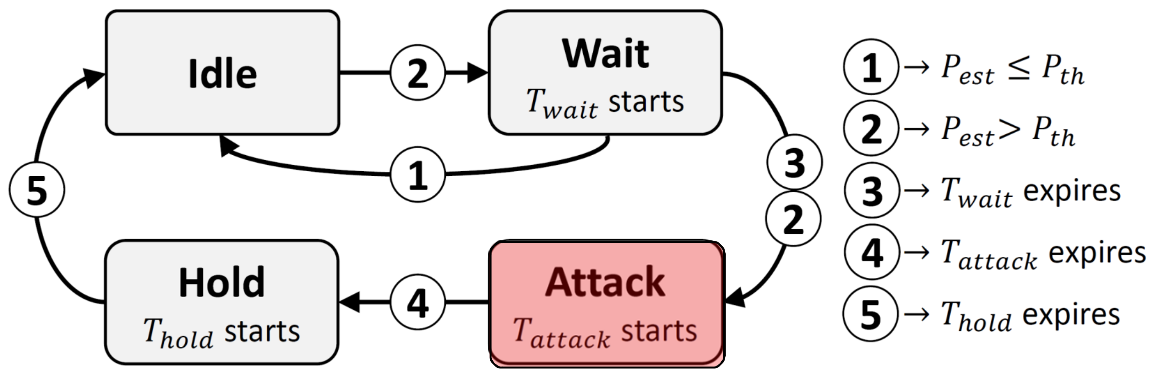

e.g., running CPU-intensive computations...

An attack strategy…

- Attack when the estimate of benign tenants' power usage is sufficiently high
- Wait for some time before attacks
- Each attack lasts no more than T_{hold}, and no consecutive attacks

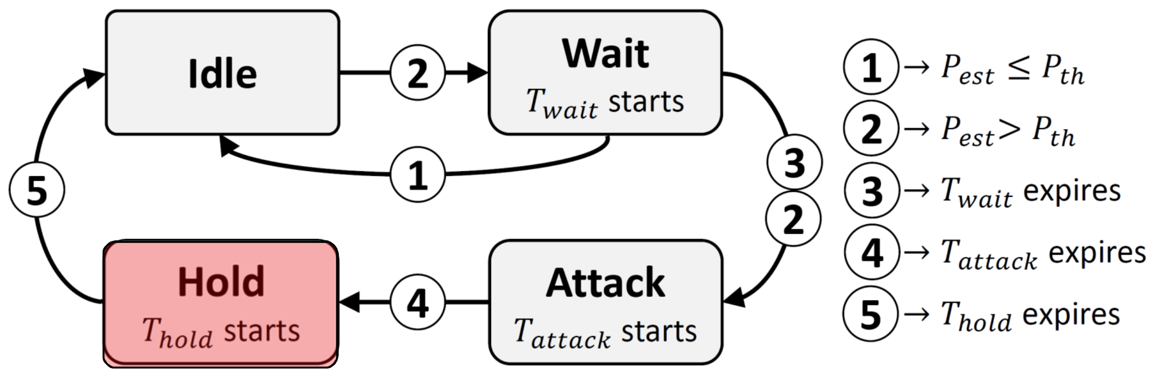

e.g., running CPU-intensive computations…

An attack strategy…

- Attack when the estimate of benign tenants' power usage is sufficiently high
- Wait for some time before attacks
- Each attack lasts no more than T_{hold}, and no consecutive attacks

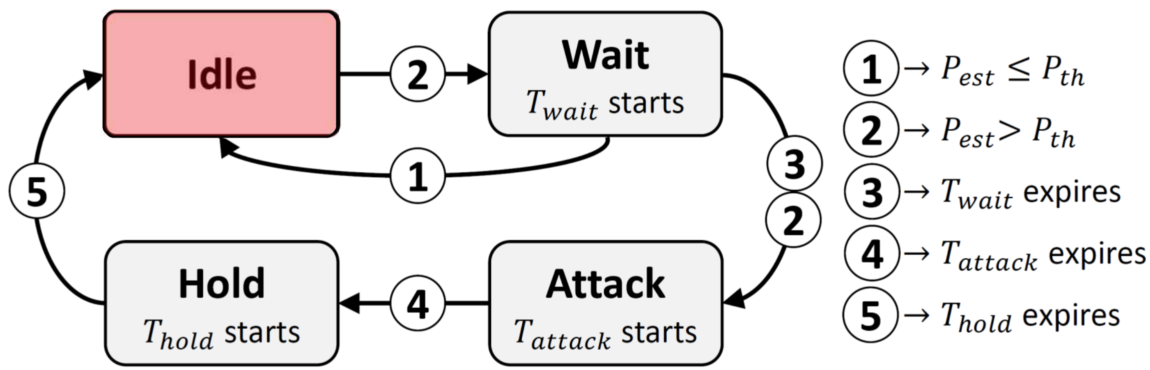

Illustration of well-timed power attacks

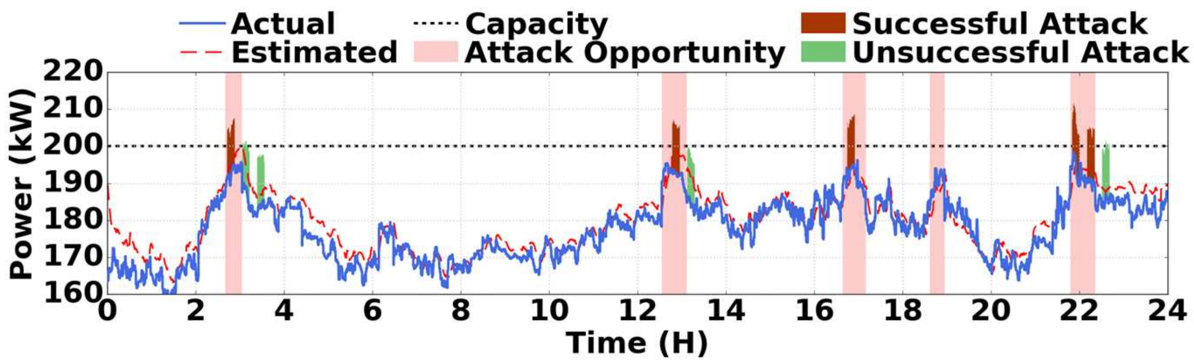

- Experimental settings
  - Simulated real workload traces based on a HP data center layout
  - Consider an attacker sharing a data center capacity of 200kW with benign tenants
  - Attack for no more than 10% of the times

Illustration of well-timed power attacks

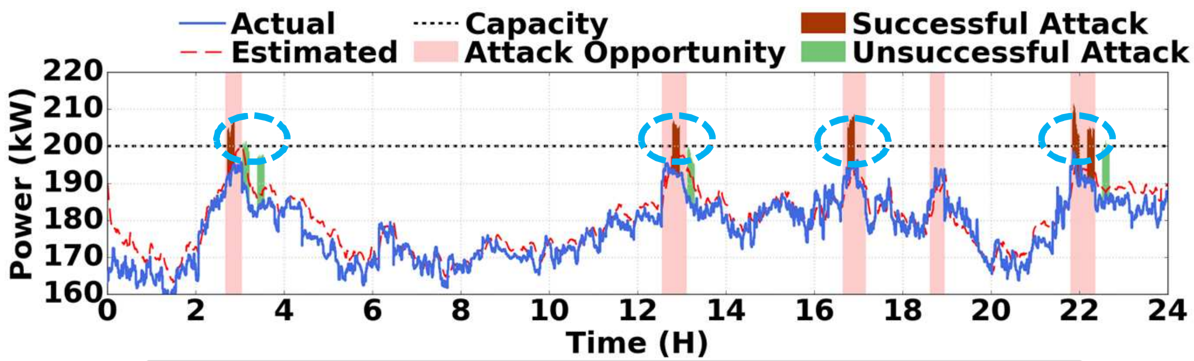

Precise timing through a thermal side channel...

- Experimental settings
  - Simulated real workload traces based on a HP data center layout
  - Consider an attacker sharing a data center capacity of 200kW with benign tenants
  - Attack for no more than 10% of the times

Timing accuracy

Attack more frequently with a lower triggering threshold

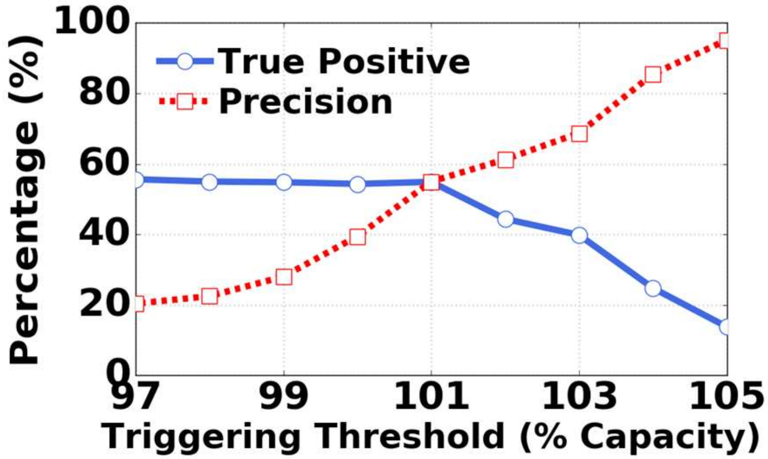

- True positive: % of attack opportunities detected
- Precision: % of attacks being successful

Timing accuracy

Attack more frequently with a lower triggering threshold

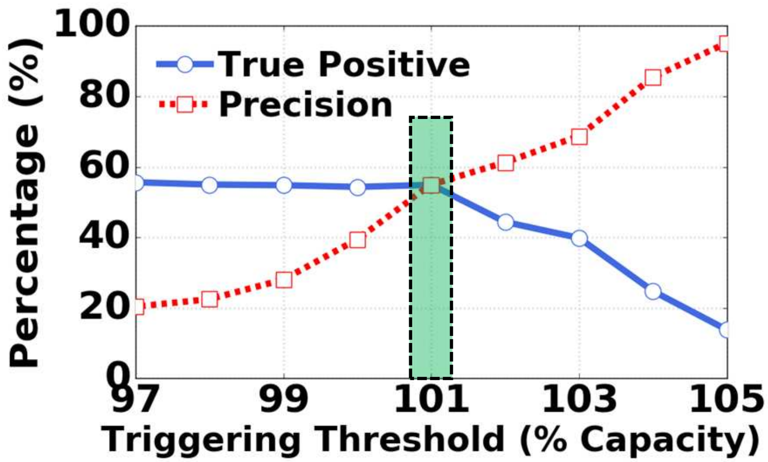

54% TP (10% for random attacks), and 53% precision

- True positive: % of attack opportunities detected
- Precision: % of attacks being successful

- Why are multi-tenant data centers vulnerable to power attacks?
  - Current safeguards are ineffective for well-timed power attacks
- What is the potential impact of power attacks?
  - Million dollar loss and service interruption
- How could an attacker mount a power attack?
  - Exploiting physical side channels (e.g., thermal/acoustic networks…)
- How to defend a data center against power attacks?

- Why are multi-tenant data centers vulnerable to power attacks?
    - Current safeguards are ineffective for well-timed power attacks
- What is the potential impact of power attacks?
    - Million dollar loss and service interruption
- How could an attacker mount a power attack?
    - Exploiting physical side channels (e.g., thermal/acoustic networks...)
- How to defend a data center against power attacks?

Randomizing physical side channels...

Thermal network model: $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$

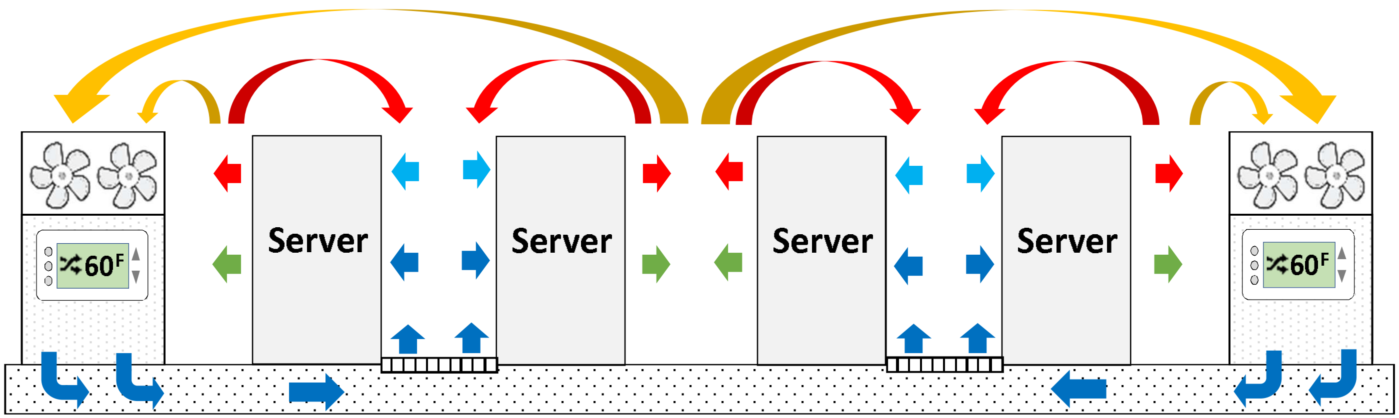

Randomizing physical side channels...

Thermal network model: $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$

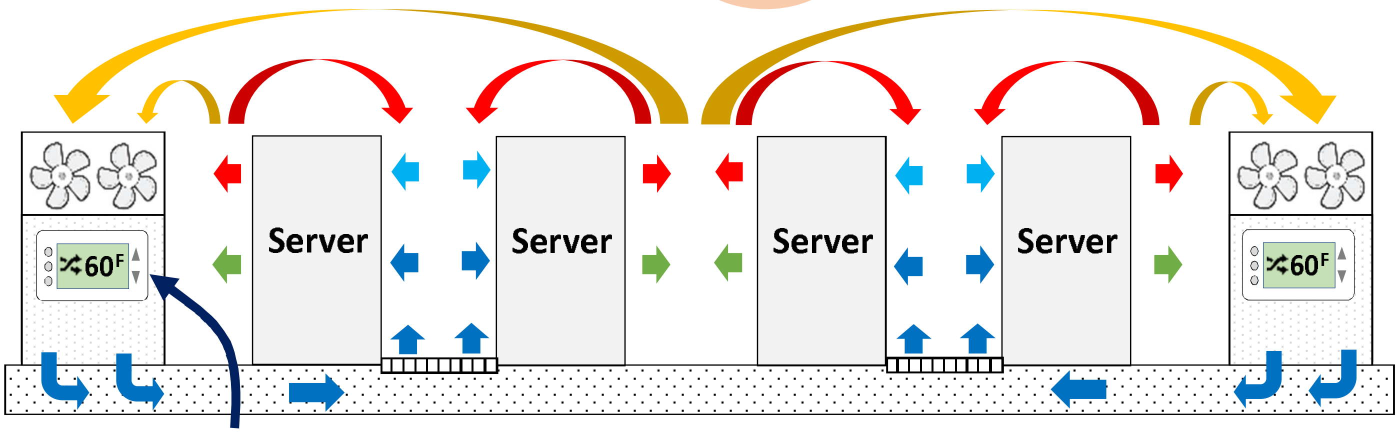

Random supply air temperature

Attacker can track the change!

Randomizing physical side channels…

Thermal network model: $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$

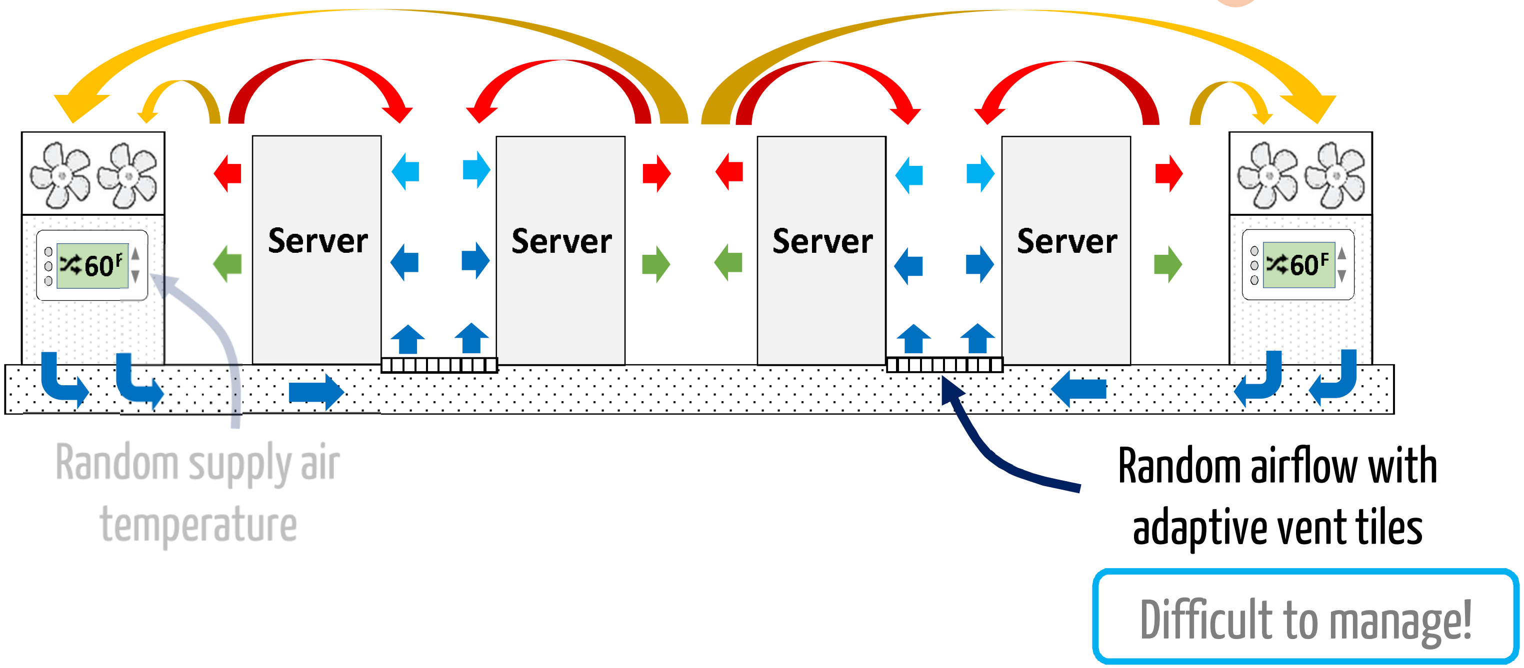

Randomizing physical side channels...

Thermal network model: $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$

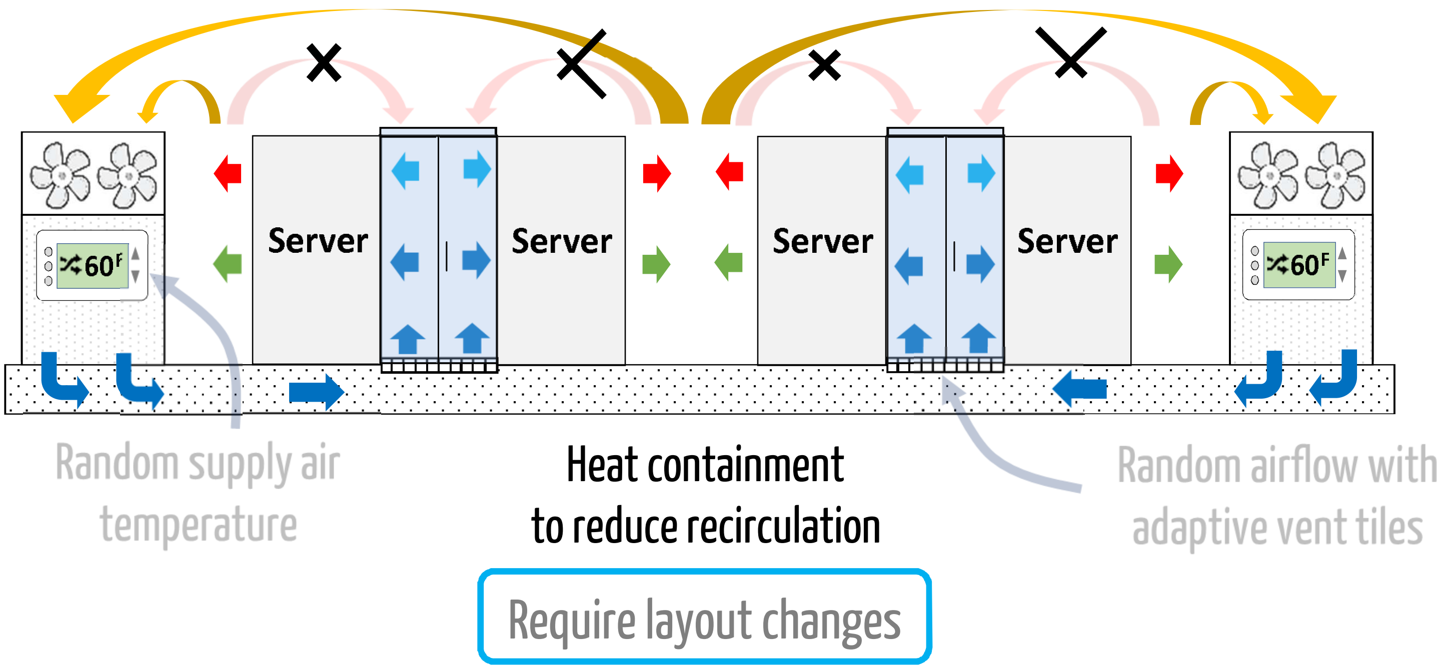

Randomizing physical side channels…

Thermal network model: $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t$

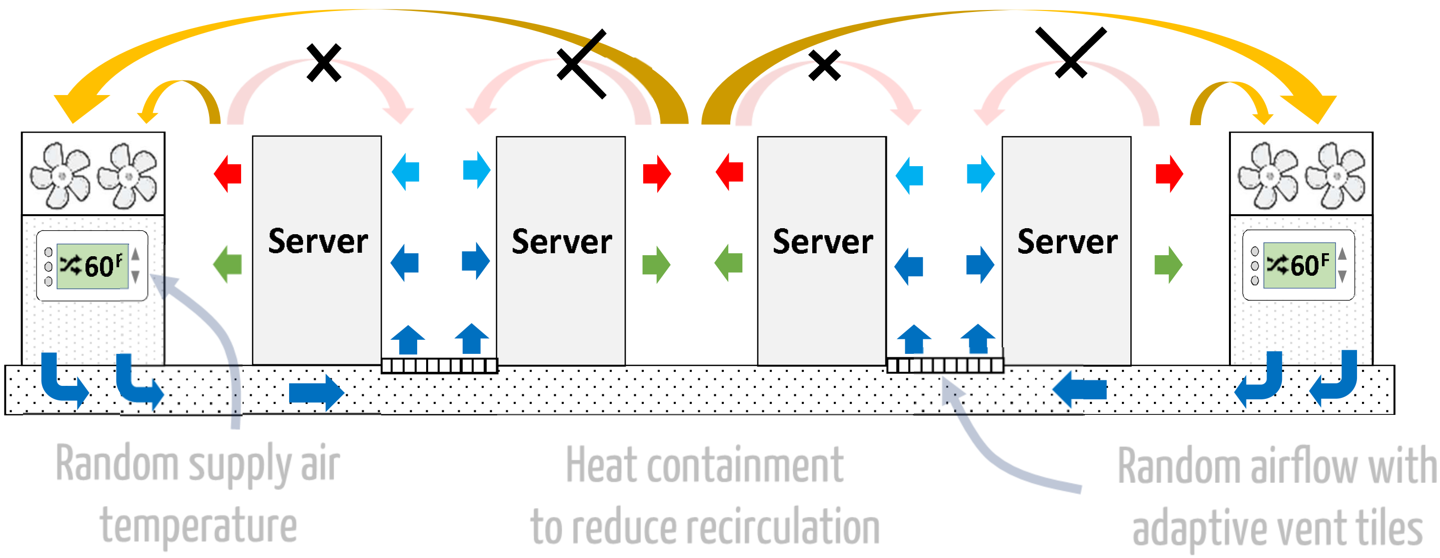

- Finding and evicting suspicious tenants
  - Intelligent power monitoring to find abnormal power usage patterns

- Why are multi-tenant data centers vulnerable to power attacks?
  - Current safeguards are ineffective for well-timed power attacks
- What is the potential impact of power attacks?
  - Million dollar loss and service interruption
- How could an attacker mount a power attack?
  - Exploiting physical side channels (e.g., thermal/acoustic networks…)
- How to defend a data center against power attacks?
  - A comprehensive investigation required

A cyber-physical view...

How about physical security?

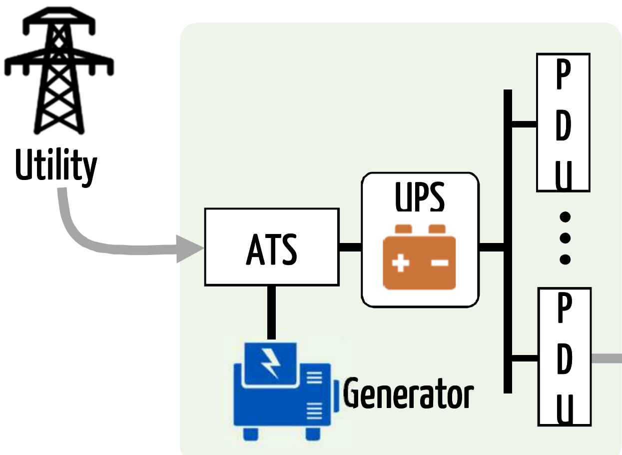

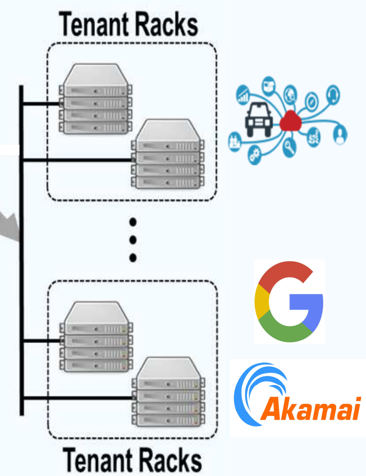

Securing the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

[Mirkovic, Sigcomm'04] [Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]...

A cyber-physical view…

How about physical security?

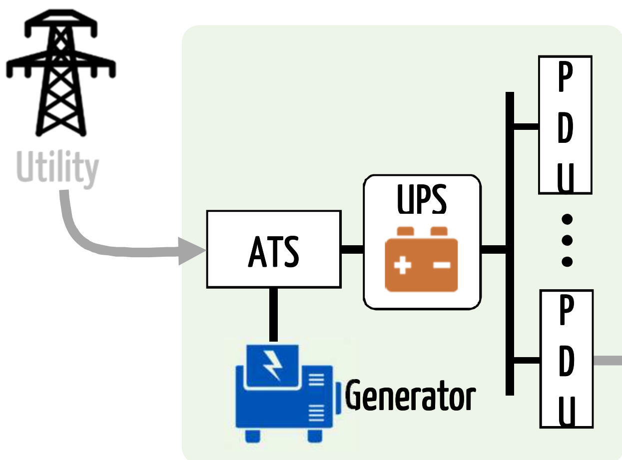

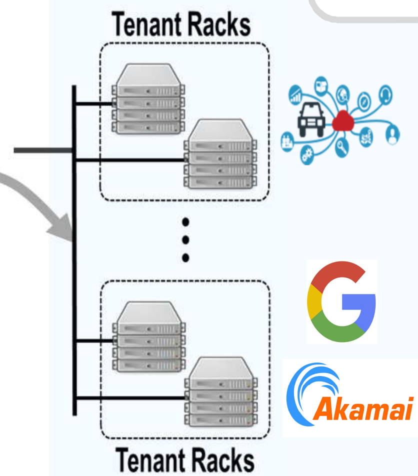

A thermal side channel can help the attacker precisely time its power attacks

Securing the cyberspace

- DDoS attack, network intrusion, privacy protection, etc.

[Mirkovic, Sigcomm'04] [Zhang CCS'12] [Moon CCS'15] [Dong CCS'17]…

Thanks!