Grid-Aware AI: Operational and Market Strategies for Large-Scale Data Centers

Vladimir Dvorkin

Department of Electrical Engineering and Computer Science

University of Michigan – Ann Arbor

"Data Center and Power Grid Interactions: Challenges and Opportunities"

Focused Research Program Symposium

Boston University

May 19, 2025

What makes data centers (DC) such a special load?

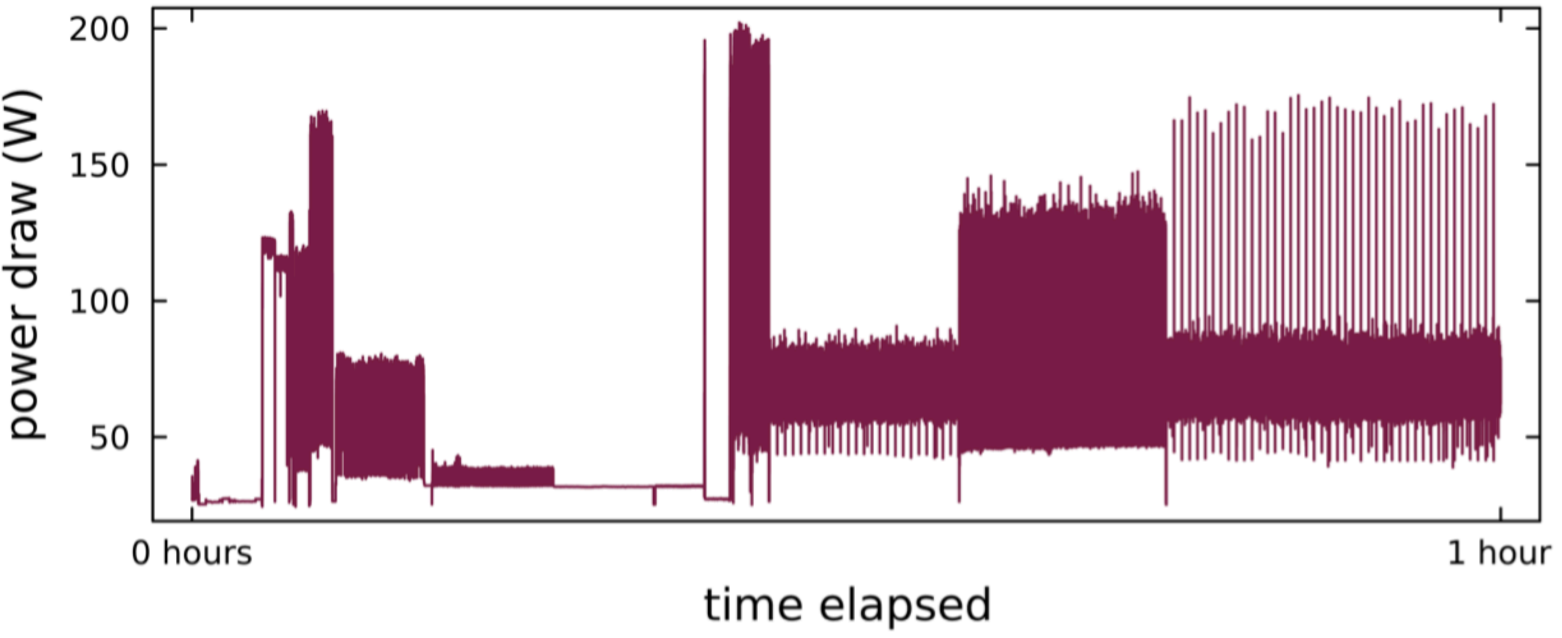

Disruptive electricity consumption

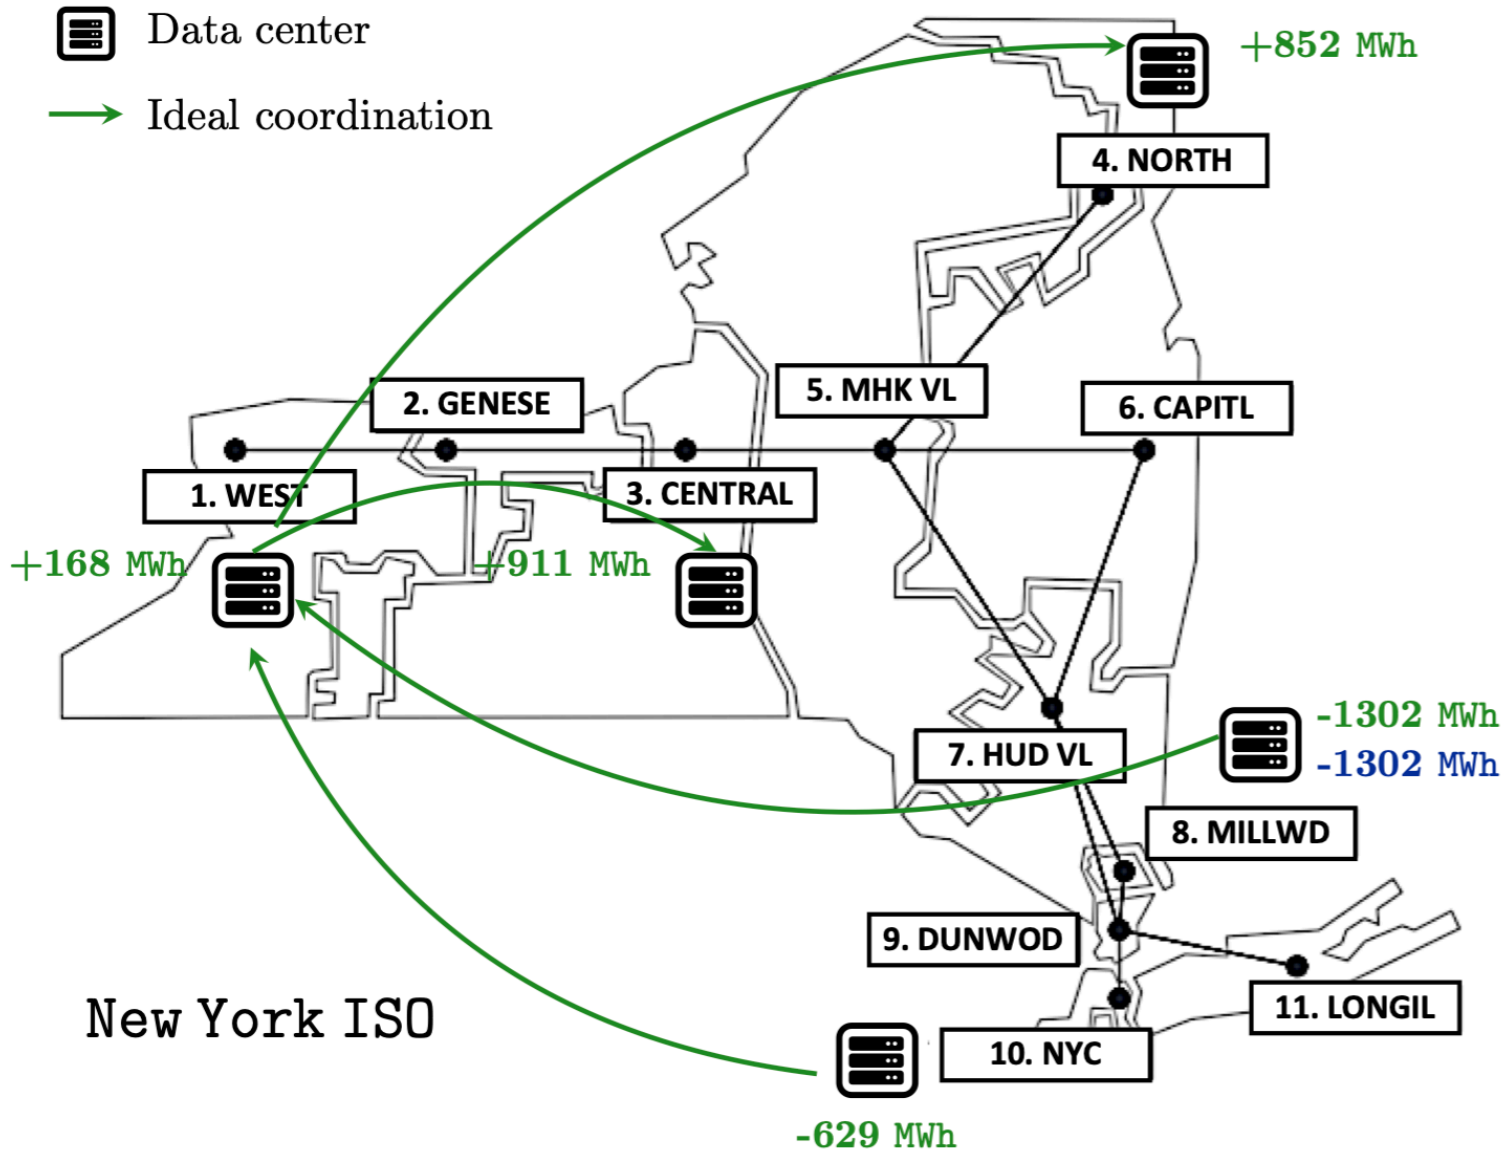

Space-time flexibility of workloads

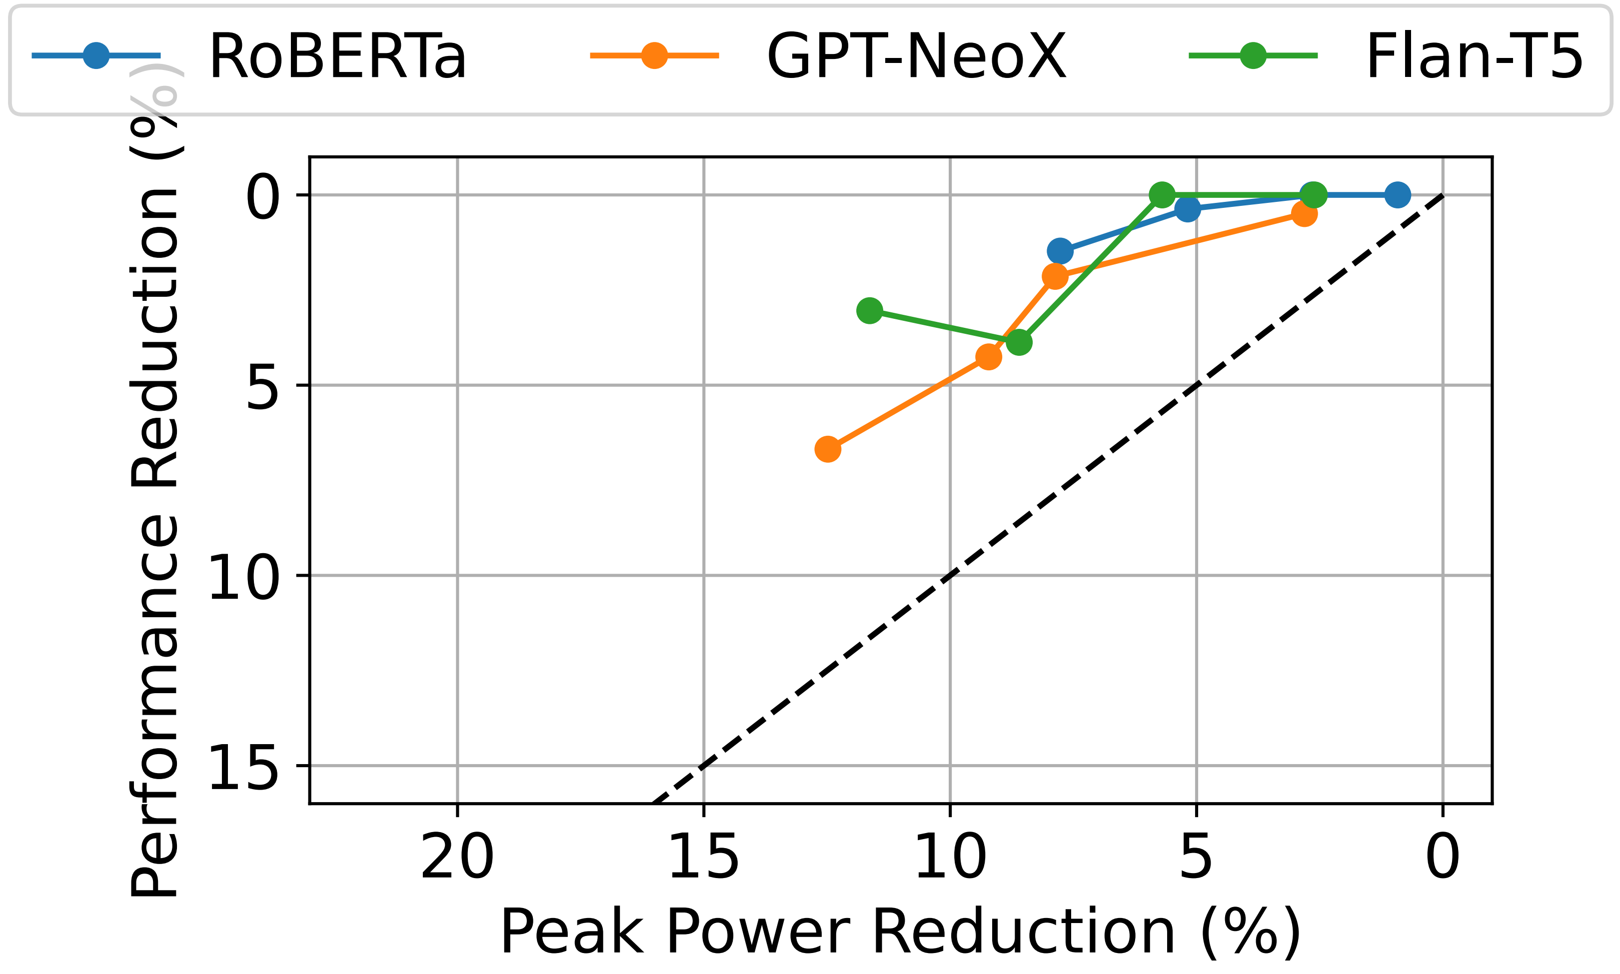

Demand engineering opportunities

Grid integration of DC at different timescales

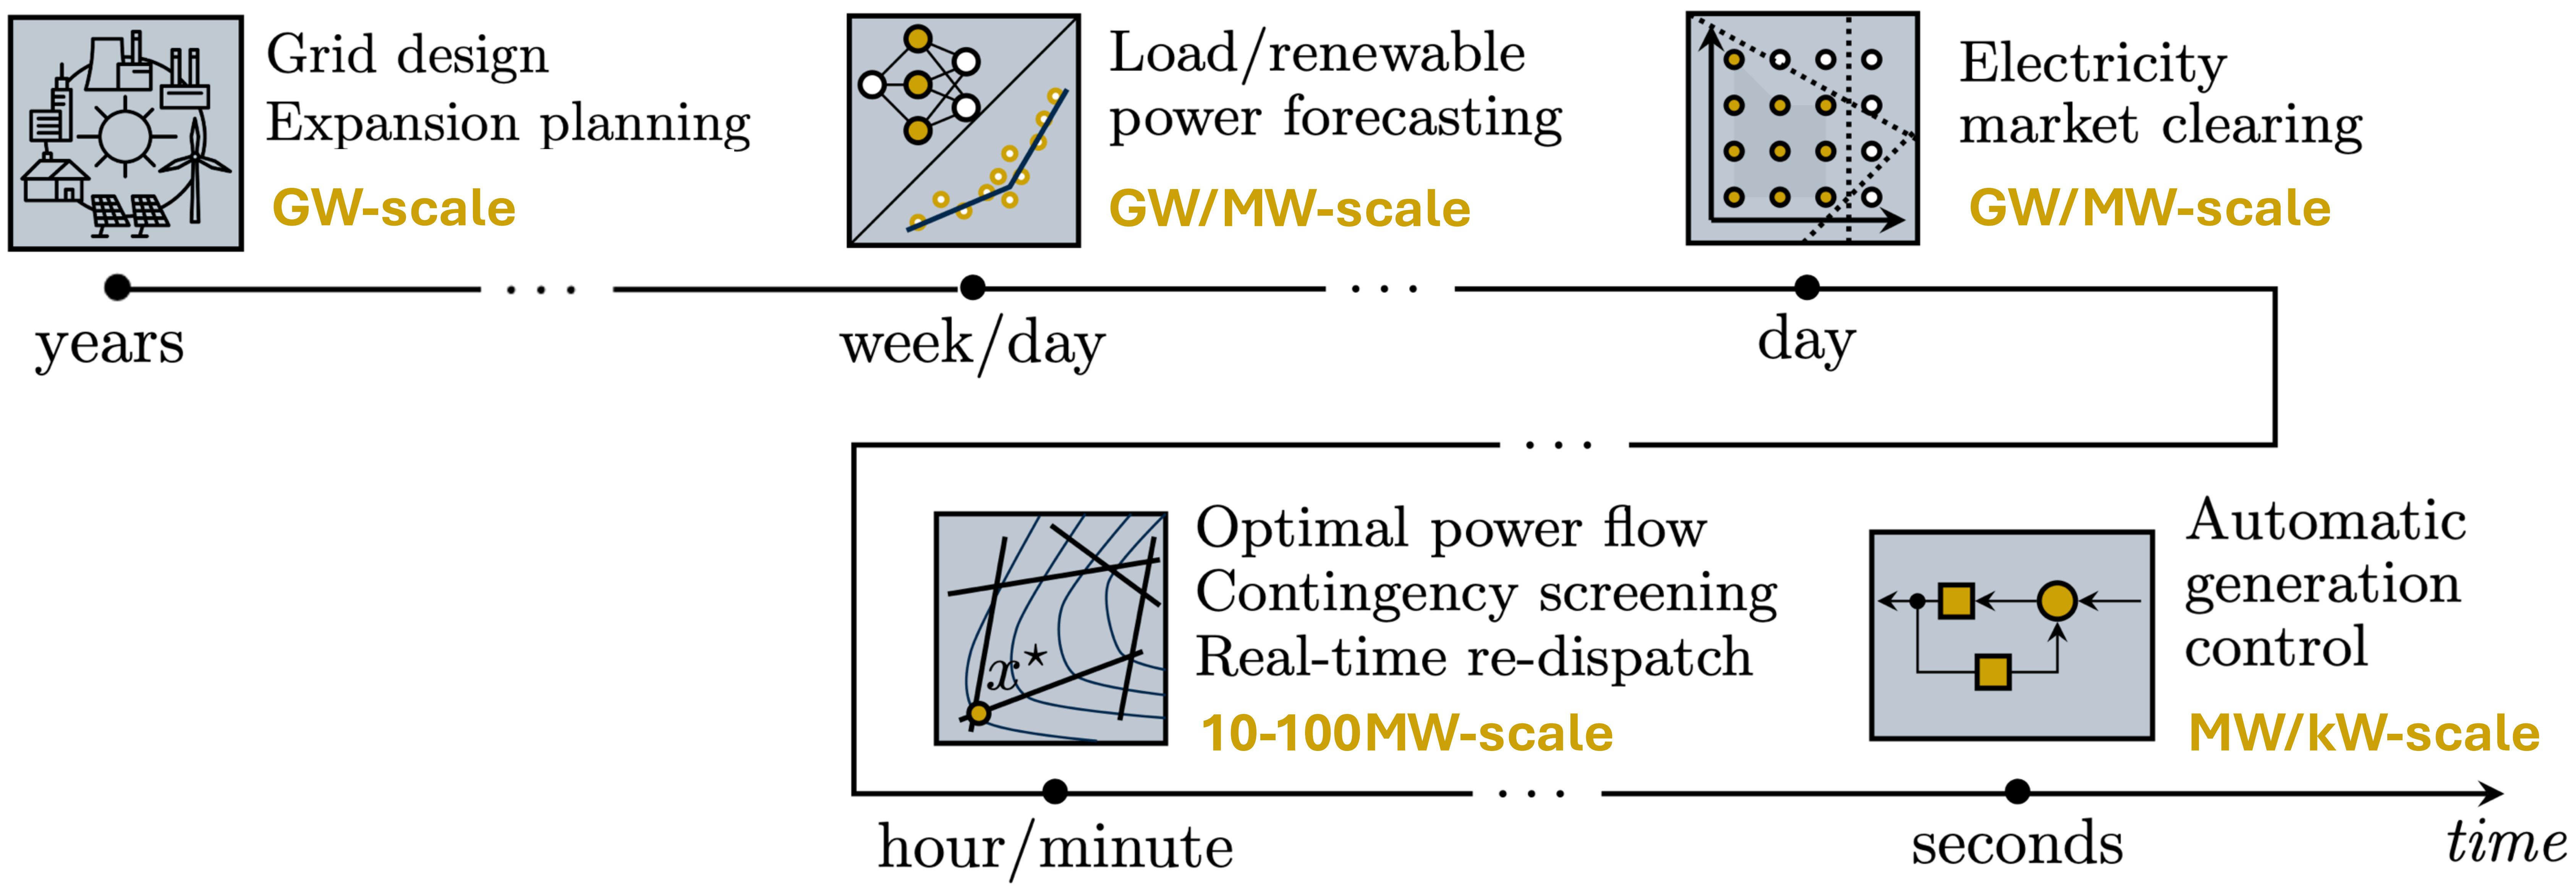

Diminishing coord. benefits

- Optimal sizing, siting and timing of data centers in the power grid
- Load forecasting and operating reserve sizing
- Day-ahead operational planning and market-clearing
- Near real-time operational coordination
- Real-time control

Grid integration of DC at different timescales

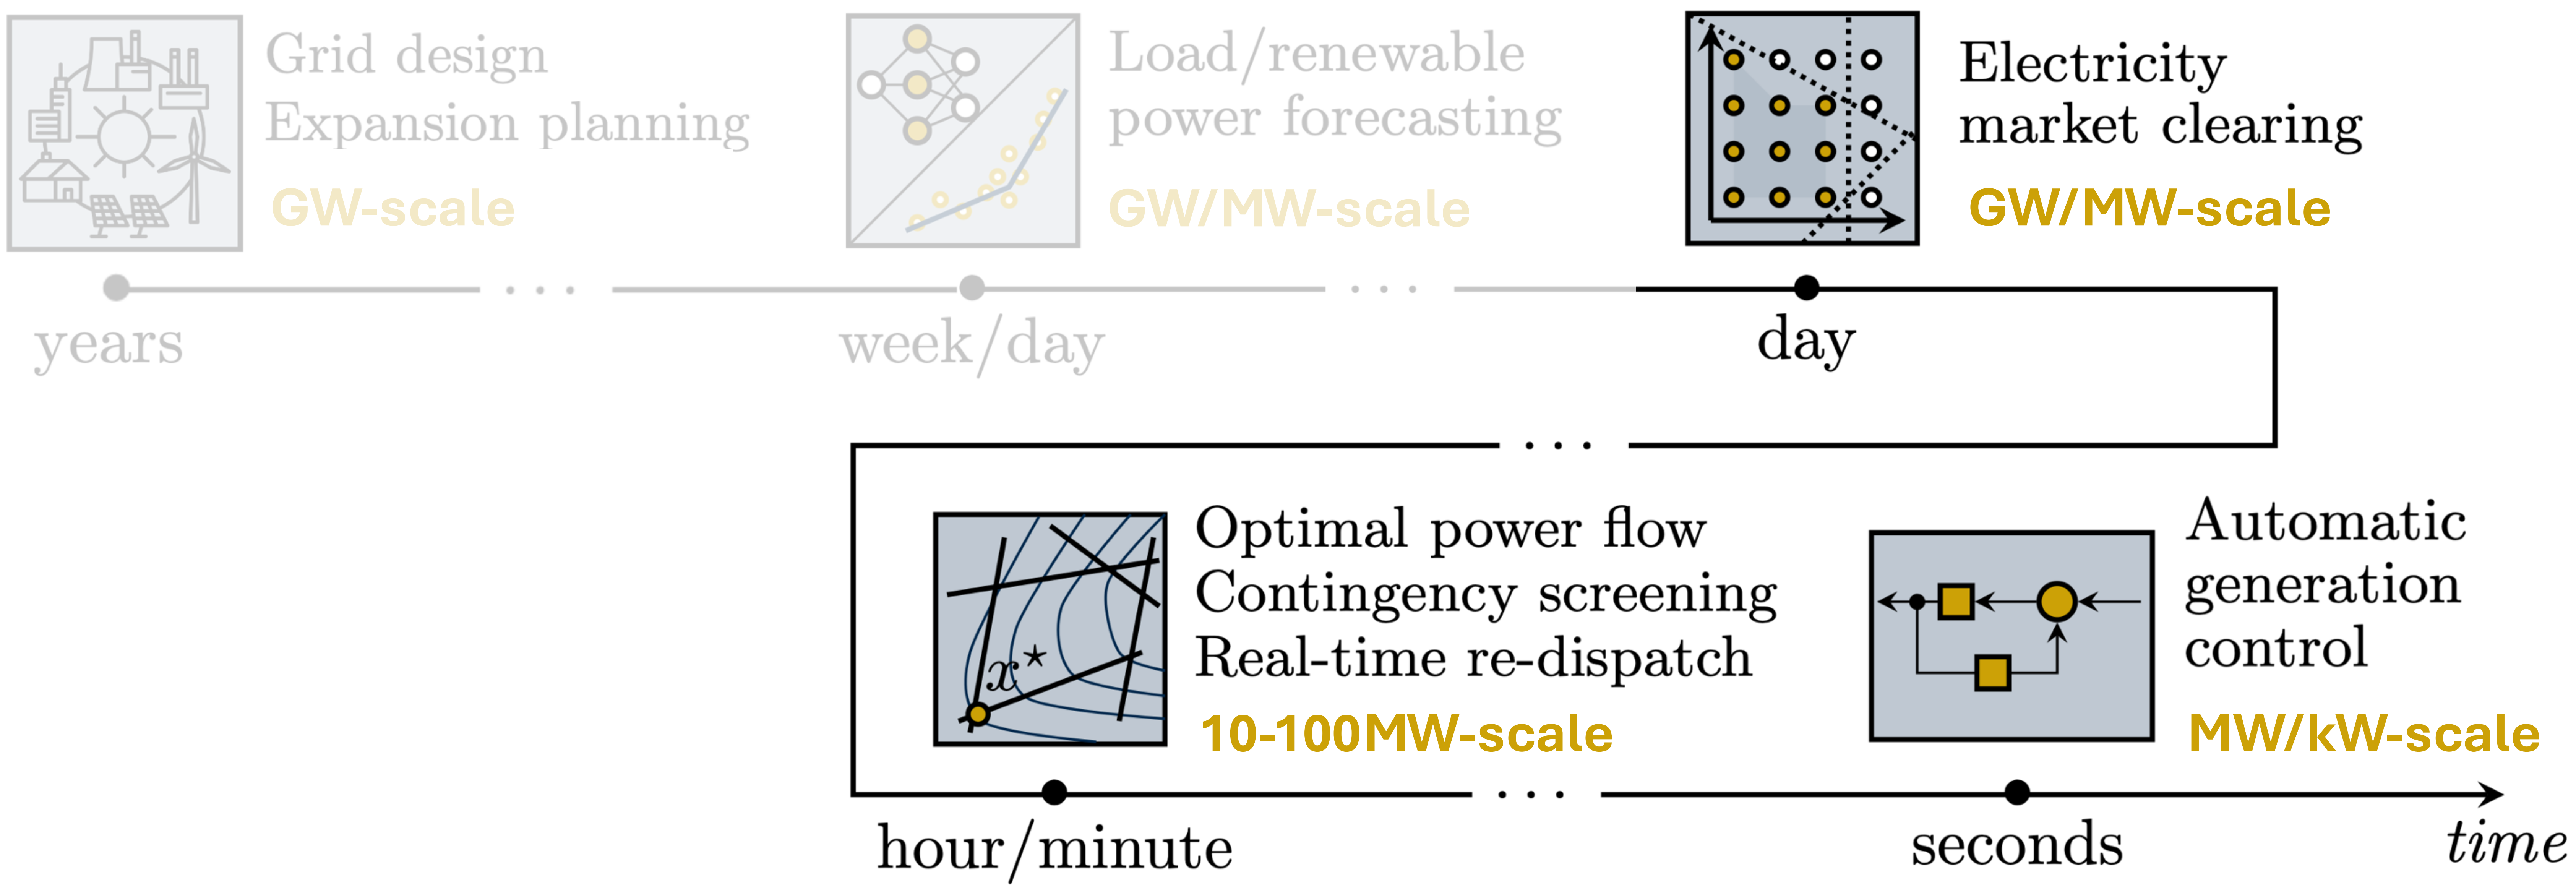

Diminishing coord. benefits

- Optimal sizing, siting and timing of data centers in the power grid
- Load forecasting and operating reserve sizing
- Day-ahead operational planning and market-clearing
- Near real-time operational coordination
- Real-time control

Agenda

1. Hierarchical optimization models for coordinating spatial DC flexibility
2. Market-based incentives to invoke spatial DC flexibility
3. Distribution-level energy storage solutions to smooth DC loads

Hierarchical optimization models for coordinating spatial DC flexibility

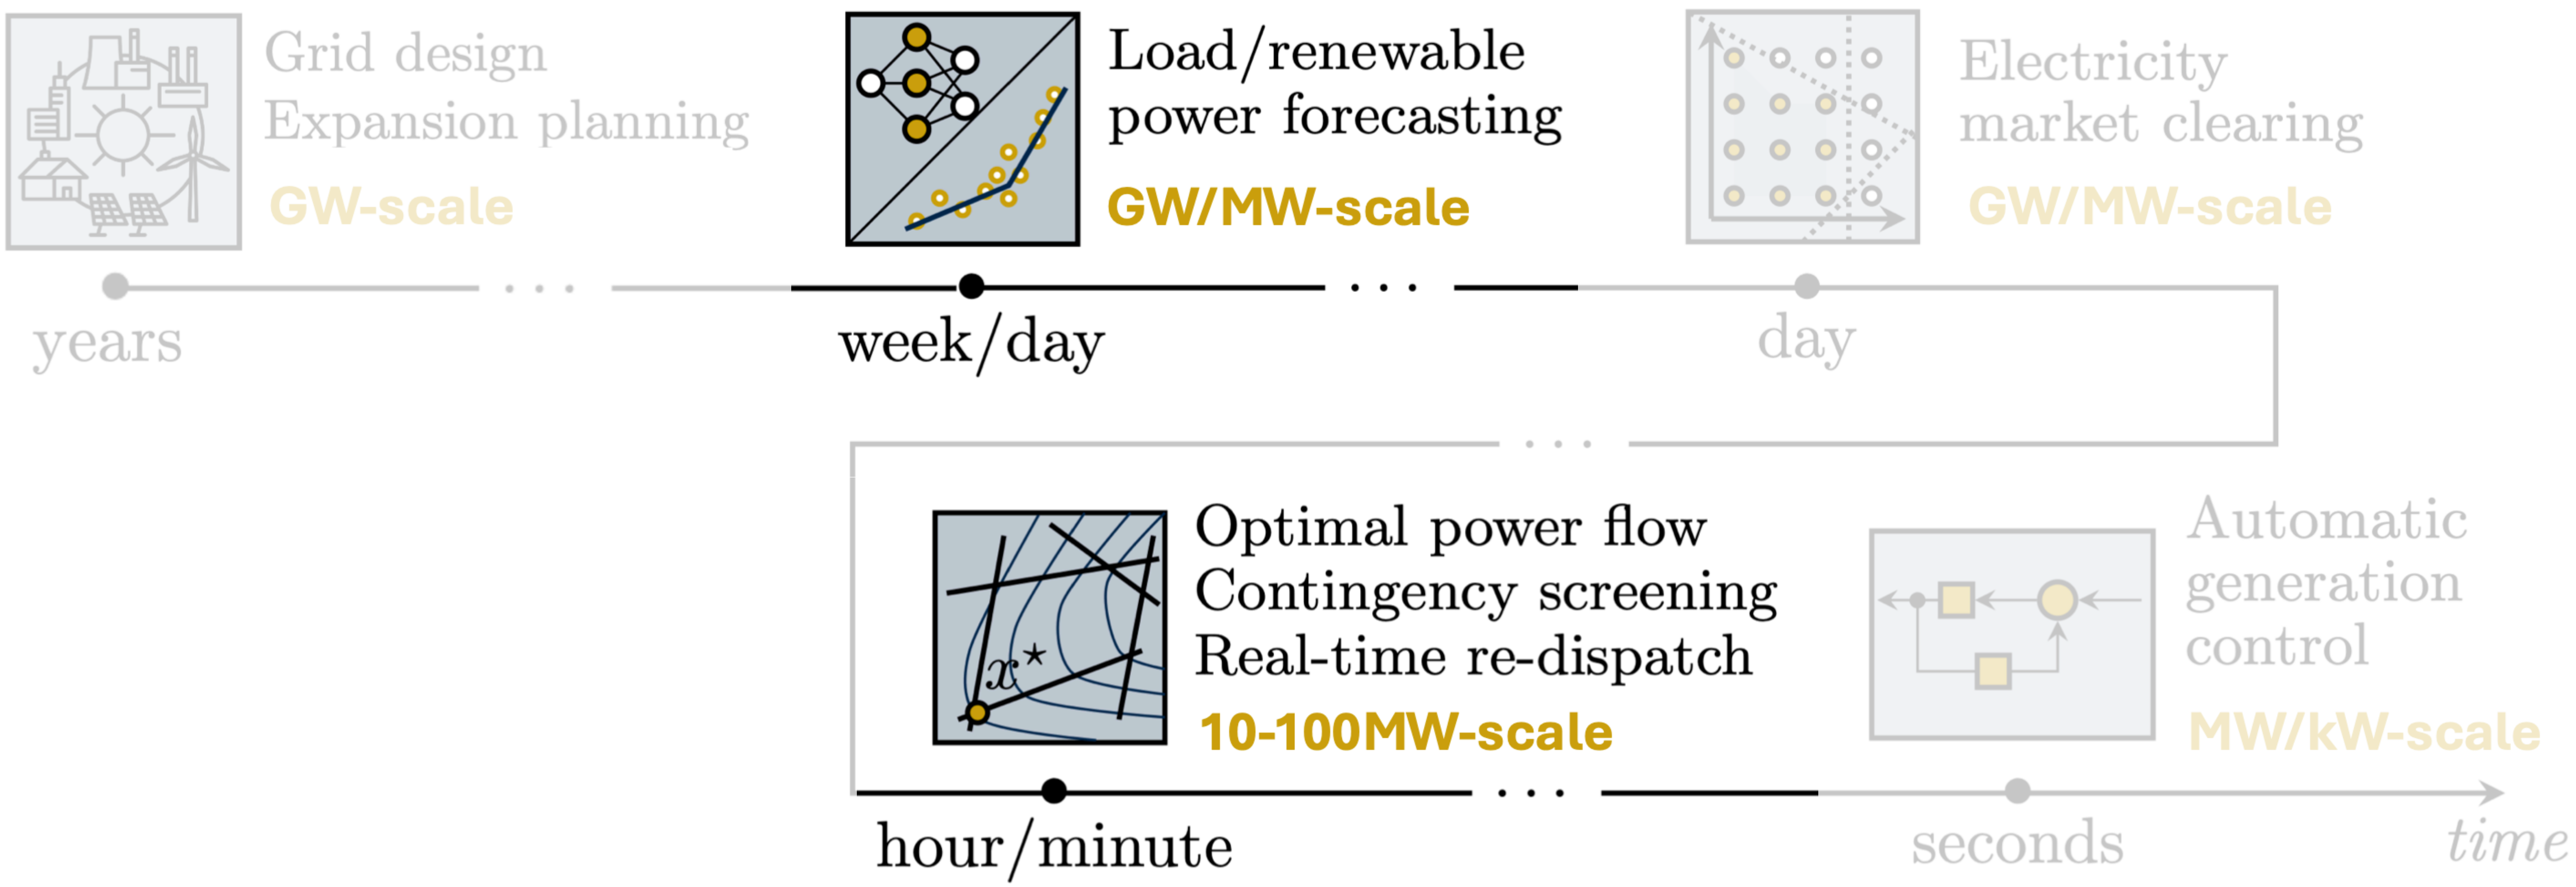

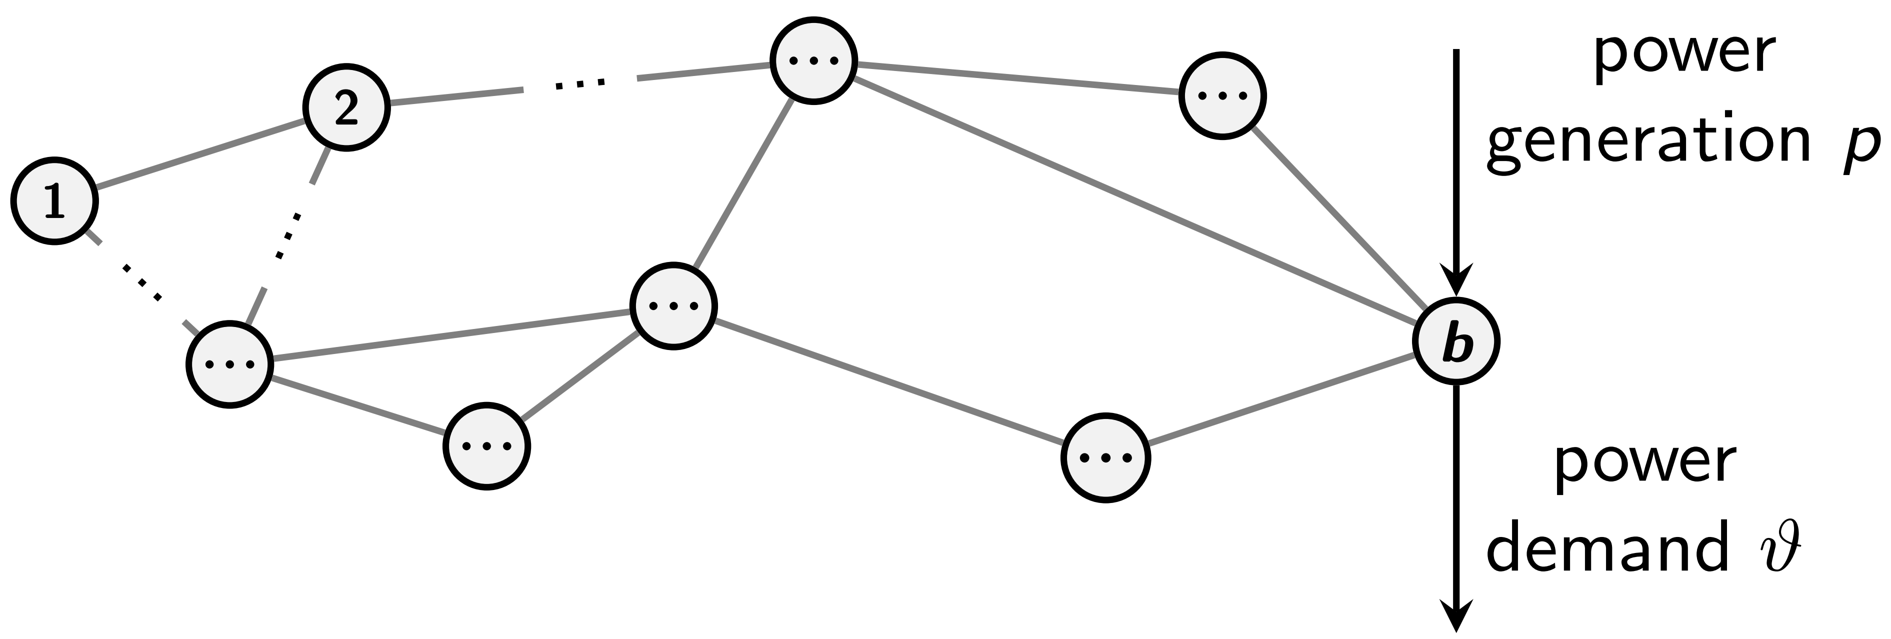

Power grid optimization problem:

$$\begin{aligned} \underset{p}{\text{minimize}} \quad & c_{\text{pwr}}(p) && \triangleright \texttt{ Dispatch cost} \\ \text{subject to} \quad & p \in \mathcal{P}_{\text{pwr}}(\vartheta) && \triangleright \texttt{ Grid feasibility} \end{aligned}$$

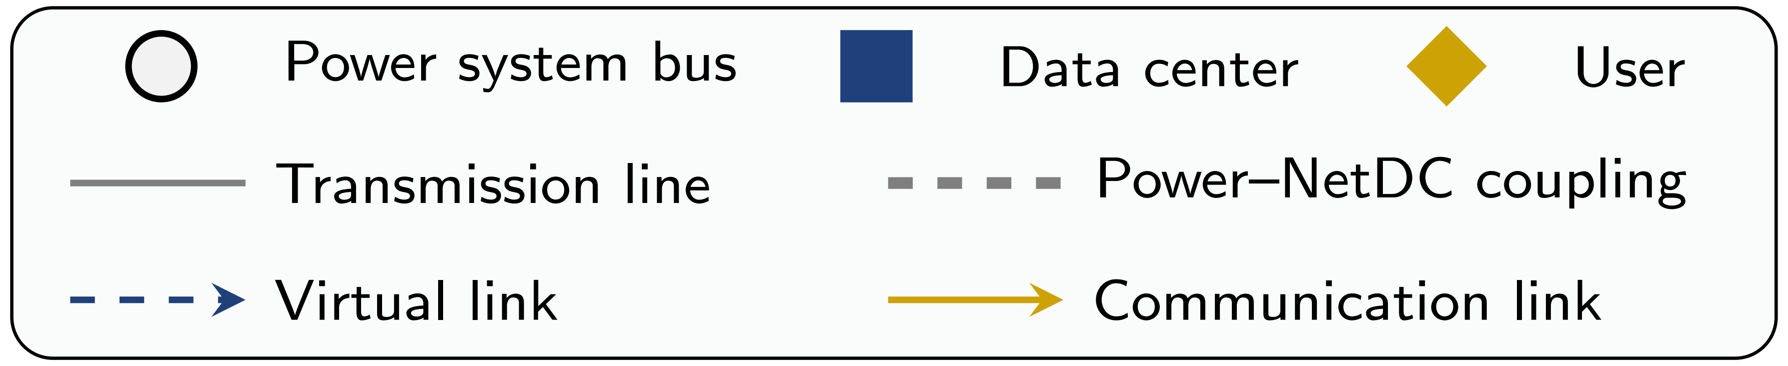

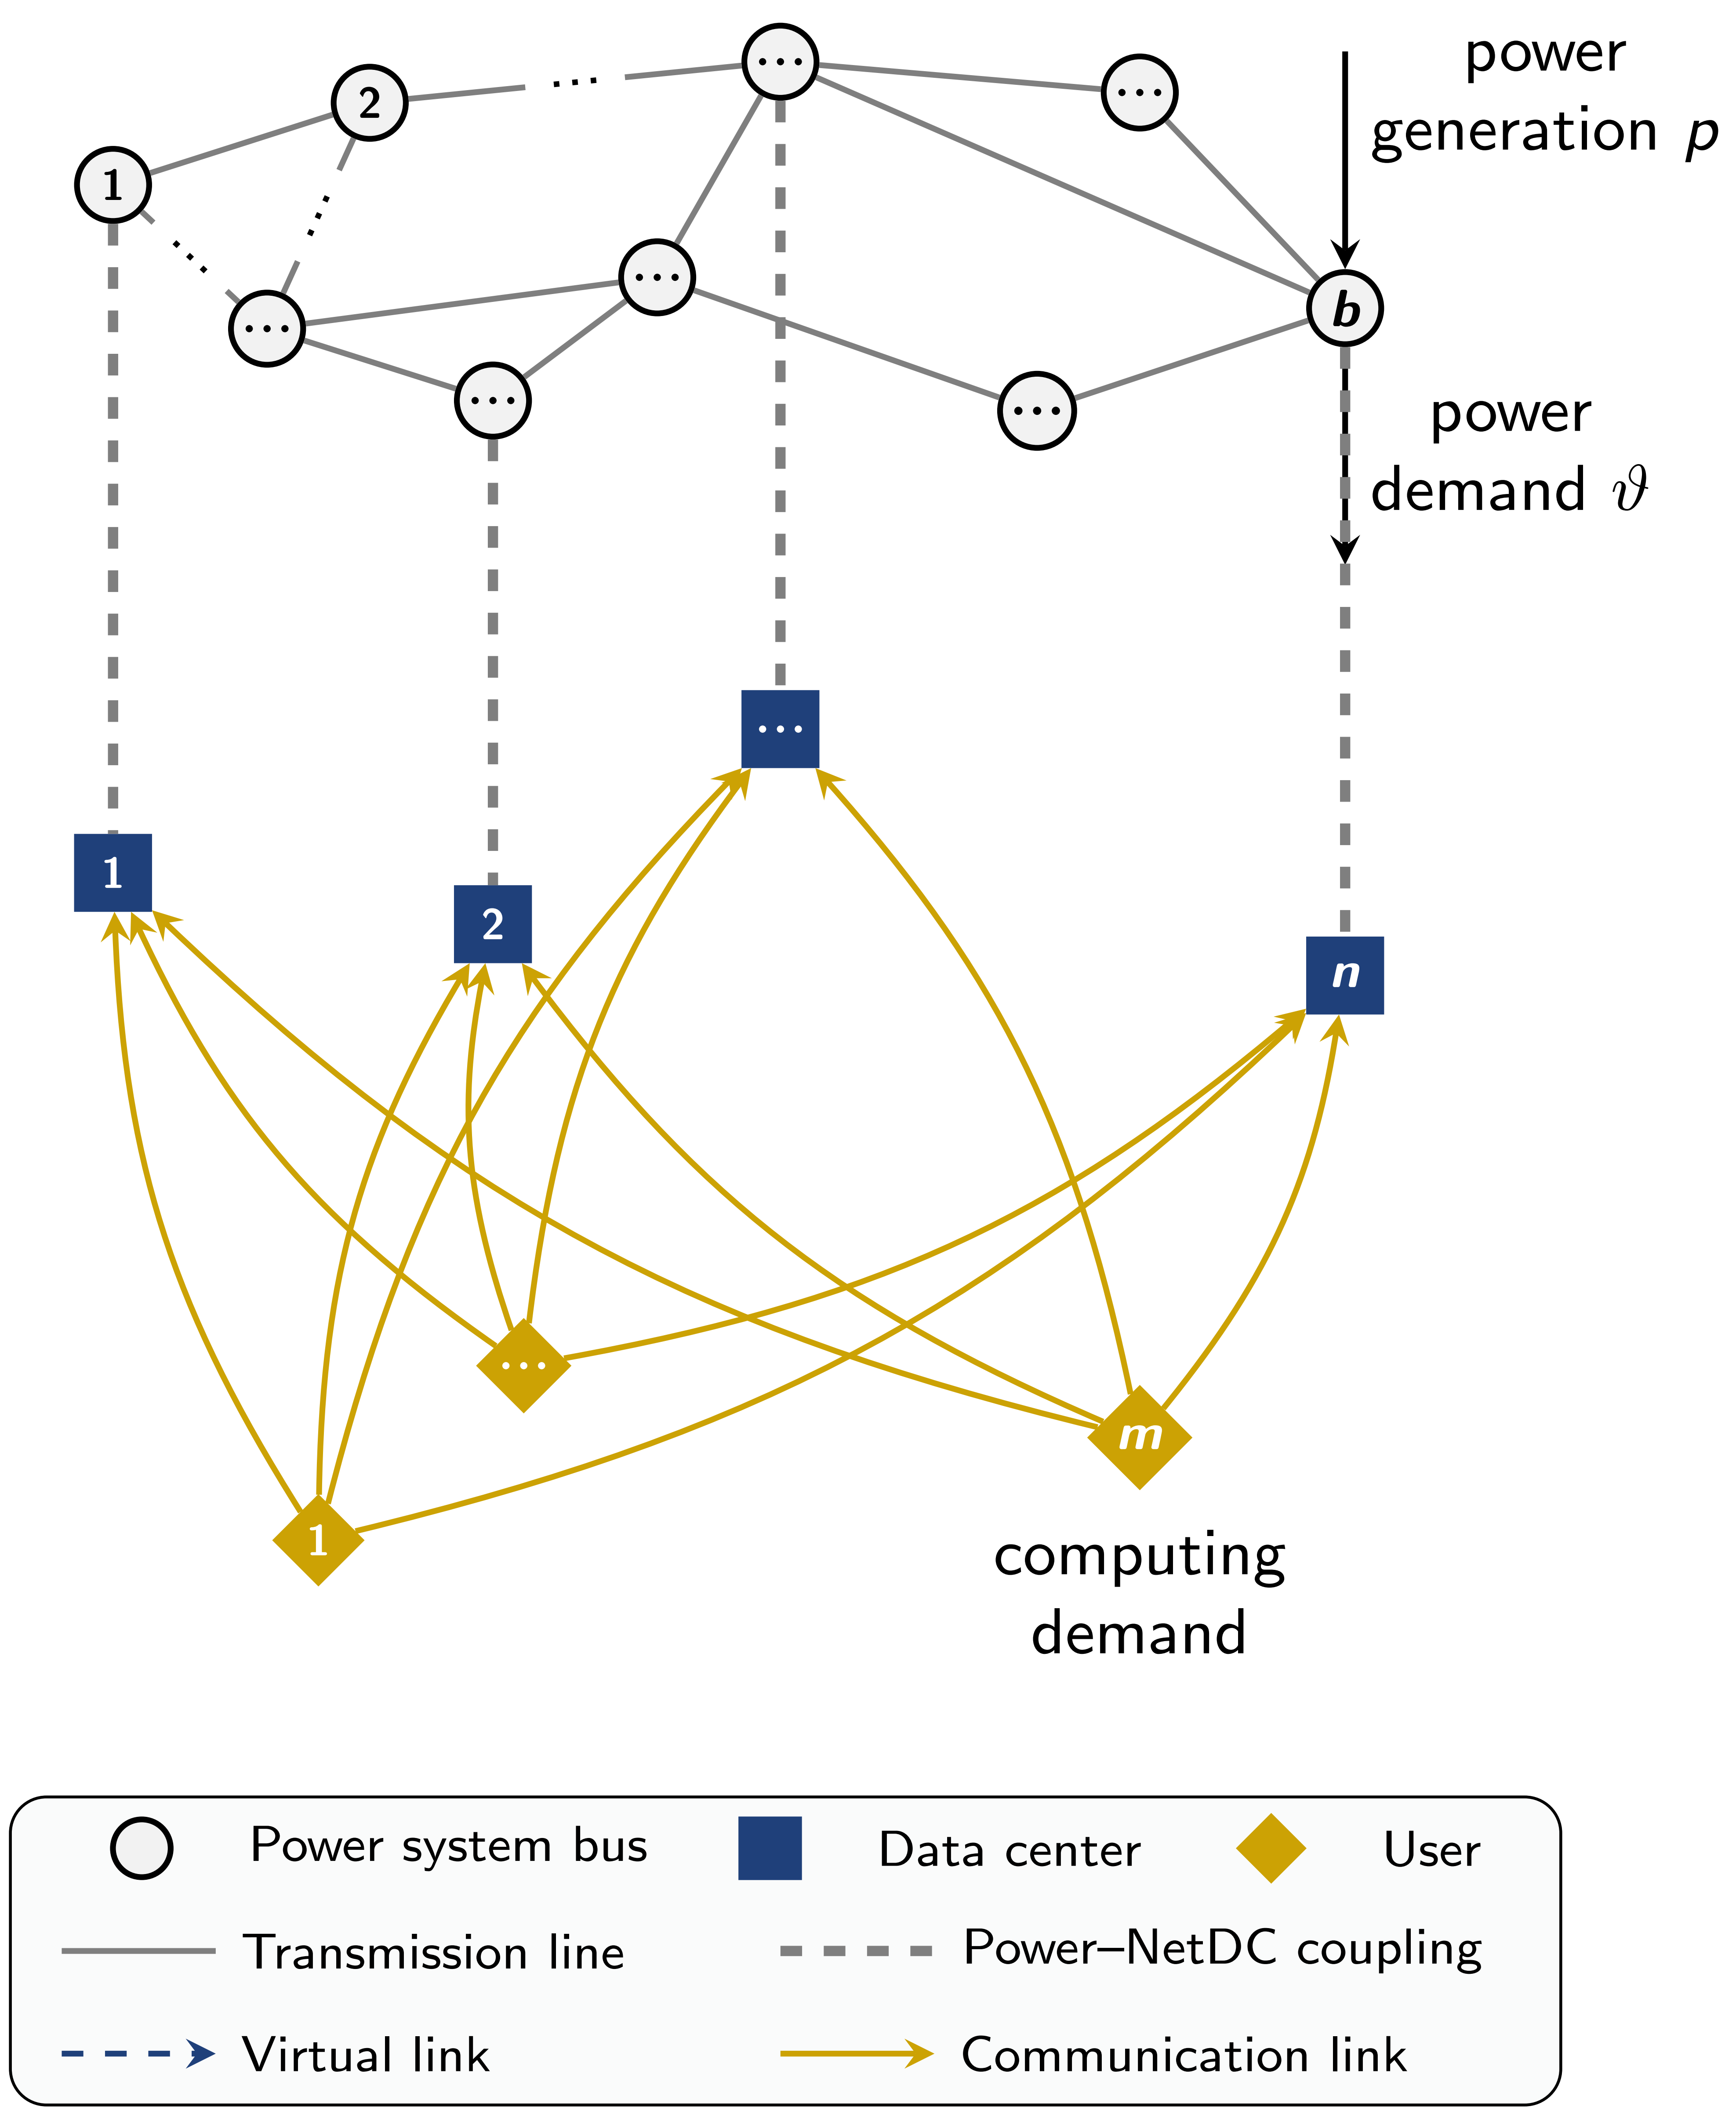

Power grid optimization problem:

$$\underset{p}{\text{minimize}} \quad c_{\text{pwr}}(p) \qquad \triangleright \texttt{ Dispatch cost}$$

$$\text{subject to} \quad p \in \mathcal{P}_{\text{pwr}}(\vartheta) \qquad \triangleright \texttt{ Grid feasibility}$$

Data centers optimization problem:

$$\underset{\vartheta}{\text{minimize}} \quad c_{\text{net-dc}}(\vartheta) \qquad \triangleright \texttt{ Latency loss}$$

$$\text{subject to} \quad \vartheta \in \mathcal{W}_{\text{net-dc}}(\varphi) \qquad \triangleright \texttt{ NetDC feasibility}$$

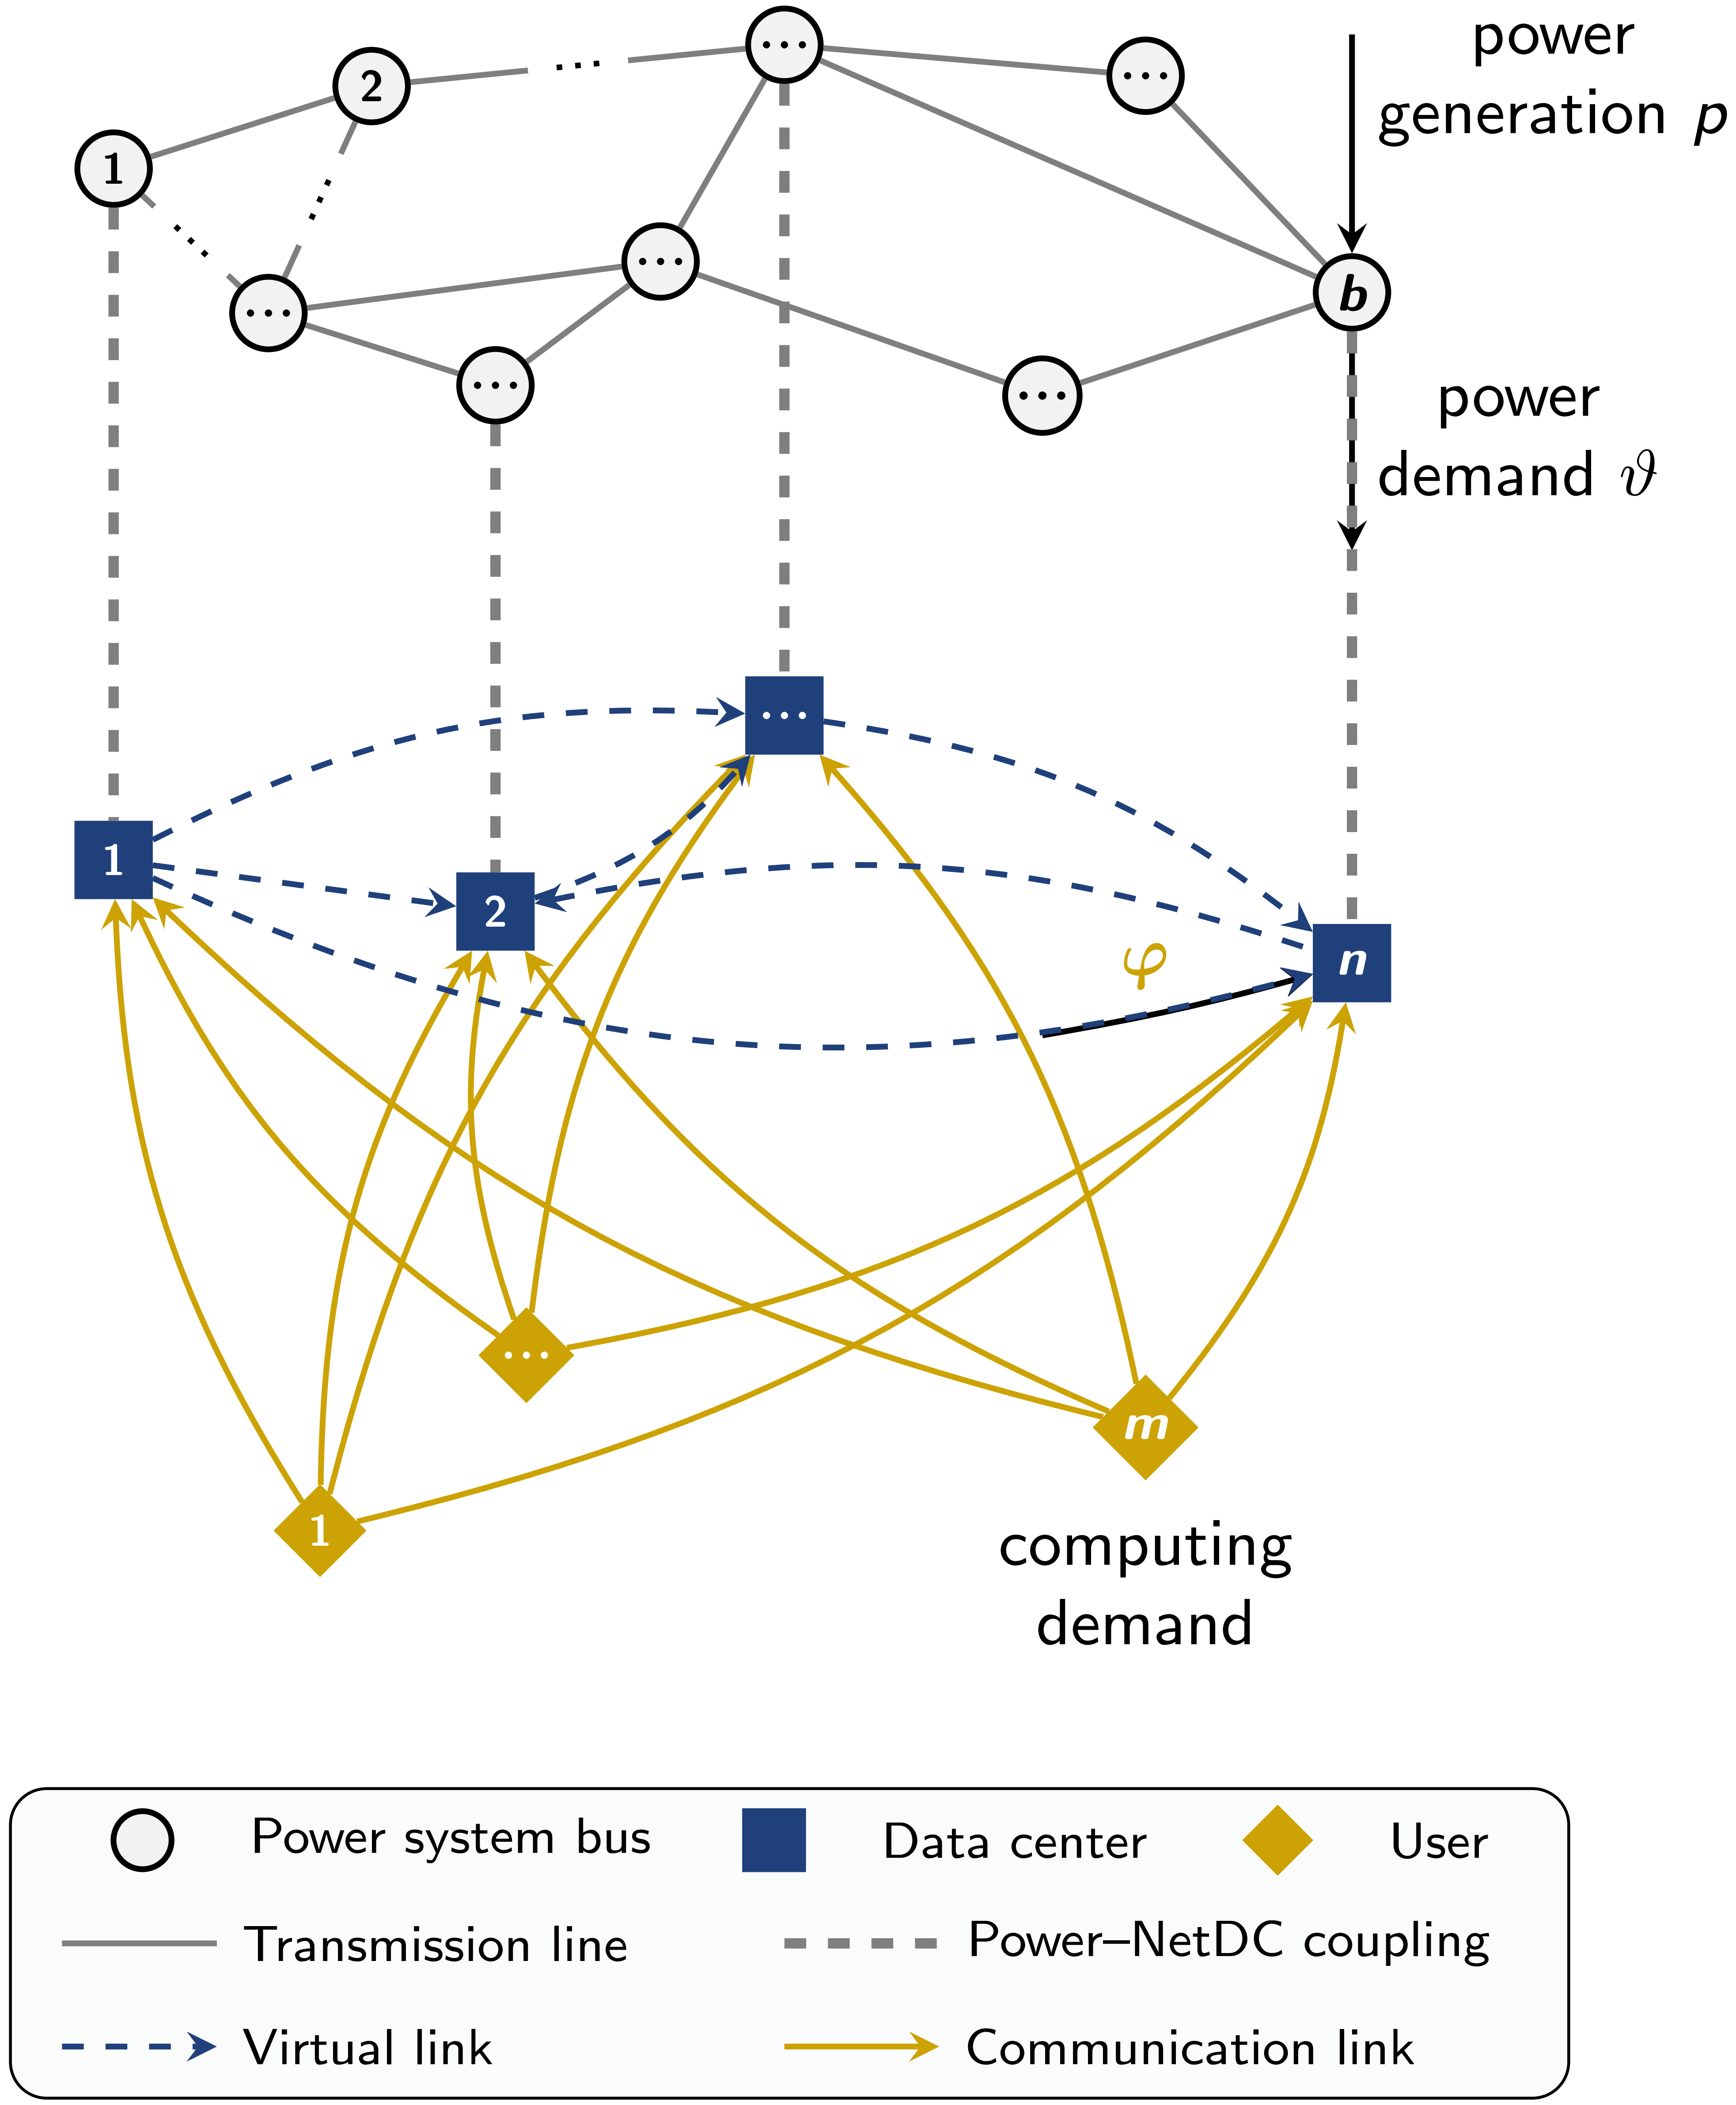

Power grid optimization problem:

$$\begin{aligned} \underset{p}{\text{minimize}} \quad & c_{\text{pwr}}(p) && \triangleright \texttt{ Dispatch cost} \\ \text{subject to} \quad & p \in \mathcal{P}_{\text{pwr}}(\vartheta) && \triangleright \texttt{ Grid feasibility} \end{aligned}$$

Data centers optimization problem:

$$\begin{aligned} \underset{\vartheta}{\text{minimize}} \quad & c_{\text{net-dc}}(\vartheta) && \triangleright \texttt{ Latency loss} \\ \text{subject to} \quad & \vartheta \in \mathcal{W}_{\text{net-dc}}(\varphi) && \triangleright \texttt{ NetDC feasibility} \end{aligned}$$

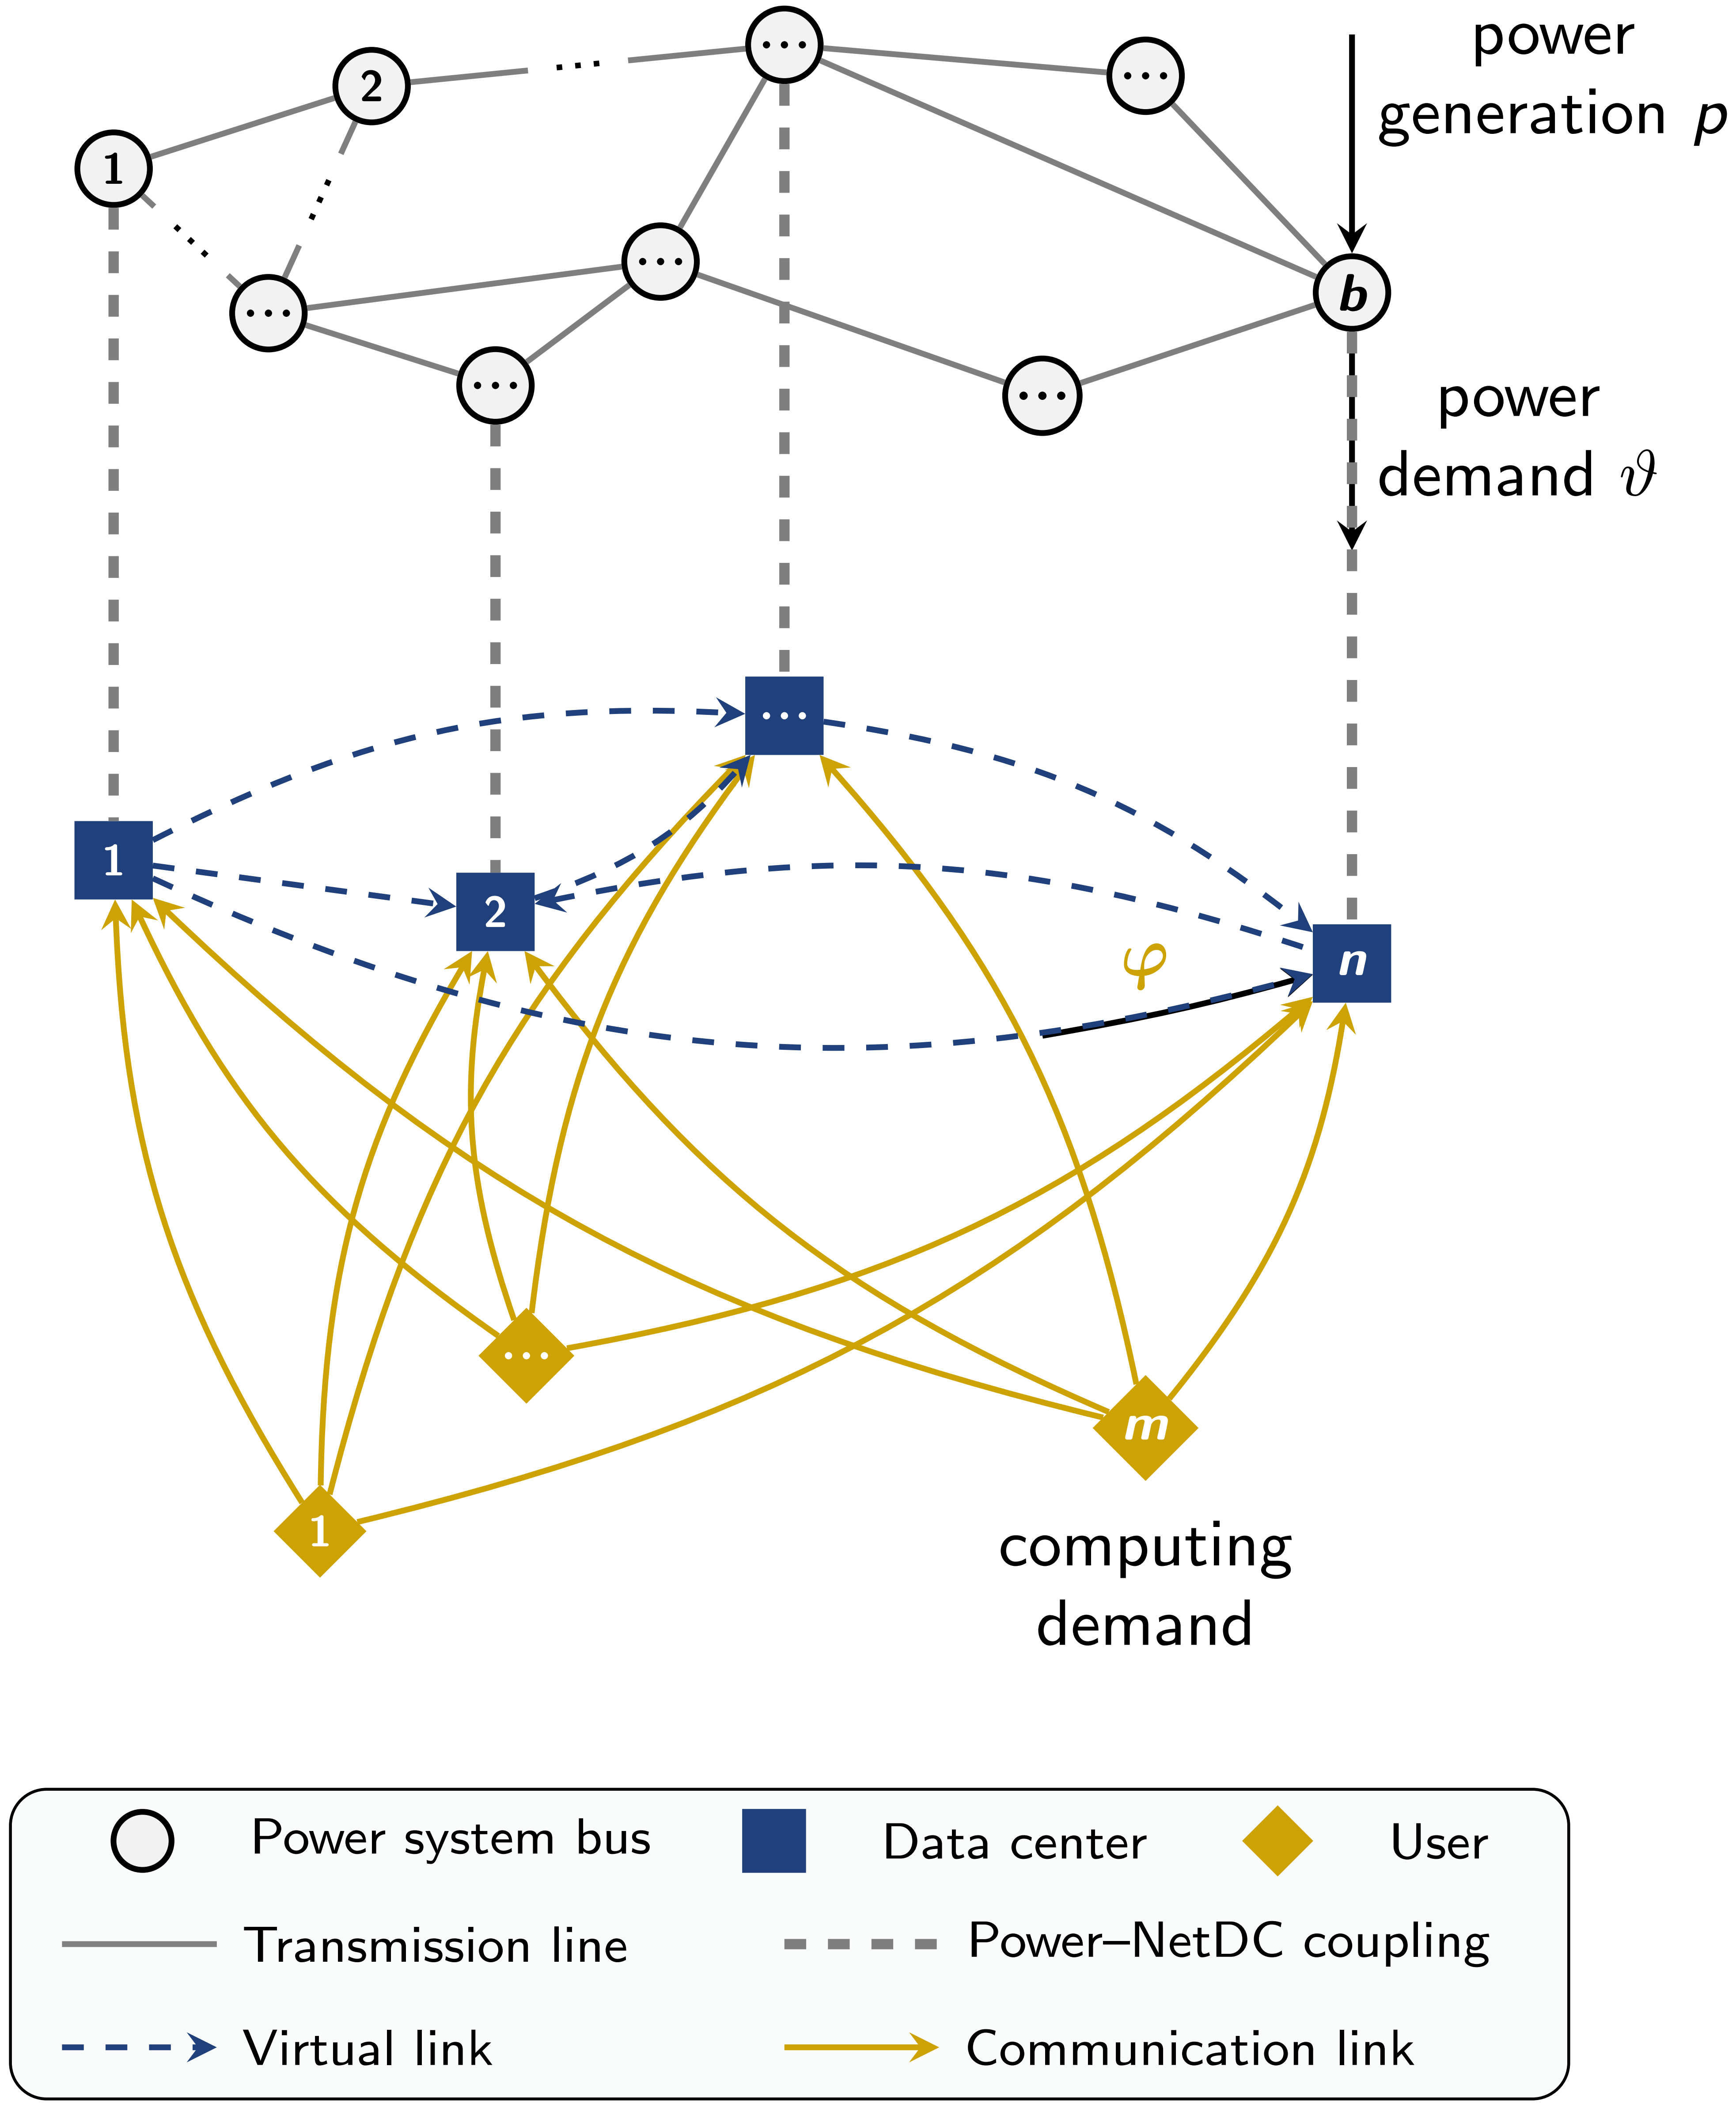

Power Grid–NetDC coordination problem:

$$
\begin{aligned}
\underset{\varphi, p}{\text{minimize}} \quad & c_{\text{pwr}}(p) && \triangleright \text{ Dispatch cost}\\
\text{subject to} \quad & p \in \mathcal{P}_{\text{pwr}}(\vartheta^\star) \xleftarrow{\text{el. demand}} && \triangleright \text{ Grid feasibility}\\
& \vartheta^\star \in \underset{\vartheta}{\text{argmin}} \quad c_{\text{net-dc}}(\vartheta) && \triangleright \text{ Latency loss}\\
& \qquad \text{subject to} \quad \vartheta \in \mathcal{W}_{\text{net-dc}}(\varphi) && \triangleright \text{ NetDC feasibility}
\end{aligned}
$$

computing task shift

- Grid optimization is constrained by NetDC optimization
- Every task shift request φ receives demand feedback ϑ
- The optimal $\varphi^\star$ minimizes the dispatch cost and satisfies both power grid and data center operational constraints

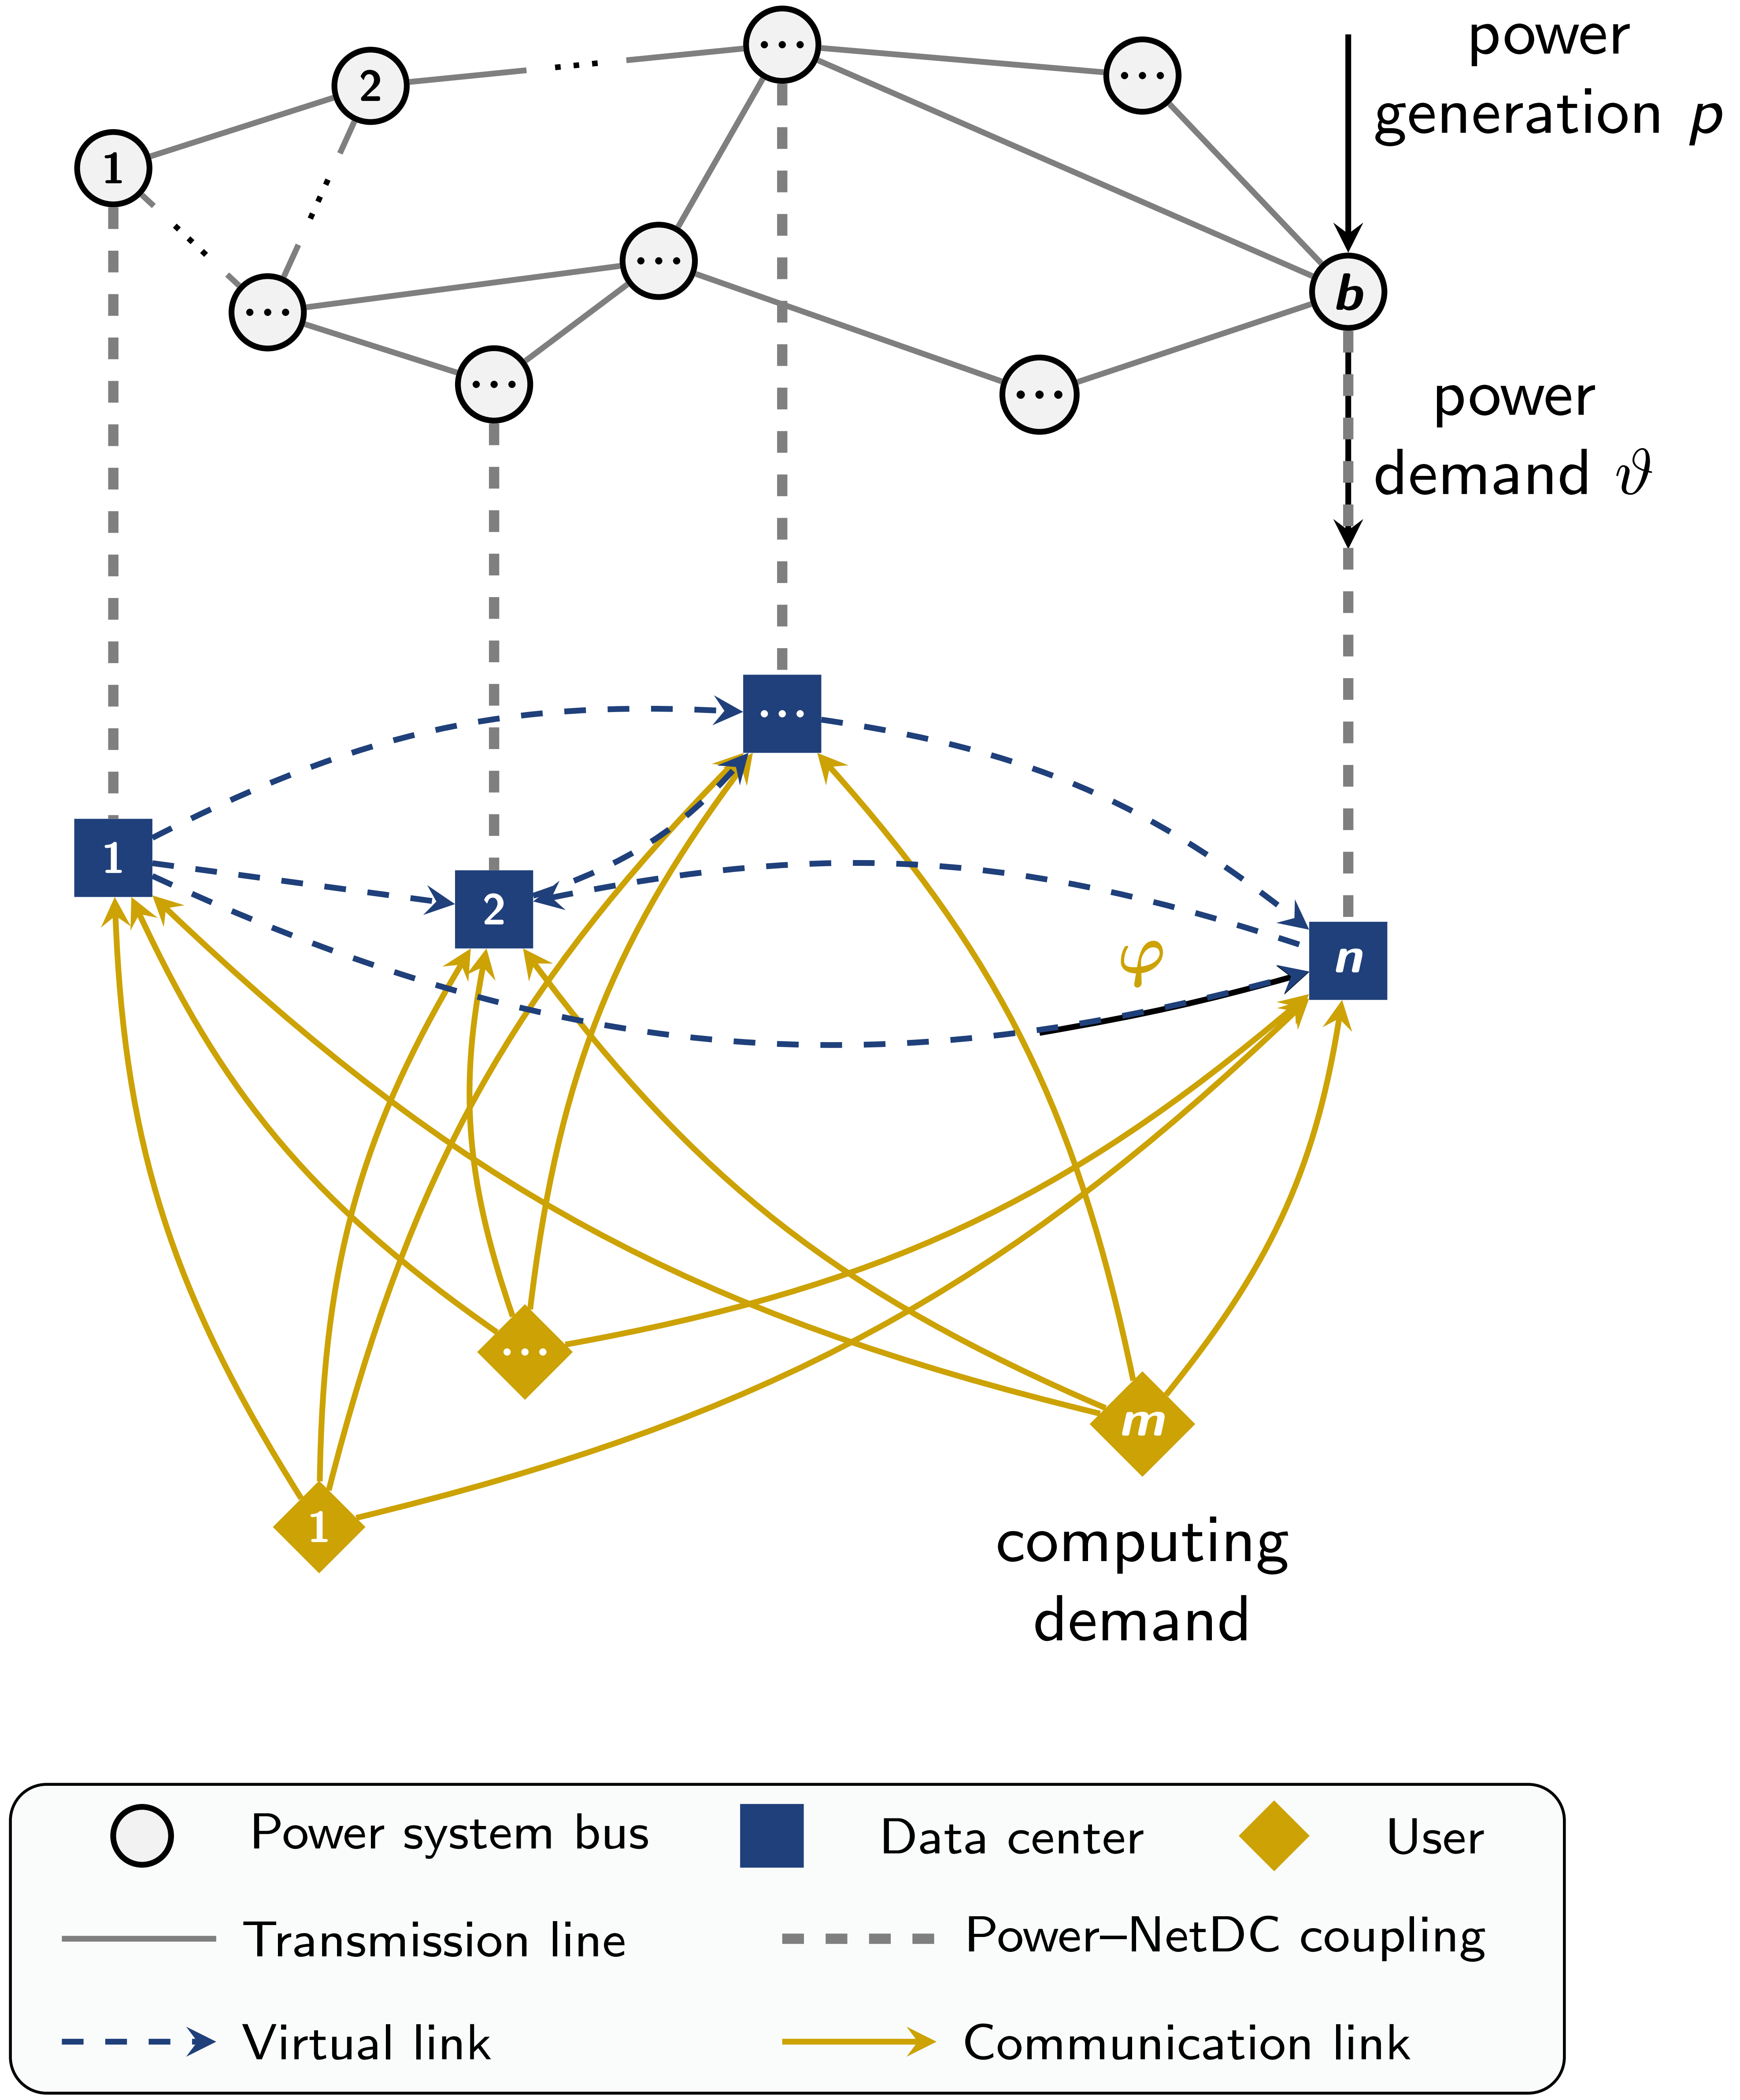

Power Grid–NetDC coordination problem:

$$
\begin{aligned}
\underset{\varphi, p}{\text{minimize}} \quad & c_{\text{pwr}}(p) && \triangleright \texttt{ Dispatch cost}\\
\text{subject to} \quad & p \in \mathcal{P}_{\text{pwr}}(\vartheta^{\star}) && \triangleright \texttt{ Grid feasibility}\\
& \vartheta^{\star} \in \underset{\vartheta}{\text{argmin}} \quad c_{\text{net-dc}}(\vartheta) && \triangleright \texttt{ Latency loss}\\
& \qquad \text{subject to} \quad \vartheta \in \mathcal{W}_{\text{net-dc}}(\varphi) && \triangleright \texttt{ NetDC feasibility}
\end{aligned}
$$

el. demand

computing task shift

$$\vartheta^{\star} \in \text{KKTs}(\varphi)$$

- Grid optimization is constrained by NetDC optimization
- Every task shift request φ receives demand feedback ϑ
- The optimal $\varphi^{\star}$ minimizes the dispatch cost and satisfies both power grid and data center operational constraints

From optimization to regression for Power-NetDC coordination

- Consider a set of features $\{x_1, \ldots, x_q\}$, each specific to a particular coordination scenario
- Contextual feature $x_i = \{$nodal prices, loads, generation statistics, ...$\}$
- **Goal:** train the coordination policy to map contextual features into the optimal task shifts

$$\phi(x) = \beta_0 + \beta_1 x$$

where $\beta = (\beta_0, \beta_1)$ are regression weights

From optimization to regression for Power-NetDC coordination

- Consider a set of features $\{x_1, \ldots, x_q\}$, each specific to a particular coordination scenario
- Contextual feature $x_i = \{\text{nodal prices, loads, generation statistics, ...}\}$
- **Goal:** train the coordination policy to map contextual features into the optimal task shifts

$$\phi(x) = \beta_0 + \beta_1 x$$

where $\beta = (\beta_0, \beta_1)$ are regression weights

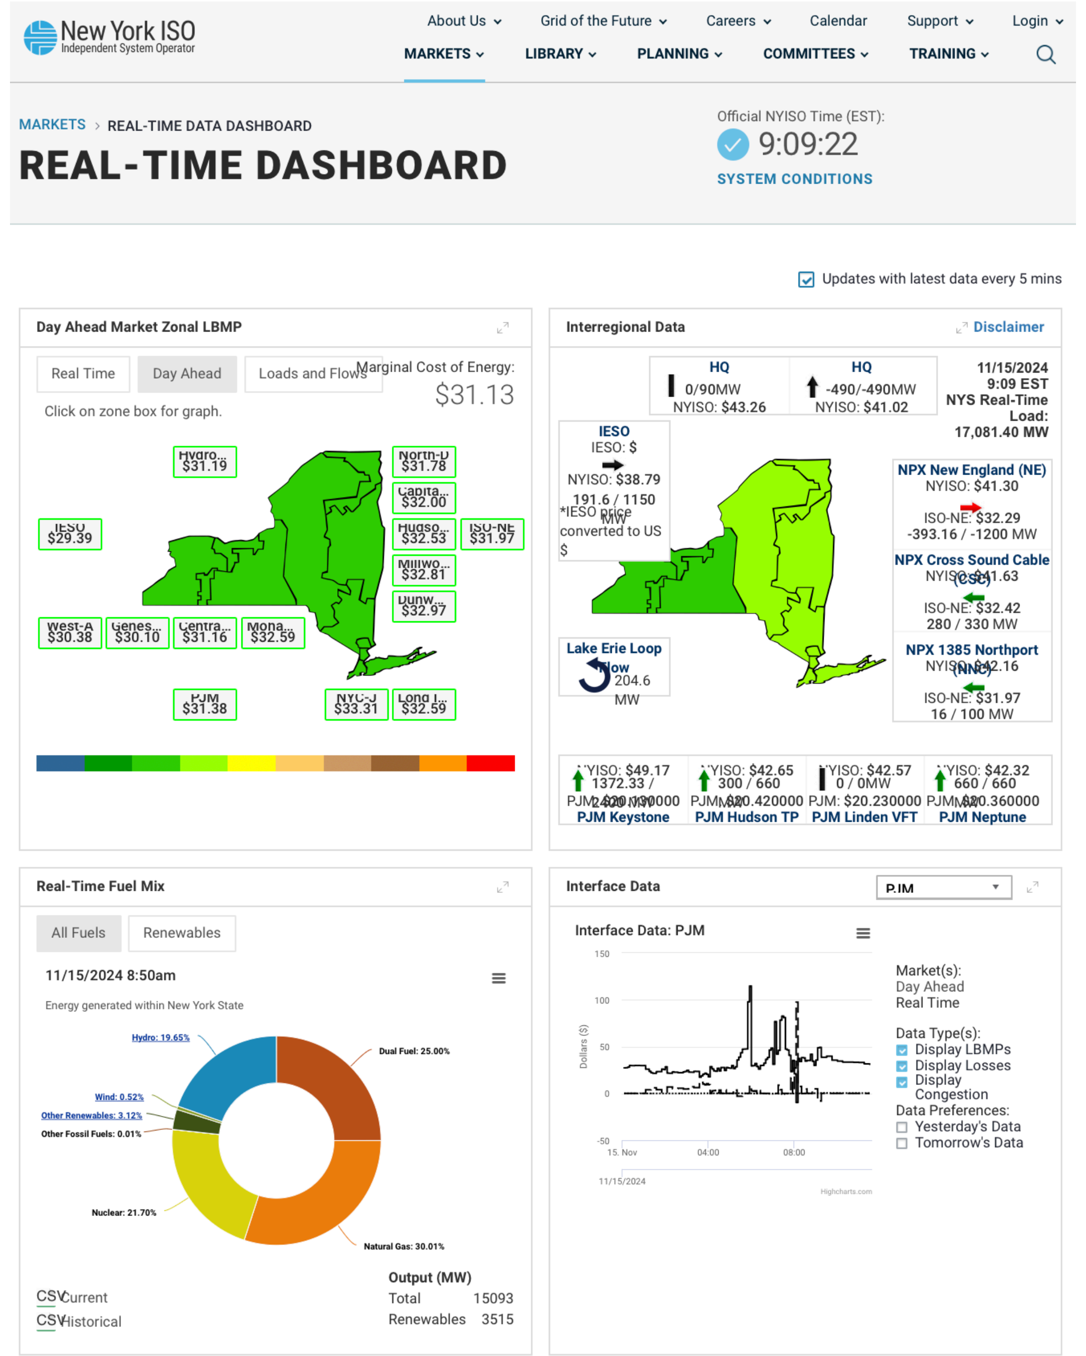

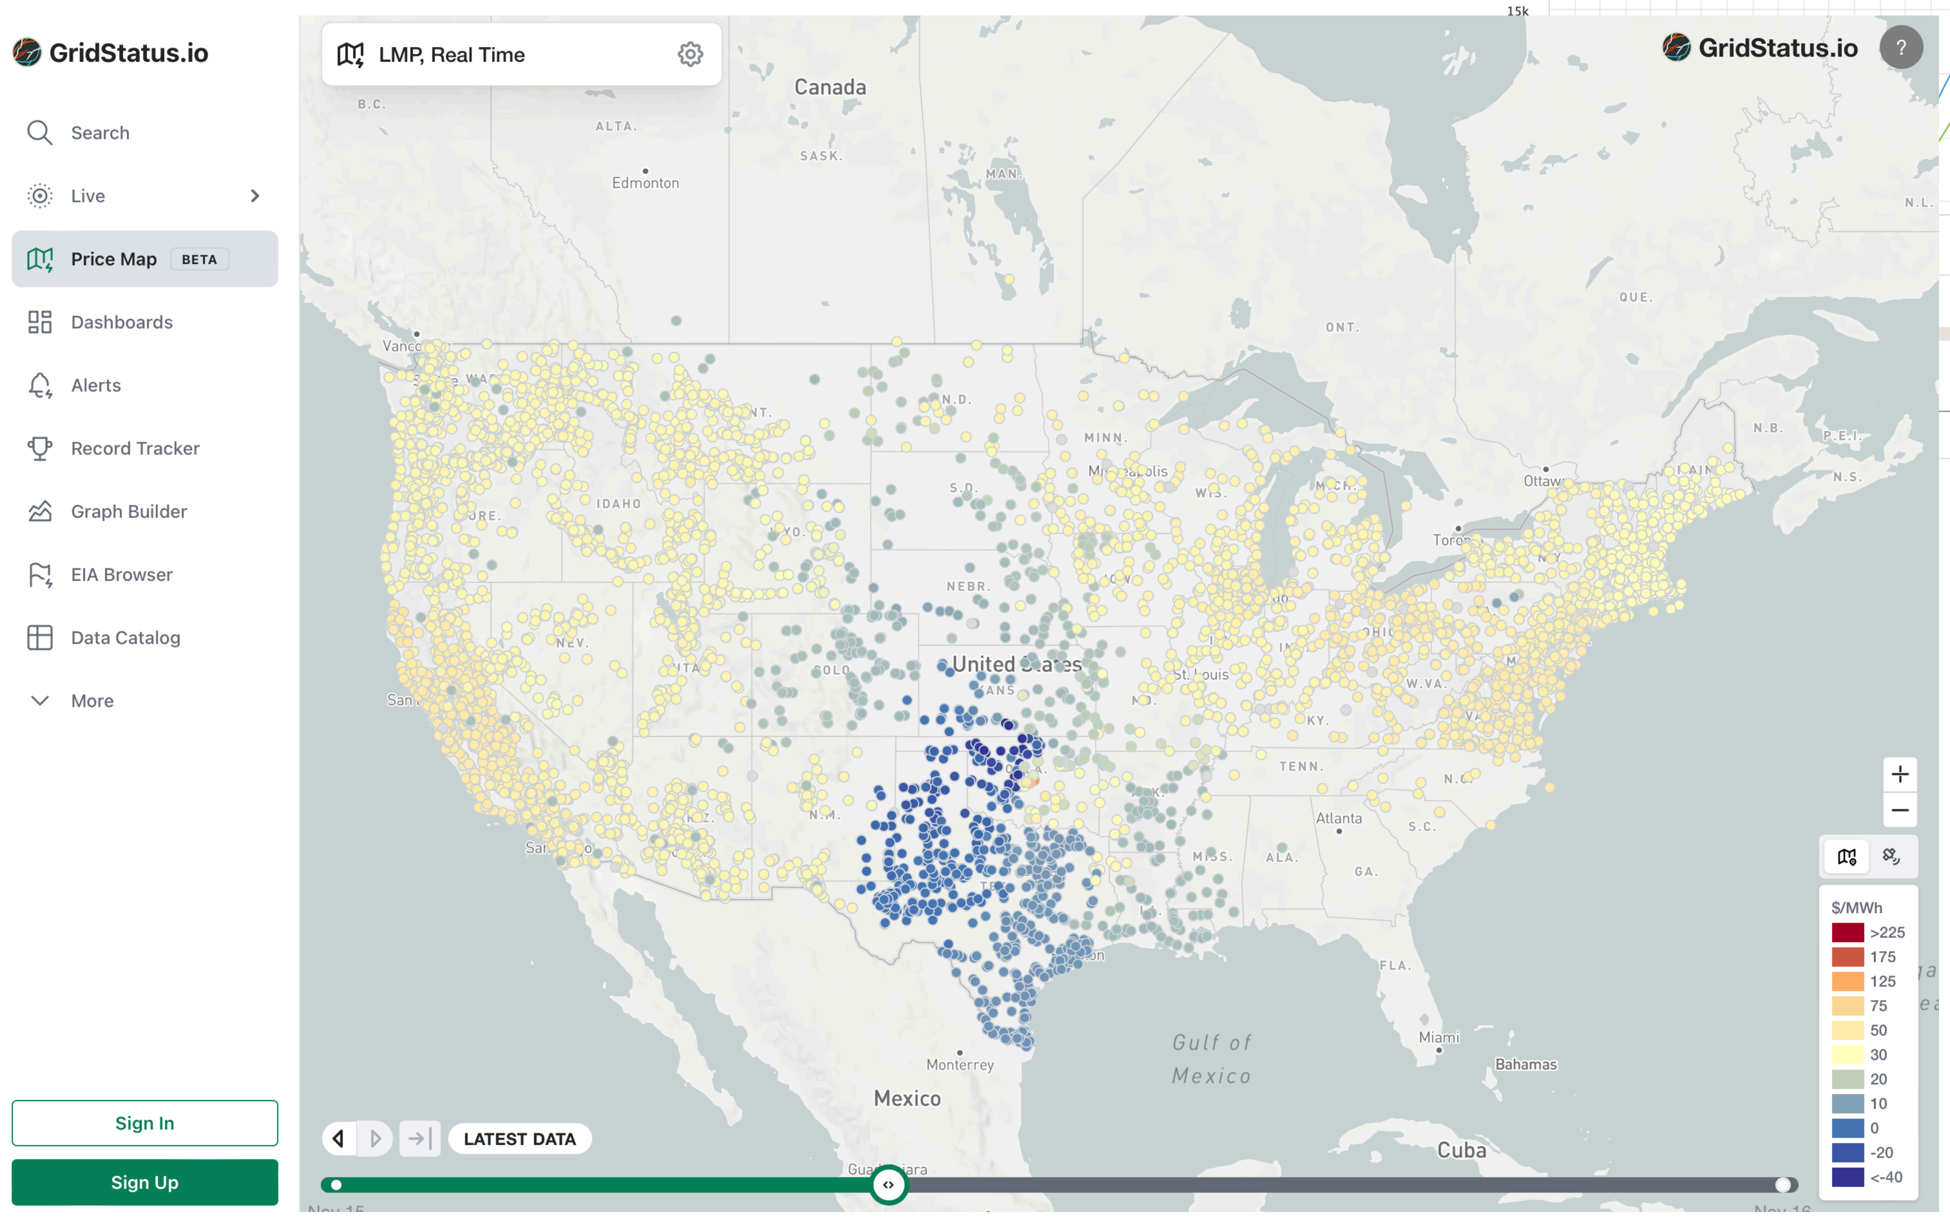

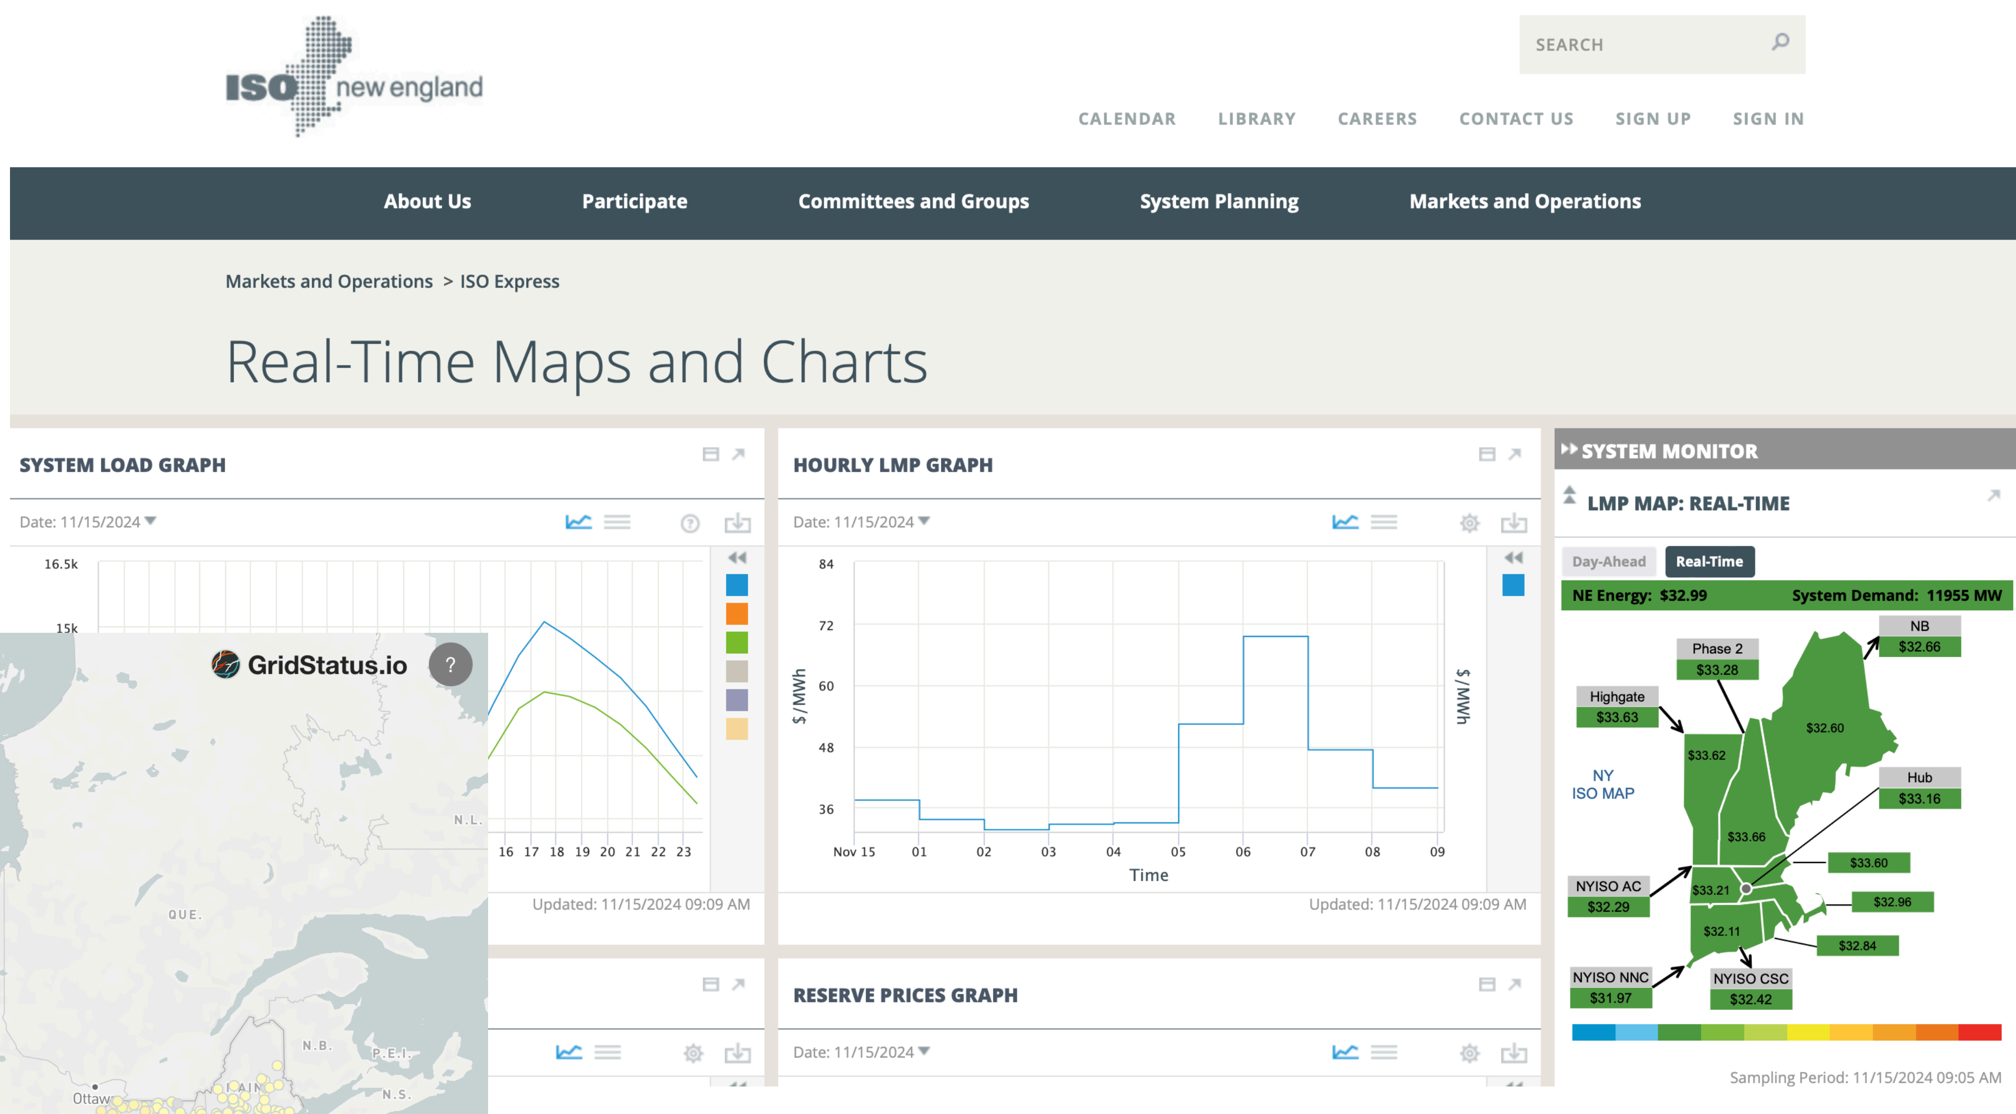

- Consider a set of features $\{x_1, \ldots, x_q\}$, each specific to a particular coordination scenario
- Contextual feature $x_i = \{$nodal prices, loads, generation statistics, ...$\}$
- **Goal:** train the coordination policy to map contextual features into the optimal task shifts

$$\phi(x) = \beta_0 + \beta_1 x$$

where $\beta = (\beta_0, \beta_1)$ are regression weights

Grid- and NetDC-Informed policy optimization

$$\begin{aligned}
&\underset{\beta,\varphi_i,\ldots,\varphi_q,p_i,\ldots,p_q}{\text{minimize}} && \frac{1}{q}\sum_{i=1}^{q} c_{\text{pwr}}(p_i) && \triangleright \text{ Average dispatch cost}\\
&\text{subject to} && p_i \in \mathcal{P}_{\text{pwr}}(\vartheta_i^\star), \quad \forall i = 1, \ldots, q && \triangleright \text{ Grid equations for each scenario}\\
& && \varphi_i = \beta_0 + \beta_1 x_i,\ \|\beta\|_1 \leqslant \varepsilon, \quad \forall i = 1, \ldots, q && \triangleright \text{ Coupling contextual regression}\\
& && \vartheta_i^\star \in \underset{\vartheta_i}{\text{argmin}} \quad c_{\text{net-dc}}(\vartheta_i) && \triangleright \text{ Latency loss}\\
& && \quad \text{subject to} \quad \vartheta_i \in \mathcal{W}_{\text{net-dc}}(\varphi_i), \quad \forall i = 1, \ldots, q && \triangleright \text{ NetDC feasibility}
\end{aligned}$$

From optimization to regression for Power-NetDC coordination

- Consider a set of features $\{x_1, \ldots, x_q\}$, each specific to a particular coordination scenario
- Contextual feature $x_i = \{$nodal prices, loads, generation statistics, ...$\}$
- **Goal:** train the coordination policy to map contextual features into the optimal task shifts

$$\phi(x) = \beta_0 + \beta_1 x$$

where $\beta = (\beta_0, \beta_1)$ are regression weights

Grid- and NetDC-Informed policy optimization

$$
\begin{aligned}
\underset{\beta, \varphi_i, \ldots, \varphi_q, p_i, \ldots, p_q}{\text{minimize}} \quad & \frac{1}{q}\sum_{i=1}^{q} c_{\text{pwr}}(p_i) && \triangleright \texttt{ Average dispatch cost} \\
\text{subject to} \quad & p_i \in \mathcal{P}_{\text{pwr}}(\vartheta_i^\star), \quad \forall i = 1, \ldots, q && \triangleright \texttt{ Grid equations for each scenario} \\
& \varphi_i = \beta_0 + \beta_1 x_i, \; \|\beta\|_1 \leqslant \varepsilon, \quad \forall i = 1, \ldots, q && \triangleright \texttt{ Coupling contextual regression} \\
& \vartheta_i^\star \in \underset{\vartheta_i}{\text{argmin}} \quad c_{\text{net-dc}}(\vartheta_i) && \triangleright \texttt{ Latency loss} \\
& \qquad \text{subject to} \quad \vartheta_i \in \mathcal{W}_{\text{net-dc}}(\varphi_i), \quad \forall i = 1, \ldots, q && \triangleright \texttt{ NetDC feasibility}
\end{aligned}
$$

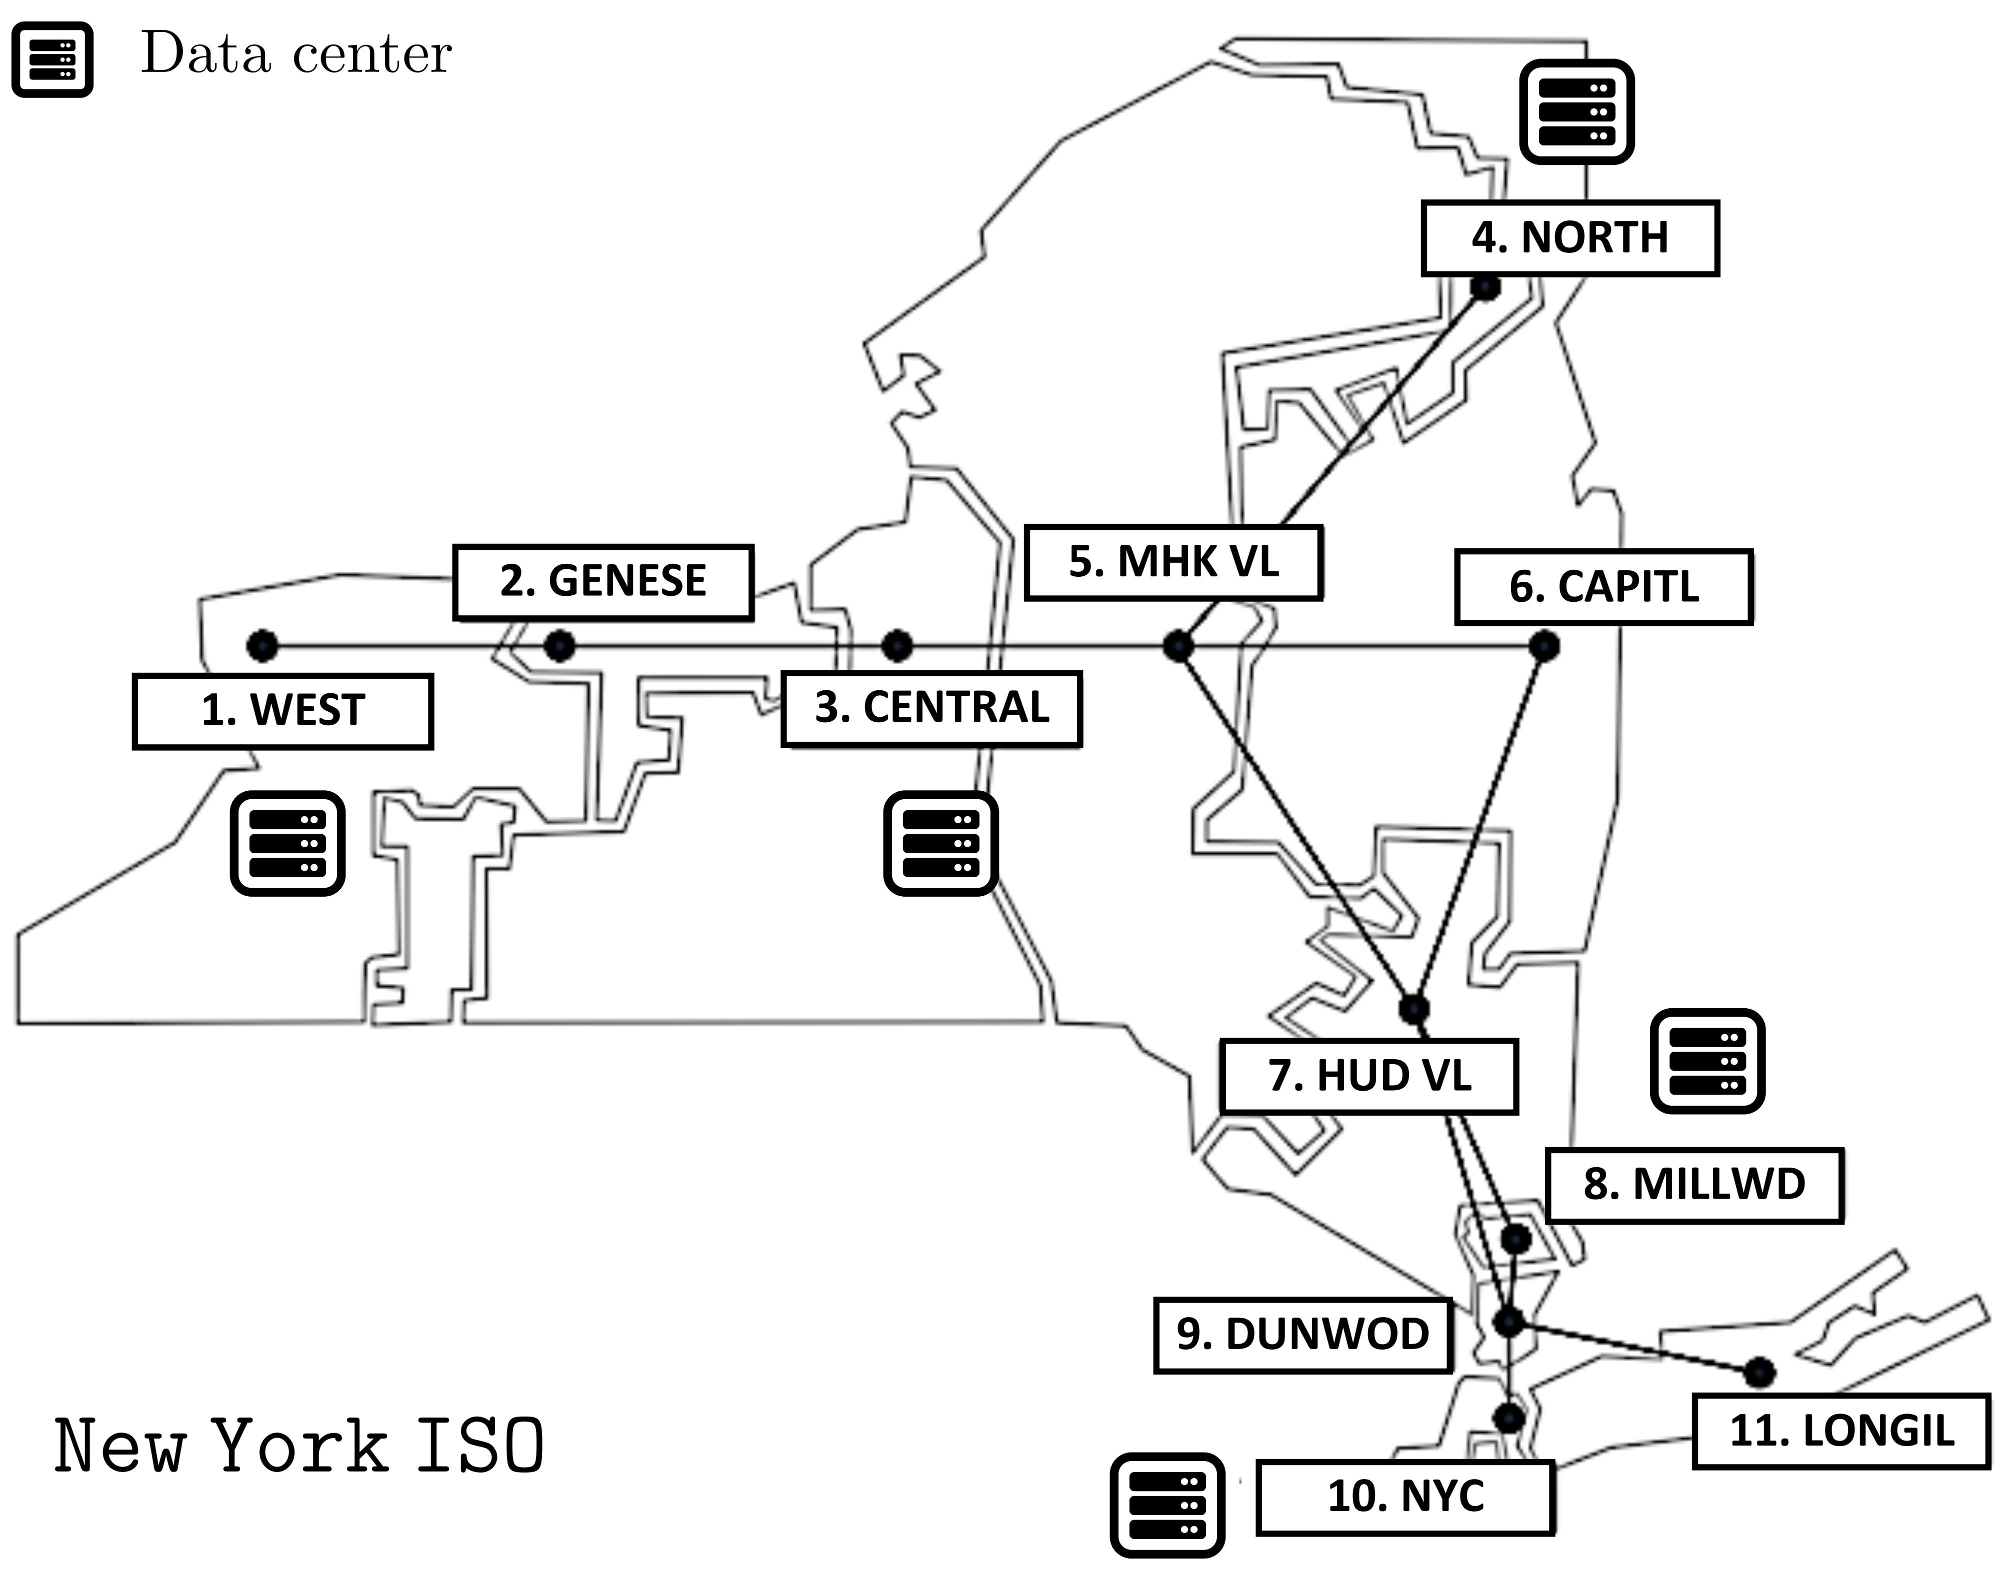

Data inputs:

- 11-zone aggregation of the New York ISO
- Network of 5 data centers (10 virtual links)
- Varying demand from 5% to 30% of the peak load

We study two coordination settings:

- Ideal day-ahead coordination with optimization
- Real-time coordination with contextual regression

NYISO: Ideal coordination at the day-ahead stage

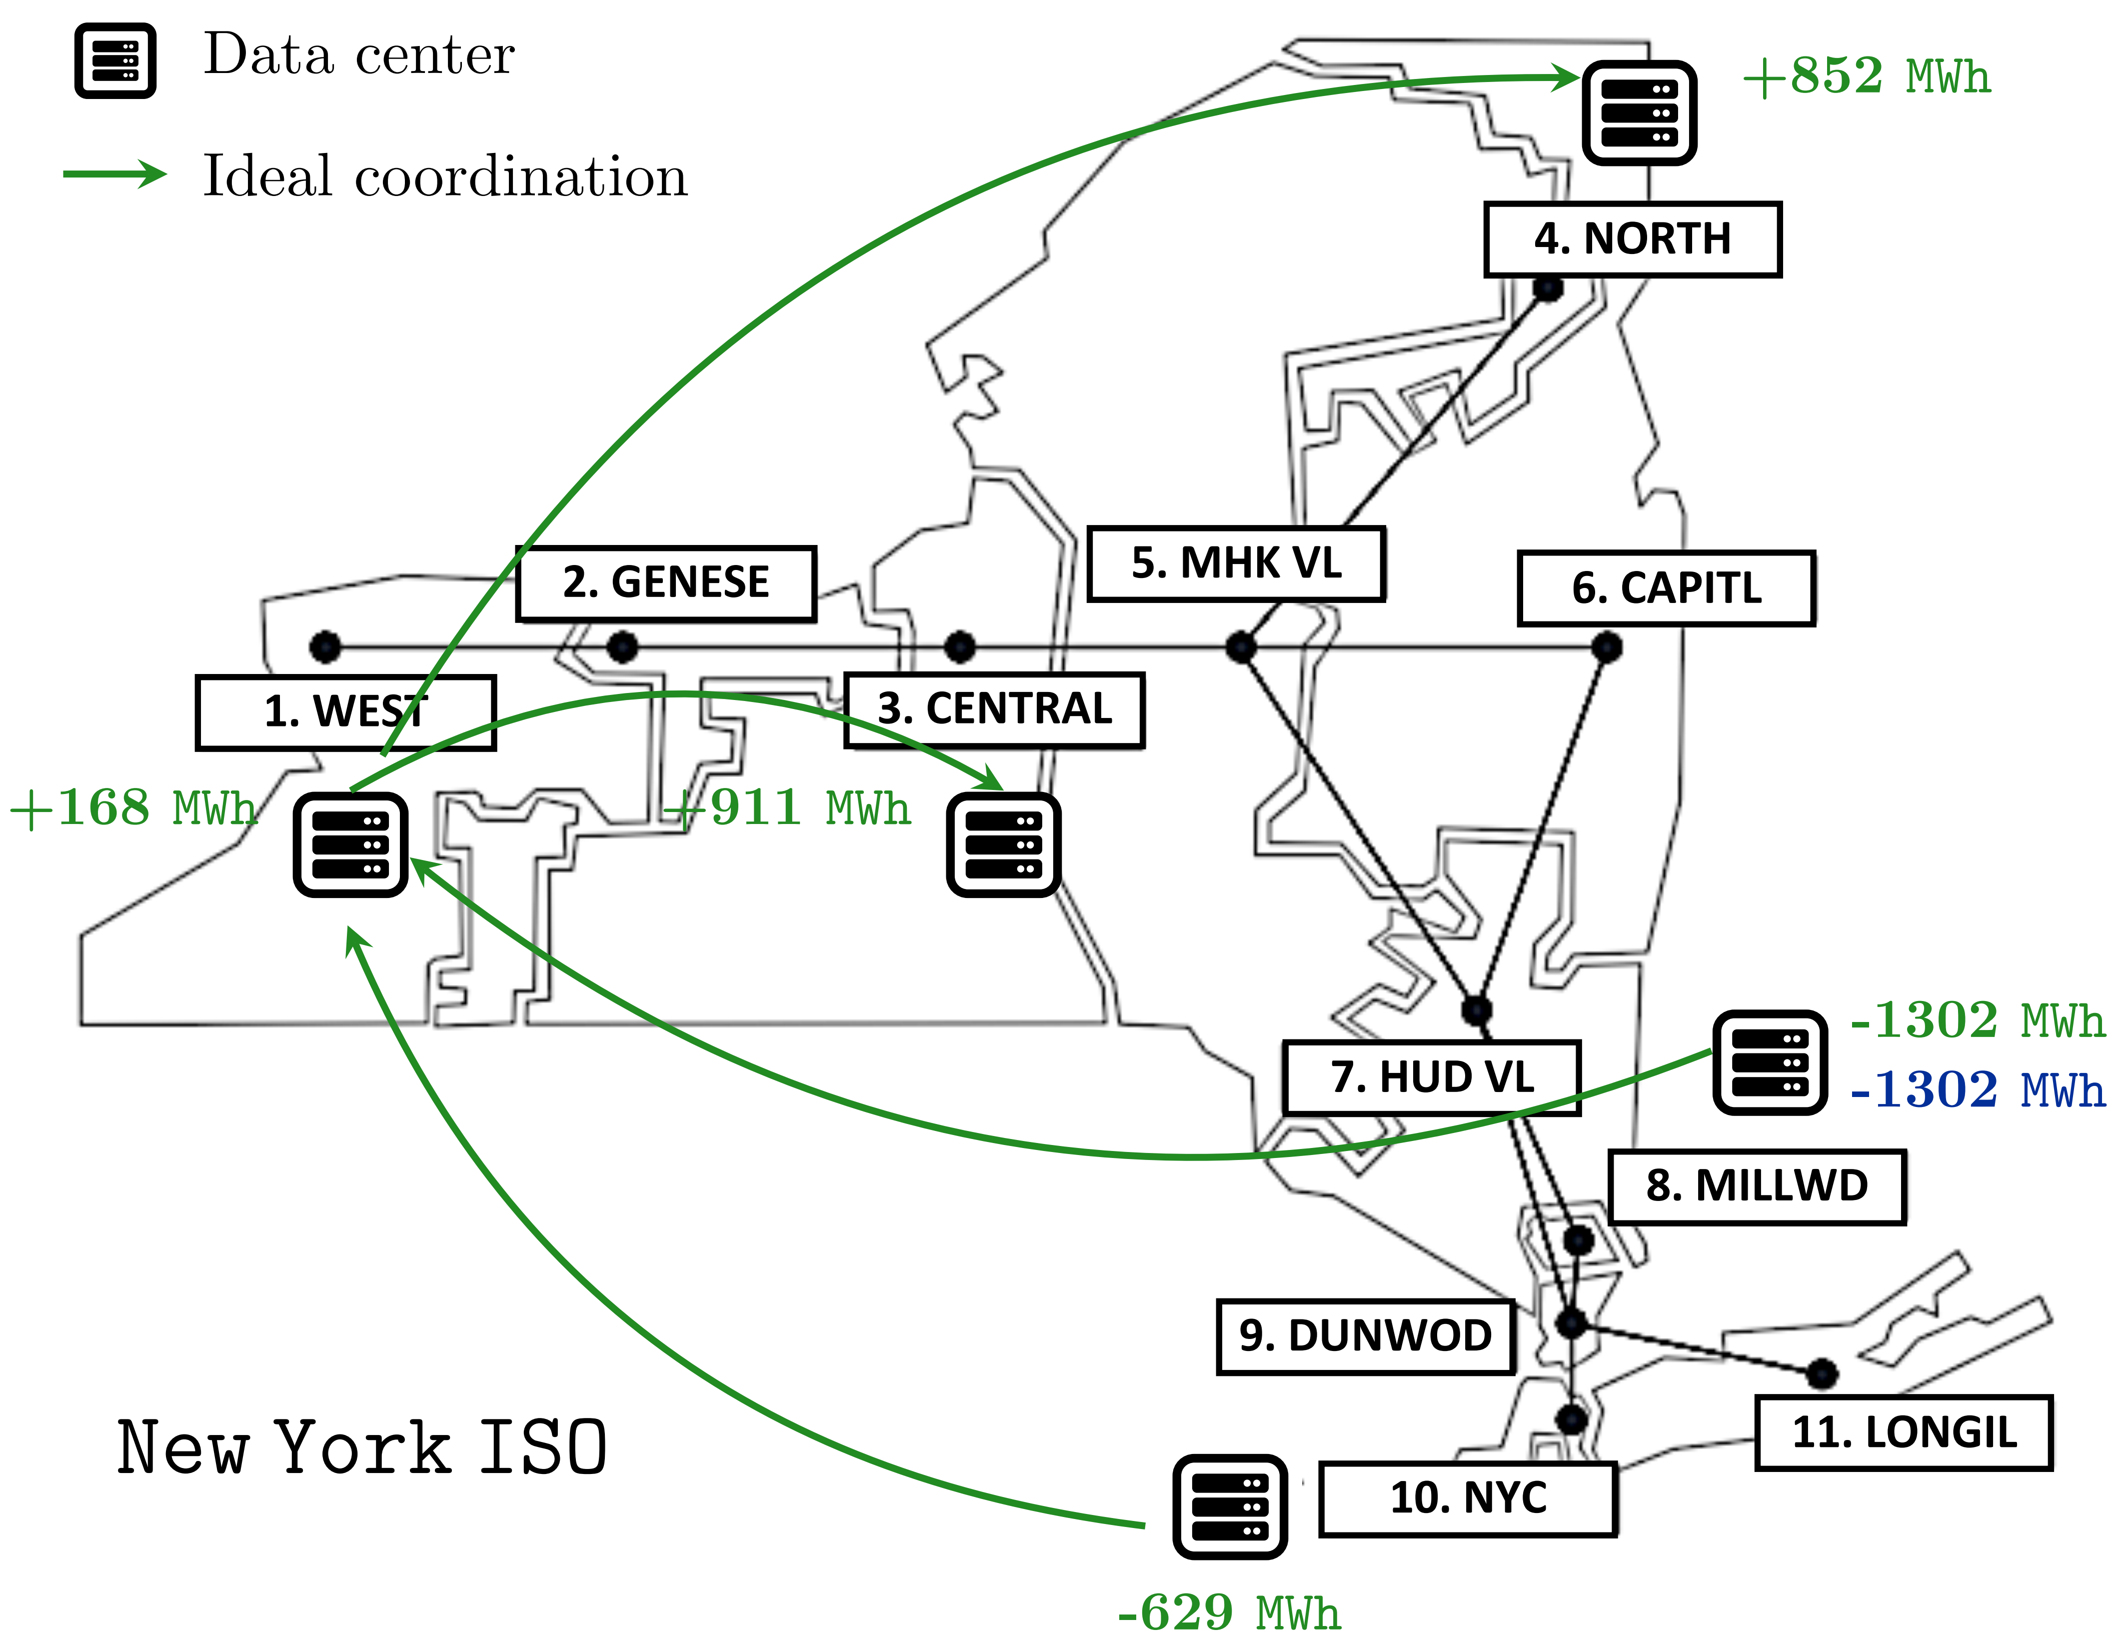

Example of task shifts across the network of data centers

Relaxation of latency constraints ⇒ greater generation cost savings

Unit commitment constraints prevent unlocking the whole NetDC flexibility

The flat (in red) loading profile is re-distributed in space and time

NYISO: Ideal coordination at the day-ahead stage

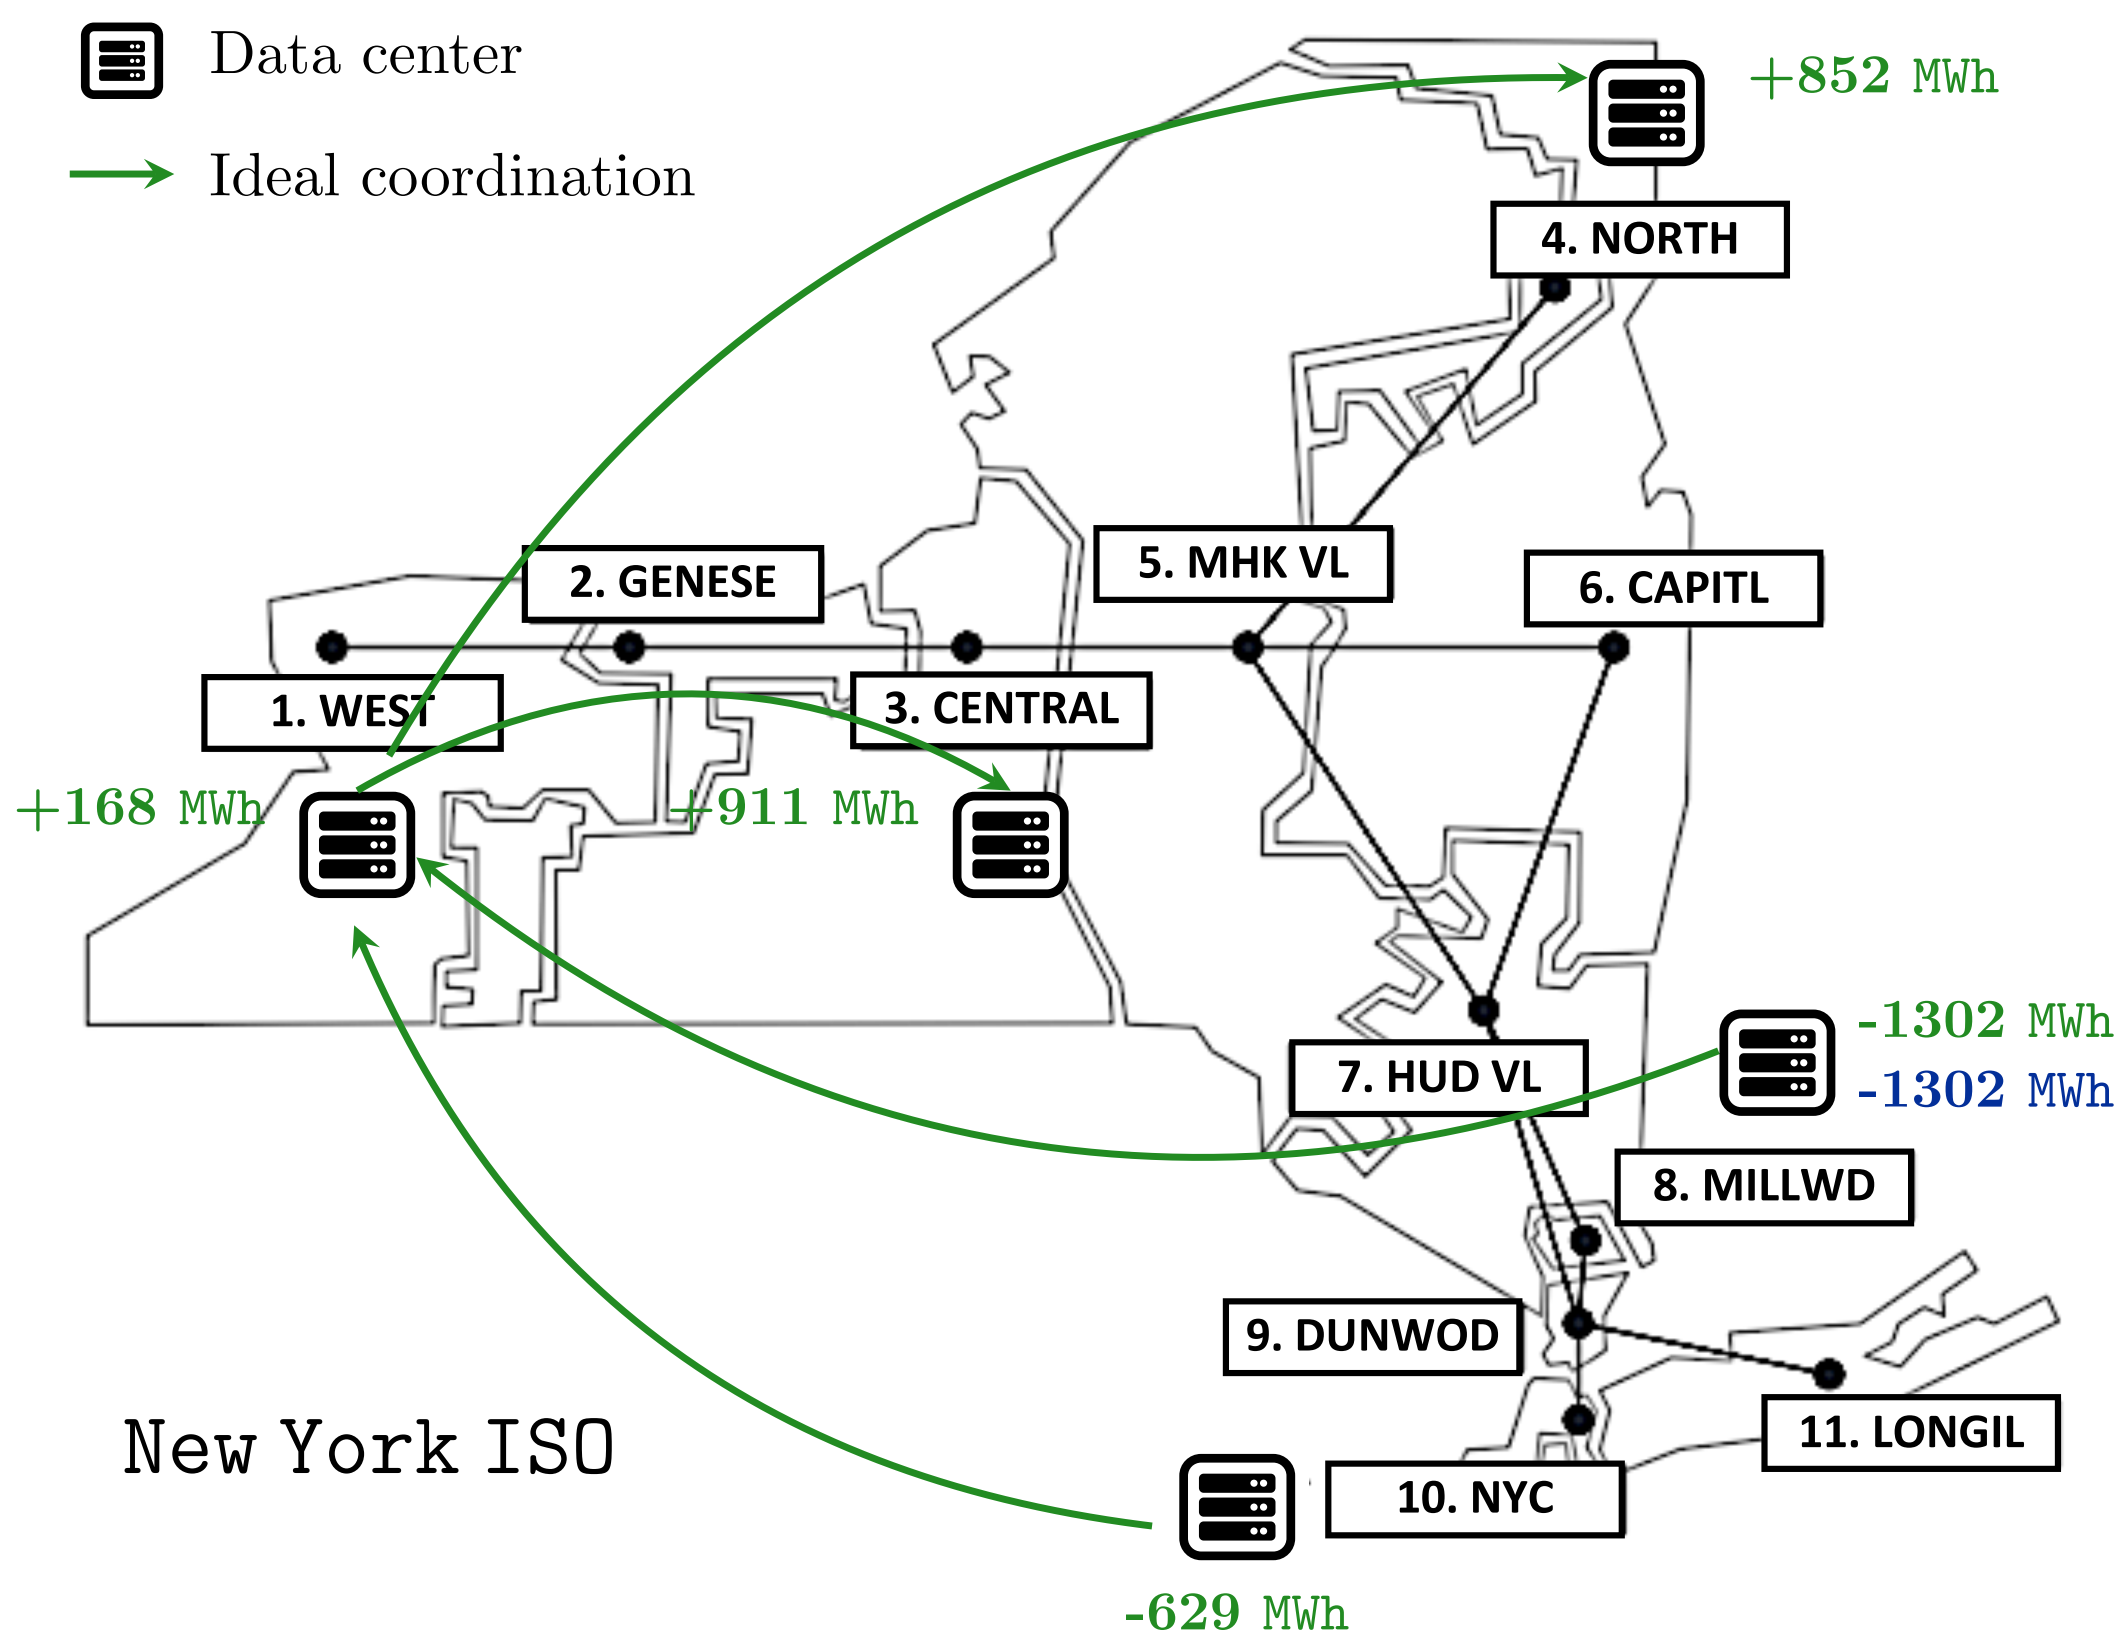

Example of task shifts across the network of data centers

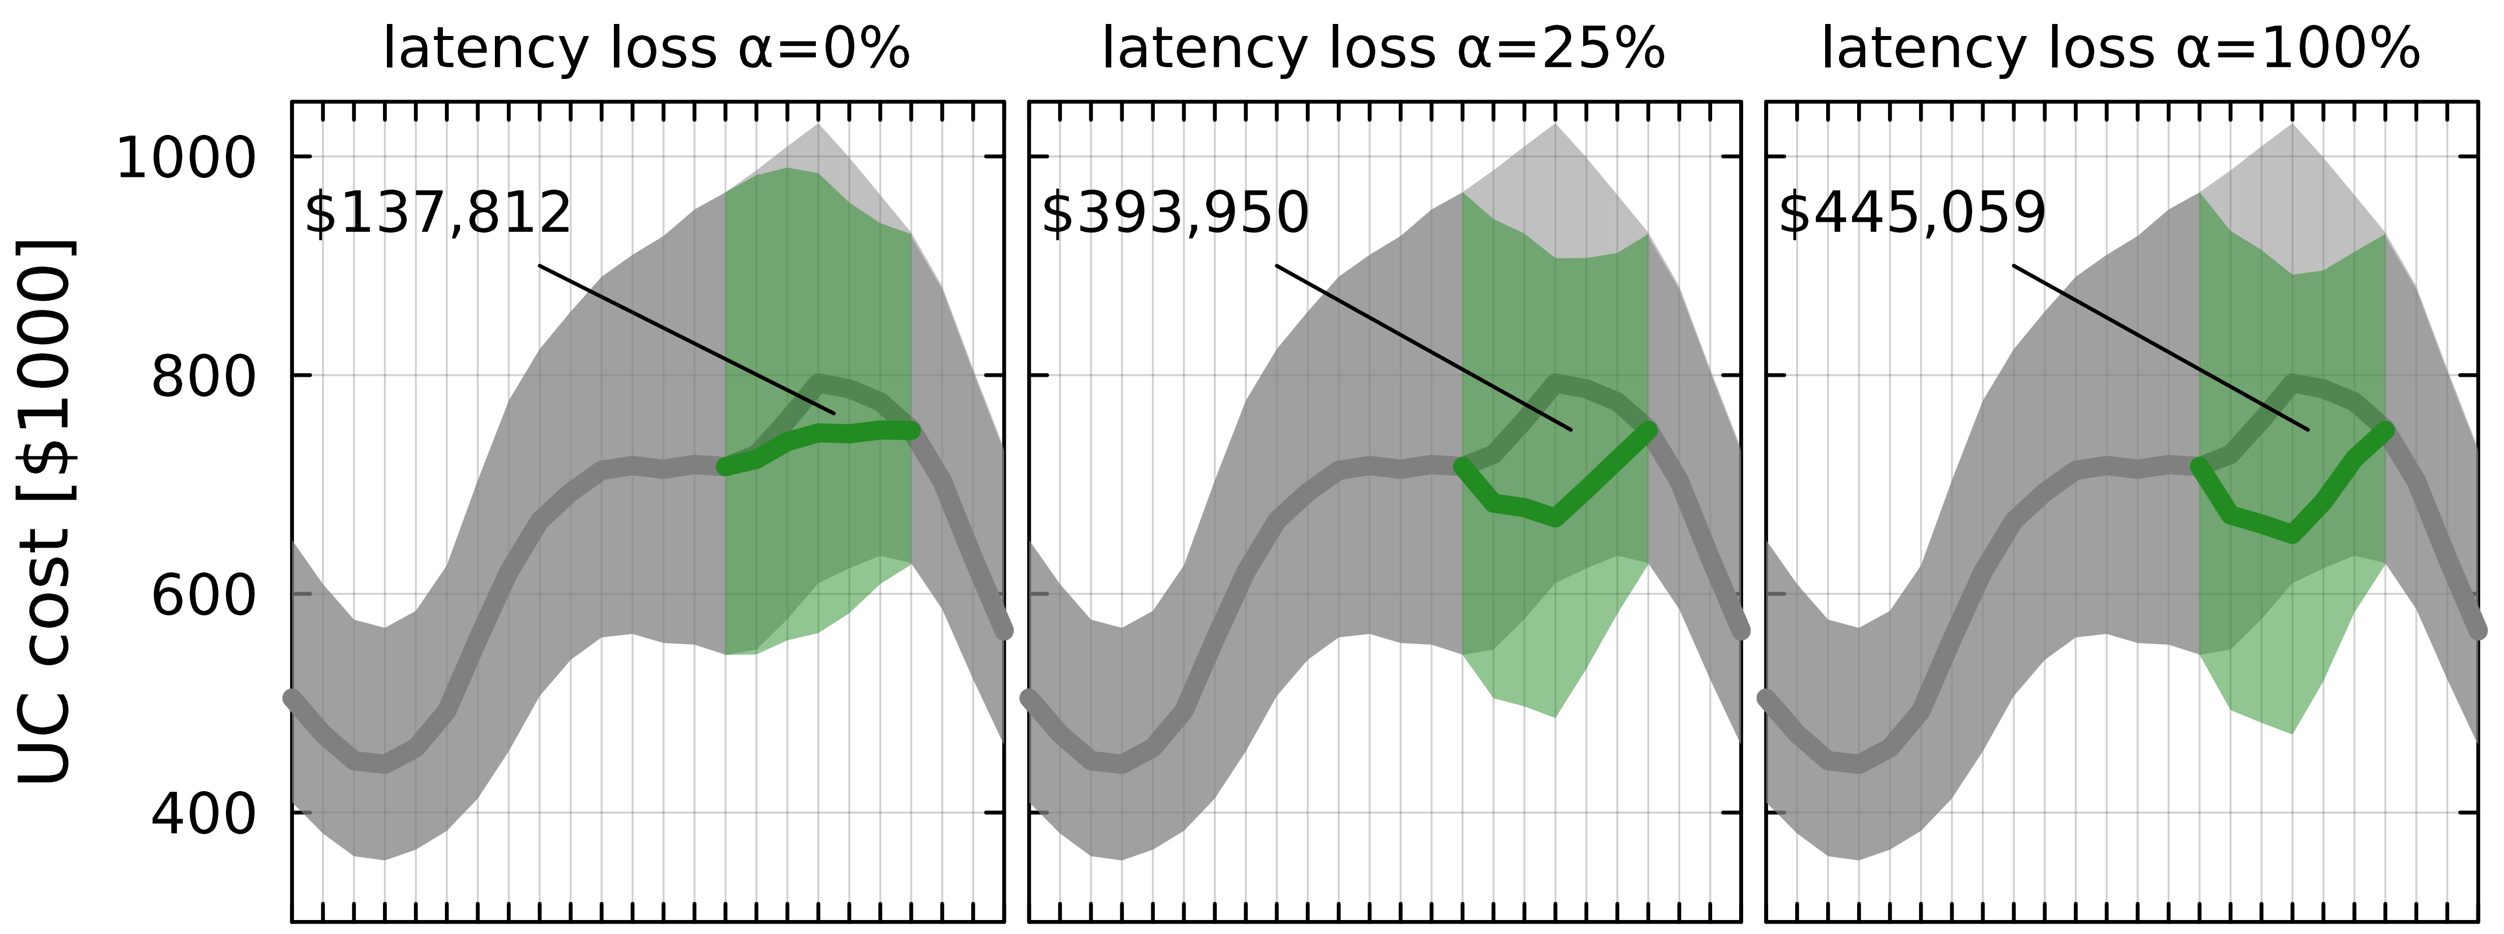

Relaxation of latency constraints ⇒ greater generation cost savings

Unit commitment constraints prevent unlocking the whole NetDC flexibility

The flat (in red) loading profile is re-distributed in space and time

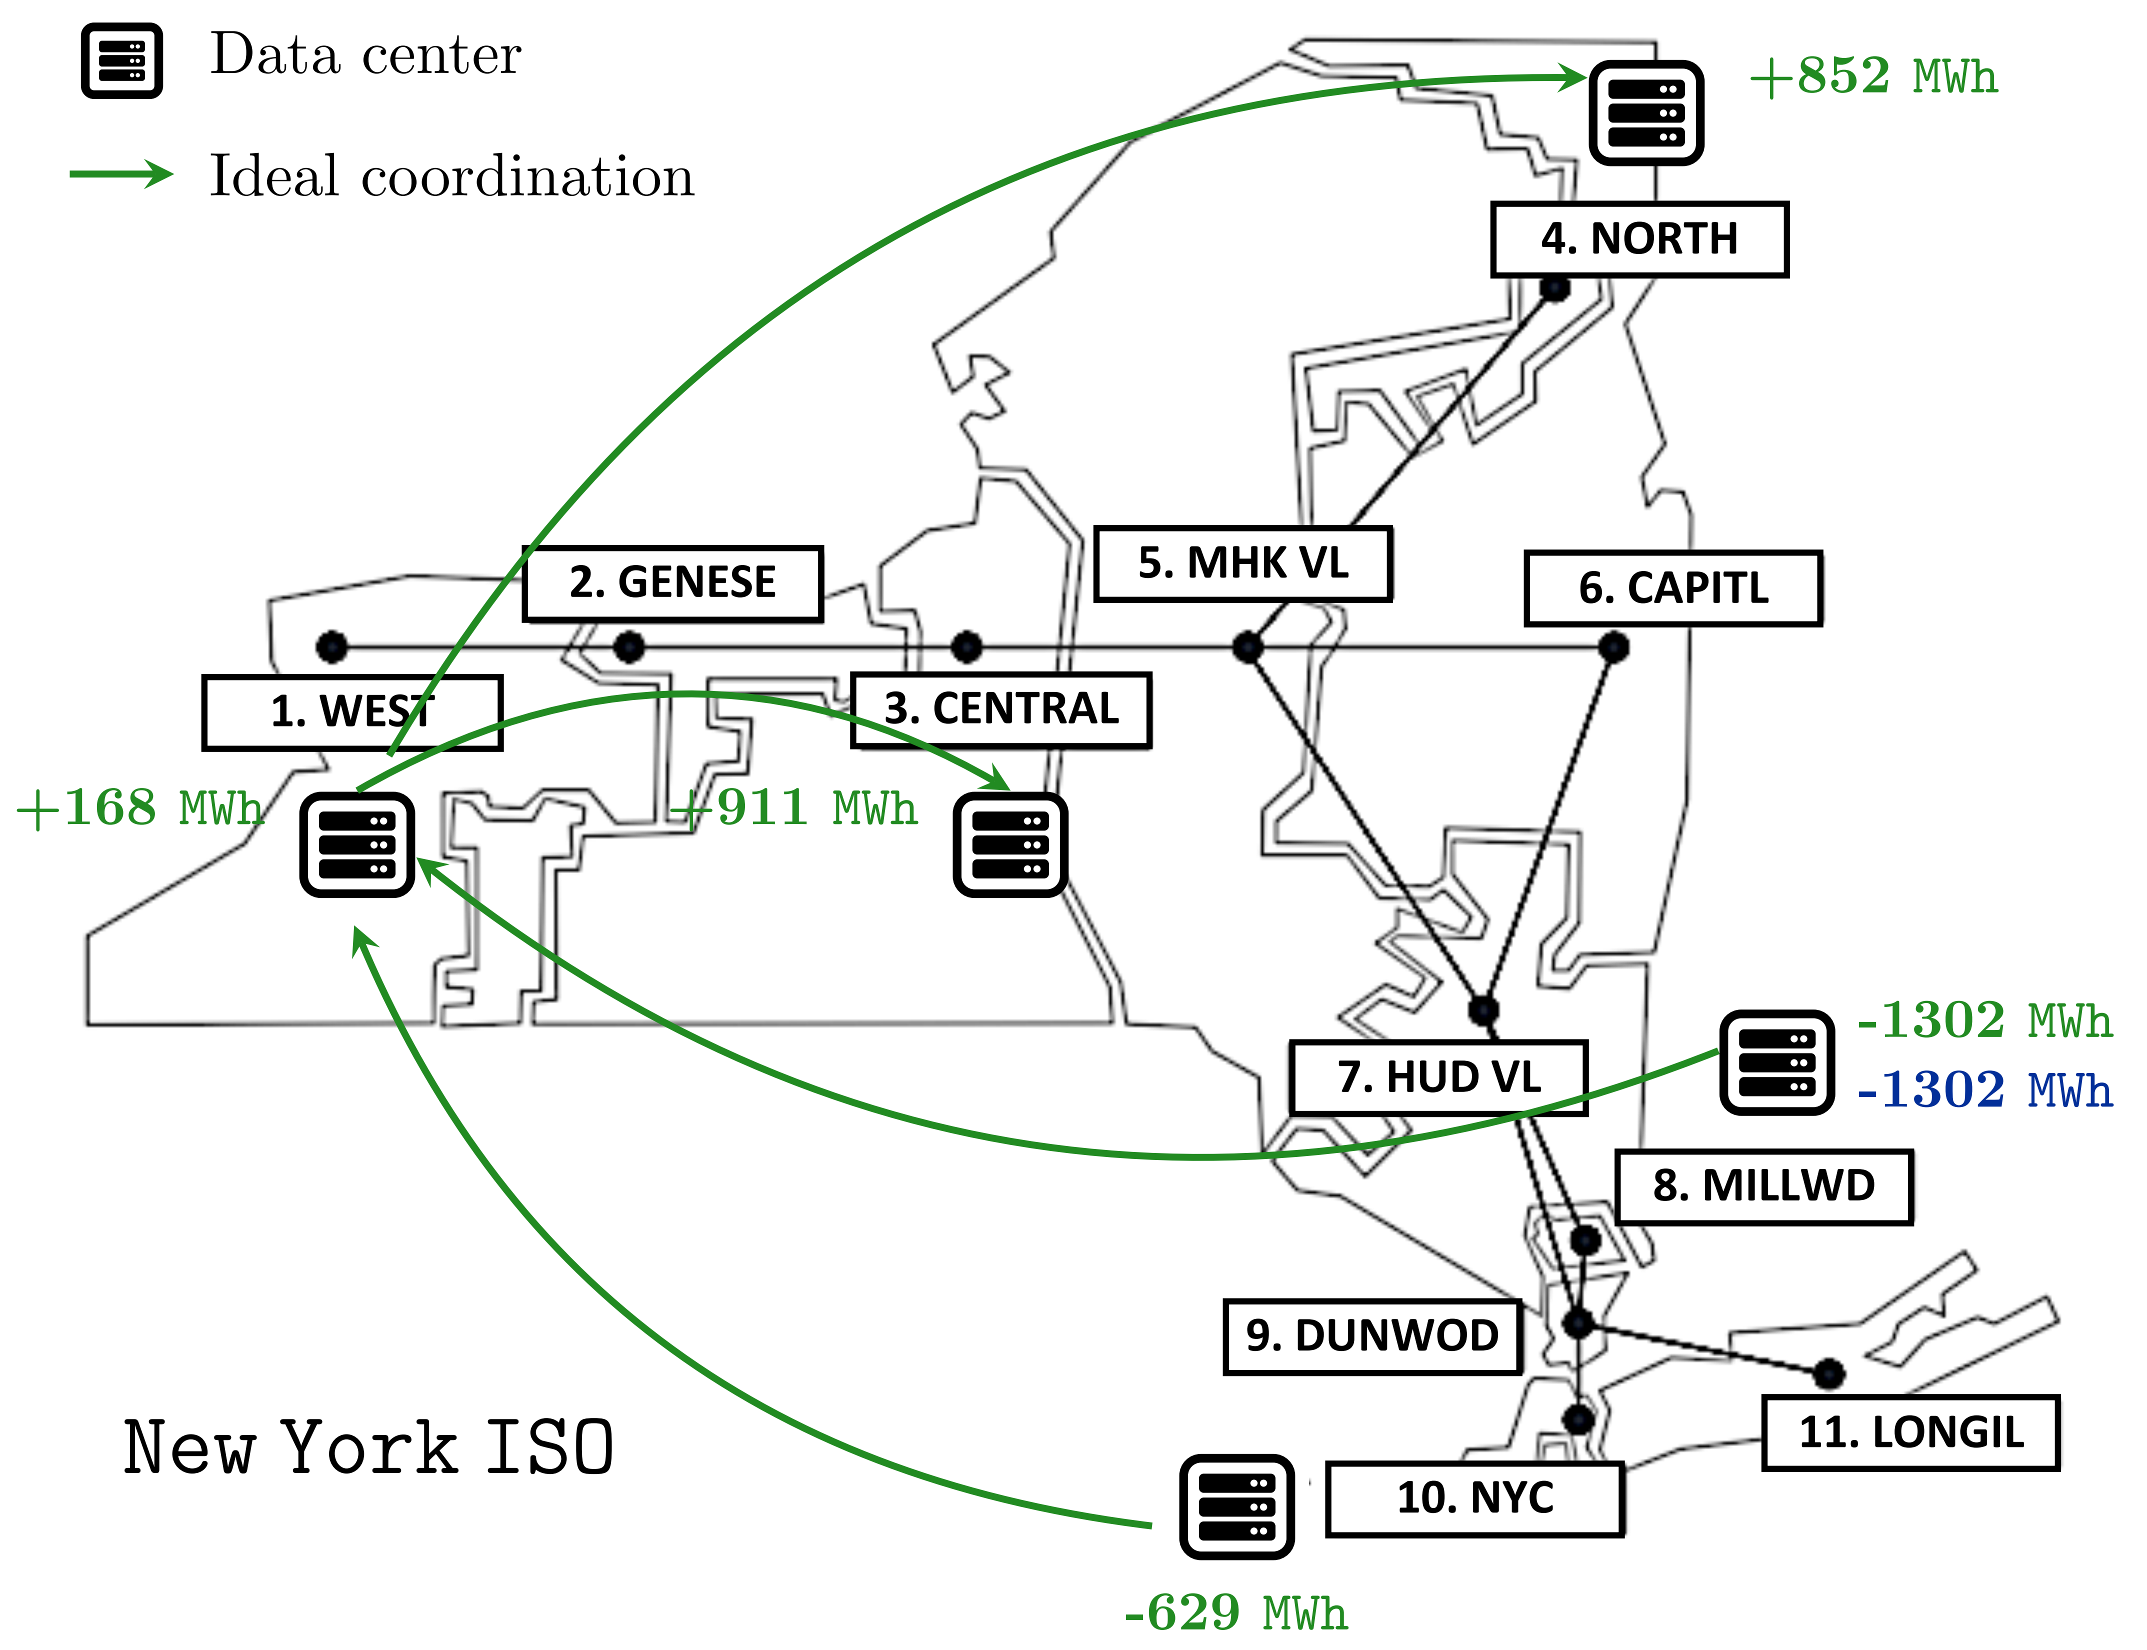

Example of task shifts across the network of data centers

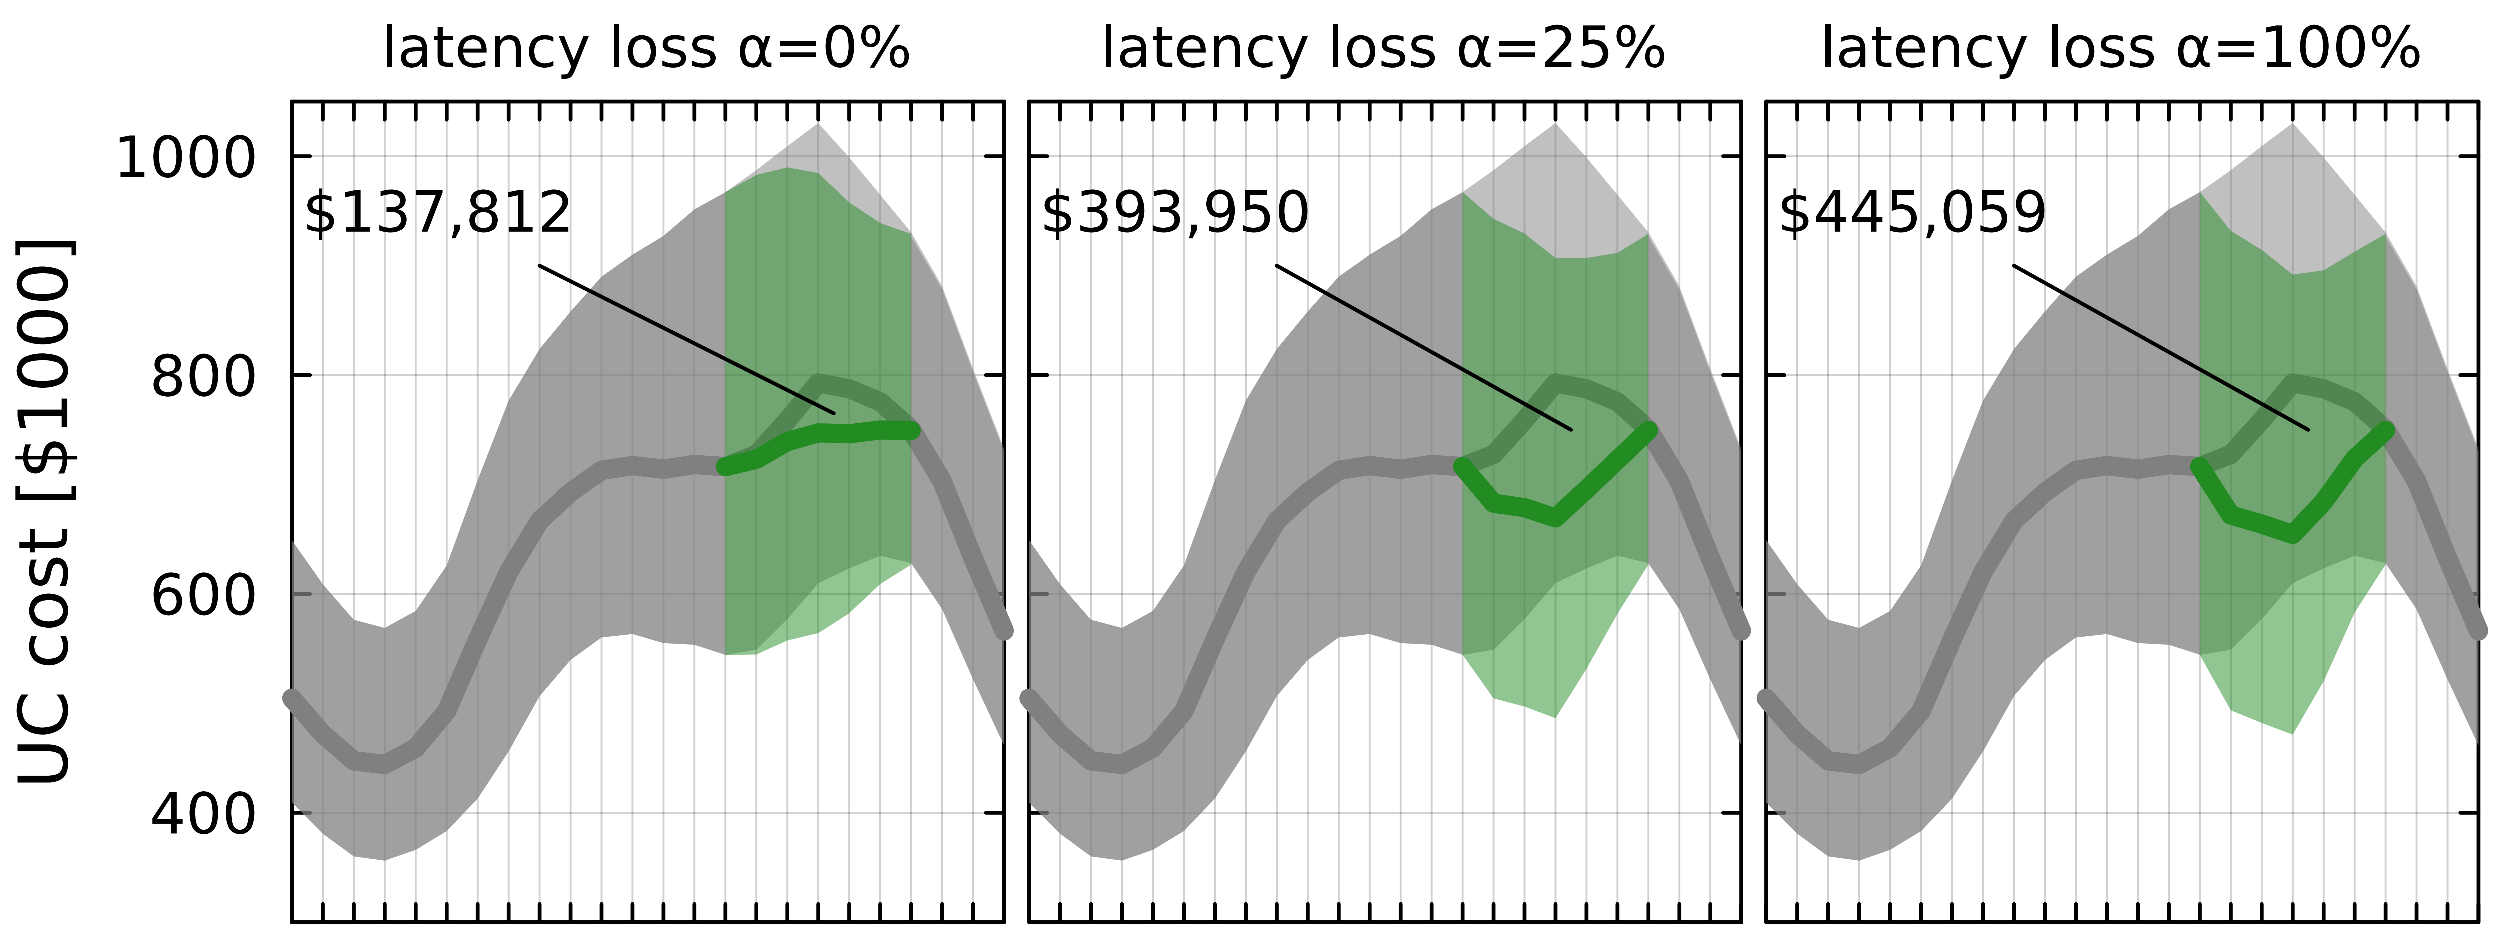

Relaxation of latency constraints ⇒ greater generation cost savings

Unit commitment constraints prevent unlocking the whole NetDC flexibility

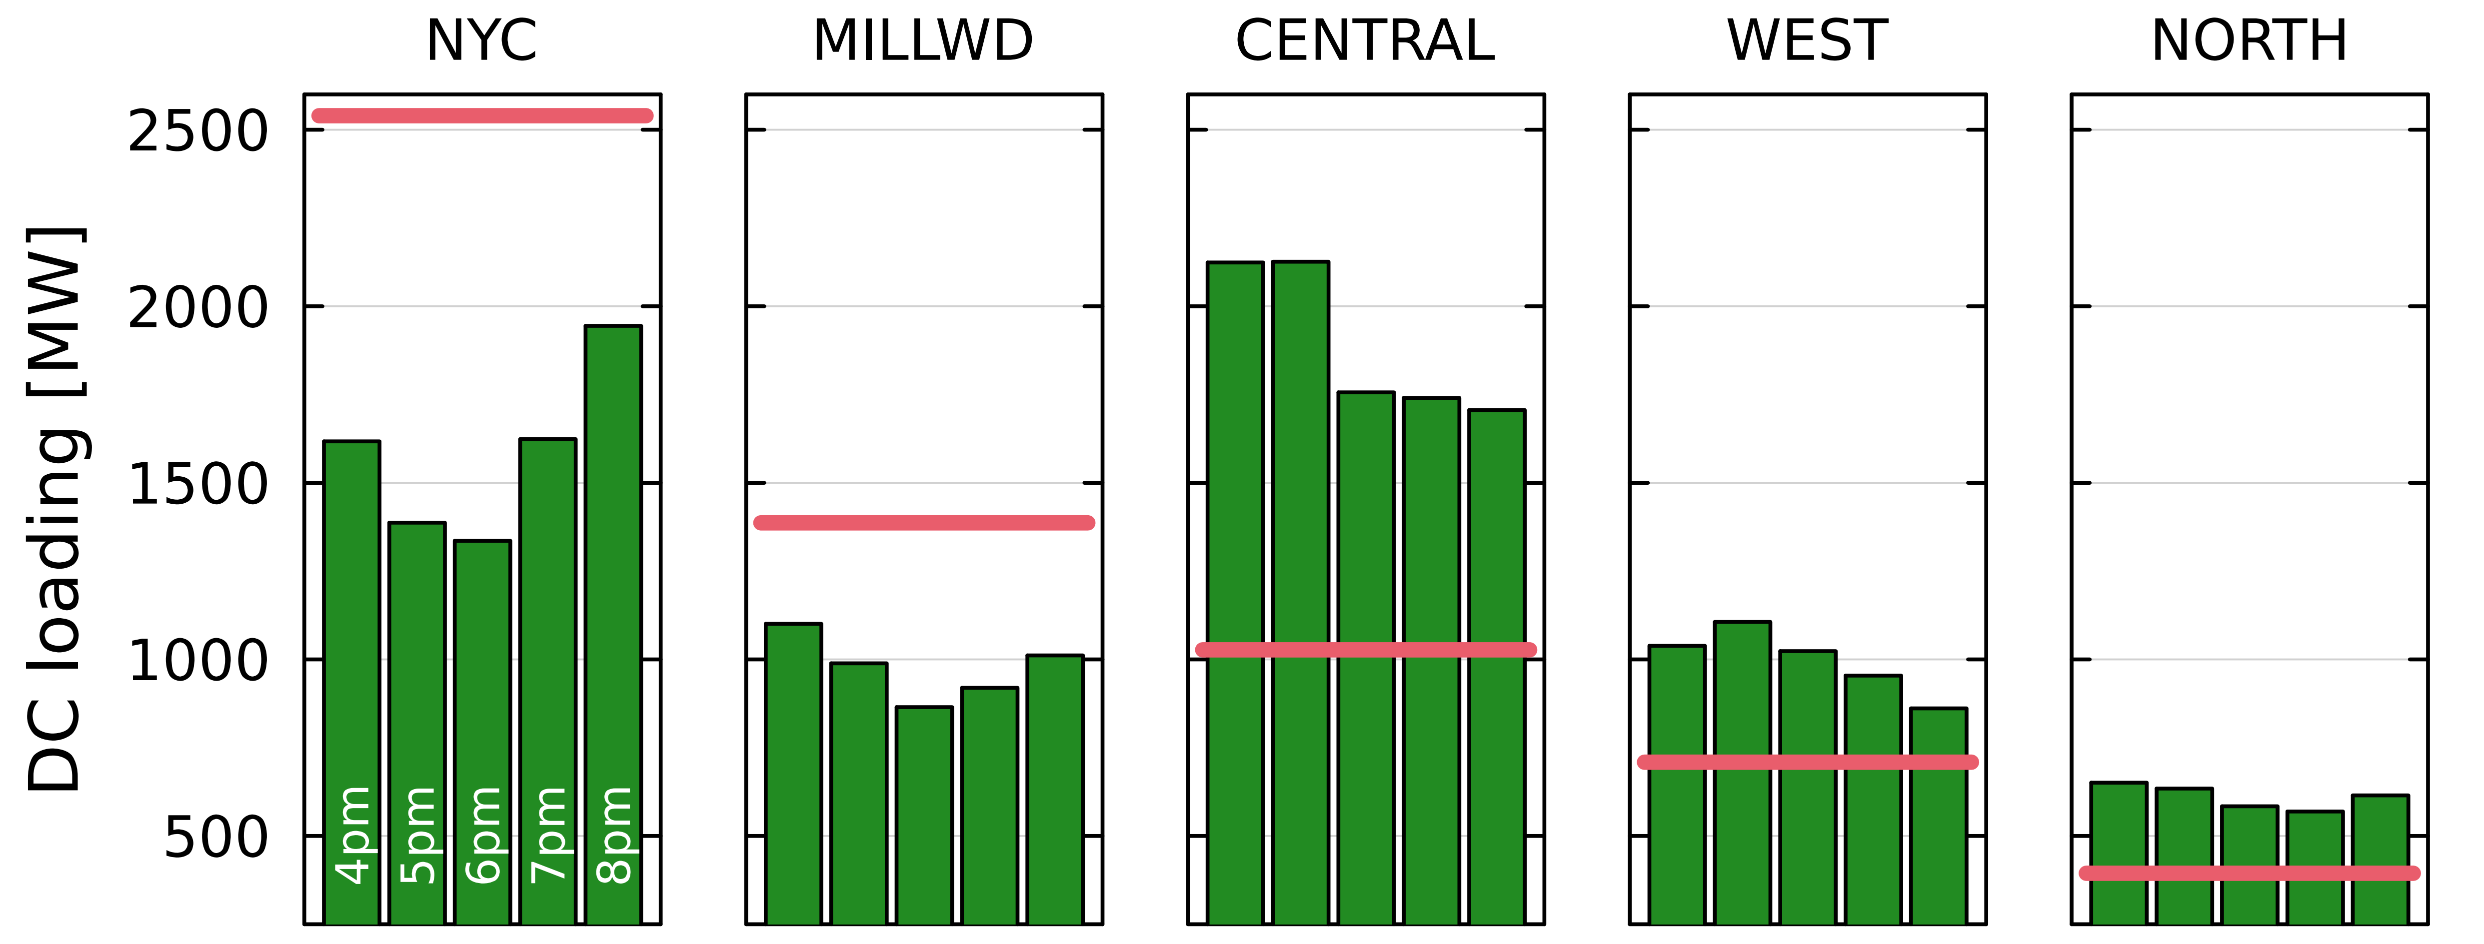

The flat (in red) loading profile is re-distributed in space and time

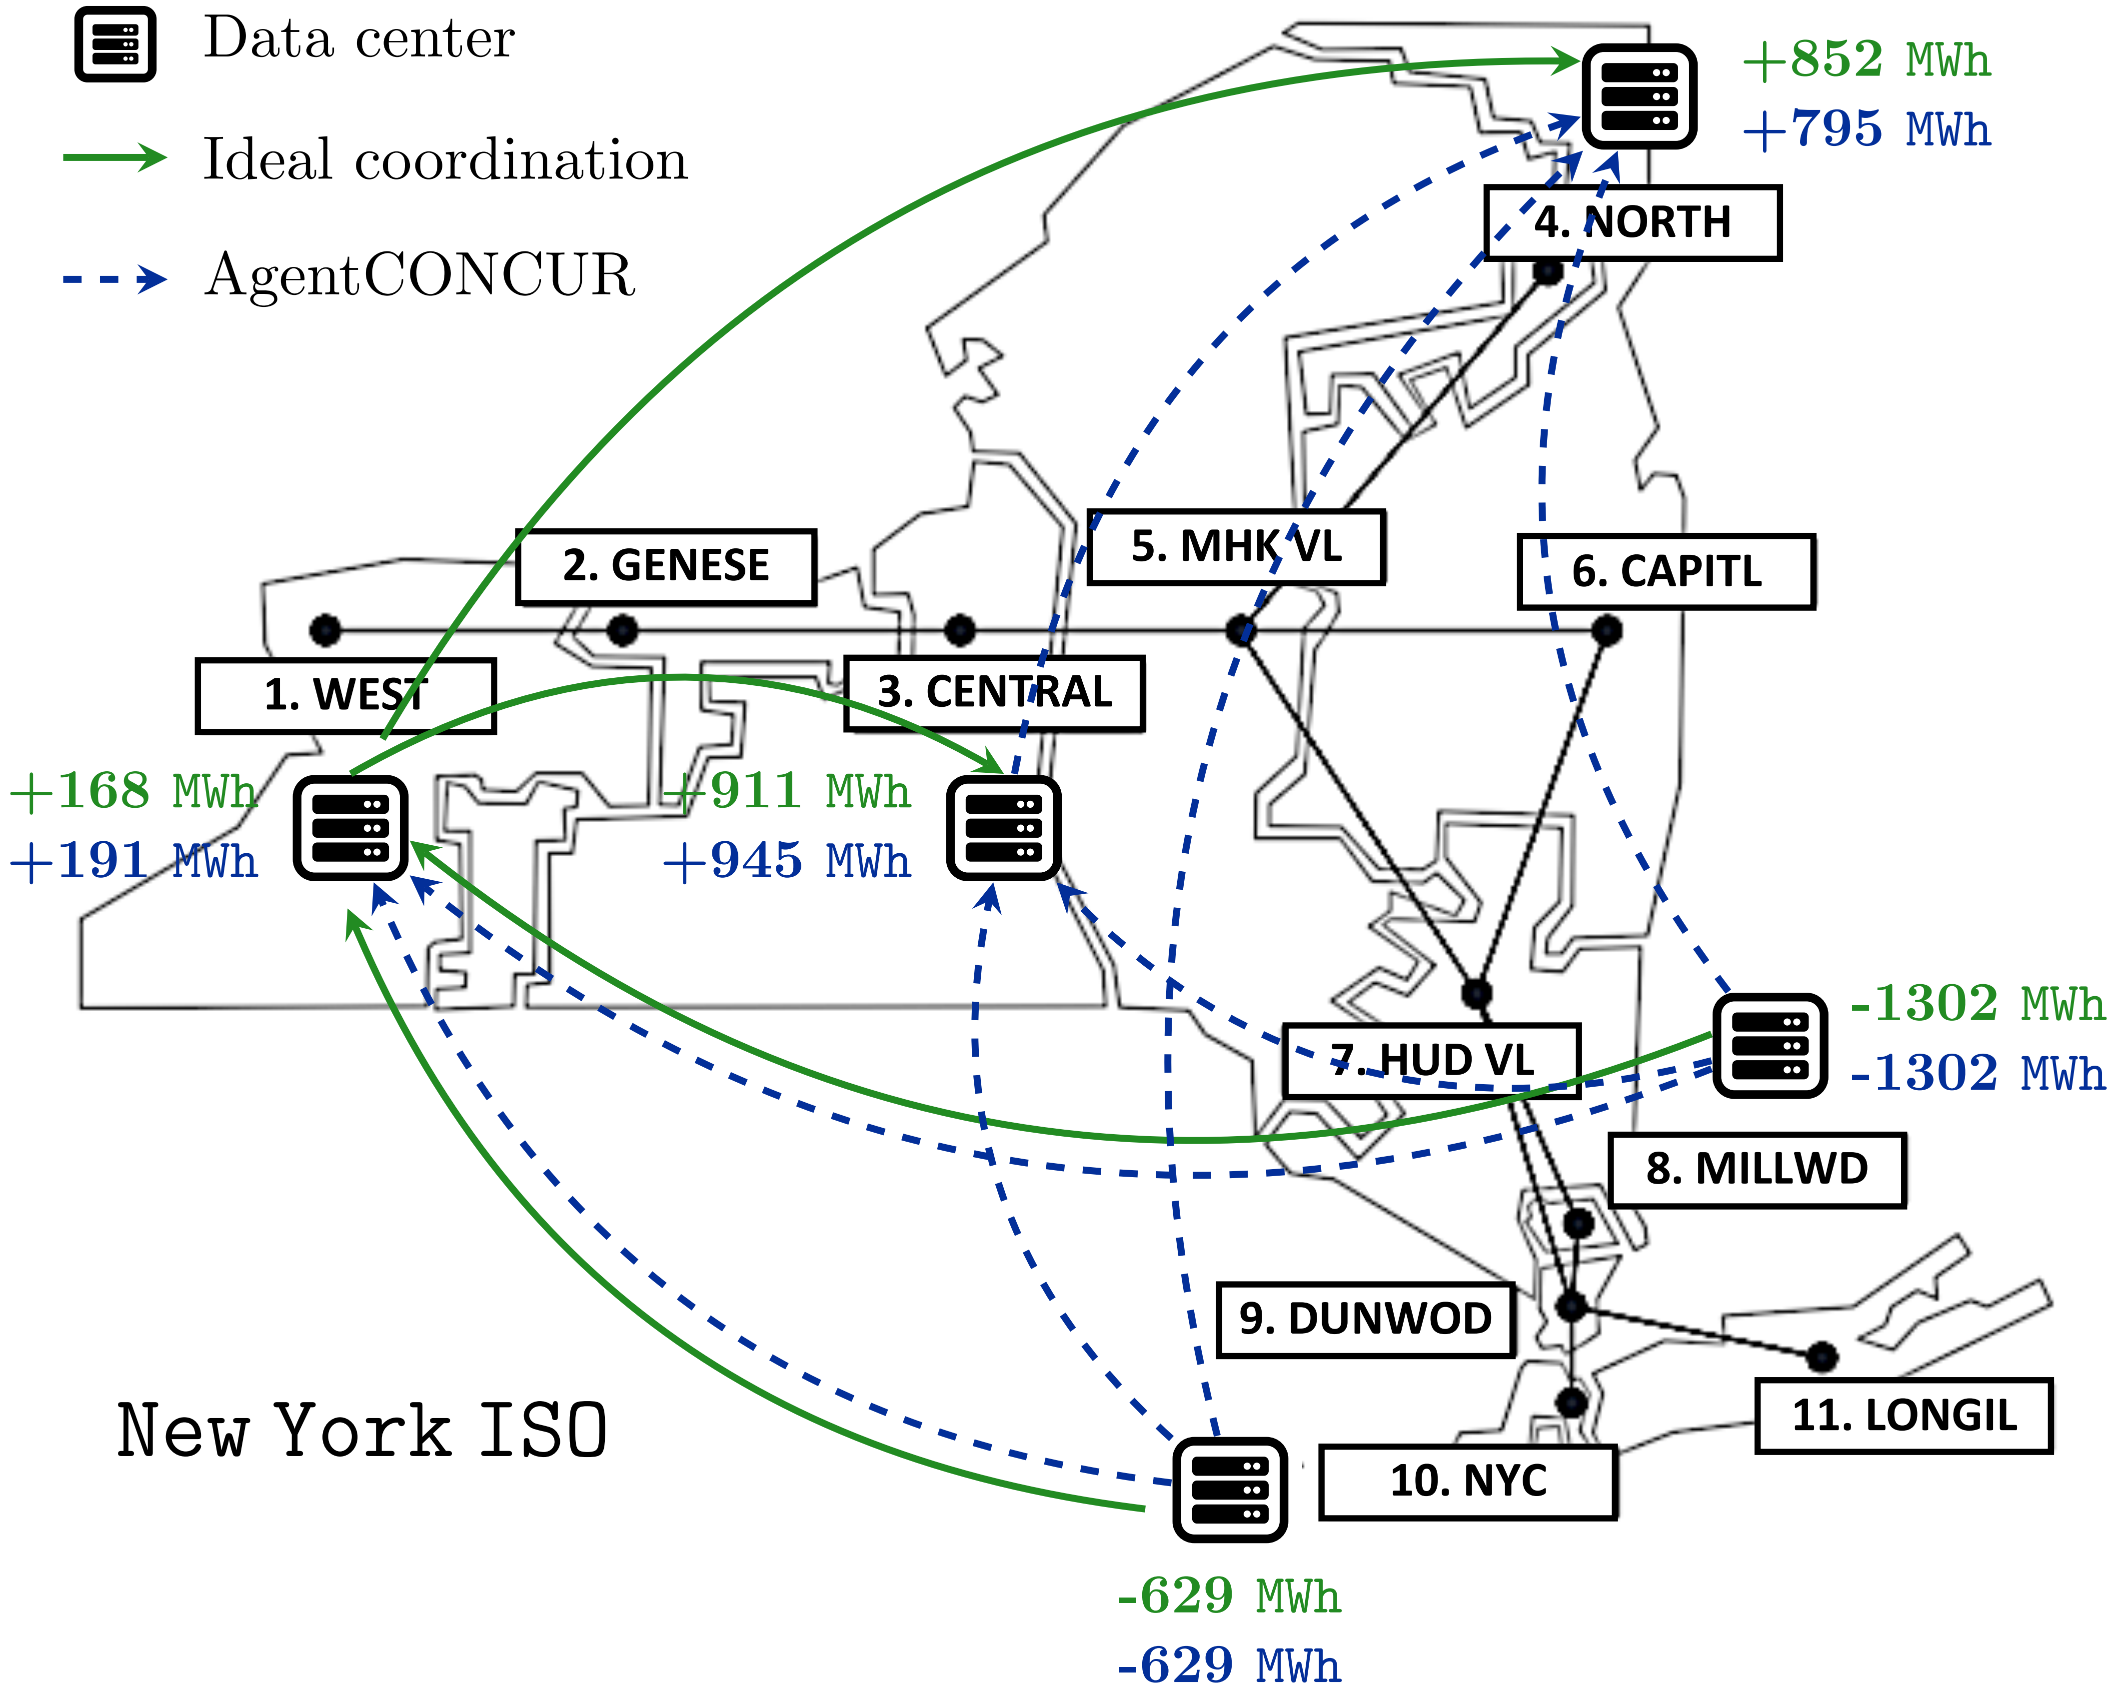

Ideal coordination versus the AgentCONCUR solution

Contextual features from NYISO website

- Zonal real-time electricity demand (d);
- Zonal electricity prices (λ);
- Zonal renewable power generation (r);
- Power flows between aggregation zones (f).

Coordination policy to be optimized offline:

$$\phi \triangleq \beta_0 + \beta_1^d \begin{bmatrix} d_1 \\ \vdots \\ d_{11} \end{bmatrix} + \beta_1^\lambda \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_{11} \end{bmatrix} + \beta_1^r \begin{bmatrix} r_1 \\ \vdots \\ r_{11} \end{bmatrix} + \beta_1^f \begin{bmatrix} f_1 \\ \vdots \\ f_{12} \end{bmatrix}$$

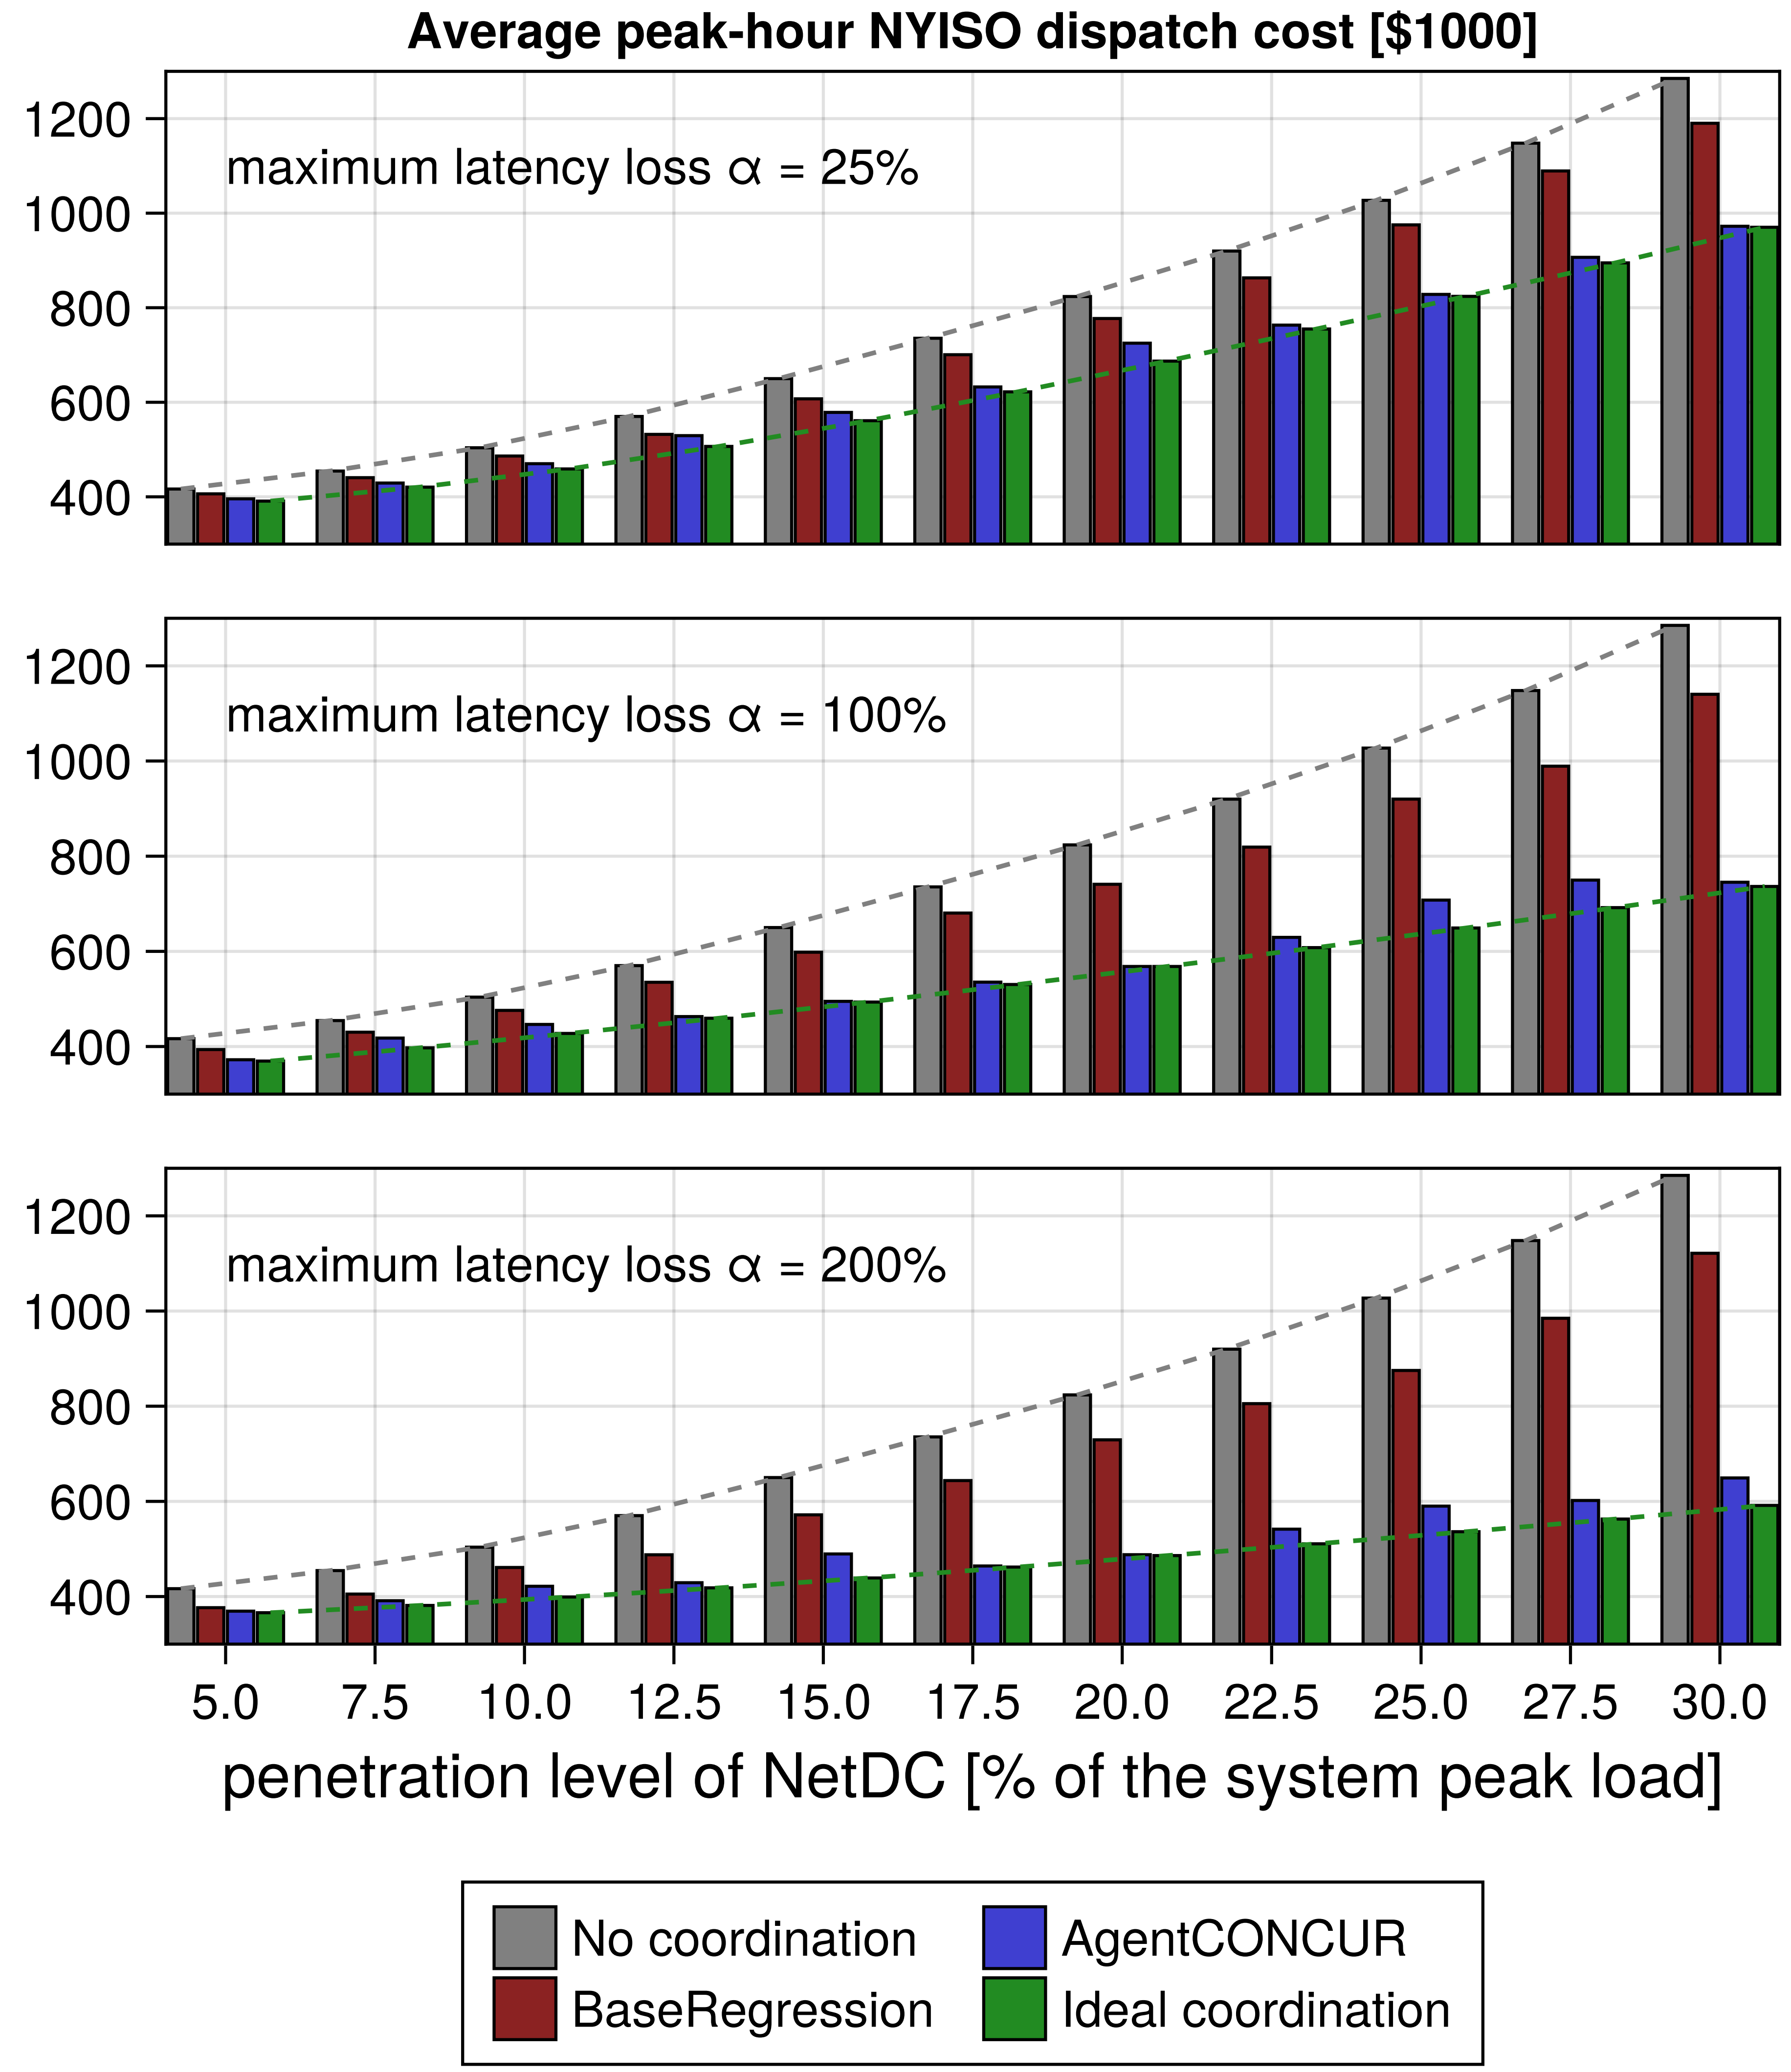

- Non-coordinated solution ⇒ quadratic cost growth
- Ideal coordination ⇒ more linear cost growth
- Baseline regression is agnostic to grid and NetDC constraints
- Informed regression guarantees feasibility ⇒ efficient approximation of the ideal coordination

NYISO: Cost-savings under regression approach

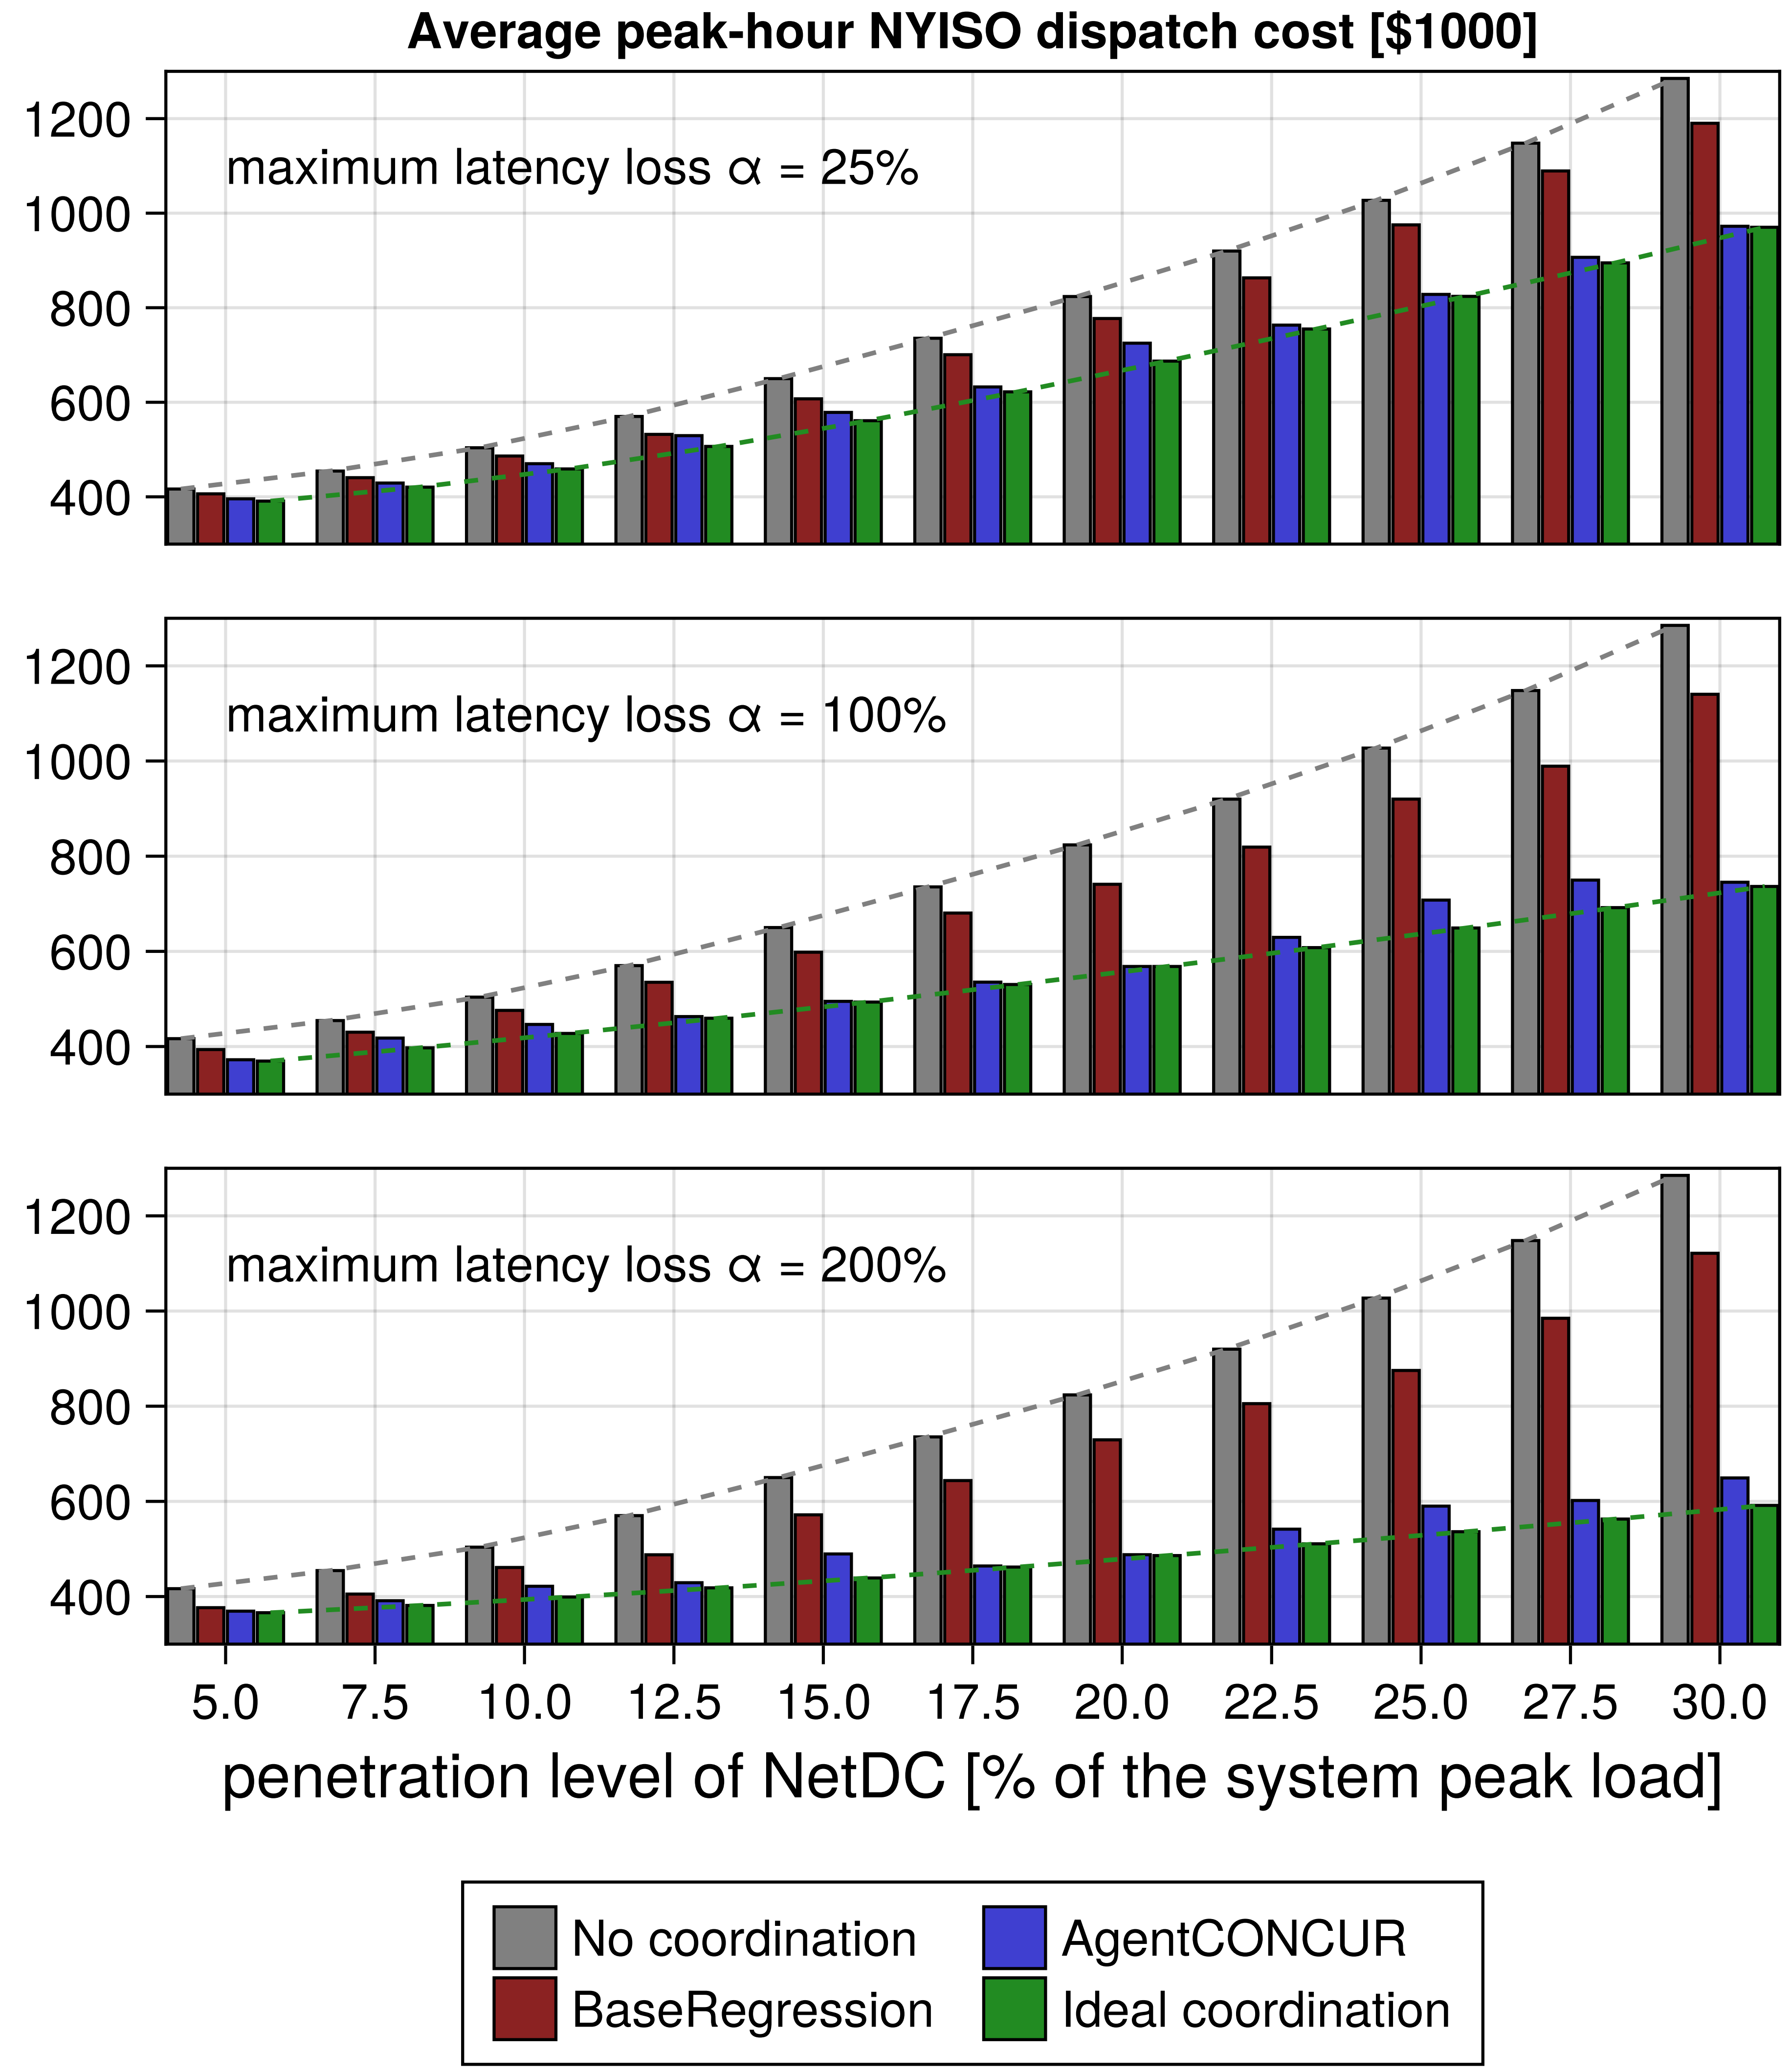

- **Non-coordinated solution** $\Rightarrow$ quadratic cost growth
- **Ideal coordination** $\Rightarrow$ more linear cost growth
- **Baseline regression** is agnostic to grid and NetDC constraints
- **Informed regression** guarantees feasibility $\Rightarrow$ efficient approximation of the ideal coordination

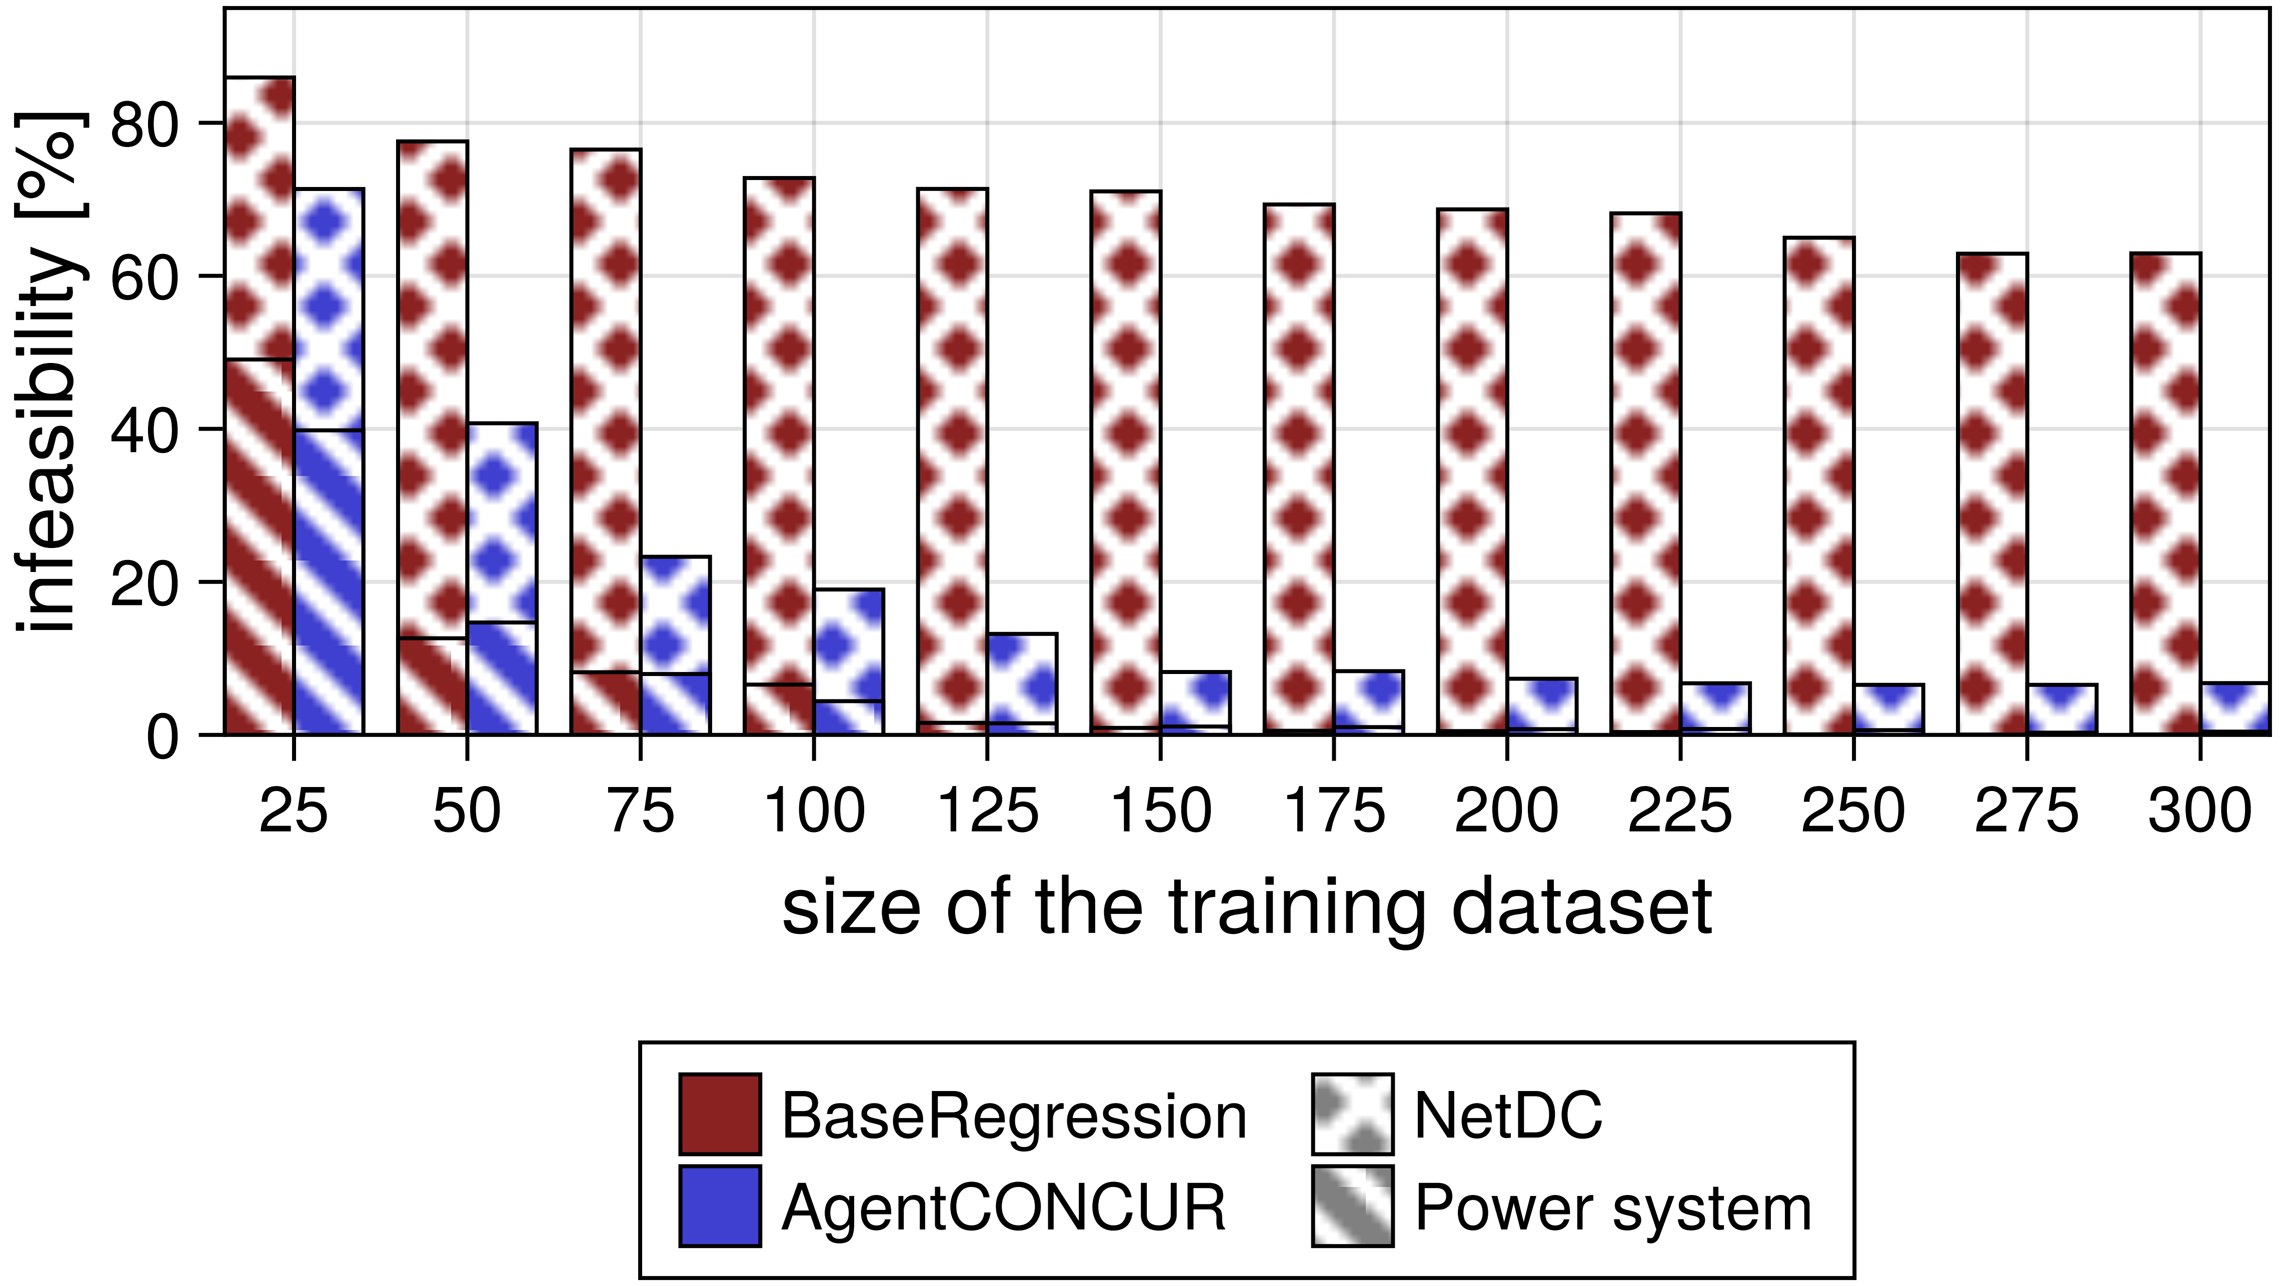

NYISO: Coordination feature selection for regression

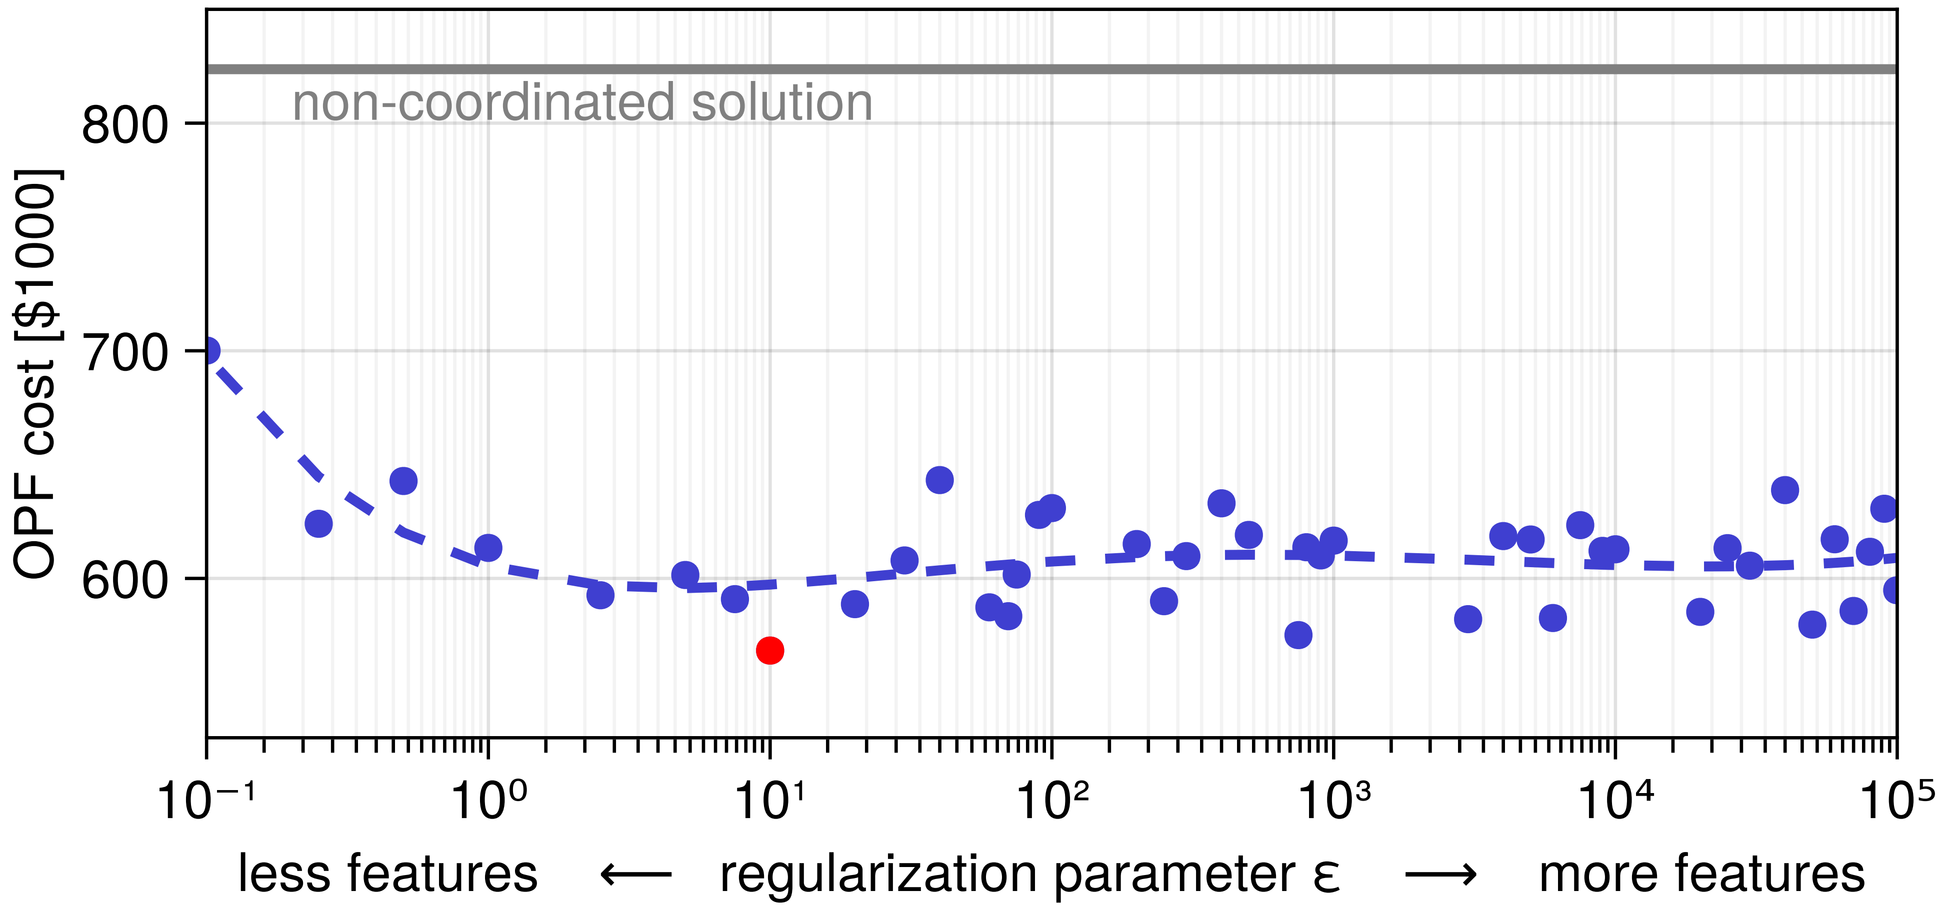

- Feature selection by means of ℓ_1-regularization
- ℓ_1-regularization also ensures coordination robustness
- Can we organize coordination using just one feature?

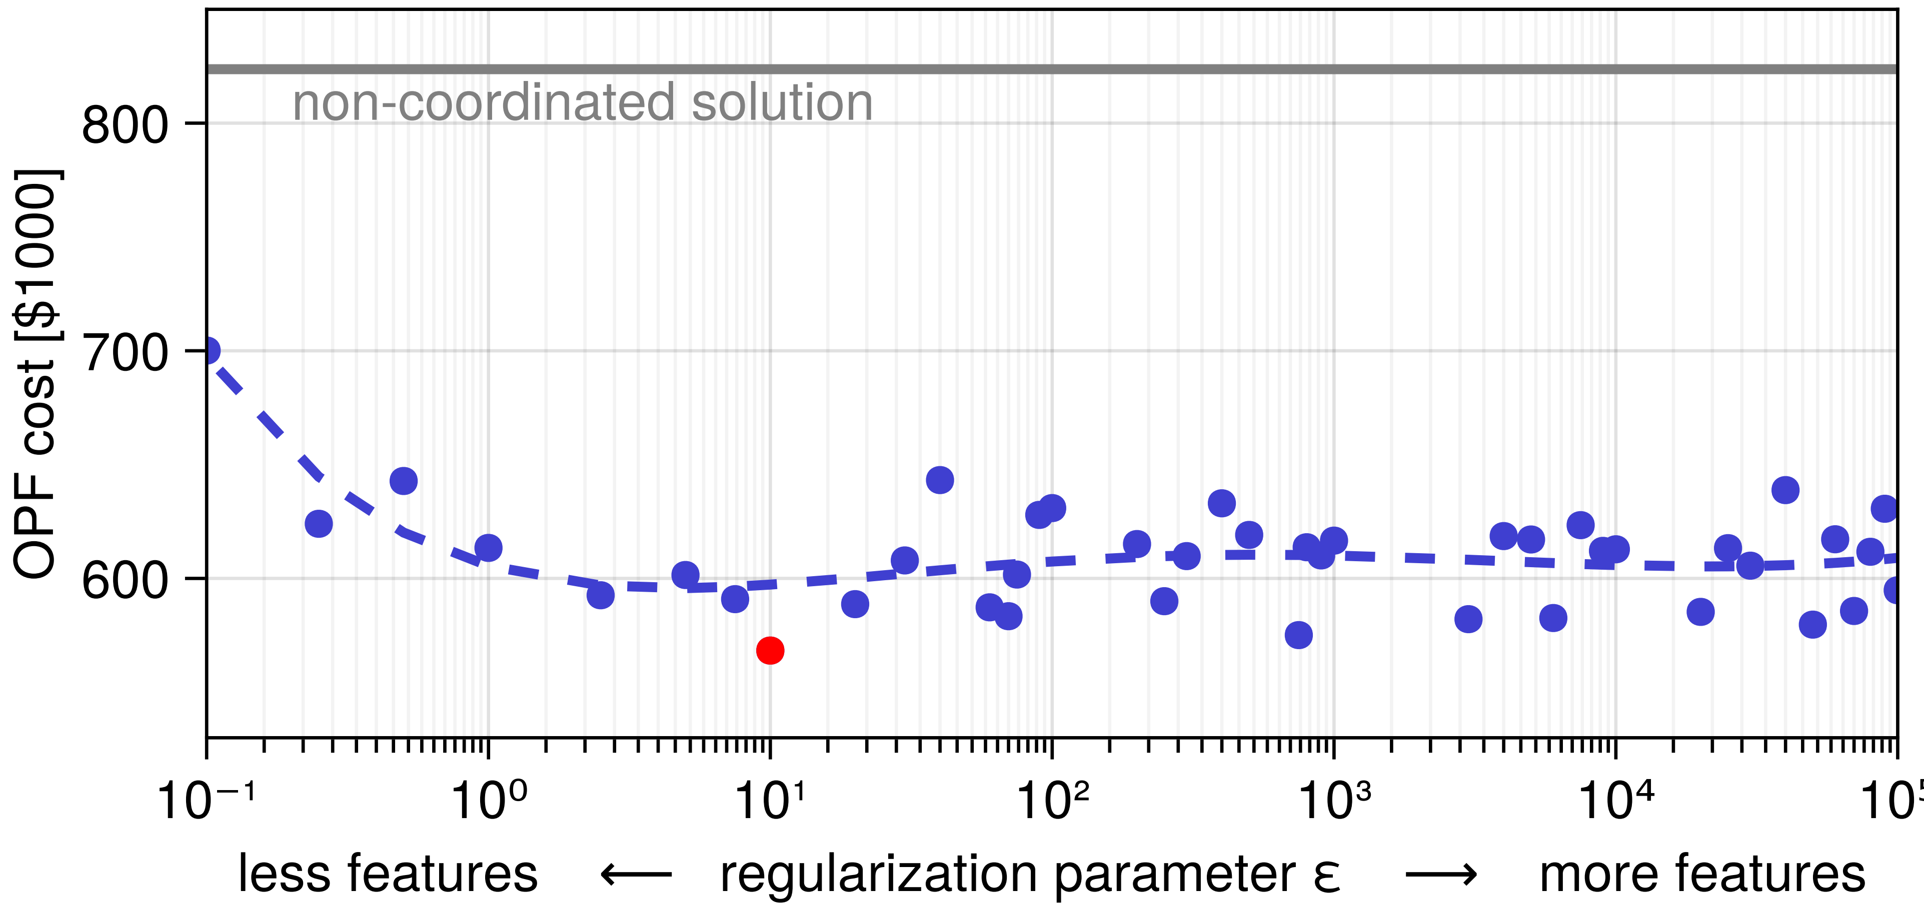

- Feature selection by means of ℓ_1-regularization
- ℓ_1-regularization also ensures coordination robustness
- Can we organize coordination using just one feature?

zonal electricity demand

ε	# of features	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}	d_{11}
1000.0	29	○	●	●	●	○	●	○	●	●	○	○
100.0	28	●	●	●	●	●	●	○	●	●	●	○
10.0	24	●	●	●	●	○	●	●	●	●	●	●
5.0	20	●	●	○	●	●	●	●	●	●	●	●
2.5	13	●	○	●	○	○	●	●	●	●	●	○
1.0	6	●	○	●	○	○	○	○	○	○	●	○
0.5	3	●	○	○	○	○	○	○	○	○	●	○
0.1	1	○	○	○	○	○	○	○	○	○	●	○

power flow

ε	f_1	f_2	f_3	f_4	f_5	f_6	λ_7	f_8	f_9	f_{10}	f_{11}	f_{12}
1000.0	●	●	●	●	●	●	●	●	●	●	●	○
100.0	●	●	●	●	●	●	●	●	○	●	○	○
10.0	●	●	●	●	○	●	●	●	●	●	○	○
5.0	●	●	●	○	●	○	●	●	●	○	○	○
2.5	●	●	○	○	○	●	○	○	●	○	○	○
1.0	●	●	○	○	○	○	○	○	○	●	○	○
0.5	○	○	●	○	○	○	○	○	○	○	○	○
0.1	○	○	○	○	○	○	○	○	○	○	○	○

zonal electricity price

ε	λ_1	λ_2	λ_3	λ_4	λ_5	λ_6	λ_7	λ_8	λ_9	λ_{10}	λ_{11}
1000.0	●	●	●	●	●	●	●	○	●	●	○
100.0	○	●	●	●	●	○	○	●	●	●	●
10.0	○	○	○	○	○	○	○	●	●	●	○
5.0	○	○	○	○	○	○	○	○	●	●	○
2.5	○	○	○	○	○	○	○	○	●	●	○
1.0	○	○	○	○	○	○	○	○	○	○	○
0.5	○	○	○	○	○	○	○	○	○	○	○
0.1	○	○	○	○	○	○	○	○	○	○	○

zonal renewable power output

ε	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}	r_{11}
1000.0	○	●	●	○	○	○	○	○	○	○	●
100.0	○	●	○	○	○	○	○	○	○	○	●
10.0	○	●	○	○	○	○	○	○	○	○	●
5.0	○	●	○	○	○	○	○	○	○	○	○
2.5	○	○	○	○	○	○	○	○	○	○	○
1.0	○	○	○	○	○	○	○	○	○	○	○
0.5	○	○	○	○	○	○	○	○	○	○	○
0.1	○	○	○	○	○	○	○	○	○	○	○

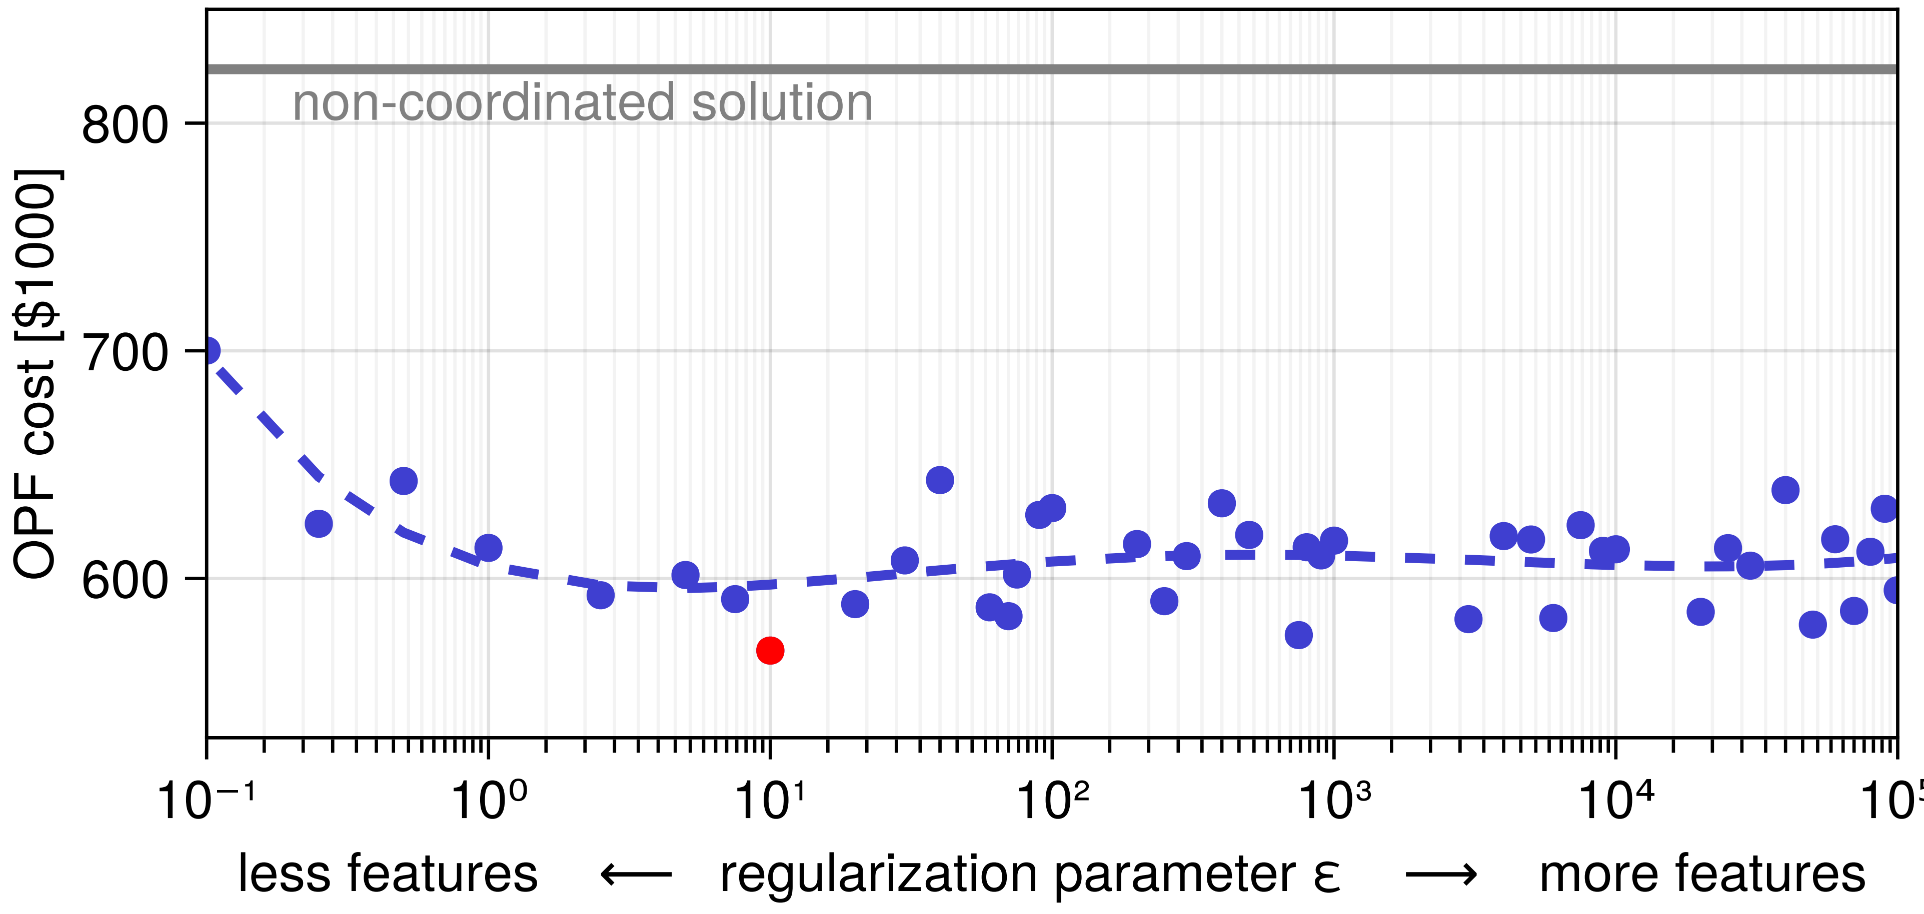

- Feature selection by means of ℓ_1-regularization
- ℓ_1-regularization also ensures coordination robustness
- Can we organize coordination using just one feature?

ε	# of features	zonal electricity demand											power flow												zonal electricity price											zonal renewable power output										
		d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}	d_{11}	f_1	f_2	f_3	f_4	f_5	f_6	λ_7	f_8	f_9	f_{10}	f_{11}	f_{12}	λ_1	λ_2	λ_3	λ_4	λ_5	λ_6	λ_7	λ_8	λ_9	λ_{10}	λ_{11}	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}	r_{11}
1000.0	29	○	●	●	●	○	●	○	●	●	○	○	●	●	●	●	●	●	●	●	●	●	●	○	●	●	●	●	●	●	●	○	●	●	○	○	●	●	○	○	○	○	○	○	○	●
100.0	28	●	●	●	●	●	●	○	●	●	●	○	●	●	●	●	●	●	●	●	○	●	○	○	○	●	●	●	●	○	○	●	●	●	●	○	●	○	○	○	○	○	○	○	○	●
10.0	24	●	●	●	●	○	●	●	●	●	●	●	●	●	●	●	○	●	●	●	●	●	○	○	○	○	○	○	○	○	○	●	●	●	○	○	●	○	○	○	○	○	○	○	○	●
5.0	20	●	●	○	●	●	●	●	●	●	●	●	●	●	●	○	●	○	●	●	●	○	○	○	○	○	○	○	○	○	○	○	●	●	○	○	●	○	○	○	○	○	○	○	○	○
2.5	13	●	○	●	○	○	●	●	●	●	●	○	●	●	○	○	○	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	●	●	○	○	○	○	○	○	○	○	○	○	○	○
1.0	6	●	○	●	○	○	○	○	○	○	●	○	●	●	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0.5	3	●	○	○	○	○	○	○	○	○	●	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0.1	1	○	○	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○

peak electricity demand in New York City

Market-based incentives to invoke spatial DC flexibility

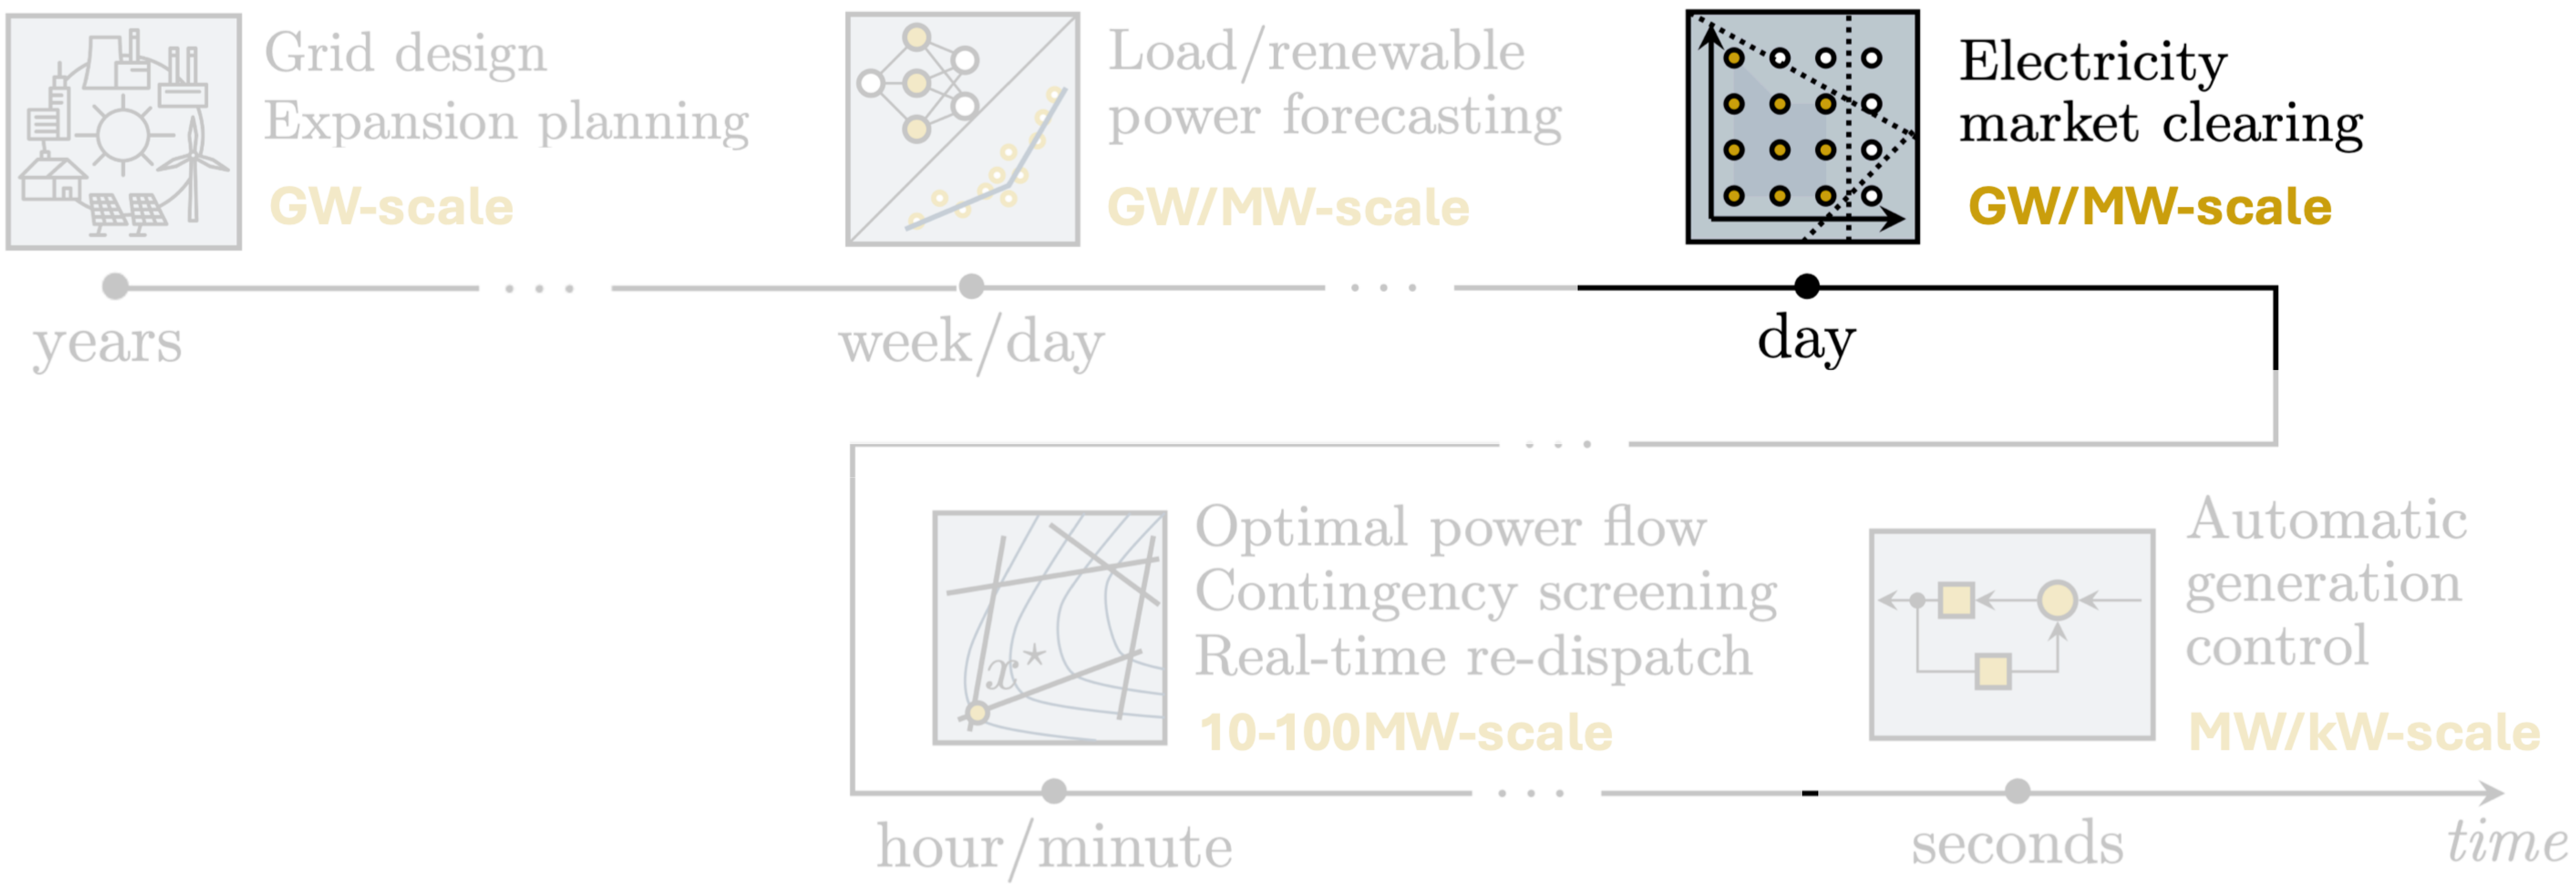

Electricity market is cleared by solving the **Optimal Power Flow (OPF)** problem:

$$
\begin{aligned}
\underset{\underline{\mathbf{p}}\leqslant\mathbf{p}\leqslant\overline{\mathbf{p}}}{\text{minimize}} \quad & \mathbf{p}^\top\mathbf{C}\mathbf{p}+\mathbf{c}^\top\mathbf{p} && \textit{generator dispatch cost}\\
\text{subject to} \quad & \mathbf{1}^\top(\mathbf{p}-\mathbf{d})=0 \,:\, \boldsymbol{\lambda}, && \textit{power balance condition}\\
& |\mathbf{F}(\mathbf{p}-\mathbf{d})|\leqslant\overline{\mathbf{f}} \,:\, \boldsymbol{\lambda}_{\overline{\mathbf{f}}},\boldsymbol{\lambda}_{\underline{\mathbf{f}}}, && \textit{power flow limits}
\end{aligned}
$$

Electricity market is cleared by solving the **Optimal Power Flow (OPF)** problem:

$$
\begin{aligned}
\underset{\underline{\mathbf{p}}\leqslant\mathbf{p}\leqslant\overline{\mathbf{p}}}{\text{minimize}} \quad & \mathbf{p}^\top\mathbf{C}\mathbf{p}+\mathbf{c}^\top\mathbf{p} && \textit{generator dispatch cost}\\
\text{subject to} \quad & \mathbf{1}^\top(\mathbf{p}-\mathbf{d})=0\;:\;\boldsymbol{\lambda}, && \textit{power balance condition}\\
& |\mathbf{F}(\mathbf{p}-\mathbf{d})|\leqslant\overline{\mathbf{f}}\;\;:\;\boldsymbol{\lambda}_{\overline{\mathbf{f}}},\boldsymbol{\lambda}_{\underline{\mathbf{f}}}, && \textit{power flow limits}
\end{aligned}
$$

Location marginal prices (LMPs) are derived from the dual solution:

$$
\boldsymbol{\pi}(\mathbf{d}) = \underbrace{\boldsymbol{\lambda}_b^\star\cdot\mathbf{1}}_{\text{uniform price}} - \underbrace{\mathbf{F}^\top(\boldsymbol{\lambda}_{\overline{\mathbf{f}}}^\star-\boldsymbol{\lambda}_{\underline{\mathbf{f}}}^\star)}_{\text{adjustment due to congestion}}
$$

Electricity market is cleared by solving the **Optimal Power Flow (OPF)** problem:

$$\begin{aligned}
\underset{\underline{\mathbf{p}}\leqslant\mathbf{p}\leqslant\overline{\mathbf{p}}}{\text{minimize}} \quad & \mathbf{p}^\top\mathbf{C}\mathbf{p}+\mathbf{c}^\top\mathbf{p} && \textit{generator dispatch cost}\\
\text{subject to} \quad & \mathbf{1}^\top(\mathbf{p}-\mathbf{d})=0\ :\ \boldsymbol{\lambda}, && \textit{power balance condition}\\
& |\mathbf{F}(\mathbf{p}-\mathbf{d})|\leqslant\overline{\mathbf{f}}\ :\ \boldsymbol{\lambda}_{\overline{\mathbf{f}}},\boldsymbol{\lambda}_{\underline{\mathbf{f}}}, && \textit{power flow limits}
\end{aligned}$$

Location marginal prices (LMPs) are derived from the dual solution:

$$\boldsymbol{\pi}(\mathbf{d}) = \underbrace{\boldsymbol{\lambda}_b^\star\cdot\mathbf{1}}_{\text{uniform price}} - \underbrace{\mathbf{F}^\top(\boldsymbol{\lambda}_{\overline{\mathbf{f}}}^\star-\boldsymbol{\lambda}_{\underline{\mathbf{f}}}^\star)}_{\text{adjustment due to congestion}}$$

By duality theory, the electricity cost of power consumption is proportional to grid dispatch costs:

$$\mathbf{p}^{\star\top}\mathbf{C}\mathbf{p}^\star+\mathbf{c}^\top\mathbf{p}^\star \propto \boldsymbol{\pi}(\mathbf{d})^\top\mathbf{d}$$

$\Longrightarrow$ financial incentives for DCs to allocate electricity demands in the least-cost way for the power grid

- Standard IEEE 118-Bus test system
- 10 DCs with 1GW of peak demand
- LMPs significantly vary across the grid

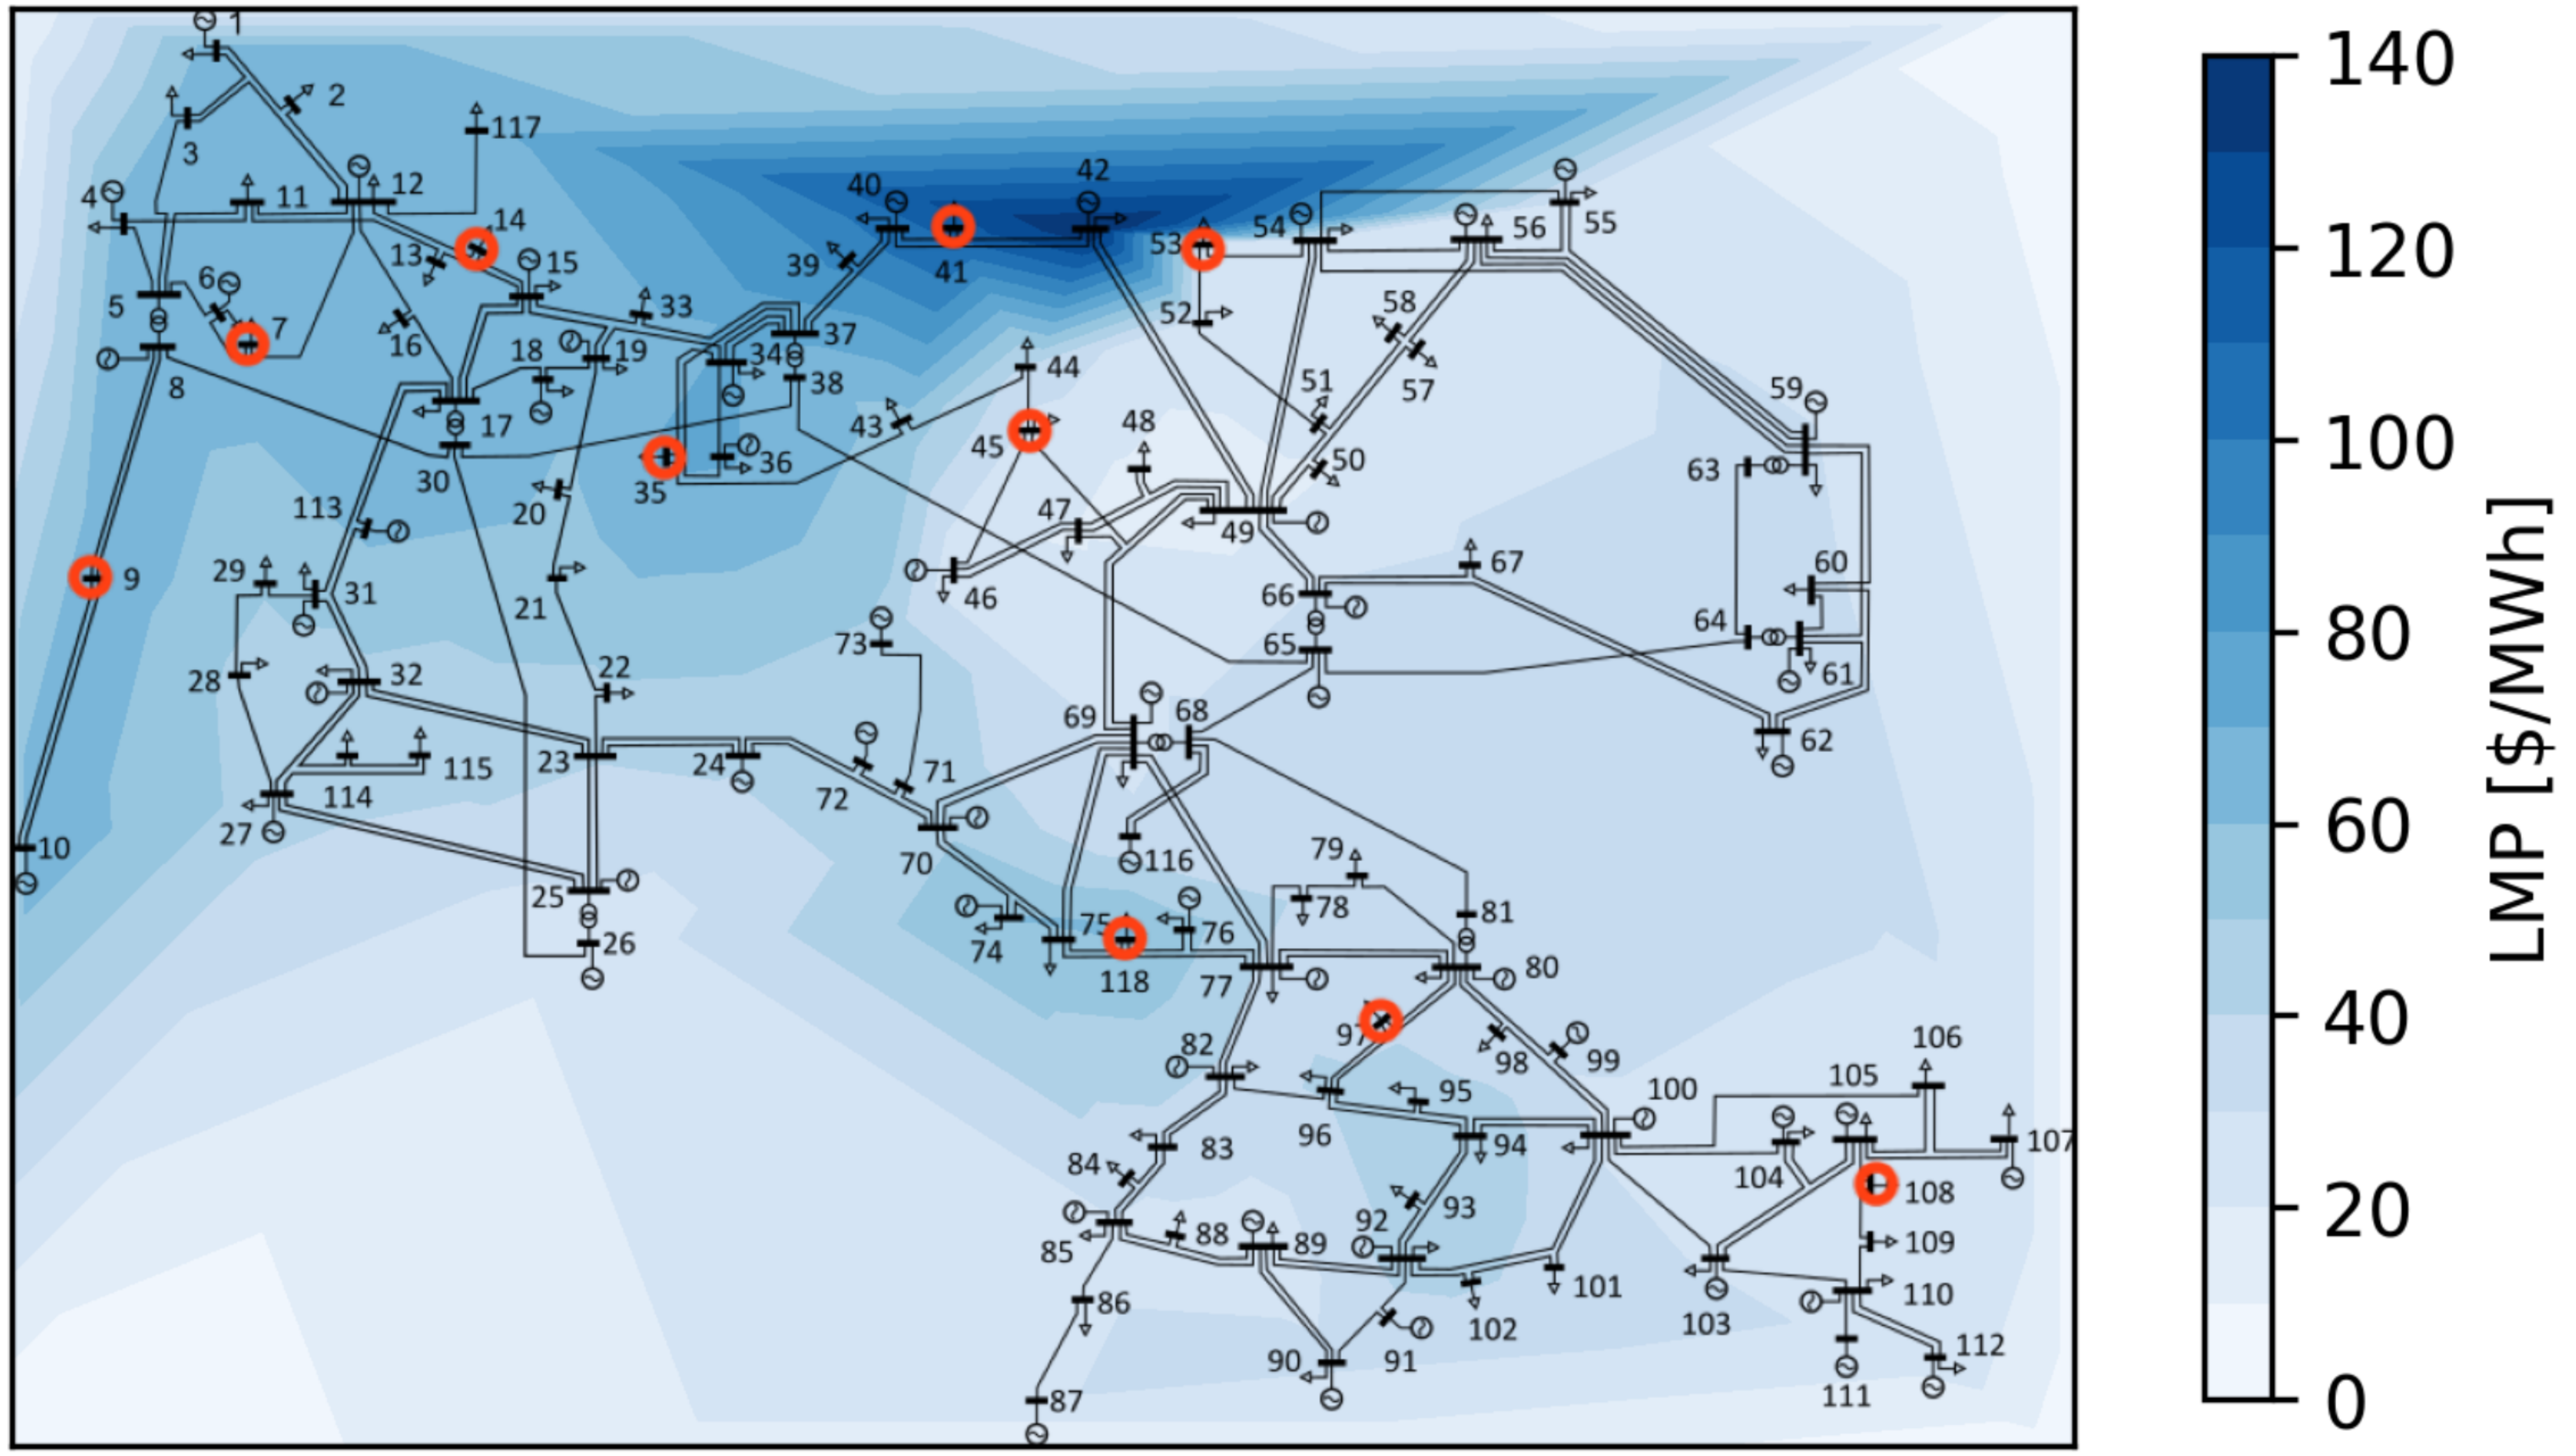

Clear incentives to shift DC loads from costly to cheaper nodes

Bilevel optimization formulation:

$$\begin{aligned}
\underset{\mathbf{d}}{\text{minimize}} \quad & \boldsymbol{\pi}^\top \mathbf{d} \\
\text{subject to} \quad & \underline{\mathbf{d}} \leqslant \mathbf{d} \leqslant \overline{\mathbf{d}} \\
& \mathbf{1}^\top \mathbf{d} = \Delta \\
& \boldsymbol{\pi} \in \text{dual sol. of OPF}
\end{aligned}$$

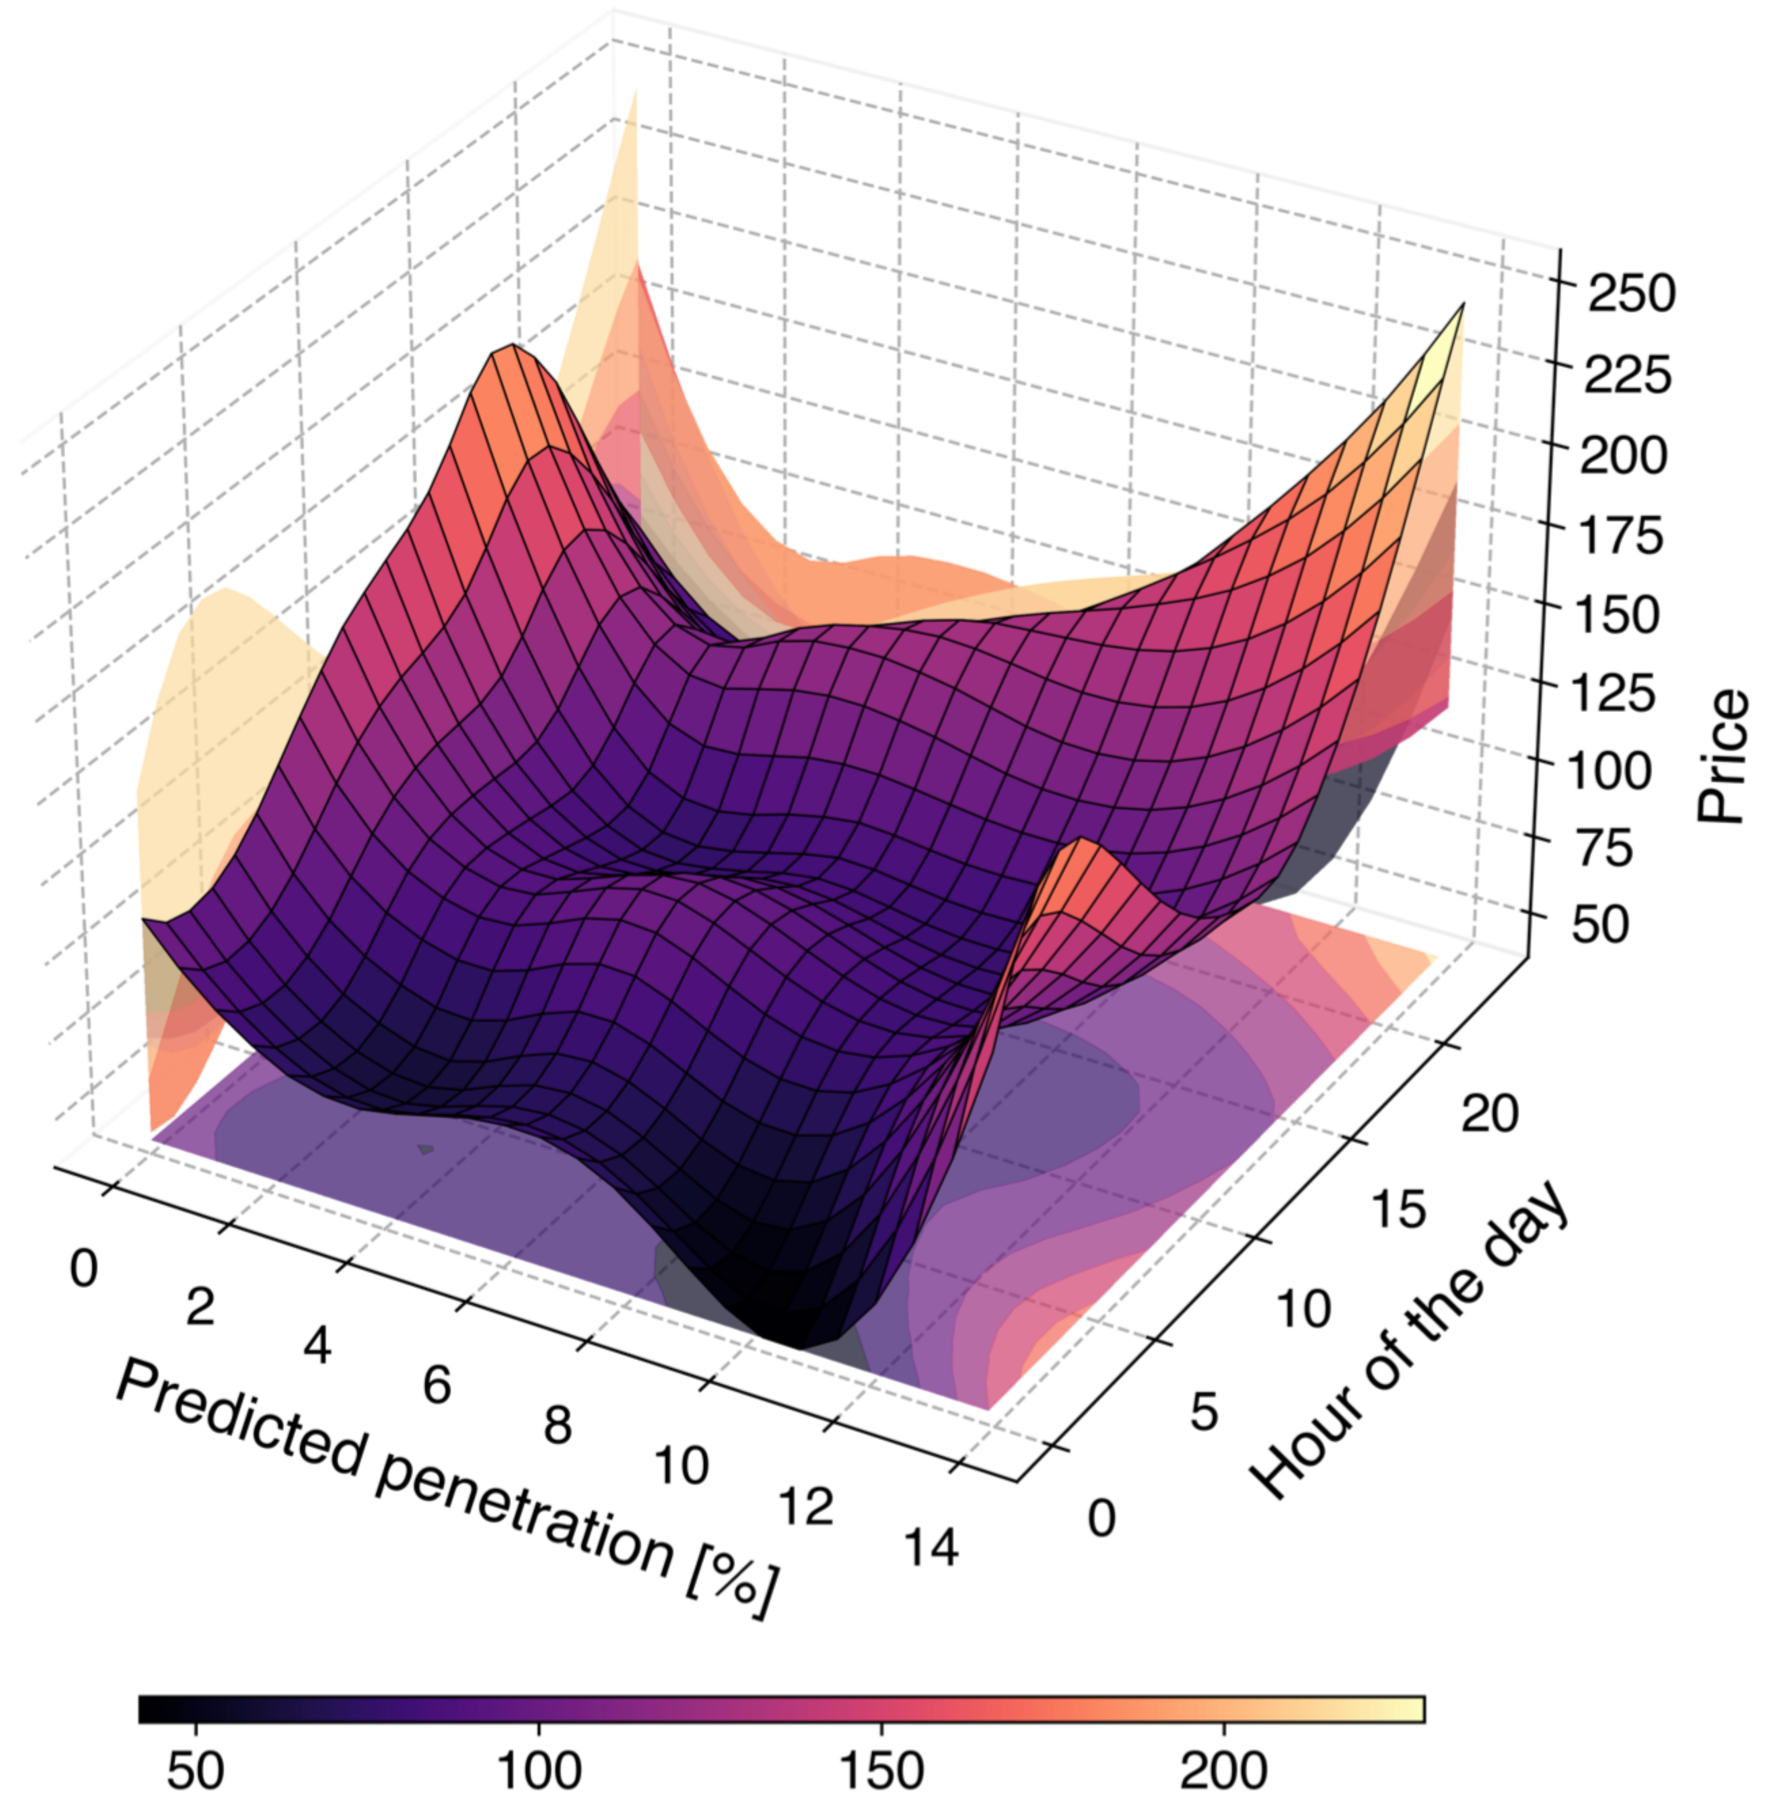

[Cacciarelli *et al.*'25]

Existing literature: Bilevel optimization of DC demand payment

- Requires the full knowledge of OPF (market) optimization
- The actual market engine can be more complicated then that

X Liu, V Dvorkin. Optimization over trained neural networks: Difference-of-convex algorithm and application to data center scheduling. 2025

Strategic participation of DCs in electricity markets

Bilevel optimization formulation:

$$\begin{aligned} \underset{\mathbf{d}}{\text{minimize}} \quad & \boldsymbol{\pi}^\top \mathbf{d} \\ \text{subject to} \quad & \underline{\mathbf{d}} \leqslant \mathbf{d} \leqslant \overline{\mathbf{d}} \\ & \mathbf{1}^\top \mathbf{d} = \Delta \\ & \boldsymbol{\pi} \in \text{dual sol. of OPF} \end{aligned}$$

Neural network approximation:

$$\begin{aligned} \underset{\mathbf{d}}{\text{minimize}} \quad & \mathbb{NN}(\mathbf{d}) \\ \text{subject to} \quad & \underline{\mathbf{d}} \leqslant \mathbf{d} \leqslant \overline{\mathbf{d}} \\ & \mathbf{1}^\top \mathbf{d} = \Delta \end{aligned}$$

Existing literature: Bilevel optimization of DC demand payment

- Requires the full knowledge of OPF (market) optimization
- The actual market engine can be more complicated then that

Proposed solution: Approximate the payment function using neural network

- Optimization over trained neural network $\mathbb{NN}$: DC loads $\mapsto$ Demand charges
- Efficient difference-of-convex algorithm to solve the problem
- All necessary data is readily available from ISO dashboards and modern analytics startups

X Liu, V Dvorkin. Optimization over trained neural networks: Difference-of-convex algorithm and application to data center scheduling. 2025

Reducing congestion in IEEE 118-Bus system

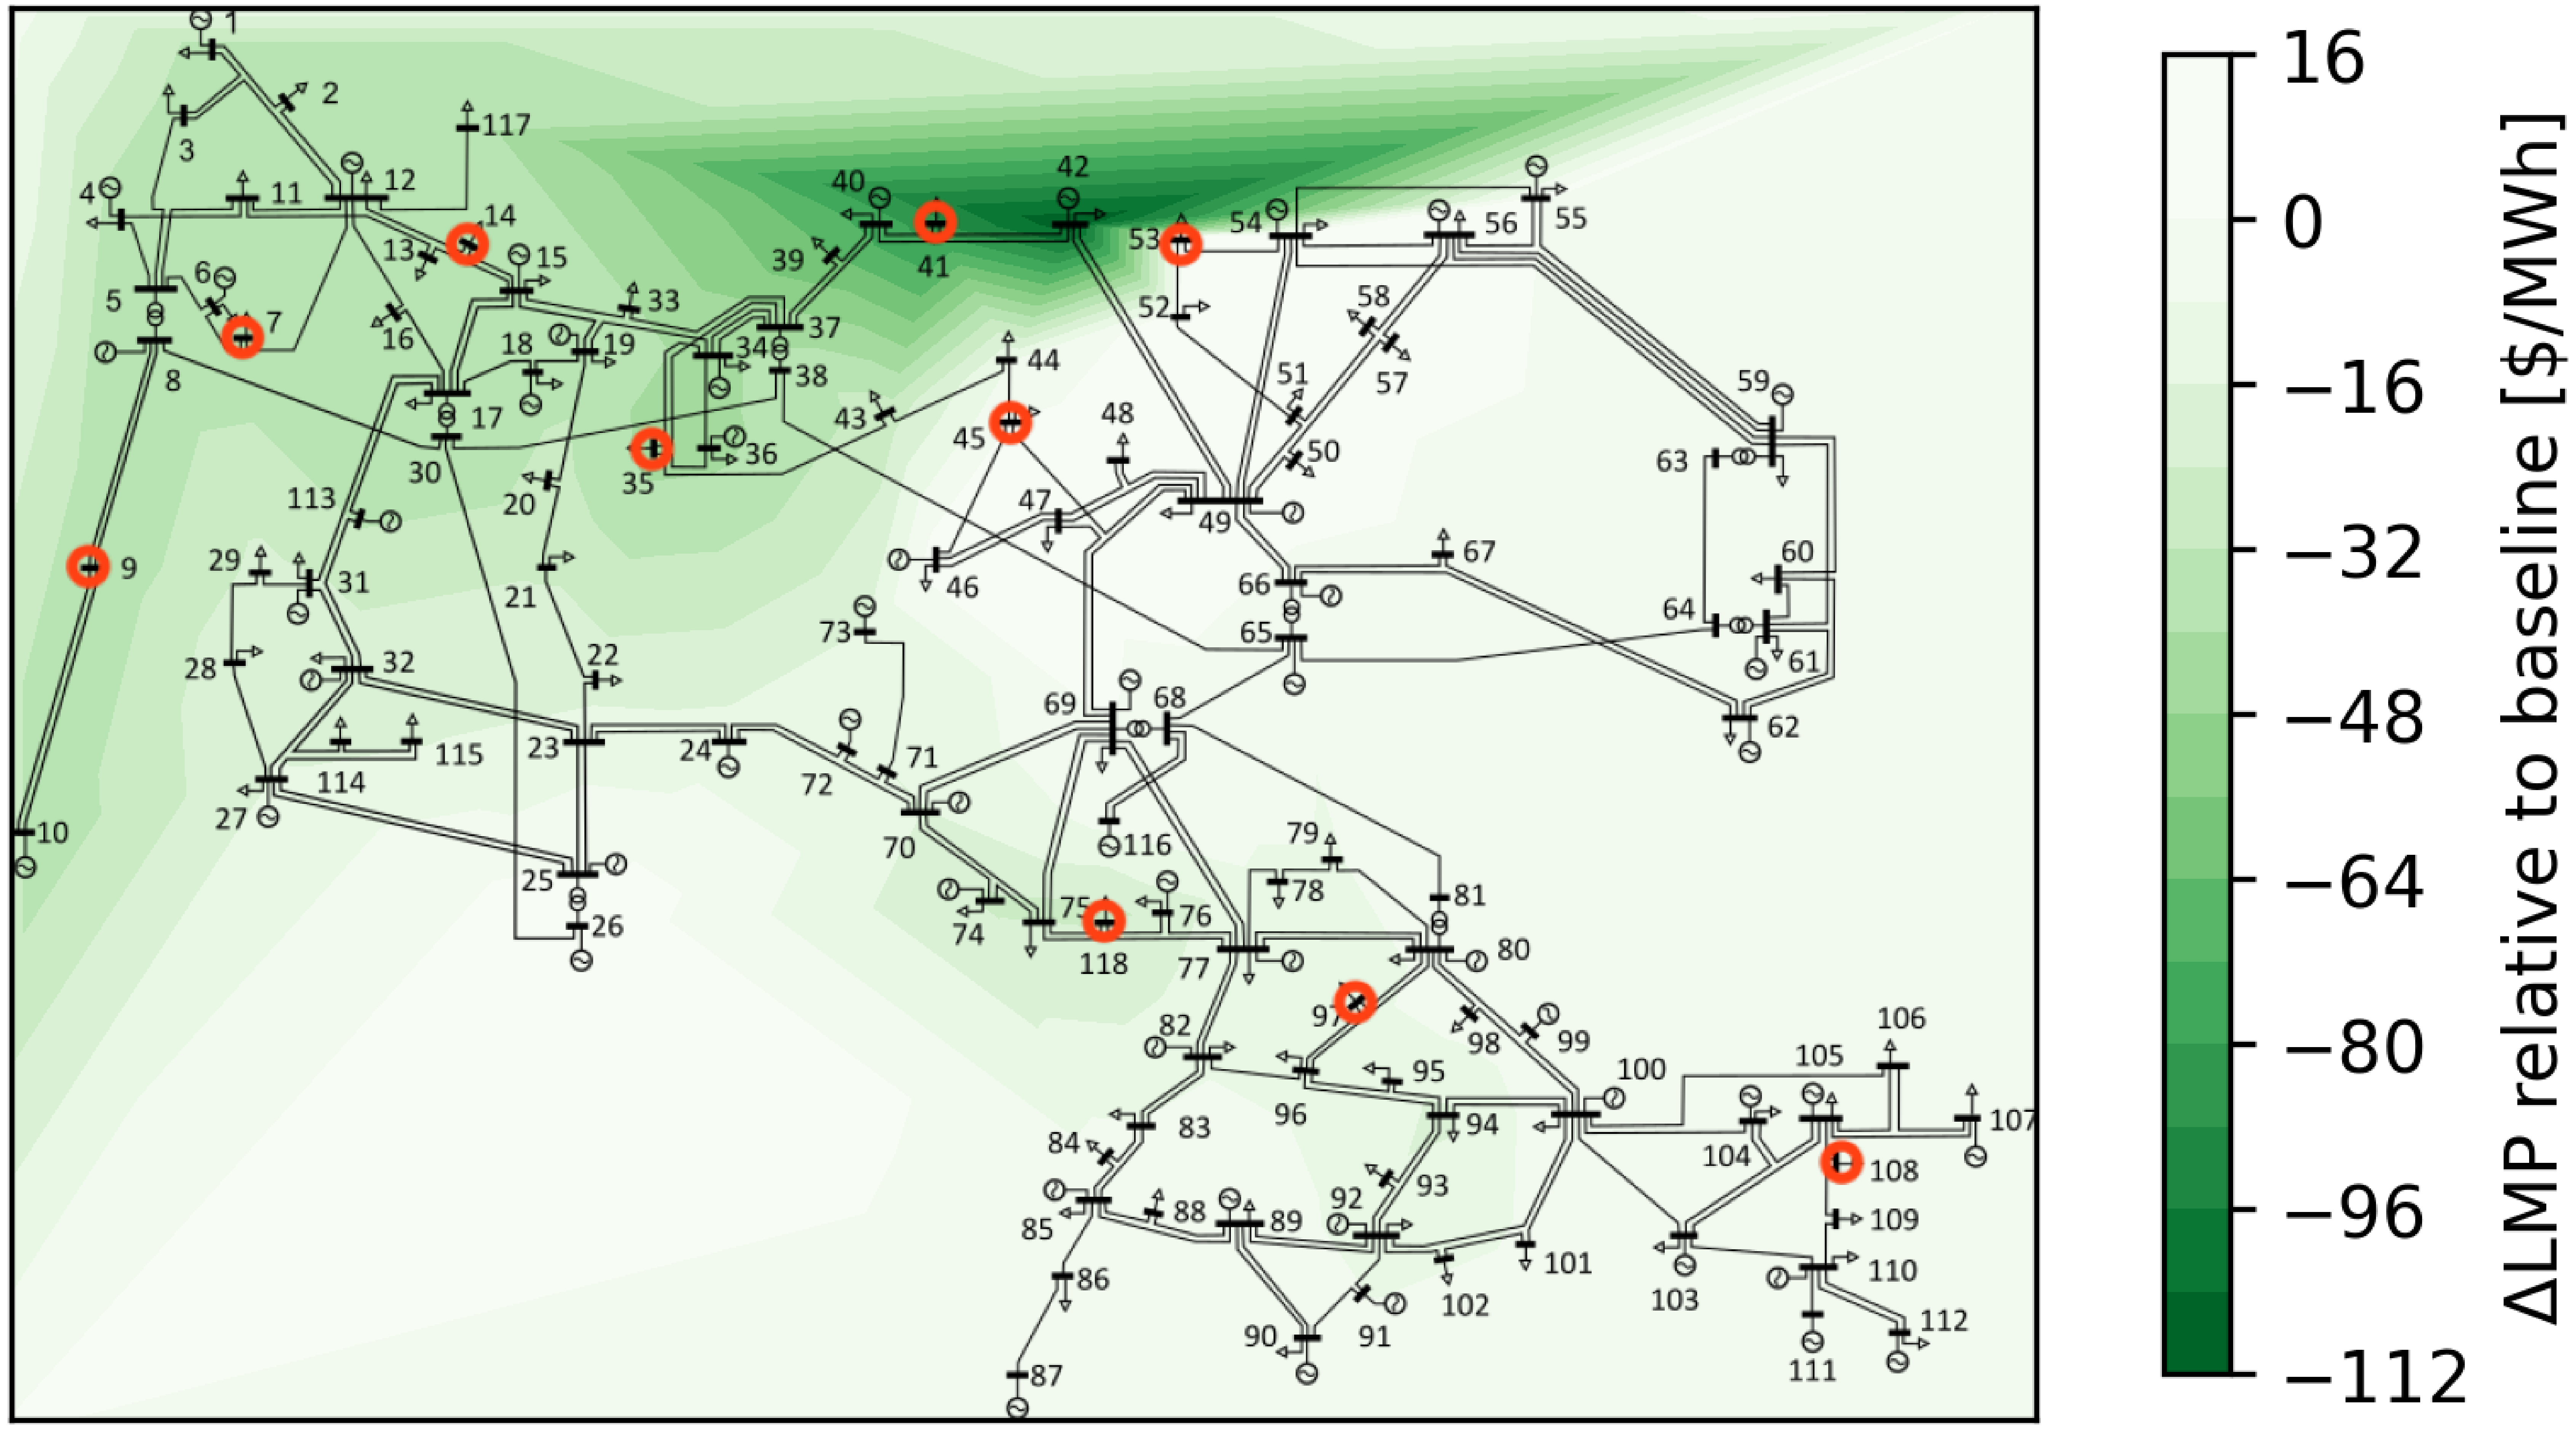

- The difference of LMPs in the baseline and strategic solutions
- LMPs were significantly reduced at the data center nodes.

Cost savings in IEEE 118-Bus system

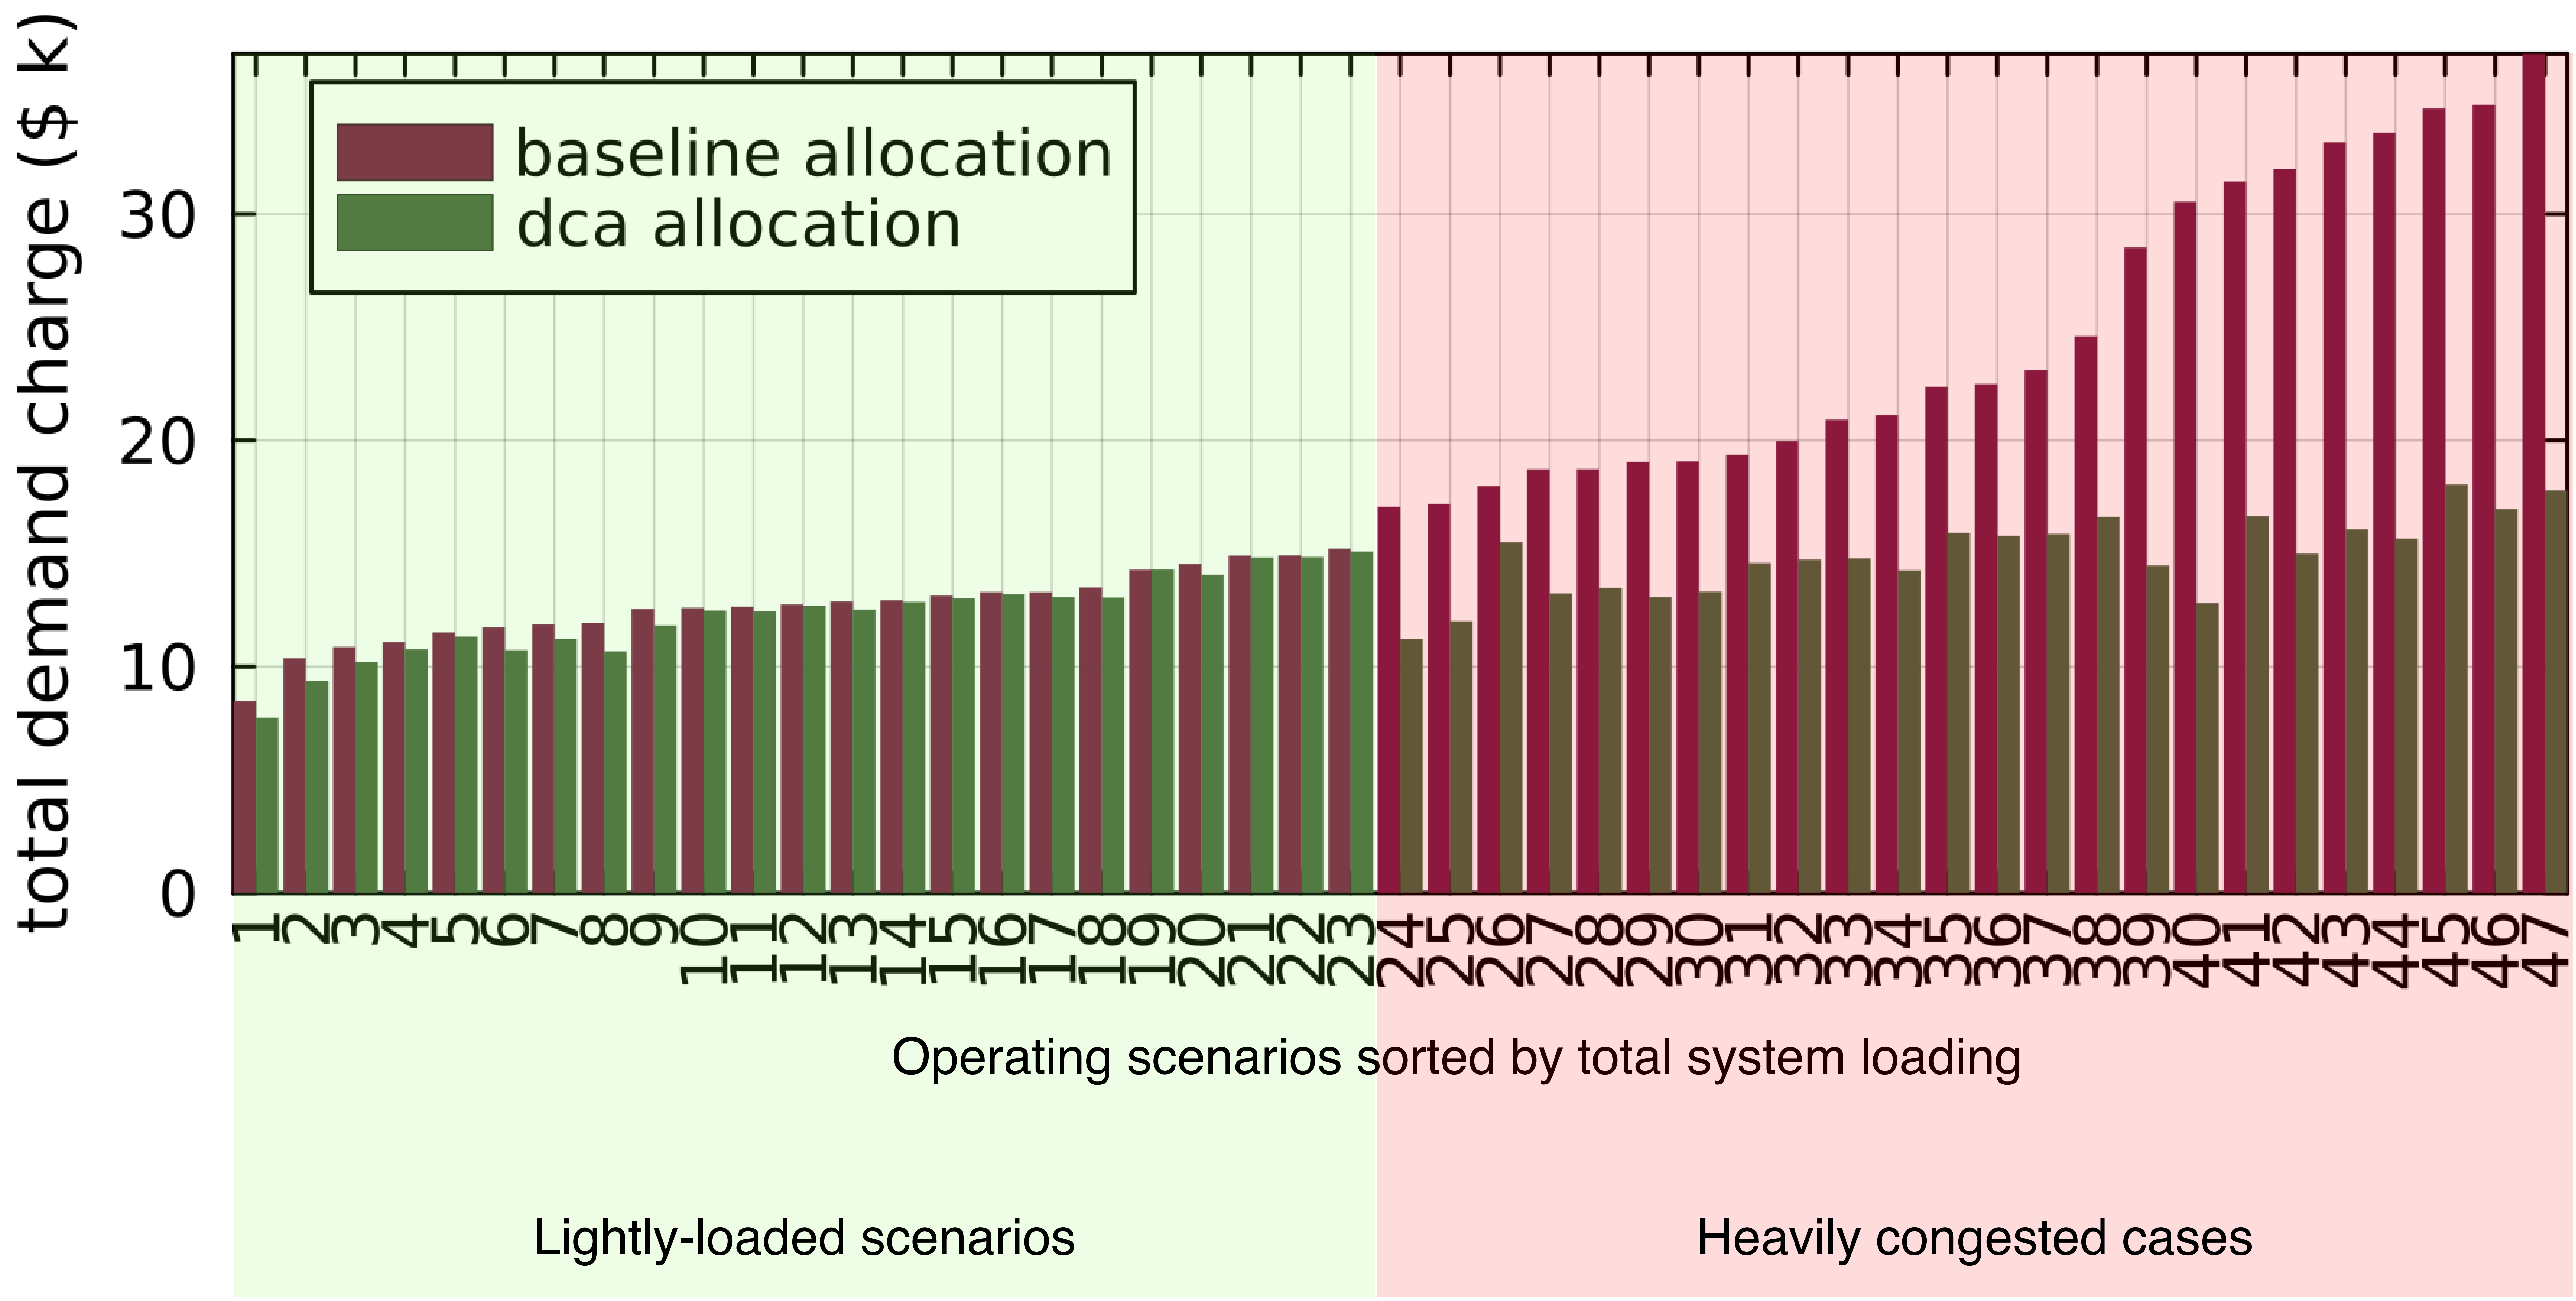

Distribution-level energy storage solutions to smooth DC loads

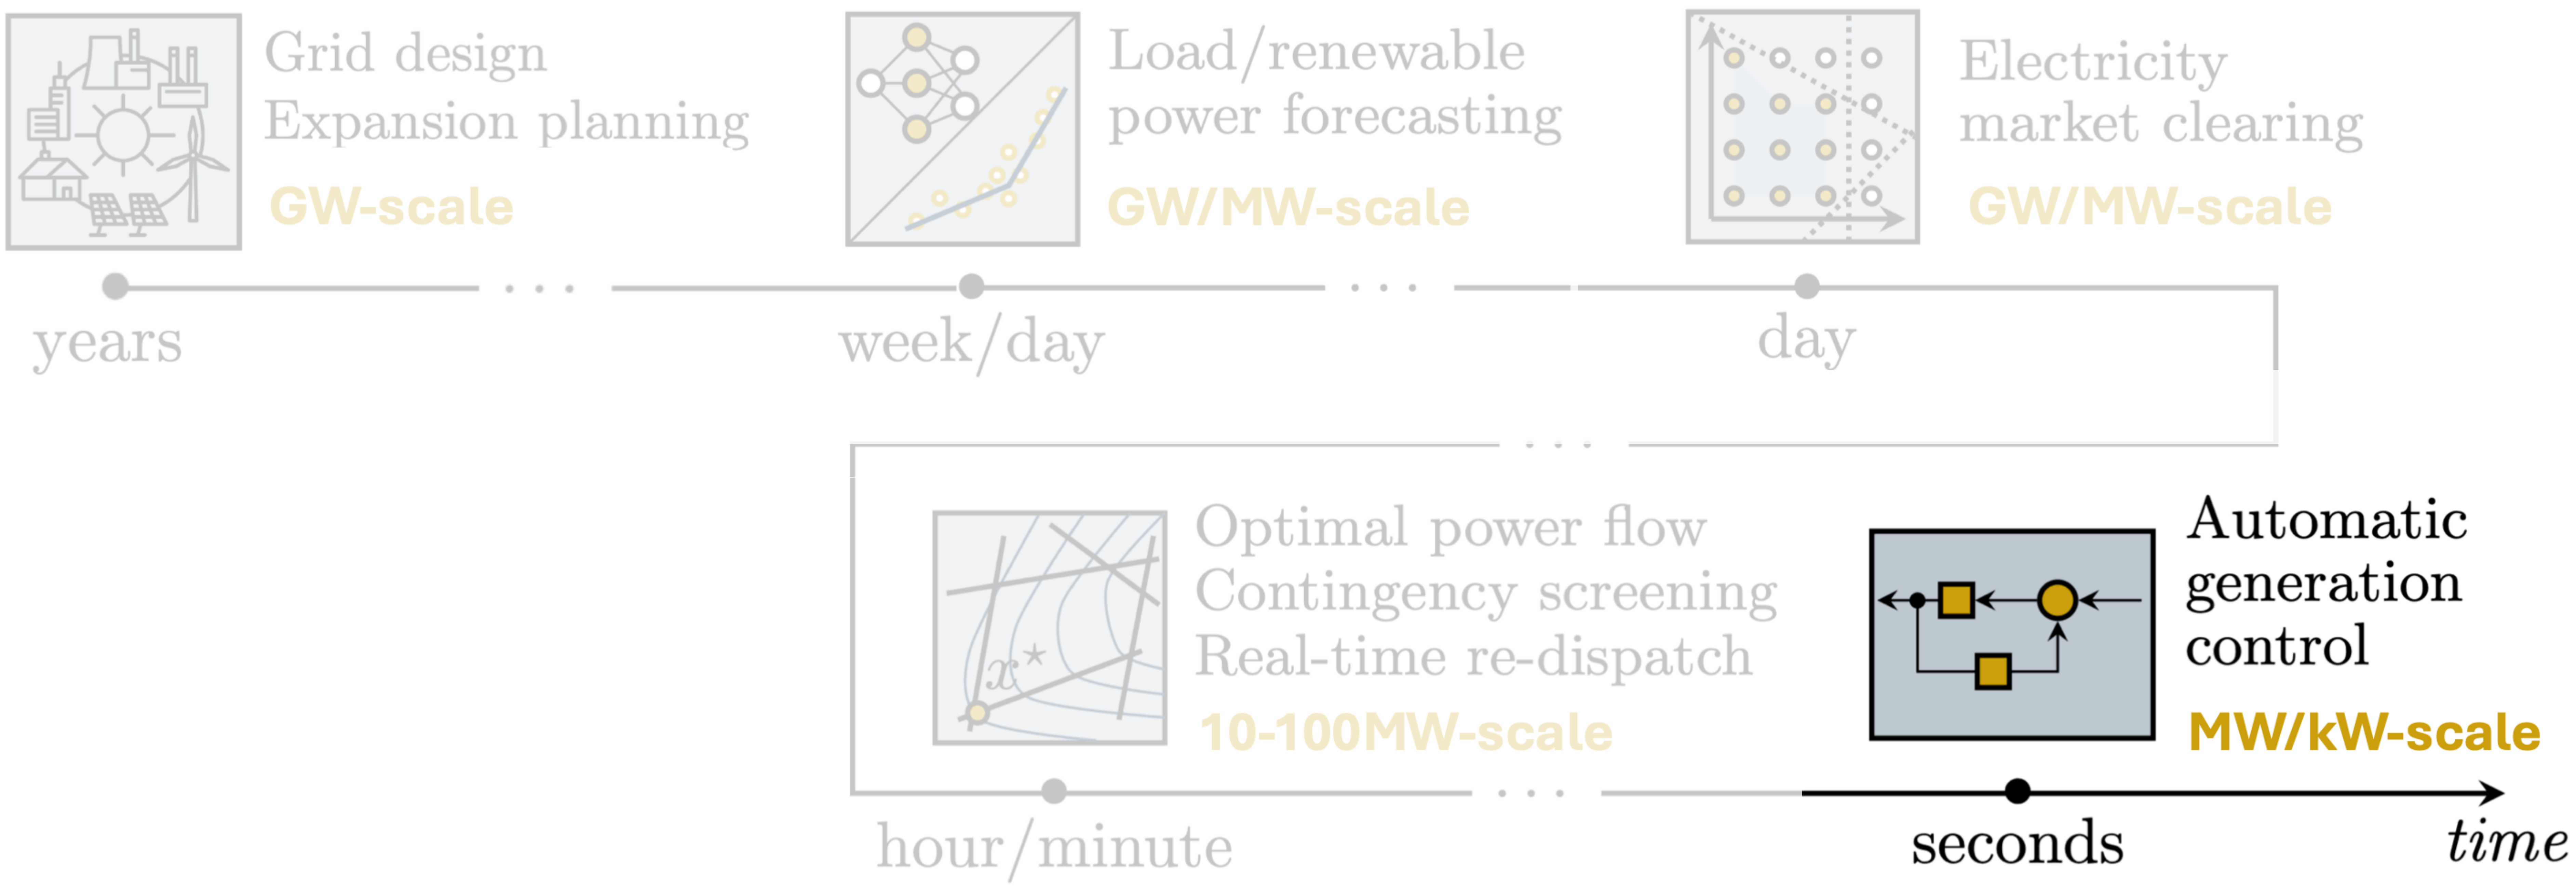

Disruptive DC loads lead to voltage issues in distribution systems

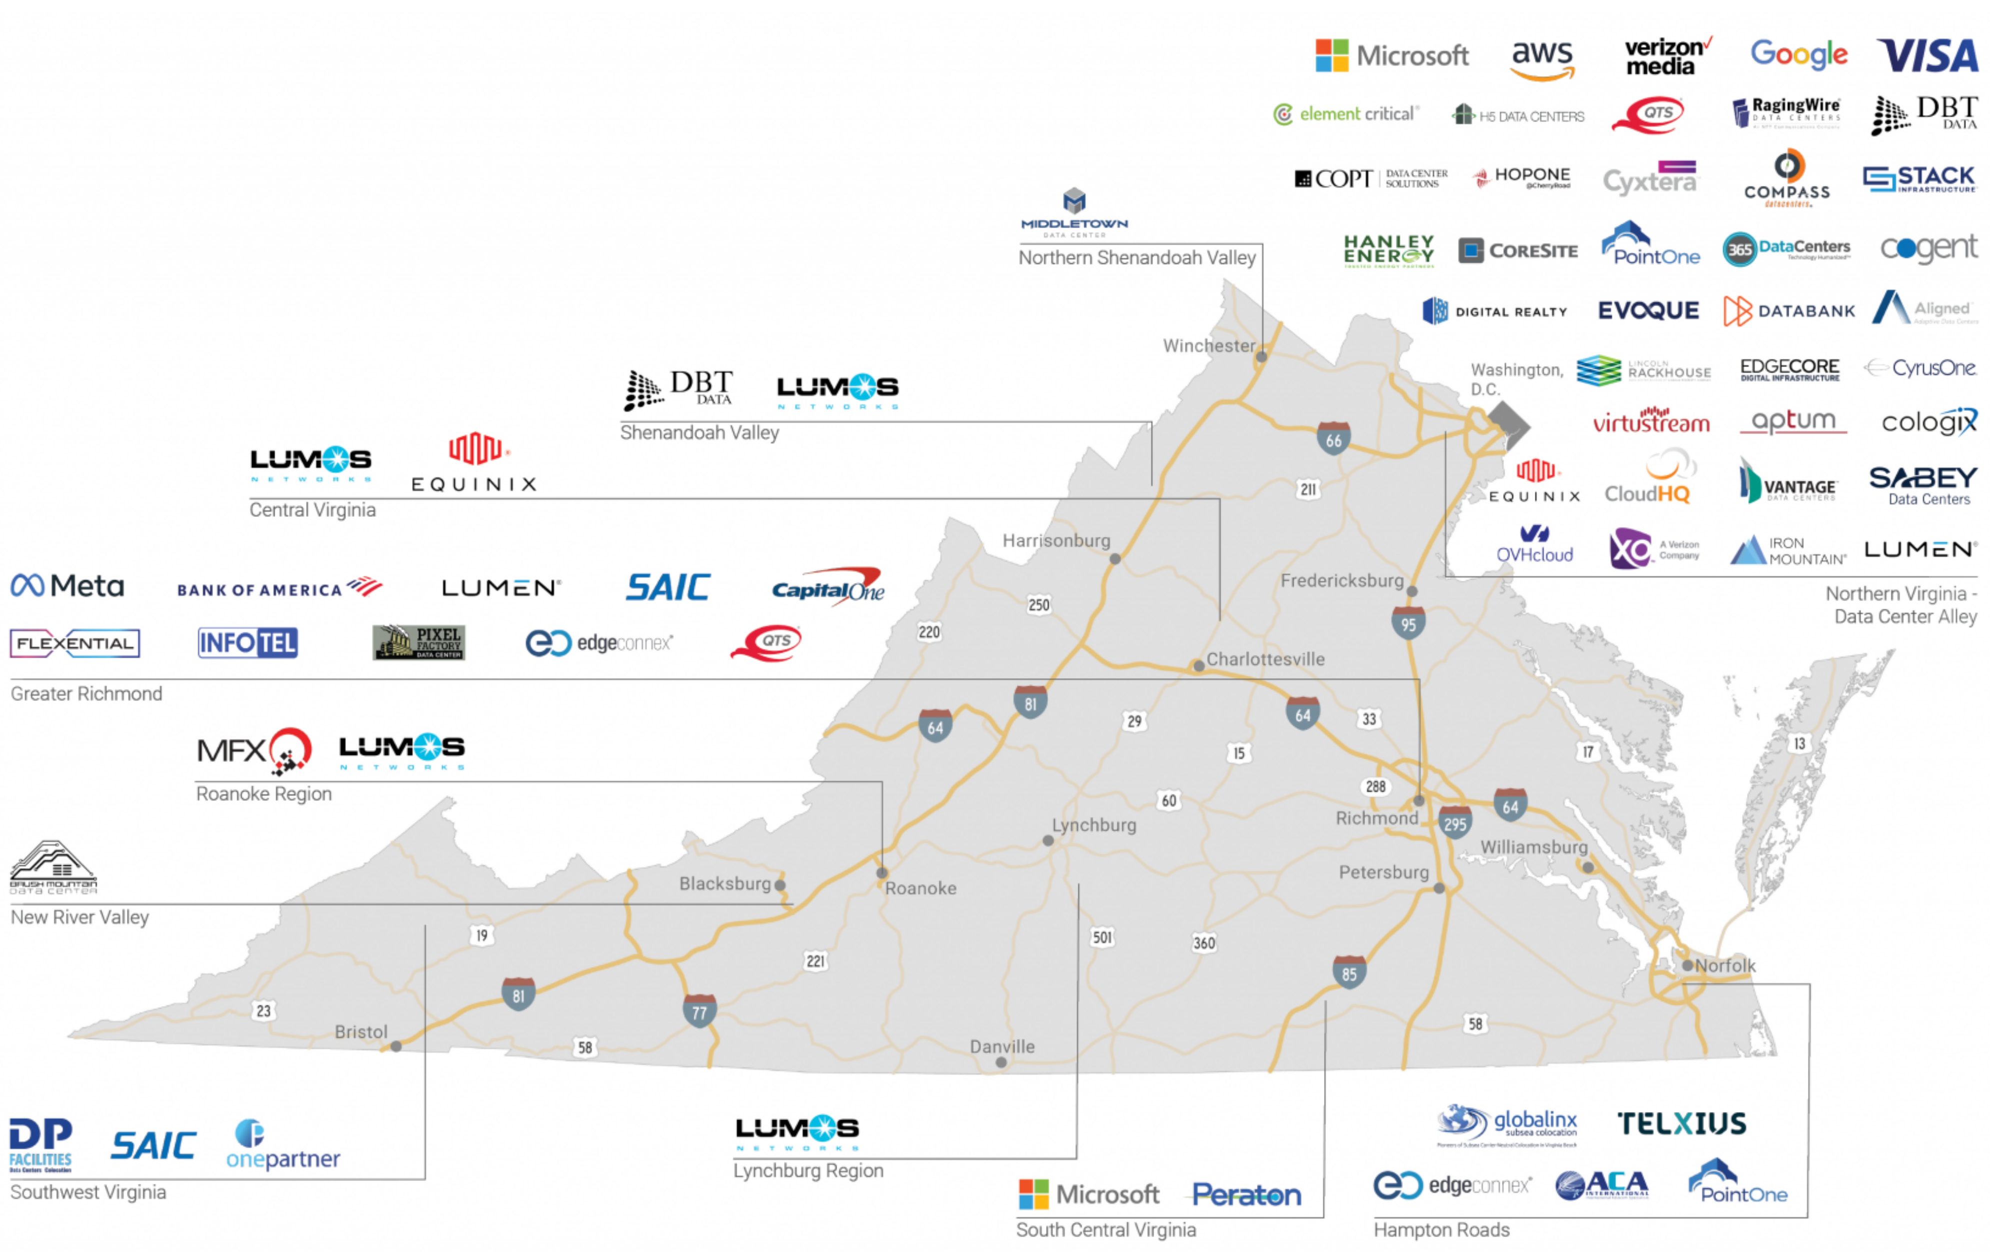

Growth of data centers in Virginia

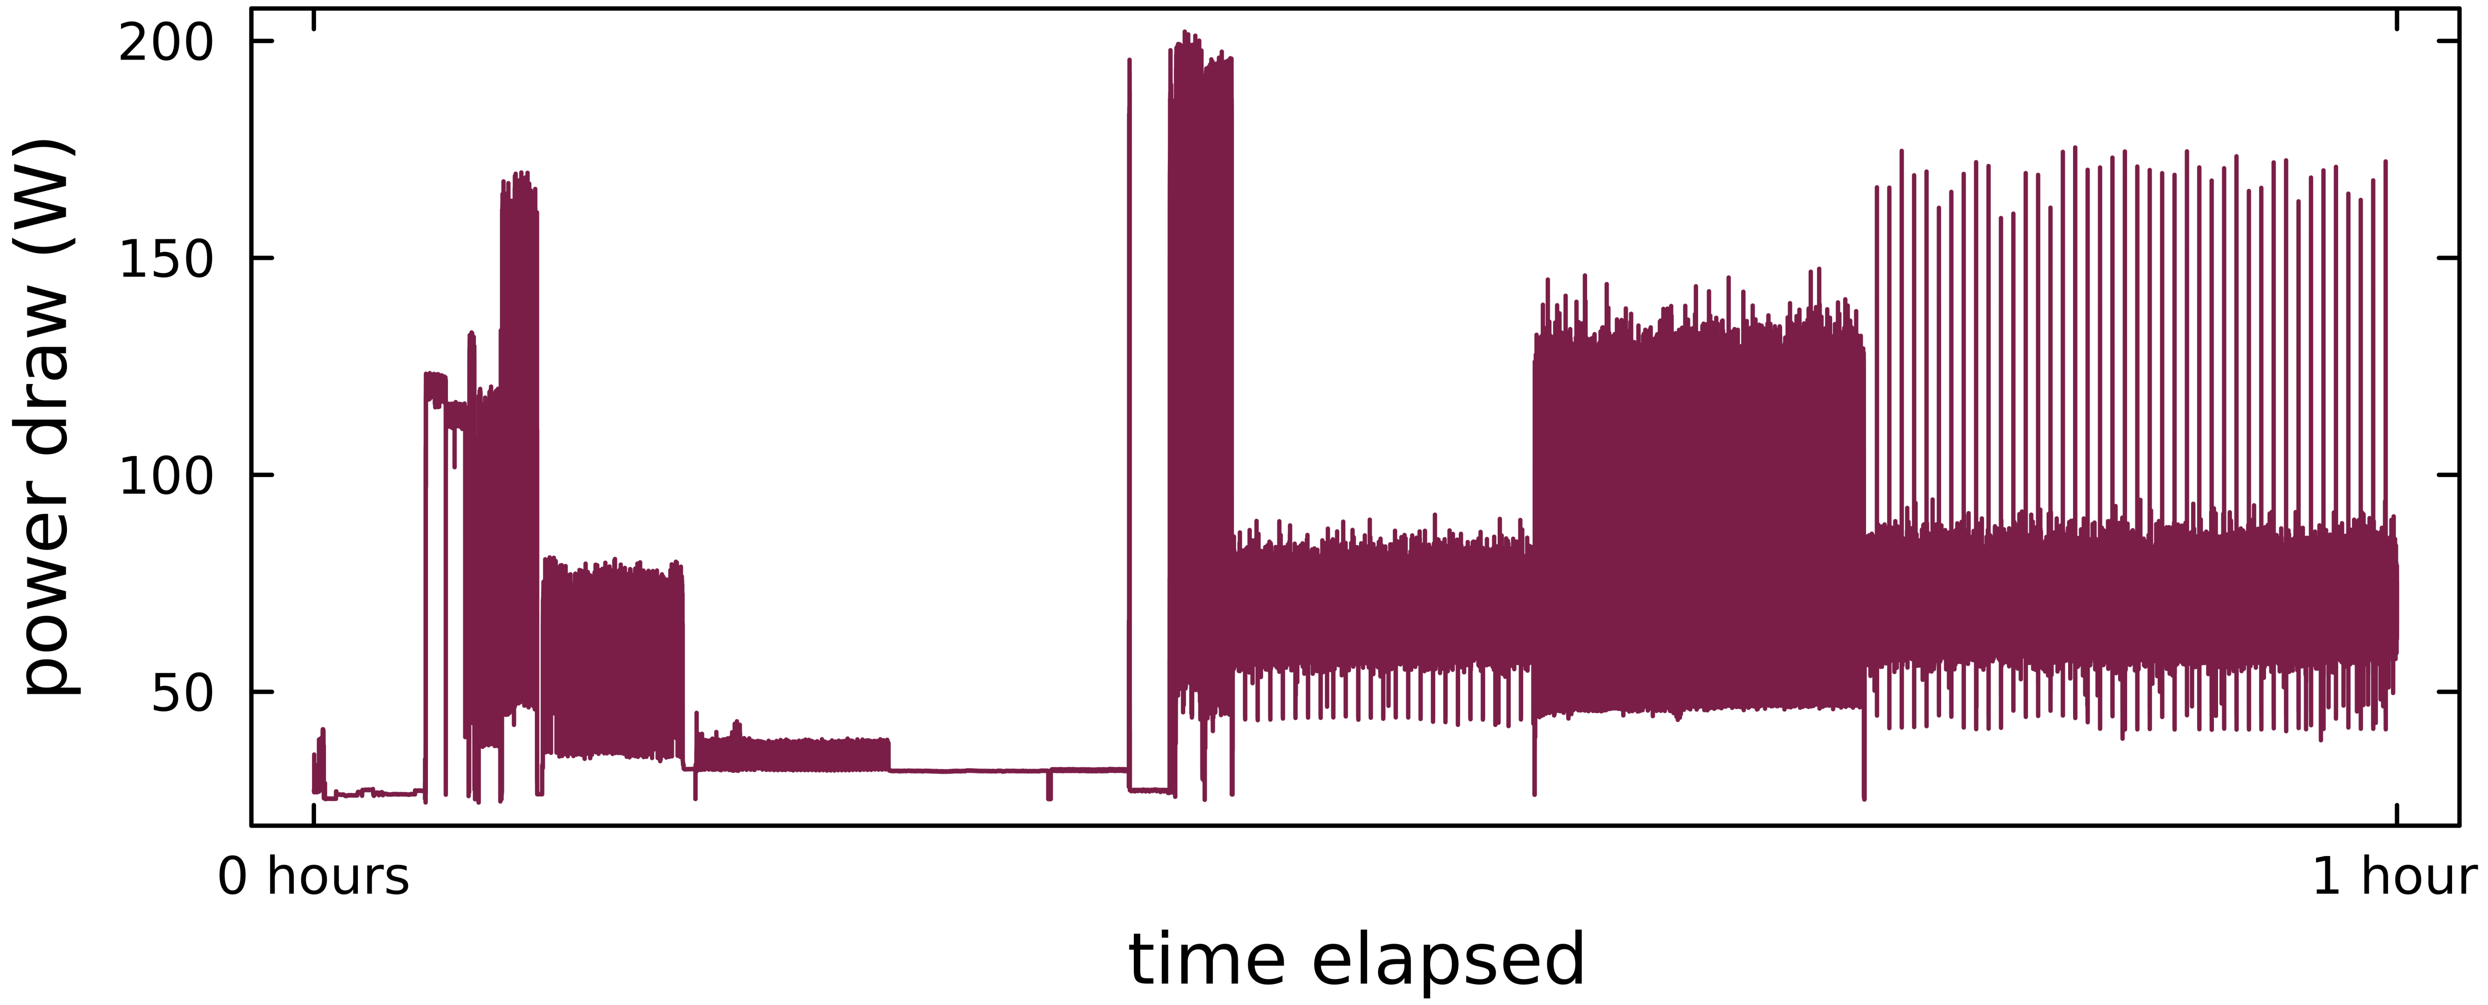

Growth of data centers in Virginia

Smoothing DC loads with energy storage

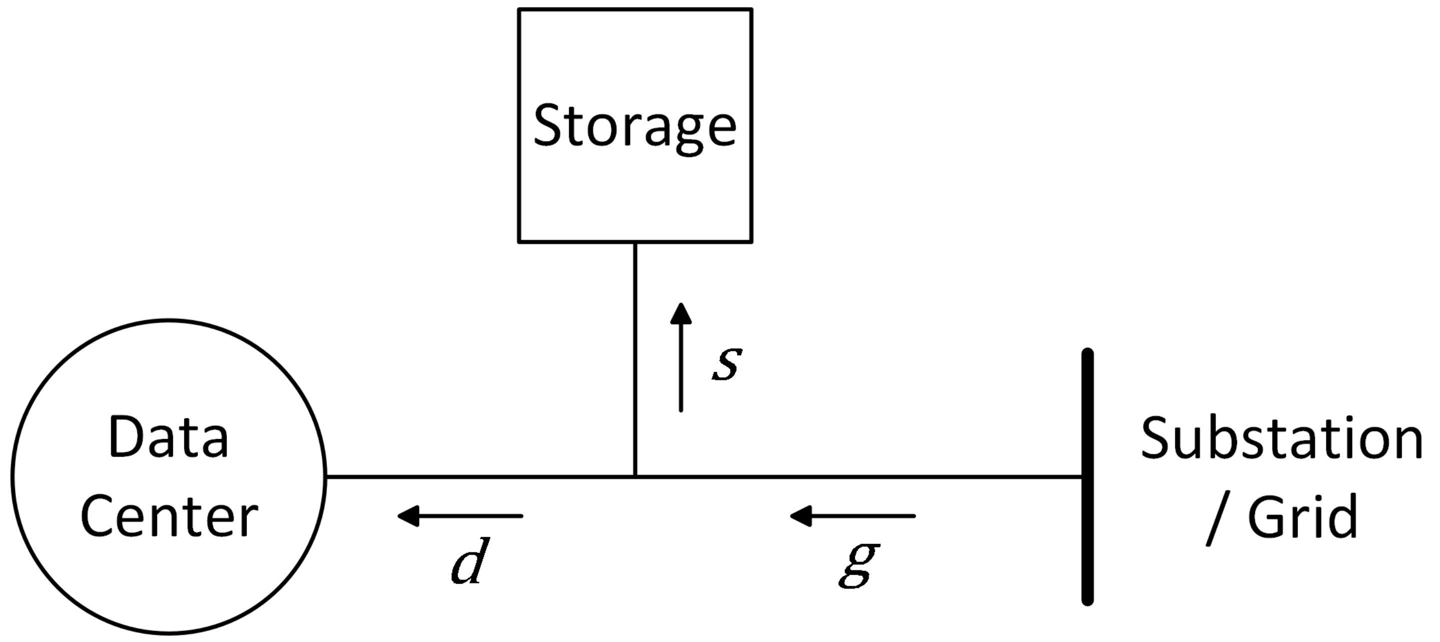

- **Goal:** Smooth the power profile observed by the gird/substation
- **Drop-control strategy** acting solely on measurements
- **Feedback optimization strategy** acting on measurements **and** the intensity of constraint violations (mixed-saddle flow)

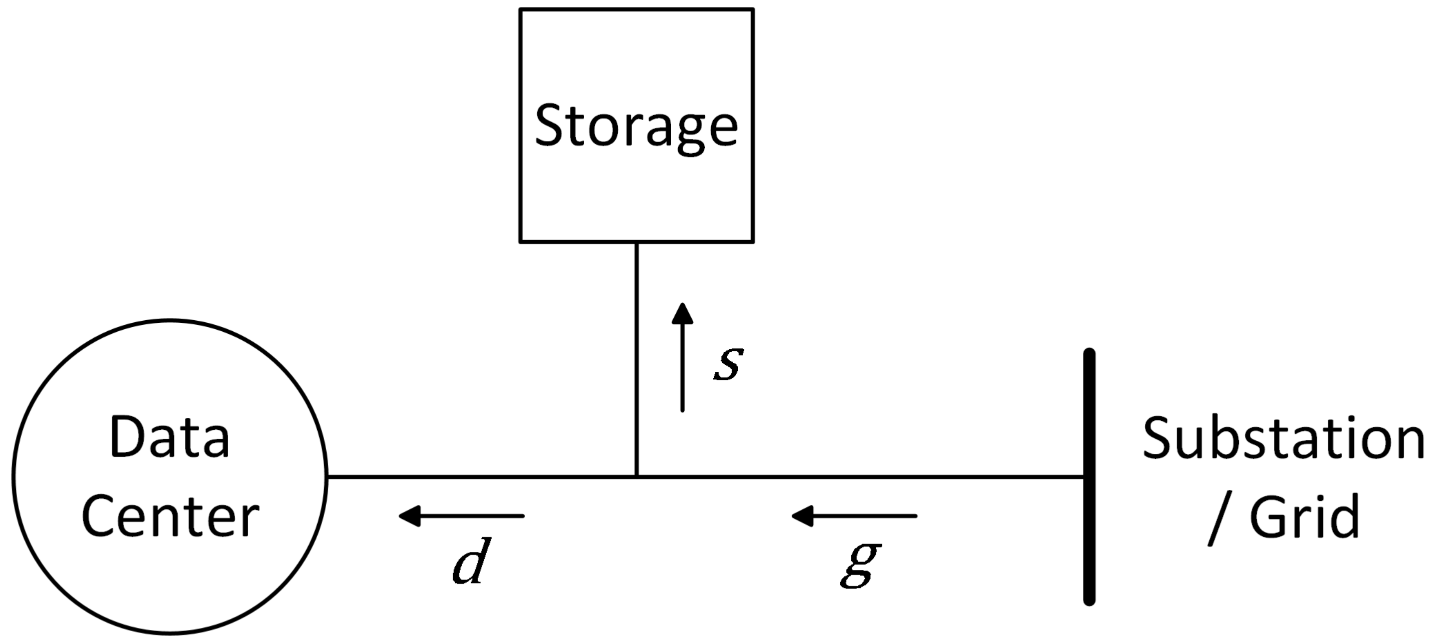

- **Goal:** Smooth the power profile observed by the gird/substation
- **Drop-control strategy** acting solely on measurements
- **Feedback optimization strategy** acting on measurements **and** the intensity of constraint violations (mixed-saddle flow)

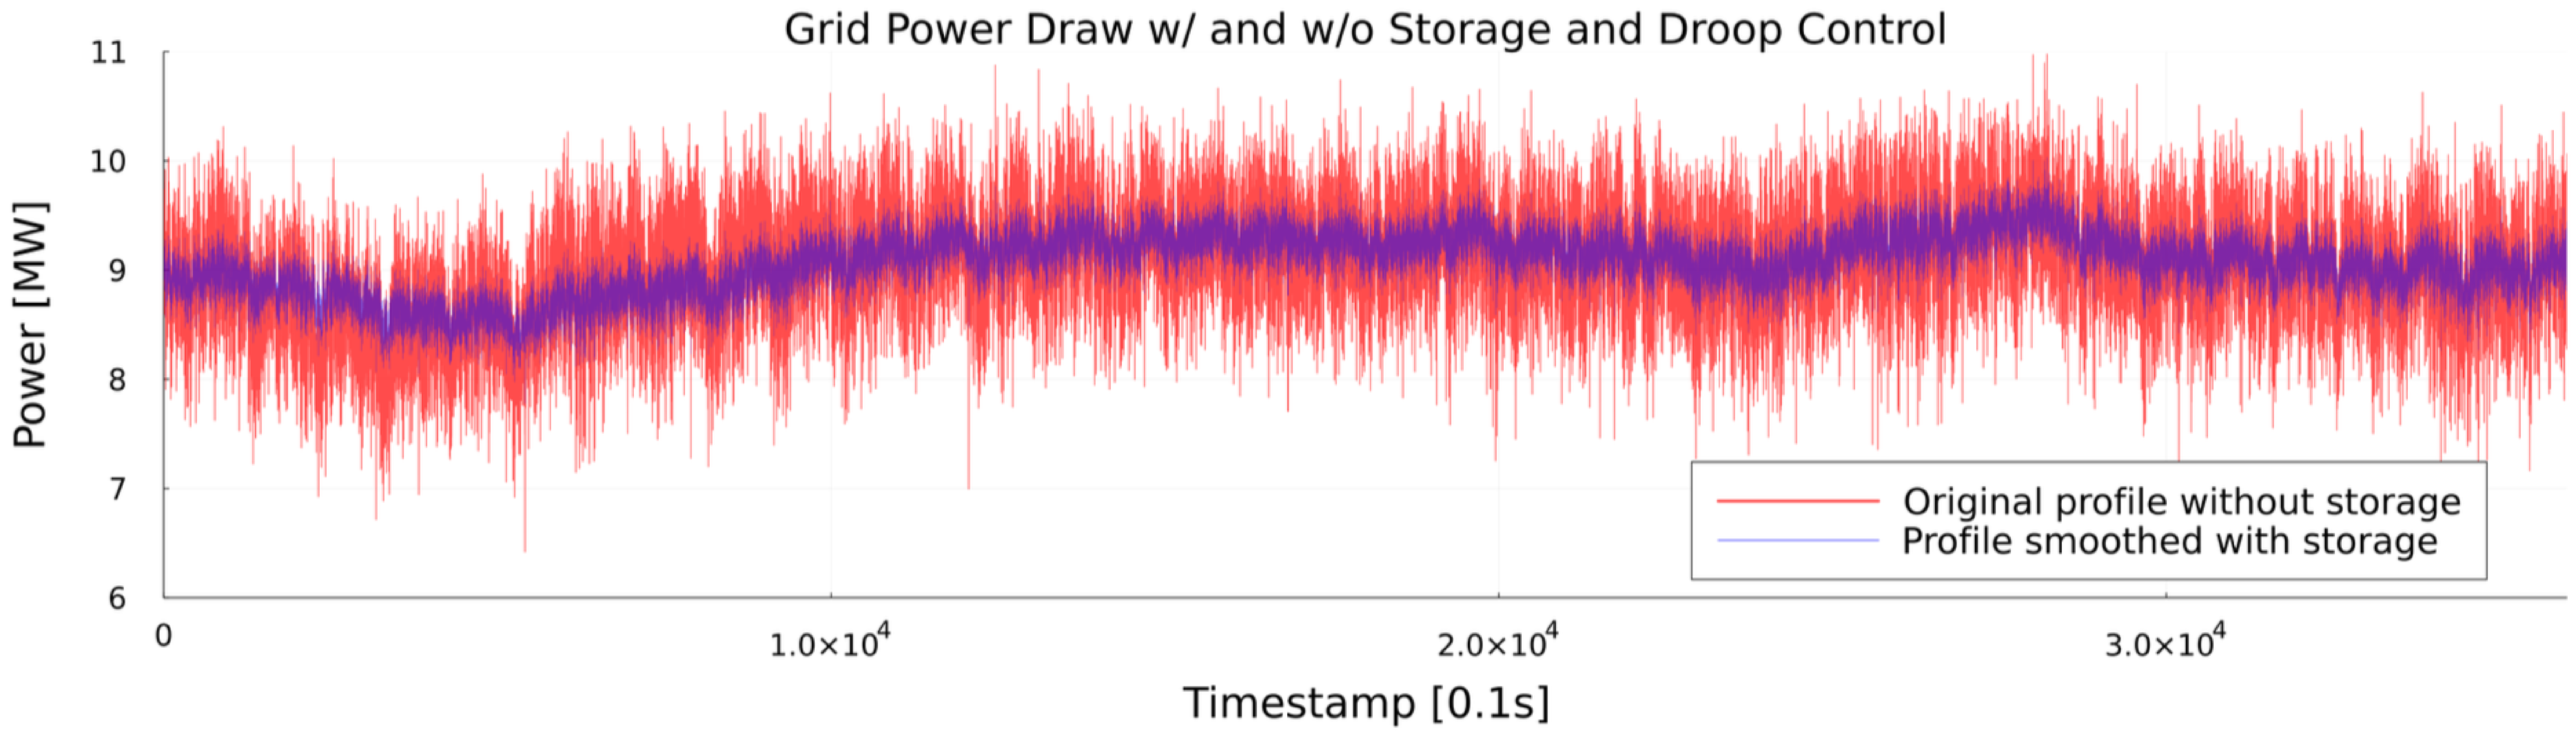

Smoothing DC loads with energy storage

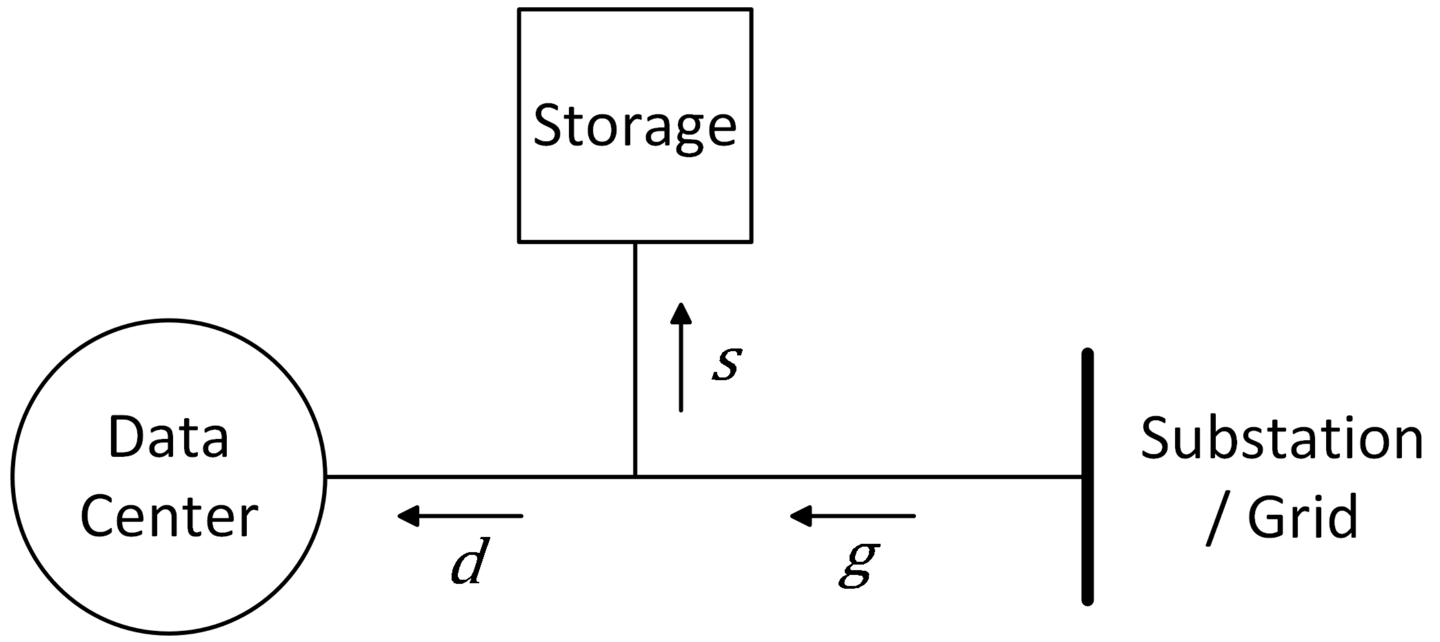

- **Goal:** Smooth the power profile observed by the gird/substation
- **Drop-control strategy** acting solely on measurements
- **Feedback optimization strategy** acting on measurements **and** the intensity of constraint violations (mixed-saddle flow)

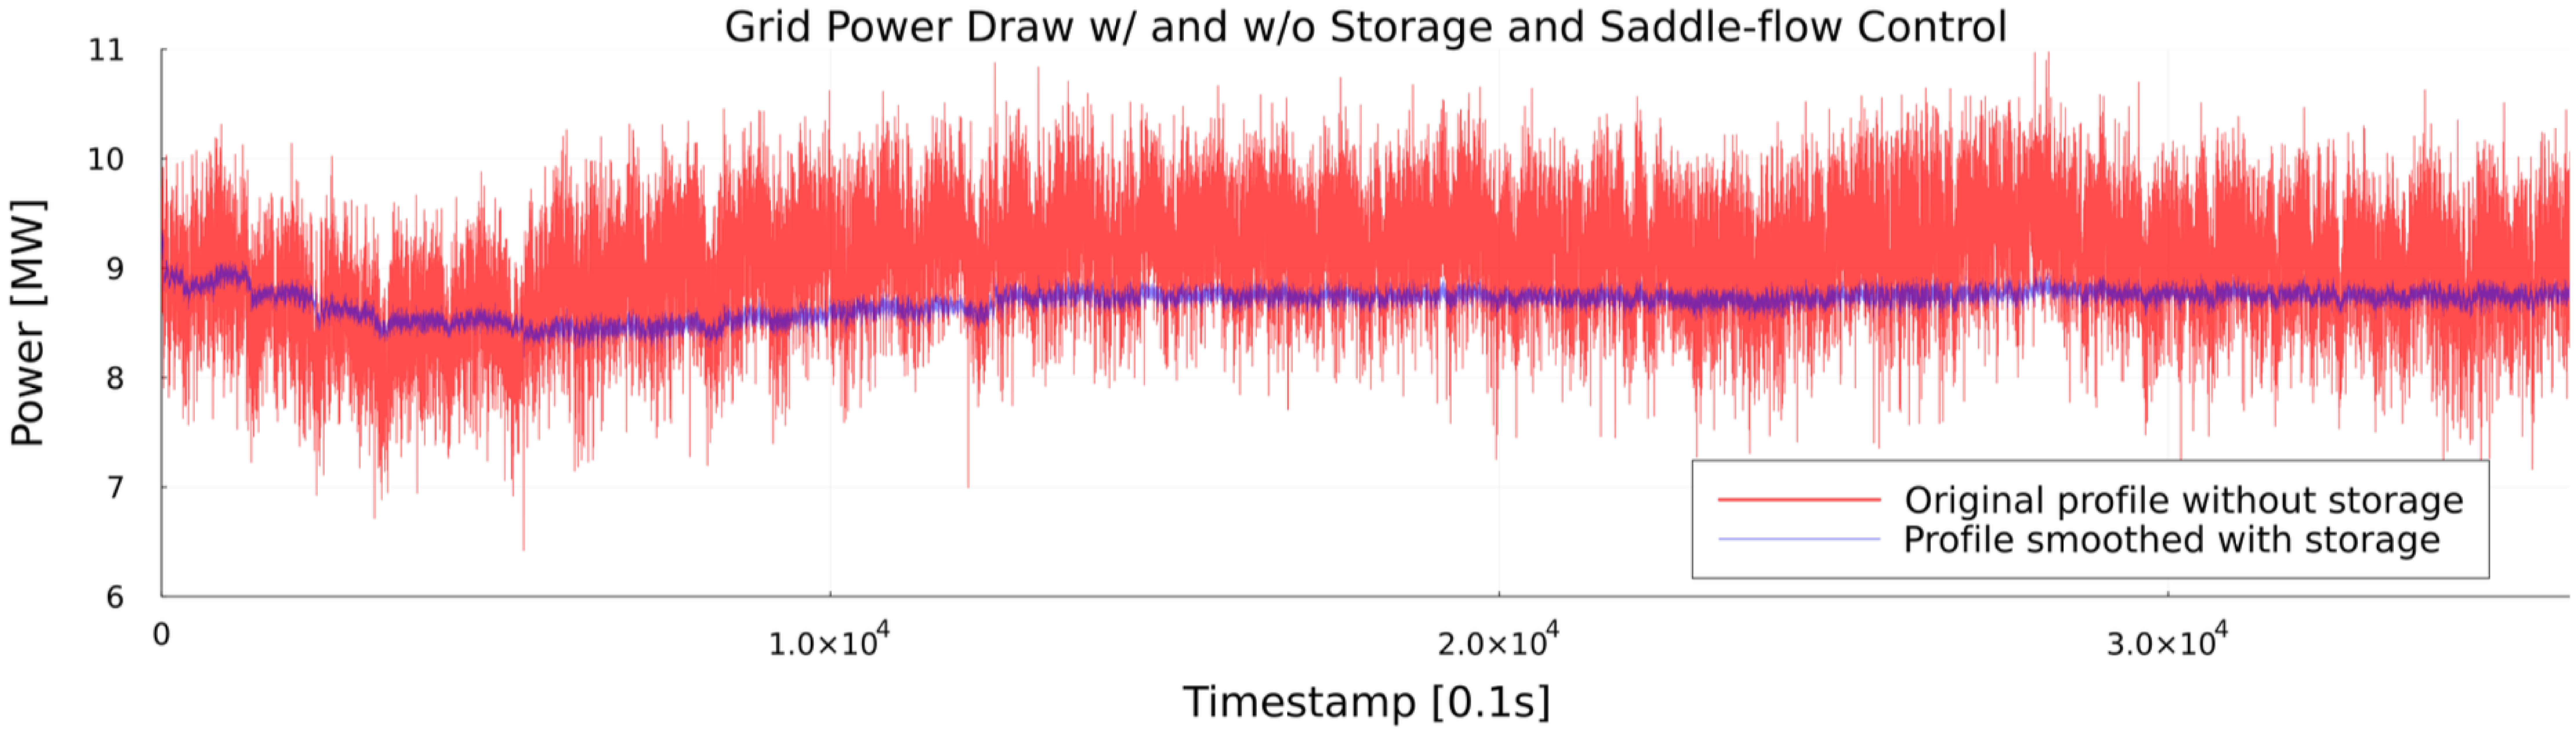

Summary

Hierarchical optimization models for coordinating spatial DC flexibility

- DC optimization is embed into grid optimization to coordinate demand shifts
- Can efficiently be approximated by contextual regression ensuring cost-optimality and feasibility

Market-based incentives to invoke spatial DC flexibility

- Difference of LMPs provide natural incentives for spatial demand redistribution
- Enabling technology: optimization over trained neural networks approximating the price/revenue function
- Depending on the grid loading, cost-savings are up to 50% for DCs and the power grid

Distribution-level energy storage solutions to smooth DC loads

- Re-purposing classic droop control to smooth smooth DC loads by energy storage
- Gradient-flow algorithms acting on measurements and constraint violations for better performance

References:

1. V. Dvorkin. Agent coordination via contextual regression (AgentConcur) for data center flexibility. 2025
2. X Liu, V Dvorkin. Optimization over trained neural networks: Difference-of-convex algorithm and application to data center scheduling. 2025