Who Gets Connected?
A Classroom Game for Large-Load Interconnection

Minghao Mou and Junjie Qin · {mmou, jq}@purdue.edu
Elmore Family School of Electrical and Computer Engineering, Purdue University

PURDUE UNIVERSITY

The challenge

Data centers ask to connect faster than grids can study them: about 474 GW in ERCOT alone, now studied in batches [2].

Our goal

Teach load and resource siting, interconnection, and AI-assisted decision making with a classroom game on a 2,000-bus synthetic Texas grid [1].

Code, data, tutorial
github.com/MhaoMou/ercot-connection-game

The activity

ECE 51032 *Computational Methods for Power System Analysis*, Purdue, Fall 2026: teams of undergraduate and master's students act as large-load developers for three weeks.

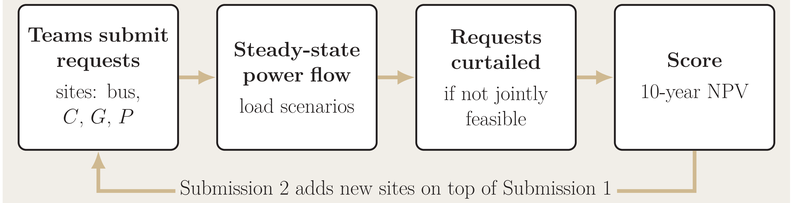


Figure 1. One round of the game.

Week 1: practice	Week 2: graded	Week 3: graded
Nov 16–29	Nov 30 – Dec 6	Dec 7–13, + market impact

What students practice

- Hands-on planning on a realistic grid
- Data analytics
- Using AI agents, and checking their results against engineering principles and optimization outputs
- Strategy against other teams
- Presenting and defending a plan

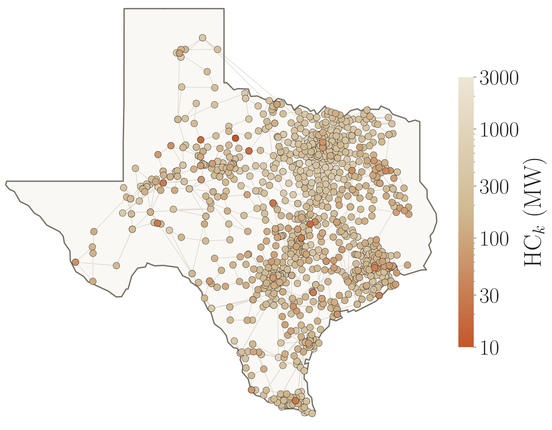


Figure 2. A first task: hosting capacity at every bus on the peak day (median 372 MW).

How the game works

Each site j of team i

A bus and a load $C_{i,j}$. Optionally bringing its own generator $G_{i,j} \le C_{i,j}$ and a 4-hour battery $P_{i,j}$ that the grid operator dispatches.

Batch allocation: serve the most MW

If the batch does not fit, every site j gets a share $\alpha_{i,j}$:

$$\max_{\boldsymbol{\alpha}\in[0,1]} \sum_i \sum_j \alpha_{i,j} C_{i,j} \quad \text{s.t. the DC-PF is feasible in every hour.}$$

ERCOT's batch studies do not specify how to split capacity when a batch does not fit [5].

Scoring

Ten-year net present value: \$1,000 per MWh served, minus fuel, capital, and energy at the LMP. In Week 3, students must simulate how their loads move wholesale prices.

What students learn: hosting capacity

Hosting capacity of bus k: the largest constant new load L the grid allows.

$$\mathrm{HC}_k := \max_{L,\,\mathbf{p}^{\mathrm{g}}_t,\,\mathbf{p}_t} L \quad \text{s.t.} \quad \mathbf{p}_t = \mathbf{p}^{\mathrm{g}}_t - \mathbf{d}_t - L\,\mathbf{e}_k, \quad \mathbf{1}^\top \mathbf{p}_t = 0,$$
$$|\mathbf{H}\,\mathbf{p}_t| \le \mathbf{F}, \quad \mathbf{p}^{\min} \le \mathbf{p}^{\mathrm{g}}_t \le \mathbf{p}^{\max} \quad \forall t.$$

It ignores N-1 contingency [3], voltage, etc.

Hosting capacities do not add up

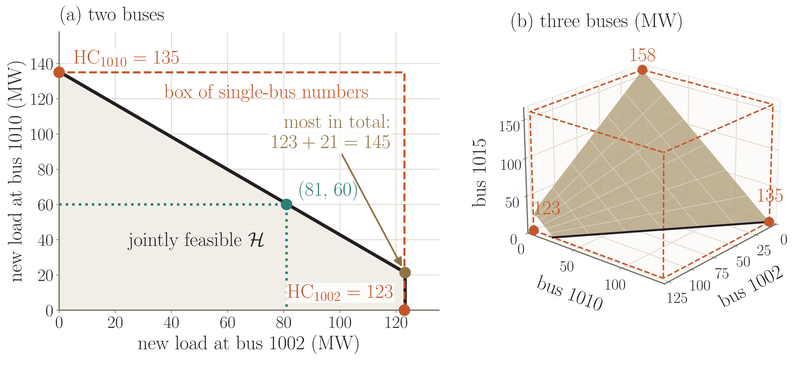


Figure 3. What the grid can host at two and three buses at once. Dashed: the single-bus numbers.

References

[1] A. B. Birchfield et al., *IEEE Trans. Power Syst.*, 32(4), 2017.
[2] ERCOT, "Assessing the Grid," Texas Senate panel, Jul. 29, 2026.
[3] H. Wan et al., arXiv:2610.18405, 2026.
[4] mhaomou.github.io/interconnection-gym.
[5] ERCOT, comments on NPRR1325 and PGRR145, Board meeting, Jun. 2026.

We thank the students of ECE 51032.

A simulated batch

Twelve scripted teams: 40 sites, **9.4** GW requested, **6.0** GW served.

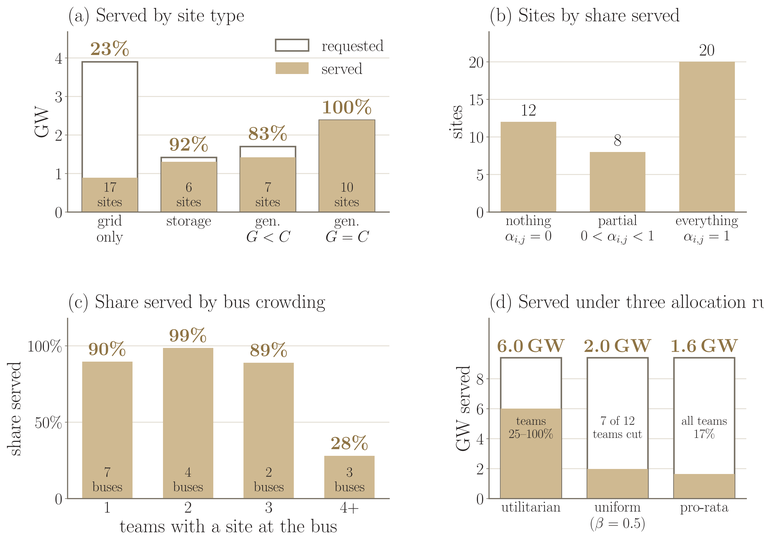


Figure 4. (a) Sites bringing generation are served while grid-only sites get 23%. (b) Most sites get all or nothing. (c) Buses shared by 4+ teams get 28%. (d) Pro-rata: two branches in one peak hour cap every site at 17%. Uniform: the largest requests are cut first.

Take-aways

- Game steps are modeled on interconnection studies (Figure 1).
- Hosting capacities do not add up (Figure 3).
- Own generation gets served: 100% vs. 23% grid-only (Figure 4a).
- Buses shared by 4+ teams: only 28% served (Figure 4c).
- The allocation rule sets what is served: 6.0, 2.0, or 1.6 GW (Figure 4d).

For utilities and ISOs

- Recruit students with hands-on interconnection practice.
- Reuse the game to train new engineers.
- Test batch allocation rules in an open sandbox [4].
- Join us: judge the final presentations, give a guest talk, or suggest a rule.

Next: (1) students will play Nov 16 – Dec 13, 2026; (2) theory and benchmark in interconnection-gym [4].