AI Data Center Duty Cycle Development

Alasdair Crawford, Jackie Hunyh, Vish Viswanathan, Vince Sprenkle

PROBLEM & MOTIVATION

- AI racks have scaled from ~5–40 kW (pre-AI) to 600 kW+ today, with next-gen systems targeting >1 MW
- This creates highly volatile rack-level power demand driven by GPU-intensive training and inference workloads
- Rack-level ESS can smooth these fluctuations to support grid code requirements and reduce stress on the genset's shaft, as well as reducing peak grid power draw and easing upstream transmission and distribution requirements
- No public data exists to generate representative battery duty cycles based on an AI duty cycle for testing and validating ESS technologies in this application

$$P_{IT\ rack} = P_{ESS} + P_{conn}$$

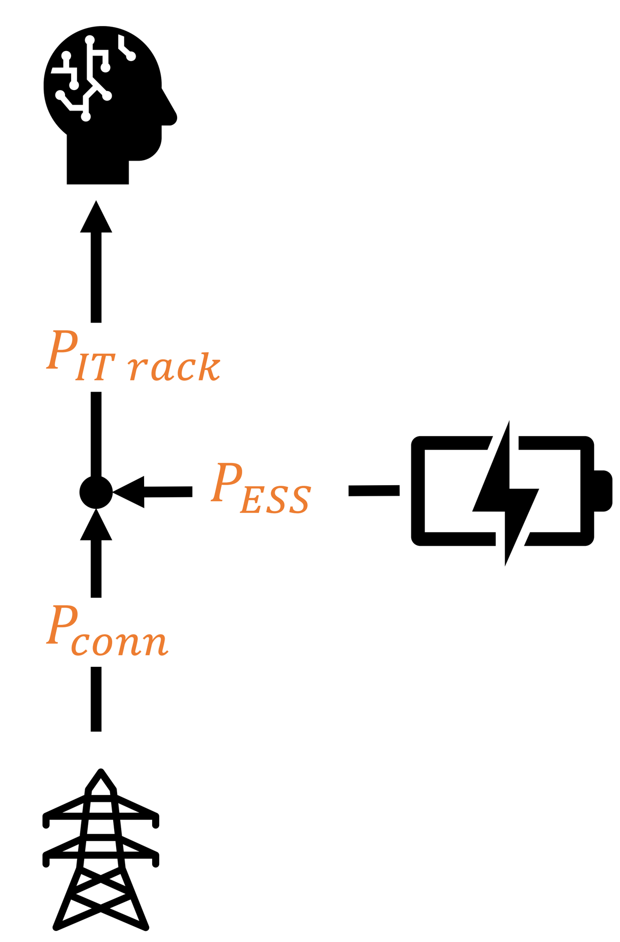

Simplified Layout

SYSTEM MODEL

- Three power components define the system, all in per-unit (1 pu = max rack power) :
- **Rack Load ($P_{IT\ rack}$):** Time-varying AI workload power consumption — always positive
- **Grid Power (P_{conn}):** Constant power from grid connection — the objective is to minimize this value and eliminate power swings
- **ESS Power (P_{ESS}):** Battery absorbs or delivers the difference

$$dSOC = \begin{cases} -P_{BESS}dt & P_{BESS} \geq 0 \\ -\eta_{RTE}P_{BESS}dt & P_{BESS} < 0 \end{cases}$$

$$SOC_{start} = SOC_{end}$$

$P_{ESS} > 0 \rightarrow$ discharge; $P_{ESS} < 0 \rightarrow$ charge

- SOC tracks energy throughput with round-trip efficiency losses
- To scale to a real installation, multiply all signals by actual max rack power (kW or MW)

INPUT: THREE AI LOAD PROFILES

- Three data center load profiles are used as inputs to generate ESS duty cycles
- Each captures AI workload power dynamics at different timescales
- All expressed in per-unit notation
- PNNL
  - EMT-derived Ross & Follum 2025, 40-sec facility-level, transformed to rack
- U of Alberta
  - GPU-derived Li et al. 2024, 2-sec GPU-level, transformed to rack
- Microsoft / OpenAI / NVIDIA
  - Choukse et al. 2025, GPU load profile

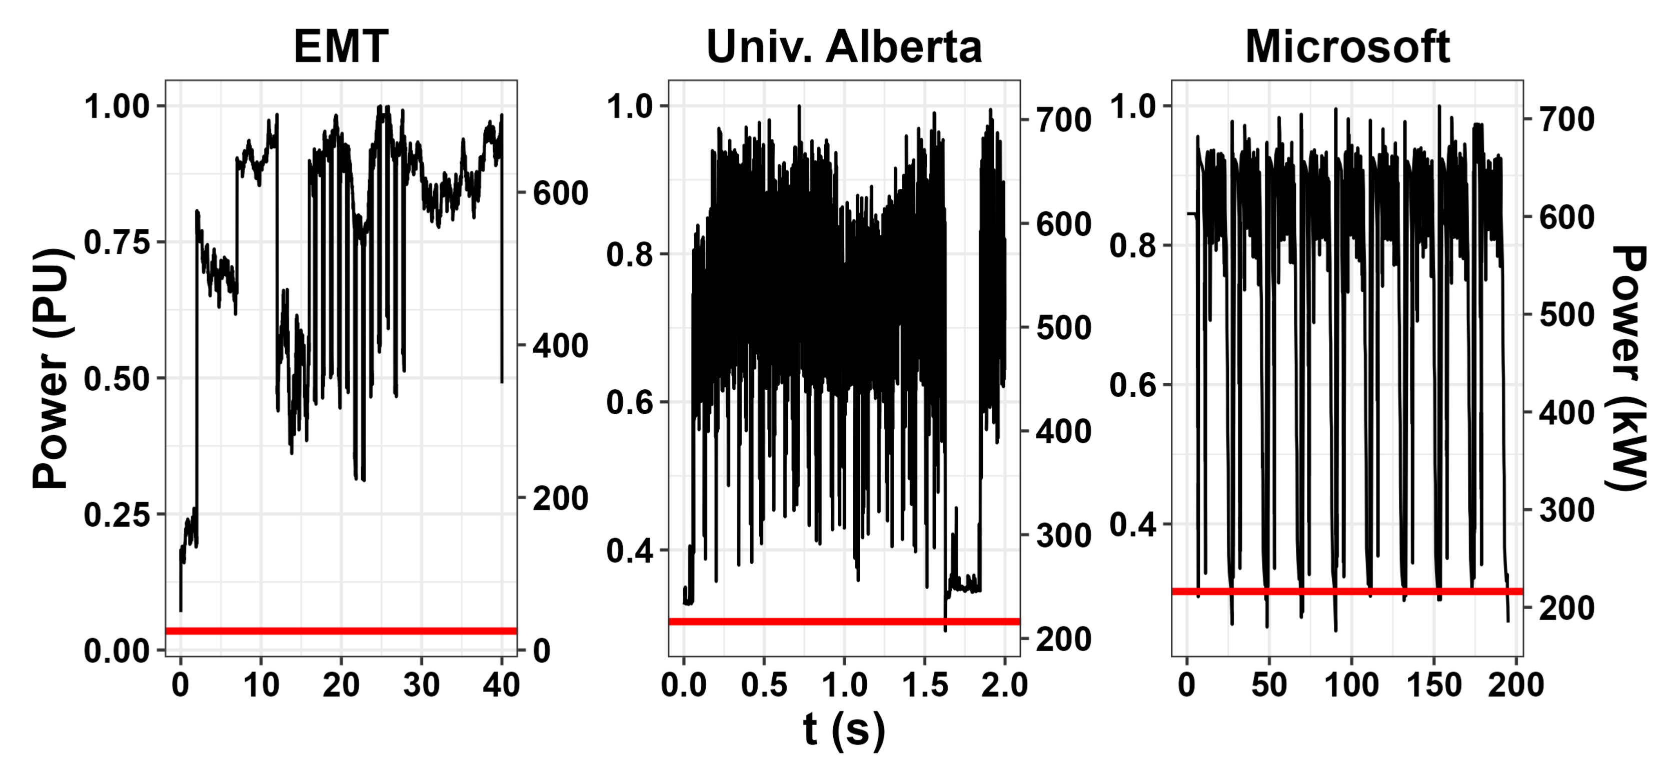

OUTPUT: ESS DUTY CYCLES

For each load profile, the methodology generates :

Steady-state duty cycle:

•Rack power consumption (the input signal)

•Required battery power profile to hold grid power constant

•The constant grid power level

•Predicted SOC trajectory for a 2-hour duration battery

Transition scenarios :

•**Start-of-task:** Rack ramps from idle to active load; battery discharges during deficit until the load profile begins

•**End-of-task:** Rack drops to idle; excess power ($P_{idle} - P_{conn}$) absorbed by battery until SOC limit is reached

•These capture the boundary conditions the ESS must handle at the beginning and end of AI workloads

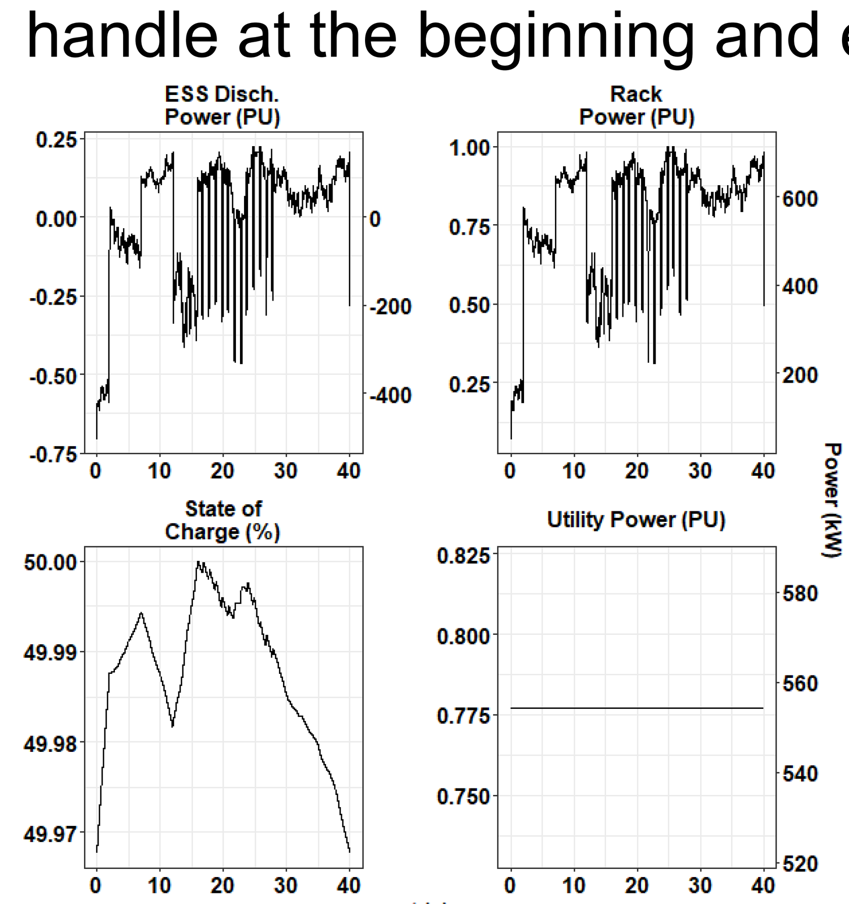

EMT steady-state repeating unit duty cycle

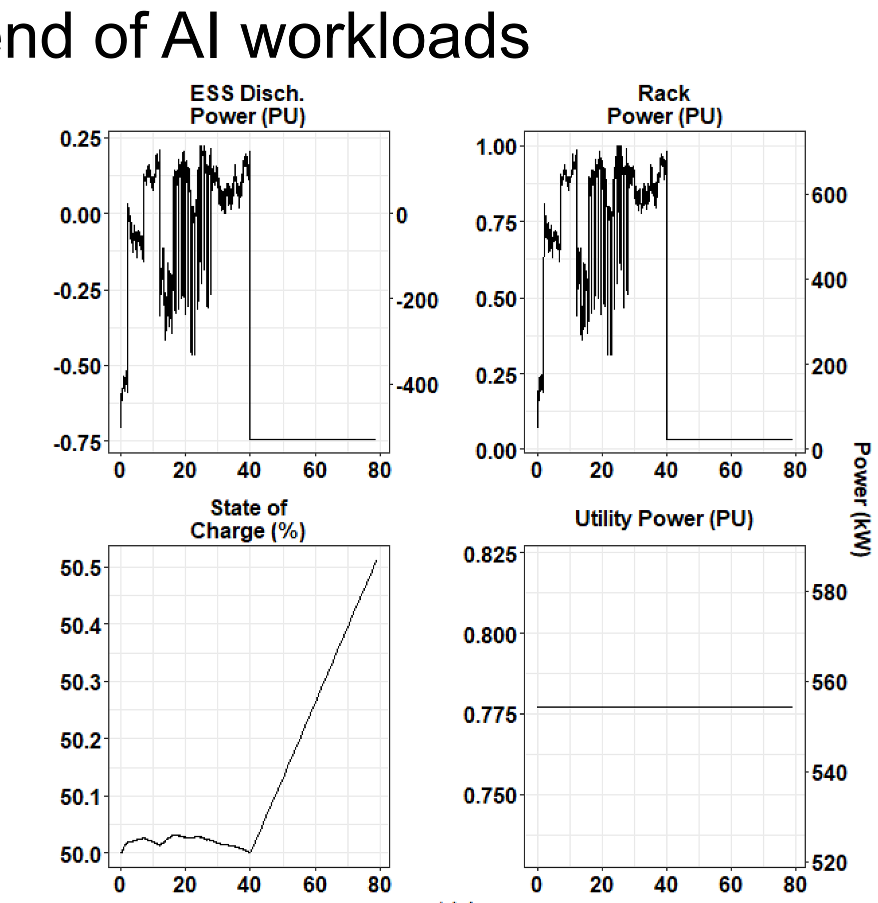

EMT End-of-task duty cycle

For additional information, contact:
Alasdair Crawford
+1 (503) 417-7551
Alasdair.Crawford@pnnl.gov