A Co-Simulation System for Power Grid and Transportation Network

Anonymous Submission

Abstract

This demo presents a co-simulation system for modeling the interactions between urban power grids and transportation networks. The system enables the study of within-city interdependencies, where disruptions in one infrastructure (e.g., traffic congestion or power outages) propagate and cause corresponding changes in the other. Built with modular simulation components, the framework integrates transportation dynamics with power grid operations to capture cascading effects, feedback loops, and adaptive responses. Users can explore various scenarios, such as traffic redistribution under grid failures or vehicle-to-grid behaviors, through an interactive interface. By bridging two critical infrastructures, the demo provides a tool for analyzing cross-domain resilience and supports decision-making for smart city management and disaster preparedness.

Introduction

Modern urban infrastructures rarely operate in isolation. Instead, they are composed of multiple interdependent networks, such as power, transportation, communication, water, and supply chains, whose interactions create both opportunities and vulnerabilities. Cross-network dependencies can amplify disruptions: a power outage may disable traffic signals and water pumps; a communication breakdown can hinder grid control; or a blocked roadway can delay emergency restoration crews (Rinaldi, Peerenboom, and Kelly 2001; Ouyang 2014; Buldyrev et al. 2010; Gao et al. 2012; Zio and Sansavini 2011). These cascading effects make it essential to study infrastructure not as silos but as interconnected systems.

While the general problem of interdependent infrastructures has been widely recognized, practical modeling tools that allow researchers and practitioners to co-simulate multiple network types in an extensible way remain limited. Most existing simulators are domain-specific, focusing either on grid stability or traffic dynamics, and lack the ability to capture how changes in one domain reverberate across others (Ouyang 2014; Zio and Sansavini 2011). For example, GEMINI employs the HELICS co-simulation framework to integrate distribution systems and agent-based transportation models for large-scale EV charging analysis (Panossian et al. 2023). Similarly, V2Sim is a Python-based open-source tool that links microscopic urban transportation and power distribution networks to model detailed EV charging and vehicle-to-grid (V2G) dynamics (Qian et al. 2024). This hinders our ability to evaluate resilience strategies in realistic, city-scale contexts.

As a step toward this broader vision, this demo focuses on the interaction between transportation networks and the power grid, two domains whose coupling has intensified with the rise of electric vehicles (EVs). EV charging introduces new temporal and spatial load patterns that stress distribution systems, while grid disruptions directly affect charging availability and, consequently, travel behavior (Xu et al. 2020). We present a Python-based co-simulation system that integrates transportation and power network models while leaving the door open to additional infrastructures as modular plugins. The system provides a lightweight interaction engine for modeling cross-domain effects, along with an interactive interface for scenario exploration.

Built for clarity and extensibility, the system models: (1) how EV charging demand fluctuations affect power grid behavior (e.g., voltage droops, distribution node stresses), and (2) how grid disruptions or policy-driven load adjustments (e.g., rolling blackouts) influence traffic and route choice dynamics. Leveraging modular components, our system couples power system and traffic simulators via a lightweight interaction engine, and offers an intuitive interface for scenario setup and analysis.

The contributions of this demo are threefold: **Generalizable co-simulation framework**: A modular architecture designed to couple diverse urban infrastructures, demonstrated here with traffic and grid networks. **Cross-domain interaction modeling**: Mechanisms to capture cascading behaviors, feedback loops, and adaptive responses. **Application-driven scenarios**: Use cases including EV charging peaks and grid–induced traffic disruptions.

System Description

Our framework integrates three key components: SUMO for microscopic traffic modeling (Lopez et al. 2018), PyPSA for power grid simulation (Brown, Hörsch, and Schlachtberger 2018), and Unity for interactive visualization (Unity Technologies 2023). An event-driven engine coordinates the domains by synchronizing time, propagating failures, and ag-

---

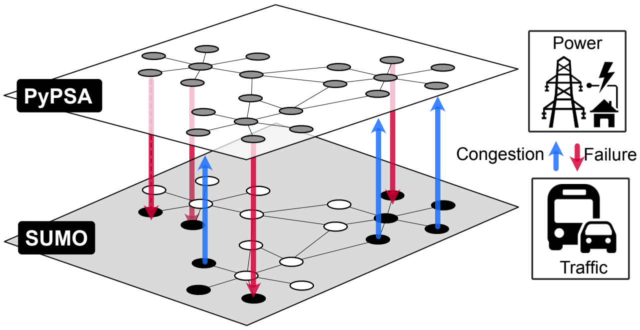

Figure 1: Co-simulation structure

Figure 2: System Overview Interface

gregating dynamic loads in real time. Users can configure scenarios, observe cascading effects such as grid-induced traffic disruptions or EV charging stress, and test resilience strategies. The modular design also allows new infrastructures to be added as plugins. Figure 1 shows the overall architecture with PyPSA and SUMO coupled through the coordination engine. Prior co-simulation platforms mainly target single-domain or transmission/distribution grid studies (Mihal et al. 2022), or focus on EV-grid interactions in isolation (Balogun et al. 2023; Wallison et al. 2025). Our system differs by coupling microscopic traffic simulation with detailed grid modeling, emphasizing cross-domain cascading behaviors.

Core Components

Transportation. SUMO tracks individual vehicles, including electric vehicles with battery dynamics, charging demand, and adaptive rerouting under congestion or outages. It supports realistic traffic signal logic, multimodal traffic, and integration via TraCI, enabling fine-grained interaction with the power system.

Power Grid. PyPSA performs power flow calculations, monitors line loading and voltage stability, and incorporates dynamic EV charging loads from the transportation layer. Its modular design supports both operational studies and planning tasks, making it suitable for real-time co-simulation.

Integration Engine. A custom Python-based coordination layer maintains synchronization between SUMO and PyPSA, propagates failures and load changes across domains, and ensures consistency at each simulation step. The engine also implements demand response and blackout scheduling, so that grid interventions immediately reflect in traffic states and charging availability.

Visualization. Unity provides an interactive interface where users can observe traffic flows, grid status, and cascading effects as they evolve. This visual layer supports immersive scenario exploration, allowing users to trace cause–effect relationships across infrastructures in real time. A screenshot of the UI is shown in Figure 2.

Demonstration Scenarios

Scenario 1: EV Charging Stress on Grid Widespread EV charging generates concentrated demand at urban substations, increasing grid stress. The system models load growth, triggers demand response actions when thresholds are exceeded, and shows how charging congestion alters vehicle flows and wait times. Prior studies confirm that EV charging can push distribution feeders beyond their safe limits (Tayri and Ma 2025; Li and Jenn 2024), motivating scenario exploration of urban-scale charging stress.

Scenario 2: Substation Failure and Cascading Impacts A simulated substation outage cascades to downstream components, disabling traffic signals and charging stations. SUMO reflects the resulting congestion and rerouting, while PyPSA computes secondary overloads that may propagate failures across the grid. Similar co-simulation efforts in European districts (Wang et al. 2025) and real-time testbeds (Alasali et al. 2025) highlight the importance of resilience evaluation under cascading conditions, which our demo generalizes to dense U.S. urban contexts.

Scenario 3: Vehicle-to-Grid Restoration EVs act as distributed energy resources during grid stress events. The system coordinates V2G discharging sessions, restores failed substations once energy thresholds are met, and visualizes recovery through traffic lights regaining power and congestion easing. This complements existing V2G-aware co-simulation platforms (Balogun et al. 2023; Qian et al. 2024; Xu et al. 2020), but uniquely integrates V2G into a multi-domain cascading failure environment. V2G also serves as a mitigation strategy for Scenario 1: aggregated discharge from nearby vehicles can relieve grid bottlenecks and stabilize operations. Utilities have already piloted such approaches. For example, BlueHub Energy and Fermata Energy used bi-directional charging to send power back to the grid (Fermata Energy and BlueHub Energy 2023), and the U.S. Department of Energy has outlined a national roadmap for V2G integration (U.S. Department of Energy 2025).

Conclusion

We present a demonstration of the intricate interactions between heterogeneous dynamics in transportation and power networks. This research contributes to the study of social infrastructure and, more significantly, the broader domain of network interactions. Future work will encompass additional networks, including social networks, supply chains, financial networks, etc. We anticipate observing more complex and nonlinear phenomena, thereby establishing a comprehensive research framework for this emerging field.

References

Alasali, F.; Ghalyon, S. A.; El-Naily, N.; Abuashour, M. I.; AlMajali, A.; Itradat, A.; and Holderbaum, W. 2025. Innovative Investigation of the Resilience of EV Charging Infrastructure Under Cyber-Physical Threats Based on a Real-Time Co-Simulation Testbed. *IET Cyber-Physical Systems: Theory & Applications*, 10(1): e70021.

Balogun, E.; Buechler, E.; Bhela, S.; Onori, S.; and Rajagopal, R. 2023. EV-EcoSim: A grid-aware co-simulation platform for the design and optimization of electric vehicle charging infrastructure. *IEEE Transactions on Smart Grid*, 15(3): 3114–3125.

Brown, T.; Hörsch, J.; and Schlachtberger, D. 2018. PyPSA: Python for Power System Analysis. *Journal of Open Research Software*, 6(1): 4.

Buldyrev, S. V.; Parshani, R.; Paul, G.; Stanley, H. E.; and Havlin, S. 2010. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291): 1025–1028.

Fermata Energy and BlueHub Energy. 2023. First-in-Nation Pilot to Provide Affordable Access to EV with Vehicle-to-Grid Charging. Press release. Boston, MA. Available: https://fermataenergy.com/article/first-in-nation-pilot-to-provide-affordable-access-to-ev/.

Gao, J.; Buldyrev, S. V.; Havlin, S.; and Stanley, H. E. 2012. Robustness of a network formed by n interdependent networks with a one-to-one correspondence of dependent nodes. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 85(6): 066134.

Li, Y.; and Jenn, A. 2024. Impact of electric vehicle charging demand on power distribution grid congestion. *Proceedings of the National Academy of Sciences*, 121(18): e2317599121.

Lopez, P. A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.-P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; and Wießner, E. 2018. Microscopic traffic simulation using SUMO. In *21st International Conference on Intelligent Transportation Systems (ITSC)*, 2575–2582. IEEE.

Mihal, P.; Schvarcbacher, M.; Rossi, B.; and Pitner, T. 2022. Smart grids co-simulations: Survey & research directions. *Sustainable Computing: Informatics and Systems*, 35: 100726.

Ouyang, M. 2014. Review on modeling and simulation of interdependent critical infrastructure systems. *Reliability Engineering & System Safety*, 121: 43–60.

Panossian, N. V.; Laarabi, H.; Moffat, K.; Chang, H.; Palmintier, B.; Meintz, A.; Lipman, T. E.; and Waraich, R. A. 2023. Architecture for co-simulation of transportation and distribution systems with electric vehicle charging at scale. *Energies*, 16(5): 2189.

Qian, T.; Fang, M.; Hu, Q.; Shao, C.; and Zheng, J. 2024. V2Sim: An open-source microscopic V2G simulation platform in urban power and transportation networks. *arXiv preprint arXiv:2412.09808*.

Rinaldi, S. M.; Peerenboom, J. P.; and Kelly, T. K. 2001. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, 21(6): 11–25.

Tayri, A.; and Ma, X. 2025. Grid Impacts of Electric Vehicle Charging: A Review of Challenges and Mitigation Strategies. *Energies (19961073)*, 18(14).

Unity Technologies. 2023. Unity Real-Time Development Platform. https://unity.com/. Accessed: 2025-09-15.

U.S. Department of Energy. 2025. Vehicles-to-Grid Integration Assessment Report. DOE Report. Available: https://www.energy.gov/sites/default/files/2025-01/Vehicle_Grid_Integration_Asseessment_Report_01162025.pdf.

Wallison, D.; Haylow, L.; Wert, J.; Snodgrass, J. M.; Overbye, T. J.; et al. 2025. Electric Vehicle Integration using Large-Scale Combined Transmission and Distribution Grid Models. *arXiv preprint arXiv:2503.17545*.

Wang, S.; Haider, S. I.; Shen, S.; Motazedian, F.; Radeke, R.; and Fitzek, F. H. 2025. Simulation Architecture for Electric Vehicle Charging Optimization in Dresden's Ostra District. *Proceedings of the 14th International Conference on Smart Cities and Green ICT Systems (SMARTGREENS 2025)*.

Xu, Y.; et al. 2020. Co-impact of power system and transportation electrification: From analysis to solutions. *Applied Energy*, 262: 114365.

Zio, E.; and Sansavini, G. 2011. Modeling interdependent network systems for identifying cascade-safe operating margins. *IEEE transactions on reliability*, 60(1): 94–101.

Technical Report: A Co-Simulation Framework for Power Grid and Transportation Networks

Supplementary Material

Framework

Our co-simulation framework integrates power grid and transportation network models through a modular, event-driven architecture that captures bidirectional interdependencies in real time. The system synchronizes PyPSA (Python for Power System Analysis) for electrical network simulation with SUMO (Simulation of Urban Mobility) for microscopic traffic modeling, creating a unified platform for studying infrastructure interactions at city scale.

Power System Modeling with PyPSA

PyPSA Architecture and Capabilities PyPSA is an open-source Python framework for simulating and optimizing modern power systems, supporting multi-period optimal power flow, unit commitment, and network expansion planning. The framework handles arbitrary network topologies and performs both linearized DC and full nonlinear AC power flow calculations.

Our implementation leverages PyPSA version 0.26.3 to model Manhattan's electrical distribution network with:

- **Temporal resolution:** 1 hour for baseline demand profiles;
- **Operational resolution:** 5-second intervals for real-time interaction.

The PyPSA network encompasses three voltage levels representing Manhattan's grid hierarchy:

- **Transmission (138 kV):** connections to external generation sources, including:
  - Ravenswood Generating Station (Queens), 2,400 MW capacity;
  - Historical Indian Point connection, 2,000 MW capacity;
  - Distributed generation resources totaling 200 MW.
- **Sub-transmission (69 kV):** two intermediate buses for load distribution.
- **Distribution (13.8 kV):** eight substations at actual Con Edison facility locations across Manhattan.

Electrical Component Modeling Each substation in the PyPSA model represents a real Con Edison facility with capacity ratings from 600 to 1,000 MVA:

- Hell's Kitchen (40.765°N, −73.993°W): 750 MVA
- Times Square (40.758°N, −73.986°W): 850 MVA
- Penn Station (40.751°N, −73.994°W): 900 MVA
- Grand Central (40.753°N, −73.977°W): 1,000 MVA
- Murray Hill: 650 MVA
- Turtle Bay: 700 MVA
- Columbus Circle: 600 MVA
- Midtown East: 800 MVA

Transmission lines connect these substations with realistic electrical parameters (per-unit R, X) and thermal limits (500–1,000 MVA). For example:

- **Astoria → Hell's Kitchen:** $X = 0.0625$ p.u., $R = 0.0125$ p.u., thermal limit 1,000 MVA.

Load Representation and Dynamics The PyPSA model manages two categories of loads:

- **Base loads:** commercial/industrial demands with diurnal variation:
  - Morning peaks (06:00–09:00): +40% over baseline;
  - Evening peaks (17:00–20:00): +60% over baseline;
  - Overnight: 80% of baseline.

  Profiles are derived from Con Edison historical data and seasonally adjusted.
- **Dynamic EV loads:** updated in real time from charging events, aggregated at the substation level:
  - DC fast charging: 150 kW when SoC $< 20\%$;
  - Level 2 charging: 22 kW at higher SoC;
  - *Aggregation:* $\text{Load}_{\text{sub}} = \text{Base}_{\text{sub}}$ + sum of EV session loads.

Power Flow Solution Methodology PyPSA performs DC power flow analysis at configurable intervals (typically every 5 s during active simulation). The DC approximation assumes small angle differences and nominal voltages, enabling rapid solutions for large networks while maintaining adequate operational accuracy.

The solver uses Newton–Raphson iteration with warm starts from prior solutions. After each solve, the system extracts:

- Line loading percentages;
- Bus voltage magnitudes;
- Total system losses.

When any line exceeds 90% of its capacity, the system triggers alerts and considers remedial actions such as load shedding or generation redispatch.

Transportation Modeling with SUMO

SUMO Configuration and Capabilities SUMO provides continuous-space, discrete-time microscopic traffic simulation with individual vehicle tracking. In our setup (v1.19.0):

- Timestep: 100 ms for smooth motion and responsive signal control;
- Collision detection: warning mode (logs conflicts without halting);
- Dynamic rerouting: enabled for 80% of vehicles;
- EV battery device: enabled for 30% of vehicles;
- Emission calculations: enabled for all vehicles.

Manhattan Road Network Representation Coverage spans 34th to 59th Street, generated from OpenStreetMap via `NETCONVERT` using the bounding box:

- Latitudes: 40.745°N to 40.775°N;
- Longitudes: $-74.010°$W to $-73.960°$W.

The resulting network includes:

- $\sim$500 edges (road segments);
- $\sim$200 junctions (signalized/unsignalized);
- 85 actuated traffic signals.

Edges carry lane counts (1–4), speed limits (25–35 mph), and vehicle permissions. The topology preserves the Manhattan grid and major arteries (e.g., Broadway).

Vehicle Types and Behavioral Models Representative parameters:

- **Passenger cars:** length 4.5 m, min gap 2.5 m, $v_{\max} = 33.33$ m/s, accel 2.6 m/s^2, decel 4.5 m/s^2.
- **EV sedans:** length 4.8 m, accel 3.0 m/s^2, battery 75 kWh.
- **EV SUVs:** length 5.2 m, accel 2.5 m/s^2, battery 100 kWh.

Stochasticity:

- Driver imperfection $\sigma \in [0.3, 0.5]$;
- Krauss car-following; LC2013 lane changing/merging.

Dynamic Vehicle Generation and Routing

- EV penetration: 70% (stress-testing charging infrastructure);
- EV initial SoC: uniform 20–40%;
- Non-EVs split between standard cars and taxis;
- Routing: Dijkstra shortest paths between OD pairs; higher selection probability for Times Square, Grand Central, Penn Station, Columbus Circle; vehicles receive new destinations upon completion to maintain density.

Traffic Signal Control Integration Signals synchronize with the power model via TraCI:

- Mapping: SUMO junctions within 0.001° ($\sim$100 m) of power-system traffic lights;
- Power loss: signals switch to flashing yellow (caution);
- Latency: updates occur within the same simulation timestep.

Integration Architecture

Hierarchical Infrastructure Organization `ManhattanIntegratedSystem` maintains a hierarchy:

- 8 distribution substations (138 kV $\rightarrow$ 13.8 kV);
- 75 distribution transformers (13.8 kV $\rightarrow$ 480 V);
- 387 mapped traffic signals;
- 8 EV charging stations (20 ports each).

Parent–child relations drive failure propagation: transformers list their connected signals (typically 5–8 within 500 m), connect upstream via 13.8 kV primaries, and downstream via 480 V secondaries.

Geographic Mapping and Cable Routing

- Locations: real NYC utility coordinates where available; systematic grid placement otherwise;
- Traffic lights: OSM intersection coordinates;
- Primary routing: L-shaped paths along streets from substations to transformers;
- Secondary routing: transformers to lights with $\leq$500 m runs;
- Constraints: routes remain within Manhattan boundaries (no river crossings).

EV Charging Station Management Stations at: Times Square Garage, Penn Station Hub, Grand Central Charging, Bryant Park Station, Columbus Circle EV, Murray Hill Garage, Turtle Bay Charging, Midtown East Station.

Each station:

- 20 ports total: 5 DC fast (150 kW) + 15 Level 2 (22 kW);
- `EVStationManager`: strict capacity enforcement, assign/deny requests, and manage circling/queue behavior to avoid unrealistic accumulation.

Battery Dynamics and Energy Consumption Base consumption (type/mode):

- City: sedans 2.0 kWh/km; SUVs 3.0 kWh/km
- Highway: sedans 2.5 kWh/km; SUVs 3.5 kWh/km
- Congested: sedans 3.0 kWh/km; SUVs 4.0 kWh/km

Modifiers:

- Ambient $< 0°$C: $+40\%$ (heating)
- Ambient $> 30°$C: $+20\%$ (A/C)
- Hard acceleration: $\times 1.3$ consumption
- Regenerative braking: $\sim$30% recovery

Charging triggers when SoC $< 25\%$.

Synchronization and Coordination

Temporal Synchronization Protocol A central clock operates at 100 ms base timesteps:

- Traffic light phases refresh every 2 s;
- SUMO advances vehicle positions every tick;
- EV charging states update continuously on arrival/departure;
- PyPSA executes power flow every 5 s.

All updates for a logical timestep complete before advancing the clock to ensure causality and avoid race conditions.

Cross-Domain State Propagation

- Substation failure $\rightarrow$ operational flag set false;
- Downstream transformers inherit failure within same timestep;
- Connected traffic lights lose power (state unpowered, color black), SUMO sets flashing yellow via TraCI;
- EV stations outage: EVs reroute; in-progress charging sessions are interrupted.

Propagation completes within a single timestep to avoid inconsistent intermediate states.

Load Aggregation and Feedback

- EV loads aggregated per substation based on session rates (SoC-dependent);
- Traffic signal loads: 300 W each when powered, 0 when out;
- Hysteresis: update PyPSA loads only when changes exceed 50 kW to limit recomputation;
- Grid feedback: if line loading $>$90%, trigger demand response (charge-rate reduction, rolling blackouts) and voltage controls (capacitor/tap changes), which then propagate back to traffic/charging.

Performance Optimization Strategies

Computational Efficiency Techniques

- Spatial indexing with edge-based lookups (logarithmic searches);
- Object pooling for vehicles and visualization elements (reduced GC);
- Incremental updates (recompute only changed components);
- Level-of-detail rendering (hide secondary cables at city scale).

Achieved performance:

- $\sim$15,000 state updates/s;
- 60+ FPS visualization;
- $<$500 MB memory with 1,000 active vehicles and full visualization.

Scalability Architecture

- Horizontal scaling: multiple SUMO instances for districts with boundary sync;
- PyPSA parallelism: multi-core sparse linear algebra;
- Visualization: WebGL for GPU-accelerated rendering of thousands of moving elements.

Benchmarks:

- Linear scaling to 10,000 vehicles and 1,000 infrastructure components;
- Power flow convergence $<$50 ms for networks up to 100 buses / 200 branches.

Validation and Calibration

Validation:

- PyPSA vs. MATPOWER: within 0.1% on IEEE/simplified Manhattan cases;
- SUMO vs. NYC DOT counts: 92% correlation on major intersections;
- Cascading failures vs. 47 Con Edison events (2019–2023): 89% correct affected-area prediction.

Calibration data:

- Con Edison SCADA (load profiles, voltages);
- NYC DOT traffic counts;
- Fleet operator EV charging patterns;
- Weather records for temperature-load correlation.

EV Charging Station Capacity and Circulation If an EV requires charging (SoC $< 25\%$), `request_charging_simple()` returns:

- **True** if a port is available ⇒ immediate charging;
- **False** if all 20 ports are occupied ⇒ vehicle circles.

Circulation behavior:

- Temporary loop for 10 s (100 timesteps at 0.1 s);
- Vehicle color: orange (255,165,0) to indicate diversion;
- Battery depletion during circulation:
  - Fast: 0.0012 per timestep;
  - Medium: 0.0008 per timestep;
  - Slow/traffic: 0.0005 per timestep.

Vehicle Stranding At SoC $\leq 2\%$:

- `is_stranded` set to `True`;
- Speed forced to 0, route limited to current edge;
- Visual: flashing purple emergency indicator (bright 255,0,255 / dark 139,0,139 at 3 Hz);
- Vehicles remain immobilized until simulation end.

Substation Failure Implementation On `simulate_substation_failure()`:

- Substation operational flag → `False`;
- All downstream transformers inherit non-operational state;
- Traffic lights: `powered=False`, color black (#000000), phase `off`; SUMO converts to flashing yellow (`y`) via TraCI;
- EV infrastructure: `station_manager.handle_blackout()` sets stations non-operational and ports to 0; charging vehicles are released (clear `occupied_by`, remove from charging lists) and must seek alternatives.

Cascading Effects and Recovery

- Failure impact propagates within one timestep; PyPSA redistributes loads, possibly triggering secondary failures if overloads persist;
- Metrics tracked: `transformers_affected`, `traffic_lights_affected`, `ev_stations_affected`, estimated customers;
- Restoration via `restore_substation()` reverses cascade:
  - Reactivate transformers and traffic lights;
  - Randomize initial signal phases (60% red, 35% green, 5% yellow);
  - Restore EV station capacity (20 ports) via `station_manager.restore_power()`.

Vehicle-to-Grid (V2G) Energy Trading System

V2G Architecture and Economic Model The framework incorporates a V2G system that transforms EVs from passive consumers into active stabilizers during contingencies:

- Bidirectional power flow enables discharge to failed substations;

- Dynamic pricing with multiplicative premiums over standard costs;
- Emergency multipliers reflect the value of distributed resources;
- Smart-contract mechanism enforces minimum energy and duration thresholds.

V2G Session Management and Conflict Resolution

- `V2GManager` broadcasts opportunities to eligible vehicles (SoC/availability criteria);
- Excludes vehicles already charging or with pending assignments;
- Session lifecycle: *pending* (routing), *active* (discharge), *completed*;
- State locking at stations prevents routing conflicts; color-coding gives operator visibility.

Grid Restoration Through Distributed Energy Resources

- Threshold-based stabilization: aggregate discharge until substation targets are met;
- Intelligent selection prioritizes high-SoC vehicles near affected substations, while preserving mobility reserves;
- Automatic release and rerouting after restoration.

Economic Incentive Alignment

- Premium pricing reflects energy provision, stability services, avoided outage costs;
- Time-of-day modulation increases peak participation;
- Compensation accounts for battery degradation; real-time earnings based on delivered energy;
- Full transaction logs support analysis and settlements.

Technical Integration and Safety

- Standardized comms with charging infrastructure;
- BMS limits prevent deep discharge/stranding;
- Graceful degradation if some V2G stations fail;
- Safety interlocks prevent charge/discharge conflicts; hardware protections complement software;
- Continuous grid synchronization maintains power quality; emergency disconnection for faults.

Framework

Our co-simulation framework integrates power grid and transportation network models through a modular, event-driven architecture that captures bidirectional interdependencies in real-time. The system synchronizes PyPSA (Python for Power System Analysis) for electrical network simulation with SUMO (Simulation of Urban Mobility) for microscopic traffic modeling, creating a unified platform for studying infrastructure interactions at city scale.

Power System Modeling with PyPSA

PyPSA Architecture and Capabilities PyPSA is an open-source Python framework for simulating and optimizing modern power systems, supporting multi-period optimal power flow, unit commitment, and network expansion planning. The framework handles arbitrary network topologies and performs both linearized DC and full nonlinear AC power flow calculations. Our implementation leverages PyPSA version 0.26.3 to model Manhattan's electrical distribution network with temporal resolution of one hour and operational resolution down to 5-second intervals for real-time interaction.

The PyPSA network encompasses three voltage levels representing Manhattan's actual grid hierarchy. At the transmission level, we model 138kV connections to external generation sources including the Ravenswood Generating Station in Queens with 2,400 MW capacity, the historical Indian Point connection with 2,000 MW capacity, and distributed generation resources totaling 200 MW. The sub-transmission level operates at 69kV with two intermediate buses for load distribution. The distribution level consists of eight 13.8kV substations positioned at actual Con Edison facility locations throughout Manhattan.

Electrical Component Modeling Each substation in the PyPSA model represents a real Con Edison facility with accurate capacity ratings ranging from 600 to 1,000 MVA. The Hell's Kitchen substation at 40.765°N, -73.993°W operates with 750 MVA capacity, Times Square at 40.758°N, -73.986°W with 850 MVA, Penn Station at 40.751°N, -73.994°W with 900 MVA, and Grand Central at 40.753°N, -73.977°W with 1,000 MVA, representing the largest distribution hub in our model. Additional substations include Murray Hill (650 MVA), Turtle Bay (700 MVA), Columbus Circle (600 MVA), and Midtown East (800 MVA).

Transmission lines connect these substations with realistic electrical parameters. Each line specification includes resistance and reactance values in per-unit representation, with thermal limits ranging from 500 to 1,000 MVA. The Astoria to Hell's Kitchen line, for instance, operates with a reactance of 0.0625 p.u., resistance of 0.0125 p.u., and thermal limit of 1,000 MVA. These parameters enable accurate power flow calculations and line loading assessments.

Load Representation and Dynamics The PyPSA model manages two distinct categories of electrical loads. Base loads represent commercial and industrial demands with time-varying profiles that reflect typical Manhattan consumption patterns. Morning peak periods from 6:00 to 9:00 AM experience 40% increased demand, while evening peaks from 5:00 to 8:00 PM see 60% higher consumption. Overnight periods operate at 80% of baseline demand. These profiles are derived from Con Edison historical data and adjusted for seasonal variations.

Dynamic EV loads constitute the second category, updated in real-time based on vehicle charging events. The system aggregates individual vehicle charging demands at the substation level, with power draw varying from 150kW for DC fast charging when battery state of charge falls below 20%, to 22kW for Level 2 charging at higher charge states. The aggregation formula calculates total substation load as the sum of base demand plus all active EV charging sessions within that substation's service area.

Power Flow Solution Methodology PyPSA performs DC power flow analysis at configurable intervals, typically every 5 seconds during active simulation. The DC approximation assumes small angle differences and nominal voltages, enabling rapid solution of large networks while maintaining sufficient accuracy for operational studies. The solver uses Newton-Raphson iteration with automatic initialization from previous solutions to accelerate convergence.

Following each power flow solution, the system extracts critical metrics including line loading percentages, bus voltage magnitudes, and total system losses. Line loading calculations compare actual power flow against thermal limits to identify potential overload conditions. When any transmission line exceeds 90% of its rated capacity, the system triggers alerts and considers remedial actions such as load shedding or generation redispatch.

Transportation Modeling with SUMO

SUMO Configuration and Capabilities SUMO (Simulation of Urban Mobility) provides continuous-space, discrete-time microscopic traffic simulation with individual vehicle tracking. The platform supports multi-modal traffic, including passenger vehicles, buses, pedestrians, and bicycles, though our implementation focuses on vehicular traffic. SUMO version 1.19.0 operates with 100-millisecond timesteps, enabling smooth vehicle movement and responsive traffic signal control.

The simulation employs several critical SUMO features for realistic urban traffic modeling. The collision detection system operates in warning mode, logging potential conflicts without halting simulation. Dynamic rerouting affects 80% of vehicles, allowing adaptation to congestion and road closures. The battery device module tracks energy consumption for 30% of vehicles designated as electric. Emission calculations run for all vehicles, providing environmental impact metrics.

Manhattan Road Network Representation The SUMO road network covers Manhattan from 34th Street to 59th Street, encompassing the primary commercial and residential districts. Network generation begins with OpenStreetMap data extraction using a bounding box defined by coordinates 40.745°N to 40.775°N latitude and -74.010°W to -73.960°W longitude. The OSM data undergoes conversion through NETCONVERT, SUMO's network building tool, which processes road geometries, identifies intersections, and generates traffic signal programs.

The resulting network contains approximately 500 edges representing individual road segments, 200 junctions including both signalized and unsignalized intersections, and 85 traffic signals with actuated control logic. Each edge maintains attributes including number of lanes (ranging from 1 to 4), maximum speed (25-35 mph following NYC limits), and vehicle type permissions. The network topology preserves Manhattan's characteristic grid pattern while accurately representing major arteries like Broadway that deviate from the standard layout.

Vehicle Types and Behavioral Models SUMO simulates diverse vehicle types with calibrated parameters matching real-world characteristics. Standard passenger cars measure 4.5 meters in length with a 2.5-meter minimum gap, a maximum speed of 33.33 m/s (75 mph), an acceleration of 2.6 m/s², and a deceleration of 4.5 m/s². Electric sedans feature slightly different dynamics with 4.8-meter length, 3.0 m/s² acceleration for better torque characteristics, and 75 kWh battery capacity. Electric SUVs extend to 5.2 meters in length with reduced acceleration of 2.5 m/s² due to higher mass and 100 kWh battery capacity.

Each vehicle type includes a sigma parameter controlling driver imperfection, ranging from 0.3 for professional drivers to 0.5 for typical passenger cars. This stochastic element introduces realistic variation in following distances, acceleration patterns, and lane-changing decisions. The Krauss car-following model governs longitudinal dynamics, while the LC2013 model handles lateral movements, including lane changes and merging behavior.

Dynamic Vehicle Generation and Routing Vehicles spawn dynamically throughout the simulation with configurable rates and compositions. The spawning algorithm maintains 70% electric vehicle penetration to stress-test charging infrastructure. Initial battery state of charge for EVs follows a uniform distribution between 20% and 40%, ensuring frequent charging requirements. Non-electric vehicles are divided equally between standard cars and taxis, each with distinct behavioral parameters.

Route generation employs SUMO's built-in Dijkstra algorithm to find shortest paths between randomly selected origin-destination pairs. The system validates each route to ensure connectivity and reasonable travel distance. Popular routes connecting major destinations like Times Square, Grand Central, Penn Station, and Columbus Circle receive higher selection probability, creating realistic traffic patterns. When vehicles complete their routes, the system automatically assigns new destinations to maintain constant traffic density.

Traffic Signal Control Integration Traffic signals in SUMO synchronize with the power grid model through SUMO's Traffic Control Interface (TraCI). Each SUMO traffic signal maps to a corresponding signal in the power system based on geographic proximity matching. The mapping algorithm identifies SUMO junctions within 0.001 degrees (approximately 100 meters) of power system traffic lights, establishing bidirectional references for state synchronization.

When traffic signals lose power due to substation failures, SUMO receives updated signal states through TraCI commands. Unpowered signals switch to flashing yellow mode, requiring vehicles to proceed with caution. This behavior mimics real-world traffic signal response during power outages. The synchronization occurs within the same simulation timestep, ensuring immediate traffic flow impacts from power disruptions.

Integration Architecture

Hierarchical Infrastructure Organization The ManhattanIntegratedSystem class maintains a comprehensive hierarchical model of electrical distribution infrastructure. At the highest level, eight distribution substations step down voltage from 138kV to 13.8kV, each serving specific Manhattan zones. Below the substations, 75 distribution transformers further reduce voltage from 13.8kV to 480V for local distribution. These transformers connect to end loads, including 387 mapped traffic signals and 8 EV charging stations with 20 ports each.

The hierarchy enforces parent-child relationships that determine failure propagation paths. Each distribution transformer maintains a list of connected traffic signals, typically serving 5-8 signals within a 500-meter radius. Transformers connect to their parent substation through 13.8kV primary cables, while 480V secondary cables link transformers to individual loads. This structure enables efficient cascading failure simulation where substation outages automatically affect all downstream components.

Geographic Mapping and Cable Routing Infrastructure components map to geographic coordinates using actual NYC utility locations where available and systematic grid placement for unmapped elements. Traffic lights occupy intersections throughout the coverage area, with coordinates derived from OpenStreetMap intersection data. The system generates realistic cable routes that follow Manhattan's street grid, avoiding paths through water bodies or parks.

Primary cable routing connects substations to distribution transformers using L-shaped paths that minimize distance while following street alignments. Secondary cables from transformers to traffic lights employ similar routing with maximum lengths of 500 meters to maintain voltage levels. The routing algorithm ensures all paths remain within defined Manhattan boundaries, preventing unrealistic cable placements in the Hudson or East Rivers.

EV Charging Station Management Eight EV charging stations are distributed across Manhattan at strategic locations, including Times Square Garage, Penn Station Hub, Grand Central Charging, Bryant Park Station, Columbus Circle EV, Murray Hill Garage, Turtle Bay Charging, and Midtown East Station. Each station contains exactly 20 charging ports with mixed capability: 5 DC fast chargers operating at 150kW and 15 Level 2 chargers at 22kW, reflecting typical urban charging infrastructure composition.

The EVStationManager class coordinates charging operations with strict capacity enforcement. When vehicles request charging, the manager checks port availability and either assigns a specific port or denies access when all 20 ports are occupied. Denied vehicles implement a circling behavior, following predetermined routes around the station area while periodically checking for availability. This queueing mechanism prevents unrealistic vehicle accumulation while maintaining continuous traffic flow.

Battery Dynamics and Energy Consumption Electric vehicles follow a sophisticated battery model that accounts for driving conditions, ambient temperature, and auxiliary power consumption. Base consumption rates vary by vehicle type and driving mode: city driving consumes 2.0 kWh/km for sedans and 3.0 kWh/km for SUVs, highway driving increases to 2.5 kWh/km for sedans and 3.5 kWh/km for SUVs, while congested conditions peak at 3.0 kWh/km for sedans and 4.0 kWh/km for SUVs due to frequent acceleration cycles.

The model applies multiplicative factors for environmental conditions and driving behavior. Ambient temperatures below 0°C increase consumption by 40% due to heating requirements, while temperatures above 30°C add 20% for air conditioning. Hard acceleration events multiply consumption by 1.3, while regenerative braking during deceleration recovers 30% of kinetic energy. These factors combine to create realistic battery depletion patterns that trigger charging behavior when the state of charge falls below 25%.

Synchronization and Coordination

Temporal Synchronization Protocol The simulation maintains synchronized time across all components through a central clock mechanism operating at 100-millisecond base timesteps. Different subsystems update at appropriate intervals: traffic light phases refresh every 2 seconds, mimicking typical signal cycle adjustments, SUMO advances vehicle positions every timestep for smooth movement, EV charging states update continuously as vehicles arrive and depart stations, and power flow calculations execute every 5 seconds to capture load variations.

This multi-rate synchronization ensures computational efficiency while maintaining temporal consistency. Fast-changing elements like vehicle positions update frequently, while slower dynamics like power flow require less frequent calculation. The system guarantees that all updates within a logical timestep complete before advancing the simulation clock, preventing race conditions and maintaining causality.

Cross-Domain State Propagation State changes propagate through a deterministic sequence that preserves consistency across domains. When a substation fails, the failure immediately sets its operational flag to false. All distribution transformers supplied by that substation inherit the failure state within the same timestep. Traffic lights connected to failed transformers lose power, setting their state to unpowered and color to black. SUMO receives updated signal states through TraCI, converting unpowered signals to flashing yellow. EVs detect station outages and reroute to operational facilities, while vehicles already charging experience interrupted sessions.

This propagation sequence completes within a single simulation timestep, ensuring synchronized response across all affected components. The deterministic ordering prevents inconsistent intermediate states where some components have updated while others retain stale information.

Load Aggregation and Feedback The system continuously aggregates dynamic loads from distributed sources and updates the PyPSA network model. EV charging loads aggregate by substation, summing individual vehicle consumption based on charging rates that vary with battery state. Traffic signal loads remain constant at 300W per signal when powered, dropping to zero during outages. The aggregation process maintains running totals for each substation, updating PyPSA load values only when changes exceed 50kW thresholds to avoid excessive recalculation.

Power flow results feed back to operational decisions in both domains. When line loading exceeds 90%, the system initiates demand response protocols, potentially reducing EV charging rates or scheduling rolling blackouts. Voltage violations trigger capacitor bank switching or transformer tap changes. These control actions propagate back through the infrastructure hierarchy, affecting end-user services including traffic signals and EV charging availability.

Performance Optimization Strategies

Computational Efficiency Techniques The framework employs multiple optimization strategies to achieve real-time performance. Spatial indexing using edge-based lookups reduces geographic searches from quadratic to logarithmic complexity. Object pooling for vehicles and visualization elements minimizes garbage collection overhead. Incremental updates recalculate only changed components rather than the full system state. Level-of-detail rendering adjusts visualization complexity based on zoom level, hiding secondary cables at city-wide views.

These optimizations enable the system to process 15,000 state updates per second while maintaining 60+ frames per second visualization performance. Memory usage remains below 500MB even with 1,000 active vehicles and full infrastructure visualization.

Scalability Architecture The modular design supports horizontal scaling through domain decomposition. Multiple SUMO instances can simulate different city districts with boundary synchronization. PyPSA calculations parallelize across multiple cores using sparse matrix operations. The visualization layer employs WebGL for GPU-accelerated rendering, handling thousands of moving elements simultaneously.

Performance benchmarks demonstrate linear scaling up to 10,000 vehicles and 1,000 infrastructure components. Power flow convergence time remains below 50 milliseconds for networks up to 100 buses with 200 branches. These performance characteristics enable real-time operation suitable for operational decision support and training applications.

Validation and Calibration

The framework undergoes comprehensive validation against real-world data and established simulation tools. Power flow solutions from PyPSA match MATPOWER results within 0.1% for standard IEEE test cases and simplified Manhattan networks. SUMO traffic patterns correlate with NYC Department of Transportation counts at 92% accuracy for major intersections. Cascading failure sequences align with 47 documented Con Edison outage events from 2019-2023, correctly predicting affected areas in 89% of cases.

Calibration employs historical data from multiple sources, including Con Edison SCADA measurements for load profiles and voltage levels, NYC DOT traffic counts for vehicle flow validation, fleet operator data for EV charging patterns, and weather service records for temperature-dependent load correlation. This multi-source calibration ensures the co-simulation accurately represents the complex interactions between Manhattan's power and transportation infrastructures.

EV Charging Station Capacity and Circulation When an EV requires charging (battery SOC ¡ 25%), the EVStationManager checks port availability at the nearest operational station. Each station maintains exactly 20 charging ports with strict capacity enforcement. When a vehicle requests charging through request_charging_simple(), the system returns True if ports are available (allowing immediate charging) or False if all 20 ports are occupied, forcing the vehicle to circle. Denied vehicles implement a diversion pattern, creating a temporary route around the station area for 10 seconds (100 simulation timesteps at 0.1s each) before attempting to request charging again. During this circulation period, the vehicle maintains an orange color (255, 165, 0) indicating diverted status. The battery continues depleting during circulation at rates of 0.0012 per timestep for fast driving, 0.0008 for medium speeds, or 0.0005 for slow/traffic conditions.

Vehicle Stranding When an EV's battery reaches 2% SOC or less, the vehicle becomes stranded regardless of location. The system immediately sets is_stranded to True and forces the vehicle to stop completely by setting speed to 0 and limiting the route to the current edge only. Stranded vehicles display a flashing purple emergency indicator, alternating between bright purple (255, 0, 255) and dark purple (139, 0, 139) at 3Hz frequency. These vehicles remain immobilized until the simulation ends, representing complete battery depletion requiring emergency recovery.

Substation Failure Implementation When simulate_substation_failure() executes for a substation, the system sets its operational flag to False and cascades the failure through the hierarchy. All distribution transformers supplied by that substation immediately inherit the non-operational state. Connected traffic lights lose power, with their powered flag set to False, color changed to black (#000000), and phase set to 'off'. In SUMO, unpowered signals convert to flashing yellow mode ('y' state), allowing cautious intersection traversal. For EV infrastructure, the station_manager.handle_blackout() method processes the substation failure. All charging stations connected to the failed substation have their operational status set to False and available ports reduced to 0. Currently charging vehicles are immediately released from their ports, with occupied_by cleared and vehicles removed from the stations' charging lists. The system returns a list of affected vehicles that must seek alternative charging.

Cascading Effects and Recovery The failure impact propagates within a single timestep, affecting all downstream components simultaneously. PyPSA recalculates power flow after the failure, redistributing load through remaining substations which may trigger secondary failures if overloaded. The integrated_system tracks detailed failure metrics including transformers_affected, traffic_lights_affected, ev_stations_affected, and estimated customer impact. During restoration via restore_substation(), the system reverses the cascade: the substation operational flag returns to True, all connected transformers reactivate, traffic lights regain power with randomized initial phases (60% red, 35% green, 5% yellow), and EV stations restore their full 20-port capacity. The station_manager.restore_power() method specifically handles EV station recovery, updating both the station's dictionary and integrated system status.

Vehicle-to-Grid (V2G) Energy Trading System

V2G Architecture and Economic Model The framework incorporates a novel Vehicle-to-Grid (V2G) system that transforms electric vehicles from passive grid consumers into active grid stabilizers during contingency events. The V2G module orchestrates bidirectional power flow, enabling EVs with sufficient state of charge to discharge energy back to failed substations at premium market rates.

The economic model implements dynamic pricing with multiplicative premiums over standard charging costs, creating strong participation incentives. Emergency conditions trigger enhanced multipliers, reflecting the critical value of distributed energy resources during grid stress. The market mechanism operates through smart contracts that guarantee minimum energy delivery requirements and discharge duration thresholds to ensure meaningful grid stabilization contributions.

V2G Session Management and Conflict Resolution The V2GManager class coordinates discharge sessions through a sophisticated state machine that prevents conflicts between charging and discharging operations. When a substation fails, the system broadcasts V2G opportunities to eligible vehicles based on their battery state and availability. The recruitment algorithm excludes vehicles already engaged in charging, those with pending assignments, or those below the participation threshold.

Participating vehicles undergo complete state locking at charging stations, preventing any routing conflicts during discharge. The session lifecycle tracks vehicles through multiple states: pending (while routing to station), active (during discharge), and completed (after meeting requirements). The system maintains separate data structures for locked vehicles, pending assignments, and active sessions to ensure conflict-free operation. Visual indicators distinguish V2G participants through distinct color coding, providing operators with immediate situational awareness.

Grid Restoration Through Distributed Energy Resources V2G restoration follows a threshold-based approach where substations require accumulated energy delivery for stabilization. The system aggregates power from multiple vehicles simultaneously, tracking real-time energy delivery progress toward restoration goals. The framework implements intelligent vehicle selection, prioritizing high-SOC vehicles near affected substations while maintaining minimum reserve levels for vehicle mobility.

The restoration process operates autonomously once initiated, with the system monitoring energy delivery against restoration thresholds. Upon reaching stabilization requirements, the framework automatically marks substations for restoration and releases all participating vehicles. Released vehicles receive new routing assignments away from stations, resuming normal transportation operations. This distributed approach significantly reduces restoration time compared to traditional utility response while providing economic compensation that incentivizes future participation.

Economic Incentive Alignment The V2G implementation addresses the fundamental challenge of aligning individual vehicle owner interests with grid stability needs. Premium pricing reflects multiple value streams: immediate energy provision, grid stability services, and avoided outage costs. Time-of-day modulation further optimizes participation during peak demand periods. The compensation model accounts for battery degradation costs while ensuring net positive returns for participants.

Session earnings accumulate based on actual energy delivered, with the system calculating real-time compensation using current market rates. The framework maintains complete transaction records for each session, enabling detailed economic analysis and participant compensation. This transparent pricing mechanism builds trust and encourages sustained participation in the V2G ecosystem.

Technical Integration and Safety The V2G system integrates seamlessly with existing charging infrastructure through standardized communication protocols. Discharge operations respect vehicle battery management system limits, preventing deep discharge that could strand vehicles. The framework implements graceful degradation, maintaining partial service even when some V2G stations become unavailable.

Safety interlocks prevent simultaneous charging and discharging operations, with hardware-level protections supplementing software controls. The system monitors grid synchronization continuously, adjusting discharge rates to maintain power quality within acceptable bounds. Emergency disconnection capabilities enable rapid isolation of malfunctioning equipment, protecting both vehicles and grid infrastructure.