

Edge-Assisted Smart Campus Energy Management using Federated Learning and Context-Aware Control

Satyam Singh

Department of Information Technology Noida Institute of Engineering and Technology
Greater Noida, Uttar Pradesh, India Email:
satyamrajput8853@gmail.com


Minhaj Nezami

Assistant Professor
Department of Information Technology Noida Institute of Engineering and Technology
Greater Noida, India minhaj.nezami@niet.co.in



<https://doi.org/10.55041/ijst.v2i5.219>

Cite this Article: Singh, S. (2026). Edge-Assisted Smart Campus Energy Management using Federated Learning and Context-Aware Control. International Journal of Science, Strategic Management and Technology, 02(05). <https://doi.org/10.55041/ijst.v2i5.219>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract--University campuses consume large amounts of electricity through classrooms, laboratories, hostels, libraries, and administrative buildings whose occupancy changes throughout the day. Conventional building management systems commonly use fixed schedules and centralized analytics, which limits their ability to adapt to local comfort needs and raises privacy concerns when occupant traces are collected at scale. This paper proposes an edge-assisted smart campus energy management system that combines short-term load forecasting, federated learning, and context-aware control. The proposed design trains local forecasting models inside each building and shares only model updates with a coordination server. A lightweight policy layer then adjusts lighting, ventilation, and noncritical loads according to occupancy, weather, tariff, and academic timetable signals. Simulated evaluation on a multi-building campus scenario shows a 17.6% reduction in energy consumption, a 21.3% reduction in peak demand, and stable comfort performance compared with rule-based scheduling.

Index Terms--Smart campus, energy management, federated learning, edge computing, occupancy forecasting, demand response.

I. INTRODUCTION

Educational campuses increasingly resemble small cities. They include academic blocks, hostels, cafeterias, laboratories, sports facilities, and data-intensive digital infrastructure. The energy profile of such an environment is difficult to manage because usage is driven by lecture schedules, examination periods, seasonal weather, special events, and irregular student activity. A classroom may be fully occupied at 10:00 a.m. and empty fifteen minutes later, while a laboratory may require stable ventilation even during low occupancy. Many institutions still rely on fixed operating schedules for lighting, air conditioning, and auxiliary equipment. These schedules are easy to maintain, but they often waste power during low-usage periods and create discomfort when rooms become unexpectedly crowded. Centralized analytics can improve this situation, yet sending fine-grained occupancy and device data to a single cloud service creates privacy and bandwidth concerns.

This paper presents a privacy-preserving energy management approach for campuses. The central idea is to move learning close to the building where data is produced. Local edge nodes forecast demand and occupancy, while a federated coordinator aggregates model updates without collecting raw traces. The resulting predictions are used by a context-aware controller that selects operational actions under comfort and safety constraints.

The main contributions of this work are threefold. First, it defines a layered architecture for campus energy optimization using edge intelligence. Second, it introduces a federated forecasting workflow for learning from heterogeneous buildings. Third, it evaluates the design against rule-based and centralized baselines using energy, peak demand, and comfort metrics.

II. RELATED WORK

Smart building research has explored occupancy detection, appliance scheduling, renewable integration, and demand response. Rule-based systems remain popular because they are transparent and inexpensive, but their performance depends heavily on manual tuning. Machine learning methods can model nonlinear patterns in energy use, especially when weather and occupancy features are available.

Edge computing has become important for applications where latency, privacy, and network cost matter. In a campus setting, edge nodes can be installed per building or per floor, allowing immediate control decisions even when cloud connectivity is limited. Federated learning extends this idea by allowing multiple buildings to improve a shared model while retaining local records.


The proposed system differs from conventional smart building controllers by treating the campus as a collaborative network of buildings. Instead of enforcing one global rule set, each building learns its own patterns and contributes to a common model. This is useful because hostels, laboratories, and classrooms have different rhythms but still share broad academic and seasonal trends.

III. PROPOSED SYSTEM ARCHITECTURE

The architecture is organized into four layers: sensing, edge analytics, federated coordination, and control. The sensing layer collects meter readings, indoor temperature, humidity, occupancy estimates, and timetable events. The edge analytics layer runs short-term prediction models inside each building. The federated coordination layer aggregates model updates at scheduled intervals. The control layer translates forecasts into actions such as adjusting set points, dimming lights, or delaying noncritical loads.

Privacy-preserving edge intelligence loop

Fig. 1. Architecture of the proposed edge-assisted campus energy



management system.

A. Sensing and Context Layer

The system uses three categories of inputs. Energy inputs include smart meter readings and circuit-level load measurements. Environmental inputs include temperature, humidity, and weather forecasts. Academic-context inputs include room schedules, holidays, examination periods, and

event bookings. These signals are synchronized into five-minute windows and normalized locally before training.

B. Federated Forecasting Model

For each building b , the edge node maintains a local dataset D_b

containing recent context and load observations. A gated recurrent model estimates the next-horizon load L_{t+k} from the feature sequence $X_{t-n:t}$. During federated training, each node computes local gradients and sends only parameter updates to the coordinator. The global parameter vector is updated as

$$\theta_{g,t+1} = \sum_b \frac{|D_b|}{|D|} \theta_{b,t}$$

This aggregation gives larger buildings proportional influence while preventing raw occupancy traces from leaving the local site. Secure transport and update clipping are applied to reduce information leakage from individual updates.

C. Context-Aware Control

The control layer converts forecasts into practical actions. It minimizes a cost function that includes expected energy cost, peak penalty, and comfort deviation. Hard constraints prevent actions that violate safety limits, equipment cycling restrictions, or laboratory ventilation requirements. For example, a classroom with low predicted occupancy can use a wider cooling band, while a scheduled seminar hall receives preconditioning before arrival.

D. Demand Response Scheduling

The controller separates campus loads into critical, flexible, and deferrable groups. Critical loads include laboratory ventilation, server rooms, and safety lighting. Flexible loads include classroom cooling, corridor lighting, and water pumping within allowed time windows. Deferrable loads include battery charging, selected laundry equipment, and nonurgent maintenance operations. This classification helps the policy layer reduce peaks without disrupting academic work.

When the tariff signal indicates an upcoming peak interval, the controller first searches for low-impact actions. It may pre-cool rooms before occupancy, slightly dim daylight-supported corridors, or shift pumping to a lower-price interval. Actions are ranked by predicted energy benefit and comfort risk, so the least disruptive choices are executed first.

IV. METHODOLOGY

A synthetic campus simulator was created using building categories commonly found in engineering institutions. The simulated campus contains academic blocks, hostels, a library, administrative offices, and laboratory spaces. Each building has distinct operating hours, occupancy volatility, and thermal sensitivity. Weather inputs include dry-bulb temperature, humidity, and solar intensity.

The baseline rule scheduler follows fixed time bands for lights and air conditioning. The centralized baseline trains a single forecasting model using all building records. The proposed method trains local models and aggregates them through federated averaging. All controllers operate under identical comfort limits so that energy savings are not achieved by simply allowing unacceptable indoor conditions.

A. Evaluation Metrics

Performance is measured using total energy consumption, peak demand, forecasting error, and comfort violation rate. Forecasting error is reported as mean absolute percentage error

(MAPE). A comfort violation occurs when the indoor condition remains outside the allowable band for more than ten consecutive minutes during an occupied period.

B. Experimental Settings

The simulation covers one academic semester of 120 days. The first 70% of the time series is used for training, 10% for validation, and 20% for testing. Edge nodes exchange model updates every six hours. The controller uses a fifteen-minute decision interval, which balances responsiveness with equipment stability.

C. Privacy and Reliability Considerations

Because campus occupancy data can reveal sensitive movement patterns, the architecture avoids transferring raw

sensor records. Local preprocessing removes personal identifiers,

and the federated workflow shares only bounded model updates. If the coordinator becomes unavailable, each building continues operating with its most recent local model and rule-based fallback limits.

D. Implementation Workflow

The implementation is divided into calibration, local training, federation, and deployment phases. During calibration, each building records baseline energy behavior and validates sensor quality. During local training, the edge node learns short-horizon demand patterns using only building-specific records. During federation, the coordinator aggregates model parameters and returns the updated global model. During deployment, the controller evaluates decisions continuously and logs only summarized performance indicators.

The workflow is intentionally modular. A campus may start with only meter-level forecasting and later add room-level occupancy or renewable generation data. This staged adoption is useful for institutions where budget and infrastructure upgrades occur gradually.


V. RESULTS AND DISCUSSION

Table I compares the forecasting performance of three approaches. The federated model obtains lower MAPE than the rule-driven predictor and approaches the performance of a centralized model while preserving local data boundaries.

TABLE I
Forecasting and Comfort Performance

Model	MAPE	Peak Error	Comfort Viol.
Rule Forecast	14.8%	18.9%	4.6%
Centralized GRU	8.1%	9.4%	2.8%
Federated Edge (Proposed)	8.7%	10.2%	2.5%

The small gap between centralized and federated prediction is expected because buildings do not expose all local correlations to the shared model. However, the privacy benefit is significant, and the performance loss is modest enough for operational deployment.



Energy saving compared with fixed scheduling

Fig. 2. Comparative energy savings across baseline and proposed control strategies.

A. Energy and Peak Demand

The proposed approach reduces total energy use by 17.6% compared with fixed scheduling. Peak demand decreases by 21.3%, mainly because the controller delays flexible loads and performs preconditioning before crowded time slots. The cloud-based centralized baseline saves 18.2% energy, but it requires continuous transfer of raw building-level data.

B. Comfort Analysis

Comfort violations remain low because control actions are constrained by occupancy and room type. Laboratories receive stricter ventilation bounds than classrooms, while hostels allow smoother set-point transitions during night hours. This differentiated treatment is important in campuses where a single comfort policy rarely fits every space.

C. Communication Overhead

Federated updates reduce network traffic because edge nodes transmit compact model parameters instead of continuous sensor streams. In the simulated deployment, update traffic is less than 6% of the raw telemetry volume. This makes the design suitable for institutions with limited network infrastructure.

D. Ablation Study

An ablation study was performed by removing one component at a time. Without timetable context, the model missed sharp occupancy transitions between lecture periods and produced a 2.9 percentage-point increase in MAPE. Without federated aggregation, small buildings such as administrative offices learned more slowly because their local datasets were limited. Without comfort-aware constraints, energy savings increased slightly, but the violation rate became unacceptable for occupied rooms.

E. Practical Deployment Notes

A practical deployment should begin with the buildings that have the highest energy intensity and reliable metering. Edge nodes can be attached to existing building management systems through standard protocols, while manual override remains available for facility staff. The human-in-the-loop design is important because maintenance teams understand local exceptions that may not appear in historical data.

VI. DISCUSSION

The results indicate that campus energy optimization benefits from both prediction and local autonomy. A building-level edge node can respond quickly to occupancy changes, while the federated coordinator helps smaller buildings learn from campus-wide trends. The method is especially useful during irregular academic periods, such as examinations or technical festivals, when fixed schedules become inaccurate.

There are still practical challenges. Sensor calibration affects forecast quality, and legacy electrical systems may not support fine-grained control. The framework therefore includes fallback rules and gradual action limits. Another concern is model drift during semester breaks, which can be handled through periodic validation and retraining.

VII. CONCLUSION

This paper proposed an edge-assisted smart campus energy management system using federated learning and context-aware control. The approach reduces energy consumption and peak demand while preserving local data privacy and maintaining comfort constraints. Simulated results show that the proposed method can achieve strong operational gains without requiring raw occupancy data to be centralized.

Future work will focus on integrating renewable generation forecasts, battery scheduling, and real deployment data from

academic buildings. Additional privacy mechanisms such as differential privacy may further strengthen the system for large-scale institutional use.

REFERENCES

- [1] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5-32, 2001.
- [2] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [3] J. Konecny et al., "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [4] H. B. Gunay, W. O'Brien, I. Beausoleil-Morrison, and B. Huchuk, "On adaptive occupant-learning window blind and lighting controls," *Building Research & Information*, vol. 42, no. 6, pp. 739-756, 2014.
- [5] A. Afram and F. Janabi-Sharifi, "Theory and applications of HVAC control systems: A review of model predictive control," *Building and Environment*, vol. 72, pp. 343-355, 2014.
- [6] K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, and H. V. Poor, "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3454-3469, 2020.